

A Distributed Calculus for Rôle-Based Access Control

Chiara Braghin

Dip. Informatica
Univ. "Ca' Foscari" di Venezia

Daniele Gorla

Dip. di Informatica
Univ. di Roma "La Sapienza"

Vladimiro Sassone

Dept. of Informatics
University of Sussex

Abstract

Rôle-based access control (RBAC) is increasingly attracting attention because it reduces the complexity and cost of security administration by interposing the notion of *rôle* in the assignment of permissions to users. In this paper, we present a formal framework relying on an extension of the π -calculus to study the behaviour of concurrent systems in a RBAC scenario. We define a type system ensuring that the specified policy is respected during computations, and a behavioural equivalence to equate systems. We then consider a more sophisticated feature that can be easily integrated in our framework, i.e., the possibility of automatically adding rôle activations and deactivations to processes to be run under a given policy (whenever possible). Finally, we show how the framework can be easily extended to express significant extensions of the core RBAC model, such as rôles hierarchies or constraints determining the acceptability of the system components.

Keywords: RBAC, Process Calculi, Type Systems, Behavioural Equivalences

Introduction

Rôle-based access control (RBAC) [10, 23] has recently emerged as an alternative to classical discretionary and mandatory access controls: a standard is currently under development by the National Institute of Standards and Technology (NIST) [11] and several commercial applications directly support some forms of RBAC, e.g., Oracle, Informix and Sybase in the field of commercial database management systems. Furthermore, the RBAC technology is finding applications in areas ranging from health-care to defence, in addition to the commerce systems for which it was originally designed.

RBAC is a flexible and policy-neutral access control technology: it regulates the access of users to information and system resources on the basis of activities they need to execute in the system. The essence of RBAC lies with the notions of *user*, *rôle* and *permission*: users are authorised to use only the permissions assigned to the rôles they belong to. More specifically, RBAC allows for a preliminary assignment of permissions to rôles (thus abstracting from which users will play the various rôles at run-time). A user may then establish multiple sessions, e.g., by signing on to the system, during which he activates a subset of rôles that he is a member of. This greatly simplifies system management, as it reduces the cost of administering access control policies, as well as making the administration process less error-prone. In fact, by assigning to users predefined rôles that naturally express the organisation's structure, the administrative process of establishing permissions is streamlined, and management time for reviewing permissions assignment is reduced. Anyway, the complexity of the models (e.g., in

large systems the number of r oles can exceed hundreds or thousands) demands a structured approach to the analysis and design of such systems.

This paper aims at developing a theory for reasoning about system behaviours in a RBAC scenario; to the best of our knowledge this is the first attempt in this direction. Our reference model is the so-called RBAC96 model, introduced by Sandhu et al. in their seminal paper [23]. More advanced RBAC models include r ole hierarchies and constraints such as r ole mutual exclusion, separation of duties, delegation of authority and negative permissions. Our starting point is the π -calculus [24], which provides well-established mathematical tools for expressing concurrent and possibly distributed systems. Essentially, our idea is to equip the π -calculus with the notion of *users*: we tag processes with a (not necessarily unique) name representing the user that activated them – this is very similar to the located threads of the $D\pi$ [14]. Moreover, we add two new constructs, that enable processes to activate/deactivate r oles in the user session where they run, and we include a way to grant permissions to r oles. Thus, each process is associated with a name (representing the user owning it) and with a set ρ recording the r oles activated during the current session. Hence, the term $r\|P\|_\rho$ represents a session of the user r , running a process P with active r oles ρ . We model r ole’s activation/deactivation by exploiting the following reductions:

$$r\| \mathbf{role} R.P\|_\rho \mapsto r\| P\|_{\rho \cup \{R\}} \qquad r\| \mathbf{yield} R.P\|_\rho \mapsto r\| P\|_{\rho \setminus \{R\}}$$

Intuitively, when a process activates a r ole R during a session, R must be added to the set of activated r oles ρ and the continuation P will be executed with the set ρ updated. Vice versa for the deactivation of R .

As an example, the following system

$$client\| \mathbf{role} \mathbf{auth_client.port_80}\langle index.html \rangle.P\|_\rho \parallel server\| port_80(x).Q\|_{\rho'}$$

models the interaction between a client and a HTTP server. The system contains two users, *client* and *server*, running in parallel. It may evolve as follows. First, user *client* activates the r ole **auth client** by exercising the **role** action, which in practice would involve to authenticate himself by means of a secure certificate. Then, he sends the request to the HTTP server along the usual port 80 by performing an output action along the channel *port_80*.

The introduction of named users immediately suggests the idea of a distributed system. In such systems, as e.g. the Internet, the notion of global, non-located channels as *port_80* is quite an abstraction over what is realistically achievable. We therefore use a notion of localised channels *à la Dπ* [14], where each channel is associated with a single user. Syntactically, we implement this feature by tagging output actions to specify the user (or location) where the exchange is supposed to take place. On the other hand, input actions are not tagged with any user name, as they are supposed to take place locally. Thus, the example above may be rewritten as:

$$client\| \mathbf{role} \mathbf{auth_client.port_80}^{server}\langle index.html \rangle.P\|_\rho \parallel server\| port_80(x).Q\|_{\rho'}$$

We also allow user names to be exchanged during communications. This feature adds flexibility and realism to the language, since in distributed systems users have only a partial and evolving knowledge of their execution environment. For example, the client above can be generalised to leave the server identity unspecified and to dynamically retrieve it with an input from channel *choose_a_server*:

$$client\| \mathbf{role} \mathbf{auth_client.choose_a_server}(x).port_80^x\langle index.html \rangle.P\|_\rho$$

More details on the calculus, together with an illustrative example, will be given in Section 1.

The mapping among users, r^oles and permissions, which controls the access of subjects to objects, is achieved by a pair of relations $(\mathcal{U}; \mathcal{P})$, called *RBAC schema*. In $(\mathcal{U}; \mathcal{P})$, the relation \mathcal{U} is the r^oles-to-users association, while \mathcal{P} is the permissions-to-r^oles association. As a first contribution of this paper, in Section 2 we define a type system which complements the dynamics of the calculus: it provides us with static guarantees that systems not respecting a given RBAC schema are rejected. In the client/server example above, a client not authenticated (i.e. interacting with the server without having previously performed a **role auth_client**) would be rejected, if the RBAC schema enabled only authorised users to perform HTTP requests.

As a second contribution of this paper, in Section 3 we study the behavioural semantics of the calculus via a standardly defined (*typed*) *barbed congruence*. The behavioural semantics allows us to study the behaviour of systems, concentrating on their functionalities while abstracting from their syntax. In particular, the barbed congruence allows us to prove some interesting algebraic laws that hold in our framework. As an example, we show how RBAC schemata may change the algebraic theory of the π -calculus. Consider the following system, adapted from the client/server example above:

$$(\nu \text{port_80}^{\text{server}}:R)(\text{client}\|\text{port_80}^{\text{server}}\langle \text{index.html} \rangle.P\|_{\emptyset} \|\ \text{server}\|\text{port_80}(x).Q\|_{\rho'})$$

where $(\nu \text{port_80}^{\text{server}}:R)$ is the standard restriction operator of a typed π -calculus (it declares $\text{port_80}^{\text{server}}$ at type R and limits the visibility of the channel to *client* and *server* only). By resuming the assumption that only authorised users can perform HTTP requests, the above system is blocked because the client has not been authenticated before performing the output. On the contrary, by removing the assumption that each action must be authorised by the activation of a proper r^ole, the term above would have been equivalent to

$$(\nu \text{port_80}^{\text{server}}:R)(\text{client}\|P\|_{\emptyset} \|\ \text{server}\|Q[\text{index.html}/x]\|_{\rho'})$$

that is the term resulting from the client/server exchange ($Q[\text{index.html}/x]$ denotes the process Q where each occurrence of x has been replaced by the value *index.html*). By the way, this is exactly what would have happened in a similar term of the π -calculus, since the name *port_80* is restricted and no authorisation is needed to perform input/output actions.

As highlighted by the example above, the essence of our calculus resides in the assumption that each action can be performed only if a privilege enabling it is available in the user session where the action is executed. Since privileges are associated to r^oles, it is fundamental to properly program r^ole activations and deactivations within user sessions. To this aim, in Section 4 we describe an algorithm to automatically add r^ole activations/deactivations within a system in such a way that the resulting system can be executed under a given schema, whenever possible.

In Section 5, we describe how our simple framework can be extended to express extensions of the core RBAC model. For example, r^oles can be hierarchically ordered to reflect in a natural way the different levels of authority, responsibility and competency of the employees working in an enterprise. Moreover, the system administrator may want to enforce constraints limiting the set of r^oles that can be activated during a session. Both extensions can be expressed in a uniform and scalable way by enriching the RBAC schema.

We conclude by comparing our approach with related work in Section 6. Appendix A contains the proofs of some results stated in the paper, while Appendix B provides a sound proof technique for barbed congruence in terms of a labelled transition system and a labelled bisimulation.

This paper is an extended and revised version of [4]; with respect to the extended abstract, in this paper we give all the technical details and proofs, we expand one of the possible examples of our

<i>Rôles</i>	$R, S, \dots \in \mathcal{R}$	
<i>Names</i>	$a, b, \dots, r, s, \dots, x, y, \dots \in \mathcal{N}$	
<i>Values</i>	$m, n, \dots \in \mathcal{N} \cup \mathcal{N} \times \mathcal{N}$	
<i>Processes</i>	$P, Q ::= \mathbf{nil}$	<i>inactive process</i>
	$P \mid Q$	<i>parallel composition</i>
	$!P$	<i>replication</i>
	$(\nu a : R)P$	<i>name restriction</i>
	$[m = n]P$	<i>value matching</i>
	$a(x).P$	<i>input</i>
	$m\langle n \rangle.P$	<i>output</i>
	$\mathbf{role} R.P$	<i>rôle activation</i>
	$\mathbf{yield} R.P$	<i>rôle deactivation</i>
<i>Systems</i>	$A, B ::= \mathbf{0}$	<i>empty system</i>
	$r \llbracket P \rrbracket_\rho$	<i>user session</i>
	$A \parallel B$	<i>parallel composition</i>
	$(\nu a^r : R)A$	<i>channel restriction</i>

Table 1: Syntax of the Calculus

framework, and we show how advanced RBAC features may be added in a modular way to the general picture.

1 The Language

In this section we formally introduce our calculus. First, we define its syntax and operational semantics; then, we formalise the RBAC schema to describe the rôles-to-users and permissions-to-rôles assignments.

1.1 Syntax

Since the calculus is an extension of the π -calculus [19, 24], we assume the reader be somehow acquainted with its basic features. The syntax of the calculus is given in Table 1; we assume two countable and pairwise disjoint sets: \mathcal{R} of *rôles*, ranged over by R, S, \dots , and \mathcal{N} of *names*. Names can serve three (logically) different purposes: they can be used as user names (in this case, we prefer letters r, s, \dots), channel names (in this case, we prefer letters a, b, \dots) or input variables (in this case, we prefer letters x, y, \dots). As we have discussed in the Introduction, channels are associated with users. Thus, the set of *values* of our calculus includes not only raw names but also pairs of names, written a^r ; such pairs are called *channels* and include the name of the channel, a , and the user it is associated with, r .

Processes \mathbf{nil} , $P \mid Q$, $!P$, $[m = n]P$, $(\nu a : R)P$, $a(x).P$, $m\langle n \rangle.P$ are derived from the corresponding π -like constructs. They represent, respectively, the inactive process, parallel composition of processes, replication (to model recursive processes), value matching, restriction of channel names and input/output actions over channels. Notice that input channels are not decorated with a user name: this is a syntactic means to localise them, as input channels implicitly belong to the user they appear in. The main novelty of the calculus resides in the actions $\mathbf{role} R$ and $\mathbf{yield} R$ that implement activations/deactivations of rôles in the user session they belong to, and modify the session rôles accordingly.

<i>System</i>	$\text{FR}(-)$	$\text{BND}(-)$
$\mathbf{0}$	\emptyset	\emptyset
$r \parallel P \parallel_{\rho}$	$\{r\} \cup \text{FR}_r(P)$	$\text{BND}_r(P)$
$A \parallel B$	$\text{FR}(A) \cup \text{FR}(B)$	$\text{BND}(A) \cup \text{BND}(B)$
$(va^r : R)A$	$\text{FR}(A) \setminus \{a^r\}$	$\text{BND}(A) \cup \{a^r\}$

<i>Process</i>	$\text{FR}_r(-)$	$\text{BND}_r(-)$
\mathbf{nil}	\emptyset	\emptyset
$a(x).P$	$\{a^r\} \cup \text{FR}_r(P)$	$\{x\} \cup \text{BND}_r(P)$
$m\langle n \rangle.P$	$\{m, n\} \cup \text{FR}_r(P)$	$\text{BND}_r(P)$
$\mathbf{role} R.P$	$\text{FR}_r(P)$	$\text{BND}_r(P)$
$\mathbf{yield} R.P$	$\text{FR}_r(P)$	$\text{BND}_r(P)$
$!P$	$\text{FR}_r(P)$	$\text{BND}_r(P)$
$P Q$	$\text{FR}_r(P) \cup \text{FR}_r(Q)$	$\text{BND}_r(P) \cup \text{BND}_r(Q)$
$(va^r : R)P$	$\text{FR}_r(P) \setminus \{a^r\}$	$\text{BND}_r(P) \cup \{a^r\}$
$[m = n]P$	$\{m, n\} \cup \text{FR}_r(P)$	$\text{BND}_r(P)$

Table 2: Free and Bound Channels

The syntax of processes we have just presented is too permissive, as it also contains meaningless terms. For example, when a name represents a channel, it cannot be transmitted as such, since it makes little sense without the indication of the user owning it. Similarly, output channels must indicate the name of the user containing the invoked channel. For example, a process like $a(x).b^x\langle n \rangle.P$ can be accepted but, in order to be executed, at run-time x must be assigned a user name r which owns an input channel b^r . One of the aims of the type system in Section 2 is to restrict the admissible language terms, thus rejecting terms that contain any kind of anomalies.

Systems consist of the parallel composition of user sessions that can share private channels (the latter ones are decorated with a role as described later, in Section 1.3). A user session $r \parallel P \parallel_{\rho}$ represents a process spawned by a user named r , with code P and with $\rho \subseteq \mathcal{R}$ recording the roles activated so far. Observe that different sessions of the same user can run in parallel within a system A , either with the same or with different activated roles: this is the usual notion of sessions in RBAC models.

The constructs $(va^r : R)P$, $(va^r : R)A$ and $a(x).P$ act as binders for a , a^r and x , respectively. Thus, we need to extend the standard notion of free and bound names of the π -calculus to encompass free and bound channels too. The formal definition of functions $\text{FR}(A)$ and $\text{BND}(A)$ is given in Table 2; it exploits the auxiliary functions $\text{FR}_r(P)$ and $\text{BND}_r(P)$ for processes of user r . Alpha-conversion, written $=_{\alpha}$, is then standardly defined and it allows the renaming of bound channels and names. Throughout the paper, we always assume that bound channels and names are pairwise distinct and different from the free ones; by using alpha-conversion, this requirement can be always satisfied.

Finally, observe that user names cannot be restricted, since the creation of a new user is a sensitive operation: it has to be performed only by the system administrator, as it may affect the RBAC policy underlying the entire system.

Notation. In this paper, we use ‘ $_$ ’ as a generic placeholder; thus, we denote with $\widetilde{_}$ a (possibly empty) tuple of entities of kind $_$. Moreover, we write $\widetilde{a^r} : \widetilde{R}$ to denote the tuple $a_1^r : R_1, \dots, a_k^r : R_k$, for $k \geq 0$. Sometimes, we shall use tuples as sets (i.e., without considering the order of their elements) and we write, e.g., $b^s \in \widetilde{a^r}$ or $b^s : S \in \widetilde{a^r} : \widetilde{R}$. We use $\prod_{i=1}^k P_i$ as a shorthand for $P_1 \mid \dots \mid P_k$. Finally, as usual, we will omit trailing inactive processes.

Axioms for Structural Congruence:	
(S-ALPHA) $A \equiv B \quad \text{if } A =_\alpha B$	(S-ID) $A \parallel \mathbf{0} \equiv A$
(S-PCOM) $A \parallel B \equiv B \parallel A$	(S-ASS) $(A \parallel B) \parallel C \equiv A \parallel (B \parallel C)$
(S-EQ) $r \llbracket [n = n]P \rrbracket_\rho \equiv r \llbracket P \rrbracket_\rho$	(S-LOCCH) $r \llbracket (\nu a^r : R)P \rrbracket_\rho \equiv (\nu a^r : R)r \llbracket P \rrbracket_\rho \quad \text{if } a \neq r$
(S-REPL) $r \llbracket !P \rrbracket_\rho \equiv r \llbracket P \mid !P \rrbracket_\rho$	(S-RCOM) $(\nu a^r : R)(\nu b^s : S)A \equiv (\nu b^s : S)(\nu a^r : R)A$
(S-SPLIT) $r \llbracket P \mid Q \rrbracket_\rho \equiv r \llbracket P \rrbracket_\rho \parallel r \llbracket Q \rrbracket_\rho$	(S-EXT) $(\nu a^r : R)A \parallel B \equiv (\nu a^r : R)(A \parallel B) \quad \text{if } a^r \notin \text{Fr}(B)$
Rules for Reduction Relation:	
(R-ROLE) $r \llbracket \mathbf{role} R.P \rrbracket_\rho \mapsto r \llbracket P \rrbracket_{\rho \cup \{R\}}$	(R-COM) $r \llbracket a(x).P \rrbracket_\rho \parallel s \llbracket a^r \langle n \rangle . Q \rrbracket_{\rho'} \mapsto r \llbracket P[n/x] \rrbracket_\rho \parallel s \llbracket Q \rrbracket_{\rho'}$
(R-YIELD) $r \llbracket \mathbf{yield} R.P \rrbracket_\rho \mapsto r \llbracket P \rrbracket_{\rho \setminus \{R\}}$	(R-PAR) $\frac{A \mapsto A'}{A \parallel B \mapsto A' \parallel B}$
(R-RES) $\frac{A \mapsto A'}{(\nu a^r : R)A \mapsto (\nu a^r : R)A'}$	(R-STRUCT) $\frac{A \equiv A' \quad A' \mapsto B' \quad B' \equiv B}{A \mapsto B}$

Table 3: Dynamic Semantics of the Calculus

1.2 Dynamic Semantics

The dynamics of the calculus is given in the form of a *reduction relation*. As customary, the reduction semantics is based on an auxiliary relation called *structural congruence* which allows to freely re-arrange systems in order to make reduction rules applicable. The key feature of the structural congruence is to equate terms that describe the same system; indeed, the syntax of the calculus provides a way to describe system behaviours, and the same behaviour can be described in different ways. For example, $A \parallel B$ describes a system of two parallel components that coincides with the system described by $B \parallel A$. The reason to split reductions and structural rules is to reserve reductions for actual computations, i.e. where the system actually performs some action, and keep them free of spurious term manipulation artifacts. In this way, reductions reflect at a glance the foundational building blocks of the computation, at the chosen abstraction level.

The structural congruence relation, \equiv , is the least congruence on systems that is closed under the rules of the upper part of Table 3. Rule (S-ALPHA) equates alpha-convertible systems. Rules (S-ID), (S-PCOM) and (S-ASS) state that ‘ \parallel ’ is a commutative monoidal operator, with ‘ $\mathbf{0}$ ’ as identity. Rules (S-LOCCH), (S-RCOM) and (S-EXT) regulate the scope of restricted names. In particular, (S-LOCCH) can be used to turn a restriction of a name inside a user into a restriction over the corresponding

channel at the system level; (S-RCom) allows to swap restrictions; (S-Ext) allows to extend the scope of a bound channel to include further parallel components, provided that this does not cause any name capture. Rule (S-EQ) states that a satisfied equality test does not affect the behaviour of the continuation process. Rule (S-REPL) allows to freely fold/unfold a replicated process. Finally, rule (S-SPLIT) states that a session of user r with r oles ρ hosting two parallel processes P and Q denotes the same system as two parallel sessions of r with r oles ρ hosting P and Q in isolation. Indeed, the key issues of a session are the user owning it and the activated r oles. Rewriting one in the other represents no system computation, but only a different way of describing the same system, exactly like $A \parallel B$ and $B \parallel A$.

The reduction relation is defined by the axioms and rules in the lower part of Table 3. Rule (R-ROLE) adds R to the r oles ρ activated in the current session, while (R-YIELD) removes R from ρ . Here and in what follows, ‘ \cup ’ and ‘ \setminus ’ denote the usual union and difference operations between sets; in particular, $\rho \setminus \{R\}$ is defined even if $R \notin \rho$. Rule (R-Com) regulates the inter-process communication. It states that, whenever a process in s sends a message n along channel a^r and a process in r is waiting for a message on such a channel, an interaction occurs; as a result, n replaces each occurrence of the input variable x in the process P prefixed by the input action. Finally, rules (R-PAR) and (R-RES) state that reductions are preserved by system contexts; rule (R-STRUCT) states that structurally equivalent systems have the same reductions.

Notice that, by exploiting rules (S-SPLIT), (R-ROLE) and (R-YIELD), the user session $r \parallel \mathbf{role} R.P \mid \mathbf{yield} S.Q \parallel_\rho$ may evolve into $r \parallel P \parallel_{\rho \cup \{R\}} \parallel r \parallel Q \parallel_{\rho \setminus \{S\}}$, i.e., actions **role/yield** only affect the process thread executing them. Moreover, a single session with r ole R activated can later split in two distinct sessions; thus, a single r ole activation can be ‘closed’ by several r ole deactivations.

1.3 RBAC Schema

So far, we discussed the way in which RBAC sessions can be modelled in our calculus. We now present a way to model in our framework the remaining features of the core RBAC96 model. To this aim, we need to define the *RBAC schema*, i.e., the r oles-to-users and permissions-to-r oles associations, where permissions enable the actions a user can perform within a system (in our framework, input and output actions only).

Managing r oles and their interrelationships is a difficult and sensitive task that is often centralised and delegated to a small team of security administrators. Traditionally, the RBAC schema consists of a pair of partial functions with finite domains $(\mathcal{U}_u; \mathcal{P})$, where \mathcal{U}_u assigns r oles to users and \mathcal{P} assigns permissions to r oles. Formally,

$$\mathcal{U}_u : \mathcal{N} \rightarrow_{\text{fin}} 2^{\mathcal{R}} \quad \mathcal{P} : \mathcal{R} \rightarrow_{\text{fin}} 2^{\mathcal{A}}$$

where $\mathcal{A} \triangleq \mathcal{R} \times \{!, ?\}$ represents the set of performable actions. For notational convenience, we write the pairs $(R, !)$ and $(R, ?)$ as $R^!$ and $R^?$, respectively. Intuitively, permissions $R^?$ and $R^!$ determine the possibility of performing input and output actions over a channel of r ole R , respectively. Thus, input/output permissions are not defined in terms of channels, but of r oles. To this aim, we assume a partial function with finite domain, \mathcal{U}_c , assigning a r ole to a channel¹, i.e.

$$\mathcal{U}_c : \mathcal{N} \times \mathcal{N} \rightarrow_{\text{fin}} \mathcal{R}$$

In this way, we are flexible enough to model both the permission to communicate over a single channel (when the relation \mathcal{U}_c maps only one channel to a r ole), and the permission to communicate over the

¹Since located channels can be considered as functionalities provided by users, it seems reasonable that each channel is associated with only one r ole.

member of a group of channels (when the relation \mathcal{U}_c maps more than one channel to the same r \hat{o} le). This case may be useful in situations where more channels can handle the same kind of requests (cf. Example 1 for a possible situation). Observe that in \mathcal{A} no permission represents actions **role** and **yield**. Indeed, we assume that a r \hat{o} le can be activated if (and only if) it is assigned by \mathcal{U} to the user willing to perform the action; a r \hat{o} le can be deactivated if (and only if) it has been activated before.

For notational convenience, we merge \mathcal{U}_u and \mathcal{U}_c together and denote with \mathcal{U} their union. Moreover, we also assume that r \hat{o} les associated with channels (grouped together in \mathcal{R}_c) are not included in the domain of \mathcal{P} , i.e., \mathcal{P} matters only for r \hat{o} les assigned to users (grouped together in \mathcal{R}_u). Clearly, \mathcal{R}_c and \mathcal{R}_u are assumed to be disjoint. To sum up, in our framework RBAC schemata are defined as follows.

Definition 1.1 (RBAC Schema). A RBAC schema \mathcal{S} is a pair of partial functions with finite domains $(\mathcal{U}; \mathcal{P})$ such that

- $\mathcal{U} : (\mathcal{N} \cup \mathcal{N} \times \mathcal{N}) \rightarrow_{\text{fin}} (2^{\mathcal{R}} \cup \mathcal{R})$ such that, for any r and a^s in the domain of \mathcal{U} , it holds that $\mathcal{U}(r) \in 2^{\mathcal{R}}$ and $\mathcal{U}(a^s) \in \mathcal{R}$;
- $\mathcal{P} : \mathcal{R}_u \rightarrow_{\text{fin}} 2^{\mathcal{A}}$, where $\mathcal{A} \triangleq \mathcal{R}_c \times \{!, ?\}$, $\mathcal{R}_u \cup \mathcal{R}_c \subseteq \mathcal{R}$ and $\mathcal{R}_u \cap \mathcal{R}_c = \emptyset$.

To conclude the presentation of our language, we now give an example using the features introduced so far.

Example 1. Let us now formalise a scenario where a bank client is waiting to be served by one of the branch cashiers available. There are two users, r and s , representing respectively the client and the bank branch, while cashiers are modelled as channels named c_1, \dots, c_n located at user s . The r \hat{o} les available are **client** and **cashier**. The relation \mathcal{U} assigns r \hat{o} le **client** to user r and **cashier** to channels c_i , while \mathcal{P} assigns to **client** the permission to communicate with any of the cashiers, i.e., $\text{cashier}^! \in \mathcal{P}(\text{client})$. In this way, r can indistinctly interact with any of the cashier available. The overall system can be described as follows:

$$\begin{aligned}
& r \parallel \text{role client} . \text{signal}^s \langle r \rangle . \text{served}(z) . z \langle \text{req}_1 \rangle . \dots . z \langle \text{req}_k \rangle . z \langle \text{stop} \rangle . \text{yield client} \parallel_{\emptyset} \parallel \\
& s \parallel (\nu \text{free} : \text{scheduling}) (! \text{signal}(x) . \text{free}(y) . \text{served}^x \langle y \rangle \mid \prod_{i=1}^n \text{free}^s \langle c_i^s \rangle \mid \\
& \quad \prod_{i=1}^n !c_i(x) . ([x = \text{withdraw_req}] < \text{handle withdraw request} > \\
& \quad \mid [x = \text{dep_req}] < \text{handle deposit request} > \\
& \quad \mid [x = \text{stop}] \text{free}^s \langle c_i^s \rangle)) \parallel_{\rho'}
\end{aligned}$$

Once the client enters the bank (i.e., he activates r \hat{o} le **client**), he signals his presence to the bank and waits to be served. When one of the cashiers becomes available (information maintained internally by the bank via the reserved channel *free* used for cashiers' scheduling), the client is notified and can make requests along the received channel z , referring to the available cashier. Then cashiers repeatedly receive and handle requests. For simplicity, we only assume functionalities to handle money withdraw and deposit. Moreover, we do not consider the order in which clients arrive; a system of queues can however be added routinely [24]. \diamond

2 Static Semantics

We now describe a type system that provides us with static guarantees that the set of actions performed by any user during the computation respects the RBAC schema. Moreover, as already discussed when presenting the syntax of the calculus, it is also used to reject ill-formed terms.

The syntax of types is defined by the following productions:

$$\begin{array}{ll} \text{Value Types} & T ::= \rho[\tilde{a} : \tilde{C}] \mid C \\ \text{Channel Types} & C ::= R(T) \end{array}$$

Type $\rho[a_1 : R_1(T_1), \dots, a_n : R_n(T_n)]$ can be assigned to a user r belonging to roles in ρ and owning channels \tilde{a}^r ordinally of type $\tilde{R}(T)$. Type $R(T)$ can be assigned to channels of role R along which values of type T can be exchanged. Notice that the base case of the recursive definition of types is when the set $\tilde{a} : \tilde{C}$ in a type $\rho[\tilde{a} : \tilde{C}]$ is empty.

Notation. Here and in the rest of the paper, $\mathcal{P}(\rho)$ denotes the set $\bigcup_{R \in \rho} \mathcal{P}(R)$. Moreover, we denote with \uplus the union of partial functions with disjoint domains.

A *typing environment* Γ is a partial mapping with finite domain from \mathcal{N} into types and it can be extended as follows:

$$\begin{aligned} \Gamma, x : T &\triangleq \Gamma \uplus \{x : T\} \\ \Gamma, a^r : C &\triangleq \Gamma', \text{ for } \Gamma'(s) = \begin{cases} \Gamma(s) & \text{if } s \neq r \\ \rho[a : C, \tilde{b} : \tilde{C}] & \text{if } s = r \text{ and } \Gamma(r) = \rho[\tilde{b} : \tilde{C}] \text{ and } a \notin \tilde{b} \end{cases} \end{aligned}$$

Remarkably, the extension of a typing environment could be undefined (e.g., if $x \in \text{dom}(\Gamma)$ in the first case or if $a \in \tilde{b}$ in the second case). A typing environment Γ can be used to type a system under a schema $(\mathcal{U}; \mathcal{P})$ only if the role information in Γ ‘respects’ the associations in \mathcal{U} . This intuition is formalised by the following definition.

Definition 2.1. Given a RBAC schema $(\mathcal{U}; \mathcal{P})$ and a typing environment Γ , we say that Γ *respects* \mathcal{U} if, for all $r \in \text{dom}(\Gamma)$, it holds that $\Gamma(r) = \rho[a_1 : R_1(T_1), \dots, a_n : R_n(T_n)]$ with $\mathcal{U}(r) = \rho$ and $\mathcal{U}(a_i^r) = R_i$, for all $i = 1, \dots, n$.

The primary judgements of the type system are of the form $\Gamma \vdash^{\mathcal{S}} A$. Such a judgement states that A is well-formed with respect to Γ and \mathcal{S} ; this implies that A respects the RBAC schema \mathcal{S} . To infer the main judgement, we rely on two auxiliary judgements, one for values and one for processes. Judgement $\Gamma \vdash n : T$ states that the value n has type T in Γ ; judgement $\Gamma; \rho \vdash_r^{\mathcal{S}} P$ states that P respects Γ and \mathcal{S} when it is run in a session of r with roles ρ activated.

The typing rules are collected in Table 4. Most of them are self-explanatory; we only comment below the most significant ones, i.e. those related to the actions in our calculus. The idea beyond these rules is that an action can be executed only if the current session has activated a role enabling the action. Rule (T-INPUT) states that, for typing $a(x).P$ in a session of r where roles ρ are activated, we need to establish that a^r has type $R(T)$ in Γ , that inputs over a channel of group R can be performed when playing roles ρ and that P is typeable once assumed that x has type T . Rule (T-OUTPUT) is similar: it checks that an output over a channel of group R is allowed when roles in ρ are activated; moreover, it also requires that the transmitted value n can be assigned type T in Γ . Rule (T-ROLE) states that, for typing process **role** $R.P$ in a session of r where roles ρ are activated, we need to check that r can assume role R and that P is typeable for r having activated $\rho \cup \{R\}$. Rule (T-YIELD) states that process **yield** $R.P$ is legal for r only when R has been previously activated and if P is typeable for r when R is off. Finally, notice that in rules (T-RES) and (T-SYSRES) the type of the restricted channel is not tracked in the restriction construct. Indeed, for type checking purposes, it suffices to ensure that the new channel is used coherently by all the processes accessing it. To this aim, we only need to invent a suitable T when applying the rules and verify that all the accesses to the channel conform to T .

Typing Values:		
$\frac{\Gamma(r) = \rho[\tilde{a} : \tilde{C}]}{\Gamma \vdash r : \rho[\tilde{a} : \tilde{C}]}$	$\frac{\Gamma(r) = \rho[\tilde{b} : \tilde{C}, a : C, \tilde{b}' : \tilde{C}']}{\Gamma \vdash a' : C}$	
Typing Processes:		
$\frac{\Gamma; \rho \vdash_r^S \mathbf{nil}}{\Gamma; \rho \vdash_r^S \mathbf{nil}}$	$\frac{\Gamma \vdash a' : R(T) \quad R^? \in \mathcal{P}(\rho) \quad \Gamma, x : T; \rho \vdash_r^{(\mathcal{U}; \mathcal{P})} P}{\Gamma; \rho \vdash_r^{(\mathcal{U}; \mathcal{P})} a(x).P}$	
$\frac{\Gamma; \rho \vdash_r^S P \quad \Gamma; \rho \vdash_r^S Q}{\Gamma; \rho \vdash_r^S P \mid Q}$	$\frac{\Gamma \vdash m : R(T) \quad \Gamma \vdash n : T \quad R^! \in \mathcal{P}(\rho) \quad \Gamma; \rho \vdash_r^{(\mathcal{U}; \mathcal{P})} P}{\Gamma; \rho \vdash_r^{(\mathcal{U}; \mathcal{P})} m\langle n \rangle.P}$	
$\frac{\Gamma; \rho \vdash_r^S P}{\Gamma; \rho \vdash_r^S !P}$	$\frac{\Gamma \vdash r : \rho'[\tilde{a} : \tilde{C}] \quad R \in \rho' \quad \Gamma; \rho \cup \{R\} \vdash_r^S P}{\Gamma; \rho \vdash_r^S \mathbf{role} R.P}$	
$\frac{\Gamma, a' : R(T); \rho \vdash_r^S P}{\Gamma; \rho \vdash_r^S (va : R)P}$	$\frac{R \in \rho \quad \Gamma; \rho \setminus \{R\} \vdash_r^S P}{\Gamma; \rho \vdash_r^S \mathbf{yield} R.P}$	$\frac{\Gamma; \rho \vdash_r^S P}{\Gamma; \rho \vdash_r^S [m = n]P}$
Typing Systems:		
$\frac{}{\Gamma \vdash^S \mathbf{0}}$	$\frac{\Gamma \vdash r : \rho'[\tilde{a} : \tilde{C}] \quad \rho \subseteq \rho' \quad \Gamma; \rho \vdash_r^S P}{\Gamma \vdash^S r \parallel P \parallel_\rho}$	
$\frac{\Gamma \vdash^S A \quad \Gamma \vdash^S B}{\Gamma \vdash^S A \parallel B}$	$\frac{\Gamma, a' : R(T) \vdash^S A}{\Gamma \vdash^S (va' : R)A}$	

Table 4: Typing Rules

Definition 2.2 (Well-typedness). Given a RBAC schema $\mathcal{S} = (\mathcal{U}; \mathcal{P})$ and a system A , we say that A is *well-typed in \mathcal{S}* if there exists a typing environment Γ respecting \mathcal{U} such that $\Gamma \vdash^S A$.

Example 2. Let us consider the banking scenario described in Example 1. To illustrate the type system introduced above, let us give a suitable typing for the system. Let

$$T_{csh} \triangleq \mathbf{cashier}(\{\mathbf{request}\}[])$$

be the type of the cashiers, i.e., channels belonging to `role cashier` and exchanging values of type `{request}[]`. Type `{request}[]` represents the possible requests of clients; values of this type are

names belonging to `role request` which do not provide any channel. Moreover, we let

$$T_{cl} \triangleq \{\text{client}\}[\text{served} : \text{cashier_get}(T_{csh})]$$

be the type of user r . This represents users belonging to `role client` and owning a channel named `served` of type `cashier_get(Tcsh)`. Then, a suitable typing environment Γ is

$$\begin{aligned} r &\mapsto T_{cl} \\ s &\mapsto \rho'[signal : \text{cashier_req}(T_{cl}), c_1 : T_{csh}, \dots, c_n : T_{csh}] \\ \text{withdrw_req} &\mapsto \{\text{request}\}[] \\ \text{dep_req} &\mapsto \{\text{request}\}[] \\ \text{stop} &\mapsto \{\text{request}\}[] \end{aligned}$$

The system of Example 1 is then well-typed in any schema $(\mathcal{U}, \mathcal{P})$ such that Γ respects \mathcal{U} and \mathcal{P} is such that

$$\begin{aligned} \{\text{cashier_req}^!, \text{cashier_get}^?, \text{cashier}^!\} &\subseteq \mathcal{P}(\text{client}) \\ A \cup \{\text{cashier_req}^?, \text{cashier_get}^!, \text{cashier}^?, \\ &\quad \text{scheduling}^?, \text{scheduling}^!\} &\subseteq \mathcal{P}(\rho') \end{aligned}$$

where $A \subseteq \mathcal{A}$ is a set of action permissions that allow the handling of client's requests. \diamond

Example 3. In the real world, it is unrealistic to allow any bank client to ask for any kind of bank operation. For instance, when a client applies for a credit card, he is always asked for some credentials. To model this finer scenario, we let each available operation to be modelled as a specific process, which can be activated through a specific channel (e.g., channel `wdrw` handles withdraw requests and activates process P_{wdrw} , `opn` handles open account requests, `cc` handles credit card requests). The communication along different channels requires different `roles` and, thus, it is a way to control the credentials of the client. In this setting, the cashier c_i of Example 1 is implemented by the following process (the remaining behaviour of the bank is implemented as in Example 1):

$$c_i(x).([x = \text{withdrw_req}] wdrw(y).P_{wdrw} \quad | \quad [x = \text{open_req}] opn(y).P_{opn} \quad | \\ [x = \text{creditcard_req}] cc(y).P_{cc} \quad | \quad [x = \text{stop}] free^s\langle c_i^s \rangle)$$

Let \mathcal{U} assign to channel `wdrw` (resp., `opn` and `cc`) the group `wdrw` (resp., `opn` and `cc`), and let \mathcal{P} be such that $cc^! \in \mathcal{P}(\text{rich_client})$, $wdrw^! \in \mathcal{P}(\text{client})$ and $opn^! \in \mathcal{P}(\text{user})$. Under this schema, the following client is not well-typed, as he has not activated the correct `role` to perform credit card requests:

$$r \parallel \text{role client}.signal^s\langle r \rangle.served(z).z\langle \text{creditcard_req} \rangle.cc^s\langle \text{signature} \rangle.z\langle \text{stop} \rangle \parallel_0$$

Indeed, the type checking fails when applying the rule (T-OUTPUT) to action $cc^s\langle \text{signature} \rangle$ because $cc^! \notin \mathcal{P}(\{\text{client}\})$. On the contrary, assuming that $\text{rich_client} \in \mathcal{U}(r_1)$, $\text{client} \in \mathcal{U}(r_2)$ and $\text{user} \in \mathcal{U}(r_3)$, and that $\{\text{cashier_req}^!, \text{cashier_get}^?, \text{cashier}^!\} \subseteq \mathcal{P}(\text{user}) \cap \mathcal{P}(\text{client}) \cap \mathcal{P}(\text{rich_client})$, the following clients are well-typed:

$$\begin{aligned} r_1 \parallel \text{role rich_client}.signal^s\langle r \rangle.served(z).z\langle \text{creditcard_req} \rangle.cc^s\langle \text{signature} \rangle.z\langle \text{stop} \rangle \parallel_0 \\ r_2 \parallel \text{role client}.signal^s\langle r \rangle.served(z).z\langle \text{withdrw_req} \rangle.wdrw^s\langle \text{sum} \rangle.z\langle \text{stop} \rangle \parallel_0 \\ r_3 \parallel \text{role user}.signal^s\langle r \rangle.served(z).z\langle \text{open_req} \rangle.opn^s\langle \text{personal_data} \rangle.z\langle \text{stop} \rangle \parallel_0 \end{aligned}$$

Indeed, actions $cc^s\langle \text{signature} \rangle$, $wdrw^s\langle \text{sum} \rangle$ and $opn^s\langle \text{personal_data} \rangle$ are all enabled by the `roles` previously activated by users, viz. `rich_client`, `client` and `user`, respectively. \diamond

(E-SESS) $\frac{\rho \notin \mathcal{U}(r)}{r \parallel P \parallel_\rho \rightsquigarrow_{(\mathcal{U}; \mathcal{P})}}$	(E-RÔLE) $\frac{R \notin \mathcal{U}(r)}{r \parallel \mathbf{role} R.P \parallel_\rho \rightsquigarrow_{(\mathcal{U}; \mathcal{P})}}$	(E-YIELD) $\frac{R \notin \rho}{r \parallel \mathbf{yield} R.P \parallel_\rho \rightsquigarrow_S}$
(E-IN) $\frac{\mathcal{U}(b^r) = S \quad S^? \notin \mathcal{P}(\rho)}{r \parallel b(x).P \parallel_\rho \rightsquigarrow_{(\mathcal{U}; \mathcal{P})}}$	(E-OUT) $\frac{\mathcal{U}(b^s) = S \quad S^! \notin \mathcal{P}(\rho)}{r \parallel b^s(n).P \parallel_\rho \rightsquigarrow_{(\mathcal{U}; \mathcal{P})}}$	
(E-PAR) $\frac{A \rightsquigarrow_S}{A \parallel B \rightsquigarrow_S}$	(E-RES) $\frac{A \rightsquigarrow_{(\mathcal{U} \uplus \{a^r; R\}; \mathcal{P})}}{(\nu a^r : R)A \rightsquigarrow_{(\mathcal{U}; \mathcal{P})}}$	(E-STRUCT) $\frac{A \equiv B \quad B \rightsquigarrow_S}{A \rightsquigarrow_S}$

Table 5: Run-time Errors

We now establish the soundness of our type system in the standard way, i.e. by proving subject reduction and type safety. The first result states that well-typedness is preserved along reductions; the second result ensures that only systems abiding by the RBAC schema are allowed (i.e., users only perform actions permitted by their duly activated rôles). The proofs are in Appendix A.1.

Theorem 2.1 (Subject Reduction). *If $\Gamma \vdash^S A$ and $A \mapsto A'$, then $\Gamma \vdash^S A'$.*

To state type safety, we first need to formally define what situations our type system wants to avoid. Thus, in Table 5, we introduce the notion of *run-time errors* and prove that they never arise in any well-typed system. Intuitively, run-time errors are generated in three possible ways: whenever a session is equipped with rôles not assigned to the user owning that session (see law (E-SESS)); whenever a rôle is activated (resp., deactivated) by a user (resp., by a session) not owning such a rôle (see laws (E-RÔLE) and (E-YIELD), resp.); whenever an input/output action is performed in a user session where no privilege enabling such an action is provided by the rôles active in that session (see laws (E-IN) and (E-OUT)).

Theorem 2.2 (Type Safety). *If A is well-typed in S , then $A \rightsquigarrow_S$ cannot hold.*

We conclude this section remarking that our type system is not powerful enough to type all legal systems. For example, the system $r \parallel a^r \langle r \rangle \parallel_\rho$ is untypeable. Similarly, we have no notion of subtyping; thus, a channel must always carry values exactly of the same type. We now sketch how these deficiencies could be remedied, by following standard techniques; full details are omitted, as they are completely well understood and orthogonal w.r.t. the new ideas of our work.

Recursive Types In order to type $r \parallel a^r \langle r \rangle \parallel_\rho$, we would need a typing environment Γ assigning to r a type T such that $T = \rho' [a^r : R(T)]$, for some $\rho' \supseteq \rho$ and R . Clearly, such a type T is not expressible with our type syntax, as it would require an infinite nesting of type constructors, i.e. $\rho' [a^r : R(\rho' [a^r : R(\rho' [a^r : R(\dots)]))]$. By following [8, 20, 24], this problem can be solved by using equations between type expressions whose solutions are infinite types, like T above. To this aim, we assume a set of type variables Ξ , ranged over by ξ , and extend the syntax of value types as follows:

$$T ::= \rho[\bar{a} : \bar{C}] \mid C \mid \xi \mid \mu\xi.T$$

Intuitively, $\mu\xi.T$ stands for the solution of the (recursive) equation $\xi = T$. However, to avoid nonsensical expressions like $\mu\xi.\xi$, we impose the constraint that in $\mu\xi.T$ the variable ξ occurs *guarded* in T ,

i.e. it occurs underneath at least one of the other type constructors. Moreover, we only consider *closed* type expressions, i.e. expressions in which each occurrence of a type variable ξ is underneath a $\mu\xi.$ -construct. By exploiting these two assumptions, it can be standardly proved that the set of types is a c.p.o. and, thus, the solution of an equation $\mu\xi.T$ is unique and can be obtained with a least fix-point construction.

If we now represent (the unfolding of) a (possibly recursive) type as its (possibly infinite) parse tree, we can deem two types *equivalent*, written $T_1 \asymp T_2$, if and only if they represent the same tree. A lot of literature on type systems is devoted to compute \asymp algorithmically (see, e.g., [2, 20]); for our purposes, it suffices to remember that \asymp is a congruence on types such that $\mu\xi.T \asymp T[\mu\xi.T/\xi]$.

If we now add the rule

$$(T\text{-ID-REC}) \frac{\Gamma \vdash n : T \quad T \asymp T'}{\Gamma \vdash n : T'}$$

to the type system of Table 4 and modify rule (T-ID₂) to be

$$\frac{\Gamma(r) \asymp \rho[\tilde{b} : \tilde{C}, a : C, \tilde{b}' : \tilde{C}']}{\Gamma \vdash a' : C}$$

it is easy to see that $\Gamma \vdash^{(U:\mathcal{P})} r \Vdash a' \langle r \rangle \Vdash_{\rho}$, whenever $\Gamma(r) \triangleq \mu\xi.\rho[a' : R(\xi)]$ and $R^! \in \mathcal{P}(\rho)$.

Subtyping Subtyping is a preorder on types that can be thought of as inclusion between the set of the values of the types. If T' is subtype of T , then a value of type T' is also of type T ; thus, any expression of type T' can always replace an expression of type T , without compromising well-typedness. In traditional programming languages, this feature is used to reduce the size of a program, as the same function can be invoked on parameters of different types, without writing a ‘copy’ of the same function for each subtype. For example, assume, as usual, that `int` is a subtype of `real` and that there is a function to multiply two reals. Then, the same function can also be used to multiply two integers.

In our setting, we can define an ordering on types to allow the passage of values of different (yet somehow related) types along the same channel. We sketch here a very basic form of subtyping inspired by [14]; for more elaborated settings see, e.g., [21, 24]. First, we need to define the *subtyping relation*, \sqsubseteq , that is the least preorder on types satisfying the following rules:

$$\frac{\rho \subseteq \rho' \quad h \leq k \quad \forall i = 1, \dots, h. C_i \sqsubseteq C'_i}{\rho[a_1 : C_1, \dots, a_h : C_h] \sqsubseteq \rho'[a_1 : C'_1, \dots, a_k : C'_k]} \quad \frac{R = R' \quad T \sqsubseteq T'}{R(T) \sqsubseteq R'(T')}$$

Then, we need to update rule (T-OUTPUT) to become

$$\frac{\Gamma \vdash m : R(T) \quad \Gamma \vdash n : T' \quad T \sqsubseteq T' \quad R^! \in \mathcal{P}(\rho) \quad \Gamma; \rho \vdash_r^{(U:\mathcal{P})} P}{\Gamma; \rho \vdash_r^{(U:\mathcal{P})} m \langle n \rangle . P}$$

In this framework, we can type the system

$$r \Vdash a(x).b^x \langle n \rangle \Vdash_{\rho} \parallel s \Vdash a' \langle r_1 \rangle \Vdash_{\rho_1} \parallel t \Vdash a' \langle r_2 \rangle \Vdash_{\rho_2} \parallel r_1 \Vdash b(y).c(z) \Vdash_{\rho'} \parallel r_2 \Vdash b(y).c'(z) \Vdash_{\rho'}$$

Indeed, it suffices to take Γ such that $\Gamma(n) = T$, $\Gamma(r_1) = \rho'[b : R(T), c : C]$, $\Gamma(r_2) = \rho'[b : R(T), c' : C']$ and $\Gamma(r) = \rho[a : S(\rho'[b : R(T)])]$, for some R and S .

3 Behavioural Semantics

One of the main advantages of process calculi is the possibility of developing over them behavioural equivalences, that abstract a term from its syntax and concentrate on its functionalities. To this aim, we consider a standardly defined typed behavioural congruence, viz. *reduction barbed congruence* [15]. This is a touchstone equivalence defined in terms of the reduction relation and of a notion of observation, and then closed under all possible system contexts. The reason to consider a typed congruence is that only well-typed contexts guarantee a reduction behaviour abiding by the RBAC policy.

In its typed version, barbed congruence is tagged with an environment Γ and RBAC schema \mathcal{S} , to signify that it equates terms that are typeable under Γ and \mathcal{S} . Moreover, only contexts typeable under Γ and \mathcal{S} are considered in the definition of the congruence. Thus, following the style of [13], we write $\Gamma \models^{\mathcal{S}} A_1 \cong A_2$ to mean that $\Gamma \vdash^{\mathcal{S}} A_i$, for $i = 1, 2$, and that A_1 and A_2 exhibit the same behaviour in all environments ‘compatible’ with Γ and \mathcal{S} .

We now formally define barbed congruence. As usual, we denote with \mapsto^* the reflexive and transitive closure of the reduction relation \mapsto .

Definition 3.1 (Barbs). The *observation predicate* $A \Downarrow \eta$ holds if

- either $\eta = a^r$ and $A \equiv (\nu \tilde{b}^s : \tilde{R})(A' \parallel r \parallel a(x).P \parallel_{\rho})$ for $a^r \notin \tilde{b}^s$,
- or $\eta = \bar{a}^r$ and $A \equiv (\nu \tilde{b}^s : \tilde{R})(A' \parallel s' \parallel a^r \langle n \rangle . P \parallel_{\rho})$ for $a^r \notin \tilde{b}^s$.

The predicate $A \Downarrow \eta$ holds if there exists A' such that $A \mapsto^* A'$ and $A' \Downarrow \eta$.

We remark that the chosen barbs only express the ability to interact over channels. Indeed, observing role activations/deactivations is not reasonable, as no context can determine whether a user performs a **role/yield**: these operations only affect the thread performing them.

Definition 3.2 (Reduction Barbed Congruence). *Reduction barbed congruence* is the largest binary and symmetric typed relation over systems such that, whenever $\Gamma \models^{\mathcal{S}} A_1 \cong A_2$, it holds that

1. (*Barb preservation*) if $A_1 \Downarrow \eta$, then $A_2 \Downarrow \eta$
2. (*Reduction closure*) if $A_1 \mapsto A'_1$, then there exists a system A'_2 such that $A_2 \mapsto^* A'_2$ and $\Gamma \models^{\mathcal{S}} A'_1 \cong A'_2$
3. (*Contextuality*) let \mathcal{S} be $(\mathcal{U}; \mathcal{P})$; then,
 - (a) for all \mathcal{P}' and $\tilde{n} : \tilde{T}$ such that $\tilde{n} \cap \text{dom}(\mathcal{U}) = \emptyset$, it holds that $\Gamma, \tilde{n} : \tilde{T} \models^{(\mathcal{U} \cup \tilde{n} : \tilde{T}; \mathcal{P} \cup \mathcal{P}')} A_1 \cong A_2$
 - (b) for all systems B such that $\Gamma \vdash^{\mathcal{S}} B$, it holds that $\Gamma \models^{\mathcal{S}} A_1 \parallel B \cong A_2 \parallel B$
 - (c) for all $a^r : R(T)$ such that $\Gamma = \Gamma'$, $a^r : R(T)$ and $\mathcal{U} = \mathcal{U}' \uplus \{a^r : R\}$, it holds that $\Gamma' \models^{(\mathcal{U}'; \mathcal{P})} (\nu a^r : R) A_1 \cong (\nu a^r : R) A_2$.

The less intuitive condition of the above definition is contextuality. Essentially, it requires that the equated systems A_1 and A_2 be equivalent in any execution context. An execution context can affect the behaviour of such systems in three ways: it can extend the RBAC schema, thus enabling more functionalities of A_1 and A_2 ; it can provide more parallel components that, by interacting with A_1 and A_2 , could change their behaviours; it can hide channels and, hence, delete observable behaviours of A_1 and A_2 . These aspects are handled by the sub-conditions (a), (b) and (c), respectively.

The problem with this definition of barbed congruence is that it must be proved by analysing all the system contexts, which makes it hardly tractable. In Appendix B, we provide a more tractable proof-technique for \cong . Since this task requires several technicalities taken from the field of process calculi,

we leave it to the interested reader; the other readers should only know that tools for establishing barbed congruence in a simpler way do exist.

To conclude, we now list a few algebraic laws that illustrate some key features of our framework. In what follows, we fix a RBAC schema \mathcal{S} and a suitable typing environment Γ . The first equation states that a terminated session of a user does not affect the evolution of a system. Indeed, it holds that

$$\Gamma \models^{\mathcal{S}} r \llbracket \mathbf{nil} \rrbracket_{\rho} \cong \mathbf{0}.$$

This is different from some distributed calculi, like e.g. the Ambient calculus [7] or KLAIM [9], where the presence of a place for computations is relevant. Moreover, differently from several distributed languages, the user performing an *output* action is irrelevant; the only relevant aspect is the set of permissions activated when performing the action. This is summarised in the following law:

$$\Gamma \models^{\mathcal{S}} r \llbracket b^s \langle n \rangle \rrbracket_{\rho} \cong t \llbracket b^s \langle n \rangle \rrbracket_{\rho}.$$

A similar law holds for the **yield** action. Notice that only for these two actions the identity of the user performing them is irrelevant. For example, relocating an input action breaks the equivalence between processes, as input channels implicitly refer the user owning them. Indeed, we have that

$$\Gamma \models^{\mathcal{S}} r \llbracket a(x).P \rrbracket_{\rho} \not\cong t \llbracket a(x).P \rrbracket_{\rho}$$

Similarly, it is possible to move a **role** R prefix between two users only when R is assigned to both of them.

By exploiting these observations, we develop in the following example a relocation procedure to establish whether a process can be moved from a user to another. This procedure can be exploited to reduce the number of users in a system, while maintaining the overall system behaviour.

Example 4. We now give simple procedure to infer judgements of the form $\Gamma \models^{\mathcal{S}} r \llbracket P \rrbracket_{\rho} \cong s \llbracket P \rrbracket_{\rho}$. This judgement says that process P can be indifferently executed by r and s without altering its observable functionalities. Thus, session $r \llbracket P \rrbracket_{\rho}$ can be replaced by $s \llbracket P \rrbracket_{\rho}$. If no other session of r is left in the system, then r itself has been removed.

The procedure is very simple. Try to infer both $\Gamma; \rho \vdash_r^{\mathcal{S}} P$ and $\Gamma; \rho \vdash_s^{\mathcal{S}} P$ without using rules (T-INPUT) and (T-RES). If you succeed, then $\Gamma \models^{\mathcal{S}} r \llbracket P \rrbracket_{\rho} \cong s \llbracket P \rrbracket_{\rho}$, otherwise $\Gamma \models^{\mathcal{S}} r \llbracket P \rrbracket_{\rho} \not\cong s \llbracket P \rrbracket_{\rho}$. We cannot use rule (T-INPUT) because, as we have already discussed, we cannot relocate processes containing input prefixes. A similar problem arises also for restriction (thus, we cannot use rule (T-RES)). Indeed, the interplay between user names, restricted channel names and restricted channels is subtle. For example, consider the system $P \triangleq (va : R)a^r \langle a^s \rangle$, try to run it in users r , s and t , and put the resulting session in an arbitrary system context. In the first case, P cannot be engaged in any reduction, as it emits a value on a channel known only by P itself; in the second case, P sends a bound value; in the third case, P sends a free value. Thus, relocating processes with restrictions breaks equivalences, in general.

It can now be easily proved that this procedure is a sound and complete proof-technique for judgement $\Gamma \models^{\mathcal{S}} r \llbracket P \rrbracket_{\rho} \cong s \llbracket P \rrbracket_{\rho}$, whenever P does not contain restrictions and input prefixes. \diamond

4 Adding Rôle Activations and Deactivations

In this section, we show how our framework can be adapted to encompass more advanced features. Usually, the task of properly putting **role/yield** operations within a system is tedious and error-prone;

moreover, it assumes a full knowledge of the RBAC schema at programming time. We now describe a way to add r ole activations/deactivations within a system in such a way that the resulting system can be executed under any given schema $(\mathcal{U}; \mathcal{P})$, whenever possible, i.e., when users are allowed to activate the r oles required by the actions they are willing to perform.

A first technique rewrites a system A without actions **role/yield** by activating at the beginning of each session of a generic user r all the r oles in $\mathcal{U}(r)$. Intuitively, the refined system contains all the legal behaviours of A with respect to the RBAC schema given. However, the fact that all the r oles assigned to a user are always activated violates a basilar RBAC design principle: a r ole should be active only when needed.

A second na ive algorithm replaces each input/output prefix $\alpha.Q$ occurring in each session of a generic user r with **role** $R.\alpha$.**yield** $R.Q$, where r ole R belongs to $\mathcal{U}(r)$ and enables action α . The algorithm is very simple but it presents several drawbacks: it always adds a pair of auxiliary actions **role/yield** for every prefix α occurring in the process, although it could be that r ole R enables also the following prefixes in Q . Furthermore, when rewriting $\alpha.Q$, it is possible that several r oles enable α : in this case, a thorough choice of which r ole to activate may minimise the number of **role/yield** actions.

We now present an algorithm that adds a smaller number of actions **role/yield**. The algorithm works on a tree-representation of the process in each session of a generic user r : first, the tree is partitioned into subtrees by collecting together nodes (i.e., actions) which require the same r ole in order to be executed correctly; then a **role** R auxiliary action is added before the root of each subtree requiring r ole R , and dually a **yield** R action is added after the leaf of each such subtree.

More specifically, a process P running inside a session of user r is translated into an annotated binary tree, where nodes represent process operators and each node is annotated with the set of r oles whose permissions enable the action associated with the node (if the node is not associated with an input/output action, any r ole available for user r will enable it). The tree is expressed in terms of a tuple $t = (V, E, rt, \phi)$, where V is a finite set of *nodes*, $E \subseteq V \times V$ is the set of the *edges* (i.e., $(v, v') \in E$ iff there is an edge from v to v'), $rt \in V$ is the *root* of the tree, and $\phi : V \rightarrow 2^{\mathcal{R}}$ is the assignment of r oles to nodes used to annotate each node. The annotated tree associated with a finite process P (without **role/yield** actions) running inside a session of user r may be generated by the function $\text{TREE}_{\Gamma, \varphi}^r(P)$ described below, with $(\mathcal{U}; \mathcal{P})$ a RBAC schema and Γ a typing environment respecting \mathcal{U} . In the definition, we write $\Gamma\{r\}$ for the set ρ such that $\Gamma \vdash r : \rho[\tilde{a} : \tilde{C}]$; thus, $R \in \Gamma\{r\}$ is a shortcut for

$$\Gamma \vdash r : \rho[\widetilde{a} : \widetilde{C}] \wedge R \in \rho.$$

$$\begin{aligned} \text{TREE}_{\Gamma, \mathcal{P}}^r(\mathbf{nil}) &\triangleq (\{v\}, \emptyset, v, \{v \mapsto \Gamma\{r\}\}) \\ \text{TREE}_{\Gamma, \mathcal{P}}^r(a(x).P) &\triangleq (V \cup \{v\}, E \cup \{(v, rt)\}, v, \phi \cup \{v \mapsto \rho\}) \\ &\quad \text{where } \text{TREE}_{\Gamma, x:T; \mathcal{P}}^r(P) = (V, E, rt, \phi), v \notin V, \Gamma \vdash a^r : S(T) \\ &\quad \text{and } \rho = \{R \in \Gamma\{r\} : S^? \in \mathcal{P}(R)\} \\ \text{TREE}_{\Gamma, \mathcal{P}}^r(m\langle n \rangle.P) &\triangleq (V \cup \{v\}, E \cup \{(v, rt)\}, v, \phi \cup \{v \mapsto \rho\}) \\ &\quad \text{where } \text{TREE}_{\Gamma, \mathcal{P}}^r(P) = (V, E, rt, \phi), v \notin V, \Gamma \vdash a^r : S(T), \\ &\quad \Gamma \vdash n : T \text{ and } \rho = \{R \in \Gamma\{r\} : S^! \in \mathcal{P}(R)\} \\ \text{TREE}_{\Gamma, \mathcal{P}}^r(P_1 \mid P_2) &\triangleq (V_1 \cup V_2 \cup \{v\}, E_1 \cup E_2 \cup \{(v, rt_1), (v, rt_2)\}, v, \\ &\quad \phi_1 \cup \phi_2 \cup \{v \mapsto \Gamma\{r\}\}) \\ &\quad \text{where } \text{TREE}_{\Gamma, \mathcal{P}}^r(P_i) = (V_i, E_i, rt_i, \phi_i) \text{ for } i \in \{1, 2\}, \\ &\quad V_1 \cap V_2 = \emptyset \text{ and } v \notin V_1 \cup V_2 \\ \text{TREE}_{\Gamma, \mathcal{P}}^r((va : R)P) &\triangleq (V \cup \{v\}, E \cup \{(v, rt)\}, v, \phi \cup \{v \mapsto \Gamma\{r\}\}) \\ &\quad \text{where } \text{TREE}_{\Gamma, a^r:R(T); \mathcal{P}}^r(P) = (V, E, rt, \phi), v \notin V \\ &\quad \text{and } \Gamma, a^r : R(T); \Gamma\{r\} \vdash_r^{(\mathcal{U}; \mathcal{P})} P \\ \text{TREE}_{\Gamma, \mathcal{P}}^r([m = n]P) &\triangleq (V \cup \{v\}, E \cup \{(v, rt)\}, v, \phi \cup \{v \mapsto \Gamma\{r\}\}) \\ &\quad \text{where } \text{TREE}_{\Gamma, \mathcal{P}}^r(P) = (V, E, rt, \phi) \text{ and } v \notin V \end{aligned}$$

Example 5. Let us consider a system consisting of a single user r running the following process

$$P = a(x).([x = b^r]a^r\langle x \rangle \mid [x = s](vc : S)(a^r\langle x \rangle \mid a^s\langle c^r \rangle))$$

and the RBAC schema $(\mathcal{U}; \mathcal{P})$ defined as

$$\begin{array}{ll} \mathcal{U} : r &\mapsto \{R_1, R_2\} \\ &a^r \mapsto R \\ &a^s \mapsto S \end{array} \quad \begin{array}{ll} \mathcal{P} : R_1 &\mapsto \{R^!, R^?\} \\ &R_2 \mapsto \{S^!\} \end{array}$$

A pictorial representation of $\text{TREE}_{\mathcal{U}, \mathcal{P}}^r(P)$ is given in Figure 1, where each node is explicitly named (the name is shown on the left-hand side of the node), the process operator associated with the node is written within the node, and the annotation (i.e., the set of rˆoles associated with the node) is depicted on the right-hand side of the node. Notice that in this case it suffices to parameterise $\text{TREE}_{\mathcal{U}, \mathcal{P}}^r$ only with a rˆoles-to-users assignment \mathcal{U} instead of a (more complex) typing environment Γ (respecting \mathcal{U}) because no received value is used as an output channel. \diamond

Once process P has been translated into an annotated binary tree, the problem of finding a minimal refinement of P (in terms of the number of actions **role/yield** added) for user r under the schema $(\mathcal{U}; \mathcal{P})$ can be reformulated as the problem of finding a partition of nodes such that:

- the partition generates the minimum number of blocks;
- each block is a subtree²;
- all the nodes v belonging to the same block have in common one of their annotating rˆoles, i.e.,

²Here, we use the term *subtree* to refer any connected subgraph of the given tree.

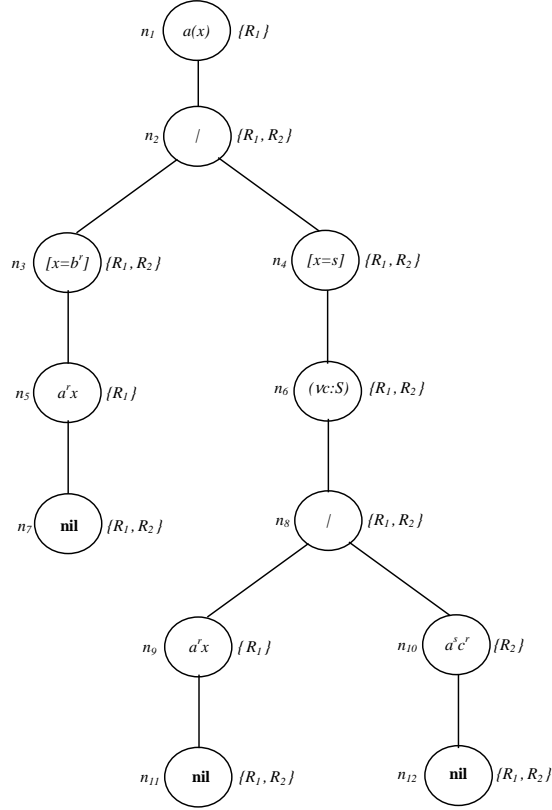


Figure 1: The annotated binary tree for process P from Example 5

We call such a problem the *minimal partition problem*.

We now describe a way to find a minimal partition of t and assign to each node v a label taken from $\phi(v)$. To this aim, for every node v and role R , the number $m[v, R]$ denotes the minimum number of blocks that can be obtained in the subtree rooted in v when v is labelled with R ; we let $m[v, R] = \infty$ if $R \notin \phi(v)$. An algorithm for calculating the quantity $m[v, R]$ is given in Table 6. Intuitively, we work in a bottom-up fashion on the tree. When we consider a leaf v and a role R that authorises it, we can use R to generate a singleton tree (hence, $m[v, R] = 1$); if, on the other hand, R cannot authorise the (action associated with) the node, then it can be ignored because it cannot induce any block in the partition (hence, $m[v, R] = \infty$). When we consider an internal node v with just one child v' and a role R suitable for v , we can either try to include v in the subtree of v' induced by R (hence, $m[v, R] = m[v', R]$), or we can put it in a new subtree (that can possibly grow up when analysing v 's ancestors). In the latter case, the new subtree is induced by R , while the subtree for v' can be induced by any other role S ; thus, $m[v, R] = 1 + \min_{S \neq R} \{m[v', S]\}$. Finally, the case for a node v with two children is similar, but it requires to examine four possible situations (according to whether v is included in both, in none, or in just one of the subtrees induced by R for v_1 and v_2).

Now, we can compute, for every node v , a role $\text{LABEL}(v)$ which represents the role common to all the nodes in the block which v belongs to. To this aim, we assume a standard function $\text{FATHER}(v)$ returning the father of node v in t (if any).

<p>Visit t in postorder</p> <p>When visiting the node v do</p> <ul style="list-style-type: none"> if v is a leaf then $m[v, R] := \begin{cases} \infty & \text{if } R \notin \phi(v) \\ 1 & \text{otherwise} \end{cases}$ if v has only one child (and let it be v') then $m[v, R] := \begin{cases} \infty & \text{if } R \notin \phi(v) \\ \min\{ m[v', R], \\ \min_{S \neq R} \{ m[v', S] \} + 1 \} & \text{otherwise} \end{cases}$ if v has children v_1 and v_2 then $m[v, R] := \begin{cases} \infty & \text{if } R \notin \phi(v) \\ \min\{ m[v_1, R] + m[v_2, R] - 1, \\ \min_{S \neq R} \{ m[v_1, S] \} + \min_{S \neq R} \{ m[v_2, S] \} + 1 \\ m[v_1, R] + \min_{S \neq R} \{ m[v_2, S] \}, \\ \min_{S \neq R} \{ m[v_1, S] \} + m[v_2, R] \} & \text{otherwise} \end{cases}$
--

Table 6: Computing function $m[v, R]$

Visit t in preorder

When visiting the node v do

$$m_v := \{ R : m[v, R] = \min_{S \in \phi(v)} \{ m[v, S] \} \}$$

if $v = rt$ or LABEL(FATHER(v)) $\notin m_v$
then LABEL(v) := R , where $R \in m_v$
else LABEL(v) := LABEL(FATHER(v))

Notice that the choice of $R \in m_v$ (in the ‘then’ branch) is totally arbitrary: any such R can be chosen, since m_v only contains rˆoles that minimise $m[v, \cdot]$. We can now formulate the soundness of the algorithm presented so far; a sketch of the proof is in Appendix A.2.

Proposition 4.1. *Function LABEL can be used to induce a partition of t ’s nodes in subtrees satisfying the requirements of the minimal partition. Moreover, the overall procedure takes $O(|V| \times K^2)$, where K is the size of the largest set annotating a node of the tree.*

A solution of the original problem of properly putting actions **role/yield** in a process P can be then extracted easily from $t = \text{TREE}_{\Gamma, \mathcal{P}}^r(P)$ and from the associated function LABEL. Each block of the partition induced by LABEL represents the set of process operators in P that are under the influence of the rˆole labelling the block. If the tree consists of a single node, then P must be **nil** and no auxiliary action is needed. Otherwise, the annotated (and partitioned) tree can be visited in preorder: depending on the value of LABEL(v) a pair of **role/yield** auxiliary actions are either added or not. In particular, the operator corresponding to the root is prefixed in P with **role** R , where $R = \text{LABEL}(rt)$. Then, no other actions are added until LABEL(v) \neq LABEL(FATHER(v)). In this case, the operator associated with node v is prefixed with **yield** R .**role** S , where $R = \text{LABEL}(\text{FATHER}(v))$ and $S = \text{LABEL}(v)$. We denote the process resulting from this procedure as $\text{REFINED}_{\Gamma, \mathcal{P}}^r(P)$.

Example 6. Consider again Example 5 and the process P running inside user r :

$$P = a(x).([x = b^r]a^r\langle x \rangle \mid [x = s](vc : S)(a^r\langle x \rangle \mid a^s\langle c^r \rangle))$$

In this case, we have that function $m[v, R]$ is

	n_1	n_2	n_3	n_4	n_5	n_6	n_7	n_8	n_9	n_{10}	n_{11}	n_{12}
R_1	2	2	1	2	1	2	1	2	1	∞	1	1
R_2	∞	3	2	2	∞	2	1	2	∞	1	1	1

Thus, a minimal partition of $\text{TREE}_{\mathcal{U};\mathcal{P}}^r(P)$ is given by the following two blocks

$$b_1 = \{n_1, n_2, n_3, n_4, n_5, n_6, n_7, n_8, n_9, n_{11}\}, \quad b_2 = \{n_{10}, n_{12}\},$$

from which we can extract the refined process

$$\text{REFINED}_{\Gamma;\mathcal{P}}^r(P) = \mathbf{role} R_1.a(x).([x = b^r]a^r\langle x \rangle \mid [x = s](vc : S)(a^r\langle x \rangle \mid \mathbf{yield} R_1.\mathbf{role} R_2.a^s\langle c^r \rangle))$$

◇

The correctness of the approach can be stated as follows; some details on the proof are in Appendix A.2. Recall from the definition of the type system that r can run P if $\Gamma; \mathcal{U}(r) \vdash_r^{(\mathcal{U};\mathcal{P})} P$.

Proposition 4.2. *Let $S = (\mathcal{U}; \mathcal{P})$ be a RBAC schema and P be a finite process without **role/yield** to be run by user r . Then*

1. $\Gamma; \mathcal{U}(r) \vdash_r^S P$ implies that $\text{REFINED}_{\Gamma;\mathcal{P}}^r(P)$ is defined;
2. symmetrically, Γ respects \mathcal{U} and $\text{REFINED}_{\Gamma;\mathcal{P}}^r(P) = P'$ imply that $\Gamma; \emptyset \vdash_r^S P'$.

Intuitively, the first implication ensures that every process P that can be run by r under the schema S can be properly annotated with actions **role** and **yield**. The second implication states that the result of the annotation procedure we have just presented is a well-typed process for r in the schema given. Notice that $\text{REFINED}_{\Gamma;\mathcal{P}}^r(P)$ is the minimal typeable process obtained from P by adding actions **role/yield**: this is an easy corollary of Propositions 4.1 and 4.2(2).

Finally, by Proposition 4.1, the overall procedure is linear in the size of P (i.e., in the number of its operators). Indeed, $|V|$ is proportional to the size of P and $\mathcal{U}(r)$ is an upper bound to the sets annotating the nodes of the tree (usually, $|\mathcal{U}(r)|$ is a small constant). This is the best asymptotic performance we could aim at, since we at least have to parse all P to properly add actions **role/yield**.

Least Privilege. Example 5 can be easily adapted to enforce the *least privilege* property [22, 26]. This is a well-known property requiring that every program and every user of the system operate using the least set of privileges necessary to complete their job.³ Primarily, this principle limits the damage that can result from an accident or error. It also reduces the number of potential interactions among privileged programs to the minimum for correct operation, so that unintentional, unwanted, or improper uses of privileges are less likely to occur.

³To be precise, one should use the term *minimal privilege* instead of *least privilege*. Indeed, imagine that you have only two roles, R_1 and R_2 , such that R_1 enables actions $\{R^2, S^1\}$ and R_2 enables actions $\{R^2, S^2\}$; then, it is unclear which one would be the role giving the least privilege to execute an input from a channel of role R . Nevertheless, the current terminology in computer security uses the word “least” instead of “minimal”; we adhere to this trend.

In our setting, we can say that a user r satisfies this property w.r.t. a schema $(\mathcal{U}; \mathcal{P})$ while running in A if, whenever $A \mapsto^* (v\tilde{a}^r : \tilde{R})(A' \parallel r \parallel \alpha.P \parallel_\rho)$, it holds that ρ is a minimal (w.r.t. $|\mathcal{P}(\cdot)|$, i.e. the cardinality of the set of privileges associated to \cdot) set of rôles assignable to r that enables α . The approach presented in this section can be easily adapted to encompass the least privilege. The only thing we need to modify in the algorithm given above is the definition of function ϕ when building the tree for a process prefixed by action α . Let v be the node associated to α . By letting $enable(v) = \{R \in \Gamma\{r\} : \mathcal{P}(R) \text{ enable action } \alpha\}$, we let $\phi(v) = \{R \in enable(v) : |\mathcal{P}(R)| = \min_{S \in enable(v)} \{|\mathcal{P}(S)|\}\}$.

Example 7. Consider a user r that connects to a mail server to read his e-mail and changes his password before quitting. Suppose that $(\mathcal{U}; \mathcal{P})$ is such that $\{\text{user}, \text{admin}\} \subseteq \mathcal{U}(r)$, $\mathcal{U}(\text{login}^{e\text{-server}}) = \text{login}$, $\mathcal{U}(\text{read_mail}^r) = \text{read_mail}$, $\mathcal{U}(\text{change_pwd}^{e\text{-server}}) = \text{change_pwd}$, $\{\text{login}^1, \text{read_mail}^2\} \subseteq \mathcal{P}(\text{user})$ and $\mathcal{P}(\text{admin}) = \mathcal{P}(\text{user}) \cup \{\text{change_pwd}^1\}$. Remarkably, the rôle admin gives r the permission to change his password. The following two systems

$$\begin{aligned} A_1 &= r \parallel \text{login}^{e\text{-server}} \langle \text{pwd} \rangle . \text{read_mail}(x) . \text{change_pwd}^{e\text{-server}} \langle \text{pwd}' \rangle \parallel_{\{\text{admin}\}} \\ A_2 &= r \parallel \text{login}^{e\text{-server}} \langle \text{pwd} \rangle . \text{read_mail}(x) . \mathbf{role} \text{ admin} . \\ &\quad \text{change_pwd}^{e\text{-server}} \langle \text{pwd}' \rangle . \mathbf{yield} \text{ admin} \parallel_{\{\text{user}\}} \end{aligned}$$

are both well-typed in $(\mathcal{U}, \mathcal{P})$. However, they differ in the auxiliary actions used: system A_2 satisfies the least-privilege requirement, since at each execution step user r owns a minimal set of permissions required to execute the action, while A_1 does not, since it activated the rôle admin also to login and read mails. \diamond

5 Possible Extensions of the Core RBAC Model

There is a wide spectrum of RBAC models differing on the operations supported. For example, [10, 23] propose various extensions of the core RBAC96 model we have used up to now, that in *loc. cit.* is referred to as $RBAC_0$. In particular, two variants are proposed: $RBAC_1$, adding rôle hierarchies, and $RBAC_2$, introducing constraints to permissions a user can exploit. In this section, we describe how these extensions can be easily expressed also in our framework.

5.1 Hierarchical RBAC

Hierarchies are a natural means for structuring rôles to reflect the organisational structure of an enterprise. A hierarchy is a partial order defining a seniority relation between rôles, whereby senior rôles acquire the permissions of their juniors, and junior rôles acquire the users authorised for their senior rôles. For example, in a health-care scenario, a rôle cardiologist is hierarchically superior to the rôle doctor , thus the cardiologist should have all the permissions of the doctor as well, and all the users that are authorised for the cardiologist rôle should be authorised also for the doctor rôle. This approach can increase the administrative efficiency of the enterprise: rather than specifying all the permissions of the junior rôle for the senior rôle, the junior rôle is specified as a permission of the senior rôle.

Our framework can be easily extended to express rôle hierarchies by adding a third component \leq , a partial order on \mathcal{R} , to the RBAC schema which becomes a triplet $(\mathcal{U}; \mathcal{P}; \leq)$. More specifically, when $R \leq S$, R is a *junior rôle* of S or, similarly, S is a *senior rôle* of R . Once a hierarchical RBAC schema

$(\mathcal{U}; \mathcal{P}; \leq)$ has been fixed, we can define the set of junior roles with respect to a given role R , or to a given set of roles ρ , as $jnr(R) \triangleq \{S : S \leq R\}$ and $jnr(\rho) \triangleq \bigcup_{R \in \rho} jnr(R)$. Then, we may re-define $\mathcal{P}(\rho)$ as $\mathcal{P}(\rho) \triangleq \bigcup_{R \in jnr(\rho)} \mathcal{P}(R)$ and adapt both the type system and the barbed congruence take into account the hierarchy relation. In particular, in Definition 3.3.3(a) we can also extend the partial order \leq with all \leq' such that $\leq \cup \leq'$ is still a partial order. These modifications suffice to let the theory presented in Sections 2 and 3 properly work.

Example 8. Consider the health-care scenario again. In a hospital there is often a strict hierarchy establishing which operations are permitted depending on the position of the different employees. For example, the role `specialist` usually contains the roles `doctor` and `intern`. This means that users activating role `specialist` are implicitly associated also with the permissions associated with the `doctor` and `intern` roles, without the administrator having to explicitly list the `doctor` and `intern` permissions. This is an example of *multiple inheritance*, which provides us with the ability to inherit permissions from two or more role sources. Indeed, a role is composed from multiple subordinate roles with fewer permissions as in the organisation and business structure which these roles are intended to represent.

This hierarchy can be expressed by having `intern` \leq `specialist` and `doctor` \leq `specialist`. Moreover, the roles `cardiologist` and `radiologist` could each contain the `specialist` role. In this case, we also let `specialist` \leq `cardiologist` and `specialist` \leq `radiologist`, leading to, e.g., $jnr(\text{radiologist}) = \{\text{radiologist}, \text{specialist}, \text{doctor}, \text{intern}\}$. Now let

$$\begin{aligned} \mathcal{U}(\text{prescr_aspirin}^{\text{patient}}) &= \text{prescr_aspirin} \\ \mathcal{U}(\text{use_XRays}^{\text{hospital}}) &= \text{use_XRays} \\ \text{prescribe_aspirin}^! &\in \mathcal{P}(\text{doctor}) \\ \text{use_XRays}^! &\in \mathcal{P}(\text{radiologist}) \end{aligned}$$

then the user

$$r \llbracket \text{role radiologist. use_XRays}^{\text{hospital}} \langle \text{patient} \rangle. \text{prescr_aspirin}^{\text{patient}} \langle \text{posology} \rangle \rrbracket_0$$

is typeable by only assuming that $\mathcal{U}(r) = \{\text{radiologist}\}$. ◇

5.2 Constrained RBAC

The core RBAC model can be further extended by requiring different kinds of constraints to be satisfied before allowing a user to activate a role, or when defining the RBAC schema. According to [11], there are two possible forms of constraints: *static* and *dynamic*. The first ones deal with the permissions-to-roles and with the roles-to-users assignments. For example, it might be required that a user cannot be assigned some specified roles at the same time, or that the same permission is not assigned to different roles. These constraints are usually enforced during the definition of the RBAC schema [1, 16, 17, 25]. On the contrary, dynamic constraints deal with user sessions. By exploiting this form of constraints, it is possible, e.g., to assign to the same user two conflicting roles, although requiring that these roles are never activated simultaneously (for most practical purposes, this kind of requirement suffices).

We can easily extend our framework to deal in a uniform way also with different dynamic constraints. In this case, another component C is added to the RBAC schema. C is a finite set of binary predicates (that in this paper we assume to be first-order logic formulae built up over the atoms \mathcal{R}) relating a role and a set of roles. Given a constrained RBAC schema $(\mathcal{U}; \mathcal{P}; C)$, we let $C(R, \rho)$ be

$\bigwedge_{\text{constr} \in C} \text{constr}(R, \rho)$. As for the hierarchy extension, both the type system and the barbed congruence must be parameterised also with respect to C . Moreover, in rule (T-R^{OLE}), the premise $C(R, \rho)$ must be added. With respect to Definition 3.3.3(a), we want to remark that extending C usually reduces the set of possible evolutions. Thus, instead of requiring that the equivalence holds in every extended schema, we now require that the equivalence is preserved when making the schema more liberal, i.e. when *reducing* the set of constraints.

By adapting the logic formulae in C to different situations, we are able to express the most typical examples of constraints:

1. *Mutual exclusion*: the same user can activate simultaneously at most one rôle in a mutually exclusive set. For example, if rôles R and S are mutually exclusive, R cannot be activated in a user session where rôle S is already active, and vice-versa. This can be formalised as

$$\text{constr}_{R \oplus S}(R', \rho) \triangleq ((R' = R) \Rightarrow (S \notin \rho)) \wedge ((R' = S) \Rightarrow (R \notin \rho))$$

2. *Prerequisite rôle*: a user can activate rôle R only if he has already activated rôle S . This can be written as

$$\text{constr}_{S \rightarrow R}(R', \rho) \triangleq (R' = R) \Rightarrow (S \in \rho)$$

3. *Cardinality constraints (1)*: at most n rôles can be activated in each user session. This can be expressed as

$$\text{constr}_{|\cdot| \leq n}(R, \rho) \triangleq |\rho \cup \{R\}| \leq n$$

4. *Cardinality constraints (2)*: each user can own at most n permissions simultaneously. This can be enforced by requiring that

$$\text{constr}_{|\mathcal{P}(\cdot)| \leq n}(R, \rho) \triangleq |\mathcal{P}(\{R\} \cup \rho)| \leq n$$

Example 9 (Prerequisite rôle). The concept of *prerequisite rôle* is based on competency: in some circumstances, one may want to require a rôle to be activated only by a user already playing a certain rôle. For example, a common feature of a bank policy is to require an authentication phase to identify clients before any sensible operation, like money withdrawal. In practice, this amounts to ask for a valid identity document or a secret code/password. In our refined framework, this scenario can be modelled by letting the RBAC schema $(\mathcal{U}; \mathcal{P}; C)$ be such that $\{\text{client}, \text{authenticated}\} \in \mathcal{U}(r)$, $\mathcal{U}(\text{wdrw}^{\text{bank}}) = \text{wdrw}, \text{wdrw}^! \in \mathcal{P}(\text{client})$ and $\text{constr}_{\text{authenticated} \rightarrow \text{client}} \in C$. Hence, client

$$r \parallel \text{role client.wdrw}^{\text{bank}} \langle \text{amount} \rangle \parallel_0$$

cannot be typed, as the activation of rôle `client` is forbidden by C , while client

$$r \parallel \text{role authenticated.role client.wdrw}^{\text{bank}} \langle \text{amount} \rangle \parallel_0$$

can be typed. ◇

Example 10 (Mutual Exclusion). For the sake of fairness, sometimes it is desirable to control the distribution of sensible permissions; e.g., a user willing to perform a sensible operation should be different from the user in charge of controlling the legality of such an operation. Consider a scenario where some scientists submit a paper to a journal. Clearly, the reviewers of that paper cannot be chosen among the authors of the paper itself. This requirement can be modelled in our framework by having two rôles, `paperP : author` and `paperP : reviewer`, such that C contains the constraint $\text{constr}_{\text{paperP:author} \oplus \text{paperP:reviewer}}$. ◇

6 Related Work

To the best of our knowledge, no previous study building on process-calculi has ever been conducted on RBAC. A number of papers have instead dealt with the formal specification and verification of RBAC schemata. In [16, 25] formal methods are used only to verify the correctness of the schema definition but not of the whole system. In [25], the ALLOY language is used to detect possible conflicts in RBAC schemata supporting simultaneously delegation of authority and separation of duties. A constraint analyser allows the schema validation to be computed automatically. In [16, 17], the authors use a graph transformation which combines an intuitive visual description of the RBAC schema with solid semantical foundations. In [1], Ahn et al. introduce a formal language for the specification of more sophisticated role-based authorisation constraints, such as prohibition and obligation constraints. These approaches are complementary to ours: they can be integrated with our technique in order to verify the consistency of a schema \mathcal{S} , but they do not give any hint about the correct execution of a system as our method does.

In [3], Bertino et al. develop a logical framework for reasoning about access control models. The framework is general enough to model discretionary, mandatory and role-based access control models. Such a framework is useful for comparing the expressive power of the models, but it cannot be used to verify the correct execution of a system under a given schema.

Probably, the most related work, although not aiming at studying RBAC systems, is [6], insofar as roles can be understood as (privilege) groups. *Groups* are introduced in *loc. cit.* as types for channels, and used to limit their visibility. A type system ensures that channels belonging to a fresh group can be only used by processes within the initial scope of the group. Thus, processes can access channels according to their physical distribution (with respect to group restrictions). In our work this feature is modified so that not only the place where the process runs (i.e., the user running the process) but also its execution history (i.e., the user session where the process runs and the associated activations/deactivations of roles) is relevant to execute an action. E.g., outputs over \bar{a} of group R can be executed only by processes whose user r is such that $R^1 \in \mathcal{P}(\mathcal{U}(r))$; moreover, such an action must be enabled by at least one of the roles active in r 's session. The set of such sessions changes according to the computation and, thus, the processes enabled to access a channel change dynamically. In this sense, this work can be seen as a calculus of *dynamic* groups.

Acknowledgements. This work has been partially supported by EU FET – Global Computing initiative, projects MIKADO IST-2001-32222 and MyThS IST-2001-32617, and by the FIRB project RBAU018RCZ.002. The funding bodies are not responsible for any use that might be made of the results presented here.

A first draft of this paper has been improved by following the valuable suggestions and comments of the CSFW'04 and of the JCS anonymous referees. The second author is very grateful to Angelo Monti for his fundamental support in the development of the algorithm in Table 6.

A Technical Proofs

In this section, we give details on the proofs omitted from the body of the paper.

A.1 Proofs of Section 2

To prove subject reduction, we first need three lemmata, that are standard results for a type system. The first one states that names can be replaced with other names of the same type. The second

one states that enlarging the assumptions in a type judgement does not compromise the inference of the judgement itself. Finally, the third result states that well-typedness is an invariant of structural congruence. To prove the latter, we formally define a *system context*, as \equiv is closed under all such contexts. Formally, a context $C[\cdot]$ is a system with an occurrence of a ‘hole’ to be filled with any system A , thus yielding $C[A]$. Formally,

$$C[\cdot] ::= [\cdot] \mid C[\cdot] \parallel B \mid (va^r : R)C[\cdot]$$

Lemma A.1 (Substitution). *If $\Gamma, x : T, \Gamma'; \rho \vdash_r^S P$ and $\Gamma \vdash n : T$, then $\Gamma, \Gamma'; \rho \vdash_r^S P[n/x]$.*

Proof. The proof is by induction on the depth of the inference of the type judgement $\Gamma, x : T, \Gamma'; \rho \vdash_r^S P$. The proof is quite standard and faithfully rephrases the corresponding result for the pure π -calculus; thus, we omit it. \square

Lemma A.2 (Weakening). *If $\Gamma \vdash^{(\mathcal{U}; \mathcal{P})} A$ then $\Gamma, \tilde{n} : \tilde{T} \vdash^{(\mathcal{U} \cup \tilde{n}; \tilde{T}; \mathcal{P} \cup \mathcal{P}')} A$ for all \mathcal{P}' and $\tilde{n} : \tilde{T}$ such that $\tilde{n} \cap \text{dom}(\mathcal{U}) = \emptyset$.*

Proof. The proof is by induction on the depth of the inference for $\Gamma \vdash^{(\mathcal{U}; \mathcal{P})} A$. \square

Lemma A.3. *If $\Gamma \vdash^S A$ and $A \equiv B$, then $\Gamma \vdash^S B$.*

Proof. By mutual induction on the depth of the inferences for $A \equiv B$ and $B \equiv A$. Let us consider how $A \equiv B$ has been inferred; the case for $B \equiv A$ is similar. The base case covers the axioms in Table 3. All the cases are simple; we consider two samples, viz. (S-LOCCH) and (S-EXT). Let $A \triangleq r \parallel (va : R)P \parallel_\rho$ and $B \triangleq (va^r : R)r \parallel P \parallel_\rho$ (the converse is similar). By hypothesis, using rules (T-SESSION) and (T-RES), we have that $\Gamma \vdash r : \rho'[\tilde{a} : \tilde{C}]$, $\rho \subseteq \rho'$ and $\Gamma, a^r : R(T); \rho \vdash_r^S P$; this suffices to conclude $\Gamma \vdash^S B$, by using rules (T-SESSION) and (T-SYSRES). The case when $A \triangleq (va^r : R)A' \parallel B'$ and $B \triangleq (va^r : R)(A' \parallel B')$ is similar (the converse is symmetric): the typing of A has been inferred via (T-SYSPAR) and (T-SYSRES); thus, B can be typed by using the same rules in the reverse order.

The inductive steps for symmetry and transitivity follow straightforwardly. For context closure, let $A \triangleq C[A_1]$ and $B \triangleq C[B_1]$, for some $A_1 \equiv B_1$. We now work by induction on the structure of $C[\cdot]$. The base case is when $C[\cdot] \triangleq [\cdot]$ and is trivial. For the inductive case, let us reason by case analysis on the outermost operator of $C[\cdot]$. If $C[\cdot] \triangleq \mathcal{D}[\cdot] \parallel \tilde{A}$, then, by using (T-SYSPAR), we know by hypothesis that $\Gamma \vdash^S \mathcal{D}[A_1]$ and $\Gamma \vdash^S \tilde{A}$. By induction hypothesis, as $\mathcal{D}[\cdot]$ is smaller than $C[\cdot]$, it holds that $\Gamma \vdash^S \mathcal{D}[B_1]$; then, by (T-SYSPAR), $\Gamma \vdash^S B$. The case for $C[\cdot] \triangleq (va^r : R)\mathcal{D}[\cdot]$ is similar, but relies on (T-SYSRES). \square

Theorem 2.1 (Subject Reduction). *If $\Gamma \vdash^S A$ and $A \mapsto A'$, then $\Gamma \vdash^S A'$.*

Proof. By induction on the depth of the derivation of $A \mapsto A'$.

Base Step: By case analysis on the axioms of the second part of Table 3.

(R-COM) By hypothesis, we have that $\Gamma \vdash^S r \parallel a(x).P \parallel_\rho \parallel s \parallel a^r \langle n \rangle . Q \parallel_{\rho'}$. Due to the form of the system involved, (T-SYSPAR) is the last rule applied to deduce the type judgement, hence we also have that $\Gamma \vdash^S r \parallel a(x).P \parallel_\rho$ and $\Gamma \vdash^S s \parallel a^r \langle n \rangle . Q \parallel_{\rho'}$. The latter two judgements must have been derived by using (T-SESSION), with $\Gamma \vdash r : \rho''[\tilde{a} : \tilde{C}]$ and $\rho \subseteq \rho''$, $\Gamma \vdash s : \rho'''[\tilde{a}' : \tilde{C}']$ and $\rho' \subseteq \rho'''$, $\Gamma; \rho \vdash_r^S a(x).P$ and $\Gamma; \rho' \vdash_s^S a^r \langle n \rangle . Q$. Judgement $\Gamma; \rho \vdash_r^S a(x).P$ must have been derived by using (T-INPUT), with $\Gamma \vdash a^r : R(T)$ and $\Gamma, x : T; \rho \vdash_r^S P$, whereas judgement $\Gamma; \rho' \vdash_s^S a^r \langle n \rangle . Q$ must have been derived by using (T-OUTPUT), with $\Gamma \vdash a^r : R(T)$, $\Gamma \vdash n : T$, and $\Gamma; \rho' \vdash_s^S Q$. By Lemma A.1, we get that $\Gamma; \rho \vdash_r^S P[n/x]$. By a double application of (T-SESSION) and of (T-SYSPAR), we get that $\Gamma \vdash^{\mathcal{P}} r \parallel P[n/x] \parallel_\rho \parallel s \parallel Q \parallel_{\rho'}$, as required.

(R-ROLE) By hypothesis, we have that $\Gamma \vdash^S r \parallel \mathbf{role} R.P \parallel_{\rho}$. Due to the form of the system involved, (T-SESSION) is the last rule applied to deduce the type judgement, hence we also have that $\Gamma \vdash r : \rho'[\tilde{a} : \tilde{C}]$, $\rho \subseteq \rho'$ and $\Gamma; \rho \vdash_r^S \mathbf{role} R.P$. Judgement $\Gamma \vdash r : \rho'[\tilde{a} : \tilde{C}]$ must have been derived by using (T-ID₁), with $\Gamma(r) = \rho'[\tilde{a} : \tilde{C}]$; judgement $\Gamma; \rho \vdash_r^S \mathbf{role} R.P$ has been derived by using (T-ROLE), with $\Gamma \vdash r : \beta[\tilde{a}' : \tilde{C}']$, $\Gamma; \rho \cup \{R\} \vdash_r^S P$ and $R \in \rho''$. Then, $\rho \cup \{R\} \subseteq \rho'$; by rule (T-SESSION), we can derive $\Gamma \vdash^S r \parallel P \parallel_{\rho \cup \{R\}}$, as required.

(R-YIELD) By hypothesis, we have that $\Gamma \vdash^S r \parallel \mathbf{yield} R.P \parallel_{\rho}$. Due to the form of the system involved, (T-SESSION) is the last rule applied to deduce the type judgement, hence we also have that $\Gamma \vdash r : \rho'[\tilde{a} : \tilde{C}]$, $\Gamma; \rho \vdash_r^S \mathbf{yield} R.P$ and $\rho \subseteq \rho'$. Judgement $\Gamma; \rho \vdash_r^S \mathbf{yield} R.P$ has been derived by using (T-YIELD), with $\Gamma; \rho \setminus \{R\} \vdash_r^S P$ and $R \in \rho$. By applying rule (T-SESSION) to $\Gamma; \rho \setminus \{R\} \vdash_r^S P$, we can derive $\Gamma \vdash^S r \parallel P \parallel_{\rho \setminus \{R\}}$, as required.

Inductive Step: By case analysis of the last applied operational rule of the second part of Table 3.

(R-RES) By definition, $A \triangleq (\nu d : R)B$ and $A' \triangleq (\nu a' : R)B'$, where $B \mapsto B'$; moreover, by hypothesis, we have that $\Gamma \vdash^S (\nu a' : R)B$. By rule (T-SYSRES), we have that $\Gamma, a' : R(T) \vdash^S B$, for some T . By induction hypothesis, $\Gamma, a' : R(T) \vdash^S B'$ that, by rule (T-SYSRES), implies $\Gamma \vdash^S A'$.

(R-PAR) By hypothesis, $A \triangleq A_1 \parallel B$ is well-typed; hence, A_1 and B are well-typed too. Moreover, $A_1 \mapsto A'_1$ and the induction hypothesis imply that A'_1 is well-typed; thus, $A'_1 \parallel B \triangleq A'$ is well-typed too.

(R-STRUCT) We now have that $A \equiv A_1 \mapsto A_2 \equiv A'$. By well-typedness of A and Lemma A.3, it follows that A_1 is well-typed; by induction hypothesis, it follows that A_2 is well-typed and, again by Lemma A.3, A' is well-typed. \square

Theorem 2.2 (Type Safety). *If A is well-typed in \mathcal{S} , then $A \rightsquigarrow_{\mathcal{S}}$ cannot hold.*

Proof. We prove the contrapositive, i.e. $A \rightsquigarrow_{\mathcal{S}}$ implies that A cannot be well-typed in \mathcal{S} ; this is done by induction on the depth of the inference for $A \rightsquigarrow_{\mathcal{S}}$. Let \mathcal{S} be $(\mathcal{U}; \mathcal{P})$. For the base case, we consider only one sample, namely when the judgement has been inferred via (E-IN); the other cases are similar. By definition, A is $r \parallel b(x).P \parallel_{\rho}$ and $S \notin \mathcal{P}(\rho)$, for $S = \mathcal{U}(b')$. Thus, for any Γ respecting \mathcal{U} , it cannot hold that $\Gamma; \rho \vdash_r^S b(x).P$: indeed, the premises of rule (T-INPUT) (that is the only applicable to infer the judgement) cannot be satisfied.

For the inductive step, we only consider the case when the last rule used is (E-RES); the cases for (E-PAR) and (E-STRUCT) are simpler (the latter one relies on Lemma A.3). By definition, A is $(\nu a' : R)B$ and $B \rightsquigarrow_{(\mathcal{U} \uplus \{a' : R\}; \mathcal{P})}$. By induction hypothesis, B cannot be well-typed in $(\mathcal{U} \uplus \{a' : R\}; \mathcal{P})$, i.e. for every Γ respecting $\mathcal{U} \uplus \{a' : R\}$, judgement $\Gamma \vdash^{(\mathcal{U} \uplus \{a' : R\}; \mathcal{P})} B$ cannot be inferred. Since Γ respects $\mathcal{U} \uplus \{a' : R\}$, it must be that $\Gamma = \Gamma', a' : R$; this easily implies that there is no Γ' such that $\Gamma' \vdash^S (\nu a' : R)B$, as desired. \square

A.2 Proofs of Section 4

Proposition 4.1. *Function LABEL can be used to induce a partition of t 's nodes in subtrees satisfying the requirements of the minimal partition. Moreover, the overall procedure takes $O(|V| \times K^2)$, where K is the size of the largest set annotating a node of the tree.*

Proof. Having computed function LABEL(\cdot), we proceed in the following way:

```

Visit  $t$  in preorder
When visiting the node  $v$  do
    if  $v = rt$  OR  $\text{LABEL}(v) \neq \text{LABEL}(\text{FATHER}(v))$ 
        then add  $v$  in a new block
    else add  $v$  in the block of  $\text{FATHER}(v)$ 

```

It should be clear that the output of this procedure is a partition of V (no block is empty and each node is inserted in exactly one block); we call it the partition induced by $\text{LABEL}(\cdot)$. We have to prove that this partition satisfies the following conditions: (a) each block is a subtree of t ; (b) each block β is such that $\exists R \in \mathcal{R} \forall v \in \beta : R \in \phi(v)$ (in this case, we call R the *pivot* for β); (c) it has the minimum number of blocks satisfying the previous two properties.

Condition (a) is proved by induction on the size of the generic block β . The base case is for $\beta = \{v\}$ and is trivial, as a single node is a subtree of t . For the inductive step, let $\beta = \beta' \cup \{v\}$, where v is the last node added by the above procedure. By construction, it must be that $\text{FATHER}(v) \in \beta'$ and, by induction hypothesis, β' is a subtree of t . Thus, easily, also β is a subtree of t . Condition (b) is simple: by construction, it holds that $\text{LABEL}(v) = \text{LABEL}(v')$ and $\text{LABEL}(v) \in \phi(v)$, for every v and v' belonging to β . Condition (c) easily follows, once we prove the following Lemma.

Lemma A.4 (Soundness of the algorithm in Table 6). *If $m[v, R] = h \neq \infty$, then there exists a partition of the subtree rooted in v with h blocks such that it satisfies conditions (a) and (b) above, and R is the pivot of v 's block; moreover, each partition satisfying these properties has at least h blocks.*

Proof. By induction on the height of the tree rooted in v . For notational convenience, we denote the subtree of t rooted in v as t_v . The base case is when v is a leaf and is trivial. For the inductive step, we only consider the case when v has just one child, v' ; the case when v has two children is similar. Since $m[v, R] \neq \infty$, it holds that $R \in \phi(v)$; thus, R can be the pivot of v 's block. Since $\text{LABEL}(\cdot)$ is defined (by the hypothesis of Proposition 4.1), there must exist $S \in \phi(v')$; thus, $m[v', S] \neq \infty$. By induction hypothesis, there exists a partition of $t_{v'}$ with $m[v', S]$ blocks that satisfies conditions (a) and (b), and with S as pivot of v' 's block; moreover, each partition satisfying these properties has at least $m[v', S]$ blocks. If R is not one of such S , then a minimal partition of t_v with R as pivot of v 's block can be obtained by putting v in a block on its own and by considering the partition of $t_{v'}$ induced by a S that minimises $m[v', _]$. Otherwise, adding v to the block of v' in the partition of $t_{v'}$ induced by R could generate a minimal partition of t_v or not. In the first case, the partition of t_v has $m[v', R]$ blocks; in the second case, we put v in a block on its own and the resulting partition has $m[v', R] + 1$ blocks. In both cases, it is easy to prove that no partition with less blocks can exist. \square

Now, let $\text{LABEL}(rt) = R$. Trivially, the partition induced by function $\text{LABEL}(\cdot)$ has $m[rt, R]$ blocks; thus, by Lemma A.4, each partition satisfying (a) and (b) has at least $m[rt, R]$ blocks. This proves (c).

We conclude with the complexity of the overall algorithm. The algorithm in Table 6 to compute $m[-, -]$ works in $O(|V| \times K^2)$. Indeed, matrix m has $|V|$ rows and K columns, and each element of this matrix is written exactly once. Moreover, to write a generic element $m[v, R]$, we have to check whether $R \in \phi(v)$ (this requires $O(K)$, as $|\phi(v)| \leq K$) and to analyse the rows associated to the children of v (if any); the latter task requires $O(K)$, that leads the overall complexity to $O(|V| \times K^2)$. The algorithm for computing function LABEL works in $O(|V| \times K)$. Indeed, for each node v , it has to compute m_v :

this requires $O(K)$, as it has to scan all the row of m associated to v . Finally, the partition induced by $\text{LABEL}(\cdot)$ is derived by a preorder visit, that is linear in $|V|$. \square

We now consider Proposition 4.2 and prove its two claims separately; for both of them, we only present a key sample, leaving the other cases (that are similar) to the interested reader.

1. Let $\mathcal{S} = (\mathcal{U}; \mathcal{P})$ be a RBAC schema and P be a finite process without **role/yield** to be run by user r . Then $\Gamma; \mathcal{U}(r) \vdash_r^{\mathcal{S}} P$ implies that $\text{REFINED}_{\Gamma; \mathcal{P}}^r(P)$ is defined.

Proof. The proof is by structural induction on P . The base step is trivial: P is **nil** and, by definition, $\text{REFINED}_{\Gamma; \mathcal{P}}^r(\mathbf{nil}) = \mathbf{nil}$. For the inductive step, we only consider the case for $P = a(x).Q$. By hypothesis and by rule (T-INPUT), we have that

$$\frac{\Gamma \vdash a^r : R(T) \quad R^? \in \mathcal{P}(\mathcal{U}(r)) \quad \Gamma, x : T; \mathcal{U}(r) \vdash_r^{\mathcal{S}} Q}{\Gamma; \mathcal{U}(r) \vdash_r^{\mathcal{S}} P}$$

By induction hypothesis, $\text{REFINED}_{\Gamma, x; T; \mathcal{P}}^r(Q)$ is defined; this implies that $\text{TREE}_{\Gamma, x; T; \mathcal{P}}^r(Q)$ is defined and is equipped with the matrix $m[\cdot, \cdot]$. Now, by construction, $\text{TREE}_{\Gamma; \mathcal{P}}^r(P)$ is defined; moreover, it can be equipped with a matrix $m'[\cdot, \cdot]$ such that

$$m'[x, R] = \begin{cases} m[x, R] & \text{if } x \text{ is not the root of } \text{TREE}_{\Gamma; \mathcal{P}}^r(P) \\ \min\{ m[v', R], \\ \min_{S \neq R} \{ m[v', S] \} + 1 \} & \text{otherwise} \end{cases}$$

where, in the second case, v' is the root of $\text{TREE}_{\Gamma, x; T; \mathcal{P}}^r(Q)$. By the premise of (T-INPUT), we know that there exists a role $S \in \mathcal{U}(r)$ such that $R \in \mathcal{P}(S)$. This fact, together with the fact that $\text{REFINED}_{\Gamma, x; T; \mathcal{P}}^r(Q)$ is defined, implies that there exists $S \in \mathcal{U}(r)$ such that $m'[v', S] \neq \infty$, where v' is the root of $\text{TREE}_{\Gamma; \mathcal{P}}^r(P)$. This fact suffices to conclude that $\text{REFINED}_{\Gamma; \mathcal{P}}^r(P)$ is defined. \square

2. Let $\mathcal{S} = (\mathcal{U}; \mathcal{P})$ be a RBAC schema, Γ be a typing environment respecting \mathcal{U} and P be a finite process without **role/yield** to be run by user r . Then $\text{REFINED}_{\Gamma; \mathcal{P}}^r(P) = P'$ implies that $\Gamma; \emptyset \vdash_r^{\mathcal{S}} P'$.

Proof. Again, the proof is by structural induction on P . The base step is trivial: $P = P' \triangleq \mathbf{nil}$ and $\Gamma; \emptyset \vdash_r^{\mathcal{S}} \mathbf{nil}$. For the inductive step, we only consider the case for $P = a(x).Q$. By construction, we have that $P' = \mathbf{role} R.a(x).Q'$, where R is the label of the root of $\text{TREE}_{\Gamma; \mathcal{P}}^r(P)$, that exists by hypothesis. By construction, this latter fact implies that $\text{TREE}_{\Gamma, x; T; \mathcal{P}}^r(Q)$ exists and that $\Gamma \vdash a^r : S(T)$, for some role S and type T . Let us now consider the matrix $m'[\cdot, \cdot]$ obtained from the matrix $m[\cdot, \cdot]$ for $\text{TREE}_{\Gamma; \mathcal{P}}^r(P)$ by deleting the row associated with the root of $\text{TREE}_{\Gamma; \mathcal{P}}^r(P)$. We can now say that there exists a role R such that $m'[v', R'] \neq \infty$, where v' is the root of $\text{TREE}_{\Gamma, x; T; \mathcal{P}}^r(Q)$.⁴ This suffices to conclude that $\text{REFINED}_{\Gamma, x; T; \mathcal{P}}^r(Q)$ is defined; let us say that it returns the process Q' . By induction hypothesis, $\Gamma, x : T; \emptyset \vdash_r^{\mathcal{S}} Q''$. Again, by construction it must be that $Q'' = \mathbf{role} R''.\hat{Q}$; we now consider the only possible cases:

⁴To see this, proceed by contradiction. Assume that, for all R' , $m'[v', R'] = \infty$; then $m[v, R'] = \infty$, where v is the root of $\text{TREE}_{\Gamma; \mathcal{P}}^r(P)$, for every role R (see Table 6, second item). Thus, $\text{REFINED}_{\Gamma; \mathcal{P}}^r(P)$ would be undefined, as function LABEL would be. Contradiction.

- if $\text{LABEL}(v) = \text{LABEL}(v')$, i.e. $R = R''$, then $Q' = \hat{Q}$. In this case, we have that

$$\frac{R \in \Gamma\{r\} \quad \frac{\Gamma \vdash a^r : S(T) \quad S^? \in \mathcal{P}(R) \quad \Gamma, x : T; \{R\} \vdash_r^S Q'}{\Gamma; \{R\} \vdash_r^S a(x).Q'}}{\Gamma; \emptyset \vdash_r^S P'}$$

This inference holds by using rules (T-R^{OLE}) and (T-INPUT). Moreover, notice that $R \in \Gamma\{r\}$ and $S^? \in \mathcal{P}(R)$ must hold, otherwise $\text{LABEL}(v)$ cannot be R . Finally, $\Gamma, x : T; \{R\} \vdash_r^S Q'$ is implied by the induction hypothesis.

- if $\text{LABEL}(v) \neq \text{LABEL}(v')$, then $Q' = \mathbf{yield} R.Q''$. This case is similar to the previous one, but $\Gamma, x : T; \{R\} \vdash_r^S Q'$ is inferred from the induction hypothesis, by using rule (T-YIELD). \square

B Alternative Characterisation of Barbed Congruence

As pointed out in Section 3, barbed congruence is hard to prove because of its universal quantification over language contexts. A standard way to overcome this problem is to reformulate the semantics of the language via a *labelled transition system* (LTS for short), that makes apparent the external interaction offered, and build up over it a bisimulation, adequate for barbed congruence. In this section, we present a possible way to adapt known techniques to our framework; however, to make the presentation lighter, most of the proofs in this section are only sketched; the interested reader is referred to a technical report [5] for full details.

The LTS allows to study system components in isolation and compositionally. Thus, in general, we cannot assume such components to be well-typed, as this would require a full knowledge of the system. Hence, we embody in the LTS some dynamic policy checks. In this way, the LTS also provides a tight operational specification for the minimal engine underlying any implementation of a RBAC-based run-time system.

The standard way to describe the interactions a system can offer externally is by labelling the system evolution with this information. Thus, we define a labelled transition system, $\xrightarrow{\mu}$, that makes apparent the action performed (and, thus, the external interaction offered). In order to account for systems' roles varying over time, the LTS relates *configurations*, i.e. pairs $\mathcal{S} \triangleright A$ made up of a RBAC schema \mathcal{S} and a system A . Configurations are ranged over by D, E, \dots . The labels of the LTS are derived from those of the π -calculus and can be described as follows.

$$\mu ::= \tau \mid a^r n \mid a^r n : R \mid \bar{a}^r n \mid \bar{a}^r n : R$$

Label τ represents an internal computation of the system. Labels $\bar{a}^r n$ and $a^r n$ describe the intention to send/receive value n , known to the environment, on/from channel a^r . Labels $\bar{a}^r n : R$ and $a^r n : R$ are similar to but the value sent/received is 'fresh' (i.e. unknown to the environment) and has group R . We now extend functions $\text{FR}(-)$ and $\text{BND}(-)$ to labels.

Label	FR(-)	BND(-)
τ	\emptyset	\emptyset
$a^r n$	$\{a^r, n\}$	\emptyset
$\bar{a}^r n$	$\{a^r, n\}$	\emptyset
$a^r n : R$	$\{a^r\}$	$\{n\}$
$\bar{a}^r n : R$	$\{a^r\}$	$\{n\}$

(LTS-R $\hat{\text{O}}\text{LE}$)	$\frac{R \in \mathcal{U}(r)}{(\mathcal{U}; \mathcal{P}) \triangleright r \parallel \mathbf{role} R.P \parallel_{\rho} \xrightarrow{\tau} (\mathcal{U}; \mathcal{P}) \triangleright r \parallel P \parallel_{\rho \cup \{R\}}}$
(LTS-YIELD)	$\frac{R \in \rho}{\mathcal{S} \triangleright r \parallel \mathbf{yield} R.P \parallel_{\rho} \xrightarrow{\tau} \mathcal{S} \triangleright r \parallel P \parallel_{\rho - \{R\}}}$
(LTS-OUTPUT)	$\frac{\mathcal{U}(a^s) = R \quad R^l \in \mathcal{P}(\rho)}{(\mathcal{U}; \mathcal{P}) \triangleright r \parallel a^s \langle n \rangle . P \parallel_{\rho} \xrightarrow{\bar{a}^s n} (\mathcal{U}; \mathcal{P}) \triangleright r \parallel P \parallel_{\rho}}$
(LTS-K-INPUT)	$\frac{\mathcal{U}(a^r) = R \quad R^? \in \mathcal{P}(\rho) \quad n \in \text{dom}(\mathcal{U})}{(\mathcal{U}; \mathcal{P}) \triangleright r \parallel a(x).P \parallel_{\rho} \xrightarrow{a^r n} (\mathcal{U}; \mathcal{P}) \triangleright r \parallel P[n/x] \parallel_{\rho}}$
(LTS-F-INPUT)	$\frac{\mathcal{U}(a^r) = R \quad R^? \in \mathcal{P}(\rho) \quad n \notin \text{dom}(\mathcal{U})}{(\mathcal{U}; \mathcal{P}) \triangleright r \parallel a(x).P \parallel_{\rho} \xrightarrow{a^r n : S} (\mathcal{U} \uplus \{n : S\}; \mathcal{P}) \triangleright r \parallel P[n/x] \parallel_{\rho}}$

Table 7: Axioms for the Labelled Transition System

The definition of $\xrightarrow{\mu}$ is given in Tables 7 and 8. The overall structure of the LTS is similar to π -calculus' early-style one (see, e.g., [24]) and implicitly assumes alpha-conversion. The premises of rules (LTS-K-INPUT), (LTS-F-INPUT), (LTS-OUTPUT), (LTS-R $\hat{\text{O}}\text{LE}$) and (LTS-YIELD) adapt respectively the premises of the typing rules (T-INPUT), (T-OUTPUT), (T-R $\hat{\text{O}}\text{LE}$) and (T-YIELD), and block the evolution of ill-typed systems. Rule (LTS-K-INPUT) can be applied when the received value is known to the schema, while (LTS-F-INPUT) is used when a fresh value (i.e. unknown to the schema) is received. In this case, the schema is extended to record the group of the fresh value. Similarly, when extruding a restricted channel b^s , rule (LTS-OPEN) enlarges the relation \mathcal{U} of the current configuration by recording that b^s has the r $\hat{\text{o}}\text{le}$ declared in the restriction. The information about a fresh/extruded channel is deleted from the schema when the channel is communicated: indeed, the restriction is pushed back in the system and closes the scope of the channel – cf. rule (LTS-CLOSE). Notice that a bound output can synchronise only with a fresh input (and vice versa), and the r $\hat{\text{o}}\text{le}$ declared for the extruded/fresh channel must be the same. Also observe that τ -moves do not modify the schema \mathcal{S} .

The semantics given in Table 3 and the LTS just presented are related by the following Proposition.

Proposition B.1. *If $\mathcal{S} \triangleright A \xrightarrow{\tau} \mathcal{S} \triangleright A'$, then $A \mapsto A'$. Conversely, if A is well-typed in \mathcal{S} , then $A \mapsto A'$ implies $\mathcal{S} \triangleright A \xrightarrow{\tau} \mathcal{S} \triangleright B$, for some $B \equiv A'$.*

Proof. The first statement is proved by a simple induction over the depth of the inference for $\xrightarrow{\tau}$. The second statement is proved by induction over the depth of the shortest inference for \mapsto . The only complicated case is when the last rule applied to infer the reduction is (R-STRUCT), i.e. $A \equiv B$, and $B \mapsto B'$ and $B' \equiv A'$. We proceed by mutual induction on the depth of the inferences for $A \equiv B$ and $B \equiv A$. We reason on $A \equiv B$, but everything can be rephrased for $B \equiv A$. Moreover, we assume that the last rule to infer $B \mapsto B'$ is not (R-STRUCT), otherwise the original inference of $A \mapsto A'$ could be shortened, thanks to transitivity of \equiv .

All the base cases are simple, except when we use rule (S-EXT) from left to right, i.e. when $A \triangleq (va^r : R)A_1 \parallel B_1$ and $B \triangleq (va^r : R)(A_1 \parallel B_1)$. By hypothesis, we have that $a^r \notin \text{FR}(B_1)$; moreover, by induction hypothesis, $(va^r : R)(A_1 \parallel B_1) \xrightarrow{\tau} B'$ (to simplify notations, in what follows we shall omit the RBAC schema \mathcal{S} whenever it does not come into the picture). By definition of the LTS,

<p>(LTS-COMM)</p> $\frac{S \triangleright A \xrightarrow{a^n} S \triangleright A' \quad S \triangleright B \xrightarrow{\bar{a}^n} S \triangleright B'}{S \triangleright A \parallel B \xrightarrow{\tau} S \triangleright A' \parallel B'}$	
<p>(LTS-RES)</p> $\frac{(\mathcal{U} \uplus \{a^r : R\}; \mathcal{P}) \triangleright A \xrightarrow{\mu} (\mathcal{U} \uplus \{a^r : R\}; \mathcal{P}) \triangleright A' \quad a^r \notin \text{FR}(\mu)}{(\mathcal{U}; \mathcal{P}) \triangleright (va^r : R)A \xrightarrow{\mu} (\mathcal{U}; \mathcal{P}) \triangleright (va^r : R)A'}$	
<p>(LTS-OPEN)</p> $\frac{(\mathcal{U} \uplus \{b^s : S\}; \mathcal{P}) \triangleright A \xrightarrow{\bar{a}^r b^s} (\mathcal{U} \uplus \{b^s : S\}; \mathcal{P}) \triangleright A' \quad a^r \neq b^s}{(\mathcal{U}; \mathcal{P}) \triangleright (vb^s : S)A \xrightarrow{\bar{a}^r b^s : S} (\mathcal{U} \uplus \{b^s : S\}; \mathcal{P}) \triangleright A'}$	
<p>(LTS-CLOSE)</p> $\frac{S \triangleright A \xrightarrow{a^r b^s : S} S' \triangleright A' \quad S \triangleright B \xrightarrow{\bar{a}^r b^s : S} S' \triangleright B' \quad b^s \notin \text{FR}(A)}{S \triangleright A \parallel B \xrightarrow{\tau} S \triangleright (vb^s : S)(A' \parallel B')}$	
<p>(LTS-PAR)</p> $\frac{S \triangleright A \xrightarrow{\mu} S' \triangleright A' \quad \text{BND}(\mu) \cap \text{FR}(B) = \emptyset}{S \triangleright A \parallel B \xrightarrow{\mu} S' \triangleright A' \parallel B}$	
<p>(LTS-REPL)</p> $\frac{S \triangleright r \parallel P \mid !P \parallel_{\rho} \xrightarrow{\mu} S' \triangleright A}{S \triangleright r \parallel !P \parallel_{\rho} \xrightarrow{\mu} S' \triangleright A}$	<p>(LTS-EXT)</p> $\frac{S \triangleright (va^r : R)r \parallel P \parallel_{\rho} \xrightarrow{\mu} S' \triangleright A \quad a \neq r}{S \triangleright r \parallel (va : R)P \parallel_{\rho} \xrightarrow{\mu} S' \triangleright A}$
<p>(LTS-EQ)</p> $\frac{S \triangleright r \parallel P \parallel_{\rho} \xrightarrow{\mu} S' \triangleright A}{S \triangleright r \parallel [n = n]P \parallel_{\rho} \xrightarrow{\mu} S' \triangleright A}$	<p>(LTS-SPLIT)</p> $\frac{S \triangleright r \parallel P \parallel_{\rho} \parallel r \parallel Q \parallel_{\rho} \xrightarrow{\mu} S' \triangleright A}{S \triangleright r \parallel P \mid Q \parallel_{\rho} \xrightarrow{\mu} S' \triangleright A}$
<p>plus the symmetric version of rules of (LTS-PAR), (LTS-COMM) and (LTS-CLOSE)</p>	

Table 8: Inference Rules for the Labelled Transition System

(LTS-RES) is the last rule to infer such judgement; hence, $A_1 \parallel B_1 \xrightarrow{\tau} A_2$, where $B' \triangleq (va^r : R)A_2$. We now examine how this τ -step could have been generated.

1. $A_1 \xrightarrow{\tau} A'_1$ and $A_2 \triangleq A'_1 \parallel B_1$. Then, trivially, $(va^r : R)A_1 \xrightarrow{\tau} (va^r : R)A'_1$; hence $A \triangleq (va^r : R)A_1 \parallel B_1 \xrightarrow{\tau} (va^r : R)A'_1 \parallel B_1 \equiv (va^r : R)(A'_1 \parallel B_1) \equiv B' \equiv A'$.
2. $B_1 \xrightarrow{\tau} B'_1$ and $A_2 \triangleq A_1 \parallel B'_1$. Similar.
3. $A_1 \xrightarrow{b^s n} A'_1$, $B_1 \xrightarrow{\bar{b}^s n} B'_1$ and $A_2 \triangleq A'_1 \parallel B'_1$. Since $a^r \notin \text{FR}(B_1)$, we have that $a^r \notin \{b^s, n\}$. Thus, $(va^r : R)A_1 \xrightarrow{b^s n} (va^r : R)A'_1$; hence $A \xrightarrow{\tau} (va^r : R)A'_1 \parallel B'_1 \equiv A'$.
4. $A_1 \xrightarrow{\bar{b}^s n} A'_1$ (for $n \neq a^r$), $B_1 \xrightarrow{b^s n} B'_1$ and $A_2 \triangleq A'_1 \parallel B'_1$. Similar.

5. $A_1 \xrightarrow{\bar{b}^s a^r} A'_1$, $B_1 \xrightarrow{b^s a^r} B'_1$ and $A_2 \triangleq A'_1 \parallel B'_1$. Again, $b^s \neq a^r$; thus, by rule (LTS-OPEN), $(\mathcal{U}; \mathcal{P}) \triangleright (va^r : R)A_1 \xrightarrow{\bar{b}^s a^r : R} (\mathcal{U} \uplus \{a^r : R\}; \mathcal{P}) \triangleright A'_1$. By using (the symmetric version of) rule (LTS-CLOSE), that can be used because $a^r \notin \text{FR}(B_1)$, we have that $A \triangleq (va^r : R)A_1 \parallel B_1 \xrightarrow{\tau} (va^r : R)(A'_1 \parallel B'_1) \equiv A'$.
6. $A_1 \xrightarrow{b^s a^r} A'_1$ and $B_1 \xrightarrow{\bar{b}^s a^r} B'_1$. This case is not possible, as $a^r \notin \text{FR}(B_1)$.

When $A \equiv B$ has been inferred by symmetry and transitivity of \equiv , the claim can be proved by an easy induction. When $A \triangleq C[A_1]$, $B \triangleq C[B_1]$ and $A_1 \equiv B_1$, we reason by induction on the structure of $C[\cdot]$. The base case is trivial. When $C[\cdot] \triangleq \mathcal{D}[\cdot] \parallel \bar{A}$, the last rule to infer $\mathcal{D}[B_1] \parallel \bar{A} \mapsto B'$ is (R-PAR); thus, by induction hypothesis and (LTS-PAR) we can conclude. The case for $C[\cdot] \triangleq (va^r : R)\mathcal{D}[\cdot]$ is similar, by using rules (R-RES) and (LTS-RES). \square

Next, we build upon this LTS a standard bisimulation. As usual, \Rightarrow denotes the reflexive and transitive closure of $\xrightarrow{\tau}$, and $\xRightarrow{\mu}$ denotes $\Rightarrow \xrightarrow{\mu} \Rightarrow$. Finally, $\xRightarrow{\hat{\mu}}$ is \Rightarrow if $\mu = \tau$, and $\xRightarrow{\mu}$ otherwise.

Definition B.1 (Bisimilarity). A *bisimulation* is a binary symmetric relation \mathfrak{R} between configurations such that, if $(D, E) \in \mathfrak{R}$ and $D \xrightarrow{\mu} D'$, there exists a configuration E' such that $E \xRightarrow{\hat{\mu}} E'$ and $(D', E') \in \mathfrak{R}$. *Bisimilarity*, \approx , is the largest bisimulation.

We now state and prove some properties of \approx . First, we consider the congruence properties of \approx ; then, we prove that it is a sound proof technique for barbed congruence.

Theorem B.2 (Congruence Properties of \approx). *The following facts hold.*

1. If $S_1 \triangleright A_1 \approx S_2 \triangleright A_2$ and $S_1 \triangleright B \approx S_2 \triangleright B$, then $S_1 \triangleright A_1 \parallel B \approx S_2 \triangleright A_2 \parallel B$.
2. If $(\mathcal{U}_1 \uplus \{a^r : R\}; \mathcal{P}_1) \triangleright A_1 \approx (\mathcal{U}_2 \uplus \{a^r : R\}; \mathcal{P}_2) \triangleright A_2$, then $(\mathcal{U}_1; \mathcal{P}_1) \triangleright (va^r : R)A_1 \approx (\mathcal{U}_2; \mathcal{P}_2) \triangleright (va^r : R)A_2$.

To prove Theorem B.2, we first need some preliminary results for the LTS and for the relation \approx . We write $a^r \in \text{FR}(A, \mu, \bar{b}^s, \text{dom}(\mathcal{U}))$ to mean that $a^r \in \text{FR}(A) \cup \text{FR}(\mu) \cup \bar{b}^s \cup \text{dom}(\mathcal{U})$; derived notations have a similar meaning. A channel a^r is said to be *fresh* for $_$ if it does not occur in $_$, i.e. if $a^r \notin \text{FR}(_) \cup \text{BND}(_)$; we shall not specify $_$ whenever it refers all the entities (i.e. systems, schemata, sets of channels, ...) involved, and simply say that a^r is fresh. Moreover, $_ b^s : S$ denotes either $a^r b^s : S$ or $\bar{a}^r b^s : S$; similarly, $a^r _$ denotes either $a^r n$ or $a^r n : S$.

Lemma B.3. *The following facts hold.*

1. If $(\mathcal{U}; \mathcal{P}) \triangleright A \xrightarrow{\mu} (\mathcal{U}'; \mathcal{P}) \triangleright A'$ and a^r fresh, then $(\mathcal{U}[a^r/b^r]; \mathcal{P}) \triangleright A[a^r/b^r] \xrightarrow{\mu[a^r/b^r]} (\mathcal{U}'[a^r/b^r]; \mathcal{P}) \triangleright A'[a^r/b^r]$.
2. Let $(\mathcal{U} \uplus \{a^r : R\}; \mathcal{P}) \triangleright A \xrightarrow{\mu} (\mathcal{U}' \uplus \{a^r : R\}; \mathcal{P}) \triangleright A'$ and $a^r \notin \text{FR}(A)$. If $a^r \in \text{FR}(\mu)$, then $\mu = b^s a^r$ and $(\mathcal{U}; \mathcal{P}) \triangleright A \xrightarrow{b^s a^r : R} (\mathcal{U} \uplus \{a^r : R\}; \mathcal{P}) \triangleright A'$; otherwise, $(\mathcal{U}; \mathcal{P}) \triangleright A \xrightarrow{\mu} (\mathcal{U}'; \mathcal{P}) \triangleright A'$.
3. If $(\mathcal{U}_1; \mathcal{P}_1) \triangleright A \approx (\mathcal{U}_2; \mathcal{P}_2) \triangleright B$ and a^r is fresh, then $(\mathcal{U}_1 \uplus \{a^r : R\}; \mathcal{P}_1) \triangleright A \approx (\mathcal{U}_2 \uplus \{a^r : R\}; \mathcal{P}_2) \triangleright B$.

Proof. The first claim is proved by an easy induction over the depth of the inference for $\xrightarrow{\mu}$.

The second claim is proved as follows. Let $a^r \in \text{FR}(\mu)$. Since a^r is fresh for A , it can only be $\mu = b^s a^r$; thus, by rule (LTS-F-INPUT), $(\mathcal{U}; \mathcal{P}) \triangleright A \xrightarrow{b^s a^r : R} (\mathcal{U} \uplus \{a^r : R\}; \mathcal{P}) \triangleright A'$. On the contrary, if $a^r \notin \text{FR}(\mu)$, then a^r is fresh for A , μ and $\text{dom}(\mathcal{U})$; Thus, it is easy to prove, by induction on the depth of the inference for $\xrightarrow{\mu}$, that $(\mathcal{U}; \mathcal{P}) \triangleright A \xrightarrow{\mu} (\mathcal{U}'; \mathcal{P}) \triangleright A'$.

Finally, the last claim is proved by showing that the relation

$$\mathfrak{R} \triangleq \approx \cup \{((\mathcal{U}_1 \uplus \{a^r : R\}; \mathcal{P}_1) \triangleright A, (\mathcal{U}_2 \uplus \{a^r : R\}; \mathcal{P}_2) \triangleright B) : (\mathcal{U}_2; \mathcal{P}_2) \triangleright A \approx (\mathcal{U}_2; \mathcal{P}_2) \triangleright B \wedge a^r \notin \text{FR}(A, B)\}$$

is a bisimulation. Consider $(\mathcal{U}_1 \uplus \{a^r : R\}; \mathcal{P}_1) \triangleright A \xrightarrow{\mu} (\mathcal{U}'_1 \uplus \{a^r : R\}; \mathcal{P}_1) \triangleright A'$. Because of case (2) of this Proposition, if a^r occurs in μ then μ must be $b^s a^r$ for some b^s and $(\mathcal{U}_1; \mathcal{P}_1) \triangleright A \xrightarrow{b^s a^r : R} (\mathcal{U}_1 \uplus \{a^r : R\}; \mathcal{P}_1) \triangleright A'$. Thus, $(\mathcal{U}_2; \mathcal{P}_2) \triangleright B \xrightarrow{b^s a^r : R} (\mathcal{U}_2 \uplus \{a^r : R\}; \mathcal{P}_2) \triangleright B'$ and $(\mathcal{U}_1 \uplus \{a^r : R\}; \mathcal{P}_1) \triangleright A' \approx (\mathcal{U}_2 \uplus \{a^r : R\}; \mathcal{P}_2) \triangleright B'$. This implies that $(\mathcal{U}_2 \uplus \{a^r : R\}; \mathcal{P}_2) \triangleright B \xrightarrow{b^s a^r} (\mathcal{U}_2 \uplus \{a^r : R\}; \mathcal{P}_2) \triangleright B'$ and $((\mathcal{U}_1 \uplus \{a^r : R\}; \mathcal{P}_1) \triangleright A', (\mathcal{U}_2 \uplus \{a^r : R\}; \mathcal{P}_2) \triangleright B') \in \mathfrak{R}$. The case when a^r does not occur in μ is simpler; just notice that, upon transitions, a^r still remains fresh. \square

Proof of Theorem B.2 Both the clauses of the theorem are proved once shown that relation

$$\mathfrak{R} \triangleq \{((\mathcal{U}_1; \mathcal{P}_1) \triangleright (v \bar{a}^r : \bar{R})(A_1 \parallel B), (\mathcal{U}_2; \mathcal{P}_2) \triangleright (v \bar{a}^r : \bar{R})(A_2 \parallel B)) : (\mathcal{U}_1 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_1) \triangleright A_1 \approx (\mathcal{U}_2 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_2) \triangleright A_2 \wedge (\mathcal{U}_1 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_1) \triangleright B \approx (\mathcal{U}_2 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_2) \triangleright B\}$$

is a bisimulation. Let $(\mathcal{U}_1; \mathcal{P}_1) \triangleright (v \bar{a}^r : \bar{R})(A_1 \parallel B) \xrightarrow{\mu} (\mathcal{U}'_1; \mathcal{P}_1) \triangleright \bar{A}_1$; by definition of the LTS, it can only be one of the following cases.

1. $(\mathcal{U}_1 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_1) \triangleright A_1 \xrightarrow{\mu} (\mathcal{U}'_1 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_1) \triangleright A'_1$ for $\text{BND}(\mu) \cap \text{FR}(B) = \emptyset$. Thus, $\bar{A}_1 \triangleq (v \bar{a}^r : \bar{R})(A'_1 \parallel B)$. By hypothesis, $(\mathcal{U}_2 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_2) \triangleright A_2 \xrightarrow{\hat{\mu}} (\mathcal{U}'_2 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_2) \triangleright A'_2$ and $(\mathcal{U}'_1 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_1) \triangleright A'_1 \approx (\mathcal{U}'_2 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_2) \triangleright A'_2$. This implies that $(\mathcal{U}_2; \mathcal{P}_2) \triangleright (v \bar{a}^r : \bar{R})(A_2 \parallel B) \xrightarrow{\hat{\mu}} (\mathcal{U}'_2; \mathcal{P}_2) \triangleright (v \bar{a}^r : \bar{R})(A'_2 \parallel B)$ and $((\mathcal{U}'_1; \mathcal{P}_1) \triangleright (v \bar{a}^r : \bar{R})(A'_1 \parallel B), (\mathcal{U}'_2; \mathcal{P}_2) \triangleright (v \bar{a}^r : \bar{R})(A'_2 \parallel B)) \in \mathfrak{R}$. Indeed, if $\text{BND}(\mu) = \emptyset$ then $\mathcal{U}'_i = \mathcal{U}_i$ and the claim is trivial. Otherwise, $\mu = - b^s : S$ for some S and $\mathcal{U}'_i = \mathcal{U}_i \uplus \{b^s : S\}$. By hypothesis, b^s is fresh for B and, thus, by Proposition B.3(3) it holds that $((\mathcal{U}'_1 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_1) \triangleright B \approx (\mathcal{U}'_2 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_2) \triangleright B$.
2. $(\mathcal{U}_1 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_1) \triangleright B \xrightarrow{\mu} (\mathcal{U}'_1 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_1) \triangleright B'$ for $\text{BND}(\mu) \cap \text{FR}(A_1) = \emptyset$. By Proposition B.3(1), the transition can be inferred also upon renaming of the (possible) bound channel in μ into a channel fresh also for A_2 . The proof then proceeds like in the case 1.
3. $\mu = \bar{a}^r b^s : S$ and $(\mathcal{U}_1 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_1) \triangleright A_1 \xrightarrow{\bar{a}^r b^s} (\mathcal{U}_1 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_1) \triangleright A'_1$ for $b^s : S \in \bar{a}^r : \bar{R}$ and $b^s \notin \text{FR}(B)$. By letting $\bar{c}^t : \bar{R}'$ be the set $\bar{a}^r : \bar{R} \setminus \{b^s : S\}$, we have that $\bar{A}_1 \triangleq (v \bar{c}^t : \bar{R}')(A'_1 \parallel B)$ and $\mathcal{U}'_1 = \mathcal{U} \uplus \{b^s : S\}$. By hypothesis, $(\mathcal{U}_2 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_2) \triangleright A_2 \xrightarrow{\bar{a}^r b^s} (\mathcal{U}_2 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_2) \triangleright A'_2$ and $(\mathcal{U}_1 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_1) \triangleright A'_1 \approx (\mathcal{U}_2 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_2) \triangleright A'_2$. This implies that $(\mathcal{U}_2; \mathcal{P}_2) \triangleright (v \bar{a}^r : \bar{R})(A_2 \parallel B) \xrightarrow{\bar{a}^r b^s : S} (\mathcal{U}'_2; \mathcal{P}_2) \triangleright (v \bar{c}^t : \bar{R}')(A'_2 \parallel B)$, for $\mathcal{U}'_2 = \mathcal{U}_2 \uplus \{b^s : S\}$, and $((\mathcal{U}'_1; \mathcal{P}_1) \triangleright (v \bar{c}^t : \bar{R}')(A'_1 \parallel B), (\mathcal{U}'_2; \mathcal{P}_2) \triangleright (v \bar{c}^t : \bar{R}')(A'_2 \parallel B)) \in \mathfrak{R}$.

4. The case when $\mu = \bar{a}^r b^s : S$ and $(\mathcal{U}_1 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_1) \triangleright B \xrightarrow{\bar{a}^r b^s} (\mathcal{U}_1 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_1) \triangleright B'$ for $b^s : S \in \bar{a}^r : \bar{R}$ is similar to case 3.
5. $\mu = \tau$, $(\mathcal{U}_1 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_1) \triangleright A_1 \xrightarrow{a^r b^s} (\mathcal{U}_1 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_1) \triangleright A'_1$ and $(\mathcal{U}_1 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_1) \triangleright B \xrightarrow{\bar{a}^r b^s} (\mathcal{U}_1 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_1) \triangleright B'$. Thus, $\bar{A}_1 \triangleq (\nu \bar{a}^r : \bar{R})(A'_1 \parallel B')$. By hypothesis, $(\mathcal{U}_2 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_2) \triangleright A_2 \xrightarrow{a^r b^s} (\mathcal{U}_2 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_2) \triangleright A'_2$ and $(\mathcal{U}_1 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_1) \triangleright A'_1 \approx (\mathcal{U}_2 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_2) \triangleright A'_2$. Similarly, $(\mathcal{U}_2 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_2) \triangleright B \xrightarrow{\bar{a}^r b^s} (\mathcal{U}_2 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_2) \triangleright B'$ and $(\mathcal{U}_1 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_1) \triangleright B' \approx (\mathcal{U}_2 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_2) \triangleright B'$. Hence, $(\mathcal{U}_2; \mathcal{P}_2) \triangleright (\nu \bar{a}^r : \bar{R})(A_2 \parallel B) \Rightarrow (\mathcal{U}_2; \mathcal{P}_2) \triangleright (\nu \bar{a}^r : \bar{R})(A'_2 \parallel B')$ and $((\mathcal{U}_1; \mathcal{P}_1) \triangleright (\nu \bar{a}^r : \bar{R})(A'_1 \parallel B'), (\mathcal{U}_2; \mathcal{P}_2) \triangleright (\nu \bar{a}^r : \bar{R})(A'_2 \parallel B')) \in \mathfrak{X}$.
6. The case for $\mu = \tau$, $(\mathcal{U}_1 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_1) \triangleright A_1 \xrightarrow{\bar{a}^r b^s} (\mathcal{U}_1 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_1) \triangleright A'_1$ and $(\mathcal{U}_1 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_1) \triangleright B \xrightarrow{a^r b^s} (\mathcal{U}_1 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_1) \triangleright B'$ is similar to case 5.
7. $\mu = \tau$, $(\mathcal{U}_1 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_1) \triangleright A_1 \xrightarrow{a^r b^s : S} (\mathcal{U}_1 \uplus \bar{a}^r : \bar{R} \uplus \{b^s : S\}; \mathcal{P}_1) \triangleright A'_1$ and $(\mathcal{U}_1 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_1) \triangleright B \xrightarrow{\bar{a}^r b^s : S} (\mathcal{U}_1 \uplus \bar{a}^r : \bar{R} \uplus \{b^s : S\}; \mathcal{P}_1) \triangleright B'$. Thus, $\bar{A}_1 \triangleq (\nu \bar{a}^r : \bar{R}, b^s : S)(A'_1 \parallel B')$. By hypothesis, $(\mathcal{U}_2 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_2) \triangleright A_2 \xrightarrow{a^r b^s : S} (\mathcal{U}_2 \uplus \bar{a}^r : \bar{R} \uplus \{b^s : S\}; \mathcal{P}_2) \triangleright A'_2$ and $(\mathcal{U}_1 \uplus \bar{a}^r : \bar{R} \uplus \{b^s : S\}; \mathcal{P}_1) \triangleright A'_1 \approx (\mathcal{U}_2 \uplus \bar{a}^r : \bar{R} \uplus \{b^s : S\}; \mathcal{P}_2) \triangleright A'_2$. Similarly, $(\mathcal{U}_2 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_2) \triangleright B \xrightarrow{\bar{a}^r b^s : S} (\mathcal{U}_2 \uplus \bar{a}^r : \bar{R} \uplus \{b^s : S\}; \mathcal{P}_2) \triangleright B'$ and $(\mathcal{U}_1 \uplus \bar{a}^r : \bar{R} \uplus \{b^s : S\}; \mathcal{P}_1) \triangleright B' \approx (\mathcal{U}_2 \uplus \bar{a}^r : \bar{R} \uplus \{b^s : S\}; \mathcal{P}_2) \triangleright B'$. Hence, $(\mathcal{U}_2; \mathcal{P}_2) \triangleright (\nu \bar{a}^r : \bar{R})(A_2 \parallel B) \Rightarrow (\mathcal{U}_2; \mathcal{P}_2) \triangleright (\nu \bar{a}^r : \bar{R}, b^s : S)(A'_2 \parallel B')$ and $((\mathcal{U}_1; \mathcal{P}_1) \triangleright (\nu \bar{a}^r : \bar{R}, b^s : S)(A'_1 \parallel B'), (\mathcal{U}_2; \mathcal{P}_2) \triangleright (\nu \bar{a}^r : \bar{R}, b^s : S)(A'_2 \parallel B')) \in \mathfrak{X}$.
8. The case for $\mu = \tau$, $(\mathcal{U}_1 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_1) \triangleright A_1 \xrightarrow{\bar{a}^r b^s : S} (\mathcal{U}_1 \uplus \bar{a}^r : \bar{R} \uplus \{b^s : S\}; \mathcal{P}_1) \triangleright A'_1$ and $(\mathcal{U}_1 \uplus \bar{a}^r : \bar{R}; \mathcal{P}_1) \triangleright B \xrightarrow{a^r b^s : S} (\mathcal{U}_1 \uplus \bar{a}^r : \bar{R} \uplus \{b^s : S\}; \mathcal{P}_1) \triangleright B'$ is similar to case 7. \square

We now turn to the main result of this section, i.e. that \approx is a sound proof technique for \cong . However, this result only holds for well-typed configurations, i.e. configurations $\mathcal{S} \triangleright A$ such that A is well-typed in \mathcal{S} . Indeed, as already said, ill-typed systems are not considered in the definition of barbed congruence. Given a typing environment Γ , we let \mathcal{U}_Γ be the roles-to-users assignment extracted from Γ , that is the least assignment such that, for any association $r : \rho[\bar{a} : \bar{R}(T)]$ in Γ , it holds that $\mathcal{U}_\Gamma(r) = \rho$ and $\mathcal{U}_\Gamma(a^r) = R$, for any $a : R(T) \in \bar{a} : \bar{R}(T)$. For notational convenience, we write $\mathcal{S} \triangleright A \approx \mathcal{S} \triangleright B$ as $\mathcal{S} \triangleright A \approx B$.

Theorem B.4 (Soundness of \approx). *Let $\mathcal{S} = (\mathcal{U}; \mathcal{P})$, $\Gamma \vdash^{\mathcal{S}} A$ and $\Gamma \vdash^{\mathcal{S}} B$. If $(\mathcal{U}_\Gamma; \mathcal{P}) \triangleright A \approx B$, then $\Gamma \models^{\mathcal{S}} A \cong B$.*

The proof relies on the following lemmata.

Proposition B.5. *\approx is an equivalence relation. Moreover, $A \equiv B$ implies that $\mathcal{S} \triangleright A \approx B$, for all RBAC schemata.*

Proof. The first part trivially follows from Definition B.1. For the second part, we show that relation

$$\mathfrak{X} \triangleq \{(\mathcal{S} \triangleright A, \mathcal{S} \triangleright B) : A \equiv B\}$$

is a bisimulation. Let $S \triangleright A \xrightarrow{\mu} S' \triangleright A'$. By mutual induction on the depth of the inferences for $A \equiv B$ and $B \equiv A$, we shall prove that $S \triangleright B \xrightarrow{\mu} S' \triangleright B'$ and $A' \equiv B'$; this suffices to prove that $(S' \triangleright A', S' \triangleright B') \in \mathfrak{K}$. The only tricky base case is when $A \triangleq (va^r : R)A_1 \parallel B_1$ and $B \triangleq (va^r : R)(A_1 \parallel B_1)$, for $a^r \notin \text{Fr}(B_1)$. Like in the proof of Theorem B.2 we have to examine eight cases to distinguish how $(va^r : R)A_1 \parallel B_1$ has generated μ . We only consider one of them; the other ones are similar. Let S be $(\mathcal{U}; \mathcal{P})$ and $A_1 \xrightarrow{\bar{b}^s a^r} A'_1$; thus, $\mu = \bar{b}^s a^r : R$, $S' = (\mathcal{U} \uplus \{a^r : R\}; \mathcal{P})$ and $A' \triangleq A'_1 \parallel B_1$. Then, easily, $S \triangleright B \xrightarrow{\bar{b}^s a^r : R} S' \triangleright A'_1 \parallel B_1$; this suffices to conclude. For the inductive step, symmetry and transitivity easily hold; thus, we only consider context closure, i.e. $A \triangleq C[A_1]$, $B \triangleq C[B_1]$ and $A_1 \equiv B_1$. We now proceed by induction over the structure of the context. The base step (when the context is $[\cdot]$) is trivial. For the inductive step, we reason by case analysis on the last operator used to build the context:

- $C[\cdot] \triangleq \mathcal{D}[\cdot] \parallel \bar{A}$. Then, $S \triangleright \mathcal{D}[A_1] \parallel \bar{A} \xrightarrow{\mu} S' \triangleright A'$ could have been inferred in eight ways (like in the proof of Theorem B.2). When μ has been originated by \bar{A} , then we trivially have that $S \triangleright B \xrightarrow{\mu} S' \triangleright B'$ and $A' \equiv B'$. In all the other cases, it holds that $S \triangleright \mathcal{D}[A_1] \xrightarrow{\mu'} S'' \triangleright A''$, for some μ' , S'' and A'' . Since context $\mathcal{D}[\cdot]$ is smaller than $C[\cdot]$, we can apply the induction hypothesis and obtain that $S \triangleright \mathcal{D}[B_1] \xrightarrow{\mu'} S'' \triangleright B''$, for some $B'' \equiv A''$. Then, by mimicking the steps leading from $S \triangleright \mathcal{D}[A_1] \xrightarrow{\mu'} S'' \triangleright A''$ to $S \triangleright \mathcal{D}[A_1] \parallel \bar{A} \xrightarrow{\mu} S' \triangleright A'$, we obtain that $S \triangleright \mathcal{D}[B_1] \parallel \bar{A} \xrightarrow{\mu} S' \triangleright B'$, for some $B' \equiv A'$.
- $C[\cdot] \triangleq (va^r : R)\mathcal{D}[\cdot]$. Then, by definition of the LTS, the last rule used to infer $S \triangleright (va^r : R)\mathcal{D}[A_1] \xrightarrow{\mu} S' \triangleright A'$ can only be (LTS-RES) or (LTS-OPEN). In both cases, we have that $S'' \triangleright \mathcal{D}[A_1] \xrightarrow{\mu'} S'' \triangleright A''$, for some μ' , S'' and A'' . The proof then proceeds like in the previous case. \square

Lemma B.6. *Let $S = (\mathcal{U}; \mathcal{P})$. Then, the following facts hold.*

1. If $S \triangleright A \xrightarrow{a^n} S \triangleright A'$, then $A \equiv (v\bar{a}^r : \bar{R})(B \parallel r\|a^r(x).P\|_\rho)$ for $\{a^r, n\} \cap \bar{a}^r = \emptyset$, $n \in \text{dom}(\mathcal{U})$ and $A' \equiv (v\bar{a}^r : \bar{R})(B \parallel r\|P[\eta/x]\|_\rho)$. The converse holds true whenever $R^2 \in \mathcal{P}(\rho)$, for $\mathcal{U}(a^r) = R$.
2. If $S \triangleright A \xrightarrow{a^n R} S' \triangleright A'$, then $A \equiv (v\bar{a}^r : \bar{R})(B \parallel r\|a^r(x).P\|_\rho)$ for $a^r \notin \bar{a}^r$, $n \notin \bar{a}^r \cup \text{dom}(\mathcal{U})$, $S' = (\mathcal{U} \uplus \{n : R\}; \mathcal{P})$ and $A' \equiv (v\bar{a}^r : \bar{R})(B \parallel r\|P[\eta/x]\|_\rho)$. The converse holds true whenever $R^2 \in \mathcal{P}(\rho)$, for $\mathcal{U}(a^r) = R$.
3. If $S \triangleright A \xrightarrow{\bar{a}^n} S \triangleright A'$, then $A \equiv (v\bar{a}^r : \bar{R})(B \parallel s\|a^r\langle n \rangle.P\|_\rho)$ for $\{a^r, n\} \cap \bar{a}^r = \emptyset$ and $A' \equiv (v\bar{a}^r : \bar{R})(B \parallel s\|P\|_\rho)$. The converse holds true whenever $R^1 \in \mathcal{P}(\rho)$, for $\mathcal{U}(a^r) = R$.
4. If $S \triangleright A \xrightarrow{\bar{a}^n R} S' \triangleright A'$, then $A \equiv (vb^s : R)(v\bar{a}^r : \bar{R})(B \parallel t\|a^r\langle b^s \rangle.P\|_\rho)$, $n = b^s$, $a^r \notin \bar{a}^r \cup \{b^s\}$, $S' = (\mathcal{U} \uplus \{n : R\}; \mathcal{P})$ and $A' \equiv (v\bar{a}^r : \bar{R})(B \parallel t\|P\|_\rho)$. The converse holds true whenever $R^1 \in \mathcal{P}(\rho)$, for $\mathcal{U}(a^r) = R$.

Proof. The four points are all proved in the same way; we just sketch point 4. The first implication is done by induction over the depth of the inference for $\xrightarrow{\bar{a}^n R}$; the second implication is based on the definition of the LTS and is very simple. The base step requires the application of rules (LTS-OUTPUT) and (LTS-OPEN), in sequence; thus, $A \triangleq (vb^s : R)t\|a^r\langle b^s \rangle.P\|_\rho$, $a^r \neq b^s$, $A' \triangleq t\|P\|_\rho$ and

$S' = (\mathcal{U} \uplus \{(vb^s : R)\}; \mathcal{P})$. For the inductive step, we reason by case analysis on the last rule used in the inference. We only explicitly consider (LTS-PAR); the case for (LTS-RES) and (LTS-EXT) is similar. By hypothesis, $A \triangleq A_1 \parallel A_2$, $\mathcal{S} \triangleright A_1 \xrightarrow{\bar{a}^r n : R} S' \triangleright A'_1$ and $n \notin \text{FR}(A_2)$. By induction hypothesis, $A_1 \equiv (vb^s : R)(v\bar{a}^r : \bar{R})(B \parallel t \parallel a^r \langle b^s \rangle . P \parallel_\rho)$, $n = b^s$, $a^r \notin \bar{a}^r \cup \{b^s\}$, $S' = (\mathcal{U} \uplus \{n : R\}; \mathcal{P})$ and $A'_1 \equiv (v\bar{a}^r : \bar{R})(B \parallel t \parallel P \parallel_\rho)$. Since $a^r \notin \text{FR}(A_2)$, we can conclude, by rule (S-EXT), that $A \equiv (vb^s : R)(v\bar{a}^r : \bar{R})(B \parallel A_2 \parallel t \parallel a^r \langle b^s \rangle . P \parallel_\rho)$ and $A' \equiv (v\bar{a}^r : \bar{R})(B \parallel A_2 \parallel t \parallel P \parallel_\rho)$. This suffices to conclude. \square

Lemma B.7 (Weakening for \approx). *If $(\mathcal{U}; \mathcal{P}) \triangleright A \approx B$ and A and B are well-typed in $(\mathcal{U}; \mathcal{P})$, then $(\mathcal{U} \uplus \mathcal{U}'; \mathcal{P} \cup \mathcal{P}') \triangleright A \approx B$ for all \mathcal{U}' and \mathcal{P}' .*

Proof. We have to prove that the relation

$$\mathfrak{R} \triangleq \{((\mathcal{U} \uplus \mathcal{U}'; \mathcal{P} \cup \mathcal{P}') \triangleright A, (\mathcal{U} \uplus \mathcal{U}'; \mathcal{P} \cup \mathcal{P}') \triangleright B) : (\mathcal{U}; \mathcal{P}) \triangleright A \approx B\}$$

is a bisimulation. Let $(\mathcal{U} \uplus \mathcal{U}'; \mathcal{P} \cup \mathcal{P}') \triangleright A \xrightarrow{\mu} (\mathcal{U} \uplus \mathcal{U}' \uplus \mathcal{U}''; \mathcal{P} \cup \mathcal{P}') \triangleright A'$. We now distinguish on μ .

1. $\mu = a^r n$. In this case, $\mathcal{U}'' = \emptyset$ and $n \in \text{dom}(\mathcal{U} \uplus \mathcal{U}')$. We now distinguish whether $n \in \text{dom}(\mathcal{U})$ or $n \in \text{dom}(\mathcal{U}')$.

(a) Let $n \in \text{dom}(\mathcal{U})$; thus, because of well-typedness, $(\mathcal{U}; \mathcal{P}) \triangleright A \xrightarrow{\mu} (\mathcal{U}; \mathcal{P}) \triangleright A'$. By hypothesis, $(\mathcal{U}; \mathcal{P}) \triangleright B \xrightarrow{\mu} (\mathcal{U}; \mathcal{P}) \triangleright B'$ and $(\mathcal{U}; \mathcal{P}) \triangleright A' \approx B'$. Then, $(\mathcal{U} \uplus \mathcal{U}'; \mathcal{P} \cup \mathcal{P}') \triangleright B \xrightarrow{\mu} (\mathcal{U} \uplus \mathcal{U}'; \mathcal{P} \cup \mathcal{P}') \triangleright B'$ and $((\mathcal{U} \uplus \mathcal{U}'; \mathcal{P} \cup \mathcal{P}') \triangleright A', (\mathcal{U} \uplus \mathcal{U}'; \mathcal{P} \cup \mathcal{P}') \triangleright B') \in \mathfrak{R}$, as required.

(b) Let $n \in \text{dom}(\mathcal{U}')$; thus, $(\mathcal{U}; \mathcal{P}) \triangleright A \xrightarrow{a^r n : R} (\mathcal{U} \uplus \{n : R\}; \mathcal{P}) \triangleright A'$ for any R ; in particular, we can choose $R \in \mathcal{U}'(n)$ and let $\bar{\mathcal{U}} = \mathcal{U}' \setminus \{(n, R)\}$. By hypothesis, $(\mathcal{U}; \mathcal{P}) \triangleright B \xrightarrow{a^r n : R} (\mathcal{U} \uplus \{n : R\}; \mathcal{P}) \triangleright B'$ and $(\mathcal{U} \uplus \{n : R\}; \mathcal{P}) \triangleright A' \approx B'$. Then, $(\mathcal{U} \uplus \mathcal{U}'; \mathcal{P} \cup \mathcal{P}') \triangleright B \xrightarrow{\mu} (\mathcal{U} \uplus \mathcal{U}' \uplus \mathcal{U}''; \mathcal{P} \cup \mathcal{P}') \triangleright B'$ and $((\mathcal{U} \uplus \mathcal{U}'; \mathcal{P} \cup \mathcal{P}') \triangleright A', (\mathcal{U} \uplus \mathcal{U}'; \mathcal{P} \cup \mathcal{P}') \triangleright B') \in \mathfrak{R}$, because $\mathcal{U} \uplus \mathcal{U}' = \mathcal{U} \uplus \{n : R\} \uplus \bar{\mathcal{U}}$.

2. $\mu = a^r n : R$. In this case, $\mathcal{U}'' = \{n : R\}$ since $n \notin \text{dom}(\mathcal{U} \uplus \mathcal{U}')$. This implies that $(\mathcal{U}; \mathcal{P}) \triangleright A \xrightarrow{\mu} (\mathcal{U} \uplus \mathcal{U}''; \mathcal{P}) \triangleright A'$. By hypothesis, $(\mathcal{U}; \mathcal{P}) \triangleright B \xrightarrow{\mu} (\mathcal{U} \uplus \mathcal{U}''; \mathcal{P}) \triangleright B'$ and $(\mathcal{U} \uplus \mathcal{U}''; \mathcal{P}) \triangleright A' \approx B'$. Then, $(\mathcal{U} \uplus \mathcal{U}'; \mathcal{P} \cup \mathcal{P}') \triangleright B \xrightarrow{\mu} (\mathcal{U} \uplus \mathcal{U}' \uplus \mathcal{U}''; \mathcal{P} \cup \mathcal{P}') \triangleright B'$ and $((\mathcal{U} \uplus \mathcal{U}' \uplus \mathcal{U}''; \mathcal{P} \cup \mathcal{P}') \triangleright A', (\mathcal{U} \uplus \mathcal{U}' \uplus \mathcal{U}''; \mathcal{P} \cup \mathcal{P}') \triangleright B') \in \mathfrak{R}$.

3. $\mu = \bar{a}^r n$. This case is dealt with similarly to case 1.

4. $\mu = \bar{a}^r n : R$. This case is dealt with similarly to case 2.

5. $\mu = \tau$. This case is dealt with similarly to case 1(a). \square

Proof of Theorem B.4 It suffices to prove that the relation

$$\mathfrak{R} \triangleq \{\Gamma \models^S (A, B) : \Gamma \vdash^S A \wedge \Gamma \vdash^S B \wedge (\mathcal{U}_\Gamma; \mathcal{P}) \triangleright A \approx B\}$$

is barb preserving, reduction closed and contextual.

1. Let $A \downarrow a^r$. By Definition 3.1 and well-typedness, it holds that $(\mathcal{U}_\Gamma; \mathcal{P}) \triangleright A \xrightarrow{a^r}$, by exploiting Proposition B.6(1)/(2)); then, $(\mathcal{U}_\Gamma; \mathcal{P}) \triangleright B \Rightarrow (\mathcal{U}_\Gamma; \mathcal{P}) \triangleright B' \xrightarrow{a^r}$. Again by Proposition B.6(1)/(2), Proposition B.1 and Definition 3.1, $B \Downarrow d$. The case for $A \downarrow \bar{a}^r$ is
2. Let $A \mapsto A'$. By Proposition B.1 and well-typedness, this implies that $(\mathcal{U}_\Gamma; \mathcal{P}) \triangleright A \xrightarrow{\tau} (\mathcal{U}_\Gamma; \mathcal{P}) \triangleright A'$. Thus, $(\mathcal{U}_\Gamma; \mathcal{P}) \triangleright B \Rightarrow (\mathcal{U}_\Gamma; \mathcal{P}) \triangleright B'$ and $(\mathcal{U}_\Gamma; \mathcal{P}) \triangleright A' \approx B'$. Again by Proposition B.1, $B \mapsto^* B'$ and $\Gamma \models^{\mathcal{P}} A' \mathfrak{R} B'$. Indeed, by Theorem 2.1, it holds that $\Gamma \vdash^{\mathcal{P}} A'$ and $\Gamma \vdash^{\mathcal{P}} B'$. Moreover, by using Proposition B.5, we have that $(\mathcal{U}_\Gamma; \mathcal{P}) \triangleright A' \approx A''$ and, thus, $(\mathcal{U}_\Gamma; \mathcal{P}) \triangleright A' \approx B'$.
3. We pick up $\Gamma \models^S A \mathfrak{R} B$ and analyse the three clauses defining the contextuality property.
 - (a) Let \mathcal{P}' be a permissions-to-roles assignment and $\tilde{n} : \tilde{T}$ be such that $\tilde{n} \cap \text{dom}(\mathcal{U}) = \emptyset$; since Γ respects \mathcal{U} , this implies that $\Gamma, \tilde{n} : \tilde{T}$ is defined. By Lemma A.2, we know that $\Gamma, \tilde{n} : \tilde{T} \vdash^{(\mathcal{U} \uplus \tilde{n} : \tilde{T}; \mathcal{P} \cup \mathcal{P}')} A$ and $\Gamma, \tilde{n} : \tilde{T} \vdash^{(\mathcal{U} \uplus \tilde{n} : \tilde{T}; \mathcal{P} \cup \mathcal{P}')} B$. Moreover, we let $\mathcal{U}_{\tilde{n} : \tilde{T}}$ to be the roles-to-users assignment such that $\mathcal{U}_{\tilde{n} : \tilde{T}}(r) = \rho$, whenever $r : \rho[\bar{a} : \bar{C}] \in \tilde{n} : \tilde{T}$, and $\mathcal{U}_{\tilde{n} : \tilde{T}}(a^r) = R$, whenever $a^r : R(T) \in \tilde{n} : \tilde{T}$. It is easy to check that $\mathcal{U}_{\Gamma, \tilde{n} : \tilde{T}} = \mathcal{U}_\Gamma \uplus \mathcal{U}_{\tilde{n} : \tilde{T}}$ (indeed, $\Gamma, \tilde{n} : \tilde{T}$ is defined if and only if names in \tilde{n} do not occur in the domain of Γ ; thus, \mathcal{U}_Γ and $\mathcal{U}_{\tilde{n} : \tilde{T}}$ have disjoint domains, and their union coincides with $\mathcal{U}_{\Gamma, \tilde{n} : \tilde{T}}$). Thus, by Lemma B.7, $(\mathcal{U}_{\Gamma, \tilde{n} : \tilde{T}}; \mathcal{P} \cup \mathcal{P}') \triangleright A \approx B$. This suffices to conclude that $\Gamma, \tilde{n} : \tilde{T} \models^{(\mathcal{U} \cup \tilde{n} : \tilde{T}; \mathcal{P} \cup \mathcal{P}')} A \mathfrak{R} B$.
 - (b) Let \bar{A} be a system such that $\Gamma \vdash^S \bar{A}$. By Theorem B.2(1), we can state that $(\mathcal{U}_\Gamma; \mathcal{P}) \triangleright A \parallel \bar{A} \approx B \parallel \bar{A}$. Moreover, by rule (T-SYSPAR), it holds that $\Gamma \vdash^S A \parallel \bar{A}$ and $\Gamma \vdash^S B \parallel \bar{A}$. Thus, $\Gamma \models^S A \parallel \bar{A} \mathfrak{R} B \parallel \bar{A}$, as required.
 - (c) Let $\Gamma = \Gamma', a^r : R(T)$ and $\mathcal{U} = \mathcal{U}' \uplus \{a^r : R(T)\}$. It is easy to check that $\mathcal{U}_\Gamma = \mathcal{U}_{\Gamma'} \uplus a^r : R$ and, thus, $(\mathcal{U}_{\Gamma'} \uplus a^r : R; \mathcal{P}) \triangleright A \approx B$. By Theorem B.2(2), this implies that $(\mathcal{U}_{\Gamma'}; \mathcal{P}) \triangleright (va^r : R)A \approx (va^r : R)B$; moreover, by rule (T-SYSRES), $\Gamma' \vdash^{(\mathcal{U}'; \mathcal{P})} (va^r : R)A$ and $\Gamma' \vdash^{(\mathcal{U}'; \mathcal{P})} (va^r : R)B$. Thus, $\Gamma' \models^{(\mathcal{U}'; \mathcal{P})} (va^r : R)A \mathfrak{R} (va^r : R)B$, as required. \square

We remark that \approx is used as a proof-technique for barbed congruence. Indeed, while the former is easy to use, the latter is very hard to handle because of the contextual closure requirement. This suffices for our purpose in the present paper, whose intention is to present the calculus and lay out its main properties. Nevertheless, for theoretical reasons, it is often important to know whether \approx is a complete characterisation of \cong . This is a laborious question to answer. We leave it as future work to follow well-known paths towards the answer (as, e.g. [12, 18]) to prove the converse of Theorem B.4, i.e. that bisimilarity is complete for barbed congruence.

To conclude, we now briefly discuss some possible use of the bisimulation, apart from proving barbed congruence. Mainly, its distinctive features are the possibility of relating ill-typed systems and/or systems under different schemata. For example, by letting α to range over action prefixes (i.e. inputs/outputs and **role/yield**), it holds that

$$\mathcal{S} \triangleright r \parallel \alpha.P \parallel_\rho \approx \mathbf{0}$$

whenever α is not legal for a session $r \parallel \cdot \parallel_\rho$ with respect to \mathcal{S} , that is, if the premises of rules (LTS-ROLE), (LTS-YIELD), (LTS-K-INPUT), (LTS-F-INPUT) and (LTS-OUTPUT) are not satisfied. This law stresses that LTS and types both enforce the same requirements (compare the run-time checks of

the LTS with the run-time errors in Table 5 and, consequently, the results in Proposition B.1). As a consequence, the following law differentiates our language from the π -calculus. Indeed, it holds that

$$\mathcal{S} \triangleright (va^r : R)(r \parallel a(x).P \parallel s \parallel a^r \langle n \rangle . Q) \approx \mathbf{0}$$

whenever $R^? \notin \mathcal{P}(\rho)$ or $R^! \notin \mathcal{P}(\rho')$.

Finally, we can use the bisimulation to find a ‘minimal schema’ to run a given system without altering its functionalities. Let A be a system well-typed in a RBAC schema \mathcal{S} . Potentially, there are many schemata under which the system can run correctly; thus, it seems reasonable to look for a ‘minimal’ such. According to the metrics chosen, several properties can be associated to this element. For example, if the metrics is the size of the schema (seen as a pair of sets), the minimal element would be one of the smallest; thus, its storage and handling would be cheaper. We now define the set of configurations whose second component is A as follows:

$$CONF_A = \{\mathcal{S}' \triangleright A : \mathcal{S}' \text{ is a RBAC schema}\}$$

We now partition $CONF_A$ with respect to \approx and consider the equivalence class containing $\mathcal{S} \triangleright A$, called $CONF_A^{\mathcal{S}}$. By fixing a metrics over schemata, the minimal schema to run the system A will be a minimal element of $CONF_A^{\mathcal{S}}$. Indeed, such an element behaves like $\mathcal{S} \triangleright A$, because they both belong to the same equivalence class $CONF_A^{\mathcal{S}}$, but its schema is smaller, as it is a minimal element of $CONF_A^{\mathcal{S}}$. Clearly, the existence of such a minimal element and the way in which it is chosen depend on the chosen metrics. Possible metrics could be based on the memory required to store the schema, on the number of rôles used to define the schema, on the weight of the permissions associated with some users (once assumed a weight function to discriminate sensible permissions from common ones), on the average number of permissions associated with each rôle, and so on.

References

- [1] G.-J. Ahn and R. Sandhu. Rôle-based authorisation constraints specification. *ACM Transactions on Information and System Security*, 3(4):207–226, 2000.
- [2] R. Amadio and L. Cardelli. Subtyping recursive types. *ACM Transactions on Programming Languages and Systems*, 15(4):575–631, 1993.
- [3] E. Bertino, B. Catania, E. Ferrari, and P. Perlasca. A logical framework for reasoning about access control models. In *Proc. of 6th SACMAT*, pages 41–52. ACM Press, 2001.
- [4] C. Braghin, D. Gorla, and V. Sassone. A distributed calculus for rôle-based access control. In *Proc. of 17th Computer Security Foundations Workshop (CSFW'04)*, pages 48–60. IEEE Computer Society, 2004.
- [5] C. Braghin, D. Gorla, and V. Sassone. A distributed calculus for rôle-based access control. Technical Report 08/2004, Dip. di Informatica, Univ. di Roma ‘La Sapienza’, 2004.
- [6] L. Cardelli, G. Ghelli, and A. D. Gordon. Secrecy and group creation. *Information and Computation*, 196(2):127–155, 2005.
- [7] L. Cardelli and A. D. Gordon. Mobile ambients. *Theoretical Computer Science*, 240(1):177–213, 2000.
- [8] F. Cardone and M. Coppo. Two extensions of Curry’s type inference system. In *Logic and Computer Science*, pages 19–75. Academic Press, 1990.
- [9] R. De Nicola, G. Ferrari, and R. Pugliese. KLAIM: a Kernel Language for Agents Interaction and Mobility. *IEEE Transactions on Software Engineering*, 24(5):315–330, 1998.

- [10] D. Ferraiolo and D. Kuhn. Rôle-based access control. In *Proc. of the NIST-NSA National Computer Security Conference*, pages 554–563, 1992.
- [11] D. Ferraiolo, R. Sandhu, S. Gavrila, D. Kuhn, and R. Chandramouli. Proposed NIST standard for rôle-based access control. *ACM Transactions on Information and System Security*, 4(3):224–274, 2001.
- [12] C. Fournet and C. Laneve. Bisimulations for the join-calculus. *Theoretical Computer Science*, 266(1-2):569–603, 2001.
- [13] M. Hennessy and J. Rathke. Typed behavioural equivalences for processes in the presence of subtyping. *Mathematical Structures in Computer Science*, 14(5):651–684, 2004.
- [14] M. Hennessy and J. Riely. Resource access control in systems of mobile agents. *Information and Computation*, 173:82–120, 2002.
- [15] K. Honda and N. Yoshida. On reduction-based process semantics. *Theoretical Computer Science*, 152(2):437–486, 1995.
- [16] M. Koch, L. Mancini, and F. Parisi-Presicce. A formal model for rôle-based access control using graph transformation. In *Proc. of 5th ESORICS*, volume 1895 of *LNCS*, pages 122–139. Springer, 2000.
- [17] M. Koch, L. Mancini, and F. Parisi-Presicce. Decidability of safety in graph-based models for access control. In *Proc. of 7th ESORICS*, volume 2502 of *LNCS*, pages 229–243. Springer, 2002.
- [18] M. Merro and D. Sangiorgi. On asynchrony in name-passing calculi. *Mathematical Structures in Computer Science*, 14(5):715–767, 2004.
- [19] R. Milner. The polyadic π -calculus: a tutorial. In *Logic and Algebra of Specification*, volume 94 of *Series F: Computer and System Sciences*. NATO Advanced Study Inst., 1993.
- [20] B. C. Pierce. *Types and programming languages*. MIT Press, 2002.
- [21] B. C. Pierce and D. Sangiorgi. Typing and subtyping for mobile processes. *Mathematical Structures in Computer Science*, 6(5):409–454, 1996.
- [22] J. H. Saltzer and M. D. Schroeder. The protection of information in computer systems. *Proceedings of the IEEE*, 63(9):1278–1308, 1975.
- [23] R. Sandhu, E. Coyne, H. Feinstein, and C. Youman. Rôle-based access control models. *IEEE Computer*, 29(2):38–47, 1996.
- [24] D. Sangiorgi and D. Walker. *The π -calculus: a theory of mobile processes*. Cambridge University Press, 2001.
- [25] A. Schaad and J. Moffett. A lightweight approach to specification and analysis of rôle-based access control extensions. In *Proc. of 7th SACMAT*, pages 13–22. ACM Press, 2002.
- [26] F. B. Schneider. Least privilege and more. *Security and Privacy*, 1(3):55–59, 2003.