Regulating data exchange in service oriented applications ⋆

Alessandro Lapadula, Rosario Pugliese and Francesco Tiezzi

Dipartimento di Sistemi e Informatica Università degli Studi di Firenze

Abstract. We define a type system for COWS, a formalism for specifying and combining services, while modelling their dynamic behaviour. Our types permit to express policies constraining data exchanges in terms of sets of service partner names attachable to each single datum. Service programmers explicitly write only the annotations necessary to specify the wanted policies for communicable data, while a type inference system (statically) derives the minimal additional annotations that ensure consistency of services initial configuration. Then, the language dynamic semantics only performs very simple checks to authorize or block communication. We prove that the type system and the operational semantics are sound. As a consequence, we have the following data protection property: services always comply with the policies regulating the exchange of data among interacting services. We illustrate our approach through a simplified but realistic scenario for a service-based electronic marketplace, and an automotive scenario studied in the SENSORIA project.

1 Introduction

Service-oriented computing (SOC) is an emerging paradigm for developing loosely coupled, interoperable, evolvable applications which exploits the pervasiveness of the Internet and its related technologies. SOC systems deliver application functionality as services to either end-user applications or other services. Current software engineering technologies for SOC, however, remain at the descriptive level and do not support analytical tools for checking that SOC applications enjoy desirable properties and do not manifest unexpected behaviors. To reason about and guarantee such properties, one must also be able to specify and enforce some security policies. Indeed, programming service oriented middlewares and the applications running on them without putting data at risk or compromising robustness of the whole platform requires services to be checked and their resource usage to be strictly put in relation to their capabilities.

Great efforts have been recently devoted to embed security mechanisms within standard programming features (some of these techniques are surveyed in [28]). Language-based mechanisms are a scalable way to provide evidence that a large number of applications enjoy some given properties. For example, by using type systems, one can prove the type soundness of the language as a whole, from which it follows that all well-typed applications do comply with the policies stated by their types. To facilitate the task of designing such a sound language for SOC, one can initially focus only on

⋆ This work has been supported by the EU project SENSORIA, IST-2 005-016004.

the mechanisms at the basis of SOC. Afterwards, this core formalism could hopefully be expanded into a full-fledged language by adding the high level, often redundant, constructs typical of effective programming languages.

Many researchers have hence put forward exploiting the studies on *process calculi*, a cornerstone of current foundational research on specification and analysis of concurrent, distributed and mobile systems through mathematical — mainly algebraic and logical — tools. Indeed, due to their algebraic nature, process calculi convey in a distilled form the compositional programming style of SOC. This is witnessed by the several process calculi like formalisms for SOC proposed in the literature by now (see, e.g., [2, 29, 11, 5, 18, 21, 13, 22]). However, although capable of describing complex systems and applications, such proposals still lack those reasoning mechanisms and analytical tools, e.g. type systems and behavioural equivalences, that process calculi usually hand down.

In this paper, we tailor the type-based approach for protecting data in global computing applications put forward in [9] to COWS, a formalism for specifying service-based applications that we introduce in [22]. We thus define a typed variant of COWS that permits expressing and forcing policies regulating the exchange of data among interacting services. Programmers can indeed settle the partners usable to exchange any given datum (and, then, the services that can share it), thus avoiding the datum be accessed (by unwanted services) through unauthorized partners. The language (static and dynamic) semantics then guarantees that well-typed services always comply with the constraints expressed by the type associated to each single datum.

The rest of the paper is organized as follows. Section 2 introduces syntax, type inference and operational semantics of (our typed variant of) COWS, while Section 3 presents our main results. Section 4 demonstrates our approach through a simplified but realistic scenario for a service-based electronic marketplace, and an automotive scenario studied in the SENSORIA project. Finally, Section 5 touches upon comparisons with more strictly related work and directions for future work.

2 COWS: Calculus for Orchestration of Web Services

Before formally defining our language, we provide some insights on its main features. We refer the interested reader to [22] for further motivations on the design of COWS, for many examples illustrating its peculiarities and expressiveness, and for comparisons with other process-based and orchestration formalisms.

The design of COWS has been influenced by the principles underlying WS-BPEL [25], the *de facto* standard language for orchestration of web services. Similarly to WS-BPEL, COWS supports service instances with shared states, allows a same process to play more than one partner role and permits programming stateful sessions by correlating different service interactions. However, COWS intends to be a foundational model not specifically tight to web services' current technology. Thus, some WS-BPEL constructs, such as e.g. fault and compensation handlers and flow graphs, do not have a precise counterpart in COWS, rather they are expressed in terms of more primitive operators (see [20], Sect. 3). The design of COWS has also taken advantage of previous work on process calculi. In fact, it combines in an original way constructs and features borrowed from well-known process calculi, e.g. asynchronous communication,

polyadic synchronization, pattern matching, protection, delimited receiving and killing activities, while however resulting different from any of them.

The basic elements of COWS are *partners* and *operations*. They can be combined to designate *communication endpoints* and can be exchanged in communication, but dynamically received names cannot form endpoints used to receive further invocations. Endpoints naming mechanism is very flexible, e.g. it permits identifying a same service by means of different logic names and separately dealing with the names composing an endpoint. This is, e.g., exploited in request-response interaction, where usually the service provider knows the name of the response operation, but not the partner name of the service it has to reply to.

COWS computational entities are called *services*. Typically, a service creates one specific instance to serve each received request. Instances may run concurrently. Each instance can be composed of concurrent threads that may offer a choice among alternative receive activities. Services could be able to receive multiple messages in a statically unpredictable order and in such a way that the first incoming message triggers creation of a service instance which subsequent messages are routed to. *Pattern matching* is the mechanism used for correlating messages logically forming a same interaction 'session' by means of their same contents. It permits locating those data that are important to identify service instances and is flexible enough for allowing a single message to participate in multiple interaction sessions, each identified by separate correlation values.

Inter-service communication give rise to substitutions of variables with values. However, to enable concurrent instances or threads within an instance to share the state (or part of it), receive activities in COWS do *not* bind variables. The range of application of the substitution generated by a communication is then regulated by the *delimitation* operator, that is the only binder of the calculus. Delimitation, additionally, can be used to generate fresh private names (as the restriction operator of the π-calculus) and to delimit the field of action of the *kill* activity, a powerful orchestration construct that can be used to force termination of whole service instances. Sensitive code can however be protected from the effect of a forced termination by using the *protection* operator.

The type system we present in this paper permits to express and enforce policies for regulating the exchange of data among services. To implement such policies, programmers can annotate data with sets of partner names characterizing the services authorized to use and exchange them; these sets are called *regions*. The language operational semantics uses these annotations to guarantee that computations proceed according to them. This property, called *soundness*, can be stated as follows

> A service s is *sound* if, for any datum v in s associated to region r and for all evolutions of s, it holds that v can be exchanged only by using partners in r.

To facilitate the task of decorating COWS terms with type annotations, we let the type system partially infer such annotations *à la* ML: service programmers explicitly write only the annotations necessary to specify the wanted policies for communicable data; then, a type inference system (statically) performs some coherence checks (e.g. the partner used by an invoke must belong to the regions of all data occurring in the argument of the activity) and derives the minimal region annotations for variable declarations that ensure consistency of services initial configuration. This allows us to define an *operational semantics with types* [12] which is simpler than a full-fledged *typed operational*

Table 1. COWS syntax

s ::=	**kill**(k)	|	$u \bullet u'!\overline{\{e(\bar{x})\}_r}$	|	g	|	$s \mid s$	|	$\{\!\{s\}\!\}$	|	$[d]\, s$	|	$*\, s$ (services)
g ::=	$\mathbf{0}$	|	$p \bullet o?\bar{w}.s$	|	$g + g$								(receive-guarded choice)

semantics, because it only performs simple checks (i.e. subset inclusion) using region annotations to authorize or block transitions. Our main results prove that the type system and the operational semantics are sound. As a consequence, we have that services always comply with the constraints expressed by the type of each single datum.

Syntax. COWS syntax is parameterized by three countable and pairwise disjoint sets: the set of *(killer) labels* (ranged over by $k, k', \ldots$), the set of *values* (ranged over by v, v', $\ldots$) and the set of 'write once' *variables* (ranged over by x, y, $\ldots$). The set of values is left unspecified; however, we assume that it includes the set of *names*, ranged over by n, m, $\ldots$, mainly used to represent partners and operations. COWS is also parameterized by a set of *expressions*, ranged over by e, whose exact syntax is deliberately omitted; we just assume that expressions contain, at least, values and variables. Notably, killer labels are *not* (communicable) values. Notationally, we prefer letters $p, p', \ldots$ when we want to stress the use of a name as a partner, $o, o', \ldots$ when we want to stress the use of a name as an operation. We will use w to range over values and variables, u to range over names and variables, and d to range over killer labels, names and variables.

Regions can be either finite subsets of partners and variables or the distinct element $\top$ (denoting the universe of partners). The set of all regions, ranged over by r, is partially ordered by the subset inclusion relation $\subseteq$, and has $\top$ as top element.

Notation $\bar{\cdot}$ stands for tuples of objects, e.g. $\bar{x}$ is a compact notation for denoting the tuple of variables $\langle x_1, \ldots, x_n \rangle$ (with $n \geq 0$). We assume that variables in the same tuple are pairwise distinct. All notations shall extend to tuples component-wise. An expression e tagged with region r will be written as $\{e\}_r$; an untagged e will stand for $\{e\}_\top$. We will write $e(\bar{x})$ to make explicit all the variables $\bar{x}$ occurring in e (we still write e when this information is not needed), and $\bar{e}$ (resp. $\bar{r}$) to denote the tuple of the expressions (resp. regions) occurring in $\overline{\{e\}_r}$.

We will call *raw* services those COWS services written according to the syntax in Table 1. Intuitively, raw services only contain those region annotations that implement the policies for data exchange settled by the programmers. *Services* are structured activities built from basic activities, i.e. the empty activity $\mathbf{0}$, the kill activity $\mathbf{kill}(_)$, the invoke activity $_ \bullet _!_$ and the receive activity $_ \bullet _?_$, by means of (receive) prefixing $_._$, guarded choice $_ + _$, parallel composition $_ \mid _$, protection $\{\!\{_\}\!\}$, delimitation $[_]\,_$ and replication $*\,_$. Notably, as in the Lπ [24], communication endpoints of receive activities are identified statically because their syntax only allows using names and not variables. We adopt the following conventions about the operators precedence: monadic operators bind more tightly than parallel composition, and prefixing more tightly than choice. We shall omit a trailing $\mathbf{0}$ and use $[d_1, \ldots, d_n]\, s$ to denote $[d_1] \ldots [d_n]\, s$.

The only *binding* construct is delimitation: $[d]\, s$ binds d in the scope s. The occurrence of a name/variable/label is *free* if it is not under the scope of a binder. We denote by $\mathrm{fd}(t)$ (resp. $\mathrm{bd}(t)$) the set of names, variables and killer labels that occur free (resp.

Table 2. Type inference system

$$\Gamma \vdash \mathbf{0} \succ \Gamma \vdash \mathbf{0} \quad \textit{(t-nil)} \qquad\qquad \Gamma \vdash \mathbf{kill}(k) \succ \Gamma \vdash \mathbf{kill}(k) \quad \textit{(t-kill)}$$

$$\frac{\forall\, r' \in \{r_i\}_{i\in\{1,\ldots,n\}} \quad u_1 \in r'}{\begin{array}{c}\Gamma \vdash u_1 \bullet u_2!\langle\{e_1(\bar{y}_1)\}_{r_1},\ldots,\{e_n(\bar{y}_n)\}_{r_n}\rangle \succ \\ (\Gamma + \{x : r_1\}_{x\in\bar{y}_1} + \ldots + \{x : r_n\}_{x\in\bar{y}_n}) \vdash u_1 \bullet u_2!\langle\{e_1(\bar{y}_1)\}_{r_1},\ldots,\{e_n(\bar{y}_n)\}_{r_n}\rangle\end{array}} \textit{(t-inv)}$$

$$\frac{\Gamma + \{x : \{p\}\}_{x\in \mathrm{fv}(\bar{w})} \vdash s \succ \Gamma' \vdash s'}{\Gamma \vdash p \bullet o?\bar{w}.s \succ \Gamma' \vdash p \bullet o?\bar{w}.s'} \textit{(t-rec)}$$

$$\frac{\Gamma \vdash g_1 \succ \Gamma_1 \vdash g_1' \qquad \Gamma \vdash g_2 \succ \Gamma_2 \vdash g_2'}{\Gamma \vdash g_1 + g_2 \succ \Gamma_1 + \Gamma_2 \vdash g_1' + g_2'} \textit{(t-sum)}$$

$$\frac{\Gamma \vdash s \succ \Gamma' \vdash s'}{\Gamma \vdash \{\!|s|\!\} \succ \Gamma' \vdash \{\!|s'|\!\}} \textit{(t-prot)} \qquad\qquad \frac{\Gamma \vdash s \succ \Gamma' \vdash s'}{\Gamma \vdash *s \succ \Gamma' \vdash *s'} \textit{(t-repl)}$$

$$\frac{\Gamma \vdash s \succ \Gamma' \vdash s' \qquad n \notin \mathrm{reg}(\Gamma')}{\Gamma \vdash [n]\, s \succ \Gamma' \vdash [n]\, s'} \textit{(t-del}_{name}\textit{)} \qquad\qquad \frac{\Gamma \vdash s \succ \Gamma' \vdash s'}{\Gamma \vdash [k]\, s \succ \Gamma' \vdash [k]\, s'} \textit{(t-del}_{lab}\textit{)}$$

$$\frac{\Gamma, \{x : \emptyset\} \vdash s \succ \Gamma', \{x : r\} \vdash s' \qquad x \notin \mathrm{reg}(\Gamma')}{\Gamma \vdash [x]\, s \succ \Gamma' \vdash [\{x\}^{r-\{x\}}]\, s'} \textit{(t-del}_{var}\textit{)}$$

$$\frac{\Gamma \vdash s_1 \succ \Gamma_1 \vdash s_1' \qquad \Gamma \vdash s_2 \succ \Gamma_2 \vdash s_2'}{\Gamma \vdash s_1 \mid s_2 \succ \Gamma_1 + \Gamma_2 \vdash s_1' \mid s_2'} \textit{(t-par)}$$

bound) in a term t, by $\mathrm{fv}(t)$ (resp. $\mathrm{bv}(t)$) the set of free (resp. bound) variables in t, and by $\mathrm{fk}(t)$ the set of free killer labels in t. Two terms are *alpha-equivalent* if one can be obtained from the other by consistently renaming bound names/variables/labels. As usual, we identify terms up to alpha-equivalence. For simplicity sake, in the sequel we assume that bound variables in services are pairwise distinct (of course, this condition is not restrictive and can always be fulfilled by possibly using alpha-conversion).

A type inference system. The annotations put by the type inference are written as superscripts, to better distinguish them from those put by the programmers. Thus, the syntax of variable delimitation becomes $[\{x\}^r]\, s$, which means that the datum that dynamically will replace x will be used at most by the partners in r. *Typed* COWS services are then generated by the syntax in Table 1 where, differently from the previous section, d ranges over killer labels, names and annotated variables as $\{x\}^r$. Notably, types may *depend* on partner variables, i.e. on parameters of receiving activities; during computation, they are therefore affected by application of substitutions that replace partner variables with partner names. We assume that the region of a partner name always contains, at least implicitly, such partner.

The type inference system is presented in Table 2. Typing judgements are written $\Gamma \vdash s \succ \Gamma' \vdash s'$, where the type environment Γ is a finite function from variables

to regions such that $\mathrm{fv}(s) \subseteq \mathrm{dom}(\Gamma)$ and $\mathrm{bv}(s) \cap \mathrm{dom}(\Gamma) = \emptyset$ (the same holds for Γ' and s'). Type environments are written as sets of pairs of the form $x : r$, where x is a partner variable and r is its assumed region annotation. The domain of an environment is defined as usual: $\mathrm{dom}(\emptyset) = \emptyset$ and $\mathrm{dom}(\Gamma, \{x : r\}) = \mathrm{dom}(\Gamma) \cup \{x\}$, where ',' denotes union between environments with disjoint domains. The *region* of Γ is the union of the regions in Γ, i.e. $\mathrm{reg}(\emptyset) = \emptyset$ and $\mathrm{reg}(\Gamma, \{x : r\}) = r \cup \mathrm{reg}(\Gamma)$. We will write $\Gamma + \Gamma'$ to denote the environment obtained by extending Γ with Γ'; + is inductively defined by

$$\Gamma + \emptyset = \Gamma$$
$$\Gamma + \{x : r\} = \begin{cases} \Gamma', \{x : r \cup r'\} & \text{if } \Gamma = \Gamma', \{x : r'\} \\ \Gamma, \{x : r\} & \text{otherwise} \end{cases}$$
$$\Gamma + (\{x : r\}, \Gamma') = (\Gamma + \{x : r\}) + \Gamma'$$

Hence, the judgement $\emptyset \vdash s \succ \emptyset \vdash s'$ can be derived only if s is a closed raw service (because the initial environment is empty); if it is derivable, then s' is the typed service obtained by decorating s with the region annotations describing the use of each variable of s in its scope. Type inference determines such regions by considering the invoking and receiving partners where the variables occur.

We now comment on the most significant typing rules. Rule *(t-inv)* checks if the invoked partner u_1 belongs to the regions of the communicated data. If it succeeds, the type environment Γ is extended by associating a proper region to each variable used in the expressions argument of the invoke activity. Rule *(t-rec)* tries to type s in the type environment Γ extended by adding the receiving partner to the regions of the variables in $\bar{w}$. Rules *(t-sum)* and *(t-par)* yield the same typing; this is due to the sharing of variables. For instance, service $[x]\,(p \bullet o?\langle x\rangle \mid p' \bullet o'!\langle\{x\}_r\rangle)$ with $p' \in r$ is annotated as $[\{x\}^{r'}]\,(p \bullet o?\langle x\rangle \mid p' \bullet o'!\langle\{x\}_r\rangle)$ with $r' = (\{p\} \cup r - \{x\})$. In rule *(t-del$_{name}$)*, premise $n \notin \mathrm{reg}(\Gamma')$ prevents a new name n to escape from its binder $[n]$ in the inference. As an example, consider the closed raw service

$$[z]\, p \bullet o?\langle z\rangle \,.\, [p']\, p'' \bullet o''!\langle\{z\}_{\{p'',p'\}}\rangle \qquad (*)$$

Without the premise $n \notin \mathrm{reg}(\Gamma')$, the service resulting from the type inference would be $[\{z\}^{\{p'',p',p\}}]\, p \bullet o?\langle z\rangle \,.\, [p']\, p'' \bullet o''!\langle\{z\}_{\{p'',p'\}}\rangle$. The problem with this service is that the name p' occurring in the annotation associated to z by the inference system escapes from the scope of its binder and, thus, represents a completely different name. Although, service $(*)$ is not typable, by a simple semantics preserving manipulation one can get a typable service as, e.g., the following one $[p']\,[z]\, p \bullet o?\langle z\rangle \,.\, p'' \bullet o''!\langle\{z\}_{\{p'',p'\}}\rangle$.

Similarly, in rule *(t-del$_{var}$)*, premise $x \notin \mathrm{reg}(\Gamma')$ prevents initially closed services to become open at the end of the inference. Otherwise, e.g., the type inference would transform the closed raw service

$$[x]\, p \bullet o?\langle x\rangle \,.\, [y]\, p' \bullet o'?\langle y\rangle \,.\, p'' \bullet o''!\langle\{x\}_{\{p'',y\}}\rangle \qquad (**)$$

into the open service $[\{x\}^{\{p,p'',y\}}]\, p \bullet o?\langle x\rangle \,.\, [\{y\}^{\{p'\}}]\, p' \bullet o'?\langle y\rangle \,.\, p'' \bullet o''!\langle\{x\}_{\{p'',y\}}\rangle$. Also in this case, we can easily modify the untypable service $(**)$ to get a typable one with a similar semantics like, e.g., the service $[y]\,[x]\, p \bullet o?\langle x\rangle \,.\, p' \bullet o'?\langle y\rangle \,.\, p'' \bullet o''!\langle\{x\}_{\{p'',y\}}\rangle$.

Furthermore, in *(t-del$_{var}$)*, x is annotated with $r - \{x\}$, rather than with r, otherwise initially closed services could become open. E.g., the closed raw service $[x]\, p \bullet o?\langle x\rangle \,.$

Table 3. Structural congruence

$*\,\mathbf{0}$	$\equiv$	$\mathbf{0}$		$*\,s$	$\equiv$	$s \mid *\,s$	$\{\!\{\mathbf{0}\}\!\}$	$\equiv$ $\mathbf{0}$
$\{\!\{\{\!\{s\}\!\}\}\!\}$	$\equiv$	$\{\!\{s\}\!\}$		$\{\!\{[d]\,s\}\!\}$	$\equiv$	$[d]\,\{\!\{s\}\!\}$	$[d]\,\mathbf{0}$	$\equiv$ $\mathbf{0}$
$[d_1]\,[d_2]\,s$	$\equiv$	$[d_2]\,[d_1]\,s$				if $d_1 \neq \{x\}^{r_1}$ and $d_2 \neq \{y\}^{r_2}$		
$[n]\,[\{x\}^r]\,s$	$\equiv$	$[\{x\}^r]\,[n]\,s$				if $n \notin r$		
$[\{x\}^{r_1}]\,[\{y\}^{r_2}]\,s$	$\equiv$	$[\{y\}^{r_2}]\,[\{x\}^{r_1}]\,s$				if $y \notin r_1$ and $x \notin r_2$		
$s_1 \mid [d]\,s_2$	$\equiv$	$[d]\,(s_1 \mid s_2)$				if $d \notin \mathrm{fd}(s_1) \cup \mathrm{fk}(s_2)$		

Table 4. Matching rules

$$\mathcal{M}(v, \{v\}_r) = \emptyset \qquad \mathcal{M}(x, \{v\}_r) = \{x \mapsto \{v\}_r\} \qquad \frac{\mathcal{M}(w_1, \{v_1\}_{r_1}) = \sigma_1 \qquad \mathcal{M}(\bar{w}_2, \overline{\{v_2\}_{r_2}}) = \sigma_2}{\mathcal{M}((w_1, \bar{w}_2), (\{v_1\}_{r_1}, \overline{\{v_2\}_{r_2}})) = \sigma_1 \uplus \sigma_2}$$

$p' \cdot o'!\langle\{x\}_{\{p',x\}}\rangle$ would be transformed into the open service $[\{x\}^{\{p,p',x\}}]\,p \cdot o?\langle x\rangle.\,p' \cdot o'!\langle\{x\}_{\{p',x\}}\rangle$ (indeed, x occurs in the annotation associated to its declaration). Notice that, although the region associated to x by the inference does never record that a service possibly transmits x with regions containing x, rule *(t-del$_{var}$)* is sound because we assumed that the region of a partner name, at least implicitly, contains the partner name.

Definition 1. *A service s is* well-typed *if $\emptyset \vdash s' \succ \emptyset \vdash s$ for some raw service s'.*

Operational semantics. COWS operational semantics is defined only for *closed* services, i.e. services without free variables/labels (similarly to many real compilers, we consider terms with free variables/labels as programming errors), but of course the rules also involve non-closed services (see e.g. the premises of rules *(del_)*). Formally, the semantics is given in terms of a structural congruence and of a labelled transition relation.

The structural congruence $\equiv$ identifies syntactically different services that intuitively represent the same service. It is defined as the least congruence relation induced by a given set of equational laws. We explicitly show in Table 3 the laws for replication, protection and delimitation, while omit the (standard) laws for the other operators stating that parallel composition is commutative, associative and has $\mathbf{0}$ as identity element, and that guarded choice enjoys the same properties and, additionally, is idempotent. All the presented laws are straightforward. Only notice that the last law can be used to extend the scope of names (like a similar law in the π-calculus), thus enabling communication of restricted names, except when the argument d of the delimitation is a free killer label of s_2 (this avoids involving s_1 in the effect of a kill activity inside s_2).

To define the labelled transition relation, we need a few auxiliary functions. First, we exploit a function $[\![_]\!]$ for evaluating *closed* expressions (i.e. expressions without variables): it takes a closed expression and returns a value. However, $[\![_]\!]$ cannot be explicitly defined because the exact syntax of expressions is deliberately not specified.

Then, through the rules in Table 4, we define the partial function $\mathcal{M}(_,_)$ that permits performing *pattern-matching* on semi-structured data thus determining if a receive

Table 5. Is there an active **kill**(k)? / Are there conflicting receives along $p \bullet o$ matching $\overline{\{v\}_r}$?

$\mathbf{kill}(k) \downarrow_{kill}$	$\dfrac{s \downarrow_{kill} \ \vee \ s' \downarrow_{kill}}{s \mid s' \downarrow_{kill}}$	$\dfrac{s \downarrow_{kill}}{\{\!\!\{s\}\!\!\} \downarrow_{kill}}$	$\dfrac{s \downarrow_{kill}}{[d]\, s \downarrow_{kill}}$	$\dfrac{s \downarrow_{kill}}{*\, s \downarrow_{kill}}$
$\dfrac{\lvert \mathcal{M}(\bar{w}, \overline{\{v\}_r}) \rvert < \ell}{p \bullet o ? \bar{w}.s \downarrow^{\ell}_{p \bullet o, \overline{\{v\}_r}}}$	$\dfrac{s \downarrow^{\ell}_{p \bullet o, \overline{\{v\}_r}} \quad d \notin \{p, o\}}{[d]\, s \downarrow^{\ell}_{p \bullet o, \overline{\{v\}_r}}}$	$\dfrac{s \downarrow^{\ell}_{p \bullet o, \overline{\{v\}_r}}}{\{\!\!\{s\}\!\!\} \downarrow^{\ell}_{p \bullet o, \overline{\{v\}_r}}}$		
$\dfrac{g \downarrow^{\ell}_{p \bullet o, \overline{\{v\}_r}} \ \vee \ g' \downarrow^{\ell}_{p \bullet o, \overline{\{v\}_r}}}{g + g' \downarrow^{\ell}_{p \bullet o, \overline{\{v\}_r}}}$	$\dfrac{s \downarrow^{\ell}_{p \bullet o, \overline{\{v\}_r}} \ \vee \ s' \downarrow^{\ell}_{p \bullet o, \overline{\{v\}_r}}}{s \mid s' \downarrow^{\ell}_{p \bullet o, \overline{\{v\}_r}}}$	$\dfrac{s \downarrow^{\ell}_{p \bullet o, \overline{\{v\}_r}}}{*\, s \downarrow^{\ell}_{p \bullet o, \overline{\{v\}_r}}}$		

and an invoke over the same endpoint can synchronize. The rules state that two tuples match if they have the same number of fields and corresponding fields have matching values/variables. Variables match any annotated value, and a value matches an annotated value only if, apart for the region annotation, they are identical. When tuples $\bar{w}$ and $\overline{\{v\}_r}$ do match, $\mathcal{M}(\bar{w}, \overline{\{v\}_r})$ returns a substitution, that also records region annotations of values exchanged in communication, for the variables in $\bar{w}$; otherwise, it is undefined. *Substitutions* (ranged over by σ) are functions mapping variables to annotated values and are written as collections of pairs of the form $x \mapsto \{v\}_r$. Application of substitution σ to s, written $s \cdot \sigma$, has the effect of replacing every free occurrence of x in s with v, for each $x \mapsto \{v\}_r \in \sigma$, by possibly using alpha-conversion for avoiding v to be captured by name delimitations within s. We use $|\sigma|$ to denote the number of pairs in σ and $\sigma_1 \uplus \sigma_2$ to denote the union of σ_1 and σ_2 when they have disjoint domains.

We also define a function, named $halt(_)$, that takes a service s as an argument and returns the service obtained by only retaining the protected activities inside s. $halt(_)$ is defined inductively on the syntax of services. The most significant case is $halt(\{\!\!\{s\}\!\!\}) = \{\!\!\{s\}\!\!\}$. In the other cases, $halt(_)$ returns $\mathbf{0}$, except for parallel composition, delimitation and replication operators, for which it acts as an homomorphism.

Finally, in Table 5, we inductively define two predicates: $s{\downarrow_{kill}}$ checks if s can immediately perform a **kill**(k); $s{\downarrow^{\ell}_{p \bullet o, \overline{\{v\}_r}}}$, with ℓ natural number, checks existence of potential communication conflicts, i.e. the ability of s of performing a receive activity matching $\overline{\{v\}_r}$ over the endpoint $p \bullet o$ that generates a substitution with fewer pairs than ℓ.

The labelled transition relation $\xrightarrow{\alpha}$ is the least relation over services induced by the rules in Table 6, where α is generated by the following grammar:

$$\alpha ::= \dagger k \quad | \quad (p \bullet o) \lhd \overline{\{v\}_r} \quad | \quad (p \bullet o) \rhd \bar{w} \quad | \quad p \bullet o \lfloor \sigma \rfloor \bar{w}\, \overline{\{v\}_r} \quad | \quad \dagger$$

In the sequel, we use $d(\alpha)$ to denote the set of names, variables and killer labels occurring in α, except for $\alpha = p \bullet o \lfloor \sigma \rfloor \bar{w}\, \overline{\{v\}_r}$ for which we let $d(p \bullet o \lfloor \sigma \rfloor \bar{w}\, \overline{\{v\}_r}) = d(\sigma)$, where $d(\{x \mapsto \{v\}_r\}) = \{x, v\} \cup r$ and $d(\sigma_1 \uplus \sigma_2) = d(\sigma_1) \cup d(\sigma_2)$. The meaning of labels is as follows: $\dagger k$ denotes execution of a request for terminating a term from within the delimitation $[k]$, $(p \bullet o) \lhd \overline{\{v\}_r}$ and $(p \bullet o) \rhd \bar{w}$ denote execution of invoke and receive

Table 6. Operational semantics

$$\mathbf{kill}(k) \xrightarrow{\dagger k} \mathbf{0} \quad (kill) \qquad\qquad p \bullet o?\bar{w}.s \xrightarrow{(p \bullet o) \triangleright \bar{w}} s \quad (rec)$$

$$\frac{[\![\bar{e}]\!] = \bar{v} \qquad \mathrm{fv}(\bar{r}) = \emptyset}{p \bullet o!\overline{\{e\}_r} \xrightarrow{(p \bullet o) \triangleleft \overline{\{v\}_r}} \mathbf{0}} \ (inv) \qquad\qquad \frac{g_1 \xrightarrow{\alpha} s}{g_1 + g_2 \xrightarrow{\alpha} s} \ (choice)$$

$$\frac{s \xrightarrow{p \bullet o \lfloor \sigma \uplus \{x \mapsto \{v\}_r\} \rfloor \bar{w}\, \overline{\{v'\}_{r'}}} s' \qquad r'' \cdot \sigma \subseteq r}{[\{x\}^{r''}]\, s \xrightarrow{p \bullet o \lfloor \sigma \rfloor \bar{w}\, \overline{\{v'\}_{r'}}} s' \cdot \{x \mapsto \{v\}_r\}} \ (del_{sub}) \qquad\qquad \frac{s \xrightarrow{\dagger k} s'}{[k]\, s \xrightarrow{\dagger} [k]\, s'} \ (del_{kill})$$

$$\frac{s \xrightarrow{\alpha} s' \quad d \neq \{x\}^r \quad d \notin d(\alpha) \quad s \downarrow_{kill} \Rightarrow \alpha = \dagger, \dagger k}{[d]\, s \xrightarrow{\alpha} [d]\, s'} \ (del_{pass}) \qquad\qquad \frac{s \xrightarrow{\alpha} s' \quad x \notin d(\alpha)}{[\{x\}^r]\, s \xrightarrow{\alpha} [\{x\}^r]\, s'} \ (x_{pass})$$

$$\frac{s_1 \xrightarrow{(p \bullet o) \triangleright \bar{w}} s_1' \qquad s_2 \xrightarrow{(p \bullet o) \triangleleft \overline{\{v\}_r}} s_2' \qquad \mathcal{M}(\bar{w}, \overline{\{v\}_r}) = \sigma \qquad \neg(s_1 \mid s_2 \downarrow^{|\sigma|}_{p \bullet o, \overline{\{v\}_r}})}{s_1 \mid s_2 \xrightarrow{p \bullet o \lfloor \sigma \rfloor \bar{w}\, \overline{\{v\}_r}} s_1' \mid s_2'} \ (com)$$

$$\frac{s_1 \xrightarrow{p \bullet o \lfloor \sigma \rfloor \bar{w}\, \overline{\{v\}_r}} s_1' \qquad \neg(s_2 \downarrow^{|\mathcal{M}(\bar{w}, \overline{\{v\}_r})|}_{p \bullet o, \overline{\{v\}_r}})}{s_1 \mid s_2 \xrightarrow{p \bullet o \lfloor \sigma \rfloor \bar{w}\, \overline{\{v\}_r}} s_1' \mid s_2} \ (par_{conf}) \qquad\qquad \frac{s_1 \xrightarrow{\dagger k} s_1'}{s_1 \mid s_2 \xrightarrow{\dagger k} s_1' \mid halt(s_2)} \ (par_{kill})$$

$$\frac{s_1 \xrightarrow{\alpha} s_1' \qquad \alpha \neq (p \bullet o \lfloor \sigma \rfloor \bar{w}\, \overline{\{v\}_r}), \dagger k}{s_1 \mid s_2 \xrightarrow{\alpha} s_1' \mid s_2} \ (par_{pass}) \qquad\qquad \frac{s \xrightarrow{\alpha} s'}{\{\!|s|\!\} \xrightarrow{\alpha} \{\!|s'|\!\}} \ (prot)$$

$$\frac{s \equiv s_1 \quad s_1 \xrightarrow{\alpha} s_2 \quad s_2 \equiv s'}{s \xrightarrow{\alpha} s'} \ (cong)$$

activities over the endpoint $p \bullet o$, respectively, $p \bullet o \lfloor \sigma \rfloor \bar{w}\, \overline{\{v\}_r}$ (if $\sigma \neq \emptyset$) denotes execution of a communication over $p \bullet o$ with receive parameters $\bar{w}$ and matching values $\overline{\{v\}_r}$ and with substitution σ to be still applied, $\dagger$ and $p \bullet o \lfloor \emptyset \rfloor \bar{w}\, \overline{\{v\}_r}$ denote *computational steps* corresponding to taking place of forced termination and communication (without pending substitutions), respectively. Hence, a *computation* from a closed service s_0 is a sequence of connected transitions of the form

$$s_0 \xrightarrow{\alpha_1} s_1 \xrightarrow{\alpha_2} s_2 \xrightarrow{\alpha_3} s_3 \ldots$$

where, for each i, α_i is either $p \bullet o \lfloor \emptyset \rfloor \bar{w}\, \overline{\{v\}_r}$ or $\dagger$, and s_i is called *reduct* of s_0.

We comment on salient points. Activity $\mathbf{kill}(k)$ forces termination of all unprotected parallel activities (rules *(kill)* and *(par_{kill})*) inside an enclosing $[k]$, that stops the killing effect by turning the transition label $\dagger k$ into $\dagger$ (rule *(del_{kill})*). Existence of such delimitation is ensured by the assumption that the semantics is only defined for closed services. Sensitive code can be protected from killing by putting it into a protection $\{\!|_|\!\}$; this way, $\{\!|s|\!\}$ behaves like s (rule *(prot)*). Similarly, $[d]\, s$ behaves like s, except when the transition

label α contains d or when a kill activity is active in s and α does not correspond to a kill activity (rules *(del$_{pass}$)* and *(x$_{pass}$)*): in such cases the transition should be derived by using rules *(del$_{kill}$)* or *(del$_{sub}$)*. In other words, kill activities are executed *eagerly*. A service invocation can proceed only if the expressions in the argument can be evaluated and their regions do not contain variables (rule *(inv)*). A receive activity offers an invocable operation along a given partner name (rule *(rec)*). The execution of a receive permits to take a decision between alternative behaviours (rule *(choice)*). Communication can take place when two parallel services perform matching receive and invoke activities (rules *(com)*). Communication generates a substitution that is recorded in the transition label (for subsequent application), rather than a silent transition as in most process calculi. If more then one matching is possible the receive that needs fewer substitutions is selected to progress (rules *(com)* and *(par$_{conf}$)*). This mechanism permits to correlate different service communications thus implicitly creating interaction sessions and can be exploited to model the precedence of a service instance over the corresponding service specification when both can process the same request. A substitution $\{x \mapsto \{v\}_r\}$ for a variable x is applied to a term (rule *(del$_{sub}$)*) when the delimitation for x is encountered, i.e. the whole scope s of x is determined, provided that the region annotations of the variable declaration and of the substituent datum v do comply i.e. $r' \cdot \sigma \subseteq r$. This condition also means that as a value is received it gets annotated with a smaller region. The substitution for x is then applied to s and x disappears from the term and cannot be reassigned a value. Execution of parallel services is interleaved (rule *(par$_{pass}$)*), but when a kill activity or a communication is performed. Indeed, the former must trigger termination of all parallel services (according to rule *(par$_{kill}$)*), while the latter must ensure that the receive activity with greater priority progresses (rule *(com)* and *(par$_{conf}$)*). The last rule states that structurally congruent services have the same transitions.

2.1 Examples

We end this section with a few examples aimed at clarifying the peculiarities of our formalism. For simplicity sake, in the examples data have no attached policies, which, according to our convention, means that they are tagged with $\top$.

Communication of private names. Communication of private names is standard and exploits scope extension as in π-calculus. Notably, receive and invoke activities can interact only if both are in the scopes of the delimitations that bind the variables argument of the receive. Thus, to enable communication of private names, besides their scopes, we must possibly extend the scopes of some variables, as in the following example:

$$\begin{array}{lcl} [x]\,(p \bullet o?\langle x\rangle.s \mid s') \mid [n]\,p \bullet o!\langle n\rangle & \equiv & (n \text{ fresh}) \\ [n]\,([x]\,(p \bullet o?\langle x\rangle.s \mid s') \mid p \bullet o!\langle n\rangle) & \equiv & \\ [n]\,[x]\,(p \bullet o?\langle x\rangle.s \mid s' \mid p \bullet o!\langle n\rangle) & \xrightarrow{p \bullet o\,\lfloor\emptyset\rfloor\,\langle x\rangle\,\langle n\rangle} & \\ [n]\,(s \mid s') \cdot \{x \mapsto n\} & & \end{array}$$

Notice that the substitution $\{x \mapsto n\}$ is applied to all terms delimited by $[x]$, not only to the continuation s of the service performing the receive. This is different from most process calculi and accounts for the global scope of variables. This very feature permits

to easily model the *delayed input* of fusion calculus [26], which is instead difficult to express in π-calculus.

Delimited killer labels. We require killer labels to be delimited to avoid a single service be capable to stop all the other parallel services which would be unreasonable in a service-oriented setting. Indeed, suppose a service s can perform a $\mathbf{kill}(k)$ with k undelimited in s. The killing effect could not be stopped, thus, due to a transition labelled by $\dagger k$, the whole service s would be terminated (but for protected activities). Moreover, the effect of $\mathbf{kill}(k)$ could not be confined to s, thus, if there are other parallel services, the whole service composition might be terminated by $\mathbf{kill}(k)$.

Protected kill activity. The following simple example illustrates the effect of executing a kill activity within a protection block:

$$[k]\,(\{\!\{s_1 \mid \{\!\{s_2\}\!\} \mid \mathbf{kill}(k)\}\!\} \mid s_3) \mid s_4 \xrightarrow{\dagger} [k]\,\{\!\{\{\!\{s_2\}\!\}\}\!\} \mid s_4$$

where, for simplicity, we assume that $halt(s_1) = halt(s_3) = \mathbf{0}$. In essence, $\mathbf{kill}(k)$ terminates all parallel services inside delimitation $[k]$ (i.e. s_1 and s_3), except those that are protected at the same nesting level of the kill activity (i.e. s_2).

Interplay between communication and kill activity. Kill activities can break communication, as the following example shows:

$$p \bullet o!\langle n\rangle \mid [k]\,([x]\,p \bullet o?\langle x\rangle.s \mid \mathbf{kill}(k)) \xrightarrow{\dagger} p \bullet o!\langle n\rangle \mid [k]\,[x]\,\mathbf{0}$$

Communication can however be guaranteed by protecting the receive activity, as in

$$\begin{array}{ll} p \bullet o!\langle n\rangle \mid [k]\,([x]\,\{\!\{p \bullet o?\langle x\rangle.s\}\!\} \mid \mathbf{kill}(k)) & \xrightarrow{\dagger} \\ p \bullet o!\langle n\rangle \mid [k]\,[x]\,\{\!\{p \bullet o?\langle x\rangle.s\}\!\} & \equiv \\ [x]\,(p \bullet o!\langle n\rangle \mid [k]\,\{\!\{p \bullet o?\langle x\rangle.s\}\!\}) & \xrightarrow{p \bullet o\lfloor\emptyset\rfloor\langle x\rangle\langle n\rangle} \\ [k]\,\{\!\{s \cdot \{x \mapsto n\}\}\!\} & \end{array}$$

Conflicting receive activities. This example shows a *persistent service* (implemented by mean of replication), that, once instantiated, enables two conflicting receives:

$$\begin{array}{ll} *\,[x]\,(\,p_1 \bullet o?\langle x\rangle.s_1 \mid p_2 \bullet o?\langle x\rangle.s_2\,) \mid p_1 \bullet o!\langle v\rangle \mid p_2 \bullet o!\langle v\rangle & \xrightarrow{p_1 \bullet o\lfloor\emptyset\rfloor\langle x\rangle\langle v\rangle} \\ *\,[x]\,(\,p_1 \bullet o?\langle x\rangle.s_1 \mid p_2 \bullet o?\langle x\rangle.s_2\,) \mid s_1 \cdot \{x \mapsto v\} \mid p_2 \bullet o?\langle v\rangle.s_2 \cdot \{x \mapsto v\} \mid p_2 \bullet o!\langle v\rangle & \end{array}$$

Now, the persistent service and the created instance, being both able to receive the same tuple $\langle v\rangle$ along the endpoint $p_2 \bullet o$, compete for the request $p_2 \bullet o!\langle v\rangle$. However, our (prioritized) semantics, in particular rule *(com)* in combination with rule *(par_{conf})*, allows only the existing instance to evolve (and, thus, prevents creation of a new instance):

$$*\,[x]\,(\,p_1 \bullet o?\langle x\rangle.s_1 \mid p_2 \bullet o?\langle x\rangle.s_2\,) \mid s_1 \cdot \{x \mapsto v\} \mid s_2 \cdot \{x \mapsto v\}$$

Message correlation. Consider now uncorrelated receive activities executed by a same instance, like in the following service:

$$*\,[x]\,p_1 \bullet o_1?\langle x\rangle.[y]\,p_2 \bullet o_2?\langle y\rangle.s$$

The fact that the messages for operations o_1 and o_2 are uncorrelated implies that, e.g., if there are concurrent instances then successive invocations for a same instance can mix up and be delivered to different instances. If one thinks it right, this behaviour can be avoided simply by correlating successive messages by means of some correlation data, e.g. the first received value as in the following service:

$$*\,[x]\,p_1 \bullet o_1?\langle x\rangle.[y]\,p_2 \bullet o_2?\langle y, x\rangle.s$$

3 Main Results

Our main results are standard and state that well-typedness is preserved along computations (*subject reduction*) and that well-typed services do respect region annotations (*type safety*). Together, these results imply the *soundness* of our theory, i.e. no violation of data regions will ever occur during the evolution of well-typed services. The formal account of these results follow.

For the proof of subject reduction, we need some standard lemmata concerning substitution and weakening. The substitution lemma handles the substitution of partner variables by partner names. Application of a substitution σ to a type environment Γ, written $\Gamma \cdot \sigma$, is defined only when $\mathrm{dom}(\sigma) \cap \mathrm{dom}(\Gamma) = \emptyset$ and, for each $x \mapsto \{v\}_r \in \sigma$, has the effect of replacing every occurrence of x in the regions of Γ with v, i.e.

$$\emptyset \cdot \{x \mapsto \{v\}_r\} = \emptyset \quad \text{and} \quad (\Gamma, \{y : r'\}) \cdot \{x \mapsto \{v\}_r\} = \Gamma \cdot \{x \mapsto \{v\}_r\}, \{y : (r' \cdot \{x \mapsto \{v\}_r\})\}.$$

Lemma 1 (Substitution Lemma). *If* $\Gamma, \{x : r\} \vdash s \succ \Gamma', \{x : r'\} \vdash s'$ *and* $\sigma = \{x \mapsto \{v\}_{r''}\}$*, then* $\Gamma \cdot \sigma \vdash s \cdot \sigma \succ \Gamma' \cdot \sigma \vdash s' \cdot \sigma$.

Proof. The proof proceeds by induction on the length of the inference used to derive the typing judgement.

The base cases *(t-nil)* and *(t-kill)* are trivial to conclude. Let us consider the base case *(t-inv)*: by definition $s = s' = u_1 \bullet u_2!\langle\{e_1(\bar{y}_1)\}_{r_1}, \ldots, \{e_n(\bar{y}_n)\}_{r_n}\rangle$ where $u_1 \in r_i$ for each $i \in \{1, \ldots, n\}$. Hence, we also have that $u_1 \cdot \sigma \in r_i \cdot \sigma$ for any r_i and σ. We now distinguish three cases:

1. $\bar{y}_1 \cup \ldots \cup \bar{y}_n = \emptyset$. In this case, $\Gamma' = \Gamma$ and $r' = r$; thus by using rule *(t-inv)*, we can conclude that $\Gamma \cdot \sigma \vdash u_1 \bullet u_2!\langle\{e_1\}_{r_1}, \ldots, \{e_n\}_{r_n}\rangle \cdot \sigma \succ \Gamma' \cdot \sigma \vdash u_1 \bullet u_2!\langle\{e_1\}_{r_1}, \ldots, \{e_n\}_{r_n}\rangle \cdot \sigma$.
2. Let $I \stackrel{def}{=} \{j \in \{1, \ldots, n\} : x \in \bar{y}_j\}$. In this case, it must be $r' = \bigcup_{i \in I} r_i \cup r$ and $\Gamma' = \Gamma + \{z : r_1\}_{z \neq x, z \in \bar{y}_1} + \ldots + \{z : r_n\}_{z \neq x, z \in \bar{y}_n}$. Now, we can to conclude like in the previous case.
3. $x \notin \bar{y}_1 \cup \ldots \cup \bar{y}_n$. In this case $r' = r$ and, again, we can conclude like before.

Let us now consider the inductive case and reason by case analysis on the last rule used to infer the judgement. We explicitly show the most significant cases, the remaining ones are easier.

(t-rec): Let $s = p \bullet o?\bar{w}.t$ and $s' = p \bullet o?\bar{w}.t'$. Then, we can suppose that $\Gamma'', \{x : r''\} \vdash t \succ \Gamma', \{x : r'\} \vdash t'$ for some Γ'' and r''. We distinguish three cases:

1. $\mathrm{fv}(\bar{w}) = \emptyset$. In this case $\Gamma'' = \Gamma$ and $r'' = r$. By induction, we have $\Gamma \cdot \sigma \vdash t \cdot \sigma \succ \Gamma' \cdot \sigma \vdash t' \cdot \sigma$. Now, by using rule *(t-rec)*, we can conclude that $\Gamma \cdot \sigma \vdash p \bullet o?\bar{w}.(t \cdot \sigma) \succ \Gamma' \cdot \sigma \vdash p \bullet o?\bar{w}.(t' \cdot \sigma)$.
2. $x \in \mathrm{fv}(\bar{w})$. In this case $\Gamma'' = \Gamma + \{z : \{p\}\}_{z \neq x, z \in \mathrm{fv}(\bar{w})}$ and $r'' = r \cup \{p\}$ and by induction, $\Gamma \cdot \sigma + \{z : \{p\}\}_{z \neq x, z \in \mathrm{fv}(\bar{w})} \vdash t \cdot \sigma \succ \Gamma' \cdot \sigma \vdash t' \cdot \sigma$. Hence by using rule *(t-rec)*, we can conclude the wanted $\Gamma \cdot \sigma \vdash (p \bullet o?\bar{w}.t) \cdot \sigma \succ \Gamma' \cdot \sigma \vdash (p \bullet o?\bar{w}.t') \cdot \sigma$.
3. $\mathrm{fv}(\bar{w}) \neq \emptyset$ and $x \notin \mathrm{fv}(\bar{w})$. In this case $\Gamma'' = \Gamma + \{z : \{p\}\}_{z \in \mathrm{fv}(\bar{w})}$ and $r'' = r$, and we can conclude like before.

(t-del$_{var}$): Let $s = [z]\, t$ and $s' = [\{z\}^{r''-\{z\}}]\, t'$ for some r''. Since $z \in \mathrm{bv}(s)$, by type environment definition, we have $z \neq x$. Then, we can suppose that $\Gamma, \{x : r\}, \{z : \emptyset\} \vdash t \succ \Gamma', \{x : r'''\}, \{z : r''\} \vdash t'$ for some r''' such that $z \notin (r''' \cup \mathrm{reg}(\Gamma'))$. Thus, by induction $\Gamma \cdot \sigma, \{z : \emptyset\} \vdash t \cdot \sigma \succ \Gamma' \cdot \sigma, \{z : r'' \cdot \sigma\} \vdash t' \cdot \sigma$. Hence, by using rule *(t-del$_{var}$)*, we can conclude the wanted $\Gamma \cdot \sigma \vdash [z]\, t \cdot \sigma \succ \Gamma' \cdot \sigma \vdash [\{z\}^{r'' \cdot \sigma - \{z\}}]\, t' \cdot \sigma$.
(t-par): Let $s = s_1 \mid s_2$ and $s' = s'_1 \mid s'_2$. Then, we can suppose that $\Gamma, \{x : r\} \vdash s_1 \succ \Gamma_1, \{x : r'_1\} \vdash s'_1$ and $\Gamma, \{x : r\} \vdash s_2 \succ \Gamma_2, \{x : r'_2\} \vdash s'_2$. By induction, $\Gamma \cdot \sigma \vdash s_1 \cdot \sigma \succ \Gamma_1 \cdot \sigma \vdash s'_1 \cdot \sigma$ and $\Gamma \cdot \sigma \vdash s_2 \cdot \sigma \succ \Gamma_2 \cdot \sigma \vdash s'_2 \cdot \sigma$. Hence by letting $\Gamma' = \Gamma_1 + \Gamma_2$ we can conclude the wanted $\Gamma \cdot \sigma \vdash (s_1 \mid s_2) \cdot \sigma \succ \Gamma' \cdot \sigma \vdash (s'_1 \mid s'_2) \cdot \sigma$. □

Lemma 2 (Weakening Lemma). *Let $\Gamma' \vdash s' \succ \Gamma \vdash s$ and $x \notin \mathrm{bd}(s)$, then $\Gamma' + \{x : r\} \vdash s' \succ \Gamma + \{x : r\} \vdash s$.*

Proof. By a straightforward induction on the length of the type derivation, with a case analysis on the last used rule, and by exploiting the fact that extending Γ by adding $\{x : r\}$ does not affect the premise of rule *(t-inv)*. □

We also need a few auxiliary results. The first one states that function *halt*(_) preserves well-typedness and can be easily proved by induction on the definition of *halt*(_).

Lemma 3. *If s is well-typed then halt(s) is well-typed.*

The other two auxiliary results establish well-typedness preservation by the structural congruence and by the labelled transition relation, respectively. We will use the following preorder $\sqsubseteq$ on type environments: we write $\Gamma \sqsubseteq \Gamma'$ if there exists a type environment Γ'' such that $\Gamma + \Gamma'' = \Gamma'$.

Lemma 4. *If $\Gamma' \vdash s'_1 \succ \Gamma \vdash s_1$ and $s_1 \equiv s_2$ then there exists a raw service s'_2 such that $\Gamma' \vdash s'_2 \succ \Gamma \vdash s_2$.*

Proof. By a straightforward induction on the derivation of $s_1 \equiv s_2$. □

Theorem 1. *If $\Gamma'_1 \vdash s'_1 \succ \Gamma_1 \vdash s_1$ and $s_1 \xrightarrow{\alpha} s_2$ then there exist a raw service s'_2 and two type environments Γ_2 and Γ'_2 such that $\Gamma_2 \sqsubseteq \Gamma_1$, $\Gamma'_1 \sqsubseteq \Gamma'_2$ and $\Gamma'_2 \vdash s'_2 \succ \Gamma_2 \vdash s_2$.*

Proof. The proof proceeds by induction on the length of the inference of $s_1 \xrightarrow{\alpha} s_2$.
Base Step: We reason by case analysis on the axioms of the typed operational semantics.

(kill) By rule *(t-kill)*, $s_1 = s'_1 = \mathbf{kill}(k)$ and $\Gamma_1 = \Gamma'_1$. Thus, since $s_2 = \mathbf{0}$ and by rule *(t-nil)*, it is trivial to conclude, by letting $\Gamma_2 = \Gamma'_2 = \Gamma'_1$.

(inv) By rule *(t-inv)*, $s_1 = s_1 = p \bullet o!\{\overline{e}\}_r$. Thus, since $s_2 = \mathbf{0}$ and by rule *(t-nil)*, it is trivial to conclude, by letting $\Gamma_2 = \Gamma'_2 = \Gamma'_1$.

(rec) By hypothesis, $s_1 = p \bullet o?\bar{w}.s_2$ and $\Gamma'_1 \vdash p \bullet o?\bar{w}.s'_2 \succ \Gamma_1 \vdash p \bullet o?\bar{w}.s_2$. By the premise of rule *(t-rec)*, we can conclude that $\Gamma'_2 \vdash s'_2 \succ \Gamma_2 \vdash s_2$ with $\Gamma'_2 = \Gamma'_1 + \{x : \{p\}\}_{x \in \mathrm{fv}(\bar{w})}$ and $\Gamma_2 = \Gamma_1$.

Inductive Step: We reason by case analysis on the last applied inference rule of the typed operational semantics.

(com) By hypothesis, $s_1 = s_l \mid s_r$, $s'_2 = s'_l \mid s'_r$ and there exist two raw services s''_l and s''_r, and two type environments Γ_l and Γ_r, such that $\Gamma_1 = \Gamma_l + \Gamma_r$ and $\Gamma'_1 \vdash s''_l \mid s''_r \succ \Gamma_l + \Gamma_r \vdash s_l \mid s_r$. By the premises of rule *(t-par)*, $\Gamma'_1 \vdash s''_l \succ \Gamma_l \vdash s_l$ and $\Gamma'_1 \vdash s''_r \succ \Gamma_r \vdash s_r$. By the premises of rule *(com)*, $s_l \xrightarrow{(p \bullet o) \rhd \bar{w}} s'_l$ and $s_r \xrightarrow{(p \bullet o) \lhd \{\overline{v}\}_r} s'_r$, then, by induction, there exist two raw services s'''_l and s'''_r, and four type environment $\Gamma'_l, \Gamma''_l, \Gamma'_r$ and Γ''_r, such that $\Gamma'_1 \sqsubseteq \Gamma''_l, \Gamma'_1 \sqsubseteq \Gamma''_r, \Gamma''_l \vdash s'''_l \succ \Gamma'_l \vdash s'_l$ and $\Gamma''_r \vdash s'''_r \succ \Gamma'_r \vdash s'_r$. By the definition of the preorder $\sqsubseteq$, $\Gamma'_1 \sqsubseteq \Gamma''_l$ implies that there exists Γ such that $\Gamma'_1 + \Gamma = \Gamma''_l$, and, similarly, $\Gamma'_1 \sqsubseteq \Gamma''_r$ implies that there exists Γ' such that $\Gamma'_1 + \Gamma' = \Gamma''_r$. By repeated applications of Lemma 2, we obtain $\Gamma''_l + \Gamma' \vdash s'''_l \succ \Gamma'_l + \Gamma' \vdash s'_l$ and $\Gamma''_r + \Gamma \vdash s'''_r \succ \Gamma'_r + \Gamma \vdash s'_r$. Since $\Gamma''_l + \Gamma' = \Gamma'_1 + \Gamma + \Gamma' = \Gamma''_r + \Gamma$, by applying rule *(t-par)*, we can conclude.

(par_{conf}) - (par_{pass}) - (par_{kill}) These cases are similar to the previous one; the latter case relies on Lemma 3.

(choice) By hypothesis, $s_1 = g_1 + g_2$ and there exist g'_1 and g'_2, and two type environments Γ and Γ', such that $\Gamma_1 = \Gamma + \Gamma'$ and $\Gamma'_1 \vdash g'_1 + g'_2 \succ \Gamma + \Gamma' \vdash g_1 + g_2$. By the premise of rule *(t-sum)*, $\Gamma'_1 \vdash g'_1 \succ \Gamma \vdash g_1$. By the premises of rule *(choice)*, $g_1 \xrightarrow{\alpha} s_2$ and, by induction, we can conclude.

(prot) By hypothesis, $s_1 = \{\!|s|\!\}$ and there exists a raw service s' such that $\Gamma'_1 \vdash \{\!|s'|\!\} \succ \Gamma_1 \vdash \{\!|s|\!\}$. By the premise of rule *(t-prot)*, $\Gamma'_1 \vdash s' \succ \Gamma_1 \vdash s$. By the premise of rule *(prot)*, $s \xrightarrow{\alpha} s''$ such that $s_2 = \{\!|s''|\!\}$. By induction, there exist s''', Γ and Γ' such that $\Gamma' \vdash s''' \succ \Gamma \vdash s''$. Then, by rule *(t-prot)*, we can conclude.

(del_{sub}) By hypothesis, $s_1 = [\{x\}^r]\, s$ and there exists a raw services s'' such that $\Gamma'_1 \vdash [x]\, s'' \succ \Gamma_1 \vdash [\{x\}^r]\, s$. By the premise of rule *($t\text{-}del_{var}$)*, we obtain $\Gamma'_1, \{x : \emptyset\} \vdash s'' \succ \Gamma_1, \{x : r\} \vdash s$ (for simplicity, we assume $x \notin r$). By the premise of rule *(del_{sub})*, $s \xrightarrow{p \bullet o\, \lfloor \sigma \uplus \{x \mapsto \{v\}_{r'}\} \rfloor\, \bar{w}\, \overline{\{v'\}_{r''}}} s'$ with $s_2 = s' \cdot \{x \mapsto \{v\}_{r'}\}$. Hence, by Lemma 1, $\Gamma'_1 \cdot \{x \mapsto \{v\}_{r'}\} \vdash s'' \cdot \{x \mapsto \{v\}_{r'}\} \succ \Gamma_1 \cdot \{x \mapsto \{v\}_{r'}\} \vdash s \cdot \{x \mapsto \{v\}_{r'}\}$. Then, by definition of function *match* and rule *(com)*, $s \cdot \{x \mapsto \{v\}_{r'}\} \xrightarrow{p \bullet o\, \lfloor \sigma \rfloor\, \bar{w}\, \overline{\{v'\}_{r''}}} s' \cdot \{x \mapsto \{v\}_{r'}\}$ and, by induction, we can conclude.

(del_{kill}) By hypothesis, $s_1 = [k]\, s$ and there exists a raw services s'' such that $\Gamma'_1 \vdash [k]\, s'' \succ \Gamma_1 \vdash [k]\, s$. By the premise of rule *($t\text{-}del_{lab}$)*, $\Gamma'_1 \vdash s'' \succ \Gamma_1 \vdash s$. By the premises of rule *(del_{kill})*, $s \xrightarrow{\dagger k} s'$ with $s_2 = [k]\, s'$, then, by induction, there exist s''', Γ and Γ' such that $\Gamma' \vdash s''' \succ \Gamma \vdash s'$. Thus, by applying rule *($t\text{-}del_{lab}$)*, we can conclude.

(del$_{pass}$) This case is similar to the previous one.

(x$_{pass}$) By hypothesis, $s_1 = [\{x\}^r]\, s$ and there exists a raw services s'' such that $\Gamma'_1 \vdash [x]\, s'' \succ \Gamma_1 \vdash [\{x\}^r]\, s$. By the premise of rule *(t-del$_{var}$)*, we obtain $\Gamma'_1, \{x : \emptyset\} \vdash s'' \succ \Gamma_1, \{x : r\} \vdash s$ (for simplicity, we assume $x \notin r$). By the premise of rule *(del$_{pass}$)*, $s \xrightarrow{\alpha} s'$ with $s_2 = [\{x\}^r]\, s'$. By induction, there exist s''', Γ and Γ' such that $\Gamma' \vdash s''' \succ \Gamma \vdash s'$ and $\Gamma \sqsubseteq \Gamma_1, \{x : r\}$. Thus, we can obtain $\Gamma'', \{x : \emptyset\} \vdash s''' \succ \Gamma''', \{x : r\} \vdash s'$ for some Γ'' and Γ''', and, by applying rule *(t-del$_{var}$)*, we can conclude.

(cong) By straightforward induction and by relying on Lemma 4.

□

We can now easily prove that well-typedness is preserved along computations.

Corollary 1 (Subject Reduction). *If s is well-typed and $s \xrightarrow{\alpha} s'$ with $\alpha \in \{\dagger, \hat{n}\, \lfloor \emptyset \rfloor\, \bar{w}\, \overline{\{v\}_r}\}$, then s' is well-typed.*

To characterize the errors that our type system can capture we use predicate $\Uparrow$: $s \Uparrow$ holds true when s can immediately generate a runtime error. This happens when in an active context there is an invoke activity on a partner not included in the region annotation of some of the expressions argument of the activity. Formally, $\Uparrow$ is defined as the least predicate closed under the following rules

$$\frac{\exists\, r' \in \bar{r}\,.\ p \notin r'}{p \bullet o!\overline{\{e\}_r} \Uparrow} \qquad \frac{s \Uparrow}{\mathbb{A}[\![s]\!] \Uparrow} \qquad \frac{s \equiv s' \quad s \Uparrow}{s' \Uparrow}$$

We remark that the runtime errors that our type discipline can capture are related to the policies for the exchange of data. We skip such runtime errors as 'unproper use of variables' (e.g. in $x \bullet o!\bar{v}$ the variable x is not replaced by a partner name) that can be easily dealt with standard type systems. We can now prove that well-typed services do respect region annotations.

Theorem 2 (Type Safety). *If s is a well-typed service then $s \Uparrow$ holds false.*

Proof. The proof is straightforward. We prove, by rule induction, that if $s \Uparrow$ then s is not well-typed. Looking at the rules in Table 6, we see that the only case to be considered is when there exists $r' \in \bar{r}$ such that $p \notin r'$: in such a case, there is no way to infer a type for $p \bullet o!\overline{\{e\}_r}$ using *(t-inv)*. The other cases follow by induction. Suppose, for example, that $[n]\, s \Uparrow$ because $s \Uparrow$. By induction, s is not well-typed and therefore we can not use *(t-del$_{name}$)*, that is the only possible rule to infer that $[n]\, s$ is well-typed. □

We can finally conclude by stating that the type system and the operational semantics are *sound*.

Corollary 2 (Type Soundness). *Let s be a well-typed service. Then $s' \Uparrow$ holds false for every reduct s' of s.*

Proof. Corollary 1 can be repeatedly applied to prove that s' is well-typed, then Theorem 2 permits to conclude. □

4 Examples

In this section we illustrate two applications of our framework. The former example is a simplified but realistic electronic marketplace scenario inspired by W3C [27], the latter one is inspired to an automotive scenario studied in the SENSORIA project [**?**].

As a matter of notation, we will write $Z \triangleq s$ to assign a symbolic name Z to s.

4.1 An electronic marketplace

To show usefulness of our approach, we focus on the central part of the protocol where sensitive data are exchanged, i.e. we omit the initial bartering and the concluding interactions, and expand the part relative to the payment process.

Suppose a service *buyer* invokes a service *seller* to purchase some goods. Once *seller* has received an order request, it sends back the partner name of the service *credit_agency* to be used for the payment. *buyer* can then check the information on *credit_agency* and, possibly, confirm the payment by sending its credit card data to *seller*. In this case, *seller* forwards the received data to *credit_agency* and passes the order to the service *shipper*. In the end, the whole system is

$$EMP \triangleq buyer \mid credit_agency \mid [p_{sh}]\ (seller \mid shipper)$$

When fixing the policies for data exchange, services can (safely) assume that, at the outset, partner names p_s, p_{ca} and p_b are publicly available for invoking *seller*, *credit_agency* and *buyer*, respectively. Instead, the partner name p_{sh} for invoking *shipper* is private and only shared with *seller*. Of course, due to the syntactical restrictions, the 'locality' condition for partner names is preserved by the semantics. Thus, the initials assumptions remain true forever.

The *buyer* service is defined as

$$\begin{aligned} buyer \triangleq [id]\,(\ & p_s \cdot o_{ord}!\langle \{id\}_{\{p_s,p_b\}}, p_b, order\rangle \\ & \mid [x_{ca}]\, p_b \cdot o_{ca_info}?\langle id, x_{ca}\rangle. \\ & \qquad [p,o]\,(\ p \cdot o!\langle\rangle \mid p \cdot o?\langle\rangle.p_s \cdot o_{pay}!\langle \{id\}_{\{p_s,p_b\}}, \{cc_data\}_{\{p_s,x_{ca}\}}\rangle \\ & \qquad\qquad + \ p \cdot o?\langle\rangle.p_s \cdot o_{canc}!\langle \{id\}_{\{p_s,p_b\}}\rangle\ )\ ) \end{aligned}$$

The endpoint $p_s \cdot o_{ord}$ is used for invoking the seller service and transmitting the order together with the *buyer*'s partner name p_b. The (restricted) name *id* represents the order identifier and is used for correlating all those service interactions that logically form a same session relative to the processing of *order*. For example, the specification of *buyer* could be slightly modified to allow the service to simultaneously make multiple orders: of course, although all such parallel threads must use the same partner p_s to interact with *seller*, they can exploit different order identifiers as a means to correlate messages belonging to different interaction sessions. The type attached to *id* only allows *buyer* and *seller* to exchange and use it, since they are the only services that can receive along p_s and p_b. Instead, p_b comes without an attached policy, since it is publicly known (it is transmitted to indicate the service making the invocation for the call-back operation). For simplicity, also *order* has no attached policy; thus, it could be later on communicated to any other service. Variable x_{ca} is used to store the partner name of the credit

agency service to be used to possibly finalize the purchase and also to implement the policy for *buyer*'s credit card data. After the information on the credit agency service are verified, *buyer* sends a message to *seller* either to confirm or to cancel the order. This is simply modelled as an internal non-deterministic choice, by exploiting the private endpoint $p \cdot o$ (a more precise model can be obtained by exploiting the encodings shown in [22]).

The *seller* service is defined as

$$
\begin{aligned}
seller \triangleq *\,[x_b, x_{id}, x_{ord}, k]\; & p_s \cdot o_{ord}?\langle x_{id}, x_b, x_{ord}\rangle. \\
& (\; x_b \cdot o_{ca_info}!\langle \{x_{id}\}_{\{x_b\}}, p_{ca}\rangle \\
& \quad \mid\; [x_{cc}]\; p_s \cdot o_{pay}?\langle x_{id}, x_{cc}\rangle.(\; p_{ca} \cdot o_{cr_req}!\langle x_{ord}, \{x_{cc}\}_{\{p_{ca}\}}\rangle \\
& \qquad\qquad\qquad\qquad \mid\; p_{sh} \cdot o_{sh_req}!\langle x_{ord}\rangle\;) \\
& \quad \mid\; p_s \cdot o_{canc}?\langle x_{id}\rangle.\mathbf{kill}(k)\;)
\end{aligned}
$$

Once *seller* receives an order along $p_s \cdot o_{ord}$, it creates one specific instance that sends back to *buyer* (via x_b) the partner name p_{ca} of the credit agency service where the payment will be made. Whenever the seller instance receives the credit card data correlated to x_{id}, it forwards them to *credit_agency* and passes the order to the (internal) shipper service. Instead, if *buyer* demands cancellation of the order, the corresponding instance of *seller* is immediately terminated. Name k is used to delimit the effect of the kill activity only to the relevant instance.

The remaining two services are defined as

$$
\begin{aligned}
credit_agency &\triangleq *\,[x, y]\; p_{ca} \cdot o_{cr_req}?\langle x, y\rangle.\text{“}\texttt{execute_the_payment}\text{”} \\
shipper &\triangleq *\,[z]\; p_{sh} \cdot o_{sh_req}?\langle z\rangle.\text{“}\texttt{process_the_order}\text{”}
\end{aligned}
$$

Let now consider the type inference phase. Service *seller* gets annotated as follows:

$$
\begin{aligned}
seller' \triangleq *\,[\{x_b\}^{\{p_s\}}]\,[\{x_{id}\}^{\{p_s, x_b\}}, \{x_{ord}\}^{\top}, k]\; & p_s \cdot o_{ord}?\langle x_{id}, x_b, x_{ord}\rangle. \\
& (\; x_b \cdot o_{ca_info}!\langle \{x_{id}\}_{\{x_b\}}, p_{ca}\rangle \\
& \quad \mid\; [\{x_{cc}\}^{\{p_s, p_{ca}\}}]\; p_s \cdot o_{pay}?\langle x_{id}, x_{cc}\rangle.(\; p_{ca} \cdot o_{cr_req}!\langle x_{ord}, \{x_{cc}\}_{\{p_{ca}\}}\rangle \\
& \qquad\qquad\qquad\qquad \mid\; p_{sh} \cdot o_{sh_req}!\langle x_{ord}\rangle\;) \\
& \quad \mid\; p_s \cdot o_{canc}?\langle x_{id}\rangle.\mathbf{kill}(k)\;)
\end{aligned}
$$

The type inference has the task of checking consistency of region annotations of the arguments occurring within invoke activities and that of deriving the annotations for variable declarations. As regards consistency, there are only two explicitly typed expressions used as arguments of invoke activities, i.e. x_{id} and x_{cc}, and their types $\{x_b\}$ and $\{p_{ca}\}$ satisfy the consistency constraint (see rule *(t-inv)*). The remaining expressions occurring as arguments of invoke activities, i.e. the only x_{ord}, have implicitly assigned type $\top$ (indeed, recall that we assumed that an untagged e stands for $\{e\}_{\top}$) and are thus trivially consistent. As regards type derivation, when a variable is put in the environment (rule *(t-del$_{var}$)*), it is assigned type $\emptyset$. Later on, when a variable is used as an argument of an invoke or receive, its type can possibly be enriched (rules *(t-inv)* and *(t-rec)*). Thus, at the end of the inference, declaration of variable x_b, that is only used in $p_s \cdot o_{ord}?\langle x_{id}, x_b, x_{ord}\rangle$, will have assigned region $\{p_s\}$ (application of rule *(t-rec)*). Instead, declaration of x_{ord} has assigned type $\top$ (rule *(t-inv)* is used) while that of x_{cc} has assigned type $\{p_s, p_{ca}\}$ and, similarly, declaration of x_{id} gets annotated with $\{p_s, x_b\}$

(in both cases rules *(t-inv)* and *(t-rec)* are used). Notably, in *seller′*, delimitation $[\{x_b\}^{p_s}]$ does not commute any longer with delimitations $[\{x_{id}\}^{\{p_s,x_b\}}, \{x_{ord}\}^{\top}, k]$ (otherwise the service would become opened).

The variable declarations of the other services are annotated in a trivial way: x_{ca} with $\{p_b\}$, x and y with $\{p_{ca}\}$, and z with $\{p_{sh}\}$ (we assume that *credit_agency* and *shipper* do not re-transmit the received data). Thus, if we call *buyer′*, *credit_agency′* and *shipper′* the other typed services, then the system resulting from the type inference is

$$buyer' \mid credit_agency' \mid [p_{sh}]\ (seller' \mid shipper')$$

After some computation steps, the system can become

$$\begin{array}{l}
[id]\,(\ p_s \cdot o_{pay}!\langle\{id\}_{\{p_s,p_b\}}, \{cc_data\}_{\{p_s,p_{ca}\}}\rangle \ \mid\ [p_{sh}]\,(\,seller' \ \mid \\
\qquad [k, \{x_{cc}\}^{\{p_s,p_{ca}\}}]\,(\ p_s \cdot o_{pay}?\langle id, x_{cc}\rangle.(\ p_{ca} \cdot o_{cr_req}!\langle order, \{x_{cc}\}_{\{p_{ca}\}}\rangle \\
\qquad\qquad\qquad\qquad\qquad\qquad \mid\ p_{sh} \cdot o_{sh_req}!\langle order\rangle\ ) \\
\qquad\qquad\qquad\qquad \mid\ p_s \cdot o_{canc}?\langle id\rangle.\mathbf{kill}(k)\ ) \\
\qquad \mid\ *\,[\{x\}^{\{p_{ca}\}}, \{y\}^{\{p_{ca}\}}]\ p_{ca} \cdot o_{cr_req}?\langle x, y\rangle.\text{“}\texttt{execute_the_payment}\text{”} \\
\qquad \mid\ *\,[\{z\}^{\{p_{sh}\}}]\ p_{sh} \cdot o_{sh_req}?\langle z\rangle.\text{“}\texttt{process_the_order}\text{”}\,)\,)
\end{array}$$

Thus, after *buyer′* sends the credit card data, we get

$$\begin{array}{l}
[id, p_{sh}]\,(\ seller' \\
\qquad \mid\ [k]\,(\ p_{ca} \cdot o_{cr_req}!\langle order, \{cc_data\}_{\{p_{ca}\}}\rangle \mid p_{sh} \cdot o_{sh_req}!\langle order\rangle \\
\qquad\qquad \mid\ p_s \cdot o_{canc}?\langle id\rangle.\mathbf{kill}(k)\ ) \\
\qquad \mid\ *\,[\{x\}^{\{p_{ca}\}}, \{y\}^{\{p_{ca}\}}]\ p_{ca} \cdot o_{cr_req}?\langle x, y\rangle.\text{“}\texttt{execute_the_payment}\text{”} \\
\qquad \mid\ *\,[\{z\}^{\{p_{sh}\}}]\ p_{sh} \cdot o_{sh_req}?\langle z\rangle.\text{“}\texttt{process_the_order}\text{”}\,)
\end{array}$$

At this point, *seller′* can safely communicate credit card data of *buyer′* to *credit_agency′* and, then, forward the order to *shipper′*.

Suppose now that *seller′* also contains such a malicious invoke as $p_{sh} \cdot o!\langle \ldots, \{x_{cc}\}_r, \ldots\rangle$. In order to successfully pass the type inference phase, it should be that $p_{sh} \in r$ (otherwise rule *(t-inv)* could not be applied). Therefore, in the resulting typed service we would have the variable declaration $[\{x_{cc}\}^{r'}]$, with $r \subseteq r'$. Now, communication with *buyer′* would be blocked by the runtime checks because the datum is tagged as $\{cc_data\}_{\{p_s,p_{ca}\}}$, and $p_{sh} \in r \subseteq r'$ implies that $r' \nsubseteq \{p_s, p_{ca}\}$.

4.2 Automotive security

A car manufacturer offers an integrated infotainment system capable of forwarding the GPS coordinates of a car to relevant recipient services. We consider the case of a system designed to allow drivers to assist each other on deserted roads until professional service can arrive at the scene. The system is based on an automobile association monitoring the location of all trustworthy members' cars and contacting the drivers nearest to an accident to ask them to provide first aid. The driver in trouble is assured that information about his location cannot become available to unauthorized users (e.g. paparazzi photographers). We suppose that customers of the automobile association service may not be trustworthy members. Thus, a driver can play both helper and customer roles, or only the latter one.

We consider the following (simplified) scenario

$$\mathrm{C} \mid \mathrm{T}_1 \mid \mathrm{T}_2 \mid \mathrm{T}_3 \mid \mathrm{A}$$

where C is the service installed in the customer's car, T_1, T_2, and T_3 are services installed in the trustworthy members' cars, and A is the service provided by the automobile association. For the sake of simplicity, we do not consider the customer services installed in the member's cars.

The *customer* service is defined as

$$\mathrm{C} \triangleq [p_{sys}]\,(\,p_{sys} \bullet o_{accident}!\langle gps\rangle \mid [y_{\mathrm{T}}]\,[y_{gps}]\,p_{sys} \bullet o_{accident}?\langle y_{gps}\rangle.\,\mathrm{C}_{\mathrm{findHelp}}\,)$$

$$\begin{aligned}\mathrm{C}_{\mathrm{findHelp}} \triangleq (\,& p_{\mathrm{A}} \bullet o_{carAccident}!\langle \{p_{\mathrm{C}}\}_{\{p_{\mathrm{A}}\}}, \{y_{gps}\}_{\{p_{\mathrm{A}}\}}\rangle \\ & \mid [y_{id}]\,(\,p_{\mathrm{C}} \bullet o_{helpNotFound}?\langle\rangle \\ & \qquad + \\ & \qquad p_{\mathrm{C}} \bullet o_{helpFound}?\langle y_{id}, y_{\mathrm{T}}\rangle.p_{\mathrm{A}} \bullet o_{position}!\langle y_{id}, \{y_{gps}\}_{\{p_{\mathrm{A}}, y_{\mathrm{T}}\}}\rangle\,)\,)\end{aligned}$$

When the diagnostic system of the customer's car detects that an accident has happened, it activates the recovery service by sending the current GPS coordinates of the car along the private endpoint $p_{sys} \bullet o_{accident}$. Then, it invokes the automobile association service A, by using the endpoint $p_{\mathrm{A}} \bullet o_{carAccident}$, and transmits its partner name (i.e. p_{C}, for the reply) and its coordinates. The type attached to those data only allows A to exchange and use them, since A is the only service that can receive along p_{A}. Afterwards, C waits a reply from A. If no trustworthy helper has been found, then the driver has to wait until professional service arrives at the scene. In case A has found a trustworthy helper T_i, C transmits its coordinates again, this time with region enlarged with the partner name of T_i. Deliberately we do not model C as a persistent service; instead, we assume that it will be reinstalled at the end of the repairing activity.

A *trustworthy member* service is defined as

$$\begin{aligned}\mathrm{T}_i \triangleq [\hat{m}]\,(\,& \hat{m}!\langle\rangle \mid *\ \hat{m}?\langle\rangle.[x_{id}]\,p_{\mathrm{T}i} \bullet o_{check}?\langle x_{id}\rangle. \\ & [\hat{n}]\,(\,\hat{n}!\langle\rangle \mid \hat{n}?\langle\rangle.(\,p_{\mathrm{A}} \bullet o_{refuse}!\langle x_{id}\rangle \mid \hat{m}!\langle\rangle\,) \\ & \qquad + \\ & \qquad \hat{n}?\langle\rangle.(\,p_{\mathrm{A}} \bullet o_{ok}!\langle x_{id}\rangle \\ & \qquad\qquad \mid [x_{gps}]\,p_{\mathrm{T}i} \bullet o_{fwdPos}?\langle x_{gps}\rangle. \\ & \qquad\qquad\quad \text{“}\texttt{go to}\ x_{gps}\ \texttt{and provide first aid}\text{”}. \\ & \qquad\qquad\quad \hat{m}!\langle\rangle\,)\,)\end{aligned}$$

This is a persistent service that, however, can serve requests only sequentially (i.e. one at time). When the automobile association contacts one of its members, firstly asks him (by invoking the operation o_{check}) if he is willing to provide first aid to a near driver. Service T_i can accept or refuse by replying on the operations o_{ok} or o_{refuse}, respectively. For the sake of simplicity, we model the above decision by means of an internal non-deterministic choice exploiting the private endpoint $\hat{n}$. In case of negative reply, T_i is immediately reactivated (by means of $\hat{m}!\langle\rangle$). Otherwise, i.e. in case of positive reply, the service waits for the position of the driver in trouble from A, goes to the accident scene, provides first aid and, finally, reactivates itself.

The *automobile association* service is defined as

$$
\begin{aligned}
\mathrm{A} \triangleq\ & [o_{find}, o_{found}, o_{notFound}] \\
& (\,*\,[z_{\mathrm{C}}, z_{gps}]\; p_{\mathrm{A}} \cdot o_{carAccident}?\langle z_{\mathrm{C}}, z_{gps}\rangle.\; [id]\; \mathrm{A}_{\mathrm{inst}} \\
& \quad \mid\ *\,[z_{id}, z_{pos}]\; p_{\mathrm{A}} \cdot o_{find}?\langle z_{id}, z_{pos}\rangle. \\
& \qquad\qquad \text{“}\texttt{find in DB and reply on } o_{found} \texttt{ or } o_{notFound}\text{”})
\end{aligned}
$$

$$
\begin{aligned}
\mathrm{A}_{\mathrm{inst}} \triangleq\ & [\hat{n}]\,(\,\hat{n}!\langle\rangle \mid *\,\hat{n}?\langle\rangle.\,(\,p_{\mathrm{A}} \cdot o_{find}!\langle id, z_{gps}\rangle \\
& \qquad \mid [z'_T]\,(\,p_{\mathrm{A}} \cdot o_{notFound}?\langle id\rangle.\, z_{\mathrm{C}} \cdot o_{helpNotFound}!\langle\rangle \\
& \qquad\qquad + \\
& \qquad\qquad p_{\mathrm{A}} \cdot o_{found}?\langle id, z'_{\mathrm{T}}\rangle.\,(\,z'_{\mathrm{T}} \cdot o_{check}!\langle id\rangle \\
& \qquad\qquad\quad \mid p_{\mathrm{A}} \cdot o_{refuse}?\langle id\rangle.\,\hat{n}!\langle\rangle \\
& \qquad\qquad\quad + \\
& \qquad\qquad\quad p_{\mathrm{A}} \cdot o_{ok}?\langle id\rangle.\,[z_{\mathrm{T}}, o]\,(\,p_{\mathrm{A}} \cdot o!\langle z'_{\mathrm{T}}\rangle \\
& \qquad\qquad\qquad \mid p_{\mathrm{A}} \cdot o?\langle z_{\mathrm{T}}\rangle. \\
& \qquad\qquad\qquad (\,z_{\mathrm{C}} \cdot o_{helpFound}!\langle id, \{z_{\mathrm{T}}\}_{\{z_{\mathrm{C}}\}}\rangle \\
& \qquad\qquad\qquad\quad \mid [z'_{gps}]\; p_{\mathrm{A}} \cdot o_{position}?\langle id, z'_{gps}\rangle. \\
& \qquad\qquad\qquad\qquad\quad z_{\mathrm{T}} \cdot o_{fwdPos}!\langle \{z'_{gps}\}_{\{z_{\mathrm{T}}\}}\rangle\,)\,)\,)\,)\,)\,)
\end{aligned}
$$

This service is composed of two persistent subservices, both capable of receiving along p_{A}. Service A is publicly invocable and can interact with customers and members services, other than with the ‘internal’ service $\mathrm{A}_{\mathrm{inst}}$. This latter service, instead, can only be invoked by A (indeed, all the operations used by the service, i.e. $o_{find}, o_{found}, o_{notFound}$, are restricted and this prevents them to be invoked from the outside) and has the task of looking up in a database the member nearest to given a location and replying accordingly.

Differently from C and T_i, A and $\mathrm{A}_{\mathrm{inst}}$ can serve requests concurrently and, thus, exploit correlation mechanisms to route each received message to its right instance. Indeed, when a customer request is received along $p_{\mathrm{A}} \cdot o_{carAccident}$, a new specific service instance is created that is uniquely identified by a fresh correlation identifier *id*; this identifier is generated by the instance itself and is communicated to every other involved service. The instance then tries to find a helper member that is near to the accident scene and is willing to provide first aid. This is done by repeatedly looking up the database ($p_{\mathrm{A}} \cdot o_{find}!\langle id, z_{gps}\rangle$) and checking willingness of the returned trustworthy member ($z'_{\mathrm{T}} \cdot o_{check}!\langle id\rangle$), until either one available member is found (i.e. the member replies on operation o_{ok}) or the looking up in the database fails (i.e. a reply on $o_{notFound}$ is returned). Technically, the cycle is implemented by means of the replication operator and the private endpoint $\hat{n}$. If no available helper is found, the instance notifies this to the customer ($z_{\mathrm{C}} \cdot o_{helpNotFound}!\langle\rangle$) and terminates; otherwise, it sends the partner name of the found helper to the customer ($z_{\mathrm{C}} \cdot o_{helpFound}!\langle id, \{z_{\mathrm{T}}\}_{\{z_{\mathrm{C}}\}}\rangle$), waits for the customer location (whose region now also include the partner name of the helper) and forwards it to the helper ($z_{\mathrm{T}} \cdot o_{fwdPos}!\langle \{z'_{gps}\}_{\{z_{\mathrm{T}}\}}\rangle$).

Notably, the only relevant types are that attached to the partner name of the found helper, which allows only C to receive this partner, and that attached to the customer location, which allows only the found helper to receive such location.

Let now consider the type inference phase. Service C gets annotated as follows:

$$
\begin{aligned}
\mathrm{C}' \triangleq\ & [p_{sys}]\,(\,p_{sys} \cdot o_{accident}!\langle gps\rangle \\
& \quad \mid [\{y_{\mathrm{T}}\}^{\{p_{\mathrm{C}}\}}]\,[\{y_{gps}\}^{\{p_{sys},p_{\mathrm{A}},y_{\mathrm{T}}\}}]\; p_{sys} \cdot o_{accident}?\langle y_{gps}\rangle.\ \mathrm{C}'_{\mathrm{findHelp}}\,)
\end{aligned}
$$

$$
\begin{aligned}
\mathrm{C}'_{\mathrm{findHelp}} \triangleq\ & (\,p_{\mathrm{A}} \cdot o_{carAccident}!\langle \{p_{\mathrm{C}}\}_{\{p_{\mathrm{A}}\}}, \{y_{gps}\}_{\{p_{\mathrm{A}}\}}\rangle \\
& \quad \mid [\{y_{id}\}^{\top}]\,(\,p_{\mathrm{C}} \cdot o_{helpNotFound}?\langle\rangle \\
& \qquad\qquad + \\
& \qquad\qquad p_{\mathrm{C}} \cdot o_{helpFound}?\langle y_{id}, y_{\mathrm{T}}\rangle.\ p_{\mathrm{A}} \cdot o_{position}!\langle y_{id}, \{y_{gps}\}_{\{p_{\mathrm{A}},y_{\mathrm{T}}\}}\rangle\,)\,)
\end{aligned}
$$

As regards consistency of region annotations, the only explicitly typed expressions used as arguments of invoke activities are p_{C} and y_{gps} (this latter is used twice); in any case, their types ($\{p_{\mathrm{A}}\}$, $\{p_{\mathrm{A}}\}$ and $\{p_{\mathrm{A}}, y_{\mathrm{T}}\}$, resp.) satisfy the consistency constraint (see rule *(t-inv)*). The remaining expressions occurring as arguments of invoke activities, i.e. gps and y_{id}, have implicitly assigned type $\top$ and, thus, the consistency constraint is trivially satisfied. As regards derivation of types for variable declarations, at the end of the inference, declaration of variable y_{T}, that is only used in $p_{\mathrm{C}} \cdot o_{helpFound}?\langle y_{id}, y_{\mathrm{T}}\rangle$ and in the region of y_{gps}, will have assigned region $\{p_{\mathrm{C}}\}$ (application of rule *(t-rec)*). Instead, y_{gps} will have assigned type $\{p_{sys}, p_{\mathrm{A}}, y_{\mathrm{T}}\}$ (rules *(t-inv)* and *(t-rec)*), while y_{id} will have assigned type $\top$ (rule *(t-inv)*). Notably, in C', delimitation $[\{y_{\mathrm{T}}\}^{\{p_{\mathrm{C}}\}}]$ does not commute any longer with delimitation $[\{y_{gps}\}^{\{p_{sys},p_{\mathrm{A}},y_{\mathrm{T}}\}}]$ (otherwise the service would become opened).

Services T_i get annotated as follows:

$$
\begin{aligned}
\mathrm{T}'_i \triangleq\ & [\hat{m}]\,(\hat{m}!\langle\rangle \mid *\ \hat{m}?\langle\rangle.[\{x_{id}\}^{\top}]\; p_{\mathrm{T}i} \cdot o_{check}?\langle x_{id}\rangle. \\
& \qquad [\hat{n}]\,(\,\hat{n}!\langle\rangle \mid \hat{n}?\langle\rangle.(\,p_{\mathrm{A}} \cdot o_{refuse}!\langle x_{id}\rangle \mid \hat{m}!\langle\rangle\,) \\
& \qquad\qquad + \\
& \qquad\qquad \hat{n}?\langle\rangle.(\,p_{\mathrm{A}} \cdot o_{ok}!\langle x_{id}\rangle \\
& \qquad\qquad\quad \mid [\{x_{gps}\}^{\{p_{\mathrm{T}i}\}}]\; p_{\mathrm{T}i} \cdot o_{fwdPos}?\langle x_{gps}\rangle. \\
& \qquad\qquad\qquad \text{“}\texttt{go to}\ x_{gps}\ \texttt{and provide first aid}\text{”}. \\
& \qquad\qquad\qquad \hat{m}!\langle\rangle\,)\,)
\end{aligned}
$$

In this case, derivation of types for variables is trivial: x_{id} has type $\top$ and x_{gps} has type $\{p_{\mathrm{T}i}\}$ (we assume that T_i does not re-transmit the GPS data store in x_{gps}).

Service A and $\mathrm{A}_{\mathrm{inst}}$ get annotated as follows:

$$
\begin{aligned}
\mathrm{A}' \triangleq\ & [o_{find}, o_{found}, o_{notFound}] \\
& (\,*\,[\{z_{\mathrm{C}}\}^{\{p_{\mathrm{A}}\}}, \{z_{gps}\}^{\{p_{\mathrm{A}}\}}]\; p_{\mathrm{A}} \cdot o_{carAccident}?\langle z_{\mathrm{C}}, z_{gps}\rangle.\ [id]\ \mathrm{A}'_{\mathrm{inst}} \\
& \quad \mid\ *\,[\{z_{id}\}^{\{p_{\mathrm{A}}\}}, \{z_{pos}\}^{\{p_{\mathrm{A}}\}}]\; p_{\mathrm{A}} \cdot o_{find}?\langle z_{id}, z_{pos}\rangle. \\
& \qquad\qquad \text{“}\texttt{find in DB and reply on}\ o_{found}\ \texttt{or}\ o_{notFound}\text{”}\,)
\end{aligned}
$$

$$
\begin{aligned}
\mathrm{A}'_{\mathrm{inst}} \triangleq\ & [\hat{n}]\,(\,\hat{n}!\langle\rangle \mid *\,\hat{n}?\langle\rangle.\,(\,p_{\mathrm{A}} \cdot o_{find}!\langle id, z_{gps}\rangle \\
& \qquad \mid [\{z'_{\mathrm{T}}\}^{\{p_{\mathrm{A}}\}}]\,(\,p_{\mathrm{A}} \cdot o_{notFound}?\langle id\rangle.\ z_{\mathrm{C}} \cdot o_{helpNotFound}!\langle\rangle \\
& \qquad\qquad + \\
& \qquad\qquad p_{\mathrm{A}} \cdot o_{found}?\langle id, z'_{\mathrm{T}}\rangle.\,(z'_{\mathrm{T}} \cdot o_{check}!\langle id\rangle \\
& \qquad\qquad\quad \mid p_{\mathrm{A}} \cdot o_{refuse}?\langle id\rangle.\ \hat{n}!\langle\rangle \\
& \qquad\qquad\quad + \\
& \qquad\qquad\quad p_{\mathrm{A}} \cdot o_{ok}?\langle id\rangle.\ [\{z_{\mathrm{T}}\}^{\{p_{\mathrm{A}},z_{\mathrm{C}}\}}, o]\,(\,p_{\mathrm{A}} \cdot o!\langle z'_{\mathrm{T}}\rangle \\
& \qquad\qquad\qquad \mid p_{\mathrm{A}} \cdot o?\langle z_{\mathrm{T}}\rangle.(\,z_{\mathrm{C}} \cdot o_{helpFound}!\langle id, \{z_{\mathrm{T}}\}_{\{z_{\mathrm{C}}\}}\rangle \\
& \qquad\qquad\qquad\quad \mid [\{z'_{gps}\}^{\{p_{\mathrm{A}},z_{\mathrm{T}}\}}]\; p_{\mathrm{A}} \cdot o_{position}?\langle id, z'_{gps}\rangle. \\
& \qquad\qquad\qquad\qquad z_{\mathrm{T}} \cdot o_{fwdPos}!\langle\{z'_{gps}\}_{\{z_{\mathrm{T}}\}}\rangle\,)\,)\,)\,)\,)\,)
\end{aligned}
$$

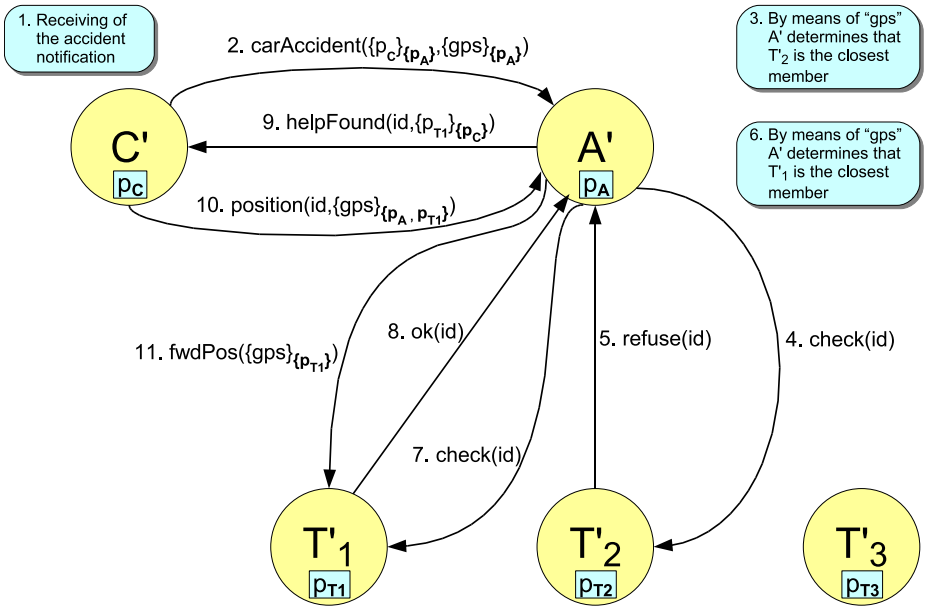

Fig. 1. Automotive security scenario

Variables z_{C}, z_{gps}, z_{id}, z_{pos}, z'_{T} are trivially annotated with type $\{p_{\mathrm{A}}\}$, because A does not re-transmit the corresponding data, while variables z_{T} and z'_{gps} are annotated with $\{p_{\mathrm{A}}, z_{\mathrm{C}}\}$ and $\{p_{\mathrm{A}}, z_{\mathrm{T}}\}$, respectively.

Therefore, the system resulting from the type inference is

$$\mathrm{C}' \mid \mathrm{T}'_1 \mid \mathrm{T}'_2 \mid \mathrm{T}'_3 \mid \mathrm{A}'$$

Figure 1 shows a possible evolution of the scenario, where T'_3 is too far from the accident scene, T'_2 is not available, and T'_1 accepts to provide first aid to C'.

For example, after the computation steps 1–9 described in Figure 1, the system becomes

$$
\begin{array}{l}
p_{\mathrm{A}} \bullet o_{position}!\langle id, \{gps\}_{\{p_{\mathrm{A}}, p_{\mathrm{T1}}\}}\rangle \\
\mid \\
[o_{find}, o_{found}, o_{notFound}] \\
(\,*\,[\{z_{\mathrm{C}}\}^{\{p_{\mathrm{A}}\}}, \{z_{gps}\}^{\{p_{\mathrm{A}}\}}]\; p_{\mathrm{A}} \bullet o_{carAccident}?\langle z_{\mathrm{C}}, z_{gps}\rangle.\ \cdots \\
\quad \mid *\,[\{z_{id}\}^{\{p_{\mathrm{A}}\}}, \{z_{pos}\}^{\{p_{\mathrm{A}}\}}]\; p_{\mathrm{A}} \bullet o_{find}?\langle z_{id}, z_{pos}\rangle.\ \cdots \\
\quad \mid [\{z'_{gps}\}^{\{p_{\mathrm{A}}, p_{\mathrm{T1}}\}}]\; p_{\mathrm{A}} \bullet o_{position}?\langle id, z'_{gps}\rangle.\ p_{\mathrm{T1}} \bullet o_{fwdPos}!\langle \{z'_{gps}\}_{\{p_{\mathrm{T1}}\}}\rangle\,) \\
\mid \\
[\hat{m}]\,(*\,\hat{m}?\langle\rangle.\cdots \\
\qquad \mid [\{x_{gps}\}^{\{p_{\mathrm{T1}}\}}]\; p_{\mathrm{T1}} \bullet o_{fwdPos}?\langle x_{gps}\rangle.\ \text{“}\texttt{go to}\ x_{gps}\ \texttt{and provide first aid}\text{”}.\,\hat{m}!\langle\rangle\,) \\
\mid \\
\mathrm{T}_2 \mid \mathrm{T}_3
\end{array}
$$

At this point, the customer can safely communicate its location to the automobile association (namely, rule *(del$_{sub}$)* of the operational semantics can be applied because the region annotations of gps and z'_{gps} are both $\{p_{\mathrm{A}}, p_{\mathrm{T1}}\}$ and, hence, do comply), and the system becomes

$$
\begin{array}{l}
[o_{find}, o_{found}, o_{notFound}] \\
(\,*\,[\{z_{\mathrm{C}}\}^{\{p_{\mathrm{A}}\}}, \{z_{gps}\}^{\{p_{\mathrm{A}}\}}]\; p_{\mathrm{A}} \bullet o_{carAccident}?\langle z_{\mathrm{C}}, z_{gps}\rangle.\ \cdots \\
\quad \mid *\,[\{z_{id}\}^{\{p_{\mathrm{A}}\}}, \{z_{pos}\}^{\{p_{\mathrm{A}}\}}]\; p_{\mathrm{A}} \bullet o_{find}?\langle z_{id}, z_{pos}\rangle.\ \cdots \\
\quad \mid p_{\mathrm{T1}} \bullet o_{fwdPos}!\langle \{gps\}_{\{p_{\mathrm{T1}}\}}\rangle\,) \\
\mid \\
[\hat{m}]\,(*\,\hat{m}?\langle\rangle.\cdots \\
\qquad \mid [\{x_{gps}\}^{\{p_{\mathrm{T1}}\}}]\; p_{\mathrm{T1}} \bullet o_{fwdPos}?\langle x_{gps}\rangle.\text{“}\texttt{go to } x_{gps} \texttt{ and provide first aid}\text{”}.\,\hat{m}!\langle\rangle\,) \\
\mid \\
\mathrm{T}_2 \;\mid\; \mathrm{T}_3
\end{array}
$$

Now, the automobile association can safely forward the customer location, because the region annotations of gps and x_{gps} do comply.

5 Concluding remarks

We have introduced a first analytical tool for checking that COWS specifications enjoy some desirable properties concerning the partners, and hence the services, that can safely access any given datum and, in that respect, do not manifest unexpected behaviors. Our type system is quite simple: types are just sets and operations on types are union, intersection, subset inclusion, etc. The language operational semantics only involves types in efficiently implementable checks, i.e. subset inclusions. While implementation of our framework is currently in progress, we are also working on the definition of a completely static variant where all dynamic checks have been moved to the static phase.

The types used in this paper are essentially inspired by the ‘region types’ for Confined-λ of [15] and for global computing calculi of [9]. There are however some noticeable differences. In fact, COWS permits describing not necessarily distributed systems and exchanging heterogeneous data along endpoints, which calls for a more dynamic typing mechanism than communication channels. Moreover, COWS permits annotating only the relevant data while Confined-λ requires typing any constant, function and channel. The group types, originally proposed for the Ambients calculus [8] and then recast to the π-calculus [7], have purposes similar to our region annotations, albeit they are only used for constraining the exchanges of ambient and channel names. Confinement has been also explored in the context of Java, and related calculi, for confining classes and objects within specific packages [30, 32].

More expressive type disciplines based, e.g., on session types and behavioural types are emerging as powerful tools for taking into account behavioural and non-functional properties of computing systems. In the case of services, they could permit to express and enforce many relevant policies for, e.g., regulating resources usage, constraining the sequences of messages accepted by services, ensuring service interoperability and compositionality, guaranteeing absence of deadlock in service composition, checking that

interaction obeys a given protocol. Some of the studies developed for the π-calculus (see e.g. [6, 31, 16, 14, 17]) are promising starting points, but they need non trivial adaptations to deal with all COWS peculiar features. For example, one of the major problems we envisage concerns the treatment of killing and protection activities, that are not commonly used in process calculi.

Many efforts have been devoted to develop analytical tools for SOC foundational languages. Some works study mechanisms for comparing global descriptions (i.e. choreographies) and local descriptions (i.e. orchestrations) of a same system. Means to check conformance of these different views have been defined in [5] and, by relying on session types, in [6]. COWS, instead, only considers service orchestration and focuses on modelling the dynamic behaviour of services without the limitations possibly introduced by a layer of choreography. Some other works [19, 23] have concentrated on modelling web transactions and on studying their properties in programming languages based on the π-calculus, while [4, 3] formalize long running transactions with special care for the *Sagas* mechanism [10]. A type system specifying security policies for orchestration has been introduced in [1] for a very basic formalism based on the λ-calculus. Finally, a type system for checking compliance between (simplified) WS-BPEL terms and the associated WSDL documents has been defined in [21].

References

1. M. Bartoletti, P. Degano, and G. Ferrari. Security Issues in Service Composition. In *Proceedings of FMOODS 2006, 8th IFIP International Conference on Formal Methods for Open Object-Based Distributed Systems*, volume 4037 of *Lecture Notes in Computer Science*, pages 1–16, 2006.
2. A. Brogi, C. Canal, E. Pimentel, and A. Vallecillo. Formalizing web service choreographies. *ENTCS*, 105:73–94, 2004.
3. R. Bruni, M. Butler, C. Ferreira, T. Hoare, H. Melgratti, and U. Montanari. Comparing two approaches to compensable flow composition. In *CONCUR*, volume 3653 of *LNCS*, pages 383–397. Springer, 2005.
4. R. Bruni, H.C. Melgratti, and U. Montanari. Theoretical foundations for compensations in flow composition languages. In *POPL*, pages 209–220. ACM, 2005.
5. N. Busi, R. Gorrieri, C. Guidi, R. Lucchi, and G. Zavattaro. Choreography and orchestration conformance for system design. In *COORDINATION*, volume 4038 of *LNCS*, pages 63–81. Springer, 2006.
6. M. Carbone, K. Honda, and N. Yoshida. A calculus of global interaction based on session types. In *DCM'06*, 2006. To appear as ENTCS.
7. L. Cardelli, G. Ghelli, and A. D. Gordon. Secrecy and group creation. *Inf. Comput.*, 196(2):127–155, 2005.
8. L. Cardelli, G. Ghelli, and A.D. Gordon. Types for the ambient calculus. *Inf. Comput.*, 177(2):160–194, 2002.
9. R. De Nicola, D. Gorla, and R. Pugliese. Confining data and processes in global computing applications. *Science of Computer Programming*, 63:57–87, 2006.
10. H. Garcia-Molina and K. Salem. Sagas. In *SIGMOD*, pages 249–259. ACM Press, 1987.
11. P. Geguang, Z. Xiangpeng, W. Shuling, and Q. Zongyan. Towards the semantics and verification of bpel4ws. In *WLFM 2005*. Elsevier, 2005.

12. H. Goguen. Typed operational semantics. In Mariangiola Dezani-Ciancaglini and Gordon D. Plotkin, editors, *Typed Lambda Calculi and Applications*, volume 902 of *LNCS*, pages 186–200. Springer, 1995.
13. C. Guidi, R. Lucchi, R. Gorrieri, N. Busi, and G. Zavattaro. SOCK: a calculus for service oriented computing. In *ICSOC*, volume 4294 of *LNCS*, pages 327–338. Springer, 2006.
14. A. Igarashi and N. Kobayashi. A generic type system for the pi-calculus. *Theor. Comput. Sci.*, 311(1-3):121–163, 2004.
15. Z. D. Kirli. Confined mobile functions. In *14th IEEE Computer Security Foundations Workshop (CSFW '01)*, pages 283–294, Washington - Brussels - Tokyo, June 2001. IEEE Computer Society.
16. N. Kobayashi. Type systems for concurrent programs. In *10th Anniversary Colloquium of UNU/IIST*, volume 2757 of *LNCS*, pages 439–453. Springer, 2003.
17. N. Kobayashi, K. Suenaga, and L. Wischik. Resource usage analysis for the π-calculus. In *VMCAI*, volume 3855 of *LNCS*, pages 298–312. Springer, 2006.
18. C. Laneve and L. Padovani. Smooth orchestrators. In *FoSSaCS*, volume 3921 of *LNCS*, pages 32–46. Springer, 2006.
19. C. Laneve and G. Zavattaro. Foundations of web transactions. In *FoSSaCS*, volume 3441 of *LNCS*, pages 282–298. Springer, 2005.
20. A. Lapadula, R. Pugliese, and F. Tiezzi. A Calculus for Orchestration of Web Services (full version). Technical report, Dipartimento di Sistemi e Informatica, Univ. Firenze, 2006. `http://rap.dsi.unifi.it/cows`.
21. A. Lapadula, R. Pugliese, and F. Tiezzi. A WSDL-based type system for WS-BPEL. In *COORDINATION*, volume 4038 of *LNCS*, pages 145–163. Springer, 2006.
22. A. Lapadula, R. Pugliese, and F. Tiezzi. A Calculus for Orchestration of Web Services. In *ESOP*, volume 4421 of *LNCS*, pages 33–47. Springer, 2007.
23. M. Mazzara and R. Lucchi. A pi-calculus based semantics for WS-BPEL. *Journal of Logic and Algebraic Programming*, 70(1):96–118, 2006.
24. M. Merro and D. Sangiorgi. On asynchrony in name-passing calculi. *Mathematical Structures in Computer Science*, 14(5):715–767, 2004.
25. OASIS. Web Services Business Process Execution Language Version 2.0. Technical report, WS-BPEL TC OASIS, August 2006. `http://www.oasis-open.org/`.
26. J. Parrow and B. Victor. The fusion calculus: Expressiveness and symmetry in mobile processes. In *Logic in Computer Science*, pages 176–185, 1998.
27. S. Ross-Talbot and T. Fletcher. Web services choreography description language: Primer (working draft). Technical report, W3C, June 2006.
28. F.B. Schneider, G. Morrisett, and R. Harper. A language-based approach to security. In *Informatics: 10 Years Ahead, 10 Years Back. Conference on the Occasion of Dagstuhl's 10th Anniversary*, number 2000 in LNCS, pages 86–101. Springer, 2000.
29. M. Viroli. Towards a formal foundational to orchestration languages. *ENTCS*, 105:51–71, 2004.
30. J. Vitek and B. Bokowski. Confined types in java. *SPE*, 31(6):507–532, 2001.
31. N. Yoshida and V. T. Vasconcelos. Language primitives and type discipline for structured communication-based programming revisited: Two systems for higher-order session communication. In *1st International Workshop on Security and Rewriting Techniques*, ENTCS, 2006.
32. T. Zhao, J. Palsber, and J. Vitek. Lightweight confinement for featherweight java. In *Proceedings of the 18th ACM SIGPLAN conference on Object-oriented programing, systems, languages, and applications*, pages 135–148. ACM Press, 2003.