Type Checking Privacy Policies in the π -calculus

Dimitrios Kouzapas, Anna Philippou

► To cite this version:

Dimitrios Kouzapas, Anna Philippou. Type Checking Privacy Policies in the π -calculus. 35th International Conference on Formal Techniques for Distributed Objects, Components, and Systems (FORTE), Jun 2015, Grenoble, France. pp.181-195, 10.1007/978-3-319-19195-9_12 . hal-01767337

HAL Id: hal-01767337

<https://hal.inria.fr/hal-01767337>

Submitted on 16 Apr 2018

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

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

Distributed under a Creative Commons Attribution| 4.0 International License

Type checking privacy policies in the π -calculus

Dimitrios Kouzapas¹ and Anna Philippou²

¹ Department of Computing, Imperial College London and
Department of Computing Science, University of Glasgow
dk208@doc.ic.ac.uk

² Department of Computer Science, University of Cyprus
annap@cs.ucy.ac.cy

Abstract. In this paper we propose a formal framework for studying privacy. Our framework is based on the π -calculus with groups accompanied by a type system for capturing privacy requirements relating to information collection, information processing and information dissemination. The framework incorporates a privacy policy language. We show that a system respects a privacy policy if the typing of the system is compatible with the policy. We illustrate our methodology via analysis of privacy-aware schemes proposed for electronic traffic pricing.

1 Introduction

The notion of privacy is a fundamental notion for society and, as such, it has been an object of study within various scientific disciplines. Recently, its importance is becoming increasingly pronounced as the technological advances and the associated widespread accessibility of personal information is redefining the very essence of the term *privacy*.

A study of the diverse types of privacy, their interplay with technology, and the need for formal methodologies for understanding and protecting privacy is discussed in [19], where the authors follow in their arguments the analysis of David Solove, a legal scholar who has provided a discussion of privacy as a taxonomy of possible privacy violations [18]. According to Solove, privacy violations can be distinguished in four categories: *invasions*, *information collection*, *information processing*, and *information dissemination*. Invasion-related privacy violations are violations that occur on the physical sphere of an individual. The authors of [19] concentrate on the latter three categories and they identify a model for studying them consisting of the *data holder* possessing information about the *data subject* and responsible to protect this information against *unauthorized adversaries* within the environment.

The motivation of this paper stems from the need of developing formal frameworks for reasoning about privacy-related concepts. Such frameworks may provide solid foundations for understanding the notion of privacy and allow to rigorously model and study privacy-related situations. More specifically, our objective is to develop a static method for ensuring that a privacy policy is satisfied by an information system using the π -calculus as the underlying theory.

To achieve this objective, we develop a meta-theory for the π -calculus that captures privacy as policy. Following the model of [19], we create a policy language that enables us to describe privacy requirements for private data over data entities. For each type of private data we expect entities to follow different policy requirements. Thus, we define

policies as objects that describe a hierarchical nesting of entities where each node/entity of the hierarchy is associated with a set of privacy permissions. The choice of permissions encapsulated within a policy language is an important issue because identification of these permissions constitutes, in a sense, a characterization of the notion of privacy. In this work, we make a first attempt of identifying some such permissions, our choice emanating from the more obvious privacy violations of Solove's taxonomy which we refine by considering some common applications where privacy plays a central role.

As an example consider a medical system obligated to protect patient's data. Inside the system a nurse may access patient files to disseminate them to doctors. Doctors are able to process the data without any right to disseminate them. Overall, the data cannot be disclosed outside the hospital. We formalize this policy as follows:

$$t \gg \text{Hospital} : \{\text{nondisclose}\} [\text{Nurse} : \{\text{access}, \text{disclose Hospital 1}\}, \\ \text{Doctor} : \{\text{access}, \text{read}, \text{write}\}]$$

where t is the type of the patient's data. The policy describes the existence of the Hospital entity at the higher level of the hierarchy associated with the nondisclose permission signifying that patient data should not be disclosed outside the system. Within this structure, a nurse may access (but not read) a patient file and disseminate the file once (disclose Hospital 1). Similarly a doctor may be given access to a patient file but is also allowed to read and write data within the files (permissions access, read and write).

Moving on to the framework underlying our study, we employ the π -calculus with groups [5]. This calculus extends the π -calculus with the notion of *groups* and an associated type system in a way that controls how data is being disseminated inside a system. It turns out that groups give a natural abstraction for the representation of entities in a system. Thus, we build on the notion of a group of the calculus of [5], and we use the group memberships of processes to distinguish their roles within systems. Information processing issues can be analysed through the use of names of the calculus in input, output and object position to identify when a channel is reading or writing private data or when links to private data are being communicated between groups.

An implementation of the hospital scenario in the π -calculus with groups would be

$$(\nu \text{ Hospital})((\nu \text{ Nurse})(\bar{a}\langle l \rangle. \mathbf{0}) \mid (\nu \text{ Doctor})(a(x).x(y).\bar{x}\langle d \rangle. \mathbf{0}))$$

In this system, $(\nu \text{ Hospital})$ creates a new group that is known to the two processes of the subsequent parallel composition while (Nurse) and (Doctor) are groups nested within the Hospital group and available to processes $\bar{a}\langle l \rangle. \mathbf{0}$ and $a(x).x(y).\bar{x}\langle d \rangle. \mathbf{0}$, respectively. The group memberships of the two processes characterize their nature while reflecting the entity hierarchy expressed in the privacy policy defined above.

The types of the names in the above process are defined as $y : t, d : t$, that is y and d are values of sensitive data, while $l : \text{Hospital}[t]$ signifies that l is a channel that can be used only by processes which belong to group Hospital to carry data of type t . Similarly, $a : \text{Hospital}[\text{Hospital}[t]]$ states that a is a channel that can be used by members of group Hospital, to carry objects of type $\text{Hospital}[t]$. Intuitively, we may see that this system conforms to the defined policy, both in terms of the group structure as well as the permissions exercised by the processes. Instead, if the nurse were able to engage in a $\bar{l}\langle d \rangle$ action then the defined policy would be violated as it would be the case if the

type of a was defined as $a : \text{Other}[\text{Hospital}[t]]$ for some distinct group Other . Thus, the encompassing group is essential for capturing requirements of non-disclosure.

Using these building blocks, our methodology is applied as follows: Given a system and a typing we perform type checking to confirm that the system is well-typed while we infer a permission interface. This interface captures the permissions exercised by the system. To check that the system complies with a privacy policy we provide a correspondence between policies and permission interfaces the intention being that: a permission interface satisfies a policy if and only if the system exercises a subset of the allowed permissions of the policy. With this machinery at hand, we state and prove a safety theorem according to which, if a system Sys type-checks against a typing Γ and produces an interface Θ , and Θ satisfies a privacy policy $\mathcal{P}$, then Sys respects $\mathcal{P}$.

2 The Calculus

Our study of privacy is based on the π -calculus with groups proposed by Cardelli et al. [5]. This calculus is an extension of the π -calculus with the notion of a group and an operation of group creation where a group is a type for channels. In [5] the authors establish a close connection between group creation and secrecy as they show that a secret belonging to a certain group cannot be communicated outside the initial scope of the group. This is related to the fact that groups can never be communicated between processes. Our calculus is based on the π -calculus with groups with some modifications.

We assume the existence of two basic entities: $\mathcal{G}$, ranged over by $G, G_1, \dots$ is the set of groups and $\mathcal{N}$, ranged over by $a, b, x, y, \dots$, is the set of names. Furthermore, we assume a set of basic types D , ranged over by t_i , which refer to the basic data of our calculus on which privacy requirements should be enforced. Specifically, we assign each name in $\mathcal{N}$ a type such that a name may either be of some base type t or of type $G[T]$, where G is the group of the name and T the type of value that can be carried on the name. Given the above, a type is constructed via the following BNF.

$$T ::= t \mid G[T]$$

Then the syntax of the calculus is defined at two levels. At the process level, P , we have the standard π -calculus syntax. At the system level, S , we include the group construct, applied both at the level of processes $(\nu G)P$, and at the level of systems, $(\nu G)S$, the name restriction construct as well as parallel composition for systems.

$$\begin{aligned} P &::= x(y:T).P \mid \bar{x}(z).P \mid (\nu a:T)P \mid P_1 \mid P_2 \mid !P \mid \mathbf{0} \\ S &::= (\nu G)P \mid (\nu G)S \mid (\nu a:T)S \mid S_1 \mid S_2 \mid \mathbf{0} \end{aligned}$$

In $(\nu a:T)P$ and $(\nu a:T)S$, name a is bound in P and S , respectively, and in process $x(y:T).P$, name y is bound in P . In $(\nu G)P$ and $(\nu G)S$, the group G is bound in P and S . We write $\text{fn}(P)$ and $\text{fn}(S)$ for the sets of names free in a process P and a system S , and $\text{fg}(S)$ and $\text{fg}(T)$, for the free groups in a system S and a type T , respectively. Note that free occurrences of groups occur within the types T of a process/system.

We now turn to defining a labelled transition semantics for the calculus. We define a labelled transition semantics instead of a reduction semantics due to a characteristic of

the intended structural congruence in our calculus. In particular, the definition of such a congruence would omit the axiom $(\nu G)(S_1 \mid S_2) \equiv (\nu G)S_1 \mid S_2$ if $G \notin \text{fg}(S_2)$ as it was used in [5]. This is due to our intended semantics of the group concept which is considered to assign capabilities to processes. Thus, nesting of a process P within some group G , as in $(\nu G)P$, cannot be lost even if $G \notin \text{fg}(P)$, since the (νG) construct has the additional meaning of group membership in our calculus and it instills P with privacy-related permissions as we will discuss in the sequel. The absence of this law renders a reduction semantics rule of parallel composition rather complex.

To define a labelled transition semantics we first define a set of labels:

$$\ell ::= \tau \mid x(y) \mid \bar{x}(y) \mid (\nu y)\bar{x}(y)$$

Label τ is the internal action whereas labels $x(y)$ and $\bar{x}(y)$ are the input and output actions, respectively. Label $(\nu y)\bar{x}(y)$ is the restricted output where the object y of the action is restricted. Functions $\text{fn}(\ell)$ and $\text{bn}(\ell)$ return the set of the free and bound names of ℓ , respectively. We also define the relation $\text{dual}(\ell, \ell')$ which relates dual actions as

$$\text{dual}(\ell, \ell') \text{ if and only if } \{\ell, \ell'\} = \{x(y), \bar{x}(y)\} \text{ or } \{\ell, \ell'\} = \{x(y), (\nu y)\bar{x}(y)\}.$$

We use the meta-notation $(F ::= P \mid S)$ to define the labelled transition semantics.

$$\begin{array}{c} x(y : T).P \xrightarrow{x(z)} P\{z/y\} \text{ (In)} \qquad \bar{x}(z).P \xrightarrow{\bar{x}(z)} P \text{ (Out)} \\[10pt] \frac{F_1 \xrightarrow{\ell} F'_1 \quad \text{bn}(\ell) \cap \text{fn}(F_2) = \emptyset}{F_1 \mid F_2 \xrightarrow{\ell} F'_1 \mid F_2} \text{ (ParL)} \qquad \frac{F_2 \xrightarrow{\ell} F'_2 \quad \text{bn}(\ell) \cap \text{fn}(F_1) = \emptyset}{F_1 \mid F_2 \xrightarrow{\ell} F_1 \mid F'_2} \text{ (ParR)} \\[10pt] \frac{F \xrightarrow{\ell} F' \quad x \notin \text{fn}(\ell)}{(\nu x : T)F \xrightarrow{\ell} (\nu x : T)F'} \text{ (ResN)} \qquad \frac{F \xrightarrow{\bar{x}(y)} F'}{(\nu y : T)F \xrightarrow{(\nu y)\bar{x}(y)} F'} \text{ (Scope)} \\[10pt] \frac{F \xrightarrow{\ell} F'}{(\nu G)F \xrightarrow{\ell} (\nu G)F'} \text{ (ResG)} \qquad \frac{F \equiv_{\alpha} F'' \quad F'' \xrightarrow{\ell} F'}{F \xrightarrow{\ell} F'} \text{ (Alpha)} \\[10pt] \frac{P \xrightarrow{\ell} P'}{!P \xrightarrow{\ell} P' \mid !P} \text{ (Repl)} \qquad \frac{F_1 \xrightarrow{\ell_1} F'_1 \quad F_2 \xrightarrow{\ell_2} F'_2 \quad \text{dual}(\ell_1, \ell_2)}{F_1 \mid F_2 \xrightarrow{\tau} (\nu \text{bn}(\ell_1) \cup \text{bn}(\ell_2))(F'_1 \mid F'_2)} \text{ (Com)} \end{array}$$

Fig. 1. The labelled transition system

The labelled transition semantics follows along the lines of standard π -calculus semantics where $\equiv_{\alpha}$ denotes α -equivalence and the rule for the group-creation construct, (ResG), captures that transitions are closed under group restriction.

3 Policies and Types

In this section we define a policy language and the appropriate type machinery to enforce policies over processes.

3.1 Policies

Typically, privacy policy languages express positive and negative norms that are expected to hold in a system. These norms distinguish what *may* happen, in the case of a positive norm, and what may not happen, in the case of a negative norm on *data attributes* which are types of sensitive data within a system, and, in particular, how the various agents, who are referred to by their *roles*, may/may not handle this data.

The notions of an attribute and a role are reflected in our framework via the notions of base types and groups, respectively. Thus, our policy language is defined in such a way as to specify the allowed and disallowed permissions associated with the various groups for each base type (sensitive data). This is achieved via the following entities:

$$\begin{aligned}
(\text{Linearities}) \quad \lambda &::= 1 \mid 2 \mid \dots \mid * \\
(\text{Permissions}) \quad P &::= \text{read} \mid \text{write} \mid \text{access} \mid \text{disclose } G\lambda \mid \text{nondisclose} \\
(\text{Hierarchies}) \quad H &::= \varepsilon \mid G : \tilde{p}[H_j]_{j \in J} \\
(\text{Policies}) \quad \mathcal{P} &::= t \gg H \mid \mathcal{P}; \mathcal{P}
\end{aligned}$$

Specifically, we define the set of *policy permissions* P : they express that data may be read (read) and written (write), that links to data may be accessed (access) or disclosed within some group G up to λ times (disclose $G\lambda$). Notation λ is either a natural number or equal to $*$ which denotes an infinite number. While the above are positive permissions, permission *nondisclose* is a negative permission and, when associated with a group and a base type, expresses that the base type cannot be disclosed to any participant who is not a member of the group.

In turn, a policy has the form $t_1 \gg H_1; \dots; t_n \gg H_n$ assigning a structure H_i to each type of sensitive data t_i . The components H_i , which we refer to as *permission hierarchies*, specify the group-permission associations for each base type. A permission hierarchy H has the form $G : \tilde{p}[H_1, \dots, H_m]$, and expresses that an entity belonging to group G has rights $\tilde{p}$ to the data in question and if additionally it is a member of some group G_i where $H_i = G_i : \tilde{p}_i [\dots]$, then it also has the rights $\tilde{p}_i$, and so on.

We define the auxiliary functions $\text{groups}(H)$ and $\text{perms}(H)$ so as to gather the sets of groups and the set of permissions, respectively, inside a hierarchy structure:

$$\begin{aligned}
\text{groups}(H) &= \begin{cases} \{G\} \cup (\bigcup_{j \in J} \text{groups}(H_j)) & \text{if } H = G : \tilde{p}[H_j]_{j \in J} \\ \emptyset & \text{if } H = \varepsilon \end{cases} \\
\text{perms}(H) &= \begin{cases} \tilde{p} \cup (\bigcup_{j \in J} \text{perms}(H_j)) & \text{if } H = G : \tilde{p}[H_j]_{j \in J} \\ \emptyset & \text{if } H = \varepsilon \end{cases}
\end{aligned}$$

We say that a policy $\mathcal{P} = t_1 \gg H_1; \dots; t_n \gg H_n$ is *well formed*, written $\mathcal{P} : \diamond$, if it satisfies the following:

1. The t_i are distinct.
2. If $H = G : \tilde{p}[H_j]_{j \in J}$ occurs within some H_i then $G \notin \text{groups}(H_j)$ for all $j \in J$, that is, the group hierarchy is acyclic.
3. If $H = G : \tilde{p}[H_j]_{j \in J}$ occurs within some H_i , $\text{nondisclose} \in \tilde{p}$ and $\text{disclose } G'\lambda \in \text{perms}(H_j)$ for some $j \in J$, then $G' \in \text{groups}(H)$. In words, no non-disclosure requirement imposed at some level of a hierarchy is in conflict with a disclosure requirement granted in its sub-hierarchy.

Hereafter, we assume that policies are well-formed policies. As a shorthand, we write $G : \tilde{p}$ for $G : \tilde{p}[\varepsilon]$ and we abbreviate G for $G : \emptyset$.

As an example, consider a hospital containing the departments of surgery (Surgery), cardiology (Cardiology), and psychotherapy (Psychotherapy), where cardiologists of group Cardio belong to the cardiology department, surgeons of group Surgeon belong to the surgery department, psychiatrists of group Psy belong to the psychotherapy department, and CarSurgeon refers to doctors who may have a joint appointment with the surgery and cardiology departments. Further, let us assume the existence of data of type MedFile which (1) should not be disclosed to any participant outside the Hospital group, (2) may be read, written, accessed and disclosed freely within both the surgery and the cardiology departments and (3) may be read, written and accessed but not disclosed outside the psychotherapy department. We capture these requirements via policy $\mathcal{P} = \text{MedFile} \gg H$ where $p_{cd} = \{\text{read}, \text{access}, \text{write}, \text{disclose Cardiology}^*\}$, $p_{sd} = \{\text{read}, \text{access}, \text{write}, \text{disclose Surgery}^*\}$, $p_{pd} = \{\text{nondisclose}, \text{read}, \text{access}, \text{write}\}$,

$$\begin{aligned} H = & \text{Hospital} : \{\text{nondisclose}\} [\text{Cardiology} : p_{cd} [\text{Cardio}, \text{CarSurgeon}], \\ & \text{Surgery} : p_{sd} [\text{Surgeon}, \text{CarSurgeon}], \\ & \text{Psychotherapy} : p_{pd} [\text{Psy}]] \end{aligned}$$

At this point we note that, as illustrated in the above policy, hierarchies need not be tree-structured: group CarSurgeon may be reached via both the Hospital, Cardiology path as well as the Hospital, Surgery path. In effect this allows to define a process $(\nu \text{Hospital})(\nu \text{Cardiology})(\nu \text{SD})(\nu \text{CarSurgeon})P$ belonging to four groups and inheriting the permissions of each one of them.

3.2 The Type System

We proceed to define a typing system for the calculus.

Typing Judgements. The environment on which type checking is carried out consists of the component Γ . During type checking we infer the two additional structures of Δ -environments and Θ -interfaces as follows

$$\begin{aligned} \Gamma &::= \emptyset \mid \Gamma \cdot x : T \mid \Gamma \cdot G \\ \Delta &::= \emptyset \mid t : \tilde{p} \cdot \Delta \\ \Theta &::= t \gg H^\theta; \Theta \mid t \gg H^\theta \end{aligned}$$

with $H^\theta ::= G[H^\theta] \mid G[\tilde{p}]$. Note that H^θ captures a special type of hierarchies where the nesting of groups is linear. We refer to H^θ as interface hierarchies. The domain of environment Γ , $\text{dom}(\Gamma)$, contains all groups and names recorded in Γ . Environment Δ assigns permissions to sensitive data types t . When associated with a base type t , permissions read and write express that it is possible to read/write data of type t along channels of type $G[t]$ for any group G . Permission access, when associated with a type t , expresses that it is possible to receive a channel of type $G[t]$ for any G and, finally, if permission disclose $G\lambda$ is associated with t then it is possible to send channels of type $G[t]$ for up to λ times. Thus, while permissions read and write are related to manipulating sensitive data, permissions access and disclose are related to manipulating links

to sensitive data. Finally, interface Θ associates a sensitive type with a linear hierarchy of groups and a set of permissions, namely, an entity of the form $G_1[G_2[\dots G_n[\tilde{p}]\dots]]$.

We define three typing judgements: $\Gamma \vdash x \triangleright T$, $\Gamma \vdash P \triangleright \Delta$ and $\Gamma \vdash S \triangleright \Theta$. Judgement $\Gamma \vdash x \triangleright T$ says that under typing environment Γ , name x has type T . Judgement $\Gamma \vdash P \triangleright \Delta$ stipulates that process P is well typed under the environment Γ and produces a permission environment Δ . In this judgement, Γ records the types of the names of P and Δ records the permissions exercised by the names in P for each base type. Finally, judgement $\Gamma \vdash S \triangleright \Theta$ defines that system S is well typed under the environment Γ and produces interface Θ which records the group memberships of all components of S as well as the permissions exercised by each component.

Typing System. We now move on to our typing system. We begin with some useful notation. We write:

$$\Delta_T^r = \begin{cases} t : \text{read} & \text{if } T = t \\ t : \text{access} & \text{if } T = G[t] \\ \emptyset & \text{otherwise} \end{cases} \quad \Delta_T^w = \begin{cases} t : \text{write} & \text{if } T = t \\ t : \text{disclose } G \mid & \text{if } T = G[t] \\ \emptyset & \text{otherwise} \end{cases}$$

Furthermore, we define the $\uplus$ operator over permissions:

$$\begin{aligned} \tilde{p}_1 \uplus \tilde{p}_2 = & \{p \mid p \in \tilde{p}_1 \cup \tilde{p}_2 \wedge p \neq \text{disclose } G \lambda \wedge p \neq \text{nondisclose}\} \\ & \cup \{\text{disclose } G(\lambda_1 + \lambda_2) \mid \text{disclose } G \lambda_1 \in \tilde{p}_1 \wedge \text{disclose } G \lambda_2 \in \tilde{p}_2\} \\ & \cup \{\text{disclose } G \lambda \mid \text{disclose } G \lambda \in \tilde{p}_1 \wedge \text{disclose } G \lambda' \notin \tilde{p}_2\} \\ & \cup \{\text{disclose } G \lambda \mid \text{disclose } G \lambda' \notin \tilde{p}_1 \wedge \text{disclose } G \lambda \in \tilde{p}_2\} \end{aligned}$$

Operator $\uplus$ adds two permission sets by taking the union of the non nondisclose permissions modulo adding the linearities of the disclose $G \lambda$ permissions. We extend the $\uplus$ operator for Δ -environments: assuming $t : \emptyset \in \Delta$ if $t : \tilde{p} \notin \Delta$, we define $\Delta_1 \uplus \Delta_2 = \{t : \tilde{p}_1 \uplus \tilde{p}_2 \mid t : \tilde{p}_1 \in \Delta_1, t : \tilde{p}_2 \in \Delta_2\}$.

Finally, we define the $\oplus$ operator as:

$$\begin{aligned} G \oplus (t_1 : \tilde{p}_1, \dots, t_m : \tilde{p}_m) &= t_1 \gg G[\tilde{p}_1], \dots, t_m \gg G[\tilde{p}_m] \\ G \oplus (t_1 \gg H_1^\theta, \dots, t_m \gg H_m^\theta) &= (t_1 \gg G[H_1^\theta], \dots, t_m \gg G[H_m^\theta]) \end{aligned}$$

Operator $\oplus$ when applied to a group G and an interface Δ attaches G to all permission sets of Δ , thus yielding a Θ interface, whereas, when applied to a group G and an interface Θ , it attaches group G to all interface hierarchies of Θ .

The typing system is defined in Fig. 2. Rule (Name) is used to type names: in name typing we require that all group names of the type are present in Γ . Process $\mathbf{0}$ can be typed under any typing environment (axiom (Nil)) to infer the empty Δ -interface.

Rule (In) types the input-prefixed process. If environment Γ extended with the type of y produces Δ as an interface of P , we conclude that the process $x(y).P$ produces an interface where the type of T is extended with the permissions Δ_T^r , where (i) if T is base type t then Δ is extended by $t : \text{read}$ since the process is reading an object of type t , (ii) if $T = T'[t]$ then Δ is extended by $t : \text{access}$, since the process has obtained access to a link for base type t and (iii) Δ remains unaffected otherwise.

$$\begin{array}{ll}
\text{(Name)} \quad \frac{\mathbf{fg}(T) \subseteq \Gamma}{\Gamma \cdot x : T \vdash x \triangleright T} & \text{(Nil)} \quad \Gamma \vdash \mathbf{0} \triangleright \emptyset \\
\text{(In)} \quad \frac{\Gamma \cdot y : T \vdash P \triangleright \Delta \quad \Gamma \vdash x \triangleright G[T]}{\Gamma \vdash x(y : T).P \triangleright \Delta \uplus \Delta_T^r} & \text{(Out)} \quad \frac{\Gamma \vdash P \triangleright \Delta \quad \Gamma \vdash x \triangleright G[T] \quad \Gamma \vdash y \triangleright T}{\Gamma \vdash \bar{x}(y).P \triangleright \Delta \uplus \Delta_T^w} \\
\text{(ParP)} \quad \frac{\Gamma \vdash P_1 \triangleright \Delta_1 \quad \Gamma \vdash P_2 \triangleright \Delta_2}{\Gamma \vdash P_1 \mid P_2 \triangleright \Delta_1 \uplus \Delta_2} & \text{(ParS)} \quad \frac{\Gamma \vdash S_1 \triangleright \Theta_1 \quad \Gamma \vdash S_2 \triangleright \Theta_2}{\Gamma \vdash S_1 \mid S_2 \triangleright \Theta_1 \cdot \Theta_2} \\
\text{(ResNP)} \quad \frac{\Gamma \cdot x : T \vdash P \triangleright \Delta}{\Gamma \vdash (\nu x : T)P \triangleright \Delta} & \text{(ResNS)} \quad \frac{\Gamma \cdot x : T \vdash S \triangleright \Theta}{\Gamma \vdash (\nu x : T)S \triangleright \Theta} \\
\text{(ResGP)} \quad \frac{\Gamma \cdot G \vdash P \triangleright \Delta}{\Gamma \vdash (\nu G)P \triangleright G \oplus \Delta} & \text{(ResGS)} \quad \frac{\Gamma \cdot G \vdash S \triangleright \Theta}{\Gamma \vdash (\nu G)S \triangleright G \oplus \Theta} \\
\text{(Rep)} \quad \frac{\Gamma \vdash P \triangleright \Delta}{\Gamma \vdash !P \triangleright \Delta^!}
\end{array}$$

Fig. 2. The Typing System

Rule (Out) is similar: If y is of type T , x of type $G[T]$ and Δ is the permission interface for P , then, $\bar{x}(y).P$ produces an interface which extends Δ with permissions Δ_T^w . These permissions are (i) $\{t : \text{write}\}$ if $T = t$ since the process is writing data of type t , (ii) $\{\text{disclose } G\}$ if $T = G[t]$, since the process is disclosing once a link to private data via a channel of group G , and (iii) the empty set of permissions otherwise.

Rule (ParP) uses the $\uplus$ operator to compose the process interfaces of P_1 and P_2 . Parallel composition of systems, rule (ParS), concatenates the system interfaces of S_1 and S_2 . For name restriction, (ResNP) specifies that if P type checks within an environment $\Gamma \cdot x : T$, then $(\nu x)P$ type checks in environment Γ . (ResNS) is defined similarly. Moving on to group creation, for rule (ResGP) we have that, if P produces a typing Δ , then system $(\nu G)P$ produces the Θ -interface $G \oplus \Delta$ whereas for rule (ResGS), we have that if S produces a typing interface Θ then process $(\nu G)S$ produces interface $G \oplus \Theta$. Thus, enclosing a system within an (νG) operator results in adding G to the group memberships of each of the components.

Finally, for replication, axiom (Rep) states that if P produces an interface Δ then $!P$ produces an interface $\Delta^!$, where $\Delta^!$ is such that if a type is disclosed $\lambda > 1$ in Δ then it is disclosed for an unlimited number of times in $\Delta^!$. That is, $\Delta^! = \{t : \bar{p}^! \mid t : \bar{p} \in \Delta\}$, where $\bar{p}^! = \{p \in \bar{p} \mid p \neq \text{disclose } G\lambda\} \cup \{\text{disclose } G* \mid \text{disclose } G\lambda \in \Delta\}$. Note that the type system never assigns the nondisclose permissions, thus interfaces are never inferred on the nondisclose permission. This is the reason the nondisclose permission is ignored in the definition of the $\uplus$ operator.

As an example consider $S = (\nu G_1)(\nu G_2)P$ where $P = !\text{get}(loc : T_l).\bar{put}(loc).\mathbf{0}$. Further, suppose the existence of a base type Loc and types $T_l = G_1[\text{Loc}]$, $T_r = G_2[T_l]$ and $T_s = G_1[T_l]$. Let us write $\Gamma = \text{get} : T_r \cdot \text{put} : T_s \cdot \text{loc} : T_l$. Then we have:

$$\begin{array}{ll}
\Gamma \vdash \mathbf{0} \triangleright \emptyset & \text{by (Nil)} \\
\Gamma \vdash \bar{put}(loc).\mathbf{0} \triangleright \{\text{Loc} : \{\text{disclose } G_1\}\} & \text{by (Out)}
\end{array}$$

$\Gamma \vdash \text{get}(\text{loc} : T_l). \overline{\text{put}}(\text{loc}). \mathbf{0} \triangleright \{\text{Loc} : \{\text{disclose } G_1 \text{ l}, \text{access}\}\}$	by (In)
$\Gamma \vdash !\text{get}(\text{loc} : T_l). \overline{\text{put}}(\text{loc}). \mathbf{0} \triangleright \{\text{Loc} : \{\text{disclose } G_1 *, \text{access}\}\}$	by (Rep)
$\Gamma \vdash (\nu G_2) P \triangleright \text{Loc} \gg G_2 : [\{\text{disclose } G_1 *, \text{access}\}]$	by (ResNP)
$\Gamma \vdash S \triangleright \text{Loc} \gg G_1 : [G_2 : [\{\text{disclose } G_1 *, \text{access}\}]]$	by (ResNS)

4 Soundness and Safety

In this section we establish soundness and safety results for our framework. Missing proofs of results can be found in the appendix. First, we establish that typing is preserved under substitution.

Lemma 1 (Substitution). If $\Gamma \cdot x : T \vdash P \triangleright \Delta$ then $\Gamma \cdot y : T \vdash P\{y/x\} \triangleright \Delta$.

The next definition defines an operator that captures the changes on the interface environment when a process executes an action.

Definition 1 ($\Theta_1 \preceq \Theta_2$).

1. $\tilde{p}_1 \preceq \tilde{p}_2$ if (i) for all $p \in \tilde{p}_1$ and $p \neq \text{disclose } G \lambda$ implies $p \in \tilde{p}_2$, and (ii) for all $\text{disclose } G \lambda \in \tilde{p}_1$ implies $\text{disclose } \lambda' G \in \tilde{p}_2$ and $\lambda' \geq \lambda$ or $\lambda' = *$.
2. $\Delta_1 \preceq \Delta_2$ if $\forall t, t : \tilde{p}_1 \in \Delta_1$ implies that $t : \tilde{p}_2 \in \Delta_2$ and $\tilde{p}_1 \preceq \tilde{p}_2$.
3. (i) $G[\tilde{p}_1] \preceq G[\tilde{p}_2]$ if $\tilde{p}_1 \preceq \tilde{p}_2$, and (ii) $G[H_1] \preceq G[H_2]$ if $H_1 \preceq H_2$.
4. $\Theta_1 \preceq \Theta_2$ if (i) $\text{dom}(\Theta_1) = \text{dom}(\Theta_2)$, and (ii) for all $t, t \gg H_1 \in \Theta_1$ implies that $t \gg H_2 \in \Theta_2$ and $H_1 \succeq H_2$.

Specifically, when a process executes an action we expect a name to maintain or lose its interface capabilities that are expressed through the typing of the name.

We are now ready to define the notion of *satisfaction* of a policy $\mathcal{P}$ by a permission interface Θ thus connecting our type system with policy compliance.

Definition 2.

- Consider a policy hierarchy $H = G : \tilde{p}[H_j]_{j \in J}$ and an interface hierarchy H^θ . We say that H^θ *satisfies* H , written $H \Vdash H^\theta$, if:

$$\frac{\begin{array}{l} \text{groups}(H^\theta) = G \cup \bigcup_{j \in J} \text{groups}(H_j^\theta) \quad \forall j \in J, H_j \Vdash H_j^\theta \\ \text{perms}(H^\theta) \preceq (\biguplus_{j \in J} \text{perms}(H_j^\theta)) \uplus \tilde{p} \end{array}}{G : \tilde{p}[H_j]_{j \in J} \Vdash H^\theta}$$

- Consider a policy $\mathcal{P}$ and an interface Θ . Θ *satisfies* $\mathcal{P}$, written $\mathcal{P} \Vdash \Theta$, if:

$$\frac{H \Vdash H^\theta}{t \gg H; \mathcal{P} \Vdash t \gg H^\theta} \quad \frac{H \Vdash H^\theta \quad \mathcal{P} \Vdash \Theta}{t \gg H; \mathcal{P} \Vdash t \gg H^\theta; \Theta}$$

According to the definition of $H \Vdash H^\theta$, an interface hierarchy H^θ satisfies a policy hierarchy H , if its groups can be decomposed into a partition $\{G\} \cup \bigcup_{j \in J} G_j$, such that there exist interface hierarchies H_j^θ referring to groups G_j , each satisfying hierarchy H_j and where the union of the assigned permissions H_j^θ with permissions $\tilde{p}$ is a superset

of the permissions of H^θ , that is, $\text{perms}(H^\theta) \preceq (\uplus_{j \in J} \text{perms}(H_j^\theta)) \uplus \tilde{p}$. Similarly, a Θ -interface satisfies a policy, $\mathcal{P} \Vdash \Theta$, if for each component $t \gg H^\theta$ of Θ , there exists a component $t \gg H$ of $\mathcal{P}$ such that H^θ satisfies H . A direct corollary of the definition is the preservation of the $\preceq$ operator over the satisfiability relation:

Corollary 1. If $\mathcal{P} \Vdash \Theta_1$ and $\Theta_2 \preceq \Theta_1$ then $\mathcal{P} \Vdash \Theta_2$.

The next definition formalises when a system satisfies a policy:

Definition 3 (Policy Satisfaction). Let $\mathcal{P} : \diamond$. We say that S satisfies $\mathcal{P}$, written $\mathcal{P} \vdash S$, if $\Gamma \vdash S \triangleright \Theta$ for some Γ and Θ such that $\mathcal{P} \Vdash \Theta$.

We may now state our result on type preservation by action execution of processes.

Theorem 1 (Type Preservation).

1. Let $\Gamma \vdash P \triangleright \Delta$ and $P \xrightarrow{\ell} P'$ then $\Gamma \vdash P' \triangleright \Delta'$ and $\Delta' \preceq \Delta$.
2. Let $\Gamma \vdash S \triangleright \Theta$ and $S \xrightarrow{\ell} S'$ then $\Gamma \vdash S' \triangleright \Theta'$ and $\Theta' \preceq \Theta$.

Corollary 2. Let $\mathcal{P} \vdash S$ and $S \xrightarrow{\ell} S'$ then $\mathcal{P} \vdash S'$.

Let $\text{countLnk}(P, \Gamma, G[t])$ count the number of output prefixes of the form $\bar{x}\langle y \rangle$ in process P where $x : G[t]$ for some base type t . (This can be defined inductively on the structure of P .) Moreover, given a policy hierarchy H and a set of groups $\tilde{G}$, let us write $H_{\tilde{G}}$ for the interface hierarchy such that (i) $\text{groups}(H_{\tilde{G}}) = \tilde{G}$, (ii) $H \Vdash H_{\tilde{G}}$ and, (iii) for all H^θ such that $\text{groups}(H^\theta) = \tilde{G}$ and $H \Vdash H^\theta$, then $\text{perms}(H^\theta) \preceq \text{perms}(H_{\tilde{G}})$. Intuitively, $H_{\tilde{G}}$ captures an interface hierarchy with the maximum possible permissions for groups $\tilde{G}$ as determined by H . We may now define the notion of the *error process* which clarifies the satisfiability relation between the policies and processes.

Definition 4 (Error Process). Consider a policy $\mathcal{P}$, an environment Γ and a system

$$S \equiv (\nu G_1)(\nu \tilde{x}_1 : \tilde{T}_1)((\nu G_2)(\nu \tilde{x}_2 : \tilde{T}_2)(\dots((\nu G_n)(\nu \tilde{x}_n : \tilde{T}_n)P \mid Q \mid S_n) \dots) \mid S_1)$$

System S is an *error process* with respect to $\mathcal{P}$ and Γ , if there exists t such that $\mathcal{P} = t \gg H$; $\mathcal{P}'$ and at least one of the following holds, where $\tilde{G} = \langle G_1, \dots, G_n \rangle$:

1. $\text{read} \notin \text{perms}(H_{\tilde{G}})$ and $\exists x$ such that $\Gamma \vdash x \triangleright G[t]$ and $P = x(y).P'$.
2. $\text{write} \notin \text{perms}(H_{\tilde{G}})$ and $\exists x$ such that $\Gamma \vdash x \triangleright G[t]$ and $P = \bar{x}\langle y \rangle.P'$.
3. $\text{access} \notin \text{perms}(H_{\tilde{G}})$ and $\exists x$ such that $\Gamma \vdash x \triangleright G[t]$ and $P = y(x).P'$.
4. $\text{disclose } G' \lambda \notin \text{perms}(H_{\tilde{G}})$ and $\exists x, y$ such that $\Gamma \vdash x \triangleright G[t]$, $\Gamma \vdash y \triangleright G'[G[t]]$ and $P = \bar{y}\langle x \rangle.P'$.
5. $\text{disclose } G \lambda \in \text{perms}(H_{\tilde{G}})$, $\lambda \neq *$ and $\text{countLnk}(P, \Gamma, G[t]) > \lambda$
6. there exists a sub-hierarchy of H , $H' = G_k : \tilde{p}[H_i]_{i \in I}$, $1 \leq k \leq n$ with $\text{nondisclose} \in \tilde{p}$ and $\exists x, y$ such that $\Gamma \vdash x \triangleright G[t]$, $\Gamma \vdash y \triangleright G'[G[t]]$ and $P = \bar{y}\langle x \rangle.P'$ with $G' \notin \text{groups}(H')$.

The first two error processes expect that a process with no read or write permissions on a certain level of the hierarchy should not have, respectively, a prefix receiving or sending an object typed with the private data. Similarly an error process with no access permission on a certain level of the hierarchy should not have an input-prefixed subject with object a link to private data. An output-prefixed process that send links through a channel of sort G' is an error process if it is found in a specific group hierarchy with no disclose $G'\lambda$ permission. In the fifth clause, a process is an error if the number of output prefixes to links in its definition (counted with the $\text{countLnk}(P, \Gamma, t)$ function) are more than the λ in the disclose $G\lambda$ permission of the process's hierarchy. Finally, if a policy specifies that no data should be disclosed outside some group G , then a process should not be able to send private data links to groups that are not contained within the hierarchy of G .

As expected, if a process is an error with respect to a policy $\mathcal{P}$ and an environment Γ its Θ -interface does not satisfy $\mathcal{P}$:

Lemma 2. Let system S be an error process with respect to well formed policy $\mathcal{P}$ and sort Γ . If $\Gamma \vdash S \triangleright \Theta$ then $\mathcal{P} \not\models \Theta$.

By Corollary 2 and Lemma 2 we conclude with our safety theorem which verifies that the satisfiability of a policy by a typed process is preserved by the semantics.

Theorem 2 (Safety). If $\Gamma \vdash S \triangleright \Theta$, $\mathcal{P} \models \Theta$ and $S \xrightarrow{\ell^*} S'$ then S' is not an error with respect to policy $\mathcal{P}$.

5 Example

Electronic Traffic Pricing (ETP) is an electronic toll collection scheme in which the fee to be paid by drivers depends on the road usage of their vehicles where factors such as the type of roads used and the times of the usage determine the toll. To achieve this, for each vehicle detailed time and location information must be collected and processed and the due amount can be calculated with the help of a digital tariff and a road map. A number of possible implementation schemes may be considered for this system [8]. In the centralized approach, all location information is communicated to the pricing authority which computes the fee to be paid based on the received information. In the decentralized approach the fee is computed locally on the car via the use of a third trusted entity such as a smart card. In the following subsections we consider these approaches and their associated privacy characteristics.

5.1 The centralized approach

This approach makes use of on-board equipment (OBE) which computes regularly the geographical position of the car and forwards it to the Pricing Authority (PA). To avoid drivers tampering with their OBE and communicating false information, the authorities may perform checks on the spot to confirm that the OBE is reliable.

We may model this system with the aid of five groups: ETP corresponds to the entirety of the ETP system, Car refers to the car and is divided into the OBE and the

GPS subgroups, and PA refers to the pricing authority. As far as types are concerned, we assume the existence of two base types: *Loc* referring to the attribute of locations and *Fee* referring to the attribute of fees. We write $T_l = \text{ETP}[\text{Loc}]$, $T_r = \text{Car}[T_l]$, $T_{pa} = \text{ETP}[T_l]$, $T_x = \text{ETP}[T_l]$ and $T_{sc} = \text{ETP}[T_x]$.

$$\begin{aligned}
O &= !\text{read}(\text{loc} : T_l).\overline{\text{topa}}\langle \text{loc} \rangle.0 \\
&\quad | !\text{spotcheck}(s : T_x).\text{read}(ls : T_l).\bar{s}\langle ls \rangle.0 \\
L &= !(\nu \text{newl} : T_l)\overline{\text{read}}\langle \text{newl} \rangle.0 \\
A &= !\text{topa}(z : T_l).z(l : \text{Loc}).0 \\
&\quad | !\overline{\text{send}}\langle \text{fee} \rangle.0 \\
&\quad | !(\nu x : T_x)\overline{\text{spotcheck}}\langle x \rangle.x(y : T_l).y(l_s : \text{Loc}).0 \\
\text{System} &= (\nu \text{ETP})(\nu \text{spotcheck} : T_{sc})(\nu \text{topa} : T_{pa}) \\
&\quad [(\nu \text{PA})A \mid (\nu \text{Car})((\nu \text{read} : T_r)((\nu \text{OBE})O \mid (\nu \text{GPS})L))]
\end{aligned}$$

In the above model we have the component of the OBE, O , belonging to group OBE, and the component responsible for computing the current location, L , belonging to group GPS. These two components are nested within the Car group and share the private name *read* on which it is possible for L to pass to O a name via which the current location may be read. The OBE O may spontaneously read on name *read* or it may enquire the current location for the purposes of a spot check. Such a check is initiated by the pricing authority A who may engage in three different activities: Firstly, it may receive a name z from the OBE via channel *topa* and then use z for reading the car location (action $z(l)$). Secondly, it may periodically compute the fee to be paid and communicate the link ($\text{fee} : \text{ETP}[\text{Fee}]$) via name *send* : $\text{ETP}[\text{ETP}[\text{Fee}]]$. Thirdly, it may initiate a spot check, during which it creates and sends the OBE a new channel via which the OBE is expected to return the current location for a verification check.

By applying the rules of our type system we may show that $\Gamma \vdash \text{System} \triangleright \Theta$, where $\Gamma = \{\text{fee} : \text{ETP}[\text{Fee}], \text{send} : \text{ETP}[\text{ETP}[\text{Fee}]]\}$ and where

$$\begin{aligned}
\Theta &= \text{Fee} \gg \text{ETP}[\text{PA}[\{\text{disclose ETP}*\}]]; \text{Loc} \gg \text{ETP}[\text{PA}[\{\text{access}, \text{read}\}]]; \\
&\quad \text{Loc} \gg \text{ETP}[\text{Car}[\text{OBE}[\{\text{access}, \text{disclose ETP}*\}]]]; \text{Loc} \gg \text{ETP}[\text{Car}[\text{GPS}[\{\text{disclose Car}*\}]]]
\end{aligned}$$

A possible privacy policy for the system might be one that states that locations may be freely forwarded by the OBE. We may define this by $\mathcal{P} = \text{Loc} \gg H$ where

$$\begin{aligned}
H &= \text{ETP} : \text{nondisclose} [\text{Car} : [\text{OBE} : \{\text{access}, \text{disclose ETP}*\}, \\
&\quad \text{GPS} : \{\text{disclose Car}*\}], \\
&\quad \text{PA} : \{\text{access}, \text{read}\}]
\end{aligned}$$

We have that $\mathcal{P} \Vdash \text{Loc} \gg \text{ETP}[\text{PA}[\{\text{access}, \text{read}\}]]$, since the permissions assigned to groups ETP and PA by the policy are equal to $\{\text{access}, \text{read}\} \preceq \{\text{access}, \text{read}\} = \text{perms}(\text{ETP}[\text{PA}[\{\text{access}, \text{read}\}]]]$. Similarly, $\mathcal{P} \Vdash \text{Loc} \gg \text{ETP}[\text{Car}[\text{OBE}[\{\text{access}, \text{disclose ETP}*\}]]]$ and $\mathcal{P} \Vdash \text{Loc} \gg \text{ETP}[\text{Car}[\text{GPS}[\{\text{disclose Car}*\}]]]$. Thus, we conclude that System satisfies $\mathcal{P}$.

This architecture is simple but also very weak in protecting the privacy of individuals: the fact that the PA gets detailed travel information about every vehicle constitutes a privacy and security threat. An alternative implementation that limits the transmission of locations is presented in the second implementation proposal presented below.

5.2 The decentralized approach

To avoid the disclosure of the complete travel logs of a system this solution employs a third trusted entity (e.g. smart card) to make computations of the fee locally on the car and send its value to the authority which in turn may make spot checks to obtain evidence on the correctness of the calculation.

The policy here would require that locations can be communicated for at most a small fixed amount of times and that the OBE may read the fee computed by the smart card but not change its value. Precisely, the new privacy policy might be:

$$\begin{array}{ll}
 \text{Loc} \gg \text{ETP} : \text{nondisclose} [& \text{Fee} \gg \text{ETP} : \text{nondisclose} [\\
 \quad \text{Car} : [& \quad \text{Car} : [\\
 \quad \quad \text{OBE} : \{\text{access}, \text{disclose ETP } 2\}, & \quad \quad \text{OBE} : \{\}, \\
 \quad \quad \text{GPS} : \{\text{disclose Car } *\}, & \quad \quad \text{GPS} : \{\}, \\
 \quad \quad \text{SC} : \{\text{access}, \text{read}\}, & \quad \quad \text{SC} : \{\text{write}, \text{disclose ETP } *\}, \\
 \quad \quad \text{PA} : \{\text{access}, \text{read}\} & \quad \quad \text{PA} : \{\text{access}, \text{read}\}
 \end{array}$$

The new system as described above may be modelled as follows, where we have a new group SC and a new component S , a smart card, belonging to this group:

$$\begin{aligned}
 S &= !\text{read}(loc : T_l).loc(l : \text{Loc}).(\nu \text{ newval} : \text{Fee})\overline{fee}\langle \text{newval} \rangle.\overline{send}\langle fee \rangle.\mathbf{0} \\
 O &= \text{spotcheck}(s_1 : T_x).\text{read}(ls_1 : T_l).\overline{s_1}\langle ls_1 \rangle.\text{spotcheck}(s_2 : T_x).\text{read}(ls_2 : T_l).\overline{s_2}\langle ls_2 \rangle.\mathbf{0} \\
 L &= !(\nu \text{ newl} : T_l)\overline{read}\langle newl \rangle.\mathbf{0} \\
 A &= !(\nu x : T_x)\overline{spotcheck}\langle x \rangle.x(y : T_l).y(l_s : \text{Loc}).\mathbf{0} \\
 &\quad | \overline{send}(fee).fee(v : \text{Fee}).\mathbf{0} \\
 \text{System} &= (\nu \text{ ETP})(\nu \text{ spotcheck} : T_{sc})(\nu \text{ topa} : T_{pa}) \\
 &\quad [(\nu \text{ PA})A \mid (\nu \text{ Car})((\nu \text{ read} : T_r)((\nu \text{ OBE})O \mid (\nu \text{ GPS})L) \mid (\nu \text{ SC})S)]
 \end{aligned}$$

We may verify that $\Gamma \vdash \text{System} \triangleright \Theta$, where $\Gamma = \{fee : \text{ETP}[\text{Fee}], \text{send} : \text{ETP}[\text{ETP}[\text{Fee}]]\}$ and interface Θ satisfies the enunciated policy.

6 Related work

There exists a large body of literature concerned with formally reasoning about privacy. To begin with, a number of languages have been proposed to express privacy policies [16, 15] and can be used to verify the consistency of policies or to check whether a system complies with a certain policy via static techniques such as model checking [15, 13], on-the-fly using monitoring, or through audit procedures [7, 2].

Related to the privacy properties we are considering in this paper is the notion of Contextual Integrity [2]. Aspects of this notion have been formalized in a logical framework and were used for specifying privacy regulations while notions of compliance of policies by systems were considered. Also related to our work is [17] where a family of models named P-RBAC (Privacy-aware Role Based Access Control) is presented that extends the traditional role-based access control to support specification of complex privacy policies. In particular, the variation thereby introduced called Hierarchical P-RBAC introduces, amongst others, the notion of role hierarchies which is reminiscent

of our policy hierarchies. However, the methodology proposed is mostly geared towards expressing policies and checking for conflicts within policies as opposed to assessing the satisfaction of policies by systems, which is the goal of our work.

Also related to our work is the research line on typed-based security in process calculi. Among these works, numerous studies have focused on access control which is closely related to privacy. For instance the work on the $D\pi$ calculus has introduced sophisticated type systems for controlling the access to distributed resources [11, 12]. Furthermore, discretionary access control has been considered in [4] which similarly to our work employs the π -calculus with groups, while role-based access control (RBAC) has been considered in [3, 9, 6]. In addition, authorization policies and their analysis via type checking has been considered in a number of papers including [10, 1]. While adopting a similar approach, our work departs from these works in the following respects: To begin with we note that role-based access control is insufficient for reasoning about certain privacy violations. While in RBAC it is possible to express that a doctor may read patient's data and send emails, it is not possible to detect the privacy violation breach executed when the doctor sends an email with the sensitive patient data. In our framework, we may control such information dissemination by distinguishing between different types of data and how these can be manipulated. Furthermore, a novelty of our approach is the concept of hierarchies within policies which allow to arrange the system into a hierarchical arrangement of disclosure zones while allowing the inheritance of permissions between groups within the hierarchy.

To conclude, we mention our previous work of [14]. In that work we employed the π -calculus with groups accompanied by a type system based on i/o and linear types for capturing privacy-related notions. In the present work, the type system is reconstructed and simplified using the notions of groups to distinguish between different entities and introducing permissions inference during type checking. Most importantly, a contribution of this work in comparison to [14] is that we introduce a policy language and prove a safety criterion that establishes policy satisfaction by typing.

7 Conclusions

In this paper we have presented a formal framework based on the π -calculus with groups for studying privacy. Our framework is accompanied by a type system for capturing privacy-related notions and a privacy language for expressing privacy policies. We have proved a type preservation theorem and a safety theorem which establishes sufficient conditions for a system to satisfy a policy.

The policy language we have proposed is a simple language that constructs a hierarchical structure of the entities composing a system and assigning permissions for accessing sensitive data to each of the entities while allowing to reason about some simple privacy violations. These permissions are certainly not intended to capture every possible privacy issue, but rather to demonstrate a method of how one might formalize privacy rights. Identifying an appropriate and complete set of permissions for providing foundations for the notion of privacy in the general context should be the result of intensive and probably interdisciplinary research that justifies each choice. To this ef-

fect, Solove’s taxonomy of privacy violations forms a promising context in which these efforts can be based and it provides various directions for future work.

Other possible directions for future work can be inspired by privacy approaches such as contextual integrity and P-RBAC. We are currently extending our work to reason about more complex privacy policies that include *conditional* permissions and the concepts of *purpose* and *obligation* as in P-RBAC. Finally, it would be interesting to explore more dynamic settings where the roles evolve over time.

References

1. M. Backes, C. Hritcu, and M. Maffei. Type-checking zero-knowledge. In *Proceedings of CCS’08*, pages 357–370, 2008.
2. A. Barth, A. Datta, J. C. Mitchell, and H. Nissenbaum. Privacy and contextual integrity: Framework and applications. In *Proceedings of S&P’06*, pages 184–198, 2006.
3. C. Braghin, D. Gorla, and V. Sassone. Role-based access control for a distributed calculus. *Journal of Computer Security*, 14(2):113–155, 2006.
4. M. Bugliesi, D. Colazzo, S. Crafa, and D. Macedonio. A type system for discretionary access control. *Mathematical Structures in Computer Science*, 19(4):839–875, 2009.
5. L. Cardelli, G. Ghelli, and A. D. Gordon. Secrecy and group creation. *Information and Computation*, 196(2):127–155, 2005.
6. A. B. Compagnoni, E. L. Gunter, and P. Bidinger. Role-based access control for boxed ambients. *Theoretical Computer Science*, 398(1-3):203–216, 2008.
7. A. Datta, J. Blocki, N. Christin, H. DeYoung, D. Garg, L. Jia, D. K. Kaynar, and A. Sinha. Understanding and protecting privacy: Formal semantics and principled audit mechanisms. In *Proceedings of ICISS’11*, pages 1–27, 2011.
8. W. de Jonge and B. Jacobs. Privacy-friendly electronic traffic pricing via commits. In *Proceedings of FAST’08*, LNCS 5491, pages 143–161. Springer, 2009.
9. M. Dezanı-Ciancaglini, S. Ghilezan, S. Jaksic, and J. Pantovic. Types for role-based access control of dynamic web data. In *Proceedings of WFLP’10*, LNCS 6559, pages 1–29. Springer, 2010.
10. C. Fournet, A. Gordon, and S. Maffei. A type discipline for authorization in distributed systems. In *Proceedings of CSF’07*, pages 31–48, 2007.
11. M. Hennessy, J. Rathke, and N. Yoshida. safedpi: a language for controlling mobile code. *Acta Informatica*, 42(4-5):227–290, 2005.
12. M. Hennessy and J. Riely. Resource access control in systems of mobile agents. *Information and Computation*, 173(1):82–120, 2002.
13. M. Koleini, E. Ritter, and M. Ryan. Model checking agent knowledge in dynamic access control policies. In *TACAS’13*, LNCS 7795, pages 448–462. Springer, 2013.
14. D. Kouzapas and A. Philippou. A typing system for privacy. In *Proceedings of SEFM Workshops 2013*, LNCS 8368, pages 56–68. Springer, 2014.
15. Y. Liu, S. Müller, and K. Xu. A static compliance-checking framework for business process models. *IBM Systems Journal*, 46(2):335–362, 2007.
16. M. J. May, C. A. Gunter, and I. Lee. Privacy APIs: Access control techniques to analyze and verify legal privacy policies. In *Proceedings of CSFW-06*, pages 85–97, 2006.
17. Q. Ni, E. Bertino, J. Lobo, C. Brodie, C. Karat, J. Karat, and A. Trombetta. Privacy-aware role-based access control. *ACM Trans. on Information and System Security*, 13(3), 2010.
18. D. J. Solove. A Taxonomy of Privacy. *University of Pennsylvania Law Review*, 154(3):477–560, 2006.
19. M. C. Tschantz and J. M. Wing. Formal methods for privacy. In *Proceedings of FM’09*, LNCS 5850, pages 1–15. Springer, 2009.