



HAL
open science

Rule-based Specification and Analysis of Security Policies

Tony Bourdier, Horatiu Cirstea, Mathieu Jaume, H el ene Kirchner

► **To cite this version:**

Tony Bourdier, Horatiu Cirstea, Mathieu Jaume, H el ene Kirchner. Rule-based Specification and Analysis of Security Policies. 2009. inria-00429240v2

HAL Id: inria-00429240

<https://inria.hal.science/inria-00429240v2>

Preprint submitted on 1 Jun 2010 (v2), last revised 30 Jul 2010 (v3)

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 destin ee au d ep ot et  a la diffusion de documents scientifiques de niveau recherche, publi es ou non,  emanant des  tablissements d'enseignement et de recherche fran ais ou  trangers, des laboratoires publics ou priv es.

Rule-based Specification and Analysis of Security Policies

Tony Bourdier¹, Horatiu Cirstea¹, Mathieu Jaume², and H el ene Kirchner³

¹ INRIA Nancy - GRAND-EST Research Center & Nancy University & LORIA

² LIP6, Universit e Paris 6

³ INRIA Bordeaux - SUD-OUEST Research Center

Abstract. We propose a formal framework for the specification and validation of security policies. A security policy responds to the authorisation requests of a system according to a certain number of rules and to the configuration of the system at the moment of the request. A system constrained by a security policy consists of two parts: on one hand, the set of rules describing the way the decisions are taken and on the other hand, the information used by the rules and the way they evolve in the system. We call the former the policy rules and the latter the security system. Policy rules are constrained rewrite rules, whose constraints are safe first-order formulas on finite domains, which provides enhanced expressive power compared to classical security policy specification approaches like the ones using Datalog, for example. Our specifications have an operational semantics based on transition and rewriting systems and are thus executable. This framework also provides a common formalism to define, compare and compose security systems and policies. We define transformations over security systems in order to perform validation of classical security properties.

1 Introduction

When addressing the field of security policies in computer science, we are faced to multiple definitions of this concept, most often based on their purpose rather than on their behavior. For instance, in a very generic way, one can say that the purpose of a security policy is to define what it means to be secure for a system, an organization or another entity. With this point of view, security policies can be seen as special procedures that deliver authorizations to perform specific actions: for instance, they decide whether or not an access is granted, whether or not a transaction may be approved, possibly taking into account the history of transactions (e.g., on a bank account, the total amount of cash withdrawal during the month should not exceed a fixed amount), or priority considerations (e.g., an emergency call is always given priority).

The additional specificity of security policies is their reactive behaviour with respect to their execution environment: on one hand, a target system may query the policy for an authorization before performing specific accesses or transactions; on the other hand, the answers of the policy not only determine the way the corresponding action is handled in the system but can also modify the (security) information of the system and consequently subsequent executions. For example, a negative authorisation from an ATM machine security policy due to an incorrect PIN not only prevents immediate

money withdrawal but can also induce a (bad PIN) counter incrementation and lead to a permanent blocking of the corresponding account after a certain number of unsuccessful attempts. So, the security information could be seen as part of the target system but it is also intrinsic to the corresponding policy whose decisions strongly depend on it.

In this paper, we formalise separately the security system that manipulates all the security information that are used for producing the authorisation decisions and the policy rules that compute the decisions. This separation is relevant not only for a conceptually clear specification and design, but also for the verification, comparison and composition of policies. In particular, this allows one to analyse separately properties related to the management of the security information (expressed as invariants of the security system) and properties related to the policy rules (consistency or completeness for example).

A security system is formalised as a transition system whose states are first-order structures generated by syntactic environments, and whose transitions are described by transition rules on environments. Each transition is triggered by an event which corresponds to an authorization given by the security policy. The policy is given as a set of rules describing the way the decisions are taken. Policy rules are constrained rewrite rules, whose constraints are safe first-order formulas that are solved in the current state of the transition system. According to the authorization, the transition rule may or may not apply. So, conceptually, the security policy restricts the possible transitions of the security system. Such specifications of security systems and policies have a well-understood operational semantics based on transition and rewriting systems and are thus executable.

A security policy is often expected to fulfill a certain security property expressed on some entities, while it is dealing with a different set of entities. A typical example is given by access control policies designed for ensuring flow properties: such policies do not deal with information flow but only with objects containing information to be traced. Intuitively, a link is needed between “what you do” (the policy) and “what you want” (the goal for which the policy is designed). We formalize this link through a transformation of environments, whose aim is to translate an environment into another one dealing with the entities we are interested in. Such a transformation allows thus the validation of a property over a system even if the property is expressed in a different specification. In practice, this approach provides a way to reuse the same specification of a security property in order to analyse or to verify several policies and systems, thus showing the benefits of a library of generic security properties, dedicated to particular domains (like information flows) and that can be considered in several contexts.

We first introduce some useful notions and notations in Section 2. Section 3 presents the different components of our specification framework: security signatures, environments, transition rules as well as security systems, policy rules, and secured systems. Section 4 addresses the validation point of view by defining environment transformations and illustrating the verification of security properties. In Section 5, we compare our approach with other works. Conclusion and future work are presented in Section 6.

2 Preliminaries

We assume the reader familiar with the standard notions of term rewriting, first order logic and Datalog. This section briefly recalls basic notions used in this paper; more

details can be found in [12] for logic considerations, [1] for rewriting considerations and [18] for Datalog ones.

A many-sorted signature $\Sigma = (\mathcal{S}, \mathcal{F}, \mathcal{P})$ is given by a set of sorts \mathcal{S} , a set of function symbols \mathcal{F} and a set of predicate symbols \mathcal{P} . A function symbol f with arity $w = s_1, \dots, s_n \in \mathcal{S}^*$ and co-arity s is written $f:w \mapsto s$. A predicate symbol p with arity $s_1, \dots, s_n \in \mathcal{S}^*$ is written $p:w$. Variables are also sorted and $x:s$ means that the variable x has sort s . The set of variables sorted by s is denoted by \mathcal{X}_s . We assume in this paper that all variables are ranging over finite sets. This condition can be relaxed under some conditions [18], especially for allowing built-in sorts such as integers. Given a set ζ extending a set of variables \mathcal{X} (possibly empty) with constants sorted by \mathcal{S} , the set of Σ -terms over ζ denoted by $\mathcal{T}_{\Sigma, \zeta}^s$ is the smallest set containing elements of ζ of sort s and all the words $f(t_1, \dots, t_n)$ such that $f:s_1, \dots, s_n \mapsto s \in \Sigma$ and $t_i \in \mathcal{T}_{\Sigma, \zeta}^{s_i}$ for $i \in [1..n]$. The sort is omitted when not important in the context and we write \mathcal{T}_{Σ}^s instead of $\mathcal{T}_{\Sigma, \emptyset}^s$. We also consider a partial ordering $<$ on the set \mathcal{S} of sorts of a signature Σ and we write $s_1 < s_2$ if $\mathcal{T}_{\Sigma, \zeta}^{s_1} \subseteq \mathcal{T}_{\Sigma, \zeta}^{s_2}$. $\mathcal{Pos}(t)$ denotes the set of positions of a term t , ε the root position, $t|_{\omega}$, resp. $t(\omega)$, the subterm of t , resp. the symbol of t , at position ω , and $t[s]_{\omega}$ the term t with the subterm at position ω replaced by s . The set of variables occurring in a term t is denoted by $\mathcal{Var}(t)$. If $\mathcal{Var}(t)$ is empty, t is called a ground term. All the following definitions are given *w.r.t.* to a set ζ whose subset of variables is denoted by \mathcal{X} . A substitution is a mapping from \mathcal{X} to $\mathcal{T}_{\Sigma, \zeta}$ which is the identity except over a finite set of variables called domain of σ and denoted by $\mathcal{Dom}(\sigma)$. σ naturally extends to an endomorphism of $\mathcal{T}_{\Sigma, \zeta}$. The set $\mathcal{Codom}(\sigma) = \{t \in \mathcal{T}_{\Sigma, \zeta} \mid \exists x \in \mathcal{Dom}(\sigma), \sigma(x) = t\}$ is called codomain of σ . If $\mathcal{Codom}(\sigma)$ contains only ground terms, σ is said to be a ground substitution. A Σ -atom is a word of the form $p(t_1, \dots, t_n)$ or $t_1 = t_2$ or $t:s$ where $t_1, \dots, t_n, t \in \mathcal{T}_{\Sigma, \zeta}, p \in \mathcal{P}$ and $s \in \mathcal{S}$. A Σ -literal is either a Σ -atom or a negated (with \neg) Σ -atom and the set of Σ -formulae built out of Σ -literals is denoted by $\mathcal{For}_{\Sigma, \zeta}$. The set of free variables of a formula ϕ (i.e. variables not in the scope of a quantifier) is denoted by $\mathcal{FVar}(\phi)$. A Σ -constraint is either a formula φ or $\mathcal{Count}(\varphi) \odot n$ where $n \in \mathbb{N}$ and \odot a comparison symbol. A Σ -action is either a Σ -literal without equality symbol or an oriented equality $lhs \leftrightarrow rhs$ with $lhs, rhs \in \mathcal{T}_{\Sigma, \zeta}$. A logical rule over Σ , denoted by $l_1 \wedge \dots \wedge l_n \mapsto a$, consists of a conjunction of Σ -literals l_i called the body and a Σ -atom a called the goal. Safety and stratification of logical rules as well as the semantics of a set of safe and stratified logical rules are defined as usual [18].

A constrained rewrite rule over a signature Σ is a 3-tuple $(l, \varphi, r) \in \mathcal{T}_{\Sigma, \mathcal{X}} \times \mathcal{For}_{\Sigma, \mathcal{X}} \times \mathcal{T}_{\Sigma, \mathcal{X}}$, denoted by $l \xrightarrow{\varphi} r$, such that $\mathcal{Var}(r) \subseteq \mathcal{Var}(l) \cup \mathcal{FVar}(\varphi)$. A constrained term rewrite system (CTRS) \mathcal{R} is a set of constrained rewrite rules. Given a signature $\Sigma' \supseteq \Sigma$, we say that $t \in \mathcal{T}_{\Sigma'}$ rewrites into a term $t' \in \mathcal{T}_{\Sigma'}$ with respect to \mathcal{R} and a Σ' -theory ϑ , which is denoted by $t \xrightarrow{\vartheta}_{\mathcal{R}} t'$ iff there exist a position $p \in \mathcal{Pos}(t)$, a rewrite rule $l \xrightarrow{\varphi} r \in \mathcal{R}$, and a ground substitution σ with $\mathcal{Dom}(\sigma) = \mathcal{Var}(l) \cup \mathcal{FVar}(\varphi)$ such that $\vartheta \models \{t|_p = \sigma(l) ; t' = t[\sigma(r)]_p ; \sigma(\varphi)\}$.

3 Secured systems

A security policy responds to the authorisation requests of a system according to a certain number of rules and to the configuration of the system at the moment of the

request. We consider thus that a system constrained by a security policy consists of two parts: on one hand, the set of rules describing the way the decisions are taken and on the other hand, the information used by the rules and the way these evolve in the system. We call the former the policy rules and the latter the security system. In our framework all objects manipulated by the security system and the policy rules are described as first order terms over a common signature called the security signature. We define the security system using transition rules over environments and the policy rules as a constrained rewrite system.

3.1 Security signature

The relationship between the policy rules and the security system is realized by some security events, which are exactly pairs consisting of an authorisation request and the associated decision. Besides the underlying signature of the security information that may vary from policy to policy, a security signature should thus always define the sorts *Query* and *Decision*.

Definition 1 (Security signature). *A security signature is a signature $\Sigma_{Sys} \cup \Sigma_{Ev}$ such that Σ_{Ev} contains two sorts *Query*, *Decision* with $Decision < Query$ and a set of function symbols of co-arity *Query* or *Decision*.*

Example 1. Along the lines of this paper, we consider an access control system, on which we define a confidentiality policy (which can be viewed as a variant of the mandatory part of the Bell and LaPadula policy [5]). This policy constrains accesses done by subjects (*S*) over objects (*O*) according to access modes (*A*) by considering levels of security (belonging to a finite lattice (L, inf)) associated with subjects and objects. Hence, we introduce the security signature $\Sigma^{lbp} = \Sigma_{Sys}^{lbp} \cup \Sigma_{Ev}^{lbp}$ as follows.

First, $\Sigma_{Sys}^{lbp} = (S, \mathcal{F}, \mathcal{P})$ consists of $S_{Sys}^{lbp} = \{S, O, A, L\}$ and

$$\mathcal{P}_{Sys}^{lbp} = \left\{ \begin{array}{l} inf:L, L \\ m:S, O, A \\ sudo:S \end{array} \right\} \quad \mathcal{F}_{Sys}^{lbp} = \left\{ \begin{array}{l} read: \mapsto A, write: \mapsto A, erase: \mapsto A, \\ root: \mapsto S, top: \mapsto L, \\ f_s: S \mapsto L, f_o: O \mapsto L \end{array} \right\}$$

The functions f_s and f_o describe security levels associated with subjects and objects; *root* (resp. *top*) is a particular subject (resp. security level). The predicate m describes current accesses over objects by subjects: $m(s, o, a)$ means that the subject s has an access over an object o in the access mode a . The predicate *sudo* describes sudoers, i.e. users with root privileges. Σ_{Ev}^{lbp} is based on the following function symbols:

$$\mathcal{F}_{Ev}^{lbp} = \left\{ \begin{array}{l} ask:S, O, A \mapsto Query, release:S, O, A \mapsto Query, deny: \mapsto Decision \\ create:S, L \mapsto Query, delete:S, L \mapsto Query, permit: \mapsto Decision \end{array} \right\}$$

$ask(s, o, a)$ (resp. $release(s, o, a)$) means that the subject s asks to get (resp. to release) an access over an object o according to the access mode a ; $create(s, l)$ means that the subject s asks to create an object o whose security level is l ; $delete(s, l)$ means that the subject s asks to delete all objects whose security levels are smaller than l .

3.2 Environments and transition rules

A security system is a transition system that describes the way security information evolve. A state of this system is defined by a set of kernel information, called envi-

ronment, that can be modified by the transition rules of the system and an immutable set of closure rules used to compute the complete security information. The result of such a computation, called in what follows the semantics of the environment, represents an extensional description of the state defined intensionally by the corresponding environment and is obtained by saturation of the environment using the closure rules.

Definition 2 (Environment). An environment η over a signature $\Sigma = \langle \mathcal{S}, \mathcal{F}, \mathcal{P} \rangle$ consists of:

- (i) a **domain**: a finite set $|\eta|$ of sorted constants such that $\mathcal{C}_\Sigma \subseteq |\eta|$;
- (ii) a **base of facts**: a finite set \mathcal{B}_η of atoms of the form $p(t_1, \dots, t_n)$ with $p \in \mathcal{P}$, $n > 0$ and $t_1, \dots, t_n \in |\eta|$;
- (iii) a **base of equalities**: a finite set \mathcal{E}_η of equalities of the form $f(t_1, \dots, t_n) = t$ with $f \in \mathcal{F}$, $n > 0$ and $t_1, \dots, t_n, t \in |\eta|$ which does not contain two equalities with the same left-hand side;
- (iv) **closure rules**: a set \mathcal{R}_η of safe and stratified logical rules over Σ .

The base of equalities gives the interpretation into the domain of the environment for any term of the signature. We denote by $t \downarrow_\eta$ the interpretation of the term t in $|\eta|$, i.e. $f(t_1, \dots, t_n) \downarrow_\eta = u$ iff $f(u_1, \dots, u_n) = u \in \mathcal{E}_\eta$ and $u_i = t_i \downarrow_\eta$ for all $i \in [1, n]$.

Example 2. If we consider the security signature Σ^{lbp} introduced in Example 1, we can define the environment η^{lbp} defined as follows. The domain $|\eta^{lbp}|$ contains the constants *Bob*, *Alice* and *Charlie* of sort S , and the constants *Secret*, *Confidential*, L_1 , L_2 , *Public*, *Sanitized* of sort L . The base of facts $\mathcal{B}_{\eta^{lbp}}$ (partially) defines the partial order *inf* and states that *Charlie* is a “suoer”:

$$\mathcal{B}_{\eta^{lbp}} = \left\{ \begin{array}{l} \text{inf}(\text{Confidential}, \text{Secret}), \text{inf}(L_1, \text{Confidential}), \text{inf}(\text{Public}, L_2), \\ \text{inf}(\text{Sanitized}, L_1), \text{inf}(L_2, \text{Confidential}), \text{inf}(\text{Sanitized}, \text{Public}) \\ \text{suo}(\text{Charlie}) \end{array} \right\}$$

The base of equalities $\mathcal{E}_{\eta^{lbp}}$ provides a definition for the security levels associated with the subjects defined in the domain:

$$\mathcal{E}_{\eta^{lbp}} = \{ f_s(\text{root}) = \text{top}, f_s(\text{Bob}) = \text{Secret}, f_s(\text{Alice}) = L_2, f_s(\text{Charlie}) = \text{Public} \}$$

The set of closure rules completes the definition of *inf*:

$$\mathcal{R}_{\eta^{lbp}} = \{ \models \text{inf}(x, x); \text{inf}(x, y) \wedge \text{inf}(y, z) \models \text{inf}(x, z) \}$$

Definition 3 (Semantics of an environment). The semantics of an environment η over a signature Σ , denoted by $\llbracket \eta \rrbracket$, is the unique least fixed point of the logic program containing \mathcal{B}_η , \mathcal{E}_η and \mathcal{R}_η . We say that a formula φ holds in η iff $\llbracket \eta \rrbracket \models \varphi$.

Due to the restrictions imposed on the domain and on the formulas, we can state:

Proposition 1. For any environment η over Σ , $\llbracket \eta \rrbracket$ exists and is unique and computable. Moreover the validity of any first-order formula in $\llbracket \eta \rrbracket$ is decidable.

The transition rules of the security system describing the evolution of the states are labeled by the events that trigger them.

Definition 4 (Transition rule). Given a security signature $\Sigma = \Sigma_{S_{ys}} \cup \Sigma_{E_v}$, a Σ -transition rule consists of an event $= \langle f_q(t_1, \dots, t_n), f_d(u_1, \dots, u_m) \rangle \in \mathcal{T}_{\Sigma, \mathcal{X}}^{\text{Query}} \times$

$\mathcal{T}_{\Sigma, \mathcal{X}}^{Decision}$ with $f_q, f_d \in \mathcal{F}_{Ev}$ and $t_1, \dots, t_n, u_1, \dots, u_m \in \mathcal{C}_{\Sigma_{Sys}} \cup \mathcal{X}$ and a sequence of pairs $\langle condition, statement \rangle$ where condition is a Σ_{Sys} -constraint and statement is a conjunction of operations of the form **new**_k s with k integer and $s \in \mathcal{S}$ or **delete** t with $t \in \mathcal{T}_{\Sigma_{Sys}, \mathcal{X}}$ and Σ_{Sys} -actions which can use a constant \square_k of sort s for each **new**_k s . A transition rule is usually written:

$$(event) \frac{condition_1}{statement_1}; \frac{condition_2}{statement_2}; \dots; \frac{condition_m}{statement_m}$$

Intuitively, a *statement* is “executed” when an *event* is triggered and the corresponding *condition* is satisfied. We require the conditions occurring in the same transition rule to be mutually exclusive, i.e. in each environment, at most one of them is satisfied. In practice, this is not a strong restriction, since it is always possible to add to the k -th condition the negation of the previous conditions. When executing a statement on an environment, we can add/remove entities to/from its domain, add/remove atoms to/from its base of facts and change the base of equalities in some specific points.

Example 3. If we consider the security signature Σ^{lbp} introduced in Example 1, we can define the following set δ^{lbp} of Σ^{lbp} -transition rules:

$$\begin{aligned} & (ask(s, o, a), permit) \frac{}{m(s, o, a)} \quad (create(s, l), permit) \frac{\text{Count}(o:O) \leq 10}{\mathbf{new}_1 O \wedge f_o(\square_1) \leftrightarrow l} \\ & (release(s, o, a), permit) \frac{m(s, o, a)}{\neg m(s, o, a)} \quad (delete(s, l), permit) \frac{inf(l', l) \wedge f_o(o) = l'}{\mathbf{delete} o} \end{aligned}$$

Intuitively, the rules indicate that when an access request is permitted, the corresponding fact is added and when the respective access is released, the fact is removed. An object can be created only if the number of existing objects has not reached a certain threshold (10 in our case) even if the corresponding request is permitted. When a request $delete(s, l)$ is permitted, all objects of a lower level than l are deleted.

Allowing a sequence of pairs $\langle condition, statement \rangle$ in transition rules leads to a more concise description of the operations to be performed for a given event. Notice that using only one pair is strictly less expressive. For example, if we suppose a security signature containing a function symbol $createS : S, L \mapsto Query$ then, specifying that the operations performed at the creation of a subject are different for *root* needs the use of two separate conditions:

$$(createS(s, l), permit) \frac{s = root}{\mathbf{new}_1 S \wedge sudo(\square_1) \wedge f_s(\square_1) \leftrightarrow l}; \frac{}{\mathbf{new}_1 S \wedge f_s(\square_1) \leftrightarrow l}$$

Alternatively, several rules with a similar event can be used but our approach allows the factorization with respect to events having the same shape and thus, we get a simpler definition for the application of a set of transition rules, given below.

Definition 5 (Relation induced by a transition rule w.r.t. an event). *The transition rule $r = \langle event, \langle condition_i, statement_i \rangle_{1 \leq i \leq n} \rangle$ is triggered in a given environment η by the ground event $evt \in \mathcal{T}_{\Sigma, |\eta|}^{Query} \times \mathcal{T}_{\Sigma, |\eta|}^{Decision}$ if there exists a ground substitution σ of the variables of event into $|\eta|$ such that $\sigma(event) = evt \downarrow_\eta$. It induces a relation \xrightarrow{evt}_r over environments such that $\eta \xrightarrow{evt}_r \eta'$ iff, when denoting by*

$\langle \text{condition}, \text{statement} \rangle$ the first pair such that $Sol_\eta(\sigma(\text{condition})) = \{\mu \mid \llbracket \eta \rrbracket \models \mu(\sigma(\text{condition}))\} \neq \emptyset$, η' is obtained by taking initially $\eta' = \eta$ and executing the following operations:

- for each **new** s occurring in statement, add to $\mid \eta' \mid$ a fresh constant c of sort s (i.e. c is not yet in $\mid \eta' \mid$);
- for each $p(t_1, \dots, t_m)$ in statement, add to $\mathcal{B}_{\eta'}$ the fact $p(\mu(t_1)\downarrow_\eta, \dots, \mu(t_m)\downarrow_\eta)$, for each $\mu \in Sol_\eta(\sigma(\text{condition}))$;
- for each $\neg p(t_1, \dots, t_m)$ in statement, remove from $\mathcal{B}_{\eta'}$ the fact $p(\mu(t_1)\downarrow_\eta, \dots, \mu(t_m)\downarrow_\eta)$, for each $\mu \in Sol_\eta(\sigma(\text{condition}))$;
- for each $f(t_1, \dots, t_m) \leftrightarrow t$ in statement, remove $f(\mu(t_1)\downarrow_\eta, \dots, \mu(t_m)\downarrow_\eta) = u$ from $\mathcal{E}_{\eta'}$ if it exists and add $f(\mu(t_1)\downarrow_\eta, \dots, \mu(t_m)\downarrow_\eta) = \mu(t)\downarrow_\eta$ to $\mathcal{E}_{\eta'}$, for each $\mu \in Sol_\eta(\sigma(\text{condition}))$;
- for each **delete** t in statement remove $\mu(t)\downarrow_\eta$ from $\mid \eta' \mid$ and all facts (from $\mathcal{B}_{\eta'}$) and equalities (from $\mathcal{E}_{\eta'}$) in which occur $\mu(t)\downarrow_\eta$, for each $\mu \in Sol_\eta(\sigma(\text{condition}))$.

We say that η is transformed by r w.r.t. evt into η' . Notice that if there exists no σ s.t. $\sigma(\text{event}) = evt\downarrow_\eta$ or no k s.t. $Sol_\eta(\sigma(\text{condition}_k)) \neq \emptyset$ then, there is no η' in relation with η . A set of Σ -transition rules $r_k = \langle \text{event}_k, \langle \text{cond}_i, \text{stat}_i \rangle_{1 \leq i \leq n_k} \rangle$, $1 \leq k \leq n$, such that for any $k \neq k'$ there exists no substitution σ s.t. $\sigma(\text{event}_k) = \sigma(\text{event}_{k'})$ is said **disjoint** and then $\xrightarrow{\text{evt}}_\delta$ is the relation such that $\eta \xrightarrow{\text{evt}}_\delta \eta'$ iff there exists k such that $\eta \xrightarrow{\text{evt}}_{r_k} \eta'$ or else $\eta = \eta'$.

Although the application of a transition rule yields a unique environment when the **new** operations are performed before any other operation, the result might depend on the strategy used for computing it and, in particular, on the order the different solutions μ are selected and on the order the operations are performed for each solution. In order to obtain the same result independently of the application strategy, we should guarantee that during the execution of the operations corresponding to the rule: (1) a fact cannot be added and removed; (2) the interpretation of a term does not change; (3) constants of the signature and constants potentially involved in the actions of the rule are not deleted.

There are strong syntactic restrictions constraining the syntax of transition rules that would enforce these conditions. For instance, we can impose that (1) a statement cannot contain p and its negation; (2) a statement cannot contain two oriented equalities whose left-hand sides can be unified and the right-hand sides of all oriented equalities in a statement are necessarily constants or variables of the event of the rule; (3) an operation **delete** t occurs alone in a statement and should not affect the constants of the signature. Indeed, these conditions are too strong, since they limit significantly the expressive power of the formalism.

For example, let us suppose the existence of the symbols $chown:S, S, O, A \mapsto \text{Query}$, $create: \mapsto \text{Query}$, and $remove:S \mapsto \text{Query}$. Then, the following transition rule specifying what happens when the subject accessing an object is replaced by another one

$$(\text{chown}(s_1, s_2, o, a), \text{permit}) \frac{m(s_1, o, a) \wedge s_1 \neq s_2}{\neg m(s_1, o, a) \wedge m(s_2, o, a)}$$

would be forbidden. Similarly, the transition rule specifying that the level of any new created object is smaller or equal to the level of any other existing object

$$(create, permit) \frac{\forall o.(y = f_o(o) \vee \inf(y, f_o(o))) \wedge \forall l.\forall o.\inf(y, l) \Rightarrow \inf(f_o(o), l)}{\mathbf{new}_1 O \wedge f_o(\square_1) \leftarrow y}$$

would not be possible. Moreover, we could not specify that the accesses of removed subjects should be deleted

$$(remove(s), permit) \frac{m(s, o, a)}{\mathbf{delete} s \wedge \neg m(s, o, a)}.$$

We replace thus the above restrictions by the following weaker and decidable assumptions \mathcal{H} on each pair $\langle condition_i, statement_i \rangle$ of a transition rule $r = \langle event, \langle condition_i, statement_i \rangle_{1 \leq i \leq n} \rangle$:

1. a fact cannot be added and removed: for any pair $p(t_1, \dots, t_n), \neg p(t'_1, \dots, t'_n)$ in $statement_i$, $t_1, \dots, t_n, t'_1, \dots, t'_n$ contain no free variable from $condition_i$ and either t_j and t'_j are two different ground terms or $condition_i$ is a conjunction containing $t_j \neq t'_j$ for some $j \in [1, n]$.
2. the interpretation of a term does not change: $statement_i$ cannot contain two oriented equalities $l \leftarrow r$ and $l' \leftarrow r'$ such that there exists a substitution σ with $\sigma(l) = \sigma(l')$; the variables of the *rhs* of oriented equalities in $statement_i$ must be instantiated in a unique way, i.e. for any $lhs \leftarrow rhs$ in $statement_i$, $condition_i$ is of the form $(\forall \vec{x} \exists! \vec{y} \forall \vec{z}.\varphi) \wedge \varphi$ with $\vec{x} = \text{Var}(event) \cup \text{Var}(lhs)$, $\vec{y} = \text{Var}(rhs) \setminus \vec{x}$ and $\vec{z} = (\mathcal{F}\text{Var}(\varphi) \cup \text{Var}(statement_i)) \setminus (\vec{x} \cup \vec{y})$.
3. constants of the signature are not deleted: for any **delete** t in $statement_i$, $condition_i$ should contain, for every $u \in \mathcal{C}_\Sigma$, the atom $u \neq t$; constants potentially involved in other actions are not deleted: for every **delete** t with t of sort S in $statement_i$, $condition_i$ should contain, for every term u of sort S occurring in an action of $statement_i$, the atom $u \neq t$.

Proposition 2. *For any disjoint set δ of transition rules satisfying \mathcal{H} , the relation \xrightarrow{evt}_δ is deterministic.*

3.3 Security systems, policy rules and secured systems

A security system is defined by a set of transition rules and by an initial environment.

Definition 6 (Security system). *Given a security signature $\Sigma = \Sigma_{Sys} \cup \Sigma_{Ev}$, a **system presentation** over Σ consists of an initial environment η_{init} over Σ_{Sys} , and a set δ of disjoint Σ -transition rules. The semantics of a system presentation $\mathfrak{S} = (\eta_{init}, \delta)$, called a **security system**, is the labelled transition system $\llbracket \mathfrak{S} \rrbracket$ whose states are environments over Σ_{Sys} , whose initial state is η_{init} and whose transitions are $\eta \xrightarrow{evt}_\delta \eta'$ for some $evt \in \mathcal{T}_{\Sigma, |\eta|}^{Query} \times \mathcal{T}_{\Sigma, |\eta|}^{Decision}$.*

Example 4. The security system \mathfrak{S}^{lbp} over the security signature Σ^{lbp} defined in Example 1 consists of the initial environment η^{lbp} defined in Example 2 and the set δ^{lbp} of transition rules defined in Example 3.

Definition 7 (Policy rules). A set of **policy rules** over a security signature Σ is an ordered constrained term rewrite system over $\Sigma = \Sigma_{Sys} \cup \Sigma_{Ev}$ with all the rules of the form $f(t_1, \dots, t_n) \xrightarrow{\varphi} r$ with $f \in \mathcal{F}_{Ev}$ and φ a Σ_{Sys} -constraint.

We write $q \xrightarrow{\eta}_{\mathfrak{R}} d$ when q is rewritten in one step w.r.t. the set of policy rules \mathfrak{R} and the environment η into d and we write $q \xrightarrow{*}_{\mathfrak{R}} d$ for a multiple steps rewriting.

Example 5. If we consider our example, the following ordered CTRS defines a policy:

$$\left\{ \begin{array}{l} ask(s, o, a) \xrightarrow{sudo(s) \wedge s \neq root} ask(root, o, a) \\ ask(s, o, read) \xrightarrow{inf(f_o(o), f_s(s)) \wedge (\forall o' m(s, o', write) \Rightarrow inf(f_o(o), f_o(o')))} permit \\ ask(s, o, write) \xrightarrow{\forall o' m(s, o', read) \Rightarrow inf(f_o(o'), f_o(o))} permit \\ ask(s, o, erase) \xrightarrow{inf(f_o(o), f_s(s))} permit \\ ask(s, o, a) \rightarrow deny \\ release(s, o, a) \rightarrow permit \\ create(s, l) \xrightarrow{s=root \vee sudo(s) \vee inf(f_s(s), l)} permit \\ create(s, l) \rightarrow deny \\ delete(s, l) \xrightarrow{sudo(s) \wedge s \neq root} delete(root, l) \\ delete(s, l) \xrightarrow{inf(l, f_s(s))} permit \\ delete(s, l) \rightarrow deny \end{array} \right.$$

These rules specify that:

- a *sudo* subject has the same access and deletion rights as *root*;
- a subject can read an object whose level of security is smaller than its level of security if it does not write into an object of a lower security level;
- a subject can write into an object if it does not read an object of a higher level;
- a subject can erase an object whose security level is smaller than its security level;
- a subject can release any of its accesses;
- except for root and sudoers which are authorized to create objects of any security level, a subject can only create objects of security levels higher than its level;
- a subject can delete objects of security levels smaller than its security level;
- in all other cases, the request is denied.

Notice that since the rules are ordered, the constraints do not need to impose explicitly the negation of the constraints of previous overlapping rules and, in particular, no constraint is needed for the “default” rules.

Definition 8. A set of policy rules \mathfrak{R} over a security signature $\Sigma = \Sigma_{Sys} \cup \Sigma_{Ev}$ is **η -consistent** (resp. **η -complete**) for an environment η over Σ_{Sys} iff for any query $q \in \mathcal{T}_{\Sigma, |\eta|}^{Query}$, there exists at most (resp. at least) one decision $d \in \mathcal{T}_{\Sigma, |\eta|}^{Decision}$ such that $q \xrightarrow{*}_{\mathfrak{R}} d$.

These properties can be proved for a large class of policy rules.

Proposition 3. A set \mathfrak{R} of policy rules is η -consistent if (1) for each rule, its left-hand side contains only one occurrence of each variable and its constraint does not involve

terms of sort $Query$; (2) for any two rules $l \xrightarrow{\varphi} r$ and $l' \xrightarrow{\varphi'} r'$ there exists no position ω and no substitution σ such that $\llbracket \eta \rrbracket \models \{\sigma(l|_{\omega}) = \sigma(l') \wedge \sigma(\varphi) \wedge \sigma(\varphi')\}$.

Proposition 4. *A set \mathfrak{R} of policy rules over a security signature $\Sigma = \Sigma_{Sys} \cup \Sigma_{Ev}$ is η -complete if (1) the reduction $\xrightarrow{*}_{\mathfrak{R}}^{\eta}$ terminates, (2) for any $f:S_1, \dots, S_n \mapsto S \in \Sigma_{Ev}$ with $S \neq Decision$ there exists a default rule $f(x_1, \dots, x_n) \rightarrow c$ for some c , (3) each rule of \mathfrak{R} is sort-preserving or sort-decreasing (i.e. the sort of its left-hand side is equal or greater than the sort of its right-hand side).*

The classical methods for proving termination of TRS can be adapted for CTRS. For example, the policy rules introduced in Example 5 can be shown terminating using an obvious polynomial interpretation [1] connected to the corresponding constraints. There is also a default rule for each symbol of sort $Query$ and the rules are sort-preserving or sort-decreasing. Consequently, the corresponding normal forms are clearly in this case *permit* or *deny*. The policy rules obviously satisfy condition (1) of Proposition 3 and, because of the order, condition (2) as well. The policy rules of Example 5 are thus η -complete and η -consistent for any environment η .

Definition 9 (Secured system). *Given a security signature $\Sigma = \Sigma_{Sys} \cup \Sigma_{Ev}$, a **presentation of a secured system** over Σ consists of a system presentation $\mathfrak{S} = (\eta_{init}, \delta)$ over Σ and a set \mathfrak{R} of η -complete and η -consistent policy rules over Σ for any environment η over Σ_{Sys} . The semantics of a secured system presentation $\wp = \langle \mathfrak{S}, \mathfrak{R} \rangle$, called a **secured system**, is the labelled transition system $\llbracket \wp \rrbracket$ whose states are environments over Σ_{Sys} , whose initial state is the initial state of \mathfrak{S} and whose transitions are $\eta \xrightarrow{\langle q, d \rangle}_{\wp} \eta'$ for some $\langle q, d \rangle \in \mathcal{T}_{\Sigma, |\eta|}^{Query} \times \mathcal{T}_{\Sigma, |\eta|}^{Decision}$ such that $\eta \xrightarrow{\langle q, d \rangle}_{\delta} \eta'$ and $q \xrightarrow{*}_{\mathfrak{R}}^{\eta} d$.*

In a secured system, by definition, (i) \mathfrak{R} is η -complete and η -consistent for any environment η over Σ_{Sys} , so $\xrightarrow{*}_{\mathfrak{R}}^{\eta}$ is computable, and moreover (ii) $\xrightarrow{evt}_{\delta}$ is computable by construction for any *evt*. So we get:

Proposition 5. *The transition relation \rightarrow_{\wp} is computable for any presentation of a secured system \wp .*

4 Checking and enforcing security properties

Our framework provides a common formalism to define security signatures, environments, systems, and policy rules. We have shown that secured systems specified in this formalism have an operational semantics based on transition and rewriting systems and are thus executable. In this section, we go one step forward and we propose a methodology based on environment transformations for the validation of security properties enforced by a policy over a system. We then use the same technique to enforce a policy given as a security property over a system. Indeed, policy rules are often expected to fulfill a certain security property expressed on some entities, while they are dealing with a different set of entities. A typical example is given by access control policies designed for ensuring flow properties: such policies do not deal with information flow but only

with objects containing information to be traced. Intuitively, a link is needed between “what you do” (the policy rules) and “what you want” (the goal for which the policy is designed). This link is formalized here through a transformation of environments, whose aim is to translate an environment into another one dealing with the entities we are interested in. Such a transformation is defined below by a signature morphism that allows the translation of the domains and of the base of equalities and by a set of transformation rules on the base of facts.

Definition 10 (Signature morphism). A signature morphism θ from $\Sigma_1 = (\mathcal{S}_1, \mathcal{F}_1, \mathcal{P}_1)$ to $\Sigma_2 = (\mathcal{S}_2, \mathcal{F}_2, \mathcal{P}_2)$ is a pair (θ_S, θ_F) such that $\theta_S : \mathcal{S}_1 \rightarrow \mathcal{S}_2$ and $\theta_F : \mathcal{F}_1 \rightarrow \mathcal{F}_2$ are (partial or total) functions such that $\forall f : s_1, \dots, s_n \mapsto s \in \text{Dom}(\theta_F)$ where $s_1, \dots, s_n, s \in \text{Dom}(\theta_S)$, $\theta_F(f) : \theta_S(s_1), \dots, \theta_S(s_n) \mapsto \theta_S(s) \in \mathcal{F}_2$. We extend θ to a morphism $\hat{\theta}$ (which will be simply denoted by θ) over terms as follows:

- $\forall x : s \in \mathcal{X}, \hat{\theta}(x : s) = x : \theta_S(s)$
- $\forall f \in \text{Dom}(\theta_F), \hat{\theta}(f(t_1, \dots, t_n) : s) = \theta_F(f)(\hat{\theta}(t_1), \dots, \hat{\theta}(t_n))$

Definition 11 (Environment transformation). Given two signatures $\Sigma_1 = (\mathcal{S}_1, \mathcal{F}_1, \mathcal{P}_1)$ and $\Sigma_2 = (\mathcal{S}_2, \mathcal{F}_2, \mathcal{P}_2)$, an environment transformation Θ is a tuple $(\theta, \delta_\Theta, \mathcal{R})$ where:

- θ is a signature morphism from Σ_1 to Σ_2 ;
- δ_Θ is a set of pairs $\langle \text{condition}, \text{conclusion} \rangle$ where condition is a Σ_1 -constraint and conclusion is a conjunction of Σ_2 -atoms such that $x:s \in \mathcal{FVar}(\text{condition})$ iff $\theta(x:s) \in \mathcal{FVar}(\text{conclusion})$ and thus variables of $\mathcal{FVar}(\text{conclusion})$ (resp. $\mathcal{FVar}(\text{condition})$) are sorted by $\text{Dom}(\theta_S)$ (resp. $\text{Im}(\theta_S)$);
- \mathcal{R} is a set of safe and stratified logical rules over Σ_2 .

Definition 12 (Application of an environment transformation). Let $\Theta = (\theta, \delta_\Theta, \mathcal{R})$ be an environment transformation from Σ_1 to Σ_2 , and η be an environment over Σ_1 . Applying Θ on η produces an environment $\Theta(\eta)$ over Σ_2 defined as follows:

- $|\Theta(\eta)| = \{c : \theta(s) \mid c : s \in |\eta| \wedge s \in \text{Dom}(\theta)\}$;
- $\mathcal{E}_{\Theta(\eta)}$ contains an equality $\theta(f(t_1, \dots, t_n)) = \theta(t)$ for each $f(t_1, \dots, t_n) = t$ in \mathcal{E}_η whose image by θ is defined;
- $\mathcal{B}_{\Theta(\eta)}$ contains all the Σ_2 -atoms $p(\mu(t_1) \downarrow_{\Theta(\eta)}, \dots, \mu(t_m) \downarrow_{\Theta(\eta)})$ for which there exists a pair $\langle \text{condition}, \text{conclusion} \rangle \in \delta_\Theta$ where $p(t_1, \dots, t_n)$ occurs in conclusion, and a mapping μ from $\mathcal{VVar}(\text{condition})$ to $|\eta|$ such that $\llbracket \eta \rrbracket \models \mu(\text{condition})$;
- $\mathcal{R}_{\Theta(\eta)} = \mathcal{R}$.

We say that η is transformed by Θ into $\Theta(\eta)$.

Notice that, due to the partial definition of θ , to the conditions to be satisfied by the translated equalities, and to the satisfiability of the condition in η , any of the components (domain, base of equalities, base of facts, closure rules) of $\Theta(\eta)$ may be empty.

Since the rules defining the environment transformation Θ are syntactically much simpler than those of a security system in Section 3, we easily get:

Proposition 6. Any environment transformation $\Theta = (\theta, \delta_\Theta, \mathcal{R})$ from Σ_1 to Σ_2 induces a total mapping $\eta \mapsto \Theta(\eta)$ from Σ_1 -environments into Σ_2 -environments.

This operational view justifies to call Θ a **transformation operator**. Thanks to the notion of environment transformations, it becomes possible to check a security property

expressed as a Σ_2 -formula ψ over reachable environments of a secured system $\llbracket \varphi \rrbracket$ over Σ_1 . Indeed, this amounts to check that for every reachable environment η of $\llbracket \varphi \rrbracket$, we have $\llbracket \Theta(\eta) \rrbracket \models \psi$. Thanks to Proposition 1, such a property is decidable.

Example 6. We consider now environment transformations that can be used to deal with information flow properties of access control policies. First, we introduce the “generic” signature $\Sigma_{FLOW} = (\{Actor, Information\}, \mathcal{F}_{FLOW}, \mathcal{P}_{FLOW})$ where:

$$\mathcal{P}_{FLOW} = \left\{ \begin{array}{l} Get : Actor, Information; \quad MoveTo : Information, Information; \\ Put : Actor, Information; \quad Trustworthy : Actor, Information; \\ Eligible : Actor, Information; \quad Gflow : Information, Information \end{array} \right\}$$

and where \mathcal{F}_{FLOW} is an arbitrary set of function symbols. $Get(a, i)$ means that the actor a knows the information i , $Put(a, i)$ means that the actor a modifies the information i (by using the information he knows), $MoveTo(i_1, i_2)$ means that the information i_2 is enriched with information i_1 , $Eligible(a, i)$ means that the actor a is granted to know the information i , $Trustworthy(s, i)$ means that the actor a is granted to modify the information i and $Gflow(i_1, i_2)$ means that the information i_1 is authorized to flow into i_2 . The predicates Get , Put and $MoveTo$ are useful for describing existing flows while the predicates $Eligible$, $Trustworthy$, and $Gflow$ are used to specify flow policies (respectively a confidentiality policy, an integrity policy and a confinement policy). Now, it is possible to define, in a generic way, confidentiality, integrity and confinement security properties as follows:

$$\begin{array}{ll} \text{Confidentiality} & \psi_{conf} \quad \forall a, i. Get(a, i) \Rightarrow Eligible(a, i) \\ \text{Integrity} & \psi_{int} \quad \forall a, i. Put(a, i) \Rightarrow Trustworthy(a, i) \\ \text{Confinement} & \psi_{info} \quad \forall i, i'. MoveTo(i, i') \Rightarrow Gflow(i, i') \end{array}$$

Let us consider the environment transformation defined from the signature Σ of Example 1 and the signature Σ_{FLOW} and consisting of the partial function $\theta_S : \mathcal{S} \rightarrow \mathcal{S}_{FLOW}$ such that $Dom(\theta_S) = \{S, O\}$ with $\theta_S(S) = Actor$ and $\theta_S(O) = Information$ together with the identity function $\theta_{\mathcal{F}}$, the following logical rules over Σ_{FLOW} :

$$\mathcal{R}_{FLOW} = \left\{ \begin{array}{l} \vdash MoveTo(i, i) \\ Get(a, i) \wedge Put(a, i') \vdash MoveTo(a, i') \\ MoveTo(i, i') \wedge Get(a, i') \vdash Get(a, i) \\ MoveTo(i, i') \wedge Put(a, i) \vdash Put(a, i') \\ MoveTo(i, i') \wedge MoveTo(i', i'') \vdash MoveTo(i, i'') \end{array} \right\}$$

and δ_{Θ} defined by

$$\begin{array}{ll} \frac{m(x, y, read) \wedge \forall x'. \neg m(x', y, erase)}{Get(x, y)} & \frac{inf(f_o(y), f_s(x))}{Eligible(x, y)} \\ \frac{m(x, y, write) \wedge \forall x'. \neg m(x', y, erase)}{Put(x, y)} & \frac{inf(f_o(y), f_o(y'))}{Gflow(y, y')} \end{array}$$

Such an environment transformation provides means for checking that our policy ensures confinement. This can be done by checking that each reachable environment η of the secured system $\llbracket \varphi \rrbracket$ is such that: $\llbracket \Theta(\eta) \rrbracket \models \psi_{info}$. However, the existence of “sdoers” may generate reachable environments that do not satisfy the confidentiality property *w.r.t.* Θ : it becomes possible to obtain a reachable environment in $\llbracket \varphi \rrbracket$ which is transformed into an environment which does not satisfy ψ_{conf} . This is for example

the case of the environment obtained by considering the following sequence of events: $create(root, L1)$ (leading to the creation of an object O_1); $ask(Charlie, O_1, read)$. Indeed, with this sequence, we have $Get(Charlie, O_1)$ but $\neg Eligible(Charlie, O_1)$. Of course, when adding the rule $\frac{sudo(s) \wedge o : O}{Eligible(s, o)}$ which gives a different semantics to the confidentiality property, it becomes possible to check that each reachable environment η of the system $\llbracket \varphi \rrbracket$ is such that $\llbracket \Theta(\eta) \rrbracket \models \psi_{info} \wedge \psi_{conf}$.

The transformation approach can also be useful when policies are expressed as security properties: it makes possible to constrain transitions of a system in order to ensure the desired property. Indeed, suppose we want to constrain a security system $\mathfrak{S} = (\eta_{init}, \delta)$ over a signature $\Sigma_1 = \Sigma_{Sys_1} \cup \Sigma_{Ev}$ in order to ensure security properties expressed as a Σ_2 -formula φ . The corresponding secured system can be obtained from an environment transformation $\Theta = (\theta, \delta_\Theta, \mathcal{R}_2)$ from Σ_{Sys_1} to Σ_2 by considering the transition relation $\rightarrow_{\delta_\varphi}$ such that $\eta \xrightarrow{evt}_{\delta_\varphi} \eta'$ iff $(\eta \xrightarrow{evt}_\delta \eta' \wedge \llbracket \Theta(\eta') \rrbracket \models \varphi)$. Of course, such construction leads to a system whose reachable states satisfy φ iff $\llbracket \Theta(\eta_{init}) \rrbracket \models \varphi$. Hence, thanks to our notion of environment transformation, it becomes possible to apply a security policy expressed as properties to several systems.

5 Related work

Among a rich literature on security policies (see for instance [9] for policy specification languages), our approach is in the line of logic-based languages providing a well-understood formalism, which is amenable to verification, proof and analysis.

Our formalism borrows inspiration from various sources. Horn clause logic has been used extensively for RBAC models [16]. But since negation and recursion are often needed, the concept of stratified theories has been used for instance in the authorisation specification language ASL [13] for access control. Integrity rules specify application dependent conditions that limit the range of acceptable policies. Stratified logic for RBAC policies is also developed in [2]. In our work, we use similar concepts but do not restrict to RBAC models.

Constraint logic programming for designing RBAC and temporal RBAC policies is considered in [3]. Their constraints are conjunctions of equational constraints over sets of constants, and arithmetic constraints over nonnegative integers. While keeping a declarative approach, CLP adds the expressive power and efficiency of constraint solving and database querying. A security administrator has then analysis capability thanks to the computation of sets of constraints as answers. Formalisation of security analysis in an abstract setting is done in [15] and exemplified for RBAC. In comparison, we allow a different class of constraints, that we keep decidable by restricting to safe theories, and we use constraints in a rewriting context. Note that it is also possible to apply constraint narrowing to get analysis power as in [14].

Whereas most existing work on reasoning about security policies model the environment only lightly, if at all, there are some exceptions. One of the closest works is [10] who represents the behavior of access control policies in a dynamic environment. Policies are written in Datalog and can refer to facts in the authorization state. Events, such

as access requests, can change the authorization state, and the changes are specified as a state machine whose transition labels are guarded by the policy. Security properties can then be analyzed by model checking formulas in first-order temporal logic. In [4], the authors introduce a logic for specifying policies where access requests can have effects on the authorization state. The effects are explicitly specified in the language, an extension of Datalog backed on transaction logic. They also propose a proof system for reasoning about sequences of user actions. In comparison, thanks to constraint rewriting, we provide a more expressive formalism, while keeping operational and decidable. The full expressive power of constraint rewriting is explored in [7].

Comparing the expressive power of access control models is a fundamental problem in computer security, already addressed in several works. In [6], different access control models are represented in C-Datalog (an object-oriented extension of Datalog) and compared using results from logic programming. In [17], the authors express access control control systems as state transitions systems as we do and introduce security-preserving mappings, called reductions, to compare security analysis based on accessibility relations in two different models. In [8, 11], the comparison mechanism is based on a notion of simulation. Thanks to the notion of environment transformation, we address this problem with an operational transition rules based approach.

6 Conclusion and future work

As we have already said, our framework allows the definition of transformations between security signatures and environment, and consequently between secured systems. In Section 4, a transformation operator between signatures and environments has been used to check security properties over reachable environments of a secured system. Hence, such a transformation operator allows us to check a property over a system even if this property is expressed on a different signature (and/or specification). In practice, this approach provides a way to reuse the same specification of a security property in order to analyse or to verify several policies and systems, thus showing the benefits that a library of generic security properties, dedicated to particular domains (as we did for information flows) and that can be considered in several contexts, could provide. Such an approach can also be useful when policies are expressed as security properties: environment transformations allow us to constrain transitions of a system in order to ensure the desired property, and it becomes possible to apply a policy on several systems. As future work, we aim to focus on the extension of the proposed transformation. For example, the environment transformation defined in Section 4 could be extended in order to consider transformations between events of a system presentation \mathfrak{S} and events of a set of policy rules \mathfrak{R} . This allows defining in an independent way policies and systems, providing thus modularity in formal developments. Furthermore, transformation operators could also be useful to compare and to compose security policies and systems. Indeed, the comparison between two policies expressed as policy rules \mathfrak{R}_1 and \mathfrak{R}_2 , respectively based on the signatures Σ_1 and Σ_2 , is often based on an embedding of Σ_1 -formulas into Σ_2 -formulas. Such an approach can also be considered for systems, using transformations between environments to define a comparison mechanism. Similarly, for composition, transformation operators could be used to translate policies and

systems into policies and systems sharing the same security signature and specification, thus easing the definition of a composition relation.

References

1. F. Baader and T. Nipkow. *Term Rewriting and All That*. Cambridge University Press, 1998.
2. S. Barker. Access control for deductive databases by logic programming. In P. J. Stuckey, editor, *Logic Programming, 18th International Conference (ICLP) 2002*, volume 2401 of *LNCS*, pages 54–69. Springer, 2002.
3. S. Barker and P. J. Stuckey. Flexible access control policy specification with constraint logic programming. *ACM Transactions on Information and System Security*, 6(4):501–546, 2003.
4. M. Y. Becker and S. Nanz. A logic for state-modifying authorization policies. In *ESORICS 2007*, volume 4734 of *LNCS*, pages 203–218. Springer, 2007.
5. D. Bell and L. LaPadula. Secure Computer Systems: a Mathematical Model. Technical Report MTR-2547 (Vol. II), MITRE Corp., Bedford, MA, May 1973.
6. E. Bertino, B. Catania, E. Ferrari, and P. P. A logical framework for reasoning about access control model. In *ACM Transactions on Information and System Security*, volume 6, pages 71–127, 2003.
7. T. Bourdier and H. Cirstea. Constrained rewriting in recognizable theories. Technical report, INRIA, 2010. <http://hal.archives-ouvertes.fr/inria-00456848/en/>.
8. A. Chander, J. Mitchell, and D. Dean. A state-transition model of trust management and access control. In *Proceedings of the 14th IEEE Computer Security Foundation Workshop CSFW*, pages 27–43. IEEE Comp. Society Press, 2001.
9. N. C. Damianou, A. K. Bandara, M. S. Sloman, and E. C. Lupu. A survey of policy specification approaches. Technical report, Department of Computing, Imperial College of Science Technology and Medicine, London, UK, 2002. <http://www.doc.ic.ac.uk/~mss/Papers/PolicySurvey.pdf>.
10. D. J. Dougherty, K. Fislser, and S. Krishnamurthi. Specifying and reasoning about dynamic access-control policies. In *3rd International Joint Conference on Automated Reasoning (IJ-CAR)*, volume 4130 of *LNCS*, pages 632–646, 2006.
11. L. Habib, M. Jaume, and C. Morisset. Formal definition and comparison of access control models. *Journal of Information Assurance and Security*, 4(4):372–381, 2009.
12. P. Hinman. *Fundamentals of mathematical logic*. A.K. Peters, Ltd., 2005.
13. S. Jajodia, P. Samarati, V. S. Subrahmanian, and E. Bertino. A unified framework for enforcing multiple access control policies. In *SIGMOD '97: Proceedings of the 1997 ACM SIGMOD international conference on Management of data*, pages 474–485, New York, NY, USA, 1997. ACM.
14. C. Kirchner, H. Kirchner, and A. Santana de Oliveira. Analysis of rewrite-based access control policies. In D. Dougherty and S. Escobar, editors, *Proceedings of the Third International Workshop on Security and Rewriting Techniques (SecReT 2008), Pittsburgh, PA, USA, 22 June 2008*, volume 234 of *Electronic Notes in Theoretical Computer Science*, pages 55–75. Elsevier, 2009.
15. N. Li and M. V. Tripunitara. Security analysis in role-based access control. *ACM Transactions on Information and System Security (TISSEC)*, 9(4):391–420, 2006.
16. R. S. Sandhu, E. J. Coyne, H. L. Feinstein, and C. E. Youman. Role-based access control models. *Computer*, 29:38–47, 1996.
17. M. V. Tripunitara and N. Li. A theory for comparing the expressive power of access control models. *Journal of Computer Security*, 15(2):231–272, 2007.
18. J. Ullman. *Database and Knowledge - Base Systems*, volume 1: Classical Database Systems. Computer Science Press, 1988.

A Proofs

Proposition 1. *For any environment η over Σ , $\llbracket \eta \rrbracket$ exists and is unique and computable. Moreover the validity of any first-order formula in $\llbracket \eta \rrbracket$ is decidable.*

Proof. By definition, η describes a Datalog program composed of ground facts and safe stratified logical rules (any equality of the form $f(t_1, \dots, t_n) = t$ being seen as the fact $f(t_1, \dots, t_n, t)$). Thus, there exists a unique least fixed point [18]. Moreover, since $|\eta|$ is finite, we can decide the validity of a formula by evaluating it naively: existential and universal quantifications being evaluated by enumerating all possible valuations and by solving the corresponding propositional satisfaction problem.

Proposition 2. *For any disjoint set δ of transition rules satisfying \mathcal{H} , the relation $\xrightarrow{\text{evt}}_{\delta}$ is deterministic.*

Proof. We prove that no contradictory operations are performed when a rule is triggered. Assumption (1) guarantees that if two actions $p(t_1, \dots, t_n)$ and $\neg p(t'_1, \dots, t'_n)$ occur in the statement then there is only one possible instantiation (given by the event) for their variables (which are not free in the condition) and $p(t_1, \dots, t_n)$ is different from $p(t'_1, \dots, t'_n)$ since there exists at least one j such that $t_j \neq t'_j$. Assumption (2) guarantees that there is only one possible instantiation for the variables in the right-hand side of an oriented equality and since, additionally, it imposes that the left-hand sides of the different oriented equalities are not unifiable, we always obtain the same interpretation. Assumption (3) obviously guarantees that only constants not affected by the rule can be deleted. The result follows immediately for disjoint rules since at most one rule can be triggered in this case.

Proposition 3. *A set \mathfrak{R} of policy rules is η -consistent if (1) for each rule, its left-hand side contains only one occurrence of each variable and its constraint does not involve terms of sort *Query*; (2) for any two rules $l \xrightarrow{\varphi} r$ and $l' \xrightarrow{\varphi'} r'$ there exists no position ω and no substitution σ such that $\llbracket \eta \rrbracket \models \{\sigma(l|_{\omega}) = \sigma(l') \wedge \sigma(\varphi) \wedge \sigma(\varphi')\}$.*

Proof. The proof is obtained by adapting the confluence proof for orthogonal TRS [1].

Proposition 4. *A set \mathfrak{R} of policy rules over a security signature $\Sigma = \Sigma_{Sys} \cup \Sigma_{Ev}$ is η -complete if (1) the reduction $\xrightarrow{*}_{\mathfrak{R}}^{\eta}$ terminates, (2) for any $f:S_1, \dots, S_n \mapsto S \in \Sigma_{Ev}$ with $S \neq \text{Decision}$ there exists a default rule $f(x_1, \dots, x_n) \rightarrow c$ for some c , (3) each rule of \mathfrak{R} is sort-preserving or sort-decreasing (i.e. the sort of its left-hand side is equal or greater than the sort of its right-hand side).*

Proof. We prove that $q \in \mathcal{T}_{\Sigma, |\eta|}^{\text{Query}}$ has a normal form of sort *Decision* by induction on the well-founded rewrite relation induced by \mathfrak{R} . If q is not of the form $f(_, \dots, _)$ for some $f:S_1, \dots, S_n \mapsto S \in \Sigma_{Ev}$ with $S \neq \text{Decision}$, then q is necessarily of sort *Decision* and thus, its normal form is of sort *Decision*. If q has such a form, then q is reducible by \mathfrak{R} (because of the default rule) into a ground term q' which, by induction hypothesis, has a normal form of sort *Decision*.