Authorization in Distributed Systems: A Formal Approach*

Thomas Y.C. Woo Simon S. Lam
Department of Computer Sciences
The University of Texas at Austin
Austin, Texas 78712-1188

Abstract

In most systems, authorization is specified using some low-level system-specific mechanisms, e.g. protection bits, capabilities and access control lists. We argue that authorization is an independent semantic concept that must be separated from implementation mechanisms and given a precise semantics. We propose a logical approach to representing and evaluating authorization. Specifically, we introduce a language for specifying *policy bases*. A policy base encodes a set of authorization requirements and is given a precise semantics based upon a formal notion of *authorization policy*. The semantics is computable, thus providing a basis for authorization evaluation. We also introduce two composition operators for policy bases, which are appropriate for modeling distributed systems with multiple administrative domains.

1 Introduction

To guarantee the security of a distributed system, many concerns need to be addressed. These include authentication, authorization, auditing, accounting and availability, among others. In this paper, we propose a new foundation for authorization, specifically, one that is appropriate for the design and implementation of distributed systems.

The problem of authorization can be divided into two related subproblems: *representation* and *evaluation*. Representation refers to the specification of authorization requirements, while evaluation refers to the actual determination of the authorities of *subjects* given the authorization requirements. The authority of a subject is its rights to access *objects*. (Thus our view of authorization is limited to access control; we do not consider issues of covert channels and secure information flow [5, 11, 16].)

Conceptually, the rights of subjects to access objects can be stored in an *access matrix* [14, 15], with rows corresponding to subjects, columns corresponding to objects, and matrix entries indicating various *access rights*. (See examples in Section 2.) Practical implementations of an access matrix usually take advantage of the sparseness of the matrix, and are based upon capabilities (access rights stored by row), access control lists (access rights stored by columns), or some hybrid combination of these approaches [5, 7].

Distributed systems and the prevalent client-server style of computing give rise to new problems in the specification of authorization requirements. For examples:

- New kinds of attributes need to be considered. For instance, an authorization requirement in a distributed system may include the *location* of a subject as an attribute in addition to the identity of the subject. That is,

*This work was supported in part by a grant from the Texas Advanced Research Program and by National Science Foundation grant no. NCR-9004464. Published in *Proceedings IEEE Symposium on Research in Security and Privacy*, Oakland, CA, May 1992. Postscript files of this and other papers of the Networking Research Laboratory are available from http://www.cs.utexas.edu/~lam/NRL.

it is possible that a subject U is authorized to update a file F from node N but not from another node N'. Other attributes include: the role a subject is assuming, the groups a subject belongs to, any delegations a subject may have, and such.

- A large-scale distributed system is typically composed of multiple independent domains, which are managed by possibly different administrative authorities. In fact, even a single domain may have several security administrators. In these situations, authorizations in one domain may affect those in other domains in unexpected ways. For instance, let X, Y and Z be three independent domains within a distributed system administered respectively by authorities A, B and C. Suppose A authorizes requests from Y to access resources in X but denies requests from Z. If B authorizes requests from Z to access resources in Y, such authorization would indirectly contradict the one by A, because a user in Z might be able to access resources in X by "going through" domain Y.

Existing models of authorization have not been designed to address these problems [17]. Furthermore, existing approaches are unsatisfactory in the following respect: authorization requirements can only be specified using some low-level system-specific mechanisms. For example, in Unix, accesses to the file system are specified by protection bits associated with each file, and authorization is determined by how these protection bits are set. Such embedding of authorization requirements into mechanisms presents serious drawbacks. First, authorization requirements are limited to those that can be specified by these low-level mechanisms. Second, the semantics of authorization is dependent on the semantics of the low-level mechanisms, which is not formally defined and indeed may vary from one implementation to another.[1] This poses problems in large-scale distributed systems with heterogeneous implementations.

For example, many people have recognized the limitations of protection bits in Unix and have proposed various ad-hoc extensions to it. Each of these extensions addresses one type of authorization requirements or another without solving the above problems as a whole. Furthermore, there can be subtle interactions among these extensions, which may render a security administrator unable to comprehend what actually has been authorized.[2] In fact, such confusion can be a major source of security violations.

We advocate the following view: Authorization is an independent semantic concept that should be separated from its implementation in system-specific mechanisms. For its representation, we need a language that is expressive enough for specifying commonly encountered authorization requirements. The language must be given a formal semantics so that the meaning of an authorization requirement stated using the language can be precisely determined. This way, a security administrator is able to reconcile easily between what he intends to authorize with what he has actually authorized.

With this approach, authorization evaluation reduces to computation of semantics. The complexity of such computation is highly dependent on the particular language used. The computation mechanism can range from a trivial table lookup (e.g. if the language is simply an access matrix) to a full-fledged theorem proving procedure (e.g. if the language is first order logic). In general, the more expressive the representation language, the more complex the computation mechanism. Thus issues of representation and evaluation must be examined hand in hand with careful consideration of various tradeoffs.

In this paper, we propose a new foundation for representing and evaluating authorization. Our contributions are as follows. We first identify three types of structural properties inherent in

[1] A vivid example of this is the assortment of *setuid*/*setgid* function calls available in different flavors of Unix.

[2] See for example [13] and the POSIX Security Draft Standard P1003.6 which discuss how to supplement Unix protection bits with access control lists.

authorization requirements. We argue that such structural properties can be effectively exploited to reduce the complexity of representing and evaluating authorization in large-scale distributed systems. We introduce a representation language in which the structural properties can be represented in a straightforward manner. The language is designed to specify *policy bases*. A policy base encodes a set of authorization requirements and is given a precise semantics based upon a formal notion of *authorization policy*. The semantics is computable via a translation to *extended logic programs* (see Theorem 2 in Section 6.3), thus providing an efficient computation mechanism based on the interpretation of extended logic programs. We also introduce two composition operators for policy bases, which are appropriate for modeling distributed systems with multiple administrative domains.

In relating our research to previous work, we observe that our concerns are orthogonal to those of others in security modeling [11, 10, 21, 22]. These references are concerned with modeling abstract security properties of a system as a whole, which includes authorization as a key component. The papers by Abadi, et al. [2] and Lunt [19] are similar in spirit to ours, in that their focus is on understanding the semantics of authorization. Concrete models such as those proposed in [6, 12, 18] address the same general concerns as ours, but for application-specific domains. Lastly, our composition operators are designed for authorization requirements, and are different from the one in [20] which is designed for a particular notion of security.

The balance of this paper is organized as follows. In the next section, we identify three types of structural properties in authorization requirements. In Section 3, we discuss the requirements of a language for representing authorization. In Section 4, we present our model of authorization. In Section 5, we introduce authorization policy as a semantic notion. In Section 6, we introduce our language for specifying policy bases and describe its syntax and semantics as well as some guidelines for its usage. In Section 7, we provide some examples of policy bases, including the *Bell-LaPadula* model [4] and some inheritance rules. In Section 8, we examine two notions of composition for policy bases. In Section 9, we outline an implementation of our model. Lastly, in section 10, we provide some concluding remarks.

2 Three Types of Structural Properties

Authorization requirements are highly structured because the set of subjects and the set of objects in a system are usually highly structured. For example, users belonging to the same working groups are likely to share similar authorizations; while objects pertaining to a common task are usually given similar authorizations.

To illustrate such structures, we look at some examples. Consider the authorization specified by the following access matrix:

	$P.src$	$P.exe$	$P.doc$
A	r, w	e, w	r, w
B		e	r

Subject A, who is the developer of software P, can read/write the source file $P.src$, execute and write the executable file $P.exe$, and read/write the documentation file $P.doc$; while subject B, who is a user of P, is only allowed to execute $P.exe$ and read $P.doc$. Certain structures in the authorization are readily apparent:

(1) A, being the developer of P, must be able to update all the files related to P, i.e. A must have write access to all three files. Similarly, B should be allowed to read the documentation $P.doc$ if he is allowed to execute $P.exe$.

(2) Denials of access rights are represented implicitly by their absence in an entry. (Thus explicit denials are not possible and moreover, a denial is indistinguishable from a lack of information about an authorization.)

We call the structures exhibited in (1) *closure* properties among authorizations. In general, a

closure property stipulates that a set of authorizations should either be simultaneously authorized or denied, because a partial authorization produces an "unusable" system. Closure properties can be used to ensure the "consistency" of authorizations as in the above example or to derive new authorizations from existing ones.

The structure exhibited in (2) is called a *default* property. A default property can be used as a convention to represent implicit knowledge as in the example above (i.e. absence implies denial) or as a deduction rule when information is incomplete. In fact, most real systems employ default properties in one way or another. For example, two kinds of policies are typically used: a *restrictive* policy is one whereby a request is denied unless explicitly authorized and a *permissive* policy is one whereby a request is always granted unless explicitly denied. Both make use of default properties.

We now turn to another example. Consider the authorization specified by the following access matrix (where, for an access right a, we use $\neg a$ to denote its explicit denial):

	$F.1$	$F.2$	F	H
A			e	
G_1	r	w		r
G_2	r	$\neg w$	$\neg e$	

Suppose A is a member of groups G_1 and G_2, and $F.1$ and $F.2$ are two components of an object F (e.g. two tables in a database). Several questions can be asked about the authorization:

(1) G_1 is authorized to read $F.1$. Is A, who is a member of G_1, also authorized to do the same? This is easy to resolve as both groups to which A belongs are allowed to read $F.1$; hence, A should be authorized to read $F.1$. Consider now file H for which only one of the two groups is authorized to read. Is A authorized to read H? The answer is not obvious. So is $F.2$ for which G_1 and G_2 have been given opposite authorizations.

(2) Consider object F. A is authorized to execute F. However, G_2 to which A belongs is explicitly denied the same access. Does the denial of G_2 revoke the authorization of A? Or does A's explicit authorization override the denial of G_2?

All of the above questions can be answered by precisely defining another kind of structural properties, called *inheritance* properties. Inheritance is especially important in large-scale distributed systems where the granularity of authorization ranges from an individual to an entire domain. Inheritance properties are used to relate authorizations specified with these different granularities.

In sum, it is important that the structural properties described in this section be exploited to obtain succinct representation and efficient evaluation of authorization requirements.

3 Language Requirements

From the above discussions, a language for representing authorization requirements should satisfy the following criteria:

- it should be declarative and have a semantics that is independent of implementation mechanisms
- the semantics should be efficiently computable, hence allowing efficient authorization evaluation
- it should allow easy expression of the closure, default and inheritance properties discussed in Section 2.

In the following, we discuss four more requirements for such a language.

First, authorization is *nonmonotonic*. That is, if a set of authorization requirements is augmented by a new requirement, a subject who was previously allowed access to an object may no longer be allowed the same access. A good example of such nonmonotonicity arises in the use of defaults. For example, suppose the set of authorization requirements includes the following default:

if s is not explicitly denied read access to o
then by default s is allowed read access to o

If the set of requirements is later augmented with an explicit requirement denying s read access to o, the previous grant should be retracted. Thus, the semantics of a language for authorization must allow such nonmonotonic behavior.

Second, authorization may be *incomplete*. That is, there may be authorization requests such that insufficient information is available to determine if they should be granted or denied. Such incompleteness should be allowed in the semantics of a language for authorization. There are two reasons:

- An incompleteness may be the result of an oversight or error on the part of security administrators. Thus when an incompleteness is detected, it can serve as an alarm signalling potentially more serious problems. Therefore, it is advantageous that such incompleteness not be masked out automatically by the language semantics.

- An incompleteness may be intentional so that it can be "filled in" later when composition is performed (see below). Thus, it is important that such intentional incompleteness be allowed by the language semantics.

Note that this strictly generalizes the idea of a *reference monitor* [7], where no incompleteness is allowed.

Third, authorization may be *inconsistent*. That is, it is possible for an authorization request to be both granted and denied. The reasoning is similar to that of incomplete authorization: An inconsistent authorization may signal errors on the part of security administrators or they can arise from the composition of authorization requirements, especially in a large scale distributed system. Therefore, the semantics of a language for authorization must be able to handle inconsistencies.

Fourth, multiple authorities may coexist in a distributed system environment. These authorities can be peers who coadminister a system or they can be hierarchically related in a supervisor-subordinate fashion. Each of them may contribute authorization requirements pertinent to the part of the system he is concerned with. The authorization of the entire system is a *composition* of these individually contributed authorization requirements. Thus a language for authorization should include operators for composing authorization requirements. (Composition is discussed in greater detail in Section 8.)

4 Our Model

Our model of authorization is shown in Figure 1. Before a subject s can perform a particular access r on an object o, s must first obtain the access right r for o. Subject s does so by submitting a request of the form $req(r, s, o)$ to the *authorization module*, which responds with $grant(r, s, o)$, $deny(r, s, o)$ or $fail(r, s, o)$. A $grant(r, s, o)$ is returned if the authorization module can determine that s is authorized to have r access to o, while a $deny(r, s, o)$ is returned if the authorization module can determine that s is denied r access to o. A $fail(r, s, o)$ is returned if the authorization module fails to establish either one of the previous two cases.

To make an authorization decision, the authorization module consults the authorization requirements and the system state. The system state is needed for authorization requirements that contain system state variables as parameters. Some examples of this kind of authorization requirements are "At most 5 copies of a program P can be running concurrently in all nodes of the system" and "User A is allowed to execute program P only if the current system load is less than 2".

In our model, an authorization requirement is stated as a rule and a collection of such rules constitutes a policy base (see Section 6). There can be multiple policy bases in a system, each corresponding to an administrative domain or the

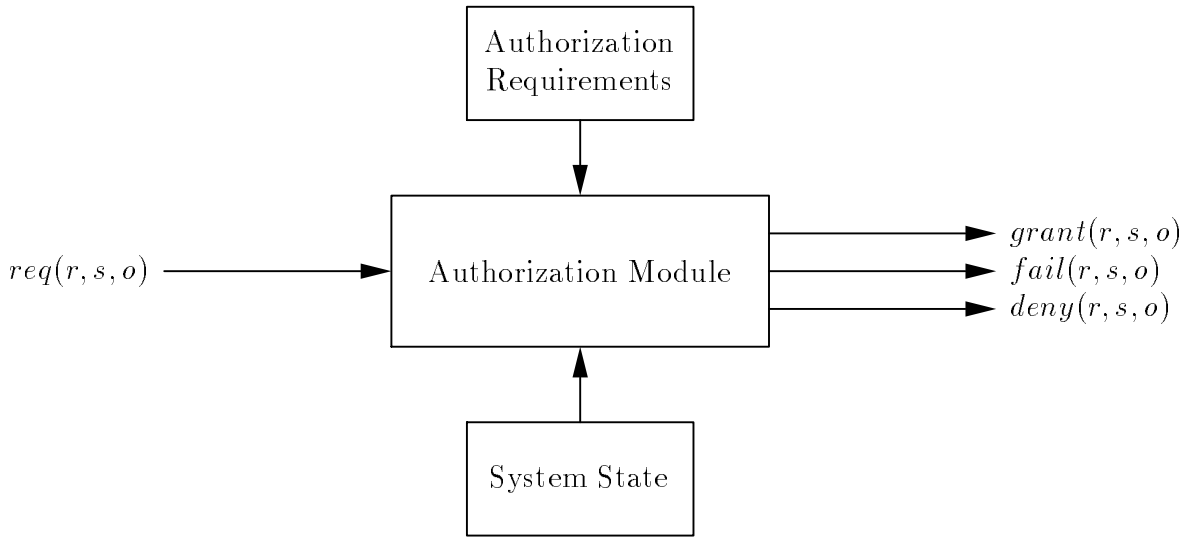

Figure 1: Model of authorization

jurisdiction of a security administrator. A composition of all policy bases defines the authorization in the entire system.

The authorization module is an interpreter which takes as input a set of policy bases $\{B_i\}$, the current system state, and a request $req(r, s, o)$, and tries to verify that either $grant(r, s, o)$ or $deny(r, s, o)$ "follows" from the composition of $\{B_i\}$ and the current system state. If $grant(r, s, o)$ follows, the request is granted. If $deny(r, s, o)$ follows, the request is denied. If neither follows, a $fail(r, s, o)$ is returned. The pathological case in which both $grant(r, s, o)$ and $deny(r, s, o)$ follow can be resolved by enforcing certain priorities between grants and denials. A precise definition of the "follows" relation is given by a formal semantics of policy bases to be presented below.

Note that Figure 1 is actually a simplified picture of our model. In general, the policy bases can be located in different parts of a distributed system, and multiple instantiations of the authorization module can be running concurrently across the system. More discussions on implementation are provided in Section 9.

5 Authorization Policy

Informally, an authorization policy is the set-theoretic equivalent of an access matrix. Its precise meaning is defined in what follows.

Definition An *authorization policy* (or *policy* in short) over a set of subjects S, a set of objects O and a set of access rights R is a 4-tuple (P^+, P^-, N^+, N^-) where each component is a subset of $\{(r, s, o) \mid r \in R, s \in S, o \in O\}$. □

The intuitive meaning of a policy $A = (P^+, P^-, N^+, N^-)$ is as follows: P^+ records the rights that are explicitly granted, i.e. if $(r, s, o) \in P^+$, subject s is explicitly granted access right r to object o. Similarly, N^+ records the rights that are explicitly denied. P^- (N^- respectively) records those rights that should not be explicitly granted (denied respectively) under this policy. P^- and N^- are useful for defining the semantics of policy composition.

A policy $A = (P^+, P^-, N^+, N^-)$ is *sound* if there does not exist a triple $t = (r, s, o)$ such that $t \in P^+ \cap P^-$ or $t \in N^+ \cap N^-$. A policy $A = (P^+, P^-, N^+, N^-)$ is *strongly sound* if it is sound and $P^+ \cap N^+ = \emptyset$.

A policy $A = (P^+, P^-, N^+, N^-)$ is *complete* if for all $s \in S$, $o \in O$ and $r \in R$, $(r, s, o) \in P^+ \cup$

$P^- \cup N^+ \cup N^-$. A policy $A = (P^+, P^-, N^+, N^-)$ is *strongly complete* if it is complete and both P^- and N^- are empty. Thus it is sufficient to represent a strongly complete policy with an ordered pair (P, N).

Given a strongly sound policy $A = (P^+, P^-, N^+, N^-)$. We can define three *authorization relations* between A and a triple (r, s, o):

$$\begin{aligned} A \textit{ grants } (r,s,o) &\quad \text{iff} \quad (r,s,o) \in P^+ \\ A \textit{ denies } (r,s,o) &\quad \text{iff} \quad (r,s,o) \in N^+ \\ A \textit{ fails } (r,s,o) &\quad \text{iff} \quad (r,s,o) \notin P^+ \cup N^+ \end{aligned}$$

Authorization evaluation can proceed as follows: Given a request from a subject s for access r to an object o, $grant(r, s, o)$ is returned if A *grants* (r, s, o), $deny(r, s, o)$ is returned if A *denies* (r, s, o) and $fail(r, s, o)$ is returned if A *fails* (r, s, o). Note that if A is also strongly complete, then $fail(r, s, o)$ would never be returned.

6 Policy Base

In this section, we present a language for stating authorization requirements in policy bases. The language is essentially a many-sorted first-order language with a *rule* construct. The rule construct is similar to the default construct in *default logic* [24]; however, we provide it with a different semantics. The rule construct is useful for stating structural properties of authorization requirements.

From some domain-specific considerations, we impose several restrictions on the kind of first-order formulas allowed. We briefly describe the restrictions and their motivations:

- We desire to have a computable semantics. As validity in an infinitary theory is typically semi-decidable, we restrict ourselves to finitary theories. To achieve this, we do not allow function symbols in our language and postulate only finite sets of access rights, subject and object constants. This also allows us to eliminate quantifications.

- We allow the use of disjunction only in highly restricted ways. For example, we cannot state in our language the authorization requirement "Subject A is either allowed to read file F or write file G". Neither can we state "There is a subject x who can read file F" in our language. Our view is that such disjunctive authorization requirements provide insufficient information for determining the exact extent of authorization.

  On a closer look, this limitation is not as restrictive as it seems. In a realistic authorization policy, disjunctive authorization requirements are stated mostly for convenience and their disjunctive nature is usually immediately resolved when other requirements are taken into consideration. This is analogous to the case in classical logic where the statement $A \vee B$ when combined with $\neg A$ yields B, which is non-disjunctive. Purely disjunctive authorization requirements are rare and counterintuitive.

6.1 Syntax

The alphabet of our language is derived from the system to be modeled. Consider a system with S as its set of subjects, O its set of objects and R its set of access rights. (Note that S, O and R are all finite sets.) We postulate the following alphabet for our language:

- a set of *ordinary* variables $\mathcal{V}$,
- a set of *propositional* variables $\mathcal{P}$,
- two propositional constants T and F
- a finite set of subject constants $\mathcal{S}$,
- a finite set of object constants $\mathcal{O}$,
- a finite set of binary predicate symbols $\mathcal{R} = \{r^+, r^- \mid r \in R\}$
- two special predicate symbols "=", and "$\in$".

The set $\mathcal{S}$ ($\mathcal{O}$) contains at least a constant symbol for each subject (object) in S (O). In other words,

each subject or object in the system is explicitly represented by a constant symbol in the language.

A *term* is an ordinary variable, a subject constant or an object constant. An *atom* is a propositional constant, a propositional variable or a predicate $p(t,t')$ where p is a predicate symbol and t,t' are terms. We adopt the convention of writing predicates involving $=$ or $\in$ in the infix form, i.e. we write $=(t,t')$ as $t = t'$ and $\in(t,t')$ as $t \in t'$. An atom formed from a predicate symbol in $\mathcal{R}$ is called a *distinguished* atom; and the rest *ordinary* atoms.

A *literal* is an atom or the negation of an atom. Negation is denoted by the symbol $\neg$. A literal formed from a distinguished atom is called a *distinguished* literal, while a literal formed from an ordinary atom is called an *ordinary* literal. A literal is *positive* if it is an atom, and *negative* if it is the negation of an atom. Let a be an atom, then the two literals a and $\neg a$ are called *complementary* literals. We define $\overline{a}$ to be $\neg a$ and $\overline{\neg a}$ to be a. Thus, ℓ and $\overline{\ell}$ are always complementary for any literal ℓ

A *formula* is a literal, a conjunction of two formulas f and f', denoted by $f \wedge f'$, or a disjunction of two formulas f and f', denoted by $f \vee f'$. A *basic formula* is a formula that only contains propositional constants and distinguished literals. A subclass of basic formulas that does not contain disjunctions is called *conjunctive formulas*. Note that in our formulas, unlike those of first order logic, negation occurs only at the level of literals. A formula is *closed* if it does not contain ordinary variables, otherwise it is *open*.

A *rule* is written in the form $\frac{f : f'}{g}$ where f is a formula, f' a basic formula and g a conjunctive formula. f, f' and g are respectively called the *prerequisite*, *assumption* and *consequent* of the rule.

Notation To simplify our presentation, we introduce a syntactic operator *neg* for basic formulas. Let f be a basic formula:

- $neg(\mathsf{T})$ is F,
- $neg(\mathsf{F})$ is T,
- f is a positive literal a, then $neg(f)$ is $\neg a$,
- f is a negative literal $\neg a$, then $neg(f)$ is a,
- f is a conjunction $f_1 \wedge f_2$, then $neg(f)$ is $neg(f_1) \vee neg(f_2)$,
- f is a disjunction $f_1 \vee f_2$, then $neg(f)$ is $neg(f_1) \wedge neg(f_2)$.

Thus the effect of the *neg* operator is similar to that of applying negation to the entire basic formula and then pushing it inward using De Morgan's law. □

Convention To be succinct, we use several abbreviations. First, if any component formula is missing from a rule, it is assumed to be T. Second, we use the notation $f \Rightarrow g$ to represent a rule of the form $\frac{f : \mathsf{T}}{g}$. Third, $\mathsf{T} \Rightarrow g$ is further abbreviated to g. □

Example 1 Let $\mathcal{V} = \{x, y, \ldots\}$, $\mathcal{P} = \{p, q\}$, $\mathcal{S} = \{\mathsf{A}, \mathsf{B}, \mathsf{G}\}$, $\mathcal{O} = \{\mathsf{X}, \mathsf{Y}, \mathsf{Z}\}$, and $R = \{\mathsf{read}, \mathsf{write}\}$. Then the following are rules:

$$\mathsf{read}^-(\mathsf{G}, x)$$
$$\mathsf{read}^+(\mathsf{A}, \mathsf{X}) \Rightarrow \mathsf{read}^+(\mathsf{A}, \mathsf{Y})$$
$$x \in \mathsf{G} \wedge \mathsf{read}^-(\mathsf{G}, \mathsf{Y}) \Rightarrow \mathsf{read}^-(x, \mathsf{Y})$$
$$\neg p \vee \mathsf{write}^+(x, \mathsf{Z}) \Rightarrow \neg\mathsf{read}^+(x, y)$$
$$\frac{p \wedge \mathsf{read}^+(x, \mathsf{Z}) : \mathsf{read}^+(x, \mathsf{Y})}{\mathsf{read}^+(x, \mathsf{Y})}$$
$$\frac{x \in G \wedge \mathsf{write}^+(\mathsf{G}, y) : \mathsf{write}^+(x, y) \wedge \neg\mathsf{write}^-(x, y)}{\mathsf{write}^+(x, y)}$$

□

Definition A *policy base* (or *base* in short) is a finite set of rules. □

A rule $\frac{f : f'}{g}$ is *closed* if f, f' and g are all closed; otherwise it is *open*. A rule is *pure* if f is also a basic formula. A base is *closed* if it contains only closed rules. A base is *pure* if it contains only pure rules.

6.2 Semantics for Closed Policy Base

We present a semantics for closed bases here. The semantics for open bases is in Section 6.4. Thus every mention of base in this subsection is taken to mean a closed base unless explicitly stated otherwise.

The semantics of a base is given by its *extensions*. An extension is a set of distinguished literals[3]. An extension provides a straightforward interpretation for distinguished literals: such a literal is true if and only if it is contained in the extension. An extension is similar in concept to a model in the standard semantics for classical logic.

An extension naturally defines a policy. More precisely, let Σ be an extension and let

$$\begin{aligned} P^+ &= \{(r,s,o) \mid r^+(s,o) \in \Sigma\} \\ P^- &= \{(r,s,o) \mid \neg r^+(s,o) \in \Sigma\} \\ N^+ &= \{(r,s,o) \mid r^-(s,o) \in \Sigma\} \\ N^- &= \{(r,s,o) \mid \neg r^-(s,o) \in \Sigma\} \end{aligned}$$

Clearly, (P^+, P^-, N^+, N^-) is a policy. We call it the *policy defined by* Σ, and denote it by $\alpha(\Sigma)$. This establishes a one-one correspondence between an extension and a policy.

To provide meanings for ordinary literals, we use an *assignment* function $\mathcal{I} : \mathcal{P} \mapsto \{\mathsf{true}, \mathsf{false}\}$ and a *group relation* $\mathcal{G}$. An assignment function provides interpretation for propositional variables, while a group relation provides interpretation for the predicate "$\in$"; they together model the system state. The equality predicate "$=$" is interpreted as the identity relation. We also adopt the *unique names assumption*, i.e. $c \neq c'$ for all c, c' in $\mathcal{S} \cup \mathcal{O}$.

Before we give our definition for extension, we first define a *satisfaction* relation between a set Σ of distinguished literals and a closed formula f with respect to an assignment $\mathcal{I}$ and a group relation $\mathcal{G}$. We denote the satisfaction relation by $\Sigma \models_{\mathcal{I},\mathcal{G}} f$. The definition is by structural induction:

- f is a propositional constant, then
  $\Sigma \models_{\mathcal{I},\mathcal{G}} f$ iff f is T
- f is a propositional variable, then
  $\Sigma \models_{\mathcal{I},\mathcal{G}} f$ iff $\mathcal{I}(f) = \mathsf{true}$
- f is $t = t'$, then
  $\Sigma \models_{\mathcal{I},\mathcal{G}} f$ iff $t \equiv t'$
- f is $t \in t'$, then
  $\Sigma \models_{\mathcal{I},\mathcal{G}} f$ iff $(t,t') \in \mathcal{G}$
- f is a distinguished literal L, then
  $\Sigma \models_{\mathcal{I},\mathcal{G}} f$ iff $L \in \Sigma$
- f is $\neg f'$ where f' is an ordinary literal, then
  $\Sigma \models_{\mathcal{I},\mathcal{G}} f$ iff $\Sigma \not\models_{\mathcal{I},\mathcal{G}} f'$
- f is $f_1 \wedge f_2$, then
  $\Sigma \models_{\mathcal{I},\mathcal{G}} f$ iff $\Sigma \models_{\mathcal{I},\mathcal{G}} f_1$ and $\Sigma \models_{\mathcal{I},\mathcal{G}} f_2$
- f is $f_1 \vee f_2$, then
  $\Sigma \models_{\mathcal{I},\mathcal{G}} f$ iff $\Sigma \models_{\mathcal{I},\mathcal{G}} f_1$ or $\Sigma \models_{\mathcal{I},\mathcal{G}} f_2$

Note that our semantics is different from the standard semantics for classical logic in several ways. First, $F \wedge \neg F$ represents a contradiction in classical logic and hence does not admit any model. In our case, we have $\{F, \neg F\} \models_{\mathcal{I},\mathcal{G}} F \wedge \neg F$. Second, in classical logic, if Σ satisfies both $F \vee G$ and $\neg F$, then it must also satisfy G. This is not true in our semantics as $\{F, \neg F\} \models_{\mathcal{I},\mathcal{G}} F \vee G$ and $\{F, \neg F\} \models_{\mathcal{I},\mathcal{G}} \neg F$, but $\{F, \neg F\} \not\models_{\mathcal{I},\mathcal{G}} G$. A semantics that exhibits such non-classical behavior is often called *paraconsistent*.

Given a base, we are now ready to define its extensions. Let B be a base, $\mathcal{I}$ an assignment and $\mathcal{G}$ a group relation. We define an operator $\Gamma_{B,\mathcal{I},\mathcal{G}}$ that given a set of distinguished literals, returns a new set of distinguished literals. The formal definition of $\Gamma_{B,\mathcal{I},\mathcal{G}}$ is as follows: Let Σ be a set of distinguished literals. Define

$$S^{\mathcal{I},\mathcal{G}}_{B,\Sigma} = \left\{ M \;\middle|\; \begin{array}{l} M \text{ is a set of} \\ \text{distinguished literals and} \\ \text{for all } \frac{f : f'}{g} \in B, \\ \text{if } M \models_{\mathcal{I},\mathcal{G}} f \text{ and } \Sigma \not\models_{\mathcal{I},\mathcal{G}} neg(f') \\ \text{then } M \models_{\mathcal{I},\mathcal{G}} g \end{array} \right\}$$

then

$$\Gamma_{B,\mathcal{I},\mathcal{G}}(\Sigma) = \text{the intersection of all elements in } S^{\mathcal{I},\mathcal{G}}_{B,\Sigma}$$

The intuitive meaning of a rule $\frac{f : f'}{g}$ is as follows: If a set Σ of distinguished literals satisfies f, and there is no evidence that the negation of f'

[3] An extension is similar to a Herbrand base except that it contains literals instead of just atoms.

is satisfied (hence it is consistent to assume that Σ satisfies f'), then Σ must also satisfy g.

Definition Let B be a base, $\mathcal{I}$ an assignment, $\mathcal{G}$ a group relation and Σ a set of distinguished literals. Σ is an *extension* of B under assignment $\mathcal{I}$ and group relation $\mathcal{G}$ if $\Sigma = \Gamma_{B,\mathcal{I},\mathcal{G}}(\Sigma)$, i.e. Σ is a fixed point of the operator $\Gamma_{B,\mathcal{I},\mathcal{G}}$. □

Since each extension of a base defines a policy, in the case that B admits a *unique* extension Σ under $\mathcal{I}$ and $\mathcal{G}$, the policy defined by Σ can be taken to be the semantics of B under $\mathcal{I}$ and $\mathcal{G}$. We formalize this in the following definition.

Definition Let B be a base, $\mathcal{I}$ an assignment and $\mathcal{G}$ group relation. Suppose B admits a unique extension Σ under $\mathcal{I}$ and $\mathcal{G}$. Then $\alpha(\Sigma)$, the *policy determined by B under $\mathcal{I}$* and $\mathcal{G}$, will be denoted by $\mathcal{E}_{\mathcal{I},\mathcal{G}}(B)$. □

The authorization relations introduced in Section 5 can be naturally extended to a base as follows. (Note that this is well defined only when $\mathcal{E}_{\mathcal{I},\mathcal{G}}(B)$ itself is well-defined and is a strongly sound policy.)

Definition Let B be a base, $\mathcal{I}$ an assignment and $\mathcal{G}$ group relation. Let $s \in S$, $o \in O$ and $r \in R$:

$$
\begin{array}{l}
B \textit{ grants } (r,s,o) \textit{ under } \mathcal{I},\mathcal{G} \text{ iff} \\
\quad \mathcal{E}_{\mathcal{I},\mathcal{G}}(B) \textit{ grants } (r,s,o) \\
B \textit{ denies } (r,s,o) \textit{ under } \mathcal{I},\mathcal{G} \text{ iff} \\
\quad \mathcal{E}_{\mathcal{I},\mathcal{G}}(B) \textit{ denies } (r,s,o) \\
B \textit{ fails } (r,s,o) \textit{ under } \mathcal{I},\mathcal{G} \text{ iff} \\
\quad \mathcal{E}_{\mathcal{I},\mathcal{G}}(B) \textit{ fails } (r,s,o)
\end{array}
$$

These three relations represent the authorization defined by a base B, and are taken to be its semantics. □

Clearly, the above semantics is well-defined only when B admits a unique extension under $\mathcal{I}$ and $\mathcal{G}$. However, as shown in the examples below, this unique extension property is not true in general.

Example 2 Consider the base

$$
B_1 = \{ \frac{: \mathsf{read}^+(\mathsf{A},\mathsf{X}) \wedge \neg \mathsf{read}^+(\mathsf{A},\mathsf{Y})}{\mathsf{read}^+(\mathsf{A},\mathsf{X})}, \frac{: \mathsf{read}^+(\mathsf{A},\mathsf{Y}) \wedge \neg \mathsf{read}^+(\mathsf{A},\mathsf{Z})}{\mathsf{read}^+(\mathsf{A},\mathsf{Y})}, \frac{: \mathsf{read}^+(\mathsf{A},\mathsf{Z}) \wedge \neg \mathsf{read}^+(\mathsf{A},\mathsf{X})}{\mathsf{read}^+(\mathsf{A},\mathsf{Z})} \}
$$

B_1 does not admit any extension under all assignments and group relations. Note that B_1 does not contain any ordinary literal, thus the assignment or group relation cannot be the cause for its lack of extension. □

Example 3 Consider the base

$$
B_2 = \{ \frac{: \neg \mathsf{write}^+(\mathsf{A},\mathsf{X})}{\mathsf{write}^+(\mathsf{A},\mathsf{Y})}, \frac{: \neg \mathsf{write}^+(\mathsf{A},\mathsf{Y})}{\mathsf{write}^+(\mathsf{A},\mathsf{X})} \}
$$

B_2 admits two extensions $\{\mathsf{write}^+(\mathsf{A},\mathsf{Y})\}$ and $\{\mathsf{write}^+(\mathsf{A},\mathsf{X})\}$ under all assignments and group relations. □

Example 4 Consider the base

$$
B_3 = \{ \mathsf{read}^+(\mathsf{A},\mathsf{X}), \frac{p \wedge \mathsf{read}^+(\mathsf{A},\mathsf{X}) : \neg \mathsf{write}^-(\mathsf{A},\mathsf{Z})}{\mathsf{write}^-(\mathsf{A},\mathsf{Z})} \}
$$

If $\mathcal{I}(p) = \mathsf{false}$ then $\{\mathsf{read}^+(\mathsf{A},\mathsf{X})\}$ is an extension. However if $\mathcal{I}(p) = \mathsf{true}$, B_3 does not admit any extension. Thus both $\mathcal{I}$ and $\mathcal{G}$ do affect the extensions (if any) of a base. □

Although these examples demonstrate that the unique extension property is not true in general, they also serve to illustrate a common underlying cause for failure. In the above examples, there is a kind of circularity in the rules involving atoms that occur both positively and negatively. For instance, in Example 3, each of the atoms $\mathsf{write}^+(\mathsf{A},\mathsf{X})$ and $\mathsf{write}^+(\mathsf{A},\mathsf{Y})$ occurs positively in the consequent of one rule but negatively in the assumption of the other rule. Thus the application of one rule would necessarily disable the application of the other rule, hence resulting in two different extensions. However, if priorities are enforced between the two rules (e.g. the derivation of $\mathsf{write}^+(\mathsf{A},\mathsf{X})$ is more "important" than the derivation of $\mathsf{write}^+(\mathsf{A},\mathsf{Y})$), then only $\{\mathsf{write}^+(\mathsf{A},\mathsf{X})\}$ would be considered an extension of B_2. This idea can indeed be generalized and a notion of *stratification* can be defined on the set

of distinguished atoms, so that a stratified base always possesses a unique extension. We omit the details here and refer the readers to [3, 8, 23].

The semantics of bases can also be given by first "factoring out" the effects of assignments and group relations. We formalize this below.

Let B be a base, $\mathcal{I}$ an assignment and $\mathcal{G}$ a group relation. Let $d \equiv \frac{f \,:\, f'}{g}$ be a rule in B. Suppose we apply the following transformation to d:

- replace all occurrences of p in f by T if $\mathcal{I}(p) = \mathsf{true}$ and F otherwise
- replace all occurrences of $t \in t'$ in f by T if $(t, t') \in \mathcal{G}$ and F otherwise
- replace all occurrences of $t = t'$ in f by T if $t \equiv t'$ and F otherwise

Clearly, the resulting rule is pure. We denote by $[\mathcal{I}, \mathcal{G}](B)$ the base obtained by applying the above transformation to each rule in B. By definition, $[\mathcal{I}, \mathcal{G}](B)$ is a pure base. Note that the extensions of a pure base are independent of the assignment and the group relation.

Theorem 1 Let B be a base, $\mathcal{I}$ an assignment and $\mathcal{G}$ a group relation. Let Σ be a set of literals. Then Σ is an extension of B under $\mathcal{I}$ and $\mathcal{G}$ iff Σ is an extension of $[\mathcal{I}, \mathcal{G}](B)$.

Proof Omitted; see [26]. □

6.3 Computation of $\mathcal{E}_{\mathcal{I},\mathcal{G}}(B)$

For our semantics, authorization evaluation reduces to the computation of $\mathcal{E}_{\mathcal{I},\mathcal{G}}(B)$. In this subsection, we present a simple semantics-preserving translation of a base B into an *extended logic program* Π_B, thus reducing the computation of $\mathcal{E}_{\mathcal{I},\mathcal{G}}(B)$ to the computation of Π_B [9].

We first introduce the concept of an extended logic program. An *extended program clause* is a statement of the form:

$$L \leftarrow L_1, \ldots, L_n, \mathsf{not}\ L_{n+1}, \ldots, \mathsf{not}\ L_m$$

where L and L_i's are literals. An *extended logic program* is a finite collection of extended program clauses. Extended logic programs are a strict superset of general logic programs, because literals rather than just atoms are allowed in the program clauses.

For extended logic programs, we have developed a paraconsistent semantics (expressed in terms of *models*) using ideas from stable model construction [8]. Our semantics is an extension to the one proposed in [9], and is similarly computable via reduction to general logic programs. Details can be found in [26].

The essence of our approach is to translate a base into an extended logic program as follows: Let B be a base and let $d \equiv \frac{f \,:\, f'}{g}$ be a rule in B. We translate d into the extended program clause

$$g \leftarrow f \wedge not(neg(f'))$$

where *not* is an operator with a definition similar to *neg*: Let h be a basic formula:

- $not(\mathsf{T})$ is F,
- $not(\mathsf{F})$ is T,
- h is a literal ℓ, then $not(h)$ is $\mathsf{not}\ \ell$,
- h is a conjunction $h_1 \wedge h_2$, then $not(h)$ is $not(h_1) \vee not(h_2)$,
- h is a disjunction $h_1 \vee h_2$, then $not(h)$ is $not(h_1) \wedge not(h_2)$.

We denote by Π_B the extended program obtained by applying the above translation to each rule in B.

Theorem 2 Let B be a pure base and Σ a set of literals. Then Σ is an extension of B iff Σ is a model of Π_B.

Proof Omitted; see [26]. □

Corollary Let B be a base, $\mathcal{I}$ an assignment and $\mathcal{G}$ a group relation. Let Σ be a set of literals. Then Σ is an extension of B under $\mathcal{I}$ and $\mathcal{G}$ iff Σ is a model of $\Pi_{[\mathcal{I},\mathcal{G}](B)}$.

Proof Omitted; see [26]. □

6.4 Semantics for Open Policy Base

Let B be an open base. We view each open rule in B as standing for all its ground instances. In other word, let $d(\bar{x})$ be a rule whose free variables are $\bar{x}$. $d(\bar{x})$ should actually be understood as representing the set of closed rules

$$\{d(\bar{c}) \mid \bar{c} \text{ is a ground substitution for } \bar{x}\}$$

For example, if $\mathcal{S} = \{\mathsf{A}, \mathsf{B}\}$ and $\mathcal{O} = \{\mathsf{X}, \mathsf{Y}\}$, then $\frac{\mathsf{write}^+(x,\mathsf{X}) \,:\, \mathsf{read}^+(x,y)}{\mathsf{read}^+(x,y)}$ stands for:

$$\{\ \frac{\mathsf{write}^+(\mathsf{A},\mathsf{X}) \,:\, \mathsf{read}^+(\mathsf{A},\mathsf{X})}{\mathsf{read}^+(\mathsf{A},\mathsf{X})},\ \frac{\mathsf{write}^+(\mathsf{A},\mathsf{X}) \,:\, \mathsf{read}^+(\mathsf{A},\mathsf{X})}{\mathsf{read}^+(\mathsf{A},\mathsf{X})},\ \frac{\mathsf{write}^+(\mathsf{B},\mathsf{X}) \,:\, \mathsf{read}^+(\mathsf{B},\mathsf{Y})}{\mathsf{read}^+(\mathsf{B},\mathsf{Y})},\ \frac{\mathsf{write}^+(\mathsf{B},\mathsf{X}) \,:\, \mathsf{read}^+(\mathsf{B},\mathsf{Y})}{\mathsf{read}^+(\mathsf{B},\mathsf{Y})}\ \}$$

Thus each open base B can be associated with an "equivalent" closed base B'. The semantics of B is defined to be the same as that for B'.

6.5 Application Guidelines

Having defined the syntax and semantics of policy bases, we now turn to the practical aspects of specifying policy bases. In particular, we provide some guidelines for representing the three kinds of structural properties discussed in Section 2.

Consider a rule $\frac{f \,:\, f'}{g}$. Its intuitive meaning is that the authorization specified by g is allowed if the authorization specified by f is allowed and no authorization contradicting f' has been specified. Informally, f specifies some prerequisite authorization required for g, while f' specifies assumptions that can be used to deduce the authorization specified by g. In the following, we discuss different forms of our rule and show how they can be used.

Consider a rule of the form $\frac{\mathsf{T} \,:\, \mathsf{T}}{g}$, or simply g. Such a rule expresses basic authorization requirements that must be satisfied in a system. There is no prerequisite nor assumption. For example, to say that a user A must be able to read and write his home directory, we write:

$$\mathsf{read}^+(\mathsf{A}, \mathsf{A.home}) \ \wedge\ \mathsf{write}^+(\mathsf{A}, \mathsf{A.home})$$

where $\mathsf{A.home}$ denotes user A's home directory. These basic authorization requirements form the core upon which other authorizations can be deduced.

A rule of the form $\frac{f \,:\, \mathsf{T}}{g}$, or simply $f \ \Rightarrow\ g$, can be used for two purposes. First, it can be used to express a closure property between authorization requirements. For example, consider the rule:

$$\mathsf{execute}^+(x, \mathsf{P.exe}) \ \Rightarrow\ \mathsf{read}^+(x, \mathsf{P.doc})$$

which says that a user who is authorized to execute a program $\mathsf{P.exe}$ should also be allowed to read its associated documentation $\mathsf{P.doc}$.

Another use of the above rule is to define new authorization requirements in terms of others. For example, in Unix, the right to delete a file is equivalent to the right to write the directory containing the file. This can be made explicit as:

$$\mathsf{write}^+(x, d) \ \wedge\ f \in d \ \Rightarrow\ \mathsf{delete}^+(x, f)$$

where f and d are variables standing for a file and a directory respectively.

Rules can also be used to represent implicit authorizations. There are several reasons why an authorization is left implicit. First, it can be a convention. For example, in general, the number of negative authorizations in a system far exceeds the number of positive ones. Thus for efficiency, a security administrator may specify only the positive ones and leave the negative ones implicit. In other words, the convention is that if a right has not been explicitly authorized, then it is denied. This convention can be formalized in a policy base with the following schema:

$$\frac{:\ r^-(s,o)}{r^-(s,o)}$$

where $r \in R$.

An inheritance property is another example of implicit authorizations that can be formalized as rules. An example is given in Section 7.

6.6 Specifying Exceptions

In the following, we explore several strategies to specify exceptions. We first introduce the concept of *virtual* rights. Virtual rights are not access

rights per se, but are introduced for stating exceptions. We explain this with an example. Suppose we have the following authorization requirements:

> *(1) User* A *is not allowed to write file* X*.*
> *(2) A user who is not allowed to write file* X *is also not allowed to read* X *except for those who belong to group* G *and those who can read file* Y*.*

As a first attempt, we can express this as two rules:

$$\mathsf{write}^-(\mathsf{A},\mathsf{X})$$
$$\mathsf{write}^-(x,\mathsf{X}) \wedge \neg(x \in \mathsf{G}) \wedge \neg\mathsf{read}^+(x,\mathsf{Y}) \Rightarrow \mathsf{read}^-(x,\mathsf{X})$$

Clearly, these rules correctly represent the requirements. However, they are inflexible and error prone in the following sense: They require every exception to be known and be included in the left hand side of the second rule. Thus for a subject whose exception status is unknown (e.g. subject A above), it will not be explicitly denied the right to read X.

A better way to represent this would be to introduce a virtual right except to represent exceptions and a rule to limit exceptions to the ones explicitly specified.

$$B_4 = \left\{ \begin{array}{l} \mathsf{write}^-(\mathsf{A},\mathsf{X}) \\ \mathsf{write}^-(x,\mathsf{X}) \wedge \neg\mathsf{except}^+(x,\mathsf{X}) \Rightarrow \mathsf{read}^-(x,\mathsf{X}) \\ x \in \mathsf{G} \Rightarrow \mathsf{except}^+(x,\mathsf{X}) \\ \mathsf{read}^+(x,\mathsf{Y}) \Rightarrow \mathsf{except}^+(x,\mathsf{X}) \\ \frac{:\neg\mathsf{except}^+(x,\mathsf{X})}{\neg\mathsf{except}^+(x,\mathsf{X})} \end{array} \right\}$$

Subject A is denied read access to X by the second rule in B_4. This is the case because A is assumed to be not an exception by the last rule in B_4. Thus this specification errs on the safe side from a security viewpoint.

Another way of stating the same requirements without using the virtual right except is the following.

$$B_5 = \left\{ \begin{array}{l} \mathsf{write}^-(\mathsf{A},\mathsf{X}) \\ x \in \mathsf{G} \Rightarrow \neg\mathsf{read}^-(x,\mathsf{X}) \\ \mathsf{read}^+(x,\mathsf{Y}) \Rightarrow \neg\mathsf{read}^-(x,\mathsf{X}) \\ \frac{\mathsf{write}^-(x,\mathsf{X}) : \mathsf{read}^-(x,\mathsf{X})}{\mathsf{read}^-(x,\mathsf{X})} \end{array} \right\}$$

The main difference between B_4 and B_5 is that in B_4, we only have sufficient conditions for exceptions while in B_5, we can conclude that $\neg\mathsf{read}^-(x,\mathsf{X})$ holds for all excepted individuals.

Yet another way to specify the requirements, one that can be viewed as a hybrid of B_4 and B_5, is the following.

$$B_6 = \left\{ \begin{array}{l} \mathsf{write}^-(\mathsf{A},\mathsf{X}) \\ x \in \mathsf{G} \Rightarrow \mathsf{except}^+(x,\mathsf{X}) \\ \mathsf{read}^+(x,\mathsf{Y}) \Rightarrow \mathsf{except}^+(x,\mathsf{X}) \\ \frac{\mathsf{write}^-(x,\mathsf{X}) : \neg\mathsf{except}^+(x,\mathsf{X})}{\mathsf{read}^-(x,\mathsf{X})} \end{array} \right\}$$

B_6 is similar to B_4 in that only sufficient conditions for exceptions are specified. However, for any subject, such as A, who is not explicitly specified as an exception, its exception status remains unknown.

Although B_4, B_5 and B_6 are different with respect to what can be concluded about the excepted individuals, they all have the same semantics with respect to write^- and read^-. In particular, $\mathsf{read}^-(\mathsf{A},\mathsf{X})$ holds in each case.

7 Examples of Policy Bases

In this section, we present two examples of using policy bases to specify authorization requirements. The first example is the Bell-LaPadula model (BLP) [4]. We present a straightforward formulation of the basic BLP model in the policy base notation and also an enhancement with *need-to-know* restrictions. The second example shows how to formalize inheritance properties (as illustrated by examples in Section 2).

The essence of the basic BLP model can be summarized by two rules, "no read up and no write down". To simplify our presentation, we consider only two security levels low and high. We specify the BLP model as follows:

$$BLP = R^- \cup W^- \cup R^+ \cup W^+$$

where

(No Read Up)
$R^- = \{s \in \mathsf{low} \wedge o \in \mathsf{high} \Rightarrow \mathsf{read}^-(s,o)\}$

(No Write Down)
$W^- = \{s \in \mathsf{high} \wedge o \in \mathsf{low} \Rightarrow \mathsf{write}^-(s,o)\}$

(Can Read Down)
$R^+ = \{o \in \mathsf{low} \Rightarrow \mathsf{read}^+(s,o)\}$

(Can Write Up)
$W^+ = \{s \in \mathsf{low} \Rightarrow \mathsf{write}^+(s,o)\}$

In the above, denials are absolute in the sense that no exception is allowed. Given a complete description of the group relation, the above policy base uniquely defines a strongly sound and strongly complete authorization policy that satisfies the *simple* and $\star$-security properties [4].

However, this basic model suffers from two drawbacks. First, the group relation must be completely defined in order to give a strongly complete authorization policy. Second, although positive authorizations that are granted do satisfy the simple and $\star$-security property, they violate the *principle of minimal privileges* [25].

We remedy this by adding need-to-know restrictions and denials by default. We modify R' and W' to be:

$$\begin{aligned} R' &= \{o \in \mathsf{low} \wedge \mathsf{need\text{-}to\text{-}know}^+(s,o) \\ &\qquad \Rightarrow \mathsf{read}^+(s,o)\} \\ W' &= \{s \in \mathsf{low} \wedge \mathsf{need\text{-}to\text{-}know}^+(s,o) \\ &\qquad \Rightarrow \mathsf{write}^+(s,o)\} \end{aligned}$$

and BLP to be

$$BLP' = R^- \cup W^- \cup R' \cup W' \cup D$$

where D is

$$\{\frac{: \mathsf{read}^-(s,o)}{\mathsf{read}^-(s,o)},\ \frac{: \mathsf{write}^-(s,o)}{\mathsf{write}^-(s,o)},\ \frac{: \neg\mathsf{need\text{-}to\text{-}know}^+(s,o)}{\neg\mathsf{need\text{-}to\text{-}know}^+(s,o)}\}$$

The virtual right need-to-know formalizes the need-to-know restrictions and can be defined in terms of *compartments* of subjects and objects using other rules.

We now turn to our second example. Consider the following inheritance properties:

> *(1) If a subject s has not been explicitly granted a right r to an object o, then s will inherit a denial of r to o if it belongs to a group that has a denial of r to o.*
> *(2) If a subject s has not been explicitly denied a right r to an object o, then s will inherit a grant of r to o if all groups to which s belongs have grants of r to o.*

These can be expressed respectively by the following schemas:

$$d_1 = \frac{s \in g \wedge r^-(g,o) \ : \neg r^+(s,o) \wedge r^-(s,o)}{r^-(s,o)}$$

and

$$d_2 = \frac{\begin{array}{c}\forall\, s,g,o[\neg(s \in g) \ \vee\ r^+(g,o)] \\ : r^+(s,o) \ \wedge\ \neg r^-(s,o)\end{array}}{r^+(s,o)}$$

where $\forall\ s,g,o[f(s,g,o)]$ in d_2 is a shorthand for the conjunction of all formulas of the form $f(\mathsf{A},\mathsf{G},\mathsf{X})$ where $\mathsf{A},\mathsf{G} \in \mathcal{S}$ and $\mathsf{X} \in \mathcal{O}$.

8 Composition of Policy Bases

There are two notions of composition for policy bases that are important in a distributed system environment. We next examine the situations that give rise to multiple policy bases, and point out the different needs in these situations.

First, a system may be administered by multiple security administrators, each responsible for a distinct part of the system. Each security administrator specifies a policy base for the part of the system he is responsible for. In this case the different policy bases complement each other, in the sense that each fills in a part that has not been specified by others. Thus a composition gives the "sum" of all authorization requirements in the policy bases. We call this type of composition *peer* or *horizontal* composition.

Second, a security administrator may delegate his responsibilities to a number of subordinate administrators. This gives rise to a *root* policy

base corresponding to the delegating administrator and a number of *leaf* policy bases corresponding to the subordinate administrators. The leaf policy bases are more specific and detailed than the root policy base and typically contain refinements of the root policy base. Composition in this case would combine all of the authorizations present in the root policy base together with their refinements in the leaf policy bases. We call this type of composition *hierarchical* or *vertical* composition.

The key difference between horizontal and vertical compositions is in their resolution of conflicts. In horizontal composition, a conflict is resolved in favor of the base with the positive authorization, while in vertical composition, a conflict is resolved in favor of the first (i.e. the more authoritative) base if the conflict involves a negative authorization in the first base and a positive authorization in the second base. The definitions of horizontal and vertical compositions are given below. Let $\tilde{1} = 2$ and $\tilde{2} = 1$.

Definition (Horizontal Composition) Let B_1 and B_2 be two bases. Apply the following to B_i $(i = 1, 2)$: For all distinguished atoms a such that $\overline{a}$ occurs in the consequent of a rule d in B_i and a occurs in the consequent of a rule d' in $B_{\tilde{i}}$, remove all occurrences of $\overline{a}$ from the consequent of d. Let B_i' be the resulting base. Then

$$B_1 \boxed{\mathbf{h}} B_2 = B_1' \cup B_2'$$

□

Definition (Vertical Composition) Let B_1, B_2 be two bases. Apply the following to B_2: For all distinguished atoms a such that $\overline{a}$ occurs in the consequent of a rule d in B_1 and a occurs in the consequent of a rule d' in B_2, remove all occurrences of a from the consequent of d'. Let B_2' be the resulting base. Then

$$B_1 \boxed{\mathbf{v}} B_2 = B_1 \cup B_2'$$

□

Because of space limitation, we relegate a study of the properties of these two types of composition to another paper [26].

9 Implementation

Our model can be implemented as follows in a distributed system. Each policy base is stored and managed by a node in the system. We call such a node a *policy server*. These policy servers are organized in a hierarchical manner. Policy servers at the same level are called *peers*. Clients submit their access requests to appropriate policy servers for authorization decisions. The policy servers communicate with each other in authorizing an access request. Thus both peer and hierarchical composition are implemented by passing messages between policy servers.

The assignment function used in interpreting propositional variables is implemented by a set of *distributed monitors* that keep track of the status of propositional variables in a distributed manner.

The group relation is implemented by a set of *group servers* that collectively maintain group membership information for all subjects and objects in the system. Thus all updates of group memberships (e.g. additions and deletions) in the system are handled by the group servers.

Both the distributed monitors and group servers are regularly queried by the policy servers in making authorization decisions. The evaluation mechanism used in each policy server is based on an interpreter for general logic programs. In fact, a suitably modified Prolog interpreter is sufficient.

10 Concluding Remarks

We have presented a new approach to representing and evaluating authorization. In our approach, a set of authorization requirements is specified declaratively by a policy base. Unlike most existing approaches, the semantics of authorization is defined independently and is separate from implementation mechanisms.

Our approach is readily extensible. First, new predicate symbols can be added to our representation language to increase its expressiveness without a significant increase in computational

requirements. Second, new notions of composition for policy bases can be defined based upon the semantic notion of authorization policy.

For future work, there are some general problems that deserve further investigation. These include: (1) the use of disjunctive information in authorization, and (2) the incorporation of structured subjects (e.g. one subject being a role or delegate of another subject) and structured objects (e.g. one object being an implementation of another object) into our representation language.

Acknowledgements

The authors would like to thank Clifford Neuman and the anonymous referees for their helpful comments on revising the paper. The authors are also grateful to Vladimir Lifschitz for several useful discussions on our paraconsistent semantics for extended logic programs.

References

[1] *IEEE Symposium on Research in Security and Privacy*, Oakland, California, April 18–21 1988.

[2] M. Abadi, M. Burrows, B.W. Lampson, and G. Plotkin. A calculus for access control in distributed systems. Technical Report 70, Systems Research Center, Digital Equipment Corporation, February 1991. An abbreviated version appeared in Advances in Cryptology — CRYPTO '91, pages 1–23, Santa Barbara, California, August 11–15 1991.

[3] K.R. Apt, H.A. Blair, and A. Walker. Towards a theory of declarative knowlege. In J. Minker, editor, *Foundations of Deductive Database and Logic Programming*, pages 89–148. Morgan Kaufmann Publishers, Inc., 1988.

[4] D.E. Bell and L.J. LaPadula. Secure computer system: Unified exposition and Multics interpretation. Technical Report ESD-TR-75-306, Mitre Corporation, March 1976.

[5] D.E. Denning. *Cryptography and Data Security*. Addison-Wesley, 1982.

[6] D.E. Denning, T.F. Lunt, R.R. Schell, M. Heckman, and W.R. Shockley. The SeaView formal security policy model. Technical Report A003, Computer Science Laboratory, SRI International, 1987.

[7] M. Gasser. *Building a Secure Computer System*. Van Nostrand Reinhold Company, New York, 1988.

[8] M. Gelfond and V. Lifschitz. The stable model semantics for logic programming. In R. Kowalski and K. Bowen, editors, *Logic Programming: Proceedings of the Fifth International Conference*, pages 1070–1080. The MIT Press, 1988.

[9] M. Gelfond and V. Lifschitz. Classical negation in logic programs and disjunctive databases. *New Generation Computing*, 9(3,4):365–385, 1991.

[10] J. Glasgow, G. MacEwen, and P. Panangaden. A logic for reasoning about security. In *Proceedings of The Computer Security Foundations Workshop III*, pages 2–13, Franconia, New Hampshire, June 12–14 1990.

[11] J.A. Goguen and J. Meseguer. Security policies and security models. In *Proceedings of 3rd IEEE Symposium on Research in Security and Privacy*, pages 11–20, Oakland, California, April 26–28 1982.

[12] P.P. Griffiths and B.W. Wade. An authorization mechanism for a relational database system. *ACM Transactions on Database Systems*, 1(3):242–255, 1976.

[13] S. Kramer. On incorporating access control lists into the Unix operating system. In *Proceedings of Usenix Unix Security Workshop*, pages 38–48, Portland, Oregon, August 29–30 1988.

[14] B.W. Lampson. Dynamic protection structures. In *Proceedings of AFIPS Fall Joint Computer Conference*, volume 35, pages 27–38, Las Vegas, Nevada, November 18–20 1969.

[15] B.W. Lampson. Protection. In *Proceedings of 5th Princeton Symposium on Information Sciences and Systems*, pages 437–443, Princeton University, March 1971. Reprinted in ACM Operating Systems Review, 8(1):18–24, January 1974.

[16] B.W. Lampson. A note on the confinement problem. *Communications of the ACM*, 16(10):613–615, October 1973.

[17] C.E. Landwehr. Formal models for computer security. *ACM Computing Surveys*, 13(3):247–278, September 1981.

[18] C.E. Landwehr, C.L. Heitmeyer, and J. McLean. A security model for military message systems. *ACM Transactions on Computer Systems*, 2(3):198–222, August 1984.

[19] T.F. Lunt. Access control policies: Some unanswered questions. *Computer & Security*, 8(1):43–54, February 1989.

[20] D. McCullough. Noninterference and the composability of security properties. In *Proceedings of 9th IEEE Symposium on Research in Security and Privacy* [1], pages 177–186.

[21] J. McLean. The algebra of security. In *Proceedings of 9th IEEE Symposium on Research in Security and Privacy* [1], pages 2–7.

[22] J. McLean. The specification and modeling of computer security. *Computer*, 23(1):9–16, January 1990.

[23] T. Przymusinski. On the declarative and procedural semantics of logic programs. *Journal of Automated Reasoning*, 5(2):167–205, June 1989.

[24] R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13(1-2):81–132, April 1980.

[25] J.H. Saltzer and M.D. Schroeder. The protection of information in computer systems. *Proceedings of IEEE*, 63(9):1278–1308, September 1975.

[26] Thomas Y.C. Woo and Simon S. Lam. Authorization in distributed systems: A new approach. Technical report, Department of Computer Sciences, The University of Texas at Austin, 1992. In preparation.