

A Calculus for Access Control in Distributed Systems

M. Abadi* M. Burrows* B. Lampson* G. Plotkin†

Abstract

We study some of the concepts, protocols, and algorithms for access control in distributed systems, from a logical perspective. We account for how a principal may come to believe that another principal is making a request, either on his own or on someone else's behalf. We also provide a logical language for access control lists, and theories for deciding whether requests should be granted.

1 The Problem

At least three ingredients are essential for security in computing systems:

- **A trusted computing base:** the hardware and systems software should be capable of preserving the secrecy and integrity of data.
- **Authentication:** it should be possible to determine who made a statement; for example, a user should be able to request that his files be deleted and to prove that the command is his, and not that of an intruder.
- **Authorization, or access control:** access control consists in deciding whether the agent that makes a statement is trusted on this statement; for example, a user may be trusted (hence obeyed) when he says that his files should be deleted.

These ingredients are fairly well understood in centralized systems.

Distributed systems pose new problems. Scale becomes an issue, as ultimately one would want to envision a global distributed system, with a global name space and global security. In addition, there are difficulties with communication, booting, loading, authentication, and authorization.

The computing base of a distributed system need not reside in a single location, under a single management. This implies, in particular, that secure communication cannot be taken for granted. Since it is hard to provide physically secure communication lines, some form of encryption is typically required. In what follows, we assume that shared-key encryption (e.g., [7]) and public-key encryption (e.g., [8, 19]) are available where needed, but we use them frugally. In addition, there are problems of secure booting and

*Digital Equipment Corporation, Systems Research Center. 130 Lytton Avenue, Palo Alto, California 94301, USA.

†Department of Computer Science, University of Edinburgh, King's Buildings, Edinburgh EH9 3JZ, UK. Part of this work was completed while at Digital Equipment Corporation, Systems Research Center.

secure loading of software. For instance, the operating system that a host runs may be obtained from a repository across the network; a mechanism is needed to guarantee that this software is an authentic release, free of viruses. This mechanism will inevitably rely on authentication and access control, as it is necessary to restrict who can make a release.

The nodes of a distributed system may act on their own behalf, or on behalf of users. Users and nodes trust one another to different extents. Moreover, a user may be trusted only when he is working at certain nodes; he may have more rights at his office than when working from a remote public terminal. Therefore, both users and nodes need to be authenticated, and their identities considered in access control decisions.

These issues give rise to a change in the nature of authentication and access control. The basic questions of authentication and access control are, always, "who is speaking?" and "who is trusted?" Typically the answer is the name of a simple principal (a user or a host). In a distributed environment, these questions can receive a variety of answers; the notion of principal can be extended, to include:

- Users and machines.
- Channels, such as input devices and cryptographic channels. There is no formal reason to distinguish channels from users and machines, and it is advantageous not to. Cryptographic channels are identified by keys, and so we may write that the key K says s when s is asserted in a message decrypted with K .
- Conjunctions of principals, of the form $A \wedge B$. If A and B make the same statement s , then $A \wedge B$ says s as well. It often happens that $A \wedge B$ is trusted on s , but neither A nor B is trusted by himself—a "joint signature" is required.
- Groups. It is often inconvenient to list explicitly all of the principals that are trusted in some respect, both because the list would be long and because it may change too often. Groups provide an indirection mechanism. The use of groups implies the need for a scheme for deciding whether a principal is a member of a group. This is not always straightforward, for example when the registrar for a group is remote; group membership certificates then become necessary.
- Principals in roles, of the form $A \text{ as } R$. The principal A may adopt the role R and act under the name $A \text{ as } R$ when he wants to diminish his powers, in particular as a protection against blunders. For example, a system manager may act in the role *normal-user* most of the time, and enable the *manager* role only on occasion. Similarly, a machine may adopt a weak role before beginning to run a piece of untrusted software, or before delegating authority to an untrusted machine.
- Principals on behalf of principals, of the form $B \text{ for } A$. The principal A may delegate authority to B , and B can then act on his behalf, using the identity $B \text{ for } A$. In the most common case, a user A delegates to a machine B . It is also possible for a machine to delegate to another machine, and for delegations to be cascaded (iterated); then we encounter expressions such as $C \text{ for } (B \text{ for } A)$.

This list raises formal questions of some practical importance. First, there is a need to determine which of the operations on principals should be viewed as primitive, and which can be defined from the others. Then one may ask what laws the operations satisfy, as for example whether $B \text{ for } (A \wedge A')$ and $(B \text{ for } A) \wedge (B \text{ for } A')$ are in some sense

equivalent; if two principals are equivalent then they should be granted the same requests. The resulting theory of principals should provide a reasonable degree of expressiveness, and it should be amenable to a mathematical justification. It is also essential that the theory of principals be simple, because users need to specify permissions on resources and programs need to make access control decisions.

Further, a variety of protocols and algorithms must be designed and analyzed. Mechanisms are required for secure booting and loading, for joining groups and for proving membership, for adopting roles, for delegations of authority and for certifying such delegations, and for deciding whether a request should be granted.

This paper is a study of some of the concepts, protocols, and algorithms for security in distributed systems, with a focus on access control. Our treatment is fairly formal, as it is based on logics. We account for how a principal may come to believe that another principal is making a request, either on his own or on someone else's behalf. We also provide a logical language for access control lists (ACLs), and theories for deciding whether requests should be granted. The logics enable us to explain a variety of protocols which can differ from one another in subtle ways.

On occasion, the formal analysis has suggested ideas useful to implementations, for example that some digital signatures could be saved. Moreover, logics make it possible to describe protocols and policies at a reasonable level of abstraction; we avoid the need for *ad hoc* arguments about particular implementations. This abstraction is important in the context of heterogeneous distributed environments, where several implementations of a design may coexist.

Our study is intended as a formal basis for parts of a security architecture, and for the Digital Distributed Systems Security Architecture (DSSA) in particular [10]; this architecture is currently under implementation. Other formal explanations of security are conceivable. We hope that this work clarifies some of the issues that these alternative accounts may address.

The next section is an overview. We describe the basic logical framework in section 3. Sections 4 and 5 extend the treatment with roles and delegation. Finally, in section 6, we consider delegation schemes and algorithms for making access control decisions.

This paper does not cover many issues in the area; some of these issues are briefly mentioned in passing. Since an implementation of the ideas presented is not complete, we leave most implementation considerations to a companion paper [14] and to future work. In particular, more thought needs to be given to the user interface and to the policies for the writing of ACLs.

2 Overview

Composite principals play a central role in informal reasoning; for example, one often talks about “ A and B ” and “ B on behalf of A .” Therefore, we start by introducing formal notations for composite principals.

Some classical constructors for composite principals such as conjunction and disjunction come to mind first:

- $A \wedge B$: A and B as co-signers. A request from $A \wedge B$ is a request that both A and B make. (It is not implied that A and B make this request in concert.)

- $A \vee B$ is the dual notation; $A \vee B$ represents the group of which A and B are the sole members.

Conjunction is important in our designs. Disjunction is often replaced with implication, in particular in dealing with groups. As discussed further below, “ A is a member of the group G ” can be written $A \Rightarrow G$. The group G can be named, as all other principals; it need not be given a definition beyond that implicit in formulas such as $A \Rightarrow G$. The ability to write out a group in full (as a disjunction) is not essential in practice.

There are also new connectives, in particular:

- A as R : the principal A in role R .
- $B|A$: the principal obtained when B speaks on behalf of A , not necessarily with a proof that A has delegated authority to B ; by definition, $B|A$ says s if B says that A says s ; we pronounce $B|A$ as “ B quoting A .”
- B for A : the principal obtained when B speaks on behalf of A , with appropriate delegation certificates; B for A says s when A has delegated to B and B says that A says s , possibly after some checking that s is reasonable.

Of these, only $|$ is primitive in our approach; *for* and *as* are important, but they are coded in terms of \wedge and $|$. Since *for* is stronger than $|$, we tend to use $|$ only for encoding *for* and *as*.

In order to define the rights of these composite principals, we develop an algebraic calculus. In this calculus, one can express equations such as $(B \wedge C)$ for $A = (B$ for $A) \wedge (C$ for $A)$, and then examine their consequences. Since \wedge is the standard meet in a semilattice, we are dealing with an ordered algebra, and we can use a partial order \Rightarrow among principals: $A \Rightarrow B$ stands for $A = A \wedge B$, and it means that A is at least as powerful as B ; we pronounce this “ A implies B ” or “ A speaks for B .”

A modal logic extends the algebra of principals. In this logic, A says s represents the informal statement that the principal A says s . Here s may function as an imperative (“the file should be deleted”) or not (“ C ’s public key is K ”); imperative modalities are not explicit in the formalism.

The modal logic is a basis for various algorithms and protocols. For example, it is possible to explain logically how two principals A and B establish a public key K_d for B for A .

The logic also underlies a theory of ACLs. We write \supset for the usual logical implication connective, and A controls s as an abbreviation for $(A$ says $s) \supset s$, which expresses trust in A on the truth of s ; in the logic, an ACL for a formula s is a list of assertions of the form A controls s . If s represents a command to a server, then the ACL entry A controls s records the server’s trust in A on s , and hence that A will be obeyed when he says s . When s is clear from context, the ACL for s may simply be presented as the list of principals trusted on s .

If $A \Rightarrow B$ and B controls s then A controls s as well. Therefore, ACLs may not mention all trusted principals explicitly; when B is listed, access should be granted to any A such that $A \Rightarrow B$. It is not always entirely trivial to decide whether a request should be granted.

Encryption is not treated formally, as its use is not essential for security—it is just present in the most reasonable implementations. The logic does not consider the problems

associated with timestamps and lifetimes, either. Therefore, we do not explain how principals come to believe that messages are recent, and not replays. In this, and in the treatment of cryptographic channels as mere principals, the logic is more abstract than the logic of authentication proposed in [4]. The logic of authentication has as its goal describing and debugging fairly low-level authentication protocols; in contrast, the logic discussed here is intended as an aid in developing and applying a general design.

3 The Basic Logic

This section describes the basic logical framework, with syntax, axioms, and semantics. The last subsection presents some axioms that we do not adopt. The laws for *as* and *for* appear in sections 4 and 5.

3.1 A calculus of principals

We study a calculus of principals, in a minimal syntax. As discussed later in the paper, this syntax suffices for expressing roles and the needed delegation connectives.

Principals form a semilattice under the operation of conjunction, and obey the usual semilattice axioms:

- \wedge is associative, commutative, and idempotent.

As usual, $A \Rightarrow B$ can be taken as an abbreviation for $A = A \wedge B$, or $=$ can be defined in terms of \Rightarrow .

Implication is often used to represent group membership. This notion of membership is slightly peculiar, and in particular it is transitive. We have found that transitivity is rather convenient for reasoning about hierarchical groups. Moreover, \Rightarrow has a fairly pleasant interaction with the other connectives of our calculus.

However, our treatment of groups is incomplete. It would be desirable to provide group subtraction and intersection operations. It would also be desirable to provide naming support for usual subgroups—so that there would be a standard operation that one applies to a group name G to obtain the group of principals with a certain position or job in G . Most probably, these operations would appear only at the lowest level of logical formulas, and would not mix with other constructs. Therefore, we do not include these additional operations here; intersection and subtraction are discussed in the full paper [2].

The principals form a semigroup under $|$:

- $|$ is associative.

The final axiom is the multiplicativity of $|$ in both of its arguments, which means:

- $|$ distributes over \wedge .

Multiplicativity implies monotonicity.

In short, the principals are organized as a multiplicative semilattice semigroup. Often, these same axioms are taken for structures with some sort of union or disjunction (instead of \wedge), and then the term additive is preferred to multiplicative.

A common example of a multiplicative semilattice semigroup is an algebra of binary relations over a set, with the operations of union and composition; a multiplicative semilattice semigroup isomorphic to an algebra of binary relations is called representable. In fact, the binary-relation example is common in a mathematical sense: every free multiplicative semilattice semigroup is representable (e.g., [3, 18]). This implies that the equational theory of multiplicative semilattice semigroups coincides with that of binary-relation algebras.

Equational theories do not suffice for reasoning about access rights, as for example group-membership assumptions are needed. Since the binary-relation model is a clear, appealing one, it is then natural to hope that every multiplicative semilattice semigroup is representable. Andr eka has recently proved that this is not the case [3]. She has also studied the distributive multiplicative semilattice semigroups, which are defined by the additional distributivity axiom:

- if $B \wedge C \Rightarrow A$ then there exist B' and C' such that $B \Rightarrow B'$, $C \Rightarrow C'$, and $A = B' \wedge C'$.

Every distributive multiplicative semilattice semigroup is representable. The class of distributive multiplicative semilattice semigroups is a quasi-variety, but neither a variety nor finitely axiomatizable.

The algorithmic applications of the distributivity axiom are hard to see, because of its existential nature. Therefore, in the remainder of the paper, we examine both the axiomatic and the model-theoretic possibilities. We tend to work axiomatically, with the theory of multiplicative semilattice semigroups, but rely on binary relations as a way of constructing models. Moreover, section 6 describes an algorithm for the binary-relation model.

Several extensions to the signature are possible, and natural. For example, a unit 1 turns the semigroup into a monoid; 1 can be viewed as a perfectly honest principal. We describe these extensions in the full paper.

The syntax for reasoning about principals is much the obvious one. We manipulate equations between principal expressions, built from atoms A_0, A_1, A_2, \dots with the symbols \wedge and $|$. We combine equations into boolean formulas, and then adopt the usual axioms from propositional logic, and the axioms for the algebraic structures of choice (multiplicative semilattice semigroups, when not stated otherwise).

Section 6 discusses decidability issues for the calculus of principals.

3.2 A logic of principals and their statements

Here we develop a modal logic based on the calculus of principals.

Syntax

The formulas are defined inductively, as follows:

- a countable supply of primitive propositions p_0, p_1, p_2, \dots are formulas;
- if s and s' are formulas then so are $\neg s$ and $s \wedge s'$;
- if A and B are principal expressions then $A \Rightarrow B$ is a formula;

- if A is a principal expression and s is a formula then A says s is a formula.

We use the usual abbreviations for boolean connectives, such as \supset , and we also treat equality between principals ($=$) as an abbreviation. In addition, A controls s stands for $(A \text{ says } s) \supset s$.

Axioms

The basic axioms are those for normal modal logics [12]:

- if s is an instance of a propositional-logic tautology then $\vdash s$;
- if $\vdash s$ and $\vdash (s \supset s')$ then $\vdash s'$;
- $\vdash A \text{ says } (s \supset s') \supset (A \text{ says } s \supset A \text{ says } s')$;
- if $\vdash s$ then $\vdash A \text{ says } s$, for every A .

The calculus of principals is included:

- if s is a valid formula of the calculus of principals then $\vdash s$.

Other axioms connect the calculus of principals to the modal logic:

- $\vdash (A \wedge B) \text{ says } s \equiv (A \text{ says } s) \wedge (B \text{ says } s)$;
- $\vdash (B|A) \text{ says } s \equiv B \text{ says } A \text{ says } s$;
- $\vdash (A \Rightarrow B) \supset ((A \text{ says } s) \supset (B \text{ says } s))$.

The last axiom is equivalent to $(A = B) \supset ((A \text{ says } s) \equiv (B \text{ says } s))$, a substitutivity property.

Let us write $A \rightarrow B$ if, for every s , if A says s then B says s . Although \rightarrow is weaker than \Rightarrow , it suffices in many of the situations where we currently use \Rightarrow , and hence we have contemplated adding it to the formalism. It may also be convenient to have second-order quantification over formulas and over principals. Second-order quantification enables us to define $A \rightarrow B$ as $\forall x.((A \text{ says } x) \supset (B \text{ says } x))$. We do not adopt any of these additional constructs, for the sake of minimality. However, there is no difficulty in giving a semantics for them, and the proof rules needed in examples would be simple.

Another important relation between principals is defined by the formula $(A \text{ says false}) \supset (B \text{ says false})$, which we abbreviate $A \mapsto B$. Intuitively, $A \mapsto B$ means that there is something that A can do (say *false*) that yields an arbitrarily strong statement by B (in fact, *false*). Thus, $A \mapsto B$ means that A is at least as powerful as B in practice. Clearly, $A \Rightarrow B$ implies $A \mapsto B$. We do not have the converse, and actually the converse does not seem attractive, as for example $B \mapsto (B|A)$ is valid but we would not want $B \Rightarrow (B|A)$. Nevertheless, the relation \mapsto serves as a point of reference in axiomatizing the algebra of principals; a careful study of \mapsto may be of some interest.

3.3 Semantics

The simplest semantics is a Kripke semantics, based on accessibility relations [12]. A structure \mathcal{M} is a tuple $\langle W, w_0, I, J \rangle$, where:

- W is a set (as usual, a set of possible worlds);
- w_0 is a distinguished element of W ;
- I is an interpretation function which maps each proposition symbol to a subset of W (the set of worlds where the proposition symbol is true);
- J is an interpretation function which maps each principal symbol to a binary relation over W (the accessibility relation for the principal symbol).

The meaning function \mathcal{R} extends J , mapping a principal expression to a relation:

$$\begin{aligned}\mathcal{R}(A_i) &= J(A_i) \\ \mathcal{R}(A \wedge B) &= \mathcal{R}(A) \cap \mathcal{R}(B) \\ \mathcal{R}(B|A) &= \mathcal{R}(A) \circ \mathcal{R}(B)\end{aligned}$$

There is no difficulty in giving a semantics to other operations on principals such as infinite conjunctions, disjunctions, 0, and 1.

Then the meaning function \mathcal{E} maps each formula to its extension, that is, to the set of worlds where it is true:

$$\begin{aligned}\mathcal{E}(p_i) &= I(p_i) \\ \mathcal{E}(\neg s) &= W - \mathcal{E}(s) \\ \mathcal{E}(s \wedge s') &= \mathcal{E}(s) \cap \mathcal{E}(s') \\ \mathcal{E}(A \text{ says } s) &= \{w \mid \mathcal{R}(A)(w) \subseteq \mathcal{E}(s)\} \\ \mathcal{E}(A \Rightarrow B) &= W \text{ if } \mathcal{R}(B) \subseteq \mathcal{R}(A) \text{ and } \emptyset \text{ otherwise}\end{aligned}$$

where $\mathcal{R}(C)(w) = \{w' \mid w\mathcal{R}(C)w'\}$.

A formula s holds in \mathcal{M} at a world w if $w \in \mathcal{E}(s)$, and it holds in \mathcal{M} if it holds at w_0 . In the latter case, we write $\mathcal{M} \models s$, and say that \mathcal{M} satisfies s . Moreover, s is valid if it holds in all models; we write this $\models s$. The axioms are sound, in the sense that if $\vdash s$ then $\models s$. Although useful for our application, the axioms are not complete. For example, the formula

$$(C \text{ says } (A \Rightarrow B)) \equiv ((A \Rightarrow B) \vee (C \text{ says } \text{false}))$$

is valid but not provable. A more interesting source of incompleteness is that the algebras of principals that underly the semantics are obviously representable, while some multiplicative semilattice semigroups are not. (We do obtain a completeness result for a sublanguage, in section 6.)

This simple relational model provides some justification for the axioms for principals. If $\mathcal{R}(A) = \mathcal{R}(B)$, then it is natural to have $A = B$, and this holds in the semantics. For example, $\mathcal{R}(C|(B|A)) = \mathcal{R}((C|B)|A)$ provides a justification for the associativity of $|$.

A more abstract model may be desirable in order to avoid the requirement of representability. Dynamic algebras suggest such a model (see Pratt's [18] for a recent discussion of dynamic algebras). The operator *says* is directly analogous to the usual partial correctness operator $[]$ of dynamic logic. It connects the two sorts of the algebras, that of principals and that of propositions.

3.4 On idempotence

Idempotence postulates are attractive. For example, when A is a user, there seems to be little advantage in distinguishing $A|A$ and A ; in fact, this distinction seems quite artificial. Therefore, we may postulate that $A|A = A$, for all A , and hence that A *says* A *says* s and A *says* s are equivalent. The idempotence of $|$ implies the idempotence of *for* (see section 5 for details). Here we discuss the idempotence of $|$ and *for*; most of the remarks below are written in terms of $|$, but exactly the same arguments apply to *for*.

The idempotence axiom is rather convenient for handling chains of principals. For instance, suppose that G represents a collection of nodes, that B and C represent members of G , and that an ACL includes $G|A$. Thanks to idempotence, $C|B|A$ obtains access. This means that multiple "hops" within a collection of nodes does not reduce rights and should not reduce security. In particular, there is no need to postulate that $G|G \Rightarrow G$, or to make sure that $G|G|A$ appears in the ACL explicitly.

In addition, idempotence yields $(A \wedge B) \Rightarrow (B|A)$, since $(A \wedge B) = (A \wedge B)|(A \wedge B)$ and $|$ is monotonic, and similarly $(A \wedge B) \Rightarrow (B \text{ for } A)$. These implications are rather pleasing, but not necessarily intuitive. They complicate the problem of making access control decisions. When $(A \wedge B)$ makes a request, one must check whether either $(B|A)$ or $(A|B)$ is trusted, and grant access if either is trusted.

Finally, the semantics of subsection 3.3 does not validate idempotence. We have been unable to find a sensible condition on binary relations that would force idempotence and would be preserved by both union and composition. A related symptom is that idempotence seems less compelling for complex principals than for simple ones.

For all these reasons, we prefer to do without idempotence; in compensation, we often rely on assumptions of the form $G|G \Rightarrow G$, or write less elegant but more explicit ACLs.

We give a similar treatment to other possible axioms related to idempotence. Although in many cases it is reasonable to assume formulas of the forms A *controls* $(B \Rightarrow A)$ and A *controls* $(t \supset A \text{ says } s)$, this does not seem well founded in general. In particular, when A is a group, if A *controls* $(B \Rightarrow A)$ for all B then any member of A can add members to A . Therefore, we choose to adopt neither A *controls* $(B \Rightarrow A)$ nor A *controls* $(t \supset A \text{ says } s)$ as general axioms.

4 Roles

The formal systems we have presented so far include no exotic connective, such as *for*. Neither do they include a primitive for roles. We discuss roles in this section and delegation in the next one. We start with an informal argument on the nature of roles and then describe their logic.

4.1 Roles versus groups

Roles are often related to groups. For example, a member of a group G can act in the role of member of G . It is tempting to establish a formal identification between the group G and a corresponding role G_role . We find that this identification is harmful.

In particular, roles make it possible to reduce privileges as a protection against accidents. The roles used for this purpose need not correspond to groups. For example, there may be a role such that adopting it makes it possible to access only a certain directory in the file system; the role corresponds naturally to a set of resources (files) rather than to any group of principals.

In what follows, then, we do not identify roles and groups. In fact, this confusion would interact somewhat strangely with our encoding of roles, given below. We do allow roles related to groups (such as G_role), but this relation is not formal; a formal operator for relating groups and roles may be a useful extension of our calculus.

It is reasonable to expect that not all roles make sense for all principals, and that a principal can act only in certain roles. For example, a principal may be allowed to act as G_role only if he is a member of G . The simplest implementation of this idea relies on the judicious writing of ACLs. Roles can be adopted freely, and any A can speak in the role G_role , with the identity A as G_role . However, it may be that no ACL in the system grants access to A as G_role , and for example the ACL G as G_role does not. It is this implementation that we choose.

4.2 The encoding

Given the view that roles can be freely adopted, it is quite satisfactory to define A as R to equal $A|R$.

In many cases, this encoding makes some intuitive sense. For example, let A be a machine and R a (possibly untrusted) piece of software that A is running. The requests that A makes while running this software should not emanate from A with full powers; rather, they should emanate from A in some restricted role. The role can be named after the piece of software, or after its class (e.g., *untrusted-software*). In any case, when A makes a request s in this role, A says that R says s . This is precisely $A|R$ says s . Similarly, if A is a user and R is one of his positions, we can think of R as a program that A is following, and apply the analysis above.

Furthermore, this presentation of roles offers formal advantages, thanks to the monotonicity, the multiplicativity, and the associativity of $|$.

- Since $|$ is monotonic in both arguments, if R is a more trusted role than R' and A a more trusted principal than A' , then A as R is more trusted than A' as R' , that is, if $R \Rightarrow R'$ and $A \Rightarrow A'$ then $(A$ as $R) \Rightarrow (A'$ as $R')$. Thus, it is possible to exploit a hierarchy of roles.
- The multiplicativity of $|$ yields that $A \wedge B$ in a role R is identical to the conjunction A and B both in role R :

$$(A \wedge B) \text{ as } R = (A \text{ as } R) \wedge (B \text{ as } R)$$

Similarly, we can exploit the conjunction operation for roles:

$$A \text{ as } (R \wedge R') = (A \text{ as } R) \wedge (A \text{ as } R')$$

Both of these properties seem reasonable for roles as we conceive them informally, and they are quite handy.

- Associativity yields a satisfactory interaction between roles and delegation. In particular, the user A in role R on a machine B is identical to the user A on a machine B , with the compound in role R :

$$B|(A \text{ as } R) = (B|A) \text{ as } R$$

and then the encoding of *for*, given in the next section, yields:

$$B \text{ for } (A \text{ as } R) = (B \text{ for } A) \text{ as } R$$

Some special properties for roles can be expected:

- We sometimes rely on the postulates that all roles are idempotent ($R|R = R$) and commute with one another ($R'|R = R|R'$). This yields:

$$A \text{ as } R \text{ as } R = A \text{ as } R$$

and

$$A \text{ as } R \text{ as } R' = A \text{ as } R' \text{ as } R$$

(We make it explicit when these properties are assumed.)

- We assume $A \Rightarrow (A \text{ as } R)$ for all A . With a unit, this could just be written $1 \Rightarrow R$. In the binary-relation model, $1 \Rightarrow R$ means that roles are subsets of the identity relation; note that the previous idempotence and commutativity properties follow from this.

5 Delegation

Delegation is a fairly ill-defined term. Here we explain our approach to delegation; some example schemes for delegation illustrate the use of the logic. Then we give an encoding of delegation.

5.1 The forms of delegation

Delegation is a basic primitive for secure distributed computing. It is the ability of a principal A to give to another principal B the authority to act on A 's behalf. When B requests a service from a third principal, B might present credentials which are supposed to demonstrate that B is making the request and that A has delegated to B .

Naturally, access control decisions need to take delegations into account. A range of designs is possible.

In the simplest case, A delegates all of his rights to B . Upon B 's request, a service will be granted if it would have been granted to A , had A requested it directly.

An obvious improvement is for A to delegate only some of his rights, such as the right to read certain files from a file server. Capability systems (e.g., [9, 20, 15]) embody this

approach in the case where delegation certificates can be transferred; the containment of capabilities is a serious concern. If delegation certificates are not transferrable, cascaded delegations are hampered, as A may have to prepare adequate certificates for any expected collaborators to B in advance. In all cases, the provider of the service never checks that B is a reasonable delegate for A .

These considerations suggest more sophisticated designs, where access control decisions may depend on the identities of both A and B . For example, the file server may not let a known malicious workstation read a user's files, even if the user was unlucky enough to log onto the workstation. The file server may also let a trusted node read files over which the user has execute-only rights, in the belief that these files will not be shown to the user, but just executed; here the node acts as a protected subsystem, and offers guarantees that its clients cannot provide by themselves. These examples illustrate a desirable flexibility. It exists in some form in a few designs and systems [21, 13], and it is a prominent component of DSSA [11]. We believe that it should and will become widespread.

Different mechanisms embody the concept of delegation in different settings. A user might rely on a smart-card for delegating to a workstation, while delegation across the address spaces of a machine might benefit from operating system support. (The use of smart-cards is discussed in [1], with a simplistic view of delegation.) Similarly, credentials can be signed with either shared-key cryptography or with public-key cryptography, or they might not be signed at all when a secure physical channel is available. Recognizing and treating all these variants as mere instances of delegation makes possible a uniform view of access control throughout an entire distributed system.

Delegation relies on authentication, which is often achieved through authentication protocols (see for example [17, 5, 16, 13, 4]). In fact, some of the messages for delegation can be combined with those for authentication. However, in our study, it is often convenient to view delegation as independent.

Delegation without certificates

The framework outlined so far enables us to describe and compare various schemes for delegation. In what follows, A delegates to B , who makes requests to C . For instance, A may be a user with a sufficiently powerful smart-card, B a workstation, and C a file server. For simplicity, the authority delegated from A to B is not limited to any particular class of requests, and A does not adopt a special role before the delegation.

We assume that synchronized clocks are available, and that appropriate timestamps, lifetimes, sequence numbers, and message type fields are included and checked for all messages. We also assume that all principals can perform digital signatures, for example with the RSA algorithm [19]. We denote by K_A and K_B the public keys for A and B , and by K_A^{-1} and K_B^{-1} the matching secret keys. The formula K says X represents the certificate X encrypted under K^{-1} , which is commonly written $\{X\}_{K^{-1}}$.

A certification authority S provides certificates for the principals' public keys as needed. The necessary certificates are K_S says $K_A \Rightarrow A$ and K_S says $K_B \Rightarrow B$, where K_S is S 's public key.

We consider three instances of delegation. In each case, we are led to ask whether composite principals, such as $B|A$, appear on C 's ACL. The simplest instance of delegation is delegation without certificates:

1. When B wishes to make a request r on A 's behalf, B sends the signed request along with A 's name, for example in the format $K_B \text{ says } A \text{ says } r$.
2. When C receives the request r , he has evidence that B has said that A requests r , but not that A has delegated to B ; then C consults the ACL for request r , and determines whether the request should be granted under these circumstances.

The logic serves in describing the reasoning of the service provider, C , who makes the access control decision. Some assumptions are needed. First, C must believe that K_S is S 's public key:

$$K_S \Rightarrow S$$

and C obtains a certificate encrypted under the inverse of K_S :

$$K_S \text{ says } (K_B \Rightarrow B)$$

These two formulas yield:

$$S \text{ says } (K_B \Rightarrow B)$$

We should assume that C trusts S for such statements:

$$S \text{ controls } (K_B \Rightarrow B)$$

and we immediately obtain that C has B 's key:

$$K_B \Rightarrow B$$

In addition, C sees a message under the inverse of K_B , requesting r on behalf of A ; in other words,

$$K_B \text{ says } A \text{ says } r$$

Immediately, C obtains:

$$B \text{ says } A \text{ says } r$$

that is,

$$(B|A) \text{ says } r$$

Now C consults an ACL for r . If $B|A$ appears in this ACL, that is, if $B|A$ is trusted on r , then C believes r —access is granted.

Delegation with certificates

The protocol just described can hardly be considered satisfactory in general. It is preferable for A to issue a delegation certificate that proves the delegation to B :

1. After mutual authentication, A issues a certificate to B , under A 's key. This certificate states that A has delegated some authority to B .
2. When B wishes to request r on A 's behalf, no further interaction with A is needed: B can present A 's certificate to C , along with $K_B \text{ says } A \text{ says } r$.
3. At this point C has evidence that B has requested r on behalf of A ; now C consults the ACL for r , and determines whether the request should be granted.

For the time being, let us use the notation B serves A to mean that B is a delegate for A . (But *serves* need not be primitive; see subsection 5.2.) Then the delegation certificate from A for B can be expressed:

$$K_A \text{ says } (B \text{ serves } A)$$

and the usual checking of public-key certificates from S yields:

$$A \text{ says } (B \text{ serves } A)$$

If C trusts A on this statement, C gets:

$$((B|A) \text{ says } r) \wedge (B \text{ serves } A)$$

In our theories, this implies:

$$(B \text{ for } A) \text{ says } r$$

Again C can consult the ACL for r , and r will be granted if the list includes B for A . In particular, r will be granted if the list mentions $B|A$, a weaker principal, but may be granted even otherwise.

This scheme is illustrated in more detail by an example in section 6. An important variant consists in omitting the authentication between B and C , leaving the responsibility of authenticating B solely to A ; we study it in the full paper.

5.2 The encoding

The delegation of all powers can be defined convincingly: A delegates to B when A says that B speaks for A (that is, when $A \text{ says } (B \Rightarrow A)$ holds). The concept of delegation that interests us here is subtler. When A delegates to B , the powers of B on behalf of A depend on the contents of the ACLs in the distributed system; A may be more reckless with delegations if these powers are small. In turn, the contents of the ACLs depend on how their writers understand the meaning of delegation. There is a circularity here. We conclude that the meaning of delegation is purely conventional—it is a convention between the writers of ACLs, the delegators, and the delegates.

These remarks suggest an approach where the notion of delegation is primitive; we now sketch this approach. The operator *serves* is taken as primitive, and axiomatized. Some straightforward axioms come to mind, and in particular that $B \text{ serves } A$ is anti-monotonic in B , and that if $B_i \text{ serves } A$ for all i , then $\bigvee_i B_i \text{ serves } A$. Let $E(A)$ be the weakest principal to which A has delegated, and let $E_B(A)$ be the weakest principal stronger than B to which A has delegated. (If disjunction is available, such principals always exist.) Then $B \text{ for } A$ can be defined to be $E_B(A)|A$, and this equals $(B \wedge E(A))|A$. Axioms for delegation yield properties of *for*; for example, the axioms given yield that *for* is multiplicative in its second argument. We have not found a satisfactory way to obtain associativity.

We prefer not having delegation as a primitive; the following thought experiment inspires an alternative approach. Imagine that there is a distinguished principal D in the distributed system who operates as a delegation server. The function of D is to certify delegations. When A wants to delegate to B , it suffices for A to tell D that if B says

that A says s , then D should back B . Thus, if $B|A$ says s then $D|A$ says s . Intuitively, D says that A says s when a delegate of A 's makes this assertion. Thus, we can take $(B \wedge D)|A$ as a formal encoding for B for A . Notice the striking similarity between this encoding of *for* and the previous formulation.

It is not actually necessary for D to be implemented, just as typical roles do not correspond to real users and machines. When A wishes to delegate to B , A says that for all s , if B says that A says s then D says that A says s ; formally, it suffices to have A says $(B|A \Rightarrow D|A)$. The statement A controls $(B|A \Rightarrow D|A)$ then represents that A can delegate to B . These two assertions together immediately yield $(B|A \Rightarrow D|A)$, and when $(B|A)$ says s , we obtain $(D|A)$ says s , and then $((B \wedge D)|A)$ says s . In this formal derivation, it is irrelevant whether D is a real server or not. (It is perhaps even better if D is not a real server, for then it cannot be attacked.)

Thus, we can take B for A to mean $(B \wedge D)|A$, where D is a distinguished fictional principal. Similarly, $(B|A \Rightarrow D|A)$ can represent B serves A ; hence, if B serves A then $(B|A \Rightarrow B$ for $A)$. This expresses the intuition that *for* is | “plus a delegation.” The logical framework also allows the possibility of weaker delegation statements, where A delegates only some of this rights; we prefer the use of roles for limiting the power of delegations.

Our encoding has sensible formal consequences:

- *for* is monotonic in both arguments.
- *for* is multiplicative in both arguments, in fact. This follows from the multiplicativity of | and the definition of *for*:

$$\begin{aligned}
 (B \wedge B') \text{ for } (A \wedge A') &= (B \wedge B' \wedge D)|((A \wedge A')) \\
 &= ((B \wedge D) \wedge (B' \wedge D))|(A \wedge A') \\
 &= ((B \wedge D)|A) \wedge ((B' \wedge D)|A') \wedge ((B' \wedge D)|A) \wedge ((B' \wedge D)|A') \\
 &= (B \text{ for } A) \wedge (B' \text{ for } A) \wedge (B' \text{ for } A) \wedge (B' \text{ for } A')
 \end{aligned}$$

- B for A is always defined, even if A has not delegated to B . In fact, we have:

$$(B|A) \wedge (C \text{ for } A) \Rightarrow ((B \wedge C) \text{ for } A)$$

and hence also

$$(B|A) \wedge (C \text{ for } A) \Rightarrow (B \text{ for } A)$$

This somewhat surprising theorem is the consequence of two desirable, basic properties: the monotonicity of *for*, and the antimonicity of delegation (which means that if A delegates to C , then it also delegates to the stronger principal $B \wedge C$).

Additional properties of D yield further consequences:

- If $D|D \Rightarrow D$ then *for* possesses a weak associativity property:

$$C \text{ for } (B \text{ for } A) \Rightarrow (C \text{ for } B) \text{ for } A$$

which follows from the associativity and the multiplicativity of |:

$$C \text{ for } (B \text{ for } A) = (C \wedge D)|((B \wedge D)|A)$$

$$\begin{aligned}
&= (C|B|A) \wedge (C|D|A) \wedge (D|B|A) \wedge (D|D|A) \\
&\Rightarrow (C|B|A) \wedge (C|D|A) \wedge (D|B|A) \wedge (D|A) \\
&\Rightarrow (C|B|A) \wedge (D|B|A) \wedge (D|A) \\
&= ((C \text{ for } B)|A) \wedge (D|A) \\
&= (C \text{ for } B) \text{ for } A
\end{aligned}$$

The other direction of the implication is not valid. The essential reason for this is that $C \text{ for } (B \text{ for } A)$ implies $C|D|A$, while $(C \text{ for } B) \text{ for } A$ does not. Intuitively, this is because C 's part in $(C \text{ for } B) \text{ for } A$ need not involve checking evidence that B is a delegate for A , or even that A exists.

- If $A \Rightarrow D|A$ then $(A \wedge (B|A)) \Rightarrow (B \text{ for } A)$, because $|$ is multiplicative:

$$\begin{aligned}
A \wedge (B|A) &\Rightarrow (D|A) \wedge (B|A) \\
&= ((B \wedge D)|A) \\
&= B \text{ for } A
\end{aligned}$$

This property means that when A makes a statement himself, there is no need to find a corresponding statement by the delegation server.

- If $A \Rightarrow D|A$ and $A = A|A$ then $A = A \text{ for } A$, that is, the idempotence of $|$ implies the idempotence of for :

$$\begin{aligned}
A \text{ for } A &= ((A \wedge D)|A) \\
&= (A|A) \wedge (D|A) \\
&= A \wedge (D|A) \\
&= A
\end{aligned}$$

The additional properties for D are reasonable, and we adopt them. These properties are reminiscent of those for roles. In the binary-relation model, for example, these properties of D and those for roles all amount to saying the same thing, that the associated binary relations are subsets of the identity.

6 Protocols and Algorithms

In this section we consider some mechanisms for delegation and for access control decisions. The last subsection presents an example.

6.1 Delegation

While delegation has a single semantics, substantially different implementations are recommended for users and for nodes.

6.1.1 Delegation from users to nodes

The schemes for delegation discussed in section 5 do not seem adequate for users. Delegation from users to nodes poses special problems. One difficulty is that it is inconvenient for users to refresh delegation certificates. Refreshing a delegation certificate may require reintroducing a smart-card, for example. Therefore, it seems desirable that delegation certificates be relatively long-lived (valid for many hours).

Long-lived delegation certificates pose a security threat, because a delegation certificate may be valid long after the user has stopped using the node. An attacker may subvert the node over a relatively long time period, and abuse the user's delegation. To prevent this from happening, it is best to introduce an auxiliary delegation key that the node can forget deliberately when the user leaves. The delegation certificate remains valid, but becomes useless.

In more detail, the protocol goes as follows. First the user A delegates to the node B and to a public key K_d provided by the node:

$$A \text{ says } ((K_d \wedge B) \text{ serves } A)$$

The node has the inverse of the public key, and it can set up a channel Ch for $(K_d \wedge B)$ for A :

$$K_d \text{ says } A \text{ says } (Ch \Rightarrow (K_d \wedge B) \text{ for } A)$$

and

$$B \text{ says } A \text{ says } (Ch \Rightarrow (K_d \wedge B) \text{ for } A)$$

Hence

$$((K_d \wedge B) \text{ for } A) \text{ says } (Ch \Rightarrow (K_d \wedge B) \text{ for } A)$$

follows logically. When this statement is believed, it yields:

$$Ch \Rightarrow (K_d \wedge B) \text{ for } A$$

and then monotonicity leads to the desired result:

$$Ch \Rightarrow B \text{ for } A$$

Hereafter B can make requests for A through the channel Ch . In practice, the channel may be obtained by multiplexing a longer-term secure channel from B ; this longer-term channel may well be a DES channel.

The delegation from K_d to the channel has a relatively short lifetime, and needs to be renewed frequently. When the node forgets K_d deliberately, it loses the ability to refresh the certificate for the channel Ch .

6.1.2 Delegation from nodes to nodes

For nodes, the schemes of section 5 are convenient enough. The one corresponding to the *for* operator is simple, reasonably secure, and it can be made efficient enough for implementation.

The same ideas work for cascaded delegations. Suppose that user A has delegated to node B , and B can operate with the identity B for A . If a further delegation is needed, to a node C , the precise delegation statement is:

$$(B|A) \text{ says } (C \text{ serves } (B \text{ for } A))$$

Since A has delegated to B , it follows logically that:

$$(B \text{ for } A) \text{ says } (C \text{ serves } (B \text{ for } A))$$

This statement is believed, and it yields:

$$C \text{ serves } (B \text{ for } A)$$

Now C can make a request r on behalf of B for A :

$$C \text{ says } (B \text{ for } A) \text{ says } r$$

and then the delegation from A yields:

$$(C \text{ for } (B \text{ for } A)) \text{ says } r$$

6.2 Access control decisions

Unfortunately, the set of valid formulas of the calculus of principals is not recursive for any useful definition of validity, but it is recursively enumerable. Undecidability can be proved by a reduction from the word problem for Thue systems, for example. On the other hand, the formulas that arise in access control are not arbitrary; the next two parts discuss decidable access control problems.

6.2.1 A general access control problem

The problem of making access control decisions is computationally complex. It is important therefore to understand the precise form of its instances. The parts of an instance are:

- An expression P in the calculus of principals represents the principal that is making the request. In particular, all appropriate delegations are taken into account in constructing this expression. The various relevant certificates are presented for checking.
- A statement s represents what is being requested or asserted; the statement is presented explicitly by the requestor, and the service provider does not need to derive it logically from other statements. The precise nature of s is ignored—it is treated as an uninterpreted proposition symbol.
- Assumptions state implications among principals; these typically represent assumptions about group memberships. They have the form $P_i \Rightarrow G_i$, where P_i is an arbitrary expression in the calculus of principals and G_i an atom. Note that this syntax is liberal enough to write $G|G \Rightarrow G$ for every appropriate G of interest, obtaining some of the benefit of the idempotence axiom.
- Certain atomic symbols R_0, \dots, R_i, \dots are known to denote roles. This may be obvious from their names.
- An ACL is a list of expressions E_0, \dots, E_i, \dots in the calculus of principals; these represent the principals that are trusted on s .

The problem of access control is deciding whether $\bigwedge_i (P_i \Rightarrow G_i)$ and $\bigwedge_i (E_i \text{ controls } s)$ imply $P \text{ controls } s$, given the special properties of roles and of the delegation server D .

As discussed in the full paper, this problem is intractable. (We have sketched a reduction from the acceptance problem for alternating pushdown automata [6].) Further restrictions are wanted.

6.2.2 A more tractable problem

A second version of the access control problem is based on some further simplifications:

- The *for* operator takes a more prominent role. Since it is expected to be common, one would want an algorithm that does not treat it by expanding it in terms of \wedge and $|$, with an exponential blow-up.
- In effect, *for* is treated as an associative operator. More precisely, we assume that all ACLs have the weakest parenthesization possible. This enables us to ignore the parenthesization of the requestor.
- The $|$ operator occurs only in the encoding of *for* and *as*, and D occurs only in the encoding of *for*.
- Roles are differentiated from other principals (called proper principals), and occur in special positions. The set of atomic symbols for proper principals is \mathcal{P} ; the set of atomic symbols for roles is \mathcal{Q} .
- Assumptions (group memberships) are restricted to be of the form $atom \Rightarrow atom$, where both atoms are either in \mathcal{P} or in \mathcal{Q} . In particular, this forbids the assumptions of the form $G|G \Rightarrow G$, which would give us some benefits of idempotence where appropriate.
- It is therefore necessary to include a construct for writing, in an ACL, that any positive number of iterations of a principal are allowed. We introduce a metalogical construct $()^+$ for this purpose; for example $F \text{ for } G^+$ is a shorthand for the list $F \text{ for } G, (F \text{ for } G) \text{ for } G, \dots$

An ACL is a list of ACL entries. In turn, an ACL entry is a conjunction of ACL entries free of \wedge , called *for-lists*. A *for-list* is a list connected by *for*'s, that is, it has the form $P_0 \text{ for } \dots \text{ for } P_m$, where each P_i is a principal in roles. A principal in roles is of the form $Q \text{ as } R_1 \text{ as } \dots R_n$, where Q is a proper principal and each R_j is a role.

This definition is summarized by a grammar:

<i>ACL</i>	=	<i>list of Entry</i>
<i>Requestor</i>	=	<i>Entry</i>
<i>Entry</i>	=	<i>conjunction of for-list</i>
<i>for-list</i>	=	<i>Principal-in-Roles for-list for Principal-in-Roles</i>
<i>Principal-in-Roles</i>	=	<i>Proper-Principal Principal-in-Roles as Role</i>
<i>Membership</i>	=	<i>Proper-Principal \Rightarrow Proper-Principal Role \Rightarrow Role</i>

The syntactic conditions can be loosened, as this normal form for ACL entries can be obtained by two syntactic transformations justified by the laws of the logic. Conjunctions can be pushed outwards, since both *as* and *for* distribute over \wedge ; this normalization is the only source of exponential behavior, and it can be expected to be unimportant in the common examples where conjunctions are rare. Roles can be pushed inwards, into the delegator argument of *for*'s; this step is efficient.

With these syntactic restrictions, a more efficient access control algorithm is possible:

- The request is granted if the requestor implies one of the ACL entries.
- Each ACL entry is a conjunction of *for*-lists, and so is the requestor. For the requestor to imply an ACL entry, it must be that for each conjunct of the ACL entry there exists some conjunct of the requestor that implies it.
- A *for*-list implies another if they are of the same length, and each principal in roles of the requestor implies the corresponding one of the entry.
- A principal in roles Q *as* R_1 *as* ... *as* R_n implies another Q' *as* R'_1 *as* ... *as* R'_n , if Q implies Q' and for each R_j there exists R'_k such that R_j implies R'_k .
- An atomic symbol P implies another atomic symbol P' if there is a chain of assumptions $P = P_0 \Rightarrow \dots \Rightarrow P_n = P'$.

It is simple to describe this algorithm to users, and it is also simple to implement it. The adequate speed of a prototype implementation suggests that the algorithm might well be practical. The metalogical construct $()^+$ does not introduce any major algorithmic difficulty. Its treatment is well understood, as it is standard in the context of regular languages.

The algorithm is sound for the binary-relation model, where we take all roles to be subsets of the identity. The algorithm is also complete for the binary-relation model. Note in particular that the algorithm treats all roles as idempotent, and they all commute with one another.

Theorem 1 (Soundness and Completeness) *Let E_0, \dots, E_i, \dots be an ACL, P a requestor, and $\bigwedge_i (B_i \Rightarrow B'_i)$ a conjunction of assumptions. The algorithm grants the request if and only if*

$$\bigwedge_i (B_i \Rightarrow B'_i) \wedge \bigwedge_i (E_i \text{ controls } p) \supset (P \text{ controls } p)$$

is valid for all p over binary-relation models.

6.3 An example

The most typical, complete example is that of a user logging into a workstation, and the workstation making requests on a server.

- The user A authenticates and delegates in some role R_A to a workstation B in role R_B . The user may rely on a smart-card, and use the scheme for user delegation outlined in subsection 6.1. The delegation certificate is:

$$K_A \text{ says } R_A \text{ says } ((K_d \wedge (B \text{ as } R_B)) \text{ serves } (A \text{ as } R_A))$$

- The workstation sets up a secure channel to a server. Logically, this requires two statements:

$$K_d \text{ says } (A \text{ as } R_A) \text{ says } (Ch \Rightarrow ((B \text{ as } R_B) \text{ for } (A \text{ as } R_A)))$$

under the delegation key, and

$$K_B \text{ says } R_B \text{ says } (A \text{ as } R_A) \text{ says } (Ch \Rightarrow ((B \text{ as } R_B) \text{ for } (A \text{ as } R_A)))$$

(For simplicity, the workstation acts in the role R_B , but any stronger role would do just as well.)

- The server needs to check that K_A is the user's key and that K_B is the workstation's key, that is, $K_A \Rightarrow A$ and $K_B \Rightarrow B$. Typically this requires looking at certificates from a certification authority; it is possible that a chain of certification authorities needs to be involved, as there may be no single universally trusted certification authority.
- With these properties of the keys, it follows from the delegation certificate that:

$$(A \text{ as } R_A) \text{ says } ((K_d \wedge (B \text{ as } R_B)) \text{ serves } (A \text{ as } R_A))$$

and this statement is believed. The channel set-up certificates yield:

$$((K_d \wedge (B \text{ as } R_B))|(A \text{ as } R_A)) \text{ says } (Ch \Rightarrow ((B \text{ as } R_B) \text{ for } (A \text{ as } R_A)))$$

and this leads to

$$((B \text{ as } R_B) \text{ for } (A \text{ as } R_A)) \text{ says } (Ch \Rightarrow ((B \text{ as } R_B) \text{ for } (A \text{ as } R_A)))$$

and then

$$Ch \Rightarrow ((B \text{ as } R_B) \text{ for } (A \text{ as } R_A))$$

as $(B \text{ as } R_B) \text{ for } (A \text{ as } R_A)$ is trusted on this matter.

- The user may indicate to the workstation that he wishes to reduce his privileges, and adopt a further role R'_A , in order to make a request r ; or the workstation may do this on behalf of the user, on its own initiative:

$$(Ch \text{ as } R'_A) \text{ says } r$$

The requestor here is $((B \text{ as } R_B) \text{ for } (A \text{ as } R_A)) \text{ as } R'_A$, which is equivalent to $(B \text{ as } R_B) \text{ for } (A \text{ as } R_A \text{ as } R'_A)$.

- The ACL at the server may contain $(G' \text{ as } R_B) \text{ for } (G \text{ as } R''_A)$. The server may have, or the workstation may present, certificates that prove that $A \Rightarrow G$, $R_A \Rightarrow R''_A$, $R'_A \Rightarrow R''_A$, and $B \Rightarrow G'$.
- Actually, the group membership certificates come signed with someone's public key. In each case it must be possible to resolve this public key into a name, and then to discover that the name is that of someone who is trusted to certify membership in the group.
- At this point, the algorithm for access control can easily determine that access should be granted.

7 Conclusions

Our main goal in this paper has been to isolate some useful and mathematically tractable concepts. We have only touched on many practical and theoretical matters that deserve more detailed treatments. In addition, the basic logic could be extended in many interesting ways. The full paper describes some of these extensions.

Acknowledgements

Morrie Gasser, Andy Goldstein, and Charlie Kaufman were at the origin of many of the ideas discussed here. Tim Mann, Garret Swart, and Ted Wobber participated in fruitful discussions about implementations and examples. Greg Nelson and Vaughan Pratt helped in understanding the algebra of principals. Roger Maddux told us about Andréka's work. Luca Cardelli and Cynthia Hibbard provided editorial suggestions.

References

- [1] M. Abadi, M. Burrows, C. Kaufman, and B. Lampson. Authentication and Delegation with Smart-Cards. Digital Equipment Corporation Systems Research Center report No. 67, October 1990.
- [2] M. Abadi, M. Burrows, B. Lampson, and G. Plotkin. A Calculus for Access Control in Distributed Systems. Digital Equipment Corporation Systems Research Center report No. 70, February 1990.
- [3] H. Andréka. Representations of Distributive Lattice-ordered Semigroups with Binary Relations. Manuscript, August 1989.
- [4] M. Burrows, M. Abadi, and R.M. Needham. A Logic of Authentication. *Proceedings of the Royal Society of London A* Vol. 426, 1989, pp. 233-271.
- [5] CCITT. CCITT Blue Book, Recommendation X.509 and ISO 9594-8: The Directory-Authentication Framework. Geneva, March 1988.
- [6] A. Chandra, D. Kozen, and L. Stockmeyer. Alternation. *JACM* Vol. 28, No. 1, January 1981, pp. 114-133.
- [7] National Bureau of Standards. Data Encryption Standard. Fed. Inform. Processing Standards Pub. 46. Washington DC, January 1977.
- [8] W. Diffie and M. Hellman. New Directions in Cryptography. *IEEE Transactions on Information Theory* IT-22, No. 6, November, 1976, pp. 644-654.
- [9] R. Fabry. Capability-based Addressing. *CACM* Vol. 17, No. 7, July 1974, pp. 403-412.
- [10] M. Gasser, A. Goldstein, C. Kaufman, and B. Lampson. The Digital Distributed System Security Architecture. Proceedings of the 1989 National Computer Security Conference, October 1989, pp. 305-319.

- [11] M. Gasser and E. McDermott. An Architecture for Practical Delegation in a Distributed System. Proceedings of the 1990 IEEE Symposium on Security and Privacy, May 1990, pp. 20–30.
- [12] G.E. Hughes and M.J. Cresswell. An Introduction to Modal Logic. Methuen Inc., New York, 1968.
- [13] J. Kohl, C. Neuman, and J. Steiner. The Kerberos Network Authentication Service (version 5, draft 3). Available by anonymous ftp from athena-dist.mit.edu as /pub/doc/kerberos/V5DRAFT3-RFC.{PS,TEXT}, October 1990.
- [14] B. Lampson, M. Abadi, M. Burrows, and E. Wobber. Authentication in Distributed Systems: Theory and Practice. To appear in the Proceedings of the Thirteenth Symposium on Operating System Principles, October 1991.
- [15] H. Levy. Capability-based Computer Systems. Digital Press, 1983.
- [16] S.P. Miller, C. Neuman, J.I. Schiller, and J.H. Saltzer. Kerberos Authentication and Authorization System. *Project Athena Technical Plan* Section E.2.1, MIT, July 1987.
- [17] R.M. Needham and M.D. Schroeder. Using Encryption for Authentication in Large Networks of Computers. *CACM* Vol. 21, No. 12, December 1978, pp. 993–999.
- [18] V. Pratt. Dynamic Algebras as a Well-behaved Fragment of Relation Algebras. In *Algebraic Logic and Universal Algebra in Computer Science*, Springer-Verlag LNCS 425, 1990.
- [19] R.L. Rivest, A. Shamir, and L. Adleman. A Method for Obtaining Digital Signatures and Public-key Cryptosystems. *CACM* Vol. 21, No. 2, February 1978, pp. 120-126.
- [20] J. Saltzer and M. Schroeder. The Protection of Information in Computer Systems. *Proceedings of the IEEE* Vol. 63, No. 9, September 1975, pp. 1278-1308.
- [21] K. Sollins. Cascaded Authentication. Proceedings of the 1988 IEEE Symposium on Security and Privacy, April 1988, pp. 156-163.