On SDSI's Linked Local Name Spaces

Martín Abadi
Digital Equipment Corporation
Systems Research Center
130 Lytton Avenue
Palo Alto, CA 94301, USA
ma@pa.dec.com

April 17, 1998

Abstract

Rivest and Lampson have recently introduced SDSI, a Simple Distributed Security Infrastructure. One of the important innovations of SDSI is the use of linked local name spaces. This paper suggests a logical explanation of SDSI's local name spaces, as a complement to the operational explanation given in the SDSI definition.

Contents

1 Linked Local Names Spaces **1**

2 Concepts and Notations **2**
2.1 Keys and Names 2
2.2 Compound Names 3
2.3 Statements 4
2.4 Bindings and Their Meanings 4
2.5 A General Form for Bindings 6

3 A Logic for Linked Local Name Spaces **6**
3.1 Syntax 6
3.2 Axioms 7
3.3 An Example 9
3.4 Simulating Name Resolution 11

4 Semantics **15**

5 `Self` **16**

6 Conclusions **17**

Appendix: Proofs **19**

References **22**

1 Linked Local Names Spaces

Rivest and Lampson have recently introduced a Simple Distributed Security Infrastructure (SDSI) [10]. One of the important innovations of SDSI is the use of linked local name spaces. In SDSI, each principal has a name space where names are bound to values, possibly by reference to the name spaces of other principals.

For example, in a particular local name space, the name `Lampson` may be associated with a public key KL. As a consequence of this binding, any statement whose signature can be verified using KL will be viewed as coming from `Lampson`. In addition, the name `Rivest` may be associated with `Lampson`'s `Rivest`. As a consequence of this binding, if `Lampson` says that `Rivest`'s public key is KR, then any statement whose signature can be verified using KR will be viewed as coming from `Rivest`. Compound names like `Lampson`'s `Rivest` allow one name space to import bindings from another.

Linked local name spaces offer the promise of combining some of the advantages of PGP-style local certification with those of hierarchical certification schemes (see, e.g., [4, 6, 12, 3, 5]). In particular, SDSI certification is egalitarian, and does not need to assume any global trust or any global notion of identity (beyond that inherent in public-key cryptography). On the other hand, SDSI can take advantage of structured trust relations and naming conventions, when such exist.

The precise rules for local names are not as simple as our first example may have suggested. Those rules deal with compound names such as `Lampson`'s `Rivest`'s `secretary` and `Lampson`'s `Rivest`'s `students`, and also with DNS names (that is, Internet e-mail names). The SDSI definition gives a rather operational account of local names. Basically, it explains local names by giving an algorithm for name resolution that maps names to their meanings.

This paper suggests an alternative account of local names. This account is based on a logic where one can express compound names such as `Lampson`'s `Rivest`, and assert that `Lampson`'s `Rivest` is bound to the public key KR. One can also ask whether `Lampson`'s `Lampson` is `Lampson`, and whether `Lampson`'s `Rivest`'s `secretary` is `Lampson`'s (`Rivest`'s `secretary`). These questions can be addressed independently of the binding for `Lampson`, and even when name resolution is not possible because `Lampson` has not yet been bound.

The purpose of the logic is to explain local names in a general, self-contained way, without requiring reference to particular implementations. Nevertheless, the SDSI name-resolution algorithm can be recast as a sound proof method within the logic. Hopefully, this work will contribute to the understanding of naming in SDSI and in related systems such as Simple Public Key Infrastructure (SPKI) [5].

The next section introduces basic concepts and notations. Section 3 develops a logic for linked local name spaces, and proves the soundness of the SDSI

name-resolution algorithm with respect to the logic. It also contains an example. Section 4 defines a semantics for the logic. The semantics serves as a mathematical tool, and as a precise counterpart to some of the informal explanations of naming. Section 5 concerns self-reference. An appendix contains all proofs.

This work can be seen as a descendant of the work of Lampson et al. in [9, 2, 11]. That work develops a calculus of principals and their statements, analogous to that of section 3; however, it does not explore the SDSI approach to naming.

At this time (September 1997 / April 1998), SDSI is still evolving. This paper is based on a comprehensive but informal description of SDSI 1.1 [10]. That document was the most recent SDSI definition available when a preliminary version of this paper was written [1], and it is still a useful reference. In the last few months, SDSI has been integrated with SPKI [5]. This paper briefly discusses some aspects of the resulting design.

The definitions of SDSI and SPKI are rather sophisticated and long. This paper does not attempt to cover them fully; ignoring many facets of those definitions, it focuses on linked local name spaces.

2 Concepts and Notations

This section is an informal discussion of basic concepts related to linked local name spaces. The main goal of this section is to review some of the ideas of SDSI at a suitable level of abstraction and to clarify their meanings. This section also introduces notations used in the rest of the paper.

2.1 Keys and Names

Hierarchical certification schemes rely on a general agreement about the roots of certification hierarchies; the names of these roots should be interpreted similarly in all name spaces. On the other hand, in egalitarian certification schemes, each principal can introduce names of its own, for private use. SDSI supports both styles of naming. In our study of SDSI, we therefore deal with both global and local ways of referring to principals:

- *Global identifiers* are understood equally in all name spaces. Global identifiers include public keys and global names. Public keys are global because any two principals should agree on the value of a public key when they see it, and should agree that a statement is signed using the inverse of a public key when they see the public key and the statement. In addition, the names of certain "special roots" are global, by special dispensation from the designers

of SDSI; for example, `DNS!!` is a global name, and represents the root of the DNS hierarchy.

- *Local names* are the auxiliary names that each principal can use in its name space. For example, `Lampson` may be used as a local name by anyone, and bound to any value. The value could be a global identifier (such as a public key), or another local name. Since two principals may disagree on the value of `Lampson`, they may not agree on whether a statement comes from `Lampson`.

We assume a set of global identifiers G and a disjoint set of local names N. The symbols f, g, and h range over G, and the symbols m and n over N.

2.2 Compound Names

In SDSI, the use of compound names is the basic mechanism for linking local name spaces. A compound name is an expression of the form (`ref`: $p_1 \ldots p_n$), which is abbreviated to p_1'`s` $\ldots$'`s` p_n. Each of $p_1, \ldots, p_n$ may be a global identifier, a local name, or a compound name. Intuitively, p'`s` q is the principal that is called q in the name space of the principal that is called p in the current name space. In general, p_1'`s` $\ldots$'`s` p_n is the principal that is called p_n in the name space $\ldots$ of the principal that is called p_1 in the current name space.

Going beyond these first intuitions about compound names, some questions immediately arise:

- Does (`ref`: $p_1 \ldots p_n$) make sense when $n = 0$? SDSI seems to allow this case. The meaning of (`ref`:) is the current principal—whoever is trying to do the name resolution.

- Does (`ref`: $p_1 \ldots p_n$) make sense when $n = 1$? SDSI explicitly allows this case, and says that (`ref`: p) is equivalent to p, since the first argument to `ref` is always interpreted in the current name space.

- Is ((p'`s` q)'`s` r) the same as (p'`s` q'`s` r) and the same as (p'`s` (q'`s` r))? This is not asserted in the SDSI paper, but is consistent with it. It is valid in the logic of this paper.

- Is (p'`s` p) the same as p? SDSI does not assume this, as it would be unreasonable, for example in the case where p is `spouse`.

In light of these properties, we consider `ref` only in the binary case, without loss of generality:

- To recover the expressiveness of the nullary case ($n = 0$), we introduce the construct `Self`, and replace (`ref:`) with `Self`. We postpone the study of `Self` to section 5.
- In the unary case, instead of writing (`ref:` p), we can simply write p.
- The case where $n > 2$ is reducible to the binary case, since for example we can regard (p'`s` q'`s` r) as an abbreviation for ((p'`s` q)$'$`s` r).

The set of *principal expressions* is the least set P containing G and N, and such that if $p, q \in P$ then (p'`s` q) $\in P$. Thus, a principal expression is a syntactic representation for a principal. The symbols p, q, and r range over P.

2.3 Statements

SDSI is concerned with statements that principals make. Those statements can be certificates and requests for service, in particular. All statements are understood with respect to a current name space.

A true statement may not be true forever. Consequently, certificates contain temporal information, indicating their periods of validity. For simplicity, we prefer to abstract away issues of time (for example, interesting questions about early versus late binding). Therefore, all statements are understood with respect to an arbitrary but fixed time, the current time.

We write p *says* s to indicate that the principal denoted by p makes statement s. When p is a public key K, the statement can be made by signing a textual representation of s using the private key that corresponds to K.

2.4 Bindings and Their Meanings

A *name space* is a set of bindings of names to values. We use the notation $n \mapsto v$ to represent the binding of the name n to the value v. The name n and the value v form a *name/value pair*. The value can be an arbitrary object, but the only case of relevance to this paper is that where the value is a principal expression. In a name space where name n is bound to principal expression p, any statement by p is taken as a statement by n.

Binding n to p is much like defining n to be p within the current name space. One may therefore view a binding as an assertion of equality. For example, one may bind the local name `lawyer` to the principal expression `Ted`$'$`s lawyer` when `lawyer` and `Ted`$'$`s lawyer` represent the same "real" principal.

Interpreting binding as equality is appealingly simple; but this interpretation runs into some trouble. Suppose, for example, that two different public keys $KL1$

and $KL2$ are values for the local name `Lampson` in some name space (perhaps because the "real" Lampson controls two keys). SDSI allows this:

> A local name can be undefined, or it can be bound to some value [...]. The principal may assign a value to a local name by issuing a corresponding certificate. If the local name already has a valid name/value certificate, the new certificate augments the old one, in the sense that a SDSI application is deemed to act correctly if it uses the name/value binding given in either certificate.

However, it is a little hard to make sense of the equalities $KL1 =$ `Lampson` $= KL2$, since $KL1$ and $KL2$ are different. These equalities may make sense if we view $KL1$, `Lampson`, $KL2$ as being three ways of referring to one "real" principal. It is not clear whether this is an appropriate semantic intuition. One of the principles of SDSI is that "principals are public keys"; the equalities $KL1 =$ `Lampson` $= KL2$ put us in an interesting situation where one principal is two different keys. Clearly, in this situation, the principle that "principals are public keys" should be understood loosely. Section 3 gives a stronger, formal argument against the equality interpretation.

In light of these difficulties, we may consider a weaker interpretation of binding. This weaker interpretation is possible because bindings need not be used in a symmetric way. For example, the binding of `lawyer` to `Ted's lawyer` may be used to turn a statement by `Ted's lawyer` into a statement by `lawyer`, not vice versa.

According to the weaker interpretation, a binding is simply an assertion of a *speaks-for* relation [9, 2]. Whenever p speaks for q, if p makes a statement then q makes the same statement. However, unlike equality, the speaks-for relation is not symmetric. With this weaker interpretation, the binding of `lawyer` to `Ted's lawyer` means that `Ted's lawyer` speaks for `lawyer`, not that `lawyer` is `Ted's lawyer`. Similarly, the binding of `Lampson` to $KL1$ and $KL2$ means that $KL1$ and $KL2$ both speak for `Lampson`, but does not imply the equality of `Lampson`, $KL1$, and $KL2$.

This paper generally favors the speaks-for interpretation of binding. However, it is not strictly necessary to commit to one specific interpretation. It suffices to assume properties of binding that seem compatible with either interpretation.

SPKI relies on the speaks-for relation, and enhances it with an elaborate, fine-grained notion of authorization. In SPKI, one can express not only that one principal speaks for another, but also that one principal speaks for another for statements in a particular class (for example, for requests for a particular service).

2.5 A General Form for Bindings

Suppose that a name space o contains the binding $\texttt{sysAdmin} \mapsto K1$ and many bindings of the form $n \mapsto (\texttt{sysAdmin}'\texttt{s CertAuthority}'\texttt{s}\ n)$. Suppose further that $K1$ says $\texttt{CertAuthority} \mapsto K2$. Within the name space o, we may deduce that $\texttt{sysAdmin}'\texttt{s CertAuthority}$ corresponds to $K2$, and then we may use that fact to deduce $n \mapsto (K2'\texttt{s}\ n)$.

In this situation, it is useful to generalize the $\mapsto$ relation so that we can write $p \mapsto q$ even when p is not a simple name. For example, we may write

$$(\texttt{sysAdmin}'\texttt{s CertAuthority}) \mapsto K2$$

in order to capture the association between the compound name $\texttt{sysAdmin}'\texttt{s CertAuthority}$ and the key $K2$. If name resolution was implemented using a cache, then the cache might record facts like $(\texttt{sysAdmin}'\texttt{s CertAuthority}) \mapsto K2$.

In general, $p \mapsto q$ means that p is bound to q in the current name space, either as the result of an explicit binding (like $\texttt{sysAdmin} \mapsto K1$), or as a consequence of other bindings (like $(\texttt{sysAdmin}'\texttt{s CertAuthority}) \mapsto K2$).

3 A Logic for Linked Local Name Spaces

In this section, we define a logic based on the concepts and notations of section 2. We also explain the SDSI name-resolution algorithm as a particular proof method within the logic, and show that it is sound.

3.1 Syntax

First we introduce a syntax for statements, building on section 2. We assume a set S of atomic statements (written $s_0, s_1, s_2, \ldots$). The set of *formulas* (or *statements*) is the least set such that:

- an atomic statement is a formula;
- if s and t are formulas then so are $\neg s$ and $s \wedge t$; we write the implication $s \supset t$ as an abbreviation for $\neg(s \wedge \neg t)$;
- if p and q are principal expressions then $p \mapsto q$ is a formula;
- if p is a principal expression and s is a formula then $p\ says\ s$ is a formula.

The symbols s and t range over the set of formulas.

The intuitive meanings of principal expressions, of $\mapsto$, and of $says$ are those explained in section 2. As discussed there, all formulas are understood with respect to a current name space and a current time.

For example, $KL\ says\ (\texttt{Rivest} \mapsto KR)$ is a formula if KL and KR are global identifiers and `Rivest` is a local name. This formula expresses an assertion signed with the inverse of the public key KL. The assertion is that `Rivest` is bound to KR. This binding is asserted for the local name space associated with KL, not universally.

3.2 Axioms

Next we give a proof system; this proof system has three components:

1. The standard axioms and rules of propositional logic, such as the axiom $(s \wedge t) \supset s$ and the rule

$$\frac{s \quad s \supset t}{t}$$

2. The standard axiom and rule of modal logic (e.g., [7]):

$$(p\ says\ (s \supset t)) \supset ((p\ says\ s) \supset (p\ says\ t))$$

$$\frac{s}{p\ says\ s}$$

   These are useful for manipulating statements within the scope of $says$, but we do not actually need them in this paper. The axiom expresses that the statements that a principal says are closed under consequence; the rule expresses that every principal says all provable statements.

3. New axioms that deal with linked local name spaces, listed in Figure 1.

   Reflexivity and Transitivity express that the binding relation is reflexive and transitive.

   Left-monotonicity means that if p is bound to q then $p'\texttt{s}\ r$ is bound to $q'\texttt{s}\ r$. We do not have that $r'\texttt{s}\ p$ is bound to $r'\texttt{s}\ q$, because p may not be bound to q in r's name space.

   Globality expresses that, if g is a global identifier and p an arbitrary principal expression, then $p'\texttt{s}\ g$ is bound to g. (Nevertheless, $p'\texttt{s}\ g$ might be bound to other values as well.) This axiom is justified by the globality of g.

Reflexivity:	$p \mapsto p$
Transitivity:	$(p \mapsto q) \supset ((q \mapsto r) \supset (p \mapsto r))$
Left-monotonicity:	$(p \mapsto q) \supset ((p'\mathtt{s}\ r) \mapsto (q'\mathtt{s}\ r))$
Globality:	$(p'\mathtt{s}\ g) \mapsto g$ if g is a global identifier.
Associativity:	$((p'\mathtt{s}\ q)'\mathtt{s}\ r) \mapsto (p'\mathtt{s}\ (q'\mathtt{s}\ r))$ $(p'\mathtt{s}\ (q'\mathtt{s}\ r)) \mapsto ((p'\mathtt{s}\ q)'\mathtt{s}\ r)$
Linking:	$(p\ says\ (n \mapsto r)) \supset ((p'\mathtt{s}\ n) \mapsto (p'\mathtt{s}\ r))$ if n is a local name.
Speaking-for:	$(p \mapsto q) \supset ((q\ says\ s) \supset (p\ says\ s))$

Figure 1: Axioms for linked local name spaces.

Associativity postulates that references are associative. Two formulas are used because we do not assume that $\mapsto$ is symmetric. (We actually need only the second formula.)

Linking expresses that if p says that n is bound to r (implicitly, in p's name space), then $p'\mathtt{s}\ n$ is bound to $p'\mathtt{s}\ r$ (implicitly, in the current name space). Thus, p is trusted on assertions about its name space.

Finally, Speaking-for expresses that if p is bound to q, then any statement by q is taken as a statement by p.

These axioms may be intuitively clear. In any case, they seem almost inevitable: it seems hard to justify SDSI's name-resolution algorithm without assuming most of the properties that the axioms express.

This axiom system could be extended in several ways. In particular, we could add:

(a) the converse of Globality: $g \mapsto (p'\mathtt{s}\ g)$ for g a global identifier;

(b) a generalization of Linking: $(p\ says\ (q \mapsto r)) \supset ((p'\mathtt{s}\ q) \mapsto (p'\mathtt{s}\ r))$ for q an arbitrary principal expression.

Both of these extensions may seem reasonable. Interestingly, however, there is a fundamental conflict between them. For instance, suppose that we have the formula $KL\ says\ (\mathtt{DNS!!} \mapsto KL)$, where KL is a public key; this formula expresses that KL says that DNS!! is bound to KL, a suspicious but possible statement. By (b), we can derive $(KL'\mathtt{s}\ \mathtt{DNS!!}) \mapsto (KL'\mathtt{s}\ KL)$; by (a), we can conclude $\mathtt{DNS!!} \mapsto KL$. This

is an undesirable consequence: it means that a strange statement by an arbitrary key KL can pollute the current name space. By Speaking-for, once we have the binding $\texttt{DNS!!} \mapsto KL$ any statement from KL will be viewed as a statement from DNS!!; thus, KL can hijack DNS!!. Therefore, (a) and (b) are incompatible.

Another possible addition is the following symmetry axiom:

(c) $(p \mapsto q) \supset (q \mapsto p)$.

This axiom is motivated by the equality interpretation of binding. Interestingly, it has undesirable consequences too. For instance, suppose that we have the formulas KL *says* $(n \mapsto KL)$ and KL *says* $(n \mapsto \texttt{DNS!!})$, where n is a local name. By Linking, Globality, and Transitivity, we obtain $(KL'\texttt{s}\ n) \mapsto KL$ and $(KL'\texttt{s}\ n) \mapsto \texttt{DNS!!}$. By the new symmetry axiom and Transitivity, we can conclude $\texttt{DNS!!} \mapsto KL$. Thus, again, statements by an arbitrary key KL can pollute the current name space. This provides a strong argument against the adoption of symmetry, and (at least in the setting of our logic) against the equality interpretation of binding.

In the arguments about (a), (b), and (c), somewhat plausible axioms lead to catastrophic results. A bit of paranoia is therefore in order. How can we know that our present axiom system does not lead to exactly the same results? We prove that it does not in Corollary 6 of section 4. We show that, no matter what KL says, one cannot derive $\texttt{DNS!!} \mapsto KL$. This proof should not be construed as a final guarantee of the security of our logic or of SDSI. It is however a useful sanity check.

3.3 An Example

In order to illustrate how the logic is used, we work through an example. Applying a SDSI abbreviation, we write smith@aol.com as a shorthand for

$$\texttt{DNS!!}'\texttt{s com}'\texttt{s aol}'\texttt{s smith}$$

Note that we have omitted some parentheses in the expression $\texttt{DNS!!}'\texttt{s com}'\texttt{s aol}'\texttt{s smith}$. By axiom Associativity, fortunately, we do not have to worry about how this expression is parenthesized. We omit parentheses in the rest of the example, with the same justification.

Consider the following bindings for the current name space and the following certificates:

1. the local name BrokersInc is bound to the public key $K1$;
2. the local name broker is bound to $\texttt{BrokersInc}'\texttt{s NYoffice}'\texttt{s Smith}$;
3. a certificate signed with the inverse of $K1$ says that the local name NYoffice is bound to the public key $K2$;

4. a certificate signed with the inverse of $K2$ says that the local name `Smith` is bound to `smith@aol.com`.

Logically, we represent these bindings and certificates with the formulas:

1. `BrokersInc` $\mapsto$ $K1$
2. `broker` $\mapsto$ (`BrokersInc`'s `NYoffice`'s `Smith`)
3. $K1$ *says* (`NYoffice` $\mapsto$ $K2$)
4. $K2$ *says* (`Smith` $\mapsto$ `smith@aol.com`)

Using these formulas as hypotheses, we construct the following semi-formal proof:

1. `BrokersInc` *says* (`NYoffice` $\mapsto$ $K2$) from hypotheses (1) and (3) by Speaking-for.
2. (`BrokersInc`'s `NYoffice`) $\mapsto$ (`BrokersInc`'s $K2$) from step (1) by Linking.
3. (`BrokersInc`'s `NYoffice`) $\mapsto$ $K2$ from step (2) by Globality and Transitivity.
4. (`BrokersInc`'s `NYoffice`) *says* (`Smith` $\mapsto$ `smith@aol.com`) from hypothesis (4) and step (3) by Speaking-for.
5. (`BrokersInc`'s `NYoffice`'s `Smith`) $\mapsto$ `smith@aol.com` from step (4) by Linking, Globality, and Transitivity, since `smith@aol.com` is a shorthand for `DNS!!`'s `com`'s `aol`'s `smith` and `DNS!!` is a global name.
6. `broker` $\mapsto$ `smith@aol.com` from hypothesis (2) and step (5) by Transitivity.

Thus, we obtain that the local name `broker` is bound to `smith@aol.com`.

Following Lampson, we can examine a proof like this one much like we would examine an audit trail. We can identify which of the hypotheses are used, and how they contribute to the conclusion. In this case, it is clear, for example, that both of the certificates from $K1$ and $K2$ are relevant, and that $K1$ and $K2$ must be trusted to have issued them appropriately.

A name like `broker` may be used in an ACL. Given a signed request, one may need to determine whether that request comes from `broker` in order to make an access-control decision. By the axiom Speaking-for and the binding `broker` $\mapsto$ `smith@aol.com`, it suffices to determine whether the signature on the request

$$
\begin{array}{lll}
\texttt{REF2}(o,\, p) & = & \text{if } p \text{ is a global identifier } f \\
 & & \text{then return } f \\
 & & \text{else if } p \text{ is a local name } n \\
 & & \qquad \text{and assumptions}(o) \text{ includes } n \mapsto q \\
 & & \text{then return } \texttt{REF2}(o, q) \\
 & & \text{else if } p \text{ is a compound name } q'\texttt{s}\ r \\
 & & \text{then return } \texttt{REF2}(\texttt{REF2}(o, q), r) \\
 & & \text{else fail}
\end{array}
$$

Figure 2: Name-resolution algorithm.

is that of `smith@aol.com`. This determination requires a chain of certificates starting from the root of the DNS hierarchy. However, the result `broker` $\mapsto$ `smith@aol.com` can be derived (and evaluated) before that chain of certificates is presented.

3.4 Simulating Name Resolution

Name resolution is the process of mapping a principal expression to a global identifier. Name resolution happens in the context of assumptions, in particular assumptions about the values of local names. The form of those assumptions may vary, depending on protocols and architecture. We consider a set of assumptions of the forms $n \mapsto q$ and $g\ says\ (n \mapsto q)$, where n is a local name, q a principal expression, and g a global identifier. The assumptions of the form $n \mapsto q$ correspond to bindings in the current name space. Those of the form $g\ says\ (n \mapsto q)$ correspond to certificates about bindings in other name spaces. Typically g is a public key K, and $g\ says\ (n \mapsto q)$ is implemented as a certificate signed with the inverse of K. We let E be the conjunction of all those assumptions. It is an easy exercise to include also bindings for global names (for example, a binding of `DNS!!` to a corresponding public key).

In SDSI, name resolution is specified as a recursive algorithm. The core of the algorithm is presented as a (nondeterministic) function `REF2`(o, p) that returns the meaning of p in the name space associated with o. The argument o may be either `current principal` (`cp`, for short) or a global identifier g.

Figure 2 gives an adaptation of SDSI's definition of `REF2`. It is fairly faithful, but simpler than the original because, for example, it does not deal with quoting or

encrypted objects. We write assumptions(o) for the set of assumptions for o:

- assumptions(cp) consists of the assumptions of the form $n \mapsto q$, and
- assumptions(g) consists of the assumptions of the form $g \ says \ (n \mapsto q)$.

The computation of REF2(o, p) is nondeterministic because assumptions(o) may include several bindings of the form $n \mapsto q$ for a given n. We say that REF2(o, p) yields f when f is one of the possible outputs of REF2(o, p).

In this section, we show how REF2 can be simulated within our logic. We carry this out in two steps. First, we define a new set of rules, called the *name-resolution rules*. These rules rely on the notation of the logic, but constitute a separate proof system. We establish that the name-resolution rules are equivalent to REF2. As a second step, we show that the name-resolution rules are sound with respect to the logic.

The name-resolution rules are simple rules designed for proving formulas of the form $E \supset (p \mapsto f)$. They are given in Figure 3. One can understand this set of rules as an algorithm. The algorithm starts out with an input p and assumptions E; it finds a suitable f, if one exists. The algorithm does a case analysis on the form of p. There is one rule (Base case 1) for when p is a global identifier f, one rule (Local name 1) for when p is a local name n, and four rules for when p is of the form $q'\mathtt{s}\ r$. Three of those four (Base case 2, Local name 2, Ref 2) do a case analysis on the form of r when q is a global identifier g. The remaining rule (Ref 1) applies whenever p is of the form $q'\mathtt{s}\ r$.

Despite their apparent differences, the name-resolution rules and REF2 are equivalent, as the following proposition shows:

Proposition 1 *Assume that a set of bindings, each of the form $(n \mapsto q)$, is given for* cp *and for each global identifier g. Let E be the conjunction of the formulas $(n \mapsto q)$, for each binding $(n \mapsto q)$ for* cp*, and $g \ says \ (n \mapsto q)$, for each binding $(n \mapsto q)$ for g. For every principal expression p and global identifiers g and f, we have:*

- *$E \supset (p \mapsto f)$ is provable with the name-resolution rules if and only if* REF2(cp, p) *yields f.*
- *$E \supset ((g'\mathtt{s}\ p) \mapsto f)$ is provable with the name-resolution rules if and only if* REF2(g, p) *yields f.*

The next proposition relates the name-resolution rules to the logic of section 3.2:

(Base case 1)

$$\frac{f \in G}{E \supset (f \mapsto f)}$$

(Base case 2)

$$\frac{f, g \in G}{E \supset ((g' \mathtt{s}\ f) \mapsto f)}$$

(Local name 1)

$$\frac{n \mapsto q \text{ is a conjunct in } E \qquad E \supset (q \mapsto f)}{E \supset (n \mapsto f)}$$

(Local name 2)

$$\frac{g\ \mathit{says}\ (n \mapsto q) \text{ is a conjunct in } E \qquad E \supset ((g' \mathtt{s}\ q) \mapsto f)}{E \supset ((g' \mathtt{s}\ n) \mapsto f)}$$

(Ref 1)

$$\frac{E \supset (q \mapsto f_0) \qquad E \supset ((f_0' \mathtt{s}\ r) \mapsto f)}{E \supset ((q' \mathtt{s}\ r) \mapsto f)}$$

(Ref 2)

$$\frac{g \in G \qquad E \supset ((g' \mathtt{s}\ q) \mapsto f_0) \qquad E \supset ((f_0' \mathtt{s}\ r) \mapsto f)}{E \supset ((g' \mathtt{s}\ (q' \mathtt{s}\ r)) \mapsto f)}$$

Figure 3: Name-resolution rules.

Proposition 2 *The name-resolution rules are sound with respect to the logic. That is, if $E \supset (p \mapsto f)$ is provable using those rules, then it is also provable in the logic.*

We conclude:

Corollary 3 *The name-resolution algorithm is sound with respect to the logic. That is, given E as in Proposition 1 and any principal expression p, if* REF2(cp, p) *yields a global identifier f, then $E \supset (p \mapsto f)$ is provable in the logic.*

The converse of this property does not hold: sometimes $E \supset (p \mapsto f)$ is provable in the logic although REF2(cp, p) does not yield f.

- The logic is more powerful than the algorithm in that it permits reasoning even when full name resolution is impossible. While it would be easy to forbid such reasoning, it seems beneficial and harmless.

  For example, the logic's Globality axiom yields (Lampson′s KR) $\mapsto$ KR, without any need for additional hypotheses. In contrast, the computation of REF2(cp, Lampson′s KR) fails when the name Lampson is not bound.

  Similarly, using the logic, we can reason that if Lampson $\mapsto$ KL and Rivest $\mapsto$ (Lampson′s Rivest) then Rivest $\mapsto$ (KL′s Rivest), independently of any certificates about KL′s Rivest. Without such certificates, the algorithm fails.

- Other differences between the logic and the algorithm may be more surprising. In particular, given the assumptions Lampson $\mapsto$ $KL1$, Lampson $\mapsto$ $KL2$, $KL1$ *says* (Ron $\mapsto$ Rivest), and $KL2$ *says* (Rivest $\mapsto$ KR), the logic enables us to derive (Lampson′s Ron) $\mapsto$ KR. (The logical proof relies on the axiom Speaking-for, for example to go from Lampson $\mapsto$ $KL1$ and $KL1$ *says* (Ron $\mapsto$ Rivest) to Lampson *says* (Ron $\mapsto$ Rivest).) In contrast, REF2(cp, Lampson′s Ron) does not yield KR.

  A recent definition of SPKI incorporates SDSI names. Following this definition, we would not write $KL1$ *says* (Ron $\mapsto$ Rivest) but rather $KL1$ *says* ($KL1$′s Ron $\mapsto$ $KL1$′s Rivest). More generally, we would consider only bindings of the form $p \mapsto g$ or $p \mapsto g$′s q, where the principal expression on the right of $\mapsto$ starts with a global identifier. Interestingly, this restriction suffices to block our proof of (Lampson′s Ron) $\mapsto$ KR. Perhaps a logic with this restriction would be worth investigating.

4 Semantics

In this section, we define a semantics for the logic of section 3. By the standards of logicians, the semantics is rather elementary. It could probably be enhanced (and complicated) using ideas from the modal-logic literature (see, e.g., [7, 2, 8]). Despite its simplicity, the semantics was helpful in refining the rules of section 3, and it can be rather useful, as we demonstrate with Corollary 6 below.

Recall that G is the set of global identifiers, N the set of local names, P the set of all principal expressions, and S the set of atomic propositions. The semantics assumes:

- a set $\mathcal{W}$,
- a mapping α from G to subsets of $\mathcal{W}$,
- a mapping ρ from $N \times \mathcal{W}$ to subsets of $\mathcal{W}$, and
- a mapping μ from $S \times \mathcal{W}$ to truth values.

In the vocabulary of modal logic, $\mathcal{W}$ is a set of possible worlds. The choice of the set $\mathcal{W}$ does not matter for our purposes. What matters is that $\mathcal{W}$ has a subset $\alpha(g)$ for each $g \in G$; a subset $\rho(n, a)$ for each $n \in N$ and $a \in \mathcal{W}$; and that $\mu(s_i, a)$ indicates whether s_i is true or false at each $a \in \mathcal{W}$.

Given $p \in P$ and $a \in \mathcal{W}$, we define $[\![p]\!]_a$ to be a subset of $\mathcal{W}$, as follows:

- $[\![g]\!]_a = \alpha(g)$ if $g \in G$,
- $[\![n]\!]_a = \rho(n, a)$ if $n \in N$,
- $[\![p' \mathtt{s}\ q]\!]_a = \bigcup\{[\![q]\!]_b \mid b \in [\![p]\!]_a\}$.

For each p, we obtain a relation on $\mathcal{W}$, namely $\{(a, b) \mid b \in [\![p]\!]_a\}$. The relation for $p' \mathtt{s}\ q$ is simply the composition of the relations for p and q.

Given a statement s and $a \in \mathcal{W}$, we write $a \models s$ to mean that s is true at a. The definition of $a \models s$ is:

- $a \models s_i$ is $\mu(s_i, a)$, for $s_i \in S$,
- $a \models (s \wedge t)$ is true iff both $a \models s$ and $a \models t$ are true,
- $a \models (\neg s)$ is true iff $a \models s$ is false,
- $a \models (p \mapsto q)$ is true iff $[\![p]\!]_a \subseteq [\![q]\!]_a$,
- $a \models (p\ says\ s)$ is true iff $b \models s$ is true for all $b \in [\![p]\!]_a$.

Here we are not using the equality interpretation of binding, as the interpretation of $p \mapsto q$ is obviously not symmetric. The interpretation guarantees that q speaks for p in the sense that if $q\ says\ s$ then $p\ says\ s$.

A statement s is *valid* if and only if $a \models s$ is true for all choices of $\mathcal{W}$, α, ρ, and μ, and all $a \in \mathcal{W}$. The following proposition establishes the soundness of the logic of section 3 with respect to the semantics of this section. The converse to this proposition (completeness) is probably false.

Proposition 4 *If s is provable in the logic, then s is valid.*

Soundness implies consistency:

Corollary 5 *If s is provable in the logic, then $\neg s$ is not.*

As a sanity check, we can now show that some undesirable formulas (discussed in section 3) are not derivable in the logic. We show that a formula $\texttt{DNS!!} \mapsto KL$ cannot be derived from two hypotheses s and t. The first hypothesis, s, lists some "normal" statement that DNS!! makes, binding local names to principal expressions; one may be able to generalize the corollary to include other "normal" statements. The second hypothesis, t, lists statements that KL makes, possibly with the goal of hijacking DNS!!; these statements are completely arbitrary. The corollary guarantees that, despite these arbitrary statements by KL, one cannot derive that DNS!! is bound to KL.

Corollary 6 *Assume that KL and* DNS!! *are distinct global identifiers. Let s be any formula of the form* DNS!! $says\ (n_1 \mapsto p_1) \wedge \ldots \wedge$ DNS!! $says\ (n_k \mapsto p_k)$*, where $n_1, \ldots, n_k$ are local names and $p_1, \ldots, p_k$ are arbitrary principal expressions. Let t be any formula of the form $KL\ says\ t_1 \wedge \ldots \wedge KL\ says\ t_k$, where $t_1, \ldots, t_k$ are arbitrary statements. Then the formula $s \wedge t \supset (\texttt{DNS!!} \mapsto KL)$ is not derivable in the logic.*

5 Self

In this section, we consider the construct Self, which is our notation for ref with no arguments, (ref:).

Intuitively, Self represents the current principal. Therefore, it seems sensible to add Self to the set of principal expressions, and to extend the logic with the following four axioms:

$$\begin{array}{lll} \text{Identity:} & (\texttt{Self's}\ p) \mapsto p & p \mapsto (\texttt{Self's}\ p) \\ & (p'\texttt{s Self}) \mapsto p & p \mapsto (p'\texttt{s Self}) \end{array}$$

It is also easy to extend the semantics, validating the new axioms. (It suffices to let $[\![\texttt{Self}]\!]_a = \{a\}$ for every $a \in \mathcal{W}$.)

Given the formal properties of `Self`, it may appear as a fairly dull construct. After all, the new axioms show how it can be eliminated from expressions; and the semantics of `Self` is routine.

Nevertheless, the use of `Self` is sometimes convenient. For example, it enables us to write the formula p *says* $(n \mapsto \texttt{Self})$. In the case in which p is a public key K and n is `Ted`, this statement corresponds to a certificate that binds `Ted` to K. The certificate would not need to mention K explicitly. Instead, it can refer to K as `Self`, that is, the public key used for verifying the signature on the certificate. Given K *says* $(\texttt{Ted} \mapsto \texttt{Self})$, we can derive $(K\texttt{'s Ted}) \mapsto (K\texttt{'s Self})$, hence $(K\texttt{'s Ted}) \mapsto K$. In this respect, K *says* $(\texttt{Ted} \mapsto \texttt{Self})$ has the same effect as K *says* $(\texttt{Ted} \mapsto K)$, but the "wire representation" of K *says* $(\texttt{Ted} \mapsto \texttt{Self})$ might be different.

Continuing this example, let us assume that $\texttt{Edward} \mapsto K$ holds in the current name space. From K *says* $(\texttt{Ted} \mapsto \texttt{Self})$ we obtain `Edward` *says* $(\texttt{Ted} \mapsto \texttt{Self})$, and then $(\texttt{Edward's Ted}) \mapsto \texttt{Edward}$. This is a nice minor result that we would not have reached from K *says* $(\texttt{Ted} \mapsto K)$.

According to our new axioms, $K\texttt{'s Self}$ and K are bound to each other. Therefore, K could say, truthfully, $K \mapsto \texttt{Self}$. Since K speaks for `Edward`, `Edward` would say $K \mapsto \texttt{Self}$ as well. This is a somewhat perplexing consequence of apparently sensible hypotheses: a true statement by K leads to a suspicious statement by `Edward`. On the other hand, `Edward`'s statement should be harmless, since `Edward` should not have the authority to convince others that it speaks for K when in fact `Edward` is bound to K but not vice versa.

In summary, the construct `Self` (in SDSI, `(ref: ))` may be useful occasionally. However, it has the potential to be surprising; like all relative names, it should be handled with care.

6 Conclusions

SDSI's linked local name spaces appear as a promising innovation. This paper attempts to contribute to their understanding, both through informal explanations and through the development of a logic. The logic generalizes the SDSI name-resolution algorithm. It permits not only the evaluation of compound names, but also reasoning about those names and their bindings. Having a logic is not synonymous with security; however, a logic can complement an operational approach, and provide another perspective from which to examine security issues.

Acknowledgements

Ron Rivest made many useful comments on a draft of this paper. Several discussions with Carl Ellison, Butler Lampson, and Ron Rivest were generally helpful.

Appendix: Proofs

Proposition 1 *Assume that a set of bindings, each of the form* $(n \mapsto q)$*, is given for* cp *and for each global identifier* g*. Let* E *be the conjunction of the formulas* $(n \mapsto q)$*, for each binding* $(n \mapsto q)$ *for* cp*, and* g *says* $(n \mapsto q)$*, for each binding* $(n \mapsto q)$ *for* g*. For every principal expression* p *and global identifiers* g *and* f*, we have:*

- $E \supset (p \mapsto f)$ *is provable with the name-resolution rules if and only if* REF2(cp, p) *yields* f.
- $E \supset ((g\text{'s } p) \mapsto f)$ *is provable with the name-resolution rules if and only if* REF2(g, p) *yields* f.

Proof The "if" direction of the proof is by induction on the execution of REF2(o, p), for both kinds of o together.

- First, let us suppose that the algorithm terminates immediately returning p when p is a global identifier f. In that case, we obtain:
  - $E \supset (f \mapsto f)$ by rule Base case 1,
  - $E \supset ((g\text{'s } f) \mapsto f)$ by rule Base case 2.
- Next, if p is a local name m, and the assumptions for o include $(m \mapsto q)$, then the algorithm invokes itself recursively, and REF2(o, q) returns f. By induction hypothesis, $E \supset (q \mapsto f)$ is provable when o is cp, and $E \supset ((g\text{'s } q) \mapsto f)$ is provable when o is g. We obtain:
  - $E \supset (m \mapsto f)$ by rule Local name 1,
  - $E \supset ((g\text{'s } m) \mapsto f)$ by rule Local name 2.
- Finally, suppose that p is $(q\text{'s } r)$. Then the algorithm invokes itself recursively twice. In the first recursive invocation, REF2(o, q) yields a result f_0. By induction hypothesis, $E \supset (q \mapsto f_0)$ is provable if o is cp, and $E \supset ((g\text{'s } q) \mapsto f_0)$ is provable if o is g. In the second recursive invocation, REF2(f_0, r) yields the final result f. By induction hypothesis, $E \supset ((f_0\text{'s } r) \mapsto f)$ is provable. We obtain:
  - $E \supset ((q\text{'s } r) \mapsto f)$ by rule Ref 1,
  - $E \supset ((g\text{'s } (q\text{'s } r)) \mapsto f)$ by rule Ref 2.

The "only if" direction of the proof is by induction on the proof of $E \supset (p \mapsto f)$ or $E \supset ((g\text{'s } p) \mapsto f)$. There is one case for each of the name-resolution rules:

- Base case 1: trivial, since $\texttt{REF2}(\texttt{cp}, f)$ yields f.
- Base case 2: trivial, since $\texttt{REF2}(\texttt{cp}, (g'\texttt{s}\ f))$ and $\texttt{REF2}(g, f)$ yield f.
- Local name 1: By induction hypothesis, $\texttt{REF2}(\texttt{cp}, q)$ yields f. This implies that $\texttt{REF2}(\texttt{cp}, n)$ yields f if n is bound to q.
- Local name 2: By induction hypothesis, $\texttt{REF2}(g, q)$ yields f. This implies that both $\texttt{REF2}(\texttt{cp}, (g'\texttt{s}\ n))$ and $\texttt{REF2}(g, n)$ yield f if n is bound to q in g's name space.
- Ref 1: By induction hypothesis, $\texttt{REF2}(\texttt{cp}, q)$ yields f_0 and $\texttt{REF2}(f_0, r)$ yields f, hence $\texttt{REF2}(\texttt{cp}, (q'\texttt{s}\ r))$ yields f. In the case where q happens to be a global identifier h, we also need to prove that $\texttt{REF2}(h, r)$ yields f. But if q is h and $\texttt{REF2}(\texttt{cp}, q)$ yields f_0, then h is f_0, and we already have that $\texttt{REF2}(f_0, r)$ yields f.
- Ref 2: By induction hypothesis, $\texttt{REF2}(\texttt{cp}, (g'\texttt{s}\ q))$ and $\texttt{REF2}(g, q)$ yield f_0, and $\texttt{REF2}(f_0, r)$ yields f. First we prove that $\texttt{REF2}(\texttt{cp}, (g'\texttt{s}\ (q'\texttt{s}\ r)))$ yields f. An invocation of $\texttt{REF2}(\texttt{cp}, (g'\texttt{s}\ (q'\texttt{s}\ r)))$ turns into an invocation of $\texttt{REF2}(\texttt{REF2}(\texttt{cp}, (g'\texttt{s}\ q)), r)$; since $\texttt{REF2}(\texttt{cp}, (g'\texttt{s}\ q))$ yields f_0, this turns into an invocation of $\texttt{REF2}(f_0, r)$, which yields f. Next we prove that $\texttt{REF2}(g, q'\texttt{s}\ r)$ yields f. This is an immediate consequence of the facts that $\texttt{REF2}(g, q)$ yields f_0 and $\texttt{REF2}(f_0, r)$ yields f.

□

Proposition 2 *The name-resolution rules are sound with respect to the logic. That is, if $E \supset (p \mapsto f)$ is provable using those rules, then it is also provable in the logic.*

Proof The proof is by induction on the proof of $E \supset (p \mapsto f)$ with the name-resolution rules. There is one case for each of the rules.

- Base case 1: by axiom Reflexivity.
- Base case 2: by axiom Globality. (Note that only a special case of Globality is needed: $(f'\texttt{s}\ g) \mapsto g$.)
- Local name 1: by axiom Transitivity.
- Local name 2: by axioms Linking and Transitivity.
- Ref 1: by axioms Left-monotonicity and Transitivity.

- Ref 2: by axioms Associativity, Left-monotonicity, and Transitivity. (Note that Associativity is used only in one direction, and in a special case, namely $(f'\mathtt{s}\ (q'\mathtt{s}\ r)) \mapsto ((f'\mathtt{s}\ q)'\mathtt{s}\ r)$.)

□

Proposition 4 *If s is provable, then s is valid.*

Proof The propositional axioms and rules are validated, since the propositional connectives are interpreted in the usual manner. So are the standard axiom and rule of modal logic. Reflexivity and Transitivity are validated, straightforwardly. Associativity holds essentially because the composition of binary relations is associative. The remaining axioms require small arguments:

- For Globality: If g is a global identifier, then $[\![g]\!]_a$ is independent of a (and is determined by α alone); therefore, $[\![p'\mathtt{s}\ g]\!]_a = \bigcup\{[\![g]\!]_b \mid b \in [\![p]\!]_a\} \subseteq [\![g]\!]_a$. Equality holds whenever $[\![p]\!]_a \neq \emptyset$; otherwise, $[\![p'\mathtt{s}\ g]\!]_a = \emptyset$.

- For Linking: If $a \models (p\ says\ (n \mapsto r))$ is true, then $b \models n \mapsto r$ is true for all $b \in [\![p]\!]_a$, and $[\![n]\!]_b \subseteq [\![r]\!]_b$ for all such b. Hence, $\bigcup\{[\![n]\!]_b \mid b \in [\![p]\!]_a\} \subseteq \bigcup\{[\![q]\!]_b \mid b \in [\![p]\!]_a\}$, that is, $[\![p'\mathtt{s}\ n]\!]_a \subseteq [\![p'\mathtt{s}\ r]\!]_a$. Therefore, $a \models (p'\mathtt{s}\ n) \mapsto (p'\mathtt{s}\ r)$ is true.

- For Speaking-for: If $a \models q \mapsto r$ is true, then $[\![q]\!]_a \subseteq [\![r]\!]_a$. Therefore, if $a \models r\ says\ s$ is true, then $b \models s$ is true for all $b \in [\![r]\!]_a$, and $b \models s$ is true for all $b \in [\![q]\!]_a$, so $a \models q\ says\ s$ is true.

□

Corollary 6 *Assume that KL and* DNS!! *are distinct global identifiers. Let s be any formula of the form* DNS!! *says* $(n_1 \mapsto p_1) \wedge \ldots \wedge$ DNS!! *says* $(n_k \mapsto p_k)$*, where $n_1, \ldots, n_k$ are local names and $p_1, \ldots, p_k$ are arbitrary principal expressions. Let t be any formula of the form KL says $t_1 \wedge \ldots \wedge KL$ says t_k, where $t_1, \ldots, t_k$ are arbitrary statements. Then the formula $s \wedge t \supset$ (DNS!! $\mapsto KL$) is not derivable in the logic.*

Proof If the formula in question was derivable, then it would be valid, according to Proposition 4. So we show that it is not valid. It suffices to give a set $\mathcal{W}$, an element $a \in \mathcal{W}$, and functions α, ρ, and μ such that $a \models s \wedge t$ is true while $a \models$ DNS!! $\mapsto KL$ is false. Trivial choices suffice. We let $\mathcal{W}$ be the singleton $\{1\}$, let $a = 1$, let $\alpha(KL) = \emptyset$ and $\alpha($DNS!!$) = \mathcal{W}$, and let $\rho(n_i, a) = \emptyset$ for all local names n_i. (Any remaining properties of α, ρ, and μ are irrelevant.) We have that $a \models s$ is true, since $\rho(n_i, a) = \emptyset$, that $a \models t$ is true, since $\alpha(KL) = \emptyset$, and that $a \models$ DNS!! $\mapsto KL$ is false, since $\alpha($DNS!!$) \not\subseteq \alpha(KL)$. □

References

[1] Martín Abadi. On SDSI's linked local name spaces. In *10th IEEE Computer Security Foundations Workshop*, pages 98–108. IEEE Computer Society, June 1997.

[2] Martín Abadi, Michael Burrows, Butler Lampson, and Gordon Plotkin. A calculus for access control in distributed systems. *ACM Transactions on Programming Languages and Systems*, 15(4):706–734, October 1993.

[3] Matt Blaze, Joan Feigenbaum, and Jack Lacy. Decentralized trust management. In *Proceedings 1996 IEEE Symposium on Security and Privacy*, pages 164–173, May 1996.

[4] CCITT. *Blue Book (Recommendation X.509 and ISO 9594-8: The directory-authentication framework)*. CCITT, 1988.

[5] Carl M. Ellison. SPKI certificate documentation. Web pages at `http://www.clark.net/pub/cme/html/spki.html`, 1997.

[6] Morrie Gasser, Andy Goldstein, Charlie Kaufman, and Butler Lampson. The Digital Distributed System Security Architecture. In *Proceedings of the 1989 National Computer Security Conference*, pages 305–319, October 1989.

[7] Robert Goldblatt. *Logics of Time and Computation*. Number 7 in CSLI Lecture Notes. CSLI, Stanford, 1987.

[8] Adam J. Grove and Joseph Y. Halpern. Naming and identity in epistemic logics, I: The propositional case. *Journal of Logic and Computation*, 3(4):345–378, 1993.

[9] Butler Lampson, Martín Abadi, Michael Burrows, and Edward Wobber. Authentication in distributed systems: Theory and practice. *ACM Transactions on Computer Systems*, 10(4):265–310, November 1992.

[10] Ronald L. Rivest and Butler Lampson. SDSI — A Simple Distributed Security Infrastructure. Version 1.1, at `http://theory.lcs.mit.edu/~rivest/sdsi11.html`, replicated at `http://www.research.digital.com/SRC/personal/Martin_Abadi/Papers/sdsi11.html`, October 2, 1996.

[11] Edward Wobber, Martín Abadi, Michael Burrows, and Butler Lampson. Authentication in the Taos operating system. *ACM Transactions on Computer Systems*, 12(1):3–32, February 1994.

[12] Philip R. Zimmermann. *The Official PGP User's Guide*. MIT Press, 1995.