

Beyond Proof-of-compliance: Safety and Availability Analysis in Trust Management

NINGHUI LI and JOHN C. MITCHELL

Stanford University

and

WILLIAM H. WINSBOROUGH

Network Associates Laboratories

Trust management is a form of distributed access control using distributed policy statements. Since one party may delegate partial control to another party, it is natural to ask what permissions may be granted as the result of policy changes by other parties. We study security properties such as safety and availability for a family of trust management languages, devising algorithms for deciding the possible consequences of certain changes in policy. While trust management is more powerful in certain ways than mechanisms in the access matrix model, and the security properties considered are more than simple safety, we find that in contrast to the classical HRU undecidability of safety properties, our primary security properties are decidable. In particular, most properties we studied are decidable in polynomial time. Containment, the most complicated security property we studied, is decidable in polynomial time for the simplest TM language in the family. The problem becomes coNP-hard when intersection or linked roles are added to the language.

Categories and Subject Descriptors: []:

1. INTRODUCTION

Trust Management (TM) is an approach to access control in decentralized distributed systems with access control decisions based on policy statements made by multiple principals. A key aspect of trust management is delegation: a principal may transfer limited authority over one or more resources to other principals. Since delegation gives a certain degree of control to a principal that may be only partially trusted, a natural security concern is whether a resource owner still has some guarantees about who can access their resources. If we think of the union of all policies of all principals as the state of a TM system, then a resource owner always has control over some part of the state, but not necessarily all parts. In this paper, we consider the *security analysis* problem, which asks what accesses may be allowed or prevented by prospective changes in the state of a TM system.

A few definitions are useful for stating the security analysis problem more precisely. In general, a TM language has a syntax for specifying *policy statements* and *queries*, together with a semantic relation \vdash . We call a set \mathcal{P} of policy statements a *state* of a TM system. Given a state \mathcal{P} and a query Q , the relation $\mathcal{P} \vdash Q$ means that Q is true in \mathcal{P} . When

Authors' addresses: Ninghui Li, Department of Computer Science, Gates 4B, Stanford, CA 94305-9045, USA; email: ninghui@stanford.edu. John C. Mitchell, Department of Computer Science, Gates 4B, Stanford, CA 94305-9045, USA; email: mitchell@cs.stanford.edu. William H. Winsborough, 15204 Omega Drive, Suite 300, Rockville, MD 20850-4601, william.winsborough@nai.com.

Q arises from an access request, $\mathcal{P} \vdash Q$ means that access Q is allowed in \mathcal{P} ; a proof demonstrating $\mathcal{P} \vdash Q$ is then called a *proof-of-compliance*.

Recognizing that a principal or a coalition of cooperating principals may control only a part of the global state, we assume there is a *restriction rule*, \mathcal{R} , that defines the uncontrolled changes of state that are possible. For example, the principal in question may consider the part of the state controlled by fully trusted principals to be fixed, while considering that other principals may remove some policy statements and/or add new ones. Given a state \mathcal{P} and a restriction rule \mathcal{R} , we write $\mathcal{P} \mapsto_{\mathcal{R}} \mathcal{P}'$ if the change from \mathcal{P} to \mathcal{P}' is allowed by \mathcal{R} , and $\mathcal{P} \mapsto^*_{\mathcal{R}} \mathcal{P}'$ if a sequence of zero or more allowed changes leads from \mathcal{P} to \mathcal{P}' . If $\mathcal{P} \mapsto^*_{\mathcal{R}} \mathcal{P}'$, we say that \mathcal{P}' is \mathcal{R} -*reachable* from \mathcal{P} , or simply \mathcal{P}' is *reachable*, when \mathcal{P} and \mathcal{R} are clear from context.

DEFINITION 1. Let \mathcal{P} be a state, \mathcal{R} a restriction rule, and Q a query. An *existential security analysis query* takes the form: Does there exist \mathcal{P}' such that $\mathcal{P} \mapsto^*_{\mathcal{R}} \mathcal{P}'$ and $\mathcal{P}' \vdash Q$? When the answer is affirmative, we say Q is *possible* given \mathcal{P} and \mathcal{R} . A *universal security analysis query* takes the form: For every \mathcal{P}' such that $\mathcal{P} \mapsto^*_{\mathcal{R}} \mathcal{P}'$, does $\mathcal{P}' \vdash Q$? If so, we say Q is *necessary* given \mathcal{P} and \mathcal{R} .

Here are some motivating examples of security analysis queries.

Simple Safety. (Existential) Does there exist a reachable state in which a (presumably untrusted) principal has access to a resource?

Simple Availability. (Universal) In every reachable state, does a (presumably trusted) principal have access to a resource?

Bounded Safety. (Universal) In every reachable state, is the set of all principals that have access to a resource bounded by a given set of principals?

Containment. (Universal) In every reachable state, does every principal that has one property (e.g., has access to a resource) also have another property (e.g., being an employee)? A query of this form can express safety or availability (e.g., by interchanging the two example properties in the previous sentence).

Simple safety queries were first formalized in [Harrison et al. 1976] in the context of the well-known access matrix model [Lampson 1971]. The model in [Harrison et al. 1976] is commonly known as the HRU model. In the general HRU model, *safety analysis*, i.e., answering simple safety queries, is undecidable [Harrison et al. 1976]. A number of protection models were developed to address this, for example, the take-grant model [Lipton and Snyder 1977], the Schematic Protection Model [Sandhu 1988], and the Typed Access Matrix model [Sandhu 1992]. The other kinds of queries listed above were not considered in these previous works. Since some of our queries are about availability rather than safety, we use the term *security analysis* rather than safety analysis.

Security analysis is especially relevant in TM systems, which focus on delegation. While the use of delegation greatly enhances flexibility, it may also reduce the control that a principal has over the resources it owns. Therefore a resource owner who is considering delegating control over access rights needs assistance in assessing the potential consequences. Who could get access? Who could be denied? These simple questions lead to the security analysis queries considered here.

Security analysis is needed when the state of a TM system may change in the future. It is also useful when the global state of a TM system is fixed, but only partially known.

Such partial knowledge is typical in distributed systems, and TM is no exception. For instance, previously unknown statements may be presented along with new access requests. Thus, although the global state does not change, one’s view of that state is changing. To understand the potential consequences of candidate delegation acts, one needs to consider as-yet-unknown policy statements that could exist.

To the best of our knowledge, security analysis for TM systems has not been investigated previously as such. In this paper, we define a precise model for security analysis in trust management. The policy language we consider is a slightly simplified (yet expressively equivalent) version of the RT_0 language [Li et al. 2003b]. We call the language SRT , for Simplified RT_0 . All the security analysis queries listed above are considered. While the TM language we are studying supports delegation and is more expressive than the access matrix model in certain ways, and the kinds of queries we are considering are more general, somewhat surprisingly, answering these queries are decidable. Simple safety, simple availability, and bounded safety queries can all be answered in time polynomial in the size of \mathcal{P} . Containment queries can be more expensive, depending on which policy language features are used. We show that, for BRT , a sub-language of SRT that has simple delegation, containment queries can be answered in polynomial time. Adding intersection to BRT increases the complexity of answering containment queries to coNP-complete. Adding a linked-role feature to BRT increases the complexity of answering containment queries to PSPACE-complete. For the case of SRT , which is obtained by adding both intersection and role linking, answering containment queries is still decidable (in coNEXP).

The rest of this paper is organized as follows. In Section 2, we define the model we use to study security analysis in TM. In Section 3, we handle simple safety, simple availability, and bounded safety queries. In Section 4, we present results about containment queries. We discuss related work in Section 5, and conclude in Section 6. An appendix contains proofs that are not included in the main body.

2. INSTANTIATING THE SECURITY ANALYSIS PROBLEM

The abstract definition of security analysis in Definition 1 has three parameters: the language used to express \mathcal{P} , the form of query \mathcal{Q} , and the restriction rule \mathcal{R} . In this section, we instantiate the security analysis problem with specific parameters, discuss our choices, and give an example that will be used throughout the paper.

2.1 Syntax of The TM Language

The policy language we consider is SRT , a slightly simplified version of RT_0 [Li et al. 2003b]. As in RT_0 , the basic constructs of SRT include *principals* and *role names*. In this paper, we use A, B, D, E, F, X, Y , and Z , sometimes with subscripts, to denote principals. A role name is a word over some given standard alphabet. We use r, u , and w , sometimes with subscripts, to denote role names. A *role* takes the form of a principal followed by a role name, separated by a dot, e.g., $A.r$ and $X.u$. A *role* defines a set of principals who are members of this role. Each principal A has the authority to define who are the members of each role of the form $A.r$. An access control permission is represented as a role as well; for example, that B is a member of the role of $A.r$ may represent that B has the permission to do action r on the object A .

There are four types of policy statements in SRT , each corresponding to a different way of defining role membership:

— *Type-1:* $A.r \longleftarrow D$

This statement means that A defines D to be a member of A 's r role.

— *Type-2:* $A.r \longleftarrow B.r_1$

This statement means that A defines its r role to include (all members of) B 's r_1 role. This represents a delegation from A to B , since B may affect who is a member of the role $A.r$ by issuing statements.

— *Type-3:* $A.r \longleftarrow A.r_1.r_2$

We call $A.r_1.r_2$ a *linked role*. This statement means that A defines $A.r$ to include $B.r_2$ for every B that is a member of $A.r_1$. This represents a delegation from A to all the members of the role $A.r_1$.

— *Type-4:* $A.r \longleftarrow B_1.r_1 \cap B_2.r_2$

We call $B_1.r_1 \cap B_2.r_2$ an *intersection*. This statement means that A defines $A.r$ to include every principal who is a member of both $B_1.r_1$ and $B_2.r_2$. This represents partial delegations from A to B_1 and to B_2 .

A *role expression* is a principal, a role, a linked role, or an intersection. We say that each policy statement *defines* the role $A.r$. Given a set \mathcal{P} of policy statements, we define the following: $\text{Principals}(\mathcal{P})$ is the set of principals in \mathcal{P} , $\text{Names}(\mathcal{P})$ is the set of role names in \mathcal{P} , and $\text{Roles}(\mathcal{P}) = \{A.r \mid A \in \text{Principals}(\mathcal{P}), r \in \text{Names}(\mathcal{P})\}$.

The language *SRT* simplifies RT_0 in that type-4 statements in *SRT* allow the intersection of only two roles; in RT_0 , the intersection may contain k components, each can be a principal, a role, or a linked role. RT_0 statements using such intersections can be equivalently expressed in *SRT* by introducing new intermediate roles and additional statements. This simplification helps simplify the proofs in this paper. When studying containment queries, we consider the following sub-languages of *SRT*: *BRT* (for Basic RT_0), which has only type-1 and 2 statements, *LRT* (for Linking RT_0), which has type-1, 2, and 3 statements, and *NRT* (for iNtersection RT_0), which has type-1, 2, and 4 statements.

The four types of statements in *SRT* cover the most common delegation relationships in existing TM languages such as the *RT* family of Role-based Trust-management languages [Li et al. 2003b; Li et al. 2002; Li and Mitchell 2003b], SPKI/SDSI [Ellison et al. 1999; Clarke et al. 2001], and KeyNote [Blaze et al. 1999]. *SRT* captures the basic forms of delegation relationships in the *RT* framework, though it does not cover manifold roles. The sub-language *LRT* can be viewed as a simplified yet expressively equivalent version of SDSI. SDSI allows long linked names, which, as observed in [Li et al. 2003b], can be broken up by introducing new role names and statements. With the exception of thresholds, the delegation relationships (though, not the s-expression-based representation of permission) in SPKI's 5-tuples, can be captured by using type-1 statements and a restricted form of type-2 statements. A SPKI 5-tuple in which A delegates a permission r to B can be represented as $A.r \longleftarrow B$. A SPKI 5-tuple in which A delegates r to B and allows B to further delegate r can be represented as two *SRT* statements: $A.r \longleftarrow B$ and $A.r \longleftarrow B.r$. Threshold structures in SPKI can be implemented by intersections, as noted in [Li 2000]. Similar analogies can be drawn for KeyNote.

Although *SRT* is limited in that role names are constants, extending role names in *SRT* to have internal structures does not change the nature of security analysis. As we will see, security analysis is mostly affected by the structure of the delegation relationships. We believe that many results and techniques developed for *SRT* can be carried over to more expressive languages, e.g., RT_1 [Li et al. 2002], which adds to RT_0 the ability to have

parameterized roles, RT_1^C [Li and Mitchell 2003a], which further adds constraints to RT_1 , and, to a certain extent, SPKI/SDSI and KeyNote.

The security analysis problem in *SRT* is quite nontrivial. Semantics and inferencing of SDSI, which is essentially the sub-language *LRT*, has been extensively studied [Abadi 1998; Clarke et al. 2001; Halpern and van der Meyden 2001; Jha and Reps 2002; Li 2000; Li et al. 2003b]. Some of these works only consider answering queries in a fixed state. Some considered universal queries where no restriction is placed on how the state may grow [Abadi 1998; Halpern and van der Meyden 2001]. The most interesting case of security analysis, i.e., answering queries with restrictions placed on state changes, has not been studied in these previous works.

2.2 Semantics of the TM Language

We give a formal characterization of the meaning of a set \mathcal{P} of policy statements by translating each policy statement into a datalog clause. (Datalog is a restricted form of logic programming (LP) with variables, predicates, and constants, but without function symbols.) We call the resulting program the *semantic program* of \mathcal{P} . We use the LP-based approach to define semantics because we will formulate safety computation rules by using a similar approach, and the LP-based approach generalizes easily to the case in which role names contain parameters (see [Li et al. 2002]).

DEFINITION 2 SEMANTIC PROGRAM. Given a set \mathcal{P} of policy statements, the *semantic program*, $SP(\mathcal{P})$, of \mathcal{P} , has one ternary predicate m . Intuitively, $m(A, r, D)$ represents that D is a member of the role $A.r$. $SP(\mathcal{P})$ is constructed as follows. (Symbols that start with “?” represent logical variables.)

For each $A.r \leftarrow D$ in \mathcal{P} , add

$$m(A, r, D) \tag{m1}$$

For each $A.r \leftarrow B.r_1$ in \mathcal{P} , add

$$m(A, r, ?Z) :- m(B, r_1, ?Z) \tag{m2}$$

For each $A.r \leftarrow A.r_1.r_2$ in \mathcal{P} , add

$$m(A, r, ?Z) :- m(A, r_1, ?Y), m(?Y, r_2, ?Z) \tag{m3}$$

For each $A.r \leftarrow B_1.r_1 \cap B_2.r_2$ in \mathcal{P} , add

$$m(A, r, ?Z) :- m(B_1, r_1, ?Z), m(B_2, r_2, ?Z) \tag{m4}$$

A datalog program is a set of datalog clauses. Given a datalog program, \mathcal{DP} , its semantics can be defined through several equivalent approaches. The model-theoretic approach views \mathcal{DP} as a set of first-order sentences and uses the minimal Herbrand model as the semantics. We write $SP(\mathcal{P}) \models m(X, u, Z)$ when $m(X, u, Z)$ is in the minimal Herbrand model of $SP(\mathcal{P})$. This semantics corresponds exactly to the set-theoretic semantics of RT_0 in [Li et al. 2003b].

We now summarize a standard fixpoint characterization of the minimal Herbrand model, which we will use in the proofs in this paper. Given a datalog program \mathcal{DP} , let \mathcal{DP}^{inst} be the ground instantiation of \mathcal{DP} using constants in \mathcal{DP} ; the *immediate consequence operator*, $T_{\mathcal{DP}}$, is defined as follows. Given a set of ground logical atoms K , $T_{\mathcal{DP}}(K)$ consists of all logical atoms, a , such that $a :- b_1, \dots, b_n \in \mathcal{DP}^{inst}$, where $n \geq 0$, and either $n = 0$ or $b_j \in K$ for $1 \leq j \leq n$. The least fixpoint of $T_{\mathcal{DP}}$, denoted by $T_{\mathcal{DP}} \uparrow^\omega$, can be constructed as follows. Define $T_{\mathcal{DP}} \uparrow^0 = \emptyset$ and $T_{\mathcal{DP}} \uparrow^{i+1} = T_{\mathcal{DP}}(T_{\mathcal{DP}} \uparrow^i)$ for $i \geq 0$. This defines an increasing sequence of subsets of a finite set. Thus there exists an N such

that $T_{\mathcal{DP}}(T_{\mathcal{DP}}\uparrow^N) = T_{\mathcal{DP}}\uparrow^N$. $T_{\mathcal{DP}}\uparrow^N$ is easily shown to be $T_{\mathcal{DP}}\uparrow^\omega$, which is identical to the minimal Herbrand model of \mathcal{DP} [Lloyd 1987]; therefore, $SP(\mathcal{P}) \models m(X, u, Z)$ if and only if $m(X, u, Z) \in T_{SP(\mathcal{P})}\uparrow^\omega$.

It has been shown that the minimal Herbrand model of \mathcal{DP} can be computed in time linear in the size of \mathcal{DP}^{inst} [Dowling and Gallier 1984]. If the total size of \mathcal{DP} is M , then there are $O(M)$ constants in \mathcal{DP} . Assuming that the number of variables in each clause is bounded by a constant, v , the number of instances of each clause is therefore $O(M^v)$, so the size of \mathcal{DP}^{inst} is $O(M^{v+1})$. Thus, the worst-case complexity of evaluating $SP(\mathcal{P})$ is $O(|\mathcal{P}|^3)$, since $|SP(\mathcal{P})| = O(|\mathcal{P}|)$ and each rule in $SP(\mathcal{P})$ has at most two variables.

2.3 Queries

In this paper, we consider the following three forms of query \mathcal{Q} :

— *Form-1:* $A.r \sqsupseteq \{D_1, \dots, D_n\}$

Intuitively, this means that all the principals D_1, \dots, D_n are members of $A.r$. Formally, $\mathcal{P} \vdash A.r \sqsupseteq \{D_1, \dots, D_n\}$ if and only if $\{Z \mid SP(\mathcal{P}) \models m(A, r, Z)\} \supseteq \{D_1, \dots, D_n\}$.

— *Form-2:* $\{D_1, \dots, D_n\} \sqsupseteq A.r$

Intuitively, this means that the member set of $A.r$ is bounded by the given set of principals. Formally, $\mathcal{P} \vdash A.r \sqsupseteq \{D_1, \dots, D_n\}$ if and only if $\{D_1, \dots, D_n\} \supseteq \{Z \mid SP(\mathcal{P}) \models m(A, r, Z)\}$.

— *Form-3 (role inclusion):* $X.u \sqsupseteq A.r$

Intuitively, this means that all the members of $A.r$ are also members of $X.u$. Formally, $\mathcal{P} \vdash X.u \sqsupseteq A.r$ if and only if $\{Z \mid SP(\mathcal{P}) \models m(X, u, Z)\} \supseteq \{Z \mid SP(\mathcal{P}) \models m(A, r, Z)\}$.

Simple safety and simple availability queries are special cases of form-1 queries; they have the form $A.r \sqsupseteq \{D\}$. Simple safety queries are existential, and simple availability queries are universal. Bounded safety queries are universal form-2 queries. Containment queries are universal form-3 queries.

2.4 Restriction Rules on State Changes

Before discussing how we model restrictions on changes in policy state, we consider one motivating scenario. Suppose that the users within an organization control certain principals, and that these principals delegate partial control to principals outside the organization. By using security analysis, the organization can ensure that these delegations do not violate desired security properties, which are specified by a collection of security analysis queries. In this usage of the analysis, roles defined by principals within the organization can be viewed as unchanging, since the analysis will be repeated before any future candidate change is made to those roles. Roles defined by principals outside the organization, however, may change in arbitrary ways, since they are beyond the organization's control.

To model such control over roles, we use restriction rules of the form $\mathcal{R} = (\mathcal{G}_{\mathcal{R}}, \mathcal{S}_{\mathcal{R}})$, which consist of a pair of finite sets of roles. (In the rest of the paper we drop the subscripts from \mathcal{G} and \mathcal{S} , as \mathcal{R} is clear from context.)

—Roles in \mathcal{G} are called *growth-restricted* (or *g-restricted*); no policy statements defining these roles can be added. Roles not in \mathcal{G} are called *growth-unrestricted* (or *g-unrestricted*).

—Roles in \mathcal{S} are called *shrink-restricted* (or *s-restricted*); policy statements defining these roles cannot be removed. Roles not in \mathcal{S} are called *shrink-unrestricted* (or *s-unrestricted*).

One example of \mathcal{R} is (\emptyset, \emptyset) , under which every role is g/s-unrestricted, i.e., both g-unrestricted and s-unrestricted. Under this \mathcal{R} , all three forms of queries have trivial answers. Another example is $\mathcal{R} = (\emptyset, \text{Roles}(\mathcal{P}))$, i.e., every role may grow with out restriction, and no statement defining roles in $\text{Roles}(\mathcal{P})$ can be removed. This models the case of having incomplete knowledge of a fixed policy state. A third example, corresponding to the scenario discussed above, is $\mathcal{R} = (\mathcal{G}, \mathcal{S})$, where $\mathcal{G} = \mathcal{S} = \{X.u \mid X \in \{X_1, \dots, X_k\}, u \in \text{Names}(\mathcal{P})\}$, i.e., X_1, \dots, X_k are trusted (controlled); therefore, every role $X.u$ such that $X \in \{X_1, \dots, X_k\}$ is restricted, all other roles are unrestricted. If a principal X does not appear in \mathcal{R} , then for every role name r , by definition $X.r$ is *g/s-unrestricted*. This models the fact that the roles of unknown principals may be defined arbitrarily.

We allow some roles controlled by one principal to be g-restricted while other roles controlled by the same principal to be g-unrestricted. This provides more flexibility than simply identifying principals as trusted and untrusted, and permits one in practice to perform security analysis only when changing certain roles. Similarly, we allow a role to be both g-restricted and s-unrestricted, which has the effect of making a safety check necessary when modifying the definition of the role only if adding a new statement.

The above kinds of restrictions are *static* in the sense that whether or not a state-change step is allowed by \mathcal{R} does not depend on the current state. A dynamic restriction could, for instance, have $B.r_2$ be g-restricted if B is a member of $A.r_1$, which depends on the current state. Security analysis with dynamic restrictions is potentially interesting future work.

2.5 An Example

EXAMPLE 1. The system administrator of a company, SA, controls access to some resource, which we abstractly denote by SA.access. The company policy is the following: managers always have access to the resource, managers can delegate the access to other principals, but only to employees of the company. HR is trusted for defining employees and managers. The state \mathcal{P} consists of the following statements:

```
SA.access ← SA.manager
SA.access ← SA.delegatedAccess ∩ HR.employee
SA.manager ← HR.manager
SA.delegatedAccess ← SA.manager.access
HR.employee ← HR.manager
HR.employee ← HR.programmer
HR.manager ← Alice
HR.programmer ← Bob
HR.programmer ← Carl
Alice.access ← Bob
```

Given the above \mathcal{P} , Alice and Bob have access, Carl does not. One possible restriction rule has $\mathcal{G} = \{ \text{SA.access}, \text{SA.manager}, \text{SA.delegatedAccess}, \text{HR.employee} \}$ and $\mathcal{S} = \{ \text{SA.access}, \text{SA.manager}, \text{SA.delegatedAccess}, \text{HR.employee}, \text{HR.manager} \}$. We now list some example queries, together with the answers:

. A simply safety query:
Is “SA.access \sqsupseteq {Eve}” possible? (Yes.)

- . A simple availability query:
Is “SA.access \sqsubseteq {Alice}” necessary? (Yes.)
- . A bounded safety query:
Is “{Alice, Bob} \sqsubseteq SA.access” necessary. (No.)
- . A containment query:
Is “HR.employee \sqsubseteq SA.access” necessary? (Yes.)

3. ANSWERING FORM-1 AND FORM-2 QUERIES

SRT and its sub-languages are all monotonic in the sense that more statements will derive more role memberships (i.e., logical atoms of the form $m(A, r, D)$). This follows from the fact that the semantic program is a positive logic program. Form-1 queries are monotonic; given a form-1 query Q , if $\mathcal{P} \vdash Q$, then for every \mathcal{P}' such that $\mathcal{P} \subseteq \mathcal{P}'$, $\mathcal{P}' \vdash Q$. Form-2 queries are anti-monotonic; given a form-2 query Q , if $\mathcal{P} \vdash Q$, then for every $\mathcal{P}' \subseteq \mathcal{P}$, $\mathcal{P}' \vdash Q$.

Intuitively, universal form-1 (simple availability) queries and existential form-2 queries can be answered by considering the lower bound of role memberships. A role’s lower bound is the set of principals that are members of the role in every reachable state. Because \mathcal{R} is static, there exists a minimal state that is reachable from \mathcal{P} and \mathcal{R} , which is obtained from \mathcal{P} by removing all statements defining s-unrestricted roles. We denote this state by $\mathcal{P}|_{\mathcal{R}}$. Clearly, $\mathcal{P}|_{\mathcal{R}}$ is reachable; furthermore, $\mathcal{P}|_{\mathcal{R}} \subseteq \mathcal{P}'$ for every reachable \mathcal{P}' . Since *SRT* is monotonic, one can compute the lower bound by computing the role memberships in $\mathcal{P}|_{\mathcal{R}}$.

Similarly, existential form-1 (simple safety) queries and universal form-2 (bounded safety) queries can be answered by computing an “upper bound” of role memberships. The upper bound of a role is the set of principals that could become a member of the role in some reachable state. Intuitively, such bounds can be computed by considering a “maximal reachable state”. However, such a “state” is not well-defined since it would contain an infinite set of policy statements, and we only allow a state to contain a finite set of policy statements. We will show that one can simulate the “maximal reachable state” by a finite state and derive correct answers.

3.1 The Lower Bound

DEFINITION 3 THE LOWER BOUND PROGRAM. Given \mathcal{P} and \mathcal{R} , the *lower bound program* for them, $LB(\mathcal{P}, \mathcal{R})$, is constructed as follows:

For each $A.r \leftarrow D$ in $\mathcal{P}|_{\mathcal{R}}$, add
 $lb(A, r, D)$ (b1)

For each $A.r \leftarrow B.r_1$ in $\mathcal{P}|_{\mathcal{R}}$, add
 $lb(A, r, ?Z) :- lb(B, r_1, ?Z)$ (b2)

For each $A.r \leftarrow A.r_1.r_2$ in $\mathcal{P}|_{\mathcal{R}}$, add
 $lb(A, r, ?Z) :- lb(A, r_1, ?Y), lb(?Y, r_2, ?Z)$ (b3)

For each $A.r \leftarrow B_1.r_1 \cap B_2.r_2$ in $\mathcal{P}|_{\mathcal{R}}$, add
 $lb(A, r, ?Z) :- lb(B_1, r_1, ?Z), lb(B_2, r_2, ?Z)$ (b4)

The worst-case complexity of evaluating the lower bound program is $O(|\mathcal{P}|^3)$.

Observe that the above lower bound program is essentially the same as the semantic program for the minimal state $\mathcal{P}|_{\mathcal{R}}$. They differ in that anywhere $LB(\mathcal{P}, \mathcal{R})$ uses the

predicate lb , $SP(\mathcal{P}|_{\mathcal{R}})$ uses the predicate m . Therefore, we have the following fact.

FACT 3.1. $LB(\mathcal{P}, \mathcal{R}) \models lb(A, r, D)$ if and only if $SP(\mathcal{P}|_{\mathcal{R}}) \models m(A, r, D)$.

PROPOSITION 3.2. $LB(\mathcal{P}, \mathcal{R}) \models lb(A, r, D)$ if and only if for every \mathcal{P}' such that $\mathcal{P} \xrightarrow{*}_{\mathcal{R}} \mathcal{P}'$, $SP(\mathcal{P}') \models m(A, r, D)$.

PROOF. The “only if” part: If $LB(\mathcal{P}, \mathcal{R}) \models lb(A, r, D)$, then from Fact 3.1, $SP(\mathcal{P}|_{\mathcal{R}}) \models m(A, r, D)$. For every \mathcal{P}' that is reachable, $\mathcal{P}|_{\mathcal{R}} \subseteq \mathcal{P}'$. Furthermore, the language *SRT* is monotonic; therefore, $SP(\mathcal{P}') \models m(A, r, D)$.

The “if” part: if for every reachable \mathcal{P}' , $SP(\mathcal{P}') \models m(A, r, D)$, then $SP(\mathcal{P}|_{\mathcal{R}}) \models m(A, r, D)$, since $\mathcal{P}|_{\mathcal{R}}$ is reachable. From Fact 3.1, $LB(\mathcal{P}, \mathcal{R}) \models lb(A, r, D)$. \square

Proposition 3.2 means that the lower bound program can be used to answer universal form-1 queries and existential form-2 queries. We have not found an intuitive security meaning of existential form-2 queries, but include answering method for them here for completeness.

COROLLARY 3.3. Given \mathcal{P} and \mathcal{R} , a form-1 query $A.r \sqsupseteq \{D_1, \dots, D_n\}$ is necessary if and only if $LB(\mathcal{P}, \mathcal{R}) \models lb(A, r, D_i)$ for every i , $1 \leq i \leq n$.

COROLLARY 3.4. Given \mathcal{P} and \mathcal{R} , a form-2 query $\{D_1, \dots, D_n\} \sqsupseteq A.r$ is possible if and only if $\{D_1, \dots, D_n\} \sqsupseteq \{Z \mid LB(\mathcal{P}, \mathcal{R}) \models lb(A, r, Z)\}$.

PROOF. For the “if” part, we must show that if $\{D_1, \dots, D_n\} \sqsupseteq \{Z \mid LB(\mathcal{P}, \mathcal{R}) \models lb(A, r, Z)\}$, then there exists a reachable \mathcal{P}' such that each D satisfying $\mathcal{P}' \models m(A, r, D)$ also satisfies $D \in \{D_1, \dots, D_n\}$. It is easily seen by using Fact 3.1 that $\mathcal{P}|_{\mathcal{R}}$ is such a \mathcal{P}' .

The “only if” part follows from Proposition 3.2 as follows. Suppose there exists Z such that $LB(\mathcal{P}, \mathcal{R}) \models lb(A, r, Z)$ and $Z \notin \{D_1, \dots, D_n\}$. By Proposition 3.2, for every reachable \mathcal{P}' , $SP(\mathcal{P}') \models m(A, r, Z)$; therefore, the query is not possible. \square

Consider Example 1. The simple availability query “is SA.access \sqsupseteq {Alice} necessary” is true when SA.access, SA.manager, and HR.manager are s-restricted, since then the statements “SA.access \leftarrow SA.manager”, “SA.manager \leftarrow HR.manager”, and “HR.manager \leftarrow Alice” exist in the minimal state. On the other hand, it is not necessary that Bob has access, even when SA.delegatedAccess, HR.employee, and HR.programmer are also s-restricted, since Alice could remove her statement “Alice.access \leftarrow Bob”.

3.2 The Upper Bound

To compute the upper bound of roles, we introduce the following notion: A role is *g-unbounded* if for every principal Z , there exists a reachable \mathcal{P}' such that $SP(\mathcal{P}') \models m(A, r, Z)$. In other words, $A.r$ could have every principal as its member. A g-unrestricted role is clearly g-unbounded. A g-restricted role may also be g-unbounded, as it may be defined to include a g-unbounded role.

The following fact about g-unbounded roles says that one only needs to consider one principal that does not occur in \mathcal{P} (instead of every principal) to determine whether a role is g-unbounded.

FACT 3.5. Given \mathcal{P}, \mathcal{R} , a role $A.r$, and a principal E that does not occur in \mathcal{P} , $A.r$ is g-unbounded if and only if there exists a reachable state \mathcal{P}' such that $SP(\mathcal{P}') \models m(A, r, E)$.

See Appendix A.1 for the proof. We now show how to compute the upper bound, which simulates an infinite state.

DEFINITION 4 THE UPPER BOUND PROGRAM. Given $\mathcal{P}, \mathcal{R} = (\mathcal{G}, \mathcal{S})$, their upper bound program, $UB(\mathcal{P}, \mathcal{R})$, is constructed as follows. (\top is a special principal symbol not occurring in \mathcal{P}, \mathcal{R} , or any query \mathcal{Q} .)

Add
 $ub(\top, ?r, ?Z)$ (u)
 For each $A.r \in \text{Roles}(\mathcal{P}) - \mathcal{G}$, add
 $ub(A, r, ?Z)$ (u0)
 For each $A.r \leftarrow D$ in \mathcal{P} , add
 $ub(A, r, D)$ (u1)
 For each $A.r \leftarrow B.r_1$ in \mathcal{P} , add
 $ub(A, r, ?Z) :- ub(B, r_1, ?Z)$ (u2)
 For each $A.r \leftarrow A.r_1.r_2$ in \mathcal{P} , add
 $ub(A, r, ?Z) :- ub(A, r_1, ?Y), ub(?Y, r_2, ?Z)$ (u3)
 For each $A.r \leftarrow B_1.r_1 \cap B_2.r_2$ in \mathcal{P} , add
 $ub(A, r, ?Z) :- ub(B_1, r_1, ?Z), ub(B_2, r_2, ?Z)$ (u4)

The computational complexity for evaluating $UB(\mathcal{P}, \mathcal{R})$ is $O(|\mathcal{P}|^3)$. Note that $\text{Roles}(\mathcal{P})$ has $O(|\mathcal{P}|^2)$ elements, since there are $O(|\mathcal{P}|)$ principals and $O(|\mathcal{P}|)$ role names in \mathcal{P} . Therefore, there are $O(N^2)$ instance rules of (u0); however, each such rule has only one variable.

PROPOSITION 3.6. *Given any $\mathcal{P}, \mathcal{R} = (\mathcal{G}, \mathcal{S})$, $A.r \in \text{Roles}(\mathcal{P})$, and $Z \in \text{Principals}(\mathcal{P}) \cup \{\top\}$, $UB(\mathcal{P}, \mathcal{R}) \models ub(A, r, Z)$ if and only if there exists \mathcal{P}' such that $\mathcal{P} \xrightarrow{*}_{\mathcal{R}} \mathcal{P}'$ and $SP(\mathcal{P}') \models m(A, r, Z)$.*

See Appendix A.1 for the proof. From Fact 3.5 and Proposition 3.6, we have the following.

COROLLARY 3.7. *A role $A.r$ is g-unbounded if and only if $UB(\mathcal{P}, \mathcal{R}) \models ub(A, r, \top)$.*

COROLLARY 3.8. *Given \mathcal{P} and $\mathcal{R} = (\mathcal{G}, \mathcal{S})$, a form-1 query $A.r \sqsupseteq \{D_1, \dots, D_n\}$ is possible if and only if one of the following three conditions hold: (1) $A.r \notin \mathcal{G}$, (2) $UB(\mathcal{P}, \mathcal{R}) \models ub(A, r, \top)$, or (3) $UB(\mathcal{P}, \mathcal{R}) \models ub(A, r, D_i)$ for every $i, 1 \leq i \leq n$.*

PROOF. When $A.r \notin \text{Roles}(\mathcal{P})$, the second and the third condition will not hold; the query is possible if and only if $A.r$ is g-unrestricted, i.e., $A.r \notin \mathcal{G}$. When $A.r \in \text{Roles}(\mathcal{P})$, the first condition implies the second condition. Condition (2) or (3) both imply that the query is possible. If none of the three conditions holds, the query is not possible. Condition (2) is needed to deal with the case that some of the D_i 's in the query do not occur in \mathcal{P} . \square

COROLLARY 3.9. *Given \mathcal{P} and $\mathcal{R} = (\mathcal{G}, \mathcal{S})$, a form-2 query $\{D_1, \dots, D_n\} \sqsupseteq A.r$ is necessary if and only if $A.r \in \mathcal{G}$ and $\{D_1, \dots, D_n\} \supseteq \{Z \mid UB(\mathcal{P}, \mathcal{R}) \models ub(A, r, Z)\}$.*

Consider Example 1 again and observe that the policy is not safe according to either the simple safety query or the bounded safety query. One reason is that the role HR.manager is g-unrestricted, meaning that new managers may be added. Another reason is that the role HR.programmer is g-unrestricted; therefore, new programmers may be added and access may be delegated to them. However, if the company knows that Eve is an enemy, then the company probably will not hire Eve as a manager or a programmer.

In fact, simple safety is quite unnatural: to use it effectively, one has to be able to identify the principals that should never have access, the number of such principals could be arbitrary large. Bounded safety is also unnatural, one does not know, for example, who in the future the company will hire as a manager. A more natural policy is to ensure that, for example, only employees of the company are allowed to access the resource. This can be done by using form-3 queries.

4. CONTAINMENT ANALYSIS: ANSWERING UNIVERSAL FORM-3 QUERIES

Form-3 queries (i.e., role inclusion queries) are neither monotonic nor anti-monotonic. Given a form-3 query $X.u \sqsupseteq Z.w$ and three states $\mathcal{P}' \subseteq \mathcal{P} \subseteq \mathcal{P}''$, it is possible that $\mathcal{P} \vdash Q$, but both $\mathcal{P}' \not\vdash Q$ and $\mathcal{P}'' \not\vdash Q$. As a result, the approach taken with form-1 and form-2 queries is not applicable here. We cannot simply look at a specific minimal (or maximal) state and answer the query.

In this paper, we restrict our attention to universal role inclusion queries. We have not found meaningful readings of existential role inclusion queries in terms of security properties. We say that a role $X.u$ *contains* another role $A.r$ if $X.u \sqsupseteq A.r$ is necessary, i.e., $X.u$ includes $A.r$ in every reachable state. And we call the problem of answering containment queries *containment analysis*.

The case that one of $X.u$ and $A.r$ is not in $\text{Roles}(\mathcal{P})$ is uninteresting. If $A.r \notin \text{Roles}(\mathcal{P})$, then $X.u$ contains $A.r$ if and only if $A.r$ is g-restricted. If $A.r \in \text{Roles}(\mathcal{P})$ and $X.u \notin \text{Roles}(\mathcal{P})$, then $X.u$ contains $A.r$ if and only if $A.r$ has an upper bound that is empty. In the rest of this section, we only consider the case that both $X.u$ and $A.r$ are in $\text{Roles}(\mathcal{P})$.

Intuitively, there are two cases in which a role $X.u$ contains a role $A.r$. The first case is that this containment is *forced* by the statements that are in \mathcal{P} . For example, if a statement $X.u \leftarrow A.r$ exists and cannot be removed, then $X.u$ contains $A.r$. A containment may be forced by a chain of credentials. Forced containment can be computed similarly to role memberships.

In the second case, containment is caused by the nonexistence of statements in \mathcal{P} . In the extreme case, if $A.r$ has no definition and is g-restricted, then $A.r$ is contained in every role, since the member set of $A.r$ is empty in every reachable state. To compute this kind of containment, observe that a g-restricted role $A.r$ is contained in another role $X.u$ if every definition of $A.r$ is contained in $X.u$. If $A.r$ has no definition at all, then it is contained in every role. However, a straightforward translation of this into a positive logic program does not work. Consider the following example: $\mathcal{P} = \{A.r \leftarrow B.r_1, A.r \leftarrow D, B.r_1 \leftarrow A.r, X.u \leftarrow D\}$ and \mathcal{R} is such that $\mathcal{G} = \{A.r, B.r_1\}$ and $\mathcal{S} = \{A.r, B.r_1, X.u\}$. In any \mathcal{P}' that is \mathcal{R} -reachable from \mathcal{P} , the member sets of $A.r$ and $B.r_1$ are always $\{D\}$, and so both roles are contained by $X.u$. A straightforward positive logic program cannot derive this, since $X.u$ contains $A.r$ only if it contains $B.r_1$ and vice versa. As a result, neither containment relationship will be in the minimal model. To deal with this problem, we take the approach to prove non-containment using the minimal model of a logic program, and derive containment using negation-as-failure. Intuitively, $X.u$ contains $A.r$ unless we can find a witness entity E that is a member of $A.r$ in some state but not a member of $X.u$ in the same state.

Intuitively, containment queries that have the flavor of availability should be proven by forced containment. That a manager always has access to a resource should be due to a credential chain forcing this. In Example 1, SA.access contains HR.manager as

long as SA.access and SA.manager are s-restricted. On the other hand, policy statements are unlikely to force everyone who has access to a resource to be an employee; the orientation of the forced containment does not naturally correspond to this practical dependency. In Example 1, HR.employee contains SA.access as long as SA.access and SA.manager are g-restricted and HR.employee is s-restricted. This is because, as long as no new rule defining SA.access or SA.mamnager is added, any member of SA.access is either a member of HR.manager or a member of HR.employee; if furthermore, the statement “HR.employee \leftarrow HR.manager” cannot be removed, then HR.employee contains SA.access.

4.1 Containment Analysis in *BRT* is in PTIME

Recall that the language *BRT* has only type-1 and type-2 statements. In this section, we give algorithm to do containment analysis in *BRT* efficiently.

DEFINITION 5. (The Role Containment Program for *BRT*) Given a *BRT* state \mathcal{P} and \mathcal{R} , the role containment program, $BCP(\mathcal{P}, \mathcal{R})$, includes the lower bound program $LB(\mathcal{P}, \mathcal{R})$ in Definition 3. In addition, it defines two predicates: $fc/4$ and $nc/4$. An atom $fc(X, u, Z, w)$ means that $X.u$ is forced to contain $Z.w$. An atom $nc(X, u, Z, w)$ means that $X.u$ does not contain $Z.w$. The program $BCP(\mathcal{P}, \mathcal{R})$ is derived from $LB(\mathcal{P}, \mathcal{R})$ as follows.

$$\begin{aligned}
& \text{Add} \\
& \quad fc(?X, ?u, ?X, ?u) \quad (c) \\
& \text{For each } A.r \leftarrow B.r_1 \text{ in } \mathcal{P}|_{\mathcal{R}}, \text{ add} \\
& \quad fc(A, r, ?Z, ?w) :- fc(B, r_1, ?Z, ?w) \quad (c2) \\
& \text{For each } A.r \in \text{Roles}(\mathcal{P}) - \mathcal{G}, \text{ add} \\
& \quad nc(?X, ?u, A, r) :- \sim fc(?X, ?u, A, r) \quad (n0) \\
& \text{For each } A.r \in \mathcal{G}, \text{ do the following:} \\
& \quad \text{For each } A.r \leftarrow D \text{ in } \mathcal{P}, \text{ add} \\
& \quad \quad nc(?X, ?u, A, r) :- \sim fc(?X, ?u, A, r), \\
& \quad \quad \quad \sim lb(?X, ?u, D) \quad (n1) \\
& \quad \text{For each } A.r \leftarrow B.r_1 \text{ in } \mathcal{P}, \text{ add} \\
& \quad \quad nc(?X, ?u, A, r) :- \sim fc(?X, ?u, A, r), \\
& \quad \quad \quad nc(?X, ?u, B, r_1) \quad (n2)
\end{aligned}$$

Rules (c) and (c2) are straightforward. The intuition behind (n0) is that for $X.u$ to contain a g-unrestricted role $A.r$, $X.u$ has to be forced to contain $A.r$, since arbitrary new members may be added to $A.r$. The intuition behind (n1) is that, since $A.r$ contains D , if $X.u$'s lower bound does not contain D , and $X.u$ is not forced to contain $A.r$, then $X.u$ does not contain $A.r$. The “ $\sim fc$ ” part is needed, since it may be the case that $A.r \leftarrow D$ can be removed and $X.u \leftarrow A.r$ exists and cannot be removed, in which case D may not be in $X.u$'s lower bound. Rule (n2) means that $X.u$ does not contain $A.r$ if it does not contain $B.r_1$ and is not forced to contain $A.r$.

We now discuss the semantics of the logic program $BCP(\mathcal{P}, \mathcal{R})$, which uses negation-as-failure, but in a stratified manner. Given a logic program \mathcal{DP} , a predicate p (directly) depends on another predicate q if p is defined using q in the body. A predicate p negatively depends on q if $\sim q$ (the negation of q) is used to define p . For example, in $BCP(\mathcal{P}, \mathcal{R})$, fc

depends on itself, nc depends on itself and negatively depends on fc and lb . A program is *stratified* if the predicates defined in it can be classified into strata such that each predicate depends only on predicates in the same or lower strata and negatively depends only on predicates in lower strata. A program without negation is trivially stratified, as no predicate depends negatively on any predicate at all. The program $BCP(\mathcal{P}, \mathcal{R})$ is also stratified. Predicates in the first stratum are lb and fc , and the only predicate in the second stratum is nc .

Most commonly accepted semantics for logic programming with negation-as-failure agree on stratified programs. Given a stratified datalog program \mathcal{DP} , let $\mathcal{DP}_1 \cup \dots \cup \mathcal{DP}_s$ be a partition of \mathcal{DP}^{Inst} such that \mathcal{DP}_j consists of clauses defining predicates in the j 'th stratum; we call $\mathcal{DP}_1 \cup \dots \cup \mathcal{DP}_s$ a stratification of \mathcal{DP}^{Inst} . The semantics is obtained by first computing the minimal Herbrand model of \mathcal{DP}_1 and then use this model to determine truthfulness of negative literals in \mathcal{DP}_2 while computing a fixpoint for $\mathcal{DP}_1 \cup \mathcal{DP}_2$, and so on. Formally, we define an operator Φ , which is parameterized by a ground logic program \mathcal{DP}' and a set of ground atoms M . Given a set of ground logical atoms K , $\Phi_{\mathcal{DP}', M}(K)$ consists of all ground logic atoms, a , such that $a : -b_1, \dots, b_n, \sim b_{n+1}, \dots, \sim b_{n+m} \in \mathcal{DP}'$ and $b_i \in K$ and $b_{n+j} \notin M$. Given a logic program \mathcal{DP} and $\mathcal{DP}_1 \cup \dots \cup \mathcal{DP}_s$ a stratification of \mathcal{DP}^{Inst} , define $\Gamma_{\mathcal{DP}}^1$ to be $\Phi_{\mathcal{DP}_1, \emptyset} \uparrow^\omega$, i.e., the least fixpoint of $\Phi_{\mathcal{DP}_1, \emptyset}$. Define $\Gamma_{\mathcal{DP}}^{k+1}$ to be $\Phi_{\mathcal{DP}_1 \cup \dots \cup \mathcal{DP}_{k+1}, \Gamma_{\mathcal{DP}}^k} \uparrow^\omega$ for $1 \leq k \leq s-1$. Then the model of \mathcal{DP} is $\Gamma_{\mathcal{DP}}^s$. Each $\Gamma_{\mathcal{DP}}^i$ can be calculated in polynomial time, so the semantics of a stratified program can also be computed in polynomial time.

The following lemma says that the fc predicate in BCP is always sound for role containment, and it is complete when the second role is g-unrestricted.

LEMMA 4.1. *Given a BRT state \mathcal{P}, \mathcal{R} , two roles $X.u$ and $A.r$, if $BCP(\mathcal{P}, \mathcal{R}) \models fc(X, u, A, r)$, then $X.u$ contains $A.r$. If $X.u$ contains $A.r$ and $A.r$ is g-unrestricted, then $BCP(\mathcal{P}, \mathcal{R}) \models fc(X, u, A, r)$.*

See Appendix A.2 for the proof. The following proposition says that role containment in *BRT* can be answered by using the program $BCP(\mathcal{P}, \mathcal{R})$.

PROPOSITION 4.2. *Given a BRT state \mathcal{P}, \mathcal{R} , and two roles $X.u$ and $A.r$ in $\text{Roles}(\mathcal{P})$, $BCP(\mathcal{P}, \mathcal{R}) \models nc(X, u, A, r)$ if and only if $X.u$ does not contain $A.r$.*

See Appendix A.2 for the proof.

4.2 Containment Analysis in *NRT* is coNP-complete

NRT adds to *BRT* type-4 statements. Intersections in type-4 statements have the effect of conjunction. A role can be defined by multiple statements, which have the effect of disjunction. As a result, *NRT* can simulate formulas in propositional logic, and answering containment queries subsumes determining validity of propositional formulas, which is coNP-complete.

THEOREM 4.3. *Containment analysis in *NRT* is coNP-complete.*

See Appendix A.3 for the proof. The coNP-hard part is by reducing the monotone 3SAT problem, which is NP-complete, to the complement of containment analysis in *NRT*. The reduction does not use s-unrestricted roles, and it is straightforward to change the proof to not use g-unrestricted roles. This shows that containment analysis in *NRT* without shrinkable roles or without growable roles is still coNP-complete.

4.3 Containment Analysis in *LRT*

LRT adds to *BRT* type-3 statements, which have linked roles. We first prove that containment analysis in *LRT* without roles growing is coNP-complete. The key observation used in that proof was that linked roles, similar to intersection, can simulate logical conjunction. Recall that the semantic rule for type-3 statements, (m3), has a conjunction in the body, similar to that for type-4 statements, (m4).

THEOREM 4.4. *Containment analysis in LRT where all roles in $\text{Roles}(\mathcal{P})$ are g-restricted is coNP-complete.*

PROOF. The coNP-hardness part is proved by the proof in [Li et al. 2003a]. Since one can nondeterministically guess a subset \mathcal{P}' of \mathcal{P} and verify that the containment does not hold, the problem is clearly in coNP. \square

Of course, the case in which some roles in $\text{Roles}(\mathcal{P})$ are g-unrestricted is more interesting. The main result of this section is the following theorem.

THEOREM 4.5. *Containment analysis in LRT where all roles in $\text{Roles}(\mathcal{P})$ are s-restricted is PSPACE-complete.*

Theorem 4.5 enables us to establish the exact complexity bound for containment analysis in *LRT*.

THEOREM 4.6. *Containment analysis in LRT is PSPACE-complete.*

PROOF. Since the cases where all roles in $\text{Roles}(\mathcal{P})$ are s-restricted are special cases of the problems here, PSPACE-hardness follows immediately from Theorem 4.5.

Given a *LRT* state \mathcal{P} , a restriction rule $\mathcal{R} = (\mathcal{G}, \mathcal{S})$, and a containment query \mathcal{Q} , we use the following algorithm. For each set \mathcal{P}' such that $\mathcal{P}' \subseteq \mathcal{P}$ and \mathcal{P}' is \mathcal{R} -reachable from \mathcal{P} , solve the containment query for \mathcal{P}' and $\mathcal{R}' = (\mathcal{G}, \text{Roles}(\mathcal{P}))$. The algorithm answers “no” if there exists a \mathcal{P}' such that containment analysis answers no. Otherwise, the algorithm answers “yes”.

From Theorem 4.5, we know that containment analysis with \mathcal{P}' and \mathcal{R}' takes space polynomial in $n = |\mathcal{P}'| + |\mathcal{R}'|$. Since n is polynomial in $|\mathcal{P}| + |\mathcal{R}|$, this algorithm takes space polynomial in the size of the original input. If the containment does hold, then clearly this algorithm answers “yes”. If the containment does not hold, then there exists a reachable state \mathcal{P}_1 and a principal E such that $SP(\mathcal{P}_1) \models m(A, r, E)$ and $SP(\mathcal{P}') \not\models m(X, u, E)$. Consider $\mathcal{P}_0 = \mathcal{P} \cap \mathcal{P}_1$, \mathcal{P}_0 is reachable; furthermore, \mathcal{P}_1 is \mathcal{R}' -reachable from \mathcal{P}_0 ; therefore, our algorithm answers “no”. \square

In the rest of this section, we prove Theorem 4.5. Our proof uses the relationships between statements in *LRT* and string rewriting systems, which we introduce now.

4.3.1 A String Rewriting Semantics for *LRT*. In the rewriting semantics, each statement in *LRT* is viewed as a rewriting rule. For example, the statement $A.r \leftarrow B.r_1$ is viewed as the rewriting rule $A.r \mapsto B.r_1$. The alphabet consists of all principals and role names. We only consider rewriting over the set of *name strings*, i.e., strings that have the form of a principal followed by zero or more role names. When rewriting a name string ω using rewriting rules resulted from statements in *LRT*, the rewriting only occurs on the left most, and the resulting string is also a name string.

DEFINITION 6 REWRITING SEMANTICS FOR *LRT*: $\text{RS}[\mathcal{P}]$. Given a set \mathcal{P} of *LRT* statements, let $\text{RS}[\mathcal{P}]$ be the rewriting system generated by viewing each statement in \mathcal{P} as a rewriting rule. Given two name strings ω_1 and ω_2 , we write $\text{RS}[\mathcal{P}] \triangleright \omega_1 \rightarrow \omega_2$ if one can rewrite ω_1 into ω_2 in one step using a rewriting rule in $\text{RS}[\mathcal{P}]$. We write $\text{RS}[\mathcal{P}] \triangleright \omega_1 \xrightarrow{*} \omega_2$ if using rewriting rules in $\text{RS}[\mathcal{P}]$, one can rewrite ω_1 into ω_2 in zero or more steps.

We define $\text{str}_{\mathcal{P}}(A.r)$ to denote the set of strings such that $\text{RS}[\mathcal{P}] \triangleright Ar \xrightarrow{*} \omega$. We define $\text{mem}_{\mathcal{P}}(A.r)$ to denote the set of principals in $\text{str}_{\mathcal{P}}(A.r)$.

The following proposition is Proposition 2 in [Li and Mitchell 2003b].

PROPOSITION 4.7 EQUIVALENCE OF REWRITING AND LP SEMANTICS FOR *LRT*.

Given a set \mathcal{P} of *LRT* statements, $SP(\mathcal{P}) \models m(A, r, D)$ if and only if $\text{RS}[\mathcal{P}] \models Ar \xrightarrow{*} D$.

LRT is essentially SDSI. Jha and Reps [Jha and Reps 2002] pointed out that SDSI string rewriting systems correspond exactly to the class of string rewriting systems modelled using push-down systems. The same applies to the rewriting systems generated by *LRT* statements.

Pushdown systems (PDSs) are similar pushdown automata; however, unlike pushdown automata they do not have an input alphabet. Thus PDSs should not be viewed as language recognizers, but as mechanisms that specify possibly infinite-state transition systems.

A pushdown system is a triple (Π, Γ, Δ) , where Π is a finite set of control locations, Γ is a finite stack alphabet, and $\Delta \subseteq (\Pi \times \Gamma) \times (\Pi \times \Gamma^*)$ is a finite set of transition rules. If $((q, \gamma), (q', \omega)) \in \Delta$, then we write it as $\langle q, \gamma \rangle \hookrightarrow \langle q', \omega \rangle$.

A configuration of a PDS is a pair $\langle q, \omega \rangle$, where $q \in \Pi$ is a control location and $\omega \in \Gamma^*$ represents the stack contents (from the top of the stack to the bottom of a stack). We say that a configuration $\langle q, \omega_1 \rangle$ can directly reach a configuration $\langle q', \omega_2 \rangle$ if $\omega_1 = \gamma\omega'_1$, $\langle q, \gamma \rangle \hookrightarrow \langle q', \omega \rangle$, and $\omega_2 = \omega\omega'_2$. We then define reachability among configurations in the straightforward manner.

Esparza et al. [Esparza et al. 2000] showed that given a PDS (Π, Γ, Δ) , the set of all configurations reachable from a regular set of configurations recognized by an NFA \mathcal{A} is recognized by an NFA that has size polynomial in $n = |\Pi| + |\Gamma| + |\Delta| + |\mathcal{A}|$ and that can be constructed in time linear in n .

4.3.2 Containment analysis in *LRT* without shrinking is in PSPACE

DEFINITION 7 CHARACTERISTIC SET. Given a *LRT* state \mathcal{P} , a restriction rule \mathcal{R} , and a role $A.r$, the characteristic set of $A.r$, denoted by $\chi_{\mathcal{P}, \mathcal{R}, [(A).r]}$ is defined to be $\text{str}_{\mathcal{P}}(A.r) \cap \kappa[\mathcal{P}, \mathcal{R}]$, where $\kappa[\mathcal{P}, \mathcal{R}] = \{B r_1 r_1 \cdots r_k \mid B.r \text{ is g-unrestricted}\} \cup \text{Principals}(\mathcal{P})$.

LEMMA 4.8. Given a *LRT* state \mathcal{P} , a restriction rule \mathcal{R} , and a role $A.r$, $\chi_{\mathcal{P}, \mathcal{R}, [(A).r]}$ is recognized by an NFA that can be constructed in time polynomial in $|\mathcal{P}| + |\mathcal{R}|$.

PROOF. $\chi_{\mathcal{P}, \mathcal{R}, [(A).r]} = \text{str}_{\mathcal{P}}(A.r) \cap \kappa[\mathcal{P}, \mathcal{R}]$. Both $\text{str}_{\mathcal{P}}(A.r)$ and $\kappa[\mathcal{P}, \mathcal{R}]$ are recognized by NFA that can be constructed in time polynomial in $|\mathcal{P}| + |\mathcal{R}|$. \square

LEMMA 4.9. Given \mathcal{P} and \mathcal{R} , $X.u$ contains $A.r$ if and only if $\chi_{\mathcal{P}, \mathcal{R}}(X.u) \supseteq \chi_{\mathcal{P}, \mathcal{R}}(A.r)$.

PROOF. We first prove that if $\chi_{\mathcal{P}, \mathcal{R}}(X.u) \not\supseteq \chi_{\mathcal{P}, \mathcal{R}}(A.r)$ then $X.u$ does not contain $A.r$. If $\chi_{\mathcal{P}, \mathcal{R}}(X.u) \not\supseteq \chi_{\mathcal{P}, \mathcal{R}}(A.r)$, then there exists a string $\omega = B r_1 \cdots r_k$ such that $\omega \in \chi_{\mathcal{P}, \mathcal{R}}(A.r)$ and $\omega \notin \chi_{\mathcal{P}, \mathcal{R}}(X.u)$. Consider $\mathcal{P}' = \mathcal{P} \cup \{B.r_1 \leftarrow C_1, C_1.r_2 \leftarrow$

$C_2, \dots, C_{k-1}.r_k \leftarrow C_k\}$, where C_1, C_2, \dots, C_k do not occur in \mathcal{P} or \mathcal{R} . \mathcal{P}' is a reachable state. $\text{RS}[\mathcal{P}'] \triangleright A.r \xrightarrow{*} B.r_1 \dots r_k \xrightarrow{*} C_k$. From Proposition 4.7, $SP(\mathcal{P}') \models m(A, r, C_k)$. We now show that $SP(\mathcal{P}') \not\models m(X, u, C_k)$. Suppose, for the sake of contradiction, that $SP(\mathcal{P}') \models m(X, u, C_k)$. From Proposition 4.7, $\text{RS}[\mathcal{P}'] \triangleright X.u \xrightarrow{*} D$. Consider the rewriting sequence, the rule applied in the last step has to be $C_{k-1}.r_k \mapsto C_k$, since that is the only rule having a C_k on its right hand side. The rule applied in the second to last step has to be $C_{k-2}.r_{k-1} \mapsto C_{k-1}$, since that is the only rule having C_{k-1} on its right-hand side, and so on. Therefore, the rewriting sequence must contain in its middle a sequence rewriting from $X.u$ to $B.r_1 \dots r_k$. Further observe that the rules in \mathcal{P}' but not in \mathcal{P} cannot be applied in this middle sequence. Therefore, $\text{RS}[\mathcal{P}] \triangleright X.u \xrightarrow{*} \omega$. This contradicts the assumption that $\omega \notin \chi_{\mathcal{P}, \mathcal{R}}(X.u)$.

We then prove that if $X.u$ does not contain $A.r$ then $\chi_{\mathcal{P}, \mathcal{R}}(X.u) \not\supseteq \chi_{\mathcal{P}, \mathcal{R}}(A.r)$. If $X.u$ does not contain $A.r$, then there exists a reachable state \mathcal{P}' and a principal E such that $SP(\mathcal{P}') \models m(A, r, E)$ and $SP(\mathcal{P}') \not\models m(X, u, E)$. From Proposition 4.7, it follows that $\text{RS}[\mathcal{P}'] \triangleright A.r \xrightarrow{*} E$. Consider the rewriting sequence that rewrites $A.r$ into E ; let the sequence be $A.r \xrightarrow{*} \omega_1 \mapsto \omega_2 \xrightarrow{*} E$, such that the step from ω_1 to ω_2 is the first step that uses a statement not in \mathcal{P} . Clearly, $\omega_1 \in \chi_{\mathcal{P}, \mathcal{R}}(A.r)$ and $\omega_1 \notin \chi_{\mathcal{P}, \mathcal{R}}(X.u)$. \square

From the above lemma, we know that whether a role contains another can be determined by checking containment among character sets.

THEOREM 4.10. *Containment analysis in LRT without shrinking is in PSPACE.*

PROOF. Follows directly from Lemma 4.9, Lemma 4.8, and the fact that determining containment of languages accepted by NFAs is in PSPACE. \square

4.3.3 Containment Analysis in LRT is PSPACE-hard. It is known that containment among regular languages that are represented using NFAs is PSPACE-complete, even when the alphabet has just two elements. (When regular languages are described using DFA, the complexity is NLOGSPACE-complete.) We consider empty-string free regular languages over $\{0, 1\}$. They can be represented by a form of right linear grammars such that each production rule has one of the following two forms: $N_1 ::= N_2b$, and $N_1 ::= b$, where N_1 and N_2 are nonterminals and $b \in \{0, 1\}$.

THEOREM 4.11. *Containment Analysis in LRT without shrinking is PSPACE-hard*

PROOF. Given two empty-string free regular languages over $\{0, 1\}$, L_1 and L_2 . They can be represented using one set G of production rules, but with different start symbols S_1 and S_2 . For any nonterminal N in G , we write $L_G(N)$ to denote the language generated by G when using N as the start symbol.

We now reduce the problem of determining whether $L_G(S_1) \supseteq L_G(S_2)$ to a containment analysis problem. We use two principals A, B , two role names u_0 and u_1 , and one role name r_N for each nonterminal N . We define \mathcal{P}_G to have the following statements. For each production $N ::= N'b$ in G , \mathcal{P}_G has $A.r_N \leftarrow A.r_{N'}.u_b$. For each production $N ::= b$, introduce a statement $A.r_N \leftarrow B.u_b$. We define \mathcal{R}_G to be such that all roles started with A are g/s-restricted, and all roles started with B are g-unrestricted.

It is not hard to see that a string $b_1b_2 \dots b_k \in L(N)$ if and only if $\text{RS}[\mathcal{P}_G] \triangleright A.r_N \xrightarrow{*} B.u_{b_1}.u_{b_2} \dots .u_{b_k}$.

One can further verify that there exists a one-to-one mapping from strings in $L_G(N)$ to $\chi_{\mathcal{P}_G, \mathcal{R}_G}(A.r_N)$ such that $L_G(N_1) \supseteq L_G(N_2)$ if and only if $\chi_{\mathcal{P}_G, \mathcal{R}_G}(A.r_{N_1}) \supseteq \chi_{\mathcal{P}_G, \mathcal{R}_G}(A.r_{N_2})$.

From Lemma 4.9, it follows that $L_G(N_1) \supseteq L_G(N_2)$ if and only if $A.r_{N_1}$ contains $A.r_{N_2}$. \square

4.4 Containment Analysis for *SRT* is in coNEXP

THEOREM 4.12. *Containment analysis in SRT is in coNEXP.*

See Appendix A.3 for the proof. The proof shows that if there exists a counter example to the containment query, then there exists a counter example of size at most exponential in the input size.

5. DISCUSSIONS AND RELATED WORK

We have shown that containment analysis is intractable in *NRT*, *LRT*, and *SRT*. This means that it is extremely unlikely that we will find an algorithm that is both sound and complete, and also has a worst-case polynomial time complexity. However, heuristic approaches are still possible. For example, it is not difficult to extend our LP-based approach for containment analysis in *BRT* to the case of *LRT* and *SRT*, such that containment relationships in Example 1 can be proved correctly. A possible approach is to use a sound but incomplete method and a complete but unsound method together to approximate the exact answer. Such a heuristic approach may be useful in practice, as it can give an exact answer in most cases. How to evaluate the effectiveness of such methods is interesting future work.

On the other hand, we have shown that in our TM model, simple safety queries can be solved efficiently. As discussed in Section 1, security analysis in the form of simple safety queries has been studied in the HRU model [Harrison et al. 1976], and shown to be undecidable there. In this section we study the relationships between the two models, arguing informally that the HRU model does not include our TM model as a special case, and showing that there is an intuitive reason why security analysis in our model is decidable. We also seek to clarify the relationship between how trusted users are modelled in the two approaches. After this discussion of related work in safety analysis, we go on to discuss related work in trust management.

5.1 Comparison with the HRU Access Matrix Model

In the HRU model [Harrison et al. 1976], a *protection system* has a finite set of rights and a finite set of commands. A *configuration* of a protection system is an access control matrix, with rows corresponding to subjects, and columns corresponding to objects; each cell in the matrix is a set of rights. A command takes the form of “if a list of conditions hold, execute a sequence of primitive operations.” Each condition tests whether a right exists in a cell in the matrix. There are six kinds of primitive operations: enter a right into a specific cell in the matrix, delete a right from a cell in the matrix, create a new subject, create a new object, destroy an existing subject, and destroy an existing object. A command may be parameterized, with parameters being subjects or objects. In [Harrison et al. 1976], Harrison et al. proved that for the HRU model, the safety question is undecidable, by showing that any Turing machine can be simulated by a protection system. For a fixed set of mono-operational commands, safety can be determined in time polynomial of the size of the access control matrix. However, if commands are a parameter to the problem, the

safety problem is NP-complete.

In our model, given a state \mathcal{P} , the minimal Herbrand model of $SP(\mathcal{P})$ is a set of ground logical atoms. An atom $m(A, r, D)$ means that D is a member of A 's r role. When A represents a resource, this can be viewed as D having the right r over A . Therefore, one can view principals as both subjects and objects and view role names as rights. This defines a straightforward mapping between the semantics of \mathcal{P} and an access matrix. If all we have are type-1 statements, then adding (or removing) $A.r \leftarrow D$ corresponds to adding (or removing) r to the cell on row D and column A . Adding a type-2 statement $A.r \leftarrow B.r_1$ can be viewed as adding a trigger program, which for each row D , use parameters A, B, D to execute the following command: “ $a2(x, y, z) \{ \text{if } r_1 \text{ in cell } (y, z), \text{ add } r \text{ to cell } (x, z) \}$ ”. Note that this trigger program needs to be executed whenever the matrix changes. For example, if after $A.r \leftarrow B.r_1$ is added, adding $B.r_1 \leftarrow E$ will need to result in r being added to the cell (A, E) . The statement $A.r \leftarrow B.r_1$ gives B the power to add things to A 's column, which represents a delegation. Similarly, adding a type-3 statement $A.r \leftarrow A.r_1.r_2$ can be viewed as adding a trigger program that executes the following command with parameters A, D, E for every D and E : “ $a3(x, y, z) \{ \text{if } r_1 \text{ in cell } (x, y), \text{ and } r_2 \text{ in cell } (y, z), \text{ add } r \text{ to cell } (x, z) \}$ ”. Adding type-4 statement can be viewed in a similar manner. It is not clear how to model removing a statement using this approach.

There might be other ways of encoding our TM model in the HRU access matrix model, but the above encoding seems quite natural. From it, we make the following observations.

It seems unlikely that the HRU model subsumes the TM model as a special case. First, in the TM model, creating and removing principals are implicit. A principal can be viewed as created if it is used. A principal is considered removed if no statement mentions it. One could view the matrix as having an infinite number of rows and columns; however, only finitely many cells are nonempty. Second, one step of change in the TM model corresponds to executing many (one for every object when adding a type-2 or 4 statement, or one for every pair of objects when adding a type-3 statement) simple commands in the HRU model. Third, triggers need to be used in order to achieve the effect of propagation. The last two are the main power of the TM model, and they do not exist in the HRU model.

That our TM model cannot subsume the HRU model is immediate from the complexity bounds. The underlying reason is that the HRU commands we use to partially simulate our TM model have fixed schemas, instead of being arbitrary programs. As a result, we can exploit the properties of these fixed schemas. This seems to be the main reason that safety analysis, or the even more powerful containment analysis, is decidable in our model, but not in the HRU model.

Handling Trusted Subjects. Intuitively, a specific protection system is “safe” if access to resources without concurrence of the owner is impossible. However, protection systems often allow the owner to share rights to the resources. In that sense, they are not safe; the HRU model uses a weaker notion of safety: a user should be able to tell whether what he is about to do can lead to the further leakage of that right to untrusted subjects. The following is quoted from [Harrison et al. 1976].

To avoid a trivial “unsafe” answer because s himself can confer generic right r , we should in most circumstances delete s itself from the matrix. It might also make sense to delete from the matrix any other “reliable” subjects who could grant r , but whom s “trusts” will not do so. It is only by using the hypothetical

safety test in this manner, with “reliable” subjects deleted, that the ability to test whether a right can be leaked has a useful meaning in terms of whether it is safe to grant a right to a subject.

Note that deleting a “reliable” subject from the matrix is stronger than stopping it from granting a right. Deleting a subject from the matrix will prevent the analysis from successfully simulating the execution of commands that check rights in the row or column corresponding the subject. However, it is inappropriate to ignore such commands: they may add undesirable rights and they may be initiated by “unreliable” subjects. In such cases, a system that is safe after the “reliable” subjects are removed is not safe in the actual system, even if “reliable” subjects do not initiate any command.

In our TM model, the restriction rule R represents the intuitive notion that certain principals are trusted. In practice, principals are controlled by users. When principals represent resources, the controller is the subject who controls access to the resource. When principals represent public keys, the controller is the user who knows the private key.

5.2 Related Work in Trust Management

To our knowledge, no prior work investigates security analysis for trust management systems in the sense of verifying security properties that consider state changes in which (parametric) restrictions are placed on allowed changes. In [Chander et al. 2001], a state transition model is used for comparing the expressive power of different access control mechanisms such as access control lists and trust management. There, security analysis is not the purpose. The language SRT is closely related to SDSI, whose semantics and evaluation has been the subject of many previous works [Abadi 1998; Clarke et al. 2001; Halpern and van der Meyden 2001; Jha and Reps 2002; Li 2000; Li et al. 2003b]. One main difference our work has is that we consider restricted state changes. We now list some similarities. The semantic approach we use is very similar to the semantics in [Halpern and van der Meyden 2001]. Both [Abadi 1998] and [Halpern and van der Meyden 2001] consider role inclusion queries in addition to membership queries. In some sense, they try to answer queries that hold when arbitrary new statements could be added, i.e., every role is g -unrestricted and s -restricted; the case that some roles are g -restricted is not considered. In [Jha and Reps 2002], evaluating queries given a set of SDSI statements is reduced to model checking pushdown systems; there, only a fixed set of SDSI statements is considered, which are encoded as transition rules in the automata. Other works [Clarke et al. 2001; Li 2000; Li et al. 2003b] do not handle role inclusion queries or consider restricted state changes.

6. CONCLUSION

Trust management systems such as RT allow independent principals to delegate partial authority over resources. While this is useful in many situations, delegation also raises the possibility of unanticipated and undesirable access. If Alice delegates access to her friend Bob, how can she be sure that Bob does not give permissions to her enemy Carol? We address this question by studying several forms of safety and availability properties, including general containment queries that capture both safety and availability.

Although the trust management primitives we consider are more expressive than some aspects of the HRU model [Harrison et al. 1976], our main results show that persistence of nontrivial safety and availability properties may be algorithmically tractable. Specifi-

cally, form-1 queries and form-2 queries, both involving containment between a role and a fixed set of principals, can be answered using datalog programs that run in polynomial time. For general role inclusion queries, we look at several cases involving different policy sub-languages. For *BRT*, which only allows membership and delegation policy statements, containment for all reachable states is computable by a stratified datalog program with negation in polynomial time. For *NRT*, which is *BRT* plus intersection, the problem becomes coNP-complete. Intuitively, the reason is that multiple statements about a role represent disjunction, while intersection of roles provides a corresponding form of conjunction. For *LRT*, which is *BRT* plus role linking, role containment for all reachable policy states is PSPACE-complete. For *SRT*, which includes role linking, the problem remains decidable; our current upper bound is coNEXP (or double-exponential time) and lower bound is PSPACE-hard.

We believe that security analysis is a critical problem for trust management. While combining policy statements from independent principals has practical appeal, the flexibility of distributed policy comes at a price. An individual or organization that owns a resource no longer has a direct way to determine who may be able to access the resource in the future. The key to providing assurance to trust management users is to develop security analysis methods. The present paper identifies and solves certain security analysis problems, but additional work remains. Exact complexity bound for containment analysis in *SRT* is still open. Although containment analysis has no efficient algorithm in the worst case, there may be tractable subcases or useful heuristics. We also leave open for future work the consequences of more intricate restriction on policy changes. For example, it may be useful to impose restrictions that depend on the current policy, possibly formulated as policy invariants in some specification language.

Acknowledgement

This work is supported by DARPA through SPAWAR contracts N66001-00-C-8015 and N66001-01-C-8005. It is also supported by DOD MURI “Semantics Consistency in Information Exchange” as ONR Grant N00014-97-1-0505 and by DOD University Research Initiative (URI) program administered by the Office of Naval Research under Grant N00014-01-1-0795.

We thank Fred Schneider and Somesh Jha for discussions that lead to our study of safety and availability properties in Trust Management. We also thank the anonymous reviewers for their encouragement.

A. PROOFS

A.1 Proofs of Fact 3.5 and Proposition 3.6

Fact 3.5: *Given \mathcal{P} , \mathcal{R} , a role $A.r$, and a principal E that does not occur in \mathcal{P} , $A.r$ is g-unbounded if and only if there exists a reachable state \mathcal{P}' such that $SP(\mathcal{P}') \models m(A, r, E)$.*

PROOF. The “only if” part follows from the definition of g-unbounded roles.

In the “if” part, because *SRT* is monotonic, we can assume without loss of generality that \mathcal{P}' is derived from \mathcal{P} by adding some statements; let $\mathcal{P}' = \mathcal{P} \cup \mathcal{P}_1$. Given any principal Z , one can replace with Z all occurrence of E in the bodies of statements in \mathcal{P}_1 , obtaining a new set of statements, \mathcal{P}_2 . Let $\mathcal{P}'' = \mathcal{P}' \cup \mathcal{P}_2$. \mathcal{P}'' is reachable from \mathcal{P} because it modifies the definitions of the same roles as does \mathcal{P}' . We show that $SP(\mathcal{P}'') \models m(A, r, Z)$ by using induction on i to show that for all $A.r$, if $m(A, r, E) \in T_{SP(\mathcal{P}')} \uparrow^i$,

then $SP(\mathcal{P}'') \models m(A, r, Z)$. The basis is trivially satisfied because $T_{SP(\mathcal{P}')} \uparrow^0 = \emptyset$. In the step, $m(A, r, E) \in T_{SP(\mathcal{P}')} \uparrow^{i+1}$. This must be due to one of the four rules in $SP(\mathcal{P}')$, (m1), (m2), (m3), or (m4), which gives us the four following cases:

Case (m1): $A.r \leftarrow E \in \mathcal{P}'$. By construction of \mathcal{P}'' , $A.r \leftarrow Z \in \mathcal{P}''$. $SP(\mathcal{P}'') \models m(A, r, Z)$ follows from (m1).

Case (m2): $A.r \leftarrow B.r_1 \in \mathcal{P}'$ and $m(B, r_1, E) \in T_{SP(\mathcal{P}')} \uparrow^i$. The induction hypothesis now gives us $SP(\mathcal{P}'') \models m(B, r_1, Z)$, from which $SP(\mathcal{P}'') \models m(A, r, Z)$ follows by (m2).

Case (m3): $A.r \leftarrow A.r_1.r_2 \in \mathcal{P}'$ and $m(A, r_1, B), m(B, r_2, E) \in T_{SP(\mathcal{P}')} \uparrow^i$ for some B . The induction hypothesis now gives us $SP(\mathcal{P}'') \models m(B, r_2, Z)$. From $m(A, r_1, B) \in T_{SP(\mathcal{P}')} \uparrow^i$, we have $SP(\mathcal{P}') \models m(A, r_1, B)$, which gives us $SP(\mathcal{P}'') \models m(A, r_1, B)$ by monotonicity of SRT . We now have $SP(\mathcal{P}'') \models m(A, r, Z)$ by (m3).

Case (m4): $A.r \leftarrow B_1.r_1 \cap B_2.r_2 \in \mathcal{P}'$ and $m(B_1, r_1, E), m(B_2, r_2, E) \in T_{SP(\mathcal{P}')} \uparrow^i$. This case proceeds similarly to case (m2) above. \square

Proposition 3.6: *Given any $\mathcal{P}, \mathcal{R} = (\mathcal{G}, \mathcal{S})$, $A.r \in \text{Roles}(\mathcal{P})$, and $Z \in \text{Principals}(\mathcal{P}) \cup \{\top\}$, $UB(\mathcal{P}, \mathcal{R}) \models ub(A, r, Z)$ if and only if there exists \mathcal{P}' such that $\mathcal{P} \xrightarrow{*}_{\mathcal{R}} \mathcal{P}'$ and $SP(\mathcal{P}') \models m(A, r, Z)$.*

PROOF. The “only if” part (Soundness): If $UB(\mathcal{P}, \mathcal{R}) \models ub(A, r, Z)$, consider $\mathcal{P}' = \mathcal{P} \cup \{X.u \leftarrow Z \mid X.u \in \text{Roles}(\mathcal{P}) - \mathcal{G}\}$. We show by induction on i that if $ub(A, r, Z) \in T_{UB(\mathcal{P}, \mathcal{R})} \uparrow^i$, then $SP(\mathcal{P}') \models m(A, r, Z)$. The basis is trivial. In the step, $ub(A, r, Z) \in T_{UB(\mathcal{P}, \mathcal{R})} \uparrow^{i+1}$, one of the rules in $UB(\mathcal{P}, \mathcal{R})$ is used to derive this. Case (u) is impossible, as $A \neq \top$. Case (u0): $A.r \in \text{Roles}(\mathcal{P}) - \mathcal{G}$, by construction of \mathcal{P}' , $A.r \leftarrow Z \in \mathcal{P}'$. So $SP(\mathcal{P}') \models m(A, r, Z)$ follows immediately by (m1). Case (u1): $A.r \leftarrow Z \in \mathcal{P} \subseteq \mathcal{P}'$. In this case, $SP(\mathcal{P}') \models m(A, r, Z)$ also follows immediately by (m1).

Case (u2): $A.r \leftarrow B.r_1 \in \mathcal{P} \subseteq \mathcal{P}'$ and $ub(B, r_1, Z) \in T_{UB(\mathcal{P}, \mathcal{R})} \uparrow^i$. The induction assumption now gives us $SP(\mathcal{P}') \models m(B, r_1, Z)$, from which $SP(\mathcal{P}') \models m(A, r, Z)$ follows by (m2).

Case (u3): $A.r \leftarrow A.r_1.r_2 \in \mathcal{P} \subseteq \mathcal{P}'$ and $ub(A, r_1, B), ub(B, r_2, Z) \in T_{UB(\mathcal{P}, \mathcal{R})} \uparrow^i$ for some B . The induction assumption now gives us $SP(\mathcal{P}') \models m(A, r_1, B), m(B, r_2, Z)$, from which $SP(\mathcal{P}') \models m(A, r, Z)$ follows by (m3).

Case (u4): $A.r \leftarrow B_1.r_1 \cap B_2.r_2 \in \mathcal{P} \subseteq \mathcal{P}'$ and $ub(B_1, r_1, Z), ub(B_2, r_2, Z) \in T_{UB(\mathcal{P}, \mathcal{R})} \uparrow^i$. The induction assumption now gives us $SP(\mathcal{P}') \models m(B_1, r_1, Z), m(B_2, r_2, Z)$, from which $SP(\mathcal{P}') \models m(A, r, Z)$ follows by (m4).

The “if” part (Completeness): Suppose that there exists a reachable state \mathcal{P}' such that $SP(\mathcal{P}') \models m(A, r, Z)$. If $A.r \notin \mathcal{G}$, then $UB(\mathcal{P}, \mathcal{R}) \models ub(A, r, Z)$ from (u0). For the case in which $A.r \in \mathcal{G}$, we use induction on i to show that if $m(A, r, Z) \in T_{SP(\mathcal{P}')} \uparrow^i$, then $UB(\mathcal{P}, \mathcal{R}) \models ub(A, r, Z)$. The basis is trivial. In the step, there are four cases. Case (m1): $A.r \leftarrow Z \in \mathcal{P}'$. From $A.r \in \mathcal{G}$, we have $A.r \leftarrow Z \in \mathcal{P}$. So $UB(\mathcal{P}, \mathcal{R}) \models ub(A, r, Z)$ follows by using (u1).

Case (m2): $A.r \leftarrow B.r_1 \in \mathcal{P}'$ and $m(B, r_1, Z) \in T_{SP(\mathcal{P}')} \uparrow^i$. The induction hypothesis gives us $UB(\mathcal{P}, \mathcal{R}) \models ub(B, r_1, Z)$, from which we obtain the desired $UB(\mathcal{P}, \mathcal{R}) \models ub(A, r, Z)$ by (u2).

Case (m3): $A.r \leftarrow A.r_1.r_2 \in \mathcal{P}'$ and $m(A, r_1, B), m(B, r_2, Z) \in T_{SP(\mathcal{P}')} \uparrow^i$ for some B . The induction hypothesis gives us $UB(\mathcal{P}, \mathcal{R}) \models ub(A, r_1, B), ub(B, r_2, Z)$, from which we obtain the desired $UB(\mathcal{P}, \mathcal{R}) \models ub(A, r, Z)$ by (u3).

Case (m4): $A.r \leftarrow B_1.r_1 \cap B_2.r_2 \in \mathcal{P}'$ and $m(B_1, r_1, Z), m(B_2, r_2, Z) \in T_{SP(\mathcal{P}')} \uparrow^i$. This case is similar to the ones above. \square

A.2 Proof of Lemma 4.1 and Proposition 4.2

We introduce the following terminology for the proof. The program $BCP(\mathcal{P}, \mathcal{R})$ has a stratification of two strata. Define BCP_1 to be the ground instantiation of clauses defining lb and fc in $BCP(\mathcal{P}, \mathcal{R})$, and BCP_2 to the ground instantiation of clauses defining nc . (We use BCP instead of $BCP(\mathcal{P}, \mathcal{R})$ for succinctness.) We write $BCP \models a$ if $a \in \Gamma_{BCP}^2$. When a is a ground instance of fc or lb , we write $BCP \models^i a$ if $a \in \Phi_{BCP_1, \emptyset} \uparrow^i$. When a is a ground instance of nc , we write $BCP \models^i a$ if $a \in \Phi_{BCP_1 \cup BCP_2, \Gamma_{BCP}^1} \uparrow^i$.

Lemma 4.1: *Given a BRT state \mathcal{P}, \mathcal{R} , two roles $X.u$ and $A.r$, if $BCP(\mathcal{P}, \mathcal{R}) \models fc(X, u, A, r)$, then $X.u$ contains $A.r$. If $X.u$ contains $A.r$ and $A.r$ is g-unrestricted, then $BCP(\mathcal{P}, \mathcal{R}) \models fc(X, u, A, r)$.*

PROOF. Soundness: If $BCP \models fc(X, u, A, r)$, then there exists an integer i such that $BCP \models^i fc(X, u, A, r)$. Induction on i . The basis is trivial, as $\Phi_{BCP_1, \emptyset} \uparrow^0 = \emptyset$. Consider the step; either c or $(c2)$ is used to deduce that $BCP \models^{i+1} fc(X, u, A, r)$. Case (c) : it must be that $X.u = A.r$, so it is trivial that $X.u$ contains $A.r$. Case $(c2)$: $X.u \leftarrow B.r_1 \in \mathcal{P}|_{\mathcal{R}}$ and $BCP \models^i fc(B, r_1, A, r)$. By induction hypothesis, $B.r_1$ contains $A.r$. Furthermore, $X.u \leftarrow B.r_1$ exists in every reachable state; therefore, $X.u$ contains $A.r$.

Completeness: Suppose $X.u$ contains $A.r$ and $A.r$ is g-unrestricted. Consider $\mathcal{P}' = \mathcal{P}|_{\mathcal{R}} \cup (A.r \leftarrow E)$, in which E does not occur in \mathcal{P} . Observe that $X.u$ includes $A.r$ is true, since \mathcal{P}' is reachable. Since $SP(\mathcal{P}') \models m(A, r, E)$, it must be that $m(X, u, E) \in T_{SP(\mathcal{P}')} \uparrow^i$ for some i . To complete the proof, we use induction on i to show that for each $Y.u$, if $m(Y, u, E) \in T_{SP(\mathcal{P}')} \uparrow^i$, then $BCP \models fc(Y, u, A, r)$. Basis is trivial. In the step, one of $(m1)$ and $(m2)$ is used to deduce that $m(Y, u, E) \in T_{SP(\mathcal{P}')} \uparrow^{i+1}$. Case $(m1)$: $Y.u \leftarrow E \in \mathcal{P}'$, it must be that $Y.u = A.r$, since E does not occur in \mathcal{P} . From (c) , $BCP \models fc(Y, u, A, r)$. Case $(m2)$: $Y.u \leftarrow Y_1.u_1 \in \mathcal{P}'$, and $m(Y_1, u_1, E) \in T_{SP(\mathcal{P}')} \uparrow^i$. By definition of \mathcal{P}' , $Y.u \leftarrow Y_1.u_1 \in \mathcal{P}|_{\mathcal{R}}$. From $(c2)$, $fc(Y, u, ?Z, ?w) :- fc(Y_1, u_1, ?Z, ?w) \in BCP$. By induction hypothesis, $BCP \models fc(Y_1, u_1, A, r)$, clearly $BCP \models fc(Y, u, A, r)$. \square

Before proving Proposition 4.2, we first prove two auxiliary lemmas. Readers may wish to read the main proof first and refer to the two lemmas when they needed. The following lemma is used to prove the soundness of $(n1)$.

LEMMA A.1. *Assume we are given \mathcal{P} in BRT, \mathcal{R} , two roles $X.u$ and $A.r$, and a principal D such that $SP(\mathcal{P}|_{\mathcal{R}}) \not\models m(X, u, D)$. Let $\mathcal{P}' = \mathcal{P}|_{\mathcal{R}} \cup \{A.r \leftarrow D\}$. If $SP(\mathcal{P}') \models m(X, u, D)$, then $BCP \models fc(X, u, A, r)$.*

PROOF. We use induction on i to prove that for any $Z.w$ such that $SP(\mathcal{P}|_{\mathcal{R}}) \not\models m(Z, w, D)$, if $m(Z, w, D) \in T_{SP(\mathcal{P}')} \uparrow^i$, then $BCP \models fc(Z, w, A, r)$.

The basis is trivial. In the step, one of $(m1)$ and $(m2)$ is used to derive $m(Z, w, D) \in T_{SP(\mathcal{P}')} \uparrow^{i+1}$. Case $(m1)$: $Z.w \leftarrow D \in \mathcal{P}'$. It must be that $Z.w = A.r$, since it cannot be that $Z.w \leftarrow D \in \mathcal{P}|_{\mathcal{R}}$. By (c) , $BCP \models fc(Z, w, A, r)$. Case $(m2)$: $Z.w \leftarrow Z_1.w_1 \in \mathcal{P}'$ and $m(Z_1, w_1, D) \in T_{SP(\mathcal{P}')} \uparrow^i$. It follows that $Z.w \leftarrow Z_1.w_1 \in \mathcal{P}|_{\mathcal{R}}$, by definition of \mathcal{P}' . And it follows that $SP(\mathcal{P}|_{\mathcal{R}}) \not\models m(Z_1, w_1, D)$, since otherwise $SP(\mathcal{P}|_{\mathcal{R}}) \models m(Z, w, D)$, which is contradictory. Now, by induction hypothesis, $BCP \models fc(Z_1, w_1, A, r)$, so the desired result holds by $(c2)$. \square

The following lemma says that (n2) is sound.

LEMMA A.2. *Assume we are given a BRT state \mathcal{P} , \mathcal{R} , and three roles $X.u$, $A.r$, $B.r_1$, such that $A.r \leftarrow B.r_1 \in \mathcal{P}$, $B\text{CP}(\mathcal{P}, \mathcal{R}) \not\models \text{fc}(X, u, A, r)$, and $X.u$ does not contain $B.r_1$. Then $X.u$ does not contain $A.r$.*

PROOF. Since $X.u$ does not contain $B.r_1$, there exists a reachable state \mathcal{P}' and a principal E such that $SP(\mathcal{P}') \models m(B, r_1, E)$ and $SP(\mathcal{P}') \not\models m(X, u, E)$. We now construct a \mathcal{P}'' such that $SP(\mathcal{P}'') \models m(A, r, E)$ and $SP(\mathcal{P}'') \not\models m(X, u, E)$. \mathcal{P}'' is obtained from \mathcal{P}' by first removing any $Z.w \leftarrow Z_1.w_1 \in \mathcal{P}' - \mathcal{P}|_{\mathcal{R}}$ such that $SP(\mathcal{P}') \not\models m(Z_1, w_1, E)$, and then adding $A.r \leftarrow B.r_1$. Clearly, \mathcal{P}'' is reachable. By induction on how $m(A, r, E)$ is proven in $SP(\mathcal{P}')$, it is easy to show that $SP(\mathcal{P}'') \models m(A, r, E)$.

To prove that $SP(\mathcal{P}'') \not\models m(X, u, E)$, we use induction on i to prove that for any $Z.w$ such that $SP(\mathcal{P}') \not\models m(Z, w, E)$, if $m(Z, w, E) \in T_{SP(\mathcal{P}'')} \uparrow^i$, then $B\text{CP} \models \text{fc}(Z, w, A, r)$. The basis is trivial. In the step, one of (m1) and (m2) is used to derive $m(Z, w, E) \in T_{SP(\mathcal{P}'')} \uparrow^{i+1}$. Case (m1): $Z.w \leftarrow E \in \mathcal{P}''$. This is impossible, as this means that $Z.w \leftarrow E \in \mathcal{P}'$, which is contradictory with $SP(\mathcal{P}') \not\models m(Z, w, E)$. Case (m2): $Z.w \leftarrow Z_1.w_1 \in \mathcal{P}''$ and $m(Z_1, w_1, E) \in T_{SP(\mathcal{P}'')} \uparrow^i$. By definition of \mathcal{P}'' , either $Z.w = A.r$ and $Z_1.w_1 = B.r_1$, or $Z.w \leftarrow Z_1.w_1 \in \mathcal{P}'$. In the former case, $\text{fc}(Z, w, A, r)$ follows from (c). In the latter case, it follows that $SP(\mathcal{P}') \not\models m(Z_1, w_1, E)$, from $SP(\mathcal{P}') \not\models m(Z, w, E)$, and, by induction hypothesis, that $B\text{CP} \models \text{fc}(Z_1, w_1, A, r)$. Now the desired result holds by (c2), provided we have $Z.w \leftarrow Z_1.w_1 \in \mathcal{P}|_{\mathcal{R}}$. This follows from the construction of \mathcal{P}'' and the case assumption that $m(Z_1, w_1, E) \in T_{SP(\mathcal{P}'')} \uparrow^i$. \square

Proposition 4.2: *Given a BRT state \mathcal{P} , \mathcal{R} , and two roles $X.u$ and $A.r$ in $\text{Roles}(\mathcal{P})$, $B\text{CP}(\mathcal{P}, \mathcal{R}) \models \text{nc}(X, u, A, r)$ if and only if $X.u$ does not contain $A.r$.*

PROOF. The “only if” part (Soundness): We use induction on i to show that if $B\text{CP} \models^i \text{nc}(X, u, A, r)$, then $X.u$ does not contain $A.r$. Basis is trivial. In the step, one of (n0), (n1), and (n2) is used to derive that $B\text{CP} \models^{i+1} \text{nc}(X, u, A, r)$. Case (n0): $A.r$ must be g-unrestricted, and $B\text{CP} \models \sim \text{fc}(X, u, A, r)$; therefore, $B\text{CP} \not\models \text{fc}(X, u, A, r)$. From Lemma 4.1, $X.u$ does not contain $A.r$. Case (n1): $A.r \leftarrow D \in \mathcal{P}$, $B\text{CP} \models \sim \text{lb}(X, u, D)$, and $B\text{CP} \models \sim \text{fc}(X, u, A, r)$. Then $SP(\mathcal{P}|_{\mathcal{R}}) \not\models m(X, u, D)$ by Fact 3.1. Let $\mathcal{P}' = \mathcal{P}|_{\mathcal{R}} \cup \{A.r \leftarrow D\}$. From Lemma A.1 it follows that $SP(\mathcal{P}') \not\models m(X, u, D)$; therefore $X.u$ does not contain $A.r$. Case (n2): $A.r \leftarrow B.r_1 \in \mathcal{P}$, $B\text{CP} \models^n \text{nc}(X, u, B, r_1)$, and $B\text{CP} \models \sim \text{fc}(X, u, A, r)$. By induction hypothesis, $X.u$ does not contain $B.r_1$; from Lemma A.2, $X.u$ does not contain $A.r$.

The “if” part (Completeness): If $X.u$ does not contain $A.r$, then we show that $B\text{CP} \models \text{nc}(X, u, A, r)$. When $A.r$ is g-unrestricted. From Lemma 4.1, $B\text{CP} \not\models \text{fc}(X, u, A, r)$, and so $B\text{CP} \models \sim \text{fc}(X, u, A, r)$. From (n0), $B\text{CP} \models \text{nc}(X, u, A, r)$. In the rest of the proof, we only need to consider the case that $A.r$ is g-restricted. If $X.u$ does not contain $A.r$, then there exists a reachable state \mathcal{P}' and a principal E such that $SP(\mathcal{P}') \models m(A, r, E)$ and $SP(\mathcal{P}') \not\models m(X, u, E)$. We use induction on i to show that if $m(A, r, E) \in T_{SP(\mathcal{P}')} \uparrow^i$, then $B\text{CP} \models \text{nc}(X, u, A, r)$. First observe that, from Lemma 4.1, it follows that $B\text{CP} \not\models \text{fc}(X, u, A, r)$, and so $B\text{CP} \models \sim \text{fc}(X, u, A, r)$. The basis is trivial. In the step, one of (m1) and (m2) is used to deduce that $m(A, r, E) \in T_{SP(\mathcal{P}')} \uparrow^{i+1}$. Case (m1): $A.r \leftarrow E \in \mathcal{P}'$, $A.r \leftarrow E$ must be in \mathcal{P} since $A.r$ is g-restricted. From Proposition 3.2 and $SP(\mathcal{P}') \not\models m(X, u, E)$, $B\text{CP} \not\models \text{lb}(X, u, E)$, and so $B\text{CP} \models \sim \text{lb}(X, u, E)$. From

(n1), $BCP(\mathcal{P}, \mathcal{R}) \models nc(X, u, A, r)$. Case (m2): $A.r \leftarrow B.r_1 \in \mathcal{P}'$ and $m(B, r_1, E) \in T_{SP(\mathcal{P}')} \uparrow^i$. Since $A.r$ is g-restricted, $A.r \leftarrow B.r_1 \in \mathcal{P}$. By induction hypothesis, $BCP(\mathcal{P}, \mathcal{R}) \models nc(X, u, B, r_1)$. Therefore, $BCP(\mathcal{P}, \mathcal{R}) \models nc(X, u, Z, w)$ by an instance of (n2). \square

A.3 Proofs of Theorems 4.3, 4.4, and 4.12

We first prove a lemma that will be used in establishing lower bounds on the complexity of containment analysis. The lemma says that if a containment does not hold, then there exists a counter-example state that only adds type-1 statements to \mathcal{P} and only uses role names in \mathcal{P} .

LEMMA A.3. *Given \mathcal{P} and \mathcal{R} , two roles $X.u$ and $A.r$ in $\text{Roles}(\mathcal{P})$, if $X.u$ does not contain $A.r$, then there exists a \mathcal{P}' such that $SP(\mathcal{P}') \models m(A, r, E)$, $SP(\mathcal{P}') \not\models m(X, u, E)$, $\mathcal{P}' - \mathcal{P}$ only has type-1 statements, and \mathcal{P}' only uses role names in \mathcal{P} .*

PROOF. If $X.u$ does not contain $A.r$, then there exists a \mathcal{P}' that $SP(\mathcal{P}') \models m(A, r, E)$ and $SP(\mathcal{P}') \not\models m(X, u, E)$. Given such a \mathcal{P}' , we first derive \mathcal{P}'' by replacing every statement $A.r \leftarrow e \in \mathcal{P}' - \mathcal{P}$, where e is a role, a linked role, or an intersection, with a set of statements $\{A.r \leftarrow Y \mid SP(\mathcal{P}') \models m(A, r, Y)\}$. Using induction, it is straightforward to show that the resulting state computes the exact same role memberships.

Now $\mathcal{P}'' - \mathcal{P}$ consists of only type-1 statements. From \mathcal{P}'' , we derive \mathcal{P}''' by removing all type-1 statements that uses role names (not roles) not appearing in \mathcal{P} . For example, a statement $A.v \leftarrow D$ in \mathcal{P}'' , where v does not appear in \mathcal{P} , will not be in \mathcal{P}''' . Using induction, it is straightforward to show that, for roles in $\text{Roles}(\mathcal{P})$, \mathcal{P}''' computes the exact same memberships as \mathcal{P}'' . Intuitively, $A.v \leftarrow D$ can affect members of roles in $\text{Roles}(\mathcal{P})$ unless some type-2, 3, or 4 statement refers to the role name v , which is impossible, since all type 2, 3, or 4 statements in \mathcal{P}' are in \mathcal{P} , and so do not use v . \square

Theorem 4.3: *Containment analysis in NRT is coNP-complete.*

PROOF. To show coNP-hardness, we reduce the monotone 3SAT problem to the complement of the universal containment problem in NRT. Monotone 3SAT is 3SAT with each clause containing either only positive literals or only negative literals; it is known to be NP-complete [Garey and Johnson 1979].

Given an instance of monotone 3SAT: $\phi = c_1 \wedge \dots \wedge c_\ell \wedge \overline{c_{\ell+1}} \wedge \dots \wedge \overline{c_n}$, in which c_1, \dots, c_ℓ are positive clauses and $\overline{c_{\ell+1}}, \dots, \overline{c_n}$ are negative clauses. Let p_1, \dots, p_s be all the propositional variables in ϕ . For each negative clause $\overline{c_k} = (\neg p_{k_1} \vee \neg p_{k_2} \vee \neg p_{k_3})$, define $d_k = (p_{k_1} \wedge p_{k_2} \wedge p_{k_3})$, then $\overline{c_k} \Leftrightarrow \neg d_k$. Then $\phi \Leftrightarrow c_1 \wedge \dots \wedge c_m \wedge \neg(d_{\ell+1} \vee \dots \vee d_n)$. The formula ϕ is satisfiable if and only if $\psi = (c_1 \wedge \dots \wedge c_\ell) \rightarrow (d_{\ell+1} \vee \dots \vee d_n)$ is not valid. We now construct \mathcal{P}, \mathcal{R} , with the goal that $A.d \sqsupseteq A.c$ is necessary if and only if ψ is valid. In the construction, we use the role $A.p_i$ to denote the propositional variable p_i , $A.c_j$ to denote the clause c_j , and $A.d_k$ to denote the clause d_k . Define $\mathcal{P} = \mathcal{P}_1 \cup \mathcal{P}_2 \cup \mathcal{P}_3 \cup \mathcal{P}_4$, in which

$$\mathcal{P}_1 = \{A.c \leftarrow A.c_1 \cap A.c'_1, A.c'_1 \leftarrow A.c_2 \cap A.c'_2, \dots, A.c'_{\ell-1} \leftarrow A.c_{\ell-1} \cap A.c_\ell\}.$$

$$\mathcal{P}_2 = \{A.c_j \leftarrow A.p_{j_1}, A.c_j \leftarrow A.p_{j_2}, A.c_j \leftarrow A.p_{j_3} \mid 1 \leq j \leq \ell, c_j = p_{j_1} \vee p_{j_2} \vee p_{j_3}\}$$

$$\mathcal{P}_3 = \{A.d \leftarrow A.d_k \mid \ell + 1 \leq k \leq n\}$$

$$\mathcal{P}_4 = \{A.d_k \leftarrow A.p_{k_1} \cap A.d'_k, A.d'_k \leftarrow A.p_{k_2} \cap A.p_{k_3} \mid \ell + 1 \leq k \leq n, d_k = p_{k_1} \wedge p_{k_2} \wedge p_{k_3}\}$$

Define R to be the restriction rule such that all the $A.p_i$'s are g-unrestricted and s-restricted, and all other roles are g/s-restricted.

We now show that $A.d \sqsupseteq A.c$ is *not* necessary if and only if ψ is *not* valid. First, the “only if” part: If $A.d \sqsupseteq A.c$ is not necessary, then there exists a reachable state \mathcal{P}' and a principal E such that $SP(\mathcal{P}') \models m(A, c, E)$ and $SP(\mathcal{P}') \not\models m(A, d, E)$. Consider the truth assignment I defined as follows, for every i such that $1 \leq i \leq s$, $I(p_i) = \text{true}$ if $SP(\mathcal{P}') \models m(A, p_i, E)$, and $I(p_i) = \text{false}$ otherwise. Then under I , $(c_1 \wedge \dots \wedge c_\ell)$ is true and $d_{\ell+1} \vee \dots \vee d_n$ is false; therefore ψ is not valid. The “if” part: If ψ is not valid, then there exists a truth assignment I such that $(c_1 \wedge \dots \wedge c_\ell)$ is true and $(d_{\ell+1} \vee \dots \vee d_n)$ is false. Consider $\mathcal{P}' = \mathcal{P} \cup \{A.p_i \leftarrow Z \mid 1 \leq i \leq s \wedge I(p_i) = \text{true}\}$. \mathcal{P}' is reachable, and $SP(\mathcal{P}') \models m(A, c, Z)$ and $SP(\mathcal{P}') \not\models m(A, d, Z)$.

We now show that containment analysis in *NRT* is in coNP. Given \mathcal{P} and \mathcal{R} , if $X.u$ does not contain $A.r$, then there exists a reachable state \mathcal{P}' and a principal E such that, $SP(\mathcal{P}') \models m(A, r, E)$ and $SP(\mathcal{P}') \not\models m(X, u, E)$. From Lemma A.3, we can assume, without loss of generality, that $\mathcal{P}' - \mathcal{P}$ consists of only type-1 statements and \mathcal{P}' uses the same role names. From \mathcal{P}' , we construct \mathcal{P}'' as follows, let $\mathcal{P}'' = \mathcal{P}' \cap (\mathcal{P} \cup \{Z.w \leftarrow E \in \mathcal{P}' \mid Z.w \in \text{Roles}(\mathcal{P})\})$. Clearly, $\mathcal{P}'' \subseteq \mathcal{P}'$ and \mathcal{P}'' is reachable. By induction on how $m(A, r, E)$ is proven in $SP(\mathcal{P}')$, it is easy to see that $SP(\mathcal{P}'') \models m(A, r, E)$. The size of \mathcal{P}'' is polynomial in \mathcal{P} . This means that if a containment does not hold, then there exists a short (polynomial in the size of the input program \mathcal{P}) counterproof such that one can check in polynomial time. This shows that the problem is in coNP. The method we use to construct the counter example \mathcal{P}'' also yields an exponential algorithm for determining containment. \square

Theorem 4.4: *Containment analysis in LRT where all roles in $\text{Roles}(\mathcal{P})$ are g-restricted is coNP-complete.*

PROOF. Since one can nondeterministically guess a subset \mathcal{P}' of \mathcal{P} and verify that the containment does not hold, the problem is clearly in coNP. To prove coNP-hardness, we reduce the monotone 3SAT problem to the complement of universal role containment in *LRT*, similar to that in the proof of Theorem 4.3. Given an instance ϕ of monotone 3SAT, we construct $\psi = (c_1 \wedge \dots \wedge c_\ell) \rightarrow (d_{\ell+1} \vee \dots \vee d_n)$ such that ϕ is satisfiable if and only if ψ is not valid.

We now construct \mathcal{P}, \mathcal{R} , such that $A.d \sqsupseteq A.c$ is necessary if and only if ψ is valid. Define \mathcal{P} to be $\mathcal{P}_1 \cup \mathcal{P}_2 \cup \mathcal{P}_3 \cup \mathcal{P}_4 \cup \mathcal{P}_5$, in which

$$\mathcal{P}_1 = \{A.c \leftarrow A.c'_1.c_1, A.c'_1 \leftarrow A.c'_2.c_2, \dots, A.c'_{\ell-2} \leftarrow A.c'_{\ell-1}.c_{\ell-1}, A.c'_{\ell-1} \leftarrow A.c_\ell\}$$

$$\mathcal{P}_2 = \{A.c_j \leftarrow A.p_{j_1}, A.c_j \leftarrow A.p_{j_2}, A.c_j \leftarrow A.p_{j_3} \mid 1 \leq j \leq \ell, c_j = p_{j_1} \vee p_{j_2} \vee p_{j_3}\}$$

$$\mathcal{P}_3 = \{A.d \leftarrow A.d_k \mid \ell + 1 \leq k \leq n\}$$

$$\mathcal{P}_4 = \{A.d_k \leftarrow A.d'_k.p_{k_1}, A.d'_k \leftarrow A.p_{k_2}.p_{k_3} \mid \ell + 1 \leq k \leq n, d_k = p_{k_1} \wedge p_{k_2} \wedge p_{k_3}\}$$

$$\mathcal{P}_5 = \{A.p_i \leftarrow A \mid 1 \leq i \leq s\}$$

Let R be the restriction rule such that all the $A.p_i$'s are g-restricted and s-unrestricted, and all other roles mentioned in \mathcal{P} are g/s-restricted.

In every reachable state, the definitions of some $A.p_i$'s are removed, which correspond to assigning false to some of the p_i 's. In every reachable state, $A.c$ and $A.d$ either includes only A or is empty. $A.c$ includes A if and only if the corresponding truth assignment makes $c_1 \wedge \dots \wedge c_m$ true, and $A.d$ includes A if and only if the corresponding truth assignment makes $(d_{m+1} \vee \dots \vee d_n)$ true. Therefore, $A.c$ contains $A.d$ if and only if ψ is valid. \square

Theorem 4.12: *Containment analysis in SRT is in coNEXP.*

PROOF. Given \mathcal{P} and \mathcal{R} , if a query $X.u \sqsupseteq A.r$ is not necessary, i.e., $X.u$ does not contain $A.r$, then there exists a reachable state \mathcal{P}' and a principal E such that $SP(\mathcal{P}') \models m(A, r, E)$ and $SP(\mathcal{P}') \not\models m(X, u, E)$. From Lemma A.3, we can assume, without loss of generality, that $\mathcal{P}' - \mathcal{P}$ consists of only type-1 statements and \mathcal{P}' uses the same role names as \mathcal{P} .

Given such a \mathcal{P}' and E , we show that one can construct another state \mathcal{P}'' that has size exponential in \mathcal{P} and $SP(\mathcal{P}'') \models m(A, r, E)$ and $SP(\mathcal{P}'') \not\models m(X, u, E)$. The way we construct \mathcal{P}'' is through collapsing equivalent principals in \mathcal{P}' into one, to be made precise as follows. Let $\text{SigRoles}(\mathcal{P}, \mathcal{P}', \mathcal{Q})$ be $\{X.u\} \cup \{A.r_1 \mid A.r \leftarrow A.r_1.r_2 \in \mathcal{P} \cap \mathcal{P}'\} \cup \{B_1.r_1, B_2.r_2 \mid A.r \leftarrow B_1.r_1 \cap B_2.r_2 \in \mathcal{P} \cap \mathcal{P}'\}$. Define a binary relation \equiv over the principals in \mathcal{P}' as follows: $Y_1 \equiv Y_2$ if one of the following two conditions are satisfied: (1) $Y_1 = Y_2$; (2) $Y_1, Y_2 \notin \text{Principals}(\mathcal{P})$ and for every role $Z.w \in \text{SigRoles}(\mathcal{P}, \mathcal{P}', \mathcal{Q})$, $SP(\mathcal{P}') \models m(Z, w, Y_1)$ if and only if $SP(\mathcal{P}') \models m(Z, w, Y_2)$. The relation \equiv is easily seen to be an equivalence relation. For each equivalence class, we pick one principal in it as a unique representative; for a given principal Y , we use $[Y]$ to denote the representative of the equivalence class of Y . We assume that $[E] = E$. \mathcal{P}'' is constructed from \mathcal{P}' as follows: for each statement, replace all the principals with their representatives, then remove duplicate statements.

Given \mathcal{P} that has size N , clearly $\text{SigRoles}(\mathcal{P}, \mathcal{P}', \mathcal{Q})$ has $O(N)$ roles. Therefore, there are in total $M = O(2^{O(N)})$ principals in \mathcal{P}' , these principals will result in $O(M^2N)$ new type-1 statements. Therefore, if a containment does not hold, there exists a counter-example state that has size exponential in \mathcal{P} . Once the state is guessed correctly, it can be verified in time polynomial in the size of the state. This shows that the problem is in coNEXP. An obvious algorithm that has double exponential time complexity is as follows: first collect $\text{SigRoles}(\mathcal{P}, \mathcal{P}, \mathcal{Q})$ from $X.u$ and all type-2 and 3 statements from \mathcal{P} , and add one principal for each subset of $\text{SigRoles}(\mathcal{P}, \mathcal{P}, \mathcal{Q})$, then enumerate all reachable sub-states to see whether a containment holds.

It remains to prove that our construction of \mathcal{P}'' works, i.e., that $SP(\mathcal{P}'') \models m(A, r, E)$ and $SP(\mathcal{P}'') \not\models m(X, u, E)$.

To prove $SP(\mathcal{P}'') \models m(A, r, E)$, we use induction to prove the following claim: For any role $Z.w$ in $\text{Roles}(\mathcal{P}')$ and Y in \mathcal{P}' , if $m(Z, w, Y) \in T_{SP(\mathcal{P}')} \uparrow^i$, then $SP(\mathcal{P}'') \models m([Z], w, [Y])$. The basis is trivial, since $T_{SP(\mathcal{P}')} \uparrow^0 = \emptyset$. Now consider the step. One of (m1), (m2), (m3), and (m4) is used to derive $m(Z, w, Y) \in T_{SP(\mathcal{P}')} \uparrow^{i+1}$. Case (m1): $Z.w \leftarrow Y \in \mathcal{P}'$. By construction of \mathcal{P}'' , $[Z].w \leftarrow [Y] \in \mathcal{P}''$; therefore, $SP(\mathcal{P}'') \models m([Z], w, [Y])$. In the next three cases, a type-2, 3, or 4 statement $A.r \leftarrow e$ exist in \mathcal{P}' . It must also exist in \mathcal{P} , since $\mathcal{P}' - \mathcal{P}$ only has type-1 statements; therefore, principals in $A.r \leftarrow e$ are each in their own equivalence class. The statement must also exist in \mathcal{P}'' , since the equivalence substitution for $A.r \leftarrow e$ will not change the statement. Case (m2): $Z.w \leftarrow Z_1.w_1 \in \mathcal{P}'$, $\mathcal{P}, \mathcal{P}''$ and $m(Z_1, w_1, Y) \in T_{SP(\mathcal{P}')} \uparrow^i$. From induction hypothesis, $SP(\mathcal{P}'') \models m([Z_1], w_1, [Y])$. It must be that $[Z_1] = Z_1$. The claim then follows from

(m2). Case (m3): $Z.w \leftarrow Z.w_1.w_2 \in \mathcal{P}', \mathcal{P}, \mathcal{P}''$ and $m(Z, w_1, F), m(F, w_2, Y) \in T_{SP(\mathcal{P}')} \uparrow^i$. It must be that $[Z] = Z$. By induction hypothesis, $SP(\mathcal{P}'') \models m([Z], w_1, [F])$, and $SP(\mathcal{P}'') \models m([F], w_2, [Y])$. The claim follows from (m3). Case (m4): $Z.w \leftarrow Z_1.w_1 \cap Z_2.w_2 \in \mathcal{P}', \mathcal{P}, \mathcal{P}''$ and $m(Z_1, w_1, Y), m(Z_2, w_2, Y) \in T_{SP(\mathcal{P}')} \uparrow^i$. This case is similar to (m2).

We now prove that $SP(\mathcal{P}'') \not\models m(X, u, E)$, by proving the following claim: for any role $Z.w \in \text{Roles}(\mathcal{P}'')$ and any principal Y in $\text{Principals}(\mathcal{P}'')$, if $m(Z, w, Y) \in T_{SP(\mathcal{P}'')} \uparrow^i$, then there exists Z', Y' such that $[Z'] = Z$ and $[Y'] = Y$ and $SP(\mathcal{P}') \models m(Z', w, Y')$. Given this claim, if $SP(\mathcal{P}'') \models m(X, u, E)$, then there exists X' and E' in $\text{Principals}(\mathcal{P}')$ such that $[X'] = X$, $[E'] = E$, and $SP(\mathcal{P}') \models m(X', u, E')$. Since $X \in \text{Principals}(\mathcal{P})$, it must be that $X' = X$. And by definition of \equiv , $[E'] = E$ means that E is also a member of $X.u$, giving us a contradiction with our assumption on \mathcal{P}' .

We now use induction to prove the claim. The basis is trivial, since $T_{SP(\mathcal{P}'')} \uparrow^0 = \emptyset$. Now consider the step. One of (m1), (m2), (m3), and (m4) is used to derive $m(Z, w, Y) \in T_{SP(\mathcal{P}'')} \uparrow^{i+1}$. Case (m1): $Z.w \leftarrow Y \in \mathcal{P}''$. By definition of \mathcal{P}'' , there exists $Z'.w \leftarrow Y' \in \mathcal{P}'$ such that $[Z'] = Z$ and $[Y'] = [Y]$. From this we have $SP(\mathcal{P}') \models m(Z', w, Y')$ by (m1). In the following three cases, a type 2, 3, or 4 statement $A.r \leftarrow e$ of \mathcal{P}'' is used; such a statement must be mapped from a type 2, 3, 4 statement in \mathcal{P}' . Since all such statements in \mathcal{P}' are also in \mathcal{P} and do not change in the mapping, $A.r \leftarrow e \in \mathcal{P} \cap \mathcal{P}'$. Case (m2): $Z.w \leftarrow Z_1.w_1 \in \mathcal{P}'', \mathcal{P}, \mathcal{P}'$ and $m(Z_1, w_1, Y) \in T_{SP(\mathcal{P}'')} \uparrow^i$. From induction hypothesis, $SP(\mathcal{P}') \models m(Z'_1, w_1, Y')$ and $[Z'_1] = Z_1$ and $[Y'] = Y$. Because $Z_1 \in \text{Principals}(\mathcal{P})$, it must be that $Z'_1 = Z_1$. The conclusion follows from (m2). Case (m3): $Z.w \leftarrow Z.w_1.w_2 \in \mathcal{P}'', \mathcal{P}, \mathcal{P}'$ and $m(Z, w_1, F), m(F, w_2, Y) \in T_{SP(\mathcal{P}'')} \uparrow^i$ for some principal F . By induction hypothesis, $SP(\mathcal{P}') \models m(Z, w_1, F'), m(F'', w_2, Y')$ and $[F'] = [F''] = F$. Since $Z.w_1 \in \text{SigRoles}(\mathcal{P}, \mathcal{P}', \mathcal{Q})$, by definition of \equiv applied to $F' \equiv F''$, $SP(\mathcal{P}') \models m(Z, w_1, F'')$. The claim follows from (m3). Case (m4): $Z.w \leftarrow Z_1.w_1 \cap Z_2.w_2 \in \mathcal{P}'', \mathcal{P}, \mathcal{P}'$ and $m(Z_1, w_1, Y), m(Z_2, w_2, Y) \in T_{SP(\mathcal{P}'')} \uparrow^i$. By induction hypothesis and the fact $Z_1, Z_2 \in \text{Principals}(\mathcal{P})$, $SP(\mathcal{P}') \models m(Z_1, w_1, Y'), m(Z_2, w_2, Y'')$ and $[Y'] = [Y''] = Y$. By definition of \equiv , $SP(\mathcal{P}') \models m(Z_2, w_2, Y')$. Therefore, $SP(\mathcal{P}') \models m(Z, w, Y')$. \square

Observe that in the proof, only roles in the body of type-3 and 4 statements need to be collected. This may be used to explain why containment in *BRT* is efficiently decidable.

REFERENCES

- ABADI, M. 1998. On SDSI's linked local name spaces. *Journal of Computer Security* 6, 1–2, 3–21.
- BLAZE, M., FEIGENBAUM, J., IOANNIDIS, J., AND KEROMYTIS, A. D. 1999. The KeyNote trust-management system, version 2. IETF RFC 2704.
- CHANDER, A., DEAN, D., AND MITCHELL, J. C. 2001. A state-transition model of trust management and access control. In *Proceedings of the 14th IEEE Computer Security Foundations Workshop*. IEEE Computer Society Press, 27–43.
- CLARKE, D., ELIEN, J.-E., ELLISON, C., FREDETTE, M., MORCOS, A., AND RIVEST, R. L. 2001. Certificate chain discovery in SPKI/SDSI. *Journal of Computer Security* 9, 4, 285–322.
- DOWLING, W. F. AND GALLIER, J. H. 1984. Linear-time algorithms for testing the satisfiability of propositional horn formulae. *Journal of Logic Programming* 1, 3, 267–284.
- ELLISON, C., FRANTZ, B., LAMPSON, B., RIVEST, R., THOMAS, B., AND YLONEN, T. 1999. SPKI certificate theory. IETF RFC 2693.
- ESPARZA, J., HANSEL, D., ROSSMANITH, P., AND SCHWOON, S. 2000. Efficient algorithms for model check-

- ing pushdown systems. In *Proceedings of the 12th International Conference on Computer Aided Verification (CAV 2000)*. LNCS, vol. 1855. Springer, 232–247.
- GAREY, M. R. AND JOHNSON, D. J. 1979. *Computers And Intractability: A Guide to the Theory of NP-Completeness*. W.H. Freeman and Company.
- HALPERN, J. AND VAN DER MEYDEN, R. 2001. A logic for SDSI's linked local name spaces. *Journal of Computer Security* 9, 1-2, 47–74.
- HARRISON, M. A., RUZZO, W. L., AND ULLMAN, J. D. 1976. Protection in operating systems. *Communications of the ACM* 19, 8 (Aug.), 461–471.
- JHA, S. AND REPS, T. 2002. Analysis of SPKI/SDSI certificates using model checking. In *Proceedings of the 15th IEEE Computer Security Foundations Workshop*. IEEE Computer Society Press, 129–144.
- LAMPSON, B. W. 1971. Protection. In *Proceedings of the 5th Princeton Conference on Information Sciences and Systems*. Reprinted in *ACM Operating Systems Review*, 8(1):18-24, Jan 1974.
- LI, N. 2000. Local names in SPKI/SDSI. In *Proceedings of the 13th IEEE Computer Security Foundations Workshop*. IEEE Computer Society Press, 2–15.
- LI, N. AND MITCHELL, J. C. 2003a. Datalog with constraints: A foundation for trust management languages. In *Proceedings of the Fifth International Symposium on Practical Aspects of Declarative Languages (PADL 2003)*. Springer, 58–73.
- LI, N. AND MITCHELL, J. C. 2003b. RT: A role-based trust-management framework. In *The Third DARPA Information Survivability Conference and Exposition (DISCEX III)*. IEEE Computer Society Press.
- LI, N., MITCHELL, J. C., AND WINSBOROUGH, W. H. 2002. Design of a role-based trust management framework. In *Proceedings of the 2002 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, 114–130.
- LI, N., WINSBOROUGH, W. H., AND MITCHELL, J. C. 2003a. Beyond proof-of-compliance: Safety and availability analysis in trust management. In *Proceedings of IEEE Symposium on Security and Privacy*. IEEE Computer Society Press. To appear.
- LI, N., WINSBOROUGH, W. H., AND MITCHELL, J. C. 2003b. Distributed credential chain discovery in trust management. *Journal of Computer Security* 11, 1 (Feb.), 35–86.
- LIPTON, R. J. AND SNYDER, L. 1977. A linear time algorithm for deciding subject security. *Journal of ACM* 24, 3, 455–464.
- LLOYD, J. W. 1987. *Foundations of Logic Programming, Second Edition*. Springer.
- SANDHU, R. S. 1988. The Schematic Protection Model: Its definition and analysis for acyclic attenuating systems. *Journal of ACM* 35, 2, 404–432.
- SANDHU, R. S. 1992. The typed access matrix model. In *Proceedings of the 1992 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, 122–136.