Comparing the Expressive Power of Access Control Models

Mahesh V. Tripunitara

tripunit@cerias.purdue.edu

Ninghui Li

ninghui@cs.purdue.edu

Center for Education and Research in Information Assurance and Security
and Department of Computer Sciences
Purdue University
656 Oval Drive, West Lafayette, IN 47907

ABSTRACT

Comparing the expressive power of access control models is recognized as a fundamental problem in computer security. Such comparisons are generally based on simulations between different access control schemes. However, the definitions for simulations that are used in the literature make it impossible to put results and claims about the expressive power of access control models into a single context and to compare such models to one another in a meaningful way. We propose a theory for comparing the expressive power of access control models. We perceive access control systems as state-transition systems and require simulations to preserve security properties. We discuss the rationale behind such a theory, apply the theory to reexamine some existing work on the expressive power of access control models in the literature and present three results. We show that: (1) RBAC with a particular administrative model from the literature (ARBAC97) is limited in its expressive power; (2) ATAM (Augmented Typed Access Matrix) is more expressive than TAM (Typed Access Matrix), thereby solving an open problem posed in the literature; and (3) a trust-management language is at least as expressive as RBAC with a particular administrative model (the URA97 component of ARBAC97).

Categories and Subject Descriptors

K.6.5 [**Management of Computing and Information Systems**]: Security and Protection; D.4.6 [**Operating Systems**]: Security and Protection — Access Controls

General Terms

Security, Theory, Verification

Keywords

Expressive Power, Reduction, State-Matching Reduction, Role-Based Access Control, Discretionary Access Control, Typed Access Matrix, Augmented Typed Access Matrix

1. INTRODUCTION

An access control system enforces a policy on who may access what resources and in what manner. Policies are generally expressed in terms of the current state of the system, and states that may result from prospective changes (e.g., "Alice should always have read access to a particular file, f"). When an access control system is perceived as a state-transition system, it consists of a set of states, rules on how state-transitions may occur and a set of properties or queries that are of interest in a given state (e.g., "Does Alice have read access to a particular file, f?") Policies may then be expressed in terms of these components, and such policies may be verified to hold notwithstanding the fact that state-transitions occur.

An access control system is an instance of an access control scheme: a scheme specifies the types of state-transition rules that may be specified in a system based on that scheme. A set of access control schemes is an access control model. An example of an access control model is the access matrix model [5]. An example of a scheme based on the access matrix model is the HRU scheme [6] which specifies that state-transition rules are commands of a particular form. A specific set of HRU commands together with a start state is an example of an access control system. The expressive power of an access control model captures the notion of whether different policies can be represented in systems based on schemes from that model.

Comparing the expressive power of access control models is recognized as a fundamental problem in information security and is studied extensively in the literature [1, 3, 4, 15, 19, 16, 18]. The expressive power of a model is tied to the expressive power of the schemes from the model. In comparing schemes based on expressive power, we ask what types of policies can be represented by systems based on a scheme. If all policies that can be represented in scheme B can be represented in scheme A, then scheme A is at least as expressive as scheme B.

A common methodology used for comparing access control models in previous work is *simulation*. When a scheme A is simulated in a scheme B, each system in A is mapped to a corresponding system in B. If every scheme in one model can be simulated by some scheme in another model, then the latter model is considered to be at least as expressive as the former. Furthermore, if there exists a scheme in the latter model that cannot be simulated by any scheme in the former, then the latter model is strictly more expressive than the former. Different definitions for simulations are used in the literature on comparing access control models. We identify two axes along which these definitions differ.

- The first axis is whether a simulation is required to preserve safety properties. In the comparison of different schemes based on the access matrix model [1, 4, 16, 18], the preservation of safety properties is required. If a scheme A is simulated in a scheme B, then a system in scheme A reaches an unsafe state if and only if the image of the system under the simulation (which is a system in scheme B) reaches an unsafe state.

  On the other hand, the preservation of safety properties is not required in the simulations used for comparing MAC (Mandatory Access Control), DAC (Discretionary Access Control), and RBAC (Role-Based Access Control) [15, 19, 13]. Nor is it required in the simulations used for the comparison of Access Control Lists (ACL), Capabilities, and Trust Management (TM) systems [3]. In these comparisons, the requirement for a simulation of A in B is that it should be possible to use an implementation of the scheme B to implement the scheme A. We call this the *implementation paradigm* of simulations.

- The second axis is whether to restrict the number of state-transitions that the simulating scheme needs to make in order to simulate one state-transition in the scheme being simulated. Chander et al. [3] define the notions of strong and weak simulations. A strong simulation of A in B requires that B makes one state-transition when A makes one state-transition. A weak simulation requires that B makes a bounded (by a constant) number of state-transitions to simulate one state-transition in A. A main result in [3] is that a specific TM scheme considered there is more expressive than ACL because there exists no (strong or weak) simulation of the TM scheme in ACL. The proof is based on the observation that an unbounded (but still finite) number of state-transitions in ACL are required to simulate one state-transition in the TM scheme.

  On the other hand, an unbounded number of state-transitions is allowed by Sandhu and Ganta [18]. They use a simulation that involves an unbounded number of state-transitions to prove that ATAM (Augmented Typed Access Matrix) is equivalent in expressive power to TAM (Typed Access Matrix).

Although significant progress has been made in comparing access control models, this current state of art is unsatisfactory for the following reasons. First, different definitions of simulations make it impossible to put different results and claims about expressive power of access control models into a single context. For example, the result that RBAC is at least as expressive as DAC [15, 13] is qualitatively different from the result that TAM is at least as expressive as ATAM [18], as the former does not require the preservation of safety properties. These results are again qualitatively different from the result that ACL is less expressive than Trust Management [3], as the latter requires a bounded number of state-transitions in simulations.

Second, some definitions of simulations that are used in the literature are too weak to distinguish access control models from one another in a meaningful way. Sandhu et al. [13, 15, 19] show that various forms of DAC (including ATAM, in which simple safety is undecidable) can be simulated in RBAC, using the notion of simulations derived from the implementation paradigm. We show in [20] that using the same notion of simulations, RBAC can be simulated in strict DAC, one of the most basic forms of DAC in which simple safety is trivially decidable. This suggests that using such a notion of simulations, it is likely that one can show that

all access control models have the same expressive power. Thus, this notion of simulations is not useful in differentiating between models based on expressive power.

Finally, the rationale for some choices made in existing definitions of simulations is often not clearly stated and justified. It is unclear why certain requirements are made or not made for simulations when comparing the expressive power of access control models. For instance, when a simulation involves an unbounded number of state-transitions, Ganta [4] considers this to be a "weak" simulation, while Chander et al. [3] do not consider this to be a simulation at all.

In this paper, we build on existing work and seek to construct uniform bases for comparing access control models. To determine the requirements on simulations in a systematic and justifiable manner, we start from the rationales and intuitions underlying different definitions for simulations. Our approach is to first identify the desirable and intuitive properties one would like simulations to have and then come up with conditions on simulations that are both sufficient and necessary to satisfy those properties. Informally, what is desired is that when one scheme can represent all types of policies that another can, then the former is deemed to be at least as expressive as the latter. This observation is made by Ganta [4] as well.

Our theory is based on definitions of simulations that preserve security properties. Examples of such security properties are availability, mutual exclusion and bounded safety. Intuitively, such security properties are the sorts of policies one would want to represent in an access control system. *Security analysis* is used to verify that desired security properties are indeed maintained across state-transitions in an access control system. It was introduced by Li et al. [11], and generalizes the notion of safety analysis [6]. In this paper, we introduce compositional security analysis, which generalizes security analysis to consider logical combinations of queries in security analysis.

We introduce two notions of simulations called *state-matching reductions* and *reductions*. We show that state-matching reductions are necessary and sufficient for preserving compositional security properties and that reductions are necessary and sufficient for preserving security properties. A state-matching reduction reduces the compositional security analysis problem in one scheme to that in another scheme. A reduction reduces the security analysis problem in one scheme to that in another scheme.

To summarize, the contributions of this paper are as follows.

- We introduce a theory for comparing access control models based on the notions of state-matching reductions and reductions, together with detailed justifications for the design decisions.

- We analyze the deficiency of using the implementation paradigm to compare access control models and show that it leads to a weak notion of simulations and cannot be used to differentiate access control models from one another based on expressive power.

- We apply our theory in three cases. We show that:

  – There exists a reduction, but no state-matching reduction from Strict DAC with Change of Ownership (SDCO) to RBAC with ARBAC97 [17] as the administrative model. To our knowledge, this is the first evidence of the limitation of the expressive power of RBAC in comparison to DAC. RBAC has been compared to various forms of DAC, including SDCO, in the literature [15, 19].

63

– There exists a state-matching reduction from RBAC with an administrative model that is a component of ARBAC97 [17] to RT [8, 9], a trust-management language.

– There exists no state-matching reduction from ATAM to TAM. This solves an open problem stated by Sandhu and Ganta [18] by formalizing the benefit of the ability to check for the absence of rights in addition to the ability to check for the presence of rights.

The rest of this paper is organized as follows. We present our theory for comparing access control models in Section 2. In Section 3, we analyze the implementation paradigm for simulations. In Section 4.1, we discuss comparisons of DAC to RBAC from the literature. In the rest of Section 4, we apply our theory to compare the expressive power of schemes in three cases. We conclude with Section 5. Proofs and precise characterizations of schemes not included in the paper appear in [20].

2. COMPARISONS BASED ON SECURITY ANALYSIS

A requirement used in the literature for simulations is the preservation of safety properties. Indeed, this is the only requirement from simulations in [1, 16, 18]. If a simulation of scheme A in scheme B satisfies this requirement, then a system in A reaches an unsafe state if and only if the system's mapping in B reaches an unsafe state. In other words, the result of safety analysis is preserved by the simulation.

Safety analysis, i.e., determining whether an access control system can reach a state in which an unsafe access is allowed, was first formalized by Harrison et al. [6] in the context of the well-known access matrix model [5, 7]. In the HRU scheme [6], a protection system has a finite set of rights and a finite set of commands. A state 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 the given conditions hold in the current state, 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. The following is an example command that allows the owner of a file to grant the read right to another user.

```
command grantRead(u1,u2,f)
  if  'own' in (u1,f)
  then enter 'read' into (u2,f)
end
```

In the example, u1, u2 and f are formal parameters to the command. They are instantiated by objects (or subjects) when the command is executed. In [6], Harrison et al. prove that in the HRU scheme, the safety question is undecidable, by showing that any Turing machine can be simulated by a protection system.

Treating the preservation of safety properties as the sole requirement of simulations is based on the implicit assumption that safety is the *only* interesting property in access control schemes, an assumption that is not valid. When originally introduced in [6], safety was described as just one class of queries one can consider. Recently, Li et al. [11] introduced the notion of security analysis, which generalizes safety to other properties such as simple safety, bounded safety, simple availability, mutual exclusion and containment.

In this section, we present a theory for comparing access control models based on the preservation of security properties.

2.1 Access Control Schemes and Security Analysis

Definition 1. (*Access Control Scheme*) An access control scheme is a state-transition system $\langle \Gamma, Q, \vdash, \Psi \rangle$, in which Γ is a set of states, Q is a set of queries, $\vdash: \Gamma \times Q \rightarrow \{true, false\}$ is called the entailment relation, and Ψ is a set of state-transition rules.

A *state*, $\gamma \in \Gamma$, contains all the information necessary for making access control decisions at a given time. The *entailment relation*, $\vdash$, determines whether a *query* is true or not in a given state. When a query, $q \in Q$, arises from an access request, $\gamma \vdash q$ means that the access request q is allowed in the state γ, and $\gamma \nvdash q$ means that q is not allowed. Some access control schemes also allow queries other than those corresponding to a specific request, e.g., whether every subject that has access to a resource is an employee of the organization. Such queries can be useful for understanding the properties of complex access control systems.

A *state-transition rule*, $\psi \in \Psi$, determines how the access control system changes state. More precisely, ψ defines a binary relation (denoted by $\mapsto_\psi$) on Γ. Given $\gamma, \gamma_1 \in \Gamma$, we write $\gamma \mapsto_\psi \gamma_1$ if the change of state from γ to γ_1 is allowed by ψ, and $\gamma \overset{*}{\mapsto}_\psi \gamma_1$ if a sequence of zero or more allowed changes leads from γ to γ_1. In other words, $\overset{*}{\mapsto}_\psi$ is the transitive closure of $\mapsto_\psi$. If $\gamma \overset{*}{\mapsto}_\psi \gamma_1$, we say that γ_1 is ψ-*reachable* from γ, or simply γ_1 is *reachable*, when γ and ψ are clear from the context.

An *access control model* is a set of access control schemes. An *access control system* in an access control scheme $\langle \Gamma, Q, \vdash, \Psi \rangle$ is given by a pair (γ, ψ), where $\gamma \in \Gamma$ is the current state of the system and $\psi \in \Psi$ is the state-transition rule that governs the system's state changes.

Similar definitions for access control schemes appear in [1, 3]; our definition from above appears also in [10], and is different from the definitions in [1, 3] in the following two respects. First, our definition is more abstract in that it does not refer to subjects, objects, and rights and that the details of a state-transition rule are not specified. We find such an abstract definition more suitable to capture the notion of expressive power especially when the models or schemes that are compared are "structurally" different (e.g., a scheme based on RBAC that has a notion of roles that is an indirection between users and permissions, and a scheme based on the access-matrix model in which rights are assigned to subjects directly). Second, our definition makes the set of queries that can be asked an explicit part of the specification of an access control scheme. In existing definitions in the literature, the set of queries is often not explicitly specified. Sometimes, the implicit set of queries is clear from context; at other times, it is not clear.

The HRU Scheme We now show an example access control scheme, the HRU scheme, that is derived from the work by Harrison et al. [6]. We assume the existence of three countably infinite sets: $\mathcal{S}$, $\mathcal{O}$, and $\mathcal{R}$, which are the sets of all possible subjects, objects, and rights. We assume further that $\mathcal{S} \subseteq \mathcal{O}$. In the HRU scheme:

- Γ is the set of all possible access matrices. Formally, each $\gamma \in \Gamma$ is identified by three finite sets, $S_\gamma \subset \mathcal{S}$, $O_\gamma \subset \mathcal{O}$, and $R_\gamma \subset \mathcal{R}$, and a function $M_\gamma[] : S_\gamma \times O_\gamma \rightarrow 2^{R_\gamma}$, where $M_\gamma[s, o]$ gives the set of rights s has over o.

- Q is the set of all queries of the form: $r \in [s, o]$, where $r \in \mathcal{R}$ is a right, $s \in \mathcal{S}$ is a subject, and $o \in \mathcal{O}$ is an

object. This query asks whether the right r exists in the cell corresponding to subject s and object o.

- The entailment relation is defined as follows: $\gamma \vdash r \in [s, o]$ if and only if $s \in S_\gamma$, $o \in O_\gamma$, and $r \in M_\gamma[s, o]$.

- Each state-transition rule ψ is given by a set of command schemas. Given ψ, the change from γ to γ_1 is allowed if there exists an instance of a command schema in ψ that when applied to γ results in γ_1.

The set of queries is not explicitly specified in [6]. It is conceivable to consider other classes of queries, e.g., comparing the set of all subjects that have a given right over a given object with another set of subjects. In our framework, HRU with different classes of queries can be viewed as different schemes in the access matrix model.

Definition 2. (*Security Analysis*) Given an access control system $\langle \Gamma, Q, \vdash, \Psi \rangle$, a *security analysis instance* has the form $\langle \gamma, q, \psi, \Pi \rangle$, where $\gamma \in \Gamma$ is a state, $q \in Q$ is a query, $\psi \in \Psi$ is a state-transition rule, and $\Pi \in \{\exists, \forall\}$ is a quantifier. An instance $\langle \gamma, q, \psi, \exists \rangle$ is said to be *existential*; it asks whether there exists γ_1 such that $\gamma \overset{*}{\mapsto}_\psi \gamma_1$ and $\gamma_1 \vdash q$. If so, we say q is *possible* (given γ and ψ). An instance $\langle \gamma, q, \psi, \forall \rangle$ is said to be *universal*; it asks whether for every γ_1 such that $\gamma \overset{*}{\mapsto}_\psi \gamma_1$, $\gamma_1 \vdash q$. If so, we say q is *necessary* (given γ and ψ).

Simple safety analysis is a special case of security analysis. A simple safety analysis instance that asks whether a system (γ, ψ) in the HRU scheme can reach a state in which the subject s has the right r over the object o is represented as the following instance: $\langle \gamma, r \in [s, o], \psi, \exists \rangle$. The universal version of this instance, $\langle \gamma, r \in [s, o], \psi, \forall \rangle$, asks whether s always has the right r over the object o in every reachable state. Thus it refers to the availability property and asks whether a particular access right is always available to the subject s. We now introduce a generalized notion of security analysis.

Definition 3. (*Compositional Security Analysis*) Given a scheme $\langle \Gamma, Q, \vdash, \Psi \rangle$, a *compositional security analysis* instance has the form $\langle \gamma, \phi, \psi, \Pi \rangle$, where γ, ψ, and Π are the same as in a security analysis instance, and ϕ is a (possibly infinite) propositional formula over Q, i.e., ϕ is constructed from queries in Q using propositional logic connectives such as $\wedge$, $\vee$, and $\neg$.

For example, the compositional security analysis instance $\langle \gamma, (r_1 \in [s, o_1]) \wedge (r_2 \in [s, o_2]), \psi, \exists \rangle$ asks whether the system (γ, ψ) can reach a state in which s has both the right r_1 over o_1 and the right r_2 over o_2. We argue that ϕ should be allowed to be inifite by considering a safety property in the context of the HRU scheme [6]. The property is whether any subject can get a particular right r over a particular object o that the subject does not have in the start-state γ. This property is represented in our formalism by letting ϕ be $\bigvee_i (r \in [s_i, o])$ where $s_i \in \mathcal{S} - \widehat{S}_\gamma$ and $\widehat{S}_\gamma$ is the set of subjects each of whom has the right r over o in the state γ.

Whether we should use security analysis or compositional security analysis is related to what types of policies we want to represent, and what types of policies we want to use as bases to compare the expressive power of different access control models or schemes. With compositional security analysis, we would be comparing models or schemes based on types of policies that are broader than with security analysis. For instance, if our set of queries Q contains queries related to users' access to files, then with compositional security analysis we can consider policies such as "Bob should never have write access to a particular file so long as his wife, Alice has a user account (and thus has some type of access to some file)."

2.2 Two Types of Reductions

In this section, we introduce the notions of reductions and state-matching reductions that we believe are adequate for comparing the expressive power of access control models. Before we introduce reductions, we discuss two types of mappings between access control schemes.

Definition 4. (*Mapping*) Given two access control schemes $A = \langle \Gamma^A, Q^A, \vdash^A, \Psi^A \rangle$ and $B = \langle \Gamma^B, Q^B, \vdash^B, \Psi^B \rangle$, a *mapping* from A to B is a function σ that maps each pair $\langle \gamma^A, \psi^A \rangle$ in A to a pair $\langle \gamma^B, \psi^B \rangle$ in B and maps each query q^A in A to a query q^B in B. Formally, $\sigma : (\Gamma^A \times \Psi^A) \cup Q^A \to (\Gamma^B \times \Psi^B) \cup Q^B$.

Definition 5. (*Security-Preserving Mapping*) A mapping σ is said to be *security-preserving* when every security analysis instance in A is true if and only if the *image* of the instance is true. Given a mapping $\sigma : (\Gamma^A \times \Psi^A) \cup Q^A \to (\Gamma^B \times \Psi^B) \cup Q^B$, the *image* of a security analysis instance $\langle \gamma^A, q^A, \psi^A, \Pi \rangle$ *under* σ is $\langle \gamma^B, q^B, \psi^B, \Pi \rangle$, where $\langle \gamma^B, \psi^B \rangle = \sigma(\langle \gamma^A, \psi^A \rangle)$ and $q^B = \sigma(q^A)$.

The notion of security-preserving mappings captures the intuition that simulations should preserve security properties. Given a security-preserving mapping from A to B and an algorithm for solving the security analysis problem in B, one can construct an algorithm for solving the security analysis problem in A using the mapping. Also, security analysis in B is at least as hard as security analysis in A, modulo the efficiency of the mapping. If an efficient (polynomial-time) mapping from A to B exists, and security analysis in A is intractable (or undecidable), then security analysis in B is also intractable (undecidable). Security preserving mappings are not powerful enough for comparisons of access control schemes based on compositional security analysis. We need the notion of a strongly security-preserving mapping for that purpose.

Definition 6. (*Strongly Security-Preserving Mapping*) Given a mapping σ from scheme A to scheme B, the image of a compositional analysis instance, $\langle \gamma^A, \phi^A, \psi^A, \Pi \rangle$, in A is $\langle \gamma^B, \phi^B, \psi^B, \Pi \rangle$, where $\langle \gamma^B, \psi^B \rangle = \sigma(\langle \gamma^A, \psi^A \rangle)$ and ϕ^B is obtained by replacing every query q^A in ϕ^A with $\sigma(q^A)$ (we abuse the terminology slightly and write $\phi^B = \sigma(\phi^A)$). A mapping σ from A to B is said to be *strongly security-preserving* when every compositional security analysis instance in A is true if and only if the image of the instance is true.

While the notions of security-preserving mappings capture the intuition that simulations should preserve security properties, they are not convenient for us to use directly. Using the definition for either type of mapping to directly prove that the mapping is (strongly) security preserving involves performing security analysis, which is often expensive. We now introduce the notions of reductions, which state structural requirements on mappings for them to be security preserving. We start with a form of reduction appropriate for compositional security analysis and then discuss weaker forms.

Definition 7. (*State-Matching Reduction*) Given a mapping from A to B, $\sigma : (\Gamma^A \times \Psi^A) \cup Q^A \to (\Gamma^B \times \Psi^B) \cup Q^B$, we say that the two states γ^A and γ^B are *equivalent* under the

mapping σ when for every $q^A \in Q^A$, $\gamma^A \vdash^A q^A$ if and only if $\gamma^B \vdash^B \sigma(q^A)$. A mapping σ from A to B is said to be a *state-matching reduction* if for every $\gamma^A \in \Gamma^A$ and every $\psi^A \in \Psi^A$, $\langle \gamma^B, \psi^B \rangle = \sigma(\langle \gamma^A, \psi^A \rangle)$ has the following two properties:

1. For every state γ_1^A in scheme A such that $\gamma^A \stackrel{*}{\mapsto}_\psi \gamma_1^A$, there exists a state γ_1^B such that $\gamma^B \stackrel{*}{\mapsto}_{\psi^B} \gamma_1^B$ and γ_1^A and γ_1^B are equivalent under σ.

2. For every state γ_1^B in scheme B such that $\gamma^B \stackrel{*}{\mapsto}_{\psi^B} \gamma_1^B$, there exists a state γ_1^A such that $\gamma^A \stackrel{*}{\mapsto}_\psi \gamma_1^A$ and γ_1^A and γ_1^B are equivalent under σ.

Property 1 says that for every state γ_1^A that is reachable from γ^A, there exists a reachable state in scheme B that is equivalent, i.e., answers all queries in the same way. Property 2 says the reverse, for every reachable state in B, there exists an equivalent state in A. The goal of these two properties is to guarantee that compositional security analysis results are preserved across the mapping. With the following theorem, we justify Definition 7.

THEOREM 1. *Given two schemes A and B, a mapping σ from A to B is strongly security-preserving if and only if σ is a state-matching reduction.*

PROOF. **The "if" direction.** When σ is a state-matching reduction, given a compositional security analysis instance $\langle \gamma^A, \phi^A, \psi^A, \Pi \rangle$ in scheme A, let $\langle \gamma^B, \psi^B \rangle = \sigma(\langle \gamma^A, \psi^A \rangle)$ and $\phi^B = \sigma(\phi^A)$, we show that $\langle \gamma^A, \phi^A, \psi^A, \Pi \rangle$ is true if and only if $\langle \gamma^B, \phi^B, \psi^B, \Pi \rangle$ is true.

First consider the case that the instance $\langle \gamma^A, q^A, \psi^A, \Pi \rangle$ is existential, i.e., Π is $\exists$. If the instance is true, then there exists a reachable state γ_1^A in which ϕ^A is true. Property 1 in Definition 7 guarantees that there exists a reachable state γ_1^B that is equivalent to γ_1^A; thus ϕ^B is true in γ_1^B; therefore, the instance in B, $\langle \gamma^B, \phi^B, \psi^B, \exists \rangle$, is also true. On the other hand, if $\langle \gamma^B, \phi^B, \psi^B, \exists \rangle$ is true, then there exists a reachable state γ_1^B in which ϕ^B is true. Property 2 in Definition 7 guarantees that there exists a state in A in which the analysis instance in A is true.

Now consider the case that the instance $\langle \gamma^A, \phi^A, \psi^A, \Pi \rangle$ is universal, i.e., Π is $\forall$. If the instance is false, then there exists a reachable state γ_1^A in which ϕ^A is false. Property 1 guarantees that the instance in B is also false. Similarly, if the instance in B is false, then the instance in A is also false.

The "only if" direction. When σ is not a state-matching reduction, then there exists $\gamma^A \in \Gamma^A$ and $\psi^A \in \Psi^A$ such that $\langle \gamma^B, \psi^B \rangle = \sigma(\langle \gamma^A, \psi^A \rangle)$ violates one of the two properties in Definition 7.

First consider the case that Property 1 is violated. There exists a reachable state γ_1^A such that no state reachable from γ^B is equivalent to γ_1^A. Construct a formula ϕ^A as follows: ϕ^A is a conjunction of queries in Q or their complement. For every query q^A in Q^A, ϕ^A includes q^A if $\gamma_1^A \vdash^A q^A$ and $\neg q^A$ if $\gamma_1^A \vdash^A \neg q^A$. Note that the length of ϕ^A may be infinite, as the total number of queries may be infinite. Clearly, ϕ^A is true in γ_1^A, but $\sigma(\phi^A)$ is false in all states reachable from γ^B. Thus, the existential compositional analysis instance involving ϕ^A has different answers, and σ is not strongly security preserving.

Then consider the case that Property 2 is violated. There exists a state γ_1^B reachable from γ^B such that no state reachable from γ^A is equivalent to γ_1^B. Construct a formula ϕ^A as follows: ϕ^A is a conjunction of queries in Q or their complement. For every query q^A in Q^A, ϕ^A includes q^A if $\gamma_1^B \vdash^B \sigma(q^A)$ and $\neg q^A$ if $\gamma_1^B \vdash^B \sigma(\neg q^A)$. Clearly, ϕ^A is false in all states reachable from γ^A, but $\sigma(\phi^A)$ is true in γ_1^B; thus, the existential compositional analysis instance involving ϕ^A has different answers, and σ is not strongly security preserving. $\square$

A state-matching reduction preserves compositional security properties. If we need only queries from Q to represent our policies and not compositions of those queries, then the following weaker notion of reductions is more suitable. However, we believe that the notion of state-matching reductions is quite natural by itself, and certainly necessary when compositional queries are of interest.

Definition 8. (*Reduction*) Given two access control schemes $A = \langle \Gamma^A, Q^A, \vdash^A, \Psi^A \rangle$ and $B = \langle \Gamma^B, Q^B, \vdash^B, \Psi^B \rangle$, a mapping from A to B, σ, is said to be a *reduction* from A to B if for every $\gamma^A \in \Gamma^A$ and every $\psi^A \in \Psi^A$, $\langle \gamma^B, \psi^B \rangle = \sigma(\langle \gamma^A, \psi^A \rangle)$ has the following two properties:

1. For every state γ_1^A and every query q^A in scheme A, if $\gamma^A \stackrel{*}{\mapsto}_\psi \gamma_1^A$, then in scheme B there exists a state γ_1^B such that $\gamma^B \stackrel{*}{\mapsto}_{\psi^B} \gamma_1^B$ and $\gamma_1^A \vdash^A q^A$ if and only if $\gamma_1^B \vdash^B \sigma(q^A)$.

2. For every state γ_1^B in scheme B and every query q^A in scheme A, if $\gamma^B \stackrel{*}{\mapsto}_{\psi^B} \gamma_1^B$, there exists a state γ_1^A such that $\gamma^A \stackrel{*}{\mapsto}_\psi \gamma_1^A$ and $\gamma_1^A \vdash^A q^A$ if and only if $\gamma_1^B \vdash^B \sigma(q^A)$.

Definition 7 differs from Definition 8 in that the former requires that for every reachable state in A (B, resp.) there exist a matching state in B (A, resp.) that gives the same answer for *every query*. Definition 8 requires the existence of a matching state for every query; however, the matching states may be different for different queries. Property 1 in Definition 8 says that for every reachable state in A and every query in A, there exists a reachable state in B that gives the same answer to (the image of) the query. Property 2 says the reverse direction. The goal of these two properties is to guarantee that security analysis results are preserved across the mapping. The fact that a reduction, as defined in Definition 8, is adequate for preserving security analysis results is formally captured by the following theorem.

THEOREM 2. *Given two schemes A and B, a mapping, σ, from A to B is security preserving if and only if σ is a reduction.*

PROOF. **The "if" direction.** When σ is a reduction, given a security analysis instance $\langle \gamma^A, q^A, \psi^A, \Pi \rangle$ in scheme A, let $\langle \gamma^B, \psi^B \rangle = \sigma(\langle \gamma^A, \psi^A \rangle)$ and $q^B = \sigma(q^A)$, we show that $\langle \gamma^A, q^A, \psi^A, \Pi \rangle$ is true if and only if $\langle \gamma^B, q^B, \psi^B, \Pi \rangle$ is true.

First consider the case that the instance $\langle \gamma^A, q^A, \psi^A, \Pi \rangle$ is existential, i.e., Π is $\exists$. If the instance is true, then there exists a reachable state γ_1^A in which q^A is true. Property 1 in Definition 8 guarantees that there exists a reachable state γ_1^B in which q^B is true. Therefore, the instance in B, $\langle \gamma^B, q^B, \psi^B, \exists \rangle$, is also true. On the other hand, if $\langle \gamma^B, q^B, \psi^B, \exists \rangle$ is true, then there exists a reachable state γ_1^B in which q^B is true. Property 2 in Definition 8 guarantees that there exists a state in A in which q^A is true; thus the analysis instance in A is true.

Now consider the case that the instance $\langle \gamma^A, q^A, \psi^A, \Pi \rangle$ is universal, i.e., Π is $\forall$. If the instance is false, then there exists a reachable state γ_1^A in which q^A is false. Property 1 guarantees that the instance in B is also false. Similarly, if the instance in B is false, then the instance in A is also false.

The "only if" direction. When σ is not a reduction, then there exists $\gamma^A \in \Gamma^A$ and $\psi^A \in \Psi^A$ such that $\langle \gamma^B, \psi^B \rangle = \sigma(\langle \gamma^A, \psi^A \rangle)$ violates one of the two properties in Definition 8.

First consider the case that Property 1 is violated. There exists a reachable state γ_1^A and a query q^A such that for every state reachable from γ^B the answer for the query $\sigma(q^A)$ in the state is different from the answer for q^A in γ_1^A. If $\gamma_1^A \vdash^A q^A$, then this means that q^B is false in every state reachable from γ^B. Thus the security analysis instance $\langle \gamma^A, q^A, \psi^A, \exists \rangle$ is true, but its image under σ is false. Thus, the mapping σ is not security-preserving. If $\gamma_1^A \nvdash^A q^A$, then this means that q^B is true in every state reachable from γ^B. Thus the security analysis instance $\langle \gamma^A, q^A, \psi^A, \forall \rangle$ is false, but its image under σ is true.

Now consider the case that Property 2 is violated. There exists a state γ_1^B reachable from γ^B and a query q^A such that for every state reachable from γ^A the answer for the query q^A in the state is different from the answer for $\sigma(q^A)$ in γ_1^B. If $\gamma_1^B \vdash^B \sigma(q^A)$, then this means that q^A is false in every state reachable from γ^A. Thus the security analysis instance $\langle \gamma^A, q^A, \psi^A, \exists \rangle$ is false, but its image under σ is true. If $\gamma_1^B \nvdash^B q^B$, then this means that q^A is true in every state reachable from γ^A. Thus the security analysis instance $\langle \gamma^A, q^A, \psi^A, \forall \rangle$ is true, but its mapping in B is false. $\square$

Comparisons of two access control models are based on comparisons among access control schemes based on those models. Comparisons of two access control schemes, in turn, are based on whether only the queries from Q need to be represented, or compositions of those queries need to be represented as well.

Definition 9. (*Comparing the Expressive Power of Access Control Models*) Given two access control models $\mathcal{M}$ and $\mathcal{M}'$, we say that $\mathcal{M}'$ is at least as expressive as $\mathcal{M}$ (or $\mathcal{M}'$ has at least as much expressive power as $\mathcal{M}'$) if for every scheme in $\mathcal{M}$ there exists a state-matching reduction (or a reduction) from it to a scheme in $\mathcal{M}'$. In addition, if for every scheme in $\mathcal{M}'$, there exists a state-matching reduction (reduction) from it to a scheme in $\mathcal{M}$, then we say that M and M' are equivalent in expressive power. If $\mathcal{M}'$ is at least as expressive as the $\mathcal{M}$, and there exists a scheme A in $\mathcal{M}'$ such that for any scheme B in $\mathcal{M}$, no state-matching reduction (reduction) from A to B exists, we say that $\mathcal{M}'$ is strictly more expressive than $\mathcal{M}$.

We compare the expressive power of two schemes based on state-matching reductions when compositional queries are needed to represent the policies of interest. Otherwise, reductions suffice. Observe that we can use the above definition to compare the expressive power of two access control schemes A and B, by viewing each scheme as an access control model that consists of just that scheme.

We emphasize that a reduction or state-matching reduction must be computable. In addition, if there exists a reduction or state-matching reduction from A to B that can be computed efficiently in the size of A, then we can use the efficiency with which security analysis can be performed in B as a tight upper bound for the analysis instance in A.

2.3 Alterative definitions for reduction

In this section, we discuss alternative definitions that differ from the ones discussed in the previous section. The first of these definitions is used by Sandhu and Ganta [16, 18] for simulations.

Definition 10. (*Form-1 Weak Reduction*) A mapping from A to B, given by $\sigma : (\Gamma^A \times \Psi^A) \cup Q^A \rightarrow (\Gamma^B \times \Psi^B) \cup Q^B$, is a *form-1 weak reduction* if for every $\gamma^A \in \Gamma^A$ and every $\psi^A \in \Psi^A$, $\langle \gamma^B, \psi^B \rangle = \sigma(\langle \gamma^A, \psi^A \rangle)$ has the following two properties:

1. For every query q^A, if there exists a state γ_1^A in scheme A such that $\gamma^A \overset{*}{\mapsto}_{\psi^A} \gamma_1^A$ and $\gamma_1^A \vdash^A q^A$, then there exists a state γ_1^B such that $\gamma^B \overset{*}{\mapsto}_{\psi^B} \gamma_1^B$ and $\gamma_1^B \vdash^B \sigma(q^A)$.

2. For every query q^A, if there exists γ_1^B in scheme B such that $\gamma^B \overset{*}{\mapsto}_{\psi^B} \gamma_1^B$ and $\gamma_1^B \vdash^B \sigma(q^A)$, then there exists a state γ_1^A such that $\gamma^A \overset{*}{\mapsto}_{\psi} \gamma_1^A$ and $\gamma_1^A \vdash^A q^A$ if and only if $\gamma_1^B \vdash^B \sigma(q^A)$.

The intuition underlying Definition 10, as stated by Sandhu [16] is, "systems are equivalent if they have equivalent worst case behavior". Therefore, simulations only need to preserve the worst-case access. Definition 10 is weaker than Definition 8 in that it requires the existence of a matching state when a query is true in the state, but does not require so when the query is false. Therefore, it is possible that a query q^A is true in all states that are reachable from γ^A, but the query $\sigma(q^A)$ is false in some states that are reachable from γ^B (the query $\sigma(q^A)$ needs to be true in at least one state reachable from γ^B). This indicates that Definition 10 does not preserve answers to universal security analysis instances. Definition 10 is adequate for the purposes in [16, 18] as only safety analysis (which is existential) was considered there.

The decision of defining a mapping to be a function from $(\Gamma^A \times \Psi^A) \cup Q^A$ to $(\Gamma^B \times \Psi^B) \cup Q^B$ also warrants some discussion. An alternative is to define a mapping from A to B to be a function that maps each state in A to a state in B, each state-transition rule in A to a state-transition rule in B, and each query in A to a query in B. Such a function would be denoted as $\sigma : \Gamma^A \cup \Psi^A \cup Q^A \rightarrow \Gamma^B \cup \Psi^B \cup Q^B$. One can verify any such function is also a mapping according to Definition 4, which gives more flexibility in terms of mapping states and state-transition rules from A to B. By Definition 4, the state corresponding to a state γ^A may also depend upon the state-transition rule being considered.

Another alternative is to define a mapping from A to B to be a function $\sigma : \Gamma^A \times \Psi^A \times Q^A \rightarrow \Gamma^B \times \Psi^B \times Q^B$. In other words, the mapping of states, state-transition rules, and queries may depend on each other. This definition also leads to a weaker notion of reduction:

Definition 11. (*Form-2 Weak Reduction*) A form-2 weak reduction from A to B is a function $\sigma : \Gamma^A \times \Psi^A \times Q^A \rightarrow \Gamma^B \times \Psi^B \times Q^B$ such that for every $\gamma^A \in \Gamma^A$, every $\psi^A \in \Psi^A$, and every $q^A \in Q^A$, $\langle \gamma^B, \psi^B, q^B \rangle = \sigma(\langle \gamma^A, \psi^A, q^A \rangle)$ has the following two properties:

1. For every state γ_1^A in scheme A such that $\gamma^A \overset{*}{\mapsto}_{\psi} \gamma_1^A$, there exists a state γ_1^B such that $\gamma^B \overset{*}{\mapsto}_{\psi^B} \gamma_1^B$ and $\gamma_1^A \vdash^A q^A$ if and only if $\gamma_1^B \vdash^B q^B$.

2. For every state γ_1^B in scheme B such that $\gamma^B \overset{*}{\mapsto}_{\psi^B} \gamma_1^B$, there exists a state γ_1^A such that $\gamma^A \overset{*}{\mapsto}_{\psi} \gamma_1^A$ and $\gamma_1^A \vdash^A q^A$ if and only if $\gamma_1^B \vdash^B q^B$.

It is not difficult to prove that a Form-2 weak reduction is also security preserving, in the sense that any security analysis instance $\langle \gamma^A, q^A, \psi^A, \Pi \rangle$ in A can be mapped to a security analysis in B. However, it is not a mapping, as the mapping of states and state-transition rules may depend on the query.

Definition 11 is used implicitly in Theorems 2 and 3 in [10] for reductions from two RBAC schemes to the RT Role-based Trust-management framework [9, 11]. As we assert in Theorem 5 in this paper, a reduction used there for one of the RBAC schemes can be changed to a security-preserving mapping in a straightforward manner.

We choose not to adopt this weaker notion of reduction for the following reason. Under this definition, given an access control system (γ^A, ψ^A), to answer n analysis instances involving different queries, one has to perform n translations of states and state-transitions, which is often time consuming. Using Definition 4 and

Definition 8, one can perform the mapping of (γ^A, ψ^A) once and use it to answer all n analysis instances.

A third weak form of reduction is introduced by Ammann et al. [1]. That work discusses the expressive power of multi-parent creation when compared to single-parent creation.

Definition 12. (*Form-3 Weak Reduction*) A mapping from A to B, given by $\sigma : (\Gamma^A \times \Psi^A) \cup Q^A \to (\Gamma^B \times \Psi^B) \cup Q^B$, is a *form-3 weak reduction* if for every $\gamma^A \in \Gamma^A$ and every $\psi^A \in \Psi^A$, $\langle \gamma^B, \psi^B \rangle = \sigma(\langle \gamma^A, \psi^A \rangle)$ has the following two properties:

1. For every state γ_1^A and every query q^A in scheme A, if $\gamma^A \overset{*}{\mapsto}_\psi \gamma_1^A$, then in scheme B there exists a state γ_1^B such that $\gamma^B \overset{*}{\mapsto}_{\psi^B} \gamma_1^B$ and $\gamma_1^A \vdash^A q^A$ if and only if $\gamma_1^B \vdash^B \sigma(q^A)$.

2. For every state γ_1^B in scheme B and every query q^A in scheme A, if $\gamma^B \overset{*}{\mapsto}_{\psi^B} \gamma_1^B$, then either (a) there exists a state γ_1^A such that $\gamma^A \overset{*}{\mapsto}_\psi \gamma_1^A$ and $\gamma_1^A \vdash^A q^A$ if and only if $\gamma_1^B \vdash^B \sigma(q^A)$, or (b) there exists a state γ_2^B such that $\gamma_1^B \overset{*}{\mapsto}_{\psi^B} \gamma_2^B$ and a state γ_1^A such that $\gamma^A \overset{*}{\mapsto}_\psi \gamma_1^A$, and $\gamma_1^A \vdash^A q^A$ if and only if $\gamma_2^B \vdash^B \sigma(q^A)$.

As pointed out by Ammann et al. [1], this form of reduction suffices for monotonic schemes — those schemes in which once a state is reached in which a query is true, in all reachable states from that state, the query remains true. Therefore, this form of reduction cannot be used to compare schemes when queries can become false after being true, or for universal analysis instances.

3. THE IMPLEMENTATION PARADIGM FOR SIMULATION: AN EXAMINATION

Several authors use the implementation paradigm for simulations, e.g., Osborn et al. [15] state that "a positive answer [to the question whether LBAC (lattice-based access control) can be simulated in RBAC] is also practically significant, because it implies that the same Trust Computing Base can be configured to enforce RBAC in general and LBAC in particular." However, in these papers [13, 15, 19], a precise definition for simulations is not given. This makes the significance of such results unclear, at least in terms of comparing the expressive power of different access control models.

In this section, we analyze the implementation paradigm and argue that it does not lead to notions of simulations that are meaningful for comparing the expressive power of different access control models. More precisely, the notions of simulations derived from this paradigm are so weak that almost all access control schemes are equivalent.

To formalize the implementation paradigm for simulation, a natural goal is to use an implementation of an access control scheme for another scheme. Intuitively, if a scheme A can be simulated in a scheme B, then there exists a *simulator* that, when given access to an interface to (an implementation of) B, can provide an interface that is exactly the same as the interface to (an implementation of) A.

When considering the interface of an access control scheme, we have to consider how state-transitions occur. Intuitively, an access control system changes its state because some actors (subjects, principals, users, etc.) initiate certain actions. Thus, an implementation of an access control scheme has an interface consisting of at least the following functions:

- $init(\gamma)$: set the current state to γ.

- $query(q)$: ask the query q and receive a yes/no response.

- $apply(a)$: apply the action a on the system, which may result in a state-transition in the system.

- functions providing other capabilities, e.g., traversing the subjects and objects in the system.

A simulator of A in B is thus a program that takes an interface of B and provides an interface of A that is indistinguishable from an implementation for A. The simulator is a blackbox that when given access to a backbox implementation of B, gives an implementation of A. This intuition seems to make sense if the goal is to use an implementation of B to implement A.

It is tempting to start formalizing the above intuition; however, there are several subtle issues that need to be resolved first.

As can be easily seen, for any two schemes A and B, a trivial simulator exists. The simulator implements all the functionalities of A by itself, without interacting with the implementation of B. Clearly, one would like to rule out these trivial simulators. A natural way to do so is to restrict the amount of space used by the simulator to be sub-linear in the size of the state of the scheme it is simulating. It seems to be a reasonable requirement that the simulator takes constant space on its own, i.e., the space used by the simulator does not depend on the size of the state. (The space used by the implementation of B is not considered here.)

Another issue is whether to further restrict a simulator's internal behavior. When the simulator receives a query in the scheme A, it may issue multiple queries to the blackbox implementation of B before answering the query; it may even perform some state-transition on B before answering the query. Similarly, the simulator may perform multiple queries and state-transitions on B to simulate one state-transition in A.

If no restriction is placed, then the notion of simulation is too weak to separate different access control models. For example, in [13], Munawer and Sandhu constructed a simulation of ATAM in RBAC. In [20], we give a simulation of RBAC in strict DAC, a discretionary model that allows only the owner of an object to grant rights over the object to another subject and disallows the tranfer of ownership. According to these results, the simplest DAC (in which security analysis is efficiently decidable) has the same expressive power as ATAM (in which safety analysis is undecidable). This illustrates that without precise requirements, simulation is not a useful concept for comparing access control models.

If one places restrictions on the simulator, then the question is what restrictions are reasonable. Our conclusion is that it is difficult to justify such restrictions. In the following, we elaborate on this.

A possibility is to restrict the internal behavior of the simulator, e.g., to restrict it to issue only one query to B in order to answer one query in A and to make a bounded number of state-transitions in B to simulate one state-transition in A. Under these restrictions, one can prove that RBAC cannot be simulated in the HRU model. The assignment of a user to a role in RBAC results in the user gaining all the accesses to objects implied by the permissions associated with that role; therefore, it changes the answers to an unbounded number of queries (queries involving those permissions.) One may argue that the assignment of a user to a role is a single "action" in RBAC, and therefore, the acquiring of those permissions by that user is accomplished in a single "action." The corresponding assignment of rights in the HRU access matrix cannot be accomplished by a single command or a bounded number of commands, as each command changes only a bounded number of cells in the matrix. Thus, any mapping of the user-assignment in RBAC involves an unbounded number of commands being executed in HRU. Nonetheless, one can argue that this is balanced by

the efficiency of checking whether a user has a particular right in the two models. A naive implementation of an RBAC model may involve collecting all roles to which that user is assigned, then collecting all permissions associated with those roles, and then checking whether one of those permissions corresponds to the object and access right for which we are checking. The time this process takes depends on the size of the current state and is unbounded. The corresponding check in HRU is simpler: we simply check whether the corresponding access right exists in the cell in the matrix. Thus, we can argue that there is a trade-off between time-to-update, and time-to-check-access between the two schemes.

Another possibility is to measure how much time the simulator takes to perform a state-transition and to answer one query in the worst case and require that there cannot be a significant slowdown. This possibility is complicated by the fact that the efficiency of these operations are not predetermined in any access control scheme, the implementation can make trade-offs between time complexity and space complexity and between query answering and state-transitions. Any comparison must involve at least three axes, query time, state-transition time, and space. Furthermore, the best approach to implementing an access control scheme is not always known. Finally, these implementation-level details do not seem to belong in the comparison of access control models; as such models by themselves are abstract models to study properties other than efficiency.

In summary, our analysis in this section suggests that the "implementation paradigm" does not seem to yield effective definitions of simulations that are useful to compare access control models. This suggests also that expressive power results proved under this paradigm should be reexamined.

An Alternate Approach Bertino et al. [2] propose a different implementation paradigm from the one discussed above. They present a framework based on logic programming within which to compare the expressive power of access control models. A library of logic facts and rules are provided, and each access control model is a collection of some facts and rules from that library. Access control models are then compared based on what facts and rules are used to represent each of them. The approach in that work is structural: if in one model we use certain facts and rules, but not in another, then the two models are incomparable. Furthermore, if one model uses more facts and rules than another, then the former is more expressive than the latter. This basis is used in arguing that RBAC is more expressive than MAC as RBAC has the notion of roles. State-transitions are not considered in this approach, and the preservation of properties across state-transitions is not part of the bases for comparison. Our theory for comparing the expressive power of access control models is based on whether schemes from one model can represent policies that schemes from another cannot. We do not have any structural restrictions in comparing two models. Thereby, our work is fundamentally different from the work by Bertino et al. [2].

4. APPLYING THE THEORY

In this section, we apply our theory from Section 2 to compare the expressive power of different access control schemes. We examine two particular results from literature using our theory: (1) that RBAC is at least as expressive as DAC (Sections 4.1 and 4.2), and (2) that TAM is at least as expressive as ATAM (Section 4.4). We show also that the trust management language $\mathsf{RT}[\cap]$ is at least as expressive as an RBAC scheme (Section 4.3). Precise characterizations of our schemes and proofs are in [20].

4.1 Examining comparisons of RBAC and DAC

Munawer and Sandhu [13] present a simulation of ATAM in RBAC and conclude that RBAC is at least as expressive as ATAM. Osborn et al. [15, 14, 19] give simulations of various MAC and DAC schemes in RBAC. The main conclusion of Osborn et al. [15, 14, 19] is that as MAC and DAC can be simulated in RBAC, a Trusted Computing Based (TCB) needs to include an implementation of RBAC only, and DAC and MAC policies can be successfully represented and enforced by the TCB.

In the simulations used in [13, 15, 14, 19], the preservation of safety (or other security) properties is not identified as an objective. From the above conclusion in [15, 14, 19], it seems that they follow the implementation paradigm. As discussed in Section 3, this paradigm leads to a weak notion of simulations, as exemplified by the simulation of RBAC in strict DAC in [20].

We observe also that the problem of comparing RBAC with DAC as stated by Osborn et al. [15, 19] is ill-defined (or at least not clearly defined). RBAC by itself only specifies the structures to store access control information, but not how to manipulate these structures, which are specified by administrative models. In other words, only the set Γ of states is precisely defined, the set Ψ of state-transition rules is not. The counterpart of RBAC is the access matrix model, and not DAC or MAC. In DAC, we specify that access control information is stored in an access matrix, and we specify also rules on how the access matrix may change. The statement that RBAC is at least as expressive as DAC (or MAC) is similar to saying that the access matrix model is at least as expressive as DAC or MAC. Comparing the RBAC model with the access matrix model is not fruitful either, as both models can include arbitrary state-transition rules.

4.2 Comparing ARBAC97 with a form of DAC

To compare any RBAC-based model with DAC, one needs to specify the administrative model (state-transition rules) for RBAC. In existing comparisons of RBAC and DAC [13, 15, 19], new and rather complicated administrative models are introduced "on the fly" to simulate the effects in DAC. In this section, we compare the expressive power of RBAC with ARBAC97 [17] as the administrative model to that of SDCO, a rather simple form of DAC. Precise characterizations of SDCO and the ARBAC97 scheme are in [20]. Osborn et al. [15] assert that SDCO can be simulated in RBAC. We assert that there does not exist a state-matching reduction from SDCO to the ARBAC97 scheme, given a natural query set for each scheme.

This result is significant as it shows that we cannot assert that RBAC is more expressive than DAC without qualifying the assertion; a strongly security-preserving mapping does not exist from SDCO to ARBAC97. Our conclusion provides the first evidence that the expressive power of RBAC (or at least some reasonable incarnation of it) is limited.

THEOREM 3. *There exists a reduction from SDCO to the AR-BAC97 scheme.*

THEOREM 4. *There exists no state-matching reduction from SDCO to the ARBAC97 scheme.*

The proofs are in [20]. One may ask whether there are other schemes based on RBAC for which there is indeed a state-matching reduction from SDCO. An approach may be to adopt a different query set for ARBAC97. We observe that for certain other query sets as well, the non-existence of a state-matching reduction holds. As an example, suppose we map the query for the presence of a

right in SDCO to a query for the absence of a permission in RBAC. In this case as well, there exists no state-matching reduction from SDCO. Whether there exists a meaningful set of state-transition rules (an administrative model) for RBAC for which there is a state-matching reduction from SDCO is an open problem.

4.3 Comparing an RBAC scheme with a Trust Management Language

In this section, we compare a particular RBAC scheme to the trust management language, $\mathsf{RT}[\cap]$. The RBAC scheme we consider is called Assignment And Revocation (AAR) [10]. In AAR, the state is an RBAC state, and state-transition rules are those from the URA97 component of ARBAC97 [17]; users may be assigned to and revoked from roles. Precise characterizations of AAR are in [10] and [20].

$\mathsf{RT}[\cap]$ is a trust management language in which a state is a set of credentials issued by the principals involved in the system. A credential denotes membership in a principal's role. A credential is one of three types: (1) A principal is asserted to be a member of another principal's role, (2) All the principals that are members of a principal's role are asserted to also be members of another principal's role, and (3) All the principals that are members of two roles (the intersection of the members of the roles) are also members of another principal's role. We refer the reader to Li et al. [9, 11, 12] for more details on $\mathsf{RT}[\cap]$.

Li and Tripunitara [10] present a form-2 weak reduction (see Definition 11) from AAR to $\mathsf{RT}[\cap]$. We assert with the following theorem that the result can be made stronger. The proof for the following theorem is in [20].

THEOREM 5. *There exists a state-matching reduction from the RBAC scheme AAR to* $\mathsf{RT}[\cap]$.

4.4 Comparing ATAM with TAM

TAM is a scheme based on the access matrix model and is similar to the HRU scheme [6] (see Section 2.1). Every object is typed, and the type cannot change once the object is created. State-transitions occur via the execution of commands that are similar to HRU commands. We specify a type for every parameter in a command. ATAM is the same as TAM, except that in ATAM, the absence of a right in a cell of the access matrix may be checked (and not just the presence of a right). See [20] for more details on the two schemes.

Sandhu and Ganta [18] present a mapping from ATAM to TAM. Based on the mapping, one may conclude that TAM is at least as expressive as ATAM. As the converse is trivially true (TAM is a special case of ATAM), one may conclude that ATAM and TAM have the same expressive power; we gain nothing from the ability to check for the absence of rights. Sandhu and Ganta [18] make the observation that the simulation of a command in ATAM may require the execution of an unbounded number of commands in TAM, and conclude with the following comment: "...practically testing for the absence of rights appears to be useful. It is an open question whether this claim can be formalized..." In this section, we formalize this claim by asserting that there is no state-matching reduction from ATAM to TAM.

THEOREM 6. *There exists no state-matching reduction from ATAM to TAM.*

The proof is in [20]. Thus, the notion of state-matching reductions formalizes the difference in expressive power between ATAM and TAM. One may ask whether there exists a reduction from ATAM to TAM. One may ask also whether reductions or state-matching

reductions exist from ATAM to TAM when we allow TAM to contain queries of the type "is $r \notin M_\gamma[s, o]$?" as well (but a command allows only checking for the presence of a right in a cell in the condition). These are open questions.

5. CONCLUSIONS AND FUTURE WORK

We have presented a theory to compare the expressive power of access control models. Our theory is based on perceiving an access control system as a state-transition system, and asking whether there exist security-preserving or strongly security-preserving mappings between two schemes. We have applied our theory in three cases and shown that: (1) RBAC with ARBAC97 as its administrative model is limited in its expressive power in comparison to a version of DAC; (2) the trust-management language $\mathsf{RT}[\cap]$ is at least as expressive as RBAC with the URA97 component of ARBAC97 as its administrative model; and (3) ATAM is more expressive than TAM. To our knowledge, (1) is the first evidence that the expressive power of RBAC is limited, and (3) solves an open problem stated in the literature [18].

As future work, we propose to use our theory to compare more models with each other. For instance, we would like to compare various versions of DAC and "layer" these versions based on their relative expressive power. Also, while our theory is based on capturing the notion of policies that can represented and verified in an access control system, we do not believe that reductions and state-matching reductions capture all types of policies we would want to consider. For instance, it is reasonable to ask a temporal query such as: "did Alice get her write access to a sensitive file only after her husband, Bob was given privileged access to the system?" Neither reductions nor state-matching reductions capture such query expressions. As part of our future work, we propose to expand our theory to include such queries.

Acknowledgements

6. REFERENCES

[1] P. Ammann, R. Lipton, and R. S. Sandhu. The expressive power of multi-parent creation in monotonic access control models. *Journal of Computer Security*, 4(2-3):149–165, Jan. 1996.

[2] E. Bertino, B. Catania, E. Ferrari, and P. Perlasca. A logical framework for reasoning about access control models. *ACM Transactions on Information and System Security*, 6(1):71–127, Feb. 2003.

[3] A. Chander, D. Dean, and J. C. Mitchell. A state-transition model of trust management and access control. In *Proceedings of the 14th IEEE Computer Security Foundations Workshop*, pages 27–43. IEEE Computer Society Press, June 2001.

[4] S. Ganta. *Expressive Power of Access Control Models Based on Propagation of Rights*. PhD thesis, George Mason University, 1996.

[5] G. S. Graham and P. J. Denning. Protection — principles and practice. In *Proceedings of the AFIPS Spring Joint Computer Conference*, volume 40, pages 417–429. AFIPS Press, May 16–18 1972.

[6] M. A. Harrison, W. L. Ruzzo, and J. D. Ullman. Protection in operating systems. *Communications of the ACM*, 19(8):461–471, Aug. 1976.

[7] B. W. Lampson. Protection. In *Proceedings of the 5th Princeton Conference on Information Sciences and Systems*, 1971. Reprinted in ACM Operating Systems Review, 8(1):18-24, Jan 1974.

[8] N. Li and J. C. Mitchell. RT: A role-based trust-management framework. In *The Third DARPA Information Survivability Conference and Exposition (DISCEX III)*. IEEE Computer Society Press, Apr. 2003.

[9] N. Li, J. C. Mitchell, and W. H. Winsborough. Design of a role-based trust management framework. In *Proceedings of the 2002 IEEE Symposium on Security and Privacy*, pages 114–130. IEEE Computer Society Press, May 2002.

[10] N. Li and M. V. Tripunitara. Security analysis in role-based access control. In *Proceedings of the Ninth ACM Symposium on Access Control Models and Technologies (SACMAT 2004)*, June 2004.

[11] N. Li, W. H. Winsborough, and J. C. Mitchell. Beyond proof-of-compliance: Safety and availability analysis in trust management. In *Proceedings of IEEE Symposium on Security and Privacy*, pages 123–139. IEEE Computer Society Press, May 2003.

[12] N. Li, W. H. Winsborough, and J. C. Mitchell. Distributed credential chain discovery in trust management. *Journal of Computer Security*, 11(1):35–86, Feb. 2003.

[13] Q. Munawer and R. S. Sandhu. Simulation of the augmented typed access matrix model (ATAM) using roles. In *Proceedings of INFOSECU99 International Conference on Information and Security*, 1999.

[14] S. Osborn. Mandatory access control and role-based access control revisited. In *Proceedings of the Second ACM Workshop on Role-Based Access Control (RBAC'97)*, pages 31–40, Nov. 1997.

[15] S. Osborn, R. S. Sandhu, and Q. Munawer. Configuring role-based access control to enforce mandatory and discretionary access control policies. *ACM Transactions on Information and System Security*, 3(2):85–106, May 2000.

[16] R. S. Sandhu. Expressive power of the schematic protection model. *Journal of Computer Security*, 1(1):59–98, 1992.

[17] R. S. Sandhu, V. Bhamidipati, and Q. Munawer. The ARBAC97 model for role-based aministration of roles. *ACM Transactions on Information and Systems Security*, 2(1):105–135, Feb. 1999.

[18] R. S. Sandhu and S. Ganta. On testing for absence of rights in access control models. In *Proceedings of the sixth Computer Security Foundations Workshop*, pages 109–118. IEEE Computer Society Press, June 1993.

[19] R. S. Sandhu and Q. Munawer. How to do discretionary access control using roles. In *Proceedings of the Third ACM Workshop on Role-Based Access Control (RBAC 1998)*, pages 47–54, Oct. 1998.

[20] M. V. Tripunitara and N. Li. Comparing the expressive power of access control models. Technical Report CERIAS-TR-2004-10, Center for Education and Research in Information Assurance and Security, Purdue University, May 2004.