

A theory for comparing the expressive power of access control models¹

Mahesh V. Tripunitara^{a,*} and Ninghui Li^b

^a Motorola Labs, Schaumburg, IL 60196, USA

E-mail: tripunit@motorola.com

^b Department of Computer Science and CERIAS, Purdue University, West Lafayette, IN 47907, USA

E-mail: ninghui@cs.purdue.edu

We present a theory for comparing the expressive power of access control models. The theory is based on simulations that preserve security properties. We perceive access control systems as state-transition systems and present two kinds of simulations, reductions and state-matching reductions. In applying the theory, we highlight four new results and discuss these results in the context of other results that can be inferred or are known. One result indicates that the access matrix scheme due to Harrison, Ruzzo and Ullman is limited in its expressive power when compared with a trust-management scheme, thereby formally establishing a conjecture from the literature. A second result is that a particular RBAC (Role-Based Access Control) scheme, ARBAC97, may be limited in its expressive power, thereby countering claims in the literature that RBAC is more expressive than DAC (Discretionary Access Control). A third result demonstrates that the ability to check for the absence of rights (in addition to the presence of rights) can cause a scheme to be more expressive. A fourth result is that a trust-management scheme is at least as expressive as RBAC with a particular administrative scheme (the URA97 component of ARBAC97).

Keywords: Access control, expressive power, reduction, state-matching reduction, access matrix, trust management, role-based access control, discretionary access control

1. Introduction

An access control system enforces a policy on who may access a resource in a certain manner (e.g., “Alice may read the file, f ”). The *protection state* (or simply, state) of the system represents all the accesses that are allowed at a given time. 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 the file, f ”). Thus, 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

¹A preliminary version of this paper appears in the proceedings of the 2004 ACM Conference on Computer and Communications Security (CCS) [23].

*Corresponding author. Address: Motorola, 1301 E Algonquin Road, IL02–2712, Schaumburg, IL 60196, USA. Tel.: +1 847 576 7883.

have read access to the 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 model* is generally associated with how the state is represented. An example of an access control model is the access matrix model [5–7], in which a state is represented by a matrix in which each cell, indexed by a *(subject, object)* pair, contains a set of rights. Formally, an access control model is a set of *access control schemes*; a scheme specifies a set of states, and a set of state-transition rules. An example of a scheme based on the access matrix model is the HRU scheme [6] for which a state is an access matrix, and a state-transition rule is a set of commands, each of which is of a particular form. An *access control system* is an instance of an access control scheme. 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 computer security and is studied extensively in the literature [1,3,4,17,19,21,22]. 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 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 three axes along which these definitions differ.

- The first axis is whether the simulation maps only the state, or also the state-change rule. The approach of Bertino et al. [2] is to map only the states of two access control models to a common language based on mathematical logic, and to compare the results to determine whether one model is at least as expressive as the other, or whether the two models are incomparable. Other work, such as [1,3,4,19,21] however, require both the state and the state-change rule to be mapped under the simulation.

An advantage with an approach such as the one that is adopted by Bertino et al. [2] is that it captures “structural” differences in how the protection state is represented in a system based on an access control model. For instance, it is observed in [2] that the existence of an indirection (the notion of a role) between

users and permissions in RBAC gives it more expressive power than an access matrix model. Such “structural” differences are not captured by our theory, or other approaches that consider both the state and the state-change rule.

We point out, however, that the state-change rule is an important component of an access control system, and therefore assert that a meaningful theory for expressive power must consider it as well. In fact, it is often the case that it is the state-change rule that endows considerable power to an access control system. Consider, for example, the access matrix schemes proposed by Graham and Denning [5] and by Harrison et al. [6]. In both schemes, the state is represented by an access matrix. However, the state-change rules are quite different: in the Graham-Denning scheme [5], there are only specific ways in which rights may be transferred, while in the HRU scheme [6], one may define arbitrary commands in a state-change rule. It has also been demonstrated [11] that safety is decidable in polynomial time in the Graham-Denning scheme, while it is known to be undecidable [6] in the HRU scheme. Such differences cannot be captured by an approach that does not consider both the state and the state-change rule.

- The second axis is whether a simulation is required to preserve safety properties. In the comparison of different schemes based on the access matrix model [1,4,19,21], 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) [15,17,22]. 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 third 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 is 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 [21]. They use a simulation that involves an unbounded num-

ber 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 [15,17] is qualitatively different from the result that TAM is at least as expressive as ATAM [21], 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. [15,17,22] 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 this paper that using the same notion of simulations, RBAC can be simulated in strict DAC, one of the most basic forms of DAC where simple safety is trivially decidable. This suggests that using such a notion of simulations, it is likely that one can show that almost 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. Neither choice is justified by Ganta [4] and Chander et al. [3].

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 the 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.

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. [10], 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 highlight four applications of our theory. We show that:
 - there exists no state-matching reduction from a rather simple trust-management scheme, $RT[\]$ [10], to the HRU scheme [6]. To our knowledge, this is the first formal evidence of the limited expressive power of the HRU scheme. Contrary to the undecidability result of safety analysis in the HRU scheme, Li et al. [10] show that safety analysis and more sophisticated security analysis in the trust management scheme, $RT[\leftarrow, \cap]$, is decidable. Li et al. [10] conjecture that these schemes cannot be encoded in the HRU scheme. In this paper, we present a formal proof for this.
 - The $RT[\]$ scheme is certainly not as expressive as the HRU scheme; we conclude that the two schemes are incomparable in expressive power with respect to state-matching reductions.
 - there exists no state-matching reduction from a rather simple DAC scheme, Strict DAC with Change of Ownership (SDCO), to RBAC with ARBAC97 [20] as the administrative model. Osborn et al. [17] and Sandhu and Munawar [22] have argued that RBAC is more expressive than various forms of DAC, including SDCO. To our knowledge, this is the first evidence of the possible limited expressive power of an RBAC scheme in comparison to DAC.

However, we show that a reduction does exist from SDCO to the ARBAC97 scheme. Consequently, whether the ARBAC97 scheme is as expressive as SDCO or not depends on the kind of reduction we consider to be appropriate. Also, this indicates that it may be possible to extend the ARBAC97 scheme so that it is indeed as expressive as SDCO with respect to state-matching reductions.

- there exists a state-matching reduction from RBAC with an administrative scheme that is a component of ARBAC97 [20] to $\text{RT}[\cap]$ [8,9], a trust-management scheme. This shows that state-matching reductions can be constructed for powerful access control schemes in the literature.
- there exists no state-matching reduction from ATAM to TAM, when we permit queries in ATAM that check for both the absence and the presence of a right in a cell. This revisits the issue addressed by Sandhu and Ganta [21] and formalizes the benefit from the ability to check for the absence of rights in addition to the ability to check for the presence of rights.

The remainder 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, we apply our theory to compare the expressive power of schemes in four cases. We summarize these and other known results in Section 4.6. We discuss future work and conclude with Section 5. Appendix 5 presents a “simulation” of RBAC in strict DAC.

2. Comparisons based on security analysis

A requirement used in the literature for simulations is the preservation of simple safety properties. Indeed, this is the only requirement on simulations in [1,19,21]. 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 simple safety analysis² is preserved by the simulation.

Simple 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.

²What we call simple safety analysis is called safety analysis in the literature. In [10], more general notions of safety analysis, for which the traditional safety analysis is just a special case, were introduced. Here we follow the terminology in [10].

```

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. Harrison et al. [6] 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 simple safety properties as the sole requirement of simulations is based on the implicit assumption that simple safety is the *only* interesting property in access control schemes, an assumption that is not valid. When originally introduced by Harrison et al. [6], simple safety was described as just one class of queries one can consider. More recently, Li et al. [10] have introduced the notion of security analysis, which generalizes simple safety to other properties such as 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 Schemes). 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 $\gamma \not\vdash q$ 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 $\gamma \mapsto_\psi^* \gamma_1$ if a sequence of zero or more allowed changes leads from γ to γ_1 . In other words, $\mapsto_\psi^*$ is the reflexive and transitive closure of $\mapsto_\psi$. If $\gamma \mapsto_\psi^* \gamma_1$, 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 the system is in and $\psi \in \Psi$ 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 also appears in [12], 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; 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 further assume that $\mathcal{S} \subseteq \mathcal{O}$, i.e., all subjects are also objects. In the HRU scheme:

- Γ is the set of all possible access matrices. Formally, each $\gamma \in \Gamma$ is identified by three 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 that are in the cell.
- Q is the set of all queries having the form: $r \in [s, o]$, where $r \in \mathcal{R}$ is a right, $s \in \mathcal{S}$ is a subject, $o \in \mathcal{O}$ 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 gets γ_1 .

The set of queries is not explicitly specified by Harrison et al. [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 $\gamma \xrightarrow{*}_{\psi} \gamma_1$ 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 $\gamma \xrightarrow{*}_{\psi} \gamma_1$, $\gamma_1 \vdash q$? 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 $\langle \gamma, \varphi, \psi, \Pi \rangle$, where γ , ψ , and Π are the same as in a security analysis instance, and φ is a propositional formula over Q , i.e., φ is constructed from queries in Q using propositional logic connectives such as $\wedge$, $\vee$, $\neg$.

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 allow the formula φ to have infinite size. For example, suppose that $\mathcal{S}$, the set of all subjects, is $\{s_1, s_2, s_3, s_4, \dots\}$, then the formula $\neg(r \in [s_2, o] \vee r \in [s_3, o] \vee r \in [s_4, o] \vee \dots)$ is true when no subject other than s_1 has the right r over object o .

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 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, $\sigma : (\Gamma^A \times \Psi^A) \cup Q^A \rightarrow (\Gamma^B \times \Psi^B) \cup Q^B$.

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 $\sigma : (\Gamma^A \times \Psi^A) \cup Q^A \rightarrow (\Gamma^B \times \Psi^B) \cup Q^B$, 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, $\langle \gamma^A, \varphi^A, \psi^A, \Pi \rangle$, in A is $\langle \gamma^B, \varphi^B, \psi^B, \Pi \rangle$, 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 $\varphi^B = \sigma(\varphi^A)$. 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 and strongly 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 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 , $\sigma : (\Gamma^A \times \Psi^A) \cup Q^A \rightarrow (\Gamma^B \times \Psi^B) \cup Q^B$, 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 $\gamma^A \xrightarrow{*}_{\psi} \gamma_1^A$, there exists a state γ_1^B such that $\gamma^B \xrightarrow{*}_{\psi^B} \gamma_1^B$ and γ_1^A and γ_1^B are equivalent under σ .

2. For every state γ_1^B in scheme B such that $\gamma_1^B \xrightarrow{\psi^B}^* \gamma_1^B$, there exists a state γ_1^A such that $\gamma_1^A \xrightarrow{\psi^A}^* \gamma_1^A$ 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.

The above definition may appear similar to the well-known notion of a bisimulation [14,18]. However, there are important differences between the two. In bisimulation, a state-change in one system must correspond to a state-change in the other system. In our context, this can be seen as more of a “lock-step” approach. The above definition for state-matching reductions is less restrictive. It requires only that for every finite sequence of state-changes in one system, there exists a finite sequence of state-changes in the other system. Also, in a bisimulation, when we talk of corresponding reachable states, we require the labels on the state-changes that lead to the states to be the same. State-matching reductions impose no such restriction; the state-changes that lead to corresponding states (states that answer corresponding queries) are not related to each other in any a-priori manner.

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 $\langle \gamma^A, \varphi^A, \psi^A, \Pi \rangle$ in scheme A , let $\langle \gamma^B, \psi^B \rangle = \sigma(\langle \gamma^A, \psi^A \rangle)$ and $\varphi^B = \sigma(\varphi^A)$, we show that $\langle \gamma^A, \varphi^A, \psi^A, \Pi \rangle$ is true if and only if $\langle \gamma^B, \varphi^B, \psi^B, \Pi \rangle$ is true.

First consider the case that the instance $\langle \gamma^A, \varphi^A, \psi^A, \Pi \rangle$ is existential, i.e., Π is $\exists$. If the instance is true, i.e., 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 , $\langle \gamma^B, \varphi^B, \psi^B, \exists \rangle$, is also true. On the other hand, if $\langle \gamma^B, \varphi^B, \psi^B, \exists \rangle$ 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 $\langle \gamma^A, \varphi^A, \psi^A, \Pi \rangle$ is universal, i.e., Π is $\forall$. If the instance is false, i.e., 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 $\sigma(\varphi^A)$ 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 $\gamma_1^B \vdash^B \neg \sigma(q^A)$. Clearly, φ^A is false in all states reachable from γ^A , but $\sigma(\varphi^A)$ 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. $\square$

Note that the proof uses a compositional analysis instance that contains a potentially infinite-length formula. If one chooses to restrict the formulas in analysis instances to be finite length, then state-matching reduction may not be necessary for being strongly security-preserving. Also, a state-matching reduction preserves compositional security properties. If we only need 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; it is 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 $\gamma^A \xrightarrow{*}_{\psi} \gamma_1^A$, then in scheme B there exists a state γ_1^B such that $\gamma^B \xrightarrow{*}_{\psi^B} \gamma_1^B$ 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 $\gamma^B \xrightarrow{*}_{\psi^B} \gamma_1^B$, there exists a state γ_1^A such that $\gamma^A \xrightarrow{*}_{\psi} \gamma_1^A$ 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, i.e., 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, i.e., 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_1^B the answer for the query $\sigma(q^A)$ under the state is different from the answer for q^A under 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 $\gamma_1^A \not\vdash^A q^A$, 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.

Then 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 under the state is different from the answer for $\sigma(q^A)$ under 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 $\gamma_1^B \not\vdash^B \sigma(q^A)$, 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 in those models.

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 $\mathcal{M}$ based on state-matching reductions (or $\mathcal{M}'$ has at least as much expressive power based on state-matching reductions as $\mathcal{M}'$) if for every scheme in $\mathcal{M}$ there exists a state-matching reduction from it to a scheme in $\mathcal{M}'$. In addition, if for every scheme in $\mathcal{M}'$, there exists a state-matching reduction from it to a scheme in $\mathcal{M}$, then we say that $\mathcal{M}$ and $\mathcal{M}'$ are equivalent in expressive power based on state-matching reductions. If $\mathcal{M}'$ is at least as expressive as $\mathcal{M}$, and there exists a scheme A in $\mathcal{M}'$ such that for any scheme B in $\mathcal{M}$, no state-matching reduction from A to B exists, we say that $\mathcal{M}'$ is strictly more expressive than $\mathcal{M}$ based on state-matching reductions. We have a similar definition for expressive power based on reductions.

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 only that scheme.

2.3. Discussions of alternative definitions for reduction

In this section, we discuss alternative definitions that differ slightly from the ones discussed in the previous section. The first of these definitions is used by Sandhu et al. [19,21] 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 $\gamma_1^A \xrightarrow{\psi^A}^* q^A$ and $\gamma_1^A \vdash^A q^A$, then there exists a state γ_1^B such that $\gamma_1^B \xrightarrow{\psi^B}^* \gamma_1^B$ 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 $\gamma_1^B \xrightarrow{\psi^B}^* \gamma_1^B$ and $\gamma_1^B \vdash^B \sigma(q^A)$, then there exists a state γ_1^A such that $\gamma_1^A \xrightarrow{\psi^A}^* \gamma_1^A$ 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 [19] 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 [19,21] as only simple 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. One 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 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 will also lead 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 $\gamma^A \xrightarrow{*}_{\psi} \gamma_1^A$, there exists a state γ_1^B such that $\gamma^B \xrightarrow{*}_{\psi^B} \gamma_1^B$ 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 $\gamma^B \xrightarrow{*}_{\psi^B} \gamma_1^B$, there exists a state γ_1^A such that $\gamma^A \xrightarrow{*}_{\psi} \gamma_1^A$ 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 by Li and Tripunitara [12] for reductions from security analysis in two RBAC schemes to that in the RT Role-based Trust-management framework [9,10]. As we state in Theorem 7 in this paper, a form-2 weak reduction used in [12] 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. Access control schemes have traditionally been specified only as the double $\langle \Gamma, \Psi \rangle$ (sets of states and state-change rules). Queries are usually introduced subsequently as part of an analysis instance. An example is the work by Harrison et al. [6], in which the query that corresponds to safety analysis is introduced as part of the analysis problem and not as part of the scheme itself. Consequently, disassociating the queries in the mapping from states and state-change rules gives us some flexibility. However, we recognize that when access control schemes have been introduced in the literature, a set of queries has been assumed, sometimes implicitly. Therefore, in our paper, we consider the set of queries to be part of a scheme.

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 $\sigma : (\Gamma^A \times \Psi^A) \cup Q^A \rightarrow (\Gamma^B \times \Psi^B) \cup Q^B$, 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 $\gamma^A \xrightarrow{\psi}^* \gamma_1^A$, then in scheme B there exists a state γ_1^B such that $\gamma^B \xrightarrow{\psi^B}^* \gamma_1^B$ 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 $\gamma^B \xrightarrow{\psi^B}^* \gamma_1^B$, then either (a) there exists a state γ_1^A such that $\gamma^A \xrightarrow{\psi}^* \gamma_1^A$ 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 $\gamma_1^B \xrightarrow{\psi^B}^* \gamma_2^B$ and a state γ_1^A such that $\gamma^A \xrightarrow{\psi}^* \gamma_1^A$, 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 preserving simple safety properties in 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. As with the reduction from Definition 10, this form of reduction cannot be used for universal queries.

3. The implementation paradigm for simulation: An examination

Several authors use the implementation paradigm for simulations, e.g., Osborn et al. [17] 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 [15,17,22], 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 this does not lead to a notion of simulations that is 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 the 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. An implementation of an access control scheme thus 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 receives 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 . In other words, the simulator is a blackbox that when given access to a blackbox 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. One 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, Munawer and Sandhu [15] constructed a simulation of ATAM in RBAC. In Appendix 5, 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 ownership cannot be transferred. According to these results, the simplest DAC (in which security analysis is efficiently decidable) has the same expressive power as ATAM (in which simple safety analysis is undecidable). This illustrates the point that, without precise requirements, simulation is not a very 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 very difficult to justify such requirements. In the following, we elaborate on this.

One possibility that we now argue to be inadequate 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 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 command for that matter, as each command only changes 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 having to collect all roles to which that user is assigned, and 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. Therefore, we argue that it does not make sense to restrict the number of steps involved in the simulation.

Another possibility that we now argue to be inadequate 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 ways to implement 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, when no restriction is placed on the simulations, the “implementation paradigm” does not separate different access control schemes. On the other hand, it seems difficult to justify the restrictions that have been considered in the literature. Therefore, 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 also suggests that expressive power results proved under this paradigm should be reexamined.

4. Applying the theory

In this section, we apply our theory from Section 2 to compare the expressive power of different access control schemes. In the following section, we show that the HRU access matrix scheme is not as expressive as a relatively simple trust management scheme, $\text{RT}[\]$. We then examine two particular results from literature using our theory: (1) that RBAC is at least as expressive as DAC (Sections 4.2 and 4.3), and (2) that TAM is at least as expressive as ATAM (Section 4.5), and in each case, assert the opposite. We show also that the trust management scheme $\text{RT}[\]$ is at least as expressive as an RBAC scheme (Section 4.4).

Proof Methodology. In this section, we prove the existence of reductions and state-matching reductions as well as the nonexistence of state-matching reductions. To prove that there exists a reduction or state-matching reduction from a scheme A to a scheme B , we constructively give a mapping and show that the mapping satisfies the requirements. To prove that there does not exist a state-matching reduction from a scheme A to a scheme B is more difficult, as we have to show that no mapping satisfies the requirements for a state-matching reduction. Our strategy is to use proof by contradiction. We find in scheme A a state γ^A , a state-transition rule ψ^A , as well as a state γ_1^A that is reachable. Suppose, for the sake of contradiction, that a state-matching reduction exists, then there exist states γ^B and γ_1^B such that γ^B is equivalent to γ^A , γ_1^B is equivalent to γ_1^A , and γ_1^B is reachable from γ^B . We show that among the sequence of states leading from γ^B and γ_1^B , there exists one for which there is no matching state that is reachable in A .

4.1. Comparing the HRU scheme to a trust management scheme

The HRU scheme [6] is based on the access matrix model, and has generally been believed to have considerable expressive power, partly because it has been shown that one can simulate a Turing Machine in the HRU scheme. In this section, we show that there does not exist a state-matching reduction from a relatively simple trust management scheme, $RT[]$ [10], to the HRU scheme. That $RT[]$ cannot be encoded in the HRU scheme is informally discussed and conjectured by Li et al. [10]. Using the theory presented in Section 2, we are able to formally prove this. To our knowledge, this is the first formal evidence of the limited expressive power of the HRU scheme.

As safety analysis is efficiently decidable in $RT[]$ but undecidable in the HRU scheme, there does not exist a state-matching reduction from the HRU scheme to the $RT[]$ scheme either. This shows that the expressive powers of the HRU scheme and of $RT[]$ are incomparable.

The fact that the HRU scheme can simulate Turing Machine shows that it can compute any computable function when used as a computation device. When used as an access control scheme, the HRU scheme may nonetheless be limited in expressive power. For example, it cannot encode an access control system where in one state a subject has no right over any object and in the next state the subject obtains rights over a potentially unbounded number of objects.

The HRU scheme

Γ We assume the existence of countably infinite sets of subjects, $\mathcal{S}$, objects $\mathcal{O}$ and rights $\mathcal{R}$, with $\mathcal{S} \subset \mathcal{O}$. Each state γ is characterized by $\langle S_\gamma, O_\gamma, R_\gamma, M_\gamma[] \rangle$ where $S_\gamma \subset \mathcal{S}$ is a finite set of subjects that exist in the state γ , $O_\gamma \subset \mathcal{O}$ is a finite set of objects that exist in the state γ , $R_\gamma \subset \mathcal{R}$ is a finite set of rights that exist in the state γ , and $M_\gamma[]$ is the access matrix, i.e., $M_\gamma[s, o] \subseteq R_\gamma$ gives the set of rights $s \in S_\gamma$ has over $o \in O_\gamma$ in the state γ . $M_\gamma[s, o]$ is defined only when $s \in S_\gamma$ and $o \in O_\gamma$. It may appear that we allow R_γ to differ across states. The definition for state-change rules precludes this possibility.

Ψ A state-change rule, ψ , in the HRU scheme is a command schema, i.e., a set of commands. Each command takes a sequence of parameters, each of which may be instantiated by an object. Each command has also an optional condition, which is a conjunction of clauses. Each clause checks whether a right is in a particular cell of $M_\gamma[]$. Following the (optional) conditions in a command is a sequence of primitive operations. The primitive operations are one of the following: (1) create an object; (2) create a subject; (3) enter a right into a cell of the access matrix; (4) remove a right from a cell of the access matrix; (5) destroy a subject; (6) destroy an object. We refer the reader to Harrison et al. [6] for more details on the syntax of commands. A state-change is the successful execution of a command.

Q We allow queries of the following two forms: (1) $r \in M[s, o]$, and (2) $r \notin M[s, o]$. In the queries, $r \in \mathcal{R}$, $s \in \mathcal{S}$ and $o \in \mathcal{O}$. To our knowledge, these are the

only kinds of queries that have been considered in the context of the HRU scheme in the literature. In particular, these are the queries that are pertinent to the safety property [6].

$\vdash$ Let q be the query $r \in M[s, o]$. Then, given a state γ , $\gamma \vdash q$ if and only if $r \in R_\gamma$, $s \in S_\gamma$, $o \in O_\gamma$ and $r \in M_\gamma[s, o]$. Otherwise, $\gamma \not\vdash q$, or equivalently $\gamma \vdash \neg q$. Let $\hat{q}$ be the query $r \notin M[s, o]$. Then $\gamma \vdash \hat{q}$ if and only if $r \in R_\gamma$, $s \in S_\gamma$, $o \in O_\gamma$ and $r \notin M_\gamma[s, o]$. Otherwise, $\gamma \not\vdash \hat{q}$, or equivalently $\gamma \vdash \neg \hat{q}$.

Observe that one should view both $r \in M_\gamma[s, o]$ and $r \notin M_\gamma[s, o]$ as atomic queries. In particular $\neg(r \in M_\gamma[s, o])$ is not equivalent to $r \notin M_\gamma[s, o]$. It is possible that $\gamma \not\vdash r \in M_\gamma[s, o]$ and $\gamma \not\vdash r \notin M_\gamma[s, o]$; this happens when either s or o does not exist in γ . Even though it is not possible that $\gamma \vdash ((r \in M_\gamma[s, o]) \wedge (r \notin M_\gamma[s, o]))$.

The $\text{RT}[\]$ scheme

Γ We assume the existence of countably infinite sets of principals (e.g., A, B, C) and role names (e.g., r, s, t, u). A role is formed by a principal and a role name, separated by a dot (e.g., $A.r, X.u$). An $\text{RT}[\]$ state consists of statements which are assertions made by principals about membership in their roles. Two types of assertions are supported. These are simple member (e.g., $A.r \leftarrow B$) and simple inclusion (e.g., $A.r \leftarrow B.r_1$). One reads the $\leftarrow$ symbol as “includes”. The example for the first kind of statement asserts that B is a member of A ’s r role. The example for the second kind of statement asserts that every member of $B.r_1$ is a member of $A.r$. The portion of a statement that appears to the left of the $\leftarrow$ symbol is called its head, and the portion that appears to the right is called the body. We refer the reader to Li et al. [9] for more details on the syntax and semantics of $\text{RT}[\]$ statements.

Ψ A state-change rule in a system based on the $\text{RT}[\]$ scheme consists of two sets, G and S . Both consist of $\text{RT}[\]$ roles. G is the set of growth-restricted roles, i.e., if $A.r \in G$, then statements with $A.r$ at the head cannot be added in future states. S is the set of shrink-restricted roles, i.e., if $A.r \in S$, then roles with $A.r$ at the head cannot be removed in future states. We refer the reader to Li et al. [10] for more details on the two sets, and the intuition behind them.

Q Li et al. [10] define three kinds of queries in $\text{RT}[\]$. (1) $\{B_1, \dots, B_n\} \supseteq A.r$ – this kind of query asks whether the role $A.r$ is bounded by the set of principals $\{B_1, \dots, B_n\}$; (2) $A.r \supseteq \{B_1, \dots, B_n\}$ – this kind of query asks whether each principal $B_1, \dots, B_n$ is a member of $A.r$; (3) $X.u \supseteq A.r$ – this kind of query asks whether the set of member of $A.r$ is included in the set of members of $X.u$.

$\vdash$ Given a state, we check if a query is entailed by first evaluating the set of members of each $\text{RT}[\]$ role in the query. This is done using credential chain discovery [13]. We then compare the two sets and check if the set to the left includes the set to the right. The first two kinds of queries are called semi-static queries as one of the sides in the query is a set of users that is independent of the state, and needs no further evaluation. We refer the reader to Li et al. [13] for more details on query-entailment in $\text{RT}[\]$.

Theorem 3. *There exists no state-matching reduction from the $\text{RT}[\]$ scheme to the HRU scheme.*

Proof. By contradiction. Assume that there exists a state-matching reduction, σ , from the $\text{RT}[\]$ scheme to the HRU scheme. We denote components of a $\text{RT}[\]$ system with the superscript R and the HRU scheme with the superscript H . We now consider a system based on the $\text{RT}[\]$ scheme. Let γ^R be the start-state in our $\text{RT}[\]$ system such that γ^R has no statements. The state-change rule in our $\text{RT}[\]$ system is $G = S = \emptyset$. We now consider the start-state in the corresponding HRU system $\sigma(\gamma^R) = \gamma^H$ and the state-change rule $\sigma(\psi^R) = \psi^H$. Let k be the number of objects in γ^H , i.e., $k = |O_{\gamma^H}|$. Let l be the maximum number of primitive operations of the form “enter right” in any of the commands in ψ^H . Let m be the maximum number of primitive operations of the form “remove right” in any of the commands in ψ^H .

Choose some $n > (k^2 + l + m) + 1$. Our choice of n is such that for any γ_1^H such that $\gamma^H \mapsto \gamma_1^H$, fewer than $n - 1$ queries that are true in γ^H (i.e., are entailed by γ^H) are false in γ_1^H (i.e., are not entailed by γ_1^H). The reason is that: (1) as γ^H has at most k objects (some or all of which may be subjects), a command may contain statements to destroy all these objects. Consequently, these statements can cause up to k^2 queries of the form $r \notin M[s, o]$ to be false in γ_1^H when they are true in γ^H ; (2) as a command in ψ^H has at most l statements to enter rights in to cells, these statements can cause up to l queries of the form $r \notin M[s, o]$ to be false in γ_1^H when they are true in γ^H ; (3) as a command in ψ^H has at most m statements to remove rights from cells, these statements can cause up to m queries of the form $r \in M[s, o]$ to be false in γ_1^H when they are true in γ^H . We emphasize that these are the only possibilities for queries to become false in a state-change from γ^H ; the number of queries that are entailed by γ^H , but not γ_1^H is fewer than $n - 1$.

Consider queries q_i^R for each integer i such that $1 \leq i \leq n$ in the $\text{RT}[\]$ system where q_i^R is of the form $\{B_i\} \sqsupseteq A.r$ for some principals $A, B_1, \dots, B_n$ and some role $A.r$. We make two observations about these queries. The first is that $\gamma^R \vdash q_1^R \wedge \dots \wedge q_n^R$. The reason is that $A.r$ is empty in γ^R and therefore is a subset of every set of the form $\{B_i\}$. The second observation is that in all states reachable from γ^R , either all queries of the form q_i^R such that $1 \leq i \leq n$ are entailed, or at most one of those queries is entailed. The reason is that for the set of users in the role $A.r$ to be a subset of $\{B_i\}$ for a particular i , it must be either empty, or contain exactly one element, B_i . Now consider the state γ_t^R such that $\gamma^R \xrightarrow{*}_{\psi} \gamma_t^R$ and $\gamma_t^R \vdash q_1^R \wedge \neg q_2^R \wedge \dots \wedge \neg q_n^R$. That is, q_1^R is true in γ_t^R , but none of the other queries of the form q_i^R is true. We use the subscript t only to demarcate the state and not as a count of the number of state-changes needed to reach it. In fact, γ_t^R can be reached from γ^R with a single state-change: we simply add the statement $A.r \leftarrow B_1$ to our $\text{RT}[\]$ system.

Now consider the corresponding states and queries in the HRU system produced as output by σ . Let $\gamma^H = \sigma(\gamma^R)$, $\gamma_t^H = \sigma(\gamma_t^R)$, and $q_i^H = \sigma(q_i^R)$ for $1 \leq i \leq n$. As we

assume that σ is a state-matching reduction, $\gamma^H \vdash q_1^H \wedge \dots \wedge q_n^H$, and there exists γ_t^H such that $\gamma^H \xrightarrow{*}_{\psi} \gamma_t^H$ and $\gamma_t^H \vdash q_1^H \wedge \neg q_2^H \wedge \dots \wedge \neg q_n^H$. Also, given the assumption that σ is a state-matching reduction, the queries $q_1^H, \dots, q_n^H$ are distinct from one another. The reason is that given any two distinct queries in the $\text{RT}[\]$ system, there exist reachable states in the $\text{RT}[\]$ system such that one of the queries is true and the other is false. Consequently, such reachable states must exist for the corresponding queries in the HRU system as well.

Consider any sequence of state-changes from γ^H to γ_t^H . Pick the first state in the sequence γ_c^H in which at least one of the queries q_i^H is false. Consider the state γ_{c-1}^H immediately preceding it. Then, $\gamma_{c-1}^H \vdash q_1^H \wedge \dots \wedge q_n^H$. Because one step of change cannot make $n - 1$ queries to go from true to false, in γ_c^H , some queries $q_1, q_2, q_3, \dots, q_n$ are false but at least 2 queries in them are true. As we argued in the previous paragraph, there cannot exist a matching state in A for γ_c^H . We now have the desired contradiction to the existence of a state-matching reduction from the $\text{RT}[\]$ scheme to the HRU scheme. $\square$

As we point out in the introduction to this section, it is easy to infer that there is no state-matching reduction from the HRU scheme to the $\text{RT}[\]$ scheme, because we know that safety is undecidable in the HRU scheme [6], but is efficiently decidable in the $\text{RT}[\]$ scheme [10]. Consequently, we conclude that the HRU and $\text{RT}[\]$ schemes are incomparable in terms of expressive power.

What underlie the non-existence of state-matching reductions in both directions are the state-transition rules in the schemes. In an $\text{RT}[\]$ system, it is possible, with a single state-change, to award a privilege or right to an unbounded number of principals. This is not possible in an HRU system, in which each state-change can award rights to only a bounded number of subjects. This is the intuition behind the non-existence of a state-matching reduction from the HRU scheme to the $\text{RT}[\]$ scheme. We point out that the first paragraph in the proof for the theorem above suggests an example of an $\text{RT}[\]$ system for which we cannot produce a corresponding HRU system. Conversely, the state-changes in $\text{RT}[\]$ cannot have preconditions; state-changes in an HRU system can have preconditions. This is the intuition behind the non-existence of a state-matching reduction from the $\text{RT}[\]$ scheme to the HRU scheme.

The class of queries is also important in this distinction. It may be possible, for example, for there to exist a state-matching reduction from $\text{RT}[\]$ to the HRU scheme if we adopt a broader class of queries in the HRU scheme. However, it is unclear what the class of queries must be. Furthermore, we must also ask whether any new kinds of queries we adopt make sense in the context of the HRU scheme.

4.2. Examining comparisons of RBAC and DAC

Munawer and Sandhu [15] present a simulation of ATAM in RBAC and conclude that RBAC is at least as expressive as ATAM. Osborn et al. [16,17,22] give simulations of various MAC and DAC schemes in RBAC. The main conclusion of Osborn

et al. [16,17,22] 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 [15–17,22], the preservation of safety (or other security) properties is not identified as an objective. From the above conclusion in [16,17,22], 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 Appendix 5.

We observe also that the problem of comparing RBAC with DAC as stated by Osborn et al. [17,22] 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, instead of DAC (or MAC). In DAC, we specify that access control information is stored in a matrix, and we also specify rules on how to change the access matrix. 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.

4.3. 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 [15,17,22], 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 [20] as the administrative model to that of SDCO, a rather simple form of DAC. We first present precise characterizations of SDCO and the ARBAC97 scheme. We then 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) may be limited.

We then show that a reduction does indeed exist from SDCO to the ARBAC97 scheme. That is, the ARBAC97 scheme captures SDCO in some limited way. Our results indicate that it may be possible to extend RBAC schemes so that they can indeed be as expressive as DAC schemes. We briefly discuss this further after we present our results; however, whether this is indeed possible and in what way are issues that are beyond the scope of this paper.

The SDCO scheme

Γ SDCO is a scheme based on the access matrix model and is a special case of the HRU scheme (see Section 2.1) and the Graham-Denning scheme [5,11]. Each state $\gamma \in \Gamma$ is $\langle S_\gamma, O_\gamma, M_\gamma[], R_\gamma \rangle$ where S_γ , O_γ and R_γ are finite, strict subsets of the countably infinite sets $\mathcal{S}$ (subjects), $\mathcal{O}$ (objects) and $\mathcal{R}$ (rights) respectively. The set of rights for the scheme is $R_\gamma = \{own, r_1, \dots, r_n\}$, where *own* is the distinguished right indicating ownership of the object. $M_\gamma[]$ is the access matrix.

Ψ The state-transition rules are the commands *createObject*, *destroyObject* and *grantOwn*, and for each $r_i \in R_\gamma - \{own\}$, a command *grant_{r_i}*.

<i>command createObject</i> (s, o)	<i>command destroyObject</i> (s, o)
<i>create object</i> o	if $own \in [s, o]$
enter own into $[s, o]$	<i>destroy</i> o
<i>command grantOwn</i> (s, s', o)	<i>command grant_{r_i}</i> (s, s', o)
if $own \in [s, o]$	if $own \in [s, o]$
enter own into $[s', o]$	enter r_i into $[s', o]$
remove own from $[s, o]$	

Q Each query is of one the following forms: (1) Is $s \in S$?; (2) Is $o \in O$?; and (3) Is $r \in M[s, o]$?

$\vdash$ The entailment relation is defined as follows for each type of query from above. In each of the following, $\gamma \in \Gamma$ is a state. (1) $\gamma \vdash s \in S$ if and only if $s \in S_\gamma$; (2) $\gamma \vdash o \in O$ if and only if $o \in O_\gamma$; (3) $\gamma \vdash r \in M[s, o]$ if and only if $r \in R_\gamma \wedge s \in S_\gamma \wedge o \in O_\gamma \wedge r \in M_\gamma[s, o]$.

The ARBAC97 scheme

Γ We assume the existence of the countably infinite sets $\mathcal{U}$ (users), $\mathcal{P}$ (permissions) and $\mathcal{R}$ (roles). An ARBAC97 state is $\langle UA, PA, RH, AR \rangle$ where UA is the user-role assignment relation that contains a pair $\langle u, r \rangle$ for every user $u \in \mathcal{U}$ that is assigned to a role $r \in \mathcal{R}$. PA is the permissions-role assignment relation that contains a pair $\langle p, r \rangle$ for every permission $p \in \mathcal{P}$ that is assigned to the role $r \in \mathcal{R}$. RH is the role-hierarchy, and for $r_1, r_2 \in \mathcal{R}$, $r_1 \succeq r_2 \in RH$ means that all users that are members of r_1 are also members of r_2 , and all permissions that are assigned to r_2 are authorized to users that are members of r_1 . $AR \subset \mathcal{R}$ is a set of administrative roles. In ARBAC97 [20], changes to AR may be made only by a central System Security Officer (SSO) who is trusted not to leave the system in an undesirable state; if the SSO effects a state-transition, then she does security analysis to ensure that the resulting state is acceptable. Therefore, in our analysis, we assume that AR does not change.

Ψ State-transitions in the ARBAC97 scheme are predicated on the relations that are part of the *URA97* (user-roles assignment), *PRA97* (permission-role assignment) and *RRA97* (role-role assignment) components. We introduce the notion of a role range that is used in the definition of the state-transitions. A role range, ξ is written

as (r_1, r_2) , where r_1 and r_2 are roles, and every role r that satisfies $r_1 \succeq r \wedge r \succeq r_2 \wedge r \neq r_1 \wedge r \neq r_2$ is in the role range ξ . We write $r \in \xi$ when r is in the role range ξ . We represent as Ξ the set of all role ranges. Role ranges in ARBAC97 satisfy some other properties, and we refer the reader to Sandhu et al. [20] for those. Those properties are not relevant to our discussion here.

$$URA97 \left\{ \begin{array}{l} can_assign \subseteq AR \times CR \times \Xi \\ can_revoke \subseteq AR \times \Xi \end{array} \right. \quad PRA97 \left\{ \begin{array}{l} can_assignp \subseteq AR \times CR \times \Xi \\ can_revokep \subseteq AR \times \Xi \end{array} \right.$$

$$RRA97 \left\{ can_modify \subseteq AR \times \Xi \right.$$

CR is a set of pre-requisite conditions. A pre-requisite condition is a propositional logic formula over regular roles. For instance, $c = r_1 \wedge \overline{r_2}$ is a pre-requisite condition that indicates: “role r_1 and not role r_2 ”, where $r_1, r_2 \in R$.

We postulate that a state-transition is the successful execution one of the following operations.

<i>assignUser</i>(a, u, r) if $\exists \langle ar, c, \xi \rangle \in can_assign$ such that a is a member of $ar \wedge u$ satisfies $c \wedge$ $r \in \xi$ then add $\langle u, r \rangle$ to UA	<i>revokeUser</i>(a, u, r) if $\exists \langle ar, \xi \rangle \in can_revoke$ such that a is a member of $ar \wedge$ $r \in \xi$ then remove $\langle u, r \rangle$ from UA
<i>assignPermission</i>(a, p, r) if $\exists \langle ar, c, \xi \rangle \in can_assignp$ such that a is a member of $ar \wedge p$ satisfies $c \wedge$ $r \in \xi$ then add $\langle p, r \rangle$ to PA	<i>revokePermission</i>(a, p, r) if $\exists \langle ar, \xi \rangle \in can_revokep$ such that a is a member of $ar \wedge$ $r \in \xi$ then remove $\langle p, r \rangle$ from PA
<i>addToRange</i>(a, ξ, r) if $\exists \langle ar, \xi \rangle \in can_modify$ such that a is a member of ar then add $r_1 \succeq r$ to RH add $r \succeq r_2$ to RH where $\xi = (r_1, r_2) \wedge r \neq r_1 \wedge r \neq r_2$	<i>removeFromRange</i>(a, ξ, r) if $\exists \langle ar, \xi \rangle \in can_modify$ such that a is a member of ar then remove $r_1 \succeq r$ from RH remove $r \succeq r_2$ from RH where $\xi = (r_1, r_2) \wedge r \neq r_1$ $\wedge r \neq r_2$
<i>addAsSenior</i>(a, r, s) if $\exists \langle ar, \xi \rangle \in can_modify$ such that a is a member of $ar \wedge r, s \in \xi$ then add $r \succeq s$ to RH	<i>removeAsSenior</i>(a, r, s) if $\exists \langle ar, \xi \rangle \in can_modify$ such that a is a member of $ar \wedge r, s \in \xi$ then remove $r \succeq s$ from RH

$Q, \vdash$ We allow queries of the following forms that are all natural for the ARBAC97 scheme: (1) given a role r , does there exist a user u such that $\langle u, r \rangle \in UA?$, (2) given user u , does there exist a role r such that $\langle u, r \rangle \in UA?$, (3) given user u and role r , is $\langle u, r \rangle \in UA?$, (4) given a permission p , does there exist a role r such that $\langle p, r \rangle \in PA?$, (5) given permission p , does there exist a role r such that $\langle p, r \rangle \in PA?$, (6) given permission p and role r , is $\langle p, r \rangle \in PA?$, (7) given roles r_1, r_2 , is $r_1 \succeq r_2 \in RH?$, and (8) give user u and permission p , is u authorized to have the permission p ? That is, do there exist roles r_1, r_2 such that $\langle u, r_1 \rangle \in UA \wedge \langle p, r_2 \rangle \in PA \wedge r_1 \succeq r_2 \in RH?$ The entailment relation, $\vdash$ is based simply on whether the conditions checked in a query hold in the given state.

Before we introduce Theorem 5, we introduce the following lemma as an intermediate result on the state-change rules in ARBAC97. The intermediate result aids in the proof of the theorem. It also provides some intuition as to why there exists no state-matching reduction from SDCO to the ARBAC97 scheme.

Lemma 4. *Let ψ be a state-transition rule, and γ and γ' be states in the ARBAC97 scheme. Then, for any two queries q_1 and q_2 , there exists no γ' such that $\gamma' \vdash t(\neg q_1 \wedge q_2)$ when $\gamma \vdash (q_1 \wedge \neg q_2)$ and $\gamma \mapsto \gamma'$.*

Proof. We observe that the operations *assignUser*, *assignPermission*, *addToRange* and *addAsSenior* can cause queries to become only true, and not false. Similarly, the operations *revokeUser*, *revokePermission*, *removeFromRange* and *removeAsSenior* cannot cause a query to become true. Therefore, given a state-transition in the ARBAC97 scheme, it cannot cause a query that is true to become false and another query that is false to become true in the new state. $\square$

Theorem 5. *There exists no state-matching reduction from SDCO to ARBAC97.*

Proof. By contradiction. Assume that there exists a state-matching reduction from SDCO to ARBAC97. Let $S = \{s_1, s_2, s_3, \dots\}$. In SDCO, adopt as γ a state with the following properties. Let $s_1 \in S_\gamma$, $o \in O_\gamma$ and $own \in M[s_1, o]$. Let q_i be the query “ $own \in [s_i, o]$ ” for each $i = 1, 2, \dots$, and q_o be the query “ $o \in O_\gamma$ ”. These queries are mapped to q_i^A and q_o^A respectively in the ARBAC97 scheme. We observe that $\gamma \vdash (q_1 \wedge \neg q_2 \wedge \neg q_3 \wedge \dots \wedge q_o)$. There exists a state $\tilde{\gamma}$ reachable from γ such that $\tilde{\gamma} \vdash (\neg q_1 \wedge q_2 \wedge \neg q_3 \wedge \dots \wedge q_o)$. And, there exists no reachable state $\hat{\gamma}$ such that $\hat{\gamma} \vdash (q_1 \wedge \neg q_2 \wedge \dots \wedge q_j \wedge \dots \wedge q_o)$ or $\hat{\gamma} \vdash (\neg q_1 \wedge \neg q_2 \wedge \dots \wedge \neg q_j \wedge \dots \wedge q_o)$ for any $j \neq 1$. (if $o \in O_\gamma$, then there must be exactly one subject that owns o). Consider the state γ^A in ARBAC97 that corresponds to γ (if there does not exist one, then we have the desired contradiction). We know that $\gamma^A \vdash (q_1^A \wedge \neg q_2^A \wedge \neg q_3^A \wedge \dots \wedge q_o^A)$. There must also exist a reachable state $\tilde{\gamma}^A$ that corresponds to $\tilde{\gamma}$ (if there does not exist one, then we have the desired contradiction). By Lemma 4, we know that $\tilde{\gamma}^A$ is not reachable from γ^A is a single state-transition. Therefore, there must exist some state $\hat{\gamma}^A$ that is reachable from γ^A such that $\hat{\gamma}^A \vdash (q_1^A \wedge \neg q_2^A \wedge \dots \wedge q_j^A \wedge \dots \wedge q_o^A)$

or $\hat{\gamma}^A \vdash (\neg q_1^A \wedge \neg q_2^A \wedge \dots \wedge \neg q_j^A \wedge \dots \wedge q_o^A)$ for at least one $j \neq 1$. As there exists no corresponding state in the SDCO scheme that is reachable from γ , we have a contradiction to the assumption that there exists a state-matching reduction from SDCO to ARBAC97. $\square$

The above theorem demonstrates that the ARBAC97 scheme is not as expressive as SDCO in the sense of the existence of a strongly security preserving mapping. The following theorem, however, demonstrates that if we weaken the desired mapping to only a security preserving mapping, then the ARBAC97 scheme is indeed at least as expressive as SDCO.

Theorem 6. *There exists a reduction from SDCO to ARBAC97.*

Proof. By construction. We present a mapping as required by Definition 8 and then prove that the mapping satisfies the two properties for it to be a reduction. Let $\gamma = \langle S_\gamma, O_\gamma, M_\gamma[], R_\gamma \rangle$ be the start-state of a given SDCO system, ψ its state-change rule and Q the set of queries. The mapping, σ , produces as output $\langle \gamma^A, \psi^A \rangle$ with input $\langle \gamma, \psi \rangle$ and output q^A for each $q \in Q$. We first define γ^A .

$$\begin{aligned} \gamma^A &= \langle UA_\gamma, PA_\gamma, RH_\gamma, AR_\gamma \rangle, \text{ where,} \\ UA_\gamma &= \{ \langle a, admin \rangle \} \cup \{ \langle s, subjectExists \rangle \mid s \in S_\gamma \} \cup \{ \langle s, o_r \rangle \mid r \in M_\gamma[s, o] \} \\ PA_\gamma &= \emptyset \\ RH_\gamma &= \{ top \succeq o_r \succeq bottom \mid o_r \in O_\gamma \times R_\gamma \} \\ AR_\gamma &= \{ admin \} \end{aligned}$$

We point out that we can infer from the above definition for γ^A that the set of all possible roles in the ARBAC97 system is $\mathcal{R} = (O \times R_\gamma) \cup \{ admin, subjectExists, bottom, top \}$, where we represent the role corresponding to $\langle o, r \rangle \in O \times R_\gamma$ as o_r . The role *admin* is used to ensure that state-changes are enabled, the role *subjectExists* is used to meaningfully map queries of the form “ $s \in S$ ” and the roles *bottom* and *top* form a role range that is used in the state-change rules below. We now present ψ^A .

$$\begin{aligned} \psi^A &= \langle can_assign, can_revoke, can_assignp, can_revokep, can_modify \rangle, \text{ where,} \\ can_assignp &= \emptyset \\ can_revokep &= \emptyset \\ can_assign &= \{ \langle admin, true, \xi \rangle \} \\ can_revoke &= \{ \langle admin, \xi \rangle \} \\ can_modify &= \{ \langle admin, \xi \rangle \} \\ &\text{where } \xi = \langle top, bottom \rangle \end{aligned}$$

Finally, the queries in Q are mapped as follows by σ .

$$\begin{aligned} \sigma("s \in S") &= "\langle s, subjectExists \rangle \in UA" \\ \sigma("o \in O") &= "\exists u \text{ such that } \langle u, o_{own} \rangle \in UA" \\ \sigma("r \in M[s, o]") &= "\langle s, o_r \rangle \in UA" \end{aligned}$$

We now show that property (1) for a reduction is satisfied by the above mapping. Let γ_0 be a start-state in SDCO. We produce the corresponding start-state γ_0^A in ARBAC97 using σ above. Given a state γ_k and query q such that $\gamma_0 \xrightarrow{\psi} \gamma_k$, we show that there exists γ_k^A and query q^A such that $\gamma_0^A \xrightarrow{\psi^A} \gamma_k^A$ where $\gamma_k^A \vdash q^A$ if and only if $\gamma_k \vdash q$. If $\gamma_k = \gamma_0$, then $\gamma_k^A = \gamma_0^A$. If q is “ $s \in S$ ”, then q^A is “ $\langle s, \text{subjectExists} \rangle \in UA$ ”. We know from the definition for UA_γ above that q^A is true if and only if q is true. If q is “ $o \in O$ ”, then q^A is “ $\exists u$ such that $\langle u, o_{own} \rangle \in UA$ ”. We know that every object that exists in SDCO has an owner associated with it (that is, $own \in M_{\gamma_0}[s, o]$ for some subject s). Consequently, from the definition for UA_γ above, there exists some s such that $\langle s, o_{own} \rangle \in UA_{\gamma_0}$ if and only if q is true. Finally, if q is “ $r \in M[s, o]$ ” then q^A is “ $\langle s, o_r \rangle \in UA$ ”. Again, by the definition of UA_γ above, we know that q is true if and only if q^A is true.

Consider some γ_k reachable from γ_0 and a query q . We show the existence of γ_k^A that is reachable from γ_0^A and that answers q^A the same way by construction. If q is of type “ $s \in S$ ”, we let $\gamma_k^A = \gamma_0^A$. If q is of type “ $o \in O$ ” or “ $r \in M[s, o]$ ”, we do the following. We consider each state-transition in the sequence $\gamma_0 \xrightarrow{\psi} \gamma_1 \xrightarrow{\psi} \dots \xrightarrow{\psi} \gamma_k$ in the SDCO system. If the state-transition is the execution of *createObject*(s, o), we execute *addToRange*(a, ξ, o_{own}) (where $\xi = \langle \text{top}, \text{bottom} \rangle$) and *assignUser*(a, s, o_{own}). If the state-transition in SDCO is the execution of *destroyObject*(s, o), we execute *revokeUser*(a, u, o_r) for every $\langle u, o_r \rangle \in UA$ for every r , and *removeFromRange*(a, ξ, o_{own}). If the state-transition in SDCO is the execution of *grantOwn*(s, s', o), we execute *revokeUser*(a, s, o_{own}) and *assignUser*(a, s', o_{own}). If the state-transition in SDCO is the execution of *grant_r*(s, s', o), we execute *assignUser*(a, s', o_{r_i}).

Now, consider each possible query q . If q is “ $s \in S$ ”, then $\gamma_k^A = \gamma_0^A$. In our SDCO scheme, the subjects are fixed at the start and never change. So $\gamma_k^A \vdash q^A$ if and only if $\gamma_0 \vdash q$. If q is “ $o \in O$ ”, then $\gamma_k \vdash q$ if and only if o exists in the state γ_k . This is the case if and only if some subject s has the *own* right over o . This is the case if and only if we have the role o_{own} in the range ξ and the user corresponding to s is a member of that role. Therefore, $\gamma_k \vdash q$ if and only if $\gamma_k^A \vdash q^A$. And finally, if q is “ $r \in M[s, o]$ ”, then $\gamma_k \vdash q$ if and only if r has been granted to s by the owner of o . This is true if and only if we have assigned the user corresponding to s to the role o_r . Thus, again, $\gamma_k \vdash q$ if and only if $\gamma_k^A \vdash q^A$.

We prove that property (2) for a reduction is satisfied by our mapping also by construction. Let γ_0^A be the start-state in ARBAC97 corresponding to γ_0 , the start-state in SDCO. Then, if γ_k^A is a state reachable from γ_0^A and q^A is a query in ARBAC97 whose corresponding query in SDCO is q , we construct γ_k , a state in SDCO reachable from γ_0 as follows. If q is “ $s \in S$ ”, we let $\gamma_k = \gamma_0$. Otherwise, for each role o_{own} that has a member s , we execute *createObject*(s, o). For each role o_r that has a member s' , if the role o_{own} has a member s , we execute *grant_r*(s, s', o). If q is $s \in S$, then q^A is $\langle s, \text{subjectExists} \rangle \in UA$, and clearly $\gamma_k^A \vdash q^A$ if and only if $\gamma_k \vdash q$, as

the subjects that exist do not change from the start-state in SDCO, and the members of *subjectExists* do not change from the start-state in ARBAC97. If q is “ $o \in O$ ”, $\gamma_k^A \vdash q^A$ if and only if $\exists s$ such that $\langle s, o_{own} \rangle \in UA_{\gamma_k}$. And if q^A is true, we would have added the *own* right to $M_{\gamma_k}[s, o]$, which means that $\gamma_k \vdash q$ if and only if $\gamma_k^A \vdash q^A$. And finally, if “ q is $r \in M[s, o]$ ”, $\gamma_k^A \vdash q^A$ if and only if $\langle s, o_r \rangle \in UA_{\gamma_k}$. The condition that q^A is true is the only one in which we would have added the right r to $M_{\gamma_k}[s, o]$, and therefore $\gamma_k \vdash q$ if and only if $\gamma_k^A \vdash q^A$. $\square$

As we discuss in the introduction to this section, the existence of the weaker notion of a reduction from SDCO to the ARBAC97 scheme may indicate that it is possible to extend the ARBAC97 scheme to be as expressive as the SDCO scheme under the stronger notion of a state-matching reduction as well. Lemma 4 gives us some intuition as to why a state-matching reduction does not exist from SDCO to the ARBAC97 scheme; the kinds of queries and the state-transition rules in the ARBAC97 scheme underlie the non-existence of a state-matching reduction from SDCO.

There may exist other schemes based on RBAC for which there is a state-matching reduction from SDCO. In extending the ARBAC97 we have adopted in this paper, an approach may be to adopt a different query set. 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. It is unclear to us how to extend the query set of the ARBAC97 scheme so that there would exist a state-matching reduction from SDCO.

Another possibility is to extend the state-transition rules. Again it is not obvious how this is to be done, and we leave the question of 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 as an open problem.

4.4. Comparing an RBAC scheme with a Trust Management Language

In this section, we compare a particular RBAC scheme to the trust management scheme, $RT[\cap]$. The RBAC scheme we consider is called Assignment And Revocation (AAR) [12]. In AAR, the state is an RBAC state, and state-transition rules are those from the URA97 component of the ARBAC97 [20]; users may be assigned to and revoked from roles.

$RT[\cap]$ is a trust management scheme 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 first present precise characterizations of the AAR scheme and $\text{RT}[\cap]$. Li and Tripunitara [12] present a form-2 weak reduction (see Definition 11) from AAR to $\text{RT}[\cap]$. We assert with the following theorem that the result can be made stronger.

The AAR scheme

Γ In AAR, a state is the RBAC state $\langle UA, PA, RH \rangle$, as discussed in the previous section for ARBAC97.

Ψ The state-transitions allowed are the operations *assignUser* and *revokeUser* from the previous section, with the exception that negation is not allowed in prerequisite conditions. In addition, in AAR, we require that for every role for which there is a *can_assign* entry, there is also a *can_revoke* entry. That is, if $\exists \langle ar, c, \xi \rangle \in \text{can_assign}$ such that *ar* has at least one member and *c* may evaluate to *true*, then $\forall r \in \xi, \exists \langle ar', \xi' \rangle \in \text{can_revoke}$ such that $r \in \xi'$ and *ar'* has at least one member.

$Q, \vdash$ Queries are of the form $s_1 \supseteq s_2$, where s_1 and s_2 are *user-sets*. A user-set is an expression that evaluates to a set of users. A set of roles, a set of permissions and a set of users are user-sets, as are unions and intersections of user-sets. We refer the reader to Li and Tripunitara [12] for more details on user-sets. Entailment involves evaluating the user-sets s_1 and s_2 to the sets of users S_1 and S_2 respectively, and determining whether $S_1 \supseteq S_2$. Several interesting queries related to safety, availability, liveness and mutual-exclusion can be posed as comparisons of user-sets.

The $\text{RT}[\cap]$ scheme

Γ An $\text{RT}[\cap]$ state is a set of credentials, each of which is one of the following types: (1) $A.r \leftarrow U$, (2) $A.r \leftarrow B.r_1$, and (3) $A.r \leftarrow B.r_1 \cap C.r_2$. Each of A, B, C, U is a principal, r, r_1, r_2 is a role name, and $A.r, B.r_1, C.r_2$ is a role. The symbol $\leftarrow$ is read as “includes”. Statement (1) asserts that *U* is a member of *A*’s *r* role. Statement (2) asserts that all members of the role $B.r_1$ are members of the role $A.r$. Statement (3) asserts that anyone that is a member of both $B.r_1$ and $C.r_2$ is a member of $A.r$.

Ψ A state-transition in $\text{RT}[\cap]$ is either the removal of a credential, or the addition of one. State-transitions are controlled by *growth* and *shrink*-restricted sets of roles — *G* and *S* respectively. A role that is in the growth-restricted set may not have any assertions added with that role at the head of the assertion, and a role that is in the shrink-restricted may not have any assertions removed. Thus, the state-transition rules are represented as $\langle G, S \rangle$.

$Q, \vdash$ We allow queries of the form $c_1 \supseteq c_2$ where each c_1 and c_2 is either an $\text{RT}[\cap]$ role, a credential, or credentials joined by union, $\cup$ or intersection, $\cap$. We observe that this is slightly different from the definition for queries in [12]. The reason is that in that work, only a form-2 weak reduction (see Definition 11) is presented, and therefore queries are processed in conjunction with each state and state-transition rule in the mapping. We seek to map queries independently of states and state-transition rules. Entailment in $\text{RT}[\cap]$ is done using credential chain discovery [13]: we find a chain of credentials that proves a (portion of a) query, if one exists.

Theorem 7. *There exists a state-matching reduction from the AAR scheme to $\text{RT}[\cap]$.*

Proof. By construction. We show that the mapping presented by Li and Tripunitara [12] from AAR to $\text{RT}[\cap]$ is a state-matching reduction. We consider each assertion from Definition 7 in turn. Each role r in AAR is associated with the role $\text{Sys}.r$ in $\text{RT}[\cap]$. We show that after a series of state-transitions, the role-memberships in AAR match the role-memberships in the corresponding state of $\text{RT}[\cap]$.

Assertion 1: Let γ be the given AAR state, and $\gamma \xrightarrow{*}_{\psi} \gamma'$. Then, $\gamma = \gamma_0 \xrightarrow{\psi} \gamma_1 \dots \xrightarrow{\psi} \gamma_m = \gamma'$. Each state-transition is either the assignment of a user to a role using *assignUser* or revocation of a user's membership in a role using *revokeUser*. Let the corresponding states in $\text{RT}[\cap]$ be $\gamma^T = \gamma_0^T, \gamma_1^T, \dots, \gamma_m^T = \gamma^{T'}$. The users that are members of any role r in γ are the same as the users that are members of the corresponding role $\text{Sys}.r$ in γ^T . If the state-transition from γ_i to γ_{i+1} is the result of the assignment of the user u to the role r , then we effect the following changes to transition from the state γ_i^T to γ_{i+1}^T : we add the two statements $\text{ASys}.r \leftarrow u$ and $\text{BSys}.r \leftarrow u$. If the state-transition is the result of the revocation of the user u from the role r , then we remove all statements that exist of the following two forms: $\text{ASys}.r \leftarrow u$ and $\text{RSys}.r \leftarrow u$. We observe that in $\gamma^{T'}$, any $\text{HSys}.r$ has as members all users that were ever members of the role r . Consequently, in $\gamma^{T'}$, each $\text{Sys}.r$ has as members those users that are members of r in γ' . Therefore, we can assert that $\gamma' \vdash q$ iff $\gamma^{T'} \vdash q^T$.

Assertion 2: In $\text{RT}[\cap]$, the only roles that can grow are the ASys and BSys roles. The only roles that can shrink are the ASys and RSys roles. Given $\gamma^T = \sigma(\gamma)$ where γ is a given AAR state and $\gamma^{T'}$ is the corresponding $\text{RT}[\cap]$ state, let $\gamma^T \xrightarrow{*}_{\psi} \gamma^{T'}$. We construct the AAR state γ' that corresponds to $\gamma^{T'}$ as follows. For each statement of the form $\text{BSys}.r \leftarrow u$ or of the form $\text{ASys}.r \leftarrow u$, we assign the user u to the role r . Now, we compare the user-role memberships of each user to the roles r and $\text{Sys}.r$. There cannot be any users in $\text{Sys}.r$ that are not in r : the reason is that we have not revoked any user membership in r (starting from the user-role membership in the state γ). There may be users in r that are not in $\text{Sys}.r$. Given the requirement that every role for which there is a *can_assign*, we also have a *can_revoke*, the only way for these extra users to be in r and not $\text{Sys}.r$ is that there exists a *can_assign* that permits those users to be assigned to r (starting at the state γ). We revoke such users' membership from r using the relevant *can_revoke* entries. Now, the memberships in r and $\text{Sys}.r$ are identical, and we can assert that for all queries q , $\gamma^{T'} \vdash \sigma(q)$ iff $\gamma' \vdash q$. $\square$

4.5. 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 to a command. ATAM is the same as TAM, except that in a condition in an ATAM command, the absence of a right in a cell of the access matrix may be checked (and not just the presence of a right). Below, we present characterizations of the two schemes.

Sandhu and Ganta [21] present a mapping from the 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 in the condition of an ATAM command. Sandhu and Ganta [21] 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.

The TAM scheme

Γ TAM is similar to the HRU scheme (see Section 2.1). Each state $\gamma \in \Gamma$ is $\langle S_\gamma, O_\gamma, M_\gamma[], R_\gamma, T_\gamma, typeOf \rangle$ where S_γ , O_γ , R_γ and T_γ are finite, strict subsets of the countably infinite sets $\mathcal{S}$ (subjects), $\mathcal{O}$ (objects), $\mathcal{R}$ (rights) and $\mathcal{T}$ (types of objects and subjects) respectively. The function $typeOf: (S_\gamma \cup O_\gamma) \rightarrow T_\gamma$, maps each subject and object to a type that cannot change once the subject or object is created. $M_\gamma[]$ is the access matrix.

Ψ A state-transition rule is a set of commands. Each command has an optional list of conditions that are joined by conjunction. A command then consists of primitive operations. Each parameter to the command is associated with a type. Each condition may check only for the presence of a right in a cell.

$Q, \vdash$ We allow queries of the form “is $r \in M[s, o]$?” Entailment is defined as follows. Given a state $\gamma \in \Gamma$, $\gamma \vdash r \in M[s, o]$ if and only if $s \in S_\gamma \wedge o \in O_\gamma \wedge r \in R_\gamma \wedge r \in M_\gamma[s, o]$.

The ATAM scheme

$\Gamma, \Psi, Q, \vdash$ An ATAM state is the same as a TAM state. State-transition rules are the same as for TAM, except that a condition in a command may check for the absence of a right (as opposed to only the presence of a right). In ATAM, we allow Q to contain queries of the following two forms: (1) Is $r \in M[s, o]$?, and (2) Is $r \notin M[s, o]$? This is consistent with the intent of Sandhu and Ganta [21] to determine whether the ability to check for the absence of rights does indeed add more expressive power. $\vdash$ is defined the same as in TAM for a query of type (1). For a query of the type (2), $\vdash$ is defined as follows. Given a state $\gamma \in \Gamma$, $\gamma \vdash r \notin M[s, o]$ if and only if $s \in S_\gamma \wedge o \in O_\gamma \wedge r \in R_\gamma \wedge r \notin M_\gamma[s, o]$.

Theorem 8. *There exists no state-matching reduction from ATAM to TAM.*

Proof. By contradiction. Assume that there exists a state-matching reduction σ from ATAM to TAM. Consider an ATAM scheme in which ψ (the state-transition rule) consists of the following commands.

$$\begin{array}{ll} \text{command } \text{createSubject}(X: t) & \text{command } \text{addRight}(Y: t, Z: t) \\ \text{create subject } X \text{ of type } t & \text{enter } r \text{ into } [Y, Z] \end{array}$$

Adopt as γ_0 (the start state) in ATAM a state with no subjects or objects. (that is, $S_{\gamma_0} = O_{\gamma_0} = \emptyset$). The set of rights, $R_{\gamma_0} = \{r\}$, and there is a single type t for all subjects (no objects other than subjects exist or can be created in our ATAM system). We denote components of the TAM system under the mapping σ with a superscript T . For example, $\sigma(\gamma_0) = \gamma_0^T$ and $\sigma(\psi) = \psi^T$.

We assume that the countably infinite set of subjects $\mathcal{S} = \{s_1, s_2, \dots\}$. In the ATAM system, we wish to consider queries of the form $q_{i,j} = r \in M[s_i, s_j]$ and $\widehat{q_{i,j}} = r \notin M[s_i, s_j]$ for some $s_i, s_j \in \mathcal{S}$. First, we make the observation that any two distinct queries $p, q \in \{q_{i,j} | s_i, s_j \in \mathcal{S}\} \cup \{\widehat{q_{i,j}} | s_i, s_j \in \mathcal{S}\}$ are mapped to distinct queries in TAM. That is, $p \neq q \Rightarrow p^T \neq q^T$. Otherwise, pick a pair p, q such that $p \neq q$ but $p^T = q^T$. For any two such queries p and q , there exists a state γ in ATAM such that $\gamma_0 \xrightarrow{\psi^*} \gamma$ and $\gamma \vdash p \wedge \neg q$. Clearly, a corresponding reachable state (that answers the queries p and q the same way) does not exist in TAM, which gives us the desired contradiction. We observe also that by the definition of a state-matching reduction, queries are mapped independent of the start state and the state-change rules.

Consider ψ^T , the command schema in TAM. As a query in TAM is of the form $r \in M[s, o]$, we can determine an upper bound, m , for the number of queries a command in the TAM system can change from false to true when executed. These are queries of both types $q_{i,j}^T$ and $\widehat{q_{i,j}}^T$. One way to determine a value for m is to count the number of “enter right” primitive operations in each command and take the maximum (even though this maximum may not be a tight upper bound). m is constant, and may be dependant on γ and ψ , but not the set of queries. Choose some $n > m$.

Now, consider the state in ATAM γ_k such that $\gamma_0 \xrightarrow{\psi^*} \gamma_k$ and $\gamma_k \vdash \neg q_{1,1} \wedge \widehat{q_{1,1}} \wedge \neg q_{1,2} \wedge \widehat{q_{1,2}} \wedge \dots \wedge \neg q_{n,n} \wedge \widehat{q_{n,n}}$ (we use the subscript k only to distinguish the state, and not as a count of the number of state-changes needed to reach it). That is, γ_k does not entail any of the queries of the type $q_{i,j}$ and entails all queries of the type $\widehat{q_{i,j}}$ for all integers i, j such that $1 \leq i, j \leq n$. The state γ_k corresponds to $S_{\gamma_k} = \{s_1, \dots, s_n\}$ with no right r in any of the cells. One way to reach this state from γ_0 is to execute the command *createSubject* n times with the parameter instantiated to s_i in the i^{th} execution.

We assume that as σ , a state-matching reduction exists, there exists a corresponding reachable state γ_k^T in TAM that answers the (mapped) queries the same way.

Consider any sequence $\gamma_0^T \mapsto_{\psi^T} \gamma_1^T \mapsto_{\psi^T} \dots \mapsto_{\psi^T} \gamma_k^T$. Pick the first state, γ_c^T in the sequence that satisfies the following condition: $\gamma_c^T \vdash q_{i,j}^T \vee \widehat{q_{i,j}}^T$ for all integers i, j such that $1 \leq i, j \leq n$. Such a state exists: γ_k^T is such a state, and may be the only state in the sequence that meets the condition. We observe also that γ_0^T does not satisfy the condition, thereby implying that the sequence has at least one state-change.

Consider the state γ_{c-1}^T in the sequence just before γ_c^T . γ_{c-1}^T has the following property: there exist integers v, w with $1 \leq v, w \leq n$, such that $\gamma_{c-1}^T \vdash \neg(q_{v,w}^T \vee \widehat{q_{v,w}}^T) \Rightarrow \gamma_{c-1}^T \vdash \neg q_{v,w} \wedge \neg \widehat{q_{v,w}}^T$. For every state in the ATAM system that entails the corresponding formula of queries $\neg q_{v,w} \wedge \neg \widehat{q_{v,w}}^T$, the state also entails at least one of the following two formulae of queries: (1) $Q_1 = \neg q_{v,1} \wedge \neg \widehat{q_{v,1}}^T \wedge \neg q_{v,2} \wedge \neg \widehat{q_{v,2}}^T \wedge \dots \wedge \neg q_{v,n} \wedge \neg \widehat{q_{v,n}}^T \wedge \neg q_{1,v} \wedge \neg \widehat{q_{1,v}}^T \wedge \dots \wedge \neg q_{n,v} \wedge \neg \widehat{q_{n,v}}^T$, or, (2) $Q_2 = \neg q_{w,1} \wedge \neg \widehat{q_{w,1}}^T \wedge \neg q_{w,2} \wedge \neg \widehat{q_{w,2}}^T \wedge \dots \wedge \neg q_{w,n} \wedge \neg \widehat{q_{w,n}}^T \wedge \neg q_{1,w} \wedge \neg \widehat{q_{1,w}}^T \wedge \dots \wedge \neg q_{n,w} \wedge \neg \widehat{q_{n,w}}^T$.

The reason is that a state in ATAM that entails $\neg q_{v,w} \wedge \neg \widehat{q_{v,w}}^T$ is one in which either the subject s_v or s_w , or both do not exist ($v = w$ is allowed, and does not affect our arguments). None of the queries of either type $q_{i,j}$ or $\widehat{q_{i,j}}^T$ corresponding to a subject that does not exist in a state is entailed by the state. Therefore, in TAM, $\gamma_{c-1}^T \vdash Q_1^T \vee Q_2^T$ (where Q_1^T and Q_2^T are obtained from Q_1 and Q_2 respectively by adding the superscript T to each query in the formula).

Consider the state-change in TAM from γ_{c-1}^T to γ_c^T . It must change (at least) n queries that appear in Q_1^T or Q_2^T from false to true. This is not possible, as each state-change can change at most $m < n$ queries from false to true. We have the desired contradiction to the existence of a state-matching reduction from the ATAM scheme to the TAM scheme. $\square$

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 also ask 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 only allows checking for the presence of a right in a cell in the condition). These are open questions.

4.6. Summary of results

In applying our theory to compare access control schemes, we have considered four broad models: DAC, RBAC, Trust Management and Access Matrix. In Fig. 1, we present results that we have shown in this paper, and that are known or can be inferred.

The results that we present in Fig. 1 are the following.

- DAC – SDCO is one of the sub-schemes discussed by Osborn et al. [17]. It is also a sub-scheme of the Graham-Denning scheme [5]. Consequently, there




Fig. 1. Comparisons between schemes in four broad access control models: DAC, RBAC, Trust Management and Access Matrix. A solid line shows the existence of a state-matching reduction (e.g., there is a state-matching reduction from AAR to $RT[\cap]$). A dotted line shows the existence of a reduction (e.g., there is a reduction from SDCO to ARBAC97). Arrows are sometimes double-headed for brevity. An “X” qualifies one of the above relationships by indicating that such a relationship does not exist. For example, there is no state-matching reduction from the Graham-Denning scheme [5] to the HRU scheme [6] and no state-matching reduction from the HRU scheme to the Graham-Denning scheme.

exist state-matching reductions from SDCO to the family of DAC schemes presented by Osborn et al. [17] and the Graham-Denning scheme [5].

- DAC and RBAC

- There exists no state-matching reduction from the ARBAC97 scheme to the Graham-Denning scheme. The intuition is that some state-changes (e.g., user to role assignment) in ARBAC97 support more sophisticated preconditions than those that are supported in the Graham-Denning scheme. This result shows that the most general DAC scheme from the literature is not as expressive as an RBAC scheme.
- There exists no state-matching reduction from SDCO to the ARBAC97 scheme. This is Theorem 5 in this paper.
- There exists a reduction from SDCO to the ARBAC97 scheme. This is Theorem 6 in this paper.
- There exists no state-matching reduction from AAR to SDCO. This follows from the fact that AAR allows for the specification of preconditions to the assignment of a user to a role. SDCO cannot capture such complex preconditions.

- RBAC – AAR is a sub-scheme of URA97, which in turn is a sub-scheme of ARBAC97. Consequently, there exists a state-matching reduction from AAR to URA97, and from URA97 to ARBAC97.
- RBAC and Trust Management
 - There exists no state-matching reduction from the URA97 scheme to $RT[\cap]$. The intuition is that URA97 supports negative preconditions in the assignment of user to roles and this cannot be captured in $RT[\cap]$.
 - There exists a state-matching reduction from AAR to $RT[\cap]$. This is Theorem 7 in this paper.
 - There exists no state-matching reduction from $RT[\cap]$ to AAR. The intuition is that AAR has the constraint that every role to which a user can be assigned must be such that users can be revoked from it (see Section 4.4 for a description of AAR). As there is no such constraint on the state-change rules of $RT[\cap]$, AAR is not as expressive as $RT[\cap]$ under state-matching reductions.
 - There exists no state-matching reduction from AAR to $RT[\]$. The intuition is that AAR supports conjunction in its preconditions in state-changes which cannot be captured in $RT[\]$. That is, if we weaken $RT[\cap]$ to $RT[\]$, it loses sufficient expressive power that it is no longer as expressive as AAR under state-matching reductions.
- Trust Management – There exists a state-matching reduction from $RT[\]$ to $RT[\cap]$; this is obvious because $RT[\]$ is a sub-scheme of $RT[\cap]$.
- Trust Management and Access Matrix
 - There exists no state-matching reduction from $RT[\]$ to the HRU scheme. This is Theorem 3 in this paper. There is also no state-matching reduction from the HRU scheme to $RT[\]$. See Section 4.1 for a discussion of the intuition behind these results.
- Access Matrix – There exists a state-matching reduction from the HRU scheme to TAM, but not vice-versa. This is because the HRU scheme is a special case of TAM in which all subjects and objects are constrained to be of a single type. There is a state-matching reduction from TAM to ATAM; TAM is a sub-scheme of ATAM. There is no state-matching reduction from ATAM to TAM; this is Theorem 8 in this paper.
- Access Matrix and DAC
 - There is no state-matching reduction from the Graham-Denning scheme to the HRU scheme. The intuition is that in the HRU scheme, commands are more “free form” than in the Graham-Denning scheme. There is also no state-matching reduction from the HRU scheme to the Graham-Denning scheme. The reason is that in the Graham-Denning scheme, when a subject is destroyed, rights to objects owned by the subject are transferred to subjects that control the subject being destroyed. That is, a potentially unbounded number

of subjects get rights to objects in a single state-change. This cannot be effected in an HRU system. We refer the reader to Li and Tripunitara [11] for more discussions on the distinction between the Graham-Denning scheme and the HRU scheme.

Open problems. Figure 1 and our discussions above suggest that there exist several open problems in the comparison of various access control schemes. We list some of them here.

- It is unknown whether the schemes discussed by Osborn et al. [17] can be fully captured by the Graham-Denning scheme [5]. For example, Osborn et al. [17] discuss a DAC scheme with the feature that when a user is granted a right, he gets it with a certain depth which is an integer. If he were to grant the right to another user, the right is granted with the depth decremented by one. A user that has the right with depth 0 cannot grant the right any further. The maximum depth is a constant specified as part of the access control system. This feature controls the delegation depth with regards to the right. It is unclear whether there exists a state-matching reduction or reduction from such a scheme to the Graham-Denning scheme.
- It is unknown whether the ARBAC97 scheme [20] is at least as expressive as all or some of the schemes in the RT family [9]. Also, we do not know whether there exists a state-matching reduction or reduction from one of the access matrix schemes to the ARBAC97 scheme. It is also unclear how and whether we can extend the ARBAC97 scheme to be as expressive as the SDCO scheme, and indeed, the full Graham-Denning scheme, with respect to state-matching reductions. These questions would continue the investigation into the issue of the expressive power of general RBAC models [15,17,22].
- As we point out in Section 4.5, we do not know whether there exists a state-matching reduction from ATAM to TAM if we allow queries that check for the absence of rights in TAM. In proving Theorem 8, we assume that such queries are not allowed in TAM. We believe there that this captures our intent of demonstrating that checking for the absence of rights does indeed add expressive power to ATAM. However, if we disallow checking for the absence of rights in only the state-change rules and keep the queries the same in TAM and ATAM, then the problem is open as to whether ATAM continues to be more expressive than TAM with respect to state-matching reductions or reductions.

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 highlighted four applications of our theory and shown that: (1) the HRU scheme and a trust management scheme, $RT[\]$, are incomparable with one another; (2) RBAC with ARBAC97 as its administrative model is at least as expressive as a particular DAC scheme under the relatively weak notion of reductions, but not under the stronger notion of state-matching reductions; (3) the trust-management scheme $RT[\cap]$ is at least as expressive as RBAC with the URA97 component of ARBAC97 as its administrative model; and (4) the higher expressive power of ATAM when compared to TAM can be formalized using the notion of state-matching reductions. To our knowledge, (1) is evidence that the expressive power of the HRU scheme is limited, (2) is the first known evidence that the expressive power of RBAC may be limited, and (4) formally demonstrates the benefit from the ability to check for the absence of a right in addition to the presence of a right.

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 be represented and verified in an access control system, we do not believe that reductions and state-matching reductions capture all the types of policies we would want to consider. For instance, a reasonable question to ask during a security audit may be: “did Alice get her write access to a sensitive file only after her husband, Bob was given privileged access to the system?” This can be perceived as a policy issue, and we may want to express this as some expression involving queries. 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 policies.

In particular, we would like to explore the use of logics beyond propositional logic; for example, it is likely the Computational Tree Logic (CTL) will be useful in capturing policies of the kind we discuss above. Adoption of more sophisticated logics as the basis for our queries will give us richer queries, and allow us to distinguish between access control schemes more meaningfully.

Acknowledgements

We thank the anonymous reviewers and the associate editor, Carl Gunter, for their feedback which have helped improve this paper. Portions of this work were supported by NSF grants CCR-0325951 and CNS-0448204, and sponsors of CERIAS.

Appendix A. A “Simulation” of RBAC in Strict DAC

We now informally describe a simulation of RBAC in strict DAC, the simplest form of DAC. The point of this simulation is to show that if precise requirements are not specified on simulations, then anything is possible.

The state of a strict DAC model is represented by an access matrix, which has one subject for each user and each role and one object for each permission. There is also one special subject **admin**, who is the creator and owner of every object in the system. All subjects are also objects. We use three rights, “**own**”, “**dc**”, and “**c**”. We assume that the implementation of the strict DAC model provides the following functionality, it internally sorts all the objects and can return the first object, given an object o , it return the object next to o . The commands implemented in the strict DAC are as follows:

```

command create(s, o)
  create o;
  enter own into (s,o);
end;
command delete(s, o)
  if own  $\in$  (s,o)
    destroy o;
end;
command grant-dc(s1, s2, o)
  if own  $\in$  (s1,o)
    enter dc into (s2,o);
    enter c into (s2,o);
end;
command grant-c(s1, s2, o)
  if own  $\in$  (s1,o)
    enter c into (s2,o);
end;
command revoke-dc(s1, s2, o)
  if own  $\in$  (s1,o)
    remove dc from (s2,o);
end;
command revoke-c(s1, s2, o)
  if own  $\in$  (s1,o)
    remove c from (s2,o);
end;

```

The addition of new users, roles, and permissions are carried out by the simulator in the straightforward way, i.e., have **admin** executes a creation command; **admin** then becomes the owner of these objects. When a new user-role assignment, (u, r) , is added, the following procedure is executed, observe that only constant space is needed for the simulation.

```

addUR(u,r) {
  run command grant-dc(admin, u, r);
  while (propagate());
}

```



```

}
propagate() {
  repeat = false;
  for every s,o1,o2 in the matrix {
    if  $\mathbf{c} \notin (s,o2)$  &&  $\mathbf{c} \in (s,o1)$  &&  $\mathbf{c} \in (o1,o2)$  {
      run command grant-c(admin, s, o2);
      repeat = true;
    }
  }
  return repeat;
}

```

The procedures for adding a role-permission assignment and a role-role inheritance relationship is similar.

Whenever a user-role assignment is removed, the simulator executes the following procedure, which first clear all the propagated rights and redo the propagation.

```

removeUR(u,r) {
  if ( $\mathbf{dc} \in (u,r)$ ) {
    run command revoke-dc(admin, u, r);
    clear();
    while (propagate());
  }
}
clear() {
  for every s,o in the matrix {
    if  $\mathbf{c} \in (s,o)$  {

      run command revoke-c(admin, s, o2);
    }
  }
}

```

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) (1996), 149–165.
- [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) (2003), 71–127.
- [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*, IEEE Computer Society Press, 2001, pp. 27–43.
- [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*, Vol. 40, AFIPS Press, 1972, pp. 417–429.

- [6] M.A. Harrison, W.L. Ruzzo and J.D. Ullman, Protection in operating systems, *Communications of the ACM* **19**(8) (1976), 461–471.
- [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) (1974), 18–24.
- [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, 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*, IEEE Computer Society Press, 2002, pp. 114–130.
- [10] N. Li, J.C. Mitchell and W.H. Winsborough, Beyond proof-of-compliance: Security analysis in trust management, *Journal of the ACM* **52**(3) (2005), 474–514. Preliminary version appeared in: *Proceedings of 2003 IEEE Symposium on Security and Privacy*.
- [11] N. Li and M.V. Tripunitara, On Safety in Discretionary Access Control, in: *Proceedings of the 2005 IEEE Symposium on Security and Privacy*, 2005.
- [12] N. Li and M.V. Tripunitara, Security analysis in role-based access control, 2006. Accepted to appear in *ACM Transactions on Information and Systems Security (TISSEC)*. Preliminary version appeared in: *Proceedings of the Ninth ACM Symposium on Access Control Models and Technologies (SACMAT 2004)*.
- [13] N. Li, W.H. Winsborough and J.C. Mitchell, Distributed credential chain discovery in trust management, *Journal of Computer Security* **11**(1) (2003), 35–86.
- [14] R. Milner, *A Calculus of Communicating Systems*, Lecture Notes in Computer Science, Vol. 92, Springer, 1980.
- [15] Q. Munawer and R.S. Sandhu, Simulation of the augmented typed access matrix model (ATAM) using roles, in: *Proceedings of INFOSEC99 International Conference on Information and Security*, 1999.
- [16] 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)*, 1997, pp. 31–40.
- [17] 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) (2000), 85–106.
- [18] D. Park, Concurrency and automata on infinite sequences, in: *Proceedings of the 5th GI Conference on Theoretical Computer Science*, Lecture Notes in Computer Science, Vol. 104, Springer, 1981.
- [19] R.S. Sandhu, Expressive power of the schematic protection model, *Journal of Computer Security* **1**(1) (1992), 59–98.
- [20] R.S. Sandhu, V. Bhamidipati and Q. Munawer, The ARBAC97 model for role-based administration of roles, *ACM Transactions on Information and Systems Security* **2**(1) (1999), 105–135.
- [21] 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*, IEEE Computer Society Press, 1993, pp. 109–118.
- [22] 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)*, 1998, pp. 47–54.
- [23] M.V. Tripunitara and N. Li, Comparing the expressive power of access control models, in: *Proceedings of 11th ACM Conference on Computer and Communications Security (CCS-11)*, ACM Press, 2004, pp. 62–71.