

Policy Administration in Tag-Based Authorization

Sandro Etalle^{1,2}, Timothy L. Hinrichs³, Adam J. Lee⁴, Daniel Trivellato¹,
and Nicola Zannone¹

¹ Eindhoven University of Technology

² University of Twente

³ University of Illinois at Chicago

⁴ University of Pittsburgh

Abstract. Tag-Based Authorization (TBA) is a hybrid access control model that combines the ease of use of extensional access control models with the expressivity of logic-based formalisms. The main limitation of TBA is that it lacks support for policy administration. More precisely, it does not allow policy-writers to specify administrative policies that constrain the tags that users can assign, and to verify the compliance of assigned tags with these policies. In this paper we introduce TBA² (Tag-Based Authorization & Administration), an extension of TBA that enables policy administration in distributed systems. We show that TBA² is more expressive than TBA and than two reference administrative models proposed in the literature, namely HRU and ARBAC97.

Keywords: access control, policy administration, auditing.

1 Introduction

Access control systems in real-world organizations are mostly based on extensional approaches to access control (e.g., access control lists), as their ease of use is preferred to the flexibility of logic-based models. Authorization policies in extensional models are based on simple assignments of rights to users, or on the characterization of users in terms of properties (e.g., roles) and the assignment of rights based on those properties. Nevertheless, the lack of expressiveness of extensional models severely limits the constraints that can be expressed in authorization policies. A hybrid approach to access control that combines the usability of extensional models and the flexibility and expressiveness of logic-based formalisms would offer great benefits for the deployment of access control systems in real-world organizations.

To accommodate this need, Tag-Based Authorization (TBA) has been proposed by Hinrichs et al. [15], based on the work by Najafian Razavi and Iverson [20] and Wang et al. [26]. TBA is a hybrid access control model that relies on formal logic for the definition of authorization policies, and on extensional models for describing a system's subjects and objects in terms of simple properties (e.g., roles). This integration allows relatively untrained users to choose descriptive tags for the system's subjects and objects; security experts then write logical policies that define access authorizations using combinations of those tags. The resulting access control model is flexible and easy to use, yet expressive enough to match the needs of complex application domains.

As an example application domain for the TBA model, consider Operation Atalanta, a military operation involving several EU navies that collaborate to prevent criminal activities (such as smuggling and pirate attacks) off the Somali coast. The information gathered and exchanged by the collaborating vessels is typically processed and classified by the vessels' operators (e.g., a transiting ship is marked as suspicious). Yet, the policies for accessing this information are regulated by complex context- and content-based conditions and must be written by higher-rank personnel. By explicitly distinguishing the task of characterizing information and the task of writing the authorization policy, TBA leads to a better separation of duties and valorization of the skills of a vessel's workforce than the existing access control models.

If on the one hand allowing low-rank users to assign tags to subjects and objects greatly enhances the usability of the access control system, on the other hand it enables low-rank users to influence the system's authorizations. In fact, inaccurate tags (whether by intention or accident) can circumvent the intended authorization policy of the system. The main limitation of TBA in this respect is the lack of support for policy administration. In particular, TBA does not allow policy-writers to: (1) define which users may assign which tags to which subjects and objects, i.e., to specify *administrative policies*; (2) hold users accountable for the tags they assign, and verify the compliance of tags with administrative policies; and (3) revoke incorrect tags or tags assigned by unauthorized users. In the scenario above, for instance, it is important to allow only operators with appropriate clearance to tag sensitive information, and to verify whether the assigned tags comply with this policy (possibly revoking the tags that do not satisfy the policy).

In a centralized system, these issues can be addressed by means of a traditional access control mechanism that regulates the tagging process and a logging mechanism that records users' actions. In a distributed system, however, these solutions are insufficient as they would require entities in one security domain to trust the enforcement and logging mechanisms of systems in different security domains. Furthermore, tracing the objects exchanged in a distributed system might be difficult, complicating or even preventing the revocation process.

In this paper, we extend the TBA model from [15] to enable policy administration in distributed systems. In particular, we introduce TBA² (Tag-Based Authorization & Administration), an extension of TBA that allows security administrators to specify constraints on the tagging process and to verify the compliance of tags with administrative policies by combining traditional "a priori" access control mechanisms with "a posteriori" verification. Technically, our extension consists of associating an *issuer* to each tag, which identifies the user who assigned the tag. For example, all the tags issued by an operator will be marked with the identifier of the operator. In contrast to previous work on a posteriori verification (e.g., [7]), this enables policy compliance verification without the need of an auditing infrastructure, by exploiting the observability of users' actions (i.e., "signed" tags). As a consequence, TBA² represents a lightweight solution for the enforcement of authorization and administrative policies in distributed systems. We show that the proposed model is more expressive than the original TBA model and than two well-known administrative models proposed in the literature, namely the Harrison, Ruzzo, and Ullman model [13] and ARBAC97 [23].

The paper is organized as follows. Section 2 reviews the TBA model and discusses its limitations. Section 3 introduces TBA², and Section 4 evaluates its expressive power. Section 5 discusses related work, and Section 6 concludes the paper.

2 Background

Tag-Based Authorization (TBA) [15] is a hybrid access control model combining the flexibility and expressiveness of logic-based formalism with the ease of use of extensional models. In this section we first present the definitions of the TBA model given in [15] that are relevant for this paper, and then we identify the main limitations of TBA with respect to its deployment in a distributed system.

2.1 The TBA Model

In this paper we use S to denote the set of subjects, O to denote the set of objects, and R to denote the set of rights that can be assigned within a system. T denotes the set of possible tags that can be assigned to S and O , and $\text{tag} : S \cup O \rightarrow \mathcal{2}^T$ is a function that maps a subject or object to the set of tags assigned to it. Tag denotes the set of all possible tag functions. A TBA authorization policy is written in some logical access control language $\langle P, L, \models \rangle$ defined as follows.

Definition 1 (TBA). For a function tag and a logical language $\langle P, L, \models \rangle$, where

- $\text{tag} : S \cup O \rightarrow \mathcal{2}^T$ maps a subject or object to the set of tags assigned to it
- P : the set of all authorization rules
- L : the set of queries including $\text{allow}(s, o, r)$ for all subjects s , objects o , rights r
- \models : a subset of $P \times \text{Tag} \times L$

$\text{auth}_{TBA}(s, o, r)$ if and only if $P', \text{tag} \models \text{allow}(s, o, r)$ for some $P' \subseteq P$.

The operator \models dictates which queries are true given the set of authorization rules P and a function tag . The following example based on the scenario introduced in Section 1 illustrates TBA using Datalog as the policy language.

Example 1 (TBA). Consider two subjects s_1 and s_2 and two objects o_1 and o_2 that are tagged as follows:

- $\text{tag}(s_1) = \{\text{uk_navy}, \text{operation_atalanta}, \text{operations_specialist}\}$
- $\text{tag}(s_2) = \{\text{fr_navy}\}$
- $\text{tag}(o_1) = \{\text{cargo_ship}, \text{radar}\}$
- $\text{tag}(o_2) = \{\text{gulf_of_aden}, \text{sat_732}, \text{high_res}\}$

Further, consider the following policy rules:

1. $\text{allow}(S_x, O_x, \text{read}) :- \text{uk_navy} \in \text{tag}(S_x), \text{cargo_ship} \in \text{tag}(O_x)$
2. $\text{allow}(S_x, O_x, \text{read}) :- \text{fr_navy} \in \text{tag}(S_x), \text{cargo_ship} \in \text{tag}(O_x)$
3. $\text{allow}(S_x, O_x, \text{read}) :- \text{operations_specialist} \in \text{tag}(S_x), \text{radar} \in \text{tag}(O_x)$
4. $\text{allow}(S_x, O_x, \text{read}) :- \text{uk_navy} \in \text{tag}(S_x), \text{operation_atalanta} \in \text{tag}(S_x),$
 $\text{high_res} \in \text{tag}(O_x), \text{sat_732} \in \text{tag}(O_x)$

This policy allows the members of the British and French Navy (denoted *uk_navy* and *fr_navy* respectively) to access documents about cargo ships (rules 1 and 2), operations specialists to access documents about radar systems (rule 3), and all members of the British Navy serving on Operation Atalanta to access high resolution satellite photographs taken by sat_732 (rule 4). As a result, subject s_1 can access objects o_1 and o_2 , while subject s_2 can only access object o_1 .

Tag-based authorization differs from standard logical access control models in that the tag function has a fixed semantics defined outside of the policy. In particular, the semantics of tag is defined by the users of a system, who assign tags to the system's subjects and objects. The fixed semantics of tag forces security administrators to define authorization policies at a higher level of abstraction than the usual $S \times O \times R$. More precisely, TBA policies are defined over the space of tags $2^T \times 2^T \times R$, i.e., subjects and objects are replaced by tag sets. This abstraction might result in a less flexible system if the tag-space is not exhaustive enough to accommodate a particular situation; however, the model can be easily adapted to define policies over $(S \cup 2^T) \times (O \cup 2^T) \times R$ (which combines the space of tags and the ones of subjects and objects) by adding to the set of tags T a tag identifying each subject and object in a system.

In [15], the authors evaluate the expressive power of TBA by expressing a range of well-known policy idioms. In particular, they show that TBA can be successfully employed to represent an access matrix, attribute-based access control policies, role-based access control policies, discretionary access control, mandatory access control, and three rule types of the RT [18] policy language (namely, all rule types except for "linked roles"). In addition, their results show that TBA is strictly more expressive than common access control models such as SDCO [21], ARBAC97 [23], and BLP [3].

2.2 Limitations of TBA

The applicability of TBA is due to a key observation: within a system, the users responsible for creating and categorizing (i.e., tagging) data are usually not in charge and do not have the expertise to define the security policies governing the system. For example, on a navy vessel some operators have the task of analyzing the surrounding maritime traffic and identifying suspicious behaviors, other operators gather intelligence and add details to these suspicious activities, etc. The policy governing the access to this information, on the other hand, is defined by the authorities in command of operations.

By assigning tags to subjects and objects, however, users directly influence the authorizations within a system. For this reason, it is necessary to enable policy-writers (e.g., security administrators) to control the tagging process. In this respect, we identify two key issues that are not addressed by the TBA model:

1. *Administrative policies*: Security administrators should be able to specify policies defining which users are allowed to assign and revoke which tags to which subjects and objects. For instance, the security administrator of a navy vessel should be able to restrict to the commanding officer the right to promote the vessel's officers to higher ranks. In addition, it should be possible to regulate the propagation of users' rights, i.e., the extent to which a user can delegate its tagging rights to other users.

2. *Tag Verification and Revocation*: It should be possible to hold users accountable for the tags they assign, and to verify the compliance of tags with administrative policies by checking who assigned them. Consider, for instance, a document containing information about a suspected pirate attack, distributed by the French Navy to its allies in Operation Atalanta. If some tag is added to the document by a user outside the navy's system, the security administrator of the French Navy should be able to verify whether the user was authorized to label the document (e.g., if the user is an operator of an allied navy, rather than an unknown subject). Accordingly, "invalid" tags identified in the verification process should be revoked.

Even though some simple administrative policies could be expressed in TBA, the fact that the "issuer" of a tag is not taken into account by the model makes it impossible to specify constraints that link tags based on the subject who assigned them (e.g., RT's linked roles [18]). Therefore, TBA does not allow the specification of rules such as "a subject can revoke only the tags that she assigned", or "a subject can mark a document as sensitive only if she is (tagged as) a senior officer by a navy that the EU labels as member of Operation Atalanta".

More than the specification of administrative policies, however, the major limitation of TBA is represented by the lack of mechanisms for verifying their enforcement. In a centralized system, administrative policies can be enforced by means of an access control system that governs the tagging process. Tag verification can be achieved by means of a logging mechanism that records all the tags assigned by the system users. Security administrators can then simply audit these logs to verify their compliance with the system's policies, and revoke the invalid tags identified in the process. In a distributed system, however, these solutions are insufficient for the following reasons:

1. They require entities in one security domain to trust the administrative policies (and their enforcement) of systems in different security domains.
2. Since tags may be assigned by users of different systems, and the issuer of a tag is not considered by the TBA model, verifying the compliance of the tags assigned by a subject with respect to administrative policies might not be feasible. In fact, this might require inspecting the logs of possibly all the systems in the distributed system. It is unlikely, however, that a system would disclose its logs to systems from a different security domain for auditing purposes. In addition, as information is exchanged between different systems, tracing the tags assigned by a certain user (e.g., to revoke them) becomes very difficult, if not impossible.

In the next section we show how the TBA model can be extended to address these limitations.

3 The TBA² Model

In this section we present the Tag-Based Authorization & Administration (TBA²) model, an extension of TBA that enables policy administration in distributed systems. TBA² extends TBA by associating an *issuer* to each tag, which identifies the user who assigned the tag. Intuitively, tags become signed statements issued by a user within a system. Formally, we modify the definition of the tag function introduced in Section 2.1 as follows:

$\text{tag} : S \cup O \rightarrow 2^{S \times T}$, which returns the set of issuer-tag pairs associated to a subject or object. Representing tags as signed statements is a first step towards addressing the limitations of TBA mentioned in Section 2.2. In the next subsections we discuss how TBA² can solve those limitations in details.

For the sake of simplicity, the solution we propose is based on the assumption that the set T of possible tags that a user can assign is common to all the systems in the distributed system. In a real-world distributed system, however, the set of assignable tags might vary from system to system, as entities in different security domains might employ different terms to denote similar concepts. Semantic alignment techniques [14, 24] could be required to align the systems' vocabularies. An additional assumption we make is that each subject and object *belongs* to one system, which represents the security domain where a subject operates, or the system that owns an object (or that has exclusive rights on it). Given a subject or object identifier, it is possible to determine the system to which the subject or object belongs.

3.1 Authorization and Administration Policies

Administrative policies constrain the tags that a user is authorized to assign or revoke. The advantage of TBA² with respect to TBA is that it links every tag to the user who assigned it, enabling security administrators to specify fine-grained administrative policies.

TBA² requires the set of rights R to include the rights *assign_tag* and *revoke_tag*. Then, TBA² is defined as follows.

Definition 2 (TBA²). For a function *tag* and a logical language $\langle P \cup A, L, \models \rangle$, where

- *tag*: $S \cup O \rightarrow 2^{S \times T}$ returns the issuer-tag pairs associated to a subject or object
- P : the set of all authorization rules
- A : the set of all administrative rules
- L : the set of queries including $\text{allow}(s, o, r)$, $\text{allow}(s, o, r', ST)$, and $\text{allow}(s, s', r', ST)$ for all subjects s and s' , objects o , rights r , right $r' \in \{\text{assign_tag}, \text{revoke_tag}\}$, and set of signed tags $ST \subseteq 2^{S \times T}$
- \models : a subset of $(P \cup A) \times \text{Tag} \times L$

$\text{auth}_{TBA^2}(s, o, r)$ if and only if $P', \text{tag} \models \text{allow}(s, o, r)$

$\text{auth}_{TBA^2}(s, o, r', ST)$ if and only if $A', \text{tag} \models \text{allow}(s, o, r', ST)$

$\text{auth}_{TBA^2}(s, s', r', ST)$ if and only if $A', \text{tag} \models \text{allow}(s, s', r', ST)$

for some $P' \subseteq P$, $A' \subseteq A$.

The following example presents TBA² authorization and administrative policies.

Example 2 (TBA²). Consider three subjects s_1 , s_2 , and s_3 and an object o belonging to a system governed by the British Navy, which are tagged as follows:

- $\text{tag}(s_1) = \{(\text{uk_navy}, \text{senior_officer})\}$
- $\text{tag}(s_2) = \{(\text{uk_navy}, \text{junior_officer})\}$
- $\text{tag}(s_3) = \{(\text{uk_navy}, \text{junior_officer})\}$
- $\text{tag}(o) = \{(\text{uk_navy}, \text{secret})\}$

The subjects' tags are issued by the British Navy (denoted as "uk_navy") and indicate that subject s_1 has rank *senior_officer*, while s_2 and s_3 have rank *junior_officer*. Object o is tagged by the British Navy as *secret*. The access to object o and the administration of rights within the system are regulated by the following rules:

1. $\text{allow}(S_x, o, \text{read}) :- (\text{eu, navy}) \in \text{tag}(S_y), (S_y, \text{senior_officer}) \in \text{tag}(S_x)$
2. $\text{allow}(S_x, O_x, \text{assign_tag}, \{(S_x, T_x)\}) :- (\text{eu, navy}) \in \text{tag}(S_y),$
 $(S_y, \text{senior_officer}) \in \text{tag}(S_x),$
 $(\text{eu, navy}) \in \text{tag}(S_z), (S_z, \text{secret}) \in \text{tag}(O_x)$
3. $\text{allow}(S_x, S_y, \text{assign_tag}, \{(S_x, \text{senior_officer})\}) :- (\text{eu, navy}) \in \text{tag}(S_z),$
 $(S_z, \text{senior_officer}) \in \text{tag}(S_x),$
 $(S_z, \text{junior_officer}) \in \text{tag}(S_y)$
4. $\text{allow}(S_x, O_x, \text{revoke_tag}, \{(S_x, T_x)\}) :- (S_x, T_x) \in \text{tag}(O_x)$
5. $\text{allow}(S_x, S_y, \text{revoke_tag}, \{(S_x, T_x)\}) :- (S_x, T_x) \in \text{tag}(S_y)$

The first rule is an authorization rule stating that object o can be read by a subject S_x if S_x is labeled as *senior officer* by an EU navy S_y . Rules 2, 3, 4 and 5 are administrative rules. Rule 2 allows senior officers of EU navies to assign tags to objects labeled as *secret* by any EU navy. Rule 3 allows senior officers of EU navies to assign a *senior_officer* tag to junior officers of the same navy (therefore delegating their rights). Finally, rules 4 and 5 allow the issuer of a tag to revoke the tag. Assuming tag $(\text{eu, navy}) \in \text{tag}(\text{uk_navy})$, the policy allows subject s_1 to read object o and to assign a *senior_officer* tag to subjects s_2 and s_3 .

3.2 Semantics of Administrative Policies

The effects of the exercise by a subject s of the administrative rights *assign_tag* and *revoke_tag* for a set of (signed) tags ST are shown in Figure 1. The effects of invoking $\text{allow}(s, o, \text{assign_tag}, ST)$ are intuitive, and imply that a tag st (for each $st \in ST$) is added to the set $\text{tag}(o)$. On the other hand, the assignment of a set of tags ST by a subject s to a subject s' can be seen as the delegation of some of the rights (or roles) of s to s' . Delegation can be implemented according to two models: *grant* or *transfer* [8]. In the grant model, after a successful delegation both s and s' are able to benefit from the delegated rights or roles. On the contrary, according to the transfer model subject s loses the delegated rights or roles. Figure 1(a) shows the effects of invoking $\text{allow}(s, s', \text{assign_tag}, ST)$ using the grant model. The transfer model would imply that all the tags in ST are removed from the set $\text{tag}(s)$ after being added to $\text{tag}(s')$.

Similarly to the tag assignment operation, also the effects of invoking $\text{allow}(s, s', \text{revoke_tag}, ST)$ depend on the revocation model employed by the system. To motivate the existence of different revocation models, we describe a scenario based on the tags and rules in Example 2. Assume that senior officer s_1 wants to temporarily delegate her rights to junior officer s_2 because of an emergency. Then, s_1 assigns a *senior_officer* tag to s_2 . Later, subject s_2 delegates her rights to s_3 , and accordingly assigns tag *senior_officer* to s_3 . When the emergency is over, s_1 returns to her regular duties and revokes the *senior_officer* tag from s_2 . Now, the question is whether s_3 's *senior_officer* tag should also be automatically revoked or not.

allow(s,o,assign_tag,ST)
$\forall st \in ST: \text{tag}(o) = \text{tag}(o) \cup \{st\}$

allow(s,s',assign_tag,ST)
$\forall st \in ST:$ $\text{tag}(s') = \text{tag}(s') \cup \{st\}$

 (a) Semantics of $\text{allow}(s, o, \text{assign_tag}, ST)$ and $\text{allow}(s, s', \text{assign_tag}, ST)$

allow(s,s',revoke_tag,ST)
$\forall st \in ST: \text{tag}(s') = \text{tag}(s') \setminus \{st\}$ let $ST = \{(s_1, t_1), \dots, (s_n, t_n)\}$ $\forall so \in S \cup O, t \in T$ such that $(s', t) \in \text{tag}(so)$ if $\exists A' \subseteq A$ such that $A', \text{tag} \models \text{allow}(s', so, \text{assign_tag}, \{(s', t)\})$ if $(s_1, t_1), \dots, (s_n, t_n) \in \text{tag}(s')$ and $\forall A'' \subseteq A$ we have that $A'', \text{tag} \not\models \text{allow}(s', so, \text{assign_tag}, \{(s', t)\})$ if $(s_1, t_1), \dots, (s_n, t_n) \notin \text{tag}(s')$ then if $so \in S$ then let ST' be the set of such tags (s', t) invoke $\text{allow}(s', so, \text{revoke_tag}, ST')$ else $\text{tag}(so) = \text{tag}(so) \setminus \{(s', t)\}$

 (b) Semantics of $\text{allow}(s, s', \text{revoke_tag}, ST)$ with Cascading Revocation

allow(s,s',revoke_tag,ST)
$\forall st \in ST: \text{tag}(s') = \text{tag}(s') \setminus \{st\}$ let $ST = \{(s_1, t_1), \dots, (s_n, t_n)\}$ $\forall so \in S \cup O, t \in T$ such that $(s', t) \in \text{tag}(so)$ if $\exists A' \subseteq A$ such that $A', \text{tag} \models \text{allow}(s', so, \text{assign_tag}, \{(s', t)\})$ if $(s_1, t_1), \dots, (s_n, t_n) \in \text{tag}(s')$ and $\forall A'' \subseteq A$ we have that $A'', \text{tag} \not\models \text{allow}(s', so, \text{assign_tag}, \{(s', t)\})$ if $(s_1, t_1), \dots, (s_n, t_n) \notin \text{tag}(s')$ then $\text{tag}(so) = \text{tag}(so) \setminus \{(s', t)\} \cup \{(s, t)\}$

 (c) Semantics of $\text{allow}(s, s', \text{revoke_tag}, ST)$ with Non-Cascade Revocation

allow(s,o,revoke_tag,ST)
$\forall st \in ST: \text{tag}(o) = \text{tag}(o) \setminus \{st\}$

 (d) Semantics of $\text{allow}(s, o, \text{revoke_tag}, ST)$
Fig. 1. Effects of the invocation of rights *assign_tag* and *revoke_tag*

Two main revocation models have been proposed in the literature. The first model, called *cascading* revocation [4, 12], aims to overturn all the changes to a system authorizations made exploiting the tags being revoked. This implies that if a subject s revokes a set of tags ST from a subject s' , then all tags subsequently assigned by s' (and by the subjects to which s' assigned a tag) without other supporting authorizations must be recursively revoked. The effects of invoking $\text{allow}(s, s', \text{revoke_tag}, ST)$ with cascading revocation are shown in Figure 1(b). A domain that typically resorts to cascading revocation is data protection. Whenever an individual revokes the consent (i.e., the right) to process her personal data to a service provider, all the authorizations on the data of the

service provider and of the subcontractors to whom the service provider delegated the processing of the data are revoked.

The dual model of cascading revocation is called *non-cascade* revocation [6] (or *simple* revocation in [4]). In non-cascade revocation, if a subject s revokes a set of tags ST from a subject s' , instead of revoking the tags that s' assigned exploiting the authorizations deriving from ST (as done by cascading revocation), these tags are modified as if they were issued by s . Intuitively, this requires s to be allowed to both revoke tags ST from subject s' and to assign tags ST in her place. The rationale behind non-cascade revocation is clarified by the following example. In most organizations, the authorizations that users possess are related to their role within the organization. Suppose there is a change in the role of a user s' . This may imply a change also in the privileges of s' : new rights will be granted to s' and some of her previous rights will be revoked. Applying cascading revocation would result in the undesirable effect of deleting all the authorizations that s' granted and, recursively, all the authorizations granted through them, which then might need to be re-issued. Moreover, all the tags assigned by s' that depend on the revoked rights would be invalidated. A better solution to this scenario is to preserve the authorizations granted by user s' , possibly substituting s' with another user as the grantor (i.e., issuer) of those authorizations. In [6], for instance, s' is replaced by the user s who is revoking her rights. The semantics of $\text{allow}(s, s', \text{revoke_tag}, ST)$ with non-cascade revocation is shown in Figure 1(c).

According to both semantics presented above, since objects cannot further delegate their rights to other subjects or objects, the effects of invoking $\text{allow}(s, o, \text{revoke_tag}, ST)$ are the same in cascading and non-cascade revocation (Figure 1(d)).

3.3 Tag Verification

Tag verification is the process of verifying the compliance of tags with administrative policies. More precisely, the goal of tag verification is to determine whether a user is (or was) authorized to assign a given tag. Typically, the enforcement of authorization and administrative policies within a system is achieved by means of a priori access control mechanisms. In a distributed system, however, relying exclusively on a priori mechanisms requires entities in one security domain to trust systems in different security domains for policy enforcement. Thanks to the observability of users' actions deriving from the signing of tags, TBA² allows security administrators to complement a priori mechanisms with a posteriori tag verification using a lightweight auditing mechanism.

Technically, we say that a tag t assigned by a subject s to an object o (respectively to a subject s') is *valid* if for a set of administrative rules A' and a set of signed tags ST such that $(s, t) \in ST$ we have that

$$A', \text{tag} \models \text{allow}(s, o, \text{assign_tag}, ST) \text{ (resp. } A', \text{tag} \models \text{allow}(s, s', \text{assign_tag}, ST))$$

Otherwise, we say that the tag is *invalid*. This validity check can be used both as an a priori mechanism for the enforcement of administrative policies and a posteriori for auditing purposes.

The verification of a tag (s, t) against the administrative policy of a system might require the verification of a set of tags $(s_1, t_1), \dots, (s_n, t_n)$ against the policies of systems

in different security domains. In fact, to verify whether tag (s, t) is valid, we need in turn to verify the validity of the *supporting tags* of (s, t) , i.e., the tags that authorized subject s to issue (s, t) . Consider, for instance, rule 3 in Example 2, which states that “senior officers of EU navies may assign a *senior_officer* tag to junior officers of the same navy”. To verify whether the *senior_officer* tag assigned by subject s_1 to subject s_2 is valid, we need first to confirm that s_1 is actually a senior officer and s_2 has a *junior_officer* tag issued by the same navy. This verification process is similar to the credential chain discovery problem in trust management. Accordingly, trust management algorithms [19, 25] can be employed to support the verification of tags’ validity.

An additional problem of tag verification is that in a distributed system entities in one security domain might not trust systems in different security domain to perform the validity check. In this respect, we identify three types of trust relationships that can be of interest for tag verification in TBA², resulting in three possible verification strategies. In what follows, we refer to the object (resp. subject) to which a tag is assigned as the *target* object (resp. subject) of the tag. The first possible verification strategy is *issuer verification* of tags, where the system to which the issuer of a tag belongs is trusted for checking the validity of the tag. The second strategy is *target verification*, which enables systems to verify the validity of the tags assigned to an object according to the intention of the object’s owner. A possible application scenario for target verification is the protection of digital media, where only the content owner is entitled to define the authorizations to access the object. Finally, *local verification* of tags can be employed in scenarios where there is no mutual trust among systems in different security domains. With local verification, the system which is interested in verifying the validity of a tag performs the check with respect to its local administrative policy.

The following example illustrates local verification of tags based on the administrative policy in Example 2.

Example 3 (Tag Verification). The French Navy distributes a document d containing the location of a suspected pirate attack to the other navies involved in Operation Atalanta. When d is received by the British Navy, it contains the following tags:

- $\text{tag}(d) = \{(\text{fr_navy}, \text{secret}), (s_4, \text{inaccurate_information})\}$

where subject s_4 is labeled as follows:

- $\text{tag}(s_4) = \{(\text{it_navy}, \text{reconnaissance_pilot})\}$

Since in Example 2 there is no rule defined by the British Navy that implies $\text{allow}(s_4, d, \text{assign_tag}, \{(s_4, \text{inaccurate_information})\})$, the tag added by subject s_4 is considered invalid. In fact, the policy of the British Navy allows only senior officers to tag documents marked as secret by an EU navy. The British Navy might thus decide to proceed with further investigations before concluding the inaccuracy of d ’s information.

For the sake of simplicity, the auditing mechanism discussed in this section verifies the compliance of tags with respect to the administrative policies that are currently in force. In other words, a tag is considered valid if its assignment is authorized by the administrative policy in force at the moment in which the tag is *verified*. An alternative verification mechanism could verify tags with respect to the administrative policy in force at the moment in which the tag was *assigned*. Intuitively, the implementation of

the latter mechanism is more complex and requires, e.g., timestamped tags and repositories containing all the administrative policies adopted by a system over time. An extension of the TBA² model in this direction is discussed in Section 6.

3.4 Tag Revocation

Whenever security administrators identify invalid tags, they should revoke them to preserve the consistency of function tag with respect to administrative policies. In a centralized system, revocation can be performed by simply deleting incorrect tags from the system. In a distributed system, revoking a set of tags is more complicated because it might not be possible to trace the tags issued by a given subject, and security administrators cannot delete tags assigned to subjects and objects residing in different security domains. TBA² enables a simple solution to this problem, where security administrators communicate the issuer-tag pairs to be revoked to other systems by broadcasting or publishing in an appropriate location *revocation lists* of tags. The recipient systems can then decide whether to revoke the listed tags or ignore the recommendation.

Definition 3 (Revocation List). *A revocation list is a triple $\langle s, so, T' \rangle$ where $T' \subseteq T$ is a set of tags, s is the issuer of the tags in T' , and so is the target subject or object.*

Intuitively, a revocation list contains the set T' of tags assigned by a subject s to subject or object so which should be revoked according to the system publishing the revocation list.

Example 4 (Revocation List). The revocation list for the invalid tag identified by the British Navy in Example 3 is the following: $\langle s_4, d, \{inaccurate_information\} \rangle$. The revocation list is published by the British Navy on its public record of invalid tags.

The decision of the security administrator of a system sys on whether to actually revoke the set of tags listed in a revocation list $\langle s, so, T' \rangle$ published by a system sys' is strictly correlated to the verification strategy employed by sys . If sys resorts to issuer (resp. target) verification, for instance, sys trusts the system to which s (resp. so) belongs to perform the validity check of tags T' . Consequently, if s (resp. so) belongs to sys' , the security administrator of sys is likely to delete tags T' from its system. On the contrary, if sys resorts to local verification, it might decide to proceed with further investigations before deleting the tags. The revocation list published by the British Navy in Example 4, for instance, might be taken into consideration by the vessels of the British Navy, but ignored by the vessels of other EU countries, because derived through local verification with respect to the British Navy's policy.

As an alternative to revocation lists, we consider the use of *negative tags*. Negative tags are signed tags that state that a certain tag assigned to a subject or object is not (or no longer) valid according to the issuer of the negative tag. The advantage of negative tags is that they enable the verification of revoked tags, which might not be possible if the tags are deleted. However, this might lead to a very large number of tags (both "positive" and negative) assigned to each subject and object.

Rather than simply deleting invalid tags (or issuing negative tags), other approaches can be employed for restoring the compliance of tags with administrative policies.

In some critical or uncertain situation, for instance, security administrators might decide to simply highlight the invalid tags and refer to a competent user for determining what to do with them. Alternatively, systems may rely on *repair constraints* [11] to determine how to handle invalid tags. A repair constraint might, for example, dictate in which conditions cascading rather than non-cascade revocation should be applied on a tag.

4 Evaluation of TBA²

In this section we evaluate the expressive power of the TBA² model. First, it is easy to demonstrate that TBA² is strictly more expressive than TBA. In fact, by not bounding the tags' issuers, TBA² can express exactly the same constraints definable by TBA. On top of this, associating an issuer to each tag enables the specification of authorization and administrative rules discriminating based on the issuer of tags, such as for instance linked roles in RT [18]. A linked role is a rule of the form $A.r \leftarrow B.r_1.r_2$, which states that subject A assigns a subject S_x (implicitly defined) to role r if S_x is labeled as a member of role r_2 by a subject S_y who is assigned to role r_1 by subject B . The reason why linked roles cannot be represented in TBA is that they require the binding of the subject of the first role r_1 to the issuer of the second role r_2 . TBA² rules 1, 2, and 3 in Example 2 are examples of RT's linked roles.

We now show how TBA² can represent policies from two reference administrative models proposed in the literature, namely the Harrison, Ruzzo, and Ullman (HRU) model [13] and ARBAC97 [23]. The HRU model [13] employs an access matrix for the specification of the rights of users on the objects in a system, and relies on a set of commands for modifying users' authorizations. The model includes three predefined commands: commands CONFER and REVOKE allow the owner of an object to respectively grant to and revoke from other subjects any right on the objects she owns; command TRANSFER allows users to delegate their rights to other users. In TBA² we use $\text{allow}(s, o, \text{assign_tag}, T')$ to define commands CONFER and TRANSFER, and $\text{allow}(s, o, \text{revoke_tag}, T')$ to define the REVOKE command. We assume the tags in T to consist of pairs $\langle s, r \rangle$, representing each possible right $r \in R$ of a subject $s \in S$. The tags associated to an object define the rights of the users of a system on that object.

Example 5 (Mapping HRU to TBA²). The following three rules define commands CONFER, TRANSFER, and REVOKE respectively:

1. $\text{allow}(S_x, O_x, \text{assign_tag}, \{(S_y, R_x)\}) :- (sys, \langle S_x, \text{own} \rangle) \in \text{tag}(O_x)$
2. $\text{allow}(S_x, O_x, \text{assign_tag}, \{(S_y, R_x)\}) :- (sys, \langle S_z, \text{own} \rangle) \in \text{tag}(O_x)$
 $(S_z, \langle S_x, R_x^* \rangle) \in \text{tag}(O_x)$
3. $\text{allow}(S_x, O_x, \text{revoke_tag}, \{(S_y, R_x)\}) :- (sys, \langle S_x, \text{own} \rangle) \in \text{tag}(O_x)$

The first rule states that the owner S_x of an object O_x can assign any right R_x on O_x to any subject S_y . The tag representing the ownership of an object is a "system tag"; we use *sys* to denote the issuer of system tags. Rule 2 represents the right of a subject S_x to delegate a right R_x on object O_x to a subject S_y , provided that R_x is a transferable right (denoted by symbol * in the HRU model) assigned to S_x by the owner S_z of object O_x . Finally, rule 3 states that the owner of an object can revoke any right on that object.

Note that the HRU model allows users to define additional commands to modify the authorizations within a system. Since those commands are arbitrary, and are not described in the model, we cannot evaluate the expressiveness of TBA² with respect to them.

Next, we show how to represent in TBA² the administrative policies supported by ARBAC97 [23], an administrative model for role-based access control. In ARBAC97, roles are divided into two classes: *administrative roles* and *regular roles*. Both classes of roles are organized into hierarchies, where each role inherits all the rights assigned to the children nodes in the hierarchy. The ARBAC97 model relies on four commands for the specification of administrative policies:

1. $can_assign(ar, \phi, \{rr_1, \dots, rr_n\})$
2. $can_revoke(ar, \{rr_1, \dots, rr_n\})$
3. $can_assignp(ar, \phi, \{rr_1, \dots, rr_n\})$
4. $can_revokep(ar, \{rr_1, \dots, rr_n\})$

where ϕ (called prerequisite condition) is a boolean expression on regular roles, which defines the requirements on the membership (or non-membership) of a user to some roles. Commands (1) and (2) are used to specify the right to assign and revoke roles, while commands (3) and (4) define the rights to assign and revoke permissions. More precisely, command (1) defines the right of a member of the administrative role ar (or a member of an administrative role above ar in the hierarchy) to assign to a user who satisfies the prerequisite conditions ϕ the membership to regular roles rr_1, \dots, rr_n . Command (2) assigns to members of the administrative role ar (or higher roles in the hierarchy) the right to revoke regular roles rr_1, \dots, rr_n . Similarly, command (3) allows members of the administrative role ar (or higher) to assign to roles rr_1, \dots, rr_n any permission whose assignment to regular roles satisfies ϕ , and command (4) enables members of ar (or higher) to revoke any right to roles rr_1, \dots, rr_n . To represent ARBAC97, we consider a set of administrative roles AR and regular roles RR to be defined as tags in T . The assignment of a user to a role is represented by the assignment of a tag from RR to the user. In addition, similarly to the previous example, we employ tags consisting of pairs $\langle s, r \rangle$ to represent a right $r \in R$ of a subject $s \in S$. Finally, for representing a prerequisite condition ϕ , we rewrite ϕ in disjunctive normal form, i.e., into a formula of the form $(CR_{11} \wedge \dots \wedge CR_{1m_1}) \vee \dots \vee (CR_{p1} \wedge \dots \wedge CR_{pm_p})$, where CR_{ij} (with $i \in \{1, \dots, p\}$, $j \in \{1, \dots, m_i\}$) is either cr_{ij} or $\neg cr_{ij}$, with $cr_{ij} \in RR$, and the negation symbol \neg denotes non-membership to a regular role. Negation as failure is employed to interpret negated roles: a user is not a member of a role cr_{ij} if she is not assigned a tag (s, cr_{ij}) , for any $s \in S$.

Example 6 (Mapping ARBAC97 to TBA²). The following TBA² rules define ARBAC97 commands (1), (2), (3), and (4) respectively:

1. $allow(S_x, S_y, assign_tag, \{(S_x, rr_1), \dots, (S_x, rr_n)\}) :- (*, ar) \in tag(S_x), (*, cr_{11}) \odot tag(S_y),$
 $\dots, (*, cr_{1m_1}) \odot tag(S_y)$
- ...
- $allow(S_x, S_y, assign_tag, \{(S_x, rr_1), \dots, (S_x, rr_n)\}) :- (*, ar) \in tag(S_x), (*, cr_{p1}) \odot tag(S_y),$
 $\dots, (*, cr_{pm_p}) \odot tag(S_y)$

2. $\text{allow}(S_x, S_y, \text{revoke_tag}, \{(*, rr_1), \dots, (*, rr_n)\}) :- (*, ar) \in \text{tag}(S_x)$
3. $\text{allow}(S_x, S_y, \text{assign_tag}, \{(S_x, \langle S_y, R_x \rangle)\}) :- (*, ar) \in \text{tag}(S_x), (*, rr_1) \in \text{tag}(S_y),$
 $(*, cr_{11}) \in \text{tag}(S_{cr_{11}}), (*, \langle S_{cr_{11}}, R_x \rangle) \odot \text{tag}(O_{cr_{11}}),$
 $\dots,$
 $(*, cr_{1m_1}) \in \text{tag}(S_{cr_{1m_1}}), (*, \langle S_{cr_{1m_1}}, R_x \rangle) \odot$
 $\text{tag}(O_{cr_{1m_1}})$
 \dots
 $\text{allow}(S_x, S_y, \text{assign_tag}, \{(S_x, \langle S_y, R_x \rangle)\}) :- (*, ar) \in \text{tag}(S_x), (*, rr_n) \in \text{tag}(S_y),$
 $(*, cr_{p1}) \in \text{tag}(S_{cr_{p1}}), (*, \langle S_{cr_{p1}}, R_x \rangle) \odot \text{tag}(O_{cr_{p1}}),$
 $\dots,$
 $(*, cr_{pm_p}) \in \text{tag}(S_{cr_{pm_p}}), (*, \langle S_{cr_{pm_p}}, R_x \rangle) \odot$
 $\text{tag}(O_{cr_{pm_p}})$
4. $\text{allow}(S_x, S_y, \text{revoke_tag}, \{(*, \langle S_y, R_x \rangle)\}) :- (*, ar) \in \text{tag}(S_x), (*, rr_1) \in \text{tag}(S_y)$
 \dots
 $\text{allow}(S_x, S_y, \text{revoke_tag}, \{(*, \langle S_y, R_x \rangle)\}) :- (*, ar) \in \text{tag}(S_x), (*, rr_n) \in \text{tag}(S_y)$

where $*$ indicates any subject in S , and \odot is either \in or \notin depending on the corresponding element in ϕ . The first set of rules allows a subject S_x with administrative role ar to assign to a subject S_y whose roles satisfy the formula $(CR_{i1} \wedge \dots \wedge CR_{im_i})$ (for any $i \in \{1, \dots, p\}$) to regular roles rr_1, \dots, rr_n . The second rule allows a member S_x of administrative role ar to revoke to a subject S_y regular roles rr_1, \dots, rr_n , independently from the subject who assigned them. The set of rules in item 3 states that a subject S_x who is a member of administrative role ar may assign a right R_x to a subject S_y , provided that S_y is a member of regular role rr_j (with $j \in \{1, \dots, n\}$), and R_x is a right whose assignment to regular roles satisfies $(CR_{i1} \wedge \dots \wedge CR_{im_i})$. Finally, the set of rules in item 4 gives to a subject S_x having administrative role ar the right to revoke any right to the members of role rr_j .

In the example above we do not consider inheritance of rights among roles in a hierarchy. Rather, we assume that a rule is defined for each administrative role ar having a certain right. A role hierarchy could be easily defined using a predicate $\text{higher_role}(ar_1, ar_2)$, and adding to each rule a condition $\text{higher_role}(ar, ar_{min})$, where ar_{min} is the minimum role in the hierarchy to which the rule applies. In addition, we slightly modify the semantics of commands (3) and (4). Whereas in ARBAC97 permissions are assigned to roles, in our representation they are assigned to the members of a role. From the practical point of view, however, the two semantics are equivalent.

The examples above demonstrate that TBA² can express the administrative constraints defined by HRU and ARBAC97. As a matter of fact, neither HRU nor ARBAC97 fully exploit the expressiveness of TBA². As shown by Example 6, for instance, ARBAC97 does not exploit the capability of TBA² of constraining the issuer of a tag and the rights that a member of an administrative role may assign. With respect to the HRU model, TBA² allows for the specification of much more complex constraints than those defined in commands *CONFER*, *TRANSFER*, and *REVOKE*, e.g., based on the properties of subjects and objects. This implies that, in terms of expressive power, TBA² represents a more comprehensive solution than the considered models.

5 Related Work

TBA has been studied informally in [20, 26], though that work allows tags on subjects but not on objects. Next to it, substantial work has been done on logical access control models, both based on Datalog (e.g., [2, 18, 22]) as well as on more expressive logics (e.g., [1, 9, 27]). While many of the existing logical access control languages can be used to encode tag-based authorization policies, it is the commitment to document and user tagging (an activity that can be carried out by users with a wide range of technical expertise) that makes TBA useful to a broad class of organizations.

The work related to the contributions of this paper spans two main topics: policy administration and auditing mechanisms. While many access control models for distributed systems have been proposed in the literature, policy administration received much less consideration. A number of administration models exist [4, 5, 10, 13, 17, 23], but they focus mainly on the expressivity of administrative policies, and do not consider the challenges associated with their enforcement in a distributed setting. The innovation of TBA² in this respect lies in the fact that it allows for an easy verification of policy compliance, thus not requiring entities in one security domain to trust systems in different security domains for the enforcement of administrative policy. In addition, we have shown that TBA² is more expressive than two reference administrative models, namely ARBAC97 [23] and HRU [13].

Similarly to TBA², the existing a posteriori solutions (e.g., [7]) perform the verification of policy compliance through auditing mechanisms. However, to achieve this, they rely on logging mechanisms that record users' actions, and trusted auditing authorities that verify the compliance of those actions with policies. Our model represents a lightweight solution for policy compliance verification that does not require the realization of such an auditing infrastructure. We propose the use of trust management algorithms [19, 25] to support the verification of policy compliance.

6 Discussion and Conclusions

In this paper we have introduced TBA², an extension of the TBA model [15] that enables policy administration in distributed systems. Similarly to TBA, TBA² allows relatively untrained users to assign descriptive tags to a system's subjects and objects; trained security experts then write logic-based authorization policies that define access rights in terms of those tags. In addition, by linking each tag to its issuer (i.e., the user who assigned it), TBA² enables the specification of fine-grained administrative policies whose enforcement can be verified through a lightweight auditing technique. We have shown that our model is more expressive than TBA and than the HRU [13] and ARBAC97 [23] administrative models. Thus, TBA² represents a flexible, easy to use, yet expressive access control solution which matches the needs of real-world organizations.

The auditing mechanism proposed in Section 3.3 verifies tags' validity with respect to the administrative policy currently in force within a system. In some situations, however, it is preferable to verify the validity of a tag with respect to the administrative policies effective when the tag was issued. For example, assume that the commanding officer of a British Navy vessel is summoned by the EU for a meeting at the Operation

Atalanta's headquarter. Then, the commanding officer would have to temporarily delegate the command of the vessel and the deriving responsibilities and authorizations to another officer until her return. During this period, the appointed officer will have to take several decisions which might lead to the granting and revocation of authorizations to the vessel's operators and to the tagging of several data objects exchanged among the collaborating navies. With the verification mechanism presented in Section 3.3, the revocation of the officer's rights by the commanding officer upon her return would have the undesirable effect of invalidating all the authorizations and tags assigned by the officer during her command. The design of an auditing mechanism verifying tags' validity with respect to the administrative policy in force when a tag was assigned would require two main extensions to the TBA² model. First, it would require the association of a timestamp to each tag to demonstrate when it was issued. Second, all the administrative policies employed by a system during its lifetime would need to be stored in a repository, together with the time interval in which they were effective. Then, whenever a tag needs to be verified, its timestamp can be used to retrieve from the repository the policy that was in force when the tag was issued, against which the validity check must be performed. The resulting enforcement mechanism is similar to those used for the enforcement of history-based access control policies [16].

To conclude, we point out that the model proposed in this paper enables security administrators to verify the compliance of users' actions with respect to the administrative policies in force within a system, but provides no guarantee that these policies are correctly specified. The verification of administrative policies with respect to the desired security properties of a system can be achieved through model checking techniques [28]. Finally, we argue that even though TBA² is presented as an access control solution for distributed systems, also centralized systems would benefit from employing the model. In fact, the association of each tag to its issuer enhances the "observability" of user's actions, simplifying the detection of policy violations, and may be used as a discriminant by other users in the system to determine whether a certain tag should be considered valid. Signed tags are currently employed by several existing web applications and social networks (e.g., Facebook).

Acknowledgments. This work has been done in the context of the THeCS project, which is supported by the Dutch national program COMMIT. Adam J. Lee was supported in part by the US National Science Foundation under awards CNS-0964295 and CNS-1228697.

References

1. Abadi, M., Burrows, M., Lampson, B.: A calculus for access control in distributed systems. *ACM Transactions on Programming Languages and Systems* 15(4), 706–734 (1993)
2. Becker, M.Y., Fournet, C.Y., Gordon, A.D.: SecPAL: Design and semantics of a decentralized authorization language. *Journal of Computer Security* 18(4), 619–665 (2010)
3. Bell, D.E.: Looking Back at the Bell-La Padula Model. In: *Proceedings of ACSAC 2005*, pp. 337–351. IEEE Computer Society (2005)

4. Ben-Ghorbel-Talbi, M., Cuppens, F., Cuppens-Boulahia, N., Bouhoula, A.: Revocation Schemes for Delegation Licences. In: Chen, L., Ryan, M.D., Wang, G. (eds.) ICICS 2008. LNCS, vol. 5308, pp. 190–205. Springer, Heidelberg (2008)
5. Ben-Ghorbel-Talbi, M., Cuppens, F., Cuppens-Boulahia, N., Bouhoula, A.: A delegation model for extended RBAC. *Int. J. Inf. Sec.* 9(3), 209–236 (2010)
6. Bertino, E., Samarati, P., Jajodia, S.: An Extended Authorization Model for Relational Databases. *IEEE Trans. Knowl. Data Eng.* 9(1), 85–101 (1997)
7. Cederquist, J.G., Corin, R., Dekker, M.A.C., Etalle, S., den Hartog, J.I., Lenzini, G.: Audit-based compliance control. *Int. J. Inf. Sec.* 6(2), 133–151 (2007)
8. Crampton, J., Khambhammettu, H.: Delegation in role-based access control. *Int. J. Inf. Sec.* 7(2), 123–136 (2008)
9. Crampton, J., Loizou, G., Oshera, G.: A logic of access control. *The Computer Journal* 44(1), 137–149 (2001)
10. Dekker, M., Crampton, J., Etalle, S.: RBAC administration in distributed systems. In: Proceedings of SACMAT 2008, pp. 93–102. ACM (2008)
11. Greco, G., Greco, S., Zumpano, E.: A logical framework for querying and repairing inconsistent databases. *IEEE Trans. Knowl. Data Eng.* 15(6), 1389–1408 (2003)
12. Griffiths, P.P., Wade, B.W.: An authorization mechanism for a relational database system. *ACM Trans. Database Syst.* 1(3), 242–255 (1976)
13. Harrison, M.A., Ruzzo, W.L., Ullman, J.D.: Protection in operating systems. *Communications of the ACM* 19(8), 461–471 (1976)
14. Heeps, S., Sventek, J., Dulay, N., Schaeffer Filho, A.E., Lupu, E., Sloman, M., Strowes, S.: Dynamic Ontology Mapping for Interacting Autonomous Systems. In: Hutchison, D., Katz, R.H. (eds.) IWSOS 2007. LNCS, vol. 4725, pp. 255–263. Springer, Heidelberg (2007)
15. Hinrichs, T.L., Garrison III, W.C., Lee, A.J., Saunders, S., Mitchell, J.C.: TBA: A Hybrid of Logic and Extensional Access Control Systems. In: Barthe, G., Datta, A., Etalle, S. (eds.) FAST 2011. LNCS, vol. 7140, pp. 198–213. Springer, Heidelberg (2012)
16. Koshutanski, H., Martinelli, F., Mori, P., Vaccarelli, A.: Fine-grained and History-based Access Control with Trust Management for Autonomic Grid Services. In: Proceedings of ICAS 2006, pp. 34–43. IEEE Computer Society (2006)
17. Li, N., Mao, Z.: Administration in role-based access control. In: Proceedings of ASIACCS 2007, pp. 127–138. ACM (2007)
18. Li, N., Mitchell, J.C., Winsborough, W.H.: Design of a Role-Based Trust-Management Framework. In: Proceedings of S&P 2002, pp. 114–130. IEEE Computer Society (2002)
19. Li, N., Winsborough, W.H., Mitchell, J.C.: Distributed credential chain discovery in trust management. *Journal of Computer Security* 11(1), 35–86 (2003)
20. Najafian Razavi, M., Iverson, L.: Supporting selective information sharing with people-tagging. In: Proceedings of CHI 2008, pp. 3423–3428. ACM (2008)
21. Osborn, S., Sandhu, R., Munawer, Q.: Configuring role-based access control to enforce mandatory and discretionary access control policies. *ACM Trans. Inf. Syst. Secur.* 3(2), 85–106 (2000)
22. Ribeiro, C., Zuquete, A., Ferreira, P., Guedes, P.: SPL: An access control language for security policies with complex constraints. In: Proceedings of NDSS 2011 (2001)
23. Sandhu, R., Bhamidipati, V., Munawer, Q.: The ARBAC97 model for role-based administration of roles. *ACM Trans. Inf. Syst. Secur.* 2(1), 105–135 (1999)
24. Trivellato, D., Spiessens, F., Zannone, N., Etalle, S.: Reputation-Based Ontology Alignment for Autonomy and Interoperability in Distributed Access Control. In: Proceedings of CSE 2009, vol. 3, pp. 252–258. IEEE Computer Society (2009)

25. Trivellato, D., Zannone, N., Etalle, S.: GEM: a Distributed Goal Evaluation Algorithm for Trust Management. *Journal of Theory and Practice of Logic Programming* (2012) (to appear)
26. Wang, Q., Jin, H., Li, N.: Usable Access Control in Collaborative Environments: Authorization Based on People-Tagging. In: Backes, M., Ning, P. (eds.) *ESORICS 2009*. LNCS, vol. 5789, pp. 268–284. Springer, Heidelberg (2009)
27. Wijesekera, D., Jajodia, S.: Policy algebras for access control - the predicate case. In: *Proceedings of CCS 2001*, pp. 171–180. ACM (2001)
28. Zhang, N., Ryan, M., Guelev, D.P.: Synthesising verified access control systems through model checking. *Journal of Computer Security* 16(1), 1–61 (2008)