

Confidentiality-Preserving Distributed Proofs of Conjunctive Queries

Adam J. Lee[†]
adamlee@cs.pitt.edu

Kazuhiro Minami[‡]
minami@cs.illinois.edu

Nikita Borisov[‡]
nikita@illinois.edu

[†]University of Pittsburgh

[‡]University of Illinois at Urbana–Champaign

ABSTRACT

Distributed proof construction protocols have been shown to be valuable for reasoning about authorization decisions in open distributed environments such as pervasive computing spaces. Unfortunately, existing distributed proof protocols offer only limited support for protecting the confidentiality of sensitive facts, which limits their utility in many practical scenarios. In this paper, we propose a distributed proof construction protocol in which the release of a fact’s truth value can be made contingent upon facts managed by other principals in the system. We formally prove that our protocol can safely prove conjunctions of facts without leaking the truth values of individual facts, even in the face of colluding adversaries and fact release policies with cyclical dependencies. This facilitates the definition of context-sensitive release policies that enable the conditional use of sensitive facts in distributed proofs.

Categories and Subject Descriptors: C.2.4 [Distributed Systems]: Distributed applications; D.4.6 [Operating Systems]: Security and Protection—*access controls, authentication*; K.6.5 [Management of Computing and Information Systems]: Security and Protection

General Terms: Security

Keywords: Consistency, distributed proof, pervasive computing

1. INTRODUCTION

Distributed proof systems allow autonomous agents to reason about their surrounding environment. In these systems, an agent can make inferences using facts stored in its own knowledge base, as well as facts stored in the knowledge bases of others. This ability to derive new local knowledge based on remote facts allows external information defining the context of a given system to be formally included in the decision making process. As a result, recent years have seen the use of distributed proof systems become popular

for making authorization decisions in open distributed environments, such as pervasive computing spaces [1, 2, 7, 22, 29]. Unfortunately, these existing protocols offer only limited support for protecting the confidentiality of facts whose release may be sensitive, as is often the case in pervasive environments.

Many distributed proof construction protocols assume that *every* principal is willing to disclose the truth value of *any* fact in its knowledge base to other principals in the system. Some systems improve upon this by allowing the use of identity-based access control lists (ACLs) to limit the disclosure of sensitive facts. However, these ACLs cannot incorporate system context to further limit information flow and are thus insufficient for many application domains. For example, consider a media controller application that controls a projector in a pervasive computing space. The controller may wish to allow access to some projector only to users who are presently located in the same room as the projector. Using identity-based ACLs, a location tracker would need to authorize the media controller to check users’ locations, which means a compromise of the media controller would result in violation of users’ location privacy. A better alternative would be to give the media controller access to users’ locations only when they are requesting access to the projector in the first place.

Motivated by practical scenarios such as the above, we develop a distributed proof construction protocol that allows such context-sensitive release policies, where the disclosure of a fact may be conditional on other facts managed in remote knowledge bases. For example, the disclosure of some fact f in a principal p ’s local knowledge base may depend upon facts f'_1, \dots, f'_n that exist in up to n remote knowledge bases. Our protocol allows access to the truth value of f only when the other facts f'_1, \dots, f'_n are also true *simultaneously* with f . If the conjunction $f \wedge f'_1 \wedge \dots \wedge f'_n$ is false, then the status of each individual fact remains hidden. We note that in our protocol, p need not be able to access the facts f'_1, \dots, f'_n itself in order to include them in the release policy for a given fact. In the context of the above example, this implies that the location tracker need not learn that a user is accessing the projector in order to place this contingency upon the release of the user’s location.

In this paper, we make the following contributions:

- We formalize the notion of confidentiality-preserving distributed proof, in which the disclosure of a fact’s status can be made contingent upon the simultaneous truth of facts maintained in remote knowledge bases.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ASIACCS '09, March 10–12, Sydney, NSW, Australia
Copyright 2009 ACM 978-1-60558-394-5/09/03 ...\$5.00.

- We develop a “best case” trusted third party (TTP) model that embodies the ideal functionality of a confidentiality-preserving distributed proof system.
- We design a distributed algorithm that can be used to construct confidentiality-preserving distributed proofs. We prove that this algorithm reveals no more information than can be learned by interacting with the TTP, even in the face of malicious and colluding protocol participants.
- We show that our protocol functions correctly even if the release policies defined by multiple principals involved in a single proof have circular dependencies.
- We use our TTP model to reason about the types of inferences that can be made regarding confidential facts if concurrent executions or multiple runs of the distributed proof process are allowed. We then show that principals in the system can make local decisions to limit the types of inference that can be made.

The remainder of this paper is organized as follows. In Section 2, we formally define both our system model and the problem of confidentiality-preserving distributed proofs. Section 3 presents cryptographic primitives and a distributed algorithm that can be used to construct confidentiality-preserving distributed proofs. In Section 4, we prove the soundness of this protocol, even in the face of malicious participants. We then explore the types of inferences made possible by concurrent executions or multiple runs of our proof protocol. We show that these risks are minimal and can be greatly reduced through either intelligent release policy design or query rate limiting. Lastly, we compare our approach to other related work in Section 5 and present our conclusions in Section 6.

2. ENVIRONMENT AND DEFINITIONS

In this section, we first describe the system model that will be used to formally reason about the distributed proof construction process. Within the context of this model, we then present a formal definition of the problem that this paper sets out to solve.

2.1 Basic System Model

We first describe the confidentiality-preserving distributed proof problem by means of an example that we will return to throughout this paper. Figure 1 shows several principals that exist within the same pervasive computing space. In this example, a user Bob wishes to access a digital projector located in a conference room. Other principals in the system include a media controller, *mc*, which mediates all access attempts to shared computing resources in the space; an inventory server, *is*, that keeps track of which resources are owned by which principals; a location service, *ls*, which tracks the locations of principals and devices in the space; and a role server, *rs*, that manages the various roles that can be taken on by principals in the space.

Each principal in the space manages a Datalog knowledge base that contains facts and derivation rules that allow the principal to reason about its surroundings. A Datalog fact is a predicate symbol followed by zero or more terms, where each term is either a lower-case or numerical constant, or a variable (denoted by an upper-case letter). For example, the location server’s knowledge base contains the base fact

$location(bob, 2124)$, which indicates that Bob is currently located in room 2124. Derivation rules are used to make inferences based on both facts in a principal’s own knowledge base, as well as facts that exist in other principals’ knowledge bases. For example, the location server’s derivation rule $colocated(P_1, P_2) \leftarrow location(P_1, L) \wedge location(P_2, L)$ relies on the use of local facts to determine whether two principals are colocated. By contrast, the media controller’s derivation rule $grant(U, P) \leftarrow “ls \text{ says } colocated(U, P)” \wedge “rs \text{ says } role(U, presenter)”$ uses *quoted facts* to determine whether some user can access a given projector based upon facts in other knowledge bases. Note that P_1, P_2, U and P are free variables that can be bound at runtime to allow these rules to apply to arbitrary principals.

Lastly, we wish to allow principals to control the disclosure of their own local facts to remote parties. This is accomplished by means of access control lists (ACLs) that can be defined by each principal. In general, the ACL for a given fact identifies the conditions under which that fact can be disclosed to certain principals. For example, the role server’s access control list indicates that the media controller is allowed to see all facts of the form $role(U, presenter)$, where U is a free variable that can be bound at runtime. In short, the media controller is always allowed to know which principals can take on the role “presenter.” The location server’s ACL for facts of the form $colocated(U, D)$ is more restrictive, however, as it will only release these types of facts to the media controller if the inventory server says that the media controller owns the device that will be bound to the free variable D and the user U indicates that they are indeed trying to access the device D . Note that this strictly limits the ability of the media controller to track the locations of users in the system. This is in contrast to systems using only identity-based ACLs (as was discussed in Section 1).

2.2 Formal System Model

We will now more formally define the concepts introduced in Section 2.1. We will let \mathcal{P} represent the set of all principals in the system. We will assume that any two principals $p_i, p_j \in \mathcal{P}$ can establish a private and authenticated communication channel between them by using a PKI or some other method of key distribution.

We use the symbol \mathcal{F} to denote the set of all facts and the symbol \mathcal{Q} to denote the set of all quoted facts of the form “ p_i says f_i ”; i.e., $\mathcal{Q} = \mathcal{P} \times \mathcal{F}$. Each principal p_i maintains a knowledge base $KB_i \subset \mathcal{F} \cup \mathcal{R}$, containing facts and derivation rules. Since any fact f_i stored in the knowledge base of a principal p_i can be written as the quoted fact “ p_i says f_i ”, we model the set \mathcal{R} of derivation rules as Horn clauses whose head is a fact and whose body is a conjunction of quoted facts; i.e., $\mathcal{R} = \mathcal{F} \times 2^{\mathcal{Q}}$. The ACL protecting facts managed by principal p_i is represented as a function $acl_{p_i} : \mathcal{F} \rightarrow 2^{2^{\mathcal{P}} \times 2^{\mathcal{Q}}}$ that maps a given fact in $f \in KB_i$ to a set of $(plist, qlist)$ pairs. Each *plist* identifies a set of principals who can access f if and only if they can also verify that each q_i in the corresponding *qlist* is also true. For example, the ACL defined by Bob in Figure 1 allows any principal P to access facts of the form $request(D)$, provided that the inventory server asserts that principal P is the owner of the device D .

In the remainder of this work, we assume that the construction of distributed proofs will take place in an asynchronous distributed system and thus we place no limita-

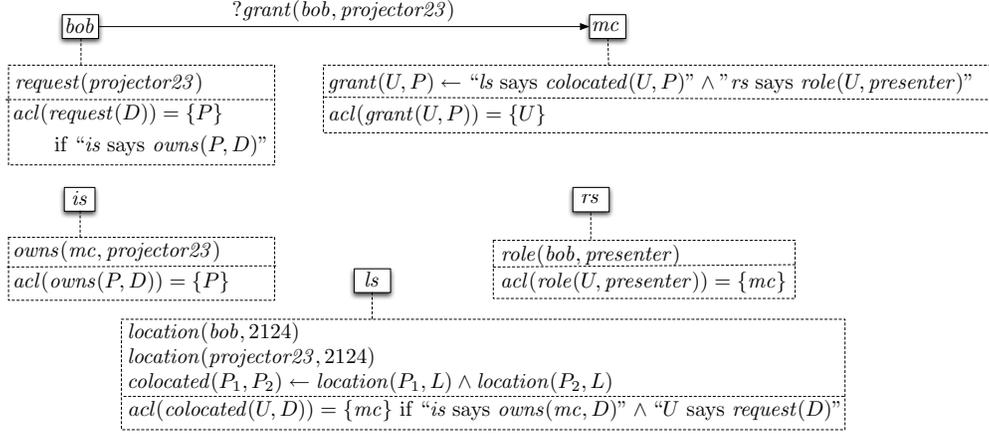


Figure 1: Knowledge bases and ACLs in an example pervasive computing system. Each principal’s knowledge bases and ACL are enclosed in a dashed rectangle.

tions on the temporal duration of the distributed proof construction process. Furthermore, we assume that each principal p_i acts autonomously and can add or remove facts from its knowledge base KB_i at any time *without* notifying other principals in the system. Lastly, we do not assume that any level of clock synchronization exists between the principals in \mathcal{P} .

2.3 Confidentiality-Preserving Distributed Proof

The high-level goal of confidentiality-preserving distributed proof is maximize the utility of a distributed proof environment while minimizing the amount of sensitive information disclosed by principals in the system. Since access control policies are often considered sensitive (e.g., see [20, 31]), we assume that ACLs are not publicly available. However, we do assume that for any two principals p_i and p_j , p_j can access the portions of p_i ’s ACL that are relevant to p_j . We call this the *restriction* of p_i ’s ACL relative to principal p_j .

DEFINITION 1 (RESTRICTED ACL). *For any two principals p_i and p_j , the restriction of p_i ’s ACL relative to p_j is represented as the function $\widehat{acl}_{p_i, p_j} : \mathcal{F} \rightarrow 2^{2^{\mathcal{Q}}}$, which is defined as:*

$$\widehat{acl}_{p_i, p_j}(f) \equiv \{qlist \mid (plist, qlist) \in acl_{p_i}(f) \wedge p_j \in plist\}$$

Intuitively, the restricted ACL $\widehat{acl}_{p_i, p_j}(f_i)$ specifies a set of dependencies used to control the disclosure of the quoted fact “ p_i says f_i ” to principal p_j . However, each of these dependencies may also have other dependencies. To address this, p_j can use repeated application of other principals’ restricted ACLs to compute all possible sets of quoted facts upon which the disclosure of “ p_i says f_i ” depends. For example, in the scenario depicted in Figure 1, the media controller can use $\widehat{acl}_{ls, mc}$ to determine that the disclosure of the fact “ ls says $colocated(bob, projector23)$ ” depends upon the truth of facts “ is says $owns(mc, projector23)$ ” and “ bob says $request(projector23)$ ”. The media controller can then check $\widehat{acl}_{is, mc}$, and $\widehat{acl}_{bob, mc}$ to verify that no other facts need to be incorporated in the conjunction to be proved.

As a result of the above observation, it is natural to formally define the ideal functionality of a confidentiality-preserving distributed proof system in terms of a trusted third party (TTP) that can be used as an oracle that proves conjunctions of facts. Prior to specifying this functionality, we first define the predicates $checkAcls : \mathcal{P} \times 2^{\mathcal{Q}} \rightarrow \mathbb{B}$ and $oracle : 2^{\mathcal{Q}} \rightarrow \mathbb{B}$ as follows.

$$checkAcls(p_0, conj) \equiv \forall “p_i \text{ says } f_i” \in conj, \quad (1)$$

$$\exists (plist, qlist) \in acl_{p_i}(f_i) : \\ p_0 \in plist \wedge qlist \subseteq conj$$

$$oracle(conj) \equiv \forall “p_i \text{ says } f_i” \in conj : \quad (2)$$

$$f_i \in KB_i$$

For a given principal p_0 trying to prove the conjunction of quoted facts $conj$, the predicate $checkAcls(p_0, conj)$ is true if and only if allowing p_0 access to any “ p_i says f_i ” $\in conj$ does not require access to any quoted fact “ p_k says f_k ” $\notin conj$. The predicate $oracle(conj)$ is satisfied if and only if each “ p_i says f_i ” $\in conj$ exists in the appropriate knowledge base. We now define the ideal functionality of a confidentiality-preserving distributed proof system.

DEFINITION 2 (IDEAL FUNCTIONALITY). *A confidentiality-preserving distributed proof system can be modeled as a trusted third party that executes conjunctive queries on behalf of the principals in the system. When a principal p_0 wishes to prove the conjunction $C = “p_1 \text{ says } f_1” \wedge \dots \wedge “p_n \text{ says } f_n”$, the TTP takes the following actions:*

- For each “ p_i says f_i ” $\in C$, notify p_i that p_0 has an interest in p_i ’s fact f_i .
- If $checkAcls(p_0, C)$ is false, send a query failure message to p_0 . This failure message is distinct from the value false.
- If $checkAcls(p_0, C)$ is true, reveal the value of $oracle(conj)$ to p_0 .

Note that the above definition of a confidentiality-preserving distributed proof system entails several desirable

properties. First, it allows each principal’s ACL to remain private, as only access to restricted ACLs is needed to determine the minimal conjunction associated with a given fact of interest. Second, the truth value of a conjunction is revealed if and only if the ACLs of each fact within the conjunction are satisfied. Third, if the conjunction is false, the querier does not learn any information about the facts comprising the conjunction. Fourth, fact providers learn which individuals are interested in the facts in their knowledge bases, which allows them to audit usage patterns. Lastly, fact providers learn *neither* what conjunctions are being proved *nor* the status of other facts in the system, and thus cannot infer portions of other principals’ ACLs. In the remainder of this paper, we describe how this functionality can be achieved in proof systems designed for use in asynchronous distributed environments.

3. PROOF CONSTRUCTION PROTOCOL

In this section, we describe our confidentiality-preserving distributed proof protocol. We begin with an overview of our protocol, which motivates discussion of the cryptographic primitives that will be used by participants to achieve the properties described in the previous section. We then present the algorithmic details of our protocol.

3.1 Protocol Intuition

Our confidentiality-preserving distributed proof protocol operates in two distinct phases. Assume that some querier p_0 wants to prove the conjunction of facts $conj = “p_1 \text{ says } f_1” \wedge \dots \wedge “p_n \text{ says } f_n”$. In the first phase of the protocol, the querier contacts each p_i to indicate an interest in their fact f_i . Each p_i generates a random number $s_{i,j}$ (henceforth called a *share*) for each “ p_j says f_j ” in $\widehat{acl}_{p_i,p_0}(f_i)$, encrypts each random share to its corresponding p_j , and returns the set of encrypted shares to the querier. Principal p_i then generates a final share, s_i , which is stored locally, such that the product of all shares is 1.

During the second phase of the protocol, p_0 homomorphically combines all of the shares encrypted for a particular principal into a single share; this hides the number of principals whose release policies depend on a given p_i ’s fact f_i from non-colluding principals participating in the protocol. Additionally, p_0 adds a blinding factor to each share in order to ensure that colluding principals cannot determine the number of principals whose release policies depend on the fact “ p_i says f_i ”. Each resulting encrypted share is then forwarded to the corresponding principal p_i . Principal p_i decrypts this ciphertext and combines it with the local secret share s_i that it generated during phase one of the protocol. If fact f_i was true from the time that phase one ended until the time that phase two began, p_i will return this newly-generated result. Otherwise, a random value will be returned. If the product of the results returned by each p_i equals the product of the blinding factors generated by p_0 , then p_0 can determine that each “ p_i says f_i ” $\in conj$ was simultaneously true at some time t during the protocol. If *any* p_i returns a random result, p_0 will determine that $conj$ is false, but cannot recover any other information regarding the truth of subclauses of $conj$. We now present a cryptographic construction that allows us to build such a protocol, and then discuss the full details of our confidentiality-preserving distributed proof construction protocol.

3.2 Cryptographic Foundations

At a first glance, the encryption scheme can be any public key scheme that supports homomorphic operations, such as ElGamal [10] or Pailler [24]. But notice that when p_i encrypts a share s_j for principal p_j , this occurs because p_i has an ACL which includes the clause “if p_j says f_j ”. However, p_i has no guarantee that p_0 will not forward the share to p_j while asking about some other fact, f'_j , thus violating the ACL constraints. Therefore, p_i needs to bind its encrypted share to the particular fact f_j . This could be done by encrypting $E_{K_j}(s_j||f_j)$ and using a non-malleable encryption scheme [8]. Non-malleability ensures that p_0 could not change an encryption of $s_j||f_j$ to an encryption of $s_j||f'_j$. However, a non-malleable encryption by definition cannot be homomorphic and thus cannot satisfy our needs.

Instead, we use a homomorphic identity-based encryption (IBE) scheme. An IBE scheme has a master public key (MPK) that can be used to encrypt to a given user ID, without knowing the user’s specific public key: $\text{Enc}(M, \text{MPK}, \text{ID})$. The master secret key can be used to derive (extract) a private key for a particular user, which can then be used to decrypt all messages destined for that ID. In our case, rather than having a single master key, with identities referring to individual principals, we have each principal create a master key for an IBE scheme, with the *facts* in its knowledge base acting as identities. In other words, to encrypt a share for p_j , p_i computes $\text{Enc}(s_{i,j}, K_j, f_j)$, where K_j is the master public key of principal j . This binds the encrypted share to the fact f_j .

The scheme we use is a homomorphic variant of the Boneh–Franklin IBE scheme [5], originally proposed by Ivan and Dodis [12]. The primitive is based on a bilinear mapping $e : G_1 \times G_1 \rightarrow G_2$ that satisfies the Bilinear Diffie–Hellman assumption [5]. We let g be a generator of G_1 , and use a hash function $h : \{0, 1\}^* \rightarrow G_1$ to define identities. The scheme is defined as follows:

Enc-Gen: Generate a secret key $s \in_R G_1$ and a master public key g^s .

Extract(ID): Extract the secret key $h(\text{ID})^s$.

Enc(M, g^s, ID): $(g^r, M \cdot e(h(\text{ID})^r, g^s))$, for a random r .

Dec(U, V, ID): $V/e(h(\text{ID})^s, U) = M$.

\otimes : $(U, V) \otimes (U', V') = (U \cdot U', V \cdot V')$

This homomorphism over the operator \otimes works because the following two equations hold:

$$U \cdot U' = g^r \cdot g^{r'} = g^{r+r'} \quad (3)$$

$$\begin{aligned} V \cdot V' &= M \cdot e(h(\text{ID})^r, g^s) \cdot M' \cdot e(h(\text{ID})^{r'}, g^s) \quad (4) \\ &= M \cdot M' \cdot e(h(\text{ID}), g^s)^r \cdot e(h(\text{ID}), g^s)^{r'} \\ &= M \cdot M' \cdot e(h(\text{ID}), g^s)^{r+r'} \\ &= M \cdot M' \cdot e(h(\text{ID})^{r+r'}, g^s) \end{aligned}$$

THEOREM 1. *The above Homomorphic IBE cryptosystem is IND-ID-CPA secure.*

The theorem can be proven by following the same argument as the original Boneh–Franklin IBE scheme [5]. This means that an adversary who has a ciphertext encrypted for a certain identity ID can perform key extraction queries for

any number of identities $ID' \neq ID$, and yet is unable to infer anything about the plaintext. Note that the scheme does not (and cannot) offer chosen ciphertext security (IND-ID-CCA) because it allows homomorphic operations. To address this issue, we add a session identifier, sid , to the fact name as the identity for encryption. In other words, each p_i produces $\text{Enc}(s_j, K_j, p_0 \| f_j \| sid)$. This ensures that while the underlying cryptography is not secure against chosen-ciphertext attacks, our protocol can guarantee that a principal will only decrypt a single message for a given identity, thereby rendering chosen ciphertext attacks impossible.

3.3 Protocol Details

We now describe the details of our confidentiality-preserving distributed proof construction protocol. In the remainder of this paper, we assume that each fact f_i in a principal p 's knowledge base KB is associated with an identifier from the set $\{0, 1\}^\ell$. Any time that a fact's status changes, its corresponding identifier will be changed at random. As we will see later in this section, this notion of fact identifiers will be used to allow queriers to ensure that all facts comprising the conjunction being proved were valid simultaneously during the proof construction protocol. Prior to explaining the full details of our protocol—which is presented in Algorithms 1 and 2—we first note the following:

- Each principal p_i acts as a key generator for an instance of the IBE scheme described in Section 3.2. The master public key of this instance of the IBE system are then made available to other principals in the system using the same mechanism used to distribute p_i 's public key K_i .
- The sets M and C represent the message and ciphertext spaces of the IBE cryptosystem, respectively. (G_2 and $G_1 \times G_2$.)
- The symbol \leftarrow_R denotes random assignment from some set. For example, $sid \leftarrow_R \mathbb{Z}_{2^{128}}$ chooses a random session identifier from the set of all 128-bit integers.
- The identifier ME is used as a placeholder for a given principal's unique identifier.
- The notation $E_{p_j \| f_i \| sid}^{p_i}(x)$ is used to denote the encryption of item x to principal p_i relative to the IBE key derived from p_j 's identity, the fact f_i , and the session ID sid . Similarly, $D_{p_j \| f_i \| sid}^{p_i}(c)$ represents the decryption of the ciphertext c by principal p_i using the IBE key derived from p_j 's identity, the fact f_i , and the session ID sid .

Phase One: Share gathering

To initiate the confidentiality-preserving distributed proof process, the querier p_0 runs the code listed in Algorithm 1. The function $\text{STARTQUERY}(conj)$ is invoked whenever p_0 wants to learn the status of some conjunction of quoted facts $conj \in 2^{\mathcal{Q}}$ managed in one or more remote knowledge bases. The querier first generates a random session identifier¹ and creates a random blinding factor for each quoted fact “ p_i says f_i ” $\in conj$. Each random blinding factor is homomorphically encrypted using the session identifier sid . These encrypted blinding factors are then stored

¹ $\mathbb{Z}_{2^{128}}$ is chosen to make the odds that two sessions will use the same identifier negligible.

Algorithm 1 Functions used by the querier to build a proof.

```

1: // Phase 1: Indicate interest in variables in conj, collect encrypted shares.
2: Function STARTQUERY(conj  $\in 2^{\mathcal{Q}}$ )
3: sid  $\leftarrow_R \mathbb{Z}_{2^{128}}$ ; expected  $\leftarrow 1$ ; encShares  $\leftarrow \emptyset$ ; contacted  $\leftarrow \emptyset$ 
4: for all “ $p_i$  says  $f_i$ ”  $\in conj$  do
5:   r  $\leftarrow_R M$ 
6:   expected  $\leftarrow expected \times r$ 
7:   encShares( $p_i, f_i$ )  $\leftarrow E_{ME \| f_i \| sid}^{p_i}(r)$ 
8:
9: for all “ $p_i$  says  $f_i$ ”  $\in conj$  do
10:  Choose depends  $\subseteq conj$  such that depends  $\in \widehat{acl}_{p_i, ME}(f_i)$ 
11:  resp  $\leftarrow \text{ASK}(p_i, f_i, depends, sid)$ 
12:  if resp = ERROR then
13:    for all ( $p_k, f_k$ )  $\in contacted$  do
14:      RECOVERSHARE( $p_k, f_k, sid, encShares(p_k, f_k)$ )
15:    return FAIL
16:  contacted.add( $p_i, f_i$ )
17:
18:  // Update encrypted secret shares
19:  for all ( $p_j, f_j, E_{ME \| f_j \| sid}^{p_j}(s_j)$ )  $\in resp$  do
20:    encShares( $p_j, f_j$ )  $\leftarrow encShares(p_j, f_j) \otimes E_{ME \| f_j \| sid}^{p_j}(s_j)$ 
21: return (sid, expected, encShares)
22:
23: // Phase 2: Send out decryption requests for encrypted shares
24: Function ENDQUERY(sid  $\in \mathbb{Z}_{2^{128}}$ , expected  $\in M$ ,
    encShares  $\in 2^{\mathcal{P} \times \mathcal{F} \times C}$ )
25: share  $\leftarrow 1$ ; failed  $\leftarrow \text{false}$ 
26: for all ( $p_i, f_i, c_i$ )  $\in encShares$  do
27:   r  $\leftarrow \text{RECOVERSHARE}(p_i, f_i, sid, c_i)$ 
28:   if r = ERROR then
29:     failed  $\leftarrow \text{true}$ 
30:   else
31:     share  $\leftarrow share \times r$ 
32:
33: // Return value
34: if failed then
35:   return FAIL
36: else
37:   return (share = expected)

```

in the *encShares* table, which is used to keep track of the secret shares collected during phase one of the protocol. For each quoted fact “ p_i says f_i ” $\in conj$, the querier then uses \widehat{acl}_{p_i, p_0} to determine the subset of $conj$ representing the release policy for “ p_i says f_i ”; in Algorithm 1, this subset is referred to as *depends*. The remote procedure call (RPC) stub ASK is then used to inquire about the status of f_i at p_i . All RPC communications are carried out over private and authenticated channels.

Principal p_0 's call to $\text{ASK}(p_i, f_i, depends, sid)$ then invokes the function $\text{ASKRESPONSE}(p_0, f_i, depends, sid)$ at the principal p_i (see Algorithm 2). If p_0 has already asked about the status of fact f during session sid or if $depends \notin \widehat{acl}_{p_i, p_0}(f_i)$, this function raises an error. If the request is legitimate, p_i generates a random secret share for each “ p_j says f_j ” $\in depends$. Each share is homomorphically encrypted to its corresponding p_j using the key whose identifier is generated by concatenating the querier's identifier p_0 , the fact f_j , and the session identifier sid . A final secret share is then generated such that the product of all secret shares is 1, and is saved in the *shares* table.

If f_i is a fact in p_i 's extensional knowledge base (i.e., it is a base fact, not the head of some rule $r \in \mathcal{R}$), the current truth value of f_i and its corresponding fact identifier are retrieved from p_i 's knowledge base and stored in the *ids* table; in the second stage of the protocol, p_i will use this information to determine whether the truth value of f_i changed during the protocol. If, on the other hand, f_i is part of p_i 's

intensional knowledge base (i.e., f_i is the head of a derivation rule), p_i executes the first stage of the distributed proof process recursively using the STARTQUERY function. The intermediate results returned by STARTQUERY are stored in the *proofState* table. At this point, p_i returns the list of quoted facts and secret shares to p_0 .

If p_0 's call to ASK($p_i, f_i, depends, sid$) returns the code ERROR, the STARTQUERY function will fail and p_0 will ask all previously contacted principals to decrypt the facts that they had previously disclosed. This ensures that the failed protocol execution appears indistinguishable from a successful protocol execution to all fact providers previously contacted. Otherwise, p_0 makes a note in the *contacted* table stating that principal p_i was successfully contacted regarding the fact f_i . The principal p_0 then uses the homomorphic property of the IBE cryptosystem to combine the secret shares returned by p_i with any corresponding secret shares or blinding factors already stored in the *encShares* table. At this point, each principal contributing a fact to *conj* has been contacted by p_0 and has set up the local state necessary to execute phase two of the protocol.

Phase Two: Consistency check and secret recovery

To start phase two of the protocol, p_0 invokes the ENDQUERY function using the intermediate state returned by the STARTQUERY function. This function uses the RECOVERSHARE RPC stub to ask each fact provider to decrypt the aggregate secret shares that p_0 has accrued. This invokes the RECOVERYRESPONSE function at each fact provider. If the fact provider has previously decrypted any other results related to this (p_0, sid, f_i) triple, an error is raised. If it is safe to proceed, the fact provider p_i decrypts the ciphertext provided by the querier, combines the result with the locally-stored secret share previously generated, and records the fact that it has now decrypted a secret share for the session described by the triple (p_0, sid, f_i) . If the fact f_i is part of p_i 's extensional knowledge base and either f_i is **false** or the identifier associated with f_i has changed since the ASKRESPONSE function was invoked, then the previously-computed secret share is replaced with a random value. Similarly, if f_i is part of p_i 's intensional knowledge base and either the first stage of the distributed proof process fails or the second stage of the distributed proof process determines that f_i is **false**, then the previously-computed secret share is replaced with a random value. The secret share (or random value) is then returned to the querier.

If the product of all the responses gathered from each fact provider is equal to the product of the blinding factors generated prior to stage one of the protocol, the querier can determine that each fact in the conjunction *conj* was simultaneously true at some point during the protocol execution. If *any* fact provider returns a random value during stage two of the protocol, the value computed by p_0 will be randomized and *conj* will be determined to be **false**.

Note that principals must remember all session identifiers they have previously observed. If loose clock synchronization is available, the amount of state to be maintained can be reduced by using the current time for the session identifier. If Δ is the largest expected difference between any two principals' clocks, a principal can remove from the *shares* table any sessions that are more than Δ old, and simultaneously forbid any queries that use a session identifier that is more than Δ time ago. Similarly, a principal that has lost

Algorithm 2 Functions used by fact providers to (conditionally) contribute information regarding locally-maintained facts to a proof being constructed by another principal.

```

1: // This function is used by principals to respond to queries
2: // issued using the ASK function.
3: Function ASKRESPONSE( $p_0 \in \mathcal{P}, f \in \mathcal{F}, depends \in 2^{\mathcal{Q}}, sid \in \mathbb{Z}_{2^{128}}$ )
4: // Make sure that this is a fresh request and  $p_0$  is
5: // (conditionally) authorized to see  $f$ 
6: if ( $shares(p_0, sid, f) \neq \perp$ )  $\vee$  ( $depends \notin \widehat{acl}_{ME, p_0}(f)$ ) then
7:   return ERROR
8:
9:  $resp \leftarrow \emptyset; product \leftarrow 1$ 
10: for all " $p_j$  says  $f_j$ "  $\in depends$  do
11:    $s_j \leftarrow_R M$ 
12:    $product \leftarrow product \times s_j$ 
13:    $resp.add(p_j, f_j, E_{p_0 || f_j || sid}^{p_j}(s_j))$ 
14: Choose  $s$  such that  $s \times product = 1$ 
15:  $shares(p_0, sid, f) \leftarrow s$ 
16:
17: if  $KB.managesFact(f)$  then
18:    $(status, id) \leftarrow KB.lookup(f)$ 
19:    $ids(p_0, sid, f) \leftarrow id$ 
20: else
21:   Use derivation rules and restricted ACLs to find a
   conjunction, conj, of all facts and dependencies needed to
   derive  $f$ 
22:    $proofState(p_0, f, sid) \leftarrow STARTQUERY(conj)$ 
23: return  $resp$ 
24:
25:
26: // This function is used by principals to respond to queries
27: // issued using the RECOVERSHARE function.
28: Function RECOVERYRESPONSE( $p_0 \in \mathcal{P}, f \in \mathcal{F}, sid \in \mathbb{Z}_{2^{128}}, c \in \mathcal{C}$ )
29: // Make sure that this is a fresh request
30: if ( $shares(p_0, sid, f) = \perp$ )  $\vee$  ( $decrypted.contains(p_0, sid, f)$ )
   then
31:   return ERROR
32:
33:  $d \leftarrow D_{p_0 || f || sid}^{ME}(c)$ 
34:  $s \leftarrow (d \times shares(p_0, sid, f)); decrypted.add(p_0, sid, f)$ 
35: if  $KB.managesFact(f)$  then
36:    $(status, id) \leftarrow KB.lookup(f)$ 
37:   if  $\neg status \vee (id \neq ids(p_0, sid, f))$  then
38:      $s \leftarrow_R M$  such that  $s \neq d \times shares(p_0, sid, f)$ 
39: else
40:   if ( $proofState(p_0, f, sid) = FAIL$ )  $\vee$ 
    $\neg ENDQUERY(proofState(p_0, f, sid))$  then
41:      $s \leftarrow_R M$  such that  $s \neq d \times shares(p_0, sid, f)$ 
42: return  $s$ 

```

its previous state can ensure security by refusing to answer queries for the first Δ period of time after it is back up.

3.4 An Example

We now revisit the example depicted in Figure 1 of Section 2.1 to explain in more detail an example execution of the protocol described above. In the interest of space and clarity of presentation, we consider only a subset of this scenario in which the media controller wishes to prove "*bob says request(projector23)*". After using its access to $\widehat{acl}_{bob, mc}$ and $\widehat{acl}_{is, mc}$, the media controller learns that it must actually prove the conjunction of quoted facts "*bob says request(projector23)*" \wedge "*is says owns(mc, projector23)*".

Figure 2 depicts the execution of phase one of the protocol from Section 3.3, using f_{bob} and f_{is} as abbreviations for *request(projector23)* and *owns(mc, projector23)*, respectively. During this stage of the protocol, the media controller first generates two blinding factors b_{bob} and b_{is} . The blinding factor b_{bob} (resp. b_{is}) is then stored in the *encShares* table

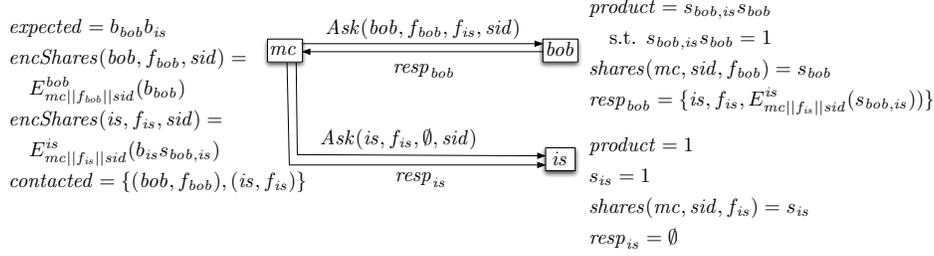


Figure 2: A sample execution of phase one of the protocol.

after being encrypted to Bob (resp. the inventory server) using a key bound to the media controller’s identifier, the fact $request(projector23)$ (resp. $owns(mc, projector23)$), and the session identifier sid . The media controller then uses the ASK function to notify Bob and the inventory server that it is interested in the status of certain facts stored in their respective knowledge bases.

Since Bob’s disclosure of the fact $request(projector23)$ depends on the status of the fact $owns(mc, projector23)$ in the knowledge base of the inventory server, Bob generates a secret share $s_{bob, is}$ that is then encrypted to the inventory server using a key derived from the media controller’s identifier, the fact $owns(mc, projector23)$, and the session identifier sid . Bob then generates a secret share s_{bob} such that $s_{bob} \times s_{bob, is} = 1$. The secret share s_{bob} is stored locally in the $shares$ table, while the encrypted share for the inventory server is disclosed to the media controller. This encrypted secret share is then homomorphically combined with the inventory server’s encrypted blinding factor that is stored in the $encShares$ table. Since the inventory server’s disclosure of the fact $owns(mc, projector23)$ has no dependencies, it stores the value 1 in its $shares$ table and discloses an empty set of dependencies to the media controller. At this point, the media controller updates its $contacted$ list and phase one of the protocol is complete.

During stage two of the protocol, the media controller asks Bob to decrypt the value $E_{mc||f_{bob}||sid}^{bob}(b_{bob})$ and asks the inventory server to decrypt the value $E_{mc||f_{is}||sid}^{is}(b_{is}s_{bob, is})$. Both principals decrypt the requested values and combine them with the secret shares stored in their $shares$ tables. The values $b_{bob}s_{bob}$ and $b_{is}s_{bob, is}$ are then returned to media controller by Bob and the inventory server, respectively. Since $b_{bob}s_{bob}b_{is}s_{bob, is} = (b_{bob}b_{is})(s_{bob, is}s_{bob}) = b_{bob}b_{is}$, which is the expected value generated by the media controller prior to the start of the protocol, the media controller can conclude that the conjunction “ bob says $request(projector23)$ ” \wedge “ is says $owns(mc, projector23)$ ” is **true**.

4. DISCUSSION

In this section, we formally discuss the properties of the confidentiality-preserving distributed proof protocol that was presented in Section 3. We begin by discussing the soundness and completeness of this protocol, and conclude by examining limitations of confidentiality-preserving distributed proof in scenarios involving multiple concurrent runs.

4.1 Protocol Characteristics

We begin our treatment of the formal properties of our confidentiality-preserving distributed proof protocol by mak-

ing the following claim regarding the correctness of this protocol. A full proof of this claim can be found in the extended version of this paper [14].

THEOREM 2 (SOUNDNESS). *Given a conjunction of quoted facts $conj \in 2^{\mathcal{Q}}$, if $ENDQUERY(STARTQUERY(conj))$ returns **true**, then there exists some time t during the execution of the protocol at which every quoted fact “ p_j says f_j ” $\in conj$ was simultaneously **true**.*

Although above theorem asserts the soundness of this proof construction protocol, it is not complete. That is, if $ENDQUERY(STARTQUERY(conj))$ returns **false**, there could still have existed some time t during the execution of the protocol at which each quoted fact in $conj$ was simultaneously **true**. In the general case, designing a sound and complete distributed proof system requires synchronized clocks, which is contrary to the definition of an asynchronous system. Due to space limitations, we do not prove this claim, but instead refer the reader to [16] where a similar claim was proven for trust negotiation systems (e.g., [28]), which are a specific type distributed proof system. However, we can make a more limited claim regarding the soundness and completeness of our proof construction protocol, the proof of which can be found in the extended version of this paper [14].

THEOREM 3 (STABLE PROOF BEHAVIOR). *Assume that a querier p_0 wants to prove the conjunction of facts $conj$, $acl_{p_i, p_0}(f_i) \subseteq conj$ for each “ p_i says f_i ” $\in conj$, and there exists a well-formed proof graph for the conjunction of facts $conj$ whose leaves are stable between times t_1 and t_2 . In this case, $ENDQUERY(STARTQUERY(conj)) \leftrightarrow conj$ if $STARTQUERY$ is invoked after time t_1 and $ENDQUERY$ returns before time t_2 .*

While completeness cannot be established in the general case, the above theorem shows that the protocol presented in this paper is both sound and complete under realistic assumptions. This is in contrast to a trivial protocol that always returns false. Such a protocol is sound, but cannot be shown to be complete under *any* realistic assumptions.

Finally, we note that our protocol preserves the necessary constraints from the ideal functionality. We do this through a series of theorems, the proofs of which can be found in the extended version of this paper [14].

THEOREM 4 (QUERY PRIVACY). *During the construction of a distributed proof, a malicious subset of providers p_1, \dots, p_m will learn which local facts each p_i provides, but will learn nothing about the other facts comprising the conjunction being proved.*

THEOREM 5 (QUERY VALIDITY). *During the construction of a distributed proof, given a subset of malicious fact providers, p_1, \dots, p_m , if one of the honest fact providers, p_i , provides a false fact, the conjunction received by the querier p_0 is false.*

Note that a malicious fact provider can always falsify a conjunction by simply performing the protocol as if its fact is false. However, the same is true inside the TTP model.

THEOREM 6 (LIMITED DISCLOSURE). *During the construction of a distributed proof, a malicious querier p_0 , colluding with set of fact providers, p_1, \dots, p_m , learns the same amount of information as it would by interacting with the TTP functionality, as long as the ACLs of honest fact providers form a strongly-connected component.*

Here, we consider ACLs to form a graph, where if p_i 's ACL includes the condition p_j says f_j , then there is a directed arrow from p_i to p_j . Note that our protocol provides a form of limited disclosure when the ACLs do not form a strongly-connected component, but we will reserve the discussion of the exact security property to Section 4.3.

The above theorems guarantee that, as long as the strongly-connected condition is satisfied, the protocol presented in Section 3 provides principals with the same guarantees afforded by the ideal functionality discussed in Section 2.3. Namely, each principal's ACLs remain private, the truth value of a conjunction is revealed if and only if each fact comprising the conjunction is *simultaneously* valid, and fact providers learn neither what conjunctions are being proved nor the status of other facts in the system.

4.2 Concurrency and Multiple Runs

It is important to note that the ideal functionality described by Definition 2 is concerned with the construction of a *single* proof. As a result, the semantics of a confidentiality-preserving distributed proof treat only the confidentiality of facts within a single run of any given proof protocol. This implies that although one failed proof does not leak any information regarding the status of the facts making up the conjunction being proved, it may be possible to use information collected during *multiple* proofs to infer the status of sensitive facts. For example, consider the case in Figure 3, in which the disclosure of the quoted fact “ p_1 says f_1 ” depends upon the quoted fact “ p_2 says f_2 ”. If a proof in the TTP model fails, the querier will learn only that the conjunction “ p_1 says f_1 ” \wedge “ p_2 says f_2 ” is false. However, if the same principal then attempts to prove “ p_2 says f_2 ” and succeeds, he can conclude that “ p_1 says f_1 ” was likely false during his previously attempted proof.

Notice that in this case, the ACLs are *weakly connected*, and thus Theorem 6 says nothing about the security of the query. However, as we will show in Section 4.3, even when ACLs are not strongly connected, a single run of our protocol reveals only as much information as several *non-overlapping* queries of the TTP (i.e., queries that do not share facts). Since combining multiple non-overlapping queries into one does not offer a querier any advantage, for the rest of this sub-section we will assume that the querier is only able to learn the truth value of a single conjunction during each run of the protocol.

Note that the above types of leakage scenario are similar to attacks in which individual records in an “anonymized” data

source are identified by combining access to the anonymized data with external knowledge (e.g., see [13] and [23]). Recently, a number of syntactic [17, 21, 26, 30] and semantic [9] approaches have been proposed to limit these types of inferences. However, the inference problem in distributed proof systems differs from this existing work in two important ways: (i) fact status values change over time and (ii) *principals* in the system control the functional dependencies used to infer data (i.e., ACL constraints). These two important differences from the more traditional data anonymization domain can be leveraged to develop realistic inference mitigation approaches.

Strict Reference Monitors. The only way to strictly prevent inferences made using the observations from multiple distributed proofs is to partition the set of allowable proofs based upon previous history. This can be accomplished by keeping track of previous queries. Each principal p_i can locally enforce the constraint that no other principal p_j is allowed query the status of fact $f_i \in KB_i$ more than once. This prevents inference, but it does so at the cost of making the proof system less useful over time. Note that if we want to guard against colluding principals, p_i must allow only a single query for f_i from *any* principal, further limiting the utility of the system.

Epoch-Based Reference Monitors. To increase system utility over the strict reference monitor model, we can leverage the observation that in many proof systems, the status of facts is expected to fluctuate over time (e.g., consider facts representing user and device locations, room occupancy predicates, or other physical phenomena). As a result, reference monitors can actually enforce the above types of constraints on a sliding-window basis. The length of such a window allows the uncertainty introduced by the transient nature of fact status values to occlude inferences to whatever level is deemed necessary.

Policy Design. In certain circumstances, principals in the system may be able to prevent inferences made during the construction of multiple proofs by intelligent policy engineering decisions. Specifically, if the dependency graph whose nodes are quoted facts in the conjunction to be proved and whose edges are entailed by the ACL entries protecting these facts is strongly-connected, then requesting any fact in the conjunction requires the disclosure of all other facts in the conjunction. For example, Figure 4 applies this principle to prevent the inference made possible in Figure 3. Obviously, this method prevents the inference problem, but can only be applied in the event that facts in a given conjunction are never used outside of the context of that conjunction.

The choice of inference control mechanism to apply in any given situation is highly dependent on the environment within which a given proof system is to be deployed. Furthermore, the three approaches to inference control described above afford a range of trade-offs that can be taken into consideration by system architects at deployment time. As a result, we argue that the notion of confidentiality-preserving distributed proof described in this paper is in fact quite reasonable. It provides much more protection for confidential facts than existing proof systems, and can be further tuned to adjust the privacy/utility trade-off as needed for any par-



Figure 3: An example ACL dependency graph that permits inference.

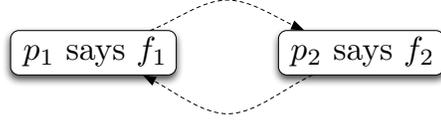


Figure 4: An example ACL dependency graph that prevents inference.

ticular deployment environment.

4.3 Multiple-Run Restrictions

Theorem 6 makes no statement about the security of the conjunction in the case where the ACLs are not strongly connected. Ideally, we would like to be able to prove the same statement in this case; however, the statement turns out to be not true.

First, it is easy to see that if the ACLs of facts in the conjunction are not weakly connected, the querier learns the truth of each weakly connected component separately. This is because there are no encrypted values that are sent from principals in one component to principals in another. Hence, the execution of the protocol is equivalent to multiple parallel executions of the protocol on each weakly connected component. In fact, the querier could use different session IDs for each component without being detected. Second, even within a weakly connected component, the querier has the ability to learn conjunctions other than the one that is ostensibly being queried. For example, if the conjunction in question is “ p_1 says f_1 ” \wedge “ p_2 says f_2 ”, and $acl_2(f_2) = \{p_0\}$, then p_0 can execute the protocol in such a way that it learns “ p_2 says f_2 ”. Both of these problems are a direct consequence of hiding the full conjunction that is being proven from the fact providers: by design, p_2 should not be able to tell whether p_0 is trying to prove “ p_1 says f_1 ” \wedge “ p_2 says f_2 ” or just “ p_2 says f_2 ”.

Fortunately, we can still state a useful security property that is provided by our protocol: the protocol allows a querier to learn as much as could be learned by multiple *non-overlapping* interactions with the TTP.

THEOREM 7. *During a single construction of a distributed proof, given a conjunction C , a malicious querier p_0 can learn the truth of a set of conjunctions $C_1, \dots, C_n \subset C$, under the condition that the ACLs of all the facts comprising each conjunction C_i are weakly connected, $C_i \cap C_j = \emptyset$ for $i \neq j$, and $checkAcls(p_0, C_i) = true$. No other conjunctions will be revealed by the construction of a distributed proof.*

This means that in the above example, p_0 can learn the truth of “ p_1 says f_1 ” \wedge “ p_2 says f_2 ” or “ p_2 says f_2 ”, but not both. Therefore, the strict or epochal reference monitor policies suggested in Section 4.2 will still work, since any time a principal learns a conjunction C , whether it be the one it ostensibly queried or not, all the principals who have facts in C will be notified and will forbid any further conjunctions that overlap with C . An ad-hoc proof of this theorem is presented in the extended version of this paper [14].

5. RELATED WORK

Although earlier distributed proof construction systems [1, 2, 7, 22, 29] are clearly related to the notion of confidentiality-

preserving distributed proof, the work presented in this paper differs in two important ways. First, most existing systems have no mechanism for protecting the release of sensitive facts. Other proof systems allow fact disclosures to be protected by identity-based ACLs [22] or more complex release policies that must be *centrally verified* by the fact provider [29]. Our proof system allows fact providers to be assured that complex conditions protecting the disclosure of their facts will be enforced *without* requiring central verification. This allows fact providers to write release policies that reference quoted facts in the system to which they themselves do not have access. Second, existing proof systems do not consider the issue of simultaneous truth of facts, and simply assume that the proof construction protocol samples a consistent system state. In [15], the authors show how a distributed proof system [22] can be extended to ensure the simultaneous truth of facts sampled in a given proof. The proof system presented in this paper considers this goal from the outset.

In trust negotiation approaches to authorization (e.g., see [3, 4, 20, 27, 28]), principals can protect access to sensitive credentials using complex release policies similar to those described in this paper. However, once again, evidence attesting to the fact that these policies are satisfied must be collected and centrally verified by the entity disclosing the credential. This has two implications: evidence of partial policy satisfaction can be learned by principals in the system and principals can only write release policies that they are authorized to verify. The proof system presented in this paper addresses both of these problems by using a two-phase approach to *distributed* policy enforcement, which has the benefit of not leaking any partial state.

Brands’s digital credential scheme [6] allows a credential owner to selectively reveal information encoded in the attributes of that credential. The credential owner can demonstrate that confidential attributes in the credential satisfy a linear relation or some expression in propositional logic without disclosing additional information about the individual attributes. This is similar to the guarantees afforded by our system, but it does not address the case where confidential attributes are maintained by multiple parties in a decentralized environment. The oblivious commitment-based envelope (OCBE) construction proposed by Li and Li [19, 18] provides similar guarantees, again, at a per-credential granularity.

Our problem can be viewed as a special case of the secure multi-party computation (SMC) problem. In this class of problems, some collection of principals wishes to calculate the result of a public function of n private inputs. General solutions exist for this class of problems [11], but can be quite inefficient for large problem sizes. These solutions also assume that all principals involved in a computation are

aware of one another, which is something that our protocol seeks to avoid. Prabhakaran and Rosulek propose a specific SMC algorithm for computing aggregated OR that is similar to our construction [25], but in their construction the fact providers communicate directly to each other.

6. CONCLUSIONS

In this paper, we develop the notion of a confidentiality-preserving distributed proof system, in which the disclosure of a sensitive fact f can be made contingent upon facts f'_1, \dots, f'_n managed by other principals in the system. In this type of proof system, the querier learns only the truth value of the conjunction $f \wedge f'_1 \wedge \dots \wedge f'_n$, which limits the knowledge leaked during the proof construction process. This enables *controlled* usage of sensitive facts, thereby increasing the utility of the proof system while minimizing the effects of this process on user privacy. We first formalized this notion of confidentiality-preserving distributed proof using a trusted third party (TTP) model of computation, and then developed a distributed algorithm for constructing confidentiality-preserving distributed proofs in asynchronous distributed systems.

We formally proved that our confidentiality-preserving distributed proof algorithm is sound, and that it reveals no more information than can be learned by interacting with the ideal functionality embodied by the TTP model. We then examined potential avenues of inference made possible in the ideal model—and thus in our proof construction algorithm—through multiple or concurrent proof constructions. Although these information leaks are worrisome, we showed that they can be limited based on local decisions made by fact providers.

Acknowledgements

We would like to thank Manoj Prabhakaran and Mike Rosulek for their discussions about the cryptographic primitives used in our protocol and the security proofs, and the anonymous reviewers for their helpful suggestions. This research was supported in part by the National Science Foundation under award CNS-0716421.

7. REFERENCES

- [1] A. W. Appel and E. W. Felten. Proof-carrying authentication. In *Proceedings of the Sixth ACM Conference on Computer and Communications Security*, Nov. 1999.
- [2] L. Bauer, S. Garriss, and M. K. Reiter. Distributed proving in access-control systems. In *Proceedings of the 2005 IEEE Symposium on Security and Privacy*, pages 81–95, 2005.
- [3] E. Bertino, E. Ferrari, and A. C. Squicciarini. Trust-X: A peer-to-peer framework for trust establishment. *IEEE Transactions on Knowledge and Data Engineering*, 16(7):827–842, July 2004.
- [4] P. Bonatti and P. Samarati. Regulating service access and information release on the web. In *Proceedings of the Seventh ACM Conference on Computer and Communications Security*, pages 134–143, 2000.
- [5] D. Boneh and M. Franklin. Identity based encryption from the Weil pairing. *SIAM Journal of Computing*, 32(3):586–615, 2003.
- [6] S. A. Brands. *Rethinking Public Key Infrastructure and Digital Certificates*. MIT Press, Cambridge, MA, USA, 2000.
- [7] J. DeTreville. Binder, a logic-based security language. In *Proceedings of the 2002 IEEE Symposium on Security and Privacy*, page 105, 2002.
- [8] D. Dolev, C. Dwork, and M. Naor. Non-malleable cryptography. In *STOC '91: Proceedings of the twenty-third annual ACM symposium on Theory of computing*, pages 542–552, New York, NY, USA, 1991. ACM.
- [9] C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Proceedings of the Third Theory of Cryptography Conference*, pages 265–284, Mar. 2006.
- [10] T. Elgamal. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE Transactions on Information Theory*, 31(4):469–472, 1985.
- [11] O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game. In *Proceedings of the 19th annual ACM Conference on Theory of Computing*, pages 218–229, New York, NY, USA, 1987. ACM Press.
- [12] A. Ivan and Y. Dodis. Proxy cryptography revisited. In *Proceedings of the 10th Annual Network and Distributed System Security Symposium (NDSS 2003)*, Feb. 2003.
- [13] T. Kohno, A. Broido, and K. Claffy. Remote physical device fingerprinting. *IEEE Transactions on Dependable and Secure Computing*, 2(2):93–108, 2005.
- [14] A. J. Lee, K. Minami, and N. Borisov. Confidentiality-preserving distributed proofs of conjunctive queries (extended version). Department of Computer Science Technical Report TR-08-161, University of Pittsburgh, Dec. 2008.
- [15] A. J. Lee, K. Minami, and M. Winslett. Lightweight consistency enforcement schemes for distributed proofs with hidden subtrees. In *Proceedings of the 12th ACM Symposium on Access Control Models and Technologies*, pages 101–110, 2007.
- [16] A. J. Lee and M. Winslett. Enforcing safety and consistency constraints in policy-based authorization systems. *ACM Transactions on Information and System Security*, to appear.
- [17] K. LeFevre, D. J. DeWitt, and R. Ramakrishnan. Mondrian multidimensional k-anonymity. In *Proceedings of the 22nd International Conference on Data Engineering (ICDE)*, Apr. 2006.
- [18] J. Li and N. Li. A construction for general and efficient oblivious commitment based envelope protocols. In *Proceedings of 8th International Conference on Information and Communications Security (ICICS)*, pages 122–138, Dec. 2006.
- [19] J. Li and N. Li. OACerts: oblivious attribute certificates. *IEEE Transactions on Dependable and Secure Computing (TDSC)*, 3(4):340–352, Oct. 2006.
- [20] J. Li, N. Li, and W. H. Winsborough. Automated trust negotiation using cryptographic credentials. *ACM Transactions on Information and System Security*, to appear.
- [21] A. Machanavajjhala, J. Gehrke, D. Kifer, and

- M. Venkatasubramanian. ℓ -diversity: Privacy beyond k-anonymity. In *Proceedings of the 22nd International Conference on Data Engineering (ICDE)*, Apr. 2006.
- [22] K. Minami and D. Kotz. Secure context-sensitive authorization. *Journal of Pervasive and Mobile Computing*, 1(1):123–156, Mar. 2005.
- [23] A. Narayanan and V. Shmatikov. Robust de-anonymization of large sparse datasets (how to break anonymity of the Netflix prize dataset). In *Proceedings of 29th IEEE Symposium on Security and Privacy*, May 2008.
- [24] P. Paillier. Public-key cryptosystems based on composite degree residuosity classes. In *Proceedings of EUROCRYPT—Advances in Cryptology*, 1999.
- [25] M. Prabhakaran and M. Rosulek. Cryptographic complexity of multi-party computation problems: Classifications and separations. *Electronic Colloquium on Computational Complexity (ECCC)*, 15(50), 2008.
- [26] L. Sweeney. k-anonymity: A model for protecting privacy. *International Journal on Uncertainty, Fuzziness and Knowledge-based Systems*, 10(5):557–570, 2002.
- [27] W. H. Winsborough and N. Li. Towards practical automated trust negotiation. In *Proceedings of the Third IEEE International Workshop on Policies for Distributed Systems and Networks*, pages 92–103, June 2002.
- [28] W. H. Winsborough, K. E. Seamons, and V. E. Jones. Automated trust negotiation. In *Proceedings of the DARPA Information Survivability Conference and Exposition*, pages 88–102, Jan. 2000.
- [29] M. Winslett, C. C. Zhang, and P. A. Bonatti. PeerAccess: a logic for distributed authorization. In *Proceedings of the 12th ACM Conference on Computer and Communications Security*, pages 168–179, 2005.
- [30] X. Xiao and Y. Tao. m-invariance: Towards privacy preserving re-publication of dynamic datasets. In *Proceedings of the ACM Conference on Management of Data (SIGMOD)*, pages 689–700, June 2007.
- [31] T. Yu, M. Winslett, and K. E. Seamons. Supporting structured credentials and sensitive policies through interoperable strategies for automated trust negotiation. *ACM Transactions on Information and System Security*, 6(1):1–42, Feb. 2003.