

Automated Trust Negotiation Using Cryptographic Credentials

Jiangtao Li
Dept. of Computer Science
Purdue University
jtli@cs.purdue.edu

Ninghui Li
Dept. of Computer Science
Purdue University
ninghui@cs.purdue.edu

William H. Winsborough
Dept. of Computer Science
Univ. of Texas at San Antonio
wwinsborough@acm.org

ABSTRACT

In automated trust negotiation (ATN), two parties exchange digitally signed credentials that contain attribute information to establish trust and make access control decisions. Because the information in question is often sensitive, credentials are protected according to access control policies. In traditional ATN, credentials are transmitted either in their entirety or not at all. This approach can at times fail unnecessarily, either because a cyclic dependency makes neither negotiator willing to reveal her credential before her opponent, because the opponent must be authorized for all attributes packaged together in a credential to receive any of them, or because it is necessary to fully disclose the attributes, rather than merely proving they satisfy some predicate (such as being over 21 years of age). Recently, several cryptographic credential schemes and associated protocols have been developed to address these and other problems. However, they can be used only as fragments of an ATN process. This paper introduces a framework for ATN in which the diverse credential schemes and protocols can be combined, integrated, and used as needed. A policy language is introduced that enables negotiators to specify authorization requirements that must be met by an opponent to receive various amounts of information about certified attributes and the credentials that contain it. The language also supports the use of uncertified attributes, allowing them to be required as part of policy satisfaction, and to place their (automatic) disclosure under policy control.

Categories and Subject Descriptors

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

General Terms

Security, Design

Keywords

Privacy, Access Control, Automated Trust Negotiation, Digital Credentials

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.

CCS'05, November 7–11, 2005, Alexandria, Virginia, USA.

Copyright 2005 ACM 1-59593-226-7/05/0011 ...\$5.00.

1. INTRODUCTION

In automated trust negotiation (ATN) [16, 28, 29, 30, 31, 32, 33, 34, 37, 38], two parties exchange digitally signed credentials that contain attribute information to establish trust and make access control decisions. In traditional ATN approaches the only way to use a credential is to send it as a whole, thus disclosing all the information in the credential. In other words, a digital credential is viewed as a black-box, and the information in a credential is disclosed in an all-or-nothing fashion. In these approaches sensitive attribute values stored in a credential are protected using access control techniques. There is an access control policy associated with each credential and a credential can be disclosed if its access control policy has been satisfied. Viewing a credential as a black-box severely limits the power of ATN. The following are some of the limitations.

- If there is a cyclic dependency among credentials and their policies, negotiations can fail unnecessarily. For example, in a negotiation between A and B , suppose A has a credential c_1 that can be disclosed only if B has c_2 , and B has c_2 , but can disclose it only if A has c_1 . Using traditional ATN techniques, the negotiation would fail because neither c_1 and c_2 can be disclosed before the other is, even though allowing A and B to exchange *both* c_1 and c_2 would not violate either negotiator's policy.
- Because attribute information is disclosed in an all-or-nothing fashion, each attribute can be disclosed only when the policy governing the credential and its entire contents is satisfied, leading to unnecessary failure. For example, suppose B would allow A to access a resource provided that A is over 21, and A has a digital driver license that includes A 's date of birth (DoB) and address. If A does not want to reveal her address (or her exact DoB) to B , the negotiation would fail, even if A would be willing to prove she is over 21.
- When one negotiator does not want to disclose detailed information about his policy and the other negotiator does not want to disclose too much information about her attributes, a negotiation can fail even though the amount of information that needs to be disclosed by each party is acceptable to both. For example, suppose B is a bank that offers a special-rate loan and A would like to know whether she is eligible for such a loan before she applies. B is willing to reveal that his loan-approval policy uses one's DoB, current salary, and the length of the current employment; however, B considers further details of this policy to be a trade secret that he is unwilling to reveal. A would like to know whether she is eligible for the loan while disclosing as little information about her attributes as possible. In particular, A does not want to

disclose the exact values of her DoB or salary level. Using traditional ATN techniques, this negotiation would fail.

A number of cryptographic credential schemes and associated protocols have been developed to address these and other problems. Oblivious signature based envelope [22], hidden credentials [6, 17], and secret handshakes [2] can be used to address the policy cycle problem. Oblivious Attribute Certificates (OACerts) [19], private credentials [7], and anonymous credentials [8, 9, 10, 26] together with zero-knowledge proof protocols can be used to prove that an attribute satisfies a policy without disclosing any other information about the attribute. Certified input private policy evaluation (CIPPE) [20] enables A and B to determine whether A 's attribute values satisfy B 's policies without revealing additional information about A 's attributes or B 's policies.

While these credential schemes and associated protocols all address some limitations in ATN, they can be used only as fragments of an ATN process. For example, a protocol that can be used to handle cyclic policy dependencies should be invoked only when such a cycle occurs during the negation process. A zero-knowledge proof protocol can be used only when one knows the policy that needs to be satisfied and is willing to disclose the necessary information to satisfy the policy. An ATN framework that harness these powerful cryptographic credentials and protocols has yet to be developed. In this paper, we develop an ATN framework that does exactly that. Our framework has the following salient features.

- The ATN framework supports diverse credentials, including standard digital credentials (such as X.509 certificates [3, 18]) as well as OACerts, hidden credentials, and anonymous credentials.
- In addition to attribute information stored in credentials, the ATN framework supports also attribute information that is not certified. For example, oftentimes one is asked to provide a phone number in an online transaction, and the phone number may not be certified in any certificate. In our framework, uncertified attribute information and certified attribute information are protected in a uniform fashion.
- The ATN framework has a logic-based policy language that we call Attribute-based Trust Negotiation Language (ATNL), which allows one to specify policies that govern the disclosure of partial information about a sensitive attribute. ATNL is based on the RT family of Role-based Trust-management languages [23, 24, 25].
- The ATN framework has a negotiation protocol that enables the various cryptographic protocols to be used to improve the effectiveness of ATN. This protocol is an extension of the Trust-Target Graph (TTG) ATN protocol [31, 32].

The rest of this paper is organized as follows. We discuss related work in Section 2, and then review several credential schemes and associated protocols that can be used in ATN in Section 3. In Section 4, we present the language ATNL. In Section 5 we present our negotiation protocol. We conclude our paper in Section 6.

2. RELATED WORK

Automated trust negotiation was introduced by Winsborough et al. [33], who presented two negotiation strategies: an eager strategy in which negotiators disclose each credential as soon as its access control policy is satisfied, and a “parsimonious” strategy in which negotiators disclose credentials only after exchanging sufficient policy content to ensure that a successful outcome is ensured.

Yu et al. [38] developed a family of strategies called the disclosure tree family such that strategies within the family can interoperate with each other in the sense that negotiators can use different strategies within the same family. Seamons et al. [28] and Yu and Winslett [37] studied the problem of protecting contents of policies as well as credentials. On the aspect of system architecture for trust negotiation, Hess et al. [16] proposed the Trust Negotiation in TLS (TNT) protocol, which is an extension to the SSL/TLS handshake protocol by adding trust negotiation features. Winslett et al. [35] introduced the TrustBuilder architecture for trust negotiation systems. The problem of leaking attribute information was recognized by Winsborough and Li [31], Seamons et al. [29], and Yu and Winslett [36]. Winsborough and Li [30, 31, 32] introduced the notion of acknowledgement policies to protect this information and provided a formal notion of safety against illegal attribute information leakage. Bonatti and Samarati [4] proposed a framework for regulating service access and information release on the web. Their framework supports both certified attributes and uncertified attributes.

Recent work on using cryptographic protocols for ATN includes hidden credentials [6, 15, 17], secret handshakes [2], oblivious signature based envelope [22], oblivious commitment based envelope [19], certified input private policy evaluation [20], and policy-based cryptography [1]. While these protocols are useful tools and building blocks for ATN, they are not general enough to solve arbitrary trust negotiation problems in a systematic way. Credential schemes that can be used in ATN include OACerts [19], private credentials [7], and anonymous credentials [8, 9, 10, 26]. We will summarize the features of these protocols and credential schemes in the next section.

3. OVERVIEW OF CRYPTOGRAPHIC CREDENTIALS AND TOOLS FOR ATN

We now give an overview of six properties that are provided by cryptographic credential schemes and their associated cryptographic tools. These properties can improve the privacy protection and effectiveness of ATN.

1. *Separation of credential disclosure from attribute disclosure:* In several credential systems, including private credentials [7], anonymous credentials [8, 9, 10, 26], and OACerts [19], a user's attribute values are not stored in the clear; instead, they are stored in a committed form in her credentials. When the commitment of an attribute value is stored in a credential, looking at the commitment does not enable one to learn anything about the attribute value. Therefore, a credential holder can disclose her credentials without revealing the attribute values in them. For example, consider a digital driver license certificate from Bureau of Motor Vehicles (BMV) consisting of name, gender, DoB, and address. In trust negotiation, a user can show that her digital driver license is valid, *i.e.*, that she is currently a valid driver, without disclosing any of her name, gender, DoB, and address.
2. *Selective show of attributes:* A credential holder can select which attributes she wants to disclose (and which attribute she does not want to disclose) to the verifier. As each attribute in a credential is in committed form, the credential holder can simply open the commitments of the attributes she wants to reveal. For instance, using the digital driver license, the credential holder can show her name and address to a verifier without disclosing her gender and DoB. Cryptographic properties of the commitment schemes ensure that

the credential holder cannot open a commitment with a value other than the one that has been committed.

3. *Zero-knowledge proof of attributes satisfying a policy:* A credential holder can use zero-knowledge proof protocols [5, 11, 12, 14] to prove that her attributes satisfy some property without revealing the actual attribute values. For example, a credential holder can prove that she is older than 21 by using her digital driver license without revealing any other information about her actual DoB.
4. *Oblivious usage of a credential:* A credential holder can use her credentials in an oblivious way to access resources using Oblivious Signature Based Envelope (OSBE) [22], hidden credentials [17], or secret handshakes [2]. In OSBE, a user sends the contents of her credential (without the signature) to a server. The server verifies that the contents satisfy his requirement, then conducts a joint computation with the user such that in the end the user sees the server's resource if and only if she has the signature on the contents she sent earlier. The hidden credentials and the secret handshakes share the similar concept; however, they assume that the server can guess the contents of the user's credentials; thus the user does not need to send the contents to the server. The oblivious usage of a credential enables a user to obtain a resource from a server without revealing the fact that she has the credential.
5. *Oblivious usage of an attribute:* A credential holder can use her attributes in an oblivious way to access resources using Oblivious Commitment Based Envelop (OCBE) [19]. In OCBE, a credential holder and a server run a protocol such that in the end the credential holder receives the server's resource if and only if the attributes in her credential satisfy the server's policy. The server does not learn anything about the credential holder's attribute values, not even whether the values satisfy the policy or not.
6. *Certified input private policy evaluation (CIPPE):* In CIPPE [20], a credential holder and a server run a protocol in which the credential holder inputs the commitments of her attribute values from her credentials, and the server inputs his private policy function. In the end, both parties learn whether the credential holder satisfies the server's policy, without revealing the attribute values to the server or the private function to the credential holder. For example, suppose that the server's policy is that age must be greater than 25 and the credential holder's age is 30. The credential holder can learn that she satisfies the server's policy without revealing her exact DoB or knowing the threshold in the server's policy.

There are other useful properties achieved in the private credentials [7] and the anonymous credentials [8, 9, 10, 26], such as multi-show unlinkable property, anonymous property, etc. Some of these properties require anonymous communication channels to be useful. In this paper, we focus on the six properties described above, because we believe they are most related to trust negotiation. Our goal is to integrate them into a coherent trust negotiation framework.

Note that we do not assume each negotiating participant supports all six properties. For instance, if one participant uses an anonymous credential system and supports properties 1–3, and the other participant supports properties 1–6, then they can use properties 1–3 when they negotiate trust. We present an ATN framework that can take advantage of these properties when they are available, but that does not require them.

4. THE LANGUAGE OF CREDENTIALS AND POLICIES

In this section, we present the Attribute-based Trust Negotiation Language (ATNL), a formal language for specifying credentials and policies. ATNL is based on *RT*, a family of Role-based Trust-management languages introduced in [23, 24, 25]. We first give an example trust negotiation scenario in ATNL, then describe the syntax of ATNL in detail in Section 4.2.

4.1 An Example

In this example, the two negotiators are BookSt (a bookstore) and Alice. We give the credentials and policies belonging to BookSt first, then give those for Alice, and then describe a negotiation process between BookSt and Alice.

BookSt's credentials and policies are given in Figure 1. BookSt has a credential ($\ell 1$) issued by the Small Business Administration (SBA) asserting that BookSt has a valid business license. BookSt is certified in ($\ell 2$) by the Better Business Bureau (BBB) to have a good security process.

BookSt offers a special discount to anyone who satisfies the policy ($m1$), which means that the requester should be certified by StateU to be a student majoring in computer science, under 21, and willing to provide a phone number. Since the discount is a resource, the head of this policy, $\text{BookSt.discount}(\text{phoneNum} = x_3)$, defines a part of the application interface provided by the ATN system using this policy; the parameter phoneNum is made available to the application through this interface. That is, the application will issue a query to determine whether the requester satisfies $\text{BookSt.discount}(\text{phoneNum} = x_3)$, and if it succeeds, the variable x_3 will be instantiated to the phone number of the requester. The body of policy ($m1$) (*i.e.*, the part to the right of \leftarrow) consists of the following two parts.

Part 1: $\text{StateU.student}(\text{program} = x_1) \cap \text{BookSt.DoB}(\text{val} = x_2) \cap \text{Any.phoneNum}(\text{val} \Rightarrow x_3)$

Part 2: $((x_1 = \text{'cs'}) \wedge (x_2 > \text{'01/01/1984'}))$

Part 1 describes the role requirement of the policy and consists of the intersection of 3 roles. To satisfy the role $\text{StateU.student}(\text{program} = x_1)$, one must provide a credential (or a credential chain) showing that one is certified by StateU to be a student; $\text{program} = x_1$ means that the value of the program field is required to satisfy additional constraints. In $\text{Any.phoneNum}(\text{val} \Rightarrow x_3)$, the keyword *Any* means that the phone number does not need to be certified by any party and the symbol \Rightarrow means that the phone number must be provided (enabling it to be returned to the application). Part 2 describes the constraints on specific field values.

BookSt's policies ($m2$) and ($m3$) mean that BookSt considers both a driver license from BMV and a passport issued by the government (Gov) to be valid documents for DoB.

Alice's credentials and policies are given in Figure 2. Alice holds three credentials. Credential ($n1$) is issued by StateU and delegates to College of Science (CoS) the authority to certify students. Credential ($n2$) is Alice's student certificate issued by CoS. Credentials ($n1, n2$) prove that Alice is a valid student from StateU. Credential ($n3$) is her digital driver license issued by BMV. For simplicity, we assume that the digital driver license contains only name and DoB. Among her credentials, Alice considers her student certificate to be sensitive, and provides it only to those who have a valid business license from SBA ($p1$). Alice does not protect the content of her driver license, except for its DoB field. She considers her date of birth and phone number to be sensitive information, thus she reveals them only to organizations whose security practices are adequate to provide reasonable privacy ($p2, p3$). For this, we assume

<i>BookSt's credentials:</i>	
$\ell 1$:	SBA.businessLicense \leftarrow BookSt
$\ell 2$:	BBB.goodSecProcess \leftarrow BookSt
<i>BookSt's policies:</i>	
$m1$:	BookSt.discount(phoneNum = x_3) \leftarrow StateU.student(program = x_1) \cap BookSt.DoB(val = x_2) \cap Any.phoneNum(val \Rightarrow x_3) ; $((x_1 = \text{'cs'}) \wedge (x_2 > \text{'01/01/1984'}))$
$m2$:	BookSt.DoB(val = x) \leftarrow BMV.driverLicense(DoB = x)
$m3$:	BookSt.DoB(val = x) \leftarrow Gov.passport(DoB = x)

Figure 1: The credentials and policies of BookSt

<i>Alice's credentials:</i>	
$n1$:	StateU.student \leftarrow CoS.student
$n2$:	CoS.student(program = 'cs', level = 'sophomore') \leftarrow Alice
$n3$:	BMV.driverLicense(name = commit('Alice'), DoB = commit('03/07/1986')) \leftarrow Alice
<i>Alice's attribute declarations:</i>	
$o1$:	phoneNum = '(123)456-7890' :: sensitive
$o2$:	DoB = '03/07/1986' :: BMV.driverLicense(DoB) :: sensitive
$o3$:	program = 'cs' :: CoS.student(program) :: non-sensitive
$o4$:	level = 'sophomore' :: CoS.student(level) :: non-sensitive
<i>Alice's policies:</i>	
$p1$:	disclose(ac, CoS.student) \leftarrow SBA.businessLicense
$p2$:	disclose(full, DoB) \leftarrow BBB.goodSecProcess
$p3$:	disclose(full, phoneNum) \leftarrow BBB.goodSecProcess
$p4$:	disclose(range, DoB, year) \leftarrow true

Figure 2: The credentials and policies possessed by Alice

that BBB provides a security process auditing service. Further, Alice is willing to reveal to everyone her year of birth ($p4$).

A negotiation between BookSt and Alice When Alice requests a discount sale from BookSt, BookSt responds with his discount policy ($m1$). Alice first discloses her driver license ($n3$), which is assumed to be an OACert, to BookSt without revealing her DoB. To protect her phone number and her student certificate, Alice wants BookSt to show a business license issued by SBA and a good security process certificate issued by BBB. After BookSt shows the corresponding certificates ($\ell 1$, $\ell 2$), Alice reveals her student certificate chain ($n1$, $n2$) and phone number ($o1$). As Alice is allowed by her policy $p4$ to reveal her year of birth to everyone, she uses a zero-knowledge proof protocol to prove to BookSt that her DoB in her driver license is between '1/1/1986' and '12/31/1986'. BookSt now knows that Alice is younger than 21, thus satisfies his discount policy. During the above interactions, Alice proves that she is entitled to obtain the discount.

The above negotiation process uses the first three properties described in Section 3.

4.2 The Syntax

Figure 3 gives the syntax of ATNL in Backus Naur Form (BNF). In the following, we explain the syntax. The numbers in the text below correspond to the numbers of definitions in Figure 3.

Each negotiation party has a *policy base* (3) that contains all information that may be used in trust negotiation. A party's policy base consists of three parts: *credentials*, *attribute declarations*, and *policy statements*. In the following, we discuss each of the three parts in detail.

4.2.1 Credentials and Roles

Two central concepts that ATNL takes from *RT* [24, 25] are principals and roles. A principal is identified with an individual or agent, and may be represented by a public key. In this sense, principals can issue credentials and make requests. A *role* designates a set of principals who are members of this role. Each principal has its own localized name space for roles in which it has sole authority to define roles. A *role* (7) takes the form of a principal followed by a role term, separated by a dot. The simplest kind of a role term consists of just a role name. As roles are parameterized, a role term may also contain fields, which will be explained later. We use A , B , D , S , and V , sometimes with subscripts, to denote principals. We use R , often with subscripts, to denote role terms. A role $A.R$ can be read as A 's R role. Only A has the authority to define the members of the role $A.R$, and A does so by issuing role-definition statements.

In ATNL, a credential can be either a membership credential or a delegation credential. A *membership credential* (5) takes the form of $A.R \leftarrow D$, where A and D are (possibly the same) principals. This means that A defines D to be a member of A 's role R . A *delegation credential* (6) takes the form of $A.R \leftarrow B.R_1$, where A and B are (possibly the same) principals, and R and R_1 are role terms. In this statement, A defines its R role to include all members of B 's R_1 role.

For example, BookSt's credential ($\ell 1$) in Figure 1 is a membership credential. It means SBA issued a business license certificate for BookSt. Alice's credential ($n1$) in Figure 2 is a delegation credential. It says that StateU delegates its authority over identifying students to CoS. Alice's credential ($n2$) in Figure 2 means that CoS asserts that Alice is a sophomore student in StateU majoring

$\langle \text{list of } X \rangle ::= \langle X \rangle \mid \langle X \rangle \text{ “,” } \langle \text{list of } X \rangle$	(1)
$\langle \text{set of } X \rangle ::= \epsilon \mid \langle X \rangle \langle \text{set of } X \rangle$	(2)
$\langle \text{policy-base} \rangle ::= \langle \text{set of credential} \rangle \langle \text{set of attr-decl} \rangle \langle \text{set of policy-stmt} \rangle$	(3)
$\langle \text{credential} \rangle ::= \langle \text{member-cred} \rangle \mid \langle \text{delegation-cred} \rangle$	(4)
$\langle \text{member-cred} \rangle ::= \langle \text{role} \rangle \text{ “—” } \langle \text{prin} \rangle$	(5)
$\langle \text{delegation-cred} \rangle ::= \langle \text{role} \rangle \text{ “—” } \langle \text{role} \rangle$	(6)
$\langle \text{role} \rangle ::= \langle \text{prin} \rangle \text{ “.” } \langle \text{role-term} \rangle$	(7)
$\langle \text{role-term} \rangle ::= \langle \text{role-name} \rangle \mid \langle \text{role-name} \rangle \text{ “(” } \langle \text{list of field} \rangle \text{ “)”}$	(8)
$\langle \text{field} \rangle ::= \langle \text{field-name} \rangle \text{ “=” } (\langle \text{var} \rangle \mid \langle \text{constant} \rangle \mid \langle \text{commitment} \rangle)$	(9)
$\langle \text{attr-decl} \rangle ::= \langle \text{attr-name} \rangle \text{ “=” } \langle \text{constant} \rangle \text{ “::” } [\langle \text{list of attr-ref} \rangle]$	(10)
$\langle \text{attr-ref} \rangle ::= \langle \text{prin} \rangle \text{ “.” } \langle \text{role-name} \rangle \text{ “(” } \langle \text{field-name} \rangle \text{ “)”}$	(11)
$\langle \text{policy-stmt} \rangle ::= \langle \text{policy-head} \rangle \text{ “—” } \langle \text{policy-body} \rangle$	(12)
$\langle \text{policy-body} \rangle ::= \langle \text{p-role-req} \rangle [\text{ “,” } \langle \text{p-constraint} \rangle] \mid \text{true}$	(13)
$\langle \text{p-role-req} \rangle ::= [\langle \text{pre-cond} \rangle \text{ “!” }] \langle \text{conj-of-p-roles} \rangle$	(14)
$\langle \text{p-constraint} \rangle ::= [\langle \text{pre-cond} \rangle \text{ “!” }] \langle \text{constraint} \rangle$	(15)
$\langle \text{pre-cond} \rangle ::= \langle \text{role} \rangle \mid \text{“false”}$	(16)
$\langle \text{conj-of-p-roles} \rangle ::= \langle \text{p-role} \rangle \mid \langle \text{p-role} \rangle \text{ “∩” } \langle \text{conj-of-p-roles} \rangle$	(17)
$\langle \text{p-role} \rangle ::= \langle \text{prin} \rangle \text{ “.” } \langle \text{p-role-term} \rangle \mid \text{Any.} \langle \text{p-role-term} \rangle$	(18)
$\langle \text{p-role-term} \rangle ::= \langle \text{role-name} \rangle \mid \langle \text{role-name} \rangle \text{ “(” } \langle \text{list of p-field} \rangle \text{ “)”}$	(19)
$\langle \text{p-field} \rangle ::= \langle \text{field-name} \rangle (\text{“=”} \mid \text{“⇒”}) (\langle \text{var} \rangle \mid \langle \text{constant} \rangle)$	(20)
$\langle \text{policy-head} \rangle ::= \langle \text{role} \rangle \mid \langle \text{dis-ack} \rangle \mid \langle \text{dis-ac} \rangle \mid \langle \text{dis-full} \rangle \mid \langle \text{dis-bit} \rangle \mid \langle \text{dis-range} \rangle$	(21)
$\langle \text{dis-ack} \rangle ::= \text{“disclose” “(” “ack” “,” } \langle \text{role} \rangle \text{ “)”}$	(22)
$\langle \text{dis-ac} \rangle ::= \text{“disclose” “(” “ac” “,” } \langle \text{role} \rangle \text{ “)”}$	(23)
$\langle \text{dis-full} \rangle ::= \text{“disclose” “(” “full” “,” } \langle \text{attr-name} \rangle \text{ “)”}$	(24)
$\langle \text{dis-bit} \rangle ::= \text{“disclose” “(” “bit” “,” } \langle \text{attr-name} \rangle \text{ “)”}$	(25)
$\langle \text{dis-range} \rangle ::= \text{“disclose” “(” “range” “,” } \langle \text{attr-name} \rangle, \langle \text{precision} \rangle \text{ “)”}$	(26)

Figure 3: Syntax of ATNL in BNF. The first two definitions $\langle \text{list of } X \rangle$ and $\langle \text{set of } X \rangle$ are macros parameterized by X . The symbol ϵ in (2) denotes the empty string. The symbols $\langle \text{var} \rangle$, $\langle \text{constant} \rangle$, and $\langle \text{prin} \rangle$ each represents a variable, a constant, and a principal respectively. The symbols $\langle \text{role-name} \rangle$, $\langle \text{field-name} \rangle$, and $\langle \text{attr-name} \rangle$ represent identifiers drawn from disjoint sets. The syntax for non-terminals $\langle \text{commitment} \rangle$, $\langle \text{precision} \rangle$, $\langle \text{constraint} \rangle$ are not defined here; they are explained in the text.

in computer science.

A *role term* (8) is a role name possibly followed by a list of fields. Each *field* (9) has a field name and a field value. A field value can be a variable, a constant, or a commitment. For example, `SBA.businessLicense` is a role without any fields, `CoS.student(program = ‘cs’, level = ‘sophomore’)` and `BMV.driverLicense(name = commit(‘Alice’), DoB = commit(‘03/07/1986’))` are roles with fields. In the preceding roles, `CoS` is a principal name, `student` is a role name, `program` is a field name, `‘cs’` is a constant of string type, and `commit(‘Alice’)` is a commitment. In ATNL, a *commitment* takes of the form of `commit(c)`, where c is a constant, and `commit` denotes the output of a commitment algorithm of a commitment scheme [13, 27]¹.

If a credential is a regular certificate, such as an X.509 certificate [18], then each field in the credential takes the form of $x = c$, where x is the field name and c is a constant. For example, Alice’s student certificate ($n2$) may be an X.509 certificate. When a

credential is implemented as a cryptographic certificate, such as an OACert or an anonymous credential, the attribute values are committed in the credential. Therefore, each field takes the form of $x = \text{commit}(c)$, where `commit(c)` is the commitment of a constant c . For example, Alice’s digital driver license ($n3$) is modeled as a cryptographic certificate.

4.2.2 Attribute declarations

Each *attribute declaration* (10) gives the name of the attribute, the value of the attribute, a list of attribute references that correspond to this attribute, and whether this attribute is considered sensitive or not. For example, Alice’s attribute declaration ($o1$) in Figure 2 means that Alice has a phone number (123)456-7890 and she considers her phone number to be sensitive information. Alice’s attribute declaration ($o3$) indicates that Alice’s major is ‘cs’ and that her program appears in her student certificate, issued by CoS. We use `attr` to denote attribute names.

Each *attribute reference* (11) corresponds to a field name in a role. The attribute reference is used to link the declared attribute to a specific role field. For example, Alice’s DoB attribute declaration has an attribute reference `BMV.driverLicense(DoB)`, it means that Alice’s DoB is documented in the DoB field of the role `BMV.driverLicense`. It is possible to have several attribute refer-

¹In order to have the hiding property, a commitment scheme usually cannot be deterministic, thus the commitment of a value also depends on a secret random value. For simplicity of presentation, we do not explicitly model the random secret in the representation of a commitment.

ences for an attribute. This means that the attribute is documented by several roles². For example, suppose Alice also has a passport, and her DoB is certified in her passport. Then the attribute declaration for her DoB looks like

DoB = '03/07/1986' :: BMV.driverLicense(DoB),
Gov.passport(BirthDate) :: **sensitive**

Because the disclosure of attribute values in a credential can be separated from the disclosure of the credential, one purpose of the attribute declarations is to uniformly manage the disclosure of an attribute value that appears in different credentials. That is, the policy author gives disclosure policies for attribute DoB, instead of assigning separate disclosure policies for BMV.driverLicense(DoB) and Gov.passport(BirthDate).

When the list of the attribute references is empty, the corresponding attribute does not appear in any role that is certified by a credential. In other words, the attribute is *uncertified* by any authorities. Unlike most prior trust negotiation systems, our framework supports uncertified attributes. In many online e-business scenarios, like the example in Section 4.1, the access control policies require some personal information about the requester, such as phone number and email address, which may not be documented by any digitally signed credentials. Like certified attributes, uncertified attributes may be sensitive, and should be protected in the same way. We treat all attributes uniformly, whether certified or not, by protecting them with disclosure policies.

If an attribute is not sensitive, then the keyword **non-sensitive** appears at the end of its corresponding attribute declaration. This means that the attribute can be revealed to anyone. There is no access control policy for this attribute. On the other hand, if an attribute is treated as a sensitive resource, the attribute owner will mark its attribute declaration with the keyword **sensitive**. In this case, if there are disclosure policy statements for this attribute, one has to satisfy the body of one of these statements to learn information about the attribute. If there is no disclosure policy statement for a sensitive attribute, it means the attribute must never be disclosed.

4.2.3 Policy statements

In ATNL, a *policy statement* (12) takes the form $\langle \text{policy-head} \rangle \leftarrow \langle \text{policy-body} \rangle$ in which $\langle \text{policy-body} \rangle$ either is true or takes the form:

pre-cond-1 ! $B_1.R_1 \cap \dots \cap B_k.R_k$;
pre-cond-2 ! $\psi(x_1, \dots, x_n)$

where B_1, \dots, B_k are principals, R_1, \dots, R_k are role terms, k is an integer greater than or equal to 1, pre-cond-1 and pre-cond-2 are two pre-conditions (which we discuss shortly), ψ is a constraint from a constraint domain Φ , and x_1, x_2, \dots, x_n are the variables in the fields of R_1, \dots, R_k . The constraint $\psi(x_1, \dots, x_n)$ is optional. We call $B_1.R_1 \cap \dots \cap B_k.R_k$ in the policy statement an *intersection*.

A *pre-condition* (16) is defined to be a role or the keyword **false**. The motivation for the pre-condition is that, oftentimes, policies may contain sensitive information. The policy enforcer does not want to reveal the policy statement to everyone. If a pre-condition is false, the pre-condition is never satisfied. If the pre-condition is a role, say $B.R$, then the negotiation opponent has to be a member

of $B.R$ for the pre-condition to be satisfied. Returning to the policy body, if pre-cond-1 is satisfied (or if pre-cond-1 is omitted), then the negotiation opponent is allowed to see $B_1.R_1 \cap \dots \cap B_k.R_k$, otherwise, she is not permitted to know the content of this policy body. Once pre-cond-1 is satisfied, if pre-cond-2 is also satisfied, then the negotiation opponent is allowed to see the constraint $\psi(x_1, \dots, x_n)$.

Verifying that a principal satisfies a policy body takes two steps. In the first step, the policy enforcer verifies that the principal has all roles and has provided all uncertified attributes given by $B_1.R_1, \dots, B_k.R_k$. In the second step, the policy enforcer verifies that the variables in the parameters of R_1, \dots, R_k satisfy the constraint $\psi(x_1, \dots, x_n)$. Such two-step policy verification process is made feasible by using cryptographic credentials and the associated cryptographic tools (see Section 3). The first step corresponds to verifying that the principal has the desired credentials. The second step corresponds to verifying that the principal's attribute values in the credentials satisfy the constraint $\psi(x_1, \dots, x_n)$. If $\psi(x_1, \dots, x_n)$ is disclosed, which happens only when the second pre-condition has been satisfied, then the principal can use zero-knowledge proof protocols to prove that her attribute values satisfy the constraint; otherwise, the principal can elect to run a private policy evaluation protocol with the policy enforcer, enabling each to determine whether she satisfies the constraint.

Using the example in Section 4.1, BookSt's policy (m_2) in Figure 1 is a policy statement with no constraint. It states that BookSt considers a driver license from BMV to provide adequate documentation of date of birth. The variable x is used in the statement to indicate that the field value of BookSt.DoB is the same as the DoB field value in BMV.driverLicense.

The BookSt policy statement (m_1) means that, in order to be a member of the role BookSt.discount, a principal has to have the roles BookSt.student($\text{program} = x_1$), BookSt.DoB($\text{val} = x_2$), and Any.phoneNum($\text{val} \Rightarrow x_3$). It further requires that the program field value x_1 in BookSt.student and the DoB field value x_2 in BookSt.DoB satisfy the constraint ($x_1 = \text{'cs'} \wedge (x_2 > \text{'01/01/1984'})$). The symbol \Rightarrow in the role Any.phoneNum($\text{val} \Rightarrow x_3$) indicates that BookSt must receive a phone number from the negotiation opponent. Where the equality symbol $=$ is used, the policy requires only proof that the associated field value satisfies any constraints given in the policy statement.

4.2.4 Policy heads

The policy head in a policy statement determines which resource is to be disclosed and how it is to be disclosed. A *policy head* (21) can be a role or a disclosure. When the policy head is a role, the statement means that if the negotiation opponent satisfies the policy body, then she is a member of the role. Roles defined in policy statements are controlled by the policy owner and are called *dummy roles* because they are not defined in signed credentials, but serve only to aid in defining local policies. If the policy head is a disclosure, then the opponent is granted a permission specified in the disclosure, once the policy body is satisfied. This section explains each type of disclosure and its associated permission.

We call (the body of) a policy statement with head $\text{disclose}(\text{ack}, A.R)$ (22) an *Ack policy* for the role $A.R$. The opponent has to satisfy one of $A.R$'s Ack policies to gain permission to learn whether the policy enforcer is a member of $A.R$. Until such satisfaction is shown, the policy enforcer's behavior should not depend in any way on whether she belongs to $A.R$.

We call a policy statement with head $\text{disclose}(\text{ac}, A.R)$ (23) an *AC policy* for the credential $A.R \leftarrow D$. We assume, in this case, that the policy enforcer is D and that D has the membership

²We assume that the attribute values from different roles are the same, however we do not require each principal to use the same field name. For example, BMV may use DoB as the field name for date of birth, whereas Gov uses BirthDate as the field name. Name agreement for different field names can be achieved using application domain specification documents [24, 25].

credential $A.R \leftarrow D$. When the negotiation opponent has satisfied an AC policy for the credential $A.R \leftarrow D$, he is authorized to receive a copy of the credential.

We call a policy statement with head $\text{disclose}(\text{full}, \text{attr})$ (24) a *full policy* for the attribute attr . If a full policy for attr is satisfied, the negotiation opponent is allowed to see the full value of attr . When attr is an uncertified attribute, the policy enforcer can simply disclose its value. When the field value linked to the attribute reference of attr is a commitment, the policy enforcer can open the commitment to the opponent.

We call a policy statement with head $\text{disclose}(\text{bit}, \text{attr})$ (25) a *bit policy* for the attribute attr . Bit policies are defined only for certified attributes. If a bit policy for attr is satisfied, the negotiation opponent has the permission to receive one bit of information about the value of attr , in the sense of receiving the answer to the question whether the value satisfies some predicate. We stress that the one bit information of attr in our context is not necessarily the value of a certain bit in the binary representation of attr , but can be the output of any predicate on attr . More specifically, the policy enforcer can run a private policy evaluation with the opponent in which the opponent learns whether attr , together with other attributes of the enforcer, satisfies the opponent's private policy. While specifying the bit disclosure policy, one should be aware that the bit disclosure of attr is vulnerable to a probing attack. If an adversarial opponent runs the private policy evaluation multiple times using different policies that constrain attr , she may learn more information about the value of attr .

We call a policy statement with head $\text{disclose}(\text{range}, \text{attr}, \text{precision})$ (26) a *range policy* for the attribute attr . Range policies are defined only for certified attributes of certain data types, such as finite integer type, finite float type, and ordered enumeration type. If the range policy for attr is satisfied, then the negotiation opponent has permission to learn that attr belongs to a range with the given precision. For example, if the negotiation opponent has satisfied the policy for $\text{disclose}(\text{range}, \text{DoB}, \text{year})$, then she is allowed to know the year of DoB, but not the exact date. How to specify a precision depends on the data type of the attribute. For example, assume credit score takes integer values from 1 to 1000, and Alice has a credit score of 722 documented in her credit report certificate using cryptographic credential schemes. If BookSt satisfies Alice's policy of $\text{disclose}(\text{range}, \text{score}, 50)$, then Alice can prove to BookSt that her credit score is between 701 and 750 using zero-knowledge proof protocols. Similarly, the policy with head $\text{disclose}(\text{range}, \text{score}, 10)$ means that if the policy is satisfied, the opponent can learn that Alice's credit score is between 721 to 730.

5. THE EXTENDED TRUST TARGET GRAPH (ETTG) PROTOCOL

In this section, we introduce a trust negotiation protocol that can take advantage of ATNL and the cryptographic protocols. This protocol extends the trust-target graph protocol introduced in [31, 32], to deal with the additional features of ATNL and cryptographic certificates.

In this protocol, a trust negotiation process involves the two negotiators working together to construct a *trust-target graph* (TTG). A TTG is a directed graph, each node of which is a trust target. Introduced below, trust targets represent questions that negotiators have about each other. When a requester requests access to a resource, the access mediator and the requester enter into a negotiation process. The access mediator creates a TTG containing one target, which we call the *primary target*. The access mediator then

tries to process the primary target by decomposing the question that it asks and expanding the TTG accordingly in a manner described below. It then sends the partially processed TTG to the requester. In each following round, one negotiator receives new information about changes to the TTG, verifies that the changes are legal and justified, and updates its local copy of the TTG accordingly. The negotiator then tries to process some nodes, making its own changes to the graph, which it then sends to the other party, completing the round. The negotiation succeeds when the primary target is satisfied; it fails when the primary target is failed, or when a round occurs in which neither negotiator changes the graph.

5.1 Nodes in a Trust-Target Graph

A node in a TTG is one of the five kinds of targets, defined as follows. We use the notation $e \leftarrow S$ for several different categories of e , meaning that S belongs to, satisfies, or has the property e . We introduce the various usages of the notation informally as they are used in the following list.

- A *role target* takes the form $\langle V : A.R \stackrel{?}{\leftarrow} S \rangle$, in which V is one of the negotiators, $A.R$ is a role, and S is a principal. S is often $\text{opp}(V)$, the negotiator opposing V , but it can be any principal. This target means that V wants to see the proof of $A.R \leftarrow S$.
- A *policy target* takes the form $\langle V : \text{policy-id} \stackrel{?}{\leftarrow} S \rangle$, in which V is one of the negotiators, S is a principal, and policy-id uniquely identifies a policy statement in V 's policy base. We assume each negotiator assigns each of her policy statements a unique identifier for this purpose. This target means that V wants to see the proof that S satisfies the body of the statement corresponding to policy-id .
- An *intersection target* takes the form $\langle V : B_1.R_1 \cap \dots \cap B_k.R_k \stackrel{?}{\leftarrow} S \rangle$, in which V is one of the negotiators, S is a principal, $B_1.R_1, \dots, B_k.R_k$ are roles, k is an integer greater than 1. This means that V wants to see the proof of $B_1.R_1 \cap \dots \cap B_k.R_k \leftarrow S$.
- A *trivial target* takes the form $\langle V : S \stackrel{?}{\leftarrow} S \rangle$, in which V is one of the negotiators, and S is a principal. Trivial targets provide placeholders for edges in the TTG.
- An *attribute goal* takes the form $\langle V : \text{attr} \stackrel{?}{\leftarrow} S \rangle$, in which attr is the name of an attribute in S 's attribute declaration. This goal means that V wants to learn some information about the value of attr , e.g., V may want to learn the full value of the attribute, or to learn partial information about the attribute, e.g., whether it satisfies a policy.

In each of the above forms of targets, we call V the *verifier*, and S the *subject* of this node. Each target has a *satisfaction state*, which has one of three values: *satisfied*, *failed*, or *unknown*. The value is determined inductively depending on the containing TTG structure and the credentials present, as discussed below.

5.2 Edges in a Trust-Target Graph

Seven kinds of edges are allowed in a trust-target graph, listed below. We use \leftarrow to represent edges in TTG's.

- A *credential edge* takes the form $\langle V : A.R \stackrel{?}{\leftarrow} S \rangle \leftarrow \langle V : e \stackrel{?}{\leftarrow} S \rangle$, in which $A.R$ is a role, and e is either a principle or a role. We call $\langle V : e \stackrel{?}{\leftarrow} S \rangle$ a *credential child* of $\langle V : A.R \stackrel{?}{\leftarrow} S \rangle$. (We use similar "child" terminology for other kinds of edges.) An edge always points from the child to the parent. Unlike the

other kinds of edges, a credential edge needs to be *justified* to be added into the TTG; a credential edge is justified if the edge is accompanied by a credential that proves $A.R \leftarrow e$.

- A *policy edge* takes the form $\langle V : A.R \stackrel{?}{\leftarrow} S \rangle \leftarrow \langle V : \text{policy-id} \stackrel{?}{\leftarrow} S \rangle$, in which policy-id is a policy identifier and $A.R$ is the role in the head of the policy statement (that corresponds to policy-id).
- A *policy control edge* takes the form $\langle V : \text{policy-id} \stackrel{?}{\leftarrow} S \rangle \leftarrow \langle V : A.R \stackrel{?}{\leftarrow} S \rangle$, in which policy-id is a policy identifier and $A.R$ is one of the pre-conditions in the policy statement.
- A *policy expansion edge* takes the form $\langle V : \text{policy-id} \stackrel{?}{\leftarrow} S \rangle \leftarrow \langle V : B_1.R_1 \cap \dots \cap B_k.R_k \stackrel{?}{\leftarrow} S \rangle$, in which policy-id is a policy identifier and $B_1.R_1 \cap \dots \cap B_k.R_k$ is the intersection in the policy statement. If $k > 1$, the policy expansion child is an intersection target; otherwise, it is a role target. Each policy expansion edge has associated with it up to one tag consisting of a constraint.
- An *intersection edge* takes the form $\langle V : B_1.R_1 \cap \dots \cap B_k.R_k \stackrel{?}{\leftarrow} S \rangle \leftarrow \langle V : B_i.R_i \stackrel{?}{\leftarrow} S \rangle$, where i is in $1..k$, and k is greater than 1.
- An *attribute edge* takes the form $\langle V : A.R \stackrel{?}{\leftarrow} S \rangle \leftarrow \langle V : \text{attr} \stackrel{?}{\leftarrow} S \rangle$, in which S is the negotiation opponent of V , attr is an attribute name, and $A.R$ is a role. This is used when the attribute attr is linked to a specific field in $A.R$ in S 's attribute declarations.
- An *attribute control edge* takes the form $\langle V : e \stackrel{?}{\leftarrow} S \rangle \leftarrow \langle \text{opp}(V) : \text{policy-id} \stackrel{?}{\leftarrow} V \rangle$, in which $\text{opp}(V)$ denotes the opponent of V , policy-id is a policy identifier, and e is the role or attribute name in the head of the policy statement. Attribute control edges are used for handling disclosure policies. Each attribute control edge has a tag consisting of one of ac, ack, full, bit, or range; in the range case, it also includes a precision parameter.

The optional tag on a policy expansion edge is used to express the constraint portion of the policy statement identified by policy-id. The tag on an attribute control edge characterizes the information that V can gain permission to learn by satisfying the body of the statement identified by policy-id.

5.3 Overview of The Extended Trust-Target Graph (ETTG) Protocol

Due to the space limit, we only sketch the ETTG protocol here and give the protocol details in the full version of this paper [21]. We begin with an example of the ETTG protocol, then briefly discuss how to process each node in TTG, and how to handle constraints in the policies.

EXAMPLE 1. This example is a simple instance of the ETTG protocol and illustrates the usage of the first three properties described in Section 3. Referring to the bookstore example in Section 4.1, we depict the final TTG in Figure 4. Alice and BookSt run the ETTG protocol as follows: As BookSt wants to see the proof of BookSt.discount \leftarrow Alice in order to grant Alice access, BookSt creates the primary target (node 1) for the negotiation and sets its satisfaction state to be unknown. If node 1 becomes satisfied, then the negotiation succeeds. In BookSt's policy base, there is a policy

statement ($m1$) for BookSt.discount, hence BookSt creates a policy target (node 2) and adds a policy edge between node 1 and node 2. As the policy statement ($m1$) has no pre-conditions, BookSt reveals the policy by adding a policy expansion child (node 3) and a constraint tag between the parent (node 2) and the child (node 3). Based on the policy ($m1$), BookSt wants to see Alice's phone number and wants to know whether Alice's program and DoB satisfy his constraint. BookSt then creates node 4, 5, 6 and adds them as intersection children to node 3. Since the role BookSt.DoB is a dummy role and there are policies ($m2, m3$) associated with it, BookSt adds a policy target (node 7) as the policy child to node 6. BookSt then adds a policy expansion child (node 8) to node 7. Similarly, BookSt adds node 9 and 10. Essentially, BookSt wants to see Alice's DoB from either a driver license or a passport. Now BookSt cannot process the TTG any more.

After receiving the TTG from BookSt, Alice begins to process the graph. Alice first discloses her credential $n1$ (as it is not sensitive) and adds a credential child (node 11). She cannot disclose her student credential ($n2$) immediately, as there exists an AC policy ($p1$) for $n2$. Therefore Alice adds a policy target (node 12) and expands it with a role target (node 13). Note that the edge between node 11 and 12 is an attribute control edge, which means that if node 12 is satisfied, then Alice can disclose her student credential ($n2$). Alice also reveals her digital driver license (without revealing her DoB) to BookSt, creates a trivial target (node 14), and adds a credential edge between node 8 and node 14. At this point, Alice notices that she needs to prove she is younger than '1/1/1984' and to reveal her phone number, she adds an attribute goal (node 15) for her DoB attribute and another attribute goal (node 19) for her phoneNum, she also expands the TTG by adding nodes 16, 17, 18, 20. As the node 16 is trivially satisfied (because the policy for $p4$ is true), Alice proves to BookSt that she is born in 1986. Alice's year of birth flows up from node 8 to node 3.

BookSt shows to Alice his StateU.businessLicense certificate and BBB.goodSecProcess certificate, which triggers the satisfaction of the nodes 12 and 20. Alice then reveals her student credential ($n2$) and her uncertified phoneNum. The values of Alice's attribute program and phoneNum flow up to node 3, where BookSt verifies that Alice's attributes satisfy the constraint. Finally, the primary target is satisfied and the negotiation succeeds.

Node Processing in ETTG We briefly explain how each node is processed in the ETTG protocol.

1. *Role target.* Suppose the role in a role target $T = \langle V : V.R \stackrel{?}{\leftarrow} S \rangle$ is a dummy role. For each of the verifier's policies that have $V.R$ as the policy head, the verifier adds a new policy child for T . The role target is satisfied if one of its children is satisfied.

Now suppose the role in a role target $T = \langle V : A.R \stackrel{?}{\leftarrow} S \rangle$ is not a dummy role. If the opponent of V has an Ack policy for $A.R$, he adds an attribute control child for T . Once the Ack policy (if any) has been satisfied, if the opponent has an AC policy for $A.R$, he adds another attribute control child. After all of T 's attribute control children have been satisfied, the opponent can disclose his credential for $A.R$ (if any), adds a credential edge, and marks T satisfied. If the credential disclosed is a membership credential, then for each field in $A.R$, if there is a sensitive attribute linked to the field in the opponent's attribute declarations, the opponent adds an attribute child for that attribute.

2. *Policy target.* Consider the policy target $T = \langle V : \text{policy-id} \stackrel{?}{\leftarrow} S \rangle$. Suppose the policy body associated

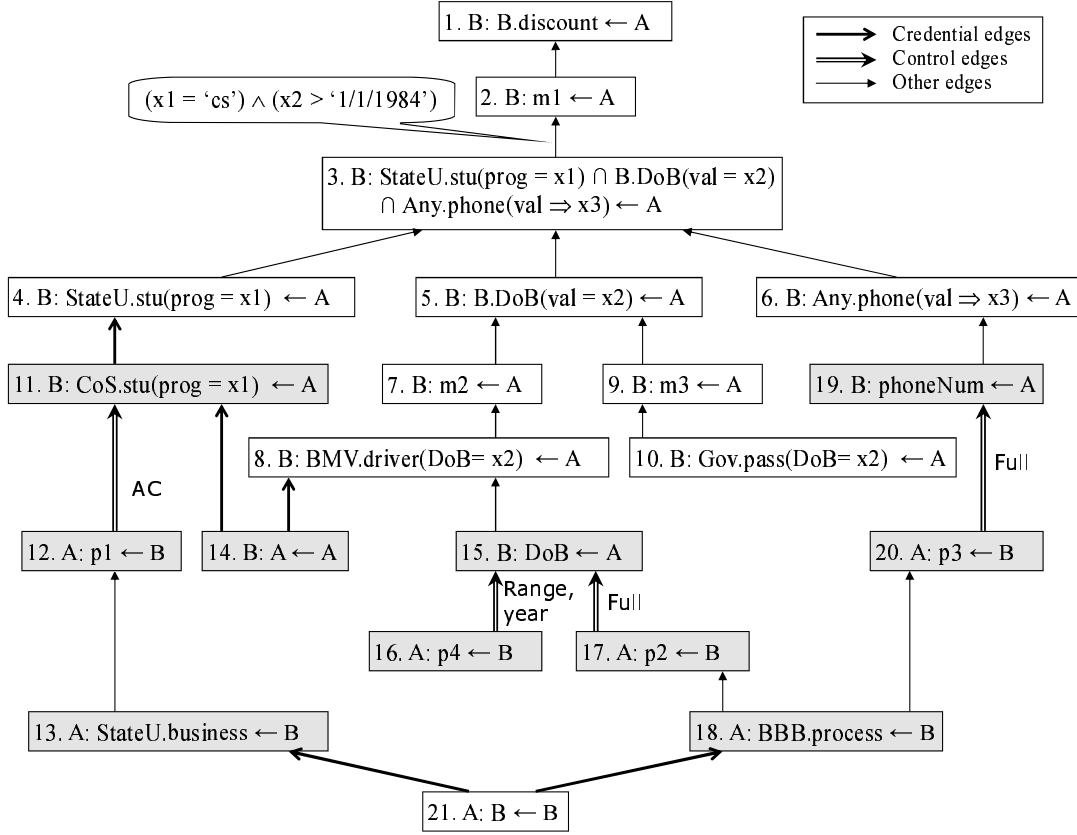


Figure 4: Final TTG for the bookstore example. In this figure, \leftarrow denotes the symbol \leftarrow , A denotes Alice, and B denotes BookSt. The white nodes are created by BookSt and the grey nodes are created by Alice.

with policy-id takes the form pre-cond-1! $B_1.R_1 \cap \dots \cap B_k.R_k$; pre-cond-2! $\psi(x_1, \dots, x_n)$. The verifier first adds a policy control child for pre-cond-1. Once the policy control child is satisfied, the verifier adds a policy expansion child for $B_1.R_1 \cap \dots \cap B_k.R_k$ and a policy control child for pre-cond-2. If the policy control child for pre-cond-2 is satisfied, the verifier adds the constraint $\psi(x_1, \dots, x_n)$ as a tag on the policy expansion edge. A policy target is satisfied if its policy expansion child is satisfied and the constraint is evaluated and satisfied. We explain how and when to evaluate a constraint below.

3. *Intersection target.* For the intersection target $T = \langle V : B_1.R_1 \cap \dots \cap B_k.R_k \xleftarrow{?} S \rangle$, the verifier adds an intersection child for each role $B_i.R_i$. The intersection target is satisfied if all of its intersection children are satisfied.
4. *Attribute goal.* For the attribute target $T = \langle V : \text{attr} \xleftarrow{?} S \rangle$, the opponent adds an attribute policy child for each disclosure policy that contains attr in the policy head.

Constraint Handling We now explain how the constraint of a policy is evaluated. If there is a constraint tag in the policy expansion edge of the policy (*i.e.*, the constraint is revealed to the opponent), it can be evaluated as follows: Whenever a full disclosure policy or a range disclosure policy for an attribute is satisfied by the verifier, the opponent reveals the attribute information accordingly. The attribute information flows from the attribute goal to the policy expansion edge where the constraint is attached. At the policy expansion edge, when the verifier receives enough information from the

opponent to determine whether or not the constraint is satisfied, he evaluates the constraint based on the attribute information received so far and outputs the result.

If there is no constraint tag in the policy expansion edge (*i.e.*, the pre-condition for the constraint has not been satisfied), the verifier can satisfy all the full disclosure policies of the attributes required in the constraint, enabling it to obtain all the attribute values from the opponent. The verifier then evaluates the constraint secretly and informs the opponent the result of the evaluation. Alternatively, the verifier can try to satisfy all the bit disclosure policies of the attributes, and then run a private policy evaluation protocol with the opponent.

5.4 Additional Examples

In this section, we give two additional examples that illustrate the ATNL language and the ETTG protocol. Example 2 deals with the scenario in which the constraint is private; Example 3 illustrates how the ETTG protocol breaks the policy cycles. Due to the space limit, the final TTG graphs for the following two examples are given in the full version of this paper [21].

EXAMPLE 2. This example illustrates the usage of properties 1, 2, and 6 described in Section 3, with a focus on property 6 (certified input private policy evaluation). Suppose BankWon, an online bank certified by National Credit Union Administration (NCUA), offers a special-rate loan. Before applying the loan, an applicant is required to show a valid driver license. The loan policy is that the applicant must have either (1) a credit score more than 680 and an income more than 55k, or (2) a credit score more than 700

and an income more than 45k. BankWon considers his loan policy as private information, and discloses (the thresholds of) the policy only to BankWon's preferred members. Carol, who is not one of BankWon's preferred members, wants to know whether she is eligible for that loan. She has a credit report from Experian and a tax certificate from Internal Revenue Service (IRS). Carol considers her credit score and her income to be sensitive attributes. BankWon and Carol's credentials and policies are given in Figure 6.

Using the ETTG protocol, BankWon and Carol can negotiate trust successfully. In the ETTG protocol, BankWon first creates a primary target for role Bank.loan, then adds a policy target for policy ($r1$). Since there are two pre-conditions for policy $r1$, BankWon creates a role target for the first pre-condition BMV.driverLicense. Essentially, BankWon wants to see a proof of BMV.driverLicense \leftarrow Carol before he discloses the intersection of ($r1$). At this time, BankWon cannot process the TTG any more. Next, Carol expands the TTG by revealing her driver license. Now the role target for BMV.driverLicense becomes satisfied, BankWon adds an intersection target for the intersection of ($r1$) and a role target for the second pre-condition Bank.preferred. After seeing the intersection target, Carol reveals her tax certificate and credit report without revealing her sensitive attributes to BankWon. As the role target for Bank.preferred cannot be satisfied (Carol is not a member of Bank.preferred), the constraint of the loan policy is not revealed to Carol. However, as the bit policies for Alice's income and score can be satisfied by BankWon, Carol and BankWon are able to run a private policy evaluation on income and score with BankWon's private constraint. After the private policy evaluation outputs true (i.e., Carol's certified attributes satisfy the constraint), the policy target for ($r1$) becomes satisfied. In the end, the primary target is also satisfied and the ETTG protocol succeeds.

EXAMPLE 3. This example illustrates the usage of properties 1, 2, 4, and 5 described in Section 3. The focus is on the property 4 and 5 (oblivious usage of credentials and attributes). Suppose Bob, a CIA agent, has a secret document to which the access is allowed by CIA agents only. Bob has a security clearance certificate from Gov with the security level committed in it. Bob can show his CIA agent credential only to his peers, and can reveal his security clearance level only to those whose security level is greater than or equal to 3. Similarly, Alice has a CIA agent credential and a security clearance certificate with certain disclosure policies. Alice shows her CIA agent credential only to CIA agents with security level greater than or equal to 2. And she discloses her security level only to CIA agents. See Figure 7 for the description of the credentials and policies in ATNL.

When Alice wants to access Bob's document, they engage in the ETTG protocol. There are two policy cycles in the TTG, see Figure 5 for the dependency relation between Alice and Bob's attributes. Without breaking the policy cycles, the negotiation between Alice and Bob would fail, because neither Alice nor Bob can update the TTG any more. As the two policy cycles share common nodes, we cannot break them separately. To break the cycles, Alice and Bob run an OSBE protocol [22] in which Bob delivers an envelope to Alice with the property that Alice can open the envelope if she has a CIA agent credential. This envelope contains Bob's CIA agent credential. In the mean time, they run an OCBE protocol [19] in which Bob delivers another envelope to Alice such that Alice can open the envelope if and only if her security level is greater than 2. In the second envelope, Bob opens the commitment of his security level. Bob learns nothing from the previous interactions. After Alice opens the two envelopes, she verifies whether the received CIA credential and security level satisfy her policies. If so, she reveals

her CIA agent credential and her security level to Bob. Now the policy cycles are broken.




Figure 5: Disclosure dependency graph for Alice's and Bob's sensitive attributes.

6. CONCLUSION AND FUTURE WORK

We have introduced a framework for ATN that supports the combined use of several cryptographic credential schemes and protocols that have been previously introduced piecemeal to provide capabilities that are useful in various negotiation scenarios. Our framework enables these various schemes to be combined flexibly and synergistically, on the fly as the need arises. The framework has two key components: ATNL, a policy language that enables negotiators to specify authorization requirements that must be met by an opponent to receive various amounts of information about certified attributes and the credentials that contain it; ETTG, an ATN protocol that organizes negotiation objectives and the use of cryptographic techniques to meet those objectives. We have shown several examples that illustrate how our framework enables negotiations to succeed that would not were they conducted using traditional ATN techniques. The details of the process of constructing the trust target graph and other aspects of the ETTG negotiation state are given in the full version of this paper [21]. In on-going research we are developing comprehensive analysis algorithms that negotiators will use to recognize all cyclic dependencies that can be resolved.

Acknowledgement

This work is supported by NSF IIS-0430274, NSF CCR-0325951, and sponsors of CERIAS. We thank the anonymous reviewers for their helpful comments.

7. REFERENCES

- [1] W. Bagga and R. Molva. Policy-based cryptography and applications. In *Proceedings of the 9th International Conference on Financial Cryptography and Data Security*, Feb. 2005.
- [2] D. Balfanz, G. Durfee, N. Shankar, D. Smetters, J. Staddon, and H.-C. Wong. Secret handshakes from pairing-based key agreements. In *Proceedings of the IEEE Symposium and Security and Privacy*, pages 180–196, May 2003.
- [3] S. Boeyen, T. Howes, and P. Richard. Internet X.509 Public Key Infrastructure LDAPc2 Schema. IETF RFC 2587, June 1999.
- [4] P. Bonatti and P. Samarati. Regulating service access and information release on the web. In *Proceedings of the 7th ACM Conference on Computer and Communications Security (CCS-7)*, pages 134–143. ACM Press, Nov. 2000.

<i>Bank's credentials and policies:</i>		
$q1$:	NCUA.member	← Bank
$r1$:	Bank.loan	← $\text{BMV.driverLicense} ! \text{IRS.tax}(\text{income} = x_1) \cap \text{Bank.credScore}(\text{val} = x_2) ;$ $\text{Bank.preferred} ! ((x_1 > 680) \wedge (x_2 > '55k')) \vee ((x_1 > 700) \wedge (x_2 > '45k'))$
$r2$:	$\text{Bank.credScore}(\text{val} = x)$	← $\text{Equifax.credReport}(\text{score} = x)$
$r3$:	$\text{Bank.credScore}(\text{val} = x)$	← $\text{Experian.credReport}(\text{score} = x)$
$r4$:	$\text{Bank.credScore}(\text{val} = x)$	← $\text{TransUnion.credReport}(\text{score} = x)$
<i>Carol's credentials:</i>		
$s1$:	$\text{Experian.credReport}(\text{score} = \text{commit}(720))$	← Carol
$s2$:	$\text{IRS.tax}(\text{income} = \text{commit}('65k'), \text{employer} = \text{commit}('Company A'))$	← Carol
$s3$:	$\text{BMV.driverLicense}(\text{name} = 'Carol', \text{DoB} = \text{commit}('06/18/1972'))$	← Carol
<i>Carol's attribute declarations:</i>		
$t1$:	$\text{DoB} = '06/18/1972'$:: $\text{BMV.driverLicense}(\text{DoB})$:: sensitive
$t2$:	$\text{score} = 720$:: $\text{Experian.credReport}(\text{score})$:: sensitive
$t3$:	$\text{income} = '48k'$:: $\text{IRS.tax}(\text{income})$:: sensitive
$t4$:	$\text{employer} = 'Company A'$:: $\text{IRS.tax}(\text{employer})$:: non-sensitive
<i>Carol's policies:</i>		
$u1$:	$\text{disclose}(\text{full}, \text{DoB})$	← $\text{BBB.goodSecProcess}$
$u2$:	$\text{disclose}(\text{bit}, \text{score})$	← NCUA.member
$u3$:	$\text{disclose}(\text{range}, \text{score}, 50)$	← true
$u4$:	$\text{disclose}(\text{bit}, \text{income})$	← true
$u5$:	$\text{disclose}(\text{range}, \text{income}, 10k)$	← $\text{BBB.goodSecProcess}$

Figure 6: The credentials and policies for Example 2

<i>Bob's credentials and attributes:</i>		
$v1$:	CIA.agent	← Bob
$v2$:	$\text{Gov.secClearance}(\text{level} = \text{commit}(3))$	← Bob
$v3$:	$\text{level} = 3$:: $\text{Gov.secClearance}(\text{level})$:: sensitive
<i>Bob's policies:</i>		
$w1$:	Bob.document	← CIA.agent
$w2$:	$\text{disclose}(\text{ack}, \text{CIA.agent})$	← CIA.agent
$w3$:	$\text{disclose}(\text{full}, \text{level})$	← $\text{Gov.secClearance}(\text{level} = x) ; x \geq 3$
<i>Alice's credentials and attributes:</i>		
$x1$:	CIA.agent	← Alice
$x2$:	$\text{Gov.secClearance}(\text{level} = \text{commit}(4))$	← Alice
$x3$:	$\text{level} = 4$:: $\text{Gov.secClearance}(\text{level})$:: sensitive
<i>Alice's policies:</i>		
$y1$:	$\text{disclose}(\text{ack}, \text{CIA.agent})$	← $\text{CIA.agent} \cap \text{Gov.secClearance}(\text{level} = x) ; x \geq 2$
$y2$:	$\text{disclose}(\text{full}, \text{level})$	← CIA.agent

Figure 7: The credentials and policies for Example 3

- [5] F. Boudot. Efficient proofs that a committed number lies in an interval. In *Advances in Cryptology: EUROCRYPT '00*, volume 1807 of *Lecture Notes in Computer Science*, pages 431–444. Springer, May 2000.
- [6] R. Bradshaw, J. Holt, and K. Seamons. Concealing complex policies with hidden credentials. In *Proceedings of 11th ACM Conference on Computer and Communications Security*, Oct. 2004.
- [7] S. A. Brands. *Rethinking Public Key Infrastructures and Digital Certificates: Building in Privacy*. MIT Press, Aug. 2000.
- [8] J. Camenisch and E. V. Herreweghen. Design and implementation of the idemix anonymous credential system. In *Proceedings of the 9th ACM Conference on Computer and Communications Security, CCS '02*, pages 21–30. ACM, nov 2002.
- [9] J. Camenisch and A. Lysyanskaya. An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In *Advances in Cryptology: EUROCRYPT '01*, volume 2045 of *Lecture Notes in Computer Science*, pages 93–118. Springer, 2001.
- [10] D. Chaum. Security without identification: Transaction systems to make big brother obsolete. *Communications of the ACM*, 28(10):1030–1044, 1985.
- [11] R. Cramer and I. Damgård. Zero-knowledge proof for finite field arithmetic, or: Can zero-knowledge be for free? In

- Advances in Cryptology: CRYPTO '98*, volume 1462 of *Lecture Notes in Computer Science*, pages 424–441. Springer, 1998.
- [12] R. Cramer, M. K. Franklin, B. Schoenmakers, and M. Yung. Multi-authority secret-ballot elections with linear work. In *Advances in Cryptology: EUROCRYPT '96*, volume 1070 of *Lecture Notes in Computer Science*, pages 72–83. Springer, 1996.
 - [13] I. Damgård and E. Fujisaki. An integer commitment scheme based on groups with hidden order. In *Advances in Cryptology: ASIACRYPT '02*, volume 2501 of *Lecture Notes in Computer Science*, pages 125–142. Springer, Dec. 2002.
 - [14] G. Durfee and M. Franklin. Distribution chain security. In *Proceedings of the 7th ACM Conference on Computer and Communications Security*, pages 63–70. ACM Press, 2000.
 - [15] K. B. Frikken, M. J. Atallah, and J. Li. Hidden access control policies with hidden credentials. In *Proceedings of the 3rd ACM Workshop on Privacy in the Electronic Society*, Oct. 2004.
 - [16] A. Hess, J. Jacobson, H. Mills, R. Wamsley, K. E. Seamons, and B. Smith. Advanced client/server authentication in TLS. In *Network and Distributed System Security Symposium*, pages 203–214, Feb. 2002.
 - [17] J. E. Holt, R. W. Bradshaw, K. E. Seamons, and H. Orman. Hidden credentials. In *Proceedings of the 2nd ACM Workshop on Privacy in the Electronic Society*, Oct. 2003.
 - [18] R. Housley, W. Ford, T. Polk, and D. Solo. Internet X.509 Public Key Infrastructure Certificate and CRL Profile. IETF RFC 2459, Jan. 1999.
 - [19] J. Li and N. Li. OACerts: Oblivious attribute certificates. In *Proceedings of the 3rd Conference on Applied Cryptography and Network Security (ACNS)*, volume 3531 of *Lecture Notes in Computer Science*, pages 301–317. Springer, June 2005.
 - [20] J. Li and N. Li. Policy-hiding access control in open environment. In *Proceedings of the 24th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 29–38. ACM Press, July 2005.
 - [21] J. Li, N. Li, and W. H. Winsborough. Automated trust negotiation using cryptographic credentials. Technical Report CERIAS-TR-2005-59, Center for Education and Research in Information Assurance and Security, Purdue University, Aug. 2005.
 - [22] N. Li, W. Du, and D. Boneh. Oblivious signature-based envelope. In *Proceedings of the 22nd ACM Symposium on Principles of Distributed Computing (PODC)*. ACM Press, July 2003.
 - [23] N. Li and J. C. Mitchell. Datalog with constraints: A foundation for trust management languages. In *Proceedings of the Fifth International Symposium on Practical Aspects of Declarative Languages (PADL 2003)*, pages 58–73. Springer, Jan. 2003.
 - [24] N. Li, J. C. Mitchell, and W. H. Winsborough. Design of a role-based trust management framework. In *Proceedings of the 2002 IEEE Symposium on Security and Privacy*, pages 114–130. IEEE Computer Society Press, May 2002.
 - [25] N. Li, W. H. Winsborough, and J. C. Mitchell. Distributed credential chain discovery in trust management. *Journal of Computer Security*, 11(1):35–86, Feb. 2003.
 - [26] A. Lysyanskaya, R. L. Rivest, A. Sahai, and S. Wolf. Pseudonym systems. In *Selected Areas in Cryptography, 6th Annual International Workshop, SAC '99*, volume 1758 of *Lecture Notes in Computer Science*, pages 184–199. Springer, 1999.
 - [27] T. P. Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In *Advances in Cryptology: CRYPTO '91*, volume 576 of *Lecture Notes in Computer Science*, pages 129–140. Springer, 1991.
 - [28] K. E. Seamons, M. Winslett, and T. Yu. Limiting the disclosure of access control policies during automated trust negotiation. In *Proceedings of the Symposium on Network and Distributed System Security (NDSS'01)*, February 2001.
 - [29] K. E. Seamons, M. Winslett, T. Yu, L. Yu, and R. Jarvis. Protecting privacy during on-line trust negotiation. In *2nd Workshop on Privacy Enhancing Technologies*. Springer-Verlag, Apr. 2002.
 - [30] W. H. Winsborough and N. Li. Protecting sensitive attributes in automated trust negotiation. In *Proceedings of the ACM Workshop on Privacy in the Electronic Society*, pages 41–51. ACM Press, Nov. 2002.
 - [31] W. H. Winsborough and N. Li. Towards practical automated trust negotiation. In *Proceedings of the Third International Workshop on Policies for Distributed Systems and Networks (Policy 2002)*, pages 92–103. IEEE Computer Society Press, June 2002.
 - [32] W. H. Winsborough and N. Li. Safety in automated trust negotiation. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 147–160, May 2004.
 - [33] W. H. Winsborough, K. E. Seamons, and V. E. Jones. Automated trust negotiation. In *DARPA Information Survivability Conference and Exposition*, volume I, pages 88–102. IEEE Press, Jan. 2000.
 - [34] M. Winslett, T. Yu, K. E. Seamons, A. Hess, J. Jacobson, R. Jarvis, B. Smith, and L. Yu. Negotiating trust on the web. *IEEE Internet Computing*, 6(6):30–37, November/December 2002.
 - [35] M. Winslett, T. Yu, K. E. Seamons, A. Hess, J. Jacobson, R. Jarvis, B. Smith, and L. Yu. Negotiating trust on the web. *IEEE Internet Computing*, 6(6):30–37, November/December 2002.
 - [36] T. Yu and M. Winslett. Policy migration for sensitive credentials in trust negotiation. In *Proceedings of the ACM Workshop on Privacy in the Electronic Society*, pages 9–20. ACM Press, Oct. 2003.
 - [37] T. Yu and M. Winslett. Unified scheme for resource protection in automated trust negotiation. In *Proceedings of IEEE Symposium on Security and Privacy*, pages 110–122. IEEE Computer Society Press, May 2003.
 - [38] 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 (TISSEC)*, 6(1):1–42, Feb. 2003.