

Towards a Logic of Privacy-Preserving Selective Disclosure Credential Protocols

Theodoros Balopoulos Stephanos Gritzalis
Dept. of Information and Communication Systems Engineering,
University of the Aegean, Samos, Greece
{tbalopoulos,sgritz}@aegean.gr

Abstract

This paper presents a first approach towards a logic suited for protocols aiming to achieve selective disclosure of credentials while preserving privacy. The analysis draws from the BAN and related logics [2, 10] that are targeted to aid reasoning about authentication protocols, as well as from formal methods on PKIs [5, 6]. The family of protocols directly covered are built using selective disclosure certificates, blind signatures and oneway hash functions as cryptographic primitives. The logic is able to prove that if the protocol's credentials are properly constructed and signed by trusted issuers, they should convince a verifier; furthermore, it provides a framework on which mechanized attacks against privacy may be attempted by an automatic theorem prover. The running example is a protocol by Holt and Seamons, specified in [4].

1. Overview

Although digital certificates are by far the most promising technique for safeguarding electronic communications and transactions, they cause a severe privacy concern: each digital certificate can be traced uniquely to its owner, and can be followed around as it moves through the system. To come around this situation, we need to build digital certificates, or more generally credentials, that preserve privacy by allowing the owner of a certificate to choose what information to disclose and to whom to disclose it [1]. This paper is about formal reasoning [7, 3] on the protocols specified by such credential systems.

2. The Primitives

This section presents the primitives that can be used to build a privacy-preserving selective disclosure protocol.

2.1. The `oneway` and \oplus Functions

We model the `oneway` function to be a bijection with the extra security property that its inverse is not easily computable, i.e. that it can *only* be evaluated by building the table of mappings of the forward function and looking for the desired image (brute force). Furthermore, we model the \oplus function to be a bijection which takes any number of arguments.

We will use the notation $\Downarrow f$ to denote that f is a bijection, f^{-1} to denote the inverse function of f , and $\Downarrow f$ to denote that f is easily computable in the sense described above. We can now state the following rules:

$$\frac{\Downarrow f}{\Downarrow f^{-1}} \quad (R1)$$

$$\frac{\Downarrow f \quad f(a) = f(b)}{a = b} \quad (R2)$$

$$\frac{\Downarrow f}{f(f^{-1}(x)) = x} \quad (R3)$$

We can also state the following axioms for our functions:

$$\overline{\Downarrow \text{oneway} \wedge \Downarrow \text{oneway}} \quad (A1)$$

$$\overline{\Downarrow \oplus \wedge \Downarrow \oplus \wedge \Downarrow \oplus^{-1}} \quad (A2)$$

Using (R1), (R2), (A1) and (A2), we can prove the following theorem:

$$\frac{\text{oneway}(\oplus_{i=1}^n a_i) = \text{oneway}(\oplus_{i=1}^n b_i)}{\forall i \in [1..n]. a_i = b_i} \quad (T1)$$

In practise, `oneway` will be replaced by a hash function and \oplus will be replaced by a function performing concatenation. A hash function may not be a bijection (because

the image is of fixed size), but it resembles one in the sense that it is (practically) impossible to find two different inputs that map to the same image. Moreover, a function performing concatenation may not be a bijection (because there are many possible ways we can separate the inputs from the image), but we only require it to be a bijection so that its image is able to hold all the information of the inputs—and the concatenation function manages to hold most of this information.

2.2. Signatures and Blind Signatures

Signatures can be modelled with a signing function `sign` that accepts as arguments a message and a private key and returns the corresponding signature, and a signature validation boolean function `sign_valid` that accepts as arguments the message, the public key and the signature. These two functions are related by the following rule (`key-1` denotes the private key that corresponds to the public key `key`):

$$\frac{}{\text{sign_valid}(\text{msg}, \text{key}, \text{sign}(\text{msg}, \text{key}^{-1}))} \quad (\text{R4})$$

Blind signatures can be modelled based on Chaum's blinding [8, 9]. We need the blinding function `blind` that accepts as arguments a message, a blinding factor and the public key of the entity that will do the signing, and returns the message blinded. We also need the unblinding function `unblind` that accepts as arguments the signature of the blinded message and the blinding factor, and returns the signature of the (unblinded) message. These functions are related as shown below:

$$\frac{\text{sign_valid}(\text{blind}(\text{msg}, \text{factor}, \text{key}), \text{key}, \text{sgn})}{\text{sign_valid}(\text{msg}, \text{key}, \text{unblind}(\text{sgn}, \text{factor}))}$$

This rule might be satisfying in providing a primitive for blind signatures; however, the mathematics behind Chaum's blinding allow us to use the more general rule given below:

$$\frac{\text{sign_valid}\left(\prod_{i=1}^n \text{blind}(\text{msg}_i, \text{factor}, \text{key}), \text{key}, \text{sgn}\right)}{\text{sign_valid}\left(\prod_{i=1}^n \text{msg}_i, \text{key}, \text{unblind}(\text{sgn}, \text{factor})\right)} \quad (\text{R5})$$

We also include an axiom that says that all these functions are easily computable:

$$\frac{}{\downarrow \text{sign} \wedge \downarrow \text{sign_valid} \wedge \downarrow \text{blind} \wedge \downarrow \text{unblind}} \quad (\text{A3})$$

2.3. The Selective Disclosure Certificate

A selective disclosure certificate `cert` will be modelled as having three fields: the issuer—denoted by `I(cert)`, the committed attribute—denoted by `A(cert)`, and the credential-id—denoted by `C(cert)`. The signature will not be part of the certificate.

The committed attribute is intended to be used in such a way that, on one hand, the real attribute cannot be deduced from it, and, on the other hand, the entity that has committed to it is able to convince other entities of the real attribute's value. In the discussion that follows, we will assume that an entity that creates a certificate `cert` and wants to commit to a real attribute denoted by `attrib(cert)`, will compute the committed attribute as illustrated by the formula

$$A(\text{cert}) = \text{oneway}(\text{attrib}(\text{cert}) \oplus \text{random}(\text{cert}))$$

where `random(cert)` denotes a random value the entity generates for this particular certificate and does not disclose to any other entity.

Notice that one disclosed attribute is enough for our modelling purposes: if more attributes are needed, one may use more certificates in his model.

2.4. Relationships between Entities and Data

Having *P* can *have* the value of an expression if he can compute the expression by himself, or if he received this value during the execution of a protocol¹.

If *P* *has* the arguments of an easily computable function, then he also *has* their image under this function. Using the notation `x1..n` or `(x1, ..., xn)` to denote the *n* arguments of a function, we have the following rules:

$$P \text{ has } \mathbf{x}_{1..n} \Leftrightarrow \forall x \in [1..n]. P \text{ has } x_i \quad (\text{R6})$$

$$\frac{P \text{ has } \mathbf{x}_{1..n} \quad \downarrow f}{P \text{ has } f(\mathbf{x}_{1..n})} \quad (\text{R7})$$

Using (R3) and (R7), we can prove the following theorem:

$$\frac{P \text{ has } f(\mathbf{x}_{1..n}) \quad \Downarrow f \quad \downarrow f^{-1}}{P \text{ has } \mathbf{x}_{1..n}} \quad (\text{T2})$$

Notice that if *P* *has* `oneway(x)`, it does not follow that *P* *has* `x`, because `↓oneway-1` is not true.

¹(M1) and (M2) formalize the latter case.

Deducing P can *deduce* that an expression is true, either because he trusts someone else to have verified it for him², or because he is able to evaluate the expression by himself. In the latter case, in order for P to deduce the expression, the expression must be true, and P must *have* the expression, in the sense of rule (R7):

$$\frac{x \quad P \text{ has } x}{P \text{ deduces } x} \quad (\text{R8})$$

The notion of this rule is that if P and Q participate in a protocol and Q deduces an expression, then P will be able to deduce it as well if Q sends P all the necessary information.

We also include the following rules about deduction in our logic:

$$P \text{ deduces } (x \wedge y) \Leftrightarrow P \text{ deduces } x \wedge P \text{ deduces } y \quad (\text{R9})$$

$$\frac{P \text{ deduces } x \quad x \Rightarrow y}{P \text{ deduces } y} \quad (\text{R10})$$

Identifying P can *identify* a piece of data. If this is the case, P can also *identify* the image of this piece of data under a bijection, assuming the bijection is easily computable:

$$\frac{P \text{ identifies } x \quad \Downarrow f \quad \Downarrow f}{P \text{ identifies } f(x)}$$

More generally, he can *identify* the image of a number of arguments under an easily computable bijection, if he can *identify* at least one of these arguments and *have* the rest:

$$\frac{\exists i \in [1..n] . P \text{ identifies } x_i \quad \Downarrow f \quad \Downarrow f \quad \forall i \in [1..n] . (P \text{ identifies } x_i \vee P \text{ has } x_i)}{P \text{ identifies } f(x_{1..n})} \quad (\text{R11})$$

Using (A1), (A2) and (R11), we can prove the following theorem:

$$\frac{\exists i \in [1..n] . P \text{ identifies } x_i \quad \forall i \in [1..n] . (P \text{ identifies } x_i \vee P \text{ has } x_i)}{P \text{ identifies } \text{oneway}(\oplus_{i=1}^n x_i)} \quad (\text{T3})$$

The notion of identification is similar to the notion of freshness in the BAN logic [2, 10]; however, we distinguish two important differences:

- Entities are able to identify images of data under any bijection, not only under conjunction (which in BAN

logic is a bijection). This is reasonable because a bijection does not throw away any information.

- Identification is a more general notion: whether it is used for freshness and/or another purpose depends on the reasoning behind the initial state's assumptions.

Trusting P can *trust* a certificate's issuer. It is implied that P *has* the public key of the issuer of the certificate he *trusts*:

$$\frac{P \text{ trusts } I(\text{cert})}{P \text{ has } \text{key}_{I(\text{cert})}} \quad (\text{R12})$$

Validating P will *validate* a certificate if he *deduces* that it is validly signed:

$$\frac{P \text{ deduces } \text{sign_valid}(\text{cert}, \text{key}_{I(\text{cert})}, \text{sgn})}{P \text{ validates } \text{cert}} \quad (\text{R13})$$

We will also define the notion of P *validating* a list of certificates signed *together*—such a list of *n* certificates will be denoted as $\text{cert}_{1..n}$. All the certificates must include the (same) signer in their issuer field. We also allow the signature to be taken on the product of the certificate images:

$$\begin{aligned} n > 1 \Rightarrow [& P \text{ validates } \text{cert}_{1..n} \Leftrightarrow \\ & P \text{ deduces } \text{sign_valid}\left(\prod_{i=1}^n f(\text{cert}_i), \text{key}_{I(\text{cert}_i)}, \text{sgn}\right) \\ & \wedge \left(\forall i, j \in [1..n] . I(\text{cert}_i) = I(\text{cert}_j)\right)] \end{aligned} \quad (\text{R14})$$

P will also *validate* a certificate if P has *validated* the list of certificates it belongs to:

$$\frac{P \text{ validates } \text{cert}_{1..n}}{\forall i \in [1..n] . P \text{ validates } \text{cert}_i} \quad (\text{R15})$$

Therefore, P can *validate* a single certificate either by checking its signature and using rule (R13), or by checking the signature of the list this certificate belongs to and using rules (R14) and (R15). Only knowing that P *validates* a certificate does not allow us to distinguish between these two cases.

3. Protocol Specification

The protocol will be specified by a list of message passings; each message passing may have one of the following forms:

²This is the case with rule (R16). Our logic needs to have such rules to handle trust, hence it is not a logic of knowledge [10] in the sense that $(P \text{ deduces } x) \Rightarrow x$ is not a theorem.

$$\begin{array}{lcl} L_i & \rightarrow & R_i : m_i \\ L_i & \rightarrow & R_i : c_i ? m_i \end{array}$$

The first form denotes that entity L_i sends to entity R_i message m_i ; the second denotes the same, subject to the condition that c_i holds. The subscript i indicates that this is the i^{th} message passing specified by the protocol.

The protocol specification is used to augment the set of valid statements, starting with the set of initial assumptions which will be denoted by \mathcal{T}_0 . When a message passing of the first form is reached, the set of valid statements may be augmented as specified below:

$$\mathcal{T}_i \equiv \mathcal{T}_{i-1} \cup \left\{ \frac{\mathcal{T}_{i-1} \vdash L_i \text{ has } m_i}{R_i \text{ has } m_i} \right\} \quad (\text{M1})$$

where $\mathcal{T} \vdash e$ denotes that e can be derived from \mathcal{T} . Similarly, when a message passing of the second form is reached, the following augmentation applies:

$$\mathcal{T}_i \equiv \mathcal{T}_{i-1} \cup \left\{ \frac{\mathcal{T}_{i-1} \vdash (c_i \wedge L_i \text{ has } m_i)}{R_i \text{ has } m_i} \right\} \quad (\text{M2})$$

The fully augmented set of valid statements—denoted by $[\mathcal{T}]$ —can be used to prove desirable or undesirable properties of the protocol.

4. The Selective Disclosure Protocol of Holt and Seamons

We will first present a formalization of the selective disclosure credentials used in Holt and Seamons' protocol in Section 4.1, we will then give a formal specification of the protocol in Section 4.2, and finally we will state in our logic desirable and undesirable properties of the protocol in Section 4.3.

4.1. The Selective Disclosure Credential

We need our logic to cover the following types of credentials: a credential consisting of a single certificate, a credential consisting of n certificates issued by the same issuer, and a credential consisting of $n \times m$ certificates issued by m issuers.

Credential consists of a single Certificate If P *trusts* an issuer of certificates and P *validates* a certificate of this issuer, then P can *deduce* that cert 's fields have appropriate values. Regarding the field denoted by $A(\text{cert})$, this translates to the following rule:

$$\frac{P \text{ trusts } I(\text{cert}) \quad P \text{ validates } \text{cert}}{P \text{ deduces } A(\text{cert}) = \text{oneway}(\text{attrib}(\text{cert}) \oplus \text{random}(\text{cert}))} \quad (\text{R16})$$

Using axioms (A1) and (A2), rules (R7), (R8), (R9), (R10), (R16), and theorem (T1), we can prove that someone can *deduce* the undisclosed attribute of a certificate if he *trusts* the certificate's issuer, if he *validates* the certificate, and if he has been given an undisclosed attribute and a random value that map correctly to the certificate's committed attribute. We therefore have the following theorem:

$$\frac{P \text{ trusts } I(\text{cert}) \quad P \text{ validates } \text{cert} \quad P \text{ has } (\text{cert}, x, y) \quad A(\text{cert}) = \text{oneway}(x \oplus y)}{P \text{ deduces } \text{attrib}(\text{cert}) = x} \quad (\text{T4})$$

Credential consists of n Certificates The undisclosed attribute of a credential consisting of n certificates is defined to be the common undisclosed attribute of these certificates, if such a common value exists:

$$\frac{\exists x. \forall i \in [1..n]. \text{attrib}(\text{cert}_i) = x}{\text{attrib}(\text{cert}_{1..n}) = x} \quad (\text{R17})$$

Using rules (R14), (R15), (R17) and theorem (T4), and splitting the two cases $n = 1$ and $n > 1$, we can prove the following theorem:

$$\frac{P \text{ trusts } I(\text{cert}_1) \quad P \text{ validates } \text{cert}_{1..n} \quad P \text{ has } x \quad \forall i \in [1..n]. P \text{ has } (\text{cert}_i, y_i) \quad \forall i \in [1..n]. A(\text{cert}_i) = \text{oneway}(x \oplus y_i)}{P \text{ deduces } \text{attrib}(\text{cert}_{1..n}) = x} \quad (\text{T5})$$

In the special case $n = 1$, (T5) is derived trivially from (T4).

Credential consists of $n \times m$ Certificates The undisclosed attribute of a credential consisting of m rows of n certificates so that the certificates in each row are issued by the same entity, is defined to be the list of the undisclosed attributes of the m credentials consisting of n certificates, if the credential-ids of the certificates in each column are mutually *identifiable* by any entity:

$$\frac{\forall i \in [1..m]. \text{attrib}(\text{cert}_{1..n,i}) = x_i \quad \forall i \in [1..n]. [Q \text{ identifies } C(\text{cert}_{i,1}) \Rightarrow \forall j \in [1..m]. Q \text{ identifies } C(\text{cert}_{i,j})]}{\text{attrib}(\text{cert}_{1..n,1..m}) = x_{1..m}} \quad (\text{R18})$$

4.2. Specifying the Protocol

Regarding the formalization of the protocol in the next subsections, the following points should be taken into account:

- In the issuing part of the protocol, we will not deal with how the attribute claimed by the client is verified by the issuer.
- We will not deal with how proof of ownership of the credential is achieved.
- The non-interactive cut-and-choose protocol will be modelled by the boolean function \mathcal{H} , which takes as arguments a *credential set request* (CSR) and a column number and returns true if and only if this column should be presented.

The rest of this Section assumes familiarity with [4].

The Issuing Sub-Protocol The issuing sub-protocol is ran once for each disclosed attribute required. I_k represents the issuer of the k^{th} disclosed attribute; A represents the issuers' client, i.e. the entity that wishes to acquire a selective disclosure credential. Note that the first message passing must be completed with each issuer before proceeding to the remaining message passings.

$$\begin{aligned}
 I_k &\rightarrow A : \text{_session}_k \\
 A &\rightarrow I_k : \text{CSR}, \text{hcred}_{1..n}, \text{hcredrnd}_{1..n}, \\
 &\quad \text{hfactor}_{1..n}, \text{hcert}_{1..n,1..m}, \\
 &\quad \text{hattrib}_{1..m}, \text{hrandom}_{1..n,1..m} \\
 I_k &\rightarrow A : \text{VLD}_k ? \text{SGN}_k
 \end{aligned}$$

where:

$$\begin{aligned}
 \text{CSR} &\equiv \text{session}_{0..m}, \text{obcred}_{1..n}, \text{blindcert}_{1..n,1..m} \\
 \text{SGN}_k &\equiv \text{sign}\left(\prod_{i=1}^n \neg \mathcal{H}(\text{CSR}, i) \text{blindcert}_{i,k}, \text{key}_{I_k}^{-1}\right) \\
 \text{VLD}_k &\equiv I_k \text{ deduces } \text{CND}_k \\
 &\quad \wedge \forall i \in [1..n]. \mathcal{H}(\text{CSR}, i) \Rightarrow \\
 &\quad (\forall j \in [1..m]. I_k \text{ identifies } \mathcal{C}(\text{hcert}_{i,j})) \\
 \text{CND}_k &\equiv \text{session}_k = \text{_session}_k \\
 &\quad \wedge \forall i \in [1..n]. \mathcal{H}(\text{CSR}, i) \Rightarrow \\
 &\quad (I(\text{hcert}_{i,k}) = I_k \wedge \forall j \in [1..m]. \\
 &\quad [A(\text{hcert}_{i,j}) = \text{oneway}(\text{hattrib}_j \oplus \text{hrandom}_{i,j}) \\
 &\quad \wedge \text{blindcert}_{i,j} = \text{blind}(\text{hcert}_{i,j}, \text{hfactor}_k, \text{key}_{I_k})])
 \end{aligned}$$

The Showing Sub-Protocol The showing sub-protocol is run only once. V represents the verifier of the selective disclosure credential.

$$A \rightarrow V : \text{cert}'_{1..n',1..m'}, \text{attrib}'_{1..m'}, \\
 \text{random}'_{1..n',1..m'}, \text{sgn}_{1..m'}$$

Initial State Assumptions In the case of a successful execution of the protocol, the set of initial assumptions \mathcal{T}_0 should contain the following:

1. Expressions that define hcred , hcredrnd , hfactor , hcert , hattrib and hrandom —representing the information disclosed by the noninteractive cut and choose protocol—in terms of cred , credrnd , factor , cert , attrib and random —representing the full information.

Here is the definition of hcert ; the others are defined in a similar way:

$$\text{hcert}_{i,j} = \begin{cases} \text{cert}_{i,j} & \text{if } \mathcal{H}(\text{CSR}, i) \\ \epsilon & \text{otherwise} \end{cases}$$

2. Each issuer owns his private key, issuers' client has obtained the corresponding public keys, and the verifier trusts the issuers:

$$I_k \text{ has } \text{key}_{I_k}^{-1} \wedge A \text{ has } \text{key}_{I_k} \wedge V \text{ trusts } I_k$$

3. Each issuer generates a fresh number and is able to identify it:

$$I_k \text{ has } \text{session}_k \wedge I_k \text{ identifies } \text{session}_k$$

4. The same applies for the issuers' client:

$$A \text{ has } \text{session}_0 \wedge A \text{ identifies } \text{session}_0$$

5. If the issuers' client is given all the fresh values, he is able to send to each issuer the information required by the protocol:

$$\begin{aligned}
 &A \text{ has } \text{session}_0 \wedge A \text{ has } \text{session}_{1..m} \Rightarrow \\
 &A \text{ has } (\text{CSR}, \text{hcred}_{1..n}, \text{hcredrnd}_{1..n}, \text{hfactor}_{1..n}, \\
 &\quad \text{hcert}_{1..n,1..m}, \text{hattrib}_{1..m}, \text{hrandom}_{1..n,1..m})
 \end{aligned}$$

6. The issuer's client will generate the information required by the protocol correctly (to be examined together with the definitions of VLD_k and CND_k , and with theorem T3):

$$\begin{aligned}
 &(\forall k \in [1..m]. \text{CND}_k) \\
 &\wedge (\forall i \in [1..n]. \mathcal{H}(\text{CSR}, i) \Rightarrow \\
 &\quad \forall j \in [1..m]. [\mathcal{C}(\text{hcert}_{i,j}) = \\
 &\quad \text{oneway}(\text{id} \oplus (\oplus_{j=0}^n \text{session}_j) \oplus \text{hcredrnd}_i) \\
 &\quad \wedge \mathcal{C}(\text{hcert}_{i,j}) = \text{hcred}_i \\
 &\quad \wedge \text{obcred}_i = \text{oneway}(\text{hcred}_i)])
 \end{aligned}$$

7. Each issuer will sign his credential:

$$\begin{aligned}
 &\forall k \in [1..m]. \\
 &\text{SGN}_k = \text{sign}\left(\prod_{i=1}^n \neg \mathcal{H}(\text{CSR}, i) \text{blindcert}_{i,k}, \text{key}_{I_k}\right)
 \end{aligned}$$

8. After deciding which attributes A wishes to show to V , we need to include:

- The expressions that define cert' , attrib' and random' —representing a subset of the information *not* disclosed by the noninteractive cut and choose protocol—in terms of cert , attrib and random —representing the full information.
- The expression that defines sgn —representing the signatures of the certificates whose attribute was selected for disclosure—in terms of SGN —representing all available signatures.

4.3. Protocol Properties

We can now prove desirable properties, such as the one that follows (to be examined together with rule R18):

$$[\mathcal{T}] \vdash \forall \text{ deduces } (\exists x_{1..m} . \text{attrib}(\text{cert}'_{1..n', 1..m'}) = x_{1..m})$$

This says that the verifier is convinced that the real attributes of the certificates' owner are indeed the ones he claims.

We can also attempt to prove undesirable properties in order to find attacks against the our protocol. Consider the following:

$$[\mathcal{T}] \vdash \exists k \in [1..m] . \exists i \in [1..n] . \neg \mathcal{H}(\text{CSR}, i) \wedge \exists j \in [1..m] . I_k \text{ identifies } C(\text{cert}_{i,j})$$

This says that there exists an issuer that can identify a certificate that was not chosen to be disclosed by the non-interactive cut-and-choose protocol. If we could prove it, we would expose a serious protocol weakness, because the owner's identity would be exposed as soon as he used the certificate.

5. Conclusions and Future Work

This work is a first attempt to reason about privacy-preserving selective disclosure credential protocols. We believe that our logic has captured and formalized the primitive methods and relationships that such a protocol will make use of. After formalizing the running example protocol, we are able to prove that properly constructed certificates signed by trusted issuers should convince a verifier. However, we are not able to prove privacy; we are only able to provide a framework on which mechanized attacks may be attempted by an automatic theorem prover.

Future work may include the following:

- Doing research on how privacy can be proved.
- Expanding the logic so that we are able to reason about the issuing protocol in more depth.

- Expressing the logic and the example protocol in an automatic theorem prover and proving desirable or undesirable properties.

We hope that our work will encourage further research into the area of privacy-preserving security systems, and into formal methods for reasoning about them.

References

- [1] S. A. Brands. *Rethinking Public Key Infrastructures and Digital Certificates—Building in Privacy*. MIT Press, 2000.
- [2] M. Burrows, M. Abadi, and R. Needham. A logic of authentication. *ACM Transactions on Computer Systems*, 8(1):18–36, Feb 1990.
- [3] S. Gritzalis, D. Spinellis, and P. Georgiadis. Security protocols over open networks and distributed systems: Formal methods for their analysis, design, and verification. In *Computer Communications Journal*, volume 22, pages 697–709. Elsevier Science, May 1999.
- [4] J. E. Holt and K. E. Seamons. Selective disclosure credential sets. <http://citeseer.nj.nec.com/541329.html>, 2002.
- [5] C. Liu, M. A. Ozols, and T. Cant. An axiomatic basis for reasoning about trust in PKIs. In *Proceedings of the 6th Australasian Conference on Information Security and Privacy (ACISP 2001)*, volume 2119 of *Lecture Notes in Computer Science*. Springer, 2001.
- [6] C. Liu, M. A. Ozols, M. Henderson, and T. Cant. A state-based model for certificate management systems. In *Proceedings of the 3rd International Workshop on Practice and Theory in Public Key Cryptosystems (PKC 2000)*, volume 1751 of *Lecture Notes in Computer Science*. Springer, 2000.
- [7] C. A. Meadows. Formal verification of cryptographic protocols: A survey. In *ASIACRYPT: Advances in Cryptology – ASIACRYPT: International Conference on the Theory and Application of Cryptology*. LNCS, Springer-Verlag, 1994.
- [8] A. J. Menezes, P. C. van Oorschot, and S. A. Vanstone. *Handbook of Applied Cryptography*. CRC Press, 1997.
- [9] B. Schneier. *Applied Cryptography*. John Wiley & Sons, 1996.
- [10] P. Syverson and I. Cervesato. The logic of authentication protocols. In *Foundations of Security Analysis and Design*, volume 2171 of *Tutorial Lectures*. Springer, 2001.