Available online at www.sciencedirect.com

SCIENCE DIRECT®

Computer Standards & Interfaces 27 (2005) 501–512

COMPUTER STANDARDS & INTERFACES

www.elsevier.com/locate/csi

Specifying Privacy-Preserving Protocols in Typed MSR

Theodoros Balopoulos*, Stefanos Gritzalis, Sokratis K. Katsikas

Laboratory of Information and Communication Systems Security, Department of Information and Communication Systems Engineering, University of the Aegean, Karlovassi, Samos, GR-83200, Greece

Available online 3 February 2005

Abstract

Privacy-preserving protocols, such as electronic cash, electronic voting and selective disclosure protocols, use special message constructors that are not widely used in other types of protocols (for example, in authentication protocols). These message constructors include blind signatures, commitments and zero-knowledge proofs. Furthermore, a standard formalization of the Dolev-Yao intruder does not take into account these message constructors, nor does it consider some types of attacks (such as privacy attacks, brute-force dictionary attacks and known-plaintext attacks) that privacy-preserving as well as other types of protocols are designed to protect against. This paper aims to present an extension of Typed MSR in order to formally specify the needed message constructors, as well as the capabilities of a Dolev-Yao intruder designed to attack such protocols.

Keywords: Specification of security protocols; Privacy; Dolev-Yao intruder; Typed MSR

1. Introduction

Formal methods are an important tool for designing and implementing secure cryptographic protocols. By applying techniques concerned with the construction and analysis of models and proving that certain properties hold in the context of these models, formal methods can significantly increase one's confidence that a protocol will meet its requirements in the real world. However, some requirements are not covered as much as others in the existing work on formal methods. A good example is the requirement for the preservation of an entity's privacy. This paper builds on the Typed MSR specification language [2,3] and aims to make it suitable for the specification of privacy-preserving protocols, as well as for the specification of a version of the Dolev-Yao intruder [5] that is designed to attack such protocols. Some aspects of these extensions are useful in other types of protocols as well.

The paper is organized as follows. In Section 2, we give an overview of the standard version of Typed MSR, as well as our extensions of the language's message constructors. In Section 3, we demonstrate how our extensions can be used to make an

* Corresponding author.
E-mail address: tbalopoulos@aegean.gr (T. Balopoulos).

0920-5489/$ - see front matter © 2005 Elsevier B.V. All rights reserved.
doi:10.1016/j.csi.2005.01.008

abstraction of three simple privacy-preserving protocols. In Section 4, we give an overview of typing in Typed MSR, present our typing extensions and apply them to our newly introduced message constructors. In Section 5, we use our syntactical and typing infrastructure to formally specify the capabilities of a Dolev-Yao intruder targeted for privacy-preserving protocols. We conclude the paper with Section 6.

2. Typed MSR

Typed MSR is a strongly typed specification language for security protocols, aiming to discover errors in their design. It is particularly suitable for privacy-preserving protocols because it features memory predicates, which enable it to faithfully encode systems consisting of a collection of coordinated subprotocols—a common characteristic of privacy-preserving protocols (consider for example the electronic cash protocol, which consists of an issuing and a showing/spending subprotocol). However, the standard language does not support the message constructors needed for privacy-preserving protocols. In Section 2.1 we give an overview of messages in the standard version of Typed MSR, and in Section 2.2 we introduce the needed message constructors.

2.1. Overview of messages in Typed MSR

In Typed MSR, messages are obtained by applying message constructors to a variety of atomic messages. Typically, the atomic messages include principals, keys, nonces and raw data. This is formalized by the following grammatical production:

$$\begin{aligned} \textit{Atomic messages}\colon a ::= {} & \mathsf{A} \quad (\textit{Principal}) \\ \mid {} & \mathsf{k} \quad (\textit{Key}) \\ \mid {} & \mathsf{n} \quad (\textit{Nonce}) \\ \mid {} & \mathsf{m} \quad (\textit{Raw data}) \end{aligned}$$

In Typed MSR A, k, n and m range over principal names, keys, nonces and raw data, respectively. Raw data denotes pieces of data whose sole function in a protocol is that they are transmitted.

The message constructors typically present in Typed MSR are those formalized by the following grammatical production:

$$\begin{aligned} \textit{Messages}\colon t ::= {} & a \quad (\textit{Atomic messages}) \\ \mid {} & x \quad (\textit{Variables}) \\ \mid {} & t_1, t_2 \quad (\textit{Concatenation}) \\ \mid {} & \{t\}_k \quad (\textit{Symmetric-key encryption}) \\ \mid {} & \{\{t\}\}_k \quad (\textit{Asymmetric-key encryption}) \\ \mid {} & [t]_k \quad (\textit{Digital Signature}) \end{aligned}$$

We will use the letter t (possibly sub-scripted) to range over messages. We will write A, k, n and m (possibly sub-scripted) for atomic constants or variables that are principals, keys, nonces and raw data, respectively. We will also use the letter B for principals and the letter S for servers (which are also principals). Note that in Typed MSR, the seriffed letters are used whenever the object we want to refer to cannot be but a constant.

In this paper we choose a different meaning for the digital signature constructor than the meaning chosen in standard MSR. Instead of $[t]_k$ denoting both the message t and its digital signature using key k, here it will denote only the latter. This will become evident in Section 3, where we present a high level view of some privacy-preserving protocols.

2.2. Adding message constructors for privacy-preserving protocols

To cope with privacy-preserving protocols we add message constructors for blinding, commitment and zero-knowledge proofs:

$$\begin{aligned} \textit{Messages}\colon t ::= {} & \ldots \quad (\textit{see above}) \\ \mid {} & \langle t \rangle_n^k \quad (\textit{Blinding}) \\ \mid {} & \|t\|_n \quad (\textit{Commitment}) \\ \mid {} & [\![t : n_s : k : n_f]\!] \quad (\textit{Zero-knowledge proof}) \end{aligned}$$

The abstraction of blinding is based on Chaum's blinding [7,4], according to which the construction of a blinded message depends on a blinding factor (which we can abstract as a nonce) and on a public key. The fundamental property is that if message $\langle t \rangle_n^k$ is signed using k' (the private key corresponding to public key k), the resulting message can be unblinded

using nonce n to produce the digital signature of message t signed using k'.

The abstraction of commitment is based on the non-interactive bit commitment using one-way hash functions [8]. According to this method, the commitment of a message is the hash of the concatenation of the message with a salt value (which we can abstract as a nonce). The fundamental properties are that observing $||t||_n$ will not reveal the values of t and n, and that there is only one commitment for each distinct message–nonce pair. Note that the latter property is implicit, because Typed MSR messages are atomic and can solely be constructed by message constructors.

The abstraction of a zero-knowledge proof is based on the non-interactive cut-and-choose protocol introduced in the selective disclosure protocol of Holt and Seamons. The interested reader can refer to section 3.2.2 of [6]. The fundamental property is that observing $[\![t : n_s : k : n_f]\!]$ reveals the values of t and $\langle ||t||_{n_s} \rangle^k_{nf}$, but not the values of n_s, k and n_f.

Notice that we have chosen to make all our new message constructors non-interactive, so that they share this property with the standard message constructors of Section 2.1.

3. Privacy-preserving protocols overview

At this point, we will demonstrate how the message constructors described above may be used to make abstractions of three simple privacy-preserving protocols: an electronic cash protocol, an electronic voting protocol and a selective disclosure protocol. The aim is not to make abstractions of real-world privacy-preserving protocols, but only to justify the introduction of our new message constructors.

3.1. Electronic cash protocol

3.1.1. Issuing

Alice wants to have some e-cash issued by her bank. To do this, Alice sends to the bank Server a zero-knowledge proof for the required amount, encrypted using their shared key. The Server verifies the proof, checks that message m has the format of an e-coin (e.g. it is equal to the message `value=$10`), debits Alice's account, signs the blinded e-coin's commitment and sends the signature to Alice.

$$A \rightarrow S : \{[\![m : s : k_s : f]\!]\}_{k_{AS}}$$

$$S \rightarrow A : \left[\langle ||m||_s \rangle^{k_s}_f\right]_{k'_S}$$

3.1.2. Showing

Alice unblinds the signature of the blinded commitment, which gives her the signature of the commitment. To spend the money at Bob's shop, she sends to Bob the signature of the commitment and the data used in the computation of the commitment, encrypted using their shared key.[1] Bob verifies the bank Server's signature and checks that the commitment is indeed computed using the data sent. He then forwards the data to the bank Server, encrypted using their shared key. The Server verifies its signature, checks again the commitment's computation, checks further that the e-coin has not been spent before (double spending) and credits Bob's account.

$$A \rightarrow B : \left\{m, s, [||m||_s]_{k'_S}\right\}_{k_{AB}}$$

$$B \rightarrow S : \left\{m, s, [||m||_s]_{k'_S}\right\}_{k_{BS}}$$

Notice that the Server does not know s, so even if Bob and the Server cooperate in an effort to disclose Alice's identity, they will fail.

3.2. Electronic voting protocol

3.2.1. Issuing

Alice wants to participate in an electronic election held by a trusted voting Server. To do this, Alice sends to the Server a zero-knowledge proof for each of the two possible votes of this election, encrypted using their shared key. The Server verifies the proofs, checks that Alice is eligible for voting and that messages m_1 and m_2 represent the possible votes,

[1] The shared key is associated with the minimum information required to complete the purchase, for example the (anonymous) post office box number of Alice.

signs the blind commitment of each vote and sends the signatures back to Alice.

$$A \rightarrow S : \{[\![m_1 : s_1 : k_S : f_1]\!], [\![m_2 : s_2 : k_S : f_2]\!]\}_{k_{AS}}$$

$$S \rightarrow A : \left[\langle ||m_1||_{s_1}\rangle_{f_1}^{k_S}\right]_{k'_S}, \left[\langle ||m_2||_{s_2}\rangle_{f_2}^{k_S}\right]_{k'_S}$$

3.2.2. Showing

Alice unblinds the signatures of the blinded commitments, which gives her the signatures of the commitments. She can now choose the commitment of the vote she wishes to cast, and send the corresponding signature to the Server via an anonymous channel, together with the data used in the computation of the commitment (one of which is the vote's representation). The Server verifies its own signature and after checking that the commitment is indeed computed using the data sent, it accepts Alice's vote.

$$A \rightarrow S : m_a, s_a, \left[||m_a||_{s_a}\right]_{k'_S}$$

Notice that the Server has no way of linking s_a to Alice.

3.3. Selective disclosure protocol

3.3.1. Issuing

Alice wants to demonstrate to Bob a certain attribute about herself, but she does not want to disclose to him any other information about her. To do so, Alice contacts the Server of a trusted certificate authority which issues selective disclosure certificates, and sends it a zero-knowledge proof about her attribute, encrypted using their shared key. The Server verifies the proof, checks that m is a proper message certifying a true attribute of Alice, signs the blind commitment and sends the signature back to Alice.

$$A \rightarrow S : \{[\![m : s : k_s : f]\!]\}_{k_{AS}}$$

$$S \rightarrow A : \left[\langle ||m||_s\rangle_f^{k_S}\right]_{k'_S}$$

3.3.2. Showing

Alice unblinds the signature of the blinded commitment, which gives her the signature of the commitment. She can now send this signature to Bob, together with the data used in the computation of the commitment, via an anonymous channel. Bob verifies the Server's signature and after checking that the commitment is indeed computed using the data sent, it accepts the attribute in message m.

$$A \rightarrow B : m, s, [||m||_s]_{k'_S}$$

Notice that the Server does not know s, so even if Bob and the Server cooperate in an effort to disclose Alice's identity, they will fail.

4. Types

Typed MSR employs types to enforce basic well-formedness conditions (e.g. that only keys can be used to encrypt a message), as well as to provide a statically checkable way to ascertain desired properties (e.g. that no principal can grab a key he is not entitled to access).

4.1. Overview of types in Typed MSR

The typing of Typed MSR is based on the notion of *dependent product types with subsorting* [1] and the basic types used are summarized in the following grammar:

Types: $\tau ::=$ principal (*Principals*)
| nonce (*Nonces*)
| shK $A\,B$ (*Shared keys*)
| pubK A (*Public keys*)
| privK k (*Private keys*)
| msg (*Messages*)

We will use the letter τ (variously decorated) to range over types. The types principal and nonce are used to classify principals and nonces, respectively. The type shK A B is used to classify the keys shared between A and B. The type pubK A is used to classify the public keys of A. The type prvK k is used to classify the private key that corresponds to the public key k. Finally, the type msg is used to classify generic messages, which include raw data, but also all the other stated types.

The notion of dependent product types with subsorting we mentioned above accommodates our need of having multiple classifications within a hierarchy. For example, everything that is of type nonce, is also of type msg—but the inverse is not true. Therefore, we say that nonce is a *subsort* of msg. We will use the notation $\tau :: \tau'$ to state that τ is a subsort of τ'. The following rules can now be presented:

$$\frac{}{\mathsf{principal} :: \mathsf{msg}} \qquad \frac{}{\mathsf{nonce} :: \mathsf{msg}} \qquad \frac{}{\mathsf{shK}\ A\ B :: \mathsf{msg}}$$

$$\frac{}{\mathsf{pubK}\ A :: \mathsf{msg}} \qquad \frac{}{\mathsf{privK}\ k :: \mathsf{msg}}$$

4.2. Adding types for privacy-preserving protocols

To better cope with privacy-preserving protocols, we add types for tractable, semitractable and intractable messages:

$$\begin{aligned} \textit{Types} : \tau :: = \ldots \quad & (\textit{see above}) \\ \mid \ \mathsf{tract} \quad & (\textit{Tractable messages}) \\ \mid \ \mathsf{semitract} \quad & (\textit{Semitractable messages}) \\ \mid \ \mathsf{intract} \quad & (\textit{Intractable messages}) \end{aligned}$$

These three types are used to classify messages according to their commonness. In other words, they qualitatively classify the number of possible values a message can have.

The type tract is used to classify messages that are very common. Because of the tractable number of their possible values, we consider that an intruder (regardless of whether these messages are publicly known or not) is able to find them out by successfully employing a brute-force dictionary attack on them. On the other hand, if a principal reveals the same (tractable) message in more than one protocol or subprotocol execution, the intruder will not be able to link these executions together (at least not because of this particular message). Therefore, this classification isolates pieces of information on the *secrecy* of which it is erroneous to base the correctness of a protocol, but on the *anonymity* of which it is safe to do so.

The type intract is used to classify messages that are extremely uncommon. These are pieces of information on the secrecy of which it is safe to base the correctness of a protocol, but on the anonymity of which it is certainly erroneous to do so.

The type semitract is used to classify messages that are common enough to be considered realistic candidates for brute-force dictionary attacks, but not common enough to be considered anonymous. It is not safe to base the correctness of a protocol either on the secrecy of such pieces of information, nor on their anonymity.

We will now classify each of the standard types according to their tractability. Private keys, shared keys and nonces should be regarded as intractable. Principals should be regarded as semitractable: we should not base the correctness of protocols on the number of available principals. Public keys should also be regarded as semitractable for the same reason. Notice that this classification conveniently enforces that everyone has access to public keys. The following rules can now be presented:

$$\frac{}{\mathsf{principal} :: \mathsf{semitract}} \qquad \frac{}{\mathsf{nonce} :: \mathsf{intract}} \qquad \frac{}{\mathsf{shK}\ A\ B :: \mathsf{intract}}$$

$$\frac{}{\mathsf{pubK}\ A :: \mathsf{semitract}} \qquad \frac{}{\mathsf{privK}\ k :: \mathsf{intract}}$$

The classification of messages that are not keys, nor nonces, nor principals will be dealt with by *signatures*, which are described in Section 4.3. To complete our subsorting rules, we add rules that classify tractable, semitractable and intractable messages as messages:

$$\frac{}{\mathsf{tract} :: \mathsf{msg}} \quad \frac{}{\mathsf{semitract} :: \mathsf{msg}} \quad \frac{}{\mathsf{intract} :: \mathsf{msg}}$$

4.3. Signatures

Typed MSR has typing rules that check whether an expression built according to the syntax of messages can be considered a ground message. These rules systematically reduce the validity of a composite message to the validity of its sub-messages. In this way, it all comes down to what the types of atomic messages are. Typed MSR uses signatures to achieve independence of rules from atomic messages. A signature is a finite sequence of declarations that map atomic messages to their type. The grammar of a signature is given below:

$$\begin{aligned} \textit{Signatures: } \Sigma ::= {} & \quad (\textit{Empty signature}) \\ & \mid \Sigma, a : \tau \quad (\textit{Atomic message declaration}) \end{aligned}$$

For our extended type system, we will need two signatures. Signature Σ will map atomic messages to one of the standard types, and signature Γ will map them to one of the extended types, i.e. classify them into tractable, semitractable or intractable. We will write $t :_{\Sigma} \tau$ to say that message t has type τ in signature Σ, and we will write $t :_{\Gamma} \tau'$ to say that message t has type τ' in signature Γ. Hence the following two rules:

$$\frac{}{(\Sigma, \alpha : \tau, \Sigma') \vdash \alpha :_{\Sigma} \tau} \quad \frac{}{(\Gamma, \alpha : \tau, \Gamma') \vdash \alpha :_{\Gamma} \tau}$$

4.4. Type rules for message constructors

We will now introduce type rules for all the message constructors presented in Sections 2.1 and 2.2 that use the new types introduced in Section 4.2 in order to further check the groundness of messages.

4.4.1. Concatenation

The concatenation of two messages of the same type will yield a message of that type.

$$\frac{\Gamma \vdash t_1 : \tau \;\; \Gamma \vdash t_2 : \tau}{\Gamma \vdash t_1, t_2 : \tau}$$

The concatenation of two messages of different types will yield a message of the least tractable type among the types of the original messages.

$$\frac{\Gamma \vdash t_1 : \mathsf{tract} \;\; \Gamma \vdash t_2 : \mathsf{semitract}}{\Gamma \vdash t_1, t_2 : \mathsf{semitract} \;\; \Gamma \vdash t_2, t_1 : \mathsf{semitract}}$$

$$\frac{\Gamma \vdash t_1 : \mathsf{tract} \;\; \Gamma \vdash t_2 : \mathsf{intract}}{\Gamma \vdash t_1, t_2 : \mathsf{intract} \;\; \Gamma \vdash t_2, t_1 : \mathsf{intract}}$$

$$\frac{\Gamma \vdash t_1 : \mathsf{semitract} \;\; \Gamma \vdash t_2 : \mathsf{intract}}{\Gamma \vdash t_1, t_2 : \mathsf{intract} \;\; \Gamma \vdash t_2, t_1 : \mathsf{intract}}$$

Note that in Typed MSR concatenated messages can be taken apart.

4.4.2. Symmetric-key and asymmetric-key encryption

The tractability of the resulting ciphertext is defined to be the same as the tractability of the plaintext.

$$\frac{\Gamma \vdash t : \tau \quad \Sigma \vdash k : \mathsf{shK}\ A\,B}{\Gamma \vdash \{t\}_k : \tau} \qquad \frac{\Gamma \vdash t : \tau \quad \Sigma \vdash k : \mathsf{pubK}\ A}{\Gamma \vdash \{\{t\}\}_k : \tau}$$

The implication is that the ciphertext of a tractable or semitractable message can now be cryptanalyzed by an intruder and the original plaintext will instantly be made available. The aim is to enforce that only intractable messages are enciphered, so that known-plaintext attacks are not possible. One way to make a tractable or semitractable message into an intractable one is to concatenate it with a nonce (see rules for concatenation).

We believe that these type rules are fully in line with the black-box view on cryptography that the Dolev-Yao abstraction adopts. The type rules only enforce a safer use of cryptography; they do not poison the abstraction with low-level details.

4.4.3. Digital signature

Similar considerations apply to digital signatures.

$$\frac{\Gamma \vdash t : \tau \quad \Sigma \vdash k' : \mathsf{privK}\ k}{\Gamma \vdash [t]_{k'} : \tau}$$

4.4.4. Commitment

Commitments may be considered to be intractable because of the nonce (salt value) used in the calculation.

$$\frac{\Gamma \vdash t : \tau \quad \Sigma \vdash n_s : \mathsf{nonce}}{\Gamma \vdash \|t\|_{n_s} : \mathsf{intract}}$$

4.4.5. Blind signatures

Blind signatures may be considered to be intractable because of the nonce (blinding factor) used in the calculation.

$$\frac{\Gamma \vdash t : \tau \quad \Sigma \vdash k : \mathsf{pubK}\ A \quad \Sigma \vdash n_f : \mathsf{nonce}}{\Gamma \vdash \langle t \rangle^{k}_{n_f} : \mathsf{intract}}$$

4.4.6. Zero-knowledge proofs

The zero-knowledge proof itself can be considered to be intractable, as two nonces are used in its calculation (a salt value and a blinding factor). However, we require that the underlying message of a zero-knowledge proof is tractable in order to enforce anonymity, and thus protect privacy. Consider for example that, if e-coins were issued at any possible denomination, the bank would be able to identify the spender in most cases.

$$\frac{\Gamma \vdash t : \mathsf{tract} \quad \Sigma \vdash n_s : \mathsf{nonce} \quad \Sigma \vdash k : \mathsf{pubK}\ A \quad \Sigma \vdash n_f : \mathsf{nonce}}{\Gamma \vdash [\![t : n_s : k : n_f]\!] : \mathsf{intract}}$$

5. The Dolev-Yao intruder

The Dolev-Yao abstraction [5] assumes that elementary data, such as keys or nonces, are atomic rather than strings of bits, and that the operations needed to assemble messages, such as concatenation or encryption, are pure constructors in an initial algebra. Typed MSR fits very well in this abstraction: elementary data are indeed atomic and messages are constructed solely by message constructors.

In this section, we present a version of the Dolev-Yao intruder which is useful in discovering more types of attacks in privacy-preserving (as well as other types of) protocols. The rules that formally describe the new capabilities of the intruder are represented in the same way as in [2], i.e. using the format shown in the following diagram:

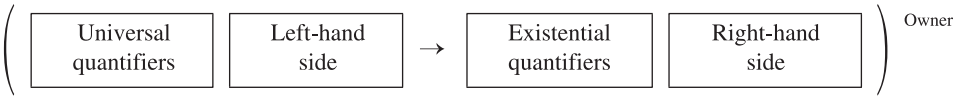

It has been proved [9] that there is no point in considering more than one Dolev-Yao intruder in any given system. Therefore, we can select a principal, I say, to represent the Dolev-Yao intruder. Furthermore, we associate I with an MSR memory predicate $\mathsf{M}_{\mathsf{I}}(\_)$, whose single argument can hold a message, to enable I to store data out of sight from other principals.

5.1. Standard version of the Dolev-Yao intruder

The standard version of the Dolev-Yao intruder can do any combination of the following operations:

- Intercept and learn messages
- Make copies of known messages
- Transmit known messages
- Decompose known (concatenated) messages
- Concatenate known messages
- Decipher encrypted messages if he knows the keys
- Encrypt known messages with known keys
- Sign messages with known keys
- Access public information
- Generate fresh data

The interested reader can refer to Ref. [2] for the formal specification of these operations in Typed MSR.

5.2. Extended version of the Dolev-Yao intruder

The version of the intruder that is presented here is an extended version in two ways. Firstly, one of the intruder's standard operations will be generalized in line with the new types introduced in Section 4.2. More specifically, we will replace the last operation, i.e. the intruder's ability to generate fresh data, with two new operations: the ability to generate fresh intractable data, and the ability to guess tractable and semitractable data. The intruder will be able either to guess the exact message required for his/her attack if this is possible, or to generate a fresh message of the required type otherwise. Secondly, the intruder will now be able to handle messages constructed using the message constructors introduced in Section 2.2.

We will now formally specify the new operations in Typed MSR.

5.2.1. Generate fresh intractable data

The intruder may generate fresh nonces, fresh private keys, fresh shared keys, as well as other intractable messages.

$$(\cdot \rightarrow \exists t :_{\Gamma} \mathsf{intract}.\ \mathsf{M}_{\mathsf{I}}(t))^{\mathsf{I}}$$

5.2.2. Guess tractable and semitractable data

The intruder may guess or get access to public keys, principals, as well as other tractable or semitractable messages.

$$(\forall t :_\Gamma \mathsf{tract}.\ \cdot\ \rightarrow\ \mathsf{M_I}(t))^{\mathsf{I}} \quad (\forall t :_\Gamma \mathsf{semitract}.\ \cdot\ \rightarrow\ \mathsf{M_I}(t))^{\mathsf{I}}$$

Notice that this rule can be used together with the previous one to allow the intruder to generate a key-pair by first generating a fresh private key, and then by 'guessing' the corresponding public key. However, the intruder is not able to guess the private keys of other principals.

5.2.3. Blind messages

The intruder may blind a message given a public key and a blinding factor (nonce).

$$\left(\begin{array}{l}\forall t :_\Sigma \mathsf{msg}.\\ \forall A :_\Sigma \mathsf{principal}.\\ \forall k :_\Sigma \mathsf{pubK}\ A.\\ \forall n :_\Sigma \mathsf{nonce}.\end{array}\ \begin{array}{l}\mathsf{M_I}(t)\\ \mathsf{M_I}(k)\\ \mathsf{M_I}(n)\end{array}\ \rightarrow\ \mathsf{M_I}\left(\langle t\rangle_n^k\right)\right)^{\mathsf{I}}$$

5.2.4. Unblind messages

The intruder may unblind a (blinded) message given the blinding factor (nonce).

$$\left(\begin{array}{l}\forall t :_\Sigma \mathsf{msg}.\\ \forall A :_\Sigma \mathsf{principal}.\\ \forall k :_\Sigma \mathsf{pubK}\ A.\\ \forall n :_\Sigma \mathsf{nonce}.\end{array}\ \begin{array}{l}\mathsf{M_I}\left(\langle t\rangle_n^k\right)\\ \mathsf{M_I}(n)\end{array}\ \rightarrow\ \mathsf{M_I}(t)\right)^{\mathsf{I}}$$

5.2.5. Unblind signatures

The intruder may unblind a (blinded) signature given the blinding factor (nonce), if the public key used in the blinding corresponds to the private key used in the signing.

$$\left(\begin{array}{l}\forall t :_\Sigma \mathsf{msg}.\\ \forall A :_\Sigma \mathsf{principal}.\\ \forall k :_\Sigma \mathsf{pubK}\ A.\\ \forall k' :_\Sigma \mathsf{privK}\ k.\\ \forall n :_\Sigma \mathsf{nonce}.\end{array}\ \begin{array}{l}\mathsf{M_I}\left(\left[\langle t\rangle_n^k\right]_{k'}\right)\\ \mathsf{M_I}(n)\end{array}\ \rightarrow\ \mathsf{M_I}\left([t]_{k'}\right)\right)^{\mathsf{I}}$$

5.2.6. Commit to a message

The intruder may commit to a message given a salt value (nonce).

$$\left(\begin{array}{ll}\forall t :_\Sigma \mathsf{msg}. & \mathsf{M_I}(t)\\ \forall n :_\Sigma \mathsf{nonce}. & \mathsf{M_I}(n)\end{array}\ \rightarrow\ \mathsf{M_I}(\|t\|_n)\right)^{\mathsf{I}}$$

5.2.7. Generate a zero-knowledge proof

The intruder may generate a zero-knowledge proof given a message, a salt value (nonce), a public key and a blinding factor (nonce).

$$\left(\begin{array}{lll} \forall t :_{\Sigma} \mathsf{msg}. & \mathsf{M_I}(t) & \\ \forall n_s :_{\Sigma} \mathsf{nonce}. & \mathsf{M_I}(n_s) & \\ \forall A :_{\Sigma} \mathsf{principal}. & \mathsf{M_I}(k) & \rightarrow \mathsf{M_I}(\llbracket t : n_s : k : n_f \rrbracket) \\ \forall k :_{\Sigma} \mathsf{pubK}\ A. & \mathsf{M_I}(n_f) & \\ \forall n_f :_{\Sigma} \mathsf{nonce}. & & \end{array}\right)^{\mathsf{I}}$$

5.2.8. Observe a zero-knowledge proof

The intruder will get the same information as anyone else who observes the zero-knowledge proof (see Section 2.2).

$$\left(\begin{array}{lll} \forall t :_{\Sigma} \mathsf{msg}. & & \\ \forall n_s :_{\Sigma} \mathsf{nonce}. & & \\ \forall A :_{\Sigma} \mathsf{principal}. & \mathsf{M_I}(\llbracket t : n_s : k : n_f \rrbracket) & \rightarrow \begin{array}{l} \mathsf{M_I}(t) \\ \mathsf{M_I}\left(\langle \| t \|_{n_s} \rangle^{k}_{n_f}\right) \end{array} \\ \forall k :_{\Sigma} \mathsf{pubK}\ A. & & \\ \forall n_f :_{\Sigma} \mathsf{nonce}. & & \end{array}\right)^{\mathsf{I}}$$

5.3. Linking protocol executions

A typical requirement for privacy-preserving protocols is that it should not be possible to link protocol or subprotocol executions together. Informally, when we say that two executions of a protocol cannot be linked to a given principal (usually the one whose privacy the protocol is supposed to protect), we mean that it is not possible for the Dolev-Yao intruder to deduce whether the same principal participated in both executions, even if the Dolev-Yao intruder manages to overtake all the other principals and get hold of their long-term secrets. Indeed, the example protocols of Section 3 are designed so that the execution of the issuing subprotocol and the execution of the showing subprotocol cannot be linked to Alice.

A protocol that allows the Dolev-Yao intruder to link two protocol executions together is not necessarily vulnerable to an attack, so we need to add extra rules that express that such a property is undesirable in a protocol. To this end, we introduce the Dolev-Yao intruder's eavesdropping memory predicate, $\mathsf{E_I}(\_)$, whose single argument can hold a message. The eavesdropping memory predicate is associated with the following intruder's operations.

5.3.1. Record messages

In general, the Dolev-Yao Intruder records messages by using three of his allowed operations: First he intercepts the messages (removing them from the network), then he makes a copy of them, and finally he transmits the copies (keeping a single instance of each message he intercepted). The eavesdropping memory predicate must not be allowed to make copies of messages (see next operation for an explanation), so it must have a rule that will allow it to record messages in one shot, as follows:

$$\left(\forall t :_{\Sigma} \mathsf{msg}.\, \mathsf{N}(t) \rightarrow \begin{array}{l} \mathsf{N}(t) \\ \mathsf{E_I}(t) \end{array}\right)^{\mathsf{I}}$$

Note that in Typed MSR, $\mathsf{N}(\_)$ is the network predicate.

5.3.2. Link two protocol executions together

We assume that two protocol executions can be linked together only because the same intractable or semitractable message is being transmitted in both of them. By definition, tractable messages cannot be used for linking, as their limited range of possible values guarantees that they will be in frequent use by most principals. We will use the letter $\mathcal{L}$ to represent a successful linking.

$$\left(\forall t :_{\Gamma} \text{intract}.\ \frac{\mathsf{E}_{\mathsf{I}}(t)}{\mathsf{E}_{\mathsf{I}}(t)} \rightarrow \mathcal{L}\right)^{\mathsf{I}} \quad \left(\forall t :_{\Gamma} \text{semitract}.\ \frac{\mathsf{E}_{\mathsf{I}}(t)}{\mathsf{E}_{\mathsf{I}}(t)} \rightarrow \mathcal{L}\right)^{\mathsf{I}}$$

Note that these rules would not stand if the eavesdropping memory predicate was allowed to make copies of the messages it records. Note further that these rules are too restrictive in the sense that: (i) they apply to all principals, not just the one whose privacy the protocol must preserve, and (ii) they apply even within the same protocol or subprotocol run. However, we believe that these restrictions would not pose a problem to real-world protocols. For example, there is no reason two principals should exchange the same intractable or semitractable message more than once in the same protocol or subprotocol execution, even if this poses no privacy risk.

6. Summary and conclusions

In this paper, we have presented an extension of Typed MSR that makes it more suitable for the specification of privacy-preserving protocols. The introduced non-interactive message constructors for blind signatures, commitments and zero-knowledge proofs make the standard language rich enough to specify protocols such as electronic cash, electronic voting and selective disclosure protocols. The introduced type rules make the standard language more capable of statically checking for desired properties in privacy-preserving as well as other types of protocols. More specifically, the introduced types can be used in the specification of protocols in order to statically check against attacks on privacy, brute-force dictionary attacks and known-plaintext attacks. Finally, the introduced version of the Dolev-Yao intruder creates a formal framework on which attacks on privacy-preserving protocols may be attempted.

Further work will include the development of a stricter and richer type system and the formal specification of real-world privacy-preserving protocols in the extended language.

References

[1] D. Aspinall, A. Compagnoni, Subtyping dependent types, in: E. Clarke (Ed.), Proceedings of the 11th Annual Symposium on Logic in Computer Science, IEEE Computer Society Press, 1996 July, pp. 86–97.

[2] Iliano Cervesato, Typed multiset rewriting specifications of security protocols, in: A. Seda (Ed.), First Irish Conference on the Mathematical Foundations of Computer Science and Information Technology—MFCSIT'00, ENTCS, vol. 40, Elsevier, Cork, Ireland, 2000 July 19–21, pp. 1–43.

[3] Iliano Cervesato, Typed MSR: syntax and examples, in: V.I. Gorodetski, V.A. Skormin, L.J. Popyack (Eds.), First International Workshop on Mathematical Methods, Models and Architectures for Computer Networks Security—MMM'01, LNCS, vol. 2052, Springer-Verlag, St. Petersburg, Russia, 2001 May 21–23, pp. 159–177.

[4] David Chaum, Security without identification: transaction systems to make big brother obsolete, Communications of the Association for Computing Machinery 28 (10) (1985 October) 1030–1044.

[5] D. Dolev, A.C. Yao, On the security of public key protocols, IEEE Transactions on Information Theory 2 (29) (1983) 198–208.

[6] Jason E. Holt, Kent E. Seamons, Selective disclosure credential sets, http://citeseer.nj.nec.com/541329.html, 2002.

[7] Alfred J. Menezes, Paul C. van Oorschot, Scott A. Vanstone, Handbook of Applied Cryptography, CRC Press, 1997.

[8] Bruce Schneier, Applied Cryptography, John Wiley and Sons, 1996.

[9] Paul Syverson, Catherine Meadows, Iliano Cervesato, Dolev-Yao is no better than Machiavelli, in: P. Degano (Eds.), First Workshop on Issues in the Theory of Security – WITS'00, 2000 July, pp. 87–92.

Theodoros Balopoulos was born in Greece in 1978. He holds a BA in Computer Science from the University of Cambridge, UK. Currently he is a PhD candidate at the Department of Information and Communications System Engineering, University of the Aegean, Greece. His research interests include information security, security protocols and privacy. His published scientific work includes two conference papers.

Stefanos Gritzalis was born in Greece in 1961. He holds a BSc in Physics, an MSc in Electronic Automation, and a PhD in Informatics all from the University of Athens, Greece. Currently he is an Associate Professor at the Department of Information and Communication Systems Engineering, University of the Aegean, Greece, and a Director of the Info-Sec-Lab. He has been involved in more than 30 national and CEC funded R and D projects in the areas of Information and Communication Systems. His published scientific work includes six books (in Greek) on Information and Communication Technologies topics, and more than 70 journal and national and international conference papers. The focus of these publications is on Information and Communication Systems Security. He has served in program and organizing committees of national and international conferences on Informatics and is a reviewer for several scientific journals.

Sokratis K. Katsikas was born in Greece in 1960. He received the Diploma in Electrical Engineering degree from the University of Patras, Greece, the MSc in Electrical and Computer Engineering from the University of Massachusetts at Amherst, USA, and the PhD in Computer Engineering from the University of Patras, Greece. He now is Professor at the Department of Information and Communication Systems Engineering and Rector of the University of the Aegean, Greece. He has authored or co-authored more than 140 technical papers and conference presentations in his areas of research interest, which include information and communication systems security, estimation theory, adaptive control, and artificial intelligence. He has served on steering, program and organizing committees of international conferences Informatics and is a reviewer for several scientific journals.