



ΠΑΝΕΠΙΣΤΗΜΙΟ ΑΙΓΑΙΟΥ
ΤΜΗΜΑ ΜΗΧΑΝΙΚΩΝ ΠΛΗΡΟΦΟΡΙΑΚΩΝ
ΚΑΙ ΕΠΙΚΟΙΝΩΝΙΑΚΩΝ ΣΥΣΤΗΜΑΤΩΝ

ΔΙΑΤΡΙΒΗ

για την απόκτηση Διδακτορικού Διπλώματος
του Τμήματος Μηχανικών Πληροφοριακών και
Επικοινωνιακών Συστημάτων

Θεόδωρου Μπαλόπουλου

ΠΡΟΔΙΑΓΡΑΦΗ ΚΑΙ ΥΛΟΠΟΙΗΣΗ
ΚΡΥΠΤΟΓΡΑΦΙΚΩΝ ΠΡΩΤΟΚΟΛΛΩΝ ΑΣΦΑΛΕΙΑΣ
ΜΕ ΑΠΑΙΤΗΣΕΙΣ ΔΙΑΣΦΑΛΙΣΗΣ ΙΔΙΩΤΙΚΟΤΗΤΑΣ
(Specifying and Implementing Privacy-Preserving Cryptographic Protocols)

Συμβουλευτική Επιτροπή:

Πρόεδρος:

Στέφανος Γκριτζαλης
Αναπληρωτής Καθηγητής
Πανεπιστημίου Αιγαίου

Μέλη:

Σωκράτης Κ. Κάτσικας
Καθηγητής
Πανεπιστημίου Πειραιώς

Κωνσταντίνος Λαμπρινουδάκης
Επίκουρος Καθηγητής
Πανεπιστημίου Αιγαίου

Εξεταστική Επιτροπή:

Πρόεδρος:

Στέφανος Γκριτζαλης
Αναπληρωτής Καθηγητής
Πανεπιστημίου Αιγαίου

Μέλη:

Σωκράτης Κ. Κάτσικας
Καθηγητής
Πανεπιστημίου Πειραιώς

Κωνσταντίνος Λαμπρινουδάκης
Επίκουρος Καθηγητής
Πανεπιστημίου Αιγαίου

Βασίλειος Χρυσικόπουλος
Καθηγητής
Ιονίου Πανεπιστημίου

Ιωάννης Δαρζέντας
Καθηγητής
Πανεπιστημίου Αιγαίου

Δημήτριος Φωτάκης
Επίκουρος Καθηγητής
Πανεπιστημίου Αιγαίου

Σπυρίδων Κωτσάκης
Καθηγητής
Πανεπιστημίου Αιγαίου

Specifying and Implementing Privacy-Preserving Cryptographic Protocols

Theodoros Balopoulos

October 26, 2007

Περίληψη

Η διατριβή αυτή ασχολείται με την προδιαγραφή και υλοποίηση πρωτοκόλλων ασφάλειας με απαιτήσεις διασφάλισης ιδιωτικότητας, όπως για παράδειγμα τα πρωτόκολλα ηλεκτρονικών μετρητών, ηλεκτρονικής ψηφοφορίας και επιλεκτικής αποκάλυψης δεδομένων. Ο στόχος, όσον αφορά την προδιαγραφή τους, είναι αυτή να γίνει με τυπική μέθοδο (formal method), και, όσον αφορά την υλοποίησή τους, να βασίζεται στην προδιαγραφή τους και να διασφαλίζει τις περί ιδιωτικότητας απαιτήσεις.

Το υπάρχον ερευνητικό έργο στη διεθνή βιβλιογραφία σε τυπικές μεθόδους δεν καλύπτει επαρκώς τα πρωτόκολλα ασφάλειας με απαιτήσεις διασφάλισης ιδιωτικότητας όσο άλλων ειδών πρωτόκολλα ασφάλειας, όπως τα πρωτόκολλα αυθεντικοποίησης. Στην παρούσα διατριβή υποστηρίζεται ότι οι λόγοι για αυτή την ανεπάρκεια μελέτης είναι οι εξής: Πρώτον, ότι τα πρωτόκολλα με απαιτήσεις διασφάλισης ιδιωτικότητας βασίζονται σε πιο εξειδικευμένη κρυπτογραφία, όπως η δέσμευση (commitment), η τυφλή υπογραφή (blind signature), η απόδειξη μηδενικής γνώσης (zero-knowledge proof), η ομομορφική κρυπτογραφία (homomorphic encryption), το mix του Chaum και το onion routing. Δεύτερον, ότι είναι απαραίτητη η διαφοροποίηση στη μοντελοποίηση της κλασικής κρυπτογραφίας (συμμετρική και ασύμμετρη κρυπτογράφηση και ψηφιακές υπογραφές) που τα πρωτόκολλα αυτά χρησιμοποιούν από κοινού με τα υπόλοιπα πρωτόκολλα.

Η διατριβή αυτή χρησιμοποιεί ως βάση τη γλώσσα προδιαγραφής πρωτοκόλλων Typed MSR [14, 15], καθώς και την προηγούμενη εργασία μας στην ίδια [10, 7, 9, 8, 6] ερευνητική κατεύθυνση και στοχεύει με τροποποιήσεις και προσθήκες να την μετατρέψει σε κατάλληλη, αφενός για την προδιαγραφή πρωτοκόλλων με απαιτήσεις διασφάλισης ιδιωτικότητας, αφετέρου για την προδιαγραφή ενός κατά Dolev-Yao επιτιθέμενου [19] σχεδιασμένου για επίθεση σε πρωτόκολλα τέτοιου είδους. Επιπλέον, χρησιμοποιεί ως βάση τη γλώσσα Jif [30, 31, 29], καθώς και την προηγούμενη εργασία μας [6] στην ίδια ερευνητική κατεύθυνση και στοχεύει να επιδείξει πως η γλώσσα αυτή, που διαθέτει σύστημα τύπων για απαιτήσεις ασφάλειας, μπορεί να χρησιμοποιηθεί με τέτοιο τρόπο ώστε οι αδυναμίες στην υλοποίηση πρωτοκόλλων ασφαλείας όσον αφορά τη συνδεσιμότητα (linkability) να μπορούν να ανιχνευτούν με ένα συνδυασμό στατικών και δυναμικών (runtime) ελέγχων.

Τα βασικά συμπεράσματα της διατριβής αυτής είναι τα ακόλουθα:

1. Προκειμένου η Typed MSR να είναι κατάλληλη για την προδιαγραφή πρωτοκόλλων ασφάλειας με απαιτήσεις ιδιωτικότητας, δεν θα πρέπει να μοντελοποιεί τη συμμετρική και την ασύμμετρη κρυπτογράφηση ως αιτιοκρατική. Μια τέτοια απλούστευση μπορεί να μη δημιουργεί προβλήματα στη μοντελοποίηση άλλων πρωτοκόλλων, αλλά οδηγεί σε ανύπαρκτες αδυναμίες διασύνδεσης στα πρωτόκολλα που μελετούμε στην παρούσα διατριβή.
2. Μπορούμε να κατασκευάσουμε υψηλού επιπέδου μοντελοποιήσεις για κρυπτογραφία πιο σύνθετη από την κλασική, όπως είναι η δέσμευση, η τυφλή υπογραφή, η απόδειξη μη-

δενικής γνώσης και η ομομορφική κρυπτογραφία.

3. Η χρήση μη διαδραστικών μοντελοποιήσεων για τις αποδείξεις μηδενικής γνώσης οδηγεί στην απλοποίηση τόσο της προδιαγραφής των πρωτοκόλλων, όσο και της μετατροπής αυτής σε υλοποίησή τους.
4. Με βάση τις προαναφερθείσες αλλαγές και προσθήκες, η Typed MSR γίνεται κατάλληλη για την προδιαγραφή πρωτοκόλλων ασφάλειας με απαιτήσεις ιδιωτικότητας, όπως δείχνει η προδιαγραφή των δύο πρωτοκόλλων ηλεκτρονικής ψηφοφορίας που περιέχονται στην παρούσα διατριβή.
5. Ένα απλό σύστημα τύπων, που χρησιμοποιείται παράλληλα με το σύστημα τύπων της Typed MSR, αποτρέπει συγκεκριμένες εσφαλμένες χρήσεις της κρυπτογραφίας που μπορεί να οδηγήσουν σε αδυναμίες συνδεσιμότητας, καθώς και να παρακολουθήσει την απειλή συνδεσιμότητας που προκύπτει από κάθε πιθανή χρήση της κρυπτογραφίας.
6. Είναι απαραίτητη η ενημέρωση του εκφρασμένου σε Typed MSR μοντέλου του κατά Dolev-Yao επιτιθέμενου με βάση τα παραπάνω, ώστε να μπορεί πλέον να επιτεθεί στα πρωτόκολλα τα οποία μελετάμε.
7. Η ενημερωμένη αυτή έκδοση του κατά Dolev-Yao επιτιθέμενου δημιουργεί ένα τυπικό (formal) περιβάλλον, στο οποίο μπορούν να εκφραστούν αδυναμίες διασύνδεσης των πρωτοκόλλων.
8. Τα παραπάνω μπορούν να αποτελέσουν τη βάση για τη χρήση της γλώσσας Jif με τέτοιο τρόπο, ώστε οι αδυναμίες στην υλοποίηση πρωτοκόλλων ασφάλειας όσον αφορά τη συνδεσιμότητα να μπορούν να ανιχνευτούν με ένα συνδυασμό στατικών και δυναμικών ελέγχων.
9. Η συνδεσιμότητα δεν είναι δυνατό να ελεγχθεί στατικά στη γενική περίπτωση, αλλά μπορεί να ελέγχεται δυναμικά κατά την εκτέλεση των πρωτοκόλλων.
10. Οι κανόνες της Typed MSR με τους οποίους παράγονται τα καινούρια μηνύματα τα οποία μπορεί να σχηματίσει ο κατά Dolev-Yao επιτιθέμενος από ένα σύνολο γνωστών μηνυμάτων χωρίζονται σε δύο κατηγορίες, ανάλογα με το αν χρησιμοποιούνται στη φάση αποδόμησης των γνωστών μηνυμάτων ή στη φάση κατασκευής των καινούριων. Συγκεκριμένα μηνύματα που βασίζονται στην κρυπτογραφία της διατριβής αυτής δεν μπορούν όμως να χωριστούν σε μία από τις δύο κατηγορίες, καθώς είναι ωφέλιμη η χρήση τους και στις δύο αυτές φάσεις.

Abstract

Formal methods are an important tool for designing and implementing secure cryptographic protocols. However, the existing work on formal methods does not cover privacy-preserving protocols as much as other types of protocols (for example, authentication protocols). Furthermore, privacy-related properties are not always easy or even possible to prove statically, but need to be checked dynamically during the protocol’s execution. This thesis: (i) proposes abstractions for some (relatively) complex cryptographic primitives used in privacy-preserving protocols, and uses these abstractions to develop suitable message constructors and a linkability-oriented type system for Typed MSR (a strongly typed specification language for security protocols), and (ii) demonstrates how these typed message constructors can be implemented in Jif (a security-oriented extension of a subset of the Java programming language dealing with information flow) in such a way that linkability vulnerabilities can be detected with a mixture of static and runtime checks.

Original Work

The symbol \diamond is used in front of a section’s title to denote that the section — in both its main body and subsections — presents original work and new results.

Contents

1	Introduction	11
1.1	Privacy Terminology	11
1.1.1	Names and Identity	11
1.1.2	Anonymity	12
1.1.3	Pseudonymity	14
1.1.4	Privacy	14
1.1.5	Trust and Reputation	15
1.1.6	Repudiation	16
1.2	Formal Methods	16
1.2.1	Formal logics	16
1.2.2	Process Calculi	18
1.2.3	Strand Spaces	19
1.2.4	Typed MSR	20
1.3	About this Thesis	20
I	Specifying Privacy-Preserving Cryptographic Protocols	21
2	The Specification Language: Typed MSR	23
2.1	Messages	24
2.2	Message Predicates	25
2.3	States	25
2.4	Types	26
2.5	Rules	26
2.6	Roles and Protocol Theories	28
2.7	⋄ A Critique of Typed MSR's Encryption	28
3	Privacy-Preserving Cryptographic Abstractions	31
3.1	⋄ Primitive Cryptographic Abstractions	31
3.1.1	Asymmetric Encryption and Digital Signatures	32
3.1.2	Commitment	32
3.1.3	Blind Signatures	32
3.1.4	Zero-Knowledge Proofs	33

3.1.5	Homomorphic Encryption	34
3.2	⋄ Derived Cryptographic Abstractions	34
3.2.1	Mixes	34
3.2.2	Onion Routing	35
4	Case Study: Specifying e-Voting Protocols	37
4.1	⋄ Protocol based on Blind Signatures and Mixes	37
4.1.1	Description	37
4.1.2	Specification	38
4.1.3	Security Analysis	38
4.2	⋄ Protocol based on Homomorphic Encryption	39
4.2.1	Description	39
4.2.2	Specification	40
4.2.3	Security Analysis	40
5	A Simple Linkability-Oriented Type System	41
5.1	⋄ Types	41
5.2	Signatures	42
5.3	⋄ Type Rules	42
6	The Dolev-Yao Intruder	45
6.1	The Standard Version	45
6.2	⋄ An Extended Version	46
6.3	⋄ Demonstrating Linkability Attacks	50
II	Implementing Privacy-Preserving Cryptographic Protocols	
	51	
7	The Implementation Language: Jif	53
7.1	The Decentralized Label Model (DLM)	53
7.2	Program Counter	53
7.3	Language Features	54
7.4	Selective Declassification	54
7.5	Class Parameterization	55
7.6	Dynamic labels	55
7.7	Handling Exceptions	55
8	Linkability-Checking in Jif	57
8.1	⋄ A Cryptographic Framework	57
8.2	⋄ Runtime Linkability-Checking	59
8.3	⋄ Message Inference	60
9	Summary and Conclusions	67

Chapter 1

Introduction

In this Chapter we introduce the necessary terminology about privacy, and give an overview of the most commonly employed formal methods for cryptographic protocols.

1.1 Privacy Terminology

In order to provide an accurate overview of this topic, it is necessary to define the basic terms. These definitions are needed in order to build more formal descriptions. This section is based on many papers in the field, but most notably on Pfitzmann and Köhntopp’s paper “Anonymity, Unobservability and Pseudonymity — A Proposal for Terminology” [34].

1.1.1 Names and Identity

Many of the concepts we examine here are in some way concerned with the obfuscation of information which relates to a principal’s identity. This information can take many forms, but the classic example is the name. The name of an individual is intended to be a unique identifier within some group that allows for that individual to be distinguished from the other members of that group. The fact that, in the increasingly large social groups in which we find ourselves, a classical name is rarely totally unique does not affect the purpose of distinguishing those around us by such labels.

When we discuss the anonymity properties of a principal, we are implicitly assuming definitions of identity. To a large extent, we can assume that a user of a system may be treated as a unique individual who performs actions that can potentially be traced by another individual. However, there is a dissociation between a user’s representation in a system and his real-world personality: multiple users can collaborate to form a single online identity, and a single user may have multiple representations online. The implications of these are not fully understood, as the assumption that a single user is linked to a single representation is almost uniformly made in the field.

1.1.2 Anonymity

Anonymity may be described at the linguistic level as the property of being nameless, or having an absence of identification. To extend this definition to a more common usage within the field, we borrow from [34]: “Anonymity is the state of being not identifiable within a set of subjects, the anonymity set.”

While this definition refers to the anonymity set, an anonymity metric which has since then fallen into disfavor, the first half of the definition expresses the main purpose of anonymity and as such is of use to us.

Anonymity is the fundamental identity hiding property, providing total removal of identifying information from its subject. As such, anonymity research comprises the vast majority of work into the field of identity hiding. This is by no means a limitation to current work in the field.

It is worth noting that any identifying information required by principals with appropriate authority can be added into a data channel within an anonymous system. As such, anonymity provides us with the choice to limit identity hiding as much or as little as desired by explicitly revealing identifying information when it is necessary.

As such, anonymity has been, and will probably remain, the focal point for research into identity hiding. Additionally, anonymity systems are based upon a small set of possible approaches, Chaum’s mix [16] being the most significant of these. It is discussed in Section 3.2.1. The most active topic of research into identity hiding is therefore the finer details of the various subtle variations of these basic ideas.

Despite this focus on anonymous systems, total anonymity is very much a two-edged sword. For certain forms of application, such as posting to mailing lists or accessing the world wide web, anonymity can be a highly desirable goal. Other systems, however, suffer greatly if there is no possibility of tracking identities. It is not good policy, for example, to entrust your life savings to an anonymous user claiming to be a reputable bank. Indeed, if you were to do so and were yourself anonymous it would be impossible in the future to track down and recover your savings even if you did happen to have stumbled across an honest anonymous banker. For this reason pseudonymous communication, which provides a certain amount of information associated with an identity, is required for a number of practical identity hiding systems.

Sender and Recipient Anonymity

We have defined anonymity as the property of being nameless, given certain assumptions concerning the meaning of a name. However, this namelessness is of use only in the context a communicating system. Communication requires two participants: a sender and a recipient, where either participant may actually be a group of individuals. To whom does anonymity apply?

There are three cases. We may ensure the anonymity of the initiator of some communication, but leave the recipient’s identity open to the world. Conversely it is possible for a sender to make available their identity, but to ensure that the recipient of their message

remains unknown. This provides a very different problem to sender anonymity and is also less focused upon. Finally, it may be desirable for both end-points of communication to be hidden from observation, by outside observers or from each other. Each of these forms of anonymity has its own set of applications, design problems and potential attacks.

Connection-Based and Message-Based Anonymity

Another important distinction in anonymity is the form of the communication which takes place between the sender and the recipient. The majority of data transmitted over the Internet takes place using the connection-based Transmission Control Protocol (TCP). However, at a higher level than the underlying protocol much of the communication that takes place, especially email, is based upon the messages themselves.

The main body of the current literature focuses strongly on message-based anonymity, which is an easier problem than that of anonymity within connection-based systems. There are, however, some applications (such as remote terminal sessions) which simply cannot be performed effectively using a message-based approach. For these systems, it is necessary to use a connection-based approach.

In recent literature, a number of approaches to connection-based anonymity have been proposed. By far the most popular of these is the onion routing system presented by Goldschlag, Reed and Syverson [22]. Onion routing is discussed in Section 3.2.2. The main concern with such systems is the requirement for low-latency connections and bandwidth restrictions. Message-based systems inherently cause delays between data transfer that are unacceptable in connection-based systems.

The Anonymity Set

The traditional method of quantifying anonymity is to utilize the cardinality of the set of all participants who could have performed an action. This approach was proposed by Chaum in his paper “Untraceable electronic mail, return addresses, and digital pseudonyms” [16], where he utilized the anonymity set quantification to analyze the anonymity provided by a mix. The point is simple: the larger the size of the set which could have performed an action, the stronger the anonymity provided by the system.

However, this quantification, while certainly of some value, is not considered ideal anymore. The most critical of the deficiencies of the anonymity set is that it assumes a uniform distribution of probabilities across the set of participants. This assumption is rather naive for a group of heterogenous users [37]. In response to this issue, a number of alternative quantification methods have been proposed which seek to deal with both this and other problems inherent in the anonymity set. These methods include the work of Serjantov and Danezis [36], as well as the work of Diaz [18], both of which describe the level of anonymity provided by a system through an information theoretic approach, making the entropy of a set of users the defining factor.

Unlinkability

An important underlying component of anonymity systems is the property of unlinkability. Pfizmann and Kohntopp [34], in reasoning about anonymity systems, propose a viewpoint defined by a set of subjects sending messages to a set of recipients. In this setting, the critical concept is an item of interest, defined as the sending or receiving of a message.

The desirable property of an anonymity system is therefore that items of interest are unlinkable to any principal in the system, and no principal in the system can be linked to a specific item of interest. This provides a basic definition of anonymity which, however, does not lend itself to any form of quantification.

However, based on this definition of linkability, we introduce a qualitative classification of messages with respect to their effect on the unlinkability of the sending principal in Chapter 5 of this thesis.

1.1.3 Pseudonymity

Pseudonymity, in terms of online systems, is achieved by having users associated with an at least semi-persistent identifier. The purpose of this is to allow types of transaction that rely on user history and behavior, and which are not possible using a truly anonymous system. This is of particular use in systems which must provide a level of confidence for users, or which seek to rely on networks of trust between users, and thus cannot rely upon a simple one-use session identifier.

Pseudonymity can be achieved through the use of an anonymous infrastructure with user information and history stored within the explicitly transmitted data. If the system upon which communication relies is inherently anonymous, then pseudonymity becomes an easier proposition, as data can be released as chosen by the user without fear of extra information leakage from the system.

Pseudonymity may therefore be seen as a problem which exists at a higher level than anonymity. An anonymous channel may have some form of persistent user identification added which is kept secret between the sender and recipient. Pseudonymity may therefore be viewed less as a primitive construction and more as a combination of other security properties such as secrecy, anonymity and authentication.

There has been relatively little practical research into pseudonymity, perhaps due to the status of anonymity as a problem yet to be definitively solved.

1.1.4 Privacy

Privacy is a less well-defined property than the others discussed here. Nonetheless, it is what provides the motivation for almost all of the research which is undertaken in the field.

There are several definitions of privacy which have appeared over the years. Arguably the first extant definition was made by the United States Supreme Court Justice Louis Brandeis and lawyer Samuel Warren in the Harvard Law Review paper “The Right to Privacy” [12], published in 1890. The notion had, however, been discussed throughout

human history as far back as Aristotle, who defined the separate spheres of public and private life rights in his “Treatise on Government” [4].

Brandeis and Warren considered the right to privacy as a natural extension of an individual’s right to liberty, stating that liberty as a right had initially been enforced with respect to preventing physical assault. As newer business models and media coverage started to have significant effects on society, intrusion into private lives for public consumption became of concern to many, and the ideal of liberty was necessarily extended to include unfair intervention into aspects of a person’s life which could be embarrassing or dangerous if publicized.

Real interest in privacy, however, appears to have begun in the second half of the twentieth century. The United Nations Universal Declaration of Human Rights [32], finalized in 1948, embodied the right to privacy in its twelfth article: “No one shall be subjected to arbitrary interference with his privacy, family, home or correspondence, nor to attacks upon his honor and reputation. Everyone has the right to the protection of the law against such interference or attacks.”

Similarly, the International Covenant on Civil and Political Rights (ICCPR) [2], a section of the International Bill of Rights which was adopted to expand upon and clarify the Universal Declaration of Human Rights, reiterates the fundamentality of privacy as a right to which humans are entitled.

In British law, the first explicit electronic privacy legislation was the 1984 Data Protection Act [1]. This act, which has since been superseded by the 1998 Act of the same name, regulated the uses to which automatically stored data could be put. This legislation places strict limitations upon individuals which store the personal data of third parties, ensuring that any information is subject to a number of data protection principles. These principles include the right of the subject to inspect any stored data and the duty of the holder of the information to provide adequate technical protection for stored data.

The growing level of personal information stored by computers has brought privacy forward as a right which must be protected. The varying privacy legislation and regulations which are being enacted in various countries across the world are a useful step in preventing the widespread storage and dissemination of personal information. There are, however, always parties which seek to bypass these measures for a variety of purposes. Rather than regulate the storage and transmission of personal information, advocates of anonymity propose to achieve the same goals by preventing such information ever becoming known. *It is the belief in this approach which provides inspiration for the research covered in this thesis.*

1.1.5 Trust and Reputation

Trust and reputation are closely linked properties, particularly within the context of anonymity systems.

Reputation is a property which associates a level of trust with a particular user in a system. This allows for future judgements to be made on the past behavior of that user. Reputation is particularly important in commerce systems where users are required to

invest real economic interests in other users of a system. In many systems, reputation statements refer to pseudonyms.

1.1.6 Repudiation

Security protocol designers often wish for a party to prove a certain fact to another party in such a way that it cannot later be retracted or denied. This is to satisfy the necessity that an agent which negotiates a contract with another party must have some method of assurance that this party will honor the contract. This property is known as non-repudiation.

In anonymity systems, however, it may often be desirable to prove a fact to another agent at some point in time, but for this information to be unusable in creating long-term profile information. The ability to present a piece of information to an agent but for this information to be valid for no longer than the course of that single transaction prevents personal information concerning the agent from being catalogued for the purposes of future identification.

An important approach towards this form of information exchange is the zero knowledge proof [21]. This form of exchange allows for individuals to prove that they hold a certain piece of information without revealing the information itself. Whilst this notion is not intrinsically a feature of anonymity systems, it is a closely related information hiding property which is of great potential use in achieving anonymity in realistic systems. Zero knowledge proofs are discussed in more detail in Section 3.1.4.

1.2 Formal Methods

Formal methods [23] 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.

In the rest of this Section we briefly present the formal methods that formed the basis, or influenced the research covered in this thesis.

1.2.1 Formal logics

Formal logics grew out of a desire to express the logical relationship between stated concepts, and to allow for generation of new (true) statements by the application of rules to existing statements. This allows for propositions to be proved, based upon facts which are already known and basic axioms which are assumed to be true.

Formal logics provide a rigorous structure in which the truth of statements may be ascertained based upon the application of given rules and axioms, within the confines of such rules. The nature of the underlying rules differ between the various forms of formal logic, dependent upon the scope and purpose of the logic in question. Different

logics express notions of belief, knowledge, uncertainty or even ignorance within specific domains.

The application of formal logics to the analysis of security protocols was one of the first approaches taken towards the verification of such systems, with the most notable example being the BAN logic [13]. Logics have been shown to detect a range of problems with protocols whilst being reasonably easy to use. However, logics suffer from being a high level abstraction of a system, and as such may allow flaws which exist in the protocol to pass undetected.

The BAN logic

The BAN logic [13], developed by Burrows, Abadi and Needham, is possibly the most well-known of all logics which have been used for analyzing protocols.

BAN is a modal logic which is applied to authentication protocols with a view to proving their correctness. To achieve this, the logic allows for the basic assumptions and goals of some protocol to be expressed as formulas in the syntax of the logic, along with the steps taken during the running of the protocol.

When this representation has been made, deduction rules are applied which provide a logical path from the steps of the protocol to the desired goals. If this can be successfully achieved, then the goals of the protocol are held to be true.

BAN relies on a syntax centered around the belief of participants, and some simple rules which allow for the beliefs of an agent to be manipulated. The BAN logic contains many constructs and deduction rules, the entirety of which is not given here. Some basic constructs include:

- $A \triangleleft X$: A sees X , or the message X has been sent to A .
- $A \mid\sim X$: A said X at some point.
- $A \equiv X$: A believes, or has justified belief in, X .
- $A \mid\Rightarrow X$: A has jurisdiction over X .
- $\sharp(X)$: The message X is fresh (i.e. it has not been seen before).
- $\{X\}_K$: The message X encrypted with the key K .

Of note is the fresh construction $\sharp(X)$. Many logics which must deal with timing information abstract away from such details. The BAN logic solves the problem of time by distinguishing between two epochs: the past and the present. This has proven sufficient to express a large number of complex protocols.

These constructions are manipulated using deduction rules which define the behavior which results from such facts. As an example, the rule of jurisdiction, which concerns an agent's belief based upon the perceived authority of another agent can be expressed in the following way:

$$\frac{A \models B \Rightarrow X \quad A \models B \models X}{A \models X}$$

which may be read as: If A believes that B has authority over X and A believes that B believes X , then A believes X .

BAN has been applied to a number of protocols, most famously the Needham-Schroeder protocol, which forms the basis of the widespread Kerberos authentication mechanism. The BAN logic was shown to detect a flaw in this protocol which allowed for an attacker to replay information due to a particular nonce being assumed fresh. It is worth mentioning that this flaw was already known before it was discovered by the BAN logic analysis.

There have been a number of logics over the years which have sought to correct some of the flaws that have been observed in the BAN logic, such as [38]. None of these have gained quite the level of acceptance that was observed in the BAN logic, and the increasing use of process calculi and automated theorem provers has caused approaches such as these to be less examined than previously.

BAN, as a logic of authentication, may provide an interesting insight into what is required of anonymity protocols which can be seen as requiring the opposite of authentication. We wish to prove that a party with whom we are communicating is unknown to us, and therefore none of the properties which the BAN logic seeks to prove should hold.

1.2.2 Process Calculi

Process calculi provide a mathematical notation which allows for the description of communicating processes in a rigorous fashion. Their focus on communicating processes makes them promising for the expression of anonymity systems, which are by their very definition concerned with the communication between entities.

This approach to describing computation has evolved relatively recently in the field of theoretical computer science, in response to the increasing view of computers as communicating entities in larger networks rather than stand-alone machines. The possible applications of process calculi for security in these networks follow as naturally as the use of the traditional formal methods in traditional security situations.

There are two major process calculi in the literature. These are CSP, which was originally described by Hoare [24], and the π -calculus of Milner [28] which was developed from Milner's earlier calculus, CCS [27]. Both CCS and the π -calculus provide a Turing complete model of computation based upon the notion of message passing. Processes in the calculus may send and receive messages along defined channels. These messages may represent data transfer or the name of a new channel, which allows for the dynamic creation of new topologies in the system. The approaches taken by each of these calculi are briefly discussed below.

CSP

CSP, or Communicating Sequential Processes, was developed by Hoare as a method for describing communicating processes operating in parallel. The original CSP paper by

Hoare [24] presented the view that: “... input and output are basic primitives of programming and that parallel composition of communicating sequential processes is a fundamental program structuring method”. Therefore, Hoare devised his calculus around the interaction between these processes.

CSP was the first published calculus which approached the formal description of software in this fashion, although it was shortly followed by Milner’s Calculus of Communicating Systems which eventually became the π -calculus. CSP has gained a wide following in the formal methods community, which has led to many interesting developments within the language such as timing and probabilistic elements. In addition, there are a number of mature model checkers which handle CSP-based proofs of systems.

CSP has unsurprisingly been applied to the analysis of security protocols, which are intuitively applicable to a calculus which is based upon the interaction between communicating parties. CSP has even been applied to a basic analysis of dining cryptographer networks, by Schneider [35].

The π -Calculus

The π -calculus was developed as an extension of Milner’s calculus of communicating systems (CCS), which was presented in 1980 [27], shortly after the publishing of Hoare’s CSP. Both CCS and CSP went about the description of communicating processes in a similar fashion and presented the same level of expressive power.

In seeking to describe communicating processes as they exist in the networked world which emerged since the development of CCS and CSP, Milner observed that [28]: “Physical systems tend to have permanent physical links; they have fixed structure. But most systems in the informatic world are not physical; their links may be virtual or symbolic... These symbolic links can be created or destroyed on the fly...”

To solve this problem, Milner extended the basic capabilities of CCS to include mobility. Mobility adds to the π -calculus the ability for agents to both form new links with other agents and to destroy old links. An agent may therefore begin life in one area of a system and, in the course of execution, relocate to an entirely new portion of a network.

1.2.3 Strand Spaces

Strand Spaces [20] is a technique for analyzing cryptographic protocols, particularly authentication and key distribution protocols, at a high level. It uses the Dolev-Yao model of cryptographic protocols, and models the real-world cryptographic algorithms as abstract operations.

A strand is a sequence of events: it represents either an execution by a legitimate party in a security protocol, or else a sequence of actions by a penetrator. A strand space is a collection of strands, equipped with a graph structure generated by causal interaction. In this framework, protocol correctness claims may be expressed in terms of the connections between strands of different kinds.

1.2.4 Typed MSR

Typed MSR [14, 15] is a strongly typed specification language for security protocols, aiming to discover errors in their design. It was developed by Iliano Cervesato. It is discussed in detail in Chapter 2. Typed MSR formed the environment for formally expressing many of the results of the research presented in this thesis.

1.3 About this Thesis

The existing work on formal methods does not cover privacy-preserving protocols as much as others types of protocols (for example, authentication protocols). Privacy-preserving protocols, such as electronic cash, electronic voting and selective disclosure protocols, need more exotic cryptographic primitives and techniques (such as blind signatures, commitments, zero-knowledge proofs, mixes and homomorphic encryption), and require a lower-level abstraction of their primitives and techniques than other types of protocols.

Furthermore, privacy-related properties are not always easy or even possible to prove statically, but need to be checked dynamically during the protocol's execution. Although such an attempt raises the effort needed to implement the protocol, it also has the advantage that it is the actual protocol's implementation that is being verified, not a possibly flawed abstraction.

This thesis builds on the Typed MSR specification language [14, 15], as well as on our previous work on privacy-preserving protocols [10], and on Typed MSR [7, 9, 8, 6], and aims to make the language suitable for the specification of privacy-preserving protocols, as well as for the specification of a version of the Dolev-Yao intruder [19] that is designed to attack such protocols.

It also builds on the Jif security-typed programming language [30, 31, 29], as well as on our previous work [6] on Jif, and aims to demonstrate how Jif can be employed in such a way that linkability vulnerabilities in the implementations of privacy-preserving protocols can be detected with a mixture of static and runtime checks.

Part I

Specifying Privacy-Preserving Cryptographic Protocols

Chapter 2

The Specification Language: Typed MSR

The design and analysis of cryptographic protocols are notoriously complex and error-prone activities. Part of the difficulty derives from subtleties of the cryptographic primitives. Another portion is due to their deployment in distributed environments plagued by powerful and opportunistic attackers.

The Dolev-Yao model of security [23, 17] tackles the first problem by promoting an abstraction that has the effect of separating the analysis of the message flow from the validation of the underlying cryptographic operations. It assumes that elementary data such as principal names, keys and nonces are atomic rather than bit strings, and views the message formation operations (e.g. concatenation, encryption and digital signature) as symbolic combinators. The cryptographic operations are therefore assumed to be flawless.

Iliano Cervesato [14] claims that a significant source of faulty designs and contradictory analyses can be traced to shortcomings in the language used to specify protocols. The popular “usual notation” relies on the Dolev-Yao model and describes a protocol as the sequence of the messages transmitted during an expected run. Besides distracting the attention from the more dangerous unexpected runs, this description expresses fundamental assumptions and requirements about message components, the operating environment and the protocol’s goals as side remarks in natural language. This is clearly ambiguous and error-prone. Strand formalizations [19], like most modern languages, represent protocols as a collection of independent roles that communicate by exchanging message. While the reference to expected runs is dropped, the reliance of this formalism on a fair amount of natural language still makes it potentially ambiguous.

Therefore, Iliano Cervesato proposed a language based on multiset rewriting, nicknamed MSR, as a formalism for unambiguously representing authentication protocols, with the aim of studying properties such as the decidability of attack detection. The actions within a role are formulated as multiset rewrite rules, threaded together by the use of dedicated role state predicates. The nature and properties of message components are expressed in a relational manner by means of persistent information predicates and to a minor extent by typing declarations. In particular, variables that ought to be instantiated to “fresh”

objects during execution are marked with an existential quantifier (this operator can indeed be used for that purpose in logical specifications).

He then proposed a thorough redesign of MSR and established this formalism as a usable specification language for security protocols (not just authentication protocols). The major innovations include the adoption of a typing methodology that subsumes persistent information predicates, and the introduction of memory predicates and of constraints on interpreted domains that significantly widen the range of applicability of this language. He called this formalism Typed MSR [14, 15].

The type annotations of the new language, drawn from the theory of dependent types with subsorting, enable precise object classifications, for example by distinguishing keys on the basis of the principals they belong to, or in function of their intended use. The typing infrastructure can point to quite subtle errors, such as a principal trying to encrypt a message with a key that does not belong to him.

Memory predicates allow a principal to remember information across role executions. Their presence opens the doors to the specification of protocols structured as a collection of coordinated subprotocols. Memory predicates can be used to give a specification of the Dolev-Yao intruder that lies completely within the syntax of MSR roles.

2.1 Messages

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:

$$\begin{array}{lcl}
 \textit{Atomic messages:} & a & ::= & A & (\textit{Principal}) \\
 & & & | & k & (\textit{Key}) \\
 & & & | & n & (\textit{Nonce}) \\
 & & & | & m & (\textit{Raw data})
 \end{array}$$

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 that are of interest to us are those formalized by the following grammatical production:

$$\begin{array}{lcl}
 \textit{Messages:} & t & ::= & a & (\textit{Atomic messages}) \\
 & & & | & x & (\textit{Variables}) \\
 & & & | & t_1 . t_2 & (\textit{Concatenation}) \\
 & & & | & \{ t \}_k & (\textit{Symmetric encryption}) \\
 & & & | & \{ \{ t \} \}_k & (\textit{Asymmetric encryption}) \\
 & & & | & [t]_{k'} & (\textit{Digital signature})
 \end{array}$$

We use the letter t (possibly sub-scripted) to range over messages. We write A , k , n and m (possibly sub-scripted) for atomic constants or variables that are principals, keys, nonces

and raw data respectively. We 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.

To be able to later support blind signatures based on Chaum's blinding (see Chapter 3), we assume that the asymmetric encryption and digital signature message constructors are based on the RSA cryptosystem.

2.2 Message Predicates

Message predicates are the fundamental ingredient of states, defined in Section 2.3. They are atomic first-order formulas with zero or more terms as their arguments. Their definition is therefore based on the concept of message tuple, defined as an ordered sequence of terms:

$$\begin{aligned} \text{Message tuples: } \vec{t} &::= . && (\text{Empty tuple}) \\ &| t, \vec{t} && (\text{Tuple extension}) \end{aligned}$$

The predicates that can enter a state or a rewrite rule are of three kinds:

- First, the predicate $N(-)$ implements the contents of the public network in a distributed fashion: for each (ground) message t currently in transit, the state will contain a component of the form $N(t)$.
- Second, active roles rely on a number of role state predicates, generally one for each rule in them, of the form $L_l(-, \dots, -)$, where l is a unique identifying label. The arguments of this predicate record the value of the known parameters of the execution of the role up to the current point.
- Third, a principal A can store data in a private memory predicate of the form $M_A(-, \dots, -)$ that survives role termination and can be used across the execution of different roles, as long as the principal stays the same.

2.3 States

States are a fundamental concept in MSR. Indeed, they are the central constituent of the snapshots of a protocol execution. They are the objects transformed by rewrite rules to simulate message exchange and information update. Finally, together with execution traces, they are the hypothetical scenarios on which protocol analysis is based.

A state is a finite collection of ground state predicates. The syntax of states is formalized by means of the following grammar:

$$\begin{aligned} \text{States: } S &::= . && (\text{Empty state}) \\ &| S, N(t) && (\text{Extension with a network predicate}) \\ &| S, L_l(\vec{t}) && (\text{Extension with a role state predicate}) \\ &| S, M_A(\vec{t}) && (\text{Extension with a memory predicate}) \end{aligned}$$

Protocol rules (see Section 2.5) transform states. They do so by identifying a number of component predicates, removing them from the state, and adding other, usually related, state elements. The antecedent and consequent of a rewrite rule embed therefore substates. However, in order to be applicable to a wide array of states, rules usually contain variables that are instantiated at application time. This calls for a parametric notion of states and message predicates.

2.4 Types

Typed MSR makes use of 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).

The typing of Typed MSR is based on the notion of *dependent product types with subsorting* [5] and the basic types are summarized in the following grammar:

$$\begin{array}{lcl}
 \text{Types: } \tau ::= & \text{principal} & (\text{Principals}) \\
 & | \text{nonce} & (\text{Nonces}) \\
 & | \text{shK } A \ B & (\text{Shared keys}) \\
 & | \text{pubK } A & (\text{Public keys}) \\
 & | \text{privK } k & (\text{Private keys}) \\
 & | \text{msg} & (\text{Messages})
 \end{array}$$

We use the letter τ (decorated in various ways) to range over types. Types **principal** and **nonce** are used to classify principals and nonces respectively. Type **shK** $A \ B$ is used to classify the keys shared between A and B . Type **pubK** A is used to classify the RSA public key of A . Type **privK** k is used to classify the private key that correspond to the RSA public key k . Finally, 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 accommodates the 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**. In fact, all the types listed above are subsorts of **msg**. We use the notation $\tau :: \tau'$ to state that τ is a subsort of τ' .

2.5 Rules

With a slight imprecision that will be corrected as the discussion proceeds, a rule has the form $lhs \rightarrow rhs$. Rules are the basic mechanism that enables the transformation of a state into another, and therefore the simulation of protocol execution: whenever the antecedent lhs matches part of the current state, this portion may be substituted with the consequent rhs .

It is convenient to make protocol rules parametric so that the same rule can be used in a number of slightly different scenarios (e.g. without fixing interlocutors or nonces). A typical rule will therefore mention variables x_1, \dots, x_n that will be instantiated to actual terms during execution. Typed universal quantifiers can conveniently express this fact so that rules assume the form $\forall x_1 : \tau_1. \dots \forall x_n : \tau_n. (lhs \rightarrow rhs)$. This idea is more precisely captured by the following grammar:

$$\begin{aligned} \text{Rule: } r & ::= lhs \rightarrow rhs && (\text{Rule core}) \\ & | \quad \forall x : \tau. r && (\text{Parameter closure}) \end{aligned}$$

Both the right-hand side and the left-hand side of a rule embed a finite collection of parametric message predicate, some ground instance of which execution will respectively add to and retract from the current state when the rule is applied.

$$\begin{aligned} \text{Predicate sequences: } \vec{P} & ::= . && (\text{Empty predicate sequence}) \\ & | \quad P, N(t) && (\text{Extension with a network predicate}) \\ & | \quad P, L(\vec{e}) && (\text{Extension with a role state predicate}) \\ & | \quad P, M_A(\vec{t}) && (\text{Extension with a memory predicate}) \end{aligned}$$

The left-hand side, or antecedent, of a rule is a finite collection of parametric message predicates guarded by finitely many constraints on interpreted data. It is therefore given by the following grammar:

$$\begin{aligned} \text{Left-Hand sides: } lhs & ::= \vec{P} && (\text{Sequence of message predicate}) \\ & | \quad lhs, \chi && (\text{Extension with a constraint}) \end{aligned}$$

The right-hand side, or consequent, of a rule consists of a predicate sequence possibly prefixed by a finite string of fresh data declarations such as nonces or short-term keys. We rely on the existential quantification symbol to express data generation. We have the following grammar.

$$\begin{aligned} \text{Right-Hand sides: } rhs & ::= \vec{P} && (\text{Sequence of message predicate}) \\ & | \quad \exists x : \tau. rhs && (\text{Fresh data generation}) \end{aligned}$$

Rules are presented using the format shown in the following diagram:

$$\left(\begin{array}{|c|} \hline \text{Universal} \\ \text{quantifiers} \\ \hline \end{array} \begin{array}{|c|} \hline \text{Left-hand} \\ \text{side} \\ \hline \end{array} \rightarrow \begin{array}{|c|} \hline \text{Existential} \\ \text{quantifiers} \\ \hline \end{array} \begin{array}{|c|} \hline \text{Right-hand} \\ \text{side} \\ \hline \end{array} \right)^{\text{Owner}}$$

2.6 Roles and Protocol Theories

Role state predicates record the information accessed by a rule. They are also the mechanism by which a rule can enable the execution of another rule in the same role. Relying on a fixed protocol-wide set of role state predicates is dangerous since it could cause unexpected interferences between different instances of a role executing at the same time. Instead, we make role state predicates local to a role by requiring that fresh names be used each time a new instance of a role is executed. As in the case of rule consequents, we achieve this effect by using existential quantifiers: we prefix a collection of rules ρ that should share the same role state predicate L by a declaration of the form $\exists L : \vec{\tau}$, where the typed existential quantifier expresses the fact that L should be instantiated with a fresh role state predicate name of type $\vec{\tau}$. With this insight, the following grammar defines the notion of rule collection:

$$\begin{array}{lcl} \text{Rule collections: } \rho & ::= & \cdot \quad (\text{Empty role}) \\ & | & \exists L : \vec{\tau} \quad (\text{Role state predicate parameter declaration}) \\ & | & r, \rho \quad (\text{Extension with a rule}) \end{array}$$

A role is given as the association between a role owner A and a collection of rules ρ . Some roles, such as those implementing a server or an intruder, are intrinsically bound to a few specific principals, often just one. We call them anchored roles and denote them as ρ^A .

Here, the role owner A is an actual principal name, a constant. Other roles can be executed by any principal. In these cases A must be kept as a parameter bound to the role. We use the syntax $\rho^{\forall A}$ to represent these generic roles, where the implicitly typed universal quantification symbol implies that A should be instantiated to a principal before any rule in ρ is executed, and sets the scope of the binding to ρ . Observe that in this case A is a variable. With a slight abuse of notation, we sometimes refer to roles of either kind with the letter ρ , variously subscripted.

A protocol theory, written \mathcal{P} , is a finite collection of roles:

$$\begin{array}{lcl} \text{Protocol theories: } \mathcal{P} & ::= & \cdot \quad (\text{Empty protocol theory}) \\ & | & \mathcal{P}, \rho^{\forall A} \quad (\text{Extension with a generic role}) \\ & | & \mathcal{P}, \rho^A \quad (\text{Extension with an anchored role}) \end{array}$$

2.7 \diamond A Critique of Typed MSR's Encryption

In Typed MSR, symmetric and public-key encryption is deterministic, i.e. the ciphertext depends only on the plaintext and the key. However, this is not the case.

Symmetric Encryption Block ciphers are indeed deterministic. However, encryption employs block ciphers in various modes of operation, of which only one — the electronic codebook (ECB) mode — retains the property of being deterministic. All the others require

an initialization vector: a random block to kick off the process for the first real block, and also to make the process non-deterministic. Moreover, the ECB mode is generally not recommended for cryptographic protocols, as it reveals patterns in the plaintext.

Asymmetric Encryption The encryption process of non-probabilistic public-key cryptosystems, such as the RSA, is indeed deterministic. However, RSA must be combined with some form of padding scheme, so that no plaintext encrypts into an insecure ciphertext. These padding schemes use random numbers in order to calculate an appropriate padding for the plaintext.

Although modeling encryption as deterministic may be acceptable for the specification of, say, authentication protocols, it is however unsuitable for the specification of privacy-preserving protocols, as it can trigger false alarms for linkability attacks.

Chapter 3

Privacy-Preserving Cryptographic Abstractions

This chapter presents Typed MSR abstractions of cryptographic techniques that are commonly used in privacy-preserving protocols. These abstractions are either primitive, i.e. they consist of one or more Typed MSR message constructors, or they are derived, i.e. they are modeled using existing message constructors.

3.1 \diamond Primitive Cryptographic Abstractions

Our proposed changes and additions are contained in the following grammatical production:

<i>Messages:</i>	
$t ::=$	a (Atomic messages)
	x (Variables)
	$t_1 . t_2$ (Concatenation)
	$\{t\}_k^n$ (Symmetric encryption)
	$\{\{t\}\}_k^n$ (Asymmetric encryption)
	$[t]_{k'}^n$ (Digital signature)
	$\#\{\{t\}\}_k^n$ (Probabilistic asymmetric encryption)
	$\#[t]_{k'}^n$ (Probabilistic digital signature)
	$\ t\ _n$ (Commitment)
	$\langle t \rangle_n^k$ (Blinding)
	$\mathcal{Z}_{BS}(t, n_s, k, n_f, h)$ (Credential proof)
	$\mathcal{Z}_{AE}(t, t', n, k, h)$ (Number's upper bound proof)
	$\mathcal{P}(\#\{\{t_1\}\}_k^{n_1}, \dots, \#\{\{t_i\}\}_k^{n_i})$ (Vote aggregation)
	$\mathcal{S}(t_1, t_2, \dots, t_i)$ (Vote tallying)

As discussed in Section 2.7, we need to modify message constructors for symmetric and public-key encryption to model the necessary non-determinism. In order to achieve this, symmetric encryption, asymmetric encryption and digital signatures are made to depend

additionally on a nonce. Whenever referring to this nonce is of little importance and we would rather not name it, we denote it using $\nu()$.

Moreover, to be able to specify privacy-preserving protocols, we need to add message constructors for probabilistic asymmetric encryption, probabilistic digital signatures, blinding, commitment, zero-knowledge proofs, vote aggregation and vote tallying.

We also add types for probabilistic public and private keys:

$$\begin{array}{lcl} \text{Types: } \tau & ::= & \dots \quad (\text{see Section 2.4}) \\ & | & \text{pubK}^P A \quad (\text{Probabilistic public keys}) \\ & | & \text{privK}^P k \quad (\text{Probabilistic private keys}) \end{array}$$

Type $\text{pubK}^P A$ is used to classify the probabilistic public key of principal A . Type $\text{privK}^P k$ is used to classify the private key that corresponds to the probabilistic public key k .

In the rest of this Section, we discuss our newly introduced message constructors and their properties.

3.1.1 Asymmetric Encryption and Digital Signatures

In order to accommodate homomorphic encryption in the context of e-voting, we assume the use of an homomorphic encryption compatible cryptosystem featuring the additive homomorphic property, such as the Paillier cryptosystem [33]. Paillier is a probabilistic cryptosystem, so the nonce used to model the non-determinism employed in its mode of operation, is assumed to model the non-determinism of the core cryptosystem as well.

3.1.2 Commitment

Cryptographic commitment allows principals to choose and commit to a value without revealing it, in such a way that they are able to prove at a later time that the value they reveal is indeed the originally committed value.

Our abstraction of commitment is based on the non-interactive bit commitment using one-way hash functions. 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 nonce n . 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.

3.1.3 Blind Signatures

A blind signature is a form of digital signature in which the content of a message is disguised (blinded) before it is signed. The resulting blind signature can be publicly verified against the original unblinded message in the manner of a regular digital signature.

Blind signatures are typically employed in privacy-preserving protocols where the signer and message author are different parties. In order to prove to the signer a statement about the blinded message without disclosing the unblinded message, blind signatures are usually used together with a zero-knowledge proof.

Our abstraction of blind signatures and blinding is based on Chaum's blinding [17], according to which the construction of a blinded message depends on a blinding factor, which we can abstract as nonce n , and on a public key k . The fundamental property is that if message $\langle t \rangle_n^k$ is signed using private key k' (which corresponds to public key k), the resulting message can be unblinded using nonce n to produce the digital signature of message t signed using k' . Chaum's blinding assumes the use of the RSA cryptosystem.

3.1.4 Zero-Knowledge Proofs

A zero-knowledge proof is a method for one principal to prove to another that a statement is true, without revealing anything other than the veracity of the statement.

Zero-knowledge proofs are not proofs in the mathematical sense of the term because there is some small probability that a cheating prover will be able to convince the verifier of a false statement. However, there are standard techniques to decrease the soundness error to any arbitrarily small value.

Credential proof using blind signatures

Suppose Alice wants to get a credential t signed by an authority. To do this she blinds the credential m times using different blinding factors, with the credential each time committed with a different salt value. She then hashes the blinded credentials, together with the hashing nonce h (chosen by the authority), to determine the one blinding factor (n_f) and the one salt value (n_s) she will *not* reveal to the authority (non-interactive cut and choose). Lastly, she sends to the authority the blinded credentials, together with the blinding factors and the salt values that must be revealed. The authority can check that everything was formed correctly, and can sign $\langle \| t \|_{n_s} \rangle_{n_f}^k$.

Our abstraction of a zero-knowledge credential proof using blind signatures uses a variation of the non-interactive cut-and-choose protocol used in the selective disclosure protocol of Holt and Seamons, as described in Section 3.2.2 of [25]. Making the cut-and-choose protocol non-interactive has the following advantages: (i) it eliminates the extra complexity of the zero-knowledge credential proof and allows us to treat it in the same way as all other message constructors, (ii) it eliminates at least an extra step from the specification of the protocol, thus making it easier to reason about, and (iii) it simplifies the protocol's implementation.

The fundamental property of message $\mathcal{Z}_{BS}(t, n_s, k, n_f, h)$ is that it can convince a principal that t was used in the construction of $\langle \| t \|_{n_s} \rangle_{n_f}^k$ without disclosing nonces n_s and n_f , as long as it was the principal who chooses the hashing nonce h .

Number’s upper bound proof using asymmetric encryption

Our abstraction of a zero-knowledge number’s upper bound proof is based on the work of Boudot [11], who devised a zero-knowledge proof which is efficient and exact in demonstrating that a committed number lies in a specific interval. The fundamental property of message $\mathcal{Z}_{AE}(t, t', n, k, h)$ is that it can convince a principal that the value of t in $\#\{\{t\}_k^n$ is no greater than t' without disclosing message t , nor nonce n , as long as it was the principal who chooses the hashing nonce h .

3.1.5 Homomorphic Encryption

Homomorphic encryption refers to certain properties of probabilistic public key cryptosystems where correspondences can be proved to exist between functions on a certain group in the message space and functions on the corresponding group in the ciphertext space.

Our abstraction of homomorphic encryption is based on properties of the Paillier cryptosystem [3], and is formalized in terms of functions \mathcal{P} and \mathcal{S} . The first property is that

$$\mathcal{P}(\#\{\{t_1\}_k^{n_1}, \dots, \#\{\{t_i\}_k^{n_i}\}) = \#\{\{\mathcal{S}(t_1, \dots, t_i)\}_k^n$$

where k is a public key and t_1, \dots, t_i are messages. The second is that, assuming t_1, \dots, t_i represent the votes to be considered, $\mathcal{S}(t_1, \dots, t_i)$ represents the result of the tallying procedure. The third is that

$$\mathcal{P}(\mathcal{P}(\#\{\{t_1\}_k^{n_1}, \dots, \#\{\{t_{i-1}\}_k^{n_{i-1}}\}), \#\{\{t_i\}_k^{n_i}\}) = \mathcal{P}(\#\{\{t_1\}_k^{n_1}, \dots, \#\{\{t_i\}_k^{n_i}\})$$

which allows for vote aggregation without keeping all the encrypted votes, but only the result of function \mathcal{P} applied to its previous result and the newly acquired encrypted vote.

3.2 \diamond Derived Cryptographic Abstractions

We now discuss mixes and onion routing, and how these cryptographic abstractions can be modeled in Typed MSR.

3.2.1 Mixes

Generally considered the father of anonymous communications, David Chaum first proposed a system for anonymous email in 1981 [16]. The system he proposed used a special mail server, called a mix, to process email. A mix is a store-and-forward device that accepts a number of fixed-length messages from numerous sources, performs cryptographic transformations on the messages, and then forwards the messages to the next destination in an order not predictable from the order of inputs. Mixes enable anonymous communication by means of cryptography, scrambling the messages, and unifying them (padding to a constant size, fixing a constant sending rate by sending dummy messages, etc.). They support sender anonymity, and protect from traffic analysis.

Chaum's mix makes use of deterministic asymmetric-key encryption and nonce creation¹. Here's how A can send a message to B using mix M :

$$\begin{aligned} A &\rightarrow M : \{ \{ n_2 . \{ n_1 . t \}_{k_B}^{\nu()} . B \}_{k_M}^{\nu()} \\ M &\rightarrow B : \{ n_1 . t \}_{k_B}^{\nu()} \end{aligned}$$

3.2.2 Onion Routing

The primary innovation in onion routing [22] is the concept of the routing onion. Routing onions are data structures used to create paths through which many messages can be transmitted. To create an onion, the router at the head of a transmission selects a number of onion routers at random and generates a message for each one, providing it with symmetric keys for decrypting messages, and instructing it which router will be next in the path. Each of these messages, and the messages intended for subsequent routers, is encrypted with the corresponding router's public key. This provides a layered structure, in which it is necessary to decrypt all outer layers of the onion in order to reach an inner layer.

The onion metaphor describes the concept of such a data structure. As each router receives the message, it peels a layer off of the onion by decrypting with its private key, thus revealing the routing instructions meant for that router, along with the encrypted instructions for all of the routers located farther down the path. Due to this arrangement, the full content of an onion can only be revealed if it is transmitted to every router in the path in the order specified by the layering.

In Typed MSR, onion routing can be specified in the same way as a mix, because the number of onion layers is irrelevant for the abstraction.

¹Chaum uses nonce creation in order to guarantee that the asymmetric encryption result is intractable. In our formalization, this is sometimes redundant.

Chapter 4

Case Study: Specifying e-Voting Protocols

At this point, we demonstrate how the cryptographic abstractions described in the previous Chapter may be used to make Typed MSR specifications of two privacy-preserving protocols: an e-voting protocol based on blind signatures and mixes, and an e-voting protocol based on homomorphic encryption.

4.1 \diamond Protocol based on Blind Signatures and Mixes

We first give an informal description of the protocol, then list its security properties, and finally provide a formal specification of the protocol in Typed MSR.

4.1.1 Description

Preparing the ballot Alice wants to participate in an electronic election held by a voting Server. To do this, Alice sends to the Server a zero-knowledge credential 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 v_1 and v_2 represent the possible votes, signs the blind commitment of each vote and sends the signatures back to Alice.

$$\begin{aligned} S &\rightarrow A : h_1 . h_2 \\ A &\rightarrow S : \{ \mathcal{Z}_{BS}(v_1, s_1, k_S, f_1, h_1) . \mathcal{Z}_{BS}(v_2, s_2, k_S, f_2, h_2) \}_{k_{AS}}^{\nu()} \\ S &\rightarrow A : [\langle \| v_1 \|_{s_1} \rangle_{f_1}^{k_S}]_{k'_S}^{\nu()} . [\langle \| v_2 \|_{s_2} \rangle_{f_2}^{k_S}]_{k'_S}^{\nu()} \end{aligned}$$

Voting Alice unblinds the signatures of the blinded commitments, which gives her the signatures of the commitments. She can now cast her vote v_B by sending the signature of the vote's commitment — together with the vote itself and the nonce used in the computation of the commitment — to the Server via a Mix. Alice must also send to the Server

all the other possible votes she possesses, so that they can be canceled. Otherwise, Alice could cast more than one (albeit different) votes, thus destroying the election. The Server verifies its own signature and, after checking that the commitment is indeed computed using the data send, it accepts Alice's vote.

$$\begin{aligned}
A \rightarrow M &: \{ \{ n_2 \cdot \{ \{ n_1 \cdot v_B \cdot s_B \cdot [\| v_B \|_{s_B}]_{k'_S}^{\nu()}. v_A \cdot s_A \cdot [\| v_A \|_{s_A}]_{k'_S}^{\nu()}. S \} \}_{k_S}^{\nu()}. S \} \}_{k_M}^{\nu()} \\
M \rightarrow S &: \{ \{ n_1 \cdot v_B \cdot s_B \cdot [\| v_B \|_{s_B}]_{k'_S}^{\nu()}. v_A \cdot s_A \cdot [\| v_A \|_{s_A}]_{k'_S}^{\nu()}. S \} \}_{k_S}^{\nu()}
\end{aligned}$$

Tallying The Server posts the commitment signatures, together with the votes and nonces used in their computation, to a world-readable bulletin board, so that every voter can verify the election result and check that his vote has been counted in.

4.1.2 Specification

Here is the Typed MSR specification of Alice's *role* in the protocol:

$$\left(\begin{array}{l}
\exists L : \text{tract} \times \text{tract} \times \text{nonce} \times \text{nonce} \times \text{pubK } S \times \text{nonce} \times \text{nonce} \times \text{shK } A \ S \times \text{privK } k_S \\
\quad \times \text{nonce} \times \text{nonce}. \\
\exists L' : \text{tract} \times \text{tract} \times \text{nonce} \times \text{nonce} \times \text{pubK } S \times \text{nonce} \times \text{nonce} \times \text{shK } A \ S \times \text{privK } k_S \\
\quad \times \text{nonce} \times \text{nonce} \times \text{nonce} \times \text{nonce} \times \text{tract} \times \text{nonce} \times \text{tract} \times \text{nonce} \times \text{principal} \times \text{pubK } M. \\
\forall A :_{\text{e}} \text{principal}. \\
\forall S :_{\text{e}} \text{principal}. \\
\forall M :_{\text{e}} \text{principal}. \\
\forall v_1 :_{\text{r}} \text{tract}. \\
\forall v_2 :_{\text{r}} \text{tract}. \\
\forall s_1 :_{\text{e}} \text{nonce}. \\
\forall s_2 :_{\text{e}} \text{nonce}. \\
\forall k_S :_{\text{e}} \text{pubK } S. \\
\forall f_1 :_{\text{e}} \text{nonce}. \\
\forall f_1' :_{\text{e}} \text{nonce}. \\
\forall k_{AS} :_{\text{e}} \text{shK } A \ S. \\
\forall k'_S :_{\text{e}} \text{privK } k_S. \\
\forall n_2 :_{\text{e}} \text{nonce}. \\
\forall n_1 :_{\text{e}} \text{nonce}. \\
\forall v_B :_{\text{r}} \text{tract}. \\
\forall s_B :_{\text{e}} \text{nonce}. \\
\forall v_A :_{\text{r}} \text{tract}. \\
\forall s_A :_{\text{e}} \text{nonce}. \\
\forall k_M :_{\text{e}} \text{pubK } M.
\end{array} \right)^{\forall A}$$

$$\begin{aligned}
& N(h_1 \cdot h_2) \\
& \downarrow \\
& N(\{ \mathcal{Z}_{BS}(v_1, s_1, k_S, f_1, h_1) \cdot \mathcal{Z}_{BS}(v_2, s_2, k_S, f_2, h_2) \}_{k_{AS}}^{\nu()}) \\
& L(v_1, v_2, s_1, s_2, k_S, f_1, f_2, k_{AS}, k'_S, h_1, h_2) \\
& \\
& N([\langle \| v_1 \|_{s_1} \rangle_{f_1}^{k_S}]_{k'_S}^{\nu()} \cdot [\langle \| v_2 \|_{s_2} \rangle_{f_2}^{k_S}]_{k'_S}^{\nu()}) \\
& L(v_1, v_2, s_1, s_2, k_S, f_1, f_2, k_{AS}, k'_S, h_1, h_2) \\
& \downarrow \\
& N(\{ \{ n_2 \cdot \{ \{ n_1 \cdot v_B \cdot s_B \cdot [\| v_B \|_{s_B}]_{k'_S}^{\nu()}. v_A \cdot s_A \cdot [\| v_A \|_{s_A}]_{k'_S}^{\nu()}. S \} \}_{k_S}^{\nu()}. S \} \}_{k_M}^{\nu()}) \\
& L'(v_1, v_2, s_1, s_2, k_S, f_1, f_2, k_{AS}, k'_S, h_1, h_2, n_2, n_1, v_B, s_B, v_A, s_A, S, k_M)
\end{aligned}$$

Together, the roles of Alice, the Mix and the Server form the *protocol theory*.

4.1.3 Security Analysis

The protocol has the following security properties:

1. Only registered voters can have votes issued for them.

2. The server cannot link an actual vote with the voter neither when issuing nor when casting the vote.
3. No voter can have votes issued more than once.
4. No voter can cast more than one vote.
5. All voters can find their vote in the election result.
6. Everyone can verify the election result given the votes.
7. Everyone can count the votes.
8. The server may infringe the elections by casting its own votes, assuming that there are voters who do not try to find their own vote in the election result.

4.2 \diamond Protocol based on Homomorphic Encryption

Again, we first give an informal description of the protocol, then list its security properties, and finally provide a formal specification of the protocol in Typed MSR.

4.2.1 Description

Voting Alice wants to participate in an electronic election held by a voting Server and a Voting authority. To do this, Alice sends to the Server a zero-knowledge number's upper bound proof of the vote v_A she wishes to cast, encrypted using their shared key. The Server checks that Alice is eligible for voting and verifies that her vote is valid, i.e. that it is no greater than the maximum allowed vote v' . However, it cannot decrypt her vote $\# \{ \{ v_A \}_{k_V}^{n_A} \}$, as it is encrypted using the Voting authority's public key.

$$\begin{aligned}
 S &\rightarrow A : h \\
 A &\rightarrow S : \{ \mathcal{Z}_{AE}(v_A, v', n_A, k_V, h) \}_{k_{AS}}^{\nu_0}
 \end{aligned}$$

Tallying The Server applies the additive homomorphic encryption property and computes the election result, encrypted using the Voting authority's public key. It then sends it to the Voting authority, encrypted using their shared key.

$$S \rightarrow V : \{ \# \{ \mathcal{S}(\dots, v_A, \dots) \}_{k_V}^n \}_{k_{SV}}^{\nu_0}$$

Furthermore, the Server posts the encrypted votes to a world-readable bulletin board, so that every voter can check that his vote has been counted in, and also verify the calculation of $\# \{ \mathcal{S}(\dots, v_A, \dots) \}_{k_V}^n$. Finally, the Voting authority posts the result of the election $\mathcal{S}(\dots, v_A, \dots)$, as well as nonce n , so that every voter can verify the election result.

4.2.2 Specification

Here is the Typed MSR specification of Alice's *role* in the protocol:

$$\left(\begin{array}{l} \exists L : \text{tract} \times \text{tract} \times \text{nonce} \times \text{pubK } V \times \text{shK } A \ S \times \text{nonce}. \\ \forall A :_{\text{e}} \text{principal}. \\ \forall S :_{\text{e}} \text{principal}. \\ \forall V :_{\text{e}} \text{principal}. \\ \forall v_A :_{\text{r}} \text{tract}. \\ \forall v' :_{\text{r}} \text{tract}. \\ \forall n_A :_{\text{e}} \text{nonce}. \\ \forall k_V :_{\text{e}} \text{pubK } V. \\ \forall k_{AS} :_{\text{e}} \text{shK } A \ S. \end{array} \right)^{\forall A}$$

$$\begin{array}{c} N(h) \\ \downarrow \\ N(\{ \mathcal{Z}_{AE}(v_A, v', n_A, k_V, h) \}_{k_{AS}}^{\nu_0}) \\ L(v_A, v', n_A, k_V, k_{AS}, h) \end{array}$$

Together, the roles of Alice, the Voting Authority and the Server form the *protocol theory*.

4.2.3 Security Analysis

The protocol has the following security properties:

1. Only registered voters can vote.
2. The server cannot link an actual vote with the voter.
3. No voter can cast more than one vote.
4. All voters can find their vote in the election result.
5. Everyone can verify the election result given the votes.
6. Everyone can count the votes.
7. The server may infringe the elections by casting its own votes, assuming that there are voters who do not try to find their own vote in the election result.

Chapter 5

A Simple Linkability-Oriented Type System

As argued by Pfitzmann and Kohntopp [34] in reasoning about unlinkability, the item of interest is the sending and receiving of messages. In this Chapter, we elaborate on this and propose simple types that qualitatively classify messages according to their effect on the linkability of the sender.

5.1 \diamond Types

To be able to reason about linkability in privacy-preserving protocols, we introduce types for tractable, semitractable and intractable messages:

$$\begin{array}{l} \text{Types: } \tau ::= \dots \quad (\textit{see Section 3.1}) \\ \quad \quad | \text{ tract} \quad (\textit{Tractable messages}) \\ \quad \quad | \text{ semitract} \quad (\textit{Semitractable messages}) \\ \quad \quad | \text{ intract} \quad (\textit{Intractable messages}) \end{array}$$

These types are used to classify messages according to their commonness, which we claim is the defining characteristic of messages being considered for linkability exposures.

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 unlinkability of which it is safe to do so.

Type `intract` is used to classify messages that are very uncommon. These are pieces of information on the secrecy of which it is safe to base the correctness of a protocol, but on the unlinkability of which it is certainly erroneous to do so.

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 unlinked. It is not safe to base the correctness of a protocol neither on the secrecy of such pieces of information, nor on their unlinkedness.

We now classify each of the standard types according to their linkability. 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.

Similarly to the standard types, *tract*, *semitract* and *intract* should be regarded as sub-sort of *msg*.

The classification of messages that are not keys, nor nonces, nor principals is dealt with by *signatures*, which are described in Section 5.2.

5.2 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:

$$\begin{array}{l} \text{Signatures: } \Sigma ::= . \quad (\text{Empty signature}) \\ \quad \quad \quad | \Sigma, a : \tau \quad (\text{Atomic message declaration}) \end{array}$$

For our extended type system, we need two signatures. Signature Σ maps atomic messages to one of the standard types, and signature Γ maps 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:

$$\frac{}{(\Sigma, \alpha : \tau, \Sigma') \vdash \alpha :_{\Sigma} \tau} \text{(SIG1)} \quad \frac{}{(\Gamma, \alpha : \tau, \Gamma') \vdash \alpha :_{\Gamma} \tau} \text{(SIG2)}$$

5.3 \diamond Type Rules

We now introduce type rules for all the message constructors presented in Section 2.1 and Chapter 3. These rules use the new types introduced in Section 5.1 to further check the groundness of messages.

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

$$\frac{\Gamma \vdash t_1 : \tau \quad \Gamma \vdash t_2 : \tau}{\Gamma \vdash t_1 . t_2 : \tau} \text{ (CONCAT)}$$

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

$$\frac{\Gamma \vdash t_1 : \text{tract} \quad \Gamma \vdash t_2 : \text{semitract}}{\Gamma \vdash t_1 . t_2 : \text{semitract} \quad \Gamma \vdash t_2 . t_1 : \text{semitract}} \text{ (CONCAT1)}$$

$$\frac{\Gamma \vdash t_1 : \text{tract} \quad \Gamma \vdash t_2 : \text{intract}}{\Gamma \vdash t_1 . t_2 : \text{intract} \quad \Gamma \vdash t_2 . t_1 : \text{intract}} \text{ (CONCAT2)}$$

$$\frac{\Gamma \vdash t_1 : \text{semitract} \quad \Gamma \vdash t_2 : \text{intract}}{\Gamma \vdash t_1 . t_2 : \text{intract} \quad \Gamma \vdash t_2 . t_1 : \text{intract}} \text{ (CONCAT3)}$$

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

Symmetric-key encryption The ciphertext may be considered to be intractable because of the nonce used in the calculation.

$$\frac{\Gamma \vdash t : \tau \quad \Sigma \vdash k : \text{shK } A \ B}{\Gamma \vdash \{t\}_k^{\nu()}: \text{intract}} \text{ (SYMENC)}$$

Asymmetric-key encryption and digital signatures Similar reasoning applies to asymmetric-key encryption and digital signatures.

$$\frac{\Gamma \vdash t : \tau \quad \Sigma \vdash k : \text{pubK } A}{\Gamma \vdash \{\!\!\{t\}\!\!\}_k^{\nu()}: \text{intract}} \text{ (ASYMENC)} \quad \frac{\Gamma \vdash t : \tau \quad \Sigma \vdash k' : \text{privK } k}{\Gamma \vdash [t]_{k'}^{\nu()}: \text{intract}} \text{ (SIGN)}$$

Probabilistic asymmetric-key encryption and digital signatures Probabilistic asymmetric-key encryption and digital signatures may be considered to be intractable because of the nonce used in their calculation.

$$\frac{\Gamma \vdash t : \tau \quad \Sigma \vdash k : \text{pubK}^P A \quad \Sigma \vdash n : \text{nonce}}{\Gamma \vdash \#\{\!\!\{t\}\!\!\}_k^n : \text{intract}} \text{ (PASYMENC)}$$

$$\frac{\Gamma \vdash t : \tau \quad \Sigma \vdash k' : \text{privK}^P k \quad \Sigma \vdash n : \text{nonce}}{\Gamma \vdash \#[t]_{k'}^n : \text{intract}} \text{ (PSIGN)}$$

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

$$\frac{\Gamma \vdash t : \tau \quad \Sigma \vdash n_s : \text{nonce}}{\Gamma \vdash \|t\|_{n_s} : \text{intract}} \text{ (COMMIT)}$$

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

$$\frac{\Gamma \vdash t : \tau \quad \Sigma \vdash k : \text{pubK } A \quad \Sigma \vdash n_f : \text{nonce}}{\Gamma \vdash \langle t \rangle_{n_f}^k : \text{intract}} \text{ (BLIND)}$$

Zero-knowledge proofs The zero-knowledge credential proof can be considered to be intractable, as three nonces are used in its calculation. However, we require that the underlying message t of a zero-knowledge credential proof is tractable in order to protect privacy.

$$\frac{\Gamma \vdash t : \text{tract} \quad \Sigma \vdash n_s : \text{nonce} \quad \Sigma \vdash k : \text{pubK } A \quad \Sigma \vdash n_f : \text{nonce} \quad \Sigma \vdash h : \text{nonce}}{\Gamma \vdash \mathcal{Z}_{BS}(t, n_s, k, n_f, h) : \text{intract}} \text{ (ZEROPS)}$$

The zero-knowledge number's upper bound proof can be considered to be intractable, as two nonces are used in its calculation. However, we require that upper bound t' of message t is tractable in order to protect privacy.

$$\frac{\Gamma \vdash t : \tau \quad \Gamma \vdash t' : \text{tract} \quad \Sigma \vdash n : \text{nonce} \quad \Sigma \vdash k : \text{pubK}^P A \quad \Sigma \vdash h : \text{nonce}}{\Gamma \vdash \mathcal{Z}_{AE}(t, t', n, k, h) : \text{intract}} \text{ (ZEROAE)}$$

Vote aggregation The vote aggregation function \mathcal{P} can be considered to be tractable when applied to zero encrypted votes (because $\mathcal{P}()$ is a known constant), and intractable when applied to a non-zero number of encrypted votes (because each probabilistically encrypted vote is intractable). Furthermore, the unencrypted votes should be considered tractable for homomorphic encryption to give meaningful results.

$$\frac{\overline{\Gamma \vdash \mathcal{P}() : \text{tract}} \text{ (AGGBASE)}}{\frac{\Gamma \vdash t_1 : \text{tract} \quad \Gamma \vdash t_2 : \text{tract} \quad \cdots \quad \Gamma \vdash t_i : \text{tract} \quad \Sigma \vdash k : \text{pubK}^P A}{\Sigma \vdash n_1 : \text{nonce} \quad \Sigma \vdash n_2 : \text{nonce} \quad \cdots \quad \Sigma \vdash n_i : \text{nonce}}}{\Gamma \vdash \mathcal{P}(\#\{\{t_1\}_k^{n_1}, \#\{\{t_2\}_k^{n_2}, \dots, \#\{\{t_i\}_k^{n_i}\}) : \text{intract}} \text{ (AGGSTP)}}$$

Vote tallying The vote tallying function \mathcal{S} can be considered to be always tractable.

$$\frac{\overline{\Gamma \vdash \mathcal{S}() : \text{tract}} \text{ (TALLYBASE)}}{\frac{\Gamma \vdash t_1 : \text{tract} \quad \Gamma \vdash t_2 : \text{tract} \quad \cdots \quad \Gamma \vdash t_i : \text{tract}}{\Gamma \vdash \mathcal{S}(t_1, t_2, \dots, t_i) : \text{tract}} \text{ (TALLYSTEP)}$$

Chapter 6

The Dolev-Yao Intruder

The Dolev-Yao abstraction [19] 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 [14], a standard version of the Dolev-Yao intruder was formalized in Typed MSR. In this Section, we present an extended version of the Dolev-Yao intruder, which is able to discover attacks in privacy-preserving protocols.

It has been proved [39] 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 $M_I(-)$, whose single argument can hold a message, to enable I to store data out of sight from other principals.

6.1 The Standard Version

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 [14] for the formal specification of these operations in

Typed MSR.

6.2 \diamond An Extended Version

The version of the intruder that is presented here is an extended version in two ways. Firstly, one of the intruder's standard operations is generalized in line with the new types introduced in Section 5.1. More specifically, we 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 is 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 is now able to handle messages constructed using the message constructors introduced in Chapter 3.

We now formally specify the new operations in Typed MSR.

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

$$(\cdot \rightarrow \exists t :_r \text{intract. } M_I(t))_{(\text{CAP-GEN-I})}^!$$

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

$$(\forall t :_r \text{tract. } \cdot \rightarrow M_I(t))_{(\text{CAP-GUESS-T})}^!$$

$$(\forall t :_r \text{semitract. } \cdot \rightarrow M_I(t))_{(\text{CAP-GUESS-S})}^!$$

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.

Probabilistically encrypt The intruder may probabilistically encrypt a message given a public key and a nonce.

$$\left(\begin{array}{l} \forall t :_x \text{msg.} \\ \forall A :_x \text{principal.} \\ \forall k :_x \text{pubK}^P A. \\ \forall n :_x \text{nonce.} \end{array} \begin{array}{l} M_I(t) \\ M_I(k) \\ M_I(n) \end{array} \rightarrow M_I(\#\{t\}_k^n) \right)_{(\text{CAP-PASYMENC})}^!$$

Probabilistically decrypt Probabilistic decryption reveals to the intruder who holds the necessary private key not only the cleartext, but also the nonce representing the probabilistic nature of encryption.

$$\left(\begin{array}{l} \forall t :_{\mathbb{E}} \text{msg.} \\ \forall A :_{\mathbb{E}} \text{principal.} \\ \forall k :_{\mathbb{E}} \text{pubK}^P A. \\ \forall k' :_{\mathbb{E}} \text{privK}^P k. \\ \forall n :_{\mathbb{E}} \text{nonce.} \end{array} \quad \begin{array}{l} M_1(\#\{t\}_k^n) \\ M_1(k') \end{array} \rightarrow \begin{array}{l} M_1(t) \\ M_1(n) \end{array} \right)^I_{(\text{CAP-PASYMDEC})}$$

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

$$\left(\begin{array}{l} \forall t :_{\mathbb{E}} \text{msg.} \\ \forall A :_{\mathbb{E}} \text{principal.} \\ \forall k :_{\mathbb{E}} \text{pubK} A. \\ \forall n :_{\mathbb{E}} \text{nonce.} \end{array} \quad \begin{array}{l} M_1(t) \\ M_1(k) \\ M_1(n) \end{array} \rightarrow M_1(\langle t \rangle_n^k) \right)^I_{(\text{CAP-BLIND})}$$

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

$$\left(\begin{array}{l} \forall t :_{\mathbb{E}} \text{msg.} \\ \forall A :_{\mathbb{E}} \text{principal.} \\ \forall k :_{\mathbb{E}} \text{pubK} A. \\ \forall n :_{\mathbb{E}} \text{nonce.} \end{array} \quad \begin{array}{l} M_1(\langle t \rangle_n^k) \\ M_1(n) \end{array} \rightarrow M_1(t) \right)^I_{(\text{CAP-UNBLIND-MSG})}$$

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.

$$\left(\begin{array}{l} \forall t :_{\mathbb{E}} \text{msg.} \\ \forall A :_{\mathbb{E}} \text{principal.} \\ \forall k :_{\mathbb{E}} \text{pubK} A. \\ \forall k' :_{\mathbb{E}} \text{privK} k. \\ \forall n :_{\mathbb{E}} \text{nonce.} \end{array} \quad \begin{array}{l} M_1([\langle t \rangle_n^k]_{k'}^{\nu()}) \\ M_1(n) \end{array} \rightarrow M_1([t]_{k'}^{\nu()}) \right)^I_{(\text{CAP-UNBLIND-SIG})}$$

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

$$\left(\begin{array}{l} \forall t :_{\mathbb{E}} \text{msg.} \\ \forall n :_{\mathbb{E}} \text{nonce.} \end{array} \quad \begin{array}{l} M_1(t) \\ M_1(n) \end{array} \rightarrow M_1(\|t\|_n) \right)^I_{(\text{CAP-COMMIT})}$$

Generate a zero-knowledge proof The intruder may generate a zero-knowledge proof, given the necessary messages.

$$\left(\begin{array}{l} \forall t :_{\Sigma} \text{msg.} \\ \forall n_s :_{\Sigma} \text{nonce.} \\ \forall A :_{\Sigma} \text{principal.} \\ \forall k :_{\Sigma} \text{pubK } A. \\ \forall n_f :_{\Sigma} \text{nonce.} \\ \forall h :_{\Sigma} \text{nonce.} \end{array} \begin{array}{l} M_I(t) \\ M_I(n_s) \\ M_I(k) \\ M_I(n_f) \\ M_I(h) \end{array} \rightarrow M_I(\mathcal{Z}_{BS}(t, n_s, k, n_f, h)) \right)_{(\text{CAP-ZEROBS})}^1$$

$$\left(\begin{array}{l} \forall t :_{\Sigma} \text{msg.} \\ \forall t' :_{\Sigma} \text{msg.} \\ \forall n :_{\Sigma} \text{nonce.} \\ \forall A :_{\Sigma} \text{principal.} \\ \forall k :_{\Sigma} \text{pubK}^P A. \\ \forall h :_{\Sigma} \text{nonce.} \end{array} \begin{array}{l} M_I(t) \\ M_I(t') \\ M_I(n) \\ M_I(k) \\ M_I(h) \end{array} \rightarrow M_I(\mathcal{Z}_{AE}(t, t', n, k, h)) \right)_{(\text{CAP-ZEROAE})}^1$$

Observe a zero-knowledge proof The intruder will get the information revealed by the zero-knowledge proof.

$$\left(\begin{array}{l} \forall t :_{\Sigma} \text{msg.} \\ \forall n_s :_{\Sigma} \text{nonce.} \\ \forall A :_{\Sigma} \text{principal.} \\ \forall k :_{\Sigma} \text{pubK } A. \\ \forall n_f :_{\Sigma} \text{nonce.} \\ \forall h :_{\Sigma} \text{nonce.} \end{array} M_I(\mathcal{Z}_{BS}(t, n_s, k, n_f, h)) \rightarrow \begin{array}{l} M_I(t) \\ M_I(k) \\ M_I(\langle \| t \|_{n_s} \rangle_{n_f}^k) \\ M_I(h) \end{array} \right)_{(\text{CAP-ZEROBS-P})}^1$$

$$\left(\begin{array}{l} \forall t :_{\Sigma} \text{msg.} \\ \forall t' :_{\Sigma} \text{msg.} \\ \forall n :_{\Sigma} \text{nonce.} \\ \forall A :_{\Sigma} \text{principal.} \\ \forall k :_{\Sigma} \text{pubK}^P A. \\ \forall h :_{\Sigma} \text{nonce.} \end{array} M_I(\mathcal{Z}_{AE}(t, t', n, k, h)) \rightarrow \begin{array}{l} M_I(t') \\ M_I(k) \\ M_I(\# \{ t \}_k^n) \\ M_I(h) \end{array} \right)_{(\text{CAP-ZEROAE-P})}^1$$

Aggregate votes The intruder may generate¹ the image of zero votes under function \mathcal{P} (induction base case).

$$(\cdot \rightarrow M_I(\mathcal{P}()))_{(\text{CAP-AGGBASE})}^1$$

Furthermore, the intruder may aggregate encrypted votes as he picks them up by holding their image under function \mathcal{P} (induction step).

¹This rule is in fact redundant, as $\mathcal{P}()$ is of type `tract`, and therefore the intruder is already able to generate it, but is included for the sake of completeness.

$$\left(\begin{array}{l} \forall A :_{\Sigma} \text{principal.} \\ \forall k :_{\Sigma} \text{pubK}^{\text{P}} A. \\ \forall t_1 :_{\Sigma} \text{msg.} \quad \text{M}_I(\mathcal{P}(\#\{\{t_1\}_k^{n_1}, \dots, \#\{t_{i-1}\}_k^{n_{i-1}}\})) \\ \dots \\ \forall t_i :_{\Sigma} \text{msg.} \quad \text{M}_I(\#\{\{t_i\}_k^{n_i}\}) \\ \quad \quad \quad \downarrow \\ \forall n_1 :_{\Sigma} \text{nonce.} \quad \text{M}_I(\mathcal{P}(\#\{\{t_1\}_k^{n_1}, \dots, \#\{t_i\}_k^{n_i}\})) \\ \dots \\ \forall n_i :_{\Sigma} \text{nonce.} \end{array} \right)^I_{(\text{CAP-AGGSTEP})}$$

Tally votes The intruder may generate² the image of zero votes under function \mathcal{S} (induction base case).

$$(\cdot \rightarrow \text{M}_I(\mathcal{S}()))^I_{(\text{CAP-TALLYBASE})}$$

Furthermore, the intruder may tally votes as he picks them up by holding their image under function \mathcal{S} (induction step).

$$\left(\begin{array}{l} \forall t_1 :_{\Sigma} \text{msg.} \\ \forall t_2 :_{\Sigma} \text{msg.} \quad \text{M}_I(\mathcal{S}(t_1, \dots, t_{i-1})) \rightarrow \text{M}_I(\mathcal{S}(t_1, \dots, t_i)) \\ \dots \\ \forall t_i :_{\Sigma} \text{msg.} \quad \text{M}_I(t_i) \end{array} \right)^I_{(\text{CAP-TALLYSTEP})}$$

Apply homomorphic encryption properties The intruder may convert the image of the encrypted votes under function \mathcal{P} to the image of the (cleartext) votes under function \mathcal{S} , and vice-versa.

$$\left(\begin{array}{l} \forall A :_{\Sigma} \text{principal.} \\ \forall k :_{\Sigma} \text{pubK}^{\text{P}} A. \\ \forall t_1 :_{\Sigma} \text{msg.} \\ \dots \\ \forall t_i :_{\Sigma} \text{msg.} \quad \text{M}_I(\mathcal{P}(\#\{\{t_1\}_k^{n_1}, \dots, \#\{t_i\}_k^{n_i}\})) \rightarrow \text{M}_I(\#\{\{\mathcal{S}(t_1, \dots, t_i)\}_k^n\}) \\ \forall n_1 :_{\Sigma} \text{nonce.} \\ \dots \\ \forall n_i :_{\Sigma} \text{nonce.} \\ \forall n :_{\Sigma} \text{nonce.} \end{array} \right)^I_{(\text{CAP-PENCH-1})}$$

²The previous footnote applies here too.

$$\left(\begin{array}{l} \forall A :_{\mathcal{E}} \text{principal.} \\ \forall k :_{\mathcal{E}} \text{pubK}^P A. \\ \forall t_1 :_{\mathcal{E}} \text{msg.} \\ \dots \\ \forall t_i :_{\mathcal{E}} \text{msg.} \quad M_I(\#\{\mathcal{S}(t_1, \dots, t_i)\}_k^n) \rightarrow M_I(\mathcal{P}(\#\{t_1\}_k^{n_1}, \dots, \#\{t_i\}_k^{n_i})) \\ \forall n_1 :_{\mathcal{E}} \text{nonce.} \\ \dots \\ \forall n_i :_{\mathcal{E}} \text{nonce.} \\ \forall n :_{\mathcal{E}} \text{nonce.} \end{array} \right)^1$$

(CAP-PENCH-2)

6.3 \diamond Demonstrating Linkability Attacks

Informally, when we say that two executions of a protocol or a subprotocol cannot be linked to a given principal (usually the principal whose privacy the protocol is supposed to protect), we mean that it is not possible for the Dolev-Yao intruder to find out whether the same principal participated in both occasions, even if he manages to overtake all the other principals and get hold of their long-term or short-term secrets.

But how can the intruder's capabilities help in demonstrating such weaknesses in protocols? Consider the case of the protocol of Section 4.1, and suppose that the intruder has access to all public data, and, additionally, he manages to overtake the Server, so he has access to all the server's data as well. If he is able to deduce the same intractable or semitractable message *separately* in the context of the two subprotocols (the ballot preparation subprotocol and the voting subprotocol), and if he couldn't deduce this message without the information sent by Alice in each context, then we can consider this an attack on Alice's privacy, as her obtaining the votes and actually voting are now linked together.

In this sense, we argue that the intruder's capabilities create a formal environment in which linkability attacks on protocols may be demonstrated.

Part II

Implementing Privacy-Preserving Cryptographic Protocols

Chapter 7

The Implementation Language: Jif

Jif [30, 31, 29] is an object-oriented, strongly-typed language capturing a large subset of the Java language. In Jif, the programmer must label types with security annotations. The compiler uses these annotations during type-checking to ensure noninterference. Jif was developed primarily by Andrew Myers.

7.1 The Decentralized Label Model (DLM)

Types in Jif are annotated with security labels based on the DLM. Similar to work in mandatory access control that tags data with complete access control lists, the DLM allows for the virtual tagging of data with owners-readers and owners-writers lists. Each label consists of a set of confidentiality or integrity policies of the form $\{o : r_1, r_2, \dots, r_n\}$ or $\{o ! : r_1, r_2, \dots, r_n\}$ respectively, where o and r_i are principals with o being the owner of the policy and r_i being the authorized readers or writers of the confidentiality or integrity policy.

Furthermore, a label can consist of multiple policies (allowing for multiple owners of a piece of data).

As an example, `int{Alice : } i;` declares an `int` owned and readable only by Alice (the owner is always implicitly included in the reader/writer list). The statement `String{Bob !: Charlie, Dana} str;` declares a `String` which is owned by Bob but also writable by Charlie and Dana. Data may also be annotated with multiple policies as in `int{Alice ;; Bob : } j;`. The policy on `j` indicates that it is owned and readable by both Alice and Bob. In Jif, when a variable is used in a security label, it refers to its own label. Thus, using `i` and `str` as defined above, `float{i; str} f;` declares a `float` that is owned by both Alice and Bob and which can be written by Charlie and Dana.

7.2 Program Counter

The label of an expression's value varies depending on the evaluation context. This is needed to prevent leaks through implicit flows: channels created by the control flow struc-

ture itself. To prevent information leaks through implicit flows, the compiler associates a program-counter label with every statement and expression, representing the information that might be learned from the knowledge that the statement or expression was evaluated.

7.3 Language Features

In a Jif method declaration, the return value, the arguments, and the exceptions may each be annotated with a label. There are additionally two optional labels in a method declaration called the begin-label and the end-label. The begin-label specifies an upper bound on the program counter at the point of invocation of the method. The begin-label allows information about the program counter of the caller to be used for statically checking the implementation, thereby preventing assignments within the method from creating implicit flows of information. When these labels are missing, some conservative rules are used to assign restrictive default labels:

- Default field label: the empty label, $\{\}$. This label is the least confidential, and the least trusted. This conservatively ensures that no confidential or trusted data can be stored in the field.
- Default argument label: the top label, $\{* : \}$. The label on the type of a formal argument is an upper bound for labels of actual arguments.
- Default method begin-label: the top label, $\{* : \}$. The method begin-label is an upper bound on the pc of the caller, and a lower bound on the side effects of the method. The default method begin label is the most restrictive label, meaning that the method has no side effects.
- Default method end-label: The join of the declared labels of any exceptions declared to be thrown. If the method does not throw any exceptions, or if the declared exceptions do not have any labels, the default method end-label is the bottom label, $\{* ! : *\}$.
- Default method return value label: the join of all the argument labels and the end-label. This is the common case, as most of the time the value returned by a method is the result of computation on all of its arguments.
- Default declared exception label: the method end-label.
- Default array base label: the empty label, $\{\}$.

7.4 Selective Declassification

Jif implements selective declassification. Principals in Jif are defined external to the program. Each one has a delegation set containing all the principals it trusts. This forms

a runtime principal hierarchy. Each process maintains an authority set which contains principals from the runtime principal hierarchy. A process is only authorized to declassify policies that are owned by principals in its authority set.

7.5 Class Parameterization

Another feature in Jif which we utilize is class parameterization. A Jif class can be parameterized by a principal or security label. This means that a class may be defined once and then be instantiated at various security levels. For example, we might want a `Vector` class which contains `secret` data and another `Vector` class that contains `public` data. Without having to implement the `Vector` class multiple times, it could be parameterized with a label and then instantiated at different levels. In Jif, such a class could be defined as follows:

```
public class Vector[principal P]
{
  Object{P:}[] {P:} elements;
}
```

Note that the member array `elements` has two labels. One is the label of the `Objects` stored in the array. The other is the label of the array itself. Since `Vector` has been parameterized by `P`, `P` can now be used throughout the body of the class to denote a principal. This principal will be instantiated when an object of type `Vector` is declared, as in the following code, where `Alice` and `Bob` are two principals:

```
Vector[Alice] vector1;
Vector[Bob] vector2;
```

7.6 Dynamic labels

Labels and principals can be used as first-class values, represented at runtime. These dynamic labels and principals can be used in the specification of other labels, and used as the parameters of parameterized classes. They can be constructed as shown below:

```
final label lb = new label {Alice: Bob; Alice!:*};
```

7.7 Handling Exceptions

One thing which makes Jif particularly challenging for programming is handling the information leaks that occur through function termination, exceptions and side-effects. For example, an encryption method that throws an `InvalidKey` exception releases information about the key (which is secret data) both by throwing the exception (indicating the key is

invalid) and by not throwing the exception, i.e. by terminating normally (indicating the key is not invalid). For this reason, it can be advantageous to catch exceptions and handle them locally in order to bound information leakage they might cause.

Chapter 8

Linkability-Checking in Jif

This Chapter demonstrates how Jif (a security-oriented extension of a subset of the Java programming language dealing with information flow) can be employed in such a way that linkability vulnerabilities in protocol implementations can be detected with a mixture of static and runtime checks. The code presented in this Chapter was compiled with Jif 3.0.

8.1 \diamond A Cryptographic Framework

In this Section, we demonstrate how the theory described in Sections 5.3 and 6.2 can be employed in a Jif cryptographic framework, which can be used to implement privacy-preserving protocols, such as those analyzed in Sections 4.1 and 4.2. Our goal is to enforce the well-formedness of message construction, as defined by the type rules of Section 5.3, and to embed the Dolev-Yao intruder’s capabilities, as described in Section 6.2. Our scope is the secure handling of intractable messages, which, as described in Section 5.1, if revealed more than once to the intruder, will be linked together, thus compromising the anonymity of the sender.

First, we introduce a dynamic label, `L`, which annotates references to objects depicting intractable information. To define this label, we need to introduce a special principal, `Linkable`, the owner and sole reader of all intractable data. We place our label `L` in the Jif interface `LinkLabel`, so that every Jif interface/class that needs to use the label can simply extend/implement this interface.

```
interface LinkLabel
{
  final label L = new label{Linkable:};
}
```

Then, we establish the interface `Message`, which all message classes must implement. To implement the interface, message classes must provide the method `getObservedMessage`, which returns the other messages one learns by observing the message. These messages are returned one-by-one by invoking this method and increasing the index argument at each

iteration. At the first invocation, the index must be 0. When the method returns `null`, there are no more messages to learn.

```
interface Message extends LinkLabel
{
    Message{L} getObservedMessage{L}(int{L} index);
}
```

We now demonstrate how a selection of message constructors can be expressed as Jif classes implementing the `Message` interface.

Nonce A nonce can be created without the knowledge of any other message. Furthermore, observing a nonce reveals no other message.

```
public class NonceMessage implements Message
{
    public NonceMessage():{L} { }
    public Message{L} getObservedMessage{L}(int{L} index)
    { return null; }
}
```

Zero-knowledge credential proof According to capability (CAP-ZEROBS), in order to create a zero-knowledge credential proof, one needs to have the underlying message, the public key of the principal who is to sign the blinded commitment, and two nonces: the blinding factor and the salt used in the commitment. Furthermore, according to capability (CAP-ZEROBS-P), observing a zero-knowledge credential proof allows one to learn the underlying message, the public key and the blinded commitment. Note that we assume the existence of the two classes `BlindMessage` and `CommitMessage` for the blinding and the commitment respectively.

```
public class ZeroBlindSigMessage implements Message
{
    final private Message{} t;
    final private NonceMessage{L} n_s;
    final private PublicKeyMessage{} k;
    final private NonceMessage{L} n_f;
    final private NonceMessage{L} h;

    /* CAP-ZEROBS */
    public ZeroBlindSigMessage(Message{} t,
                               NonceMessage{L} n_s,
                               PublicKeyMessage{} k,
                               NonceMessage{L} n_f,
                               NonceMessage{L} h):{L}
    { this.t = t; this.n_s = n_s; this.k = k;
```

```

    this.n_f = n_f; this.h = h; }

/* CAP-ZEROBS-P */
public Message{L} getObservedMessage{L}(int{L} index)
{
    switch(index)
    {
        case 0: return t;
        case 1: return k;
        case 2: return new BlindMessage(new CommitMessage(t, n_s),
                                         k,
                                         n_f);
        case 3: return h;
        default: return null;
    }
}
}
}

```

Conditional messages According to capability (*CAP-UNBLIND-SIG*), the observation of a blinded signature reveals the signature only if the observer knows the blinding factor. To be able to incorporate such dependencies in method `getObservedMessage`, we introduce the class `ConditionalMessage`. The constructor of this class takes as arguments the conditionally observed message, and the message the observer needs to know in order to observe the first. If multiple messages are needed to specify the condition, they can be concatenated into a single message.

This new message constructor is denoted in Typed MSR as $\mathcal{C}(t_1, t_2)$, and in order to incorporate it into the extended version of the Dolev-Yao Intruder (discussed in Section 6.2), we need to include the following intruder capability:

$$\left(\begin{array}{l} \forall t_1 :_{\mathbb{E}} \text{msg. } M_I(t_1) \\ \forall t_2 :_{\mathbb{E}} \text{msg. } M_I(\mathcal{C}(t_2, t_1)) \end{array} \right) \xrightarrow{(\text{CAP-COND})} M_I(t_1, t_2)$$

Notice that this message constructor does not model any new cryptographic primitive or technique. Its sole purpose in our model is to support the `Message` interface of our Jif framework.

8.2 \diamond Runtime Linkability-Checking

Having in place a mechanism to annotate linkable messages, we now focus on message sending with runtime checking of possible linkability exposures.

We first argue that, in the general case, this checking cannot be done at compile time, as the following pseudocode illustrates:

```

send_message(m);
if (theorem)
    send_message(m);

```

Suppose that message m is linkable. If this is the case, then sending it twice would cause a linkability exposure. But, if it was possible to have this check done at compile time, then we would have a compiler capable of proving any theorem. Therefore, in the general case, we can do no better than runtime checking for linkability exposures. Although such an attempt raises the effort needed to implement the protocol, it also has the advantage that it is the actual protocol's implementation that is being verified, not a possibly flawed abstraction.

Our proposed approach is based on having a single point in the Jif program that is able to send messages to the network, which distinguishes between linkable and other messages, is able to store and search through all previously sent messages, and will throw an exception when an attempt is made to send a message that could be linked with a previous one.

```

interface Network
{
    void send(Message m) throws PossibleLinkabilityException;
    Message receive();
}

```

If the message is not linkable, it is freely send on the network. However, if it is linkable, then it is send only if the following conditions are met:

1. The message has not been previously sent.
2. The linkable messages inferred by the intruder from this, and all previously sent messages, have not been sent before.

Otherwise the `PossibleLinkabilityException` is thrown.

Regarding the latter condition, the linkable messages inferred by the intruder can be found using the method described in Section 7 of [14]. The method is presented (and expanded to include rules for our introduced cryptographic primitives) in Section 8.3.

Therefore, implementing the `send` method involves: (i) storing all previously sent messages, and (ii) using a theorem prover that works on these stored messages (axioms) and on the rules of Section 8.3 to check whether the message to be sent can be inferred from the stored messages. If the last statement is proved (i.e. it is a theorem), the message is not sent.

8.3 \diamond Message Inference

The intruder's capabilities were formalized as Typed MSR rules in Section 6.2. However, these rules are too non-deterministic to use in a model checking simulation to uncover

attacks on protocols. What is needed is an operational version of the intruder, which retains the intruder's capabilities, while not allowing him to undo his own work.

As proposed in [26], to construct this version we have the intruder decompose messages he learns into atomic messages, store these atomic messages, and then use them to construct new complex messages. In other words, we partition the actions of the intruder into three distinct constituents: message decomposition, storage of elementary information, and message construction. The memory predicate $M_I(-)$ is replaced by the following memory predicates:

- $D_I(-)$: intended to contain messages while they are disassembled into their elementary constituents
- $A_I(-)$: intended to contain atomic messages learnt this way
- $C_I(-)$: intended to contain messages while they are used to construct more complex ones

Therefore, we have the following rules for interception and transmission:

$$(\forall t :_x \text{msg}. N(t) \rightarrow D_I(t))_{(\text{DECOMPOSE-INIT})}^I$$

$$(\forall t :_x \text{msg}. C_I(t) \rightarrow N(t))_{(\text{COMPOSE-END})}^I$$

We always keep a copy of intercepted messages, so that we can use information we acquire in the future to break them into more atomic messages:

$$\left(\forall t :_x \text{msg}. D_I(t) \rightarrow \begin{array}{l} D_I(t) \\ C_I(t) \end{array} \right)_{(\text{DECOMPOSE-COPY})}^I$$

Atomic messages acquired during the decomposition process, will be moved in the appropriate memory predicate:

$$(\forall a :_x \text{atm}. D_I(t) \rightarrow A_I(t))_{(\text{MEMORIZE-ATOMICS})}^I$$

Finally, the atomic messages are used to feed the composition process (copying is required as these messages may be needed later):

$$\left(\forall a :_x \text{atm}. A_I(a) \rightarrow \begin{array}{l} A_I(a) \\ C_I(a) \end{array} \right)_{(\text{COMPOSE-COPY})}^I$$

We now express the intruder capabilities of Section 6.2 using our new memory predicates.

Generate fresh intractable data The intruder may generate fresh nonces, fresh private keys, fresh shared keys, as well as other intractable messages, as long as they are atomic.

$$(\cdot \rightarrow \exists t :_{\mathcal{I}} \text{intract}. \exists t :_{\mathcal{E}} \text{atm}. A_I(t))_{(L-CAP-GEN-I)}^!$$

Guess tractable and semitractable data The intruder may guess or get access to public keys, principals, as well as other tractable or semitractable messages, as long as they are atomic.

$$(\forall t :_{\mathcal{I}} \text{tract}. \forall t :_{\mathcal{E}} \text{atm}. \cdot \rightarrow A_I(t))_{(L-CAP-GUESS-T)}^!$$

$$(\forall t :_{\mathcal{I}} \text{semitract}. \forall t :_{\mathcal{E}} \text{atm}. \cdot \rightarrow A_I(t))_{(L-CAP-GUESS-S)}^!$$

Probabilistically encrypt We consider this to be part of the construction phase.

$$\left(\begin{array}{l} \forall t :_{\mathcal{E}} \text{msg}. \\ \forall A :_{\mathcal{E}} \text{principal}. \\ \forall k :_{\mathcal{E}} \text{pubK}^P A. \\ \forall n :_{\mathcal{E}} \text{nonce}. \end{array} \begin{array}{l} C_I(t) \\ C_I(k) \\ C_I(n) \end{array} \rightarrow C_I(\#\{t\}_k^n) \right)_{(L-CAP-PASYMENC)}^!$$

Probabilistically decrypt We consider this to be part of the destruction phase.

$$\left(\begin{array}{l} \forall t :_{\mathcal{E}} \text{msg}. \\ \forall A :_{\mathcal{E}} \text{principal}. \\ \forall k :_{\mathcal{E}} \text{pubK}^P A. \\ \forall k' :_{\mathcal{E}} \text{privK}^P k. \\ \forall n :_{\mathcal{E}} \text{nonce}. \end{array} \begin{array}{l} D_I(\#\{t\}_k^n) \\ D_I(k') \\ D_I(n) \end{array} \rightarrow \begin{array}{l} D_I(t) \\ D_I(n) \end{array} \right)_{(L-CAP-PASYMDEC)}^!$$

Blind messages We consider this to be part of the construction phase.

$$\left(\begin{array}{l} \forall t :_{\mathcal{E}} \text{msg}. \\ \forall A :_{\mathcal{E}} \text{principal}. \\ \forall k :_{\mathcal{E}} \text{pubK} A. \\ \forall n :_{\mathcal{E}} \text{nonce}. \end{array} \begin{array}{l} C_I(t) \\ C_I(k) \\ C_I(n) \end{array} \rightarrow C_I(\langle t \rangle_n^k) \right)_{(L-CAP-BLIND)}^!$$

Unblind messages We consider this to be part of the destruction phase.

$$\left(\begin{array}{l} \forall t :_{\mathcal{E}} \text{msg}. \\ \forall A :_{\mathcal{E}} \text{principal}. \\ \forall k :_{\mathcal{E}} \text{pubK} A. \\ \forall n :_{\mathcal{E}} \text{nonce}. \end{array} \begin{array}{l} D_I(\langle t \rangle_n^k) \\ D_I(n) \end{array} \rightarrow D_I(t) \right)_{(L-CAP-UNBLIND-MSG)}^!$$

Unblind signatures This is not so obvious. We believe that this capability can be used both in the destruction and in the construction phase.

$$\left(\begin{array}{l} \forall t :_{\Sigma} \text{msg.} \\ \forall A :_{\Sigma} \text{principal.} \\ \forall k :_{\Sigma} \text{pubK } A. \\ \forall k' :_{\Sigma} \text{privK } k. \\ \forall n :_{\Sigma} \text{nonce.} \end{array} \begin{array}{l} D_1 \left([\langle t \rangle_n^k]_{k'}^{\nu_0} \right) \\ D_1(n) \end{array} \rightarrow D_1 \left([t]_{k'}^{\nu_0} \right) \right)^!_{(\text{L-CAP-UNBLIND-SIG-1})}$$

$$\left(\begin{array}{l} \forall t :_{\Sigma} \text{msg.} \\ \forall A :_{\Sigma} \text{principal.} \\ \forall k :_{\Sigma} \text{pubK } A. \\ \forall k' :_{\Sigma} \text{privK } k. \\ \forall n :_{\Sigma} \text{nonce.} \end{array} \begin{array}{l} C_1 \left([\langle t \rangle_n^k]_{k'}^{\nu_0} \right) \\ C_1(n) \end{array} \rightarrow C_1 \left([t]_{k'}^{\nu_0} \right) \right)^!_{(\text{L-CAP-UNBLIND-SIG-2})}$$

Commit to a message We consider this to be part of the construction phase.

$$\left(\begin{array}{l} \forall t :_{\Sigma} \text{msg.} \\ \forall n :_{\Sigma} \text{nonce.} \end{array} \begin{array}{l} C_1(t) \\ C_1(n) \end{array} \rightarrow C_1(\|t\|_n) \right)^!_{(\text{L-CAP-COMMIT})}$$

Generate a zero-knowledge proof We consider this to be part of the construction phase.

$$\left(\begin{array}{l} \forall t :_{\Sigma} \text{msg.} \\ \forall n_s :_{\Sigma} \text{nonce.} \\ \forall A :_{\Sigma} \text{principal.} \\ \forall k :_{\Sigma} \text{pubK } A. \\ \forall n_f :_{\Sigma} \text{nonce.} \\ \forall h :_{\Sigma} \text{nonce.} \end{array} \begin{array}{l} C_1(t) \\ C_1(n_s) \\ C_1(k) \\ C_1(n_f) \\ C_1(h) \end{array} \rightarrow C_1(\mathcal{Z}_{BS}(t, n_s, k, n_f, h)) \right)^!_{(\text{L-CAP-ZEROBS})}$$

$$\left(\begin{array}{l} \forall t :_{\Sigma} \text{msg.} \\ \forall t' :_{\Sigma} \text{msg.} \\ \forall n :_{\Sigma} \text{nonce.} \\ \forall A :_{\Sigma} \text{principal.} \\ \forall k :_{\Sigma} \text{pubK}^P A. \\ \forall h :_{\Sigma} \text{nonce.} \end{array} \begin{array}{l} C_1(t) \\ C_1(t') \\ C_1(n) \\ C_1(k) \\ C_1(h) \end{array} \rightarrow C_1(\mathcal{Z}_{AE}(t, t', n, k, h)) \right)^!_{(\text{L-CAP-ZEROAE})}$$

Observe a zero-knowledge proof We consider this to be part of the destruction phase.

$$\left(\begin{array}{l} \forall t :_{\Sigma} \text{msg.} \\ \forall n_s :_{\Sigma} \text{nonce.} \\ \forall A :_{\Sigma} \text{principal.} \\ \forall k :_{\Sigma} \text{pubK } A. \\ \forall n_f :_{\Sigma} \text{nonce.} \\ \forall h :_{\Sigma} \text{nonce.} \end{array} \begin{array}{l} D_1(\mathcal{Z}_{BS}(t, n_s, k, n_f, h)) \\ D_1(t) \\ D_1(k) \\ D_1(\langle \|t\|_{n_s} \rangle_{n_f}^k) \end{array} \right)^!_{(\text{L-CAP-ZEROBS-P})}$$

$$\left(\begin{array}{l} \forall t :_{\mathbb{E}} \text{msg.} \\ \forall t' :_{\mathbb{E}} \text{msg.} \\ \forall n :_{\mathbb{E}} \text{nonce.} \\ \forall A :_{\mathbb{E}} \text{principal.} \\ \forall k :_{\mathbb{E}} \text{pubK}^{\text{P}} A. \\ \forall h :_{\mathbb{E}} \text{nonce.} \end{array} \quad \text{D}_1(\mathcal{Z}_{AE}(t, t', n, k, h)) \rightarrow \begin{array}{l} \text{D}_1(t') \\ \text{D}_1(k) \\ \text{D}_1(\#\{\!\{ t \}\!\}_k^n) \end{array} \right)_{(\text{L-CAP-ZEROAE-P})}^!$$

Aggregate votes We consider the induction base case to be part of the construction phase.

$$(\cdot \rightarrow \text{C}_1(\mathcal{P}()))_{(\text{L-CAP-AGGBASE})}^!$$

We also consider the induction step to be part of the construction phase.

$$\left(\begin{array}{l} \forall A :_{\mathbb{E}} \text{principal.} \\ \forall k :_{\mathbb{E}} \text{pubK}^{\text{P}} A. \\ \forall t_1 :_{\mathbb{E}} \text{msg.} \\ \dots \\ \forall t_i :_{\mathbb{E}} \text{msg.} \\ \forall n_1 :_{\mathbb{E}} \text{nonce.} \\ \dots \\ \forall n_i :_{\mathbb{E}} \text{nonce.} \end{array} \quad \begin{array}{l} \text{C}_1(\mathcal{P}(\#\{\!\{ t_1 \}\!\}_k^{n_1}, \dots, \#\{\!\{ t_{i-1} \}\!\}_k^{n_{i-1}})) \\ \text{C}_1(\#\{\!\{ t_i \}\!\}_k^{n_i}) \\ \downarrow \\ \text{C}_1(\mathcal{P}(\#\{\!\{ t_1 \}\!\}_k^{n_1}, \dots, \#\{\!\{ t_i \}\!\}_k^{n_i})) \end{array} \right)_{(\text{L-CAP-AGGSTP})}^!$$

Tally votes We consider the induction base case to be part of the construction phase.

$$(\cdot \rightarrow \text{C}_1(\mathcal{S}()))_{(\text{L-CAP-TALLYBASE})}^!$$

We also consider the induction step to be part of the construction phase.

$$\left(\begin{array}{l} \forall t_1 :_{\mathbb{E}} \text{msg.} \\ \forall t_2 :_{\mathbb{E}} \text{msg.} \\ \dots \\ \forall t_i :_{\mathbb{E}} \text{msg.} \end{array} \quad \begin{array}{l} \text{C}_1(\mathcal{S}(t_1, \dots, t_{i-1})) \\ \text{C}_1(t_i) \end{array} \rightarrow \text{C}_1(\mathcal{S}(t_1, \dots, t_i)) \right)_{(\text{L-CAP-TALLYSTEP})}^!$$

Apply homomorphic encryption properties We believe that these capabilities can be used both in the destruction and in the construction phase.

$$\left(\begin{array}{l} \forall A :_{\mathbb{E}} \text{principal.} \\ \forall k :_{\mathbb{E}} \text{pubK}^{\text{P}} A. \\ \forall t_1 :_{\mathbb{E}} \text{msg.} \\ \dots \\ \forall t_i :_{\mathbb{E}} \text{msg.} \\ \forall n_1 :_{\mathbb{E}} \text{nonce.} \\ \dots \\ \forall n_i :_{\mathbb{E}} \text{nonce.} \\ \forall n :_{\mathbb{E}} \text{nonce.} \end{array} \quad \text{D}_1(\mathcal{P}(\#\{\!\{ t_1 \}\!\}_k^{n_1}, \dots, \#\{\!\{ t_i \}\!\}_k^{n_i})) \rightarrow \text{D}_1(\#\{\!\{ \mathcal{S}(t_1, \dots, t_i) \}\!\}_k^n) \right)_{(\text{L-CAP-PENCH-1})}^!$$

$$\left(\begin{array}{l} \forall A :_{\mathbb{E}} \text{principal.} \\ \forall k :_{\mathbb{E}} \text{pubK}^P A. \\ \forall t_1 :_{\mathbb{E}} \text{msg.} \\ \dots \\ \forall t_i :_{\mathbb{E}} \text{msg.} \\ \forall n_1 :_{\mathbb{E}} \text{nonce.} \\ \dots \\ \forall n_i :_{\mathbb{E}} \text{nonce.} \\ \forall n :_{\mathbb{E}} \text{nonce.} \end{array} \quad D_1(\#\{\mathcal{S}(t_1, \dots, t_i)\}_k^n) \rightarrow D_1(\mathcal{P}(\#\{t_1\}_k^{n_1}, \dots, \#\{t_i\}_k^{n_i})) \right) \quad \text{(L-CAP-PENCH-2)}$$

$$\left(\begin{array}{l} \forall A :_{\mathbb{E}} \text{principal.} \\ \forall k :_{\mathbb{E}} \text{pubK}^P A. \\ \forall t_1 :_{\mathbb{E}} \text{msg.} \\ \dots \\ \forall t_i :_{\mathbb{E}} \text{msg.} \\ \forall n_1 :_{\mathbb{E}} \text{nonce.} \\ \dots \\ \forall n_i :_{\mathbb{E}} \text{nonce.} \\ \forall n :_{\mathbb{E}} \text{nonce.} \end{array} \quad C_1(\mathcal{P}(\#\{t_1\}_k^{n_1}, \dots, \#\{t_i\}_k^{n_i})) \rightarrow C_1(\#\{\mathcal{S}(t_1, \dots, t_i)\}_k^n) \right) \quad \text{(L-CAP-PENCH-3)}$$

$$\left(\begin{array}{l} \forall A :_{\mathbb{E}} \text{principal.} \\ \forall k :_{\mathbb{E}} \text{pubK}^P A. \\ \forall t_1 :_{\mathbb{E}} \text{msg.} \\ \dots \\ \forall t_i :_{\mathbb{E}} \text{msg.} \\ \forall n_1 :_{\mathbb{E}} \text{nonce.} \\ \dots \\ \forall n_i :_{\mathbb{E}} \text{nonce.} \\ \forall n :_{\mathbb{E}} \text{nonce.} \end{array} \quad C_1(\#\{\mathcal{S}(t_1, \dots, t_i)\}_k^n) \rightarrow C_1(\mathcal{P}(\#\{t_1\}_k^{n_1}, \dots, \#\{t_i\}_k^{n_i})) \right) \quad \text{(L-CAP-PENCH-4)}$$

The important point here is that the implicit assumption found in the message inference mechanism of Typed MSR that each capability of the Dolev-Yao intruder can either be used constructively or destructively is not valid in the case of the capability to unblind signatures and to apply the homomorphic encryption properties. This is because: (i) in the destruction phase it may be beneficial to use these capabilities in order to continue the search for atomic messages, and (ii) in the construction phase it is not (always) the case that using these capabilities means that the intruder undoes his own work.

Chapter 9

Summary and Conclusions

In this thesis we demonstrate how, starting from an informal description of a privacy-preserving protocol in natural language, one may use a modified and extended version of the Typed MSR language to create a formal specification of this protocol, typed in a linkability-oriented type system, and then use this specification to reach an implementation of this protocol in Jif, in such a way that linkability vulnerabilities can be detected with a mixture of static and runtime checks. This, of course, assumes that a suitable abstraction of the protocol’s cryptographic primitives is included in our proposed Typed MSR extensions.

More specifically, in Chapters 2 and 3 we present our proposed modifications and extensions to Typed MSR, in Chapter 4 we show how protocols described in formal language can be formally specified, in Chapter 5 we introduce the linkability-oriented type system that type-checks the protocol’s specification, and in Chapters 6, 7 and 8 we present how this specification can lead to a Jif implementation that detects linkability.

More analytically, in this thesis we make the following points:

1. [Section 2.7] We argue that the Typed MSR specification language should not model encryption (both symmetric and asymmetric) as deterministic, not only because this is not how encryption is usually employed in practice, but more importantly because this makes the language particularly unsuitable for the specification of privacy-preserving protocols.
2. [Chapter 3] We propose high-level abstractions for blind signatures, commitments, zero-knowledge proofs and homomorphic encryption.
3. [Section 3.1.4] We argue that making the zero-knowledge proofs non-interactive results in having protocol specifications that are easier to reason about and to implement.
4. [Chapter 4] We demonstrate that augmenting the standard language with the new message constructors make it expressive enough to specify privacy-preserving protocols, such as e-voting protocols.

5. [Chapter 5] We argue that a simple type system that exists in parallel with Typed MSR's type system is able to enforce basic message well-formedness, and to track message linkability.
6. [Chapter 6] We propose an extended version of the original Dolev-Yao intruder based on our linkability-oriented type system and the new message constructors.
7. [Section 6.3] We argue that the extended version of the Dolev-Yao intruder creates a formal environment in which linkability attacks on privacy-preserving protocols may be demonstrated.
8. [Chapter 8] We illustrate how the above may be incorporated into a framework built on the Jif language, in such a way that linkability vulnerabilities in protocol implementations may be detected with a mixture of static and runtime checks.
9. [Section 8.2] We argue that some privacy-related properties are not always easy or even possible to prove statically, but need to be checked dynamically during the protocol's execution. Although such an attempt raises the effort needed to implement the protocol, it also has the advantage that it is the actual protocol's implementation that is being verified, not a possibly flawed abstraction.
10. [Section 8.3] We argue that the implicit assumption found in the message inference mechanism of Typed MSR that each capability of the Dolev-Yao intruder can either be used constructively or destructively is not valid in the case of our introduced message constructors.

Chapter 10

Future Work

In our view, further work based on the results obtained from this thesis can follow these major directions:

1. The implementation of the underlying cryptography of the linkability-oriented cryptographic framework, as described in Section 8.1. We need to be able to convert `Message` objects to and from byte arrays. Some of the cryptography, like encryption and digital signatures, can easily be implemented using suitable providers for the Java Cryptography Architecture (JCA). The rest of the cryptography, like blind signatures and zero-knowledge proofs, will be harder to implement, although there is a lot of source code available on the internet to make the task easier.
2. The integration of the cryptographic framework with a theorem prover, fed with the axioms and rules of Section 8.3, and the implementation of the runtime linkability checks of Section 8.2. The first can be achieved using, for example, the Java Theorem Prover (JTP), developed by Gleb Frank at Stanford University. The later, given that the theorem prover is in place, is just a matter of invoking it and storing the results. This task may be challenging performance-wise.
3. The creation of a library of cryptographic primitives, abstracted as Typed MSR messages, and implemented in our Jif cryptographic framework. This will significantly increase the domain of protocols this work will apply to.
4. The extension of this work to make it able to deal with additional privacy concerns. This will probably require comprehensive modifications.

Bibliography

- [1] Data protection act. UK legislation. 1984.
- [2] International covenant on civil and political rights. January 1997.
- [3] Alessandro Acquisti. Receipt-free homomorphic elections and write-in ballots. *Technical Report 2004/105, International Association for Cryptologic Research*, May 2004.
- [4] Aristotle. A treatise on government. 322 BC.
- [5] D. Aspinall and A. Compagnoni. Subtyping dependent types. In E. Clarke, editor, *Proceedings of the 11th Annual Symposium on Logic in Computer Science*, pages 86–97. IEEE Computer Society Press, July 1996.
- [6] Theodoros Balopoulos, Stefanos Gritzalis, and Sokratis K. Katsikas. Specifying and implementing privacy-preserving protocols. Under revision, *International Journal of Information Security*, Springer.
- [7] Theodoros Balopoulos, Stefanos Gritzalis, and Sokratis K. Katsikas. An extension of Typed MSR for specifying esoteric protocols and their Dolev-Yao intruder. In D. Chadwick B. Preneel, editor, *Proceedings of the CMS'2004 IFIP TC6/TC11 8th International Conference on Communications and Multimedia Security*, volume 175, pages 209–221. Springer, September 2004.
- [8] Theodoros Balopoulos, Stefanos Gritzalis, and Sokratis K. Katsikas. Specifying electronic voting protocols in Typed MSR. In S. De Capitani di Vimercati and R. Dingle-dine, editors, *Proceedings of the 2005 Computer and Communications Security Conference — Workshop on Privacy in the Electronic Society*, pages 35–39. ACM Press, November 2005.
- [9] Theodoros Balopoulos, Stefanos Gritzalis, and Sokratis K. Katsikas. Specifying privacy-preserving protocols in Typed MSR. *Computer Standards & Interfaces*, 27(5):501–512, June 2005.
- [10] Theodoros Balopoulos and Stephanos Gritzalis. Towards a logic of privacy-preserving selective disclosure credential protocols. In J. Lopez and G. Pernul, editors, *Proceedings of the DEXA 2003 — TRUSTBUS'03 2nd International Workshop on Trust and*

- Privacy in Digital Business*, pages 396–401, Prague, Czech Republic, September 2003. IEEE Computer Society Press.
- [11] Fabrice Boudot. Efficient proofs that a committed number lies in an interval. In *EUROCRYPT*, pages 431–444, 2000.
- [12] Louis Brandeis and Samuel Warren. The right to privacy. In *Harvard Law Review*, volume 4, December 1890.
- [13] Michael Burrows, Martín Abadi, and Roger Needham. A logic of authentication. *ACM Transactions on Computer Systems*, 8(1):18–36, Feb 1990.
- [14] Iliano Cervesato. Typed Multiset Rewriting Specifications of Security Protocols. In A. Seda, editor, *First Irish Conference on the Mathematical Foundations of Computer Science and Information Technology — MFCSIT’00*, pages 1–43, Cork, Ireland, 19–21 July 2000. Elsevier ENTCS 40.
- [15] Iliano Cervesato. Typed MSR: Syntax and Examples. In V.I. Gorodetski, V.A. Skormin, and L.J. Popyack, editors, *First International Workshop on Mathematical Methods, Models and Architectures for Computer Networks Security — MMM’01*, pages 159–177, St. Petersburg, Russia, 21–23 May 2001. Springer-Verlag LNCS 2052.
- [16] David Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 4(2), February 1981.
- [17] David Chaum. Security without identification: transaction systems to make big brother obsolete. *Communications of the Association for Computing Machinery*, 28(10):1030–1044, October 1985.
- [18] Claudia Diaz, Stefaan Seys, Joris Claessens, and Bart Preneel. Towards measuring anonymity. In Roger Dingledine and Paul Syverson, editors, *Proceedings of Privacy Enhancing Technologies Workshop (PET 2002)*, volume LNCS 2482. Springer-Verlag, April 2002.
- [19] D. Dolev and A. C. Yao. On the security of public key protocols. *IEEE Transactions on Information Theory*, 2(29):198–208, 1983.
- [20] F. Javier Thayer Fabrega, Jonathan C. Herzog, and Joshua D. Guttman. Strand spaces: Why is a security protocol correct? In *Proceedings of the IEEE Symposium on Security and Privacy*, May 1998.
- [21] Uriel Feige, Amos Fiat, and Adi Shamir. Zero knowledge proofs of identity. In *Proceedings of the 19th ACM Symposium on Theory of Computing*, pages 210–217, May 1987.

- [22] David M. Goldschlag, Michael G. Reed, and Paul F. Syverson. Hiding routing information. In R. Anderson, editor, *Proceedings of the 1st International Workshop on Information Hiding*, volume 1174 of Lecture Notes in Computer Science, pages 137–150. Springer-Verlag, 1996.
- [23] Stefanos Gritzalis, Diomidis Spinellis, and Panagiotis Georgiadis. Security protocols over open networks and distributed systems: Formal methods for their analysis, design, and verification. In *Computer Communications*, volume 22, pages 697–709. Elsevier Science, May 1999.
- [24] C. A. R. Hoare. Communicating sequential processes. In *j-CACM*, pages 21(8):666–677, August 1978.
- [25] Jason E. Holt and Kent E. Seamons. Selective disclosure credential sets. *Accessible as <http://citeseer.nj.nec.com/541329.html>*, 2002.
- [26] W. Marrero, E. M. Clarke, and S. Jha. Model checking for security protocols. In *Proceedings of the 1997 DIMACS Workshop on Design and Formal Verification of Security Protocols*, 1997.
- [27] Robin Milner. A calculus of communicating systems. In *Lecture Notes in Computer Science*, volume 92, 1980.
- [28] Robin Milner. *Communicating and mobile systems: the π -calculus*. Cambridge University Press, 1999.
- [29] Andrew C. Myers. Practical mostly-static information flow control. In *Proceedings of the 26th ACM Symposium on Principles of Programming Languages (POPL)*, pages 228–241, January 1999.
- [30] Andrew C. Myers and Barbara Liskov. A decentralized model for information flow control. In *Proceedings of the 16th ACM Symposium on Operating Systems Principles (SOSP)*, pages 129–142, October 1997.
- [31] Andrew C. Myers and Barbara Liskov. Complete, safe information flow with decentralized labels. In *Proceedings of the 1998 IEEE Symposium on Security and Privacy*, pages 186–197, May 1998.
- [32] United Nations. Universal Declaration of Human Rights. 1948.
- [33] P. Paillier. Public-key cryptosystems based on discrete logarithms residues. In *Advances in Cryptology - Eurocrypt '99*, pages 223–238. Springer-Verlag LNCS 1592, 1999.
- [34] Andreas Pfitzmann and Marit Köhntopp. Anonymity, unlinkability, undetectability, unobservability, pseudonymity, and identity management — a consolidated proposal for terminology. Draft, version 0.29. *Accessible as http://dud.inf.tu-dresden.de/Anon_Terminology.shtml*, July 2007.

- [35] Steve Schneider and Abraham Sidiropoulos. CSP and anonymity. In *ESORICS*, pages 198–218, 1996.
- [36] Andrei Serjantov and George Danezis. Towards an information theoretic metric for anonymity. In Roger Dingledine and Paul Syverson, editors, *Proceedings of Privacy Enhancing Technologies Workshop (PET 2002)*, volume LNCS 2482. Springer-Verlag, April 2002.
- [37] Andrei Serjantov, Roger Dingledine, and Paul Syverson. From a trickle to a flood: Active attacks on several mix types. In Fabien Petitcolas, editor, *Proceedings of Information Hiding Workshop (IH 2002)*, volume LNCS 2578. Springer-Verlag, October 2002.
- [38] Paul Syverson and Iliano Cervesato. The logic of authentication protocols. In *Foundations of Security Analysis and Design*, volume 2171 of *Tutorial Lectures*. Springer, 2001.
- [39] Paul Syverson, Catherine Meadows, and Iliano Cervesato. Dolev-Yao is no better than Machiavelli. In P. Degano, editor, *First Workshop on Issues in the Theory of Security — WITS'00*, pages 87–92, July 2000.