

# Type-checking Zero-knowledge

Michael Backes<sup>1,2</sup>, Cătălin Hrițcu<sup>1</sup>, and Matteo Maffei<sup>1</sup>

<sup>1</sup> Saarland University, Saarbrücken, Germany

<sup>2</sup> MPI-SWS

## Abstract

This paper presents the first type system for statically analyzing security protocols that are based on zero-knowledge proofs. We show how certain properties offered by zero-knowledge proofs can be characterized in terms of authorization policies and statically enforced by a type system. The analysis is modular and compositional, and provides security proofs for an unbounded number of protocol executions. We develop a new type-checker that conducts the analysis in a fully automated manner. We exemplify the applicability of our technique to real-world protocols by verifying the authenticity properties of the Direct Anonymous Attestation (DAA) protocol. The analysis of DAA takes less than three seconds.

## 1 Introduction

The design of cryptographic protocols is notoriously difficult and error-prone, and manual security proofs for such protocols are difficult to do. The multitude of attacks on existing cryptographic protocols reported lately (e.g., [35, 15, 24, 18]) demonstrate the need for formalizing the intended security properties and developing automated techniques for automatically verifying these properties. Logic-based authorization policies constitute a well-established and expressive framework for describing a wide range of security properties of cryptographic protocols, varying from authenticity properties to access control policies [2]. Type systems are particularly salient tools to statically and automatically enforce authorization policies on abstract protocol specifications [25, 26] and on concrete protocol implementations [13]. Type systems require little human effort and provide security proofs for an unbounded number of protocol executions. Furthermore, the analysis is modular, compositional, and usually guaranteed to terminate.

One of the central challenges in the verification of authorization policies for modern applications is the expressiveness of the analysis and its ability to statically characterize the security properties guaranteed by complex cryptographic operations. For instance, current analysis techniques support traditional cryptographic primitives such as encryption and digital signatures, but until recently [11] they could not cope with the most prominent and innovative modern cryptographic primitive: zero-knowledge proofs [27]

A zero-knowledge proof combines two seemingly contradictory properties. First, it constitutes a proof of a statement that cannot be forged, i.e., it is impossible, or at least computationally infeasible, to produce a zero-knowledge

proof of a wrong statement. Second, a zero-knowledge proof does not reveal any information besides the bare fact that the statement is valid. Early general-purpose zero-knowledge proofs were primarily designed for showing the existence of such proofs for the class of statements under consideration. These proofs were very inefficient and consequently of only limited use in practical applications. The recent advent of efficient zero-knowledge proofs for special classes of statements is rapidly changing this scenario. The unique security features that zero-knowledge proofs offer combined with the possibility to efficiently implement some of these proofs non-interactively have paved the way for their deployment in modern cryptographic applications. In fact, many anonymity protocols (e.g., Direct Anonymous Attestation (DAA) [16] and Pseudo Trust [34]) and electronic voting protocols (e.g., [32, 5, 19]) heavily rely on zero-knowledge proofs.

Statically analyzing protocols that use zero-knowledge proofs is conceptually and technically challenging. Zero-knowledge proofs provide security properties that go far beyond the traditional and well-understood secrecy and authenticity properties. For instance, zero-knowledge proofs may guarantee authentication yet preserve the anonymity of protocol participants, as in the Pseudo Trust protocol [34], or they may prove the reception of a certificate from a trusted server without revealing the actual content, as in the Direct Anonymous Attestation (DAA) protocol [16]. The existing techniques for type-checking cryptographic protocols do not apply to zero-knowledge proofs, since these techniques usually rely on the type of keys for typing cryptographic messages. Zero-knowledge proofs, however, do not depend on any key infrastructure.

## 1.1 Our Contributions

This paper presents the first type system for statically analyzing the security of protocols based on non-interactive zero-knowledge proofs. We show how the safety properties guaranteed by zero-knowledge proofs can be formulated in terms of authorization policies and statically enforced by a type system. Our approach extends the state-of-the-art type system for authorization policies proposed by Fournet et al. [26]. Zero-knowledge proofs are given dependent types where the messages kept secret by the proof are existentially quantified in the logic. The fundamental idea is to express zero-knowledge statements as logical formulas and to define the type of zero-knowledge proofs using these formulas. The user still has the possibility to extend such types with additional logical formulas describing protocol-dependent security properties.

We develop a new type-checker that automates the analysis. The tool verifies that protocol specifications are well-typed and relies on the first-order logic automated theorem prover SPASS [36] to discharge proof obligations. The analysis is modular and compositional, and provides security proofs for an unbounded number of protocol executions.

We exemplify the applicability of our technique to real-world protocols by verifying the Direct Anonymous Attestation protocol (DAA) [16]. We formalize the authenticity properties of this protocol in terms of authorization policies and we apply our type system to statically verify them. Our type-checker analyzed this sophisticated protocol in less than three seconds. This promising result indicates that our static analysis technique has the potential to scale up to industrial-size protocols.

## 1.2 Related Work

Dating back to the seminal work by Abadi on secrecy by typing [1], type systems were successfully used to analyze a wide range of security properties of cryptographic protocols, ranging from authenticity properties [28, 30, 31, 17, 7], to security despite compromised participants [29, 17, 26, 20], to authorization policies [25, 26, 13]. As mentioned before, type-checking is efficient, modular, compositional, and usually guaranteed to terminate. None of the existing type systems is, however, capable of dealing with zero-knowledge proofs.

To the best of our knowledge, ProVerif [14, 4] is the only automatic tool that has been applied to the analysis of protocols that use zero-knowledge proofs [11, 8, 22]. This tool is based on Horn-clause resolution and can analyze trace-based security properties as well as behavioral properties. The analysis with ProVerif is not compositional and often has unpredictable termination behaviour, with seemingly harmless code changes leading to divergence. Also, as argued in [13], type systems scale better to large protocols and more efficiently analyze protocol implementations. In terms of expressiveness, ProVerif can deal with behavioral properties that are generally out of scope for current type systems (e.g., privacy and coercion-resistance in electronic-voting protocols [21, 8]), but is restricted to cryptographic primitives that can be expressed as a confluent rewriting system. Our analysis does not pose any constraint on the semantics of cryptographic primitives and, as opposed to ProVerif, can deal with authorization policies using arbitrary logical structure (e.g., arbitrarily nested quantifiers).

## 1.3 Outline

Section 2 illustrates our approach on a simple protocol for anonymous trust. Section 3 describes the process calculus we use to model security protocols that use zero-knowledge proofs. Section 4 presents our type system for zero-knowledge. Section 5 discusses the implementation of our type-checker and the experimental evaluation of our technique. Due to space constraints, we defer all proofs, some of the technical details of our type system, the verification of DAA, and a discussion about security despite compromise to an extended version of the paper [9].

# 2 Illustrative Example

This section introduces the types of zero-knowledge proofs and highlights the fundamental ideas of our type system, which will be elaborated in more detail in the following sections. As a running example, we consider a simple protocol for anonymous trust that is inspired by the Pseudo Trust protocol proposed by Lu et al. in [34]. The goal of this protocol is to allow parties to exchange data proving each other's trust level while preserving anonymity. These two seemingly conflicting requirements are met by an authentication scheme based on zero-knowledge proofs.

## 2.1 A Protocol for Anonymous Trust

Each party has a public pseudonym, which is the hash of a secret  $m_s$ . This pseudonym replaces the actual identity of the party in the protocol. An arbitrary trust-management system like EigenTrust [33] or XenoTrust [23] can be used to

certify the trust level of each pseudonym. Whenever a party (*prover*) wants to send a message  $m_p$  to another party (*verifier*), she has to bind  $m_p$  to her own pseudonym  $\text{hash}(m_s)$  and, in order to avoid impersonation, she has to prove the knowledge of  $m_s$  without revealing it. This authentication scheme is realized by a non-interactive zero-knowledge proof that is sent from the prover to the verifier. The zero-knowledge proof guarantees that the prover knows  $m_s$  and additionally provides the non-malleability of  $m_p$ , i.e., changing  $m_p$  requires the adversary to redo the proof and thus to know  $m_s$ . The goal of this protocol is allowing the verifier to associate the trust level of the prover to  $m_p$ .

## 2.2 Zero-knowledge Proofs and Authorization Policies

Following [11], we represent our zero-knowledge proof by the following applied pi-calculus term:  $\text{zk}_{1,2,\beta_1=\text{hash}(\alpha_1)}(m_s; \text{hash}(m_s), m_p)$ . In our setting, a zero-knowledge proof  $\text{zk}_{n,m,S}(N_1, \dots, N_n; M_1, \dots, M_m)$  has  $n + m$  arguments. The first  $n$  arguments  $N_1, \dots, N_n$  form the *private component* of the proof and are kept secret ( $m_s$  in the example), while the other  $m$  arguments  $M_1, \dots, M_m$  form the *public component* and are revealed to the verifier ( $\text{hash}(m_s)$  and  $m_p$  in the example). The *statement*  $S$  of the zero-knowledge proof is a Boolean formula over the placeholders  $\alpha_i$  and  $\beta_j$  (with  $i \in [1, n]$  and  $j \in [1, m]$ ), which stand for the argument  $N_i$  in the private component and the argument  $M_j$  in the public component, respectively. The verification of a zero-knowledge proof succeeds if and only if the statement obtained by replacing the place-holders by the corresponding private and public arguments holds true.

In order to express the security property guaranteed by this protocol as an authorization policy, we decorate the security-related protocol events as follows:

$$\begin{array}{ccc}
 \text{assume Authenticate}(m_p, m_s) & & \text{assume Trust}(pseudo, k) \\
 P \xrightarrow{\text{zk}_{1,2,\beta_1=\text{hash}(\alpha_1)}(m_s; \text{hash}(m_s), m_p)} & & V \\
 & & \text{assert Associate}(m_p, k)
 \end{array}$$

Before generating the zero-knowledge proof to authenticate  $m_p$  using the knowledge of the secret  $m_s$ , the prover  $P$  assumes  $\text{Authenticate}(m_p, m_s)$ . Before receiving the zero-knowledge proof, the verifier  $V$  knows that the trust level associated to the pseudonym  $pseudo$  is  $k$  and assumes  $\text{Trust}(pseudo, k)$ . As discussed before, this information can be obtained by means of an external trust-management system. The verifier receives and verifies the zero-knowledge proof, checking that the first argument in the public component is  $pseudo$ . This guarantees that the prover knows the secret  $m_s$  such that  $pseudo = \text{hash}(m_s)$  and allows the verifier to assert  $\text{Associate}(m_p, k)$ . The authorization policy can be expressed as follows:

$$Policy := \forall m_p, k, m_s. (\text{Authenticate}(m_p, m_s) \wedge \text{Trust}(\text{hash}(m_s), k)) \Rightarrow \text{Associate}(m_p, k)$$

Since this is the only place where the  $\text{Associate}$  predicate occurs, the verifier is allowed to give  $m_p$  trust level  $k$  (assertion  $\text{Associate}(m_p, k)$ ) only if the prover wants to authenticate  $m_p$  (assumption  $\text{Authenticate}(m_p, m_s)$ ) and knows some  $m_s$  for which the trust level of the pseudonym  $\text{hash}(m_s)$  is  $k$  (assumption  $\text{Trust}(\text{hash}(m_s), k)$ ). Intuitively, everything that is not explicitly allowed by the authorization policy is prohibited.

### 2.3 The Type of Zero-knowledge Proofs

To illustrate our technique, let us consider the type associated to the zero-knowledge proof  $\text{zk}_{1,2,\beta_1=\text{hash}(\alpha_1)}(m_s; \text{hash}(m_s), m_p)$ :

$$\text{ZKProof}_{1,2,\beta_1=\text{hash}(\alpha_1)}(\langle y_1 : \text{Hash}(\text{Private}), y_2 : \text{Un} \rangle \\ \{ \exists x. y_1 = \text{hash}(x) \wedge \text{Authenticate}(y_2, x) \})$$

This dependent type indicates that the public component is composed of two messages. The first message  $y_1$  is of type  $\text{Hash}(\text{Private})$ , i.e., it is the hash of a secret message. The type  $\text{Private}$  describes messages that are not known to the adversary. The second message  $y_2$  is of type  $\text{Un}$  (untrusted), i.e., it may come from and be sent to the adversary. The logical formula  $\exists x. y_1 = \text{hash}(x) \wedge \text{Authenticate}(y_2, x)$  says that  $y_1$  is the hash of some secret  $x$  such that  $\text{Authenticate}(y_2, x)$  has been assumed by the prover. This formula contains an equality constraint on the structure of messages as well as a logical predicate.

After the verification of the zero-knowledge proof, the verifier can safely assume that the formula holds true. The constraint  $\exists x. y_1 = \text{hash}(x)$  is guaranteed by the semantics of the zero-knowledge proof, while the assumption  $\text{Authenticate}(y_2, x)$  is enforced by our type system. The verifier can thus logically derive  $\exists m_s. \text{pseudo} = \text{hash}(m_s) \wedge \text{Authenticate}(m_p, m_s) \wedge \text{Trust}(\text{pseudo}, k)$ . By a standard logical property that allows replacing equals by equals, the authorization policy allows the verifier to derive  $\text{Associate}(m_p, k)$ .

## 3 Calculus

We consider a variant of the applied pi-calculus with constructors and destructors similar to the one in [3], and we extend it with zero-knowledge proofs. Following [26, 13], the calculus also includes special operators to assume and assert logical formulas. This section overviews the syntax and semantics of the calculus.

### 3.1 Constructors and Terms

*Constructors* are function symbols that are used to build terms. The set of constructors includes  $\text{pk}$  that yields the public encryption key corresponding to a decryption key;  $\text{enc}$  for public-key encryption;  $\text{vk}$  that yields the verification key corresponding to a signing key;  $\text{sign}$  for digital signatures; and  $\text{hash}$  for hashes. The constant  $\text{true}$  represents the respective Boolean value, and has its canonical meaning in the authorization logic.

The set of *terms* (ranged over by  $K, L, M$  and  $N$ ) is the free algebra built from names ( $a, b, c, m, n$ , and  $k$ ), variables ( $x, y, z, v$ , and  $w$ ), tuples ( $\langle M_1, \dots, M_n \rangle$ ), and constructors applied to other terms ( $f(M_1, \dots, M_n)$ ). We let  $u$  range over both names and variables.

### 3.2 Destructors

*Destructors* are partial functions that processes can apply to terms, and are ranged over by  $g$ . The semantics of destructors is specified by the reduction relation  $\Downarrow$ : given the terms  $M_1, \dots, M_n$  as arguments, the destructor  $g$  can either succeed and provide a term  $N$  as a result (which we denote as

$g(M_1, \dots, M_n) \Downarrow N$ ) or it can fail (denoted as  $g(M_1, \dots, M_n) \not\Downarrow$ ). The destructor  $\text{dec}$  decrypts an encrypted message given the corresponding decryption key ( $\text{dec}(\text{enc}(M, \text{pk}(K)), K) \Downarrow M$ ). The check destructor checks a signed message using a verification key, and if this succeeds returns the message without the signature ( $\text{check}(\text{sign}(M, K), \text{vk}(K)) \Downarrow M$ ). The application of  $\text{eq}$  succeeds, yielding the constant  $\text{true}$ , if the two arguments are syntactically the same ( $\text{eq}(M, M) \Downarrow \text{true}$ ). The destructors  $\wedge$  and  $\vee$  model conjunctions and disjunctions, respectively ( $\wedge(\text{true}, \text{true}) \Downarrow \text{true}$ ,  $\vee(M, \text{true}) \Downarrow \text{true}$ , and  $\vee(\text{true}, M) \Downarrow \text{true}$ ).

### 3.3 Representing Zero-knowledge Proofs

**Constructing Zero-knowledge Proofs.** In a very similar way to what is proposed in [11], a non-interactive zero-knowledge proof of a statement  $S$  is represented as a term of the form  $\text{zk}_{n,m,S}(N_1, \dots, N_n; M_1, \dots, M_m)$ , where  $N_1, \dots, N_n$  and  $M_1, \dots, M_m$  are two sequences of terms. The proof keeps the terms in  $N_1, \dots, N_n$  secret, while the terms  $M_1, \dots, M_m$  are revealed. For clarity we will use semicolons to separate the secret terms from the public ones, and we will often write  $\tilde{N}$  instead of  $N_1, \dots, N_n$  if  $n$  is clear from the context.

**Statements.** In order to express a wide class of zero-knowledge proofs, comprising for instance proofs of signature verifications and decryptions, we need to use destructors inside logical formulas. Since destructors cannot occur inside terms we need to define a larger class of objects, called *statements*, that also contains destructors. The set of statements (ranged over by  $S$ ) is the free algebra built from names, variables, the placeholders  $\alpha_i$  and  $\beta_j$ , as well as tuples, constructors and destructors applied to other statements. It is easy to see that all terms are also statements. For clarity we distinguish an actual destructor  $g$  from its counterpart used within statements by writing the latter as  $g^\#$ . The statement  $S$  used in a term  $\text{zk}_{n,m,S}(N_1, \dots, N_n; M_1, \dots, M_m)$  is called an  $(n, m)$ -*statement*. It does not contain names or variables, and uses the placeholders  $\alpha_i$  and  $\beta_j$ , with  $i \in [1, n]$  and  $j \in [1, m]$ , to refer to the secret terms  $N_i$  and public terms  $M_j$ . For instance, the zero-knowledge term

$$\text{zk}_{1,2,\text{eq}^\#(\beta_1, \text{dec}^\#(\text{enc}(\beta_1, \beta_2), \alpha_1))}(k ; m, \text{pk}(k))$$

proves the knowledge of the decryption key  $k$  corresponding to the public encryption key  $\text{pk}(k)$ . More precisely, the statement reads: “There exists a secret key  $k$  such that the decryption of the ciphertext  $\text{enc}(m, \text{pk}(k))$  with this key yields  $m$ ”. As mentioned before,  $m$  and  $\text{pk}(k)$  are revealed by the proof while  $k$  is kept secret.

**Verifying Zero-knowledge Proofs.** The destructor  $\text{ver}_{n,m,l,S}$  verifies the validity of a zero-knowledge proof. It takes as arguments a proof together with  $l$  terms that are matched against the first  $l$  arguments in the public component of the proof. If the proof is valid, then  $\text{ver}_{n,m,l,S}$  returns the other  $m - l$  public arguments. A proof is valid if and only if the statement obtained by substituting all  $\alpha_i$ 's and  $\beta_j$ 's in  $S$  with the corresponding values  $N_i$  and  $M_j$  evaluates to  $\text{true}$ . This is formalized as follows:

$$\begin{aligned} \text{ver}_{n,m,l,S}(\text{zk}_{n,m,S}(\tilde{N}; M_1, \dots, M_l, \dots, M_m), M_1, \dots, M_l) \Downarrow \langle M_{l+1}, \dots, M_m \rangle, \\ \text{iff } S\{\tilde{N}/\tilde{\alpha}\}\{\tilde{M}/\tilde{\beta}\} \Downarrow_{\#} \text{true} \end{aligned}$$

The evaluation relation  $S \Downarrow_{\#} M$  is defined in terms of the reduction rules for the other destructors<sup>1</sup>. For example  $g^{\#}(S_1, \dots, S_n) \Downarrow_{\#} M$  if  $\forall i \in [1, n]. S_i \Downarrow_{\#} M_i$  and  $g(M_1, \dots, M_n) \Downarrow M$ . The  $\Downarrow_{\#}$  relation can in fact be seen as the extension of the  $\Downarrow$  relation to statements. For instance, in the protocol from Section 2 we have

$$\text{ver}_{1,2,1,\beta_1=\text{hash}(\alpha_1)}(\text{zk}_{1,2,\beta_1=\text{hash}(\alpha_1)}(m_s; \text{hash}(m_s), m_p), \text{pseudo}) \Downarrow \langle m_p \rangle$$

since  $(\beta_1 = \text{hash}(\alpha_1))\{m_s/\alpha_1\}\{\text{hash}(m_s)/\beta_1\} \equiv (\text{hash}(m_s) = \text{hash}(m_s)) \Downarrow_{\#} \text{true}$ . We remark that defining the operational semantics of zero-knowledge verification using a destructor is much simpler than using an infinite equational theory [11], which needs to be compiled into a confluent rewrite system in order to make it suitable to ProVerif. Since we use a type-system to analyze protocols we can rely on destructors such as `ver` that are not directly expressible in ProVerif.

The destructor `publicm` yields the public component of a zero-knowledge proof (`publicm(zkn,m,s( $\tilde{N}$ ,  $\tilde{M}$ ))  $\Downarrow$   $\langle \tilde{M} \rangle$ ). Note that the private component is not revealed by any destructor, which intuitively guarantees the zero-knowledge property of the proofs.`

### 3.4 Processes

Processes are essentially the same as in [26]. The process `out(M, N).P` outputs message  $N$  on channel  $M$  and then behaves as  $P$ ; the process `in(M, x).P` receives a message  $N$  from channel  $M$  and then behaves as  $P\{N/x\}$ ; the process `!in(M, x).P` behaves as an unbounded number of copies of `in(M, x).P` executed in parallel; `new a : T.P` generates a fresh name  $a$  of type  $T$  and then behaves as  $P$ ;  $P \mid Q$  behaves as  $P$  executed in parallel with  $Q$ ; `0` is a process that does nothing; let  $x = g(\tilde{M})$  then  $P$  else  $Q$  applies the destructor  $g$  to the terms  $\tilde{M}$  and if this succeeds and produces the term  $N$  ( $g(\tilde{M}) \Downarrow N$ ) then the process behaves as  $P\{N/x\}$ , otherwise ( $g(\tilde{M}) \not\Downarrow$ ) it behaves as  $Q$ ; the process `let  $\langle x_1, \dots, x_n \rangle = M$  in  $P$`  splits the tuple  $M$  into its components.

The processes `assume C` and `assert C`, where  $C$  is a logical formula, are used to express authorization policies, and do not have any computational significance. Assumptions are used to mark security-related events in processes, such as the intention to authenticate message  $m_p$  by the party knowing the secret value  $m_s$  (`assume Authenticate( $m_p, m_s$ )`), and also to express global policies such as:

$$\text{assume } \forall m_p, k, m_s. (\text{Authenticate}(m_p, m_s) \wedge \text{Trust}(\text{hash}(m_s), k)) \Rightarrow \text{Associate}(m_p, k)$$

The scope of assumptions is global, i.e., once an assumption becomes active it affects all processes that run in parallel.

Assertions specify logical formulas that are supposed to be entailed at runtime by the currently active assumptions. For instance, in the protocol from Section 2 the verifier asserts that it can associate trust level  $k$  to the message  $m_p$  (`assert Associate( $m_p, k$ )`). In principle it might be possible to implement such assertions as (distributed!) dynamic checks. As in [26], we take a totally different approach here. Our type system guarantees statically that in well-typed protocols all asserted formulas are valid at runtime, even in the presence of an arbitrary adversary.

<sup>1</sup> This is not circular since the `ver` destructor cannot appear inside statements.

### 3.5 Authorization Logic

Our calculus and type system are largely independent of the exact choice of authorization logic. The logic is required to fulfill some standard properties, such as monotonicity, closure under substitution and allowing the replacement of equals by equals. Additionally, statements are not only used by zero-knowledge terms, but they also have a close connection with the formulas in our authorization logic. For this reason, we require that all statements are also formulas in the logic, and we assume that  $\text{eq}^\sharp$  corresponds to equality in the authorization logic (simply denoted by “=”), while  $\wedge^\sharp$  and  $\vee^\sharp$  correspond to conjunction and disjunction in the logic. Furthermore, we add axioms in the logic that correspond to the semantics of the destructors (e.g., for decryption we add the formula  $\forall m, k. \text{dec}^\sharp(\text{enc}(m, \text{pk}(k)), k) = m$  as an axiom). This ensures that if a statement  $S$  evaluates to a term  $M$  ( $S \Downarrow_\sharp M$ ) then  $S = M$  holds in the logic ( $\models S = M$ ). Under this assumption, from the semantics of the  $\text{ver}$  destructor we can immediately infer that if  $\text{ver}_{n,m,l,S}(\text{zk}_{n,m,S}(\tilde{N}, M_1, \dots, M_l, \dots, M_m), M_1, \dots, M_l) \Downarrow \langle M_{l+1}, \dots, M_m \rangle$  then  $\models S\{\tilde{N}/\tilde{\alpha}\}\{\tilde{M}/\tilde{\beta}\}$ , which captures the soundness of our construction for zero-knowledge.

In our implementation we consider first-order logic with equality as the authorization logic and we use the automated theorem prover SPASS [36] to discharge the proof obligations generated by our type system.

### 3.6 Notations and Conventions

Throughout the paper, we identify any phrase  $\phi$  of syntax up to consistent renaming of bound names and variables. We let  $fn(\phi)$  denote the set of free names in  $\phi$ ,  $fv(\phi)$  the set of free variables, and  $free(\phi)$  the set of free names and variables. We say that  $\phi$  is closed if it does not have any free variables. We write  $\phi\{\phi'/x\}$  for the outcome of the capture-avoiding substitution of  $\phi'$  for each free occurrence of  $x$  in  $\phi$ .

A context is a process with a hole where other processes can be plugged in. An evaluation context  $\mathcal{E}$  is a context of the form  $\mathcal{E} = \text{new } \tilde{a} : \tilde{T}.([\ ] \mid P)$  for some process  $P$ . We use  $\text{new } \tilde{a} : \tilde{T}$  to denote a sequence  $\text{new } a_1 : T_1 \dots \text{new } a_k : T_k$  of typed name restrictions and, for the sake of readability, we sometimes use  $\text{let } x := M \text{ in } P$  to denote  $P\{M/x\}$ .

### 3.7 Modeling the Protocol for Anonymous Trust

With this setup in place we can formally model the protocol for anonymous trust from Section 2 as follows:

$$\begin{aligned}
P &:= \text{new } m_p : \text{Un.}(\text{assume } \text{Authenticate}(m_p, m_s) \mid \\
&\quad \text{out}(c, \text{zk}_{1,2,\beta_1=\text{hash}(\alpha_1)}(m_s; \text{hash}(m_s), m_p))) \\
V &:= \text{assume } \text{TrustLevel}(pseudo, k) \mid \text{in}(c, w). \\
&\quad \text{let } x = \text{ver}_{1,2,1,\beta_1=\text{hash}(\alpha_1)}(w, pseudo) \text{ then} \\
&\quad \text{let } \langle y_p \rangle = x \text{ then assert } \text{Associate}(y_p, k) \\
PseudoTrust &:= \text{new } m_s : \text{Private.} \\
&\quad P \mid \text{let } pseudo := \text{hash}(m_s) \text{ in } V \mid \text{assume } Policy
\end{aligned}$$

### 3.8 Operational Semantics and Safety

As for the pi-calculus, the operational semantics of our calculus is defined in terms of *structural equivalence* ( $\equiv$ ) and *internal reduction* ( $\rightarrow$ ). Structural equivalence captures rearrangements of parallel compositions and restrictions. Internal reduction defines the semantics of communication and destructor application.

A process is safe if and only if all its assertions are satisfied in every protocol execution.

**Definition 3.1 (Safety)** *A closed process  $P$  is safe if and only if for every  $C$  and  $Q$  such that  $P \rightarrow^* \text{new } \tilde{a} : \tilde{T}.(\text{assert } C \mid Q)$ , there exists an evaluation context  $\mathcal{E} = \text{new } \tilde{b} : \tilde{U}.[ ] \mid Q'$  such that  $Q \equiv \mathcal{E}[\text{assume } C_1 \mid \dots \mid \text{assume } C_n]$ ,  $\text{fn}(C) \cap \tilde{b} = \emptyset$ , and we have that  $\{C_1, \dots, C_n\} \models C$ .*

A process is robustly safe if it is safe when run in parallel with an arbitrary opponent. As we will see, our type system guarantees that if a process is well-typed, then it is also robustly safe.

**Definition 3.2 (Opponent)** *A closed process is an opponent if it does not contain any `assert` and if the only type occurring therein is `Un`.*

**Definition 3.3 (Robust Safety)** *A closed process  $P$  is robustly safe if and only if  $P \mid O$  is safe for every opponent  $O$ .*

## 4 Type System

The type system presented in this paper extends a recent type system for statically enforcing authorization policies in distributed systems [26]. This basic type system is described in Sections 4.1 and 4.2. The most important novelty of our type system is the ability to reason about zero-knowledge proofs and Section 4.3 presents this in detail. Section 4.4 explains how our type system works on a fragment of the protocol for anonymous trust. Finally, Section 4.5 summarizes the security guarantees offered by the type system.

### 4.1 Basic Types

Our type system has the following types: `Private` is the type of messages that are not revealed to the adversary, while `Un` is the type of messages possibly known to the adversary; `Ch( $T$ )` is the type of channels carrying messages of type  $T$ ; As in [26, 13], tuples are given refinement types of the form  $\langle x_1 : T_1, \dots, x_n : T_n \rangle \{C\}$ . The formula  $C$  can depend on the variables  $x_1, \dots, x_n$ . For example, in the protocol of Section 2, the verification of the zero-knowledge proof yields the term  $\langle m_p \rangle$  of type  $\langle y : \text{Un} \rangle \{ \exists x. \text{pseudo} = \text{hash}(x) \wedge \text{Authenticate}(y, x) \}$ .

In addition to these base types, we also consider types for the different cryptographic primitives. For digital signatures, `SigKey( $T$ )` and `VerKey( $T$ )` denote the types of the signing and verification keys for messages of type  $T$ , while `Signed( $T$ )` is the type of signed messages of type  $T$ . We remark that a key of type `SigKey( $T$ )` can only be used to sign messages of type  $T$ , where the type  $T$  is in general annotated by the user. Similarly, `PubKey( $T$ )` and `PrivKey( $T$ )` denote the types of the public encryption keys and of the private decryption keys for messages of type  $T$ , while `PubEnc( $T$ )` is the type of a public-key encryption of a

message of type  $T$ . The type  $\text{Hash}(T)$  denotes the type of a hashed message of type  $T$ . In all these cases the type  $T$  is usually a refinement type conveying a logical formula. For instance,  $\text{SigKey}(\langle x : \text{Private} \rangle \{OK(x)\})$  is the type of keys that can be used to sign private messages  $M$  for which we statically know that  $OK(M)$  holds.

## 4.2 Typing Judgments

The type system relies on four typing judgments: well-formed environment ( $\Gamma \vdash \diamond$ ), subtyping ( $\Gamma \vdash T <: U$ ), term typing ( $\Gamma \vdash M : T$ ), and process typing ( $\Gamma \vdash P$ ), which are described in the following.

**Well-formed Environment.** The type system relies on a *typing environment*, which is a list containing name and variable bindings of the form  $u : T$ , together with formulas of the authorization logic. We denote the formulas in a typing environment  $\Gamma$  by  $\text{forms}(\Gamma)$ . Intuitively, these formulas constitute a safe approximation of the formulas assumed at run-time.

A typing environment is *well-formed*, written  $\Gamma \vdash \diamond$ , if no name or variable is bound more than once, and if all free names and variables inside the types and formulas appearing in the environment are bound beforehand. All the other typing judgments check that the environment they use is well-formed.

**Subtyping.** All messages sent to and received from an untrusted channel have type  $\text{Un}$ , since such channels are considered under the complete control of the adversary. However, a system in which only names and variables of type  $\text{Un}$  could be communicated over the untrusted network would be too restrictive to be useful. We therefore consider a subtyping relation on types, which allows a term of a subtype to be used in all contexts that require a term of a supertype. This preorder is used to compare types with the special type  $\text{Un}$ . In particular, we allow messages having a type  $T$  that is a subtype of  $\text{Un}$ , denoted  $T <: \text{Un}$ , to be sent over the untrusted network, and we say that the type  $T$  is *public* in this case. In addition, we allow messages of type  $\text{Un}$  that are received from the untrusted network to be used as messages of type  $U$ , provided that  $\text{Un} <: U$ , and in this case we say that type  $U$  is *tainted*.

For example, in our type system the types  $\text{PubKey}(T)$  and  $\text{VerKey}(T)$  are public, meaning that public-key encryption keys as well as signature verification keys can always be sent over an untrusted channel without compromising the security of the protocol. On the other hand,  $\text{PrivKey}(T)$  is public only if  $T$  is also public, since sending to the adversary a private key that decrypts confidential messages will most likely compromise the security of the protocol. Finally, the  $\text{Private}$  type we used in Section 2 is neither public nor tainted.

**Typing Terms.** The judgment  $\Gamma \vdash M : T$  checks that message  $M$  has type  $T$ . The type of variables and names is simply looked up in the typing environment. A tuple  $\langle M_1, \dots, M_n \rangle$  has the refinement type  $\langle x_1 : T_1, \dots, x_n : T_n \rangle \{C\}$  if each  $M_i$  has type  $T_i$  and if additionally the formulas in the typing environment entail  $C\{\widetilde{M}/\widetilde{x}\}$ . The other cases are as one would expect. For instance, if the message  $M$  has type  $T$  and the key  $K$  has type  $\text{SigKey}(T)$  then we can derive that the signature  $\text{sign}(M, K)$  has type  $\text{Signed}(T)$ .

**Typing Processes.** The typing judgment  $\Gamma \vdash P$  checks whether the process  $P$  is well-typed. As we will show in Section 4.5, this guarantees that  $P$  is secure

against an arbitrary adversary. The output process  $\text{out}(M, N).P$  is well-typed, if the term  $M$  has a channel type  $\text{Ch}(T)$ ,  $N$  is of type  $T$  and the process  $P$  is well-typed. For instance, this guarantees that the adversary can only receive messages of type  $\text{Un}$  at run-time, since it is initially given only channels of type  $\text{Ch}(\text{Un})$ . Similarly, the input process  $\text{in}(M, x).P$  is well-typed only if  $M$  has type  $\text{Ch}(T)$  and  $P$  is well-typed assuming  $x$  of type  $T$ . The process  $\text{new } n : T.P$  is well-typed if  $P$  is well-typed assuming  $n$  of type  $T$ . When type-checking a parallel composition  $P \mid Q$  the top-level assumptions in  $P$  can be added to the typing environment in which  $Q$  is typed, and the top-level assumptions in  $Q$  can be added to the environment in which  $P$  is typed. This ensures that assumptions have global scope.

The process  $\text{assert } C$  is well-typed in a typing environment  $\Gamma$ , only if  $\text{forms}(\Gamma) \models C$ . Intuitively, this guarantees the safety property of well-typed processes, since the formulas in  $\Gamma$  represent a safe approximation of the formulas assumed at run-time.

Type-checking the process  $\text{let } x = g(M_1, \dots, M_n) \text{ then } P \text{ else } Q$  differs significantly from [26]. As usual, we need to check whether the arguments  $M_1, \dots, M_n$  have the types required by the destructor, and obtain a new type  $T$  for the result of the destructor application. The continuation process  $P$  is, however, type-checked in a typing environment extended not only with the binding  $x : T$ , but also with the logical formula “ $x = g^\sharp(M_1, \dots, M_n)$ ”. This can be used for further reasoning in the logic. For instance, when checking that  $\Gamma \vdash \text{let } x = \text{check}(M, K) \text{ then } P$  we first need to ensure that  $M$  has type  $\text{Signed}(T)$  and  $K$  has type  $\text{VerKey}(T)$  for the same  $T$ . Then we can type-check the process  $P$  in the environment  $\Gamma, x : T, x = \text{check}^\sharp(M, K)$ . This treatment of destructors is simpler and more elegant than the one in [26], and appears to be similar to the typing rules for splitting and matching from [13].

When type-checking a process that splits a tuple  $\text{let } \langle x_1, \dots, x_n \rangle = M$  in  $P$  we need to ensure that  $M$  has a refinement type  $\langle y_1 : T_1, \dots, y_n : T_n \rangle \{C\}$ . Then the continuation process  $P$  is checked in an environment extended with the bindings  $x_1 : T_1, \dots, x_n : T_n$  and two logical formulas. First, we assume the formula  $\langle x_1, \dots, x_n \rangle = M$ , which is helpful if the same pair is split again in  $P$ . More important, we assume that the formula in the refinement type holds, after appropriate variable replacement, i.e., the environment is extended with  $C\{\tilde{x}/\tilde{y}\}$  when checking  $P$ . This is sound since the type system guarantees that when creating the tuple  $\langle \tilde{M} \rangle$  the formula  $C\{\tilde{M}/\tilde{y}\}$  is entailed.

As an example consider the final part of the verifier process in Section 3.7. The variable  $x$  holds a tuple with one component. The process splits this tuple and assigns the component to a variable  $y_p$ , and then asserts that  $y_p$  satisfies the predicate  $\text{Associate}(y_p, k)$ . The process  $\text{let } \langle y_p \rangle = x \text{ in assert Associate}(y_p, k)$  is type-checked in the environment  $\Gamma_1$  that contains:

$$\begin{aligned} & (\forall m_p, k, m_s. (\text{Authenticate}(m_p, m_s) \wedge \text{TrustLevel}(\text{hash}(m_s), k)) \Rightarrow \text{Associate}(m_p, k)), \\ & \text{TrustLevel}(\text{pseudo}, k), x : \langle w : \text{Un} \rangle \{ \exists v. \text{pseudo} = \text{hash}(v) \wedge \text{Authenticate}(w, v) \}. \end{aligned}$$

After the tuple is split the environment becomes:

$$\Gamma_2 = \Gamma_1, y_p : \text{Un}, \langle y_p \rangle = x, \exists v. \text{pseudo} = \text{hash}(v) \wedge \text{Authenticate}(y_p, v)$$

In order to type-check the  $\text{assert Associate}(y_p, k)$  we need to ensure that  $\text{forms}(\Gamma_2) \models \text{Associate}(y_p, k)$ , which holds in the authorization logic.

### 4.3 Type-checking Zero-knowledge

The main novelty of our type system is the treatment of zero-knowledge proofs.

**The Zero-knowledge Type.** We give zero-knowledge proofs of the form  $\text{zk}_{n,m,S}(\tilde{N}; \tilde{M})$  type  $\text{ZKProof}_{n,m,S}(\langle y_1 : T_1, \dots, y_m : T_m \rangle \{ \exists x_1, \dots, x_n. C \})$ . This type contains a tuple type listing the types of the arguments in the public component. The logical formula associated to this type is of the form  $\exists x_1, \dots, x_n. C$ , where the arguments in the private component are existentially quantified. The type system guarantees that  $C\{\tilde{N}/\tilde{x}\}\{\tilde{M}/\tilde{y}\}$  is entailed by the formulas in the typing environment.

**Type Annotations.** Typing all the other cryptographic primitives we consider relies on the type of some key, which the user has to annotate explicitly. Zero-knowledge proofs, however, do not depend in general on any key. This poses a problem since type-checking the verification of zero-knowledge proofs should propagate logical formulas in the typing environment of the verifier, and it is not clear what formulas to consider. For instance, when type-checking a process  $\text{let } \langle \tilde{y} \rangle = \text{ver}_{n,m,0,S}(z) \text{ then } P$ , we can safely assume that the formula  $\exists \tilde{x}. S\{\tilde{x}/\tilde{\alpha}\}\{\tilde{y}/\tilde{\beta}\}$  holds for the continuation process  $P$ . This is in fact guaranteed by the operational semantics of the `ver` destructor<sup>2</sup> (see Section 3.3). Such a formula, however, does not suffice to type-check most examples we have tried, since it does not mention any logical predicate.

In order to solve this problem we allow the user to provide type annotations for each statement used in the process. For each  $(n, m)$ -statement  $S$  this annotation is modeled as a distinguished free variable  $s_{n,m,S}$  in the initial typing environment, which is bound to a type of the form  $\text{Stm}(\langle y_1 : T_1, \dots, y_m : T_m \rangle \{ \exists x_1, \dots, x_n. C \})$ . Additionally, the environment contains an implicit binding  $s_{n,m,S}^{un} : \text{Stm}(\text{Un})$  used to type-check proofs of  $S$  generated by the adversary.

**Typing Zero-knowledge Proofs.** With this setup in place, we can formalize the typing rule for zero-knowledge proofs:

$$\frac{\begin{array}{l} \Gamma(s_{n,m,S}) = \text{Stm}(\langle y_1 : T_1, \dots, y_m : T_m \rangle \{ \exists x_1, \dots, x_n. C \}) \\ \forall i \in [1, n]. \Gamma \vdash N_i : U_i \\ \Gamma \vdash \langle M_1, \dots, M_m \rangle : \langle y_1 : T_1, \dots, y_m : T_m \rangle \{ C\{\tilde{N}/\tilde{x}\} \} \end{array}}{\Gamma \vdash \text{zk}_{n,m,S}(N_1, \dots, N_n; M_1, \dots, M_m) : \text{ZKProof}_{n,m,S}(\langle y_1 : T_1, \dots, y_m : T_m \rangle \{ \exists x_1, \dots, x_n. C \})}$$

Note that we require that  $\Gamma \vdash \langle M_1, \dots, M_m \rangle : \langle y_1 : T_1, \dots, y_m : T_m \rangle \{ C\{\tilde{N}/\tilde{x}\} \}$ , which not only makes sure that the public arguments have the required types, but it also effectively checks that  $C\{\tilde{N}/\tilde{x}\}\{\tilde{M}/\tilde{y}\}$  logically follows from the formulas in the prover's typing environment. In this way we ensure that honest participants can only generate proofs for statements that are associated to formulas already entailed by the environment.

**Typing Zero-Knowledge Verification.** Suppose that we are given the process  $\text{let } x = \text{ver}_{n,m,l,S}(N, M_1, \dots, M_l) \text{ then } P \text{ else } Q$ , a typing environment  $\Gamma$  so that  $\Gamma(s_{n,m,S}) = \text{Stm}(\langle y_1 : T_1, \dots, y_m : T_m \rangle \{ \exists x_1, \dots, x_n. C \}) = T$  (where  $T$  is annotated by the user), and a type  $T'$  such that  $\Gamma \vdash N : T'$ . Zero-knowledge proofs are typically received from channels controlled by the attacker and in this

<sup>2</sup>Notice that we did not match any of the public arguments of the proof.

case  $T' = \text{ZKProof}_{n,m,S}(\langle y_1 : \text{Un}, \dots, y_m : \text{Un} \rangle \{\text{true}\})$ , which is equivalent to  $\text{Un}$  by subtyping. In order to type-check this process, we first check that the terms  $M_1, \dots, M_l$  have type  $T_1, \dots, T_l$ , since these arguments are matched against the first  $l$  terms in the public component of the verified proof. Moreover, we need to check that  $N$  can be safely given the stronger type  $T$ .

The main idea for obtaining stronger guarantees than those given by the semantics of the `ver` destructor is to use the types of the matched public arguments to derive the type of the other arguments of the proof, even the private ones. For instance, if a matched public argument is a hash of type  $\text{Hash}(U)$  and the statement proves the knowledge of the value inside, then we can derive that this value has type  $U$ . Similarly, if a matched public argument is a key of type  $\text{VerKey}(U)$  and the statement proves the verification of a signature using this key, then the message that is signed has type  $U$ . This kind of reasoning can be exploited to infer both type information and logical formulas. Furthermore, if we can statically verify that at least one of the arguments of the proof is neither public nor tainted, then we know that the zero-knowledge proof has been generated by a honest participant (since the adversary knows only terms that are public or tainted). This immediately implies that the proof has the stronger type  $T$ , since zero-knowledge proofs constructed by honest participants are type-checked against the type specified by the user.

This intuitive reasoning is formalized by the predicate  $\langle\langle S \rangle\rangle_{\Gamma, n, m, l, T, T'}$  (for details see [9]). If this predicate holds, then the zero-knowledge proof  $N$  of type  $T'$  is guaranteed to have the stronger type  $T$  and we can safely give the last  $m-l$  arguments of the proof type  $\langle y_{l+1} : T_{l+1}, \dots, y_m : T_m \rangle \{C\{M_i/y_i\}_{i \in [1, l]}\}$ . With this setup in place, the typing rule for the `ver` destructor is defined as follows:

$$\frac{\begin{array}{l} \Gamma(s_{n,m,S}) = \text{Stm}(T), \text{ where } T = \langle y_1 : T_1, \dots, y_m : T_m \rangle \{\exists x_1, \dots, x_n. C\} \\ \Gamma \vdash N : \text{ZKProof}_{n,m,S}(T'), \text{ where } T' = \langle y_1 : U_1, \dots, y_m : U_m \rangle \{\text{true}\} \\ \forall i \in [1, l]. \Gamma \vdash M_i : T_i \quad \langle\langle S \rangle\rangle_{\Gamma, n, m, l, T, T'} \text{ holds} \end{array}}{\Gamma, x : \langle y_{l+1} : T_{l+1}, \dots, y_m : T_m \rangle \{\exists x_1, \dots, x_n. C\{M_i/y_i\}_{i \in [1, l]}\} \vdash P \quad \Gamma \vdash Q} \\ \Gamma \vdash \text{let } x = \text{ver}_{n,m,l,S}(N, M_1, \dots, M_l) \text{ then } P \text{ else } Q$$

#### 4.4 Type-checking the Protocol for Anonymous Trust

We explain how to type-check the prover and the verifier processes from Section 3.7, in the initial typing environment:  $\Gamma_0 = s_{1,2,\beta_1=\text{hash}(\alpha_1)} : \text{Stm}(T_{zk}), m_s : \text{Private}, c : \text{Un}, k : \text{Un}$ , where  $T_{zk} = \langle y_1 : \text{Hash}(\text{Private}), y_2 : \text{Un} \rangle \{\exists x. y_1 = \text{hash}(x) \wedge \text{Authenticate}(y_2, x)\}$ .

For the prover, we need to check that the term  $\text{zk}_{1,2,\beta_1=\text{hash}(\alpha_1)}(m_s; \text{hash}(m_s), m_p)$  has type  $\text{ZKProof}_{1,2,\beta_1=\text{hash}(\alpha_1)}(T_{zk})$ , in the extended environment  $\Gamma_1 = \Gamma_0, m_p : \text{Un}, \text{Authenticate}(m_p, m_s)$ . For this we need to show that  $\Gamma_1 \vdash \text{hash}(m_s) : \text{Hash}(\text{Private})$ ,  $\Gamma_1 \vdash m_p : \text{Un}$  and  $\text{forms}(\Gamma_1) \models \exists x. \text{hash}(m_s) = \text{hash}(x) \wedge \text{Authenticate}(m_p, x)$ , which holds by instantiating  $x$  to  $m_s$ .

The verifier receives the zero-knowledge proof from the untrusted<sup>3</sup> channel  $c$ . The zero-knowledge proof is thus bound to a variable  $x$  of type  $\text{Un}$  and, by subtyping, also of type  $\text{ZKProof}_{1,2,\beta_1=\text{hash}(\alpha_1)}(\langle y_1 : \text{Un}, y_2 : \text{Un} \rangle \{\text{true}\})$ . Note that this type is not strong enough to type-check the continuation process. Type-checking the destructor application  $\text{ver}_{1,2,1,\beta_1=\text{hash}(\alpha_1)}(x, \text{pseudo})$  can however

<sup>3</sup>We consider all free names to have type  $\text{Un}$ .

rely on the fact that  $\Gamma \vdash \textit{pseudo} : \text{Hash}(\text{Private})$ , and therefore *pseudo* is the hash of some message  $m_s$  of type `Private`, which is neither public nor tainted. The statement guarantees that the prover knows  $m_s$  and this is enough to ensure that the zero-knowledge proof is generated by a honest participant. Therefore  $\langle\langle S \rangle\rangle_{\Gamma, n, m, l, T, T'}$  holds, and the type system gives the result of the destructor application type  $\langle y : \text{Un} \rangle \{ \exists v. \textit{pseudo} = \text{hash}(v) \wedge \text{Authenticate}(y, v) \}$ . This allows type-checking the continuation process as discussed at the end of Section 4.2.

## 4.5 Security Guarantees

Our type system statically guarantees that in well-typed processes all asserted formulas are valid at runtime (safety), even in the presence of an arbitrary adversary (robust safety). We use  $\Gamma \vdash_{\text{Un}} P$  to denote  $\Gamma, u_1 : \text{Un}, \dots, u_n : \text{Un} \vdash P$ , where  $\{u_1, \dots, u_n\} = \text{free}(P)$ .

**Theorem 4.1 (Robust Safety)** *Let  $\Gamma = s_{n_1, m_1, S_1} : \text{Stm}(T_1), \dots, s_{n_k, m_k, S_k} : \text{Stm}(T_k)$ . For every closed process  $P$ , if  $\Gamma \vdash_{\text{Un}} P$  then  $P$  is robustly safe.*

Due to space constraints, the proof of this theorem and of all the necessary lemmas are given in the extended version of the paper [9].

## 5 Implementation

We implemented an automatic type-checker for the type system presented in this paper. The type-checking phase is guaranteed to terminate and generates proof obligations that are discharged independently, leading to a modular and robust analysis. We use first-order logic with equality as the authorization logic and we employ the automated theorem prover SPASS [36] to discharge the proof obligations. SPASS internally uses superposition for equational reasoning [6]. Our tool is written in Objective Caml, comprises approximately 3000 lines of code, and is available at [10].

The implementation uses an algorithmic version of our type system. Devising a variant of the type system that is suitable for an implementation required us to eliminate the subsumption rule and to deal with the facts that the constructors and destructors are in fact polymorphic and that the instantiation of the type variables needs to be made automatically. The latter problem is not trivial since our type system features subtyping. On the other hand, asking the user to annotate every constructor and destructor application would have been unacceptable from a usability perspective.

We tested our tool on the DAA protocol (details are given in the extended version of the paper [9]) and several simpler examples including the anonymous trust management protocol given in Section 2. The analysis of DAA terminated in less than three seconds and discharged 39 proof obligations, while for the simpler examples the time needed was less than half a second. These promising results show that our static analysis technique has the potential to scale up to industrial-size protocols.

## 6 Conclusions

This paper shows how certain security properties of zero-knowledge proofs can be characterized as authorization policies and statically enforced by a novel type system. We developed a type-checker for this type system and use an automated theorem-prover to discharge proof obligations. We applied our technique to verify the authenticity properties of the Direct Anonymous Attestation (DAA) protocol.

The analysis technique is very efficient. Furthermore, the combination of types and authorization logics constitutes a truly expressive framework to model and analyze a variety of trace-based security properties. Zero-knowledge proofs perfectly fit into this framework, offering the possibility to implement fine-grained authorization policies that rely on the existential quantification in the logic. This is particularly well-suited for protocols for privacy and anonymity.

The type system for authorization proposed by Fournet et al. [26], on which our approach is grounded, has been recently applied to the analysis of protocol implementations written in F# [13]. We are confident that our technique could be incorporated into such a framework in order to verify implementations of protocols based on zero-knowledge proofs.

In the extended version of this paper [9] we explore using zero-knowledge proofs in the design of systems that guarantee security despite partial compromise. We believe that zero-knowledge proofs are the natural candidate for strengthening protocol specifications against compromised participants, since they can be used to verify the correct behavior of remote parties and to safely derive authorization decisions. A formal elaboration of these ideas is an interesting direction for future research.

Finally, while there is no restriction on the shape of the statements that can be proved in our abstract setting, finding a sound and efficient cryptographic implementation for these symbolic zero-knowledge proofs is far from trivial. In [12] the authors identify the properties a concrete zero-knowledge proof system should satisfy in order to soundly implement a simpler abstraction of zero-knowledge, and mention two existing cryptographic constructions that satisfy these properties. The efficiency and expressivity of these constructions have, however, not been thoroughly studied and constitute interesting topics for further investigation.

**Acknowledgments.** The prototype implementation of our type system was developed by Stefan Lorenz and Kim R. Pecina. We thank the anonymous reviewers for their valuable comments and Catherine Howell for her help on technical writing.

## References

- [1] M. Abadi. Secrecy by typing in security protocols. *Journal of the ACM*, 46(5):749–786, 1999.
- [2] M. Abadi. Logic in access control. In *Proc. 18th IEEE Symposium on Logic in Computer Science (LICS)*, pages 228–233. IEEE Computer Society Press, 2003.
- [3] M. Abadi and B. Blanchet. Analyzing security protocols with secrecy types and logic programs. *Journal of the ACM*, 52(1):102–146, 2005.
- [4] M. Abadi, B. Blanchet, and C. Fournet. Automated verification of selected equivalences for security protocols. In *Proc. 20th Annual IEEE Symposium on*

- Logic in Computer Science (LICS)*, pages 331–340. IEEE Computer Society Press, 2005.
- [5] A. Acquisti. Receipt-free homomorphic elections and write-in ballots. Cryptology ePrint Archive, Report 2004/105, 2004. <http://eprint.iacr.org/>.
  - [6] L. Bachmair and H. Ganzinger. Rewrite-based equational theorem proving with selection and simplification. *Journal of Logic and Computation*, 4(3):217–247, 1994.
  - [7] M. Backes, A. Cortesi, R. Focardi, and M. Maffei. A calculus of challenges and responses. In *Proc. 5th ACM Workshop on Formal Methods in Security Engineering (FMSE)*, pages 101–116. ACM Press, 2007.
  - [8] M. Backes, C. Hrițcu, and M. Maffei. Automated verification of remote electronic voting protocols in the applied pi-calculus. To appear in 21th IEEE Symposium on Computer Security Foundations (CSF), 2008.
  - [9] M. Backes, C. Hrițcu, and M. Maffei. Type-checking zero-knowledge. Long version available at <http://www.infsec.cs.uni-sb.de/~hritcu/publications/zk-types-full.pdf>, 2008.
  - [10] M. Backes, C. Hrițcu, and M. Maffei. Type-checking zero-knowledge. Implementation available at <http://www.infsec.cs.uni-sb.de/projects/zk-typechecker>, 2008.
  - [11] M. Backes, M. Maffei, and D. Unruh. Zero-knowledge in the applied pi-calculus and automated verification of the direct anonymous attestation protocol. In *Proc. 29th IEEE Symposium on Security and Privacy*, pages 202–215. IEEE Computer Society Press, 2008.
  - [12] M. Backes and D. Unruh. Computational soundness of symbolic zero-knowledge proofs against active attackers. To appear in 21th IEEE Symposium on Computer Security Foundations (CSF), 2008.
  - [13] J. Bengtson, K. Bhargavan, C. Fournet, A. D. Gordon, and S. Maffei. Refinement types for secure implementations. To appear in 21th IEEE Symposium on Computer Security Foundations (CSF), 2008.
  - [14] B. Blanchet. An efficient cryptographic protocol verifier based on Prolog rules. In *Proc. 14th IEEE Computer Security Foundations Workshop (CSFW)*, pages 82–96. IEEE Computer Society Press, 2001.
  - [15] D. Bleichenbacher. Chosen ciphertext attacks against protocols based on the RSA encryption standard PKCS. In *Advances in Cryptology: CRYPTO '98*, volume 1462 of *Lecture Notes in Computer Science*, pages 1–12. Springer-Verlag, 1998.
  - [16] E. Brickell, J. Camenisch, and L. Chen. Direct anonymous attestation. In *Proc. 11th ACM Conference on Computer and Communications Security*, pages 132–145. ACM Press, 2004.
  - [17] M. Bugliesi, R. Focardi, and M. Maffei. Dynamic types for authentication. *Journal of Computer Security*, 15(6):563–617, 2007.
  - [18] F. Butler, I. Cervesato, A. D. Jaggard, A. Scedrov, and C. Walstad. Formal analysis of Kerberos 5. *Theoretical Computer Science*, 367(1):57–87, 2006.
  - [19] M. R. Clarkson, S. Chong, and A. C. Myers. Civitas: A secure voting system. In *Proc. 29th IEEE Symposium on Security and Privacy*, pages 354–368. IEEE Computer Society Press, 2008.
  - [20] R. Corin, P. Denielou, C. Fournet, K. Bhargavan, and J. Leifer. Secure implementations for typed session abstractions. In *Proc. 20th IEEE Symposium on Computer Security Foundations (CSF)*, pages 170–186. IEEE Computer Society Press, 2007.
  - [21] S. Delaune, S. Kremer, and M. Ryan. Coercion-resistance and receipt-freeness in electronic voting. In *Proc. 19th IEEE Computer Security Foundations Workshop (CSFW)*, pages 28–42. IEEE Computer Society Press, 2006.
  - [22] S. Delaune, M. Ryan, and B. Smyth. Automatic verification of privacy properties in the applied pi calculus. To appear in 2nd Joint iTrust and PST Conferences on Privacy, Trust Management and Security (IFIPTM'08), 2008.

- [23] B. Dragovic, E. Kotsovinos, S. Hand, and P. R. Pietzuch. Xenotrust: Event-based distributed trust management. In *Proc. 14th International Workshop on Database and Expert Systems Applications (DEXA'03)*, pages 410–414. IEEE Computer Society Press, 2003.
- [24] D. Fisher. Millions of .Net Passport accounts put at risk. *eWeek*, May 2003. (Flaw detected by Muhammad Faisal Rauf Danka).
- [25] C. Fournet, A. D. Gordon, and S. Maffeis. A type discipline for authorization policies. In *Proc. 14th European Symposium on Programming (ESOP)*, Lecture Notes in Computer Science, pages 141–156. Springer-Verlag, 2005.
- [26] C. Fournet, A. D. Gordon, and S. Maffeis. A type discipline for authorization in distributed systems. In *Proc. 20th IEEE Symposium on Computer Security Foundations (CSF)*, pages 31–45. IEEE Computer Society Press, 2007.
- [27] O. Goldreich. *Foundations of Cryptography: Basic Tools*. Cambridge University Press, 2001.
- [28] A. D. Gordon and A. Jeffrey. Types and effects for asymmetric cryptographic protocols. *Journal of Computer Security*, 12(3):435–484, 2004.
- [29] A. D. Gordon and A. Jeffrey. Secrecy despite compromise: Types, cryptography, and the pi-calculus. In *Proc. 16th International Conference on Concurrency Theory (CONCUR)*, volume 3653, pages 186–201. Springer-Verlag, 2005.
- [30] C. Haack and A. Jeffrey. Timed spi-calculus with types for secrecy and authenticity. In *Proc. 16th International Conference on Concurrency Theory (CONCUR)*, volume 3653, pages 202–216. Springer-Verlag, 2005.
- [31] C. Haack and A. Jeffrey. Pattern-matching spi-calculus. *Information and Computation*, 204(8):1195–1263, 2006.
- [32] A. Juels, D. Catalano, and M. Jakobsson. Coercion-resistant electronic elections. In *Proc. 4th ACM Workshop on Privacy in the Electronic Society (WPES)*, pages 61–70. ACM Press, 2005.
- [33] S. D. Kamvar, M. Schlosser, and H. Garcia-Molina. The eigentrust algorithm for reputation management in p2p networks. In *In Proc. 12th International World Wide Web Conference (WWW)*, pages 640–651. ACM Press, 2003.
- [34] L. Lu, J. Han, L. Hu, J. Huai, Y. Liu, and L. M. Ni. Pseudo trust: Zero-knowledge based authentication in anonymous peer-to-peer protocols. In *Proc. 2007 IEEE International Parallel and Distributed Processing Symposium*, page 94. IEEE Computer Society Press, 2007.
- [35] D. Wagner and B. Schneier. Analysis of the SSL 3.0 protocol. In *Proc. 2nd USENIX Workshop on Electronic Commerce*, pages 29–40, 1996.
- [36] C. Weidenbach, R. Schmidt, T. Hillenbrand, R. Rusev, and D. Topic. System description: SPASS version 3.0. In *Automated Deduction – CADE-21 : 21st International Conference on Automated Deduction*, volume 4603 of *Lecture Notes in Artificial Intelligence*, pages 514–520. Springer, 2007.