

Type-checking Implementations of Protocols Based on Zero-knowledge Proofs – Work in Progress –

Michael Backes^{1,2}, Cătălin Hrițcu¹, Matteo Maffei¹, Thorsten Tarrach¹

¹Saarland University, Saarbrücken, Germany ²MPI-SWS

Abstract. We present the first static analysis technique for verifying implementations of cryptographic protocols based on zero-knowledge proofs. Protocols are implemented in $\text{RCF}_{\wedge, \vee}^{\forall}$, a core calculus of ML with support for concurrency. Cryptographic primitives are considered as fully reliable building blocks and represented symbolically using a sealing mechanism. Zero-knowledge proofs, in particular, are specified in a high-level language and automatically compiled down to a symbolic implementation using seals. Security properties are formalized as authorization policies and statically enforced by a type system featuring refinement, union, intersection, and polymorphic types. The expressiveness of our type system allows us to analyze not only protocols based on zero-knowledge proofs but also other important protocol classes that were out of the scope of existing static analysis techniques.

1 Introduction

Proofs of security protocols are known to be error-prone and awkward to make for humans. In fact, vulnerabilities have accompanied the design of such protocols ever since early authentication protocols like Needham-Schroeder, over carefully designed de-facto standards like SSL and PKCS, up to current widely deployed products like Microsoft Passport and Kerberos. Hence work towards the automation of such proofs has started soon after the first protocols were developed, focusing on high-level protocol descriptions that are suited to the formalization of security properties and to the automation of security proofs.

Despite this considerable progress on protocol analysis, the verification of security protocol implementations is still a widely unexplored field. The aforementioned high-level protocol descriptions abstract away from most troublesome implementation details, and there is no guarantee that a protocol that has been proven secure within an abstract model stays secure when implemented.

Devising automated techniques for proving the security of protocol implementations is a highly non-trivial task. First, modern applications such as trusted computing [12] and electronic voting [14] rely on *complex cryptographic primitives*, such as zero-knowledge proofs. Automated verification of such applications requires to symbolically abstract these primitives. Process calculi provide convenient mechanisms to define such abstractions, for instance flexible equational theories [3,2], which rendered an automated analysis of such applications feasible in abstract models [6,4]. This is not the case for standard programming languages,

where one needs to encode these abstractions using the primitives provided by the language. These primitives, however, were not designed for this purpose, which makes providing encodings that are suitable for automatic analysis a difficult task.

Second, high-level protocol specifications are typically compact, since many implementation details are abstracted away, while protocol implementations are much larger. Therefore *efficiency* and *scalability* of the proposed techniques are (even more) crucial when dealing with industrial-size applications.

Third, conducting *security proofs* within abstract protocol models is known to be a difficult task. Providing security proofs for protocol implementations is typically even more difficult since they require extensive reasoning about the behavior of programs in the presence of features such as data structures, recursion, and state.

1.1 Contributions

We present the first static analysis technique for verifying implementations of security protocols based on non-interactive zero-knowledge proofs. Protocols are implemented in RCF [7], a core calculus of ML, and cryptographic primitives are considered fully reliable building blocks and represented symbolically using a sealing mechanism [16,19]. The zero-knowledge proofs are specified in a high-level language and automatically compiled down to a symbolic implementation. The verification of the resulting code relies on a type system with union, intersection, and refinement types. This expressive type system extends the scope of type-based analyses of protocol implementations to important protocol classes not covered so far. In particular, we use union and intersection types to precisely characterize asymmetric cryptography, which allows us to verify protocols based on nested cryptography, signatures of private data, and public-key encryption of authenticated data. The analysis is automated and modular, and was efficient on the examples that we tried.

1.2 Related Work

Our type system builds on the work by Bengtson et al. on the type-based analysis of protocol implementations in F# [7]. Their type system supports most of the features of F# by translation to RCF. The analysis is modular and promises to scale up to large implementations [8]. Their type system, however, is not expressive enough to deal with zero-knowledge proofs and poses some serious restrictions even on the usage of standard cryptographic primitives. For instance, if a key is used to sign a secret message, then the corresponding verification key cannot be made public. These limitations prevent the analysis of many interesting cryptographic applications, such the Direct Anonymous Attestation protocol [12], which relies on zero-knowledge proofs as well as on digital signatures on secret TPM identifiers. We extend this type system with union and intersection types, which significantly increases the expressiveness of the analysis and enables the verification of protocols based on zero-knowledge proofs.

Goubault-Larrecq and Parrennes developed a static analysis technique [15] based on pointer analysis and clause resolution for cryptographic protocols implemented in C. The analysis provides security proofs and is also useful to find bugs in the implementation, but it is limited to secrecy properties, it deals only with standard cryptographic primitives, and it does not offer scalability since the number of generated clauses is very high even on small protocol examples.

Chaki and Datta have recently proposed a technique [13] based on software model checking for the automated verification of secrecy and authentication properties of protocols implemented in C. The analysis provides security proofs for a bounded number of sessions and is effective in discovering attacks. It was used to check secrecy and authentication properties of the SSL handshake protocol for configurations of up to three servers and three clients. The analysis only deals with standard cryptographic primitives, and offers only limited scalability.

Bhargavan et al. proposed a technique [10,9] for the verification of F# protocol implementations by automatically extracting ProVerif models [11]. The analysis provides security proofs and, despite its non-compositional nature, scales remarkably well and was successfully used to verify implementations of real-world cryptographic protocols such as TLS [9]. The considered fragment of F# is, however, very restrictive: it does not include higher-order functions, and it allows only for a very limited usage of recursion and state. These restrictions impose serious limitations on the programming style, and as a consequence the technique can only be used to analyze protocol implementations developed from scratch.

1.3 Outline

Section 2 introduces the calculus and Section 3 the type system we consider. Section 4 illustrates our symbolic implementation of asymmetric cryptography. Section 5 discusses our encoding of zero-knowledge proofs in $\text{RCF}_{\wedge\vee}^{\forall}$ and applies it to a simplified version of the DAA protocol. Section 6 describes our implementation. Section 7 concludes and discusses future work. Due to space constraints, we defer many technical details to an extended version of the paper [5].

2 $\text{RCF}_{\wedge\vee}^{\forall}$

The Refined Concurrent FPC (RCF) [7] is a simple programming language extending the Fixed Point Calculus with refinement types and concurrency. Although very simple, this core calculus is expressive enough to encode a considerable fragment of F#. In this paper, we further increase the expressivity of the calculus by adding intersection types [17], union types, and parametric polymorphism. We call the extended calculus $\text{RCF}_{\wedge\vee}^{\forall}$ and describe it in this and in the following section.

2.1 Syntax and Informal Semantics

The set of *values* is composed of variables, the unit value, functions, pairs, and introduction forms for recursive and polymorphic types (cf. Table 1). Names are

Table 1 Syntax of $\text{RCF}_{\wedge\vee}^{\forall}$ values and expressions

a, b, c	name	$A, B ::=$	expression
x, y, z	variable	M	value
$M, N ::=$	value	$M N$	function application
x	variable	$M\langle T \rangle$	type instantiation
$()$	unit	if $M = N$ as x then A else B	equality check with type-cast
$\lambda x : T. A$	function	let $x = A$ in B	let
(M, N)	pair	let $(x, y) = M$ in A	pair split
$\Lambda\alpha. A$	polymorphic val.	unfold $_{\mu\alpha.T} M$	use recursive value
fold $_{\mu\alpha.T} M$	recursive val.	for $\tilde{\alpha}$ in $\tilde{T}; \tilde{U}$ do A	introduction of intersections
for $\tilde{\alpha}$ in $\tilde{T}; \tilde{U}$ do M ,		case $x = M : T \vee U$ in A	elimination of union types
where $\{\tilde{\alpha}\} \cap \text{free}(M) \neq \emptyset$		$(\nu a \downarrow T)A$	restriction
intersection val.		$A \uparrow B$	fork
		$a!M$	send M on channel a
		$a?$	receive message from a
		assume C	assumption of formula C
		assert C	assertion of formula C

Notation: Given a phrase of syntax ϕ , let $\phi\{M/x\}$ denote the substitution of each free occurrence of the variable x in ϕ with the value M . We use $\tilde{\phi}$ to denote the sequence ϕ_1, \dots, ϕ_n for some n .

generated at run-time and are only used as channel identifiers, while variables are place-holders for values.

An *expression* represents a concurrent computation that may reduce to a value (or may diverge, get “stuck” or deadlock). The *reduction relation* is defined on *computation states* consisting of several expressions being evaluated in parallel; the messages sent on channels but not yet received; and a global log containing assumed formulas¹. Values are irreducible. The function application $(\lambda x:T. A) M$ reduces to $A\{M/x\}$. A type instantiation $(\Lambda\alpha. A)\langle T \rangle$ reduces to $A\{T/\alpha\}$. The conditional if $M = N$ as x then A else B reduces to $A\{M/x\}$ if M is syntactically equal to N , and to B otherwise (more details are given in Section 3.4). The introduction form for intersection types for $\tilde{\alpha}$ in $\tilde{T}; \tilde{U}$ do A allows A to reduce to a value M , and if M does not contain $\tilde{\alpha}$ then returns M , otherwise it returns the value for $\tilde{\alpha}$ in $\tilde{T}; \tilde{U}$ do M . The elimination form for union types case $x = M : T \vee U$ in A reduces to $A\{M/x\}$ (more details about intersection and union types are given in Section 3). Intuitively, the restriction $(\nu a \downarrow T)A$ generates a globally fresh channel a that can only be used in A to convey values of type T . The expression $A \uparrow B$ evaluates A and B in parallel, and returns the result of B (the result of A is discarded). The expression $a!M$ outputs M on channel a and reduces to the unit value $()$. The evaluation of $a?$ blocks until some message M is available on channel a , removes M from the channel, and then returns M . Expression assume C , where C is a logical formula, adds C to the global log. The assertion assert C reduces to $()$. If C is entailed by the multiset S of formulas in the global log, we say the assertion *succeeds*; otherwise, we say the assertion *fails*.

¹ We consider first-order logic with equality as the authorization logic.

Table 2 A direct implementation of a sign-then-encrypt protocol in $\text{RCF}_{\wedge\vee}^{\forall}$

```

( $\nu c \downarrow T_{\text{chan}}$ )
let  $y_b = \text{mkId}()$  in
let  $sk = \text{mkSK} \langle T_{\text{sign}} \rangle ()$  in let  $vk = \text{mkVK} \langle T_{\text{sign}} \rangle sk$  in
let  $dk = \text{mkDK} \langle T_{\text{enc}} \rangle ()$  in let  $ek = \text{mkEK} \langle T_{\text{enc}} \rangle dk$  in
( let  $y = \text{getPriv} ()$  in
  assume  $Ok(y)$ ;
  let  $x = (\text{sign} \langle T_{\text{sign}} \rangle sk) (y_b, y)$  in
  let  $z = (\text{encrypt} \langle T_{\text{enc}} \rangle ek) (x, (y_b, y))$  in
   $c!z$  )†
( let  $z = c?$  in
  let  $xy = (\text{decrypt} \langle T_{\text{enc}} \rangle dk) z$  in
  let  $(x, y) = xy$  in
  let  $y' = (\text{check} \langle T_{\text{sign}} \rangle vk x y)$  in
  let  $(y_1, y_2) = y'$  in
  if  $y_1 = y_b$  then assert  $Ok(y_2)$  )

```

Definition 1 (Safety). A closed expression A is safe if and only if, in all evaluations of A , all assertions succeed.

When reasoning about implementations of cryptographic protocols, we are interested in the safety of programs executed in parallel with an arbitrary attacker. This property is called *robust safety*.

Definition 2 (Opponent and Robust Safety). A closed expression O is an opponent if and only if O contains no assertions and O only contains Un for the type annotations. A closed expression A is robustly safe if and only if the application $O A$ is safe for all opponents O .

2.2 Example: Sign-then-Encrypt

We illustrate the calculus, and later the type system, on the following very simple protocol, in which A first signs and then encrypts a private message for B :

$$A \longrightarrow \left\{ \left[\left[(B, m) \right]_{k_A^-}, (B, m) \right] \right\}_{k_B^+} \longrightarrow B$$

The $\text{RCF}_{\wedge\vee}^{\forall}$ implementation of this protocol is reported in Table 2. We first generate a fresh public channel c and B 's identifier y_b . We then create the signing and encryption key-pairs. The sender gets a private value y from the user and assumes the predicate $Ok(y)$. The sender then signs this message together with the receiver's identifier y_b , encrypts this signature together with the signed pair (y_b, y) , and outputs the resulting ciphertext on channel c .

The receiver reads the message from channel c , decrypts it, splits the pair, and checks the signature. If all these checks succeed and additionally the second component of the pair is his own identifier, the receiver asserts $Ok(y_2)$, where y_2 is bound to the private value sent by A .

Table 3 Syntax of Types

$T, U, V ::= \top$	top type	$x : T \rightarrow U$	dependent function type
Un	untrusted type	$x : T * U$	dependent pair type
Private	private type	α	type variable
$T \wedge U$	intersection type	$\mu\alpha. T$	iso-recursive type
$T \vee U$	union type	$\forall\alpha. T$	polymorphic type
$\{x : T \mid C\}$	refinement type	unit	unit type

Notations:

Let $T \rightarrow U \triangleq x : T \rightarrow U$ and $T * U \triangleq x : T * U$, where in both cases $x \notin \text{free}(U)$. Let $\{C\}$ denote $\{x : \text{unit} \mid C\}$ where $x \notin \text{free}(C)$.

3 Type System

This section presents a type system for enforcing authorization policies on $\text{RCF}_{\wedge\vee}^{\forall}$ code. This extends the type system proposed by Bengtson et al. [7] with union types, intersection types [17] and unrestricted parametric polymorphism. This extension enhances the expressivity of the analysis and, in particular, allows us to obtain a more precise characterization of asymmetric cryptography and to provide a faithful encoding of zero-knowledge proofs.

3.1 Types

The syntax of types is reported in Table 3. Our type system has a top type \top that is supertype of all the others (the subtyping relation is introduced in Section 3.3). A value is given the intersection type $T \wedge U$ if it has both type T and type U . A value is given a union type $T \vee U$ if it has type T or if it has type U , but we do not necessarily know what its precise type is. As in [7], we use refinement types to associate logical formulas to messages. The refinement type $\{x : T \mid C\}$ describes values M of type T for which the formula $C\{M/x\}$ is entailed by the current typing environment. Functions $\lambda x : T. A$ taking as input values of type T and returning values of type U are given the dependent type $x : T \rightarrow U$, where the result type U can depend on the input value x . Pairs are given dependent types of the form $x : T * U$, where the type U of the second component of the pair can depend on the value x of the first component. In the following we explain the typing judgements and present the most important typing rules.

3.2 Typing Environment and Entailment

A typing environment E is a list of bindings for type variables (α), names ($a \downarrow T$, where a stands for a channel conveying values of type T), and variables ($x : T$).

A crucial judgment in the type system is $E \vdash C$, which states that the formula C is derivable from E . Intuitively, our type system ensures that whenever $E \vdash C$ we have that C is logically entailed by the global formula log at execution time. This is used for instance when type-checking `assert C`: type-checking succeeds only if C is entailed in the current typing environment. If E binds a variable y to

a refinement type $\{x : T \mid C\}$, we know that the formula $C\{y/x\}$ is entailed in the system and therefore $E \vdash C\{y/x\}$.

3.3 Subtyping and Kinding

To type-check programs interacting with the attacker, we consider a universal type Un (untrusted), which is the type of data possibly known to the attacker. For instance, all data sent to and received from an untrusted channel have type Un , since such channels are considered under the complete control of the adversary.

However, a system in which only data of type Un can be communicated over the untrusted network would be too restrictive, e.g., a value of type $\{x : \text{Un} \mid \text{Ok}(x)\}$ could not be sent over the network. 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 most often used to compare types with type Un . In particular, we allow values having type T that is a subtype of Un , denoted $T <: \text{Un}$, to be sent over the untrusted network, and we say that the type T has *kind public* in this case. Similarly, we allow values of type Un that are received from the untrusted network to be used as values of type U , provided that $\text{Un} <: U$, and in this case we say that type U has *kind tainted*. In the following, we outline some important rules for kinding and subtyping.

Refinement types. The refinement type $\{x : T \mid C\}$ is a subtype of T . This allows us to discard logical formulas when they are not needed. For instance, a value of type $\{x : \text{Un} \mid \text{Ok}(x)\}$ can be sent on a channel of type Un . Conversely, the type T is a subtype of $\{x : T \mid C\}$ only if $\forall x.C$ is entailed in the current typing environment. By the transitivity of subtyping, a refinement type $\{x : T \mid C\}$ is public if T is public, since then $\{x : T \mid C\} <: T <: \text{Un}$. Conversely, type $\{x : T \mid C\}$ is tainted if T is tainted and additionally $\forall x.C$ holds.

Union and Intersection Types. The rules for subtyping and kinding union and intersection types are reported in Table 4. The type $T_1 \wedge T_2$ is a subtype of U if T_1 or T_2 is a subtype of U (cf. SUB-AND-LB), while T is a subtype of $U_1 \wedge U_2$ if it is subtype of both U_1 and U_2 (cf. SUB-AND-GREATEST). Intuitively, the set of values of type $T_1 \wedge T_2$ is the intersection between the set of values of type T_1 and the set of values of type T_2 . The type $T_1 \wedge T_2$ is public if T_1 or T_2 is public, and it is tainted if both T_1 and T_2 are tainted. The rules for union types are dual.

3.4 Typing Values and Expressions

The main judgment of the type system we consider is $E \vdash A : T$, which states that the expression A returns a value of type T . The most important typing rules are reported in Table 5. Most of them are standard, so we focus the explanation only on the rules that are new with respect to [7].

Conditionals. The rule EXP IF exploits intersection types for strengthening the type of the values tested for equality in the conditional if $M = N$ as x then A else B .

Table 4 Kinding and Subtyping Unions and Intersections

Subtyping

$\frac{\text{SUB AND LB}}{\Gamma \vdash T_i <: U}$ $\frac{\Gamma \vdash T_i <: U}{\Gamma \vdash T_1 \wedge T_2 <: U}$	$\frac{\text{SUB AND GREATEST}}{\Gamma \vdash T <: U_1 \quad \Gamma \vdash T <: U_2}$ $\frac{\Gamma \vdash T <: U_1 \quad \Gamma \vdash T <: U_2}{\Gamma \vdash T <: U_1 \wedge U_2}$
$\frac{\text{SUB OR SMALLEST}}{\Gamma \vdash T_1 <: U \quad \Gamma \vdash T_2 <: U}$ $\frac{\Gamma \vdash T_1 <: U \quad \Gamma \vdash T_2 <: U}{\Gamma \vdash T_1 \vee T_2 <: U}$	$\frac{\text{SUB OR UB}}{\Gamma \vdash T <: U_i}$ $\frac{\Gamma \vdash T <: U_i}{\Gamma \vdash T <: U_1 \vee U_2}$

Kinding

$\frac{\text{KIND AND PUB}}{\Gamma \vdash T_i :: \text{pub}}$ $\frac{\Gamma \vdash T_i :: \text{pub}}{\Gamma \vdash T_1 \wedge T_2 :: \text{pub}}$	$\frac{\text{KIND AND TNT}}{\Gamma \vdash T_1 :: \text{tnt} \quad \Gamma \vdash T_2 :: \text{tnt}}$ $\frac{\Gamma \vdash T_1 :: \text{tnt} \quad \Gamma \vdash T_2 :: \text{tnt}}{\Gamma \vdash T_1 \wedge T_2 :: \text{tnt}}$
$\frac{\text{KIND OR PUB}}{\Gamma \vdash T_1 :: \text{pub} \quad \Gamma \vdash T_2 :: \text{pub}}$ $\frac{\Gamma \vdash T_1 :: \text{pub} \quad \Gamma \vdash T_2 :: \text{pub}}{\Gamma \vdash T_1 \vee T_2 :: \text{pub}}$	$\frac{\text{KIND OR TNT}}{\Gamma \vdash T_i :: \text{tnt}}$ $\frac{\Gamma \vdash T_i :: \text{tnt}}{\Gamma \vdash T_1 \vee T_2 :: \text{tnt}}$

If M is of type T_1 and N is of type T_2 , then we type-check A under the assumption that $x = M \wedge M = N$, and x is of type $T_1 \wedge T_2$. This corresponds to a type-cast that is always safe, since the conditional succeeds only if M is syntactically equal to N , in which case the common value has indeed both the type of M and the type of N . Additionally, if the conditional succeeds the types T_1 and T_2 cannot be disjoint. However, certain types such as `Un` and `Private` have common values only if the environment is inconsistent, i.e., $E \vdash \text{false}$. Therefore, when comparing values of disjoint types it is safe to add `false` to the environment when type-checking A , which makes checking A always succeed. Intuitively, if T_1 and T_2 are disjoint the conditional cannot succeed, so the expression A will not be executed. A similar idea has been applied in [1] for verifying secrecy properties of nonce handshakes.

Union Types are introduced by subtyping (T_1 is a subtype $T_1 \vee T_2$ for any T_2), and eliminated using the `EXP CASE` rule. Suppose that M is of type $T_1 \vee T_2$ and M is used in A . Intuitively, since we do not know whether M is of type T_1 or T_2 , we have to type-check A under each of the assumptions. Therefore, the expression case $x = M : T_1 \vee T_2$ in A , where x takes the place of M in A , is of type U only if A is of type U both when x is of type T_1 and when x is of type T_2 . This is useful when type-checking code interacting with the attacker. For instance, suppose that a party receives a value M encrypted with a public-key that is used by honest parties to encrypt messages of type T . After decryption, M is given type $T \vee \text{Un}$ since it might come from a honest party as well as from the attacker. We have thus to type-check the receiver's code twice, once under the assumption that x is of type T , and once assuming that x is of type `Un`.

Table 5 Selected Rules for Values and Expressions $E \vdash A : T$

$\frac{\text{VAL REFINE} \quad E \vdash M : T \quad E \vdash C\{M/x\}}{E \vdash M : \{x : T \mid C\}}$	$\frac{\text{EXP SUBSUM} \quad E \vdash A : T \quad E \vdash T <: T'}{E \vdash A : T'}$	
$\frac{\text{EXP IF} \quad \begin{array}{l} E \vdash M : T_1 \quad E \vdash N : T_2 \\ E, x : T_1 \wedge T_2, _ : \{x = M \wedge M = N \wedge \text{non-disj}(T_1, T_2)\} \vdash A : U \\ E, _ : \{M \neq N\} \vdash B : U \end{array}}{E \vdash \text{if } M = N \text{ as } x \text{ then } A \text{ else } B : U}$		
$\frac{\text{EXP ASSUME} \quad E \vdash \diamond \quad \text{free}(C) \subseteq \text{dom}(E)}{E \vdash \text{assume } C : \{_ : \text{unit} \mid C\}}$	$\frac{\text{EXP ASSERT} \quad E \vdash C}{E \vdash \text{assert } C : \text{unit}}$	$\frac{\text{EXP SEND} \quad E \vdash M : T \quad (a \uparrow T) \in E}{E \vdash a!M : \text{unit}}$
$\frac{\text{EXP RECV} \quad E \vdash \diamond \quad (a \uparrow T) \in E}{E \vdash a? : T}$	$\frac{\text{EXP FOR1} \quad E \vdash A\{\tilde{T}/\tilde{\alpha}\} : V}{E \vdash \text{for } \tilde{\alpha} \text{ in } \tilde{T}; \tilde{U} \text{ do } A : V}$	$\frac{\text{EXP FOR2} \quad E \vdash A\{\tilde{U}/\tilde{\alpha}\} : V}{E \vdash \text{for } \tilde{\alpha} \text{ in } \tilde{T}; \tilde{U} \text{ do } A : V}$
$\frac{\text{EXP AND} \quad E \vdash A : T \quad E \vdash A : U}{E \vdash A : T \wedge U}$	$\frac{\text{EXP CASE} \quad \begin{array}{l} E \vdash M : T_1 \vee T_2 \\ E, x : T_1 \vdash A : U \quad E, x : T_2 \vdash A : U \end{array}}{E \vdash \text{case } x = M : T_1 \vee T_2 \text{ in } A : U}$	

Intersection Types are introduced using the EXP AND and EXP FOR $_i$ rules (and eliminated by subtyping). Intuitively, the expression for $\tilde{\alpha}$ in $\tilde{T}; \tilde{U}$ do A is of type $T_1 \wedge T_2$ if $A\{\tilde{T}/\tilde{\alpha}\}$ is of type T_1 and $A\{\tilde{U}/\tilde{\alpha}\}$ is of type T_2 . Thus in order to introduce an intersection type, we have to type-check A twice. The type annotations in A can differ between the two checks. The introduction of intersection types is useful, for instance, when type-checking a function that can be used by honest participants as well as by the attacker. For instance, the function for verifying digital signatures has a type of the form $\text{Un} \rightarrow ((T \vee \text{Un}) \rightarrow T) \wedge (\text{Un} \rightarrow \text{Un})$. Honest participants pass a signature together with the signed message, either at type T or at type Un , and get back the signed message with the stronger type T . The attacker that has only access to messages of type Un can still call this function, but the returned type is just Un .

The main result we expect for the type system is that well-typed programs are robustly safe. We are currently in the process of completing the proofs.

Claim (Robust Safety). If $\emptyset \vdash A : \text{Un}$ then A is robustly safe.

3.5 Example

Let us consider again the protocol implementation from Section 2.2. Since we want to model a scenario in which the attacker has control over the network we give channel c type Un (i.e., $T_{\text{chan}} = \text{Un}$). The function $mkId$ returns a public

identifier and it is thus of type $\text{unit} \rightarrow \text{Un}$. The function *getPriv* is instead of type $\text{unit} \rightarrow \text{Private}$, since it returns a private value. Notice that the sender signs a pair composed of a public identifier and a private message y for which the predicate $Ok(y)$ holds. In order to convey the predicate $Ok(y)$ from the sender to the receiver, we instantiate the polymorphic function for creating signing keys with the type $T_{\text{sign}} = \text{Un} * \{z : \text{Private} \mid Ok(z)\}$. Since the sender encrypts the signature together with the two signed values, the encryption function is instantiated with $T_{\text{enc}} = \text{Un} * T_{\text{sign}}$.

We first analyze the sender’s code. The message y returned by *getPriv* is of type Private . After assume $Ok(y)$, we can give y type $\{z : \text{Private} \mid Ok(z)\}$ by applying `VAL REFINE`. The type of the receiver’s identifier y_b and the type of y comply with the type of the signing and encryption keys. The sender’s code is thus well-typed.

The receiver decrypts the ciphertext and the result is stored in variable xy . This variable is of type $T_{\text{enc}} \vee \text{Un}^2$, which is equivalent by subtyping to $\text{Un} * (T_{\text{sign}} \vee \text{Un})$. The receiver splits the pair, obtaining a variable x of type Un bound to the signature and a variable y of type $T_{\text{sign}} \vee \text{Un}$ bound to the signed pair. We anticipate that these are precisely the types expected by the signature verification function, which returns a value y' of type T_{sign} , since the signing key is only used to sign messages of this type. The signed pair y' is split, obtaining a variable y_1 of type Un bound to the receiver’s identifier and a variable y_2 of type $\{z : \text{Private} \mid Ok(z)\}$ bound to the private message. This refinement type allows us to successfully type-check the assert $Ok(y_2)$. The receiver’s code is thus also well-typed, so the whole program is well-typed and hence we expect it is robustly safe.

4 Implementation of Symbolic Cryptography

In this section, we illustrate the typed interface of the functions for digital signatures and public-key cryptography. Section 4.1 overviews the dynamic sealing mechanism used in [7] to encode symbolic cryptography, while sections 4.2 and 4.3 present our improvements to this encoding.

4.1 Dynamic Sealing

The notion of *dynamic sealing* was initially introduced by Morris [16] as a protection mechanism for programs. Later, Sumii and Pierce [19,18] studied the semantics of dynamic sealing within a λ -calculus, observing a close correspondence with symmetric-key cryptographic primitives.

In RCF [7] seals are encoded using pairs, functions, references and lists. A seal is a pair of a *sealing function* and an *unsealing function*, having type:

$$\text{Seal } \langle \alpha \rangle = (\alpha \rightarrow \text{Un}) * (\text{Un} \rightarrow \alpha).$$

² The ciphertext might have been generated by a honest user or by the attacker.

The sealing function takes as input a term M of type α and returns a fresh value a of type Un , after adding the pair (M, a) to a secret list that is stored in a reference. The unsealing function takes as input a value a of type Un , scans the list in search of a pair (M, a) , and returns M . Only the sealing function and the unsealing function can access this secret list. In RCF, each key-pair is (symbolically) implemented by means of a seal. In the case of public-key cryptography, for instance, the sealing function is used for encrypting, the unsealing function is used for decrypting, and the sealed value a represents the ciphertext.

Let us take a look at the type $\text{Seal } \langle \alpha \rangle$. If α is neither public nor tainted, as it is usually the case for symmetric-key cryptography, neither the sealing function nor the unsealing function are public, meaning that the symmetric key is kept secret. If α is tainted but not public, as usually the case for public-key cryptography, the sealing function is public but the unsealing function is not, meaning that the encryption key may be given to the adversary but the decryption key is kept secret. If α is public but not tainted, as typically the case for digital signatures, the sealing function is not public and the unsealing function is public, meaning that the signing key is kept secret but the verification key may be given to the adversary.

Although this unified interpretation of cryptography as sealing and unsealing functions is conceptually appealing, it actually exhibits some undesired side-effects when modeling asymmetric cryptography. If the type of a signed message is not public, then the verification key is not public either and cannot be given to the adversary. This is unrealistic, since in most cases verification keys are public no matter what is signed. As an example, in the Direct Anonymous Attestation protocol [12] the issuer signs the secret TPM's identifier and the TPM proves the knowledge of this certificate without revealing it by means of a zero-knowledge proof. The verifier, however, needs to have access to the TPM's verification key in order to verify the zero-knowledge proof.

Moreover, if the type of a message encrypted with a public key is not tainted, then the public key is not public and cannot be given to the adversary. This may be problematic, for instance, when modeling authentication protocols based on public keys and nonce handshakes, such as the Needham-Schroeder-Lowe protocol, where the type of the encrypted messages is neither public nor tainted.

These issues are not due to the sealing itself but to the type of seals, which is simple but not expressive enough to characterize some important usages of asymmetric cryptography.

4.2 Digital Signatures

In this section, we show how union and intersection types can be used to enhance the expressiveness of the type system and solve the aforementioned problems. The type of the seal used in our library for digital signatures is reported below:

$$\begin{aligned} \text{Seal}_{\text{sign}} \langle \alpha \rangle &= s : \text{Un} * \text{Sealing}_{\text{sign}} \langle \alpha \rangle * \text{Unsealing}_{\text{sign}} \langle \alpha \rangle \\ \text{Sealing}_{\text{sign}} \langle \alpha \rangle &= \alpha \rightarrow \text{Un} \\ \text{Unsealing}_{\text{sign}} \langle \alpha \rangle &= \text{Un} \rightarrow ((x : (\alpha \vee \text{Un}) \rightarrow \{y : \alpha \mid x = y\}) \wedge (\text{Un} \rightarrow \text{Un})) \end{aligned}$$

The seal is a triple composed of a seal identifier s of type Un , the sealing function, and the unsealing function. The identifier s is a “ghost variable” that is only used in the refinement types given to the sealing and unsealing function. The logical formulas occurring therein link the sealing function to the unsealing function and sealed values to the corresponding seal. This serves to characterize in the logic some important properties of cryptographic primitives, such as the determinism of the functions implementing signature verification and decryption. For easing the presentation, we will present simpler types obtained by removing such logical formulas by subtyping.

The implementation and the type of the sealing function for digital signatures is similar to the one in [7]. The unsealing function, however, differs both in the implementation and in the type. This function takes as input two arguments: the sealed value a representing a signature and the message M that was signed. Unsealing succeeds only if the pair (M, a) is in the secret list associated to the seal, and in this case it returns M . As described by the type $\text{Unsealing}_{\text{sign}}\langle\alpha\rangle$, the first argument of the unsealing function is of type Un and corresponds to the signature. The second part of this type is an intersection of two types: The type $x : (\alpha \vee \text{Un}) \rightarrow \{y : \alpha \mid x = y\}$ is used to type-check honest callers: the signed message x passed to the unsealing function has either type α (e.g., if the message is received encrypted, as in our running example, or from a private channel) or of type Un (e.g., if the signature is received from a public channel), and the message y returned by the unsealing function has the stronger type α , which means that the unsealing function casts the type of the signed message from $(\alpha \vee \text{Un})$ down to α . This is safe since the sealing function is not public and can only be used to sign messages of type α . The type $\text{Un} \rightarrow \text{Un}$ makes type $\text{Unsealing}_{\text{sign}}\langle\alpha\rangle$ always public, which allows the attacker to call the unsealing function. Since, in contrast to [7], the unsealing function and hence the verification key can be made public, we can additionally model protocols where the signing key is used to sign private messages while the verification key is public, such as DAA.

Finally, we present our typed interface for digital signatures:

$$\begin{aligned}
mkSK\langle\alpha\rangle &: \text{unit} \rightarrow \text{Seal}_{\text{sign}}\langle\alpha\rangle \\
mkVK\langle\alpha\rangle &: (x_{sk} : \text{Seal}_{\text{sign}}\langle\alpha\rangle \rightarrow \{x_{vk} : \text{Unsealing}_{\text{sign}}\langle\alpha\rangle \mid SKPair(x_{vk}, x_{sk})\}) \wedge \text{Un} \\
sign\langle\alpha\rangle &: (x_{sk} : \text{Seal}_{\text{sign}}\langle\alpha\rangle \rightarrow y : \alpha \rightarrow \{z : \text{Un} \mid Signed(x_{sk}, y, z)\}) \wedge \text{Un} \\
check\langle\alpha\rangle &: (x_{vk} : \text{Unsealing}_{\text{sign}}\langle\alpha\rangle \rightarrow z : \text{Un} \rightarrow x : (\alpha \vee \text{Un}) \rightarrow \\
&\quad \{y : \alpha \mid y = x \wedge \exists S. SKPair(x_{vk}, S) \wedge Signed(S, x, z)\}) \wedge \text{Un}
\end{aligned}$$

The $mkSK$ function generates a signing key, which is just a seal, while $mkVK$ takes this seal and returns the verification key, which is just the unsealing component of the seal. The $sign$ function takes as input the signing key x_{sk} and a message y , and it applies the sealing function to the message obtaining a sealed value z of type $\{z : \text{Un} \mid Signed(x_{sk}, y, z)\}$. The predicate in this refinement type records the signing operation. The $check$ function takes as input the verification key x_{vk} , a sealed value z , and the signed message x , and it returns a value of type $\{y : \alpha \mid y = x \wedge \exists S. SKPair(x_{vk}, S) \wedge Signed(S, x, z)\}$. The formula in this refinement type says that the value returned by this function is the signed mes-

sage x and that there exists a signing key S associated to x_{vk} such that z is the value obtained by signing x with S . Notice that we give $mkVK$, $sign$ and $check$ intersection types similar to $Unsealing_{sign}\langle\alpha\rangle$.

4.3 Public-Key Encryption

For public-key encryption we use a seal of type $Seal\langle\alpha \vee Un\rangle$. The sealing function takes as input a value of type α or one of type Un , depending on whether the encryption is performed by a honest user or by the attacker, and it returns a sealed value of type Un . In contrast to [7], the sealing function is always public, even if the type α of the encrypted message is not tainted. Finally, we present the typed interface of the functions implementing public-key cryptography. The formulas in the refinement types are similar to the ones for digital signatures.

$$\begin{aligned}
mkDK\langle\alpha\rangle &: unit \rightarrow Seal\langle\alpha \vee Un\rangle \\
mkEK\langle\alpha\rangle &: (x_{dk} : Seal\langle\alpha \vee Un\rangle) \rightarrow \{x_{ek} : Sealing\langle\alpha \vee Un\rangle \mid EKPair(x_{ek}, x_{dk})\} \wedge Un \\
encrypt\langle\alpha\rangle &: (x_{ek} : Sealing\langle\alpha \vee Un\rangle) \rightarrow y : (\alpha \vee Un) \rightarrow \{x : Un \mid Encrypted(x_{ek}, y, x)\} \wedge Un \\
decrypt\langle\alpha\rangle &: (x_{dk} : Seal\langle\alpha \vee Un\rangle) \rightarrow x : Un \rightarrow \\
&\quad \{y : \alpha \vee Un \mid \exists E. EKPair(E, x_{dk}) \wedge Encrypted(E, y, x)\} \wedge Un
\end{aligned}$$

5 Encoding of Zero-knowledge

This section describes how we automatically generate the symbolic implementation of non-interactive zero-knowledge proofs, starting from a high-level specification. Intuitively, this implementation resembles an oracle that provides three operations: one for creating zero-knowledge proofs, one for verifying such proofs, and one for obtaining the public values used to create the proofs.

For creating a zero-knowledge proof the caller needs to provide values for the variables in the specification. Some of these values are revealed by the proof to the verifier and to any eavesdropper, while the others (i.e., the witnesses of the proof) are kept secret. A zero-knowledge proof does not reveal any information about these witnesses, other than the validity (or invalidity) of the statement that is being proved. A zero-knowledge proof is valid if the values that are used to create it satisfy the statement of the proof. When creating a proof, however, we do not require that it is valid. Instead, a second operation is provided for verifying the validity of zero-knowledge proofs. The third operation allows one to obtain the public values of zero-knowledge proofs. Note that these three operations need to be available not only to the honest participants, but also to the attacker.

We implement such a zero-knowledge oracle in $RCF_{\wedge\vee}^{\forall}$ as three public functions that share a secret seal. In order to create a zero-knowledge proof the first function seals the witnesses and public values provided by the caller all together and returns a sealed value representing the non-interactive zero-knowledge proof, which can be sent to the verifier. The verification function unseals the sealed values, and checks if they indeed satisfy the statement by performing the corresponding cryptographic and logical operations. If verification succeeds then the verification function returns the public values of the proof. The public values can also be obtained with the third function, without checking the validity of the proof.

5.1 Example: Simplified DAA

As a running example throughout this section, we consider a simplified version of the Direct Anonymous Attestation (DAA) protocol [12]. The goal of the DAA protocol is to enable the TPM to sign arbitrary messages and to send them to an entity called the verifier in such a way that the verifier will only learn that a valid TPM signed that message, but without revealing the TPM's identity. The DAA protocol is composed of two sub-protocols: the *join protocol* and the *DAA-signing protocol*. The join protocol allows a TPM to obtain a certificate x_{cert} from an entity called the issuer. This certificate is just a signature on the TPM's secret identifier x_f . The DAA-signing protocol enables a TPM to authenticate a message y_m by proving to the verifier the knowledge of a valid certificate, but without revealing the TPM's identifier or the certificate. In this section, we focus on the DAA-signing protocol and we assume that the TPM has already completed the join protocol and received the certificate from the issuer. In the DAA-signing protocol the TPM sends to the verifier the following zero-knowledge proof:

$$\begin{array}{ccc}
 \text{TPM} & & \text{Verifier} \\
 \text{assume Send}(x_f, y_m) & & \\
 \xrightarrow{\text{zk}_{\mathcal{S}_{d\text{aa}}}(x_f, x_{cert}, y_{vki}, y_m, y_\zeta, y_N)} & \longrightarrow & \text{assert Authenticate}(y_m)
 \end{array}$$

where the statement $S_{\mathcal{S}_{d\text{aa}}}$ is of the form $x_f = \text{check } y_{vki} \ x_{cert} \ x_f \wedge y_N = \text{exp } y_\zeta \ x_f$. Intuitively, in the first conjunct the TPM proves the knowledge of a certificate x_{cert} of its identifier x_f that can be verified with the verification key y_{vki} of the issuer. The TPM identifier and the certificate are kept secret, while the verification key of the issuer is revealed to the verifier. This zero-knowledge proof additionally conveys a public message y_m that the TPM wants to authenticate with the verifier. Notice that this proof guarantees non-malleability, i.e., the attacker cannot change y_m without redoing the proof, since the certificate and the identifier are kept secret. The second conjunct states that y_N is obtained by exponentiating y_ζ to x_f . If y_ζ is freshly chosen for each DAA-signature then the anonymity of the user is preserved. If, however, y_ζ is fixed for each verifier then two DAA-signatures done by the same TPM can be linked by the verifier, yielding the pseudonymous version of DAA.

Before sending the zero-knowledge proof, the TPM assumes $\text{Send}(x_f, y_m)$. After verifying the zero-knowledge proof, the verifier asserts $\text{Authenticate}(y_m)$. The authorization policy we consider for the DAA-sign protocol is

$$\forall x_f, x_{cert}, y_m. \text{Send}(x_f, y_m) \wedge \text{OkTPM}(x_f) \Rightarrow \text{Authenticate}(y_m)$$

where the predicate $\text{OkTPM}(x_f)$ is assumed by the issuer before signing x_f .

5.2 High-level Specification

Our high-level specification of non-interactive zero-knowledge proofs is similar in spirit to the symbolic representation of zero-knowledge proofs in a process

calculus [6,4]. For a specification \mathcal{S} the user needs to provide: (1) variables representing the witnesses and public values of the proof, (2) a Boolean formula over these variables representing the statement of the proof, (3) types for the variables, and, if desired, (4) a promise, i.e., a logical formula that is conveyed by the proof only if the prover is honest.

Variables. We use variables to stand for the *witnesses* and public values of a zero-knowledge proof. For the purpose of typing, we further make a distinction between the public values that are checked for equality by the verifier – represented by *matched* public variables, and the ones that are obtained as the result of the verification – represented by *returned* public variables.

In the DAA example, the variables x_f and x_{cert} stand for witnesses. The value of y_{vki} is matched against the signature verification key of the issuer, which the verifier of the zero-knowledge proof already knows. The payload message y_m , as well as the values of y_N and y_ζ are returned to the verifier of the proof, so they are returned variables.

Statement. We assume that the statement conveyed by a zero-knowledge proof for specification \mathcal{S} is a positive Boolean formula $S_{\mathcal{S}}$. Statements are formed using equalities between variables and $\text{RCF}_{\wedge\vee}^{\forall}$ functions applied to variables, as well as conjunctions and disjunctions of such basic statements.

$$S, R ::= (x = f\langle\tilde{T}\rangle x_1 \dots x_n) \mid S_1 \wedge S_2 \mid S_1 \vee S_2$$

Intuitively, a statement is valid for a certain instantiation of the variables if after substituting all variables with the corresponding values and applying all $\text{RCF}_{\wedge\vee}^{\forall}$ functions to their arguments we obtain a valid Boolean formula. We assume that the $\text{RCF}_{\wedge\vee}^{\forall}$ functions occurring in the statement have deterministic behaviour³, i.e., when called twice with the same arguments they return the same value.

For example, the statement of the zero-knowledge proof in the DAA-signing protocol is $S_{\mathcal{S}_{daa}} = (x_f = \text{check}\langle T_{vki}\rangle y_{vki} x_{cert} x_f \wedge y_N = \text{exp } y_\zeta x_f)$. This statement is valid for a certain instantiation if the check function returns the value of x_f when the values of y_{vki} , x_{cert} , and x_f are passed as arguments, and if $\text{exp } y_\zeta x_f$ returns y_N . Note that although the payload message y_m does not occur in the statement, the proof guarantees non-malleability so an attacker cannot change y_m without redoing the proof.

Types. The user also needs to provide a type for all specified variables. In the following we assume a function $t_{\mathcal{S}}$ that assigns a type to each variable in specification \mathcal{S} . The DAA-sign protocol does not preserve the secrecy of the signed message, so $t_{\mathcal{S}_{daa}}(y_m) = \text{Un}$. Similarly, we have that $t_{\mathcal{S}_{daa}}(y_\zeta) = t_{\mathcal{S}_{daa}}(y_N) = \text{Un}$. On the other hand, the TPM identifier x_f is given a secret and untainted type: $t_{\mathcal{S}_{daa}}(x_f) = T_{vki} = \{z_f : \text{Private} \mid \text{OkTPM}(z_f)\}$. This ensures that x_f is not known to the attacker and that it is certified by the issuer (i.e., the predicate $\text{OkTPM}(x_f)$ holds). The verification key of the issuer is used to check signed

³ In order to model randomized functions one can take the random seed as an explicit argument.

messages of type T_{vki} , so it is given type $\text{Unsealing}_{\text{sign}} \langle T_{vki} \rangle$. Finally the certificate x_{cert} is a signature, so it has type Un . Even though it has type Un , it would break the anonymity of the user to make the certificate a public value, since the verifier could then always distinguish if two consecutive requests come from the same user or not (as in the pseudonymous version of DAA).

Promise. The user can additionally specify a *promise*: an arbitrary formula in the authorization logic that holds in the typing environment of the prover. If the statement is strong enough to identify the prover as a honest (type-checked) protocol participant⁴, then the promise can be safely transmitted to the typing environment of the verifier. For a specification \mathcal{S} we denote the promise by $P_{\mathcal{S}}$. In the DAA example we have that $P_{\mathcal{S}_{\text{daa}}} = \text{Send}(x_f, y_m)$, since this predicate holds true in the typing environment of a honest TPM.

5.3 Automatic Code Generation

We automatically generate both a typed interface and a symbolic implementation for the oracle corresponding to a zero-knowledge specification.

Typed Interface The interface generated for a specification \mathcal{S} contains three functions⁵ that share hidden state (a seal for values of type $\tau_{\mathcal{S}}$):

$$\begin{aligned} \text{create}_{\mathcal{S}} &: \tau_{\mathcal{S}} \rightarrow \text{Un}, \text{ where } \tau_{\mathcal{S}} = \text{Un} \vee \sum_{x \in \text{vars}_{\mathcal{S}}} x : t_{\mathcal{S}}(x). \{P_{\mathcal{S}}\} \\ \text{public}_{\mathcal{S}} &: \text{Un} \rightarrow \text{Un} \\ \text{verify}_{\mathcal{S}} &: \text{Un} \rightarrow (\text{Un} \wedge \prod_{y \in \text{matched}_{\mathcal{S}}} y : t_{\mathcal{S}}(y). \\ &\quad \sum_{z \in \text{returned}_{\mathcal{S}}} z : t_{\mathcal{S}}(z). \{ \exists \tilde{x}. F(S_{\mathcal{S}}, E) \wedge P_{\mathcal{S}} \}) \end{aligned}$$

The function $\text{create}_{\mathcal{S}}$ is used to create zero-knowledge proofs for specification \mathcal{S} . It takes as argument a tuple containing values for all variables of the proof, or an argument of type Un if it is called by the adversary. In case a protocol participant calls this function, we check that the values have the types provided by the user. Additionally, we check that the promise $P_{\mathcal{S}}$ provided by the user holds in the typing environment of the prover. The returned zero-knowledge proof is given type Un so that it can be sent over the public network. For instance, in the DAA example we have that: $\tau_{\mathcal{S}_{\text{daa}}} = \text{Un} \vee ((y_{vki} : \text{Unsealing}_{\text{sign}} \langle T_{vki} \rangle * y_m : \text{Un} * y_C : \text{Un} * y_N : \text{Un} * x_f : T_{vki} * x_{\text{cert}} : \text{Un}) * \{\text{Send}(x_f, y_m)\})$, where $T_{vki} = \{z_f : \text{Private} \mid \text{OkTPM}(z_f)\}$.

The function $\text{public}_{\mathcal{S}}$ is used to read the public values of a zero-knowledge proof for \mathcal{S} , so it takes as input the sealed proof of type Un and returns the tuple of public values, also at type Un .

⁴ Signature proofs of knowledge have this property [12].

⁵ We use $\sum_{x \in \text{vars}_{\mathcal{S}}} x : t_{\mathcal{S}}(x). \{P_{\mathcal{S}}\}$ to denote the nested dependent pair type $x_1 : t_{\mathcal{S}}(x_1) * \dots * x_n : t_{\mathcal{S}}(x_n) * \{P_{\mathcal{S}}\}$ where $\tilde{x} = \text{vars}_{\mathcal{S}}$, and $\prod_{y \in \text{matched}_{\mathcal{S}}} y : t_{\mathcal{S}}(y)$. T to denote the dependent function type $y_1 : t_{\mathcal{S}}(y_1) \rightarrow \dots \rightarrow y_m : t_{\mathcal{S}}(y_m)$, where $\tilde{y} = \text{matched}_{\mathcal{S}}$.

The function $\text{verify}_{\mathcal{S}}$ is used for verifying zero-knowledge proofs. This function can be called by the attacker in which case it returns a value of type Un . When called by a protocol participant, however, it takes as argument a candidate zero-knowledge proof of type Un and the values for the matched variables, which have the user-specified types. On successful verification, this function returns a tuple containing the values of the public variables, again with their respective types. The function guarantees that the formula $\exists \tilde{x}. F(\mathcal{S}_{\mathcal{S}}, E) \wedge P_{\mathcal{S}}$ holds, where the public variables are free and the witnesses are existentially quantified. The first conjunct, $F(\mathcal{S}_{\mathcal{S}}, E)$, guarantees that if verification succeeds then the statement indeed holds, no matter what the origin of the proof is. Since the statement itself is not a formula in the logic (as it was for instance the case in [4]), we use a transformation function F that computes the formula conveyed by the statement. This transformation is straightforward: it extracts the formulas guaranteed by the dependently-typed cryptographic functions (the post-conditions) and combines them using the corresponding logical connectives of the authorization logic.

For instance, in the DAA example, we have that

$$F(\mathcal{S}_{\text{daa}}, E_{\text{std}}) = F(x_f = \text{check}\langle T_{vki} \rangle y_{vki} x_{\text{cert}} x_f, E_{\text{std}}) \wedge F(y_N = \text{exp } y_{\zeta} x_f, E_{\text{std}})$$

The second conjunct is equal to $\text{Exp}(y_{\zeta}, x_f, y_N)$, since the exponentiation function is typed to $x : \top \rightarrow y : \top \rightarrow \{z : \text{Un} \mid \text{Exp}(x, y, z)\}$ in our library. As explained in Section 4, we have that $E_{\text{std}} \vdash \text{check}\langle T_{vki} \rangle : xvk : \text{Unsealing}_{\text{sign}}\langle T_{vki} \rangle \rightarrow z : \text{Un} \rightarrow x : (T_{vki} \vee \text{Un}) \rightarrow \{y : T_{vki} \mid y = x \wedge \exists SK. \text{SKPair}(xvk, SK) \wedge \text{Signed}(SK, x, z)\}$. So for the first conjunct we obtain the formula: $(x_f = x_f) \wedge (\exists sk. \text{SKPair}(y_{vki}, sk) \wedge \text{Signed}(sk, x_f, x_{\text{cert}})) \wedge \text{OkTPM}(x_f)$. The predicate $\text{OkTPM}(x_f)$ was obtained from the nested refinement type T_{vki} . Finally, after removing the trivial equality we obtain that:

$$F(\mathcal{S}_{\text{daa}}, E_{\text{std}}) = (\exists sk. \text{SKPair}(y_{vki}, sk) \wedge \text{Signed}(sk, x_f, x_{\text{cert}})) \wedge \text{OkTPM}(x_f) \wedge \text{Exp}(y_{\zeta}, x_f, y_N)$$

Generated Implementation The generated implementation for specification \mathcal{S} creates a fresh seal $k_{\mathcal{S}}$ for values of type $\tau_{\mathcal{S}} = \text{Un} \vee \sum_{x \in \text{vars}_{\mathcal{S}}} x : t_{\mathcal{S}}(x). \{P_{\mathcal{S}}\}$. The sealing function of $k_{\mathcal{S}}$ is directly used to implement the $\text{create}_{\mathcal{S}}$ function. The unsealing function of $k_{\mathcal{S}}$ is used to implement the $\text{public}_{\mathcal{S}}$ and $\text{verify}_{\mathcal{S}}$ functions. The implementation of $\text{public}_{\mathcal{S}}$ is also simple: since the zero-knowledge proof is just a sealed value, $\text{public}_{\mathcal{S}}$ unseals it and returns the public values as a tuple. The witnesses are discarded, and the validity of the statement is not checked.

The implementation of the $\text{verify}_{\mathcal{S}}$ function is more involved. This function takes a candidate zero-knowledge proof z of type Un as input, and values for the matched variables. Since the type of $\text{verify}_{\mathcal{S}}$ contains an intersection type (Un is one of the branches and this makes the type of $\text{verify}_{\mathcal{S}}$ public) we use a for construct to introduce this intersection type. If the proof is verified by the attacker we can assume that the matched variables have type Un and need to type the return

value to Un . On the other hand, if the proof is verified by a protocol participant we can assume that for all $y' \in \text{matched}_{\mathcal{S}}$ we have $y' : t_{\mathcal{S}}(y')$, and need to give the returned value type $\sum_{y \in \text{returned}_{\mathcal{S}}} y : t_{\mathcal{S}}(y) \cdot \{\exists \tilde{x} = \text{witness}_{\mathcal{S}} \tilde{x}. F(S_{\mathcal{S}}, E) \wedge P_{\mathcal{S}}\}$. Intuitively, the strong types of the matched values should allow us to guarantee the strong types of the returned public values, as well as the promise $P_{\mathcal{S}}$.

The generated $\text{verify}_{\mathcal{S}}$ function performs the following five steps: (1) it unseals z using “ $\text{snd } k_{\mathcal{S}}$ ” and obtains z' ; (2) since z' has a union type, it does case analysis on it, and assigns its value to z'' ; (3) it splits the tuple z'' into the public values $(y_{vki}, y_m, y_{\zeta}, y_N)$ and the witnesses $(x_f$ and $x_{cert})$. (4) it tests if the matched variable y_{vki} is equal to the argument y'_{vki} , and in case of success assigns the value to the variable y''_{vki} – since y''_{vki} has a stronger type than y'_{vki} and y_{vki} we use this new variable to stand for y_{vki} in the following; (5) it tests if the statement is true by applying the functions in $S_{\mathcal{S}_{\text{daa}}}$ and checking the results for equality with the corresponding values. This last step is slightly complicated by the fact that the statement can contain disjunctions, so we use decision trees. However, for the DAA example the decision tree has a very simple (linear) structure.

```

verify $\mathcal{S}_{\text{daa}}$  =  $\lambda z : \text{Un}$ .
  for  $\alpha$  in  $\text{Un}$ ;  $\text{Unsealing}_{\text{sign}} \langle T_{vki} \rangle$  do  $\lambda y'_{vki} : \alpha$ .
    let  $z' = (\text{snd } k_{\mathcal{S}}) z$  in
      case  $z'' = z' : \text{Un} \vee (\text{Unsealing}_{\text{sign}} \langle T_{vki} \rangle * y_m : \text{Un} * \text{Un} * \text{Un} * x_f : T_{vki} * \text{Un} * \{\text{Send}(x_f, y_m)\})$  in
        let  $(y_{vki}, y_m, y_{\zeta}, y_N, x_f, x_{cert}) = z''$  in
          if  $(y_{vki}) = (y'_{vki})$  as  $(y''_{vki})$  then
            if  $x_f = \text{check} \langle T_{vki} \rangle y''_{vki} x_{cert} x_f$  as  $x'_f$  then
              if  $y_N = \text{exp } y_{\zeta} x_f$  as  $y'_N$  then
                 $(y_m, y_{\zeta}, y'_N, ())$ 
              else  $\text{failwith } ()$ 
            else  $\text{failwith } ()$ 
          else  $\text{failwith } ()$ 

```

Checking the Generated Implementation Since the automatically generated implementation of zero-knowledge proofs relies on types and formulas provided by the user, which may both be wrong, the generated implementation is not guaranteed to fulfill its interface. We use our type-checker to check whether this is indeed the case. If type-checking the generated code against its interface succeeds, then this code can be safely used in protocol implementations. Note that because of the for and case constructs the body of $\text{verify}_{\mathcal{S}}$ is type-checked four times, corresponding to the following four scenarios: honest prover / honest verifier, honest prover / dishonest verifier, dishonest prover / honest verifier, and dishonest prover / dishonest verifier.

In general, there are two situations in which type-checking the generated implementation fails. First, the types provided by the user for the public witnesses are not public. In this case the implementation of $\text{public}_{\mathcal{S}}$ cannot match its de-

defined type $Un \rightarrow Un$. Second, the promise $P_{\mathcal{G}}$ is not justified by the statement and the types of the witnesses. In this case $\text{verify}_{\mathcal{G}}$ cannot match its defined type.

6 Implementation

We have implemented a complete tool-chain for $\text{RCF}_{\wedge\vee}^{\forall}$: it includes a type-checker for the type system described in Section 3, an automatic code generator for zero-knowledge as defined in Section 5, an interpreter, and a visual debugger.

The type-checker supports an extended syntax with respect to the one in the paper, including: a simple module system, algebraic data types, recursive functions, type definitions, and mutable references. We use first-order logic with equality as the authorization logic and the type-checker invokes an automatic theorem prover to discharge proof obligations. We have tested the type-checker on the simplified DAA protocol from Section 5.1 and on other simple examples such as nonce handshakes; the analysis only took several seconds in each case. The type-checker produces an XML log file containing the complete type derivation in case of success, and a partial derivation that leads to the typing error in case of failure. This can be inspected using our visualizer to easily detect and fix flaws in the protocol implementation. The type-checker also performs limited type inference: it can infer the instantiation of some polymorphic functions from the type of the arguments, however, the user has to provide all the other typing annotations.

The type-checker, the code generator for zero-knowledge, the interpreter are command-line tools implemented in F#, while the GUIs of the visual debugger and the visualizer for type derivations are specified using WPF (Windows Presentation Foundation). The type-checker consists of around 2000 lines of code, while the whole tool-chain has over 4500 lines of code.

7 Conclusions and Future Work

We have presented a general technique for verifying implementations of cryptographic protocols based on zero-knowledge proofs. The security of $\text{RCF}_{\wedge\vee}^{\forall}$ protocol implementations is verified by a type system combining refinement, union, and intersection types. The analysis is automated and modular, and was efficient on the examples that we tried.

As future work, we plan to investigate the automated generation of concrete cryptographic implementations of zero-knowledge proofs, and thus to complement the generation of symbolic implementations considered in this paper.

Finally, we intend to apply our framework to analyze modern cryptographic applications, such as the full implementation of the Direct Anonymous Attestation protocol and the recently proposed Civitas electronic voting system [14].

Acknowledgments. We thank Cédric Fournet and Andrew D. Gordon for many constructive discussions.

References

1. M. Abadi and B. Blanchet. Secrecy types for asymmetric communication. In *Proc. 4th International Conference on Foundations of Software Science and Computation Structures (FOSSACS)*, pages 25–41. Springer-Verlag, 2001.
2. M. Abadi and B. Blanchet. Analyzing security protocols with secrecy types and logic programs. *Journal of the ACM*, 52(1):102–146, 2005.
3. M. Abadi and C. Fournet. Mobile values, new names, and secure communication. In *Proc. 28th Symposium on Principles of Programming Languages (POPL)*, pages 104–115. ACM Press, 2001.
4. M. Backes, C. Hrițcu, and M. Maffei. Type-checking zero-knowledge. In *15th ACM Conference on Computer and Communications Security (CCS 2008)*, pages 357–370. ACM Press, 2008.
5. M. Backes, C. Hrițcu, M. Maffei, and T. Tarrach. Type-checking implementations of protocols based on zero-knowledge proofs. Long version available at <http://www.infsec.cs.uni-sb.de/~hritcu/publications/zk-rcf-full.pdf>, 2009.
6. 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.
7. J. Bengtson, K. Bhargavan, C. Fournet, A. D. Gordon, and S. Maffei. Refinement types for secure implementations. In *Proc. 21th IEEE Symposium on Computer Security Foundations (CSF)*, pages 17–32, 2008.
8. K. Bhargavan, R. Corin, P.-M. D nielou, C. Fournet, and J. J. Leifer. Cryptographic protocol synthesis and verification for multiparty sessions. In *Proc. 22th IEEE Symposium on Computer Security Foundations (CSF)*, 2009. To appear.
9. K. Bhargavan, R. Corin, C. Fournet, and E. Z linescu. Cryptographically verified implementations for TLS. In *15th ACM Conference on Computer and Communications Security (CCS 2008)*, pages 459–468. ACM Press, 2008.
10. K. Bhargavan, C. Fournet, A. D. Gordon, and S. Tse. Verified interoperable implementations of security protocols. In *Proc. 19th IEEE Computer Security Foundations Workshop (CSFW)*, pages 139–152. IEEE Computer Society Press, 2006.
11. 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.
12. 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.
13. S. Chaki and A. Datta. ASPIER: An automated framework for verifying security protocol implementations. Technical report, CMU CyLab, October 2008.
14. 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.
15. J. Goubault-Larrecq and F. Parrennes. Cryptographic protocol analysis on real C code. In *6th International Conference on Verification, Model Checking, and Abstract Interpretation, (VMCAI 2005)*, pages 363–379. Springer, 2005.
16. J. Morris. Protection in programming languages. *Communications of the ACM*, 16(1):15–21, 1973.
17. B. C. Pierce. Intersection types and bounded polymorphism. *Mathematical Structures in Computer Science*, 7(2):129–193, 1997.
18. E. Sumii and B. Pierce. A bisimulation for dynamic sealing. *Theoretical Computer Science*, 375(1-3):169–192, 2007.
19. E. Sumii and B. C. Pierce. Logical relations for encryption. *Journal of Computer Security*, 11(4):521–554, 2003.