

Union and Intersection Types for Secure Protocol Implementations

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

¹ Saarland University

² Max Planck Institute for Software Systems (MPI-SWS)

Abstract. We present a new type system for verifying the security of cryptographic protocol implementations. The type system combines prior work on refinement types, with union, intersection, and polymorphic types, and with the novel ability to reason statically about the disjointness of types. The increased expressivity enables the analysis of important protocol classes that were previously out of scope for the type-based analyses of protocol implementations. In particular, our types can statically characterize: *(i)* more usages of asymmetric cryptography, such as signatures of private data and encryptions of authenticated data; *(ii)* authenticity and integrity properties achieved by showing knowledge of secret data; *(iii)* applications based on zero-knowledge proofs. The type system comes with a mechanized proof of correctness and an efficient type-checker.

1 Introduction

Modern applications are mostly distributed and they rely on complex cryptographic protocols to transmit data over potentially insecure networks (e.g., e-banking, e-commerce, social networks, and mobile applications). Protocol designers struggle to keep pace with the variety of possible security vulnerabilities, which have affected early authentication protocols like Needham-Schroeder [26, 37], carefully designed de facto standards like SSL and PKCS [17, 45], and even widely deployed products like Microsoft Passport [30] and Kerberos [19]. Even if the underlying cryptographic protocols are properly designed, security vulnerabilities may still arise due to flaws in the implementation. Manual security analyses of cryptographic protocols, and even more so protocol implementations, are extremely difficult and error-prone. Therefore, it is important to devise automated analysis techniques that can provide security guarantees for protocol implementations and, more generally, for the source code of distributed applications.

An effective approach for analyzing protocol implementations is to rely on software verification techniques, such as model checking and type theory, and to adapt them to the security problem. Type systems, in particular, proved successful in the automated analysis of both cryptographic protocol models [1, 2, 31] and protocol implementations [12, 14]. Type systems provide security proofs for an unbounded number of runs. Furthermore, the analysis is modular and has a predictable termination behavior. Finally, type systems were designed from the beginning to efficiently deal with programming language features such as data structures, recursion, state, and higher-order functions: consequently, type systems are more efficient and scale better than many

state-of-the-art protocol verifiers (e.g., ProVerif [16] used as a back end by fs2pv [15]) in the analysis of source code [14].

Despite these promising features, the type-based analysis of the source code of modern distributed applications is still an open issue. The first problem is that many of these applications (e.g., trusted computing [18], electronic voting [22], and social networks [9]) rely on complex cryptographic schemes, such as zero-knowledge proofs. Although the automated verification of protocols based on some of these schemes is possible in process calculi for abstract protocol specifications, which provide convenient mechanisms to symbolically abstract these schemes (e.g., flexible equational theories), 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 were, however, not designed for abstractly representing cryptographic primitives, which makes providing encodings that are suitable for automatic analysis and capture all potential usages of cryptographic schemes a challenging task. The second, somewhat similar, problem is that some interesting security properties are obtained by specific cryptographic patterns that are difficult to encode in type systems for programming languages. For instance, authenticity and integrity properties can be achieved by showing the knowledge of secret data, as in the Needham-Schroeder-Lowe public-key protocol [37] that relies on the exchange of secret nonces to authenticate the participants or as in most authentication protocols based on zero-knowledge proofs (e.g., Direct Anonymous Attestation [18] and Civitas [22]).

1.1 Contributions

This paper presents a new type system for the verification of the source code of protocol implementations. The underlying type theory combines refinement types [12] with union, intersection, and polymorphic types. Additionally, we introduce a novel relation for statically reasoning about the disjointness of types. This expressive type system extends the scope of existing type-based analyses of protocol implementations [12, 14] to important protocol classes that were not covered so far. In particular, our types statically characterize: (i) more usages of asymmetric cryptography, such as signatures of private data and encryptions of authenticated data; (ii) authenticity and integrity properties achieved by showing knowledge of secret data; (iii) applications based on zero-knowledge proofs.

Protocols are implemented in $\text{RCF}_{\wedge\vee}^{\forall}$ [12], a concurrent lambda-calculus, and cryptographic primitives are considered fully reliable building blocks and represented symbolically using a sealing mechanism [12, 35, 43]. In addition to hashes, symmetric cryptography, public-key encryption, and digital signatures, our approach supports zero-knowledge proofs. Since the realization of zero-knowledge proofs changes according to the statement to be proven, we provide a tool that, given a statement, automatically generates a symbolic implementation of the corresponding zero-knowledge primitive.

Our type-based analysis is automated, modular, efficient, and provides security proofs for an unbounded number of sessions. We have implemented a type-checker that performed very well in our experiments: it type-checks all our symbolic libraries and samples totaling more than 1500LOC in around 12 seconds, on a normal laptop. The

type-checker features a user-friendly graphical interface for examining typing derivations. The tool-chain we have developed additionally contains an automatic code generator for zero-knowledge proofs, an interpreter, and a visual debugger.

We have formalized the type system, and all the important parts of the soundness proof in the Coq proof assistant. The formalization and the implementation are available online [7].

1.2 Related Work

Our type system extends the refinement type system by Bengtson et al. [12] with union, intersection, and polymorphic types. We also encode a novel type `Private`, which is used to characterize data that are not known to the attacker. A crucial property is that the set of values of type `Private` is disjoint from the set of values of type `Un`, which is the type of the messages known to the attacker. This property allows us to prune typing derivations following equality tests between values of type `Private` and values of type `Un`. This technique was first proposed by Abadi and Blanchet in their seminal work on secrecy types for asymmetric cryptography [1], but later disappeared in the more advanced type systems for authorization policies. Our extension is necessary to deal with protocols based on zero-knowledge proofs and to verify integrity and authenticity properties obtained by showing knowledge of secret data (e.g., the Needham-Schroeder-Lowe public-key protocol). In addition, our extension removes the restrictions that the type system proposed in [12] poses 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 were preventing the analysis of many interesting cryptographic applications, such as the Direct Anonymous Attestation protocol [18], which involves digital signatures on secret TPM identifiers.

In recent parallel work, Bhargavan et al. [14] have developed an additional cryptographic library for a simplified version of the type system proposed in [12]. This library does not rely on sealing but on datatype constructors and inductive logical invariants that allow for reasoning about symmetric and asymmetric cryptography, hybrid encryption, and different forms of nested cryptography. The aforementioned logical invariants are, however, fairly complex and have to be proven manually. Moreover, these logical invariants are global, which means that adding new cryptographic primitives could require re-proving the previous established invariants. Therefore, extending a symbolic cryptographic library in the style of [14] to new primitives requires expertise and a considerable human effort. In contrast, extending our sealing-based library does not involve any additional proof: one has just to find a well-typed encoding of the desired cryptographic primitive, which is relatively easy¹.

The main simplification Bhargavan et al. [14] propose over [12] is the removal of the kinding relation, which classifies types as public or tainted, and allows values of public types to also be given any tainted type by subsumption. While this simplification removes the last security-specific part of the type system, therefore making it more

¹ A master's student encoded the sophisticated cryptographic schemes used in the Civitas [22] electronic voting protocol (i.e., distributed decryption, plaintext equivalence tests, homomorphic encryptions, mix nets, and a variety of zero-knowledge proofs) in about three weeks [29].

standard, this change also requires attackers to be well-typed with respect to a carefully constructed attacker interface. In contrast, by retaining the kinding relation from [12] we also retain the property that *all* attackers are well-typed with respect to our type system (this property is usually called opponent typability). Despite these disadvantages, Bhargavan et al. [14] manage to solve some of the problems we address in this paper, without relying on union and intersection types, but instead using the logical connectives inside the refinement types. It would be interesting future work to try to combine the advantages of both approaches in a unified framework.

Backes et al. [11] have recently established a semantic correspondence for asymmetric cryptography between a library based on sealing and one based on constructors, showing that both libraries enjoy computational soundness guarantees.

Backes et al. [8] proposed a type system for statically analyzing security protocols based on zero-knowledge proofs in the setting of the Spi calculus. Zero-knowledge proofs are modeled using constructors and destructors. In an extension of this type system [6], union and intersection types are used to infer precise type information about the secret witnesses of zero-knowledge proofs. This is captured in a separate relation called statement verification, which is fairly complex and tailored to zero-knowledge proofs. In contrast, in our paper we encode zero-knowledge proofs symbolically using standard programming language primitives, and we type-check them using general typing rules.

Goubault-Larrecq and Parrennes developed a static analysis technique [32] based on pointer analysis and clause resolution for cryptographic protocols implemented in C. The analysis 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 proposed a technique [21] based on software model checking for the automated verification of protocols implemented in C. The analysis provides security guarantees 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 [15] for the verification of F# protocol implementations by automatically extracting ProVerif models [16]. The technique was successfully used to verify implementations of real-world cryptographic protocols such as TLS [13]. The analysis, however, is not compositional and is significantly less scalable than type-checking [14]. Furthermore, the considered fragment of F# is restrictive: it does not include higher-order functions, and it allows only for a very limited usage of recursion and state.

The more technical discussion about the related work on union and intersection types is postponed to §8.

1.3 Outline

The remainder of the paper is structured as follows. §2 gives an intuitive overview of our type system and exemplifies the most important concepts on a simple authentication protocol. §3 introduces the syntax of $\text{RCF}_{\wedge, \vee}^{\forall}$, the language supported by our type-checker. §4 presents the type system. §5 and §6 show how our type system can

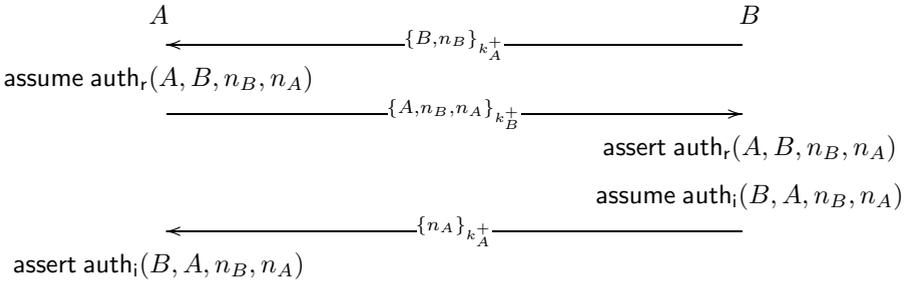
be used to obtain an expressive characterization of asymmetric cryptography and zero-knowledge proofs, respectively. §7 describes our implementation and experiments. §8 discusses some related work on union and intersection types. §9 concludes and gives some interesting research directions. We refer to the long version for the results of our Coq formalization, a more technical presentation of our encoding for zero-knowledge proofs, and other details [7].

2 Our Type System at Work

Before giving the details of the calculus and the type system, we illustrate the main concepts of our static analysis technique on the Needham-Schroeder-Lowe public-key protocol [37] (NSL), which could not be analyzed with previous refinement type systems for protocol implementations [12, 14]. For convenience, throughout this section we use some syntactic sugar that is supported by our type-checker and can be obtained from the core calculus presented in §3 by standard encodings [12].

2.1 Protocol Description and Security Annotations

The Needham-Schroeder-Lowe protocol is depicted below:



The goal of this protocol is to allow A and B to authenticate with each other and to exchange two fresh nonces, which are meant to be private and be later used to construct a session key. B creates a fresh nonce n_B and encrypts it together with his own identifier with A 's public key. A decrypts the ciphertext with her private key. At this point of the protocol, A does not know whether the ciphertext comes from B or from the opponent as the encryption key used to create the ciphertext is public. A continues the protocol by creating a fresh nonce n_A , and encrypts this nonce together with n_B and her own identifier with B 's public key. B decrypts the ciphertext and, although the encryption key used to create the ciphertext is public, if the nonce he received matches the one he has sent to A then B does indeed know that the ciphertext comes from A , since the nonce n_B is *private* and only A has access to it. Finally, B encrypts the nonce n_A received from A with A 's public key, and sends it back to A . After decrypting the ciphertext and checking the nonce, A knows that the ciphertext comes from B as the nonce n_A is *private* and only B has access to it.

Following [12], we decorate the code with assumptions and assertions. Intuitively, assumptions introduce new hypotheses, while assertions declare formulas that should

logically follow from the previously introduced hypotheses. A program is safe if in all program runs the assertions are entailed by the assumptions. The assumptions and assertions of the NSL protocol capture the standard mutual authentication property.

2.2 Types for Cryptography

Before illustrating how we can type-check this protocol, let us introduce the typed interface of our library for public-key cryptography. Intuitively, since encryption keys are public, they can be used by honest principals to encrypt data as specified by the protocol, or by the attacker to encrypt arbitrary data. This intuitive reasoning is captured by the following typed interface:

$$\begin{aligned} \text{encrypt} &: \forall \alpha. \text{PubKey}\langle \alpha \rangle \rightarrow \alpha \vee \text{Un} \rightarrow \text{Un} \\ \text{decrypt} &: \forall \alpha. \text{Un} \rightarrow \text{PrivKey}\langle \alpha \rangle \rightarrow \alpha \vee \text{Un} \end{aligned}$$

Like many of the functions in our cryptographic library, the *encrypt* and *decrypt* functions are polymorphic. Their code is type-checked only once and given an universal type. The type variable α stands in this case for the type of the payload that is encrypted, and can be instantiated with an arbitrary type when the functions are used.

Type Un describes those values that may be known to the opponent, i.e., data that may come from or be sent to the opponent. The type $\text{PubKey}\langle \alpha \rangle$ describes public keys. Since the opponent has access to the public key and to the encryption function, the type system has to take into account that the library may be used by honest principals to encrypt data of type α or by the opponent to encrypt data of type Un . The *encrypt* function takes as input a public key of type $\text{PubKey}\langle \alpha \rangle$ a message of type $\alpha \vee \text{Un}$, and returns a ciphertext of type Un . The *decrypt* function takes as input a ciphertext of type Un , a private key of type $\text{PrivKey}\langle \alpha \rangle$ and returns a payload of type $\alpha \vee \text{Un}$. Without union types, the type of the payload is constrained to be Un or a supertype thereof [12], which severely limits the expressiveness of the type system and prevents the analysis of a number of protocols, including this very simple example.

2.3 Type-Checking the NSL Protocol

We first introduce the type definitions² for the content of the three ciphertexts:

$$\begin{aligned} \text{msg1} &= (\text{Un} * \text{Private}) \\ \text{msg2}[x_B] &= (x_A : \text{Un} * x_{nB} : \text{Private} \vee \text{Un} * \{x_{nA} : \text{Private} \mid \text{auth}_r(x_A, x_B, x_{nB}, x_{nA})\}) \\ \text{msg3} &= \{x_{nA} : \text{Private} \mid \exists x_A, x_B, x_{nB}. \\ &\quad \text{auth}_r(x_A, x_B, x_{nB}, x_{nA}) \wedge \text{auth}_i(x_B, x_A, x_{nB}, x_{nA})\} \end{aligned}$$

The first ciphertext contains a pair composed of a public identifier of type Un and a nonce of type Private . Type Private describes values that are not known to the attacker: the set of values of type Un is disjoint from the set of values of type Private . Type $\text{msg2}[x_A]$ is a combination of two dependent pair types and one refinement type. This type describes a triple composed of an identifier x_A of type Un , a first nonce x_{nB} of type $\text{Private} \vee \text{Un}$, and a second nonce x_{nA} of type Private such that the predicate $\text{auth}_r(x_A, x_B, x_{nB}, x_{nA})$ is entailed by the assumptions in the system (A assumes

² Type definitions are syntactic sugar, and are inlined by the type-checker.

Table 1. NSL Initiator Code and Responder Code

<pre> init = $\lambda x_B : \text{Un}. \lambda x_A : \text{Un}. \lambda k_B : \text{PrivKey}\langle \text{payload}[x_B] \rangle. \lambda pk_A : \text{PubKey}\langle \text{payload}[x_A] \rangle. \lambda ch : \text{Ch}(\text{Un}). \text{let } n_B = \text{mkPriv}() \text{ in} \text{let } p_1 = (\text{Msg1 } (x_B, n_B)) \text{ in} \text{let } m_1 = \text{encrypt}\langle \text{payload}[x_A] \rangle pk_A p_1 \text{ in} \text{send}\langle \text{Un} \rangle ch m_1; \text{let } z = \text{recv}\langle \text{Un} \rangle ch \text{ in} \text{let } x = \text{decrypt}\langle \text{payload}[x_B] \rangle k_B z \text{ in} \text{case } x_1 = x : \text{payload}[x_B] \vee \text{Un} \text{ in} \text{match } x_1 \text{ with } \text{Msg2 } x_2 \Rightarrow \text{let } (y_A, y_{nB}, y_{nA}) = x_2 \text{ in} \text{if } y_A = x_A \text{ then} \text{if } y_{nB} = n_B \text{ then} \text{assert } \text{auth}_r(x_A, x_B, y_{nB}, y_{nA}); \text{assume } \text{auth}_i(x_B, x_A, y_{nB}, y_{nA}); \text{let } p_3 = (\text{Msg3 } y_{nA}) \text{ in} \text{let } m_3 = \text{encrypt}\langle \text{payload}[x_A] \rangle pk_A p_3 \text{ in} \text{send}\langle \text{Un} \rangle ch m_3$</pre>	<pre> resp = $\lambda x_A : \text{Un}. \lambda x_B : \text{Un}. \lambda pk_B : \text{PubKey}\langle \text{payload}[x_B] \rangle. \lambda k_A : \text{PrivKey}\langle \text{payload}[x_A] \rangle. \lambda ch : \text{Ch}(\text{Un}). \text{let } m_1 = \text{recv}\langle \text{Un} \rangle ch \text{ in} \text{let } x_1 \text{ in } \text{decrypt}\langle \text{payload}[x_A] \rangle k_A m_1 \text{case } y_1 = x_1 : \text{payload}[x_A] \vee \text{Un} \text{ in} \text{match } y_1 \text{ with } \text{Msg1 } z_1 \Rightarrow \text{let } (y_B, x_{nB}) = z_1 \text{ in} \text{if } y_B = x_B \text{ then} \text{let } n_A = \text{mkPriv}() \text{ in} \text{assume } \text{auth}_r(x_A, x_B, x_{nB}, n_A); \text{let } p_2 = \text{Msg2}(x_A, x_{nB}, n_A) \text{ in} \text{let } m_2 = \text{encrypt}\langle \text{payload}[x_B] \rangle pk_B p_2 \text{ in} \text{send}\langle \text{Un} \rangle ch m_2; \text{let } m_3 = \text{recv}\langle \text{Un} \rangle ch \text{ in} \text{let } x_3 = \text{decrypt}\langle \text{payload}[x_A] \rangle k_A m_3 \text{ in} \text{case } y_3 = x_3 : \text{payload}[x_A] \vee \text{Un} \text{ in} \text{match } y_3 \text{ with } \text{Msg3 } y_{nA} \Rightarrow \text{if } y_{nA} = n_A \text{ then} \text{assert } \text{auth}_i(x_B, x_A, x_{nB}, n_A)$</pre>
---	---

$\text{auth}_r(A, B, n_B, n_A)$ before creating the second ciphertext). The free occurrence of x_B is bound in the type definition. Notice that x_{nB} is given type $\text{Private} \vee \text{Un}$ since A does not know whether the nonce received in the first ciphertext comes from B or from the opponent. Type msg3 is a refinement type describing a nonce x_{nA} of type Private such that the formula $\exists x_A, x_B, x_{nB}. \text{auth}_r(x_A, x_B, x_{nB}, x_{nA}) \wedge \text{auth}_i(x_B, x_A, x_{nB}, x_{nA})$ is entailed by the assumptions in the system. Indeed, before creating the third ciphertext, B has asserted $\text{auth}_r(A, B, n_B, n_A)$ and assumed $\text{auth}_i(B, A, n_B, n_A)$. Since the payload of the third message only contains x_{nA} we existentially quantify the other variables. The overall type of the payload is obtained by combining the three previous types:

$$\text{payload}[x] = \text{Msg1 of msg1} \mid \text{Msg2 of msg2}[x] \mid \text{Msg3 of msg3}$$

The type of A 's public key is defined as $\text{PubKey}\langle \text{payload}[A] \rangle$ and the type of B 's public key is defined as $\text{PubKey}\langle \text{payload}[B] \rangle$.

The code of the initiator (B in our diagram) and the code of the responder (A) abstract over the principal's identity and they are type-checked independently of each other.

Since library functions such as *encrypt*, *decrypt*, *send* and so on are polymorphic, they are instantiated with a concrete types in the code (e.g., the encryptions in the initiator's code are instantiated with type $\text{payload}[x_A]$ since they take as argument x_A 's public key). The initiator creates a fresh private nonce by means of the function *mkPriv*. The nonce is encrypted together with B 's identifier and sent on the network. The message x obtained by decrypting the second ciphertext is given type $\text{payload}[x_B] \vee \text{Un}$, which reflects the fact that B does not know whether the first ciphertext comes from

A or from the attacker. Since we cannot statically predict which of the two types is the right one, we have to type-check the continuation code twice, once under the assumption that x has type $\text{payload}[x_B]$ and once assuming that x has type Un . This is realized by the expression $\text{case } x_1 = x : \text{payload}[x_B] \vee \text{Un}$ in \dots

If x has type $\text{payload}[x_B]$, then its components are given strong types: y_A is given type Un , y_{nB} is given type $\text{Private} \vee \text{Un}$, and y_{nA} is given the refinement type $\{y_{nA} : \text{Private} \mid \text{auth}_r(x_A, x_B, y_{nB}, y_{nA})\}$. This refinement type ensures that $\text{auth}_r(x_A, x_B, y_{nB}, y_{nA})$ will be entailed at run-time by the assumptions in the system and thus justifies the assertion $\text{assert } \text{auth}_r(x_A, x_B, y_{nB}, y_{nA})$. Finally, the assumption $\text{assume } \text{auth}_i(x_B, x_A, y_{nB}, y_{nA})$ allows us to give y_{nA} type msg3 and thus to type-check the final encryption.

If x has type Un then y_A , y_{nB} , and y_{nA} are also given type Un . The following equality check between the value y_{nB} of type Un and the nonce n_B of type Private makes type-checking the remaining code superfluous: since the set of values of type Un is disjoint from the set of values of type Private , it cannot be that the equality test succeeds. So type-checking the initiator's code succeeds.

Type-checking the responder's code is similar. The code contains two case expressions to deal with the union types introduced by the two decryptions. In particular, the code after the second decryption has to be type-checked under the assumption that the variable y_{nA} has type msg3 and under the assumption that y_{nA} has type Un .

In the former case, the assertion $\text{assert } \text{auth}_i(x_B, x_A, x_{nB}, n_A)$ is justified by the previously assumed formula $\text{auth}_r(x_A, x_B, x_{nB}, n_A)$, the formula in the above refinement type, and the following global assumption, stating that there cannot be two different assumptions $\text{auth}_r(x_A, x_B, x'_{nB}, x'_{nA})$ and $\text{auth}_r(x'_A, x'_B, x'_{nB}, x'_{nA})$ with the same nonce x_{nB} .

$$\begin{aligned} \text{assume } & \forall x_A, x_B, x'_A, x'_B, x_{nA}, x'_{nA}, x_{nB}. \\ & \text{auth}_r(x_A, x_B, x_{nB}, x_{nA}) \wedge \text{auth}_r(x'_A, x'_B, x_{nB}, x'_{nA}) \\ & \Rightarrow x_A = x'_A \wedge x_B = x'_B \wedge x_{nA} = x'_{nA} \end{aligned}$$

This assumption is justified by the fact that the predicate auth_r is assumed only in the responder's code, immediately after the creation of a fresh nonce x_{nB} .

If y_{nA} is given type Un then type-checking the following code succeeds because the equality check between y_{nA} and the value n_A of type Private cannot succeed.

The functions init and resp take private keys as input, so they are not available to the attacker. We provide two public functions that capture the capabilities of the attacker.

Attacker's Interface for NSL

```
createPrincipal =  $\lambda x : \text{Un}$ .
  let  $k = \text{mkPrivKey}(\text{payload}[x]) ()$  in  $\text{addToDB } x \ k$ ;  $\text{getPubKey}(\text{payload}[x]) \ k$ 
startNSL =  $\lambda(\text{role} : \text{Un})(x_A : \text{Un})(x_B : \text{Un})(c : \text{Un})$ .
  let  $k_A = \text{getFromDB } x_A$  in let  $pk_A = \text{getPubKey}(\text{payload}[x_A]) \ k_A$  in
  let  $k_B = \text{getFromDB } x_B$  in let  $pk_B = \text{getPubKey}(\text{payload}[x_B]) \ k_B$  in
  match  $\text{role}$  with inl  $\_ \Rightarrow (\text{init } x_A \ x_B \ k_A \ pk_B \ c)$ 
  | inr  $\_ \Rightarrow (\text{resp } x_B \ x_A \ pk_A \ k_B \ c)$ 
```

We allow the attacker to create arbitrarily many new principals using the createPrincipal function. This generates a new encryption key-pair, stores it in a private

database, and then returns the corresponding public key to the attacker. The second function, `startNSL`, allows the attacker to start an arbitrary number of sessions of the protocol, between principals of his choice. When calling `startNSL`, the attacker chooses whether he wants to start an initiator or a responder, the principals to be involved in the session, and the channel on which the communication occurs. One principal can be involved in many sessions simultaneously, in which it may play different roles.

The two functions above express the capabilities of the attacker for verification purposes, and would not be exposed in a production setting. However, they can also be useful for testing and debugging the code of the protocol: for instance we can execute a protocol run using the following code.

Test Setup for NSL

```
createPrincipal "Alice"; createPrincipal "Bob";
let c = mkChan(Un) () in
(startNSL (inl ()) "Alice" "Bob" c) † (startNSL (inr ()) "Alice" "Bob" c)
```

Since the code of the NSL protocol is well-typed, the soundness result of the type system ensures that in all program runs the assertions are entailed by the assumptions, i.e., the code is safe when executed by an arbitrary attacker. In addition, the two nonces are given type `Private` and thus they are not revealed to the opponent.

3 The $\text{RCF}_{\wedge\vee}^{\forall}$ Calculus

The Refined Concurrent FPC (RCF) [12] is a simple programming language extending the Fixed Point Calculus with refinement types and concurrency [4]. This core calculus is expressive enough to encode a considerable fragment of an ML-like programming language [12]. In this paper, we further increase the expressivity of the calculus by adding intersection types [40], union types [39], and parametric polymorphism. We call the extended calculus $\text{RCF}_{\wedge\vee}^{\forall}$ and describe it in this and the following section.

We start by presenting the surface syntax of $\text{RCF}_{\wedge\vee}^{\forall}$, which is a subset of the syntax supported by our type-checker. In the surface syntax of $\text{RCF}_{\wedge\vee}^{\forall}$ variables are named, which makes programs human-readable. The surface syntax also contains explicit typing annotations that guide type-checking. It is given semantics by translation (i.e., type erasure) into a core implicitly-typed calculus, which we have formalized in Coq [7]. The syntax comprises the four mutually-inductively-defined sets of values, types, expressions, and formulas. We mark with star (*) the constructs that are completely new with respect to RCF [12].

Surface syntax of $\text{RCF}_{\wedge\vee}^{\forall}$ values

x, y, z	variable
$h ::= \text{inl} \mid \text{inr}$	constructor for sum types
$M, N ::=$	value
x	variable
$()$	unit
$\lambda x : T. A$	function (scope of x is A)
(M, N)	pair

$h M$	value of sum type
$\text{fold}_{\mu\alpha. T} M$	recursive value
$\Lambda\alpha. A$	type abstraction* (scope of α is A)
for $\tilde{\alpha}$ in $\tilde{T}; \tilde{U}. M$	value of intersection type* (scope of $\tilde{\alpha} = \alpha_1, \dots, \alpha_n$ is M)

The set of *values* is composed of variables, the unit value, functions, pairs, and introduction forms for disjoint union, recursive, polymorphic, and intersection types.

Surface syntax of $\text{RCF}_{\wedge\vee}^{\forall}$ types

α, β	type variable
$T, U, V ::=$	type
unit	unit type
$x : T \rightarrow U$	dependent function type (x bound in U)
$x : T * U$	dependent pair type (x bound in U)
$T + U$	disjoint sum type
$\mu\alpha. T$	iso-recursive type (α bound in T)
α	type variable
$\{x : T \mid C\}$	refinement type (x bound in C)
$T \wedge U$	intersection type*
$T \vee U$	union type*
\top	top type*
$\forall\alpha. T$	polymorphic type* (α bound in T)

The unit value $()$ is given type unit. 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. If U does not depend on x , then we use the abbreviations $T \rightarrow U$ and $T * U$. The sum type $T + U$ describes values $\text{inl}(M)$ where M is of type T and values $\text{inr}(N)$ where N is of type U . The iso-recursive type $\mu\alpha. T$ is the type of all values $\text{fold}_{\mu\alpha. T} M$ where M is of type $T\{\mu\alpha. T/\alpha\}$. We use refinement types [12] 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. 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. The top type \top is supertype of all the other types, and contains all well-typed values. The universal type $\forall\alpha. T$ describes polymorphic values $\Lambda\alpha. A$ such that $A\{U/\alpha\}$ is of type $T\{U/\alpha\}$ for all types U .

Surface syntax of $\text{RCF}_{\wedge\vee}^{\forall}$ expressions

a, b	name
$A, B ::=$	expression
M	value
$M N$	function application
$M\langle T \rangle$	type instantiation*
let $x = A$ in B	let (scope of x is B)
let $(x, y) = M$ in A	pair split (scope of x, y is A)
match M with $\text{inl } x \Rightarrow A \mid \text{inr } y \Rightarrow B$	pattern matching (scope of x is A , of y is B)

$\text{unfold}_{\mu\alpha. T} M$	use recursive value
$\text{case } x = M : T \vee U \text{ in } A$	elimination of union types* (scope of x is A)
$\text{if } M = N \text{ as } x \text{ then } A \text{ else } B$	equality check with type cast* (scope of x is A)
$(\nu a \uparrow T)A$	restriction (scope of a is A)
$A \uparrow B$	fork off parallel expression
$a!M$	send M on channel a
$a?$	receive on channel a
$\text{assume } C$	add formula C to global log
$\text{assert } C$	formula C must hold

The syntax of expressions is mostly standard [12, 39]. A type instantiation $M\langle T \rangle$ specializes a polymorphic value M with the concrete type T . The elimination form for union types $\text{case } x = M : T \vee U \text{ in } A$ substitutes the value M in A . The conditional $\text{if } M = N \text{ as } x \text{ then } A \text{ else } B$ checks if M is syntactically equal to N , if this is the case it substitutes x with the common value. Syntactic equality is defined up to alpha-renaming of binders and the erasure of typing annotations and constructs such as for . During type-checking the variable x is given the intersection of the types of M and N . When the variable x is not necessary we omit the as clause, as we did in §2. The restriction $(\nu a \uparrow T)A$ 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 returns the unit value $()$. Expression $a?$ blocks until some message M is available on channel a , removes M from the channel, and then returns M . Expression $\text{assume } C$ adds the logical formula C to a global log. The assertion $\text{assert } C$ returns $()$ when triggered. If at this point C is entailed by the multiset S of formulas in the global log, written as $S \models C$, we say the assertion *succeeds*; otherwise, we say the assertion *fails*.

Intuitively, an expression A is *safe* if, once it is translated into $\text{Formal-RCF}_{\wedge\vee}^{\forall}$, all assertions succeed in all evaluations. 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* and is statically enforced by our type system from §4.

We consider a variant of first-order logic with equality as the authorization logic. We assume that $\text{RCF}_{\wedge\vee}^{\forall}$ values are the terms of this logic, and equality $M = N$ is interpreted as syntactic equality between values.

4 Type System

This section presents our type system for enforcing authorization policies on $\text{RCF}_{\wedge\vee}^{\forall}$ code. This extends the type system proposed by Bengtson et al. [12] with union, intersection, and polymorphic types. Additionally, we encode a new type *Private*, which is used to characterize data that are not known to the attacker, and introduce a novel relation for statically reasoning about the disjointness of types. In the following we explain the typing judgements and present the most important typing rules.

4.1 Typing Environment and Entailment

A typing environment E is a list of bindings for variables ($x : T$), type variables (α or $\alpha :: k$), names ($a \downarrow T$, where the name a stands for a channel conveying values of type T), and formulas (bindings of the form $\{C\}$). An environment is well-formed ($E \vdash \diamond$) if all variables, names, and type variables are defined before use, and no duplicate definitions exist. A type T is well-formed in environment E (written $E \vdash T$) if all its free variables, names, and type variables are defined in E .

A crucial judgment in the type system is $E \vdash C$, which states that the formula C is derivable from the formulas in 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 judgment is used for instance when type-checking assert C using (Exp Assert): type-checking succeeds only if C is entailed in the current typing environment.

4.2 Subtyping and Kinding

Intuitively, 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 $\{x : \text{Un} \mid Ok(x)\}$ 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*. We outline some important rules for kinding and subtyping (let k range over pub and tnt).

Kinding and subtyping for refinement types

(Kind Refine Pub)	(Kind Refine Tnt)
$\frac{E \vdash \{x : T \mid C\} \quad E \vdash T :: \text{pub}}{E \vdash \{x : T \mid C'\} :: \text{pub}}$	$\frac{E \vdash T :: \text{tnt} \quad E, x : T \vdash C}{E \vdash \{x : T \mid C\} :: \text{tnt}}$
(Sub Refine Left)	(Sub Refine Right)
$\frac{E \vdash \{x : T \mid C\} \quad E \vdash T <: T'}{E \vdash \{x : T \mid C\} <: T'}$	$\frac{E \vdash T <: T' \quad E, x : T \vdash C}{E \vdash T <: \{x : T' \mid C\}}$

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 $\{x : \text{Un} \mid Ok(x)\}$ 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, so by subtyping we can only add universally valid formulas.

Kinding for pair and function types

(Kind Pair)	(Kind Fun)
$\frac{E \vdash T :: k \quad E, x : T \vdash U :: k}{E \vdash (x : T * U) :: k}$	$\frac{E \vdash T :: \bar{k} \quad E, x : T \vdash U :: k}{E \vdash (x : T \rightarrow U) :: k}$

A pair type $T * U$ is public (or tainted) only if both T and U are public (respectively tainted). On the other hand, a function type $T \rightarrow U$ is public only if the return type U is public (otherwise $\lambda x:\text{unit}. M_{\text{secret}}$ would be public) and the argument type T is tainted (otherwise $\lambda k : \text{PrivKey}(\text{Private}). \text{let } x = \text{encrypt}(\text{Private}) k M_{\text{secret}}$ in $a_{\text{pub}}!x$ would be public).

Kinding and subtyping for union and intersection types (*)

(Kind And Pub 1) $\frac{E \vdash T_1 :: \text{pub} \quad E \vdash T_2}{E \vdash T_1 \wedge T_2 :: \text{pub}}$	(Kind And Pub 2) $\frac{E \vdash T_1 \quad E \vdash T_2 :: \text{pub}}{E \vdash T_1 \wedge T_2 :: \text{pub}}$	(Kind And Tnt) $\frac{E \vdash T_1 :: \text{tnt} \quad \Gamma \vdash T_2 :: \text{tnt}}{\Gamma \vdash T_1 \wedge T_2 :: \text{tnt}}$
(Kind Or Pub) $\frac{E \vdash T_1 :: \text{pub} \quad E \vdash T_2 :: \text{pub}}{E \vdash T_1 \vee T_2 :: \text{pub}}$	(Kind Or Tnt 1) $\frac{E \vdash T_1 :: \text{tnt} \quad E \vdash T_2}{E \vdash T_1 \vee T_2 :: \text{tnt}}$	(Kind Or Tnt 2) $\frac{E \vdash T_1 \quad E \vdash T_2 :: \text{tnt}}{E \vdash T_1 \vee T_2 :: \text{tnt}}$
(Sub And LB 1) $\frac{E \vdash T_1 <: U \quad E \vdash T_2}{E \vdash T_1 \wedge T_2 <: U}$	(Sub And LB 2) $\frac{E \vdash T_1 \quad E \vdash T_2 <: U}{E \vdash T_1 \wedge T_2 <: U}$	(Sub And Greatest) $\frac{E \vdash T' <: T_1 \quad E \vdash T' <: T_2}{E \vdash T' <: T_1 \wedge T_2}$
(Sub Or Least) $\frac{E \vdash T_1 <: U \quad E \vdash T_2 <: U}{E \vdash T_1 \vee T_2 <: U}$	(Sub Or UB 1) $\frac{E \vdash T <: U_1 \quad E \vdash U_2}{E \vdash T <: U_1 \vee U_2}$	(Sub Or UB 2) $\frac{E \vdash U_1 \quad E \vdash T <: U_2}{E \vdash T <: U_1 \vee U_2}$

The intersection type $T_1 \wedge T_2$ can intuitively be seen as a³ greatest lower bound of the types T_1 and T_2 . Rules (Sub And LB 1) and (Sub And LB 2) ensure that $T_1 \wedge T_2$ is a lower bound: by using reflexivity in the premise we obtain that $T_1 \wedge T_2 <: T_1$ and $T_1 \wedge T_2 <: T_2$. Rule (Sub And Greatest) ensures that $T_1 \wedge T_2$ is greater than any other lower bound: if T' is another lower bound of T_1 and T_2 then T' is a subtype of $T_1 \wedge T_2$. As far as kinding is concerned, the type $T_1 \wedge T_2$ is public if T_1 is public or T_2 is public, and it is tainted if both T_1 and T_2 are tainted.

The union type $T_1 \vee T_2$ intuitively corresponds to a least upper bound of T_1 and T_2 . The rules for union types are exactly the dual of the ones for intersection types.

Our type system has no distributivity rules between union and intersection types and the primitive type constructors. Some distributivity rules are derivable from the primitive rules above: for instance we can prove that $T \rightarrow (U_1 \wedge U_2)$ is a subtype of $(T \rightarrow U_1) \wedge (T \rightarrow U_2)$, but not the other way around. In fact adding a subtyping rule in the other direction would be unsound [24], since in our system functions can have side-effects and such distributivity rules would allow circumventing the value restriction on the introduction of intersection types (see §4.4 and §8).

Kinding and subtyping rules for universal types

(Kind Univ*) $\frac{E, \alpha \vdash T :: k}{E \vdash \forall \alpha. T :: k}$	(Sub Univ*) $\frac{E, \alpha \vdash T <: U}{E \vdash \forall \alpha. T <: \forall \alpha. U}$
---	--

³ The subtyping relation of RCF is not anti-symmetric, so least and greatest elements are not necessarily unique.

Finally, the rule for subtyping polymorphic types (Sub Univ*) is simple: the type $\forall\alpha. T$ is subtype of $\forall\alpha. U$ if T is a subtype of U . Similarly, $\forall\alpha. T$ has kind k if T has kind k in an environment extended with a binding for α . Note that α can be substituted by any type, so we cannot assume anything about α when checking that $T :: k$ and $T <: U$ respectively.

Kinding and subtyping rules for recursive types

(Kind Rec) $\frac{E, \alpha :: k \vdash T :: k}{E \vdash (\mu\alpha. T) :: k}$	(Sub Pos Rec*) $\frac{E, \alpha \vdash T <: U \quad \alpha \text{ only occurs positively in } T \text{ and } U}{E \vdash \mu\alpha. T <: \mu\alpha. U}$	(Sub Refl*) $\frac{E \vdash T}{E \vdash T <: T}$
---	--	---

The rule (Sub Pos Rec*) for subtyping recursive types is new, and differs significantly from Cardelli’s Amber rule [5, 20], which is used by the original RCF:

Cardelli’s Amber rule (used by the original RCF)

(Sub Rec) $\frac{E, \alpha <: \alpha' \vdash T <: T' \quad \alpha \neq \alpha' \quad \alpha \notin \text{ftv}(T') \quad \alpha' \notin \text{ftv}(T)}{E \vdash \mu\alpha. T <: \mu\alpha'. T'}$
--

The soundness of the Amber rule (Sub Rec) is hard to prove syntactically [12] – in particular proving the transitivity of subtyping in the presence of the Amber rule requires a complicated inductive argument, which only works for “executable” environments (see [12]), as well as spurious restrictions on the usage of type variables in the rules (Sub Refl*), (Kind And Pub 1), (Kind And Pub 2), (Kind Or Tnt 1), (Kind Or Tnt 2), (Sub And LB 1), (Sub And LB 2), (Sub Or UB 1), (Sub Or UB 2). We use the simpler (Sub Pos Rec*) rule, which is much easier to prove sound and requires no restrictions on the other rules. It resembles (Sub Univ*), our rule for subtyping universal types, with the additional restriction that the recursive variable is not allowed to appear in a contravariant position (such as $\alpha \rightarrow T$). While this positivity restriction is crucial for the soundness of the (Sub Pos Rec*) rule, this does not pose any problem in practice, where most of the time only positive recursive types [38, 44] are used. Moreover, this positivity restriction only affects subtyping, so programs involving negative occurrences of recursion variables that do not involve subtyping can still be properly type-checked (e.g., we can still type-check the encodings of fixpoint combinators on expressions [12]).

4.3 Encoding Types Un and Private in RCF

In RCF [12] the type Un is in fact not primitive. By the (Sub Pub Tnt) rule that relates kinding and subtyping, any type that is both public and tainted is equivalent to Un. Since type unit is both public and tainted, Un is actually encoded as unit.

The (Sub Pub Tnt) rule and kinding for type unit

(Sub Pub Tnt) $\frac{E \vdash T :: \text{pub} \quad E \vdash U :: \text{tnt}}{E \vdash T <: U}$	(Kind Unit) $\frac{E \vdash \diamond}{E \vdash \text{unit} :: k}$
--	--

The (Sub Pub Tnt) rule equates many of the types in the system. For instance in RCF all the following types are equivalent: Un , $\text{Un} \rightarrow \text{Un}$, $\text{Un} * \text{Un}$, $\text{Un} + \text{Un}$, $\mu\alpha. \text{Un}$, and $\forall\alpha. \text{Un}$. As a consequence it is hard to come up with RCF types that do not share any values with type Un , a property we want for our `Private` type. Perhaps unintuitively, it is not enough that a type is not public and not tainted to make it disjoint from Un . A final observation is that, in $\text{RCF}_{\wedge, \vee}^{\forall}$, in an inconsistent environment ($E \vdash \text{false}$) *all* types are equivalent and all values inhabit all types. This means that `Private` being disjoint from Un is relative to the formulas in the environment.

Encoding type `Private`

$$\{C\} \triangleq \{x : \text{unit} \mid C\} \quad x \notin \text{free}(C)$$

$$\text{Private}_C \triangleq \{f : \{C\} \rightarrow \text{Un} \mid \exists x. f = \lambda y : \{C\}. \text{assert } C; x\} \quad \text{Private} \triangleq \text{Private}_{\text{false}}$$

We therefore encode a more general type Private_C , read “private unless C ”. The values in this type are not known to the attacker, unless the formula C is entailed by the environment. Intuitively, if the attacker would know a value of this type, then he could call it (values of type Private_C have to be functions), which would exercise the `assert C` and invalidate the safety of the system, unless C can be derived from the formula log. Type Private_C resembles a singleton type, in that it contains only values of a very specific form. We use an existential quantifier over values to ensure that there are infinitely many values of this type. The type `Private` is obtained as $\text{Private}_{\text{false}}$.

4.4 Typing Values and Expressions

The main judgments of the type system we consider are $E \vdash M : T$, which states that value M has type T , and $E \vdash A : T$, stating that expression A returns a value of type T . These two judgements are mutually-inductively defined, and the most important typing rules are reported below. Most of them are standard, so we focus the explanation only on the rules that are new with respect to [12].

Selected rules for typing values $E \vdash M : T$

(Val Lam)	(Val TLam*)	(Val Refine)
$\frac{E, x : T \vdash A : U}{E \vdash \lambda x : T. A : (x : T \rightarrow U)}$	$\frac{E, \alpha \vdash A : T}{E \vdash \lambda\alpha. A : \forall\alpha. T}$	$\frac{E \vdash M : T \quad E \vdash C\{M/x\}}{E \vdash M : \{x : T \mid C\}}$
(Val And*)	(Val For 1*)	(Val For 2*)
$\frac{E \vdash M : T \quad E \vdash M : U}{E \vdash M : T \wedge U}$	$\frac{E \vdash M\{\tilde{T}/\tilde{\alpha}\} : V}{E \vdash \text{for } \tilde{\alpha} \text{ in } \tilde{T}; \tilde{U}. M : V}$	$\frac{E \vdash M\{\tilde{U}/\tilde{\alpha}\} : V}{E \vdash \text{for } \tilde{\alpha} \text{ in } \tilde{T}; \tilde{U}. M : V}$

Rule (Val And*) allows us to give value M an intersection type $T \wedge U$, if we can give M both type T and type U . As discovered by Davies and Pfenning [24] the value restriction is crucial for the soundness of this introduction rule in the presence of side-effects (also see §8). Also, unrelated to the value restriction, this rule is not very useful on its own: since we are in a calculus with typing annotations, it is hard to give one annotated value two different types. For instance, if we want to give the identity function type $(\text{Private} \rightarrow \text{Private}) \wedge (\text{Un} \rightarrow \text{Un})$ we need to annotate the argument with type Private (i.e., $\lambda x : \text{Private}. x$) in order to give it type $\text{Private} \rightarrow \text{Private}$, but then we cannot give

this value type $\text{Un} \rightarrow \text{Un}$. Following Pierce [39, 40] and Reynolds [41] we use the `construct` to explicitly alternate type annotations. For instance, the identity function of type $(\text{Private} \rightarrow \text{Private}) \wedge (\text{Un} \rightarrow \text{Un})$ can be written as (for α in $\text{Private}; \text{Un}$. $\lambda x:\alpha. x$). By rule (Val For 1*) we can give this value type $\text{Private} \rightarrow \text{Private}$ if we can give value $\lambda x:\text{Private}. x$ the same type, which is trivial. Similarly, by (Val For 2*) we can give the for value type $\text{Un} \rightarrow \text{Un}$, so by (Val And*) we can also give it the desired intersection type.

Selected rules for typing expressions $E \vdash A : T$

(Exp Appl) $\frac{E \vdash M : (x : T \rightarrow U) \quad E \vdash N : T}{E \vdash M N : U\{N/x\}}$	(Exp Inst*) $\frac{E \vdash M : \forall \alpha. U}{E \vdash M\langle T \rangle : U\{T/\alpha\}}$	(Exp Assert) $\frac{E \vdash C}{E \vdash \text{assert } C : \text{unit}}$
---	---	--

(Exp If*)

$$\frac{E \vdash M : T_1 \quad E \vdash N : T_2 \quad \vdash \text{NonDisj } T_1 T_2 \rightsquigarrow C \quad E, x : T_1 \wedge T_2, \{x = M \wedge M = N \wedge C\} \vdash A : U \quad E, \{M \neq N\} \vdash B : U}{E \vdash \text{if } M = N \text{ as } x \text{ then } A \text{ else } B : U}$$

(Exp Case*)

$\frac{E \vdash M : T_1 \vee T_2 \quad E, x : T_1 \vdash A : U \quad E, x : T_2 \vdash A : U}{E \vdash \text{case } x = M : T_1 \vee T_2 \text{ in } A : U}$	(Exp Subsum) $\frac{E \vdash A : T \quad E \vdash T <: T'}{E \vdash A : T'}$
--	---

Union Types are introduced by subtyping (T_1 is a subtype of $T_1 \vee T_2$ for any T_2), and eliminated by a case $x = M : T_1 \vee T_2$ in A expression [39] using the (Exp Case*) rule⁴. Given a value M of type $T_1 \vee T_2$, we do not know whether M is of type T_1 or of type T_2 , so we have to type-check A under each of these assumptions. This is useful when type-checking code interacting with the attacker. For instance, suppose that a party receives a value encrypted with a public-key that is used by honest parties to encrypt messages of type T (as in the protocol from §2). After decryption, the obtained plaintext 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 remaining code twice, once under the assumption that x is of type T , and once assuming that x is of type Un .

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` . 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 . This is useful for type-checking the symbolic implementations of digital signatures (see §5.2) and zero-knowledge (see §6). Additionally, if the equality test of the conditional succeeds then the types T_1 and T_2 are not 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

⁴ As pointed out by Dunfield and Pfenning [28] eliminating union types for expressions that are not in evaluation contexts is unsound in the presence of non-determinism (this is further discussed in §8).

conditional cannot succeed, so the expression A will not be executed. This idea has been applied in [1] for verifying secrecy properties of nonce handshakes, but later disappeared in the more advanced type systems for authorization policies.

Non-disjointness of types (*) $\vdash \text{NonDisj } T \ U \rightsquigarrow C$

(ND Private Un) $\frac{fv(C) = \emptyset}{\vdash \text{NonDisj } \text{Private}_C \ \text{Un} \rightsquigarrow C}$	(ND True) $\frac{}{\vdash \text{NonDisj } T_1 \ T_2 \rightsquigarrow \text{true}}$	(ND Sym) $\frac{\vdash \text{NonDisj } T_2 \ T_1 \rightsquigarrow C}{\vdash \text{NonDisj } T_1 \ T_2 \rightsquigarrow C}$
(ND Refine) $\frac{\vdash \text{NonDisj } T_1 \ T_2 \rightsquigarrow C}{\vdash \text{NonDisj } \{x : T_1 \mid C_1\} \ T_2 \rightsquigarrow C}$	(ND Rec) $\frac{\vdash \text{NonDisj } (T\{\alpha/\mu\alpha.T\}) \ (U\{\beta/\mu\beta.U\}) \rightsquigarrow C}{\vdash \text{NonDisj } (\mu\alpha.T) \ (\mu\beta.U) \rightsquigarrow C}$	
(ND Pair) $\frac{\vdash \text{NonDisj } T_1 \ U_1 \rightsquigarrow C_1 \quad \vdash \text{NonDisj } T_2 \ U_2 \rightsquigarrow C_2}{\vdash \text{NonDisj } (T_1 * T_2) \ (U_1 * U_2) \rightsquigarrow C_1 \wedge C_2}$	(ND Sum) $\frac{\vdash \text{NonDisj } T_1 \ U_1 \rightsquigarrow C_1 \quad \vdash \text{NonDisj } T_2 \ U_2 \rightsquigarrow C_2}{\vdash \text{NonDisj } (T_1 + T_2) \ (U_1 + U_2) \rightsquigarrow (C_1 \vee C_2)}$	
(ND And) $\frac{\vdash \text{NonDisj } T_1 \ U \rightsquigarrow C_1 \quad \vdash \text{NonDisj } T_2 \ U \rightsquigarrow C_2}{\vdash \text{NonDisj } (T_1 \wedge T_2) \ U \rightsquigarrow C_1 \wedge C_2}$	(ND Or) $\frac{\vdash \text{NonDisj } T_1 \ U \rightsquigarrow C_1 \quad \vdash \text{NonDisj } T_2 \ U \rightsquigarrow C_2}{\vdash \text{NonDisj } (T_1 \vee T_2) \ U \rightsquigarrow C_1 \vee C_2}$	

We take this idea a lot further: we inductively define a ternary relation, which relates two types with a logical formula. If $\vdash \text{NonDisj } T_1 \ T_2 \rightsquigarrow C$ holds then any environment E in which T_1 and T_2 have a common value, has to entail the condition C (i.e., $E \vdash C$). The base case of this relation is $\vdash \text{NonDisj } \text{Private}_C \ \text{Un} \rightsquigarrow C$, in particular $\vdash \text{NonDisj } \text{Private} \ \text{Un} \rightsquigarrow \text{false}$. We call two types *provably disjoint* if $\vdash \text{NonDisj } T_1 \ T_2 \rightsquigarrow C$ for some formula C that logically entails false, so Private and Un are provably disjoint. Intuitively, two provably disjoint types have common values only in an inconsistent environment.

The other inductive rules lift the NonDisj relation to refinement, pair, sum, recursive, union, and intersection types. We explain two of them in terms of provable disjointness. In order to show that two (non-dependent) pair types $(T_1 * T_2)$ and $(U_1 * U_2)$ are provably disjoint, we apply rule (ND Pair) and we need to show that T_1 and U_1 are provably disjoint, or that T_2 and U_2 are provably disjoint (a conjunction is false if at least one of the conjuncts is false). On the other hand, in order to show that two sum types $(T_1 + T_2)$ and $(U_1 + U_2)$ are disjoint using (ND Sum) we need to show both that T_1 and U_1 are disjoint and that T_2 and U_2 are disjoint.

To illustrate the expressivity of this definition we consider a type for binary trees: $\text{tree}(\alpha) \triangleq \mu\beta. \alpha + (\alpha * \beta * \beta)$. Each node in the tree is either a leaf or has two children, and both kind of nodes store some information of type α . We can show that $\text{tree}(\text{Private})$ and $\text{tree}(\text{Un})$ are provably disjoint. By (ND Rec) we need to show that the unfolded types $\text{Private} + (\text{Private} * \text{tree}(\text{Private}) * \text{tree}(\text{Private}))$ and $\text{Un} + (\text{Un} * \text{tree}(\text{Un}) * \text{tree}(\text{Un}))$ are disjoint. By (ND Sum) we need to show both that Private and Un are disjoint, which is immediate by (ND Private Un), and that the pair types $(\text{Private} * \text{tree}(\text{Private}) * \text{tree}(\text{Private}))$ and $(\text{Un} * \text{tree}(\text{Un}) * \text{tree}(\text{Un}))$ are disjoint.

For the latter, by (ND Pair) it suffices to show that the types of the first components of the pair are disjoint, which follows again by (ND Private Un).

We have proved in Coq that our type system enforces robust safety; for details we refer to the long version [7].

5 Implementation of Symbolic Cryptography

In contrast to process calculi for cryptographic protocols [4, 3], $\text{RCF}_{\wedge\vee}^{\forall}$ does not have any built-in construct to model cryptography. Cryptographic primitives are instead encoded using a dynamic sealing mechanism [35], which is based on standard $\text{RCF}_{\wedge\vee}^{\forall}$ constructs. The resulting symbolic cryptographic libraries are type-checked using the regular typing rules. The main advantage is that, adding a new primitive to the library does not involve changes in the calculus or in the soundness proofs: one has just to find a well-typed encoding of the desired cryptographic primitive. In addition, Backes et al. have recently [11] shown that sealing-based libraries for asymmetric cryptography are computationally sound and semantically equivalent to the more traditional Dolev-Yao libraries based on datatype constructors. §5.1 overviews the dynamic sealing mechanism used in [12] to encode symbolic cryptography, while §5.2 and §5.3 show how our expressive type system can be used to improve this encoding and extend the class of supported protocols.

5.1 Dynamic Sealing

The notion of *dynamic sealing* was initially introduced by Morris [35] as a protection mechanism for programs. Later, Sumii and Pierce [43] studied the semantics of dynamic sealing in a λ -calculus, observing a close correspondence with symmetric encryption.

In RCF [12] seals are encoded using pairs, functions, references and lists. A seal is a pair of a *sealing function* and an *unsealing function*, having type:

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

The sealing function takes as input a value M of type T and returns a fresh value N of type Un , after adding the pair (M, N) to a secret list that is stored in a reference. The unsealing function takes as input a value N of type Un , scans the list in search of a pair (M, N) , 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 N represents the ciphertext.

Let us take a look at the type $\text{Seal } \langle T \rangle$. If T 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 T is tainted but not public, as usually the case for *public-key encryption*, 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 T 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 even if the message to be signed is not (as in DAA, see §6.1). 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 as the NSL protocol (see §2), where the type of the encrypted messages is neither public nor tainted.

5.2 Digital Signatures

In this section, we focus on digital signatures and show how union and intersection types can be used to solve the aforementioned problems. The signing key consists of the seal itself and is given type $\text{SigKey}\langle T \rangle \triangleq \text{Seal}\langle T \rangle$, as in the original RCF library [12]. The verification key, instead, is encoded as a function that (i) takes the signature x and the signed message t as input; (ii) calls the unsealing function to retrieve the message y bound to x in the secret list; and (iii) returns y if y is equal to t and fails otherwise. In this encoding, the verifier has to know the signed message in order to verify the signature. This is reasonable as, for efficiency reasons, one usually signs a hash of the message as opposed to the message in plain.

Symbolic implementation of signing-verification key pair

```

mkSigPair : ∀α. unit → SigKey⟨α⟩ * VerKey⟨α⟩
mkSigPair = Λα. λu : unit.
  let (seal, unseal) = mkSeal⟨α⟩ in
  let vk = λx : Un. for β in T; Un. λt : β.
    if t = (unseal x) as z then z else failwith "verification failed"
  in (k, vk)

```

The type $\text{VerKey}\langle T \rangle$ of a verification key is defined as $\text{Un} \rightarrow ((x : \top \rightarrow \{y : T \mid x = y\}) \wedge (\text{Un} \rightarrow \text{Un}))$. The verification key takes the signature of type Un as first argument. The second part of this type is an intersection of two types: The type $x : \top \rightarrow \{y : T \mid x = y\}$ is used to type-check honest callers: the signed message x has any type (top type) and the message y returned by the unsealing function has the stronger type T , which means that the unsealing function casts the type of the signed message from \top down to T . This is safe since the sealing function is not public and can only be used to sign messages of type T . The type $\text{Un} \rightarrow \text{Un}$ makes $\text{VerKey}\langle T \rangle$ always public⁵. Hence, in contrast to [12], we can reason about protocols where the signing key is used to sign private messages while the verification key is public (e.g., in DAA [18]).

⁵ A type of the form $\text{Un} \rightarrow (T_1 \wedge T_2)$ is public if T_1 or T_2 are public, and in our case $T_2 = \text{Un} \rightarrow \text{Un}$ is public.

Finally, we present the typed interface of the functions to create and check signatures:

$$\begin{aligned} \text{sign} &: \forall \alpha. (x_{sk} : \text{SigKey}(\alpha) \rightarrow \alpha \rightarrow \text{Un}) \wedge \text{Un} \\ \text{check} &: \forall \alpha. (x_{vk} : \text{VerKey}(\alpha) \rightarrow \text{Un} \rightarrow \top \rightarrow \alpha) \wedge \text{Un} \end{aligned}$$

We type-check *sign* and *check* twice, to give them intersection types whose right-hand side is Un . While making these functions available to the adversary is not necessary (the attacker can directly use the signing and verification keys to which he has access), this is convenient for the encoding of zero-knowledge we describe in §6 (dishonest verifier cases).

5.3 Public-Key Encryption

For public-key encryption we simply use a seal of type $\text{Seal}(T \vee \text{Un})$, i.e., $\text{PrivKey}(T) \triangleq \text{Seal}(T \vee \text{Un})$ and $\text{PubKey}(T) \triangleq (T \vee \text{Un}) \rightarrow \text{Un}$. This allows us to obtain the types described in §2.2. In contrast to [12], the encryption key is always public, even if the type T of the encrypted message is not tainted⁶.

6 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. Some of the values used to create a zero-knowledge proof are revealed by the proof to the verifier and to any eavesdropper, while the others (which we call witnesses) are kept secret. A zero-knowledge proof does not reveal any information about these witnesses, other than the validity of the statement being proved.

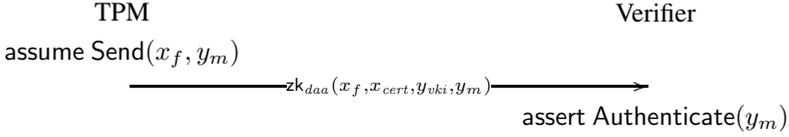
6.1 Illustrative Example: Simplified DAA

We are going to illustrate our technique on a simplified version⁷ of the Direct Anonymous Attestation (DAA) protocol [18]. 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

⁶ A type of the form $(T_1 \vee T_2) \rightarrow \text{Un}$ is public if T_1 or T_2 is tainted, and in our case $T_2 = \text{Un}$ is tainted.

⁷ The long version describes the general code generation routine in more detail [7].

protocol and received the certificate from the issuer. In the DAA-signing protocol the TPM sends to the verifier a zero-knowledge proof.



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. 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. 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

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

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

6.2 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 [10, 8]. For a specification 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.

High-level specification of simplified DAA

```
zkdef daa =
  witness = [x_f : T_vki, x_cert : Un]
  matched = [y_vki : VerKey⟨T_vki⟩]
  public = [y_m : Un]
  statement = [x_f = check⟨T_vki⟩ y_vki x_cert x_f]
  promise = [Send(x_f, y_m)]
where T_vki = {z_f : Private | OkTPM(z_f)}
```

Variables. 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 is already known to the verifier of the zero-knowledge proof. The payload message y_m is returned to the verifier of the proof, so it is public.

Statement. The statement conveyed by a zero-knowledge proof is in general a positive Boolean formula over equality checks. In our simplified DAA example this is just $x_f = \text{check}\langle T_{\text{vki}}\rangle y_{\text{vki}} x_{\text{cert}} x_f$.

Types. The user also needs to provide types for the variables. The DAA-sign protocol does not preserve the secrecy of the signed message, so y_m has type Un . On the other hand, the TPM identifier x_f is given a secret and untainted type $T_{vki} = \{z_f : \text{Private} \mid \text{OkTPM}(z_f)\}$. This type ensures that x_f is not known to the attacker and that the predicate $\text{OkTPM}(x_f)$ holds. The verification key of the issuer is used to check signed messages of type T_{vki} , so it is given type $\text{VerKey}\langle T_{vki} \rangle$. 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.

Promise. The user can additionally specify a *promise*: an arbitrary authorization logic formula that holds in the typing environment of the prover. If the statement is strong enough to identify the prover as an honest (type-checked) protocol participant (signature proofs of knowledge such as DAA-signing have this property), then the promise can be safely transmitted to the typing environment of the verifier. In the DAA example we have the promise $\text{Send}(x_f, y_m)$, since this predicate holds in the typing environment of a honest TPM.

6.3 Automatic Code Generation

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

Generated typed interface for simplified DAA

```

createdaa : Tdaa ∨ Un → Un           publicdaa : Un → Un
verifydaa : Un → ((yvki : VerKey⟨Tvki⟩ → Udaa) ∧ Un → Un)
where Tdaa = yvki : VerKey⟨Tvki⟩ * ym : Un * xf : Tvki * xcert : Un * {Send(xf, ym)}
and Udaa = {ym : Un | ∃xf, xcert. OkTPM(xf) ∧ Send(xf, ym)}

```

The *generated interface* for DAA contains three functions that share a hidden seal of type $T_{daa} \vee \text{Un}$. The function create_{daa} is used to create zero-knowledge proofs. 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 specified types. Additionally, we check that the promise $\text{Send}(x_f, y_m)$ 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.

The function public_{daa} is used to read the public values of a proof, so it takes as input the sealed proof of type Un and returns y_m , also at type Un .

The function verify_{daa} is used for verifying zero-knowledge proofs. Because of the second part of the intersection type, 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 verification key of the issuer with type $\text{VerKey}\langle T_{daa} \rangle$. On successful verification, verify_{daa} returns y_m , the only public variable, but with a stronger type than in public_{daa} . The function guarantees that the formula $\exists x_f, x_{cert}. \text{OkTPM}(x_f) \wedge \text{Send}(x_f, y_m)$ holds, where the witnesses are existentially quantified. The first conjunct, $\text{OkTPM}(x_f)$, guarantees

that if verification succeeds then the statement indeed holds, no matter what the origin of the proof is. This predicate is automatically extracted from the return type of the $\text{check}\langle T_{vki} \rangle$ function (see §5.2). The second conjunct $\text{Send}(x_f, y_m)$ is the promise of the proof.

The *generated implementation* for this interface creates a fresh seal k_{daa} for values of type $T_{daa} \vee \text{Un}$. The sealing function of k_{daa} is directly used to implement the create_{daa} function. The unsealing function of k_{daa} is used to implement the public_{daa} and verify_{daa} functions. The implementation of public_{daa} is very simple: since the zero-knowledge proof is just a sealed value, public_{daa} unseals it and returns y_m . The witnesses are discarded, and the validity of the statement is not checked.

The implementation of the verify_{daa} function is more interesting. This function takes a candidate zero-knowledge proof z of type Un as input, and a value for the matched variable y_{vki} . Since the type of verify_{daa} contains an intersection type we use a for construct to introduce this intersection type. If the proof is verified by the attacker we can assume that the y_{vki} has 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 y_{vki} has the type $\text{VerKey}\langle T_{vki} \rangle$. In general, it is the strong types of the matched values that allow us to guarantee the strong types of the returned public values, as well as the promise.

Generated symbolic implementation for simplified DAA

```

verifydaa = λz : Un.
  for α in Un; VerKey⟨Tvki⟩. λy'vki : α.
    let z' = (snd kdaa) z in (1)
    case z'' = z' : Un ∨ Tdaa in (2)
    let (yvki, ym, xf, xcert, _) = z'' in (3)
    if yvki = y'vki as y''vki then (4)
      if xf = check⟨Tvki⟩ y''vki xcert xf then ym (5)
    else failwith "statement not valid"
    else failwith "yvki does not match"

```

The generated verify_{daa} function performs the following five steps: (1) it unseals z using “ $\text{snd } k_{daa}$ ” 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} and y_m) 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 $\text{check}\langle T_{vki} \rangle$ function, and checking the result for equality with the value of x_f . In general, this last step is slightly complicated by the fact that the statement can contain conjunctions and disjunctions, so we use decision trees. However, for the DAA example the decision tree has a trivial structure with only one node.

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 verify_{daa} 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 DAA the most interesting case is dishonest prover / honest verifier, when z'' and hence x_f are given type Un , while the result of the signature verification is of type T_{vki} . Since $\vdash \text{NonDisj} \{z_f : \text{Private} \mid \text{OkTPM}(z_f)\} \text{Un} \rightsquigarrow \text{false}$ by rules (ND Refine) and (ND Private Un), false is added to the environment in which y_m is type-checked. The variable y_m has type Un in this environment, but since this environment is inconsistent y_m can also be given type U_{daa} .

7 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 §4, the automatic code generator for zero-knowledge described in §6, an interpreter, and a visual debugger.

The type-checker supports an extended syntax with respect to the one from §3, 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 the Z3 SMT solver [25] to discharge proof obligations. The type-checker performed very well in our experiments: it type-checks all our symbolic libraries and samples totaling more than 1.5kLOC in around 12 seconds, on a normal laptop. 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 very 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 – we would like to improve the amount of type inference in the future (see §9 for a discussion).

The type-checker, the code generator for zero-knowledge, and 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 2.5kLOC, while the whole tool-chain has over 5kLOC. All the tools and samples are available at [7].

8 Related Work on Unions and Intersections

The `for` construct for explicitly alternating type annotations was introduced by Pierce [39, 40] as a generalization of an idea Reynolds [41] used in Forsythe for giving intersection types to annotated lambda abstractions of the form $\lambda x:\tau_1.. \tau_n. e$. In a Church-style system, however, the `for` construct does not have a clear operational semantics. Compagnoni [23] gives an operational semantics to function application expressions of the form $((\text{for } \alpha \text{ in } T; U. \lambda x:V. e_1) e_2)$ by pushing the application inside the `for` – i.e., this expression reduces in one step to $(\text{for } \alpha \text{ in } T; U. ((\lambda x:V. e_2) e_2))$. It is unclear if this can be generalized to anything other than function applications. Moreover, this reduction rule does not respect the value restriction for the introduction of intersection

types (our rule (Val And*) in §4). As discovered by Davies and Pfenning [24] the value restriction on intersection introduction is crucial for soundness in the presence of side-effects. The counterexample they give is in fact very similar to the one used to illustrate the unsoundness of ML, in the absence of the value restriction, due to the interaction of polymorphism with side-effects [33]. Moreover, Davies and Pfenning [24] observed that some standard distributivity laws of subtyping are unsound in a setting with side-effects, since they basically allow one to circumvent the value restriction. We obtain all the benefits of the for construct in $\text{RCF}_{\wedge, \vee}^{\forall}$, but erase it completely when translating values into $\text{Formal-RCF}_{\wedge, \vee}^{\forall}$, and use the value restriction on both levels to ensure soundness.

The case construct for eliminating union types was introduced by Pierce [39] as a way to make type-checking more efficient, by asking the programmer to annotate the position in the code where union elimination should occur. Dunfield and Pfenning [28] later pointed out that unrestricted elimination of union types is unsound in the presence of non-determinism. This observation is crucial for us, since our calculus, as opposed to the one studied by Dunfield and Pfenning, is in fact non-deterministic. They propose an evaluation context restriction that recovers soundness, but this is not enough to make type-checking efficient. In recent work, Dunfield [27], shows that carefully transforming programs into let-normal form improves efficiency. This is encouraging, since our expressions are already in let normal form, so we can hope to replace the case construct by a normal let in the future, and still preserve efficient type-checking.

9 Conclusions and Future Work

We have presented a new type system that combines refinement types with union types, intersection types, and polymorphic types. A novelty of the type system is its ability to reason statically about the disjointness of types. This extends the scope of the existing type-based analyses of protocol implementations to important classes of cryptographic protocols that were not covered so far, including protocols based on zero-knowledge proofs. Our type system comes with a mechanized proof of correctness and an efficient implementation [7].

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. Also, we intend to apply our framework to analyze implementations of more complex protocols, such as the Civitas electronic voting system [22].

The type-checker we implemented had very good efficiency in our experiments, however, the amount of typing annotations it requires is at the moment quite high. This issue is more pronounced in our symbolic cryptography library, where intersection and union types are pervasive. This is less of a problem in the code that links against these libraries, and in the case of zero-knowledge even the code in the library is automatically generated together with all the necessary annotations. In the future we would like to perform more type inference, maybe leveraging some of the recent progress on type inference for refinement types [42, 34]. The good news is that intersection and union types can be very useful when devising precise type inference algorithms [8, 36].

Acknowledgments. We thank Cédric Fournet, Andy Gordon, Jan Schwinghammer, and Pierre-yves Strub for the constructive discussions. Thorsten Tarrach implemented the original F5 prototype. Stefan Lorenz helped us with the cryptographic implementation of the DAA protocol. Joshua Dunfield and Kim Pecina commented on a draft. Cătălin Hrițcu is supported by a fellowship from Microsoft Research and the International Max Planck Research School for Computer Science. Matteo Maffei is partially supported by the initiative for excellence of the German federal government, by DFG Emmy Noether program, and by MIUR project “SOFT”.

References

1. Abadi, M., Blanchet, B.: Secrecy types for asymmetric communication. *Theoretical Computer Science* 3(298), 387–415 (2003)
2. Abadi, M., Blanchet, B.: Analyzing security protocols with secrecy types and logic programs. *Journal of the ACM* 52(1), 102–146 (2005)
3. Abadi, M., Fournet, C.: Mobile values, new names, and secure communication. In: *Proc. 28th Symposium on Principles of Programming Languages (POPL)*, pp. 104–115. ACM Press, New York (2001)
4. Abadi, M., Gordon, A.D.: A calculus for cryptographic protocols: The spi calculus. *Information and Computation* 148(1), 1–70 (1999)
5. Amadio, R.M., Cardelli, L.: Subtyping recursive types. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 15(4), 575–631 (1993)
6. Backes, M., Grochulla, M.P., Hrițcu, C., Maffei, M.: Achieving security despite compromise using zero-knowledge. In: *22th IEEE Symposium on Computer Security Foundations (CSF 2009)*. IEEE Computer Society Press, Los Alamitos (July 2009)
7. Backes, M., Hrițcu, C., Maffei, M.: Union and intersection types for secure protocol implementations. Long version, formalization and implementation, <http://www.infsec.cs.uni-sb.de/projects/F5/>
8. Backes, M., Hrițcu, C., Maffei, M.: Type-checking zero-knowledge. In: *15th ACM Conference on Computer and Communications Security (CCS 2008)*, pp. 357–370. ACM Press, New York (2008)
9. Backes, M., Maffei, M., Pecina, K.: A security API for distributed social networks. In: *18th Annual Network & Distributed System Security Symposium (NDSS 2011)*, pp. 35–51. Internet Society, San Diego (2011)
10. Backes, M., Maffei, M., Unruh, D.: Zero-knowledge in the applied pi-calculus and automated verification of the direct anonymous attestation protocol. In: *Proc. of 29th IEEE Symposium on Security and Privacy*, pp. 202–215. IEEE Computer Society Press, Los Alamitos (2008)
11. Backes, M., Maffei, M., Unruh, D.: Computationally sound verification of source code. In: *Proc. 17th ACM Conference on Computer and Communications Security (CCS)*, pp. 387–398. ACM Press, New York (2010)
12. Bengtson, J., Bhargavan, K., Fournet, C., Gordon, A.D., Maffei, S.: Refinement types for secure implementations. In: *Proc. 21th IEEE Symposium on Computer Security Foundations (CSF)*, pp. 17–32. IEEE Computer Society Press, Los Alamitos (2008), long version appeared as MSR-TR-2008-118. November 2010 revision that fixes the problems we pointed out is <http://research.microsoft.com/en-us/um/people/adg/Publications/MSR-TR-2008-118-SP2.pdf>
13. Bhargavan, K., Corin, R., Fournet, C., Zălinescu, E.: Cryptographically verified implementations for TLS. In: *15th ACM Conference on Computer and Communications Security (CCS 2008)*, pp. 459–468. ACM Press, New York (2008)

14. Bhargavan, K., Fournet, C., Gordon, A.D.: Modular verification of security protocol code by typing. In: Proc. 37th Symposium on Principles of Programming Languages (POPL 2010), pp. 445–456 (2010)
15. Bhargavan, K., Fournet, C., Gordon, A.D., Tse, S.: Verified interoperable implementations of security protocols. In: Proc. 19th IEEE Computer Security Foundations Workshop (CSFW), pp. 139–152. IEEE Computer Society Press, Los Alamitos (2006)
16. Blanchet, B.: An efficient cryptographic protocol verifier based on Prolog rules. In: Proc. 14th IEEE Computer Security Foundations Workshop (CSFW), pp. 82–96. IEEE Computer Society Press, Los Alamitos (2001)
17. Bleichenbacher, D.: Chosen ciphertext attacks against protocols based on the RSA encryption standard PKCS. In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, pp. 1–12. Springer, Heidelberg (1998)
18. Brickell, E., Camenisch, J., Chen, L.: Direct anonymous attestation. In: Proc. 11th ACM Conference on Computer and Communications Security, pp. 132–145. ACM Press, New York (2004)
19. Butler, F., Cervesato, I., Jaggar, A.D., Sedrov, A., Walstad, C.: Formal analysis of Kerberos 5. *Theoretical Computer Science* 367(1), 57–87 (2006)
20. Cardelli, L.: Type systems. In: *The Computer Science and Engineering Handbook*, pp. 2208–2236 (1997)
21. Chaki, S., Datta, A.: ASPIER: An automated framework for verifying security protocol implementations. Technical report, CMU CyLab (October 2008)
22. Clarkson, M.R., Chong, S., Myers, A.C.: Civitas: A secure voting system. In: Proc. 29th IEEE Symposium on Security and Privacy, pp. 354–368. IEEE Computer Society Press, Los Alamitos (2008)
23. Compagnoni, A.B.: Subject reduction and minimal types for higher order subtyping. Technical Report ECS-LFCS-97-363, LFCS, University of Edinburgh (August 1997)
24. Davies, R., Pfenning, F.: Intersection types and computational effects. In: Proc. International Conference on Functional Programming (ICFP 2000), pp. 198–208 (2000)
25. de Moura, L., Bjørner, N.: Z3: An efficient SMT solver. In: Ramakrishnan, C.R., Rehof, J. (eds.) TACAS 2008. LNCS, vol. 4963, pp. 337–340. Springer, Heidelberg (2008)
26. Denning, D.E., Sacco, G.M.: Timestamps in key distribution protocols. *Communications of the ACM* 24(8), 533–536 (1981)
27. Dunfield, J.: Untangling typechecking of intersections and unions. In: Workshop on Intersection Types and Related Systems (ITRS) (July 2010)
28. Dunfield, J., Pfenning, F.: Tridirectional typechecking. In: Proc. 31st Symposium on Principles of Programming Languages (POPL 2004), pp. 281–292. ACM Press, New York (2004)
29. Eigner, F.: Type-based verification of electronic voting systems. Master’s thesis, Saarland University (2009)
30. Fisher, D.: Millions of .Net Passport accounts put at risk. eWeek (May 2003) (Flaw detected by Muhammad Faisal Rauf Danko)
31. Fournet, C., Gordon, A.D., Maffei, S.: A type discipline for authorization in distributed systems. In: Proc. 20th IEEE Symposium on Computer Security Foundations (CSF), pp. 31–45. IEEE Computer Society Press, Los Alamitos (2007)
32. Goubault-Larrecq, J., Parrennes, F.: Cryptographic protocol analysis on real C code. In: Cousot, R. (ed.) VMCAI 2005. LNCS, vol. 3385, pp. 363–379. Springer, Heidelberg (2005)
33. Harper, B., Lillibridge, M.: ML with callcc is unsound. Post to TYPES mailing list (July 8, 1991), archived at <http://www.seas.upenn.edu/~sweirich/types/archive/1991/msg00034.html>
34. Jhala, R., Majumdar, R., Rybalchenko, A.: HMC: Verifying functional programs using abstract interpreters. In: Gopalakrishnan, G., Qadeer, S. (eds.) CAV 2011. LNCS, vol. 6806, pp. 470–485. Springer, Heidelberg (2011), <http://arxiv.org/abs/1004.2884v2>

35. Morris Jr., J.H.: Protection in programming languages. *Communications of the ACM* 16(1), 15–21 (1973)
36. Kobayashi, N.: Types and higher-order recursion schemes for verification of higher-order programs. In: *Proc. 36th Symposium on Principles of Programming Languages (POPL 2009)*, pp. 416–428 (2009)
37. Lowe, G.: Breaking and fixing the Needham-Schroeder public-key protocol using FDR. In: Margaria, T., Steffen, B. (eds.) *TACAS 1996*. LNCS, vol. 1055, pp. 147–166. Springer, Heidelberg (1996)
38. Mendler, N.P.: Inductive types and type constraints in the second-order lambda calculus. *Annals of Pure and Applied Logic* 51(1-2), 159–172 (1991)
39. Pierce, B.C.: Programming with intersection types, union types, and polymorphism. Technical Report CMU-CS-91-106, Carnegie Mellon University (1991)
40. Pierce, B.C.: Intersection types and bounded polymorphism. *Mathematical Structures in Computer Science* 7(2), 129–193 (1997)
41. Reynolds, J.C.: Design of the programming language Forsythe. Technical Report CMU-CS-96-146, Carnegie Mellon University (June 1996); Reprinted in O’Hearn, Tennent: *ALGOL-like Languages*, vol. 1, pp. 173–233. Birkhäuser, Basel (1997)
42. Rondon, P.M., Kawaguchi, M., Jhala, R.: Liquid types. In: *Proc. ACM SIGPLAN 2008 Conference on Programming Language Design and Implementation (PLDI 2008)*, pp. 159–169 (2008)
43. Sumii, E., Pierce, B.C.: A bisimulation for dynamic sealing. *Theoretical Computer Science* 375(1-3), 169–192 (2007)
44. Urzyczyn, P.: Positive recursive type assignment. In: Hájek, P., Wiedermann, J. (eds.) *MFCS 1995*. LNCS, vol. 969, pp. 382–391. Springer, Heidelberg (1995)
45. Wagner, D., Schneier, B.: Analysis of the SSL 3.0 protocol. In: *Proc. 2nd USENIX Workshop on Electronic Commerce*, pp. 29–40 (1996)