

Universally Composable Symbolic Analysis of Cryptographic Protocols (The Case of Encryption-Based Mutual Authentication and Key Exchange)*

Ran Canetti and Jonathan Herzog

December 2, 2004

Abstract

Symbolic analysis of cryptographic protocols is dramatically simpler than full-fledged cryptographic analysis. In particular, it is readily amenable to automation. However, symbolic analysis does not *a priori* carry any cryptographic soundness guarantees. We demonstrate how Dolev-Yao style symbolic analysis can be used to assert the security of cryptographic protocols within the universally composable (UC) security framework. Consequently, our methods enable security analysis that is completely symbolic, and at the same time cryptographically sound with strong composability properties.

More specifically, we define a mapping from a class of cryptographic protocols to Dolev-Yao style symbolic protocols. For this mapping, we show that the symbolic protocol satisfies a certain symbolic criterion if and only if the corresponding cryptographic protocol is UC-secure. We concentrate on mutual authentication and key-exchange protocols that use public-key encryption as their only cryptographic primitive. For mutual authentication, our symbolic criterion is similar to the traditional Dolev-Yao criterion. For key exchange, we demonstrate that the traditional Dolev-Yao style symbolic criterion is insufficient, and formulate an adequate symbolic criterion.

Unlike prior work, we concentrate on analyzing single-instance protocols *even in the symbolic model*, relying on the universal composition theorem to guarantee security of the multi-instance case. This further simplifies the analysis and potential automation. It also enables more modular analysis of higher-level protocols.

*This work was presented at the DIMACS workshop on protocol security analysis, June 2004.

Contents

1	Introduction	1
2	Background	5
2.1	The UC framework	5
2.1.1	The Basic Framework	5
2.2	The Dolev-Yao model for symbolic encryption	8
3	The Certified Public-Key Encryption Functionality	12
4	Simple Protocols	13
4.1	The Definition of Simple Protocols	14
4.2	Example: The Needham-Schroeder-Lowe protocol	17
4.3	Example: The Dolev-Dwork-Naor protocol	19
5	The Mapping Lemma	19
6	Symbolic analysis of UC mutual authentication	24
6.1	The ideal 2-party mutual authentication functionality	25
6.2	Dolev-Yao mutual authentication	25
6.3	Soundness and completeness of the symbolic criterion	25
7	Symbolic analysis of UC key exchange	27
7.1	The ideal key exchange functionality	28
7.2	Dolev-Yao key exchange	28
7.2.1	The traditional symbolic criterion	29
7.2.2	The new symbolic criterion	30
7.3	Soundness and completeness of the symbolic criterion	32
8	Future research	34
A	Realizing F_{CPKE}	38
B	Index of notation	41
B.1	Standard mathematical/cryptographic notation	41
B.2	The UC framework	42
B.3	The Dolev-Yao Model	43
B.4	Executions and Traces	44
B.5	Simple protocols	44

1 Introduction

Analyzing the security of cryptographic protocols is a slippery business. One of the main reasons for the complexity and subtlety of this task is the need to capture adversarial behavior which is very powerful in terms of its control of the communication environment but at the same time bounded in its computational resources. Indeed, the security of most interesting cryptographic protocols holds only in a probabilistic sense and against resource-bounded adversaries, and relies in an inherent way on computational intractability assumptions.

An essential first step towards meaningful security analysis is to formulate a mathematical model for representing cryptographic protocols, and a method for translating the security requirements of a given task into mathematical criteria in that model. Indeed, devising such mathematical models and security requirements has been a main focal point in modern cryptography from its early stages, beginning with the notions of pseudo-randomness and semantic security of encryption [14, 15, 42, 29], through zero-knowledge, non-malleability, general cryptographic protocols, and many other examples (e.g., [31, 32, 26, 30, 43, 11, 18]). In particular, we now have a variety of mathematical models (including [18, 48, 19]) where one can represent cryptographic protocols, specify the security requirements of cryptographic tasks, and then potentially prove that (the mathematical representation of) a protocol meets the specification in a way that is believed to faithfully represent the security of actual protocols in actual systems.

However, proving security of cryptographic protocols in any of these models, typically under a set of complexity hardness assumptions, is highly non-trivial. Proofs are hard to write and verify, and become intricate even for relatively simple protocols for simple tasks. Moreover, proofs typically require “human creativity” (e.g., in coming up with a reduction to the underlying hard problem) and are hard to mechanize. Consequently, directly and rigorously proving security of large scale cryptographic protocols — or, equivalently, of large scale systems that make use of cryptographic protocols — is a daunting prospect.

One way to mitigate this analytical complexity is to use *security-preserving protocol composition* (as in e.g. [18, 19, 10]). Here each component in a complex system can be analyzed as stand-alone, while still guaranteeing security of the overall system. This method indeed mitigates some of the complexities in protocol analysis. Still, even when analyzing a single protocol instance as stand-alone, one has to work in a model that directly represents computational hardness assumptions and cryptographic primitives. Among other issues, this requires dealing directly with randomized algorithms and asymptotic formulations such as polynomial time computation. Furthermore, any system that uses cryptography would have to be analyzed in the same complex model, even if the cryptography is only a small part of the system.

Several alternatives to the cryptographic approach to protocol security analysis have been proposed over the years. Examples include the Dolev-Yao model [27] and its many derivatives (e.g. [52, 25]), the BAN logic [17], and a number of process calculi and other models, e.g. [3, 35, 36, 39]. The common theme in these approaches is that cryptographic primitives are represented either as symbolic operations or as inherent model rules that guarantee a set of idealized security properties *by fiat*. (For instance, transmission of encrypted data is modeled as communication that is inaccessible to the adversary, e.g. [17], or as a symbolic operation that completely hides the message, e.g. [27].) Consequently, the modeling becomes dramatically simpler. There is no need to resort to computational assumptions or to model computationally bounded computation; randomization can be replaced by non-determinism; and protocols can be modeled by simple finite constructs without asymptotics. Indeed, protocol analysis in these models is much simpler, more mechanical, and amenable to automation (see e.g. [38, 41, 51, 45, 13]). These are desirable properties when

attempting to analyze large-scale systems.

However, these “formal models” for protocol analysis (so called for their genesis in formal methods) are limited in a number of ways. First, they are typically designed specifically to capture the security of protocols using either one or a small number of cryptographic primitives (e.g., encryption, or digital signatures). Consequently, a new model is needed to capture each new primitive. More importantly, at least until very recently, there was essentially no concrete justification to the highly idealized and abstract modeling of the underlying cryptographic primitives. Consequently, these models could not be used to assert the security of “concrete” protocols, where the abstract security primitives are realized by actual algorithms. (It should be stressed, however, that many of these models were instrumental in finding flaws in protocols.)

Recently some efforts were made towards devising formal models that enjoy the relative simplicity obtained by “abstracting out” the cryptography, while at the same time providing some form of “cryptographic soundness”. One attractive approach towards this goal was initiated in Abadi and Rogaway [4] (with several follow-up works), which provide relatively simple formal criteria for when distribution ensembles constructed from repeated applications of semantically secure symmetric encryption are computationally indistinguishable. Micciancio and Warinschi [44] extend this approach to analyzing the security of protocols in face of active adversaries. Specifically, they provide formal criteria for when two-party protocols that use public-key encryption secure against the chosen-ciphertext attack (as in [49, 26]) are secure *mutual authentication* protocols (as defined in [12]).

An alternative approach was taken in the works of Backes, Pfitzmann and Waidner, and Canetti (e.g. [48, 19, 5, 7, 21]). Here the idea is to define idealized abstractions of cryptographic primitives directly in a full-fledged cryptographic model. These abstractions are realizable by actual concrete protocols in a cryptographic setting, and at the same time can be used by higher-level protocols as abstract primitives. The soundness of this approach relies on a strong composition theorem. This approach is attractive due to its generality and the strong compositional security properties it provides; furthermore, in some cases analyzing the high-level protocols based on the abstract primitives is more straightforward and mechanical. However, it still requires the analysis to directly reason about protocols within a full-fledged cryptographic model with its asymptotics, negligible probabilities, and other subtleties, rather than in a more abstract, formal model. Consequently, this approach retains much of the original complexity of the problem.

This work demonstrates how formal and symbolic reasoning in a simple finite model can be used to assert the security of protocols within a full-fledged cryptographic model with strong composability properties. (Specifically, we use the universally composable (UC) security framework [19]). In a way, we obtain the best of both approaches described above: On the one hand, our approach guarantees the strong security and composition properties of the cryptographic model. On the other hand, we end up with a relatively simple, finite formal model, and symbolic criteria to verify within that model. In fact, our formal criteria are even simpler than the existing ones mentioned above in that they are stated in terms of a *single instance* of the protocol. In contrast, existing formal criteria (e.g., [27, 52, 25]) directly address the multi-instance case; such criteria typically involve a large (sometimes infinite) number of states, and are thus harder to verify mechanically.

We exemplify our approach via two concrete and relatively simple examples. In both examples we concentrate, as in [44], on protocols whose only cryptographic primitive is public-key encryption secure against the chosen-ciphertext attack (CCA-secure). Moreover, as there, we consider only protocols with a strict and quite restrictive format, which nonetheless allows the expression of interesting and known protocols. (The goal of the restrictive format is to allow the use of standard symbolic protocol syntax. Our results can be easily extended to richer formats.) We show how to

translate such protocols to symbolic protocols in a relatively simple formal model. We then provide two symbolic criteria in the formal model: a symbolic mutual authentication property much like the standard symbolic property in the literature, and a symbolic key exchange property which is actually quite different from the standard symbolic property of the literature. Finally we show that:

1. The original, concrete protocol obtains UC mutual authentication (i.e., it securely realizes the mutual authentication ideal functionality in the UC framework) *if and only if* the symbolic protocol obtains symbolic mutual authentication.
2. The original, concrete protocol obtains UC key exchange (i.e., it securely realizes the key exchange ideal functionality in the UC framework) *if and only if* the symbolic protocol obtains symbolic key exchange.

It is stressed that the symbolic criterion is formulated in terms of a finite and relatively simple model, whereas the concrete criterion is formulated in the standard asymptotic terms of cryptographic security. Still, equivalence holds. Furthermore, the entire analysis, including the symbolic criteria, is formulated in terms of a single instance of the protocol in question. Security in a setting where an unbounded number of instances may run concurrently with arbitrary other protocols, is guaranteed via the UC and UC with joint state theorems [19, 24].

The translation from the “bare” cryptographic model to the symbolic model is done in two steps. The first step is to translate the original protocol into a protocol where the use of the encryption scheme is replaced by use of the ideal public-key encryption functionality, F_{PKE} as in [19, 23]. This step is justified by the equivalence between F_{PKE} and CCA-secure encryption (see [19]). For convenience, we will use in the body of this paper a variant of F_{PKE} , called F_{CPKE} , which guarantees ideal binding between public keys and the identities of the corresponding decryptors. However, we show (in Appendix A) that F_{CPKE} can be realized in a straightforward way, using F_{PKE} and a “registration” functionality F_{REG} .

The second step of the translation is to translate a protocol in the F_{CPKE} -hybrid model to a “symbolic protocol” in our formal model. In this latter model internal operations are implicit, protocol messages are modeled as parse-trees of compound symbolic expressions, and the adversary is represented as a non-deterministic sequence of operations on the message space.

Our symbolic mutual authentication criterion is quite standard, requiring that if a party successfully completes a session then its intended peer has initiated this session. (Essentially the same criterion appears in [7, 8] and [44], except that there it relates to a multi-session execution whereas here it relates to a single-session execution.) In contrast, our symbolic key-exchange criterion is stronger than the standard criterion in the symbolic-encryption literature. The standard criterion only requires that it is impossible to derive the symbol representing the session-key from the adversary’s view. We demonstrate that this criterion is insufficient for obtaining reasonably secure concrete key-exchange protocols. Instead, our symbolic criterion has the flavor of a “real-or-random” test, where the adversary can continue interacting with the system after the test value is known, as is common in cryptographic modeling (see e.g. [49, 50, 47, 22]). Still, the criterion is cast in a symbolic model, and relates to single session of the protocol.

A main technical tool that enables our result is a mapping from probability ensembles that describe executions of a concrete protocol in the UC framework to “traces” of the corresponding symbolic protocol. These traces are finite sequences of symbols, and do not involve probabilities or asymptotics. Still, it is guaranteed that any ensemble describing an execution of the concrete protocol is mapped to *some* valid trace of the symbolic protocol.

Related work. Pfitzmann and Waidner [48] provide a general definition of integrity properties and prove that such properties are preserved under protocol composition in their framework. Our symbolic mutual authentication criterion can be cast as such an integrity property. In addition, Backes, Pfitzmann and Waidner [8, 6], building on the idealized cryptographic library in [5], demonstrate that several known protocols satisfy a property that is similar to our symbolic mutual authentication criterion. However, these results do not shed light on the question which is the focus of this work, namely whether a given protocol (say, the Needham-Schroeder-Lowe protocol) realizes an ideal functionality (say, the mutual authentication functionality) in a cryptographic model (say, the UC framework.) Furthermore, since the [5] abstraction is inherently multi-session, the [7] analysis has to directly address the more complex multi-session case.

Our results for mutual authentication protocols follow the lines of Micciancio and Warinschi [44]. However, since we use the UC abstraction of idealized encryption, our characterization results are unconditional (rather than computational), can be stated in the simpler terms of a single session, and provide the stronger security guarantees of the UC framework.

Blanchet [13] provides a symbolic criterion (cast in a variant of the spi-calculus [2]) that captures a secrecy property, called “strong secrecy”, that is similar to our symbolic secrecy criterion for the exchanged key. Essentially, the criterion says that the view of any adversarial environment remains unchanged (modulo renaming of variables) when the symbol representing the secret key is replaced by a fresh symbol that’s unrelated to the protocol execution. Indeed, Blanchet’s ability to automatically verify his criterion gives hope that our criterion, within our model, is also automatically verifiable.

Very recently, Backes and Pfitzmann [9] have proposed a secrecy criterion for abstract key-exchange protocols that use their cryptographic library, and demonstrated that this criterion suffices for guaranteeing cryptographic security. However, it seems that their criterion is overly restrictive in that it rejects some natural key exchange protocols (such as, say, Version 1 of the Needham-Schroeder-Lowe key-exchange protocol as described in Section 4.2, that satisfies the secrecy condition in this work).

Finally, Herzog, Liskov and Micali [34] provide an alternative cryptographic realization of the Dolev-Yao abstraction of public-key encryption. Their realization makes stronger cryptographic requirements from encryption scheme in use, and assumes a model where both the sender and the receiver have public keys. Herzog later relaxes this requirement to standard CCA-2 security[33], but this work (lacking any composition theorems) still considers the multi-session case. Furthermore, it only connects executions of protocols in the concrete setting to executions of protocols in the symbolic setting. It fails to discuss whether security in the symbolic setting implies or is implied by security in the concrete setting, which the main focus of this work.

Organization. We begin with some background (Section 2) on the frameworks of Universal Composition (Section 2.1) and the Dolev-Yao model of symbolic cryptography (Section 2.2). We then introduce and describe the certified public-key functionality, F_{CPKE} (Section 3). Section 4 presents the class of “simple protocols,” which will be meaningful in both the Dolev-Yao and UC models. We illustrate this definition by expressing within it the Needham-Schroeder-Lowe protocol (Section 4.2) and the Dolev-Dwork-Naor protocol (Section 4.3). Section 5 presents and proves validity of the mapping from executions of simple protocols in the UC framework to executions in the symbolic model. We then define the security goal of mutual authentication in the UC framework (Section 6.1) and the symbolic model (Section 6.2) and prove their equivalence theorem. Lastly, we define the goal of secure key-exchange in both the UC framework (Section 7.1) and the symbolic model (Section 7.2), and prove their equivalence. We finish with some directions for future work.

2 Background

2.1 The UC framework

We review of the UC security framework [19]. The presentation is informal for clarity and brevity. Full details (as well as a history of works leading to that framework) appear there. We first present the model of computation, ideal protocols, and the general definition of securely realizing an ideal functionality. Next we present hybrid protocols and the composition theorem.

Protocol syntax. Following [31, 28], a protocol is represented as a system of interactive Turing machines (ITMs), where each ITM represents the program to be run within a party when participating in some protocol instance. Specifically, the `INPUT` and `OUTPUT TAPES` model inputs and outputs that are received from and given to other programs running within the same party (say, the same physical computer), and the `COMMUNICATION TAPES` model messages sent to and received from the network. Adversarial entities are also modeled as ITMs. Each ITM has a `session-identifier` (SID) that identifies which session (or, protocol instance) the ITM belongs to. It also has a `party identifier` (PID) that typically identifies the party within which the program is running. The pair (SID,PID) is a unique identifier of the ITM in the system. (That is, no two ITMs in the system can share both the same SID and the same PID.) In the case of mutual authentication and key exchange protocols, each ITM will also have a `role identifier` (RID) which will determine whether the ITM is an initiator or a responder in the protocol.

We assume that all ITMs run in probabilistic polynomial time (PPT). An ITM is PPT if there exists a constant $c > 0$ such that, at any point during its run, the overall number of steps taken by M is at most n^c , where n is the overall number of bits written on the input tape of M in this run, plus k , where k is the security parameter.

2.1.1 The Basic Framework

Protocols that securely carry out a given task (or, protocol problem) are defined in three steps, as follows. First, the process of executing a protocol in the presence of an adversarial environment is formalized. Next, an “ideal process” for carrying out the task at hand is formalized. In the ideal process the parties do not communicate with each other. Instead they have access to an “ideal functionality,” which is essentially an incorruptible “trusted party” that is programmed to capture the desired functionality of the task at hand. A protocol is said to securely realize an ideal functionality if the process of running the protocol amounts to “emulating” the ideal process for that ideal functionality. We overview the model for protocol execution (called the *real-life model*), the ideal process, and the notion of protocol emulation.

The model for protocol execution. We sketch the process of executing a given protocol p (run by some set of parties) with an adversary A and an environment machine Z with input z . All parties have a security parameter $k \in \mathbf{N}$. The execution consists of a sequence of *activations*, where in each activation a single participant (either Z , A , or some party) is activated, and may write on a tape of at most *one* other participant, subject to the rules below. Once the activation of a participant is complete (i.e., once it enters a special waiting state), the participant whose tape was written on is activated next. (If no such party exists then the environment is activated next.)

The environment is the first to be activated. In each activation it may read the contents of the output tapes of all parties, it may invoke a new party that runs the current instance of the protocol, or it may write information on the input tape of either one of the parties or of the adversary.

Once the adversary is activated, it may read its own tapes and the outgoing communication tapes of all parties. It may either deliver a message to some party by writing this message on the party's incoming communication tape or report information to Z by writing this information on its output tape. (The full UC model as described in [19, 24] also allows the adversary to adaptively corrupt honest participants. However, in order to be compatible with existing symbolic analysis methods we consider only static party corruptions in this work.) The delivered messages need not bear any relation to the messages sent by the parties. (This essentially means that the underlying communication model is unauthenticated.)

Once a party is activated (either due to an input given by the environment or due to a message delivered by the adversary), it follows its code and possibly writes a local output on its output tape or an outgoing message on its outgoing communication tape.

The protocol execution ends when the environment halts. The output of the protocol execution is the output of the environment. Without loss of generality we assume that this output consists of only a single bit.

Let $\text{EXEC}_{\mathbf{p},\mathbf{A},\mathbf{Z}}(k, z, \vec{r})$ denote the output of environment Z when interacting with parties running protocol \mathbf{p} on security parameter k , input z and random input $\vec{r} = r_Z, r_A, r_1, r_2, \dots$ as described above (z and r_Z for Z ; r_A for A , r_i for party P_i). Let $\text{EXEC}_{\mathbf{p},\mathbf{A},\mathbf{Z}}(k, z)$ denote the random variable describing $\text{EXEC}_{\mathbf{p},\mathbf{A},\mathbf{Z}}(k, z, \vec{r})$ when \vec{r} is uniformly chosen. Let $\text{EXEC}_{\mathbf{p},\mathbf{A},\mathbf{Z}}$ denote the ensemble $\{\text{EXEC}_{\mathbf{p},\mathbf{A},\mathbf{Z}}(k, z)\}_{k \in \mathcal{N}, z \in \{0,1\}^*}$.

Ideal functionalities and ideal protocols. Security of protocols is defined via comparing the protocol execution to an *ideal protocol* for carrying out the task at hand. A key ingredient in the ideal protocol is the *ideal functionality* that captures the desired functionality, or the specification, of that task. The ideal functionality is modeled as another ITM (representing a “trusted party”) that interacts with the parties and the adversary. More specifically, in the ideal protocol for functionality F all parties simply hand their inputs to an ITM running F . (We will simply call this ITM F . The SID of F is the same as the SID of the ITMs running the ideal protocol.) In addition, F can interact with the adversary according to its code. Whenever F outputs a value to a party, the party immediately copies this value to its own output tape. We call the parties in the ideal protocol *dummy parties*. The adversary interacting with the ideal protocols is called the *simulator* and denoted S .

Let I_F denote the ideal protocol for functionality F . Let $\text{IDEAL}_{F,S,Z}(k, z)$ denote the random variable describing $\text{EXEC}_{I_F,S,Z}(k, z)$. Let $\text{IDEAL}_{F,S,Z}$ denote the ensemble $\{\text{IDEAL}_{F,S,Z}(k, z)\}_{k \in \mathcal{N}, z \in \{0,1\}^*}$.

Securely realizing an ideal functionality. We say that a protocol \mathbf{p} *securely realizes* an ideal functionality F if there exists an ideal-process adversary S such that no environment Z , on any input, can tell with non-negligible probability whether it is interacting with \S and parties running \mathbf{p} , or it is interacting with S and parties running I_F . This means that, from the point of view of the environment, running protocol \mathbf{p} is ‘just as good’ as interacting with an ideal protocol for F .

A distribution ensemble is called *binary* if it consists of distributions over $\{0, 1\}$. We have:

Definition 1 *Two binary distribution ensembles $\{X(k, a)\}_{k \in \mathcal{N}, a \in \{0,1\}^*}$ and $\{Y(k, a)\}_{k \in \mathcal{N}, a \in \{0,1\}^*}$ are called indistinguishable (written $X \approx Y$) if for any $c \in \mathcal{N}$ there exists $k_0 \in \mathcal{N}$ such that for all $k > k_0$ and for all a we have*

$$|\Pr(X(k, a) = 1) - \Pr(Y(k, a) = 1)| < k^{-c}.$$

Definition 2 Let $n \in \mathcal{N}$. Let F be an ideal functionality and let p be an n -party protocol. We say that p securely realizes F if there exists an ideal-process simulator S such that for any environment Z we have

$$\text{IDEAL}_{F,S,Z} \approx \text{EXEC}_{p,A,Z}.$$

This definition is quite strong, and implies many properties that one might desire in a protocol. In particular, it implies that a protocol is *universally composable*: that it can be used as a sub-protocol by an arbitrary higher-level protocol with no ill effects. More specifically, suppose that a protocol p securely realizes a functionality F as defined above, and P is some higher-level protocol. Then the case where P calls p as a sub-protocol is indistinguishable from the case where P uses the functionality F , even when P uses multiple instances of p or F .

Hybrid protocols. Hybrid protocols are protocols where, in addition to communicating as usual as in the standard model of execution, the parties also have access to multiple copies of an ideal functionality. Hybrid protocols represent protocols that use idealizations of underlying primitives, or alternatively make *set-up assumptions* on the underlying network. They are also instrumental in stating the universal composition theorem. Specifically, in an F -hybrid protocol (i.e., in a hybrid protocol with access to an ideal functionality F), the parties may give inputs to and receive outputs from an unbounded number of copies of F .

The communication between the parties and each one of the copies of F mimics the ideal process. That is, giving input to a copy of F is done by directly writing the input value on the input tape of that copy. Similarly, each copy of F writes the output values directly to a tape of the corresponding party. It is stressed that the adversary does not see the interaction between the copies of F and the honest parties.

The copies of F are differentiated using their SIDs (the PIDs of all copies of F are null.) All inputs to each copy and all outputs from each copy carry the corresponding SID. The model does not specify how the SIDs are generated, nor does it specify how parties “agree” on the SID of a certain protocol copy that is to be run by them. These tasks are left to the hybrid protocol. This convention seems to simplify formulating ideal functionalities—and designing protocols that securely realize them—by freeing the functionality from the need to choose the SIDs and guarantee their uniqueness. In addition, it seems to reflect common practice of protocol design in existing networks.

The definition of a protocol securely realizing an ideal functionality is extended to hybrid protocols in the natural way.

The universal composition operation. Let p be an F -hybrid protocol, and let r be a protocol that securely realizes F . The composed protocol p^r is constructed by modifying the code of each ITM in p so that the first message sent to each copy of F is replaced with an invocation of a new copy of r with fresh random input, with the same SID, and with the contents of that message as input. Each subsequent message to that copy of F is replaced with an activation of the corresponding copy of r , with the contents of that message given to r as new input. Each output value generated by a copy of r is treated as a message received from the corresponding copy of F . The copy of r will start sending and receiving messages as specified in its code. Notice that if r is a G -hybrid protocol model (i.e., r uses ideal evaluation calls to some functionality G) then so is p^r .

The universal composition theorem. Let F be ideal functionality. In its general form, the composition theorem basically says that if r is a protocol that securely realizes F then, for any F -

hybrid protocol p , we have that an execution of the composed protocol p^r “emulates” an execution of protocol p . That is, for any adversary A there exists a simulator S such that no environment machine Z can tell with non-negligible probability whether it is interacting with A and protocol p^r or with S and protocol p . As a corollary, we get that if protocol p securely realizes some ideal functionality G , then so does protocol p^r .

Theorem 1 (universal composition [19]) *Let F, G be ideal functionalities. Let p be a protocol in the F -hybrid model, and let r be a protocol that securely realizes F . Then for any adversary A there exists a simulator S such that for any environment Z we have*

$$\text{EXEC}_{p,S,Z}^F \approx \text{EXEC}_{p^r,A,Z}^G.$$

In particular, if p securely realizes functionality G then so does p^r .

We do not discuss the general implications of Theorem 1 in this work. Still, let us mention two implications that are particularly relevant to this work. First, the composition theorem implies that it is sufficient to analyze the security of a protocol in the single-instance case. That is, it is unnecessary to consider any interaction between one protocol and another, or between concurrent runs of a single protocol. If a single run of the protocol (in isolation) securely realizes a single instance of the functionality, then it does so when composed with multiple instances of itself or other protocols.

The second implication of the composition theorem is that we, in this work, do not need to consider specific implementations of encryption. Instead, we can define protocols in terms of access to an ideal encryption functionality. The composition theorem implies that our protocols retain all their security properties even when the functionality is replaced by an actual encryption scheme, so long as the encryption scheme securely realizes the ideal encryption functionality.

2.2 The Dolev-Yao model for symbolic encryption

There are several variations on the Dolev-Yao model, each tailored to a specific tool or application. We provide and discuss a variant which is appropriate for our needs. In particular, we limit our attention to public-key encryption (ignoring e.g., symmetric encryption, signatures and so on) and we add the notion of a *local output*. We first define the basic symbolic algebra that underlies the model.

Definition 3 (The Message Algebra) *The messages of the Dolev-Yao algebra are assumed to be elements of an algebra A of values. There are several types of atomic message:*

- *Identifiers (\mathcal{M}) which are thought of as public and predictable. These are denoted by P_1, P_2, \dots and so on. We assume there to be a finite number of names in the algebra. With each identifier P_i we associate a role o_i which is a value in some finite set \mathcal{O} . In the present work we are concerned with the case $\mathcal{O} = \{\text{Initiator}, \text{Responder}\}$.¹*
- *Random-string symbols (\mathcal{R}) which are thought of as private and unpredictable. These symbols are denoted by R_1, R_2, \dots and so on. They have two purposes: First, they can be used as nonces to ensure freshness of messages and responses. Second, they can also be used as symmetric keys (i.e., output by key-exchange protocols). We assume there to be a finite number of random strings in the algebra.*

¹In general, a party carrying a given identity can of course participate in multiple interactions, where in different interactions it plays different roles. However, since we are using the symbolic model to analyze a single protocol instance, we can assume that each party has a single role.

- *Public keys* (\mathcal{K}_{Pub}) of which we assume there to be a finite number,
- *A garbage term*, written \mathcal{G} , to represent ill-formed messages,
- \perp , to represent an error or failure,
- *Starting*, to indicate that a protocol execution has begun, and
- *Finished*, to indicate that a protocol execution has ended.

Compound messages are created by two operations:

- *encrypt* : $\mathcal{K}_{Pub} \times \mathcal{A} \rightarrow \mathcal{A}$
- *pair* : $\mathcal{A} \times \mathcal{A} \rightarrow \mathcal{A}$

We write $\{\!|m|\!\}_K$ for $enc(K, m)$ and $m_1|m_2$ for $pair(m_1, m_2)$.² We require that there be a mapping $keyof : \mathcal{M} \rightarrow \mathcal{K}_{Pub}$ which maps each name to one public key.

The algebra is assumed to be *free*: every compound message has a unique representation. (Consequently, each compound message can be equated with the “parse tree” which describes the unique way in which the message was constructed.)

Symbolic Protocols Figuratively, parties running a symbolic protocol (i.e., a protocol in the symbolic model) have four components:

- An input port, where the party obtains its initial input and state.
- A communications port, by which it can send and receive messages in the algebra \mathcal{A} .
- A local output port, by which it can output elements of \mathcal{A} .
- Some internal state. Without loss of generality the internal state consists of all the messages received in the execution so far.
- An identity.
- A role in the protocol.

In keeping with the UC framework, we will assume that in each activation participants either output a local output or send a message via the communications port, but not both. More formally:

Definition 4 A symbolic protocol \mathcal{P} is a mapping from the set of states $\mathcal{S} = (\mathcal{A})^*$, the set \mathcal{O} of roles, an element in the algebra \mathcal{A} representing the incoming message, and the set \mathcal{M} of identities, to a value in $\{\text{output}, \text{message}\}$ (indicating whether the party generates output or an outgoing message), plus a value in the algebra (describing the output of outgoing message generated in this activation) plus a new state (which is the old state with the addition of the new incoming message). That is:

$$\mathcal{P} : \mathcal{S} \times \mathcal{O} \times \mathcal{A} \times \mathcal{M} \rightarrow \{\text{output}, \text{message}\} \times \mathcal{A} \times \mathcal{S}.$$

²When three or more terms are written together, such as $m_1|m_2|m_3$, we assume they are grouped to the left. That is, $m_1|m_2|m_3 = pair(pair(m_1, m_2), m_3)$.

The intended semantics is that, when an honest participant receives a message in the algebra \mathcal{A} , it produces, based on its identity and role, either a local output or an outgoing message that goes on the network, and transitions to a new state. It might also terminate, which is represented as a special state which transitions only to itself and outputs only \perp .

Remark: As is usual in Dole-Yao style modeling, the definition of a symbolic protocol does not require that the protocol be efficiently implementable. For example, it is perfectly valid to define a protocol in which a participant can receive the encryption of any message under any public key and output the plaintext as its next action. In this work, we wish to only consider protocols that have an efficient implementation in the UC framework. Thus, we will limit our attention to protocols that are derived from some concrete, efficiently implementable protocol (see Section 4).

The Symbolic Adversary. As in the UC framework, the adversary in the Dolev-Yao model completely controls the network: the honest participants send their messages only to the adversary and receive messages only from the adversary. However, here the adversary has less power over the creation of new messages. In particular, every message transmitted by the adversary must be derivable from the adversary’s initial knowledge and the messages previously received from honest participants. The initial knowledge of the adversary includes:

1. All the public keys (\mathcal{K}_{Pub}),
2. The identifiers of all the principals (\mathcal{M}), and
3. The random-string symbols which the adversary itself generates ($\mathcal{R}_{Adv} \subseteq \mathcal{R}$), which are assumed to be distinct from all the random-string symbols generated by honest participants. (The fact that the adversary cannot generate the random-string symbols assigned to uncorrupted parties represents the fact that in a concrete protocol the adversary’s probability of guessing random strings generated by an honest party is negligible.)
4. The decryption keys of all the parties in \mathcal{M}_{Adv} , which includes all the parties in \mathcal{M} except for the legitimate (and honest) participants in the protocol.

To derive new message, the adversary has access only to a small number of re-write rules:

- decryption of messages with known private keys.
- encryption with public keys,
- pairing of two known elements, and
- separation of a pair into its components.

The following notion of a *closure* captures more formally the expressions derivable by the adversary given a set S of symbolic expressions. Informally, the meaning of a closure is that, if the adversary has “seen” the expressions (or, messages) in $S \subseteq \mathcal{A}$, then it can only create messages in $C[S]$.

Definition 5 (Closure) *Let $\mathcal{R}_{Adv} \subset \mathcal{R}$ denote the set of random-string symbols associated with the adversary, and let $\mathcal{K}_{Adv} = \{K : \exists A \in \mathcal{M}_{Adv} \text{ s.t. } \text{keyof}(A) = K\}$ be the set of private keys known to the adversary. Then the closure of a set $S \subseteq \mathcal{A}$ of expressions, written $C[S]$, is the smallest subset of \mathcal{A} such that:*

1. $S \subseteq C[S]$,

2. $\mathcal{M} \cup \mathcal{K}_{Pub} \cup \mathcal{K}_{Adv} \cup \mathcal{R}_{Adv} \subseteq C[S]$,
3. If $\{\mathbf{m}\}_K \in C[S]$ and $K \in \mathcal{K}_{Adv}$, then $\mathbf{m} \in C[S]$,
4. If $\mathbf{m} \in C[S]$ and $K \in C[S]$, then $\{\mathbf{m}\}_K \in C[S]$,
5. If $\mathbf{m}|\mathbf{m}' \in C[S]$, then $\mathbf{m} \in C[S]$ and $\mathbf{m}' \in C[S]$, and
6. If $\mathbf{m} \in C[S]$ and $\mathbf{m}' \in C[S]$, then $\mathbf{m}|\mathbf{m}' \in C[S]$, where \mathbf{m}, \mathbf{m}' are elements of \mathcal{A} .

Execution of a symbolic protocol. The execution of a protocol in the Dolev-Yao model needs to first provide an initial state to each participant involved. Then, each participant sends messages to the adversary and reacts to messages from the adversary according to their current state and transition function. The adversary, on the other hand, can non-deterministically perform any sequence of actions, including computation of new messages and sending of known/computed messages. Essentially, a Dolev-Yao trace of an execution contains three things:

1. The initial states/inputs to the participants,
2. The sequence of messages exchanged between the symbolic adversary and the honest participants, and
3. The internal operations of the adversary in deriving the delivered messages.

More formally:

Definition 6 A Dolev-Yao trace for protocol \mathcal{P} is a sequence of events

$$H_0 \ H_1 \ H_2 \ H_3 \ \dots \ H_{n-2} \ H_{n-1} \ H_n$$

where H_i is either

- An event of the form $[\text{“input”}, P, o_i, P', S]$, which indicates the input of participant P (here o_i is the role of P in the protocol, P' is the identity of the peer with which to interact, and S is a general additional internal state),
- an event of one of the following forms, called adversary events (where $j, k < i$):
 - $[\text{“enc”}, j, k, \mathbf{m}_i]$, in which case $\mathbf{m}_k \in \mathcal{K}_{Pub}$ and $\mathbf{m}_i = \{\mathbf{m}_j\}_{\mathbf{m}_k}$,
 - $[\text{“dec”}, j, k, \mathbf{m}_i]$, in which case $\mathbf{m}_k \in \mathcal{K}_{Pub}$, $\mathbf{m}_k^{-1} \in \mathcal{K}_{Adv}$, and $\mathbf{m}_j = \{\mathbf{m}_i\}_{\mathbf{m}_k}$,
 - $[\text{“pair”}, j, k, \mathbf{m}_i]$, in which case $\mathbf{m}_i = \mathbf{m}_j|\mathbf{m}_j$
 - $[\text{“extract-l”}, j, \mathbf{m}_i]$, in which case $\mathbf{m}_j = \mathbf{m}_i|\mathbf{m}_k$ for some $\mathbf{m}_k \in \mathcal{A}$,
 - $[\text{“extract-r”}, j, \mathbf{m}_i]$, in which case $\mathbf{m}_j = \mathbf{m}_k|\mathbf{m}_i$ for some $\mathbf{m}_k \in \mathcal{A}$,
 - $[\text{“random”}, \mathbf{m}_i]$, in which case $\mathbf{m}_i = R$ for some $R \in \mathcal{R}_{Adv}$,
 - $[\text{“name”}, \mathbf{m}_i]$, in which case $\mathbf{m}_i = A$ for some $A \in \mathcal{M}$,
 - $[\text{“pubkey”}, \mathbf{m}_i]$, in which case $\mathbf{m}_i = K$ for some $K \in \mathcal{K}_{Pub}$,
 - $[\text{“deliver”}, j, P_i]$, in which case the message \mathbf{m}_j is delivered to party P_i .
- An event of the form (P_i, L_i, \mathbf{m}_i) , called participant events, where

- $P_i \in \mathcal{M}$ is a participant name,
- m_i is an element of \mathcal{A} , and
- $L_i \in \{\text{output}, \text{message}\}$ indicates whether m_i is a local output or a message,

in which case

$$\mathcal{P}(S_j, o_i, m, P_i) = (L_i, m', S_i)$$

where

1. [“deliver”, k, P_i] is the most recent adversary event in the trace, for some k .
2. m is the second element of the k th adversary event in the current trace.
3. S_j is the current state of P_i . Without loss of generality, we think of the current state of P_i as the sequence of inputs and messages received by P_i , as appears in the transcript so far.
4. o_i is the role of P_i in this instance of the protocol. (The value of o_i is taken from the transcript of the execution so far.)

3 The Certified Public-Key Encryption Functionality

In this section we present the certified public-key encryption ideal functionality, F_{CPKE} . This functionality can be viewed as providing a “bridge” between the cryptographic notion of CCA-security and the Dolev-Yao abstraction of public-key encryption.

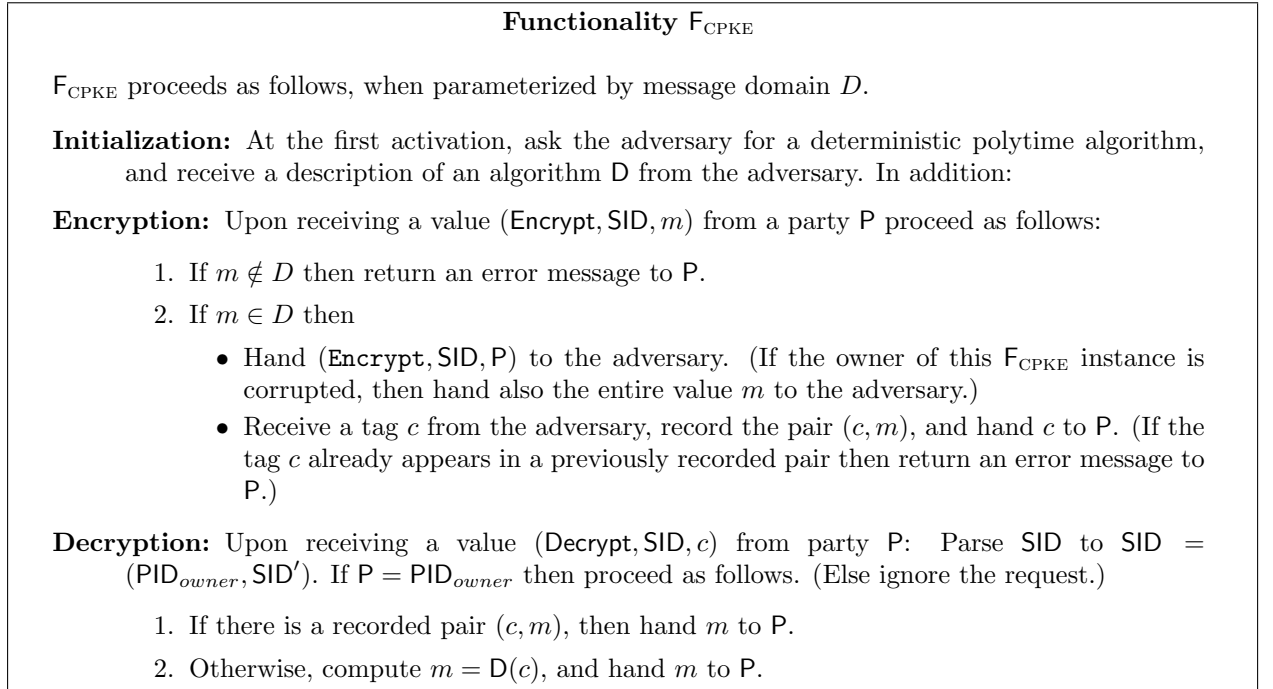


Figure 1: The certified public-key encryption functionality, F_{CPKE}

Functionality F_{CPKE} is presented in Figure 1. The formalization here is based on past formalizations of the public-key encryption functionality, F_{PKE} [19, 23], but differs from those formulations in

a number of respects, outlined below. We assume familiarity with the formulations of [19, 23] and the motivating discussions for the approach that appear there. The main difference from F_{PKE} is that in F_{CPKE} there is no explicit key generation. Instead, the session identifier (SID) of F_{CPKE} explicitly contains the identity of the legitimate receiver of encrypted messages. That is, if an instance of F_{CPKE} has session identifier SID , then SID is of the form $\langle \text{PID}_{\text{owner}}, \text{SID}' \rangle$ where $\text{PID}_{\text{owner}}$ is the party identifier (PID) of the legitimate decryptor. Thus, it is guaranteed that only the legitimate receiver can decrypt messages. This represents an instance of a public-key encryption scheme together with a “key registration service” which provides ideal binding between the public-key and the identity of the owner of the corresponding decryption key. Providing a level of abstraction that hides the process of key generation and registration is done in order to provide a closer correspondence to the Dolev-Yao model. Indeed, in that model the key generation and registration process is not part of the model. Instead, it is assumed that the parties know in advance their private keys and the public keys of all other parties.³

In addition, F_{CPKE} differs from F_{PKE} in another, more technical aspect. In F_{PKE} , if a decryption request arrives where the ciphertext c is not registered (i.e., it was not generated via running the protocol), then the adversary is allowed to determine the decryption value on the spot. In contrast, F_{CPKE} requires the adversary to provide a deterministic algorithm, D in advance; the decryption value of unregistered ciphertexts is then determined by applying D to the ciphertext. This extra restrictiveness of F_{CPKE} makes sure that a ciphertext gets decrypted to a value that was uniquely determined no later than at time of encryption. While this distinction is moot when dealing with the standard case where the decryption algorithm is local and depends only on the ciphertext and the decryption key, it becomes important when moving to the Dolev-Yao model where the specification is completely functional and does not depend on such implementation details. Indeed, the fact that the decryption value of a ciphertext is determined as soon as the ciphertext is generated (either legitimately or illegitimately) is necessary for our equivalence theorems to hold.

Recall that a public-key encryption scheme securely realizes F_{PKE} (with respect to non-adaptive corruptions) if and only if it is CCA-secure [19, 23]. Using techniques similar to those of [21] we have that F_{CPKE} is realizable given any CCA-secure encryption scheme, plus an ideal “registration service” that allows parties to register their public keys, and obtain in an ideally authenticated way the values registered by other parties. See Appendix A for more details on how F_{CPKE} can be realized.

4 Simple Protocols

Both the Dolev-Yao model and the UC framework allow protocols a great deal of flexibility, though in different ways. The Dolev-Yao model strictly regulates the forms of legal messages, restricting messages to the algebra \mathcal{A} . However, protocols can consist of any sequence of messages, including sequences that cannot be efficiently computed. The UC framework, on the other hand, requires only that participants run efficiently. So long as it obeys this one restriction, the protocol can consist of any sequence of bit-strings (or distribution on bit-strings). To find a common denominator for these two models, we must restrict our attention to those protocols that are allowed in both settings. That is, the protocol must both be efficiently executable and use only Dolev-Yao-style messages.

Inspired by [44], we restrict our attention to a class of protocols that use only operations from

³The distinction between F_{PKE} and F_{CPKE} is analogous to the distinction between F_{SIG} and F_{CERT} in [21]. There, F_{SIG} represents a “bare” signature scheme, and F_{CERT} represents a signature scheme augmented with a trusted registration service that allows parties to register their public keys. Indeed, F_{CPKE} can be realized given F_{PKE} and a “registration authority” in the same way that F_{CERT} is realized thre given F_{SIG} and a registration authority.

a small set. In particular, we will provide a “programming language” which can be used to define the behavior of participants in the UC framework model. This programming language provides only small number of operations: randomness generation; encryption and decryption; joining and separation; sending, receiving, and outputting of messages; equality testing; and branching.

Because our language has no loops, all programs must be bound by some fixed number of operations which is constant with respect to the security parameter. Furthermore, all of the allowed operations will execute in time polynomial in the security parameter. Hence, any “simple” protocol (a UC protocol that can be described by our language) will be efficiently implementable. However, we also show that such protocols have a valid interpretation in the Dolev-Yao model. That is, we show how any “simple” protocol \mathbf{p} can be mapped to its symbolic interpretation $\widehat{\mathbf{p}}$.

It is stressed that the restrictive format of simple protocols is put in place for the sole reason of complying with existing Dolev-Yao style symbolic models. It can be easily extended in a number of ways (e.g., adding randomness in the control, allowing public-keys that are generated on the fly, addressing adaptive party corruptions.) Also, to demonstrate the expressive power of the present format, we use it to represent two popular protocols: the Needham-Schroeder-Lowe public-key protocol and the Dolev-Dwork-Naor protocol.

4.1 The Definition of Simple Protocols

A simple protocol is basically a program from a language (Figure 2) that allows concrete cryptographic operations and unsecured inter-process communication. Still, the encryption operation remains abstract even in the concrete setting, being simply a call to the certified public-key encryption functionality F_{CPKE} . The composition theorem allows us to replace these abstract operations with actual encryption schemes while maintaining soundness.

For simplicity we restrict the presentation to two-party protocols, where a participant in a protocol instance has only one peer. Extending the formalism to the case of multiple peers is straightforward. We also restrict the presentation to “semi-deterministic” protocols, or protocols where the only random choices are in choosing random nonces and in realizing F_{CPKE} . In particular, the “control” is deterministic. Extending the treatment to protocols which make other random choices, including random choices at the control level, is also straightforward.

We enforce a type-scheme like that of the Dolev-Yao framework by “tagging” bit-string values with their type. That is, we assume that all well-formed values of a simple protocol have, as a prefix, a string describing their type. A PID is tagged as “name,” the SID for an instance of F_{CPKE} is tagged as “pubkey,” a random value is tagged as a “random,” and the output of a F_{CPKE} instance is tagged both as an “ciphertext” and with the SID of the F_{CPKE} instance that produced it. Furthermore, messages that represent pairs also contain enough information to uniquely determine the two components of the pair.

Definition 7 (Simple protocols) *A simple protocol is a pair of interactive Turing machines (ITMs) $\{M_1, M_2\}$, one for each role, where each machine M_i implements an algorithm described by a pair (Σ, Π) :*

- Σ is a store, a mapping from variables to tagged values (explained further below) and
- Π is a program that expects as input
 - The security parameter k ,
 - Its SID SID, its PID PID, and its RID RID,


```

II ::= BEGIN; STATEMENTLIST
BEGIN ::= input(SID, RID, PID0, PID1, RID2, ...);
        (Store ⟨“role”, RID⟩, ⟨“name”, PID0⟩, ⟨“name”, PID1⟩, ⟨“name”, PID2⟩, ...
        in local variables MyRole, MyName, PeerName, OtherName2, ... respectively.
STATEMENTLIST ::= STATEMENT STATEMENTLIST
                | FINISH
STATEMENT ::= newrandom(v)
            (generate a k-bit random string r and store ⟨“random”, r⟩ in v)
            | encrypt(v1, v2, v3)
            (Send (Encrypt, ⟨PID, SID⟩, v2) to FCPKE where v1 = ⟨“pid”, PID⟩,
            receive c, and store ⟨“ciphertext”, c, ⟨PID1, SID⟩⟩ in v3)
            | decrypt(v1, v2)
            (If the value of v1 is ⟨“ciphertext”, c′⟩ then send (Decrypt, ⟨PID0, SID⟩, c′) to
            FCPKE instance ⟨PID0, SID⟩ receive some value m, and store m in v2
            Otherwise, end.
            | send(v)
            (Send value of variable v)
            | receive(v)
            (Receive message, store in v)
            | output(v)
            (send value of v to local output)
            | pair(v1, v2, v3)
            (Store ⟨“pair”, σ1, σ2⟩ in v3, where σ1 and σ2 are the values of v1 and v2,
            respectively.)
            | separate(v1, v2, v3)
            (if the value of v1 is ⟨“join”, σ1, σ2⟩, store σ1 in v2 and σ2 in v3 (else end))
            | if (v1 == v2 then STATEMENTLIST else STATEMENTLIST
            (where v1 and v2 are compared by value, not reference)
FINISH ::= output(⟨“finished”, v⟩); end.

```

The symbols v , $v1$, $v2$ and $v3$ represent program variables. It is assumed that ⟨“pair”, σ_1, σ_2 ⟩ encodes the bit-strings σ_1 and σ_2 in such a way that they can be uniquely and efficiently recovered. A party’s input includes its own PID, the PID of its peer, and other PIDs in the system. Recall that the SID of an instance of F_{CPKE} is an encoding ⟨PID, SID⟩ of the PID and SID of the legitimate recipient.

Figure 2: The grammar of simple protocols

- PID_1 which represents the name for the other participant of this protocol execution.

The program tags the input values, binds them to variables in the store, and then acts according to a sequence of commands consistent with the grammar in Figure 2.

The structure of simple protocols makes it simple to find the Dolev-Yao counterpart to a simple protocol:

Definition 8 (Mapping of simple protocols to symbolic protocols) Let $p = \{M_0, M_1\}$ be a simple protocol. Then \hat{p} is the Dolev-Yao protocol

$$P_i : \mathcal{S} \times \mathcal{O} \times \mathcal{A} \times \mathcal{M} \rightarrow \{\text{output, message}\} \times \mathcal{A} \times \mathcal{S}$$

which implements $ITM_{\mathbf{M}}$, except that:

- The variables of \mathbf{M} are interpreted as elements of the symbolic message algebra \mathcal{A} .
- Instead of receiving as input SID, PID_0, PID_1, RID , the store is initialized with its own name P_0 , its own key K_{P_0} , and a name P_1 and public key K_{P_1} of the other participant. The symbols P_0 and P_1 represent PID_0 and PID_1 , respectively. Similarly, the symbols K_0 and K_1 represent $\langle PID_0, SID \rangle$ and $\langle PID_1, SID \rangle$, respectively.
- Instead of creating a new random bit-string, the symbolic protocol returns $R_{(i,n)}$ and increments n (which starts at 0),
- Instead of sending $(\text{Encrypt}, \langle PID, SID \rangle, M)$ to F_{CPKE} and storing the result, the composed symbol $\{\{\Sigma(M)\}_{K_{P_1}}\}$ is stored instead (where $\Sigma(M)$ is the value bound to the variable M in the store Σ).
- Instead of sending $(\text{Decrypt}, \langle PID_0, SID \rangle, C)$ to F_{CPKE} and storing the result, the value stored depends on the form of $\Sigma(C)$. If $\Sigma(C)$ is of the form $\{M\}_{K_{P_0}}$ then the value M is stored. Otherwise, the garbage value \mathcal{G} is stored instead.
- Pairing and separation use the symbolic pairing operator.
- Lastly, the bit-strings “starting” and “finished” are mapped to the Dolev-Yao symbols *Starting* and *Finished*, respectively.

Notice that simple protocols make no use of their session identifiers, other than to index the copies of F_{CPKE} . This is somewhat limiting, since using the session identifier in the protocol (e.g., including it in encrypted texts) is often very useful. We impose this restriction for the same reason that we impose all other restrictions on simple protocols: We want to keep our variant of the Dolev-Yao model as close as possible to existing models, and the existing models do not use session identifiers. (Instead, binding messages to sessions is done directly via nonces.) Still, we point out that it is quite easy to extend our results also to protocols that make full use of the SIDs, at the price of ending up with a slightly more complex variant of the Dolev-Yao model.

4.2 Example: The Needham-Schroeder-Lowe protocol

The Needham-Schroeder (public-key) protocol was originally proposed by Needham and Shroeder in 1978 [46]. However, it was later broken by Lowe in 1995, who then proposed a fix and proved his fix correct [37, 38]. By the “Needham-Schroeder-Lowe” protocol, we mean the Needham-Schroeder protocol as fixed by Lowe.

At its most basic form, this protocol consists of three messages between an initiator A and a responder B . Both parties are assumed to have the public key of the other party. We present the version of the protocol that is geared towards obtaining mutual authentication, i.e. the each party only verifies that its peer is “alive” and wishes to interact. A somewhat more intricate version of the protocol, geared towards obtaining key exchange, is discussed below.

Presented in a symbolic form, the protocol begins with the initiator A generating a new random-string symbol (denoted N_a for it’s original designation as “ A ’s nonce”). The first message consists of A ’s name and this string, encrypted in B ’s public key (K_B)⁴:

$$A \rightarrow B : \{A|N_a\}_{K_B}$$

B , upon receiving this message, creates a random string N_b of its own. It then sends back its name, the received random string, and the new random string—all encrypted in A ’s public key:

$$B \rightarrow A : \{B|N_a|N_b\}_{K_A}$$

A , upon receiving this message, checks two things: that the first component is the name of the intended responder, and that second component is the random string that it recently created. If so, it re-encrypts the third component in B ’s public key:

$$A \rightarrow B : \{N_b\}_{K_B}$$

At this point, A terminates and signals a successful protocol execution. B will do the same upon receiving the third message iff the plaintext is its recently-generated random string.

Figure 3 holds the simple programs for the Needham-Schroeder-Lowe initiator and responder roles. (It actually holds two different versions of each, for reasons we explain in a moment.) These programs also specify outputs for these roles, raising the issue exactly which security goals this protocol is supposed to achieve.

It is unclear what formal security goals the protocol was originally proposed to fulfill: few formal, appropriate definitions existed in 1975. In the intervening years, however, consensus has settled upon the *mutual authentication* security goal. Informally speaking, this goal requires that A only successfully complete a run of the protocol with input B only if B has begun a run of the protocol with input A , and vice-versa. (We define this security goal more formally in Section 6.) However, it has also been noted that the protocol could also achieve the *key-exchange* security goal. Very roughly, this protocol should also guarantee secrecy of the encrypted random strings, which can then be used for generating a symmetric key. We incorporate this observation into our simple programs by having both roles output one of the two random strings to be used as a symmetric key. Thus, the two versions. In the first version, the participants output the random string chosen by the initiator, and in the second they output the random string chosen by the responder. In Section 7 we show that Version 1 is a secure key-exchange protocol but that Version 2 is insecure.

⁴For convenience, we will use the Dolev-Yao notation for this exposition. It should be noted, however, that the protocol was not originally expressed in any particular model.

On input $(p1 : \text{PID}; r1 : \text{RID}; s : \text{SID}), (p2 : \text{PID}; r2 : \text{RID}),$ do:

Initiator (M_{init}):

```

send((p1; r1; s), (p2; r2));
newrandom(na);
pair(p1, na, a_na);
encrypt(p2, s, r2, a_na, a_na_enc);
send(a_na_enc);
receive(b_na_nb_enc);
decrypt(b_na_nb_enc, b_na_nb);
separate(b_na_nb, b, na_nb);
if (b == p2) then
  separate(na_nb, na2, nb);
  if (na == na2) then
    encrypt(p2, s, r2, nb, nb_enc);
    send(nb_enc);
    pair(p1, p2, a_b);
    pair(a_b,  $\boxed{x}$ , output);
    output(("finished", output));
    end.
  else send(("finished",  $\perp$ )); end.
else send(("finished",  $\perp$ )); end.

```

Responder (M_{resp}):

```

receive(a_na_enc);
decrypt(a_na_enc,  $\{ \}, \{ \}_a$  na);
separate(a_na, a, na);
if (b == p2) then
  newrandom(nb);
  pair(p1, na, b_na);
  pair(b_na, nb, b_na_nb);
  encrypt(p2, s, r2, b_na_nb, b_na_nb_enc);
  send(b_na_nb_enc);
  receive(nb_enc);
  decrypt(nb_enc, nb2);
  if (nb == nb2) then
    pair(p1, p2, b_a);
    pair(b_a,  $\boxed{x}$ , output);
    output(("finished", output));
    end.
  else send(("finished",  $\perp$ )); end.
else send(("finished",  $\perp$ )); end.

```

Version 1: $x=na$ (Initiator's nonce output as secret key)

Version 2: $x=nb$ (Responder's nonce output as secret key)

Figure 3: The Needham-Schroeder-Lowe protocol

On input $(p1 : \text{PID}; r1 : \text{RID}; s : \text{SID}), (p2 : \text{PID}; r2 : \text{RID}), (m)$, do:

Initiator (M_{init}): <code>send(m);</code> <code>receive(m2_r_enc);</code> <code>decrypt(m2_r_enc, m2_r);</code> <code>separate(m2_r, m2, r)</code> <code>if (m == m2) then</code> <code>send(r);</code> <code>output(⟨“finished”, m⟩)</code> <code>else output(⟨“finished”, ⊥⟩)</code>	Responder (M_{resp}): <code>newrandom(r);</code> <code>pair(m, r, m_r);</code> <code>encrypt(p2, s, r2, m_r, m_r_enc);</code> <code>send(m_r_enc);</code> <code>receive(r2);</code> <code>if (r == r2) then output(⟨“finished”, m⟩)</code> <code>else output(⟨“finished”, ⊥⟩)</code>
---	--

Figure 4: The Dolev-Dwork-Naor protocol

4.3 Example: The Dolev-Dwork-Naor protocol

By the “Dolev-Dwork-Naor protocol” we mean the “message authentication” protocol from the Dolev-Dwork-Naor paper on non-malleable encryption [26]. Although we do not consider the specific goal of message authentication in this paper, we still find this protocol to be a useful demonstration of simple protocols.

In this protocol, the initiator (A) wishes to authenticate⁵ a message (m) to the responder (B). It begins with the initiator sending the message in question to the responder:

$$A \rightarrow B : m$$

The responder chooses a random bit-string (r) and encrypts the message and this string in the initiator’s public key:

$$B \rightarrow A : \{r|m\}_{K_A}$$

The initiator decrypts and verifies the ‘message’ component of the plaintext. If the message component is the message in question, it releases the ‘random’ component:

$$A \rightarrow B : r$$

The simple protocols for the Dolev-Dwork-Naor protocol are in Figure 4.

5 The Mapping Lemma

This section states and proves a main technical tool in our analysis. In particular, we define a translation from executions in the concrete model to executions in the Dolev-Yao model. This translation operates by completely parsing each concrete message as it is generated and replacing successfully parsed messages with symbolic expressions. We then show a useful property about executions: An environment in the concrete model has only a negligible chance of producing an execution which does not translate to a valid symbolic trace of the corresponding symbolic protocol.

⁵In a way that guarantees to the responder message integrity and sender identification.

Phrased another way, we show that the concrete environment has only a negligible chance of creating an execution which is not also available to the symbolic adversary. This property plays a central role in our equivalence theorems of Sections 6 and 7.

We proceed to define the mapping from executions of a concrete simple protocol to Dolev-Yao traces of the corresponding symbolic protocol. As a first step, we define the trace of an execution of a protocol in the concrete setting:

Definition 9 (Traces of concrete protocols) *Let \mathbf{p} be a F_{CPKE} -hybrid protocol. Then inductively define $\text{TRACE}_{\mathbf{p},\mathbf{A},\mathbf{Z}}(k, z, \vec{r})$, the trace of executing protocol \mathbf{p} in conjunction with environment \mathbf{Z} and adversary \mathbf{A} on random inputs \vec{r} , input z , and security parameter k . Initially, the trace is the null string. Then:*

Suppose that the trace of the execution is \mathbf{t} at the beginning of an activation of either the environment or the adversary.

- *If this is an activation of the environment and the environment provides m as input to party (SID, RID) , then the trace at the end of the activation is $\mathbf{t}||\mathbf{G}$ where $\mathbf{G} = \langle \text{"input"}, (\text{SID}, \text{RID}), m \rangle$.*
- *If this is an activation of the adversary and the adversary delivers a message m to party PID , then the trace at the end of the activation is $\mathbf{t}||\mathbf{G}$ where $\mathbf{G} = \langle \text{"adv"}, (\text{PID}, m) \rangle$.*
- *Suppose that the trace of the execution is \mathbf{t} at the beginning of the activation of party (SID, RID) . The activation produces one of two things:*
 - *A local output m , in which case the trace at the end of the activation is $\mathbf{t}||\mathbf{G}$ where*

$$\mathbf{G} = \langle \text{"output"}, \text{PID}, m \rangle,$$
 - *A message m on the communication tape, in which case the trace at the end of the activation is $\mathbf{t}||\mathbf{G}$ where $\mathbf{G} = \langle \text{"message"}, \text{PID}, m \rangle$.*
- *Suppose that the trace of the execution is \mathbf{t} at the beginning of an activation of F_{CPKE} by party PID with call $(\text{Encrypt}, \langle \text{PID}, \text{SID} \rangle, m)$. Then the next event (possibly after the retrieval of D , above) will be a return of c . The trace at the end of F_{CPKE} 's activation is $\mathbf{t}||\mathbf{G}$ where $\mathbf{G} = \langle \text{"ciphertext"}, \langle \text{PID}, \text{SID} \rangle, m, c \rangle$. (If F_{CPKE} returns \perp , then the trace is simply \mathbf{t} .)*
- *Suppose that the trace of the execution is \mathbf{t} at the beginning of a subsequent activation of F_{CPKE} by party PID with call $(\text{Decrypt}, \langle \text{PID}, \text{SID} \rangle, c)$. Then if the next event (possibly after the retrieval of D , above) is return of m then the trace at the end of F_{CPKE} 's return is $\mathbf{t}||\mathbf{G}$ where $\mathbf{G} = \langle \text{"dec"}, \langle \text{PID}, \text{SID} \rangle, c, m \rangle$. (If F_{CPKE} returns \perp , then the trace is simply \mathbf{t} .)*

Let $\text{TRACE}_{\mathbf{p},\mathbf{Z}}(k, z)$ denote the random variable describing $\text{TRACE}_{\mathbf{p},\mathbf{Z}}(k, z, \vec{r})$ when \vec{r} is uniformly chosen. Let $\text{TRACE}_{\mathbf{p},\mathbf{Z}}$ denote the ensemble $\{\text{TRACE}_{\mathbf{p},\mathbf{Z}}(k, z)\}_{k \in \mathcal{N}, z \in \{0,1\}^}$.*

Next, we define a mapping from concrete traces to symbolic traces. This mapping clearly produces something with the same basic form of a Dolev-Yao trace: a sequence of either participant or adversary events. Furthermore, we will show that, if we start with a trace of an execution of concrete protocol, then except with negligible probability we obtain a valid Dolev-Yao trace of the corresponding symbolic protocol.

Definition 10 (The mapping from concrete traces to symbolic traces) Let \mathfrak{p} be a concrete F_{CPKE} -hybrid protocol and let \mathfrak{t} be a trace of an execution of \mathfrak{p} with security parameter k , environment Z with input z , and random input vector \vec{r} . We determine the mapping of \mathfrak{t} to a Dolev-Yao trace in two steps. (These steps can be thought of as two “passes” on the string \mathfrak{t} .)

(I.) First, we read through the string \mathfrak{t} character by character, in order, and inductively define the following partial mapping f from $\{0,1\}^*$ to elements of the algebra \mathcal{A} . (Note that the patterns in \mathfrak{t} addressed below may be nested and overlapping. That is, the same substring may be part of multiple patterns. A pattern is recognized as soon as the last character in the pattern is read.)

- Whenever we encounter a pattern of the form $\langle \text{“name”}, \sigma \rangle$ for some string σ and $f(\langle \text{“name”}, \sigma \rangle)$ is not yet defined then set $f(\langle \text{“name”}, \sigma \rangle) = P$ for some new symbol $P \in \mathcal{M}$ not in the range of f so far. (Recall that the pattern $\langle \text{“name”}, \sigma \rangle$ is generated by a simple protocol when it receives its input.)
- Whenever we encounter in some event a pattern of the form $\langle \text{“random”}, \sigma \rangle$ for some string σ and $f(\langle \text{“random”}, \sigma \rangle)$ is not yet defined then set $f(\langle \text{“random”}, \sigma \rangle) = N$ for some new symbol $N \in \mathcal{R}$ that is not in the range of f so far. (Recall that the pattern $\langle \text{“random”}, \sigma \rangle$ is generated by a simple protocol when it chooses a random string. See Figure 2.)
- Whenever we encounter a pattern of the form $\langle \langle \text{“pid”}, \text{PID} \rangle, \langle \text{“sid”}, \text{SID} \rangle \rangle$ for some strings PID, SID , and $f(\langle \text{“key”}, \langle \text{PID}, \text{SID} \rangle \rangle)$ is not yet defined, then set $f(\langle \text{“key”}, \langle \text{PID}, \text{SID} \rangle \rangle) = K$ for some new $K \in \mathcal{K}_{\text{Pub}}$ not in the range of f .
- Whenever we encounter a pattern of the form $\langle \text{“pair”}, \sigma_1, \sigma_2 \rangle$, then proceed as follows. First, if $f(\sigma_1)$ is not yet defined then set $f(\sigma_1) = \mathcal{G}$, where \mathcal{G} is the garbage symbol. Similarly, if $f(\sigma_2)$ is not yet defined then set $f(\sigma_2) = \mathcal{G}$. Finally, set $f(\langle \text{“pair”}, \sigma_1, \sigma_2 \rangle) = f(\sigma_1) | f(\sigma_2)$.
- Whenever we encounter a pattern of the form $\langle \text{“ciphertext”}, \langle \text{PID}, \text{SID} \rangle, m, c \rangle$ for some strings $\text{PID}, \text{SID}, m, c$, then f is expanded so that $f(\langle \text{“ciphertext”}, \langle \text{PID}, \text{SID} \rangle, c \rangle) = \{f(m)\}_{f(\langle \text{“key”}, \langle \text{PID}, \text{SID} \rangle \rangle)}$. (Recall that such a pattern is generated whenever an encryption call to F_{CPKE} is made. Also, at this point both $f(m)$ and $f(\langle \text{“pubkey”}, \langle \text{PID}, \text{SID} \rangle \rangle)$ must already be defined, since this is an encryption call made by a party running a simple protocol.)
- Whenever we encounter a pattern of the form $\langle \text{“dec”}, \langle \text{PID}, \text{SID} \rangle, c, m \rangle$, then proceed as follows. First, if $f(m)$ is not yet defined, then set $f(m) = \mathcal{G}$, where \mathcal{G} is the garbage symbol. Next, set $f(\langle \text{“dec”}, \langle \text{PID}, \text{SID} \rangle, c \rangle) = \{f(m)\}_{f(\langle \text{“key”}, \langle \text{PID}, \text{SID} \rangle \rangle)}$. (Recall that such a pattern is generated whenever a decryption call to F_{CPKE} is made. The case where $f(m) = \mathcal{G}$ occurs when a ciphertext was not generated via the encryption algorithm. It includes both the case where the decryption algorithm fails and the case where the decryption algorithm outputs a message that cannot be parsed by simple protocols.)

(II.) In the second step, we construct the actual Dolev-Yao trace. Let $\mathfrak{t} = G_1 || G_2 || \dots || G_n$ be the concrete trace. Then construct the Dolev-Yao trace $\hat{\mathfrak{t}}$ by processing each G in turn, as follows:

- If $G_i = \langle \text{“input”}, \langle \text{SID}, \text{RID} \rangle, m \rangle$, then we find $\mathfrak{m} = f(m)$, and generate the symbolic event $H = [\text{“input”}, P, \mathfrak{m}]$ (where P is the symbolic name of the input recipient).
- If $G_i = \langle \text{“ciphertext”}, \langle \text{PID}, \text{SID} \rangle, m, c \rangle$ or $G_i = \langle \text{“dec”}, \langle \text{PID}, \text{SID} \rangle, c, m \rangle$, then no symbolic event is generated.

- If $G_i = \langle \text{"output"}, \text{PID}, m \rangle$ then G_i is mapped to the symbolic participant event

$$(f(\langle \text{"name"}, \text{PID} \rangle), \text{output}, f(m)).$$

- If $G_i = \langle \text{"message"}, \text{PID}, m \rangle$ then G_i is mapped to the symbolic participant event

$$(f(\langle \text{"name"}, \text{PID} \rangle), \text{message}, f(m)).$$

- If $G = \langle \text{"adv"}, \text{PID}, m \rangle$, let $m = f(m)$. Then there are two cases:

1. m is in the closure of the symbolic interpretations of the messages sent by the parties in the execution so far, i.e.

$$m \in C[\{m' : m' = f(m') \text{ and the event } \langle \text{"message"}, \text{PID}, m' \rangle \text{ is a prior event in } t\}].$$

In this case there exists a finite sequence of adversary events that produces m_i . Then G is mapped to this sequence of events $H_{i,1}, H_{i,2} \dots H_{i,n'}$ so that the message of $H_{i,n'-1}$ is m_i and $H_{i,n'} = [\text{"deliver"}, (i, n' - 1), P']$ (where P' is the Dolev-Yao name of the concrete participant who received the message from the concrete adversary).

2. Otherwise, m is not in the above closure. In this case, G maps to the Dolev-Yao event $[\text{"fail"}, m_i]$.

We now show that the mapping defined above is valid. That is, we show that if t is a trace of a simple protocol p then \hat{t} is a Dolev-Yao trace of the symbolic protocol \hat{p} , except for negligible probability. We need to take care of several issues. First, we need to show that the actions of the concrete participants map to valid actions of the symbolic participants. That is, the messages from the execution of concrete protocol p map only to symbolic messages that are compatible with the symbolic protocol \hat{p} . Secondly, we need to show that the concrete adversary is no more powerful than the symbolic adversary. That is, the trace t does not contain adversary messages whose symbolic interpretations are beyond the ability of the symbolic adversary to produce.

In the following lemma, we show that these failures occur with negligible probability, and thus any concrete execution is almost always mapped to a valid symbolic interpretation. In fact, the only potential cause for error is the event where the concrete adversary or environment guess the value of a random string chosen by an uncorrupted party.

Lemma 2 For all simple protocols p , adversaries A , environments Z , and inputs z of length polynomial in the security parameter k ,

$$\Pr [t \leftarrow \text{TRACE}_{p,A,Z}(k, z) : \hat{t} \text{ is a valid DY trace for } \hat{p}] \geq 1 - \text{neg}(k)$$

Proof. Let t be a trace of a simple protocol p . We first show that the probability that \hat{t} includes an event of the form $[\text{"fail"}, m_i]$ is negligible. Next, we show that whenever \hat{t} does not include such an event, it is a valid DY trace of protocol \hat{p} .

Let m_1, m_2, \dots denote the messages that appear in the $\langle \text{"message"}, \dots \rangle$ events in \hat{t} , in order of appearance. Suppose that adversary event of the form $[\text{"fail"}, m_i]$ occurs, which means that the concrete adversary created a message m which is translated to $f(m) = m_i \notin C[m_j : j < i]$. We show that the odds of such an event are negligible: Examine the parse tree of m_i .⁶ By definition,

⁶The parse tree of a message is the tree whose root is the message, and there is an edge from node m to nodes m' and m'' if there is a derivation rule in the algebra A that derives message m from messages m', m'' . Note that the leaves of the tree are the basic symbols in A . Also, the out-degree of a node is at most 2.

membership in $C[\{\mathbf{m}_j : j < i\}]$ is closed under pairing and encryption. Thus, if two siblings in the parse tree are both in $C[\{\mathbf{m}_j : j < i\}]$, then their parent is in $C[\{\mathbf{m}_j : j < i\}]$ as well. Consequently, if every path from root to leaf in the parse tree of \mathbf{m}_i has a node in $C[\{\mathbf{m}_j : j < i\}]$, then $\mathbf{m}_i \in C[\{\mathbf{m}_j : j < i\}]$ as well—a contradiction. Thus, there exists some leaf \mathbf{m}_* in the parse tree of \mathbf{m}_i such that the path to \mathbf{m}_* has no node in $C[\{\mathbf{m}_j : j < i\}]$.

However, since \mathbf{m}_i is in $\hat{\mathbf{t}}$, we have that in \mathbf{t} the adversary has generated a bit-string m that was mapped to \mathbf{m}_i . Call this adversary \mathbf{A}_0 . Then there exists another adversary, \mathbf{A}_1 that produce a bit-string m_* which parses to \mathbf{m}_* . \mathbf{A}_1 first simulates \mathbf{A}_0 to produce m , and then recursively walks down the parse tree of \mathbf{m}_i to \mathbf{m}_* and applying the following deconstructors to m :

- If \mathbf{m}_l is a pair $\mathbf{m}_1|\mathbf{m}_2$, then \mathbf{A}_1 separates $m = \langle \text{“pair”}, \sigma_1 || \sigma_2 \rangle$ into σ_1 and σ_2 and recursively operates on one of them (depending on whether \mathbf{m}_* is a leaf of \mathbf{m}_1 or \mathbf{m}_2).
- If $m = \langle \text{“ciphertext”}, \sigma \rangle$, then \mathbf{A}_1 adversary must “decrypt” σ to continue down to \mathbf{m}_l . That is, \mathbf{A}_1 must produce what F_{CPKE} would return to the appropriate honest party when called with $(\text{Decrypt}, \sigma)$. As mentioned before, $\mathbf{m}_i \notin C[\{\mathbf{m}_j : j < i\}]$. Thus, if the pair (m, σ) is stored in F_{CPKE} (for some m) it must be that the call $(\text{Encrypt}, m)$ was made to F_{CPKE} by \mathbf{A}_0 . Hence, \mathbf{A}_1 (which simulates \mathbf{A}_0) records and recalls the message m which is the “decryption” of σ , and recursively operates on it. If, on the other hand, no pair (m, σ) is stored in F_{CPKE} , then the decryption of σ is the result of $D(\sigma)$, where D is an algorithm supplied by \mathbf{A}_0 . Thus, \mathbf{A}_1 runs D itself to learn the decryption of σ .

By recursively applying the above deconstruction operations, \mathbf{A}_1 produces a string m_* that maps to \mathbf{m}_* . But \mathbf{m}_* is an atomic symbol of the Dolev-Yao algebra which is not in $C[\{\mathbf{m}_j : j < i\}]$. Notice that the only atomic symbols that are not in the initial view of the adversary are the random-number symbols (i.e., symbols in \mathcal{R}) that were not generated by the adversary. However, if \mathbf{m}_* is not in the closure of the adversary’s view, then the view of \mathbf{A}_1 is completely independent from the string m_* . (Independence is argued as follows. If m_* is never included in the parse tree of a message seen by \mathbf{A}_1 then independence is trivial. The only way for m_* to be included in the parse tree of a message seen by \mathbf{A}_1 and still not be in the closure of \mathbf{A}_1 ’s view is if m_* is sent encrypted. However, in this case independence is guaranteed by the code of F_{CPKE} .) Also, since m_* was generated by a protocol participant, then we know it is chosen uniformly from $\{0, 1\}^k$. Thus the probability that \mathbf{A}_1 generates the string m_* is 2^{-k} . Since there are at most a polynomial number of k -bit strings in the view of \mathbf{A}_1 , we have that the overall probability that \mathbf{A}_1 generates a string that maps to \mathbf{m}_* is $\text{poly}(k)2^{-k}$. This means that the probability that $\hat{\mathbf{t}}$ includes an event of the form $[\text{“fail”}, \mathbf{m}_i]$ is $\text{poly}(k)2^{-k}$, which is a negligible function.

It remains to show that, whenever event $[\text{“fail”}, \mathbf{m}_i]$ does not occur, the trace $\hat{\mathbf{t}}$ is a valid trace for $\hat{\mathbf{p}}$. By definition of the fail event, we have that all the adversary events in $\hat{\mathbf{t}}$ are valid. We now show that the participant events in $\hat{\mathbf{t}}$ are valid as well. Suppose that a participant event of the form $(P'_i, L_i, \mathbf{m}_i)$ occurs. Then we need to show that

$$\mathcal{P}(S_j, o_i, \mathbf{m}, P_i) = (L_i, \mathbf{m}', S_i),$$

where

1. $[\text{“deliver”}, k, P_i]$ for some k is the most recent adversary event in the trace $\hat{\mathbf{t}}$.
2. \mathbf{m} is the second element in the k th adversary event in the current trace.
3. S_j is the sequence of inputs and messages received by P_i , as they appear in the current prefix of $\hat{\mathbf{t}}$.

4. o_i is the role of P_i as appears in the current prefix of \hat{t} .

However, this fact follows immediately from the definition of the symbolic counterpart of a simple protocol (Definition 8). Indeed, the transition function \mathcal{P} mimics the instructions of the concrete protocol p . The only difference between the protocols is the naming of variables; furthermore, the structure of simple protocols makes sure that any one-to-one renaming of variables does not affect the messages or outputs generated by the protocol. \square

6 Symbolic analysis of UC mutual authentication

This section demonstrates how one can use symbolic analysis to assert whether a given (concrete, simple) protocol is a UC mutual authentication protocol. More precisely, we first formulate the ideal mutual authentication functionality, F_{2MA} . Next, we formulate a criterion (expressible in terms of symbolic protocols), such that a concrete protocol p securely realizes F_{2MA} in the F_{CPKE} -hybrid model if and only the corresponding symbolic protocol \hat{p} satisfies the symbolic criterion. We note once more that this symbolic criterion is amenable to verification via automated tools, e.g. [51, 41, 40, 45].

As discussed in the Introduction, the results in this section follow in the footsteps of Micciancio and Warinschi [44] with three exceptions. First, our concrete protocols are protocols in the F_{CPKE} -hybrid model, rather than protocols in the plain real-life model. This allows our analysis to be simple and unconditional, yet still concrete due to the composition theorem. Secondly, the composition theorem allows us to simplify our analysis even further by considering only a single execution of the protocol execution, while the analysis of [44] directly deals with the much more complex multi-execution case. Lastly, our result provides a strong security and composition guarantee which is not present in the results of [44].

Formalizing the definition of *mutual authentication* requires surprisingly many choices. The intuition is simple: when one participant in a protocol terminates, it should be the case that the other participant has at least begun and it has the right peer in mind. Some protocols may guarantee more. For example, a protocol might guarantee that when one participant has finished the protocol, the other participant has finished as well as begun. Although the protocol might guarantee this for one of the two participants, however, it cannot guarantee this for both.

Other subtleties exist. For example, should the two participants agree on their respective roles? For example, is it an error if two participants successfully complete the protocol, but both are running the role “initiator?”

Also, should the protocol enforce a *bijection* between successful outputs of the protocol made by the two participants? That is, if the initiator started the protocol, is it an error for there to be two or more outputs by the responder?

For simplicity, we will use in this paper the most basic variant of mutual authentication. That is, we only guarantee that if a party P outputs **success** then the other party at least began the protocol with peer P . There is no guarantee that the roles are different, and multiple outputs are allowed. Our equivalence results, however, can be easily applied to more complex forms of mutual authentication.

6.1 The ideal 2-party mutual authentication functionality

The UC definition of 2-party mutual authentication is embodied in the functionality F_{2MA} (Figure 5). The functionality simply waits until two parties P_0 and P_1 have provided input

$$(\text{Authenticate}, \text{SID}, P_0, P_1) \text{ and } (\text{Authenticate}, \text{SID}, P_1, P_0)$$

respectively. Then, upon request of the simulator, it sends a (**Finished**) output to either party. Note that it is possible for a party to get multiple (**Finished**) output, and that there is no requirement that a (**Finished**) message is received by both parties.

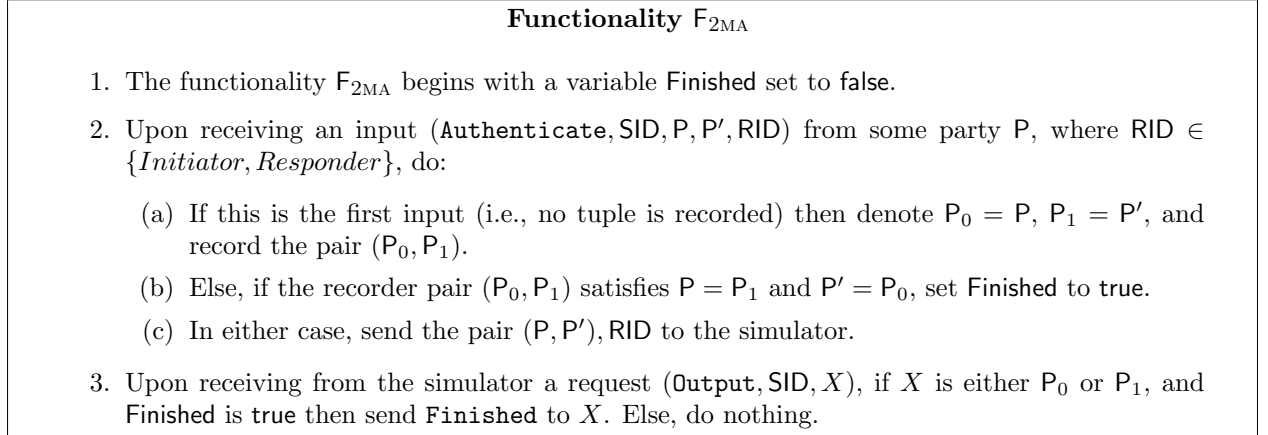


Figure 5: The 2-party mutual authentication functionality

6.2 Dolev-Yao mutual authentication

We specify the symbolic mutual authentication criterion:

Definition 11 (Dolev-Yao 2-party mutual authentication) *A Dolev-Yao protocol \mathcal{P} provides Dolev-Yao mutual authentication (DY-MA) if all Dolev-Yao traces for \mathcal{P} that include an output message $\langle \text{Finished} | P_0 | P_1 | m \rangle$ by participant P_0 , where $P_0, P_1 \notin \mathcal{M}_{Adv}$, include also a previous input message $\langle \text{Starting} | P_1 | P_0 | m' \rangle$ by P_1 .*

That is, if one party outputs a “finished” message indicating a successful execution, the other party has at least output a “starting” message indicating that an execution (with matching values for the participants) has been at least initiated. (The “starting” message also contains the keys of the participants to match the input format expected by a simple protocol in Definition 7, and both messages contain additional messages m and m' to represent additional arbitrary output.)

6.3 Soundness and completeness of the symbolic criterion

Theorem 3 *Let p be a simple two-party protocol. Then p realizes F_{2MA} if and only if the corresponding symbolic protocol \hat{p} satisfies Dolev-Yao 2-party mutual authentication.*

Proof. Assume first that \hat{p} does not achieve Dolev-Yao secure key exchange. Then there exists a valid Dolev-Yao trace where one party outputs (locally) $\langle \text{Finished} | P | P' | m \rangle$ before P' outputs

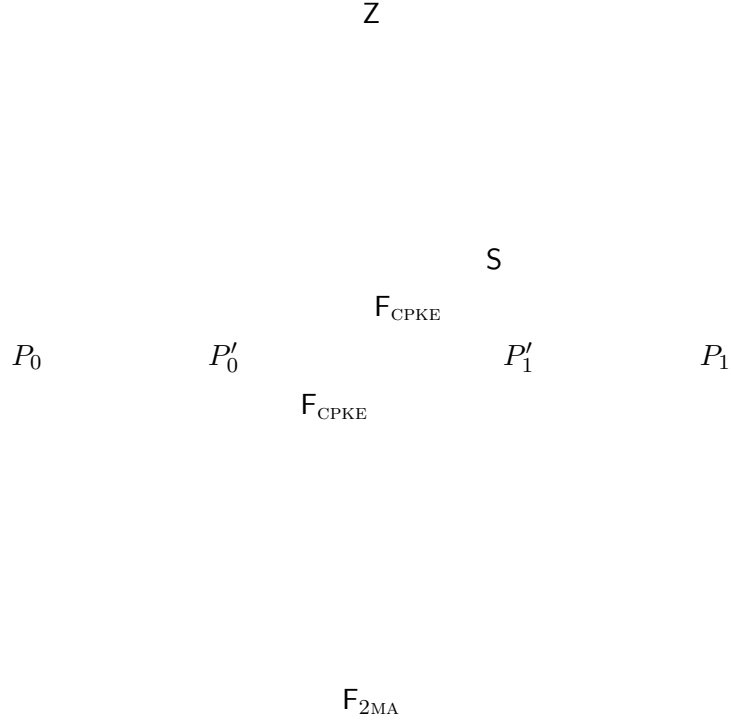


Figure 6: Simulator for $F_{2\text{MA}}$

$\langle \text{Starting} | P' | P, | \mathbf{m}' \rangle$. Given this trace, we construct an environment Z in the UC framework that simply follows the adversary instructions in the given Dolev-Yao trace. (More precisely, Z follows the concrete operations that correspond to the given Dolev-Yao trace.)

When the environment is in the F_{CPKE} -hybrid model, this strategy will produce the same result as in the Dolev-Yao model: one participant will output `Finished` before the other party outputs `(Authenticate, SID, A, B)`. However, this same behavior is simply impossible when interacting with the ideal protocol for $F_{2\text{MA}}$, since in that protocol no simulator can force the dummy parties to produce unmatched output. Thus, this environment can distinguish the ideal execution model from the real execution model with probability 1, and so the protocol \mathbf{p} cannot securely realize $F_{2\text{MA}}$. Finally, we point out that this environment is polynomial in the security parameter: the number of operations in the Dolev-Yao trace is constant, and performing each operation takes polynomial time in the security parameter.

For the other direction, we need to show that if $\hat{\mathbf{p}}$ satisfies Dolev-Yao mutual authentication then \mathbf{p} securely realizes $F_{2\text{MA}}$. To show this, we must show that there exists a simulator such that no environment can distinguish between the concrete protocol and the ideal protocol for $F_{2\text{MA}}$.

The simulator S acts as follows (shown in Figure 6):

- S internally simulates the participants P'_0, P'_1 , and a copy of F_{CPKE} for each. At the beginning, neither of these simulated participants are running.
- When the simulator receives a message (P_i, P_{1-i}) from the functionality $F_{2\text{MA}}$ (indicating that the external dummy participant P_i has received input from the environment and has passed it on to $F_{2\text{MA}}$) it activates both instances of F_{CPKE} with SIDs $\langle P_0, \text{SID} \rangle$ and $\langle P_1, \text{SID} \rangle$ respectively

(if they have not been activated yet). It then activates the simulated participant P'_i on input $(\text{Authenticate}(\text{SID}, P'_i, P'_{1-i}, \text{RID}))$.

- When the simulator receives a message from the environment sent to participant P_i , it forwards that input to the simulated copy P'_i .
- Likewise, when simulated participant P'_i produces a message to send on its communication tape, the simulator sends this message to the environment.
- When simulated participant P'_i produces a message to send an instance of F_{CPKE} , the simulator forwards this to the appropriate instance of F_{CPKE} that it is simulating.
- Likewise, when either instance of simulated F_{CPKE} requires communication with its environment, this communication is sent directly to the external environment. Communication from the external environment to either F_{CPKE} -instance is forwarded directly to the simulated F_{CPKE} -instance.
- When the simulated participant P'_i produces local output **Finished**, the simulator sends the message $(\text{Output}, \text{SID}, P_i)$ to the functionality $F_{2\text{MA}}$.

It remains to show that the simulation is valid. Let Z be an environment. We show that if Z distinguishes between the interaction with the simulator and $F_{2\text{MA}}$, and the interaction with the protocol, then Z violates the mapping lemma (Lemma 2), in the sense that it generates with non-negligible probability a trace ensemble that translates to a symbolic trace that is not valid for the symbolic version of the protocol:

Fix a value k for the security parameter. We observe that the above simulator produces a perfect simulation except when the following bad event occurs: the simulator sends $(\text{Output}, \text{SID}, P_i)$ to the functionality $F_{2\text{MA}}$, but $F_{2\text{MA}}$ does not send **Finished** to dummy party P_i . Furthermore, this event only occurs if the functionality did not previously receive both $(\text{Authenticate}, \text{SID}, P_i, P_{1-i}, \text{RID})$ and $(\text{Authenticate}, \text{SID}, P_{1-i}, P_i, \text{RID})$. Because the simulator sends $(\text{Output}, \text{SID}, P_i)$ to $F_{2\text{MA}}$, the simulated participant P_i produces an output indicating success. Hence, simulated participant P_i must have been started by the simulator, which means that the simulator must have received $(P_i, P_{1-i}, \text{RID})$ from $F_{2\text{MA}}$. Thus, the functionality must have received $(\text{Authenticate}, \text{SID}, P_i, P_{1-i}, \text{RID})$. Hence, it must have been $(\text{Authenticate}, \text{SID}, P_{1-i}, P_i)$ that was not received by the functionality $F_{2\text{MA}}$.

Thus, the simulated party P_{1-i} was not initialized in the simulator. If we look at the trace t of the simulated parties of the execution, it must be that party P_i output **Finished** before party P_j output $(\text{Authenticate}, \text{SID}, P_{1-i}, P_i, \text{RID})$. Thus, in the Dolev-Yao trace \hat{t} that is constructed from the trace t , participant P_i outputs $\langle \text{Finished} | P_i | P_{1-i} | m \rangle$ before participant P_{1-i} outputs $\langle \text{Starting} | P_{1-i} | P_i | m' \rangle$. However, the protocol \hat{p} satisfies Dolev-Yao mutual authentication, and so this trace cannot be valid. Thus the probability that the the environment distinguishes the ideal execution from the real one is the same as its probability to generate traces that translate to invalid symbolic traces—i.e., negligible. \square

7 Symbolic analysis of UC key exchange

This section demonstrates how one can use symbolic analysis to assert whether a given simple protocol is also a UC key exchange protocol. At high level, the structure of this section is similar to that of the previous section: We first recall the ideal key exchange functionality, $F_{2\text{KE}}$; Next, we

demonstrate a criterion (expressible in terms of abstract protocols), such that a concrete protocol p securely realizes F_{2KE} in the F_{CPKE} -hybrid model if and only if the corresponding abstract protocol \hat{p} satisfies the criterion.

It turns out that in the case of key exchange, formulating the abstract criterion is non-trivial and requires an approach that is quite different from existing ones. We demonstrate why the simplistic approach fails and motivate our approach.

7.1 The ideal key exchange functionality

The UC notion of secure two-party key exchange is embodied in the functionality F_{2KE} (Figure 7). As with mutual authentication, many variants exist and we chose to study one basic variant. The functionality generates and stores a single random key κ . Each party sends a simple message indicating its willingness to establish a secret key. Once a given party has done so, the simulator can cause the functionality to distribute the random key κ to that party. It is worth noting that the functionality does *not* require both parties to activate before distributing κ , and thus our version of key exchange does *not* guarantee mutual authentication. (Indeed, key exchange and mutual authentication are arguably quite different tasks.) Also, parties may receive the key κ from F_{2KE} more than once.

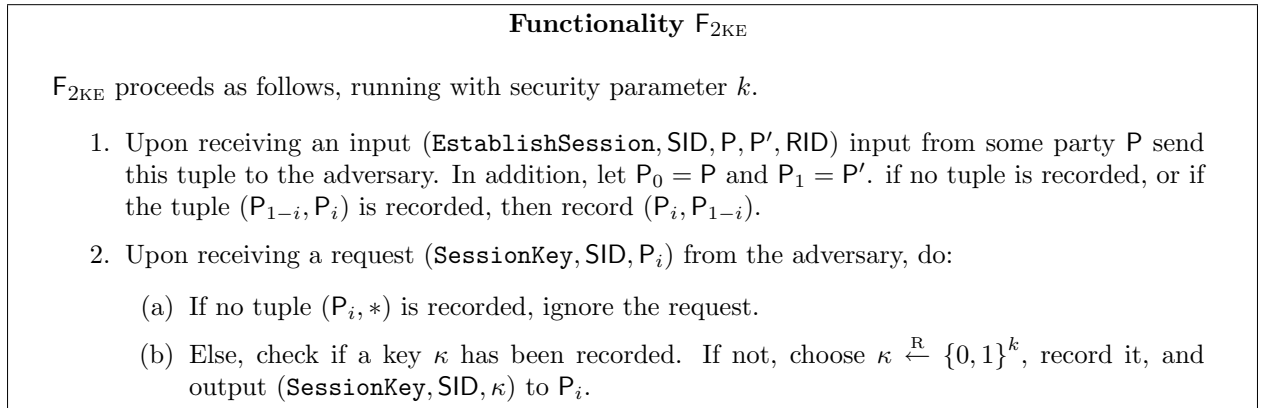


Figure 7: The Key Exchange functionality

7.2 Dolev-Yao key exchange

Our definition of Dolev-Yao key exchange is somewhat more complex. As in the case of mutual authentication, the intuition is simple: if both participants terminate, then they must output the same secret symmetric key R . (Recall that in our version of the Dolev-Yao model, random strings can also be used as symmetric keys.) However, our definition of “secrecy” will differ from existing ones. Most previous work in the Dolev-Yao model postulate a definition of secrecy which only requires the adversary to be unable to reproduce the secret in totality. This stands in contrast to the standard definitional approach in cryptographic security, where it is typically required that a secret value be indistinguishable from random. Indeed, it is tempting to believe that, since in the symbolic model the security guarantees are “all or nothing” in flavor, the ability to reproduce a secret and the ability to distinguish it from random should be equivalent. However, it turns out that this is not the case: We show that the traditional criterion is insufficient for guaranteeing security of key exchange protocols, even if all the cryptography is “perfect”.

We will thus use for our symbolic criterion the approach which requires that the adversary be unable to distinguish the *real* secret key from a *random* one (from the same distribution) even when presented with both during the protocol.

We first demonstrate the weakness of the traditional symbolic criterion, via an example. Next, we formulate the new criterion.

7.2.1 The traditional symbolic criterion

Most previous attempts to formalize a symbolic security goal for key-exchange protocols have the same basic intuition: A protocol is a secure key exchange protocol if there is no run of the protocol between two honest and uncorrupted participants in which the adversary also learns the secret key. In our terminology, this could be stated as follows.

A symbolic key-exchange protocol \mathcal{P} is secure if there is no Dolev-Yao trace t valid for \mathcal{P} , where participant P_i produces an output message $\langle \text{finished} | P_i | P_{1-i} | R \rangle$, and where there is an adversary event containing the symbolic expression R .

That is, the trace may contain adversary events that have messages in which R appears as an (encrypted) element, but none of these messages will *be* the expression R .

We show that this criterion is insufficient for guaranteeing that key exchange protocol will remain secret in reasonable protocol environments. Specifically, we show an example of a protocol that satisfies the above symbolic criterion, but which is arguably insecure in any reasonable sense. The protocol is Version 2 of the Needham-Schroeder-Lowe protocol in Figure 3. (In this version, the key is the random string N_b chosen by the responder.) It has been shown that this protocol satisfies the above symbolic criterion and that N_b is not in the closure of any symbolic adversary [52]. However, assume that the initiator completes an exchange, and locally outputs N_b as its session key. Next, another protocol within the initiator uses N_b to send to the responder and encrypted message M . Furthermore, the encryption method is one-time-pad, and M can be one out of only two possible values (say, “buy” or “sell”). Then, the adversary obtains the ciphertext $C = M \oplus N_b$. If the adversary can distinguish N_b from a random nonce, then it can tell if M is “buy” or “sell.”

Note that the initiator outputs the key N_b before the responder has received its third message. The adversary knows that the plaintext of the third message is either $C = N_b \oplus \text{“buy”}$ or $C = N_b \oplus \text{“sell”}$. Consequently, the adversary can generate two possible third messages in the protocol, such that the value of M determines which of the two messages is a valid message (i.e., which of the two messages will cause the responder to successfully complete the protocol). Now, assume that the system is such that the adversary can tell whether the responder completed the protocol successfully. Then, the adversary can deliver one of these two messages to the responder, see whether it successfully completed the protocol, and thus determine the value of both M and N_b .

We note that the above attack scenario can be translated to an attack in the UC framework. Indeed, it is possible to show that this protocol (or, rather, its concrete counterpart) does not UC-realize $F_{2\text{KE}}$. Specifically, suppose the environment lets the protocol execute normally until the point where the initiator terminates and outputs N . The environment wishes to learn whether N is the session key of the protocol or a random key generated by $F_{2\text{KE}}$. It therefore follows the following strategy:

- It chooses a random bit.
- If it chose 1, it encrypts N in the responders public key and sends the ciphertext to the responder.

- If it chose 0, it chooses a random value N' , encrypts that in the responders public key, and sends that ciphertext to the responder.

If the protocol is being executed and N is the valid key, then the ciphertext sent to the responder will be valid exactly half the time — and the environment knows which half. If the protocol execution is being simulated, on the other hand, then both ciphertexts will be completely independent of the simulated execution. Therefore, no simulator will be able to guess the expected response with probability greater than $\frac{1}{2}$. Using this simple strategy, the environment will be able to distinguish between the real and ideal settings with high advantage.

We remark that a slight variant of the protocols results in a secure one: Simply let the session key be N_a rather than N_b . (This is Version 2 of the protocol, as discussed in Section 4.)

7.2.2 The new symbolic criterion

Our symbolic criterion essentially translates to the symbolic model the approach of “real or random security” that is typical in cryptographic notions of security. That is, we consider two worlds. In the first, *real*, world, the adversary is given the real (symbolic) session key as soon as one participant outputs it. In the other, *fake*, world, the adversary is given a symbolic key at the same point, but the symbol so provided is a new fresh symbol. The key is ‘secret’ if the two situations looks exactly the same to the adversary, no matter how the adversary behaves. That is, let an *adversary strategy* be the sequence of adversary deductions and transmissions made by an adversary in one execution. Then we wish to require that any adversary strategy will produce the same trace in the both scenarios. While this captures the desired intuition, there are two technical complications that must be considered:

1. Traces in the fake world will include a key symbol not found in the first world, and so direct equivalence will be impossible. What we require instead is that, for any adversary strategy, the trace produced in the real world and the trace produced in the fake world be the same when the fresh key is *ex post facto* renamed to the real session key.
2. Although equality after renaming is strong enough to imply security of a protocol in the UC framework, it is too strong to allow the converse. That is, a Dolev-Yao protocol that satisfies this definition will correspond to a concrete protocol that securely realizes F_{2KE} , but the opposite is not true. The reason for this is that the definition, as stated above, requires that the ‘real’ trace and ‘fake’ trace be exactly the same (after renaming). This prohibits the possibility that the two traces might differ, but only in a way that is unobservable by the adversary. (For example, the two traces might have different encrypted messages, as long as the encryption key is the same and is not known to the adversary.) Thus, our final definition requires only that the two traces be equivalent in their *observable* behavior. Fortunately, previous work by Abadi and Rogaway [4] (expanded upon by Herzog[33]) has already captured the observable part of a trace in their definition of a *pattern*. Thus, we will only require that the *patterns* of the ‘real’ trace and the ‘fake’ trace be the same.

We formalize this criterion as follows.

Definition 12 (Variable Renaming) *Let R_1, R_2 be random-strings symbols, and let t be an expression in the algebra \mathcal{A} . Then $t_{[R_1 \mapsto R_2]}$ is the expression where every instance of R_1 is replaced by R_2 .*

Definition 13 (Adversary Strategy) Let an adversary strategy be a sequence of adversary events that respect the Dolev-Yao assumptions. That is, a strategy Ψ is a sequence of instructions $I_1, I_2 \dots I_n$, where each I_i has one of the following forms, where i, j, k are integers:

- [“receive”, i]
- [“enc”, j, k, i]
- [“dec”, j, k, i]
- [“pair”, j, k, i]
- [“extract-l”, j, i]
- [“extract-r”, j, i]
- [“random”, i]
- [“name”, i]
- [“pubkey”, i]
- [“deliver”, j, P_i]

When executed against protocol \mathcal{P} , a strategy Ψ produces the following Dolev-Yao trace $\Psi(\mathcal{P})$. Go over the instructions in Ψ one by one, and:

- For each [“receive”, i] instruction, if this is the first activation of party P_i , or P_i was just activated with a delivered message m , then add to the trace a participant event (P'_j, L, \mathbf{m}) which is consistent with the protocol \mathcal{P} . Else output the trace \perp .
- For any other instruction, add the corresponding event to the trace, where the index i is replaced by m_i , the message expression in the i th event in the trace so far. (If adding the event results in an invalid trace then output the trace \perp .)

Definition 14 (Public-key pattern[4, 33]) Let $T \subseteq \mathcal{K}_{Pub}$ and $\mathbf{m} \in \mathcal{A}$. We recursively define the function $p(\mathbf{m}, T)$ to be:

- $p(K, T) = K$ if $K \in \mathcal{K}$
- $p(A, T) = A$ if $A \in \mathcal{M}$
- $p(N, T) = N$ if $N \in \mathcal{R}$
- $p(N_1|N_2, T) = p(N_1, T)|p(N_2, T)$
- $p(\{\mathbf{m}\}_K, T) = \begin{cases} \{\{p(\mathbf{m}, T)\}\}_K & \text{if } K \in T \\ \{\mathcal{T}\}_K & \text{(where } \mathcal{T} \text{ is the type tree of } \mathbf{m} \text{) otherwise} \end{cases}$

Then $pattern_{pk}(\mathbf{m}, T)$, the public-key pattern of an Dolev-Yao message \mathbf{m} relative to the set T , is

$$p(\mathbf{m}, \mathcal{K}_{Pub} \cap C[\{\mathbf{m}\} \cup T]).$$

If $t = H_1, H_2, \dots H_n$ is a Dolev-Yao trace where event H_i contains message \mathbf{m}_i then $pattern_{pk}(t, T)$ is exactly the same as t except that each \mathbf{m}_i is replaced by $p(\mathbf{m}_i, \mathcal{K}_{Pub} \cap C[S \cup T])$ where $S = \{\mathbf{m}_1, \mathbf{m}_2, \dots \mathbf{m}_n\}$. The base pattern of a message \mathbf{m} , denoted $pattern_{pk}(\mathbf{m})$, is defined to be $pattern_{pk}(\mathbf{m}, \emptyset)$, and $pattern_{pk}(t)$ is defined to be $pattern_{pk}(t, \emptyset)$.

Definition 15 (Symbolic Criterion for Key Exchange) A Dolev-Yao protocol \mathcal{P} provides Dolev-Yao two-party secure key exchange (DY-2SKE) if

1. (Agreement) For all P_0 and $P_1 \notin \mathcal{M}_{Adv}$ and Dolev-Yao traces valid for \mathcal{P} , if participant P_0 produces output message $\langle \text{Finished} | \mathbf{m}_0 \rangle$ and participant P_1 produces output message $\langle \text{finished} | \mathbf{m}_1 \rangle$, then $\mathbf{m}_0 = P_0 | P_1 | R$ and $\mathbf{m}_1 = P_1 | P_0 | R$ for some $R \in \mathcal{R}$.
2. (Real-or-random secrecy) Let \mathcal{P}_f be the protocol \mathcal{P} except that a fresh fake key R_f is released instead. Then for every adversary strategy Ψ ,

$$\text{pattern}_{pk}(\Psi(\mathcal{P})) = \text{pattern}_{pk}\left(\Psi(\mathcal{P}_f)_{[R_f \mapsto R_r]}\right)$$

Note that this criterion neither implies nor is implied by our symbolic mutual-authentication criterion. The MA criterion does not imply secrecy of any values. This criterion, on the other hand, allows successful termination of one party without any participation by the peer.

7.3 Soundness and completeness of the symbolic criterion

Theorem 4 Let \mathfrak{p} be a simple protocol. Then \mathfrak{p} securely realizes F_{2KE} . if and only if $\hat{\mathfrak{p}}$ achieves Dolev-Yao secure key-exchange.

Proof. Suppose that $\hat{\mathfrak{p}}$ does not satisfy Dolev-Yao key exchange. Then one of two possibilities occur, each of which allows the environment to distinguish the ideal (F_{2KE}) model from the real (F_{CPKE} -hybrid) model. Either:

- There exists a DY trace where participants P_i and P_{1-i} produce outputs $\langle \text{finished} | P_i | P'_{1-i} | R \rangle$ and $\langle \text{finished} | P_{1-i} | P_i | R' \rangle$, respectively, where either $R' \neq R$ or $P'_i \neq P_i$ or $P'_{1-i} \neq P_{1-i}$. The strategy of this trace is easily mapped to an environment that simply performs the same sequence of calculations, receptions, and transmissions. (Recall that this sequence has some constant, finite length.) Thus, some environment can produce the same behavior in the real (F_{CPKE} -hybrid) model. However, this behavior will never arise in the F_{2KE} case as the functionality will always distribute the same κ to both dummy parties. Thus, the environment can easily tell whether it is in the F_{2KE} or F_{CPKE} -hybrid model.
- Or, it might be the case that there exists an adversary strategy Ψ so that

$$\text{pattern}_{pk}(\Psi(\mathcal{P})) \neq \text{pattern}_{pk}\left(\Psi(\mathcal{P}_f)_{[R_f \mapsto R_r]}\right).$$

Also here, there exists an environment that distinguishes between an execution of the concrete protocol and the ideal protocol for F_{2KE} : the environment simply performs the finite sequence of calculations, receptions, and transmissions that is described in the strategy Ψ . It then translates its view to a symbolic trace (using the transformation in Definition 8), and checks whether the pattern of the symbolic trace equals $\text{pattern}_{pk}(\Psi(\mathcal{P}))$. If the patterns are equal then the environment outputs “real”. Otherwise, it outputs “ideal”.

To show that this environment is a good distinguisher between the real and the ideal cases, we observe that: **(a)**. If the environment interacts with the protocol and the adversary in the F_{CPKE} -hybrid model, then we are guaranteed that the patterns will be equal. **(b)**. In contrast, if the environment interacts with the ideal protocol for F_{2KE} , then the key output by the participants is independent from the simulator’s view, thus it will be translated to a fresh key in the derived DY trace. Thus here the pattern of this trace cannot be equal to $\text{pattern}_{pk}(\Psi(\mathcal{P}))$.

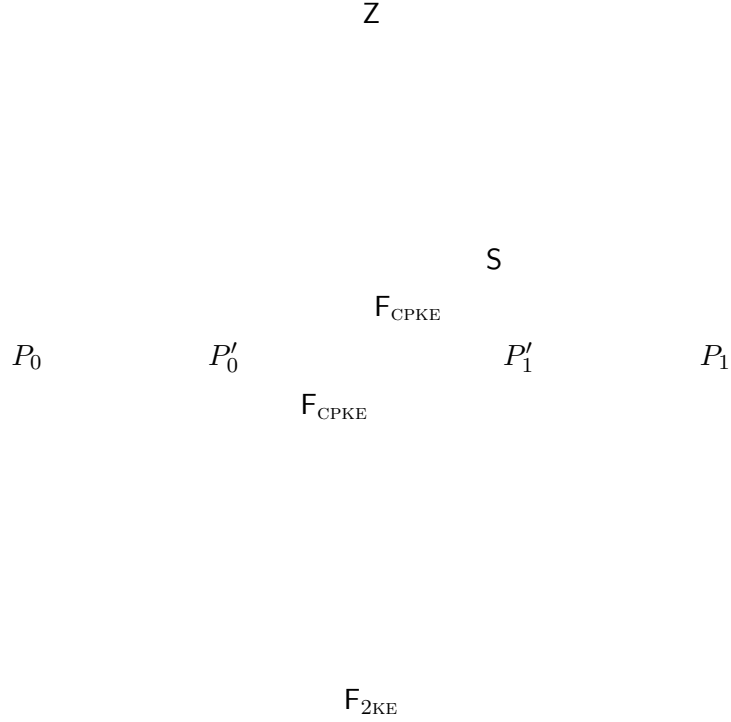


Figure 8: The simulator for $F_{2\text{KE}}$

Thus, if the Dolev-Yao protocol $\hat{\mathbf{p}}$ does not satisfy Dolev-Yao key exchange, then there exists an environment which can distinguish the $F_{2\text{KE}}$ model from the F_{CPKE} -hybrid model, and the protocol does not securely realize $F_{2\text{KE}}$.

To show the other direction, we need to provide a simulator such that no environment can tell whether it is interacting with the ideal protocol for $F_{2\text{KE}}$ or with the concrete protocol. The simulator proceeds as follows. (The simulator is similar to the simulator used in the proof of soundness of the symbolic mutual authentication criterion.)

- S internally simulates the participants P_0 , P_1 , and two instances of F_{CPKE} . At the beginning, none of these simulated participants are running.
- When the simulator receives a message $(P_i, P_{1-i}, \text{RID})$ from the functionality $F_{2\text{KE}}$ (indicating that the external dummy participant P_i has received input from the environment and has passed it on the $F_{2\text{KE}}$) it activates the simulated participant P_i with input

$$(\text{SID}, P_j, \text{RID}), \text{SID}_i, (P_{1-i}, \text{RID}_i, \text{SID}_i)$$

where SID is some arbitrary SID and SID_j is the SID of the appropriate simulated instance of F_{CPKE} .

- When the simulator receives input for the dummy adversary to send to participant P_i , it sends that input to the copy P_i that it is simulating internally.
- Likewise, when simulated participant P_i produces a message to send on its communication tape, the simulator sends this message to the environment.

- When simulated participant P_i produces a message to send an instance of F_{CPKE} , the simulator forwards this to the appropriate instance of F_{CPKE} that it is simulating.
- Likewise, when either instance of F_{CPKE} requires communication with the adversary, this communication is provided to the simulated dummy adversary (and sent directly to the environment). Communication from the environment to either F_{CPKE} -instance is forwarded directly to the simulated F_{CPKE} -instance.
- When the simulated participant P_i produces local output $\langle \text{"finished"}, P_i || P_{1-i} || v \rangle$, the simulator sends the message $(\text{Session-key}, \text{sid}, P_i, v)$ to the functionality $F_{2\text{KE}}$.

It remains to show that the simulation is valid. Consider an environment Z that distinguishes with advantage $\epsilon(k)$ between an interaction with the above simulator and the ideal protocol for $F_{2\text{KE}}$, and an interaction with the concrete protocol. We show that $\epsilon(k)$ is negligible in k . To see this, fix some k . Assume for simplicity that Z is deterministic. (No generality is lost here since Z gets arbitrary, non-uniform input.) By the mapping lemma (Lemma 2) we know that, except for probability $\epsilon'(k)$ which is negligible in k , traces of the interaction between Z and the concrete protocol are mapped to a single symbolic trace $\hat{\mathbf{t}}$ of the corresponding symbolic protocol, \mathcal{P} . Let Ψ be the adversary strategy derived from $\hat{\mathbf{t}}$, i.e. the strategy Ψ such that $\Psi(\mathcal{P}) = \hat{\mathbf{t}}$. Then, it can be verified that:

1. The trace of the interaction between Z , the above simulator, and the ideal protocol for $F_{2\text{KE}}$, is mapped to the symbolic trace $\Psi(\mathcal{P}_f)$. (This is so since in this interaction the view of Z is identical to its view of an interaction with the concrete protocol, except that the keys in the local outputs of the parties are random and independent from the rest of the view.)
2. Let \mathbf{t} and \mathbf{t}' denote the random variables describing the trace of the interaction of Z with the concrete protocol and the ideal protocol, respectively. Then, conditioned on the event that $\text{pattern}_{pk}(\hat{\mathbf{t}}) = \text{pattern}_{pk}(\hat{\mathbf{t}}'_{[R_f \mapsto R_r]})$, the distributions of \mathbf{t} and \mathbf{t}' are identical. (This is so since F_{CPKE} makes sure that Z 's view is statistically independent of the values encrypted by the honest parties, thus Z 's view is independent of whatever is not in the pattern of $\hat{\mathbf{t}}$.)

We conclude that, whenever the trace \mathbf{t} of Z 's interaction with the protocol is mapped to a valid trace $\hat{\mathbf{t}}$ of the corresponding symbolic protocol, then Z 's view of the interaction is distributed identically to its view of the interaction with the ideal protocol for $F_{2\text{KE}}$. Consequently, Z 's distinguishing probability $\epsilon(k)$ is at most $\epsilon'(k)$, the error probability in the mapping lemma. \square

8 Future research

This work demonstrates that completely symbolic analysis of security properties within a simulation-based, compositional cryptographic framework is possible. Furthermore, the chosen symbolic framework is one that is very close to the language of known automated verification tools. As such, it opens the door to a number of questions and challenges. A first challenge is to build tools that will automatically verify whether a given abstract protocol satisfies the symbolic criteria we provide. Some indications that this is feasible are (a) the Athena tool [51], which was built explicitly to verify criteria such as our symbolic mutual authentication criterion in a model that is very similar to ours, and (b) the tool of Blanchet [13], that is able to verify requirements that are closely related to the secrecy requirement of our symbolic key-exchange criterion.

A second research direction is to generalize our results to a richer and less restrictive “programming language” for protocols. One direction is to enlarge the set of allowed operations and to incorporate other cryptographic primitives, while retaining the ability to analyze only a single session of the protocol in question. Natural candidates include the Diffie-Hellman exchange, signatures schemes, pseudo-random functions, and message authentication codes. Other generalizations include adaptive security (i.e. security against adversaries that corrupt parties throughout the computation), and protocols where even their symbolic counterparts are randomized.

A third direction is to apply a similar analytical methodology to other cryptographic tasks, and even tasks that were never before addressed using formal tools. For instance, it may be possible to come up with a symbolic representation of, say, two-party protocols that use commitment schemes, and provide a symbolic criterion for when such protocols are zero-knowledge protocols (e.g., satisfy the ideal zero-knowledge functionality). Similarly, one can potentially come up with symbolic criteria as to when a protocol securely realizes an arbitrary given ideal functionality.

References

- [1] *Proceedings of the 12th IEEE Computer Security Foundations Workshop (CSFW 12)*. IEEE Computer Society, June 1999.
- [2] Martín Abadi and Bruno Blanchet. Analyzing security protocols with secrecy types and logic programs. In *Conference Record of POPL 2002: The 29th SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 33–44, January 2002.
- [3] Martín Abadi and Andrew Gordon. A calculus for cryptographic protocols: the spi calculus. *Information and Computation*, 148(1):1–70, 1999.
- [4] Martín Abadi and Phillip Rogaway. Reconciling two views of cryptography (the computational soundness of formal encryption). *Journal of Cryptology*, 15(2):103–127, 2002.
- [5] M. Backes, B. Pfitzmann, and M. Waidner. A composable cryptographic library with nested operations (extended abstract). In *Proceedings, 10th ACM conference on computer and communications security (CCS)*, October 2003. Full version available at <http://eprint.iacr.org/2003/015/>.
- [6] M. Backes, B. Pfitzmann, and M. Waidner. Low-level ideal signatures and general integrity idealization. In *Proceedings, 7th Information Security Conference (ISC)*, number 3225 in Lecture Notes in Computer Science, pages 39–51. Springer-Verlag, 2004.
- [7] Michael Backes and Birgit Pfitzmann. A cryptographically sound security proof of the Needham-Schroeder-Lowe public-key protocol. In *Proceedings of the 23rd Conference on Foundations of Software Technology and Theoretical Computer Science – FSTTCS*, volume 2914 of *Lecture Notes in Computer Science*, pages 140–152. Springer-Verlag, December 2003.
- [8] Michael Backes and Birgit Pfitzmann. A cryptographically sound security proof of the Needham-Schroeder-Lowe public-key protocol. *IEEE Journal on Selected Areas in Communications*, 2004. To appear.
- [9] Michael Backes and Birgit Pfitzmann. Relating symbolic and cryptographic secrecy. *Cryptology ePrint Archive*, Report 2004/300, November 2004. <http://eprint.iacr.org/>.

- [10] Michael Backes, Birgit Pfitzmann, and Michael Waidner. A general composition theorem for secure reactive systems. In *1st Theory of Cryptography Conference (TCC)*, volume 2951 of *Lecture Notes in Computer Science*, pages 336–354. Springer-Verlag, February 2004.
- [11] Donald Beaver. Secure multiparty protocols and zero-knowledge proof systems tolerating a faulty minority. *Journal of Cryptology*, 4(2):75–122, 1991.
- [12] Mihir Bellare and Phillip Rogaway. Entity authentication and key distribution. In D. Stinson, editor, *Advances in Cryptology - CRYPTO 1993*, volume 773 of *Lecture Notes in Computer Science*, pages 232–249. Springer-Verlag, August 1993. Full version of paper available at <http://www-cse.ucsd.edu/users/mihir/>.
- [13] Bruno Blanchet. Automatic proof of strong secrecy for security protocols. In *Proceedings, 2004 IEEE Symposium on Security and Privacy*, pages 86–102, 2004.
- [14] Manuel Blum and Silvio Micali. How to generate cryptographically strong sequences of pseudo random bits. In *Proceedings, 22th Annual Symposium on Foundations of Computer Science (FOCS 1982)*, pages 112–117, 1982.
- [15] Manuel Blum and Silvio Micali. How to generate cryptographically strong sequences of pseudo-random bits. *SIAM Journal on Computing*, 13(4):850–864, 1984.
- [16] Dan Boneh, editor. *Advances in Cryptology - CRYPTO 2003*, volume 2729 of *Lecture Notes in Computer Science*. Springer-Verlag, August 2003.
- [17] Michael Burrows, Martín Abadi, and Roger Needham. A logic of authentication. *ACM Transactions in Computer Systems*, 8(1):18–36, February 1990.
- [18] Ran Canetti. Security and composition of multiparty cryptographic protocols. *Journal of Cryptology*, 13(1):143–202, 2000.
- [19] Ran Canetti. Universal composable security: A new paradigm for cryptographic protocols. In *42nd Annual Symposium on Foundations of Computer Science (FOCS 2001)*, pages 136–145. IEEE Computer Society, October 2001.
- [20] Ran Canetti. Universally composable signatures, certification, and authentication. Cryptology ePrint Archive, <http://eprint.iacr.org/2003/239>, 2003.
- [21] Ran Canetti. Universally composable signature, certification, and authentication. In *Proceedings of the 17th IEEE Computer Security Foundations Workshop (CSFW 16)*, pages 219–233. IEEE Computer Society, June 2004.
- [22] Ran Canetti and Hugo Krawczyk. Analysis of key-exchange protocols and their use for building secure channels. In Birgit Pfitzmann, editor, *Advances in Cryptology - Eurocrypt 2001*, volume 2045 of *Lecture Notes in Computer Science*, pages 453–474. Springer-Verlag, May 2001.
- [23] Ran Canetti, Hugo Krawczyk, and Jesper Buus Nielsen. Relaxing chosen-ciphertext security. In Boneh [16], pages 565–582.
- [24] Ran Canetti and Tal Rabin. Universal composition with joint state. In Boneh [16], pages 265–281.

- [25] I. Cervesato, N. A. Durgin, P. D. Lincoln, J. C. Mitchell, and A. Scedrov. A meta-notion for protocol analysis. In *Proceedings of the 12th IEEE Computer Security Foundations Workshop (CSFW 12)* [1].
- [26] D. Dolev, C. Dwork, and M. Naor. Non-malleable cryptography. *SIAM Journal of Computing*, 30(2):391–437, 2000.
- [27] D. Dolev and A. Yao. On the security of public-key protocols. *IEEE Transactions on Information Theory*, 29:198–208, 1983.
- [28] Oded Goldreich. *Foundations of Cryptography—Basic Tools*. Cambridge University Press, 2001.
- [29] S. Goldwasser and S. Micali. Probabilistic encryption. *Journal of Computer and System Sciences*, 28(2):270–299, 1984.
- [30] Shafi Goldwasser and Leonid Levin. Fair computation of general functions in presence of immoral majority. In Alfred Menezes and Scott A. Vanstone, editors, *CRYPTO*, volume 537 of *Lecture Notes in Computer Science*, pages 77–93. Springer, August 1990.
- [31] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 18(1):186–208, 1989.
- [32] Shafi Goldwasser, Silvio Micali, and Ronald L. Rivest. A digital-signature scheme secure against adaptive chosen-message attacks. *SIAM J. Computing*, 17(2):281–308, April 1988.
- [33] Jonathan Herzog. *Computational Soundness for Standard Assumptions of Formal Cryptography*. PhD thesis, Massachusetts Institute of Technology, May 2004.
- [34] Jonathan Herzog, Moses Liskov, and Silvio Micali. Plaintext awareness via key registration. In Boneh [16], pages 548–564.
- [35] P. D. Lincoln, J. C. Mitchell, M. Mitchell, and A. Scedrov. A probabilistic poly-time framework for protocol analysis. In *Proceedings of the 5th ACM Conference on Computer and Communication Security (CCS '98)*, pages 112–121, November 1998.
- [36] P. D. Lincoln, J. C. Mitchell, M. Mitchell, and A. Scedrov. Probabilistic polynomial-time equivalence and security protocols. In Jeannette M. Wing, Jim Woodcock, and Jim Davies, editors, *World Congress on Formal Methods*, volume 1708 of *Lecture Notes in Computer Science*, pages 776–793. Springer, September 1999.
- [37] Gavin Lowe. An attack on the Needham–Schroeder public-key authentication protocol. *Information Processing Letters*, 56:131–133, 1995.
- [38] Gavin Lowe. Breaking and fixing the Needham–Schroeder public-key protocol using FDR. In Margaria and Steffen, editors, *Tools and Algorithms for the Construction and Analysis of Systems*, volume 1055 of *Lecture Notes in Computer Science*, pages 147–166. Springer–Verlag, 1996.
- [39] Nancy Lynch. I/O automaton models and proofs for shared-key communication systems. In *Proceedings of the 12th IEEE Computer Security Foundations Workshop (CSFW 12)* [1], pages 14–29.

- [40] P. Maggi and R. Sisto. Using SPIN to verify security protocols. In *Proceedings of the 9th International SPIN Workshop on Model Checking of Software*, number 2318 in Lecture Notes in Computer Science, pages 187–204, 2002.
- [41] Catherine Meadows. Applying formal methods to the analysis of a key management protocol. *The Journal of Computer Security*, 1(1), January 1992.
- [42] Silvio Micali, Charles Rackoff, and Bob Sloan. The notion of security for probabilistic cryptosystems. *SIAM Journal on Computing*, 17(2):412–426, April 1988.
- [43] Silvio Micali and Phillip Rogaway. Secure computation (abstract). In Joan Feigenbaum, editor, *CRYPTO*, volume 576 of *Lecture Notes in Computer Science*, pages 392–404. Springer, August 1991.
- [44] Daniele Micciancio and Bogdan Warinschi. Soundness of formal encryption in the presence of active adversaries. In *Proceedings, Theory of Cryptography Conference*, number 2951 in Lecture Notes in Computer Science, pages 133–151. Springer, February 2004.
- [45] John C. Mitchell, Mark Mitchell, and Ulrich Stern. Automated analysis of cryptographic protocols using Mur φ . In *Proceedings, 1997 IEEE Symposium on Security and Privacy*, pages 141–153. IEEE, Computer Society Press of the IEEE, 1997.
- [46] Roger Needham and Michael Schroeder. Using encryption for authentication in large networks of computers. *Communications of the ACM*, 21(12):993–999, 1978.
- [47] Birgit Pfitzmann, Matthias Schunter, and Michael Waidner. Cryptographic security of reactive systems. *Electronic Notes in Theoretical Computer Science*, 32, 2000.
- [48] Birgit Pfitzmann and Michael Waidner. Composition and integrity preservation of secure reactive systems. In *Proceedings of the 7th ACM Conference on Computer and Communication Security (CCS 2000)*, pages 245–254. ACM Press, November 2000.
- [49] C. Rackoff and D. Simon. Noninteractive zero-knowledge proof of knowledge and the chosen-ciphertext attack. In *Advances in Cryptology—CRYPTO 91*, number 576 in Lecture Notes in Computer Science, pages 433–444, 1991.
- [50] Victor Shoup. On formal models for secure key exchange. IBM research report, IBM Zurich Research Lab, November 1999.
- [51] D. Song. Athena, an automatic checker for security protocol analysis. In *Proceedings of the 12th IEEE Computer Security Foundations Workshop (CSFW 12)* [1], pages 192–202.
- [52] F. Javier THAYER Fábrega, Jonathan C. Herzog, and Joshua D. Guttman. Strand spaces: Proving security protocols correct. *Journal of Computer Security*, 7(2/3):191–230, 1999.

A Realizing F_{CPKE}

We show how to realize F_{CPKE} given the plain, unauthenticated public-key functionality F_{PKE} , and a registration functionality F_{REG} . Functionality is given in Figure 9. (Our formulation of F_{PKE} is the same as functionality F_{PKE} from [23], except that here unregistered ciphertexts are decrypted using a decryption algorithm that was provided by the adversary ahead of time. See more discussion in

Functionality F_{PKE}

F_{PKE} proceeds as follows, when parameterized by message domain ensemble $D = \{D_k\}_{k \in \mathcal{N}}$.

Key Generation: Upon receiving a value $(\text{KeyGen}, \text{SID})$ from some party P , verify that $\text{SID} = (P, \text{SID}')$ for some SID' . If not then ignore the input. Else:

1. Hand $(\text{KeyGen}, \text{SID})$ to the adversary.
2. Receive a value e and a description of an algorithm D from the adversary.
3. record (e, D) , and output e to P .

Encryption: Upon receiving from any party P' a value $(\text{Encrypt}, \text{SID}, e', m)$ proceed as follows:

1. If $m \notin D_k$ then return an error message to P' .
2. If $m \in D_k$ then hand $(\text{Encrypt}, \text{SID}, e', P')$ to the adversary. (If $e' \neq e$ or e is not yet defined then hand also the entire value m to the adversary.)
3. Receive a “ciphertext” c from the adversary, record the pair (c, m) , and hand c to P' . (If the tag c already appears in a previously recorded pair then return an error message to P' .)

Decryption: Upon receiving a value $(\text{Decrypt}, \text{SID}, c)$ from P (and P only), proceed as follows:

1. If there is a recorded pair (c, m) then hand m to P .
2. Otherwise, compute $m = D(c)$, record the pair (m, c) , and hand m to P .

Figure 9: The public-key encryption functionality, F_{PKE}

Section 3.) The treatment here closely mimics the the treatment in [20] of realizing certification given signature schemes and a registration service.

The primary difference between F_{CPKE} and F_{PKE} is that in F_{CPKE} there are no encryption keys; instead messages are encrypted directly to the identity of the recipient. (In other words, F_{CPKE} provides ideal binding between a public key and its “owner.”) In contrast, F_{PKE} does not provide any binding between the public key and the identity of the intended decryptor. In particular, there is no security guarantee regarding messages that were encrypted with public keys other than the key given to the legitimate decryptor.

Recall that F_{PKE} can be realized in a simple way (with respect to non-adaptive party corruptions) given any CCA-secure encryption scheme [19, 23]. Here we show how to realize F_{CPKE} given F_{PKE} , so long as one has access to an additional, simple “service” that ties public values to principals. Such a ‘registration’ functionality, F_{REG} , is given in Figure 10. Our protocol, P_{RENC} , for realizing F_{CPKE} given ideal access to both F_{PKE} and F_{REG} , is given in Figure 11. We show:

Claim 5 *Protocol P_{RENC} securely realizes F_{CPKE} , when given ideal access to F_{PKE} and F_{REG} .*

Proof. The proof proceeds along the lines of the proof in [20] for the case of constructing a certification service from signature schemes and a registration service. To summarize, let A be an adversary that interacts with parties running P_{RENC} in the $(F_{\text{PKE}}, F_{\text{REG}})$ -hybrid model. Then we construct a simulator S such that no environment Z can distinguish between an interaction with the simulator S and F_{CPKE} , and an interaction with the adversary A and the protocol P_{RENC} . As usual, the simulator contains within it a copy of A , the parties, and for each party it simulates a

Functionality F_{REG}

F_{REG} proceeds as follows:

- Upon receiving the first message $(\text{Register}, \text{SID}, v)$ from party P , send $(\text{Registering}, \text{SID}, v)$ to the adversary. Upon receiving (ok) from the adversary, and if this is the first message from P , then record the pair (P, v) .
- Upon receiving the message $(\text{Retrieve}, \text{PID})$ from a party P' , send message $(\text{Retrieve}, \text{PID})$ to the adversary and wait for the adversary to return a message ok . Then, if there is a recorded pair (PID, v) output $(\text{Retrieve}, \text{PID}, v)$ to P' . Otherwise, if there is no recorded pair, return $(\text{Retrieve}, \text{PID}, \perp)$.

Figure 10: The registration functionality, F_{REG}

Protocol P_{RENC}

Initialization: Party $\text{PID}_{\text{owner}}$ sends message $(\text{KeyGen}, (\text{PID}_{\text{owner}}, \text{SID}_{\text{owner}}))$ to F_{PKE} , gets message e . Next, $\text{PID}_{\text{owner}}$ sends $(\text{Register}, (\text{PID}_{\text{owner}}, \text{SID}_{\text{owner}}), e)$ to F_{REG} .

Encryption: When activated with message $(\text{Encrypt}, (\text{PID}, \text{SID}), m)$, party P does:

- First, P sends $(\text{Retrieve}, \text{SID})$ to F_{REG} and waits. It receives a message $(\text{Retrieve}, \text{SID}, v)$ checks that $v \neq \perp$. If it is, return \perp .
- Otherwise, send $(\text{Encrypt}, (\text{PID}, \text{SID}), v, m)$ to F_{PKE} . Receive value c from F_{PKE} and output it. Also, record the pair (m, c) .

Decryption: Upon being activated with message $(\text{Decrypt}, \text{SID}, c)$ where $\text{SID} = (P, s)$, party P sends $(\text{Decrypt}, \text{SID}, c)$ to F_{PKE} . It waits for and outputs the response.

Figure 11: The registered encryption protocol, P_{RENC}

copy of F_{PKE} and F_{REG} . All messages from Z to A and back are forwarded. In addition, S acts as follows:

Initialization The simulator, upon initialization, simulates the initialization process. That is, the simulator simply forwards the messages between the parties, the F_{PKE} and F_{REG} functionalities, and the adversary. By the end of initialization, the adversary has created (via the simulated functionalities F_{PKE}) a “public key” e , and this party is registered with F_{REG} . This entire process is visible to the environment, who can see all messages received or sent by the simulated adversary A .

Encryption. Upon receiving message $(\text{Encrypt}, \text{SID}, P)$ from F_{CPKE} , the simulator has the simulated party P retrieve the public key of party SID . Let e be the public key. The simulator then passes message $(\text{Encrypt}, \text{SID}, e, m)$ to the appropriate copy of F_{PKE} . When the F_{PKE} functionality returns a ciphertext c , this is given to the F_{CPKE} functionality.

It is simple to verify that this simulator is a perfect simulation of the adversary A in the $(F_{\text{PKE}}, F_{\text{REG}})$ -hybrid model. The initialization phase is exactly the same as that in the hybrid model, and one can easily confirm that the simulated adversary A sees exactly the same sequence of messages as it would in the hybrid model. This, protocol P_{RENC} securely realizes F_{CPKE} . □

B Index of notation

B.1 Standard mathematical/cryptographic notation

D	A domain of messages
\vec{r}	Participant randomness
$x \xleftarrow{R} X$	x is drawn randomly from distribution X
$X \approx Y$	Distributions X and Y are computationally indistinguishable
κ	A key
σ	A bitstring
k	The security parameter
m	A message
c	A ciphertext

B.2 The UC framework

SID	A session identifier
RID	A role identifier
PID	A participant identifier
p	A concrete protocol (i.e., one in the UC framework)
Z	The environment
z	Environment input
P	UC participant
F	Ideal functionality
S	The simulator
P	A higher-level protocol
F_{PKE}	The public-key encryption functionality
F_{CPKE}	The certified public-key functionality
$\text{PID}_{\text{owner}},$ $\text{SID}_{\text{owner}},$ $\text{RID}_{\text{owner}}$	PID, SID and RID, respectively, of a F_{CPKE} instance's legitimate owner
D	A computational decryption algorithm
$F_{2\text{MA}}$	The 2-party mutual-authentication functionality
$F_{2\text{KE}}$	The 2-party key-exchange functionality
D	A computational decryption algorithm
$(\text{Encrypt}, \text{SID}, m)$	A call to F_{CPKE} to encrypt m
$(\text{Decrypt}, \text{SID}, c)$	A call to F_{CPKE} to decrypt c

B.3 The Dolev-Yao Model

\mathcal{A}	The formal algebra of terms
$m, m_1, m_2 \dots$	Dolev-Yao messages; elements of \mathcal{A}
\mathcal{M}	The set of names
A, B	Names; elements of \mathcal{M}
\mathcal{R}	The set of (symbolic) random strings
$R, R_1, R_2 \dots$	Nonces; elements of \mathcal{R}
\mathcal{K}_{Pub}	The set of public keys
K	A symbolic key; element of \mathcal{K}_{Pub}
K_A	Public key of A
\mathcal{G}	Symbolic garbage term
\perp	Symbol for error or failure
<i>Starting, Finished</i>	Symbols for protocol beginning and end
$\{m\}_K$	Encryption of m under key K
$m_1 m_2$	Pair of messages m_1 and m_2
<i>keyof</i>	Mapping from names to public keys
\mathcal{R}_{Adv}	Random strings chosen by the adversary
\mathcal{M}_{Adv}	Aliases of the adversary
\mathcal{K}_{Adv}	Private keys associated with adversary aliases
\mathcal{P}	Symbolic protocol
\mathcal{P}_i	Transition function for role i in symbolic protocol
\mathcal{S}	The set of states for the transition function \mathcal{P}
S	State of symbolic participant
<i>output</i>	Constant to represent local output
<i>message</i>	Constant to represent communication
$C[S]$	Closure of set S ; all messages symbolic adversary can make from S
$pattern_{pk}(m)$	Base pattern of a symbolic message

B.4 Executions and Traces

t	A symbolic trace
H	Event on a symbolic trace
$[\text{“enc”}, j, k, M_i]$	Symbolic adversary encryption event
$[\text{“dec”}, j, k, M_i]$	Symbolic adversary decryption event
$[\text{“extract-l”}, j, M_i]$	Symbolic adversary separation (left) event
$[\text{“extract-r”}, j, M_i]$	Symbolic adversary separation (right) event
$[\text{“nonce”}, M_i]$	Symbolic adversary creation of random string
$[\text{“name”}, M_i]$	Symbolic adversary name creation
$[\text{“pubkey”}, M_i]$	Symbolic adversary public key creation
$[\text{“privkey”}, M_i]$	Symbolic adversary private key creation
$[\text{“deliver”}, j, P_i]$	Symbolic adversary transmission
(P'_i, L_i, m_i, S_i)	Symbolic participant event
$[\text{“fail”}, m_i]$	Special symbolic event to indicate mapping failure
Ψ	Adversary strategy
t	Trace of concrete execution
G	Event on a concrete trace
\hat{t}	Symbolic interpretation of concrete trace t
$\langle \text{“output”}, \text{PID}, l \rangle$	Concrete local output event
$\langle \text{“message”}, \text{PID}, m \rangle$	Concrete participant communication event
$\langle \text{“adv”}, \text{PID}, m \rangle$	Concrete adversary transmission event
$t_{[M \mapsto N]}$	Symbolic trace t with every instance of message M replaced with message N

B.5 Simple protocols

M_1	Machine for one role of simple protocol
Σ	The store for a simple protocol
Π	The program for a simple protocol
$\text{SID}, \text{PID}_0, \text{RID}_i,$ PID_1	Program’s SID, PID and RID and peer’s PID
v	Variable of a simple protocol’s program
σ	bit-string
p	A simple protocol
\hat{p}	symbolic interpretation of protocol p