

Symmetric Encryption in a Simulatable Dolev-Yao Style Cryptographic Library*

Michael Backes

IBM Zurich Research Laboratory
mbc@zurich.ibm.com

Birgit Pfizmann

IBM Zurich Research Laboratory
bpf@zurich.ibm.com

Abstract

Recently we showed how to justify a Dolev-Yao type model of cryptography as used in virtually all automated protocol provers under active attacks and in arbitrary protocol environments. The justification was done by defining an ideal system handling Dolev-Yao-style terms and a cryptographic realization with the same user interface, and by showing that the realization is as secure as the ideal system in the sense of reactive simulatability. This definition encompasses arbitrary active attacks and enjoys general composition and property-preservation properties. Security holds in the standard model of cryptography and under standard assumptions of adaptively secure primitives.

A major primitive missing in that library so far is symmetric encryption. We show why symmetric encryption is harder to idealize in a way that allows general composition than existing primitives in this library. We discuss several approaches to overcome these problems. For our favorite approach we provide a detailed provably secure idealization of symmetric encryption within the given framework for constructing nested terms.

1. Introduction

Automated proofs of security protocols with model checkers or theorem provers typically abstract from cryptography by deterministic operations on abstract terms and by simple cancellation rules. An example term is $E_{pk_{e_w}}(E_{pk_{e_v}}(\text{sign}_{sk_{s_u}}(m, N_1), N_2))$, where m denotes an application message and N_1, N_2 two nonces. A typical cancellation rule is $D_{sk_e}(E_{pk_e}(m)) = m$ for corresponding keys. The proof tools handle these terms symbolically, i.e., they never evaluate them to bitstrings. In other words, they perform abstract algebraic manipulations on trees consisting of operators and base messages, using the cancellation

rules, the transition rules of a particular protocol, and abstract models of networks and adversaries. Such abstractions, although different in details, are called the Dolev-Yao model after the first authors [19].

For many years there was no cryptographic justification for such abstractions. The problem lies in the assumption, implicit in the adversary model, that actions that cannot be expressed with the abstract operations are impossible, and that no relations hold between terms unless derivable by the cancellation rules. It is not hard to make artificial counterexamples to these assumptions. Nevertheless, no counterexamples against the method for protocols proved in the literature were found so far. Further, the overall approach of abstracting from cryptographic primitives once with rigorous hand-proofs, and then using tools for proving protocols using such primitives, is highly attractive: Besides the cryptographic aspects, protocol proofs have many distributed-systems aspects, which make proofs tedious and error-prone even if they weren't interlinked with the cryptographic aspects. To use existing efficient automated proof tools for security protocols, cryptography must indeed be abstracted into simple, deterministic ideal systems. The closer one can stay to the Dolev-Yao model, the easier the adaptation of the proof tools will be.¹

Cryptographic underpinnings of a Dolev-Yao model were first addressed by Abadi and Rogaway in [2]. However, they only handled passive adversaries and symmetric encryption. The protocol language and security properties handled were extended in [1, 24], but still only for passive adversaries. This excludes most of the typical ways of attacking protocols, e.g., man-in-the-middle attacks and

* An extended version of this work is available as IACR Cryptology ePrint Archive Report 2004/059, <http://eprint.iacr.org/>.

In Proceedings of the 16th IEEE Computer Security Foundations Workshop, Pacific Grove, CA, USA, pages 204–218, June 2004. IEEE Computer Society.

© 2004 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

¹ Efforts are also under way to formulate syntactic calculi for dealing with probabilism and polynomial-time considerations, in particular [28, 25, 29, 22] and, as a second step, to encode them into proof tools. However, this approach can not yet handle protocols with any degree of automation. Generally it is complementary to, rather than competing with, the approach of proving simple deterministic abstractions of cryptography and working with those wherever cryptography is only used in a blackbox way.

attacks by reusing a message part in a different place or a concurrent protocol run. A full cryptographic justification for a Dolev-Yao model, i.e., for arbitrary active attacks and within arbitrary surrounding interactive protocols, was first given recently in [5]. Based on the specific Dolev-Yao model whose soundness was proven in [5], the well-known Needham-Schroeder-Lowe protocol was proved in [3]. This shows that in spite of adding certain operators and rules compared with simpler Dolev-Yao models (in order to be able to use arbitrary cryptographically secure primitives without too many changes in the cryptographic realization), such a proof is possible in the style already used in automated tools, only now with a sound cryptographic basis. Another cryptographically sound proof of this protocol was concurrently developed by Warinschi [32]. The proof establishes a stronger security property but is conducted from scratch in the cryptographic approach which takes it out of the scope of formal proof tools. In [6] it was shown how the library, in other words the term algebra and rules, can be modularly extended by additional cryptographic primitives, using the example of symmetric authentication [6].

Nevertheless, symmetric encryption is still missing in this framework, while it is the most common cryptographic primitive in typical proofs with Dolev-Yao models. The goal of this paper is to add symmetric encryption to this framework. Concurrently to our work, Laud [23] has presented a cryptographic underpinning for a Dolev-Yao model of symmetric encryption under active attacks. His work enjoys a direct connection with a formal proof tool, but it is specific to certain confidentiality properties, restricts the surrounding protocols to straight-line programs in a specific language, and does not address a connection to the remaining primitives of the Dolev-Yao model. Herzog et al. [20, 21] and Micciancio and Warinschi [26] have recently also given a cryptographic underpinning under active attacks. Their results are considerably weaker than the one in [5] since they are specific for public-key encryption; moreover, the former relies on a stronger assumption whereas the latter severely restricts the classes of protocols and protocol properties that can be analyzed using this primitive. Section 6 of [26] further points out several possible extensions of their work which all already exist in the earlier work of [5].

There are intrinsic difficulties in providing a sound abstraction from symmetric encryption in the strong sense of security used in [5]. This strong notion is the concept of simulatability. Essentially, it is the cryptographic notion of secure implementation. Very roughly, a real system is called as secure as an ideal system in this sense if everything that can happen to honest users of the real system can also happen to the same honest users with the ideal system. This is typically proved by providing a simulator that, interacting with the ideal system and the honest users, and using an

adversary on the real system as a blackbox subsystem, simulates all visible actions of the real system online (i.e., at the time they occur).

For symmetric encryption, there is the following so-called *commitment problem* if one wants to achieve simulatability.² The ideal encryption system must somehow allow that secret keys are sent from one participant to another, because many protocols to be proven using such an ideal system are key-exchange protocols. This is the main difference to public-key systems, where an ideal system can assume that only public keys are sent around, because this is sufficient for all standard protocols. If the ideal system simply allows keys to be sent at any time (and typical Dolev-Yao models do allow all valid terms to be sent at any time), the following problem can occur: An honest participant first sends a ciphertext such that the adversary can see it, and later sends both the contained cleartext and the key. This behavior may even be reasonably designed into protocols, e.g., the ciphertext might be an encrypted bet that is later opened. The simulator will first learn in some abstract way that a ciphertext was sent and has to simulate it by some bitstring, which the adversary sees. Later the simulator sees abstractly that a key becomes known and that the ciphertext contains a specific application message. It cannot change the application message, thus it must simulate a key that decrypts the old ciphertext bitstring (produced without knowledge of the application message) to this specific message.

We discuss several ways of dealing with this problem. Our preferred one, for which we actually present the ideal and real symmetric encryption system, is to leave it to the surrounding protocol to guarantee that the commitment problem does not occur. Essentially, this means that the surrounding protocol must guarantee that keys are no longer sent in a form that might make them known to the adversary once an honest participant has started using them. Alternatives would be to build such a guarantee into the ideal and the real system, or to restrict oneself to the few encryption systems where this problem does not occur, or to work in models of cryptography that still have some ideal, unrealizable aspect, in particular the random-oracle model. We discuss these possibilities and our choice in more detail in Section 3. The most important argument for our choice is that, depending on the timing assumptions possible in the environment and on the protocol goals, a range of different measures are conceivable for guaranteeing the necessary order between the sending of keys and ciphertexts. Further, existing formal methods and automated tools are well suited to arguing about such properties. Instead, if we proposed measures in the underlying idealization, we would need a

2 Given that one wants to achieve simulatability, the problem is independent of a surrounding framework for nested terms, i.e., of our specific goal of making the ideal encryption system a subsystem in the library of [5].

once-and-for-all measure, and we would at present need to prove it by hand. To show the applicability of our choice for modeling the protocols typically analyzed in Dolev Yao models, we investigated the 50 protocols of the Clark-Jacob library [17] with respect to the commitment problem, and only one of them raises this problem.

Other design decisions to be taken with symmetric encryption are whether, given a ciphertext, the adversary may obtain information about the key used, and whether the ideal system prescribes that every decryption of a ciphertext (or a message of a different type) with the wrong key produces an error, or whether it may sometimes produce another message. Such questions are similar to the passive case in [2] or to the treatment of symmetric authentication in [6], but have to be combined in a consistent way into the overall ideal encryption system.

2. Underlying Definitions

Before discussing the commitment problem in more detail, we present the exact definition of simulatability, the strong security notion that causes this problem. For this, we first briefly sketch the underlying definitions from [30, 8]. This is the model used in the cryptographic library from [5] into which we embed our ideal encryption system.

A *system* consists of several possible *structures*. A structure consists of a set \hat{M} of connected correct machines and a subset S of free ports, called *specified ports*. A machine is a probabilistic IO automaton (extended finite-state machine) in a slightly refined model to allow complexity considerations. For these machines Turing-machine realizations are defined, and the complexity of those is measured in terms of a common security parameter k , given as the initial work-tape content of every machine. Readers only interested in using the ideal cryptographic library (or even only the ideal encryption system) in larger protocols only need normal, deterministic IO automata.

In a *standard real cryptographic system*, the structures are derived from one intended structure and a trust model consisting of an access structure ACC and a channel model χ . Here ACC contains the possible sets \mathcal{H} of indices of uncorrupted machines among the intended ones, and χ designates whether each channel is secure, authentic (but not private) or insecure. In a typical ideal system, each structure contains only one machine TH called *trusted host*.

Each structure is complemented to a *configuration* by an arbitrary *user* machine H and *adversary* machine A . H connects only to ports in S and A to the rest, and they may interact. The set of configurations of a system Sys is called $\text{Conf}(Sys)$. The general scheduling model in [30] gives each connection c (from an output port $c!$ to an input port $c?$) a buffer, and the machine with the corresponding clock port $c!$ can schedule a message there when it makes a tran-

sition. In real asynchronous cryptographic systems, network connections are typically scheduled by A . A configuration is a runnable system, i.e., for each k one gets a well-defined probability space of *runs*. The *view* of a machine in a run is the restriction to all in- and outputs this machine sees and its internal states. Formally, the view $view_{conf}(M)$ of a machine M in a configuration $conf$ is a *family of random variables* with one element for each security parameter value k .

2.1. Simulatability

Simulatability is the cryptographic notion of secure implementation. For reactive systems, it means that whatever might happen to an honest user in a real system Sys_{real} can also happen in the given ideal system Sys_{id} : For every structure $(\hat{M}_1, S) \in Sys_{real}$, every polynomial-time user H , and every polynomial-time adversary A_1 , there exists a polynomial-time adversary A_2 on a corresponding ideal structure $(\hat{M}_2, S) \in Sys_{id}$ such that the view of H is computationally indistinguishable in the two configurations. This is illustrated in Figure 1. Indistinguishability is a well-known cryptographic notion from [33].

Definition 2.1 (*Computational Indistinguishability*) *Two families $(var_k)_{k \in \mathbb{N}}$ and $(var'_k)_{k \in \mathbb{N}}$ of random variables on common domains D_k are computationally indistinguishable (" \approx ") iff for every algorithm D (the distinguisher) that is probabilistic polynomial-time in its first input,*

$$|P(D(1^k, var_k) = 1) - P(D(1^k, var'_k) = 1)| \in NEGL,$$

where $NEGL$ denotes the set of all negligible functions, i.e., $g: \mathbb{N} \rightarrow \mathbb{R}_{\geq 0} \in NEGL$ iff for all positive polynomials Q , $\exists k_0 \forall k \geq k_0: g(k) \leq 1/Q(k)$. \diamond

Intuitively, given the security parameter and an element chosen according to either var_k or var'_k , D tries to guess which distribution the element came from.

Definition 2.2 (*Simulatability*) *For two systems Sys_{real} and Sys_{id} , one says $Sys_{real} \geq Sys_{id}$ (at least as secure as) iff for every polynomial-time configuration $conf_1 = (\hat{M}_1, S, H, A_1) \in \text{Conf}(Sys_{real})$, there exists a polynomial-time configuration $conf_2 = (\hat{M}_2, S, H, A_2) \in \text{Conf}(Sys_{id})$ (with the same H) such that $view_{conf_1}(H) \approx view_{conf_2}(H)$. \diamond*

For the cryptographic library, this is even shown with black-box simulatability, i.e., A_2 consists of a simulator Sim that depends only on (\hat{M}_1, S) and uses A_1 as a blackbox sub-machine. An essential feature of this definition of simulatability is a composition theorem [30, 7], which roughly says that one can design and prove a larger system based on the ideal system Sys_{id} , and then securely replace Sys_{id} by the real system Sys_{real} .




Figure 1. Simulatability: The two views of H must be indistinguishable.

2.2. Notation

We write “ $:=$ ” for deterministic and “ \leftarrow ” for probabilistic assignment, and “ $\xleftarrow{\mathcal{R}}$ ” for uniform random choice from a set. By $x := y++$ for integer variables x, y we mean $y := y + 1$; $x := y$. The length of a message m is denoted as $\text{len}(m)$, and \downarrow is an error element available as an addition to the domains and ranges of all functions and algorithms. The list operation is denoted as $l := (x_1, \dots, x_j)$, and the arguments are unambiguously retrievable as $l[i]$, with $l[i] = \downarrow$ if $i > j$. A database D is a set of functions, called entries, each over a finite domain called attributes. For an entry $x \in D$, the value at an attribute att is written $x.att$. For a predicate $pred$ involving attributes, $D[pred]$ means the subset of entries whose attributes fulfill $pred$. If $D[pred]$ contains only one element, we use the same notation for this element. Adding an entry x to D is abbreviated $D := x$.

2.3. Overview of the Ideal Cryptographic Library

The ideal cryptographic library as defined in [5] offers its users abstract cryptographic operations, such as commands to encrypt or decrypt a message, to make or test a signature, and to generate a nonce. All these commands have a simple, deterministic behavior in the ideal system. In a reactive scenario, this semantics is based on state, e.g., of who already knows which terms. State is stored in a “database”. Each entry of the database has a type (e.g., “signature”), and pointers to its arguments (e.g., a key and a message). This corresponds to the top level of a Dolev-Yao term; an entire term can be found by following the pointers. Further, each entry contains handles for those participants who already know it. The reason for using handles to make an entry accessible for higher protocols is that an idealized cryptographic term and the corresponding real message have to be presented in the same way to higher protocols to allow for a provably secure implementation in the sense of simulatability. In the ideal library, handles essentially point to Dolev-Yao-like terms, while in the real library they point to cryptographic messages.

The ideal cryptographic library does not allow cheating by construction. For instance, if it receives a command to encrypt a message m with a certain key, it simply makes

an abstract database entry for the ciphertext. Another user can only ask for decryption of this ciphertext if he has handles to both the ciphertext and the secret key. Similarly, if a user issues a command to sign a message, the ideal system looks up whether this user should have the secret key. If yes, it stores that this message has been signed with this key. Later tests are simply look-ups in this database. A send operation makes an entry known to other participants, i.e., it adds handles to the entry. The underlying model does not only cover crypto operations, but it is an entire reactive system and therefore contains an abstract network model.

3. Design Decisions for Symmetric Encryption in Simulatability Proofs

In this section, we discuss several approaches to solve the commitment problem sketched in the introduction. We further elaborate on the main design decisions that we made to provide a suitable deterministic abstractions of symmetric encryption.

3.1. The Commitment Problem and Solution Approaches

As the name suggests, a “commitment problem” in simulatability proofs captures a situation where the simulator commits itself to a certain message and later has to change this commitment to allow for a correct simulation.

In the case of symmetric encryption, the commitment problem occurs if the simulator has to construct an indistinguishable ciphertext, knowing neither the secret key nor the plaintext used for the corresponding ciphertext in the real world. To simulate the missing key, the simulator will create a new secret key, or rely on an arbitrary, fixed key if the encryption systems guarantees indistinguishable keys, see [2]. Instead of the unknown plaintext, the simulator will encrypt an arbitrary message of the correct length, relying on the indistinguishability of ciphertexts of different messages. So far, the simulation is fine. It even stays fine if the message becomes known later because secure encryption still guarantees that it is indistinguishable that the simulator’s ciphertext contains a wrong message. However, if the secret key becomes known later, the simulator runs into trouble, because, learning abstractly about this fact, it has to produce a suitable key that decrypts its ciphertext into the correct message. It cannot cheat with the message because it has to produce the correct behavior towards the honest users. This is typically not possible.

There is one perfect exception to the commitment problem, the one-time pad. For this specific encryption system, the simulator can open an arbitrary ciphertext string c to an arbitrary message m by selecting the key as $c \oplus m$. However, we do not want to restrict the ideal encryption system

to modeling one-time pads, and for standard encryption systems and standard modes of operation, certainly no similar process is known. We can even show that no encryption system with fixed-length keys and a deterministic decryption algorithm can have this property. Assume there are x possible keys. Now let a protocol send fresh random messages whose overall length allows $2x$ possibilities. Hence if the simulator later has to produce a key, it has to be able to provide for $2x$ different cases of what the messages were with just x keys. Thus some of the keys it produces must fit more than one message tuple for the same given ciphertext tuple. “Fit” means in particular that the assumed deterministic decryption algorithm used by honest participants will in fact decrypt this ciphertext tuple to these messages. This is impossible.

The reason why the commitment problem did not occur in the cryptographic library before is that for public-key encryption it was enforced that secret keys are never sent, while for symmetric authentication it could be enforced that an authenticator never becomes known without the message it authenticates.

Related problems are known from uncoercible encryption and from adaptively secure multi-party function evaluation. However, none of the solutions provided there fits our case: Uncoercible encryption has even stronger requirements that need a physical assumption to be fulfilled at all [14]. Adaptively secure multi-party computation can either use deletion of old keys instead [10], or concentrates on public-key schemes [15, 9, 18], simply assuming the one-time pad for the symmetric case.

In the following, we introduce possible approaches to solve the commitment problem.

3.1.1. Assumptions about Sending Keys Our aim is an abstraction that is as simple as possible and works for the cases typically analyzed in Dolev-Yao models. It turns out that for these typical cases, the commitment problem does not occur since the overall protocol ensures that keys are not sent after having been used.

- Protocols with pre-distributed keys, as often assumed for authentication protocols, clearly fulfill this assumption. Formally, this can be seen as a synchronization assumption stating that the pre-distribution phase is over, at least per key, before one of the participants sharing this key starts using it.
- Synchronous protocols can make similar assumptions even if key exchange is part of the protocols, e.g., by indicating time bounds for the exchange phase and the usage phase for each key in the key exchange messages.
- Two-party protocols for exchanging a session key (using a symmetric or asymmetric master key) clearly ful-

fill the assumption if the party who generates the key sends it before using it. This is typically true.

- Three-party protocols where a key-distribution center S helps parties A and B to exchange a secret key sk come in several flavors: If S sends sk to A and B , and it can do so in one step (this depends on the detailed model of asynchrony), then the assumption is automatically fulfilled, and so it is if S sends sk to A , and A sends it to B and only then starts to use it. If S sends sk in two different steps, then the recipient of the first of these messages (or both if the first and second message look equal) has to wait for a confirmation from its partner before using the key. Many protocols have such a confirmation anyway.
- Many group key distribution protocols already have confirmation phases that can be used to fulfill the assumption.

To get a representative assessment of the restrictiveness of the commitment problem, we further investigated the protocols of the Clark-Jacob library [17]. From the 50 protocols of the library, only one—the (flawed) Wide Mouthed Frog protocol—raises the commitment problem. Further, avoiding the commitment problem is an integrity property that seems well within the scope of current automated protocol proof tools, so that it can be verified together with the application properties of a protocol. (This will become even clearer with the formal definition in Figure 2 and Definition 6.1.)

Hence our approach is to define a simple abstraction of symmetric encryption, and to show that it can be securely implemented provided that the commitment problem does not occur. The alternative approaches discussed below either exclude many more protocols, or require a much more complex abstraction, or rely on unrealistic assumptions like the random oracle model, or only work for special non-committing encryption schemes. An additional benefit of our solution is that we expect that our ideal encryption system also works with the random oracle model and non-committing encryption schemes, even for protocols that *do* have the commitment problem.

3.1.2. Internal Restrictions on Sending Keys Another solution is to guarantee in the ideal and real system that the commitment problem cannot occur. This means that the systems only permit operations that do not cause the commitment problem, e.g., if a key has already been used for encrypting, it may no longer be sent. The problem is that this is a distributed property and thus not trivial to enforce in the real system. To implement it without imposing further restrictions on the patterns of how keys can be passed on, we might need Byzantine agreement before any participant first uses a key. This seems a highly unnatural underlying implementation for the authentication and key ex-

change protocols typically proved with Dolev-Yao models. It seems more natural to enforce restrictions on the patterns of how keys can be passed on. This certainly means that both the ideal system and the real system have to keep track of the current status of each key for each participant, e.g., whether it may still be sent. Furthermore, we either have to provide general rules for confirmation messages (compare Section 3.1.1) or to enforce patterns where no confirmation messages are needed. As we saw, the former already excludes some important cases, while the latter pulls distributed-systems aspects down into the cryptographic primitives. We therefore decided not to follow this approach.

3.1.3. Random Oracle Approach and Special Encryption Schemes The commitment problem can be circumvented by conducting the proof in the random oracle model [12]. In a nutshell, including a random oracle in the encrypted message prevents the simulator from committing itself to a fixed value since the oracle can still be suitably instantiated when the commitment is opened respectively the secret key is sent. However, idealizations like random oracles do not capture cryptographic realities and protocols are known which are provably secure in these idealizations but insecure for any instantiation of the oracle [16], so that the benefit over simply using a Dolev-Yao model is not as great as we desire.

Further, the commitment problem does not occur for non-committing encryption schemes, but as we showed above, this currently only leaves the one-time pad.

3.2. The Need for Authenticated Encryption

Assume that a user encrypts a message m with one symmetric encryption key sk_1 , and decrypts the resulting ciphertext with a key sk_2 . In the ideal system, the result is the error symbol \downarrow because no equation for this case is defined. In the real world, however, some encryption schemes yield another message m' . In particular, the one-time pad always yields a result. Similar problems are known from normal Dolev-Yao models, e.g., see [27].

We solve this problem by only considering encryption schemes that answer decryption requests with wrong keys with \downarrow , i.e., encryption schemes that provide a certain kind of authenticity. Formally, we use *authenticated symmetric encryption schemes* as defined in [13, 11]. They intuitively guarantee that if one does not know a specific key, it is infeasible to compute a ciphertext that can be validly decrypted with this key; see the definition in Section 5.1. This definition implies that decryption with a wrong key will always output \downarrow except with negligible probability.

Instead of restricting the encryption schemes used, one could try to define the ideal system such that it allows non-error outputs for decryptions with wrong keys. However,

a deterministic abstraction cannot achieve this because in the real system, decryption with different wrong keys will yield different messages if any, while in the ideal system all such wrong keys have a common abstraction. A non-deterministic choice could be achieved by letting the adversary make the choice of the resulting message, but this seems a somewhat undesirable ideal system, given that authenticated encryption is efficiently implementable under normal cryptographic assumptions.

3.3. Modeling Special Adversary Capabilities

Our idealization finally has to reflect special capabilities that the adversary may have with respect to symmetric encryption schemes in the real world.

First, we allow for checking whether encryptions have been created with the same secret key, as the definition of authenticated encryption schemes does not exclude that this can happen in the real system. For public-key encryption, this was achieved in [5] by tagging ciphertexts with the corresponding public key. For symmetric encryption, this is not possible as no public key exists. We solve this problem by tagging abstract ciphertexts with an otherwise meaningless “public key” solely used as an identifier for the secret key. An alternative approach was taken in [2] by only considering those encryption schemes that guarantee indistinguishable keys; for these schemes, this problem does not occur. However, if we wanted to restrict ourselves to this case we would first need to extend it to authenticated encryption.

Secondly, as encryption keys can also come from the adversary, it might happen that an encryption can be validly decrypted with several keys for incorrectly chosen keys. (The security definition only considers correct keys.) Hence it must be possible to tag encryptions with additional key identifiers during the execution of the ideal system. Encryptions without key identifiers model encryptions from the adversary for which no suitable key is known yet.

4. Ideal System

In the following, we present our ideal encryption system. We do this as an addition to the ideal cryptographic library reviewed in Section 2.3 for capturing symmetric encryption primitives. We stress that for modeling and proving cryptographic protocols using our abstraction, it is sufficient to understand and use the ideal system described in this section. Later sections only justify the cryptographic faithfulness of this ideal library.

4.1. Structures and Parameters

The ideal system consists of a trusted host $\text{TH}_{\mathcal{H}}$ for every subset \mathcal{H} of a set $\{1, \dots, n\}$ of users, denoting the pos-

sible honest users. It has a port $\text{in}_u?$ for inputs from and a port $\text{out}_u!$ for outputs to each user $u \in \mathcal{H}$ and for $u = \mathbf{a}$, denoting the adversary.

The ideal system keeps track of the length of messages using a tuple L of abstract length functions. We add functions $\text{skse_len}^*(k)$ and $\text{symenc_len}^*(k, l)$ to L for the length of symmetric encryption keys and ciphertexts, depending on a security parameter k and the length l of the message. Each function has to be polynomially bounded and efficiently computable.

4.2. States

The state of $\text{TH}_{\mathcal{H}}$ consists of a database D and variables size , curhnd_u for $u \in \mathcal{H} \cup \{\mathbf{a}\}$. The database D contains abstractions from real cryptographic objects which correspond to the top levels of Dolev-Yao terms. An entry has the following attributes:

- $x.\text{ind} \in \mathcal{INDS}$, called index, consecutively numbers all entries in D . We use the index as a primary key attribute of the database, i.e., we write $D[i]$ for the selection $D[\text{ind} = i]$.
- $x.\text{type} \in \text{typeset}$ identifies the type of x . We add types skse , pkse , and symenc to typeset from [5], denoting secret symmetric encryption keys, corresponding “public keys”, and symmetric encryptions. The type pkse is a so-called *secret type*, i.e., it must not be put into lists and hence cannot be transferred.
- $x.\text{arg} = (a_1, a_2, \dots, a_j)$ is a possibly empty list of arguments. Many values a_i are indices of other entries in D and thus in \mathcal{INDS} . We sometimes distinguish them by a superscript “ind”.
- $x.\text{hnd}_u \in \mathcal{HANDS} \cup \{\downarrow\}$ for $u \in \mathcal{H} \cup \{\mathbf{a}\}$ are handles by which a user or adversary u knows this entry. $x.\text{hnd}_u = \downarrow$ means that u does not know this entry. We use a superscript “hnd” for handles.
- $x.\text{len} \in \mathbb{N}_0$ denotes the “length” of the entry, which is computed by applying the functions from L .

Initially, D is empty. $\text{TH}_{\mathcal{H}}$ has a counter $\text{size} \in \mathcal{INDS}$ for the current number of elements in D . New entries always receive $\text{ind} := \text{size}++$, and $x.\text{ind}$ is never changed. For the handle attributes, it has counters curhnd_u (current handle) initialized with 0, and each new handle for u will be chosen as $i^{\text{hnd}} := \text{curhnd}++$.

$\text{TH}_{\mathcal{H}}$ further maintains explicit bounds on the length of messages and the number of activations to achieve polynomial runtime independent of the environment. The bounds from [5] can be used without modification except that the number of permitted inputs from the adversary has to be enlarged. This is just a technical detail to allow for a correct proof of simulatability. We omit further details.

4.3. New Inputs and their Evaluation

The ideal system has several types of inputs: *Basic commands* are accepted at all ports $\text{in}_u?$; they correspond to cryptographic operations and have only local effects, i.e., only an output at the port $\text{out}_u?$ for the same user occurs and only handles for u are involved. *Local adversary commands* are of the same type, but only accepted at $\text{in}_\mathbf{a}?$; they model tolerated imperfections, i.e., possibilities that an adversary may have, but honest users do not. *Send commands* output values to other users. The notation $j \leftarrow \text{algo}(i)$ for a command algo of $\text{TH}_{\mathcal{H}}$ means that $\text{TH}_{\mathcal{H}}$ receives an input $\text{algo}(i)$ and outputs j if the input and output port are clear from the context. We only allow lists to be encrypted and transferred following a general convention in [5].

For symmetric encryption we add new basic commands and local adversary commands; the send commands are unchanged. We now define the precise new inputs and how $\text{TH}_{\mathcal{H}}$ evaluates them based on its abstract state. Handle arguments are tacitly required to be in \mathcal{HANDS} and existing, i.e., $\leq \text{curhnd}_u$, at the time of execution. The underlying model further bounds the length of each input to ensure polynomial runtime; these bounds are not written out explicitly, but can easily be derived from the domain expectations given for the individual inputs.

The algorithm $i^{\text{hnd}} \leftarrow \text{ind2hnd}_u(i)$ (with side effect) denotes that $\text{TH}_{\mathcal{H}}$ determines a handle i^{hnd} for user u to an entry $D[i]$: If $i^{\text{hnd}} := D[i].\text{hnd}_u \neq \downarrow$, it returns that, else it sets and returns $i^{\text{hnd}} := D[i].\text{hnd}_u := \text{curhnd}_u++$. On non-handles, it is the identity function. ind2hnd_u^* applies ind2hnd_u to each element of a list.

4.3.1. Basic Commands First we consider basic commands. This comprises operations for key generation, encryption, and decryption. We assume the current input is made at port $\text{in}_u?$, and the result goes to $\text{out}_u!$.

- **Key generation:** $\text{skse}^{\text{hnd}} \leftarrow \text{gen_symenc_key}()$. Set $\text{skse}^{\text{hnd}} := \text{curhnd}_u++$ and

$$D \quad := \leftarrow \quad (\text{ind} := \text{size}++, \text{type} := \text{pkse}, \\ \text{arg} := (), \text{len} := 0);$$

$$D \quad := \leftarrow \quad (\text{ind} := \text{size}++, \text{type} := \text{skse}, \\ \text{arg} := (\text{ind} - 1), \text{hnd}_u := \text{skse}^{\text{hnd}}, \\ \text{len} := \text{skse_len}^*(k)).$$

The first entry, an “empty” public key without handle, serves as the mentioned key identifier for the secret key. The argument of the secret key “points” to the empty public key.

- **Encryption:** $c^{\text{hnd}} \leftarrow \text{sym_encrypt}(\text{skse}^{\text{hnd}}, l^{\text{hnd}})$.
Let $\text{skse} := D[\text{hnd}_u = \text{skse}^{\text{hnd}} \wedge \text{type} = \text{skse}].\text{ind}$ and $l := D[\text{hnd}_u = l^{\text{hnd}} \wedge \text{type} = \text{list}].\text{ind}$. Return \downarrow if either of these is \downarrow , or if $\text{length} :=$

$\text{symenc_len}^*(k, D[l].\text{len}) > \text{max_len}(k)$. Otherwise, set $c^{\text{hnd}} := \text{curhnd}_u++$, $\text{pkse} := \text{skse} - 1$ and

$$D := (\text{ind} := \text{size}++, \text{type} := \text{symenc}, \\ \text{arg} := (l, \text{pkse}), \text{hnd}_u := c^{\text{hnd}}, \\ \text{len} := \text{length}).$$

The general argument format for entries of type `symenc` is $((l_1, \text{pkse}_1), \dots, (l_j, \text{pkse}_j))$. The arguments $\text{pkse}_1, \dots, \text{pkse}_j$ are pairwise disjoint key identifiers of those secret keys for which the encryption validly decrypts into messages l_1, \dots, l_j , respectively. We will see in Section 4.3.2 that additional key identifiers for an encryption can be added during the execution, e.g., since the adversary has created a suitable key. Such arguments are appended at the end of the existing list. An empty sequence of arguments models encryptions from the adversary for which no suitable secret key has been received yet.

- **Decryption:** $l^{\text{hnd}} \leftarrow \text{sym_decrypt}(\text{skse}^{\text{hnd}}, c^{\text{hnd}})$.
If $c := D[\text{hnd}_u = c^{\text{hnd}} \wedge \text{type} = \text{symenc}].\text{ind} = \downarrow$ or $\text{skse} := D[\text{hnd}_u = \text{skse}^{\text{hnd}} \wedge \text{type} = \text{skse}].\text{ind} = \downarrow$, return \downarrow . Otherwise, let $((l_1, \text{pkse}_1), \dots, (l_j, \text{pkse}_j)) := D[c].\text{arg}$ (where j may be 0). If $\text{skse} - 1 = \text{pkse}_i$ for some $1 \leq i \leq j$, set $l^{\text{hnd}} := \text{ind2hnd}_u(l_i)$ else $l^{\text{hnd}} := \downarrow$.

4.3.2. Local Adversary Commands The following local commands are only accepted at the port $\text{in}_a?$. They model special capabilities of the adversary, see Section 3.3. For dealing with symmetric encryptions from the adversary for which no suitable key has been received yet, we provide a command for generating an *unknown symmetric encryption*. Later, suitable secret keys may be received. A command for *fixing symmetric encryptions* takes care of this. Finally, we allow the adversary to retrieve all information that we do not explicitly require to be hidden, e.g., arguments and the type of a given handle. For this, we extend the general command for *parameter retrieval* for the symmetric encryption system. For entries of type `symenc`, only the length of the encrypted message is output instead of the message itself unless the adversary has the corresponding secret-key handle.

- **Unknown symmetric encryption:**
 $c^{\text{hnd}} \leftarrow \text{adv_unknown_symenc}(l)$ with $l \in \mathbb{N}$.
Return \downarrow if $l > \text{max_len}(k)$. Set $c^{\text{hnd}} := \text{curhnd}_a++$ and
$$D := (\text{ind} := \text{size}++, \text{type} := \text{symenc}, \\ \text{arg} := (), \text{hnd}_a := c^{\text{hnd}}, \text{len} := l).$$
- **Fixing symmetric encryption:**
 $v \leftarrow \text{adv_fix_symenc_content}(\text{skse}^{\text{hnd}}, c^{\text{hnd}}, l^{\text{hnd}})$.
Return \downarrow if $c := D[\text{hnd}_a = c^{\text{hnd}} \wedge \text{type} = \text{symenc}].\text{ind} = \downarrow$, if $\text{skse} := D[\text{hnd}_u = \text{skse}^{\text{hnd}} \wedge$

$\text{type} = \text{skse}].\text{ind} = \downarrow$, if $l := D[\text{hnd}_a = l^{\text{hnd}} \wedge \text{type} = \text{list}].\text{ind} = \downarrow$, or if $\text{symenc_len}^*(k, D[l].\text{len}) \neq D[c].\text{len}$.

Let $((l_1, \text{pkse}_1), \dots, (l_j, \text{pkse}_j)) := D[c].\text{arg}$ (where j may be 0) and $\text{pkse} := \text{skse} - 1$. If $\text{pkse} \notin \{\text{pkse}_1, \dots, \text{pkse}_j\}$ then set $D[c].\text{arg} := ((l_1, \text{pkse}_1), \dots, (l_j, \text{pkse}_j), (l, \text{pkse}))$ and $v := \text{true}$, else set $v := \text{false}$.

- **Parameter retrieval:** $(\text{type}, \text{arg}) \leftarrow \text{adv_parse}(m^{\text{hnd}})$.
This existing command always sets $\text{type} := D[\text{hnd}_a = m^{\text{hnd}}].\text{type}$, and for most types $\text{arg} := \text{ind2hnd}_a^*(D[\text{hnd}_a = m^{\text{hnd}}].\text{arg})$. This also applies to the new types `skse` and `pkse`. For $\text{type} = \text{symenc}$, let $((l_1, \text{pkse}_1), \dots, (l_j, \text{pkse}_j)) := D[\text{hnd}_a = m^{\text{hnd}}].\text{arg}$. For $i \in \{1, \dots, j\}$, let $\text{pkse}_i^{\text{hnd}} := \text{ind2hnd}_a(\text{pkse}_i)$ and $\text{skse}_i := \text{pkse}_i + 1$. Then if $D[\text{skse}_i].\text{hnd}_a \neq \downarrow$, let $l'_i := \text{ind2hnd}_a(l_i)$, else $l'_i := D[l_i].\text{len}$. Finally let $\text{arg} := ((l'_1, \text{pkse}_1^{\text{hnd}}), \dots, (l'_j, \text{pkse}_j^{\text{hnd}}))$.

For unknown encryptions, neither a key identifier nor a message exists. Note further that parsing a symmetric encryption yields handles to the “empty” public keys. For an encryption generated by an honest user, the first public key always corresponds to the secret key with which the encryption was generated. If the adversary wants to know whether two encryptions were created using the same secret key, it parses them and compares the resulting public keys.

4.3.3. Send Commands The ideal cryptographic library offers commands for virtually sending messages to other users. Sending of a message is modeled by adding a new handle for the intended recipient and possibly one for the adversary to the database entry modeling the message. These handles always point to a list entry, which can contain arbitrary application data, ciphertexts, public keys, etc., and now also symmetric encryptions and the corresponding secret keys. These commands are unchanged from [5]; as an example we present those modeling insecure channels, which are the most commonly used ones, and omit secure channels and authentic channels.

- $\text{send}_i(v, l^{\text{hnd}})$, for $v \in \{1, \dots, n\}$. Intuitively, the list l shall be sent to user v . Let $l^{\text{ind}} := D[\text{hnd}_u = l^{\text{hnd}} \wedge \text{type} = \text{list}].\text{ind}$. If $l^{\text{ind}} \neq \downarrow$, then output $(u, v, \text{ind2hnd}_a(l^{\text{ind}}))$ at $\text{out}_a!$.
- $\text{adv_send}_i(u, v, l^{\text{hnd}})$, for $u \in \{1, \dots, n\}$ and $v \in \mathcal{H}$ at port $\text{in}_a?$. Intuitively, the adversary wants to send list l to v , pretending to be u . Let $l^{\text{ind}} := D[\text{hnd}_a = l^{\text{hnd}} \wedge \text{type} = \text{list}].\text{ind}$. If $l^{\text{ind}} \neq \downarrow$ output $(u, v, \text{ind2hnd}_v(l^{\text{ind}}))$ at $\text{out}_v!$.

5. Real System

The real cryptographic library offers its users the same commands as the ideal one, i.e., honest users operate on cryptographic objects via handles. There is one separate machine with a database for each honest user in the real system, containing real cryptographic keys, real encryptions, etc.. Real bitstrings are actually sent between machines. The commands are implemented by real cryptographic algorithms, and the simulatability proof will show that nevertheless, everything a real adversary can achieve can also be achieved by an adversary in the ideal system, or otherwise the underlying cryptography can be broken. We now present our additions and modifications to the real system of [5], starting with a description of the underlying cryptographic definitions.

5.1. Underlying Cryptographic Operations

We denote a symmetric encryption scheme by a tuple $\mathcal{SE} = (\text{gen}_{\text{SE}}, \text{sym_encrypt}, \text{sym_decrypt}, \text{skse_len}, \text{symenc_len})$ of polynomial-time algorithms. Key generation for a security parameter $k \in \mathbb{N}$ is written as

$$sk \leftarrow \text{gen}_{\text{SE}}(1^k).$$

The length of sk is $\text{skse_len}(k) > 0$. We denote the encryption of a message $m \in \{0, 1\}^+$ by

$$c \leftarrow \text{sym_encrypt}_{sk}(m)$$

and decryption by

$$m := \text{sym_decrypt}_{sk}(c).$$

The result may be \downarrow ; then we call the ciphertext invalid for this key. A correctly generated ciphertext for a key of the correct length always has to be valid for this key.

The length of c is $\text{symenc_len}(k, \text{len}(m)) > 0$. This is also true for every c' with $\text{sym_decrypt}_{sk}(c') \neq \downarrow$ for a value $sk \in \{0, 1\}^{\text{skse_len}(k)}$. The functions skse_len and symenc_len must be bounded by multivariate polynomials. Our requirement that such functions exist is without loss of generality due to standard padding techniques.

Our security definition is the standard definition for authenticated symmetric encryption schemes from [13, 11]. It consists of two parts: The scheme must ensure confidentiality of messages under chosen-ciphertext attacks, and it must guarantee integrity of ciphertexts. In the following, we formulate these notions using the notation for interacting machines.

Definition 5.1 (*Security against Chosen-Ciphertext Attacks*) *Given a symmetric encryption scheme, the symmetric decryptor machine SymDec is defined as follows: It has one input and one output port, a variable sk , initialized with \downarrow , an initially empty set C and the following transition rules:*

- First generate a key as $sk \leftarrow \text{gen}_{\text{SE}}(1^k)$ and set $b \xleftarrow{\mathcal{R}} \{0, 1\}$.
- On input $(\text{symenc}, m_0, m_1)$ (intuitively a pair of messages an adversary hopes to be able to distinguish), and if $\text{len}(m_0) = \text{len}(m_1)$, set $c \leftarrow \text{sym_encrypt}_{sk}(m_b)$, $C := C \cup \{c\}$, and output c .
- On input (symdec, c_j) and if $c_j \notin C$, return $\text{sym_decrypt}_{sk}(c_j)$.

The encryption scheme is called indistinguishable under chosen-ciphertext attack if for every probabilistic polynomial-time machine A_{SD} that interacts with SymDec and finally outputs a bit b^* (meant as a guess at b), the probability of the event $b^* = b$ is bounded by $1/2 + g(k)$ for a negligible function g . \diamond

The machine for defining integrity of ciphertexts is defined similarly.

Definition 5.2 (*Integrity of Ciphertexts*) *Given a symmetric encryption scheme, we define the symmetric integrity machine SymInt as follows: It has one input and one output port, a variable sk initialized with \downarrow , and the following transition rules:*

- First generate a key as $sk \leftarrow \text{gen}_{\text{SE}}(1^k)$.
- On input (symenc, m_j) , return $c_j \leftarrow \text{sym_encrypt}_{sk}(m_j)$.
- On input (symdec, c'_j) , return $m'_j := \text{sym_decrypt}_{sk}(c'_j)$.

The encryption scheme is said to have integrity of ciphertexts if for every probabilistic polynomial-time machine A_{SI} that interacts with SymInt the probability is negligible (in k) that SymInt outputs $m \neq \downarrow$ on any input (symdec, c) where c was not output by SymInt upon a command (symenc, \cdot) until that time, i.e., not among the c_j 's. \diamond

Symmetric encryptions schemes provably secure with respect to these two definitions exist under reasonable assumptions [31]. Bellare and Namprepre even showed in [11] that such encryption schemes can be derived from any symmetric encryption scheme that is provably secure under adaptive chosen-plaintext attacks together with any strongly unforgeable message authentication code by first encrypting a plaintext and then appending a MAC to the obtained ciphertext.

5.2. Structures

The intended structure of the real cryptographic library consists of n machines $\{M_1, \dots, M_n\}$. Each M_u has ports $\text{in}_u?$ and $\text{out}_u!$, so that the same honest users can connect to the ideal and the real system. Each M_u has connections to each M_v exactly as in [5], in particular an insecure

connection called $\text{net}_{u,v,i}$ for normal use. They are called network connections and the corresponding ports network ports. Any subset \mathcal{H} of $\{1, \dots, n\}$ can denote the indices of correct machines. The resulting actual structure consists of the correct machines with modified channels according to a channel model. In particular, an insecure channel is split in the actual structure so that both machines actually interact with the adversary. Details of the channel model are not needed here. Such a structure then interacts with honest users H and an adversary A .

5.3. States of a Machine

The state of each machine M_u consists of a database D_u and a variable curhnd_u . Each entry x in D_u has the following attributes:

- $x.\text{hnd}_u \in \mathcal{H}\mathcal{N}\mathcal{D}\mathcal{S}$ consecutively numbers all entries in D_u . We use it as a primary key attribute, i.e., we write $D_u[i^{\text{hnd}}]$ for the selection $D_u[\text{hnd}_u = i^{\text{hnd}}]$.
- $x.\text{word} \in \{0, 1\}^+$ is the real representation of x .
- $x.\text{type} \in \text{typeset} \cup \{\text{null}\}$ identifies the type of x , where the value null denotes an unparsed entry.
- $x.\text{add_arg}$ is a list of (“additional”) arguments. For entries of our new types it is always $()$.

Initially, D_u is empty. M_u has a counter $\text{curhnd}_u \in \mathcal{H}\mathcal{N}\mathcal{D}\mathcal{S}$ for the current size of D_u . The subroutine

$$(i^{\text{hnd}}, D_u) := (i, \text{type}, \text{add_arg})$$

determines a handle for certain given parameters in D_u : If an entry with the word i already exists, i.e., $i^{\text{hnd}} := D_u[\text{word} = i \wedge \text{type} \notin \{\text{sks}, \text{ske}\}].\text{hnd}_u \neq \downarrow$,³ it returns i^{hnd} , assigning the input values type and add_arg to the corresponding attributes of $D_u[i^{\text{hnd}}]$ only if $D_u[i^{\text{hnd}}].\text{type}$ was null . Else if $\text{len}(i) > \text{max_len}(k)$, it returns $i^{\text{hnd}} = \downarrow$. Otherwise, it sets and returns $i^{\text{hnd}} := \text{curhnd}_u++$, $D_u := (i^{\text{hnd}}, i, \text{type}, \text{add_arg})$.

Similar to the machine $\text{TH}_{\mathcal{H}}$, M_u maintains explicit bounds on the length of messages and number of activations to achieve polynomial runtime. We omit further details.

5.4. Inputs and their Evaluation

Now we describe how M_u evaluates individual new inputs.

³ The restriction $\text{type} \notin \{\text{sks}, \text{ske}\}$ (abbreviating secret keys of signature and public-key encryption schemes) is included for compatibility to the original library. Similar statements will occur some more times, but no further knowledge of such types is needed for understanding the new work.

5.4.1. Constructors and One-level Parsing The stateful commands are defined via functional constructors and parsing algorithms for each type. A general functional algorithm

$$(\text{type}, \text{arg}) \leftarrow \text{parse}(m),$$

then parses arbitrary entries as follows: It first tests if m is of the form $(\text{type}, m_1, \dots, m_j)$ with $\text{type} \in \text{typeset} \setminus \{\text{pkse}, \text{pka}, \text{sks}, \text{ske}, \text{garbage}\}$ and $j \geq 0$. If not, it returns $(\text{garbage}, ())$. Otherwise it calls a type-specific parsing algorithm $\text{arg} \leftarrow \text{parse_type}(m)$. If the result is \downarrow , parse again outputs $(\text{garbage}, ())$. By

“parse m^{hnd} ”

we abbreviate that M_u calls $(\text{type}, \text{arg}) \leftarrow \text{parse}(D_u[m^{\text{hnd}}].\text{word})$, assigns $D_u[m^{\text{hnd}}].\text{type} := \text{type}$ if it was still null , and may then use arg . By

“parse m^{hnd} if necessary”

we mean the same except that M_u does nothing if $D_u[m^{\text{hnd}}].\text{type} \neq \text{null}$.

5.4.2. Basic Commands and parse_type First we consider basic commands. They are again local. In M_u this means that they produce no outputs at the network ports. The term “tagged list” means a valid list of the real system. We assume that tagged lists are efficiently encoded into $\{0, 1\}^+$.

- **Key constructor:** $sk^* \leftarrow \text{make_symenc_key}()$.
Let $sk \leftarrow \text{gen}_{\text{SE}}(1^k)$, $sr \stackrel{\mathcal{R}}{\leftarrow} \{0, 1\}^{\text{nonce_len}(k)}$, and return $sk^* := (\text{skse}, sk, sr)$.
- **Key generation:** $skse^{\text{hnd}} \leftarrow \text{gen_symenc_key}()$.
Let $sk^* \leftarrow \text{make_symenc_key}()$, $skse^{\text{hnd}} := \text{curhnd}_u++$, and $D_u := (skse^{\text{hnd}}, sk^*, \text{skse}, ())$.
- **Key parsing:** $\text{arg} \leftarrow \text{parse_skse}(sk^*)$.
If sk^* is of the form (skse, sk, sr) with $sk \in \{0, 1\}^{\text{skse_len}(k)}$ and $sr \in \{0, 1\}^{\text{nonce_len}(k)}$, return $()$, else \downarrow .
- **Encryption constructor:** $c^* \leftarrow \text{make_symenc}(sk^*, l)$, for $sk^*, l \in \{0, 1\}^+$.
Set $r \stackrel{\mathcal{R}}{\leftarrow} \{0, 1\}^{\text{nonce_len}(k)}$, $sk := sk^*[2]$, $sr := sk^*[3]$. Encrypt as $c \leftarrow \text{sym_encrypt}_{sk}((r, l))$, and return $c^* := (\text{symenc}, sr, r, c)$.
- **Encryption:** $c^{\text{hnd}} \leftarrow \text{sym_encrypt}(skse^{\text{hnd}}, l^{\text{hnd}})$.
Parse $skse^{\text{hnd}}$ and l^{hnd} if necessary. If $D_u[skse^{\text{hnd}}].\text{type} \neq \text{skse}$ or $D_u[l^{\text{hnd}}].\text{type} \neq \text{list}$, then return \downarrow . Otherwise set $sk^* := D_u[skse^{\text{hnd}}].\text{word}$, $l := D_u[l^{\text{hnd}}].\text{word}$, and $c^* \leftarrow \text{make_symenc}(sk^*, l)$. If $\text{len}(c^*) > \text{max_len}(k)$, return \downarrow , else set $c^{\text{hnd}} := \text{curhnd}_u++$ and $D_u := (\text{symenc}^{\text{hnd}}, c^*, \text{symenc}, ())$.
- **Encryption parsing:** $\text{arg} \leftarrow \text{parse_symenc}(c^*)$.
If c^* is not of the form $(\text{symenc}, sr, r, c)$ with $sr, r \in \{0, 1\}^{\text{nonce_len}(k)}$ and $c \in \{0, 1\}^+$, return \downarrow , else set $\text{arg} := ()$.

- *Decryption*: $l^{\text{hnd}} \leftarrow \text{sym_decrypt}(c^{\text{hnd}}, skse^{\text{hnd}})$.
Parse c^{hnd} and $skse^{\text{hnd}}$. If $D_u[c^{\text{hnd}}].\text{type} \neq \text{symenc}$ or $D_u[skse^{\text{hnd}}].\text{type} \neq \text{skse}$, return \downarrow . Else let $(\text{symenc}, sr, r, c) := D_u[c^{\text{hnd}}].\text{word}$ and $sk := D_u[skse^{\text{hnd}}].\text{word}[2]$. Let $l^* := \text{sym_decrypt}_{sk}(c)$ and $l := l^*[2]$. If $sr \neq D_u[skse^{\text{hnd}}].\text{word}[3]$, or $l^* = \downarrow$, or $l^*[1] \neq r$, or if l is not a tagged list, return $l^{\text{hnd}} := \downarrow$. Otherwise let $(l^{\text{hnd}}, D_u) := (l, \text{list}, ())$.

5.4.3. Send Commands and Network Inputs Similar to the ideal system, there is a command $\text{send}_i(v, l^{\text{hnd}})$ for sending a list l from u to v , but now using the port $\text{net}_{u,v,i}!$, i.e., using the real insecure network: On input $\text{send}_i(v, l^{\text{hnd}})$ for $v \in \{1, \dots, n\}$, M_u parses l^{hnd} if necessary. If $D_u[l^{\text{hnd}}].\text{type} = \text{list}$, M_u outputs $D_u[l^{\text{hnd}}].\text{word}$ at port $\text{net}_{u,v,i}!$.

Inputs at network ports are simply tested for being tagged lists and stored as in [5].

6. Security Proof

Our security claim is that the real cryptographic library extended with symmetric encryption is as secure as the ideal cryptographic library with symmetric encryption in the sense of Definition 2.2 provided that the commitment problem is avoided by the surrounding protocol.

We first have to define what it means that the commitment problem does not occur. We formalize the following event NoComm: if there exists an input at a specified port that causes a symmetric encryption to be generated such that the corresponding key is not known to the adversary, then future inputs may only cause this key to be sent within an encryption that cannot be decrypted by the adversary. Note that this property could still be marginally weakened by restricting it to those cases where the symmetric encryption is actually sent to the adversary; however, our variant is easier to verify for actual protocols since one does not have to additionally parse every sent term to look for a contained encryption. For technical reasons, we further exclude encryption cycles (such as encrypting a key with itself) within the definition of NoComm, which had to be required even for acquiring properties weaker than simulatability. We refer to [2] for further discussions.

To capture the event NoComm formally, we first define the tree of contained terms of a database entry $D[i]$, written $\text{tree}(D[i])$, by defining that $D[i]$ is the root of the tree, and $D[j]$ is a child of $D[k]$ if and only if $j \in D[k].\text{arg}$. We recall that symmetric encryptions do not maintain the secret keys used for generating these encryptions as arguments but only the corresponding public-key identifiers. To capture the absence of encryption cycles, we define a function order on honestly generated secret encryption keys that are not known to the adversary when they are first used. The

function order then assigns each key a number corresponding to the order in which the keys are first used for encryption. We also define that honestly generated secret keys of public-key encryption schemes are always of order 0. Later on, we will require that a key of order i may only be encrypted by keys of order $j < i$.

The event NoComm is formally defined in Figure 2. Here, a statement of the form “ $t : p?.\text{send}_A(l^{\text{hnd}}, v)$ ” means that a send command is input at port $p?$ of $\text{TH}_{\mathcal{H}}$ at time t so that the sent term will be received by the adversary. Formally, this means that v can be arbitrary for sending on insecure or authentic channels, and that v has to be dishonest for sending on secure channels. We further write $t : D$ to describe the contents of database D at time t . A statement of the form $t : \text{wrapped}(j, i)$ is true if and only if for every occurrence of the node $D[j]$ in $\text{tree}(t : D[i])$ with $t : D[j].\text{type} = \text{skse}$ there exists a node $D[k]$ in $\text{tree}(t : D[i])$ such that $t : D[k].\text{type} \in \{\text{symenc}, \text{enc}\}$, $D[j]$ is a descendant of $D[t : D[k].\text{arg}[1]]$ (i.e., of the encrypted message), $D[k].\text{hnd}_a = \downarrow$ and $t : \text{order}(sk) < t : \text{order}(j)$ where sk denotes the secret key used for encrypting the message, i.e., $sk := t : D[k].\text{arg}[1][2] + 1$ if $t : D[k].\text{type} = \text{symenc}$ respectively $sk := t : D[k].\text{arg}[2] - 1$ if $t : D[k].\text{type} = \text{enc}$.

Is is easy to see that one could as well define the event NoComm only in terms of the inputs that $\text{TH}_{\mathcal{H}}$ obtains from the honest users, i.e., independent of the state of $\text{TH}_{\mathcal{H}}$ and solely depending on the interaction with the surrounding protocol. However, this description would be very lengthy and is hence omitted for reasons of readability.

We now define those configurations to be *commitment-free* in which the event NoComm holds independent of the considered adversary, i.e., where the honest user already guarantees the validity of the event. As the event can be restated in terms of the inputs obtained from the user, commitment-free configurations are naturally also defined for the real library as it offers the same ports and commands to the honest users as the ideal library.

Definition 6.1 *A user H is commitment-free with respect to symmetric encryption and the machine $\text{TH}_{\mathcal{H}}$ if for all configurations $\text{conf} = (\text{TH}_{\mathcal{H}}, S_{\mathcal{H}}, H, A)$, the property NoComm as defined in Figure 2 holds. Configurations with a commitment-free user are called commitment-free configurations. The restriction of simulatability to the set of commitment-free configurations is denoted by \geq^{Comm} , i.e., for all commitment-free configurations of the real system, there exists a commitment-free configuration of the ideal system with the same honest user that achieves indistinguishable views for the honest user. \diamond*

Let $RPar$ be the set of valid parameter tuples for the real system, consisting of the number $n \in \mathbb{N}$ of participants, secure signature, encryption, and symmetric encryption schemes \mathcal{S} , \mathcal{E} , and \mathcal{SE} , and length functions and bounds

If there exists $t_1 \in \mathbb{N}$, $i \in \mathcal{INDS}$, $u_1 \in \mathcal{H}$ such that for $skse_{u_1}^{\text{hnd}} := D[i].\text{hnd}_{u_1}$, we have

$t_1 : \text{in}_{u_1}?.\text{sym_encrypt}(skse_{u_1}^{\text{hnd}}, l_1^{\text{hnd}})$ and	# If a term is encrypted at time t_1
$t_1 : D[i].\text{type} = \text{skse}$ and	# with a secret key
$t_1 : D[i].\text{hnd}_a = \downarrow$	# that is not known to the adversary

then the following must hold. For every $t_2 > t_1$, $v_2, u_2 \in \mathcal{H}$ we have

$t_2 : \text{in}_{u_2}?.\text{send}_A(l_2^{\text{hnd}}, v_2)$	# If another term is sent at time t_2 and
$D[i] \in \text{tree}(t_2 : D[\text{hnd}_{u_2} = l_2^{\text{hnd}}])$	# and the secret key is contained in this term
\implies	# then
$t_2 : \text{wrapped}(i, t_2 : D[\text{hnd}_{u_2} = l_2^{\text{hnd}}].\text{ind})$.	# the secret key is sufficiently wrapped

Figure 2. The property NoComm.

L' . For $(n, \mathcal{S}, \mathcal{E}, \mathcal{SE}, L') \in RPar$, let $Sys_{n, \mathcal{S}, \mathcal{E}, \mathcal{SE}, L'}^{\text{cry-sym, real}}$ be the resulting real cryptographic library. Further, let the corresponding length functions and bounds of the ideal system be formalized by a function $L := R2lpar(\mathcal{S}, \mathcal{E}, \mathcal{SE}, L')$, and let $Sys_{n, L}^{\text{cry-sym, id}}$ be the ideal cryptographic library with parameters n and L . The extension of $R2lpar$ to the newly added length functions for symmetric encryption, i.e., skse_len^* and symenc_len^* is straightforward and can be found in the long version of this paper [4]. Using the notation of Definition 2.2 and 6.1, we have

Theorem 6.1 (*Security of Cryptographic Library*) *For all parameters $(n, \mathcal{S}, \mathcal{E}, \mathcal{SE}, L') \in RPar$, we have*

$$Sys_{n, \mathcal{S}, \mathcal{E}, \mathcal{SE}, L'}^{\text{cry-sym, real}} \geq^{\text{Comm}} Sys_{n, L}^{\text{cry-sym, id}},$$

where $L := R2lpar(\mathcal{S}, \mathcal{E}, \mathcal{SE}, L')$. □

For proving this theorem for the original library without symmetric encryption, a simulator $\text{Sim}_{\mathcal{H}}$ has been defined in [5] such that even the combination of arbitrary polynomial-time users H and an arbitrary polynomial-time adversary A cannot distinguish the combination of the real machines M_u from the combination $\text{TH}_{\mathcal{H}}$ and $\text{Sim}_{\mathcal{H}}$ (for all sets \mathcal{H} indicating the correct machines). We sketch how we extend the simulator and then the proof of correct simulation to deal with symmetric encryption. A fully rigorous definition of $\text{Sim}_{\mathcal{H}}$ can be found in [4].

6.1. Simulator

Basically $\text{Sim}_{\mathcal{H}}$ has to translate real messages from the real adversary A into handles as $\text{TH}_{\mathcal{H}}$ expects them at its adversary input port $\text{in}_a?$ and vice versa. In both directions, $\text{Sim}_{\mathcal{H}}$ has to parse an incoming message completely because it can only construct the other version (abstract or real) bottom-up. This is done by recursive algorithms. The state of $\text{Sim}_{\mathcal{H}}$ mainly consists of a database D_a , similar to

the databases D_u , but storing the knowledge of the adversary. The behavior of $\text{Sim}_{\mathcal{H}}$ is sketched as follows.

Inputs from $\text{TH}_{\mathcal{H}}$. Assume that $\text{Sim}_{\mathcal{H}}$ receives an input $(u, v, x, l^{\text{hnd}})$ from $\text{TH}_{\mathcal{H}}$. If a bitstring l for l^{hnd} already exists in D_a , i.e., this message is already known to the adversary, the simulator immediately outputs l at port $\text{net}_{u, v, x}!$. Otherwise, it first constructs such a bitstring l with a recursive algorithm id2real . This algorithm decomposes the abstract term using basic commands and the adversary command adv_parse . At the same time, id2real builds up a corresponding real bitstring using real cryptographic operations and enters all new message parts into D_a to recognize them when they are reused, both by $\text{TH}_{\mathcal{H}}$ and by A .

We sketch how the simulator is extended to deal with symmetric encryption keys respectively symmetric encryptions. If the entry corresponding to l^{hnd} is a symmetric encryption key, id2real creates a new secret key by applying the function make_symenc_key and uses this key whenever an abstract encryption has to be simulated under the abstract key entry l^{hnd} . If the entry corresponding to l^{hnd} is a symmetric encryption, $\text{Sim}_{\mathcal{H}}$ first determines the corresponding secret key by means of the public key identifier of the encryption. After that, it checks whether the designated recipient of the handle is a dishonest or an honest party. In the first case, adv_parse reveals the plaintext of the encrypted message, so id2real only has to encrypt this plaintext with the determined secret key and output this encryption. If the designated recipient is honest, then adv_parse only outputs the length of the encrypted message. In this case, id2real encrypts a fixed message of equal length.

Inputs from A . Now assume that $\text{Sim}_{\mathcal{H}}$ receives a bitstring l from A at a port $\text{net}_{u, v, x}?$. If l is not a valid list, $\text{Sim}_{\mathcal{H}}$ aborts the transition. Otherwise it translates l into a corresponding handle l^{hnd} by an algorithm real2id , and outputs the abstract sending command $\text{adv_send}_x(w, u, l^{\text{hnd}})$ at port $\text{in}_a!$.

If a handle l^{hnd} for l already exists in D_a , then real2id reuses that. Otherwise it recursively parses a real bitstring

using the functional parsing algorithm. At the same time, it builds up a corresponding abstract term in the database of $\text{TH}_{\mathcal{H}}$. This finally yields the handle l^{hnd} . Furthermore, `real2id` enters all new subterms into D_a . For building up the abstract term, `real2id` makes extensive use of the special capabilities of the adversary modeled in $\text{TH}_{\mathcal{H}}$. In the real system, the bitstring may, e.g., contain an encryption which no encryption key is known yet that could validly decrypt this encryption. Therefore, the simulator has to be able to insert such an encryption with unknown key and unknown plaintext into the database of $\text{TH}_{\mathcal{H}}$, which explains the need for the command `adv_unknown_symenc`. Similarly, the adversary might send a new encryption key which has to be added to existing symmetric encryption entries for which this key is valid. All these and similar cases for symmetric encryption can be covered by using the special adversary capabilities that we offered in Section 4.3.2.

6.2. Proof of Correct Simulation

In the proof of the extended cryptographic library, now including symmetric encryption, we retain the proof structure from [5] as far as possible. The basic structure of that proof is that a combined system $C_{\mathcal{H}}$ is defined that essentially contains all aspects of both the real and the ideal system, and then bisimulations are proved between $C_{\mathcal{H}}$ and the combination $M_{\mathcal{H}}$ of the real machines, and between $C_{\mathcal{H}}$ and the combination $\text{THSim}_{\mathcal{H}}$ of the trusted host and the simulator. A bisimulation, however, cannot deal with computational indistinguishability. Hence at the beginning of the proof, the real asymmetric encryptions were replaced by simulated ones as made in the simulator. This could be done in one replacement step, using a low-level idealization of asymmetric encryption and the composition theorem. The overall proof is illustrated in Figure 3 where Steps 1 and 2 depict the treatment of public-key encryption, and where Step 4 and the system $C_{\mathcal{H}}^*$ were not present.

Symmetric encryption is more complicated because we also allow symmetric keys to be sent around. However, a typical low-level idealization would assume, like the original cryptographic definitions of encryption security, that the keys are only used for correct en- and decryption. Intuitively, this is OK in our case because the simulator treats keys that the adversary learns perfectly correctly, and if the adversary does not learn a key, i.e., the key is never sent at all or only encrypted, then it should be as good as if it had never been used apart from en- and decryption. However, here we argue with the security of encryption while trying to show the security of encryption, and we must ensure that the argument is not circular. Fortunately, our assumptions guarantee that we always argue with the security of encryption with another key when treating one key, and that the keys can be arranged in non-circular order for this

treatment.

We therefore perform a successive exchange of real encryptions for simulated encryptions by a so-called hybrid argument. We do this in the combined system because there we have all information easily available, in particular, which keys are ideally known to the adversary. In the overall proof depicted in Figure 3, Step 4 and the fact that there are multiple indexed combined systems $C_{\mathcal{H}}^{(i)}$ are the new aspects for symmetric encryption. Due to lack of space, we only briefly sketch the individual steps of the proof and refer to the long version [4] for details.

6.2.1. Initial and Final Combined Systems The initial combined system $C_{\mathcal{H}}$ is defined indirectly from the real and ideal system exactly as in [5]. In particular, it contains a database D^* that extends the database D of TH by an attribute `word` containing real word entries as in $M_{\mathcal{H}}$ or $\text{Sim}_{\mathcal{H}}$. These real words are computed as in $M_{\mathcal{H}}$ for entries generated by basic commands, i.e., by the honest users, while they are computed as in $\text{Sim}_{\mathcal{H}}$ for entries resulting from network inputs, i.e., values coming from the adversary. This implies that all symmetric encryptions produced by honest users contain a real plaintext message.

The final combined system $C_{\mathcal{H}}^*$ is equal to $C_{\mathcal{H}}$ except for symmetric encryptions: For encryptions made by honest users and with keys of honest users, a simulated message 1^{len^*} defined as in $\text{Sim}_{\mathcal{H}}$ is encrypted instead of a real plaintext message.

6.2.2. Hybrid Combined Systems Two successive hybrid combined systems differ only in the behavior for one symmetric encryption key $sk^{(i)}$: While $C_{\mathcal{H}}^{(i)}$ still encrypts real messages with this key, $C_{\mathcal{H}}^{(i+1)}$ encrypts simulated messages with it. The selection of $sk^{(i)}$ must guarantee that $sk^{(i)}$ is only encrypted with keys $sk^{(j)}$ for $j < i$, so that these encryptions have already been replaced by encryptions of fixed messages 1^{len^*} . We guarantee this by numbering the keys in the order in which they are first used for encryption. (The combined system has global knowledge of this.) This corresponds to the function order introduced for the definition of the NoComm property.

6.2.3. Low-level Combined Symmetric Encryption Machine Within the hybrid argument, we do not want to argue individually with the secrecy and integrity of each ciphertext. We therefore first define a machine `SymComb` that corresponds almost precisely to the entire action of a hybrid system with one key. Similar to the machine `SymDec`, the machine `SymComb` chooses a bit b denoting whether ciphertexts are decrypted via the actual decryption algorithm or in the idealized way via table look-up. We then show that every successful attack against `SymComb` implies a successful attack on one of the machines `SymDec` and `SymInt`.

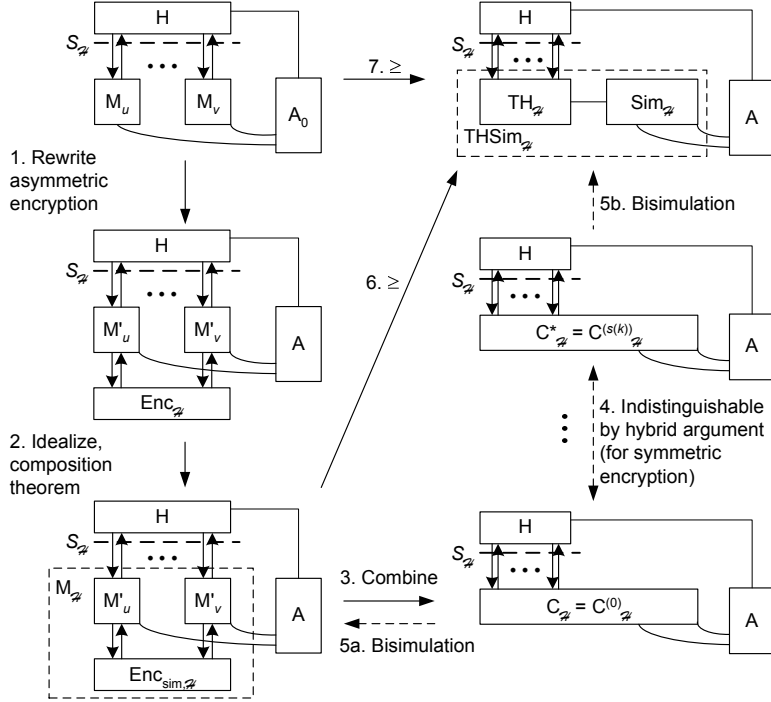


Figure 3. Proof with hybrid systems.

6.2.4. The Hybrid Argument We now show that the combined systems $C_{\mathcal{H}}$ and $C_{\mathcal{H}}^*$ are indistinguishable. The overall structure of this hybrid argument is standard for the case of a polynomially growing number of hybrids. The special aspects of our usage of symmetric encryption come in when we treat the cases of how a secret key can and cannot occur in larger terms, and why the possible occurrences do no harm.

The core of the hybrid argument is to show how the encryption machine SymComb can be used to simulate either $C_{\mathcal{H}}^{(i)}$ or $C_{\mathcal{H}}^{(i+1)}$, depending on the bit b in SymComb . We call the rest of this simulation $C'_{\mathcal{H}}^{(i)}$, i.e., the combination of $C'_{\mathcal{H}}^{(i)}$ and SymComb should yield $C_{\mathcal{H}}^{(i)}$ or $C_{\mathcal{H}}^{(i+1)}$ depending on the bit b in SymComb . Clearly, we use SymComb for encryption and decryption with the i -th used key. The main problem in showing that the simulation is correct is what we do if the two hybrid systems (both or none) use the key in other operations. In spite of our assumptions this is not impossible, e.g., they may put it into a list and send it over a secure channel. Thus we let $C'_{\mathcal{H}}^{(i)}$ choose its own key for these operations, independently of the key chosen in SymComb . The main task is to show that this does not make the simulation indistinguishable.

6.2.5. The Bisimulations Finally we have to show how the bisimulations of the original cryptographic library are extended for symmetric encryption. This corresponds to Steps 5a and 5b in Figure 3. The bisimulations are now map-

plings from $C_{\mathcal{H}}$ to $M_{\mathcal{H}}$ and from $C_{\mathcal{H}}^*$ to $\text{THSim}_{\mathcal{H}}$, called derivations in [5] because they essentially extract a part of the combined system again. As the initial and final combined systems both equal the original combined system on entries not belonging to symmetric encryption, these can both be extensions of the derivations from [5]. This is a tedious part of the proof without much novelty for symmetric encryption.

7. Conclusion

We have presented a provably secure idealization of symmetric encryption within the Dolev-Yao style cryptographic library from [5], which allows for cryptographically sound security proofs in an entirely abstract way accessible to current automated proof tools. Security holds under arbitrary attacks and in arbitrary contexts, and is based on the standard definition of authenticated encryption.

The benefit of adding symmetric encryption to the cryptographic library is impressive: Now 42 of the 50 protocols of the Clark-Jacob library can be expressed with the operations and constraints of the cryptographic library, while only 12 protocols could be expressed before. Among the remaining eight protocols, only one is excluded because of the commitment problem, five require hash functions (although one might already model some of them by message authentication codes), and two require number-theoretic operations like exponentiation and exclusive or.

References

- [1] M. Abadi and J. Jürjens. Formal eavesdropping and its computational interpretation. In *Proc. 4th International Symposium on Theoretical Aspects of Computer Software*, pages 82–94, 2001.
- [2] M. Abadi and P. Rogaway. Reconciling two views of cryptography: The computational soundness of formal encryption. In *Proc. 1st IFIP International Conference on Theoretical Computer Science*, volume 1872 of *LNCS*, pages 3–22. Springer, 2000.
- [3] M. Backes and B. Pfitzmann. A cryptographically sound security proof of the Needham-Schroeder-Lowe public-key protocol. In *Proc. 23rd Conference on Foundations of Software Technology and Theoretical Computer Science*, pages 1–12, 2003.
- [4] M. Backes and B. Pfitzmann. Symmetric encryption in a simulatable Dolev-Yao style cryptographic library. IACR Cryptology ePrint Archive 2004/059, Feb. 2004.
- [5] M. Backes, B. Pfitzmann, and M. Waidner. A composable cryptographic library with nested operations (extended abstract). In *Proc. 10th ACM Conference on Computer and Communications Security*, pages 220–230, 2003. Full version in IACR Cryptology ePrint Archive 2003/015.
- [6] M. Backes, B. Pfitzmann, and M. Waidner. Symmetric authentication within a simulatable cryptographic library. In *Proc. 8th European Symposium on Research in Computer Security*, volume 2808 of *LNCS*, pages 271–290. Springer, 2003.
- [7] M. Backes, B. Pfitzmann, and M. Waidner. A general composition theorem for secure reactive systems. In *Proc. 1st Theory of Cryptography Conference*, pages 336–354, 2004.
- [8] M. Backes, B. Pfitzmann, and M. Waidner. Secure asynchronous reactive systems. IACR Cryptology ePrint Archive 2004/082, Mar. 2004.
- [9] D. Beaver. Plug and play encryption. In *Proc. CRYPTO '97*, volume 1294 of *LNCS*, pages 75–89. Springer, 1997.
- [10] D. Beaver and S. Haber. Cryptographic protocols provably secure against dynamic adversaries. In *Proc. EUROCRYPT '92*, volume 658 of *LNCS*, pages 307–323. Springer, 1992.
- [11] M. Bellare and C. Namprempre. Authenticated encryption: Relations among notions and analysis of the generic composition paradigm. In *Proc. ASIACRYPT 2000*, volume 1976 of *LNCS*, pages 531–545. Springer, 2000.
- [12] M. Bellare and P. Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *Proc. 1st ACM Conference on Computer and Communications Security*, pages 62–73, 1993.
- [13] M. Bellare and P. Rogaway. Encode-then-encipher encryption: How to exploit nonces or redundancy in plaintexts for efficient constructions. In *Proc. ASIACRYPT 2000*, volume 1976 of *LNCS*, pages 317–330. Springer, 2000.
- [14] J. Benaloh and D. Tuinstra. Uncoercible communication. Computer Science Technical Report TR-MCS-94-1, Clarkson University, 1994.
- [15] R. Canetti, U. Feige, O. Goldreich, and M. Naor. Adaptively secure multi-party computation. In *Proc. 28th Annual ACM Symposium on Theory of Computing*, pages 639–648, 1996.
- [16] R. Canetti, O. Goldreich, and S. Halevi. The random oracle methodology, revisited. In *Proc. 30th Annual ACM Symposium on Theory of Computing*, pages 209–218, 1998.
- [17] J. Clark and J. Jacob. A survey of authentication protocol literature: Version 1.0, Nov. 1997. <http://www.cs.york.ac.uk/~jac/papers/drareview.ps.gz>.
- [18] I. Damgård and J. B. Nielsen. Improved non-committing encryption schemes based on a general complexity assumption. In *Proc. CRYPTO 2000*, volume 1880 of *LNCS*, pages 432–450. Springer, 2000.
- [19] D. Dolev and A. C. Yao. On the security of public key protocols. *IEEE Transactions on Information Theory*, 29(2):198–208, 1983.
- [20] J. Herzog. *Computational Soundness of Formal Adversaries*. PhD thesis, MIT, 2002.
- [21] J. Herzog, M. Liskov, and S. Micali. Plaintext awareness via key registration. In *Proc. CRYPTO 2003*, volume 2729 of *LNCS*, pages 548–564. Springer, 2003.
- [22] R. Impagliazzo and B. M. Kapron. Logics for reasoning about cryptographic constructions. In *Proc. 44th Symposium on Foundations of Computer Science*, pages 372–381, 2003.
- [23] P. Laud. Symmetric encryption in automatic analyses for confidentiality against active adversaries. Manuscript, 2004.
- [24] P. Laud. Semantics and program analysis of computationally secure information flow. In *Proc. 10th European Symposium on Programming*, pages 77–91, 2001.
- [25] P. Lincoln, J. Mitchell, M. Mitchell, and A. Scedrov. A probabilistic poly-time framework for protocol analysis. In *Proc. 5th ACM Conference on Computer and Communications Security*, pages 112–121, 1998.
- [26] D. Micciancio and B. Warinschi. Soundness of formal encryption in the presence of active adversaries. In *Proc. 1st Theory of Cryptography Conference*, volume 2951 of *LNCS*, pages 133–151. Springer, 2004.
- [27] J. Millen. On the freedom of decryption. *Information Processing Letters*, 86(6):329–333, June 2003.
- [28] J. Mitchell, M. Mitchell, and A. Scedrov. A linguistic characterization of bounded oracle computation and probabilistic polynomial time. In *Proc. 39th Symposium on Foundations of Computer Science*, pages 725–733, 1998.
- [29] J. Mitchell, M. Mitchell, A. Scedrov, and V. Teague. A probabilistic polynomial-time process calculus for analysis of cryptographic protocols (preliminary report). *Electronic Notes in Theoretical Computer Science*, 47:1–31, 2001.
- [30] B. Pfitzmann and M. Waidner. A model for asynchronous reactive systems and its application to secure message transmission. In *Proc. 22nd IEEE Symposium on Security & Privacy*, pages 184–200, 2001.
- [31] P. Rogaway, M. Bellare, J. Black, and T. Krovetz. OCB: A block-cipher mode of operation for efficient authenticated encryption. In *Proc. 8th ACM Conference on Computer and Communications Security*, pages 196–205, 2001.
- [32] B. Warinschi. A computational analysis of the Needham-Schroeder-(Lowe) protocol. In *Proc. 16th IEEE Computer Security Foundations Workshop*, pages 248–262, 2003.
- [33] A. C. Yao. Theory and applications of trapdoor functions. In *Proc. 23rd Symposium on Foundations of Computer Science*, pages 80–91, 1982.