

A Model for Asynchronous Reactive Systems and its Application to Secure Message Transmission

Birgit Pfitzmann

Universität des Saarlandes
Saarbrücken, Germany
pfitzmann@cs.uni-sb.de

Michael Waidner

IBM Zurich Research Laboratory
Rüschlikon, Switzerland
wmi@zurich.ibm.com

Abstract

We present a rigorous model for secure reactive systems in asynchronous networks with a sound cryptographic semantics, supporting abstract specifications and the composition of secure systems. This enables modular proofs of security, which is essential in bridging the gap between the rigorous proof techniques of cryptography and tool-supported formal proof techniques.

The model follows the general simulatability approach of modern cryptography. A variety of network structures and trust models can be described, such as static and adaptive adversaries; some examples of this are given.

As an example of our specification methodology we provide an abstract and complete specification for Secure Message Transmission, improving on recent results by Lynch, and verify one concrete implementation. Our proof is based on a general theorem on the security of encryption in a reactive multi-user setting, generalizing a recent result by Bellare et al.

1. Introduction

In the early days of security research, cryptographic protocols were designed using a simple iterative process: someone proposed a protocol, someone else found an attack, an improved version was proposed, and so on, until no further attacks were found. Today it is commonly accepted that this approach gives no security guarantee. Too many seemingly simple and secure protocols have been found flawed over the years. Moreover, problems like n -party key agreement, fair contract signing, anonymous communication, electronic auctions or payments are just too complex for this approach. Secure protocols—or more generally, secure reactive systems—need a proof of security before being acceptable.

Both the cryptographic and the formal-methods communities are working on such proofs. The former aims at complete and mathematically rigorous proofs, while the latter aims at proofs in some formal proof system that can be automatically verified or even generated. Unfortunately, current formal methods in security cannot be applied directly to cryptographic proofs. Instead they abstract from most cryptographic details, and therefore there is no guarantee that a formally proven protocol is actually secure if implemented with a cryptographically secure primitive [1, 26].

One of our goals is to link both approaches to get the best overall results: proofs that allow abstraction and the use of formal methods, but retain a sound cryptographic semantics. Thus we provide a model that allows us to split reactive systems into two layers: The lower layer is a cryptographic system whose security can be rigorously proven using standard cryptographic arguments. To the upper layer it provides an abstract (and typically deterministic) service that hides all cryptographic details. Relative to this abstract service one can verify the upper layer using existing formal methods. Since our model allows secure composition (as shown in Theorem 4.1) one can conclude that the overall system is secure if the formally verified upper layer is put on top of a cryptographically verified lower layer ([26] provides more motivation for this approach).

In the following, we carry out this approach specifically for asynchronous reactive systems. Reactive means that the system interacts with its user many times (e.g., multiple subprotocol executions). Essentially, we describe a system by sets of asynchronously communicating probabilistic state machines, connected via buffered channels. Honest users and the adversary are explicitly represented by two arbitrary machines, H and A , which can interact arbitrarily. A reactive system, Sys_0 , is considered *at least as secure as* another system, Sys_1 , written as $Sys_0 \geq Sys_1$, if whatever any adversary A_0 can do to any honest user H in Sys_0 ,

some adversary A_1 can do to the same H in Sys_1 essentially with the same probability. System Sys_0 is often a real system using concrete cryptographic primitives, while Sys_1 is an ideal system, i.e., a specification, that does not depend on any specific cryptographic implementation details and is not realistic (e.g., one trusted machine), but secure by construction.

The mechanics of our model are defined in Section 2. Section 3 shows how to represent typical trust models (or adversary structures), such as static threshold models and adaptive adversaries, with secure, authenticated and insecure channels. In Section 4 we show that secure systems can be composed, i.e., our model supports modular security proofs as outlined above.

In Section 5 we study Secure Message Transmission as an example. We follow two main design principles:

1. *The ideal system should provide abstract interfaces, hiding all cryptographic details.* This keeps the specification independent of the implementation, which is desirable when higher-layer protocols are based on the service. For instance, in order to send messages secretly from one user to another there is no need to ask the user to input cryptographic keys or to output ciphertexts to him; those can be generated, exchanged and processed completely within the system.¹

2. *The ideal system needs to explicitly specify all tolerable imperfections.* In order to improve efficiency one often accepts certain imperfections. For instance, a typical practical implementation of secure message transmission only conceals the contents of messages, but does not hide who communicates with whom, which would be much more costly to implement. In a simulatability-based approach one has to include all such tolerable imperfections in the specification of an ideal system, and they should also be abstract.

The proof of the real system in Section 5 uses a theorem (Theorem 5.2) that extends the security of public-key encryption in multi-user settings, which might be of independent interest. It captures what is often called a “standard hybrid argument” in cryptography, and generalizes a result from [4].

Related Literature. Several researchers pursue the goal of providing security proofs that allow the use of formal methods, but retain a sound cryptographic semantics: In [21, 22] the cryptographic security of specific systems is directly defined and verified using a formal language (π -calculus), but without providing abstractions (their specifications essentially comprise the actual protocols including all cryptographic details) and without tool support (as even

¹In a simulatability-based definition it would not even be possible to have keys in the interface and to be implementation-independent because keys of different implementations are distinguishable.

the specifications involve ad-hoc notations, e.g., for generating random primes). [24] has quite a similar motivation to our paper. However, cryptographic systems are restricted to the usual equational specifications (following [11]) and the semantics is not probabilistic. Hence the abstraction from cryptography is no more faithful than in other papers on formal methods in security. Moreover, only passive adversaries are considered and only one class of users (“environment”). The author actually remarks that the model of what the adversary learns from the environment is not yet general, and that general theorems for the abstraction from probabilism would be useful. Our model solves both these problems. In [1] it is shown that a slight variation of the standard Dolev-Yao abstraction [11] is cryptographically faithful specifically for symmetric encryption, but only under passive attacks.

Our security definitions follow the general simulatability approach of modern cryptography, which was first used in secure function evaluation [3, 15, 25, 34], and subsequently also for specific reactive problems (e.g., [5, 10, 12]) and for the construction of generic solutions for large classes of reactive problems [19, 14, 20] (usually yielding inefficient solutions and assuming that all parties take part in all subprotocols). General models for reactive systems have been proposed (after some earlier sketches, in particular in [19, 29, 8]) in [21, 22, 20, 27, 30]. The last three are synchronous, while the first two are in a somewhat simplified timing model with uniform choice among certain classes of unrelated possible events. Among the reactive models, the only composition theorem so far is in [30], i.e., we present the first asynchronous one in the current paper. Our model is based on [27, 30], except for the timing aspects. Those can be seen as extensions of [33, 7, 22], see Section 2.

Independently and concurrently to this work, Ran Canetti developed a model that roughly corresponds to standard cryptographic systems as discussed in Section 3.2, with polynomial-time users and adversaries, and authenticated channels only [9]. The model is less rigorously defined than the model presented here. Security is defined in terms of universal simulatability only (see Definition 2.12), which allows to securely compose a polynomial number of identical systems.

Several specifications for secure message transmissions have been proposed, as examples of general models. The specification in [21] is formal but specific for one concrete protocol and comprises all cryptographic details, i.e., it is not abstract and its intuitive correctness is relatively difficult to verify. Our concrete specification is quite close to that in [24], but we had to introduce the tolerable imperfections. Actually, the implementation in [24] has the same imperfections. They do not show up in the proof because the definition of “a system implements another one,” used in place of our “as secure as” is weaker: Secrecy is defined as

a property that the adversary cannot learn certain messages, but here the information leaked is not entire messages.

2. Asynchronous Reactive Systems

In this section, we present our model for secure reactive systems in an asynchronous network.

Our machine model is probabilistic state-transition machines, similar to probabilistic I/O automata as sketched in [23] (more details in [33]). A distinguishing feature in our model of asynchronous executions is distributed scheduling.

The standard way to fix the order of events in an asynchronous system of probabilistic I/O automata is a probabilistic scheduler that has full information about the system [33]. The “standard” understanding in cryptology (closest to a rigorous definition in [7]) is that the adversary schedules everything, but only with realistic information. This corresponds to making a certain subclass of schedulers explicit for the model from [33]. However, if one splits up a machine into local submachines, or defines intermediate systems for the purposes of proof only, this may introduce many schedules that do not correspond to a schedule of the original system and therefore just complicate the proofs. (The proof in Section 5 is of this type.) Our solution to this is a distributed definition of scheduling which allows machines that have been scheduled to schedule certain (statically fixed) other machines themselves. This does not weaken the adversary’s power in real systems, because our definition of standard cryptographic systems in Section 3 will not use this feature except for scheduling local submachines.

Similar problems with purely adversarial scheduling were already noted in [22]. They distinguish secure channels and schedule all those with uniform probability before adversary-chosen events. However, that introduces a certain amount of global synchrony. Furthermore, we do not require “local” scheduling for all secure channels; they may be blindly scheduled by the adversary (i.e., without even seeing if there are messages on the channel). For instance, this models cases where the adversary has a global influence on relative network speed.

2.1. General System Model

We now define a model of machines and executions of collections of machines.

Let a finite alphabet Σ be given, let Σ^* denote the strings over it, ϵ the empty string, and $\Sigma^+ := \Sigma^* \setminus \{\epsilon\}$. We assume that $!, ?, \leftrightarrow, \triangleleft \notin \Sigma$.

Definition 2.1 (Ports) A port p is a triple $(n, l, d) \in \Sigma^+ \times \{\epsilon, \leftrightarrow, \triangleleft\} \times \{!, ?\}$. We call $\text{name}(p) := n$ its name,

$\text{label}(p) := l$ its label, and $\text{dir}(p) := d$ its direction. We can write the triples as concatenations without ambiguity.

We call a port (n, l, d) an in-port or out-port iff $d = ?$ or $d = !$, respectively. We call it a simple port, buffer port or clock port iff $l = \epsilon, \leftrightarrow$, or \triangleleft , respectively.

For a set P of ports let $\text{out}(P) := \{p \in P \mid \text{dir}(p) = !\}$ and $\text{in}(P) := \{p \in P \mid \text{dir}(p) = ?\}$. We use the same notation for sequences of ports (retaining the order).

By p^c , the (low-level) *complement* of a port p , we denote the port with which it connects according to Figure 1, i.e., $n^{\triangleleft!c} := n^{\triangleleft?}$, $n^{!c} := n^{\leftrightarrow?}$, and $n^{\leftrightarrow!c} := n^?$ and vice versa. Accordingly we define the complement of a set or sequence of ports.

For simple ports, we also define p^C , the high-level complement, as the port connected to p without counting the buffer, i.e., $n^{?C} := n!$ and vice versa. \diamond

Definition 2.2 (Machines and Schedulers) A machine is a tuple $M = (\text{name}_M, \text{Ports}_M, \text{States}_M, \delta_M, \text{Ini}_M, \text{Fin}_M)$ of a name $\text{name}_M \in \Sigma^+$, a finite sequence Ports_M of ports, a set $\text{States}_M \subseteq \Sigma^*$ of states, a probabilistic state-transition function δ_M , and sets $\text{Ini}_M, \text{Fin}_M \subseteq \text{States}_M$ of initial and final states.² The inputs are tuples $I = (I_i)_{i=1, \dots, |\text{in}(\text{Ports}_M)|}$, where $I_i \in \Sigma^*$ is the input for the i -th in-port. Analogously, the outputs are tuples $O = (O_i)_{i=1, \dots, |\text{out}(\text{Ports}_M)|}$.³ The empty word, ϵ , denotes “no in- or output.”

δ_M maps each pair (s, I) of a state and an input to a finite distribution over pairs (s', O) . If $s \in \text{Fin}_M$ or $I = (\epsilon, \dots, \epsilon)$, then $\delta_M(s, I) = (s, (\epsilon, \dots, \epsilon))$ deterministically.

Let $\text{ports}(M)$ denote the set of ports in Ports_M , and for a set \hat{M} of machines, let $\text{ports}(\hat{M}) = \bigcup_{M \in \hat{M}} \text{Ports}_M$. A machine M is *simple* if it has only simple ports and clock out-ports. A machine M is called *master scheduler* if it has only simple ports and clock out-ports and the special *master-clock* in-port $\text{clk}^{\triangleleft?}$.

For computational aspects, a machine is regarded as implemented by a probabilistic interactive Turing machine [18], where each port is a communication tape. Its complexity is measured in terms of the length of its initial state, represented as initial worktape content (often a security parameter). Termination respects the atomicity of transactions, at least with respect to outputs. \diamond

Definition 2.3 (Buffers) For each name $q \in \Sigma^+$ we define a specific machine \tilde{q} , called a *buffer*: \tilde{q} has three ports, $q^{\triangleleft?}$, $q^{\leftrightarrow?}$, $q^{\leftrightarrow!}$ (clock, in, and out) (see Figure 1). The internal state of \tilde{q} is a queue over Σ^+ with random access. Initially it is empty, and the set of final states is empty.

²We often use “ M ” also as the name name_M of M .

³This representation makes δ_M independent of the port names. Below, we also define the view of a machine independently of the port names. Hence we can rename ports in some proofs without changing the views.




Figure 1. Ports and buffers.

For each state transition, if $q^{\leftarrow?}$ carries a non-empty input x then $\delta_{\tilde{q}}$ appends this input x to the queue. If $q^{\leftarrow!}$ carries a non-empty input, it is interpreted as a number $i \in \mathbb{N}$ and the i -th element is retrieved (where 1 indicates the oldest one), removed from the queue, and output at $q^{\leftarrow!}$. (This might be the element just appended.) If there are fewer elements, the result is ϵ .

A polynomial-time buffer \tilde{q} accepts only a polynomial number of inputs, and truncates each input after a polynomial number of symbols before appending it to its queue.⁴ \diamond

Definition 2.4 (Collections) A collection C is a finite set of machines with pairwise different machine names, disjoint sets of ports, and where all machines are simple or master schedulers or buffers. It is called polynomial-time if all its machines are polynomial-time.

Each set of low-level complementary ports $\{p, p^c\} \subseteq \text{ports}(C)$ is called a *low-level connection*, and the set of them the *low-level connection graph* $\text{gr}(C)$. By $\text{free}(C)$ we denote the *free* ports in this graph, i.e., $\text{ports}(C) \setminus \text{ports}(C)^c$. A set of high-level complementary simple ports $\{p, p^C\} \subseteq \text{ports}(C)$ is called a *high-level connection*, and the set of them the *high-level connection graph* $\text{Gr}(C)$.

A collection is *closed* if $\text{free}(C) = \{\text{clk}^{\leftarrow?}\}$.

The *completion* $[C]$ of a collection C is the smallest collection that comprises C and the corresponding buffer for each simple or clock out-port $q \in \text{ports}(C)$.

If $\tilde{q}, M \in C$ and $q^{\leftarrow!} \in \text{ports}(M)$ then we call M the scheduler for buffer \tilde{q} (in C). \diamond

Now we define the probability space of runs (or “executions” or “traces”) of a closed collection.

Definition 2.5 (Runs) Given a closed collection C with master scheduler X and a tuple $\text{ini} \in \text{Ini} := \times_{M \in C} \text{Ini}_M$ of initial states, the probability space of *runs* is defined as follows:

⁴The size of messages received from the adversary might be unlimited. Hence the buffer’s view is not naturally of polynomial size. By letting \tilde{q} truncate these inputs, all other machines receive only polynomial-length inputs, i.e., this technical problem arises only for buffers.

Each run is a sequence of *steps* inductively defined by the following algorithm. The algorithm uses a variable curr_sched over machine names; initially $\text{curr_sched} := X$. It also treats each port like a variable over Σ^* . All ports are initialized with ϵ , except $\text{clk}^{\leftarrow?} := 1$. Probabilistic choices occur in Phase (2) only and are made by the machines themselves.

1. **Termination:** If X is in a final state, the run *stops*.
2. **Switch current scheduler:** We switch machine $M := \text{curr_sched}$ and then assign ϵ to all in-ports of M (which might include $\text{clk}^{\leftarrow?}$).
3. **Save messages from M in buffers:** We switch each buffer \tilde{p} where $p!$ is a simple out-port of M , in the given order of these ports, with inputs $p^{\leftarrow?} := \epsilon$ and $p^{\leftarrow!} := p!$. Then we assign ϵ to all these ports $p!$ and $p^{\leftarrow?}$.
4. **Determine next scheduler:** If at least one clock out-port of M carries a value $\neq \epsilon$, then let $q^{\leftarrow!}$ denote the first such port, and M' the unique machine with $q^{\leftarrow?} \in \text{ports}(M')$. We set $\text{curr_sched} := M'$ and assign ϵ to all other clock out-ports of M .
Otherwise we set $\text{curr_sched} := X$ and $\text{clk}^{\leftarrow?} := 1$ and go back to Phase (1).
5. **Retrieve scheduled message for M' :** We switch \tilde{q} with input $q^{\leftarrow?} := q^{\leftarrow!}$ and $q^{\leftarrow!} := \epsilon$, set $q^{\leftarrow?} := q^{\leftarrow!}$ and assign ϵ to all ports of \tilde{q} , and to $q^{\leftarrow!}$. We go back to Phase (1).

Whenever a machine (this may be a buffer) with name name_M is switched from (s_M, I_M) to (s'_M, O_M) with non-final s_M and non-empty input I_M , we add $(\text{name}_M, s_M, I_M, s'_M, O_M)$ to the run. This gives a family of random variables

$$\text{run}_C = (\text{run}_{C, \text{ini}})_{\text{ini} \in \text{Ini}}.$$

For a number $l \in \mathbb{N}$, l -step prefixes $\text{run}_{C, \text{ini}, l}$ of runs are defined in the obvious way. For a function $l : \text{Ini} \rightarrow \mathbb{N}$, this gives a family $\text{run}_{C, l} = (\text{run}_{C, \text{ini}, l(\text{ini})})_{\text{ini} \in \text{Ini}}$. \diamond

Definition 2.6 (View) The view of a subset \hat{M} of a closed collection C in a run r is the restriction of r to \hat{M} . Restriction means that all steps $(\text{name}, s, I, s', O)$ where name is the name of a machine $M \notin \hat{M}$ are deleted from r .

This gives a family of random variables

$$\text{view}_C(\hat{M}) = (\text{view}_{C, \text{ini}}(\hat{M}))_{\text{ini} \in \text{Ini}},$$

and similarly for l -step prefixes. \diamond

2.2. Security-Specific System Model

Now we define specific collections for security purposes, first the system part and then the environment, i.e., users and adversaries. Typically, a cryptographic system is described by an intended structure, and the actual structures are derived using a trust model (see Section 3). However, as a wide range of trust models is possible, we keep the remaining definitions independent of them.

Definition 2.7 (*Structures and Systems*) A *structure* is a pair $struc = (\hat{M}, S)$ where \hat{M} is a collection of simple machines called *correct machines*, and $S \subseteq \text{free}([\hat{M}])$ is called *specified ports*. Let $\bar{S} := \text{free}([\hat{M}]) \setminus S$ and $\text{forb}(\hat{M}, S) := \text{ports}(\hat{M}) \cup \bar{S}^c$. A *system* Sys is a set of structures. A system is polynomial-time iff all its collections \hat{M} are. \diamond

The separation of the free ports into specified ports and others is an important feature of our particular reactive simulatability definitions. The specified ports are those where a certain service is guaranteed. Typical examples of inputs at specified ports are “send message m to id ” for a message transmission system or “pay amount x to id ” for a payment system. The ports in \bar{S} are additionally available for the adversary. The ports in $\text{forb}(\hat{M}, S)$ will therefore be forbidden for an honest user to have. In the simulatability definition below, only the events at specified ports have to be simulated one by one. This allows *abstract* specifications of systems with *tolerable imperfections*. See Section 5 for an example (and [27, 28] for more).

Definition 2.8 (*Configuration*) A *configuration* $conf$ of a system Sys is a tuple (\hat{M}, S, H, A) where $(\hat{M}, S) \in Sys$ is a structure, H is a simple machine without forbidden ports, $\text{ports}(H) \cap \text{forb}(\hat{M}, S) = \emptyset$, and the completion $C := [\hat{M} \cup \{H, A\}]$ is a closed collection with master scheduler A .

The set of all configurations is written $\text{Conf}(Sys)$, and those with polynomial-time user H and adversary A are called $\text{Conf}_{\text{poly}}(Sys)$. “poly” is omitted if it is clear from the context. Runs and views of a configuration are defined as those of C , see Definitions 2.5 and 2.6.

Typically, the initial states of all machines are only a security parameter k (in unary representation). Then we consider the families of runs and views restricted to the subset $\text{Ini}^l = \{(1^k)_{M \in C} \mid k \in \mathbb{N}\}$ of Ini , and write run_{conf} and $\text{view}_{conf}(\hat{M})$ for run_C and $\text{view}_C(\hat{M})$ restricted to Ini^l , and similar for l -step prefixes. Furthermore, Ini^l is identified with \mathbb{N} ; hence one can write $\text{run}_{conf,k}$ etc. \diamond

Remark 2.1. The condition on the ports of H in a configuration is equivalent to $\text{ports}(H) \cap \text{ports}(\hat{M}) = \emptyset$ and $\text{ports}(H)^c \cap \text{ports}([\hat{M}]) \subseteq S$. \circ

2.3. Simulatability

We now define the security of a system Sys_1 relative to another system Sys_2 . Typically, we only want to compare each structure of Sys_1 with certain corresponding structures in Sys_2 . What “corresponding” means can be specified by a mapping f ; we just require that only structures with the same set of specified ports correspond. An instantiation of f is usually derived from the trust model, see Section 3.

Definition 2.9 (*Valid Mapping*) A function f from a system Sys_1 to the powerset of a system Sys_2 is called a *valid mapping* if $S_1 = S_2$ for all structures (\hat{M}_1, S_1) and $(\hat{M}_2, S_2) \in f(\hat{M}_1, S_1)$.

Given Sys_2 and f , the set $\text{Conf}^f(Sys_1)$ contains those configurations $(\hat{M}_1, S, H, A_1) \in \text{Conf}(Sys_1)$ where $\text{ports}(H) \cap \text{forb}(\hat{M}_2, S) = \emptyset$ for all $(\hat{M}_2, S) \in f(\hat{M}_1, S)$; we call them *suitable configurations*. \diamond

Remark 2.2. For a valid mapping f , we have $S^c \cap \text{forb}(\hat{M}_i, S) = \emptyset$ for $i = 1, 2$, i.e., the ports that users are intended to use are not at the same time forbidden (also not in the corresponding structures of the other system).⁵

Remark 2.3. With regard to Sys_1 alone, the restriction to suitable configurations is w.l.o.g.: For every $conf_1 = (\hat{M}_1, S, H, A_1) \in \text{Conf}(Sys_1) \setminus \text{Conf}^f(Sys_1)$, there is a configuration $conf_{f,1} = (\hat{M}_1, S, H_f, A_{f,1}) \in \text{Conf}^f(Sys_1)$ such that $\text{view}_{conf_{f,1}}(H_f) = \text{view}_{conf_1}(H)$. \circ

For two families $(\text{var}_k)_{k \in \mathbb{N}}$ and $(\text{var}'_k)_{k \in \mathbb{N}}$ of random variables (or probability distributions) let “=”, “ \approx_{SMALL} ”, “ \approx_{poly} ” denote perfect indistinguishability, statistical indistinguishability (for a class $SMALL$ of functions from \mathbb{N} to $\mathbb{R}_{\geq 0}$), and computational indistinguishability, respectively [35]. We write \approx if we want to treat all cases together. The following definition captures that whatever an adversary can achieve in the real system against certain honest users, another adversary can achieve against the same honest users in the ideal system. A typical situation is illustrated in Figure 2.

Definition 2.10 (*Simulatability*) Let systems Sys_1 and Sys_2 with a valid mapping f be given.

- a) We say $Sys_1 \geq_{\text{sec}}^{f, \text{perf}} Sys_2$ (*perfectly at least as secure as*) if for every configuration $conf_1 = (\hat{M}_1, S, H, A_1) \in \text{Conf}^f(Sys_1)$, there exists a configuration $conf_2 = (\hat{M}_2, S, H, A_2) \in \text{Conf}(Sys_2)$ with $(\hat{M}_2, S) \in f(\hat{M}_1, S)$ (and the same H) such that

$$\text{view}_{conf_1}(H) = \text{view}_{conf_2}(H).$$

⁵In [30], the condition on a valid mapping is a generalization of this condition to mappings that allow different sets S_1 and S_2 and more general users. The stronger requirements here simplify the presentation and are sufficient for all cryptographic examples we considered. See [27] for a non-cryptographic example with $S_1 \neq S_2$.




Figure 2. Example of simulatability.

- b) We say $Sys_1 \geq_{\text{sec}}^{f, \text{SMALL}} Sys_2$ (statistically at least as secure as) for a class *SMALL* if the same as in a) holds with statistical indistinguishability of all families $view_{conf_1, l}(H)$ and $view_{conf_2, l}(H)$ of l -step prefixes of the views for polynomials l .
- c) We say $Sys_1 \geq_{\text{sec}}^{f, \text{poly}} Sys_2$ (computationally at least as secure as) if the same as in a) holds with configurations from $Conf_{\text{poly}}^f(Sys_1)$ and $Conf_{\text{poly}}(Sys_2)$ and computational indistinguishability of the families of views.

In all cases, we call $conf_2$ an indistinguishable configuration for $conf_1$. Where the difference between the types of security is irrelevant, we simply write \geq_{sec}^f , and we omit the indices f and sec if they are clear from the context. \diamond

2.4. Basic Lemmas and Blackbox Simulation

An essential ingredient in the composition theorem and other uses of the model is a notion of combining several machines into one, and a lemma that this makes no essential difference in views.

Definition 2.11 (*Combination of Machines*) Let C be a collection without buffers, and $\hat{D} \subseteq C$. We define the *combination* of \hat{D} into one machine D :

- The name of D is chosen arbitrarily, but different from the names of the machines in C , and $Ports_D := ports(\hat{D})$ (in an arbitrary order).
- D has all machines in \hat{D} as sub-machines: $States_D := \times_{M \in \hat{D}} States_M$, and δ_D is defined by applying the transition functions of all sub-machines to the corresponding substates and inputs, unless D has reached a final state (see below). In that case, δ_D does not change the state anymore and produces no output.
- The initial states are $Ini_D := \times_{M \in \hat{D}} Ini_M$. If there is a master scheduler $X \in \hat{D}$ then Fin_D is the set of all states of D where X is in a state from Fin_X . Otherwise D stops as soon as *all* sub-machines stopped: $Fin_D := \times_{M \in \hat{D}} Fin_M$. \diamond

Lemma 2.1 (*Combination of Machines*) With the notation of Definition 2.11:

1. If $[C]$ is closed then $[C \setminus \hat{D} \cup \{D\}]$ is closed as well.
2. The view of any set of original machines in $C \setminus \hat{D} \cup \{D\}$ is the same as in C . This includes the view of the submachines in D , which is well-defined (given C and \hat{D}).
3. If all machines in \hat{D} are poly-time, then so is D . \square

Definition 2.11 allows us to add the notion of blackbox simulatability to Definition 2.10:

Definition 2.12 (*Universal and Blackbox Simulatability*) Universal simulatability means that A_2 in Definition 2.10 does not depend on H (only on \hat{M}_1 , S , and A_1).

Blackbox simulatability means that A_2 is a simulator Sim with a machine A'_1 as a blackbox sub-machine, where A'_1 is identical to A_1 except for renamed and relabeled ports, and Sim depends at most on \hat{M}_1 , S and $ports(A_1)$. \diamond

Lemma 2.2 (*Transitivity*) If $Sys_1 \geq^{f_1} Sys_2$ and $Sys_2 \geq^{f_2} Sys_3$, then $Sys_1 \geq^{f_3} Sys_3$, where $f_3 := f_2 \circ f_1$ is defined in a natural way as follows: $f_3(\hat{M}_1, S)$ is the union of the sets $f_2(\hat{M}_2, S)$ with $(\hat{M}_2, S) \in f_1(\hat{M}_1, S)$. This holds for perfect, statistical and computational security, and also for universal and blackbox simulatability. \square

These lemmas are proven in the full version [31].

3. Standard Cryptographic Systems

We are now ready to define the specific class of standard cryptographic systems. The intuition behind this class is that in a real system Sys , there is one machine per human system participant, and each machine is correct if and only if its owner is honest. The system is derived from an *intended structure* (\hat{M}^*, S^*) and a *trust model*. We consider two variants: static and adaptive adversaries.

3.1. Systems with Static Adversary Model

We define that all buffers that connect different machines are scheduled by the adversary. We only allow M_u to schedule buffers that transport messages from itself to itself, and require all these connections to be secure: this allows us to define a machine M_u as a combination of (local) sub-machines.

Definition 3.1 (*Standard Cryptographic Systems*) A *standard cryptographic structure* is a structure (\hat{M}^*, S^*) where $\hat{M}^* = \{M_1, \dots, M_n\}$ and $S^{*c} = \{in_u!, out_u? \mid u =$

$1, \dots, n\}$, where $\text{in}_u?$ and $\text{out}_u!$ are ports of machine M_u .⁶ Each machine M_u is simple, and for all names p , if $p^a! \in \text{ports}(M_u)$ then $p?, p! \in \text{ports}(M_u)$.

A *standard trust model* for such a structure consists of an access structure, \mathcal{ACC} , and a channel model, χ . \mathcal{ACC} is a set of subsets \mathcal{H} of $\{1, \dots, n\}$, closed under insertion, and denotes the possible sets of correct machines.⁷ χ is a mapping $\chi : \text{Gr}(\hat{M}^*) \rightarrow \{s, a, i\}$. It characterizes each high-level connection as secure (private and authentic), authenticated (only authentic), or insecure (neither private nor authentic). We require that if a connection c connects a machine M_u with itself then $\chi(c) = s$.

Given such a structure and trust model, the corresponding *standard cryptographic system* is $Sys := \{(\hat{M}_{\mathcal{H}}, S_{\mathcal{H}}) | \mathcal{H} \in \mathcal{ACC}\}$ with $S_{\mathcal{H}}^c := \{\text{in}_u!, \text{out}_u? | u \in \mathcal{H}\}$ and $\hat{M}_{\mathcal{H}} := \{M_{u, \mathcal{H}} | u \in \mathcal{H}\}$, where $M_{u, \mathcal{H}}$ is derived from M_u as follows:

- The ports $\text{in}_u?$ and $\text{out}_u!$ and all clock ports are unchanged.
- Consider a simple port $p \in \text{ports}(M_u) \setminus \{\text{in}_u?, \text{out}_u!\}$, where $p^C \in \text{ports}(M_v)$ with $v \in \mathcal{H}$, i.e., $c = \{p, p^C\}$ is a high-level connection between two correct machines:
 - If $\chi(c) = s$ (secure), p is unchanged.
 - If $\chi(c) = a$ (authenticated) and p is an output port, $M_{u, \mathcal{H}}$ gets an additional new port p^d , where it duplicates the outputs at p .⁸ This port automatically remains free, and thus the adversary connects to it. If p is an input port, it is unchanged.
 - If $\chi(c) = i$ (insecure) and p is an input port, p is replaced by a new port p^a . (Thus the adversary can get the outputs from p^C and make the inputs to p^a and thus completely control the connection.) If p is an output port, it is unchanged.
- Consider a simple port $p \in \text{ports}(M_u) \setminus \{\text{in}_u?, \text{out}_u!\}$, where $p^C \notin \text{ports}(M_v)$ for all $v \in \mathcal{H}$: If p is an output port, it is unchanged. If it is an input port, it is renamed into p^a . (In both cases the adversary can connect to it.)
◇

A typical ideal system is of the form $Sys_2 = \{(\{\text{TH}_{\mathcal{H}}\}, S_{\mathcal{H}}) | \mathcal{H} \in \mathcal{ACC}\}$ with the same sets $S_{\mathcal{H}}$ as in the corresponding real system Sys_1 . The *canonical mapping* f between such systems is defined by $f(\hat{M}_{\mathcal{H}}, S_{\mathcal{H}}) = \{(\{\text{TH}_{\mathcal{H}}\}, S_{\mathcal{H}})\}$ for all \mathcal{H} .

⁶We have specified the complement of S^* because that is independent of the buffer notation.

⁷Typical examples are threshold structures $\mathcal{ACC}_t := \{\mathcal{H} \subseteq \{1, \dots, n\} \mid |\mathcal{H}| \geq t\}$ for some t .

⁸This can be done by a trivial blackbox construction. We assume w.l.o.g. that there is a systematic naming scheme for such new ports (e.g., appending d) that does not clash with prior names.

3.2. Systems with Adaptive Adversary Model

Standard cryptographic systems as defined in the previous section are based on the intuition that it is a priori clear who are the “bad guys” and who are the “good guys.” In *adaptive* (or *dynamic*) adversary models the set of corrupted machines can increase over time, e.g., because there is a “master adversary” who has to hack into machines in order to corrupt them [6, 8]. Adaptive adversary models are more powerful than static ones, i.e., there are examples of systems secure against static adversaries that are insecure against adaptive adversaries who can corrupt the same sets of machines [8].

Standard cryptographic systems with adaptive adversary can easily be defined within our model: Such a system has only one structure, (\hat{M}, S) , derived from a single standard cryptographic structure (\hat{M}^*, S^*) . As before, \hat{M}^* is a set $\{M_1, \dots, M_n\}$ and $S^{*c} = \{\text{in}_u!, \text{out}_u? | u = 1, \dots, n\}$.

Let M'_u denote the machine derived from M_u . The derivation is done with the access structure $\mathcal{ACC} = \{\mathcal{M}\}$ for $\mathcal{M} = \{1, \dots, n\}$ (all intended machines are present), and an arbitrary channel model χ satisfying that all connections from an M_u to itself are secure.

Each M'_u receives an additional specified port, $\text{corrupt}_u?$, intended for corruption requests. (Those must be made via specified ports because the service will change at least at the corresponding ports $\text{in}_u?$, $\text{out}_u!$ also in the ideal system.) Each M'_u also has two specific ports $\text{cor_out}_u!$, $\text{cor_in}_u?$ for communication with A after corruption. A newly corrupted machine sends a predefined “corruption response” (corruption, σ) to A via $\text{cor_out}_u!$, and from then on becomes “transparent.” Any input m on $p?$ is translated into output $(p?, m)$ on $\text{cor_out}_u!$, any input $(p!, m)$ on $\text{cor_in}_u?$ is translated into output m on $p!$.⁹

Two main types of corruption responses σ are natural:

- σ is the current state of M'_u .
- σ is the entire view of M'_u . This corresponds to the assumption that nothing can be erased reliably. Thus every transition of δ_{M_u} is modified in $\delta_{M'_u}$ to store the current step.

One may also extend this to different classes of storage, e.g., to model the different vulnerability of session keys and long-term keys.

In the ideal system, the trusted host TH also accepts corruption requests. Tolerable sets of corrupted machines are defined by an access structure \mathcal{ACC}^* within TH: if the set of corruption requests is no longer in \mathcal{ACC}^* , i.e., there were “too many,” then TH sends its state to A and gives all control

⁹Knowing the lifetimes of all intended machines in \hat{M} one can easily determine the lifetime of M'_u so that it can forward messages as long as some non-corrupted machine would still be able to receive or send them.

to A. Thus after this, the ideal system no longer guarantees anything and simulation becomes trivial.

4. Composition

In this section, we show that the relation “at least as secure as” is consistent with the composition of systems. The basic idea is the following: Assume that we have proven that a system Sys_0 is as secure as another system Sys'_0 (typically an ideal system used as a specification). Now we would like to use Sys_0 as a secure replacement for Sys'_0 , i.e., as an implementation of the specification Sys'_0 .

Usually, replacing Sys'_0 means that we have another system Sys_1 that uses Sys'_0 ; we call this composition Sys^* . Inside Sys^* we want to use Sys_0 instead, which gives a composition $Sys^\#$. Hence $Sys^\#$ is typically a completely real system, while Sys^* is partly ideal. Intuitively we expect $Sys^\#$ to be at least as secure as Sys^* . The situation is shown in the left and middle part of Figure 3.




Figure 3. Composition theorem and its use in a modular proof: The left and middle part show the statement of Theorem 4.1, the right part Corollary 4.1.

The remainder of this section is quite similar to the corresponding section of [30] for the synchronous case.

We define composition for every number n of systems Sys_1, \dots, Sys_n . We do not provide a composition operator that produces one specific composition. The reason is that one typically does not want to compose every structure of one system with every structure of the others, but only with certain matching ones. For instance, if the individual machines of Sys_2 are implemented on the same physical devices as those of Sys_1 , as usual in a layered distributed system, we only compose structures corresponding to the same set of corrupted machines. However, this is not the only conceivable situation. Hence we allow many different compositions.

Definition 4.1 (Composition) The composition of structures and of systems is defined as follows: We call structures $(\hat{M}_1, S_1), \dots, (\hat{M}_n, S_n)$ *composable* if $\text{ports}(\hat{M}_i) \cap \text{forb}(\hat{M}_j, S_j) = \emptyset$ and $S_i \cap \text{free}([\hat{M}_j]) = S_j \cap \text{free}([\hat{M}_i])$

for all $i \neq j$.¹⁰ We then define their composition as $(\hat{M}_1, S_1) \parallel \dots \parallel (\hat{M}_n, S_n) := (\hat{M}, S)$ with $\hat{M} = \hat{M}_1 \cup \dots \cup \hat{M}_n$ and $S = (S_1 \cup \dots \cup S_n) \cap \text{free}([\hat{M}])$.

We call a system Sys a composition of Sys_1, \dots, Sys_n and write $Sys \in Sys_1 \times \dots \times Sys_n$ if each structure $(\hat{M}, S) \in Sys$ has a unique representation $(\hat{M}, S) = (\hat{M}_1, S_1) \parallel \dots \parallel (\hat{M}_n, S_n)$ with composable structures $(\hat{M}_i, S_i) \in Sys_i$ for $i = 1, \dots, n$.

We then call (\hat{M}_i, S_i) the restriction of (\hat{M}, S) to Sys_i and write $(\hat{M}_i, S_i) = (\hat{M}, S) \upharpoonright_{Sys_i}$. \diamond

Remark 4.1. For all compositions of structures, we have $[\hat{M}] = [\hat{M}_1] \cup \dots \cup [\hat{M}_n]$ and $\text{free}([\hat{M}]) \subseteq \text{free}([\hat{M}_1]) \cup \dots \cup \text{free}([\hat{M}_n])$. We also have $S = \text{free}([\hat{M}]) \setminus (\bar{S}_1 \cup \dots \cup \bar{S}_n)$. \circ

The following theorem shows that modular proofs are indeed possible. Recall that the situation is shown in the left and middle part of Figure 3. The main issue in formulating the theorem is to characterize $Sys^\#$, i.e., to formulate what it means that Sys_0 replaces Sys'_0 .

Theorem 4.1 (Secure Two-system Composition) Let systems Sys_0, Sys'_0, Sys_1 and a valid mapping f_0 be given with $Sys_0 \geq^{f_0} Sys'_0$. Let compositions $Sys^\# \in Sys_0 \times Sys_1$ and $Sys^* \in Sys'_0 \times Sys_1$ be given that fulfil the following structural conditions:

For every structure $(\hat{M}^\#, S) \in Sys^\#$ with restrictions $(\hat{M}_i, S_i) = (\hat{M}^\#, S) \upharpoonright_{Sys_i}$ and every $(\hat{M}'_0, S_0) \in f_0(\hat{M}_0, S_0)$, the composition $(\hat{M}'_0, S_0) \parallel (\hat{M}_1, S_1)$ exists, lies in Sys^* , and fulfils $\text{ports}(\hat{M}'_0) \cap S_1^c = \text{ports}(\hat{M}_0) \cap S_1^c$.

Let $f^\#$ denote the function that maps each $(\hat{M}^\#, S)$ to the set of these compositions. Then we have $Sys^\# \geq^{f^\#} Sys^*$. This holds for perfect, statistical and, if Sys_1 is polynomial-time, for computational security, and also for the universal and blackbox definitions. \square

Proof (sketch). First we have to show that $f^\#$ is a valid mapping. This will be done in Step 0 below.

Then let a configuration $\text{conf}^\# = (\hat{M}^\#, S, H, A^\#) \in \text{Conf}^{f^\#}(Sys^\#)$ be given and $(\hat{M}_i, S_i) := (\hat{M}^\#, S) \upharpoonright_{Sys_i}$ for $i = 0, 1$. We have to show that there is an indistinguishable configuration $\text{conf}^* \in \text{Conf}(Sys^*)$. The outline of the proof is as follows; it is illustrated in Figure 4.

1. We combine H and \hat{M}_1 into a user H_0 to obtain a configuration $\text{conf}_0 = (\hat{M}_0, S_0, H, A_0) \in \text{Conf}(Sys_0)$ where the view of H as a submachine of H_0 is the same as that in $\text{conf}^\#$.

¹⁰The first condition makes one system a valid user of another. The second one excludes cases where $p \in \text{free}([\hat{M}_i]) \cap \text{free}([\hat{M}_j])$ (e.g., a clock port for a high-level connection between these systems) and $p \in S_i$ but $p \notin S_j$.

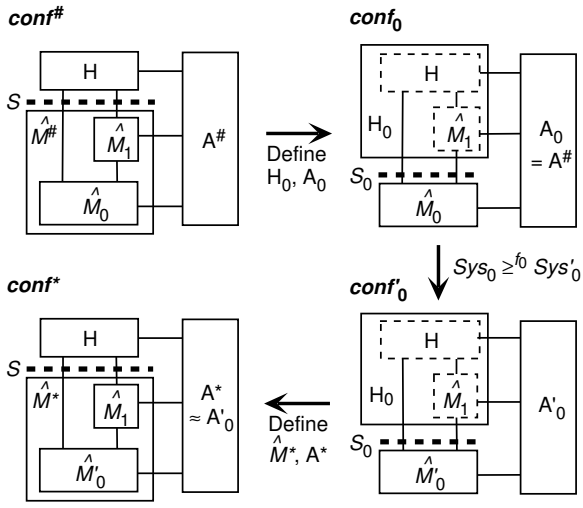


Figure 4. Configurations in the composition theorem. Dashed machines are internal submachines. (The connections drawn inside H_0 are not dashed because the combination does not hide them.)

2. We show that $conf_0 \in Conf^{f_0}(Sys_0)$. Then by the precondition $Sys_0 \geq^{f_0} Sys'_0$, there is a configuration $conf'_0 = (M'_0, S_0, H_0, A'_0) \in Conf(Sys'_0)$ with $(\hat{M}'_0, S_0) \in f_0(\hat{M}_0, S_0)$ where the view of H_0 is indistinguishable from that in $conf_0$.
3. We decompose H_0 into H and \hat{M}_1 again and derive a configuration $conf^* = (\hat{M}^*, S, H, A^*) \in Conf(Sys^*)$ where the view of H equals that of H as a submachine of H_0 in $conf'_0$.
4. We conclude that $conf^*$ is an indistinguishable configuration for $conf^\#$.

■

The following corollary finishes the definition and proof of the situation shown in Figure 3: We now assume that there is also a specification Sys' for the system Sys^* , as shown in the left part of the figure.

Corollary 4.1 Consider five systems satisfying the preconditions of Theorem 4.1, and a sixth one, Sys' , with $Sys^* \geq^{f_1} Sys'$. Then $Sys^\# \geq^f Sys'$ where $f := f_1 \circ f^\#$ as in the transitivity lemma. □

5. Secure Message Transmission

In the following we present an ideal and a real system for *secure message transmission*. The real system sends

messages over insecure channels, but for the initial key exchange authenticated channels might be used as well.

The specified notion of secure message transmission tolerates some imperfections: The adversary learns who communicates with whom and the length of the messages. He can delay messages, and thus change their order arbitrarily or even suppress them completely. He can replay messages, i.e., if an honest user u received message m from another honest user v , the adversary can trick u into accepting m arbitrarily often (although the adversary might have no idea what m looks like).

We allow this to keep the possible real machines simple (stateless except for keys, and without dummy traffic). Clearly, more complicated real systems can avoid some imperfections, although others cannot be avoided in a totally asynchronous scenario. In particular one could add a layer on top of our ideal system that detects replays or even ensures that messages are delivered in the correct sequence (by discarding messages that are out of order). Security of this new, composed system could easily be proven via the composition theorem.

Notation for data structures. For $m \in \Sigma^*$ let $\text{len}(m)$ denote the length of m . A *list*, $l = (x_1, \dots, x_k)$, is a sequence of words from Σ^* , itself encoded as a word from Σ^* . We also call fixed-length lists tuples. The exact encoding does not matter as long as the number $\text{size}(l)$ of elements in l and these elements are efficiently retrievable. Furthermore, we assume that the length of a list is (efficiently) computable from the length of its elements. For a list $l = (x_1, \dots, x_j)$ we define $l[i] := x_i$ for $1 \leq i \leq j$, and $l[i] := \downarrow$ for $i > j$, where \downarrow is a distinct error symbol, i.e., $\downarrow \notin \Sigma$. $()$ denotes an empty list. By “adding an element to a list” and similar formulations we mean appending it at the end. By $\exists x \in l$ we mean that $l[i] = x$ for some i . If we write this in a retrieval operation, the first such x is used.

5.1. The Ideal System

The ideal system is of the standard cryptographic type described at the end of Section 3.1.

Scheme 5.1 [Ideal System] Let $n \in \mathbb{N}$ and polynomials L, s over \mathbb{N} be given. Here n denotes the number of intended participants and $s(k)$ the maximum number of messages each user can send and $L(k)$ the maximum message length for the security parameter k . Let $\mathcal{M} := \{1, \dots, n\}$. An ideal system for secure message transmission is then defined as

$$Sys_{n,s,L}^{\text{secmsg,ideal}} := \{(\{\text{TH}_{\mathcal{H}}\}, S_{\mathcal{H}}) \mid \mathcal{H} \subseteq \mathcal{M}\},$$

where $S_{\mathcal{H}}$ is given by the standard definition $S_{\mathcal{H}}^c := \{\text{in}_u!, \text{out}_u? \mid u \in \mathcal{H}\}$, and $\text{TH}_{\mathcal{H}}$ is defined as follows.

When \mathcal{H} is clear from the context, let $\mathcal{A} := \mathcal{M} \setminus \mathcal{H}$ denote the set of corrupted participant indices.

The ports of $\text{TH}_{\mathcal{H}}$ are $\{\text{in}_u?, \text{out}_u! \mid u \in \mathcal{H}\} \cup \{\text{in}_{\text{sim}}?, \text{out}_{\text{sim}}!, \text{out}_{\text{sim}}^{\triangleleft}!\}$.¹¹ $\text{TH}_{\mathcal{H}}$ maintains three data structures: An array $(\text{init}_{u,v}^*)_{u,v \in \mathcal{M}}$ over $\{0, 1\}$ and an array $(\text{sc}_u^*)_{u \in \mathcal{M}}$ over $\{0, \dots, s\}$, both initialized with 0 everywhere, and an array $(\text{deliver}_{u,v}^*)_{u,v \in \mathcal{M}}$ of lists, all initially empty.

The state-transition function of $\text{TH}_{\mathcal{H}}$ is defined by the following rules. Inputs where no rule applies are ignored. If an input triggers an output $\neq \epsilon$ we say that $\text{TH}_{\mathcal{H}}$ *accepts* this input.

Initialization.

- **Send initialization:** If $\text{TH}_{\mathcal{H}}$ receives (init) at $\text{in}_u?$ and $\text{init}_{u,u}^* = 0$, it outputs (initialized, u) at $\text{out}_{\text{sim}}!$, 1 at $\text{out}_{\text{sim}}^{\triangleleft}!$, and sets $\text{init}_{u,u}^* := 1$.
- **Receive initialization from honest u :** If $\text{TH}_{\mathcal{H}}$ receives (initialize, u, v) at $\text{in}_{\text{sim}}?$, $u, v \in \mathcal{H}$ and if $\text{init}_{u,v}^* = 0$, $\text{init}_{u,u}^* = 1$, then it sets $\text{init}_{u,v}^* := 1$ and outputs (initialized, u) at $\text{out}_v!$.
- **Receive initialization from dishonest u :** If $\text{TH}_{\mathcal{H}}$ receives (initialize, u, v) at $\text{in}_{\text{sim}}?$, $u \in \mathcal{A}$, $v \in \mathcal{H}$, $\text{init}_{u,v}^* = 0$ then it sets $\text{init}_{u,v}^* := 1$ and outputs (initialized, u) at $\text{out}_v!$.

Sending and receiving messages.

- **Send:** If $\text{TH}_{\mathcal{H}}$ receives an input (send, m, v) at $\text{in}_u?$ and $m \in \Sigma^+$, $\text{len}(m) \leq L(k)$, $v \in \mathcal{M} \setminus \{u\}$, $\text{init}_{u,u}^* = 1$, $\text{init}_{v,u}^* = 1$, and $\text{sc}_u^* < s(k)$, then it first sets $\text{sc}_u^* := \text{sc}_u^* + 1$.
If $v \in \mathcal{H}$, it determines $l := \text{len}(m)$ and $i := \text{size}(\text{deliver}_{u,v}^*) + 1$, sets $\text{deliver}_{u,v}^*[i] := m$, and outputs (busy, u, i, l, v) at $\text{out}_{\text{sim}}!$ and 1 at $\text{out}_{\text{sim}}^{\triangleleft}!$.
If $v \in \mathcal{A}$, it outputs (msg, u, m, v) at $\text{out}_{\text{sim}}!$ and 1 at $\text{out}_{\text{sim}}^{\triangleleft}!$.
- **Receive from honest party u :** If $\text{TH}_{\mathcal{H}}$ receives (select, u, i, v) at $\text{in}_{\text{sim}}?$ and $u, v \in \mathcal{H}$, $\text{init}_{v,v}^* = 1$, $\text{init}_{u,v}^* = 1$, and $m := \text{deliver}_{u,v}^*[i] \neq \downarrow$, then it outputs (received, u, m) at $\text{out}_v!$.
- **Receive from dishonest party u :** If $\text{TH}_{\mathcal{H}}$ receives (send, u, m, v) at $\text{in}_{\text{sim}}?$ and $u \in \mathcal{A}$, $m \in \Sigma^+$, $\text{len}(m) \leq L(k)$, $v \in \mathcal{H}$, $\text{init}_{v,v}^* = 1$ and $\text{init}_{u,v}^* = 1$, then it outputs (received, u, m) at $\text{out}_v!$. \diamond

¹¹The order of the ports in the sequence $\text{Ports}_{\text{TH}_{\mathcal{H}}}$ does not matter below, we can assume any canonical one; similar for the following machines.

$\text{TH}_{\mathcal{H}}$ is as abstract as we hoped for: It is deterministic and contains no cryptographic objects at all. Its state-transition function should be easy to express in any formal language for automata provided it allows certain data structures, which most such languages do. We could have been even more abstract by omitting initialization (the abstraction of key exchange) and the bound s (and thus the counters). However, in practice users will have to know about initialization to provide authentic channels, and several cryptographic schemes require upper bounds on the usage. Moreover, one can choose L and s constant to make k invisible.

5.2. The Real System

The real system uses asymmetric encryption and digital signatures as cryptographic primitives; notation for them is briefly introduced in Section 5.2.1. The scheme itself is described in Section 5.2.2.

5.2.1. Primitives Used

The algorithms ($\text{gen}_S, \text{sign}, \text{test}$) denote a digital signature scheme secure against existential forgery under adaptive chosen-message attacks [17]. Let the overall message space be Σ^+ and, w.l.o.g., $\text{false} \notin \Sigma^+$. We write $(\text{sk}_s, \text{pk}_s) \leftarrow \text{gen}_S(1^k, 1^s)$ for the generation of a secret signing key and a public test key based on a security parameter, $k \in \mathbb{N}$, and the desired maximum number of signatures, $s \in \mathbb{N}$. By $\text{sig} \leftarrow \text{sign}_{\text{sk}_s, \text{sc}}(m)$ we denote the (probabilistic) signing of a message $m \in \Sigma^+$, where $\text{sc} \in \{1, \dots, s\}$ is a counter value that has to be unique among the signatures generated with sk_s . We assume that sig is of the form (m, sig') . Verification $\text{test}_{\text{pk}_s}(\text{sig})$ returns either m (then we say that the signature is valid) or false. We assume that the length of a signature on m is a function $\text{sig_len}(k, s, \text{len}(m))$.

By (gen_E, E, D) , we denote a public-key encryption scheme secure against adaptive chosen-ciphertext attacks, as introduced in [32] and formalized as “IND-CCA2” in [2]. We write $(\text{sk}_e, \text{pk}_e) \leftarrow \text{gen}_E(1^k)$ for the generation of an encryption key and a decryption key. We denote the (probabilistic) encryption of a message m by $c \leftarrow E_{\text{pk}_e}(m)$, and decryption by $m \leftarrow D_{\text{sk}_e}(c)$. The result may be false for wrong ciphertexts. We assume that messages of arbitrary length can be encrypted, but the length need not be hidden, and that the length of c is a function of k and $\text{len}(m)$.

We also need that upper bounds on the length of the results of all the algorithms (e.g., key generation) for every k are efficiently computable.

5.2.2. Real System for Secure Message Transmission

Scheme 5.2 [Real System] Let n, s, L and consequently \mathcal{M} be given as in Scheme 5.1. Also let an encryption scheme




Figure 5. Sketch of real system for secure message transmission. All connections are clocked by A. Indices \mathcal{H} are omitted; also in the following figures.

and a signature scheme be given as above. The system is a standard cryptographic system (Definition 3.1). Thus \hat{M}^* is a set $\{M_u | u \in \mathcal{M}\}$ and $S^{*c} := \{in_u!, out_u? | u \in \mathcal{M}\}$. Here \mathcal{ACC} is the powerset of \mathcal{M} . Hence the system is of the form

$$Sys_{n,s,L}^{\text{secmsg,real}} = \{(\hat{M}_{\mathcal{H}}, S_{\mathcal{H}}) | \mathcal{H} \subseteq \mathcal{M}\}$$

with $\hat{M}_{\mathcal{H}} = \{M_u, \mathcal{H} | u \in \mathcal{H}\}$.

The ports of machine M_u are $\{in_u?, out_u!\} \cup \{net_{u,v}!, net_{v,u}? | v \in \mathcal{M}\} \cup \{aut_{u,v}!, aut_{v,u}? | v \in \mathcal{M}\}$. The first ones are for the user, the second ones for normal messages to and from each M_v , and the last ones for key exchange. High-level connections $\{net_{u,v}!, net_{u,v}?\}$ are insecure and $\{aut_{u,v}!, aut_{u,v}?\}$ are authenticated. The resulting high-level connection graph is sketched in Figure 5.

Each M_u maintains an array $(init_{v,u})_{v \in \mathcal{M}}$ of lists, which are initially empty, and a counter sc_u initialized with 0. The state-transition function of M_u is defined by the following rules. Inputs where no rule applies are ignored. If an input triggers an output $\neq \epsilon$ we say that M_u *accepts* this input.

Initialization.

- **Send initialization:** If M_u receives (init) at $in_u?$ and $init_{u,u} = ()$, then it generates two key pairs, $(sk_s, pks_u) \leftarrow \text{gen}_S(1^k, 1^{s(k)})$ and $(ske_u, pke_u) \leftarrow \text{gen}_E(1^k)$, outputs (pks_u, pke_u) at all ports $aut_{u,v}!$, and sets $init_{u,u} := (sk_s, ske_u)$.
- **Receive initialization:** If M_u receives (pks_v, pke_v) (within the maximum length for a pair of public keys) at $aut_{v,u}?$ and $init_{v,u} = ()$, then it sets $init_{v,u} := (pks_v, pke_v)$ and outputs (initialized, v) at $out_u!$.

Sending and receiving messages.

- **Send:** If M_u receives an input (send, m, v) at $in_u?$, with $m \in \Sigma^+$, $\text{len}(m) \leq L(k)$, and $v \in \mathcal{M} \setminus \{u\}$, and if $init_{u,u} = (sk_s, ske_u)$, $init_{v,u} = (pks_v, pke_v)$, and $sc_u < s(k)$, then it sets $sc_u := sc_u + 1$ and

$$c \leftarrow E_{pke_v}(\text{sign}_{sk_s, sc_u}(u, m, v)).$$

It outputs c at $net_{u,v}!$.

- **Receive:** If M_u receives c at $net_{v,u}?$ (within the maximum length for a network message as above), and $init_{u,u} = (sk_s, ske_u)$ and $init_{v,u} = (pks_v, pke_v)$, then M_u tries to parse c in the form $E_{pke_u}(\text{sign}_{sk_s, sc_v}(v, m, u))$. More precisely:

- $sig := D_{ske_u}(c)$. Abort if the result is false.
- $m' := \text{test}_{pks_v}(sig)$. Abort if the result is false.
- $(v', m, u') := m'$. Abort if this fails or $u' \neq u$ or $v' \neq v$.

If this succeeds, then M_u outputs (received, v, m) at $out_u!$. \diamond

5.3. Security of the Real System

We show that Scheme 5.2 is at least as secure as the ideal Scheme 5.1.

Theorem 5.1 (Security of Secure Message Transmission)
For all $n \in \mathbb{N}$ and $L, s \in \mathbb{N}[x]$, $Sys_{n,s,L}^{\text{secmsg,real}} \geq_{\text{sec}}^{f, \text{poly}} Sys_{n,s,L}^{\text{secmsg,ideal}}$ for the canonical mapping f from Section 3.1, provided the signature and encryption schemes used are secure. This holds with blackbox simulatability (Definition 2.12). \square

The proof of Theorem 5.1 is given in Section 5.3.2. It is based on Theorem 5.2, presented in Section 5.3.1.

5.3.1. General Simulatability of Public-key Encryption in a Reactive Multi-user Setting

An essential cryptographic part of the proof of Theorem 5.1 is captured by Theorem 5.2, which extends the standard notion of chosen-ciphertext security of public-key encryption to a reactive multi-user setting, using a simulatability definition. A similar multi-user scenario has been considered in [4], but no decryption request for any of the ciphertexts produced by a correct machine is allowed there. However, in a reactive scenario like ours, most secret messages are also decrypted by some correct machine and partial knowledge may leak; hence the theorem is not immediately applicable. We therefore define ideal machines $\text{Enc}_{\text{sim}, \mathcal{H}}$ that encrypt

simulated messages, but honor *all* decryption requests by table look-up of the intended messages. Then we can show simulatability in our usual sense.¹²

Scheme 5.3 [*Encryption Systems*] Let an encryption scheme (gen_E, E, D) and parameters $n \in \mathbb{N}$ and $s_{keys}, s_{encs} \in \mathbb{N}[x]$ be given, where $s_{keys}(k)$ denotes the maximum number of keys to be generated in the system and $s_{encs}(k)$ the maximum number of encryptions per key, for security parameter k . For every $l \in \mathbb{N}$, let $m_{sim,l}$ denote a fixed message from Σ^l , say 0^l . We define two systems

- $Sys_{n,s_{keys},s_{encs}}^{\text{enc,real}} := \{(\{\text{Enc}_{\mathcal{H}}\}, S_{\text{enc},\mathcal{H}}) \mid \mathcal{H} \subseteq \mathcal{M}\}$,
- $Sys_{n,s_{keys},s_{encs}}^{\text{enc,sim}} := \{(\{\text{Enc}_{\text{sim},\mathcal{H}}\}, S_{\text{enc},\mathcal{H}}) \mid \mathcal{H} \subseteq \mathcal{M}\}$.

For every \mathcal{H} , the ports are defined as follows:

- $Ports_{\text{Enc}_{\mathcal{H}}} := Ports_{\text{Enc}_{\text{sim},\mathcal{H}}} := \{\text{in}_{\text{enc},u}?, \text{out}_{\text{enc},u}!, \text{out}_{\text{enc},u}^{\triangleleft}! \mid u \in \mathcal{H}\}$,
- $S_{\text{enc},\mathcal{H}}^c := \{\text{in}_{\text{enc},u}!, \text{in}_{\text{enc},u}^{\triangleleft}!, \text{out}_{\text{enc},u}?\mid u \in \mathcal{H}\}$.¹³

Both machines have a security parameter, k , and maintain a key counter $kc \in \mathbb{N}$, initially $kc := 0$, and initially empty lists $keys$ and $ciphers$. The latter is used for the look-up of intended cleartexts in the ideal system. The transition functions are given as follows, where we assume that the current input is made at port $\text{in}_{\text{enc},u}?$. The resulting output goes to $\text{out}_{\text{enc},u}!$, with $\text{out}_{\text{enc},u}^{\triangleleft}! := 1$.

- Input (generate) for $\text{Enc}_{\mathcal{H}}$ and $\text{Enc}_{\text{sim},\mathcal{H}}$: If $kc < s_{keys}(k)$ then $\{kc := kc + 1; (ske, pke) \leftarrow \text{gen}_E(1^k); \text{add } (u, kc, ske, pke, 0) \text{ to } keys; \text{output } pke\}$ else output \downarrow .
- Input (encrypt, pke, m) with $pke, m \in \Sigma^+$ and within the length bounds:
 - for $\text{Enc}_{\mathcal{H}}$: If $\exists v, kc, ske, s_{pke}: (v, kc, ske, pke, s_{pke}) \in keys \wedge s_{pke} < s_{encs}(k)$ then $\{s_{pke} := s_{pke} + 1; \text{output } c \leftarrow E_{pke}(m)\}$ else output \downarrow .
 - for $\text{Enc}_{\text{sim},\mathcal{H}}$: If $\exists v, kc, ske, s_{pke}: (v, kc, ske, pke, s_{pke}) \in keys \wedge s_{pke} < s_{encs}(k)$ then $\{s_{pke} := s_{pke} + 1; \text{output } c \leftarrow E_{pke}(m_{sim,\text{len}(m)}); \text{add } (m, pke, c) \text{ to } ciphers\}$ else output \downarrow .

¹²For cryptographers, our theorem can also be seen as a formalization of the notion of “a standard hybrid argument.” For a passive setting this was done by Theorem 3.6 of [13]. However, in a reactive setting one has to switch over from a real state to a “corresponding” ideal state, and there is no general definition for this. In particular, it must be made clear how decryption is handled in the hybrids. This is now well-defined at least for those systems that use an encryption system only such that they can be rewritten with our real encryption system.

¹³Thus $\bar{S}_{\text{enc},\mathcal{H}} = \emptyset$, i.e., the adversary is not connected with the correct machines except via or in the place of H . The fact that inputs are scheduled by H , and outputs for H are scheduled by the system, allows H to use the system like a local subsystem which always returns a result immediately.

- Input (decrypt, pke, c) with $pke, c \in \Sigma^+$ and within the length bounds (note that pke is used as a designator of the desired private key):

- for $\text{Enc}_{\mathcal{H}}$: If $\exists kc, ske, s_{pke}: (u, kc, ske, pke, s_{pke}) \in keys$ then $\{m \leftarrow D_{ske}(c); \text{output } m\}$ else output \downarrow .
- for $\text{Enc}_{\text{sim},\mathcal{H}}$: If $\exists kc, ske, s_{pke}: (u, kc, ske, pke, s_{pke}) \in keys$ then $\{\text{If } \exists m: (m, pke, c) \in ciphers \text{ then output } m \text{ else output } m \leftarrow D_{ske}(c)\}$ else output \downarrow . \diamond

Note that $\text{Enc}_{\text{sim},\mathcal{H}}$ limits the capability to decrypt: If $(u, kc, ske, pke, 0)$ was added to $keys$ due to an input (generate) at port $\text{in}_{\text{enc},u}?$, then (decrypt, pke, c) has an effect only if it is entered at the same port.

Theorem 5.2 (*General Simulatability of Public-key Encryption*) For all $n \in \mathbb{N}$, $s_{keys}, s_{encs} \in \mathbb{N}[x]$, we have

$$Sys_{n,s_{keys},s_{encs}}^{\text{enc,real}} \geq_{\text{SEC}}^{f,\text{poly}} Sys_{n,s_{keys},s_{encs}}^{\text{enc,sim}}$$

for the canonical mapping f , provided the encryption scheme used is secure against adaptive chosen-ciphertext attacks. This holds with blackbox simulatability. \square

Proof. Let n, s_{keys}, s_{encs} and \mathcal{H} be fixed. We use a simulator that executes its blackbox A_1 without change, i.e., always $A_2 = A_1 =: A$. The proof is a hybrid argument as first used in [16], i.e., we construct intermediate systems that differ only in one encryption each.

For every $k \in \mathbb{N}$ let $\mathcal{I}_k := (\{1, \dots, s_{keys}(k)\} \times \{1, \dots, s_{encs}(k)\}) \cup \{\alpha\}$, let $<_k$ be the lexicographic order on $\mathcal{I}_k \setminus \{\alpha\}$, and $\alpha \leq_k t$ for all $t \in \mathcal{I}_k$. Let $\text{pred}_k(t)$ be the predecessor of $t \in \mathcal{I}_k$ relative to $<_k$ and $\omega(k) := (s_{keys}(k), s_{encs}(k))$.

For every $k \in \mathbb{N}$ and $t \in \mathcal{I}_k$, we define a hybrid machine $\text{Enc}_{k,t,\mathcal{H}}$. It is like $\text{Enc}_{\text{sim},\mathcal{H}}$ with fixed initial input 1^k , except where $\text{Enc}_{\text{sim},\mathcal{H}}$ carries out $c \leftarrow E_{pke}(m_{sim,\text{len}(m)})$: Let $t' := (kc, s_{pke})$ for the values kc, s_{pke} at that moment.

- If $t' \leq_k t$, it sets $c \leftarrow E_{pke}(m)$ like $\text{Enc}_{\mathcal{H}}$;
- if $t' >_k t$, it sets $c \leftarrow E_{pke}(m_{sim,\text{len}(m)})$ like $\text{Enc}_{\text{sim},\mathcal{H}}$.

Clearly, each $\text{Enc}_{k,\alpha,\mathcal{H}}$ works like $\text{Enc}_{\text{sim},\mathcal{H}}$ on input 1^k . Furthermore, $\text{Enc}_{k,\omega(k),\mathcal{H}}$ works like $\text{Enc}_{\mathcal{H}}$ on input 1^k : $\text{Enc}_{\mathcal{H}}$ and $\text{Enc}_{k,\omega(k),\mathcal{H}}$ produce identical outputs for inputs (generate) and (encrypt, pke, m). Now consider an input (decrypt, pke, c) at $\text{in}_{\text{enc},u}?$ such that $\exists kc, ske, s_{pke}: (u, kc, ske, pke, s_{pke}) \in keys$ (otherwise both output \downarrow). If there is no tuple (m, pke, c) in $ciphers$, both output $D_{ske}(c)$. If there is, then c has been generated by $\text{Enc}_{k,\omega(k),\mathcal{H}}$ as $E_{pke}(m)$. Thus $D_{ske}(c) = m$, and both machines output m .

Assume for contradiction that the theorem is wrong for the given parameters and machines A and H. Let $conf_{\text{real}} := (\{\text{Enc}_{\mathcal{H}}\}, S_{\text{enc}, \mathcal{H}}, H, A)$ and $conf_{\text{sim}}$ similarly with $\text{Enc}_{\text{sim}, \mathcal{H}}$, and let $coll_{k,t}$ denote the collection $\{\text{Enc}_{k,t, \mathcal{H}}, H, A\}$. (The initial inputs in these collections are always 1^k .) Thus we assume $view_{conf_{\text{real}}}(H) \not\approx_{\text{poly}} view_{conf_{\text{sim}}}(H)$, and this implies

$$(view_{coll_{k,\omega(k)}}(H))_{k \in \mathbb{N}} \not\approx_{\text{poly}} (view_{coll_{k,\alpha}}(H))_{k \in \mathbb{N}}.$$

We abbreviate $view_{k,t} := view_{coll_{k,t}}(H)$. The distinguishability means that a probabilistic polynomial-time distinguisher algorithm Δ and $p \in \mathbb{N}[x]$ exist such that for all k in an infinite set $\mathcal{K} \subseteq \mathbb{N}$,

$$|P(\Delta(view_{k,\omega(k)}) = 1) - P(\Delta(view_{k,\alpha}) = 1)| > \frac{1}{p(k)}.$$

We construct two almost identical adversaries $A_{\text{enc}+}$ and $A_{\text{enc}-}$ on the encryption scheme; we write A_{enc} where statements hold for both. On input 1^k , A_{enc} randomly selects $t \in_{\mathcal{R}} \mathcal{I}_k \setminus \{\alpha\}$, say, $t = (kc, s')$, receives a correctly chosen public key pke , and can then interact with a decryption oracle Dec for this key. A_{enc} simulates $coll_{k,t}$ with the following exceptions:

- If H makes the kc -th input (generate), say at port $in_{\text{enc},u}?$, then A_{enc} only adds $(u, kc, 0, pke, 0)$ to $keys$. Instead of operations $m \leftarrow D_{ske}(c)$ corresponding to this entry (identified by pke), it uses the decryption oracle Dec.
- If H makes an input (encrypt, pke, m) which gets the index $t = (kc, s')$, i.e., A_{enc} finds the entry $(u, kc, 0, pke, s' - 1)$ in $keys$, then A_{enc} sends $(m_0, m_1) := (m, m_{\text{sim}, \text{len}(m)})$ to Dec. This oracle flips a bit $b \in_{\mathcal{R}} \{0, 1\}$ and returns $c \leftarrow E_{pke}(m_b)$. A_{enc} adds (m, pke, c) to $ciphers$.

At the end, A_{enc} runs Δ on the resulting view of H, which yields a bit b_k^* . $A_{\text{enc}+}$ outputs b_k^* , whereas $A_{\text{enc}-}$ outputs $1 - b_k^*$ (intuitively, this is a guess at b).

Note that A_{enc} never asks Dec to decrypt the ciphertext c from Dec (which would not be allowed), because A_{enc} will find (m, pke, c) in $ciphers$ and output m .

Let $view_k^{(b)}$ denote the random variable of the view of H in A_{enc} for parameter k and a specific bit b . Abbreviate $pr_{k,t} := P(\Delta(view_{k,t}) = 1)$ and $pr_k^{(b)} := P(\Delta(view_k^{(b)}) = 1)$. For $b = 0$ the simulated run is generated like a run of $coll_{k,t}$ and for $b = 1$ like a run of $coll_{k, \text{pred}_k(t)}$. With $w(k) := s_{\text{keys}}(k)s_{\text{encs}}(k)$ we get:

$$\begin{aligned} pr_k^{(0)} &= \frac{1}{w(k)} \sum_{t \in \mathcal{I}_k \setminus \{\alpha\}} pr_{k,t}, \\ pr_k^{(1)} &= \frac{1}{w(k)} \sum_{t \in \mathcal{I}_k \setminus \{\alpha\}} pr_{k, \text{pred}_k(t)}. \end{aligned}$$

For all $k \in \mathcal{K}$, this implies

$$\begin{aligned} |pr_k^{(0)} - pr_k^{(1)}| &= \frac{1}{w(k)} |pr_{k,\omega(k)} - pr_{k,\alpha}| \\ &> \frac{1}{w(k)p(k)}. \end{aligned}$$

Thus

$$\begin{aligned} P(b_k^* = b) &= P(b = 0 \wedge \Delta(view_k^{(b)}) = 0) \\ &\quad + P(b = 1 \wedge \Delta(view_k^{(b)}) = 1) \\ &= \frac{1}{2} (P(\Delta(view_k^{(0)}) = 0) + P(\Delta(view_k^{(1)}) = 1)) \\ &= \frac{1}{2} + \frac{1}{2} (pr_k^{(1)} - pr_k^{(0)}). \end{aligned}$$

Thus the success probability of either $A_{\text{enc}+}$ or $A_{\text{enc}-}$ is larger than $1/2 + 1/(2w(k)p(k))$ for all $k \in \mathcal{K}$. This is the desired contradiction to the security of the encryption system. \blacksquare

5.3.2. Proof (Sketch) of Theorem 5.1

We prove Theorem 5.1 in the following steps: In Part A we rewrite the real system to use a reactive encryption system $Sys_{n,s,L}^{\text{enc}, \text{real}}$ from Scheme 5.3 instead of performing the original encryption algorithms. This gives an intermediate system $Sys_{n,s,L}^{\text{secmsg}, \text{Enc}}$, and we show that

$$Sys_{n,s,L}^{\text{secmsg}, \text{real}} \geq_{\text{sec}} Sys_{n,s,L}^{\text{secmsg}, \text{Enc}}.$$

In Part B, we replace each machine $\text{Enc}_{\mathcal{H}}$ by $\text{Enc}_{\text{sim}, \mathcal{H}}$ to get a new system $Sys_{n,s,L}^{\text{secmsg}, \text{Encsim}}$. We show that

$$Sys_{n,s,L}^{\text{secmsg}, \text{Enc}} \geq_{\text{sec}} Sys_{n,s,L}^{\text{secmsg}, \text{Encsim}}.$$

In Part C we define a simulator $\text{Sim}_{\mathcal{H}}$ that uses an adversary from a configuration of $Sys_{n,s,L}^{\text{secmsg}, \text{Encsim}}$ as a blackbox, resulting in a configuration of the ideal system $Sys_{n,s,L}^{\text{secmsg}, \text{ideal}}$. We show that this simulation is correct, i.e., that

$$Sys_{n,s,L}^{\text{secmsg}, \text{Encsim}} \geq_{\text{sec}} Sys_{n,s,L}^{\text{secmsg}, \text{ideal}}.$$

In this part, we omit details for brevity. Theorem 5.1 follows by the transitivity of \geq_{sec} .

Part A, Intermediate System: Real System Using Encryption Subsystem. We first define the intermediate system $Sys_{n,s,L}^{\text{secmsg}, \text{Enc}}$. Its structures are of the form $(\hat{M}'_{\mathcal{H}}, S_{\mathcal{H}})$ with $\hat{M}'_{\mathcal{H}} = \{\text{Enc}_{\mathcal{H}}\} \cup \{M'_{u, \mathcal{H}} | u \in \mathcal{H}\}$, see Figure 6.¹⁴

¹⁴This intermediate system is not “real” and only used in our proof; hence it is no problem that $M'_{u, \mathcal{H}}$ makes explicit distinctions using \mathcal{H} below.

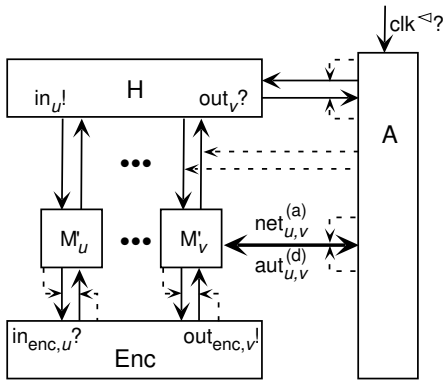


Figure 6. Real system rewritten with encryption subsystem. The dotted arrows connect clock ports.

The parameter n of the encryption subsystem is as in the overall system and $s_{keys}(k) := n$, $s_{encs}(k) := ns(k)$. Each machine $M'_{u,\mathcal{H}}$ equals $M_{u,\mathcal{H}}$ with the following modifications:

- It has additional ports $in_{enc,u}!$, $in_{enc,u}^{\triangleleft!}$, $out_{enc,u}?$ to connect to $Enc_{\mathcal{H}}$.
- In “Send initialization.” Instead of calling gen_E , $M'_{u,\mathcal{H}}$ outputs (generate) at $in_{enc,u}!$ and 1 at $in_{enc,u}^{\triangleleft!}$ and waits until it receives a result at $out_{enc,u}?$.¹⁵ It uses this result as pke_u . In $init_{u,u}$ it stores pke_u instead of ske_u .
- In “Send.” If $v \in \mathcal{H}$, instead of the computation of c , it computes $sig \leftarrow \text{sign}_{sks_u, sc_u}(u, m, v)$ and outputs (encrypt, pke_v, sig) at $in_{enc,u}!$ and 1 at $in_{enc,u}^{\triangleleft!}$. It waits until it receives a result at $out_{enc,u}?$, which it uses as c .
- In “Receive.” $M'_{u,\mathcal{H}}$ replaces Step a) of parsing by an output (decrypt, pke_u, c) at $in_{enc,u}!$ and 1 at $in_{enc,u}^{\triangleleft!}$. It waits until it receives a result at $out_{enc,u}?$, which it uses as sig .

The views of A and H in a configuration $(\hat{M}_{\mathcal{H}}, S_{\mathcal{H}}, H, A)$ and $(\hat{M}'_{\mathcal{H}}, S_{\mathcal{H}}, H, A)$ are clearly identical (perfectly indistinguishable), because the actions of $Enc_{\mathcal{H}}$ on the given inputs are precisely what $M_{u,\mathcal{H}}$ would have done at this point. In particular $s_{keys}(k)$ is not exceeded because each $M_{u,\mathcal{H}}$ inputs (generate) at most once, controlled by $init_{u,u}$, and $s_{encs}(k)$ is not exceeded for any key because each $M_{u,\mathcal{H}}$

¹⁵More precisely it enters a state (wait, init) where it only reacts on this input. As one easily sees from the scheduling that only this input can arrive next, we treat the wait state together with the previous state; also in the following cases.

inputs (encrypt, ...) at most $s(k)$ times. The interactions between $Enc_{\mathcal{H}}$ and the machines $M'_{u,\mathcal{H}}$ are invisible for A . Thus we have

$$Sys_{n,s,L}^{\text{secmsg,real}} \geq_{\text{sec}} Sys_{n,s,L}^{\text{secmsg,Enc}}.$$

Part B, Replacing the Encryption System. In $Sys_{n,s,L}^{\text{secmsg,Enc}}$, we want to replace each machine $Enc_{\mathcal{H}}$ by $Enc_{\text{sim},\mathcal{H}}$ to get a new system $Sys_{n,s,L}^{\text{secmsg,Encsim}}$.

We consider $Sys_{n,s,L}^{\text{secmsg,Enc}}$ as a composition of $Sys_0 := Sys_{n,n,ns}^{\text{enc,real}}$ and a system Sys_1 that is naturally defined as the structures without $Enc_{\mathcal{H}}$: The specified ports are those of $Sys_{n,s,L}^{\text{secmsg,real}}$ plus the low-level complements of the ports of $Enc_{\mathcal{H}}$. Then the conditions of Definition 4.1 are fulfilled. As each machine $Enc_{\text{sim},\mathcal{H}}$ has the same ports as $Enc_{\mathcal{H}}$, the definition of $Sys_{n,s,L}^{\text{secmsg,Encsim}}$ as a composition of $Sys'_0 := Sys_{n,n,ns}^{\text{enc,sim}}$ and the same Sys_1 is clear and the preconditions of the composition theorem, Theorem 4.1, are fulfilled. In particular, all machines $M'_{u,\mathcal{H}}$ are polynomial-time in k . Hence Theorem 5.2 and Theorem 4.1 imply $Sys_{n,s,L}^{\text{secmsg,Enc}} \geq_{\text{sec}} Sys_{n,s,L}^{\text{secmsg,Encsim}}$ (again with blackbox simulatability).

Part C, Simulator. It remains to be shown that $Sys_{n,s,L}^{\text{secmsg,Encsim}} \geq_{\text{sec}} Sys_{n,s,L}^{\text{secmsg,ideal}}$. Intuitively, one remaining aspect is to show that the real messages m , which are still inputs to $Enc_{\text{sim},\mathcal{H}}$, but which are not output by $TH_{\mathcal{H}}$, are indeed not needed. This is a perfectly indistinguishable rewriting. The other aspect is to show that the use of signatures guarantees authenticity as specified in the ideal system.

Simulator $Sim_{\mathcal{H}}$: We construct $Sim_{\mathcal{H}}$ as the combination of several machines, see Figure 7. It uses the given A as a submachine, without any port renaming. (Although all figures show H as using all the specified ports, the proof is general.)

- Mux acts as a multiplexer/demultiplexer between $TH_{\mathcal{H}}$ and other machines. It has ports $in_{sim}!$, $in_{sim}^{\triangleleft!}$, $out_{sim}?$, and in_u^* , $in_u^{\triangleleft!}$ and out_u^* for each $u \in \mathcal{H}$. Basically, it translates from the “top” to the “bottom” and vice versa, immediately scheduling its outputs so that the values simply “pass through:”
 - Input (initialized, u) at $out_{sim}?$ / Output (init) at in_u^* and 1 at $in_u^{\triangleleft!}$.
 - Input (busy, u, i, l, v) at $out_{sim}?$ / Output (send_blindly, i, l, v) at in_u^* and 1 at $in_u^{\triangleleft!}$.
 - Input (msg, u, m, v) at $out_{sim}?$ / Output (send, m, v) at in_u^* and 1 at $in_u^{\triangleleft!}$.

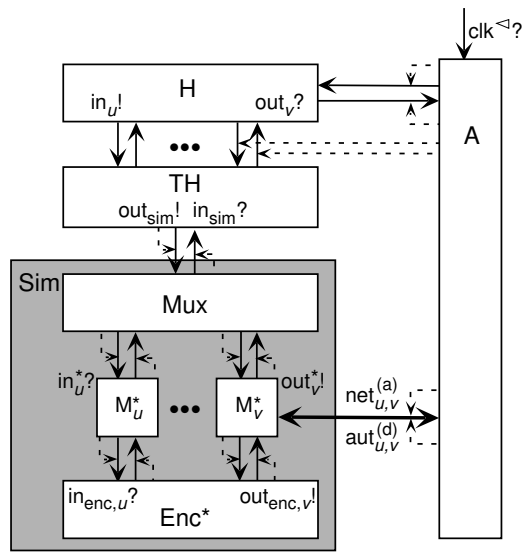


Figure 7. Simulator for secure message transmission.

- Input (initialized, u) at $out_v^*?$ / Output (initialize, u, v) at $in_{sim}!$ and 1 at $in_{sim}^{\triangleleft}!$.
 - Input (received_blindly, u, i) at $out_v^*?$ / Output (select, u, i, v) at $in_{sim}!$ and 1 at $in_{sim}^{\triangleleft}!$.
 - Input (received, u, m) at $out_v^*?$ / Output (send, u, m, v) at $in_{sim}!$ and 1 at $in_{sim}^{\triangleleft}!$.
- $M_{u,\mathcal{H}}^*$, for $u \in \mathcal{H}$, equals $M'_{u,\mathcal{H}}$ with the following modifications:
 - It has an additional port $out_u^*!$ which it sets to 1 for every output at $out_u^*!$ (so that the outputs “pass up” through Mux).
 - Upon input (send_blindly, i, l, v) at $in_u^*?$, it behaves like $M'_{u,\mathcal{H}}$ on input (send, m, v) with $len(m) = l$, but instead of computing $sig \leftarrow \text{sign}_{sks_u, sc_u}(u, m, v)$ it just computes the length l^* of any such sig using the algorithm sig_len . It outputs (encrypt_blindly, $pke_v, (u, i, v), l^*$) at $in_{enc,u}!$ and 1 at $in_{enc,u}^{\triangleleft}!$. Then it waits and continues like $M'_{u,\mathcal{H}}$.
 - In “receive”, it acts like $M'_{u,\mathcal{H}}$ until it receives a result at $out_{enc,u}?$. If this is of the form (decrypted, sig), then $M_{u,\mathcal{H}}^*$ treats it as $M'_{u,\mathcal{H}}$ treats sig . If it is of the form (decrypted_blindly, (v', i, u')), then $M_{u,\mathcal{H}}^*$ verifies that $v' = v$ and $u' = u$. If yes, it outputs (received_blindly, v, i) at $out_u^*!$.
 - $Enc_{\mathcal{H}}^*$ equals $Enc_{sim,\mathcal{H}}$ with the following modification:
 -

- An array $blind_ciphers$ replaces $ciphers$.
- Instead of inputs (encrypt, ...), it accepts inputs (encrypt_blindly, pke, mid, l^*). Here $mid \in \Sigma^*$ is a message identifier. If it finds the desired tuple in $keys$ (otherwise the result is \perp), it behaves like $Enc_{sim,\mathcal{H}}$ on input (encrypt, pke, m) with $len(m) = l^*$, and it stores (mid, pke, c) in $blind_ciphers$.
- In “decrypt”, it looks for a tuple (mid, pke, c) in $blind_ciphers$. If it finds one, it outputs (decrypted_blindly, mid). Otherwise, it decrypts and outputs the result m as (decrypted, m).

To prove the correctness of this simulator we have to compare configurations $conf_{sr} := (\hat{M}'_{\mathcal{H}}, S_{\mathcal{H}}, H, A)$ and $conf_{id} := (\{TH_{\mathcal{H}}\}, S_{\mathcal{H}}, H, Sim_{\mathcal{H}}(A))$, called semi-real and ideal configuration. The overall idea of this proof is to define a mapping ϕ from runs of $conf_{sr}$ to runs of $conf_{id}$, except for negligible subsets on both sides; we call them “error sets.” ϕ must respect probabilities and the views of A and H in runs ρ and $\phi(\rho)$ must be equal. In our case, we can simply define ϕ state-wise, and only for states before and after switching steps of H and A; this is sufficient because only the views of H and A must be identical.

We show that $\phi(\delta_{sr}(\sigma_{sr})) = \delta_{id}(\phi(\sigma_{sr}))$ for all states σ_{sr} reachable in $conf_{sr}$, except for the error sets. Here δ_{sr} and δ_{id} denote the overall probabilistic transition functions. The error sets will consist of the runs where the adversary successfully forges a signature.

Mapping ϕ : Let a state σ_{sr} of $conf_{sr}$ be given; we define the components of $\sigma_{id} := \phi(\sigma_{sr})$. Large parts of the mapping are trivial:

- The states of H and A are mapped identically.
- The states of all buffers with the same name in both systems are mapped identically, and so is the scheduled port.¹⁶
- The remaining buffers in σ_{id} are always empty.¹⁷
- Security parameters are mapped identically and not mentioned again.

Now we map the joint states of $M'_{u,\mathcal{H}}$ for all $u \in \mathcal{H}$ and $Enc_{sim,\mathcal{H}}$ to states of $TH_{\mathcal{H}}$ and $Sim_{\mathcal{H}}$. The state of $Sim_{\mathcal{H}}$ consists of those of Mux, the machines $M_{u,\mathcal{H}}^*$, and $Enc_{\mathcal{H}}^*$.

- Mux is state-less.

¹⁶These are $\widetilde{in}_u, \widetilde{out}_u, \widetilde{in}_{enc,u}, \widetilde{out}_{enc,u}, \widetilde{net}_{u,v}, \widetilde{net}_{v,u}^a, \widetilde{aut}_{u,v}, \widetilde{aut}_{v,u}^d$ for all $u \in \mathcal{H}, v \in \mathcal{M}$, and the buffers between H and A.

¹⁷Recall that we only map states before or after H or A switches. The buffers are $\widetilde{in}_{sim}, \widetilde{out}_{sim}$ and $\widetilde{in}_u^*, \widetilde{out}_u^*$ for all $u \in \mathcal{H}$.

- Key-related variables:
 - The array $init^*$ of $TH_{\mathcal{H}}$ is derived from the arrays $init_{\bullet,u}$ of the machines $M'_{u,\mathcal{H}}$: We set $init_{v,u}^* := 1$ whenever $init_{v,u} \neq ()$, else $init_{v,u}^* := 0$.
 - Each array $init_{\bullet,u}$ of a machine $M_{u,\mathcal{H}}^*$ equals that in $M'_{u,\mathcal{H}}$.
 - The counter kc and the list $keys$ of $Enc_{\mathcal{H}}^*$ equal those in $Enc_{sim,\mathcal{H}}$.
- Message-related variables: In the semi-real configuration, $M'_{u,\mathcal{H}}$ contains a counter sc_u , and $Enc_{sim,\mathcal{H}}$ the list $ciphers$. In the ideal configuration, $TH_{\mathcal{H}}$ contains counters sc_u^* and an array $deliver^*$, while $M_{u,\mathcal{H}}^*$ contains counters sc_u , and $Enc_{\mathcal{H}}^*$ the list $blind_ciphers$.
 - The counters sc_u in $M_{u,\mathcal{H}}^*$ and sc_u^* in $TH_{\mathcal{H}}$ equal sc_u in $M'_{u,\mathcal{H}}$.
 - Each entry $e := ciphers[j] \neq \downarrow$ is of the form $e = (sig, pke, c)$ and sig of the form $((u, m, v), sig')$ with $u, v \in \mathcal{H}$. Given $ciphers$, let $ind_{u,v}(j)$ denote the number of entries up to (and including) $ciphers[j]$ with the given values u, v . Hence for each such entry e we can set $i := ind_{u,v}(j)$ and
 - * $blind_ciphers[j] := ((u, i, v), pke, c)$ and
 - * $deliver_{u,v}^*[i] = m$.

The proof of the correctness of this mapping is a relatively straightforward but tedious exercise; see [31]. More precisely, this proof shows that all steps of the runs are matched by ϕ , and thus the views of A and H are identical, except if the run of the ideal system belongs to a set $Forgeries_{v,k}$ with $v \in \mathcal{H}$. In each such run, the adversary has produced a signature with a key sk_{s_v} of the correct machine $M_{v,\mathcal{H}}^*$ under a message that this machine had not signed. These runs and the forged signatures in them are efficiently recognizable.

Runs with forged signatures. The overall statement now follows from the security of the signature scheme.

The proof is a standard reduction: We show that each sequence $(Forgeries_{v,k})_{k \in \mathbb{N}}$ has negligible probability. Then the sequence of the overall sets of exceptions, $(Forgeries_k)_{k \in \mathbb{N}}$ with $Forgeries_k := \cup_{v \in \mathcal{H}} Forgeries_{v,k}$ is also negligible.

Assume the contrary for some v . We then construct an adversary A_{sig} against the signature scheme: On input a public key pks , it runs $conf_{id}$ with pks as pks_v , using the signature oracle for all signatures with the now unknown sk_{s_v} . It verifies on-line whether the run belongs to $Forgeries_{v,k}$.

If yes, it outputs the forged signature. The success probability of A_{sig} for a security parameter k is precisely the probability of $Forgeries_{v,k}$.

6. Summary

We have presented a rigorous model for secure reactive systems with cryptographic parts in asynchronous networks (Section 2) and a composition theorem for it (Section 4). Common types of cryptographic systems and their trust models can be expressed in this model, in particular systems with static or adaptive adversaries (Section 3).

As design principles, we propose to keep specifications *abstract*, i.e., free of all implementation details, and to explicitly include all *tolerable imperfections*. This allows the cryptographically verified systems to be used as building blocks for systems that are then subject to formal verification. As an example of this specification methodology we provided the first abstract and complete specification for Secure Message Transmission, and verified a concrete implementation (Section 5). This is based on a theorem about the security of encryption in a reactive setting, which may be of independent interest.

Future work will include actually using a formal language to express the abstract specifications and examples where the composition theorem is applied based on such specifications. We already used it in the example, but based on Theorem 5.2 with a non-abstract ideal system; we only regard the final ideal system of Theorem 5.1 as abstract. Nevertheless, it might be interesting to see how far even the lower-level proofs could be supported by tools.

Acknowledgments. We thank *Ran Canetti, Martin Hirt, Paul Karger, Matthias Schunter, Victor Shoup* and *Michael Steiner* for interesting discussions.

This work was supported by the European IST Project MAFTIA. However, it represents the view of the authors. The MAFTIA project is partially funded by the European Commission and the Swiss Federal Office for Education and Science (BBW).

References

- [1] M. Abadi, P. Rogaway, *Reconciling Two Views of Cryptography (The Computational Soundness of Formal Encryption)*, IFIP Intern. Conf. on Theoretical Computer Science (TCS 2000), LNCS 1872, Springer-Verlag, 2000, 3–22
- [2] M. Bellare, A. Desai, D. Pointcheval, P. Rogaway, *Relations Among Notions of Security for Public-Key Encryption Schemes*, Crypto '98, LNCS 1462, Springer-Verlag, 1998, 26–45
- [3] D. Beaver, *Secure Multiparty Protocols and Zero Knowledge Proof Systems Tolerating a Faulty Minority*, J. of Cryptology 4/2 (1991) 75–122

- [4] M. Bellare, A. Boldyreva, S. Micali, *Public-Key Encryption in a Multi-user Setting: Security Proofs and Improvements*, Eurocrypt 2000, LNCS 1807, Springer-Verlag, 2000, 259–274
- [5] M. Bellare, R. Canetti, H. Krawczyk, *A Modular Approach to the Design and Analysis of Authentication and Key Exchange Protocols*, 13th Symp. on Theory of Computing (STOC), ACM, 1998, 419–428
- [6] D. Beaver, S. Haber, *Cryptographic Protocols Provably Secure Against Dynamic Adversaries*, Eurocrypt '92, LNCS 658, Springer-Verlag, 1993, 307–323
- [7] R. Canetti, *Studies in Secure Multiparty Computation and Applications*, Department of Computer Science and Applied Mathematics, The Weizmann Institute of Science, June 1995, revised March 1996
- [8] R. Canetti, *Security and Composition of Multiparty Cryptographic Protocols*, J. of Cryptology 13/1 (2000) 143–202
- [9] R. Canetti, *A Unified Framework for Analyzing Security of Protocols*, IACR Cryptology ePrint Archive 2000/067, December 2000, <http://eprint.iacr.org/>
- [10] R. Canetti, S. Goldwasser, *An Efficient Threshold Public Key Cryptosystem Secure Against Adaptive Chosen Ciphertext Attack*, Eurocrypt '99, LNCS 1592, Springer-Verlag, 1999, 90–106
- [11] D. Dolev, A. C. Yao, *On the Security of Public Key Protocols*, IEEE Transactions on Information Theory 29/2 (1983) 198–208
- [12] R. Gennaro, S. Micali, *Verifiable Secret Sharing as Secure Computation*, Eurocrypt '95, LNCS 921, Springer-Verlag, 1995, 168–182
- [13] O. Goldreich, *Foundations of Cryptography (Fragments of a Book)*, Department of Computer Science and Applied Mathematics, Weizmann Institute of Science, Rehovot, Israel, February 23, 1995, available with updates at <http://theory.lcs.mit.edu/~oded/>
- [14] O. Goldreich, *Secure Multi-Party Computation*, Working Draft, Version 1.1, September 21, 1998, available from <http://www.wisdom.weizmann.ac.il/users/oded/pp.htm>
- [15] S. Goldwasser, L. Levin, *Fair Computation of General Functions in Presence of Immoral Majority*, Crypto '90, LNCS 537, Springer-Verlag, 1991, 77–93
- [16] S. Goldwasser, S. Micali, *Probabilistic Encryption*, J. of Computer and System Sciences 28 (1984) 270–299
- [17] S. Goldwasser, S. Micali, R. L. Rivest, *A Digital Signature Scheme Secure Against Adaptive Chosen-Message Attacks*, SIAM J. on Computing 17/2 (1988) 281–308
- [18] S. Goldwasser, S. Micali, C. Rackoff, *The Knowledge Complexity of Interactive Proof Systems*, SIAM J. on Computing 18/1 (1989) 186–207
- [19] O. Goldreich, S. Micali, A. Wigderson, *How to Play any Mental Game—Or—A Completeness Theorem for Protocols with Honest Majority*, 19th Symp. on Theory of Computing (STOC), ACM, 1987, 218–229
- [20] M. Hirt, U. Maurer, *Player Simulation and General Adversary Structures in Perfect Multiparty Computation*, J. of Cryptology 13/1 (2000) 31–60
- [21] P. Lincoln, J. Mitchell, M. Mitchell, A. Scedrov, *A Probabilistic Poly-Time Framework for Protocol Analysis*, 5th Conf. on Computer and Communications Security, ACM, 1998, 112–121
- [22] P. Lincoln, J. Mitchell, M. Mitchell, A. Scedrov, *Probabilistic Polynomial-Time Equivalence and Security Analysis*, Formal Methods '99, LNCS 1708, Springer-Verlag, 1999, 776–793
- [23] N. Lynch, *Distributed Algorithms*, Morgan Kaufmann, San Francisco 1996
- [24] N. Lynch, *I/O Automaton Models and Proofs for Shared-Key Communication Systems*, 12th Computer Security Foundations Workshop (CSFW), IEEE, 1999, 14–29
- [25] S. Micali, P. Rogaway, *Secure Computation*, Crypto '91, LNCS 576, Springer-Verlag, 1992, 392–404
- [26] B. Pfitzmann, M. Schunter, M. Waidner, *Cryptographic Security of Reactive Systems*, Electronic Notes in Theoretical Computer Science (ENTCS) 32 (2000), <http://www.elsevier.nl/cas/tree/store/tcs/free/noncas/pc/menu.htm>
- [27] B. Pfitzmann, M. Schunter, M. Waidner, *Secure Reactive Systems*, IBM Research Report RZ 3206 (#93252), IBM Research Division, Zürich, May 2000
- [28] B. Pfitzmann, M. Schunter, M. Waidner, *Provably Secure Certified Mail*, IBM Research Report RZ 3207 (#93253), IBM Research Division, Zürich, August 2000
- [29] B. Pfitzmann, M. Waidner, *A General Framework for Formal Notions of “Secure” System*, Hildesheimer Informatik-Berichte 11/94, Universität Hildesheim, April 1994, available at <http://www.semper.org/sirene/lit/abstr94.html#Pfw94>
- [30] B. Pfitzmann, M. Waidner, *Composition and Integrity Preservation of Secure Reactive Systems*, 7th Conf. on Computer and Communications Security, ACM, 2000, 245–254
- [31] B. Pfitzmann, M. Waidner, *A Model for Asynchronous Reactive Systems and its Application to Secure Message Transmission*, IBM Research Report RZ 3304 (#93350) 12/11/2000, IBM Research Division, Zurich, December 2000 and IACR Cryptology ePrint Archive 2000/066, December 2000, <http://eprint.iacr.org/>
- [32] C. Rackoff, D. R. Simon, *Non-Interactive Zero-Knowledge Proof of Knowledge and Chosen Ciphertext Attack*, Crypto '91, LNCS 576, Springer-Verlag, 1992, 433–444
- [33] R. Segala, N. Lynch, *Probabilistic Simulations for Probabilistic Processes*, Concur '94, LNCS 836, Springer-Verlag, 1994, 481–497
- [34] A. C. Yao, *Protocols for Secure Computations*, 23rd Symp. on Foundations of Computer Science (FOCS), IEEE, 1982, 160–164
- [35] A. C. Yao, *Theory and Applications of Trapdoor Functions*, 23rd Symp. on Foundations of Computer Science (FOCS), IEEE, 1982, 80–91