

Unifying Simulatability Definitions in Cryptographic Systems under Different Timing Assumptions^{*}

Michael Backes

IBM Zurich Research Laboratory, Rüschlikon, Switzerland

Abstract

The cryptographic concept of simulatability has become a salient technique for faithfully analyzing and proving security properties of arbitrary cryptographic protocols. We investigate the relationship between simulatability in synchronous and asynchronous frameworks by means of the formal models of Pfitzmann et al., which are seminal in using this concept in order to bridge the gap between the formal-methods and the cryptographic community. We show that the synchronous model can be seen as a special case of the asynchronous one with respect to simulatability, i.e., we present an embedding from the synchronous model into the asynchronous one that we show to preserve simulatability. We show that this result allows for carrying over lemmas and theorems that rely on simulatability from the asynchronous model to its synchronous counterpart without any additional work, hence future work on enhancing simulatability-based models can concentrate on the more general asynchronous case.

Key words: Probabilistic systems, security, simulatability, cryptography, synchronous / asynchronous

1 Introduction

In recent times, the analysis of cryptographic protocols has been getting more and more attention, and the demand for general frameworks for representing cryptographic protocols and the security requirements of cryptographic tasks has been rising. Existing frameworks are either motivated by the complexity-theoretic view on cryptography, which aims at proving cryptographic protocols with respect to the cryptographic semantics, or they are motivated by the view of the formal-methods

^{*} An extended abstract of this work appeared in *Proc. 14th international conference on concurrency theory (CONCUR 2003)*.

Email address: mbc@zurich.ibm.com (Michael Backes).

community, which aims at capturing abstractions of cryptography in order to make such protocols accessible for formal verification. Frameworks built on abstractions of cryptography will be further dealt with in the related literature along with a discussion on the cryptographic justification of these abstractions.

For living up to the probabilistic nature of cryptography, a framework for dealing with actual cryptography necessarily has to be able to deal with probabilistic behaviors. The standard understanding in well-known, non security-specific probabilistic frameworks like [1,2] is that the order of events is fixed by means of a probabilistic scheduler that has full information about the system. In contrast to that, the standard understanding in cryptology (closest to a rigorous definition in [3]) 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 [1]. However, if one splits 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 typical solution is a distributed definition of scheduling which allows machines that have been scheduled to schedule certain (statically fixed) other machines themselves.

Based on these requirements, several general definitions of secure protocols were developed over the years, e.g. [4–13], which are all potential candidates for such a framework. To allow for a faithful analysis of cryptographic protocols, it is well-known that such models not only have to capture probabilistic behaviors, but also complexity-theoretically bounded adversaries as well as a reactive environment of the protocol, i.e., continuous interaction with the users and the adversary. Unfortunately, most of the above work does not live up to these requirements in spite of its generality, mainly since it concentrates on the task of secure function evaluation, which does not capture a reactive environment. Currently, the models of Pfitzmann et al. [8,11,13] and Canetti [12], which have been developed concurrently but independently, seem to be establishing themselves as the standard models for sound protocol analysis and design.

Regarding the underlying definition of time, such models can be split into synchronous and asynchronous ones. In synchronous models [8], time is assumed to be expressible in rounds, whereas asynchronous scenarios [11–13] do not impose any assumption on time. This makes asynchronous scenarios attractive since no assumption is made about network delays and the relative execution speed of the parties. Moreover, the notion of rounds is difficult to justify in practice as it seems to be very difficult to establish them for the Internet for example. This attractiveness is substantiated by a large body of literature on asynchronous cryptographic protocols, e.g., [14,15]. However, time guarantees are sometimes explicitly desired, e.g., on when a process can abort. Hence assumptions have to be made in this case, which induce a certain amount of synchrony again. This sometimes makes a synchronous assumption of time nevertheless necessary in practice, e.g., in Ker-

beros [16].

Hence researchers usually restrict their attention to one definition of time, or they are driving double-tracked by maintaining two separate models which, however, presupposes proving every theorem for both models. This is not nice. An alternative approach, taken in this work, is to show that the synchronous model can be regarded as a special case of an asynchronous one, and hence does not have to be further advanced separately, but still can be used to conveniently express synchronous protocols.

Although this idea might not be surprising, it is very difficult to achieve since it turns out that carrying over general results from the asynchronous to the embedded synchronous model presupposes the possibility of (at least partially) reversing the considered embedding. Recall that suitable frameworks, especially the framework of Pfitzmann et al., have a distributed scheduling which significantly complicates this reversion.

Formally, a special case means that there is an embedding from the synchronous model into the asynchronous model that preserves a desired property. Which property has to be preserved depends on the goals to strive for. For cryptographic protocols, the property of *simulatability* stands out. Simulatability captures the notion of a cryptographically secure implementation and serves as a link to the formal-methods community, which typically only hold a top-level view of cryptography, where cryptographic primitives are replaced by deterministic abstractions. A more comprehensive discussion of simulatability and its relationship to protocol verification work done by the formal-methods community is given in the paragraph on related literature below.

In the following, we investigate the synchronous and asynchronous models of Pfitzmann et al. [8,11,13], which are seminal in using the concept of simulatability to bridge the gap between the formal-methods and the cryptographic community. We show that the synchronous model can be embedded in the asynchronous model such that simulatability is preserved by this embedding, i.e., if two systems fulfill the simulatability relation in the synchronous model, their respective images fulfill the relation in the asynchronous model and vice versa. We show that this result allows for carrying over lemmas and theorems from the asynchronous case to the synchronous case without proving them twice, hence future work on enhancing simulatability-based models can concentrate on the more general asynchronous case. We are confident that this result helps to make future protocol analysis in these models more convenient and more efficient.

Moreover, we believe that our approach for establishing the embedding and its properties can be successfully used for other models with only minor changes. Especially the asynchronous model of Canetti is surely worth to be investigated. However, his corresponding synchronous model [10] is still specific for secure function

evaluation; hence adopting it to a reactive environment is a necessary prerequisite for this future work. The lack of such a reactive synchronous model was – besides the fact that the models of Pfitzmann et al. are more rigorously defined than the one of Canetti – our main reason why we decided to base our work on the model of Pfitzmann et al.

Related Literature. If cryptographic protocols should be verified using formal methods, some kind of abstraction is needed as the underlying reduction proofs of cryptography are still out of scope of current verification techniques.¹ This abstraction is usually based on the so-called Dolev-Yao abstraction [20], which considers cryptographic primitives, e.g., E for encryption and D for decryption, as operators in a free algebra where only predefined cancellation rules hold. For instance, twofold encryption of a message m does not yield another message from the basic message space but the term $E(E(m))$. A typical cancellation rule is $D(E(m)) = m$. This abstraction simplifies proofs of larger protocols considerably, and it gave rise to a large body of literature on analyzing the security of protocols using techniques for formal verification of computer programs (a very partial list of work includes [21–30]).

Since this line of work turned out to be very successful, the interesting question arose whether these abstractions are indeed justified from the view of cryptography, i.e., whether properties proved for the abstractions are still valid for the cryptographic implementation. Such cryptographic underpinnings of a Dolev-Yao model were first addressed by Abadi and Rogaway in [31]. However, they only handled passive adversaries and symmetric encryption. The protocol language and security properties handled were extended in [32,33], but still only for passive adversaries. This excludes most of the typical ways of attacking protocols, e.g., man-in-the-middle attacks and 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 in [34]. Based on the specific Dolev-Yao model whose soundness was proven in [34], the well-known Needham-Schroeder-Lowe protocol was proved in [35]. 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

¹ Efforts are also under way to formulate syntactic calculi for dealing with probabilism and polynomial-time considerations, in particular [17,7,18,19] 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 should be seen as complementary to, rather than competing with, the approach of getting simple deterministic abstractions of cryptography and working with those wherever cryptography is only used in a blackbox way.

this protocol was concurrently developed by Warinschi [36]. 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 [37,38] 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 and symmetric encryption. Laud [39] 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. [40,41] and Micciancio and Warinschi [42] have recently also given a cryptographic underpinning under active attacks. Their results are considerably weaker than the one in [34] 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 [42] further points out several possible extensions of their work which all already exist in the earlier work of [34]. Guttman et al. [43] show that the probability of two executions of the same protocol – either executed in a Dolev-Yao-like framework or using real cryptographic primitives – may deviate from each other at most for a certain bound. However, their results are specific for the Wegman-Carter system so far. Moreover, as this system is information-theoretically secure, its security proof is much easier to handle than primitives with security guarantees only against computationally bounded adversaries since no reduction proofs against underlying number-theoretic assumptions have to be made. Some further approaches for special security goals or primitives are [44,33].

The first full justification of a Dolev-Yao model presented in [34] was achieved by exploiting the concept of simulatability, which serves as a cryptographically sufficient relationship between abstract specifications and cryptographic implementations, i.e., abstractions which can be shown to simulate a given implementation in a particular sense are known to be sound with respect to the security definitions of cryptography. Simulatability was first invented for multi-party function evaluation [45,4,6,5,10], i.e., systems with only one initial input set and only one output set. An extension to a reactive scenario, where participants can make new inputs many times, e.g., start new sessions like key exchanges, was first fully defined in [46], with extensions to asynchronous systems in [11–13]. Each of the three considered models was already successfully used to build up sound abstractions of various cryptographic primitives like secure channels [11,12], certified mail [47], or key exchange [48,49].

Comparing the models of Canetti and Pfitzmann et al., we can first state that both models enjoy very general composition theorems (where the first composition theorems in [50,11] were superseded by the theorem in [12], and again by the one in [51]). Now on the one hand, Canetti’s model has been used to address more

abstractions of stand-alone cryptographic primitives so far like secure multi-party computation [52] or commitments [53]. On the other hand, the asynchronous model of Pfitzmann et al. was used to solve the long-standing open problem of justifying a Dolev-Yao type model of cryptography as used in virtually all automated protocol provers: the aforementioned *cryptographic library* from [34]. This library is a flexible toolbox for constructing abstract nested cryptographic terms and for using them in arbitrary protocols, together with a cryptographic realization provably secure under arbitrary active attacks in the standard model of cryptography. Together with composition and preservation theorems of the underlying model, the library serves as the foundation for machine-assisted reasoning about cryptographic protocols while nevertheless providing a provably secure implementation. Furthermore, the models of Pfitzmann et al. are more rigorously defined and early examples of tool-supported proofs in their models exist [54,55], using PVS [56].

Outline. In Section 2 we review the reactive models for synchronous and asynchronous time. In Section 3, we explain how the embedding works and give a rigorous definition. Starting with a proof sketch of the first embedding theorem in Section 4 (there will be two of them) and some lemmas capturing essential steps in the theorem’s proof, we fade to the embedding theorems in Section 5. In conjunction, both theorems allow for carrying over theorems from the asynchronous to the synchronous case, which is shown in Section 6 by means of an example.

2 Review of the Reactive Models in Synchronous and Asynchronous Networks

In this section we review the synchronous and the asynchronous model for probabilistic reactive systems as introduced in [8] and [11,13], respectively. Several definitions are only sketched, whereas those that are essential for understanding our upcoming results are given in full detail. To simplify the basic understanding of these models, we start with an informal overview of the more complex asynchronous model and the distributed scheduling scheme.

2.1 Informal Overview of the Asynchronous Model

We consider sets of asynchronously communicating probabilistic state machines; such sets are called *collections* of machines. The left-hand side of Figure 1 sketches a collection of three machines connected via channels represented by solid arrows. To model asynchronous timing, messages sent between the machines stay on their respective channel until they are scheduled. Technically, each channel contains an additional machine called a buffer, which stores messages in transit. This is shown




Fig. 1. A collection of three machines is shown on the left. Solid arrows represent channels. The dashed arrow depicts that M_1 schedules messages on the channel from M_2 to M_3 . Each channel implicitly contains a buffer for storing messages in transit, shown on the right.

on the right-hand side of Figure 1. When M_2 sends a message to M_3 , this message is stored in the buffer. An incoming message at a clock channel for the buffer, represented by the dashed arrow, is interpreted as a number i , and the i -th message in the buffer is removed and output to M_3 . Buffers need not be specified explicitly; a completion operator adds all necessary buffers to a collection of normal machines.

A distributed scheduling scheme that allows for expressing all realistic scenarios is achieved by allowing a machine that has been scheduled to schedule certain other machines itself. This is done by giving the machine the control over the clock channels of certain buffers. In Figure 1, the machine M_1 can schedule messages sent from M_2 to M_3 , while the channels between M_1 and M_2 show procedure-call-style local interaction. Where one wants to express that an adversary schedules everything, one simply gives the adversary all the scheduling rights. Problems with purely adversarial scheduling were already noted in [57]; hence they schedule secure channels with uniform probability before adversary-chosen events. However, that introduces a certain amount of global synchrony. Furthermore, the considered model does 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.

Probability spaces for runs are defined in detail for such collections of machines, as well as the view of a subset of the machines. These definitions are useful beyond the more security-specific system classes considered later. Further, the Turing-machine realization and runtime considerations are defined in this generality.

Security-specific *structures* are defined as collections of machines with distinguished service ports for the honest users. Such structures are augmented by arbitrary machines H and A representing the honest users and the adversary, who can interact. One then speaks of a *configuration*. In the presence of adversaries, the structure of correct machines running may not be the *intended structure* that the designer originally planned. For instance, some machines might have been corrupted; hence they are missing from the actual structure and the adversary took over their connections. This is modeled by defining a *system* as a set of possible actual structures.




Fig. 2. Ports and buffers.

2.2 General System Model

In the following we consider a finite alphabet Σ and some special symbols $!, ?, \leftrightarrow, \triangleleft \notin \Sigma$ that will be used to express different ports of machines. For $s \in \Sigma^*$ and $l \in \mathbb{N}_0$, we define $s \upharpoonright_l$ to be the l -letter prefix of s .

Machines can exchange messages with each other via *ports*. Intuitively, a port is a possible attachment point for a channel when a machine of Figure 1 is considered in isolation. As in many other models, channels in collections of machines are specified implicitly by naming conventions on the ports. Figure 2 gives an overview of the naming scheme; it can be seen as a yet more detailed view of the right-hand side of Figure 1. The name of a port (here q) serves as an identifier and will later be used to define which ports are connected to each other. Inspired by the CSP-notation [58], input and output ports are represented by the symbols $?$ and $!$, respectively. These ports are used for “usual” message transmission, whereas the ports $q \leftrightarrow ?$, $q \triangleleft !$, $q \triangleleft ?$, and $q \triangleleft !$ are used for the distributed scheduling. In the following, we call the port $q \triangleleft !$ the *clock-out port* for buffer q . In the synchronous model, buffers do not exist nor do the “scheduling” ports $q \leftrightarrow ?$, $q \triangleleft !$, $q \triangleleft ?$, and $q \triangleleft !$.

As the *low-level complement* q^c of a port q (either in- or output port) we denote the port with which it connects according to Figure 2, i.e., $q \triangleleft !^c := q \triangleleft ?$, $q !^c := q \leftrightarrow ?$, $q \leftrightarrow !^c := q ?$, and vice versa. The *high-level complement* q^C of a port q denotes the connecting port without the buffer, i.e., $q !^C = q ?$ and vice versa. For a set or a sequence P of ports, let $\text{in}(P)$ and $\text{out}(P)$ denote the subset or subsequence of P consisting of the input ports or the output ports of P , respectively.

After introducing ports, we now define *machines*. The primary machine model is probabilistic state-transition machines, similar to probabilistic I/O automata as in [59,60]. Other terms for such machines are extended finite-state automata or state-transition machines. For the computational complexity aspects, implementations of such machines are defined by probabilistic interactive Turing machines. Turing machines are not used as the sole or primary model, in contrast to prior cryptographic literature, because the I/O automata allows for expressing non-cryptographic protocol parts and abstractions from cryptography in a well-defined way unencumbered with Turing-machine details. This is important for the desired accessibility of the resulting model to existing theorem provers and model checkers. The model makes

one addition to individual machines compared with other I/O automata models, in order to enable machines to have polynomial runtime independent of their environment without being automatically vulnerable to denial-of-service attacks by long messages: It allows state-dependent *length bounds* on the inputs that a machine will read from each channel.

A machine has a *sequence of ports*, containing both input ports and output ports, and a set of *states*, comprising sets of *initial* and *final states*. If a machine is switched, it receives an input tuple at its input ports and performs its *transition function* yielding a new state and an output tuple in the deterministic case, or a finite distribution over the set of states and possible outputs in the probabilistic case. Furthermore, each machine has state-dependent bounds on the length of the inputs accepted at each port to enable flexible enforcement of runtime bounds. The parts of an input that are beyond the length bound are ignored. The value ∞ denotes that arbitrarily long inputs are accepted.

Definition 1 (Machines) A *machine* is a tuple

$$M = (\text{name}_M, \text{Ports}_M, \text{States}_M, \delta_M, l_M, \text{Ini}_M, \text{Fin}_M)$$

of a *name* $\text{name}_M \in \Sigma^+$, a *finite sequence of ports* Ports_M , a set $\text{States}_M \subseteq \Sigma^*$ of *states*, a *probabilistic state-transition function* δ_M , a *length function* $l_M : \text{States}_M \rightarrow (\mathbb{N} \cup \{\infty\})^{|\text{in}(\text{Ports}_M)|}$, and sets $\text{Ini}_M, \text{Fin}_M \subseteq \text{States}_M$ of *initial* and *final states*. Its *input set* is $\mathcal{I}_M := (\Sigma^*)^{|\text{in}(\text{Ports}_M)|}$; the i -th element of an input tuple denotes the input at the i -th input port. Its *output set* is $\mathcal{O}_M := (\Sigma^*)^{|\text{out}(\text{Ports}_M)|}$. The empty word, ϵ , denotes no in- or output at a port. δ_M maps each pair $(s, I) \in \text{States}_M \times \mathcal{I}_M$ to a finite distribution over $\text{States}_M \times \mathcal{O}_M$. If $s \in \text{Fin}_M$ then $l_M(s) = (0, \dots, 0)$; if $I = (\epsilon, \dots, \epsilon)$ then $\delta_M(s, I) = (s, (\epsilon, \dots, \epsilon))$ deterministically. Inputs are ignored beyond the length bounds, i.e., $\delta_M(s, I) = \delta_M(s, I \upharpoonright_{l_M(s)})$ for all $I \in \mathcal{I}_M$, where $(I \upharpoonright_{l_M(s)})_i := I_i \upharpoonright_{l_M(s)_i}$ for all i .

In the text, we often write “M” also for name_M . For a set M of machines, let $\text{ports}(M)$ denote the set of ports of all machines $M \in M$. We call a machine M a *black-box submachine* of a machine M' if the machine M' has access to the state-transition function δ_M of M , i.e., it can execute δ_M for the current state of the machine and arbitrary inputs. In order to concisely denote specific input and output tuples of a machine M , we introduce some additional notation. Let $(p_1?, \dots, p_n?) := \text{in}(\text{Ports}_M)$, let $P := (p_{i_1}?, \dots, p_{i_j}?)$ denote a subsequence of $(p_1?, \dots, p_n?)$, and let $(v_i)_{i \in \{1, \dots, j\}} \in (\Sigma^*)^j$. If the sequence of input ports of M is clear from the context, we define $\mathcal{I}_{p_{i_1}?, \dots, p_{i_j}?, v_1, \dots, v_j}$ to be the tuple \mathcal{I} of length n with $\mathcal{I}_{i_l} = v_l$ for all $l \in \{1, \dots, j\}$ and $\mathcal{I}_l = \epsilon$ for all $l \in \{1, \dots, n\} \setminus \{i_1, \dots, i_j\}$. In the special case $P = ()$ or $v_i = \epsilon$ for all i , i.e., in case of an all-empty input, we write \mathcal{I}_ϵ . Outputs are defined similarly.

For computational aspects, a machine M is regarded as implemented by a probabilistic interactive Turing machine as introduced in [61]. We refer to [13] for the

precise definition of the implementation. The main reason to introduce a Turing-machine realization of the machine model is to define complexity notions. A machine is called *polynomial-time* if its Turing machine implementation only needs time polynomial in its initial worktape content, independent of all inputs on communication tapes, i.e., if there exists a polynomial Q such that all possible runs of the Turing machine are of length at most $Q(k)$, where k is the length of the initial worktape content.

After introducing individual machines, we now focus on *collections* of finitely many machines, with the intuition that these machines interact. A collection \mathcal{C} of machines is a finite set of machines with pairwise different machine names and disjoint sets of ports. A port of a collection is called *free* if its connecting port is not in the collection. The free ports of a collection \mathcal{C} are denoted as $\text{free}(\mathcal{C})$. Given a collection of machines in the asynchronous model, we want to add buffers for all channels to model asynchronous timing. This is modeled by the *completion* $[\mathcal{C}]$ of a collection \mathcal{C} . The completion is the union of all machines of \mathcal{C} and the buffers needed for every channel. In the asynchronous model, a collection \mathcal{C} is called *closed* if its completion $[\mathcal{C}]$ has no free ports except a special master clock-in port $\text{clk}^?$, i.e., $\text{free}([\mathcal{C}]) = \{\text{clk}^?\}$. When we define the interaction of several machines, the master clock-in port will belong to a distinguished machine called the *master scheduler* which is used to resolve situations where the interaction cannot proceed. In the synchronous case, we demand $\text{free}(\mathcal{C}) = \emptyset$.

For security purposes, special collections are needed, because an adversary may have taken over parts of the initially intended system, e.g., different situations have to be captured depending on which and how many users are considered as being malicious. Therefore, a system consists of several possible remaining structures.

Definition 2 (Structures and Systems) A *structure* is a pair $\text{struc} = (M, S)$ where M is a collection of non-buffer machines called *correct machines*, and $S \subseteq \text{free}(M)$ is called *specified ports*. If M is clear from the context, let $\bar{S} := \text{free}(M) \setminus S$. We call $\text{forb}(M, S) := \text{ports}(M) \cup \bar{S}^C$ the *forbidden ports*, i.e., those ports that the honest user is forbidden to have. A *system* Sys is a set of structures. It is polynomial-time iff all machines in all its collections M are polynomial-time.

The separation of the free ports into specified ports and others is an important feature of the upcoming security definitions. The specified ports are those where a certain service is guaranteed. Concretely, specified ports will later be used to connect a user machine to the structure.

Note that this definition is valid for both the synchronous and the asynchronous case. In particular, buffers do not have to be explicitly included in the specification of a system, e.g., in the specification of a cryptographic protocol that one wants to analyze, but the completion operator will be used instead. The different timing assumptions stem from the different definitions of runs which we will introduce in

the following.

A structure can be completed to a *configuration* by adding machines H and A, modeling the joint honest users and the adversary, respectively. The machine H is restricted to the specified ports S , A connects to the remaining free ports of the structure and both machines can interact, e.g., in order to model active attacks. In the asynchronous case, buffers are additionally added to close the collection. Moreover, the initial state of all machines is isomorphic to the natural numbers which allows for letting the machines run on input the same security parameter in the subsequently described run algorithm.

Definition 3 (Configurations) A *configuration* of a system Sys is a tuple $conf = (M, S, H, A)$ where $(M, S) \in Sys$ is a structure, $M \cup \{H, A\}$ is a closed collection, $ports(H) \cap forb(M, S) = \emptyset$, and $Ini_M = \{1\}^*$ for all $M \in M \cup \{H, A\}$. The set of configurations is written $Conf(Sys)$. The set of configurations of Sys with polynomial-time user H and adversary A is called $Conf_{poly}(Sys)$. The index $_{poly}$ is omitted if it is clear from the context.

2.3 Capturing Asynchronous Runs

For a configuration, both models define a probability space of runs (sometimes called *traces* or *executions*). In the asynchronous model, the collection contains a unique master scheduler X since the collection is closed. Machines switch sequentially, i.e., we have exactly one active machine M at any time. If this machine has clock-out ports, it can select the next message to be delivered by scheduling a buffer via one of these clock-out ports. If a message exists at the respective position of the buffer's internal queue, it is delivered by the buffer and the unique receiving machine is the next active machine. If M tries to schedule multiple messages, only one is taken, and if it schedules none or the message does not exist, the master scheduler X becomes active.

Definition 4 (Asynchronous Runs and Views) For a given configuration $conf = (M, S, H, A)$ with master scheduler X (which is uniquely determined by having the master-clock in-port $clk^{cl?}$), set $C := [M \cup \{H, A\}]$. Runs and their probability spaces are defined inductively by the following algorithm for each tuple $ini \in \{(1^k)_{M \in M \cup \{H, A\}}\} \subset \times_{M \in C} Ini_M$ of initial states of the machines of C.

The probability space of *runs* is defined inductively by the following algorithm. It has a variable r for the resulting *run*, an initially empty list, a variable M_{CS} (“current scheduler”) over machine names, initially $M_{CS} := X$, and treats each port as a variable over Σ^* , initialized with ϵ except for $clk^{cl?} := 1$. Probabilistic choices only occur in Phase (1).

- (1) *Switch current scheduler:* Switch machine M_{CS} , i.e., set $(s', O) \leftarrow \delta_{M_{CS}}(s, I)$

for its current state s and input port values I . Then assign ϵ to all input ports of M_{CS} .

- (2) *Termination*: If X is in a final state, the run stops.
- (3) *Buffer messages*: For each simple output port $q!$ of M_{CS} , in their given order, switch buffer q with input $q^{\leftrightarrow?} := q!$, cf. Figure 2. Then assign ϵ to all these ports $q!$ and $q^{\leftrightarrow?}$.
- (4) *Clean up scheduling*: If at least one clock-out port of M_{CS} has a value $\neq \epsilon$, let $q^{\triangleleft!}$ denote the first such port and assign ϵ to the others. Otherwise let $clk^{\triangleleft?} := 1$ and $M_{CS} := X$ and go back to Phase (1).
- (5) *Scheduled message*: Switch q with input $q^{\triangleleft?} := q^{\triangleleft!}$ (cf. Figure 2), set $q^? := q^{\triangleleft!}$ and then assign ϵ to all ports of q and to $q^{\triangleleft!}$. Let $M_{CS} := M'$ for the unique machine M' with $q^? \in \text{ports}(M')$. Go back to Phase (1).

Whenever a machine (this may be a buffer) with name $name_M$ is switched from (s, I) to (s', O) , we add a *step* $(name_M, s, I, s', O)$ with $I' := I \upharpoonright_{l_M(s)}$ to the run r , except if s is final or $I' = (\epsilon, \dots, \epsilon)$. This gives a random variable for each tuple ini , i.e., for each value k of the security parameter denoted as $run_{conf,k}$. Hence we obtain a family of random variables

$$run_{conf} = (run_{conf,k})_{k \in \mathbb{N}}.$$

The *view* of a subset $M \subset C$ in a run r is the restriction of r to M , i.e., the subsequence of all steps $(name_M, s, I, s', O)$, where $name_M$ is the name of a machine $M \in M$. This gives a family of random variables

$$view_{conf}(M) = (view_{conf,k}(M))_{k \in \mathbb{N}}.$$

For a singleton $M = \{H\}$ we write $view_{conf}(H)$ instead of $view_{conf}(\{H\})$.

This still rather informal definition of runs can naturally be formalized using transition probabilities, which induce probability spaces over the finite sequences of steps similar to Markov Chains. The extension to infinite sequences can then be achieved using well-established results of measure theory and probability theory, cf. Section 5 of [62]. It is further easy to show that views of polynomial-time machines are of polynomial size, i.e., that the length of any trace that occurs with non-zero probability according to the considered view is bounded by a polynomial in the security parameter.

2.4 Capturing Synchronous Runs

In the synchronous model, ports, machines, collections, structures, and systems are defined similar to the asynchronous model. The only exception is that there are no clock ports and no buffers, which have only been included to model asynchronous timing, i.e., corresponding ports $p^?$ and $p!$ are directly connected. The main differ-

ence is the definition of runs. Instead of our asynchronous run algorithm (cf. Definition 4), runs are defined using *rounds* which is the usual concept in synchronous scenarios. Every global round is again divided into n so-called subrounds, and there is a mapping κ , called *clocking scheme*, from the set $\{1, \dots, n\}$ into the powerset of considered machines, i.e., the machines of the structure, the user, and the adversary. $\kappa(i)$ denotes which machines switch in subround i . After finishing the n -th subround, the run starts the first subround of the next global round. At the beginning of each subround, all messages from the previous subround are transported from the output ports to the connected input ports. After that, each machine of $\kappa(i)$ switches with its current inputs yielding a finite distribution over the set of states and the set of possible outputs.

Definition 5 (Clocking Scheme) A clocking scheme κ for a configuration (M, S, H, A) and $n \in \mathbb{N}$ is a mapping from the set $\{1, \dots, n\}$ to the powerset of $M \cup \{H, A\}$, i.e., it assigns each number a subset of the machines.

Definition 6 (Synchronous Runs and Views) Given a configuration $conf = (M, S, H, A)$ along with a clocking scheme κ for $n \in \mathbb{N}$, runs are defined as follows: Each global round i has n subrounds, where we denote the j -th subround of global round i by $[i.j]$. In subround $j \in \{1, \dots, n\}$ all machines $M \in \kappa(j)$ switch simultaneously, i.e., each state-transition function δ_M is applied to M 's current input yielding a new state and output (probabilistically). The output at a port $p!$ is available as input at $p?$ until the machine with port $p?$ is switched. If several inputs arrive until that time, they are concatenated. Similar to the asynchronous case, this gives a family of random variables

$$run_{conf, \kappa} = (run_{conf, \kappa, k})_{k \in \mathbb{N}}.$$

More precisely, each *run* is a function mapping each triple $(M, i, j) \in M \cup \{H, A\} \times \mathbb{N} \times \{1, \dots, n\}$ to a quadruple (s, I', s', O) of the old state, inputs (with $I' := I \upharpoonright_{l_M(s)}$ again), new state, and outputs of machine M in subround $[i.j]$, with $I' \equiv \epsilon$, $O \equiv \epsilon$, and $s = s'$ if M is not switched in this subround. The *view* of a subset $M \subset M \cup \{H, A\}$ in a run r is the restriction of r to $M \times \mathbb{N} \times \{1, \dots, n\}$. This gives a family of random variables

$$view_{conf, \kappa}(M) = (view_{conf, \kappa, k}(M))_{k \in \mathbb{N}},$$

and we omit the subscript κ if it is clear from the context.

Again, the view of a polynomial-time machine can easily be shown to be of polynomial size.

Remark 7 *Alternatively, we can consider runs as a sequence of seven-tuples (M, i, j, s, I', s', O) for ascending values of i and j . More formally, we first have all tuples $(M, 1, 1, s, I', s', O)$ for $M \in \kappa(1)$. The order of these tuples can be chosen arbitrary since they switch simultaneously and do not influence each other.*

After that, we have the steps $(M, 1, 2, s, I', s', O)$ for all $M \in \kappa(2)$ and so on, until we finally have steps of the form $(M, 1, n, s, I', s', O)$ for all $M \in \kappa(n)$. We then continue with $(M, 2, 1, s, I', s', O)$ etc. Obviously, this characterization of runs is equivalent to the original one (we just expanded the function), but it is better suited for relating synchronous runs and “corresponding” asynchronous runs, which we will do in our upcoming proofs.

Instead of arbitrary clocking schemes as in the above definition of runs, the authors of [8] focus on only one special clocking scheme κ , given by $(M \cup \{H\}, \{A\}, \{H\}, \{A\})$. Clocking the adversary between the correct machines is the well-known model of “rushing adversaries”, where [63] is the earliest reference that we are aware of. In [8], it has been shown that this clocking scheme does not restrict the possibilities of the adversary, hence we can use it without loss of generality. Moreover, we restrict ourselves to those configurations where the honest user and the adversary are only connected via one duplex channel. This is indeed no restriction to generality in the synchronous model, because outputs at several ports to the same machine can simply be concatenated using a separation symbol and decomposed again, respectively. In the following, we give these two channels fixed names $p_{A,H}$ and $p_{H,A}$, i.e., $p_{A,H}$! sends messages from A to H and vice versa.

2.5 Simulatability

The definition of one system securely implementing another one is based on the common concept of *simulatability*. Simulatability essentially means that whatever might happen to an honest user in a real system Sys_{real} can also happen in an ideal (abstract) system Sys_{id} : For every structure $struc_1 \in Sys_{\text{real}}$, every user H, and every adversary A_1 , there exists an adversary A_2 on a corresponding ideal structure $struc_2$ such that the view of H is indistinguishable in the two configurations. Indistinguishability (“ \approx ”) is a well-defined cryptographic notion from [64]. We only give the definition of computational indistinguishability; a more comprehensive definition is given in the Appendix.

Definition 8 (Computational Indistinguishability) Two families $(\text{var}_k)_{k \in \mathbb{N}}$ and $(\text{var}'_k)_{k \in \mathbb{N}}$ of random variables (or probability distributions) on common domains D_k are *computationally indistinguishable* (“ \approx ”) if for every algorithm Dis (the distinguisher) that is probabilistic polynomial-time in its first input,

$$|P(\text{Dis}(1^k, \text{var}_k) = 1) - P(\text{Dis}(1^k, \text{var}'_k) = 1)| \in \text{NEGL}.^2$$

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

² The class *NEGL* denotes the set of all negligible functions, i.e., $g: \mathbb{N} \rightarrow \mathbb{R}_{\geq 0} \in \text{NEGL}$ if for all positive polynomials Q , $\exists k_0 \forall k \geq k_0: g(k) \leq 1/Q(k)$.

Corresponding structures in the simulatability definition are designated by a function f from Sys_{real} to the powerset of Sys_{id} . The function f is called *valid* if it maps structures with the same set of specified ports, so that the same user can connect. For many systems there is only one possible mapping that meets this requirement, because the service ports of the structures correspond one-to-one to different sets of non-corrupted machines. This mapping is then called *canonical*. We only give the definition of simulatability based on computational indistinguishability, which captures the most common case when applying simulatability to cryptographic protocols. A more comprehensive definition based on the remaining notions of indistinguishability is again postponed to the Appendix; our results hold as well for this more general definition.

Definition 9 (Simulatability) Let systems Sys_1 and Sys_2 with a valid mapping f be given. We say $Sys_1 \geq^f Sys_2$ (*at least as secure as*) if for every polynomial-time configuration $conf_1 = (M_1, S, H, A_1) \in \text{Conf}(Sys_1)$, there exists a polynomial-time configuration $conf_2 = (M_2, S, H, A_2) \in \text{Conf}(Sys_2)$ with $(M_2, S) \in f(M_1, S)$ (and the same H) such that $\text{view}_{conf_1}(H) \approx \text{view}_{conf_2}(H)$.

This is shown in Figure 3. In the following, we augment \geq with a subscript sync




Fig. 3. Overview of the simulatability definition.

or async to distinguish the definition of the synchronous and asynchronous case. In a typical ideal system, each structure contains only one machine TH called trusted host, which serves as an ideal functionality of the real system. The machine TH is usually deterministic and maintains a very simple transition function, hence validation based on this ideal functionality is in scope of current verification techniques.

3 Idea and Definition of the Embedding

The informal idea of the embedding φ_{Sys} is to add an explicit master scheduler that should simulate the synchronous run induced by the given clocking scheme. However, due to the general distributed scheduling (cf. Definition 4), leaving the actual machines unmodified leads to non-simulatable situations, as these machines can clock themselves without ever giving control to this explicit master scheduler.

Hence, we first define a mapping φ_M from “synchronous” machines, i.e., from machines that do not have any of the scheduling ports but only ports for usual message

transmission, to “asynchronous” machines, i.e., to machines which might additionally have clock-out ports.

Intuitively, this mapping surrounds single synchronous machines with an “asynchronous coat”. More precisely, if a synchronous machine makes a transition, it obtains all inputs at once that arrived since its last scheduling, whereas in asynchronous scenarios, these inputs come one by one and have to be processed in several transitions. Thus, the surrounding asynchronous machine stores all inputs internally, until it is asked to perform the transition of its synchronous submachine. It then schedules this submachine with the collected inputs and forwards its outputs. As these asynchronous machines do not produce any clock outputs, the master scheduler can try to simulate the synchronous time by a suitable scheduling strategy.

Definition 10 (Mapping φ_M) φ_M is a mapping on single synchronous machines that assigns every machine M_{sync} an asynchronous machine $M_{\text{async}} := \varphi_M(M_{\text{sync}})$ by the following rules:

- The ports of M_{async} are given by $Ports_{M_{\text{async}}} := Ports_{M_{\text{sync}}} \circ (p_{M_{\text{sync}}})$, where \circ denotes concatenation of sequences.
- Internally, M_{async} maintains arrays $(input_store_{M_{\text{sync}}, p?})_{p? \in in(Ports_{M_{\text{sync}}})}$ over Σ^* initialized with ϵ everywhere, which are used for storing incoming messages at each port of M_{sync} .
- M_{async} has the machine M_{sync} as a blackbox submachine, i.e., it has its transition function $\delta_{M_{\text{sync}}}$.
- Internally, M_{async} maintains a superset of the states of M_{sync} (the additional states are only used to model appropriate length functions). Moreover, the initial and final states of both machines are equal.
- On input i at $p? \neq p_{M_{\text{sync}}}$: It concatenates i to the element of $input_store_{M_{\text{sync}}, p?}$, i.e., it stores all inputs until the machine M_{sync} is eventually switched. The length function for such a port $p?$ is defined as $l_{M_{\text{sync}}}(s)_{p?} - |input_store_{M_{\text{sync}}, p?}|$, where $l_{M_{\text{sync}}}(s)_{p?}$ is the length function of the machine M_{sync} at port $p?$ in its current state s and $|input_store_{M_{\text{sync}}, p?}|$ is the number of elements in $input_store_{M_{\text{sync}}, p?}$.
- On input i at $p_{M_{\text{sync}}}$: It applies the state transition function $\delta_{M_{\text{sync}}}$ on the contents of the arrays $input_store_{M_{\text{sync}}, p?}$ yielding a tuple (s', \mathcal{O}) . M_{async} now assigns ϵ to $input_store_{M_{\text{sync}}, p?}$ for all $p? \in in(Ports_{M_{\text{sync}}})$, switches to the state s' and outputs the tuple \mathcal{O} . The length function for this port is defined to be zero if the lists $input_store_{M_{\text{sync}}, p?}$ are empty for all ports $p?$; otherwise it is the runtime of M_{sync} . This case corresponds to the scheduling of the synchronous machine; the port $p_{M_{\text{sync}}}$ will be connected to the explicit master scheduler.

For a set M of synchronous machines, we define $\varphi_M(M) := \bigcup_{M_{\text{sync}} \in M} \varphi_M(M_{\text{sync}})$.

M_{async} is polynomial-time by construction iff M_{sync} is polynomial-time, since the machine M_{async} only performs a polynomially bounded number of steps between

two transitions of M_{sync} (which is especially ensured by the used length functions), since both machines always stay in the same state after a transition of the blackbox submachine, and finally since their final states are equal. We stress that making the outer machine M_{sync} polynomial-time for a polynomial submachine is not as easy as one might expect, e.g., as the outer machine may be triggered exponentially often at one port without causing the submachine to switch. Note further that the length functions of M_{async} are always large enough by construction that inputs of M_{async} are not ignored respectively truncated if they would be fully read by the machine M_{sync} .

Based on this definition, we now formalize the desired mapping φ_{Sys} on synchronous systems.

Definition 11 (Mapping φ_{Sys}) Let an arbitrary synchronous system $\text{Sys}_{\text{sync}} = \{(M_{\text{sync}}, S_{\text{sync}}) \mid \text{sync} \in I\}$ for a finite index set I and a clocking scheme κ be given. We then define

$$\varphi_{\text{Sys}}(\text{Sys}_{\text{sync}}) := \{(\varphi_M(M_{\text{sync}}) \cup \{X_{\text{sync}, \kappa}\}, S_{\text{sync}}) \mid \text{sync} \in I\}.$$

The machine $X_{\text{sync}, \kappa}$ is an explicit master scheduler that has to be added to the considered structure to model the synchronous clocking scheme κ in the asynchronous system. Its ports are given by

- $\{\text{clk}^{\leftarrow?}\}$: The master clock-in port.
- $\{\text{p}^{\leftarrow!} \mid \text{p}! \in \text{Ports}_{M_{\text{sync}}}\}$: Ports for clocking all output ports of the given structure.
- $\{\text{p}^{\leftarrow!} \mid \text{p}^? \in \text{free}(M_{\text{sync}})\}$: Ports for clocking inputs of the systems (either made by H or A).
- $\{\text{p}_{A,H}^{\leftarrow!}, \text{p}_{H,A}^{\leftarrow!}\}$: Ports for clocking the connection between A and H.³
- $\{\text{p}_M^{\leftarrow!}, \text{p}_M^{\leftarrow!} \mid M \in (M_{\text{sync}} \cup \{H, A\})\}$: Ports for clocking, i.e., giving control to, each machine.

The length functions are always set to infinity for all ports. Internally, it maintains a variable local_rnd over $\{1, \dots, n\}$ and a variable global_rnd over \mathbb{N} both initialized with 1. For the sake of readability, we describe the behavior of $X_{\text{sync}, \kappa}$ using “for”-loops. This is just a notational convention that should be understood as follows: every time $X_{\text{sync}, \kappa}$ is scheduled, it performs the next step of the loop.

- (1) **Schedule Current Machines:** For all machines $M \in \kappa(\text{local_rnd})$ output $(\text{global_rnd}, \text{local_rnd})$ at $\text{p}_M^{\leftarrow!}$, 1 at $\text{p}_M^{\leftarrow!}$. The order of the switched machines can be chosen arbitrary.
- (2) **Schedule Outgoing Buffers:** For all $M \in \kappa(\text{local_rnd})$ output 1 at every port $\text{p}^{\leftarrow!}$ with $\text{p}! \in \text{Ports}_M$. Here, the order of the switched machines can only

³ Note, that $X_{\text{sync}, \kappa}$ is defined independent from the honest user H and the adversary A, so it cannot know their ports. We therefore restricted the configuration to a fixed number and fixed names of ports between H and A (cf. Section 2.4).

be chosen arbitrary with the restriction that output ports of the adversary are scheduled first if $A \in \kappa(\text{local_rnd})$.⁴

- (3) **Switch to next Round:** Set $\text{local_rnd} := \text{local_rnd} + 1$. If $\text{local_rnd} > n$, set $\text{global_rnd} := \text{global_rnd} + 1$ and $\text{local_rnd} := 1$. Go to Phase (1).

To put it all into a nutshell, the specific master scheduler simulates the clocking scheme κ by first scheduling the machines that ought to switch in the particular subround (Step 1) and afterwards scheduling all buffers that could be influenced by outputs of these machines (Step 2). Finally, it switches to the next subround (Step 3) and continues with the first step again.

Moreover, we define a mapping φ_{conf} on synchronous configurations of a system Sys , i.e., configurations which consist of synchronous machines only, by

$$\varphi_{\text{conf}}(M_{\text{sync}}, S_{\text{sync}}, H, A) := (\varphi_M(M_{\text{sync}}) \cup \{X_{\text{sync}, \kappa}\}, S_{\text{sync}}, \varphi_M(H), \varphi_M(A)),$$

with $X_{\text{sync}, \kappa}$ given as in φ_{Sys} for the particular structure. We will in the following simply write φ instead of φ_{Sys} , φ_M , and φ_{conf} if its meaning is clear from the context.

4 Preliminary Work for the Embedding Theorems

We now have to prove that the function φ has the desired properties with respect to simulatability, i.e.,

$$\varphi_{Sys}(Sys_{\text{sync},1}) \geq_{\text{async}} \varphi_{Sys}(Sys_{\text{sync},2}) \Rightarrow Sys_{\text{sync},1} \geq_{\text{sync}} Sys_{\text{sync},2}.$$

This captures the content of our first embedding theorem. Unfortunately, the converse direction does not hold, but our second embedding theorem will state a weaker version that is still sufficient for our purpose.

4.1 Proof Overview

Before we turn our attention to the auxiliary lemmas for the embedding theorems we exemplarily present an informal description of the proof of the first embedding theorem. The proof consists of four steps. A graphical illustration is given in Figure 4.

⁴ Without this restriction, the behavior of the adversary at its switching time could depend on outputs of machines scheduled in the same subround, which would lead to non-simulatable situations.




Fig. 4. Synchronous Simulatability derived by Asynchronous Simulatability.

- (1) Starting with a synchronous configuration $conf_{sync,1} \in \text{Conf}(Sys_{sync,1})$, we apply our embedding function φ_{conf} which yields an asynchronous configuration $conf_{async,1} \in \text{Conf}(\varphi_{Sys}(Sys_{sync,1}))$. We now define a mapping σ on the runs of the asynchronous system yielding runs of the synchronous system. Intuitively, σ “compresses” an asynchronous run to its synchronous counterpart, which consists of fewer steps. We then show in Theorem 14 that

$$view_{conf_{sync,1}}(H_{sync}) = \sigma(view_{conf_{async,1}}(\varphi(H_{sync}))).$$

- (2) We now apply our precondition $\varphi_{Sys}(Sys_{sync,1}) \geq_{async}^f \varphi_{Sys}(Sys_{sync,2})$ yielding an indistinguishable configuration $conf_{async,2} \in \text{Conf}(\varphi_{Sys}(Sys_{sync,2}))$, i.e., $view_{conf_{async,1}}(\varphi(H_{sync})) \approx view_{conf_{async,2}}(\varphi(H_{sync}))$. We then show that

$$\sigma(view_{conf_{async,1}}(\varphi(H_{sync}))) \approx \sigma(view_{conf_{async,2}}(\varphi(H_{sync}))).$$

- (3) We finally reverse the function φ by removing the coating of the user and that of the machines of the structure. Since we do not know anything about the newly derived adversary $A_{async,2}$, i.e., in particular it is not required that it fits the structure imposed by the mapping φ , we define a new adversary $A_{sync,2}$ using $A_{async,2}$ as a black-box submachine, and we will show in Theorem 15 that

$$\sigma(view_{conf_{async,2}}(\varphi(H_{sync}))) = view_{conf_{sync,2}}(H_{sync}).$$

- (4) Altogether, transitivity of the relation \approx implies

$$view_{conf_{sync,1}}(H_{sync}) \approx view_{conf_{sync,2}}(H_{sync}).$$

We first take a look at the runs in a synchronous system Sys_{sync} and in its asynchronous counterpart $Sys_{async} := \varphi(Sys_{sync})$. In the following, we will simply write

S instead of S_{sync} , because the set of specified ports is not influenced by the mapping φ .

4.2 Compressing asynchronous runs to synchronous counterparts

In the following, let an arbitrary synchronous system Sys_{sync} with a clocking scheme κ and an arbitrary configuration $conf_{\text{sync}} = (M_{\text{sync}}, S, H_{\text{sync}}, A_{\text{sync}}) \in \text{Conf}(Sys_{\text{sync}})$ be given. Moreover, let an asynchronous configuration $conf_{\text{async}}$ be given which fits the form $conf_{\text{async}} = (\varphi(M_{\text{sync}}) \cup \{X_{\text{sync}, \kappa}\}, S, \varphi(H_{\text{sync}}), A')$ (i.e., $\varphi(conf_{\text{sync}})$ but with an arbitrary adversary).⁵

First of all, note that runs of $conf_{\text{async}}$ always have a prescribed structure induced by the behavior of the master scheduler $X_{\text{sync}, \kappa}$: they are built by “blocks”. The steps $(M_{\text{sync}}, i, j, s, \mathcal{I}, s', \mathcal{O})$ of the machines $M_{\text{sync}} \in M_{\text{sync}} \cup \{H_{\text{sync}}\}$ switched in round $[i, j]$ in the synchronous run are represented by the following two blocks in the asynchronous run.

- (1) The first block consists of the steps induced by clocking the machines $\varphi(M_{\text{sync}})$ with $M_{\text{sync}} \in \kappa(j)$ and A' if $A_{\text{sync}} \in \kappa(j)$, i.e., Step (1) in the definition of $X_{\text{sync}, \kappa}$. More precisely, the block is built by $|\kappa(j)|$ sub-blocks, one for every switched machine. Every sub-block is built by the following steps.
 - The first step of the sub-block is always of the form $(X_{\text{sync}, \kappa}, s_1, \mathcal{I}_{\text{clk}^?=(i,j)}, s'_1, \mathcal{O}_{\text{PM}_{\text{sync}}^!=(i,j), \text{PM}_{\text{sync}}^?=(i,j)}}$ for two arbitrary states s_1, s'_1 of $X_{\text{sync}, \kappa}$, i.e., the master scheduler schedules the machine $\varphi(M_{\text{sync}})$ respectively A' .
 - After that, we have the transition of the scheduled buffer.
 - We now have to distinguish the following two cases:
 - If $M_{\text{sync}} \neq A_{\text{sync}}$, there is a step $(\varphi(M_{\text{sync}}), s, \mathcal{I}_{\text{PM}_{\text{sync}}^?=(i,j)}, s', \delta_{M_{\text{sync}}}(input_store_{M_{\text{sync}}}))$ and steps for the receiving buffers.
 - If $M_{\text{sync}} = A_{\text{sync}}$, we have a step $(A', s, \mathcal{I}_{\text{PA}_{\text{sync}}^?=(i,j)}, s', \mathcal{O})$. If $\mathcal{O} \neq \mathcal{O}_\epsilon$ we have steps for the receiving buffers. If there are nonempty outputs at ports $p!$ and $p^!$ (which has to be a self-loop because there are no free clock-in ports in the system), there is furthermore a clocking step for this particular buffer. In this case, the adversary is scheduled again, so this sub-point of the block is repeated until the self-loop of the adversary either ends or it is repeated forever in case of divergence, i.e., we obtain a step $(A', s', \mathcal{I}', s'', \mathcal{O})$ where \mathcal{I}' is now given by $\mathcal{I}' := \mathcal{I}_{p^?=\mathcal{O}_{p!}}$ and so on.
- (2) The second block consists of the steps induced by clocking the outgoing messages of the switched machines, i.e., Step (2) in the definition of $X_{\text{sync}, \kappa}$. Now the buffers of the output ports are switched by the master scheduler. This is

⁵ Note that we investigate the more general case here that A' can be chosen arbitrarily instead of being the embedded adversary $\varphi_M(A)$. This generality will be helpful during the upcoming proofs.

done similar as in the first part with the restriction that output ports of A' are clocked first if $A_{\text{sync}} \in \kappa(j)$. The block again has $|\kappa(j)|$ sub-blocks built by the following steps.

- The first step of the sub-block is given by $(X_{\text{sync},\kappa}, s_1, \mathcal{I}_{\text{clk}^?}=1, s'_1, \mathcal{O}_{p^?=1})$ for the first output port $p! \in \text{ports}(M_{\text{sync}})$ and two arbitrary states s_1, s'_1 of $X_{\text{sync},\kappa}$.
- The step of the clocked buffer.
- In case of a nonempty output let M' denote the unique machine with $p? \in \text{ports}(M')$. We now have to distinguish two cases:
 - If $M' \neq A'$, there is a step $(M', s, \mathcal{I}', s', \mathcal{O}_\epsilon)$, where \mathcal{I}' consists of the output of $\varphi(M_{\text{sync}})$ at $p!$.
 - If $M' = A'$, we obtain a step $(A', s, \mathcal{I}', s', \mathcal{O})$, where \mathcal{I}' consists of the output of $\varphi(M_{\text{sync}})$ respectively A' at $p!$. If $\mathcal{O} \neq \mathcal{O}_\epsilon$ we have steps for the receiving buffers. If \mathcal{O} has a clocked self-loop, we proceed identical to the first block.
- The three previous steps are repeated for every output port of every machine $M_{\text{sync}} \in \kappa(j)$.

After this detailed description of the run, (i.e., its blocks) the mapping σ can be defined. Informally, it combines the blocks of all machines $M_{\text{sync}} \in \kappa(j)$ yielding the synchronous steps of every machine M_{sync} that switches in the j -th subround of the particular global round.

Definition 12 (Mapping σ) Let an arbitrary synchronous system Sys_{sync} with a clocking scheme κ and an arbitrary configuration $conf_{\text{sync}} = (M_{\text{sync}}, S, H_{\text{sync}}, A_{\text{sync}}) \in \text{Conf}(Sys_{\text{sync}})$ be given. For a given asynchronous configuration $conf_{\text{async}}$ which fits the form $conf_{\text{async}} = (\varphi(M_{\text{sync}}) \cup \{X_{\text{sync},\kappa}\}, S, \varphi(H_{\text{sync}}), A')$, we define the mapping σ on the runs of $conf_{\text{async}}$ by the following algorithm. The algorithm has internal arrays $(inputs_{M,p?})$ for $M \in \varphi(M_{\text{sync}}) \cup \{\varphi(H_{\text{sync}}), A'\}$ and $p? \in \text{in}(Ports_M)$. It goes from block to block modifying them as follows.

- (1) Every step of a buffer is deleted from the run.
- (2) The two remaining steps of the first block are modified as follows. If the scheduled machine is $\varphi(M_{\text{sync}}) \neq A'$, then the block is replaced by $(M_{\text{sync}}, i, j, s, inputs_{M_{\text{sync}}}, s', \delta_{M_{\text{sync}}}(inputs_{M_{\text{sync}}}))$. If A' is scheduled, the block is replaced by $(A', i, j, s, inputs_{A_{\text{sync}}}, s', \mathcal{O}_{A'})$. Here, s denotes the state of A' when it is switched by $X_{\text{sync},\kappa}$, and s' and $\mathcal{O}_{A'}$ are the state and the output of the last step of the block, respectively (In case of divergence, the algorithm for defining the mapping σ diverges, too.).
- (3) The algorithm starts searching through the second block doing the following. If a machine M' receives a message i at $p?$ in the second block, i is concatenated to the array $inputs_{M',p?}$.
- (4) Finally, every step of the second block is deleted from the run.

Note that all necessary information (e.g., $M_{\text{sync}}, i, j, s, s'$ etc.) is already given by

the block except for the inputs of each machine in the synchronous case. At this point, it also becomes clear why we defined the master scheduler to schedule each machine specifically with a tuple (i, j) indicating the current global and local round, since this information would otherwise not be contained in the asynchronous run.

To overcome the absence of the gathered inputs in the run, the algorithm has to collect all “partial” inputs itself in its third step, and it can use this information to calculate the outputs of each machine (although for this, it could as well use the information contained in the run). Moreover, the new blocks built by the mapping σ in one particular subround do not depend on the second block of this subround. The mapping σ is obviously also defined on the view of arbitrary subsets of machines, because the step in the first block, carrying the information of the step, and the message-receiving steps in the second block will also be part of the view of the considered machine. Furthermore, note that the mapping σ is explicitly defined for arbitrary adversaries A' (not only for $\varphi(A_{\text{sync}})$) which we will need in Theorem 15. Furthermore, the following lemma establishes a computational bound on the mapping σ in polynomial-time configurations:

Lemma 13 *If $\text{conf}_{\text{async}}$ is a polynomial-time configuration that fits the form required by Definition 12, then σ applied to the view of the honest user and the adversary is computable in polynomial-time.*

Proof In case of a polynomial configuration, especially the adversary has to be polynomial-time. This implies that there cannot be any infinite successive clocked self-loops, so the steps of every sub-block are bounded by a polynomial in the security parameter k . Moreover, both the adversary and the honest user will reach final state after a polynomial number of blocks, so the algorithm for σ applied to the view either of the honest user or the adversary only makes a polynomial number of transition, each one with a polynomial number of steps. This implies that σ is computable in polynomial-time when applied to the view of the honest user and the adversary if it is used in a polynomial-time configuration. \square

4.3 Auxiliary Theorems

The following theorem captures the first step of our proofs sketch of Section 4.1.

Theorem 14 *Let a synchronous system Sys_{sync} , a clocking scheme κ , and a configuration $\text{conf}_{\text{sync}} = (M_{\text{sync}}, S, H_{\text{sync}}, A_{\text{sync}}) \in \text{Conf}(\text{Sys}_{\text{sync}})$ be given, and set $\text{conf}_{\text{async}} := \varphi(\text{conf}_{\text{sync}})$. Then*

$$\text{view}_{\text{conf}_{\text{sync}}}(\mathbf{M}_{\text{sync}}) = \sigma(\text{view}_{\text{conf}_{\text{async}}}(\varphi(\mathbf{M}_{\text{sync}})))$$

for every $M_{\text{sync}} \in (M_{\text{sync}} \cup \{H_{\text{sync}}, A_{\text{sync}}\})$. $\text{conf}_{\text{async}}$ is polynomial-time iff $\text{conf}_{\text{sync}}$ is polynomial-time.

Proof Note that the view of $\varphi(M_{\text{sync}})$ only contains the steps of its internal black-box function-call after being modified by the mapping σ . Thus, it is sufficient to show that the inputs of the blackbox call in $\text{conf}_{\text{async}}$ and the original inputs of H_{sync} in $\text{conf}_{\text{sync}}$ are equal. It is quite easy to see that the arrays $\text{input_store}_{M_{\text{sync}}}$ and $\text{inputs}_{M_{\text{sync}}}$ are always equal if the machine M_{sync} is switched. This can easily be proven by induction over the number of (sub-)rounds. In the first round, both arrays are empty yielding a correct start of the induction. Starting with the second round, the contents of these arrays are totally determined by the inputs at the ports of M_{sync} . However, these inputs only depend on *prior* outputs of other machines M . Moreover, these outputs have to be equal because these machines used the same input tuple in both configurations, since we have $\text{input_store}_M = \text{inputs}_M$ for all $M \in M$ by induction hypothesis. Therefore, the arrays $\text{inputs}_{M_{\text{sync}}}$ and $\text{input_store}_{M_{\text{sync}}}$ must be equal at replacing the block by construction of the algorithm, so $\delta_{M_{\text{sync}}}(s, \text{inputs}_{M_{\text{sync}}}) = \delta_{M_{\text{sync}}}(s, \text{input_store}_{M_{\text{sync}}})$ also holds. We do not have to worry about the arrangement of the blocks because of the following reasons. First of all, note that we first switch all machines in a subround and schedule the outgoing messages afterwards. Moreover, messages sent by the adversary are always scheduled first if the adversary is scheduled in the considered subround. This prevents that machines which should switch simultaneously in the synchronous system influence each other in the asynchronous system in the same subround. If we did not consider this restriction, the adversary would be able to create a message that is scheduled in this particular subround, but nevertheless depends on inputs arriving in this subround.

Putting it all together, the runs induced by the mapping σ in $\text{conf}_{\text{async}}$ and the original runs are equal by definition of σ , so we finally obtain

$$\text{view}_{\text{conf}_{\text{sync}}}(M_{\text{sync}}) = \sigma(\text{view}_{\text{conf}_{\text{async}}}(\varphi(M_{\text{sync}})))$$

for an arbitrary configuration $\text{conf}_{\text{sync}} \in \text{Conf}(\text{Sys}_{\text{sync}})$, $\text{conf}_{\text{async}} := \varphi(\text{conf}_{\text{sync}})$, and an arbitrary $M_{\text{sync}} \in (M_{\text{sync}} \cup \{H_{\text{sync}}, A_{\text{sync}}\})$. As a special case, this implies

$$\text{view}_{\text{conf}_{\text{sync}}}(H_{\text{sync}}) = \sigma(\text{view}_{\text{conf}_{\text{async}}}(\varphi(H_{\text{sync}})))$$

which finishes our proof. \square

After performing this first step of the proof, asynchronous simulatability can now be applied. In order to convert the derived asynchronous configuration into a synchronous configuration again (cf. Step 3 of our proofs sketch), we present the following theorem.

Theorem 15 *Let an arbitrary synchronous system Sys_{sync} and a clocking scheme κ be given such that every machine and the honest user are clocked at most once between two successive clockings of the adversary. Furthermore, let an arbitrary configuration $conf_{\text{async}} \in \text{Conf}(\varphi(Sys_{\text{sync}}))$ of the form $conf_{\text{async}} = (\varphi(M_{\text{sync}}) \cup \{X_{\text{sync},\kappa}\}, S, \varphi(H_{\text{sync}}), A_{\text{async}})$ be given. Then there exists an adversary A_{sync} using A_{async} as a blackbox such that for $conf_{\text{sync}} := (M_{\text{sync}}, S, H_{\text{sync}}, A_{\text{sync}})$, it holds*

$$view_{conf_{\text{sync}}}(M_{\text{sync}}) = \sigma(view_{conf_{\text{async}}}(\varphi(M_{\text{sync}})))$$

for every $M_{\text{sync}} \in (M_{\text{sync}} \cup \{H_{\text{sync}}\})$. $conf_{\text{async}}$ is polynomial-time iff $conf_{\text{sync}}$ is polynomial-time.

Note, that the standard clocking scheme $(M \cup \{H\}, \{A\}, \{H\}, \{A\})$ fulfills the postulated requirement. The proof of Theorem 15 is quite technical and hence postponed to Appendix B for the sake of readability.

5 The Embedding Theorems

This section contains our two main theorems. We start with a lemma capturing some simple properties of indistinguishable random variables. The lemma is well-known and easily proved.

Lemma 16 *Indistinguishability of two families of random variables implies indistinguishability of any function σ of them. For the polynomial case, the function σ has to be polynomial-time computable. Moreover, identically distributed variables are indistinguishable and indistinguishability is an equivalence relation.*

Theorem 17 (First Embedding Theorem) *Let two arbitrary synchronous systems $Sys_{\text{sync},1}$ and $Sys_{\text{sync},2}$ with clocking schemes κ_1 and κ_2 be given such that κ_2 fulfills the property that every machine of the system and the user is clocked at most once between two successive clockings of the adversary. Furthermore, $\varphi(Sys_{\text{sync},1}) \succeq_{\text{async}}^f \varphi(Sys_{\text{sync},2})$ should hold for a valid mapping f . Then*

$$Sys_{\text{sync},1} \succeq_{\text{sync}}^{f'} Sys_{\text{sync},2},$$

where f' is derived from f by $(M_2, S_2) \in f'(M_1, S_1) \Leftrightarrow \varphi(M_2, S_2) \in f(\varphi(M_1, S_1))$.

Using the result of the previous theorems, the proof will be rather simple, cf. the illustration in Figure 4.

Proof Let an arbitrary configuration $conf_{\text{sync},1} = (M_{\text{sync},1}, S, H_{\text{sync}}, A_{\text{sync},1}) \in \text{Conf}(Sys_{\text{sync},1})$ be given.

- (1) We apply φ_{conf} on $conf_{sync,1}$ which yields a configuration $conf_{async,1} = (\varphi(M_{sync,1}) \cup \{X_{sync,1,\kappa_1}\}, S, \varphi(H_{sync}), \varphi(A_{sync,1})) \in \text{Conf}(Sys_{async,1})$. According to Theorem 14, applying the mapping σ to the runs of $conf_{async,1}$ yields

$$view_{conf_{sync,1}}(H_{sync}) = \sigma(view_{conf_{async,1}}(\varphi(H_{sync}))). \quad (1)$$

Moreover, if $conf_{sync,1}$ is polynomial-time then $conf_{async,1}$ is also polynomial-time, and the mapping σ is polynomial-time computable.

- (2) Thus, the precondition $\varphi(Sys_{sync,1}) \geq_{async}^f \varphi(Sys_{sync,2})$ can be applied which yields a configuration $conf_{async,2} = (\varphi(M_{sync,2}) \cup \{X_{sync,2,\kappa_2}\}, S, \varphi(H_{sync}), A_{sync,2}) \in \text{Conf}(Sys_{async,2})$ with

$$view_{conf_{async,1}}(\varphi(H_{sync})) \approx view_{conf_{async,2}}(\varphi(H_{sync}))$$

and $\varphi(M_{sync,2}, S) \in f(\varphi(M_{sync,1}, S))$. Moreover, in the computational case, $conf_{async,2}$ is polynomial-time, so the mapping σ is polynomial-time computable. Using Lemma 16, this yields

$$\sigma(view_{conf_{async,1}}(\varphi(H_{sync}))) \approx \sigma(view_{conf_{async,2}}(\varphi(H_{sync}))). \quad (2)$$

- (3) We now apply Theorem 15 to the configuration $conf_{async,2}$, which yields a configuration $conf_{sync,2} = (M_{sync}, S, H_{sync}, A_{sync,2}) \in \text{Conf}(Sys_{sync,2})$ with

$$\sigma(view_{conf_{async,2}}(\varphi(H_{sync}))) = view_{conf_{sync,2}}(H_{sync}). \quad (3)$$

According to Theorem 15, $conf_{sync,2}$ is a polynomial-time configuration iff $conf_{async,2}$ is polynomial.

- (4) Now Equation 1-3 together with Lemma 16 imply $view_{conf_{sync,1}}(H_{sync}) \approx view_{conf_{sync,2}}(H_{sync})$. Hence, $conf_{sync,2}$ is an indistinguishable configuration for $conf_{sync,1}$. Moreover, we have $\varphi(M_{sync,2}, S) \in f(\varphi(M_{sync,1}, S))$, i.e., $(M_{sync,2}, S) \in f'(M_{sync,1}, S)$ which yields the desired result $Sys_{sync,1} \geq_{sync}^{f'} Sys_{sync,2}$. \square

Note that the theorem is applicable to the standard clocking scheme. So far, we have shown that asynchronous simulatability among these asynchronous representations implies synchronous simulatability, i.e.,

$$\varphi_{Sys}(Sys_{sync,1}) \geq_{async} \varphi_{Sys}(Sys_{sync,2}) \Rightarrow Sys_{sync,1} \geq_{sync} Sys_{sync,2}.$$

We already briefly stated in the previous section that the converse implication does not hold in general. We had to show that for each configuration $conf_{async,1} \in \text{Conf}(\varphi_{Sys}(Sys_{sync,1}))$ there exists an indistinguishable configuration $conf_{async,2} \in \text{Conf}(\varphi_{Sys}(Sys_{sync,2}))$ provided that $Sys_{sync,1} \geq_{sync} Sys_{sync,2}$.

However, both the honest user and the adversary may have clock-out ports and they can alternately schedule each other (and also the system erratically), which we

cannot capture by a fixed synchronous clocking scheme, so we cannot exploit our assumption $Sys_{\text{sync},1} \geq_{\text{sync}} Sys_{\text{sync},2}$.

Anyhow, it is sufficient for our purpose to show that the claim holds for at least those configurations where the honest user H_{sync} fits the form $\varphi_M(H_{\text{sync}})$ for a synchronous machine H_{sync} . We denote this version of simulatability for the restricted class of users by $\geq_{\text{async},H}$ in the sequel. Looking at the proof of the first embedding theorem, it is immediately clear that the theorem also holds for the weaker precondition $\varphi_{Sys}(Sys_{\text{sync},1}) \geq_{\text{async},H} \varphi_{Sys}(Sys_{\text{sync},2})$, since we only need to derive an indistinguishable configuration for users of the special form $\varphi(H_{\text{sync}})$, and the user remains unchanged at simulatability. We can now capture the content of the second embedding theorem as

$$Sys_{\text{sync},1} \geq_{\text{sync}} Sys_{\text{sync},2} \Rightarrow \varphi_{Sys}(Sys_{\text{sync},1}) \geq_{\text{async},H} \varphi_{Sys}(Sys_{\text{sync},2}).$$

Theorem 18 (Second Embedding Theorem) *Let two arbitrary synchronous systems $Sys_{\text{sync},1}$ and $Sys_{\text{sync},2}$ with clocking schemes κ_1 and κ_2 be given such that κ_1 fulfills the property that every machine of the system and the user is clocked at most once between two successive clockings of the adversary. Furthermore, $Sys_{\text{sync},1} \geq_{\text{sync}}^f Sys_{\text{sync},2}$ should hold for a valid mapping f . Then*

$$\varphi(Sys_{\text{sync},1}) \geq_{\text{async},H}^{f'} \varphi(Sys_{\text{sync},2})$$

where f' is derived from f by $\varphi(M_2, S_2) \in f'(\varphi(M_1, S_1)) \Leftrightarrow (M_2, S_2) \in f(M_1, S_1)$.

Before we turn our attention to the actual proof, we state the following lemma which captures that we can “locally reverse” the function σ for the honest user.

Lemma 19 *Let a synchronous system Sys_{sync} , a clocking scheme κ and a configuration $conf_{\text{sync}} = (M_{\text{sync}}, S, H_{\text{sync}}, A_{\text{sync}}) \in \text{Conf}(Sys_{\text{sync}})$ be given. Let $conf_{\text{async}} = (\varphi(M_{\text{sync}}) \cup \{X_{\text{sync},\kappa}\}, S, \varphi(H_{\text{sync}}), A')$ be an arbitrary asynchronous configuration. If we now have given $\sigma(\text{view}_{conf_{\text{async}}}(\varphi(H_{\text{sync}})))$ then we can “locally reverse” the function σ for the view of the user, i.e., we can define a function σ_H^{-1} on the runs of the synchronous configuration, such that*

$$\text{view}_{conf_{\text{async}}}(\varphi(H_{\text{sync}})) = \sigma_H^{-1}(\sigma(\text{view}_{conf_{\text{async}}}(\varphi(H_{\text{sync}}))))$$

holds. If $conf_{\text{async}}$ is polynomial-time, then σ_H^{-1} is polynomial-time computable.

The proof of the lemma is postponed to Appendix B.

Proof of Theorem 18 For readability, we again set $Sys_{\text{async},1} := \varphi(Sys_{\text{sync},1})$ and $Sys_{\text{async},2} := \varphi(Sys_{\text{sync},2})$. Let now an arbitrary configuration $conf_{\text{async},1} = (\varphi(M_{\text{sync},1}) \cup \{X_{\text{sync},1,\kappa_1}\}, S, \varphi(H_{\text{sync}}), A_{\text{async},1}) \in \text{Conf}(Sys_{\text{async},1})$ be given.

- (1) We apply Theorem 15 on $conf_{\text{async},1}$ which yields a synchronous configuration $conf_{\text{sync},1} = (M_{\text{sync},1}, S, H_{\text{sync}}, A_{\text{sync},1}) \in \text{Conf}(Sys_{\text{sync},1})$ with

$$\sigma(\text{view}_{conf_{\text{async},1}}(\varphi(H_{\text{sync}}))) = \text{view}_{conf_{\text{sync},1}}(H_{\text{sync}}).$$

Moreover, if $conf_{\text{async},1}$ is polynomial-time then $conf_{\text{sync},1}$ is also polynomial-time, and the mapping σ is polynomial-time computable.

- (2) Now the precondition $Sys_{\text{sync},1} \geq_{\text{sync}} Sys_{\text{sync},2}$ can be applied yielding a configuration $conf_{\text{sync},2} = (M_{\text{sync},2}, S, H_{\text{sync}}, A_{\text{sync},2}) \in \text{Conf}(Sys_{\text{sync},2})$ with

$$\text{view}_{conf_{\text{sync},1}}(H_{\text{sync}}) \approx \text{view}_{conf_{\text{sync},2}}(H_{\text{sync}})$$

and $(M_{\text{sync},2}, S) \in f(M_{\text{sync},1}, S)$. Moreover, in the computational case, $conf_{\text{sync},2}$ is polynomial-time.

- (3) We now apply Theorem 14 to the configuration $conf_{\text{sync},2}$ which yields a configuration $conf_{\text{async},2} = (\varphi(M_{\text{sync},2}) \cup \{X_{\text{sync},2,\kappa_2}\}, S, \varphi(H_{\text{sync}}), \varphi(A_{\text{sync},2}))$ with

$$\text{view}_{conf_{\text{sync},2}}(H_{\text{sync}}) = \sigma(\text{view}_{conf_{\text{async},2}}(\varphi(H_{\text{sync}}))).$$

Moreover, $conf_{\text{async},2}$ is a polynomial configuration iff $conf_{\text{sync},2}$ is polynomial, according to Theorem 14.

- (4) Putting it all together, we have

- $\sigma(\text{view}_{conf_{\text{async},1}}(\varphi(H_{\text{sync}}))) = \text{view}_{conf_{\text{sync},1}}(H_{\text{sync}})$
- $\text{view}_{conf_{\text{sync},1}}(H_{\text{sync}}) \approx \text{view}_{conf_{\text{sync},2}}(H_{\text{sync}})$
- $\text{view}_{conf_{\text{sync},2}}(H_{\text{sync}}) = \sigma(\text{view}_{conf_{\text{async},2}}(\varphi(H_{\text{sync}})))$

Using Lemma 16, we obtain

$$\sigma(\text{view}_{conf_{\text{async},1}}(\varphi(H_{\text{sync}}))) \approx \sigma(\text{view}_{conf_{\text{async},2}}(\varphi(H_{\text{sync}}))).$$

We now finally apply our “reversing” function σ_H^{-1} (cf. Lemma 19) on the above equation. Together with Lemma 16, we obtain

$$\text{view}_{conf_{\text{async},1}}(\varphi(H_{\text{sync}})) \approx \text{view}_{conf_{\text{async},2}}(\varphi(H_{\text{sync}})),$$

i.e., $conf_{\text{async},2}$ is an indistinguishable configuration for $conf_{\text{async},1}$. Moreover, we have $(M_{\text{sync},2}, S) \in f(M_{\text{sync},1}, S)$, i.e., $\varphi(M_{\text{sync},2}, S) \in f'(\varphi(M_{\text{sync},1}, S))$, which yields the desired result $\varphi(Sys_{\text{sync},1}) \geq_{\text{async},H}^{f'} \varphi(Sys_{\text{async},2})$. \square

6 Deriving Synchronous Theorems from Asynchronous Ones

Recall that our long-term goal is to avoid proving each and every theorem and lemma for both models. We now briefly show how our two embedding theorems can be used for circumventing this problem. One of the most important theorems of both models is transitivity of the relation \geq .

Lemma 20 (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 as $f_3(M_1, S)$ being the union of the sets $f_2(M_2, S)$ with $(M_2, S) \in f_1(M_1, S)$.*

This has been proven in [8] for the synchronous and in [11] for the asynchronous model. We now exemplarily show how to derive the synchronous version from the asynchronous one using our previous results.

Lemma 21 *Assume that the asynchronous version of the transitivity lemma (Lemma 20) has already been proven, then the synchronous version holds as well.*

Proof We omit the superscripts f_i for the sake of readability. Let arbitrary synchronous systems Sys_1 , Sys_2 , and Sys_3 be given such that $Sys_1 \geq_{\text{sync}} Sys_2$ and $Sys_2 \geq_{\text{sync}} Sys_3$. We have to show that $Sys_1 \geq_{\text{sync}} Sys_3$ holds, provided that asynchronous transitivity has already been proven. According to our second embedding theorem, we know that

$$\varphi(Sys_1) \geq_{\text{async,H}} \varphi(Sys_2) \quad \text{and} \quad \varphi(Sys_2) \geq_{\text{async,H}} \varphi(Sys_3).$$

Obviously, the asynchronous version of transitivity is applicable to the relation $\geq_{\text{async,H}}$ instead of \geq_{async} as well, since it is a special case only, and the honest user remains unchanged at simulatability. Thus, we can apply our (already proven) asynchronous version of the transitivity lemma, which yields

$$\varphi(Sys_1) \geq_{\text{async,H}} \varphi(Sys_3).$$

Now, we use our first embedding theorem in conjunction with its subsequent remarks (stating that the theorem holds as well for the restricted version $\geq_{\text{async,H}}$ of simulatability) yielding $Sys_1 \geq_{\text{sync}} Sys_3$. \square

This proof technique is applicable to almost all theorems that rely on simulatability. As the most important example, we name the preservation theorem [50,55], which states that integrity properties expressed in linear-time logic are preserved under simulatability. The proof of this theorem is difficult and comprises several pages for both models. Using our work, the synchronous proof could as well be omitted.

Acknowledgements

This work benefited from fruitful discussions with *Dennis Hofheinz*, *Birgit Pfizmann*, and *Michael Waidner*.

References

- [1] R. Segala, N. Lynch, Probabilistic simulation for probabilistic processes, *Nordic Journal of Computing* 2 (2) (1995) 250–273.
- [2] S.-H. Wu, S. A. Smolka, E. W. Stark, Composition and behaviors of probabilistic I/O automata, *Theoretical Computer Science* 176 (1–2) (1997) 1–38.
- [3] R. Canetti, *Studies in secure multiparty computation and applications*, Department of Computer Science and Applied Mathematics, The Weizmann Institute of Science, revised March 1996 (Jun. 1995).
- [4] S. Goldwasser, L. Levin, Fair computation of general functions in presence of immoral majority, in: *Advances in Cryptology: CRYPTO '90*, Vol. 537 of *Lecture Notes in Computer Science*, Springer, 1990, pp. 77–93.
- [5] S. Micali, P. Rogaway, Secure computation, in: *Advances in Cryptology: CRYPTO '91*, Vol. 576 of *Lecture Notes in Computer Science*, Springer, 1991, pp. 392–404.
- [6] D. Beaver, Secure multiparty protocols and zero knowledge proof systems tolerating a faulty minority, *Journal of Cryptology* 4 (2) (1991) 75–122.
- [7] P. Lincoln, J. Mitchell, M. Mitchell, A. Scedrov, A probabilistic poly-time framework for protocol analysis, in: *Proc. 5th ACM Conference on Computer and Communications Security*, 1998, pp. 112–121.
- [8] B. Pfitzmann, M. Schunter, M. Waidner, Secure reactive systems, Research Report RZ 3206, IBM Research, http://www.semper.org/sirene/publ/PfSW1_00ReactSimulIBM.ps.gz (May 2000).
- [9] M. Hirt, U. Maurer, Player simulation and general adversary structures in perfect multiparty computation, *Journal of Cryptology* 13 (1) (2000) 31–60.
- [10] R. Canetti, Security and composition of multiparty cryptographic protocols, *Journal of Cryptology* 3 (1) (2000) 143–202.
- [11] B. Pfitzmann, M. Waidner, A model for asynchronous reactive systems and its application to secure message transmission, in: *Proc. 22nd IEEE Symposium on Security & Privacy*, 2001, pp. 184–200.
- [12] R. Canetti, Universally composable security: A new paradigm for cryptographic protocols, in: *Proc. 42nd IEEE Symposium on Foundations of Computer Science (FOCS)*, 2001, pp. 136–145, extended version in *Cryptology ePrint Archive*, Report 2000/67, <http://eprint.iacr.org/>.
- [13] M. Backes, B. Pfitzmann, M. Waidner, Secure asynchronous reactive systems, *IACR Cryptology ePrint Archive* 2004/082 (Mar. 2004).
- [14] M. Bellare, R. Canetti, H. Krawczyk, A modular approach to the design and analysis of authentication and key exchange protocols, in: *Proc. 30th Annual ACM Symposium on Theory of Computing (STOC)*, 1998, pp. 419–428.

- [15] C. Dwork, M. Naor, A. Sahai, Concurrent zero-knowledge, in: Proc. 30th Annual ACM Symposium on Theory of Computing (STOC), 1998, pp. 409–418.
- [16] B. Neuman, T. Ts'o, Kerberos: An authentication service for computer networks, *IEEE Communications Magazine* 32 (9) (1994) 33–38.
- [17] J. Mitchell, M. Mitchell, A. Scedrov, A linguistic characterization of bounded oracle computation and probabilistic polynomial time, in: Proc. 39th IEEE Symposium on Foundations of Computer Science (FOCS), 1998, pp. 725–733.
- [18] J. Mitchell, M. Mitchell, A. Scedrov, V. Teague, A probabilistic polynomial-time process calculus for analysis of cryptographic protocols (preliminary report), *Electronic Notes in Theoretical Computer Science* 47 (2001) 1–31.
- [19] R. Impagliazzo, B. M. Kapron, Logics for reasoning about cryptographic constructions, in: Proc. 44th IEEE Symposium on Foundations of Computer Science (FOCS), 2003, pp. 372–381.
- [20] D. Dolev, A. C. Yao, On the security of public key protocols, *IEEE Transactions on Information Theory* 29 (2) (1983) 198–208.
- [21] J. K. Millen, The interrogator: A tool for cryptographic protocol security, in: Proc. 5th IEEE Symposium on Security & Privacy, 1984, pp. 134–141.
- [22] C. Meadows, Using narrowing in the analysis of key management protocols, in: Proc. 10th IEEE Symposium on Security & Privacy, 1989, pp. 138–147.
- [23] R. Kemmerer, Analyzing encryption protocols using formal verification techniques, *IEEE Journal on Selected Areas in Communications* 7 (4) (1989) 448–457.
- [24] M. Burrows, M. Abadi, R. Needham, A logic for authentication, Technical Report 39, SRC DIGITAL (1990).
- [25] C. Meadows, Formal verification of cryptographic protocols: A survey, in: Proc. ASIACRYPT '94, Vol. 917 of Lecture Notes in Computer Science, Springer, 1994, pp. 135–150.
- [26] R. Kemmerer, C. Meadows, J. Millen, Three systems for cryptographic protocol analysis, *Journal of Cryptology* 7 (2) (1994) 79–130.
- [27] G. Lowe, Breaking and fixing the Needham-Schroeder public-key protocol using FDR, in: Proc. 2nd International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS), Vol. 1055 of Lecture Notes in Computer Science, Springer, 1996, pp. 147–166.
- [28] L. Paulson, The inductive approach to verifying cryptographic protocols, *Journal of Cryptology* 6 (1) (1998) 85–128.
- [29] F. J. Thayer Fabrega, J. C. Herzog, J. D. Guttman, Strand spaces: Why is a security protocol correct?, in: Proc. 19th IEEE Symposium on Security & Privacy, 1998, pp. 160–171.
- [30] M. Abadi, A. D. Gordon, A calculus for cryptographic protocols: The spi calculus, *Information and Computation* 148 (1) (1999) 1–70.

- [31] M. Abadi, P. Rogaway, Reconciling two views of cryptography: The computational soundness of formal encryption, in: Proc. 1st IFIP International Conference on Theoretical Computer Science, Vol. 1872 of Lecture Notes in Computer Science, Springer, 2000, pp. 3–22.
- [32] M. Abadi, J. Jürjens, Formal eavesdropping and its computational interpretation, in: Proc. 4th International Symposium on Theoretical Aspects of Computer Software (TACS), 2001, pp. 82–94.
- [33] P. Laud, Semantics and program analysis of computationally secure information flow, in: Proc. 10th European Symposium on Programming (ESOP), 2001, pp. 77–91.
- [34] M. Backes, B. Pfitzmann, M. Waidner, A composable cryptographic library with nested operations (extended abstract), in: Proc. 10th ACM Conference on Computer and Communications Security, 2003, pp. 220–230, full version in IACR Cryptology ePrint Archive 2003/015, Jan. 2003, <http://eprint.iacr.org/>.
- [35] M. Backes, 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 (FSTTCS), 2003, pp. 1–12, full version in IACR Cryptology ePrint Archive 2003/121, Jun. 2003, <http://eprint.iacr.org/>.
- [36] B. Warinschi, A computational analysis of the Needham-Schroeder-(Lowe) protocol, in: Proc. 16th IEEE Computer Security Foundations Workshop (CSFW), 2003, pp. 248–262.
- [37] M. Backes, B. Pfitzmann, M. Waidner, Symmetric authentication within a simulatable cryptographic library, in: Proc. 8th European Symposium on Research in Computer Security (ESORICS), Vol. 2808 of Lecture Notes in Computer Science, Springer, 2003, pp. 271–290, extended version in IACR Cryptology ePrint Archive 2003/145, Jul. 2003, <http://eprint.iacr.org/>.
- [38] M. Backes, B. Pfitzmann, Symmetric encryption in a simulatable Dolev-Yao style cryptographic library, in: Proc. 17th IEEE Computer Security Foundations Workshop (CSFW), 2004, full version in IACR Cryptology ePrint Archive 2004/059, Feb. 2004, <http://eprint.iacr.org/>.
- [39] P. Laud, Symmetric encryption in automatic analyses for confidentiality against active adversaries, manuscript, 2004.
- [40] J. Herzog, Computational soundness of formal adversaries, Ph.D. thesis, MIT (2002).
- [41] J. Herzog, M. Liskov, S. Micali, Plaintext awareness via key registration, in: Advances in Cryptology: CRYPTO 2003, Vol. 2729 of Lecture Notes in Computer Science, Springer, 2003, pp. 548–564.
- [42] D. Micciancio, B. Warinschi, Soundness of formal encryption in the presence of active adversaries, in: Proc. 1st Theory of Cryptography Conference (TCC), Vol. 2951 of Lecture Notes in Computer Science, Springer, 2004, pp. 133–151.

- [43] J. D. Guttman, F. J. Thayer Fabrega, L. Zuck, The faithfulness of abstract protocol analysis: Message authentication, in: Proc. 8th ACM Conference on Computer and Communications Security, 2001, pp. 186–195.
- [44] D. Volpano, G. Smith, Verifying secrets and relative secrecy, in: Proc. 27th Symposium on Principles of Programming Languages (POPL), 2000, pp. 268–276.
- [45] A. C. Yao, Protocols for secure computations, in: Proc. 23rd IEEE Symposium on Foundations of Computer Science (FOCS), 1982, pp. 160–164.
- [46] B. Pfitzmann, M. Schunter, M. Waidner, Cryptographic security of reactive systems, Presented at the *DERA/RHUL Workshop on Secure Architectures and Information Flow*, 1999, Electronic Notes in Theoretical Computer Science (ENTCS), March 2000, <http://www.elsevier.nl/cas/tree/store/tcs/free/noncas/pc/menu.htm>.
- [47] B. Pfitzmann, M. Schunter, M. Waidner, Provably secure certified mail, Research Report RZ 3207, IBM Research, <http://www.semper.org/sirene/publ/PfSW2CertMail.ps.gz> (Aug. 2000).
- [48] M. Steiner, Secure group key agreement, Ph.D. thesis, Universität des Saarlandes, http://www.semper.org/sirene/publ/Stein_02.thesis-final.pdf (2002).
- [49] R. Canetti, H. Krawczyk, Universally composable notions of key exchange and secure channels (extended abstract), in: Advances in Cryptology: EUROCRYPT 2002, Vol. 2332 of Lecture Notes in Computer Science, Springer, 2002, pp. 337–351, extended version in IACR Cryptology ePrint Archive 2002/059, <http://eprint.iacr.org/>.
- [50] B. Pfitzmann, M. Waidner, Composition and integrity preservation of secure reactive systems, in: Proc. 7th ACM Conference on Computer and Communications Security, 2000, pp. 245–254, extended version (with Matthias Schunter) IBM Research Report RZ 3206, May 2000, http://www.semper.org/sirene/publ/PfSW1_00ReactSimulIBM.ps.gz.
- [51] M. Backes, B. Pfitzmann, M. Waidner, A general composition theorem for secure reactive systems, in: Proc. 1st Theory of Cryptography Conference (TCC), Vol. 2951 of Lecture Notes in Computer Science, Springer, 2004, pp. 336–354.
- [52] R. Canetti, Y. Lindell, R. Ostrovsky, A. Sahai, Universally composable two-party and multi-party secure computation, in: Proc. 34th Annual ACM Symposium on Theory of Computing (STOC), 2002, pp. 494–503.
- [53] D. Hofheinz, J. Müller-Quade, Universally composable commitments using random oracles, in: Proc. 1st Theory of Cryptography Conference (TCC), Vol. 2951 of Lecture Notes in Computer Science, Springer, 2004, pp. 58–76.
- [54] M. Backes, C. Jacobi, B. Pfitzmann, Deriving cryptographically sound implementations using composition and formally verified bisimulation, in: Proc. 11th Symposium on Formal Methods Europe (FME 2002), Vol. 2391 of Lecture Notes in Computer Science, Springer, 2002, pp. 310–329.

- [55] M. Backes, C. Jacobi, Cryptographically sound and machine-assisted verification of security protocols, in: Proc. 20th Annual Symposium on Theoretical Aspects of Computer Science (STACS), Vol. 2607 of Lecture Notes in Computer Science, Springer, 2003, pp. 675–686.
- [56] S. Owre, N. Shankar, J. M. Rushby, PVS: A prototype verification system, in: Proc. 11th International Conference on Automated Deduction (CADE), Vol. 607 of Lecture Notes in Computer Science, Springer, 1992, pp. 748–752.
- [57] P. Lincoln, J. Mitchell, M. Mitchell, A. Scedrov, Probabilistic polynomial-time equivalence and security analysis, in: Proc. 8th Symposium on Formal Methods Europe (FME 1999), Vol. 1708 of Lecture Notes in Computer Science, Springer, 1999, pp. 776–793.
- [58] C. A. R. Hoare, Communicating Sequential Processes, International Series in Computer Science, Prentice Hall, Hemel Hempstead, 1985.
- [59] N. Lynch, Distributed Algorithms, Morgan Kaufmann Publishers, San Francisco, 1996.
- [60] R. Segala, N. Lynch, Probabilistic simulation for probabilistic processes, in: Proc. 5th International Conference on Concurrency Theory (CONCUR), Vol. 836 of Lecture Notes in Computer Science, Springer, 1994, pp. 481–497.
- [61] S. Goldwasser, S. Micali, C. Rackoff, The knowledge complexity of interactive proof systems, SIAM Journal on Computing 18 (1) (1989) 186–207.
- [62] J. Neveu, Mathematical Foundations of the Calculus of Probability, Holden-Day, 1965.
- [63] A. Z. Broder, D. Dolev, Flipping coins in many pockets (byzantine agreement on uniformly random values), in: Proc. 16th Annual ACM Symposium on Theory of Computing (STOC), 1984, pp. 157–170.
- [64] A. C. Yao, Theory and applications of trapdoor functions, in: Proc. 23rd IEEE Symposium on Foundations of Computer Science (FOCS), 1982, pp. 80–91.

A Postponed Definitions

The following definition for indistinguishability of random variables is essentially from [64].

Definition 22 (Indistinguishability) Two families $(\text{var}_k)_{k \in \mathbb{N}}$ and $(\text{var}'_k)_{k \in \mathbb{N}}$ of random variables (or probability distributions) on common domains D_k are

- a) *perfectly indistinguishable* (“=”) if for each k , the two distributions var_k and var'_k are identical.

- b) *statistically indistinguishable* (“ \approx_{SMALL} ”) for a suitable class *SMALL* of functions from \mathbb{N} to $\mathbb{R}_{\geq 0}$ if the distributions are discrete and their statistical distances

$$\Delta(\text{var}_k, \text{var}'_k) := \frac{1}{2} \sum_{d \in D_k} |P(\text{var}_k = d) - P(\text{var}'_k = d)| \in \text{SMALL}$$

(as a function of k). *SMALL* should be closed under affine addition, and with a function g also contain every function $g' \leq g$.

- c) *computationally indistinguishable* (“ \approx_{poly} ”) if for every algorithm *Dis* (the distinguisher) that is probabilistic polynomial-time in its first input,

$$|P(\text{Dis}(1^k, \text{var}_k) = 1) - P(\text{Dis}(1^k, \text{var}'_k) = 1)| \in \text{NEGL}.$$

Intuitively, given the security parameter and an element chosen according to either var_k or var'_k , *Dis* tries to guess which distribution the element came from. The class *NEGL* denotes the set of all negligible functions, i.e., $g: \mathbb{N} \rightarrow \mathbb{R}_{\geq 0} \in \text{NEGL}$ if for all positive polynomials Q , $\exists k_0 \forall k \geq k_0: g(k) \leq 1/Q(k)$.

We write \approx if we want to treat all three cases simultaneously.

For reasons of completeness, we now present the extended definition of simulatability, based on the three different kinds of indistinguishability. Definition 9 was simplified in the sense that only computational indistinguishability of views was covered, which represents the most common case when applying simulatability to cryptographic protocols.

Definition 23 (Simulatability, extended version with three variants) 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 $\text{conf}_1 = (M_1, S, H, A_1) \in \text{Conf}(Sys_1)$, there exists a configuration $\text{conf}_2 = (M_2, S, H, A_2) \in \text{Conf}(Sys_2)$ with $(M_2, S) \in f(M_1, S)$ (and the same H) such that

$$\text{view}_{\text{conf}_1}(\mathbf{H}) = \text{view}_{\text{conf}_2}(\mathbf{H}).$$

- 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 $\text{view}_{\text{conf}_1, l}(\mathbf{H}) \approx_{\text{SMALL}} \text{view}_{\text{conf}_2, l}(\mathbf{H})$ for all polynomials l , i.e., statistical indistinguishability of all families of l -step prefixes of the views.
- 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 $\text{Conf}_{\text{poly}}(Sys_1)$ and $\text{Conf}_{\text{poly}}(Sys_2)$ 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.

B Postponed Proofs

Proof of Theorem 15 We first reverse our function φ on the structure $(\varphi(M_{\text{sync}}) \cup \{X_{\text{sync},\kappa}\}, S)$ and on the user $\varphi(H_{\text{sync}})$ yielding the structure (M_{sync}, S) of $Sys_{\text{sync},2}$ and the original honest user H_{sync} . Note, that we cannot reverse the function φ on the new adversary A_{async} in the same way, because we did not demand it to have a similar internal structure, so we construct a new adversary A_{sync} for the synchronous configuration as follows. The ports of A_{sync} are given by

$$\{p \mid p^C \in (\text{ports}(M_{\text{sync}}) \cup \text{ports}(H_{\text{sync}})) \wedge p \notin (\text{ports}(M_{\text{sync}}) \cup \text{ports}(H_{\text{sync}}))\},$$

i.e., it connects to all remaining free ports of M_{sync} and H_{sync} . Internally, A_{sync} maintains an array $(\text{output_store}_{p!})_{p! \in \text{out}(\text{ports}(A_{\text{async}}))}$ of lists over Σ^* all initially empty.

A_{sync} has the adversary A_{async} as a blackbox submachine and its behavior is defined as follows. If A_{sync} is clocked in the synchronous system, it gets an input tuple $\mathcal{I} = (\mathcal{I}_{p?})_{p? \in \text{in}(\text{ports}(A_{\text{sync}}))}$. It now tries to restore the order in which these messages would have arrived in the asynchronous system. More precisely, it knows the clocking scheme κ , so it know which machines have been clocking after the last clocking of A_{sync} . Moreover, it knows the order in which machines are switched by $X_{\text{sync},\kappa}$ in one particular subround. Using the order on the ports of the asynchronous machines, it can finally decide in which order messages sent by one machine on different ports would have arrived in the asynchronous system. The only problem which might arise is that a machine has been clocked more then once since the last clocking of the adversary. This might result in two inputs at the same port of A_{sync} which would be concatenated without any separation symbol. Such an input would not be restorable into its original form, so we had to include the restriction to the considered clocking scheme that every machine and the user are at most clocked once between two successive clockings of the adversary. Note, that our usually used clocking scheme $(M \cup \{H\}, \{A\}, \{H\}, \{A\})$ fulfills this requirement.

After restoring both the usual messages and their order, A_{sync} uses the blackbox function $\delta_{A_{\text{async}}}$ on the first input yielding an output tuple \mathcal{O} . This tuple \mathcal{O} is appended to the array output_store , i.e. each component $\mathcal{O}_{p!}$ is appended to $\text{output_store}_{p!}$. If there is a nonempty output c at a clock-out port $p^!$, we would have a clocked self-loop in $\text{conf}_{A_{\text{async}}}$ if $\text{output_store}_{p!}[c] \neq \epsilon$. In this case, this component is removed from the array and $\delta_{A_{\text{async}}}$ is called again with the new state and $\mathcal{I} := \mathcal{I}_{p? = \text{output_store}_{p!}[c]}$ and so on.

The above steps are repeated with the second input and the new state of A_{sync} and so on until all inputs have been considered. Finally, the blackbox function is used with $\mathcal{I}_{p_{A_{\text{sync}}}? = (i,j)}$ where i denotes the global round and j denotes the subround the adversary is clocked in. (The adversary obviously knows both i and j because he knows the clocking scheme κ , so he may simply maintain two counters that he

adapts every time he is clocked.) This correspond to the clocking signal of $X_{\text{sync},\kappa}$ in the asynchronous system. The output tuple is again concatenated to the same array and possible clocked self loops are considered again. Finally, A_{sync} outputs the first elements of each list of $output_store_{p_i}$ with $p_i^C \in \text{ports}(M_{\text{sync}} \cup \{H_{\text{sync}}\})$ as its output tuple \mathcal{O} and removes these elements from the lists.

Note, that this newly defined adversary A_{sync} is polynomial iff A_{async} is polynomial by construction. Thus, if the original configuration $conf_{\text{async}}$ has been polynomial-time (i.e., the user $\varphi(H_{\text{sync}})$ and the adversary A_{async} must be polynomial-time) then the configuration $conf_{\text{sync}} = (M_{\text{sync}}, S, H_{\text{sync}}, A_{\text{sync}})$ will also be polynomial-time, since the runtime of H_{sync} is always bounded by $\varphi(H_{\text{sync}})$.

A_{sync} “reverse” the function φ by construction. The asynchronous adversary would receive many single inputs, and it would produce outputs every time which would be stored in the outgoing buffers. Possible clocked self-loops are handled by repeated calls of the transition function with correct inputs. If A_{async} is scheduled by $X_{\text{sync},\kappa}$ it again performs an arbitrary transition and the first element of its outgoing buffer would be clocked. The synchronous adversary first splits its input messages into their original order and uses the blackbox function one by one storing the outputs in $output_store$. The split inputs correspond to the original inputs of the asynchronous system, so the output tuples are also equal after every step. Therefore, the contents of $output_store$ always correspond to the outgoing buffers in the asynchronous system after a clocking step of A_{async} . If the synchronous adversary is clocked it again calls its blackbox function with the correct input and stores the output in the array. After that, it outputs the first element of each list of the array and removes these elements from the lists. In the asynchronous system messages stored in the outgoing buffers are treated in the same way. More formally we can show the following lemma.

Lemma 24 *We denote this “reversion” of φ_M by $\bar{\varphi}_M$ and the reversion of the whole configuration by $\bar{\varphi}_{conf}$ for the moment. Then for an arbitrary configuration $conf_{\text{async}} = (\varphi(M_{\text{sync}}) \cup \{X_{\text{sync},\kappa}\}, S, \varphi(H_{\text{sync}}), A_{\text{async}})$ we have*

$$view_{\varphi_{conf}(\bar{\varphi}_{conf}(conf_{\text{async}}))}(\varphi(M)) = view_{conf_{\text{async}}}(\varphi(M))$$

for every $M \in (M_{\text{sync}} \cup \{H_{\text{sync}}\})$ and

$$view_{\varphi_{conf}(\bar{\varphi}_{conf}(conf_{\text{async}}))}(A_{\text{async}}) = view_{conf_{\text{async}}}(A_{\text{async}})$$

where the view of A_{async} in the first configuration is given as a submachine of $\varphi_M(\bar{\varphi}_M(A_{\text{async}}))$.

Proof of Lemma 24 The proof is illustrated in Figure B.1. We first show that $A'_{\text{async}} := \varphi_M(\bar{\varphi}_M(A_{\text{async}}))$ behaves exactly as A_{async} , i.e., both machines are perfectly indistinguishable for their environment. This is already sufficient to show




Fig. B.1. Overview of the proof of Lemma 24.

that the views of $\varphi(M)$ for every $M \in (M_{\text{sync}} \cup \{H_{\text{sync}}\})$ are equal in both configurations because they remain unchanged. We will also show that the view of A_{async} is equal in both configurations which finishes our proof.

We show that both adversaries A'_{async} and A_{async} behave identically between two successive clockings. Moreover, we show that the content of array $output_store_{p!}$ of A'_{async} always equal the outgoing buffers p in the corresponding asynchronous configuration at every clocking of A_{async} as a submachine of A'_{async} if we identify clockings of A_{async} in both configurations in the natural way. More precisely, this means that we identify the i -th clocking of A_{async} in $conf_{\text{async}}$ with the i -th call of $\delta_{A_{\text{async}}}$ by A'_{async} in $\varphi_{conf}(\bar{\varphi}_{conf}(conf_{\text{async}}))$. Furthermore, we show that outputs made by the adversary are always equal in both configurations.

At the start of the run both buffers and arrays are empty which fulfills our claim. Now assume that A'_{async} receives an arbitrary input at $p? \neq p_{A_{\text{async}}}$. It stores the message in its array $input_store_{p?}$ and gives the control to the master scheduler. If A'_{async} receives a non-empty input at $p_A?$ it applies the state transition function $\delta_{\bar{\varphi}_M(A_{\text{async}})}$ on the arrays $input_store$. Now, the arrays $input_store$ are decomposed into single inputs again preserving their original order, and the function $\delta_{A_{\text{async}}}$ is applied to every such input. Since the inputs are obviously equal in both configuration, we obtain identical outputs, and moreover identical views for A_{async} . By precondition, the arrays $output_store$ are mapped to the outgoing buffers. After one call of $\delta_{A_{\text{async}}}$, every output at $p!$ is stored either in $output_store_{p!}$ or in p at the same position, so they remain validly mapped. Now, either the first component of $output_store_{p!}$ or the first entry of p for $p!^C \in (\text{ports}(M_{\text{sync}}) \cup \{H_{\text{sync}}\})$ are output yielding identical outputs and therefore identical views for the environment in both configurations, i.e.,

$$view_{\varphi_{conf}(\bar{\varphi}_{conf}(conf_{\text{async}}))}(\varphi(M)) = view_{conf_{\text{async}}}(\varphi(M))$$

for $M \in (M_{\text{sync}} \cup \{H_{\text{sync}}\})$. We already showed that the views of A_{async} are equal in both configurations which finishes our proof. \square

According to Lemma 24, the function $\varphi_{conf} \circ \bar{\varphi}_{conf}$ yields identical views for $\varphi(\mathbf{M})$ for every $\mathbf{M} \in (M_{\text{sync}} \cup \{H_{\text{sync}}\})$ and the asynchronous adversary, i.e.,

- $view_{\varphi_{conf}(\bar{\varphi}_{conf}(conf_{\text{async}}))}(\varphi(\mathbf{M})) = view_{conf_{\text{async}}}(\varphi(\mathbf{M}))$ and
- $view_{\varphi_{conf}(\bar{\varphi}_{conf}(conf_{\text{async}}))}(\mathbf{A}_{\text{async}}) = view_{conf_{\text{async}}}(\mathbf{A}_{\text{async}})$.

We already showed in Theorem 14 that $view_{conf_{\text{sync}}}(\mathbf{M}) = \sigma(view_{\varphi(conf_{\text{sync}})}(\varphi(\mathbf{M})))$ holds for every synchronous configuration $conf_{\text{sync}} = (M_{\text{sync}}, S, H_{\text{sync}}, \mathbf{A}_{\text{sync}})$ and for every machine $\mathbf{M} \in (M_{\text{sync}} \cup \{H_{\text{sync}}, \mathbf{A}_{\text{sync}}\})$. If we now set $conf_{\text{sync}} := \bar{\varphi}_{conf}(conf_{\text{async}})$, we obtain

- $view_{conf_{\text{sync}}}(\mathbf{M}) = \sigma(view_{\varphi_{conf}(\bar{\varphi}_{conf}(conf_{\text{async}}))}(\varphi(\mathbf{M})))$

Moreover, this implies

- $view_{conf_{\text{sync}}}(\mathbf{A}_{\text{sync}}) = \sigma(view_{\varphi_{conf}(\bar{\varphi}_{conf}(conf_{\text{async}}))}(\mathbf{A}_{\text{async}}))$

since the views of $\mathbf{A}_{\text{async}}$ and $\varphi(\bar{\varphi}(\mathbf{A}_{\text{async}}))$ are identical. We apply the mapping σ on the first two equations and, using Lemma 16, we obtain

- $\sigma(view_{\varphi_{conf}(\bar{\varphi}_{conf}(conf_{\text{async}}))}(\varphi(\mathbf{M}))) = \sigma(view_{conf_{\text{async}}}(\varphi(\mathbf{M})))$ and
- $\sigma(view_{\varphi_{conf}(\bar{\varphi}_{conf}(conf_{\text{async}}))}(\mathbf{A}_{\text{async}})) = \sigma(view_{conf_{\text{async}}}(\mathbf{A}_{\text{async}}))$

Note, that σ is in fact defined on runs of these configuration because both the machines of the structure and the honest user have the prescribed form. Using transitivity, we immediately obtain the desired result

$$view_{conf_{\text{sync}}}(\mathbf{M}) = \sigma(view_{conf_{\text{async}}}(\varphi(\mathbf{M})))$$

and

$$view_{conf_{\text{sync}}}(\mathbf{A}_{\text{sync}}) = \sigma(view_{conf_{\text{async}}}(\mathbf{A}_{\text{async}}))$$

As a special case we set $\mathbf{M} := H_{\text{sync}}$ which yields

$$view_{conf_{\text{sync}}}(H_{\text{sync}}) = \sigma(view_{conf_{\text{async}}}(\varphi(H_{\text{sync}}))). \quad \square$$

Proof of Lemma 19 In order to prove the claim, we present an algorithm which undoes the changes of the algorithm for deriving the mapping σ : It has an internal list over Σ^+ initially empty, which will be used to construct the desired view. For every subround j , it goes through all tuples $(M_{\text{sync}}, i, j, s, \mathcal{I}, s', \mathcal{O}')$ modifying them as follows: If $M_{\text{sync}} = H_{\text{sync}}$ for one machine of this subround, it appends $(\varphi(H_{\text{sync}}), s, \mathcal{I}_{\mathcal{P}_{H_{\text{sync}}}}(i,j), s', \mathcal{O}')$ to its internal list. Note that this tuple precisely matches the original asynchronous tuple for switching the honest user $\varphi(H_{\text{sync}})$ by the master scheduler. After that, it proceed through all tuples of this subround in precisely the same order they have been scheduled by the master scheduler (the algorithm is surely allowed to know the clocking scheme). For a given tuple of the form $(M_{\text{sync}}, i, j, s, \mathcal{I}, s', \mathcal{O}')$, it checks, whether there is a non-empty output at

a port $p!$ in \mathcal{O}' with $p? \in \text{ports}(\varphi(H_{\text{sync}}))$. In this case, the honest user would be clocked in the second asynchronous block, so we use the state transition function $\delta_{\varphi(H_{\text{sync}})}$ on the current state s of $\varphi(H_{\text{sync}})$ and input $\mathcal{I}_{p?=O'_{p!}}$ which yields a new state s' and an (all-empty) output \mathcal{O}_ϵ . We then add a step $(\varphi(H_{\text{sync}}), s, \mathcal{I}_{p?=O'_{p!}}, s', \mathcal{O}_\epsilon)$. This is done for all ports of M_{sync} according to their order and for all machines that switch in the consider subround. Obviously, this algorithm reverses the mapping σ for the honest user by construction. In case of a polynomial configuration, especially the adversary has to be polynomial-time. This implies that there cannot be any infinite successive clocked self-loops. Moreover, both the adversary and the honest user will reach final state after a polynomial number of blocks, so the algorithm for σ_H^{-1} applied to the view of the honest user will only makes a polynomial number of transition, each one with a polynomial number of steps. This implies that σ is computable polynomial-time applied to the view of the honest user if it is used in a polynomial-time configuration. \square