

The IITM Model: a Simple and Expressive Model for Universal Composability

Ralf Küsters Max Tuengerthal

University of Trier, Germany

`{kuesters,tuengerthal}@uni-trier.de`

January 18, 2013

Abstract

The universal composability paradigm allows for the modular design and analysis of cryptographic protocols. It has been widely and successfully used in cryptography. However, devising a coherent yet simple and expressive model for universal composability is, as the history of such models shows, highly non-trivial. For example, several partly severe problems have been pointed out in the literature for the UC model.

In this work, we propose a coherent model for universal composability, called the IITM model (“Inexhaustible Interactive Turing Machine”). A main feature of the model is that it is stated without a priori fixing irrelevant details, such as a specific way of addressing of machines by session and party identifiers, a specific modeling of corruption, or a specific protocol hierarchy. In addition, we employ a very general notion of runtime. All reasonable protocols and ideal functionalities should be expressible based on this notion in a direct and natural way, and without tweaks, such as (artificial) padding of messages or (artificially) adding extra messages.

Not least because of these features, the model is simple and expressive. Also the general results that we prove, such as composition theorems, hold independently of how such details are fixed for concrete applications.

Being inspired by other models for universal composability, in particular the UC model and because of the flexibility and expressivity of the IITM model, conceptually, results formulated in these models directly carry over to the IITM model.

Contents

| | | |
|----------|---|-----------|
| 1 | Introduction | 4 |
| 2 | The IITM Model in a Nutshell | 6 |
| 3 | The General Computational Model | 10 |
| 3.1 | Inexhaustible Interactive Turing Machines | 10 |
| 3.1.1 | Syntax | 10 |
| 3.1.2 | Computation | 10 |
| 3.2 | Systems of IITMs | 11 |
| 3.3 | Running a System | 11 |
| 3.4 | Probability Space and Relevant Random Variables | 14 |
| 3.5 | Equivalence/Indistinguishability of Systems | 15 |
| 4 | Polynomial Time and Properties of Systems | 15 |
| 4.1 | Further Notation and Terminology | 15 |
| 4.2 | Polynomially Bounded Systems | 16 |
| 4.3 | Environments and Environmental Indistinguishability | 17 |
| 4.4 | Protocols and Environmentally Bounded Systems | 18 |
| 4.5 | Properties of Systems | 18 |
| 5 | Composition Theorems for Environmental Indistinguishability | 23 |
| 5.1 | Composition Theorem for a Constant Number of Systems | 23 |
| 5.2 | Composition Theorem for Unbounded Self-Composition | 25 |
| 5.2.1 | Session Versions of Systems | 25 |
| 5.2.2 | The Composition Theorem for Session Versions | 26 |
| 5.3 | Composition Theorem for Unbounded Self-Composition of SID Dependent Systems | 28 |
| 5.3.1 | Generalized Session Versions | 29 |
| 5.3.2 | A Composition Theorem for σ -Session Versions | 30 |
| 6 | Notions of Universal Composability | 32 |
| 6.1 | Further Notation and Terminology | 32 |
| 6.1.1 | Network and I/O Tapes | 32 |
| 6.1.2 | Adversarial Systems | 33 |
| 6.2 | Defining the Notions of Universal Composability | 33 |
| 6.3 | Relationships Between the Notions of Universal Composability | 35 |
| 6.4 | Reflexivity and Transitivity | 37 |
| 7 | Composition Theorems for the Realization Relations | 37 |
| 7.1 | Composition Theorem for a Constant Number of Protocol Systems | 38 |
| 7.2 | Composition Theorem for Unbounded Self-Composition | 38 |
| 7.3 | Composition Theorem for Unbounded Self-Composition of SID Dependent Protocols | 39 |
| 7.4 | Composition Theorem for More Complex Systems | 41 |
| 8 | On the Composability of Runtime Notions | 42 |
| 8.1 | On the Composability of Environmentally Strictly Bounded Systems | 42 |
| 8.2 | On the Composability of Environmentally Almost Bounded Systems | 44 |

| | | |
|-----------|--|-----------|
| 9 | On Basing Universal Composability on Environmentally Strictly Bounded Systems | 46 |
| 9.1 | Strict Simulatability | 47 |
| 9.2 | No Universal Composability for a Constant Number of Protocol Systems | 47 |
| 9.3 | No Transitivity | 49 |
| 9.4 | Strict SS does not Imply Strict UC (Incompleteness of the Dummy Adversary) | 50 |
| 10 | Instantiation of the IITM Model | 50 |
| 10.1 | Modeling Protocols | 51 |
| 10.2 | A Concrete Example | 57 |
| 11 | Related Work | 65 |
| 11.1 | UC Model | 65 |
| 11.2 | GNUC model | 67 |
| A | Proof of Lemma 9 | 69 |
| B | Proof of Lemma 27 for Uniform Environments | 75 |
| C | Proof of Theorem 7 | 77 |
| C.1 | UC \Rightarrow dummyUC | 77 |
| C.2 | dummyUC \Rightarrow SS | 77 |
| C.3 | SS \Rightarrow BB | 78 |
| C.4 | BB \Rightarrow SS | 78 |
| C.5 | SS \Leftrightarrow RS | 78 |

1 Introduction

In the universal composability paradigm [2, 27] the security of protocols is defined in such a way that security is preserved even if the protocols are used as components of an arbitrary (polynomially bounded) distributed system. This strong composability property allows for the modular design and analysis of protocols. More specifically, the security of a protocol is defined in terms of an ideal protocol (also called an ideal functionality). A real protocol securely realizes the ideal protocol if every attack on the real protocol can be translated into an “equivalent” attack on the ideal protocol, where equivalence is specified based on an environment trying to distinguish the real attack from the ideal one. That is, for every real adversary on the real protocol there must exist an ideal adversary (also called a simulator) on the ideal protocol such that no environment can distinguish whether it interacts with the real protocol and the real adversary or the ideal protocol and the ideal adversary. So the real protocol is as secure as the ideal protocol (which, by definition, is secure) in all environments.

At the core of the universal composability paradigm are composition theorems which say that if a protocol uses one or more (independent) instances of an ideal functionality, then all instances of the ideal functionality can be replaced by instances of the real protocol that realizes the ideal functionality. In this way, more and more complex protocols can be designed and analyzed in a modular way based on ideal functionalities. Often different protocol instances share some state, such as long-term keys. In this case, composition theorems with joint state (joint state theorems) can be employed [10, 20]. Moreover, in some cases it is necessary that all components of a system have access to some global state information, such as a common reference string (CRS). For this purpose, composition theorems with global state have been proposed [6].

The universal composition paradigm has been widely and successfully used in cryptography to design and analyze complex protocols in a modular way (see, e.g., [4] for an overview).

However, devising a coherent yet simple and expressive model for this paradigm has turned out to be highly non-trivial (see Section 11 for a more detailed discussion). For example, several partly severe problems have been pointed out in the literature for the UC model, concerning, for instance, the validity of the composition and joint state theorems, the way corruption is handled, the notion of runtime, and the expressivity of the model.

Contribution. In this paper, we propose a coherent model for universal composability that is both simple and expressive. Our model coincides with the IITM model proposed in previous work [18] (where IITM stands for “Inexhaustible Interactive Turing Machine”), except that we now use a more general notion of runtime than the one used in the original version of the IITM model, namely one that is based on a runtime notion proposed by Hofheinz et al. in [16]. We therefore stick to the name IITM model also for the model proposed here. The main features of the IITM model are as follows.

1. The IITM model provides a very flexible and generic mechanism to address copies of machines. Unlike other models, in which a specific way of addressing of machines by session identifiers (SIDs) and party identifiers (PIDs) is fixed, the IITM model does not hard-wire how machines are addressed.
2. Unlike other models, the IITM model also does not hard-wire corruption into the model. Corruption instead can be specified in a very general and flexible way as part of the protocol specification.
3. Unlike other models, the IITM model does not impose a specific structure on protocols, such as a hierarchical structure with protocols and subroutines. This is, for example, important for seamlessly dealing with joint and global state.
4. The runtime of machines and systems of machines is defined in a very general way. All reasonable (real and ideal) protocols should be expressible in a very natural way based on this runtime notion, without tweaks, such as (artificial) padding of messages or (artificially) adding extra messages, as necessary in other models.

5. All common notions of universal composability, including (dummy) UC, strong simulatability, black-box simulatability, and reactive simulatability are equivalent.¹
6. Compared to other models, the IITM model is very simple, not least because it is formulated in a very general way without fixing irrelevant details.
7. For the same reason, the IITM model is very expressive. One can, for example, easily specify different ways of addressing machines and corrupting machines/protocol instances/parties. Also, protocols can be structured in a very flexible way. In particular, it is straightforward to mimic the way these things are fixed in other models.
8. The composition, joint state, and global setup theorems are very general: The class of protocols they talk about is very large, both in terms of the structure of protocols and in terms of their runtime. Also, these theorems are stated and proven independently of many details fixed in other models (such as addressing of machines and corruption). Hence, they hold true no matter how these details are fixed in concrete applications. The generality of the theorems and the IITM model is also apparent in the fact that to state the joint state and global setup theorems the IITM does not have to be changed or extended, unlike other models. These theorems can smoothly be stated and proven within the model. Moreover, the joint state theorem is even merely a direct consequence of our composition theorems. In other models, new notation is required to state the theorem and the theorem requires a (non-trivial) proof.
9. Since the IITM model conceptually follows other models for universal composability, in particular the UC model, and because of its high expressivity and flexibility, results established in other models easily carry over to the IITM model.

Structure of the paper. In Section 2, we present the IITM model in a nutshell in order to provide a first impression of the model. We then, in Section 3, define the computational model on which the IITM model is based. This model is stated independently of the application to universal composability and is of independent interest. The runtime notions that we use are introduced in Section 4, along with basic properties. While these properties are as expected, some of the problems in other models stem from the fact that not all of these properties are satisfied in these models. For example, not in all models it is possible to specify a dummy machine which simply forwards messages back and forth between different system components. Also, in some models it is not possible to simulate arbitrary subsystems by one machine. In Section 5, we present general composition theorems for systems that are computationally indistinguishable by environments. These composition theorems are the core of the composition theorems for universal composition, presented in Section 7. The common notions for universal composability, namely (dummy) UC, strong simulatability, black-box simulatability, and reactive simulatability are introduced in Section 6, along with basic properties. It is shown that all these notions are equivalent, where, as already mentioned, for reactive simulatability this requires environments with external input. The general composition theorems are then, as mentioned, stated and proven in Section 7. They follow easily from those established in Section 5. The composability of the runtime notions that we use are discussed in Section 8. An alternative runtime notion and why it is unsuitable is explained in Section 9. In Section 10, we illustrate how the IITM model can be used to design and analyze (multi-party) protocols in a modular way. In particular, we present one way of addressing multiple sessions and modeling corruption of machines/instances/parties. Related work is discussed in Section 11. Details omitted in the main body of the paper can be found in the appendix.

Acknowledgments. We would like to thank Ran Canetti, Dennis Hofheinz, Dominique Unruh, Anupam Datta, John Mitchell, Olivier Pereira, Birgit Pfitzmann, and Michael Backes for many interesting discussions on models for universal composability.

¹For reactive simulatability this is true for environments with external input (non-uniform environments).

2 The IITM Model in a Nutshell

In this section, we provide a brief introduction to the IITM model, with full details presented in the subsequent sections. In the IITM model, security notions and composition/joint state theorems are formalized based on a very simple and at the same time very expressive general computational model, in which IITMs (inexhaustible interactive Turing machines) and systems of IITMs are defined. We first sketch the general computational model and then, based on it, formulate notions of universal composability and state composition theorems.

The General Computational Model. The general computational model is defined in terms of systems of IITMs. An *inexhaustible interactive Turing machine (IITM)* is a probabilistic Turing machine with named input and output tapes as well as an associated polynomial. The tape names determine how different machines are connected in a system of IITMs (see below). An IITM runs in one of two modes, **CheckAddress** and **Compute**. The **CheckAddress** mode is used as a generic mechanism for addressing copies of IITMs in a system of IITMs, as explained below. In this mode, an IITM may perform, in every activation, a deterministic polynomial time computation in the length of the security parameter plus the length of the current input plus the length of its current configuration, where the polynomial is the one associated with the IITM. The IITM is supposed to output “accept” or “reject” at the end of the computation in this mode, indicating whether the received message is processed further or ignored. The actual processing of the message, if accepted, is done in mode **Compute**. In mode **Compute**, a machine may only output at most one message on an output tape (and hence, only at most one other machine is triggered). The runtime in this mode is not a priori bounded. Later the runtime of systems and their subsystems will be defined in such a way that the overall runtime of a system of IITMs is polynomially bounded in the security parameter plus the length of the external input. We note that in both modes, an IITM cannot be exhausted (hence, the name): in every activation it can perform actions and cannot be forced to stop. This property, while not satisfied in all other models, is crucial to obtain a reasonable model for universal composability (see also Section 11).

A *system* \mathcal{S} of IITMs is of the form $\mathcal{S} = M_1 \mid \cdots \mid M_k \mid !M'_1 \mid \cdots \mid !M'_{k'}$, where M_i , $i \in \{1, \dots, k\}$, and M'_j , $j \in \{1, \dots, k'\}$, are IITMs such that, for every tape name c , at most two of these IITMs have a tape named c and if two IITMs have a tape named c , then c is an input tape in one of the machines and an output tape in the other. We say that the IITMs M'_j are in the scope of a bang operator. This operator indicates that in a run of a system an unbounded number of (fresh) copies of a machine can be generated. Conversely, machines which are not in the scope of a bang operator may not be copied. Systems in which multiple copies of a machine may be generated are often needed, e.g., in case of multi-party protocols or in case a system describes the concurrent execution of multiple instances of a protocol.

In a run of a system \mathcal{S} at any time only one IITM is active and all other IITMs wait for new input; the first IITM to be activated in a run of \mathcal{S} is the so-called master IITM, of which a system has at most one and which may get external input. By the definition of IITMs, the active machine may output only at most one message on one of its output tapes, and hence, at most one other machine is triggered after the activation of the currently active machine. To illustrate runs of systems, consider, for example, the system $\mathcal{S} = M_1 \mid !M_2$ and assume that M_1 has an output tape named c , M_2 has an input tape named c , and M_1 is the master IITM. (There may be other tapes connecting M_1 and M_2 .) Furthermore, assume that in the run of \mathcal{S} executed so far, two copies of M_2 , say M'_2 and M''_2 , have been generated, with M'_2 generated before M''_2 , and that M_1 just sent a message m on tape c . This message is delivered to M'_2 (as the first copy of M_2). First, M'_2 runs in mode **CheckAddress** with input m ; as mentioned, this is a deterministic polynomial time computation which outputs “accept” or “reject”. If M'_2 accepts m , then M'_2 gets to process m in mode **Compute** and could, for example, send a message back to M_1 . Otherwise, m is given to M''_2 which then runs in mode **CheckAddress** with input m . If M''_2 accepts m , then M''_2 gets to process m in mode **Compute**. Otherwise (if both M'_2 and M''_2 do not accept m), a new copy M'''_2 of M_2 with fresh randomness is generated and M'''_2 runs in mode **CheckAddress** with input m . If M'''_2 accepts m , then M'''_2 gets to process m . Otherwise, M'''_2 is removed again, the message m is dropped, and the master IITM is activated, in this case M_1 , and so on. The master IITM is also activated if the currently active IITM does not produce output, i.e., stops in this activation without writing to any output tape. A run stops if the master IITM, after being activated, does not produce output

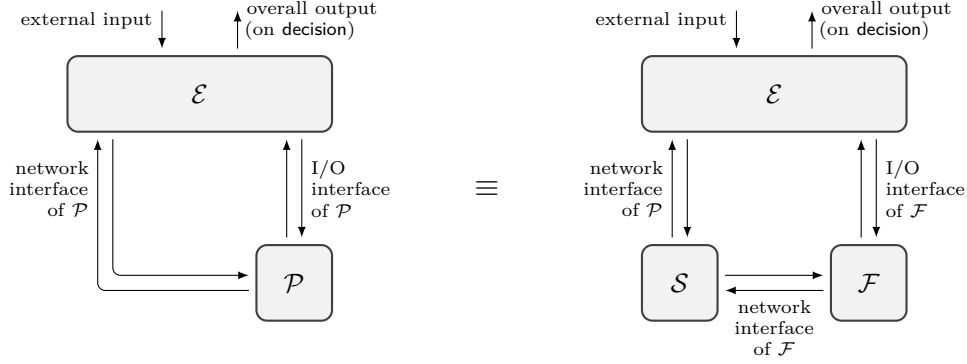


Figure 1: Strong simulatability (SS). We note that \mathcal{P} and \mathcal{F} have the same I/O interface.

(and hence, does not trigger another machine) or an IITM outputs a message on a tape named *decision*. Such a message is considered to be the *(overall) output of the system*.

Two systems \mathcal{P} and \mathcal{Q} are called *indistinguishable* ($\mathcal{P} \equiv \mathcal{Q}$) if and only if the difference between the probability that the output of \mathcal{P} is 1 and the probability that the output of \mathcal{Q} is 1 is negligible.

Note that we do not fix details such as addressing of machines by party/session IDs or corruption in this model. We also do not impose any specific structure, e.g., a hierarchical structure with protocols and subroutines, on systems. This makes the model both simpler and more expressive.

Notions of Universal Composability. We need the following terminology. For a system \mathcal{S} , the input/output tapes of IITMs in \mathcal{S} that do not have a matching output/input tape in \mathcal{S} are called *external*, where an input/output tape of one IITM in \mathcal{S} matches an output/input tape of another IITM in \mathcal{S} if both tapes have the same name. External tapes are grouped into *I/O* and *network tapes*. We often refer to the sets of I/O and network tapes of \mathcal{S} by *I/O* and *network interface*, respectively. We consider three different types of systems: protocol systems, adversarial systems, and environmental systems, modeling (i) real and ideal protocols/functionalities, (ii) adversaries and simulators, and (iii) environments, respectively. *Protocol systems*, *adversarial systems*, and *environmental systems* are systems which have an I/O and network interface, i.e., they may have I/O and network tapes. Adversarial systems may only connect to the network interface of a protocol system but not to its I/O interface. Environmental systems may contain a master machine and may produce output on the tape *decision*.

So far, we have not restricted the runtime of IITMs in mode *Compute* in any way. The following constraints will be used to enforce that systems run in polynomial time (except with negligible probability): (i) Every environmental system \mathcal{E} has to be *universally bounded*, i.e., there exists a polynomial p such that for every system \mathcal{S} that only connects to the external tapes of \mathcal{E} , security parameter η , and external input a , the overall runtime of \mathcal{E} in mode *Compute* is bounded by $p(\eta + |a|)$ in every run of $\mathcal{E} | \mathcal{S}$ with security parameter η and external input a . (ii) Every protocol system \mathcal{P} has to be *environmentally bounded*, i.e., for every environmental system \mathcal{E} there exists a polynomial p such that for every η and a the overall runtime of \mathcal{P} in mode *Compute* is bounded by $p(\eta + |a|)$ in every run of $\mathcal{E} | \mathcal{P}$ with security parameter η and external input a (except for a negligible set of runs). Since the runtime in mode *CheckAddress* is polynomially bounded, this guarantees that, for every protocol system \mathcal{P} and environmental system \mathcal{E} , the overall runtime of $\mathcal{E} | \mathcal{P}$ is polynomially bounded in the security parameter plus the length of the external input (except with negligible probability). This runtime notion for protocol systems is very general. We claim that it includes all reasonable protocol systems that occur in applications, as further explained in [16] and subsequent sections. In most other models, such as Canetti's UC model, the runtime notions are more restricted and more complex.

We now informally define strong simulatability; other equivalent security notions, such as universal composability (UC) and dummy UC, are defined in a similar way in subsequent sections. The systems considered in the following definition are illustrated in Figure 1.

Definition 1 (informal). Let \mathcal{P} and \mathcal{F} be protocol systems with the same I/O interface (i.e., with the same set of I/O tapes), the real and the ideal protocol, respectively. Then, \mathcal{P} realizes \mathcal{F} ($\mathcal{P} \leq^{SS} \mathcal{F}$) if and only if there exists an adversarial system \mathcal{S} (called a simulator or an ideal adversary) such that \mathcal{S} connects only to the network interface of \mathcal{F} , the systems \mathcal{P} and $\mathcal{S}|\mathcal{F}$ have the same external interface, $\mathcal{S}|\mathcal{F}$ is environmentally bounded, and for all environmental systems \mathcal{E} , connecting only to the external interface of \mathcal{P} (and hence, $\mathcal{S}|\mathcal{F}$), it holds that $\mathcal{E}|\mathcal{P} \equiv \mathcal{E}|\mathcal{S}|\mathcal{F}$.

We note that this relation is reflexive and transitive. We also emphasize that still details such as addressing of machines by party/session IDs, corruption, and the structure of protocols are not, and do not need to be fixed in order to define this realization relation.

Composition Theorems. Composition theorems allow for the modular analysis and design of systems and are one of the main features of the universal composability paradigm. Our first composition theorem handles concurrent composition of a fixed number of (different) protocol systems. The second one guarantees secure composition of an unbounded number of copies of a protocol system.

Theorem 1 (informal). Let $k \geq 1$. Let $\mathcal{Q}, \mathcal{P}_1, \dots, \mathcal{P}_k, \mathcal{F}_1, \dots, \mathcal{F}_k$ be protocol systems such that they connect only via their I/O interfaces, $\mathcal{Q}|\mathcal{P}_1|\dots|\mathcal{P}_k$ is environmentally bounded, and $\mathcal{P}_i \leq^{SS} \mathcal{F}_i$, for $i \in \{1, \dots, k\}$. Then, $\mathcal{Q}|\mathcal{P}_1|\dots|\mathcal{P}_k \leq^{SS} \mathcal{Q}|\mathcal{F}_1|\dots|\mathcal{F}_k$.

Note that this theorem does not require that the protocols $\mathcal{P}_i/\mathcal{F}_i$ are subprotocols of \mathcal{Q} , i.e., that \mathcal{Q} has matching external I/O tapes for all of these protocols. How these protocols connect to each other via their I/O interfaces is not restricted in any way, even the environment could connect directly to the parts of the I/O interfaces of these protocols that are not taken by another protocol.

For the following composition theorem, we introduce the notion of a session version of a protocol in order to be able to address copies of the protocol. Given an IITM M , the *session version* \underline{M} of M is an IITM which internally simulates M and acts as a “wrapper” for M . More precisely, in mode `CheckAddress`, \underline{M} accepts an incoming message m' only if the following conditions are satisfied: (i) \underline{M} has not accepted a message yet (in mode `CheckAddress`), m' is of the form (id, m) , and m is accepted by the simulated M in mode `CheckAddress`. (In this case, later when activated in mode `Compute`, the ID id will be stored by \underline{M} .) (ii) \underline{M} has accepted a message before, m' is of the form (id', m) , id' coincides with the ID id that \underline{M} has stored before (in mode `Compute`), and m is accepted by M when simulated in mode `CheckAddress`. In mode `Compute`, if \underline{M} is activated for the first time in this mode, i.e., the incoming message, say $m' = (id, m)$, was accepted in mode `CheckAddress` for the first time, then first id is stored and then M is simulated with input m . Otherwise (if \underline{M} was activated in mode `Compute` before), M is directly simulated with input m . If the simulated M produces output on some tape, then \underline{M} prefixes this output with id and then outputs the resulting message on the corresponding tape.

The ID id typically is some session ID (SID) or some party ID (PID) or a combination of both. Clearly, it is not essential that messages are of the form (id, m) . Other forms are possible as well. In fact, everything checkable in polynomial time works. We sometimes require the ID to belong to a specific (polynomially decidable) domain.

To illustrate the notion of a session version of an IITM, assume that M specifies some ideal functionality. Then $!\underline{M}$ denotes the multi-session version of M , i.e., a system in which an unbounded number of copies of M can be created where every copy of M can be addressed by a unique ID, where the ID could be a PID (then an instance of \underline{M} might model one party running M), an SID (then an instance of \underline{M} models one session of M), or it could have a more complex structure, e.g., (sid, pid) (then \underline{M} models an instance of party pid running M in session sid).

Given a system \mathcal{S} , its *session version* $\underline{\mathcal{S}}$ is obtained by replacing all IITMs in \mathcal{S} by their session version. For example, we obtain $\underline{\mathcal{S}} = \underline{M}|\dots|\underline{M}'$ for $\mathcal{S} = M|\dots|M'$.

Now, the following composition theorem says that if a protocol \mathcal{P} realizes \mathcal{F} , then the multi-session version of \mathcal{P} realizes the multi-session version of \mathcal{F} .

Theorem 2 (informal). Let \mathcal{P} and \mathcal{F} be protocol systems such that $!\underline{\mathcal{P}}$ is environmentally bounded and $\mathcal{P} \leq^{SS} \mathcal{F}$. Then, $!\underline{\mathcal{P}} \leq^{SS} !\underline{\mathcal{F}}$.

Theorems 1 and 2 can be applied iteratively to construct more and more complex systems. For example, as a corollary of the above theorems, we immediately obtain that for any protocol system \mathcal{Q} : $\mathcal{P} \leq^{SS} \mathcal{F}$ implies $\mathcal{Q} | \mathcal{P} \leq^{SS} \mathcal{Q} | \mathcal{F}$, provided that $\mathcal{Q} | \mathcal{P}$ is environmentally bounded. In words: \mathcal{Q} using an unbounded number of copies of \mathcal{P} realizes \mathcal{Q} using an unbounded number of copies of \mathcal{F} .

When addressing a session version \underline{M} of a machine M , then the machine M simulated within \underline{M} is not aware of its ID and cannot use it. For example, it cannot put the ID into a message that M creates. However, sometimes this is desirable. We therefore also consider another, more general composition theorem where machines are aware of their IDs. (Maybe a little surprisingly, this theorem is a corollary of the above theorem.) While these IDs can, as already mentioned above, be interpreted in different ways, we will often refer to them as SIDs.

To this end, we first generalize the notion of a session version. We consider a (polynomially computable) *session identifier function* σ which, given a message and a tape name, outputs an SID (a bit string) or \perp . For example, the following function takes the prefix of a message as its SID: $\sigma_{\text{prefix}}(m, c) := s$ if $m = (s, m')$ for some s, m' and $\sigma_{\text{prefix}}(m, c) := \perp$ otherwise, for all m, c . Clearly, many more examples are conceivable. The reason that σ , besides a message, also takes a tape name as input is that the way SIDs are extracted from messages may depend on the tape name a message is received from.

Now, we say that an IITM M is a σ -*session machine* (or a σ -*session version*) if the following conditions are satisfied: (i) M rejects (in mode `CheckAddress`) a message m on tape c if $\sigma(m, c) = \perp$. (ii) If m_0 is the first message that M accepted (in mode `CheckAddress`), say on tape c_0 , in a run, then, M will reject all messages m received on some tape c (in mode `CheckAddress`) with $\sigma(m, c) \neq \sigma(m_0, c_0)$. (iii) Whenever M outputs a messages m on tape c (in mode `Compute`), then $\sigma(m, c) = \sigma(m_0, c_0)$, with m_0 and c_0 as before.

It is easy to see that session versions are specific forms of σ -session versions: given an IITM M , we have that \underline{M} is a σ_{prefix} -session version. The crucial difference is that while σ -session versions look like session version from the outside, inside they are aware of their SID.

We call an environmental system \mathcal{E} σ -*single session* if it only outputs messages with the same SID according to σ . Hence, when interacting with a σ -session version, such an environmental system invokes at most one protocol session.

Let \mathcal{P} and \mathcal{F} be protocol systems, which in the setting considered here would typically describe multiple sessions of a protocol. Moreover, we assume that \mathcal{P} and \mathcal{F} are σ -session versions. Now, we define what it means that a single session of \mathcal{P} realizes a single session of \mathcal{F} . This is defined just as $\mathcal{P} \leq^{SS} \mathcal{F}$, with the difference that we consider only σ -single session environments, and hence, environments that invoke at most one session of \mathcal{P} and \mathcal{F} .

Definition 2 (informal). Let \mathcal{P} , \mathcal{F} , and σ be as above. Then, \mathcal{P} *single-session realizes* \mathcal{F} w.r.t. σ ($\mathcal{P} \leq_{\sigma\text{-single}}^{SS} \mathcal{F}$) if and only if there exists an adversarial system \mathcal{S} (a simulator or an ideal adversary) such that $\mathcal{E} | \mathcal{P} \equiv \mathcal{E} | \mathcal{S} | \mathcal{F}$ for every σ -single session environmental system \mathcal{E} . (The details concerning runtime and interfaces are similar to Definition 1.)

Now, analogously to Theorem 2, the following theorem says that if \mathcal{P} realizes \mathcal{F} w.r.t. a single session, then \mathcal{P} realizes \mathcal{F} w.r.t. multiple sessions. As mentioned before, in the setting considered here \mathcal{P} and \mathcal{F} would typically model multi-session versions of a protocol/functionality.

Theorem 3 (informal). *Let σ , \mathcal{P} , and \mathcal{F} be as above. Then, $\mathcal{P} \leq_{\sigma\text{-single}}^{SS} \mathcal{F}$ implies $\mathcal{P} \leq^{SS} \mathcal{F}$.*

As pointed out before, the proof of this theorem is in fact a corollary of Theorem 2. Clearly, this theorem can be combined with the other composition theorems to construct more and more complex systems.

We finally emphasize that still details such as addressing of machines by party/session IDs, corruption, structure of protocols are not, and do not need to be fixed in order to prove the above composition theorems. In other words, these theorems hold true for all specific choices, and hence, are very general.

Joint state theorems and composition theorems with global setup are discussed in Section 10.1. As already mentioned in the introduction, unlike other models, these theorems do not require to extend or change the IITM model. The general joint state theorem is even a trivial consequence of the composition theorems.

At first reading, the reader may want to skip the detailed description of the IITM model and jump directly to Section 10, where the use of the model is illustrated.

3 The General Computational Model

In this section, we define our general computational model. This model is defined independently of the application to universal composability and is of independent interest. It will, however, form the basis for our treatment of universal composability. We introduce single interactive Turing machines and systems of such machines, define runs of systems, and introduce further notation and terminology.

3.1 Inexhaustible Interactive Turing Machines

We first introduce the syntax of (inexhaustible) interactive Turing machines and then the way these machines perform their computations.

3.1.1 Syntax

An (*inexhaustible*) *interactive Turing machine (IITM, for short)* M is a probabilistic Turing machine with the following tapes and a polynomial q associated with it, where q will be used as a bound in computations of M in mode `CheckAddress` (see below): a read-only tape on which the mode the IITM M is supposed to run is written (the *mode tape*)—the possible modes are `CheckAddress` and `Compute` (see below)—, a read-only tape on which the random coins are written (the *random tape*), a read-only tape on which the security parameter is written (the *security parameter tape*), a write-only tape (the *address decision tape*, used in mode `CheckAddress`), zero or more *input* and *output tapes*, and *work tapes*. The input and output tapes have names and we require that different tapes of M have different names.

The set of (names of) input and output tapes of M is denoted by $\mathcal{T}(M)$, the set of input tapes by $\mathcal{T}_{in}(M)$, and the set of output tapes by $\mathcal{T}_{out}(M)$.

As further explained below, the names of input and output tapes will determine how IITMs are connected in a system of IITMs: if an IITM sends a message on an output tape named c , then only (a copy of) an IITM with an input tape named c can receive this message.

Tapes named `start` and `decision` will serve a particular purpose. We require that only input tapes can be named `start` and only output tapes can be named `decision`. We will later use `start` to provide a system with external input and to trigger an IITM if no other IITM was triggered. An IITM is triggered by another IITM if the latter sends a message to the former. An IITM with an input tape named `start` will be called *master IITM*. On tapes named `decision` the final output of a system of IITMs will be written.

As mentioned, an IITM M runs in one of two modes, `CheckAddress` or `Compute`. The mode in which M is supposed to run is written on the mode tape of M .

3.1.2 Computation

We describe the computation of an IITM M in mode `CheckAddress` and `Compute`, respectively. Informally speaking, in mode `CheckAddress` an IITM M checks whether the incoming message is in fact addressed to it. Typically, this mode is used for the following purpose: In a system of IITMs there may be several copies of M (belonging to different parties in a multi-party protocol and/or to different sessions of a protocol). To address the different copies one can, for example, prefix messages with identifiers (for example, session identifiers (SIDs) and/or party identifiers (PIDs)). Now, in mode `CheckAddress`, M checks whether the incoming message has the expected form (e.g., whether it is prefixed with the expected identifier), and either accepts or rejects that message. In mode `Compute`, an IITM can, depending on the incoming message and its current configuration, actually process incoming messages and write output on one of its output tapes, i.e., send a message to another IITM.

More formally, the computation in the two modes is defined as follows.

Mode `CheckAddress`: If M is activated in mode `CheckAddress`, then the following will always be the case: `CheckAddress` is written on the mode tape of M , the security parameter, say 1^n , is written on the security parameter tape, and one message, say m , is written on one of the input tapes, say c ; the other input tapes

and the output tapes are empty—or otherwise will be emptied before M starts to run—and the contents on the work tapes and the random tape represent the current configuration of M . We require that in mode `CheckAddress` (i) M always terminates and at the end of the activation has written `accept` or `reject` on the address decision tape, (ii) the computation performed by M in this mode is deterministic, i.e., it is independent of the content on the random tape, and (iii) the number of transitions taken in the activation is bounded from above by $q(\eta + |m| + l)$, where q is the polynomial associated with M and l is the length of the content of all work tapes at the beginning of the activation.

We emphasize that in mode `CheckAddress`, M cannot be exhausted. That is, whenever M is activated in this mode, M is able to “scan” its complete current configuration (except for the random tape, on which an infinite string is written), including the incoming message.

Mode Compute: If M is activated in mode `Compute`, then the following will always be the case: `Compute` is written on the mode tape of M , the security parameter, say 1^η , is written on the security parameter tape, and one message, say m , is written on one of the input tapes, say c ; the other input tapes and the output tapes are empty—or otherwise will be emptied before M starts to run—and the contents on the work tapes and the random tape represent the current configuration of M . The computation of M in mode `Compute` may be probabilistic (i.e., depend on the content on the random tape) and it might be the case that M does not terminate. However, if M does terminate in this activation, then we require that M has written *at most one* message on one of its output tapes (i.e., only one message can be sent to another IITM at a time).

We note that at this point, we do not restrict the runtime of IITMs in mode `Compute`. IITMs are a priori unbounded w.r.t. their runtime (number of transitions taken) and space (length of the content of all tapes). Later, in Section 4, we will introduce a notion of polynomial runtime which will guarantee that the overall runtime of a system of IITMs is polynomially bounded in the security parameter.

3.2 Systems of IITMs

A *system* of IITMs can be built according to the following grammar, where M ranges over (descriptions of) IITMs:

$$\mathcal{S} ::= M \mid (\mathcal{S} \parallel \mathcal{S}) \mid !\mathcal{S} .$$

We require that for every tape name c , at most two of the machines in \mathcal{S} have a tape named c and if two IITMs have a tape named c , then c is an input tape in one of the machines and an output tape in the other. This implies that in \mathcal{S} only at most one IITM may be a master IITM, i.e., may have `start` as input tape; there may be several copies of such a machine in a run of a system though. We say that \mathcal{S}' is a *subsystem* of \mathcal{S} if \mathcal{S} contains a subexpression (modulo commutativity and associativity of \parallel) of the form \mathcal{S}' . For example, $!M_3$ and $M_1 \parallel !M_3$ are subsystems of $\mathcal{S} = M_1 \parallel M_2 \parallel !M_3$, but $!M_1$ is not.

Intuitively, $\mathcal{S}_1 \parallel \mathcal{S}_2$ stands for the concurrent composition of the systems \mathcal{S}_1 and \mathcal{S}_2 , and $!\mathcal{S}$ stands for the concurrent composition of an unbounded number of copies of (machines in) the system \mathcal{S} , where the actual number of copies generated during a run of the system is determined by external or internal machines invoking (machines of) \mathcal{S} (see Section 3.3). We call ‘!’ the *bang operator*, borrowing terminology from process calculus [12, 25].

We say that an IITM M *occurs in the scope of a bang* in \mathcal{S} if \mathcal{S} contains a subexpression of the form $!\mathcal{S}'$ such that M occurs in \mathcal{S}' .

It will be clear from the semantics of systems, i.e., the way a system of IITMs runs, that every system \mathcal{S} can equivalently be written as $\mathcal{S} = M_1 \parallel \dots \parallel M_k \parallel !M'_1 \parallel \dots \parallel !M'_{k'}$, where M_1, \dots, M_k and $M'_1, \dots, M'_{k'}$ are IITMs, i.e., every system consists of a set of machines, where some are and other are not in the scope of a bang.

3.3 Running a System

Throughout the rest of this paper, we denote by Rand the set of all mappings from the set of natural numbers \mathbb{N} to the set of infinite bit strings $\{0, 1\}^\omega$. We refer to $\alpha \in \text{Rand}$ as *random coins*.

We now define how a system \mathcal{S} runs given a security parameter η , a bit string a as external input, and random coins $\alpha \in \text{Rand}$.² We denote such a system by $\mathcal{S}^{(\alpha)}(1^\eta, a)$.

Informal description. Informally speaking, in the run of $\mathcal{S}^{(\alpha)}(1^\eta, a)$ at any time only one (copy of an) IITM is active and all other IITMs wait for new input. The active copy, say M' , which is a copy of a machine M defined in \mathcal{S} , may write at most one message, say m , on one of its output tapes, say c . This message is then delivered to another (copy of an) IITM with an input tape named c , say N is the machine specified in \mathcal{S} with an input tape named c .³ In the current configuration of the system, there may be several copies of N . In the order of creation, the copies of N are run in mode `CheckAddress` with input m . Once one copy accepts m , this copy gets to process m , i.e., it runs in mode `Compute` with input m , and in particular, may produce output on one output tape, which is then sent to another copy and so on. If no copy of N accepts m and N is in the scope of a bang, a fresh copy of N is created and run in mode `CheckAddress`. If this copy accepts m , random coins $\alpha(i)$ for some new i are written on the random tape of N and it gets to process m in mode `Compute`. Otherwise, the new copy of N is deleted, m is dropped, and a master IITM is activated (with empty input on tape `start`). If N is not in the scope of a bang (and—the only copy of— N does not accept m), then too a master IITM is activated. The first IITM to be activated in a run is a master IITM. It gets a as external input (on tape `start`) and is run with random coins $\alpha(1)$ written on its random tape. A master IITM is also activated if the currently active machine does not produce output. A run stops if a master IITM, after being activated, does not produce output or output was written by some machine on an output tape named `decision`. An informal example of a run of a system was provided in Section 2.

Formal definition. Formally, the run of $\mathcal{S}^{(\alpha)}(1^\eta, a)$ is defined as follows: The current (global) configuration of a system in the run is described by a tuple (A, i, P) where (i) A is a sequence of configurations of IITMs, the sequence of (*previously*) *activated machines*, (ii) $i \leq |A|$ (where $|A|$ denotes the length of the sequence A) is the index of the last active machine in A (where $i = 0$ if A is the empty sequence), and (iii) P is a system. The IITMs occurring in P are called *passive*. We emphasize that the configurations in A are not the configurations of the machines that are currently active, i.e., currently performing some computation—only the i -th configuration (machine) was just active. The configurations in A rather belong to those machines that were active at some point in the run so far. As usual, the configuration of (an instance of) a machine is the content of all of its tapes, the position of the heads on these tapes, and its state. In what follows, we often do not distinguish between an (instance of an) IITM M and its current configuration: by abuse of notation, we write M for both the machine and its current configuration.

In order to define the run of the system $\mathcal{S}^{(\alpha)}(1^\eta, a)$, we first need to introduce some more notation. Given (a configuration of) an IITM M , we write

$$M(\text{CheckAddress}, c, m) = \text{accept}$$

to say that when running the IITM M in mode `CheckAddress` starting from its current configuration with m written on the input tape c and the empty bit string written on all other input tapes, on all output tapes, and the address decision tape, then M returns `accept` on its address decision tape. (Note that by definition, see Section 3.1.2, M always terminates when running in mode `CheckAddress` and runs in polynomial time. Also recall that the run of M in this mode is deterministic.) Analogously, we write

$$M(\text{CheckAddress}, c, m) = \text{reject}.$$

Similarly, given (a configuration of) an IITM M , we write

$$M(\text{Compute}, c, m) \rightarrow M'$$

²We note that the external input can be omitted. All our results also hold true in the uniform case, that is, the setting without external input. Most proofs directly carry over to the uniform case and otherwise we explicitly provide the proof for the uniform case.

³Recall that by our convention on the names of input tapes in systems of IITMs, there can be at most one such machine.

to say that when running the IITM M in mode `Compute` starting from its current configuration with m written on the input tape c and the empty bit string written on all other input tapes, on all output tapes, and the address decision tape, M terminates in configuration M' . (Note that we assume that random coins have been written on the random tape of M . By this, the run of M is fully determined.) If in the above setting M does not terminate, we write

$$M(\text{Compute}, c, m) \rightarrow \infty .$$

We now describe how a configuration (A, i, P) evolves when a message m which was output on an output tape c is read by one of the IITMs in the system, given random coins α . We will write $(A, i, P) \rightarrow_{(c,m)}^\alpha (A', i', P')$ to say that we obtain (A', i', P') as a successor configuration of (A, i, P) after m was read on c (by some IITM). Accordingly, we call (A', i', P') a $\rightarrow_{(c,m)}^\alpha$ -*successor* of (A, i, P) . Note that such a successor (if any) is uniquely determined.

Let $c \neq \text{decision}$ be a name of a tape, m be a message, α be random coins, (A, i, P) and (A', i', P') be configurations such that A does not contain ∞ (which would mean that some of the machines in A did not terminate in an activation). Then, we have

$$(A, i, P) \rightarrow_{(c,m)}^\alpha (A', i', P')$$

if one of the following conditions is satisfied where we assume that $A = M_1, \dots, M_n$.

1. *One of the activated machines accepts m on tape c :* It holds that $c \in \mathcal{T}_{in}(M_{i'})$, $M_{i'}(\text{CheckAddress}, c, m) = \text{accept}$, and i' is minimal with this property, i.e., $M_j(\text{CheckAddress}, c, m) = \text{reject}$ for all $j < i'$ with $c \in \mathcal{T}_{in}(M_j)$. Furthermore, there exists a configuration $M'_{i'}$ (possibly ∞) such that $M_{i'}(\text{Compute}, c, m) \rightarrow M'_{i'}$ and A' is obtained from A by replacing the content of every input and output tape of a configuration in A by the empty bit string and then replacing $M_{i'}$ by $M'_{i'}$ where, if $M'_{i'} \neq \infty$, the content of all input tapes of $M'_{i'}$ are replaced by the empty bit string. (Output tapes of $M'_{i'}$ are not emptied. One such tape may contain a non-empty message.) Moreover, $P' = P$.
2. *None of the activated machines accepts m on tape c , but a fresh copy of a machine does:* It holds that $i' = n + 1$ and for all $j \leq n$ with $c \in \mathcal{T}_{in}(M_j)$ it holds that $M_j(\text{CheckAddress}, c, m) = \text{reject}$ but there is an IITM M in P such that $c \in \mathcal{T}_{in}(M)$ and $M(\text{CheckAddress}, c, m) = \text{accept}$, where we identify M with its initial configuration, with 1^n written on its security parameter tape and with $\alpha(n + 1)$ written on the random tape of M . Furthermore, there exists a configuration M' (possibly ∞) such that $M(\text{Compute}, c, m) \rightarrow M'$ and A' is obtained from A by replacing the content of every input and output tape of a configuration of A by the empty bit string and appending M' at the end of A where the contents of all input tapes of M' are also deleted (if $M' \neq \infty$). (Output tapes of $M'_{i'}$ are not emptied. One such tape may contain a non-empty message.) If M is in the scope of a bang in P , then $P' = P$. Otherwise P' is obtained from P by removing M from P .

If neither 1. nor 2. is applicable, i.e., no machine (not even a fresh one) accepted m on tape c , (A, i, P) does not have a $\rightarrow_{(c,m)}^\alpha$ -successor.

We emphasize that both in 1. and 2. the configuration in which the mode `Compute` is performed is the same as the configuration in which the machine was *before* running in mode `CheckAddress`: in 1., both the mode `CheckAddress` and `Compute` are executed in the configuration $M_{i'}$, and in 2., both the mode `CheckAddress` and `Compute` are executed in the configuration M . We also note that one of the output tapes of $M'_{i'}/M'$ may contain a non-empty bit string and all other output tapes, including those of other IITMs, are empty. Finally, we point out that in 2., if some M occurs in P with $c \in \mathcal{T}_{in}(M)$, then M is uniquely determined. This is so because by definition of systems we assume that the set of names of input tapes of different occurrences of IITMs in a system are disjoint.

Having defined how a configuration (A, i, P) evolves when a machine reads a message m from tape c , we now define how a configuration (A, i, P) evolves in general. For this purpose, let α be random coins, (A, i, P)

and (A', i', P') be configurations such that A does not contain ∞ and all output tapes of the configurations occurring in A are empty, except for at most one output tape $c \neq \text{decision}$. We write

$$(A, i, P) \rightarrow^\alpha (A', i', P')$$

if one of the following conditions is satisfied where we assume that $A = M_1, \dots, M_n$.

1. *The last active IITM did not produce output and was not a master IITM. Then a master IITM is triggered:* All output tapes of the configurations in A are empty, $\text{start} \notin \mathcal{T}_{in}(M_i)$, and $(A, i, P) \rightarrow_{(\text{start}, \varepsilon)}^\alpha (A', i', P')$ where ε denotes the empty bit string. (If a master IITM did not produce output ($\text{start} \in \mathcal{T}_{in}(M_i)$), then (A, i, P) does not have a successor configuration.)
2. *The last active IITM produced output and this output is accepted by another machine:* The content of some output tape $c \neq \text{decision}$ is some non-empty message m and $(A, i, P) \rightarrow_{(c, m)}^\alpha (A', i', P')$.
3. *The last active IITM produced output which, however, is not accepted by the intended machine. Then a master IITM is triggered:* The content of some output tape $c \neq \text{decision}$ is some non-empty message m , but (A, i, P) does not have a $\rightarrow_{(c, m)}^\alpha$ -successor, and $(A, i, P) \rightarrow_{(\text{start}, \varepsilon)}^\alpha (A', i', P')$ where ε denotes the empty bit string. (Note that this means that the master IITM accepted the empty message. Otherwise (A, i, P) does not have a successor.)

We refer to (A', i', P') as a \rightarrow^α -successor of (A, i, P) . If none of the above cases applies, then (A, i, P) does not have a \rightarrow^α -successor. Also, if A contains ∞ or a non-empty message is written on the output tape **decision**, and hence, there is not non-empty message on another output tape, (A, i, P) does not have a \rightarrow^α -successor either.

The *(complete) run* ρ of a system \mathcal{S} given the security parameter η , external input a , and random coins α (the run of $\mathcal{S}^{(\alpha)}(1^\eta, a)$, for short) is the finite sequence of configurations $(A_0, i_0, P_0), (A_1, i_1, P_1), \dots, (A_k, i_k, P_k)$ or the infinite sequence of configurations $(A_0, i_0, P_0), (A_1, i_1, P_1), \dots$ such that the following conditions are satisfied.

1. (A_0, i_0, P_0) is the initial configuration, i.e., A_0 is the empty sequence, $i_0 = 0$, and $P_0 = \mathcal{S}$.
2. $(A_0, i_0, P_0) \rightarrow_{(\text{start}, a)}^\alpha (A_1, i_1, P_1)$.
3. $(A_j, i_j, P_j) \rightarrow^\alpha (A_{j+1}, i_{j+1}, P_{j+1})$ for every $j \in \{1, \dots, k-1\}$ if ρ is finite and for every $j \geq 1$ if ρ is infinite.
4. If the sequence is finite, then (A_k, i_k, P_k) does not have a \rightarrow^α -successor.

For a finite run ρ , we call k the *length* of ρ . The *overall output* of a run ρ is undefined if ρ is infinite or if in ρ some IITM does not terminate (i.e., A_k contains a configuration ∞). Otherwise, the overall output of a run ρ is the message written on a tape named **decision** at the end of the run. (This message could be empty.) Note that since a run stops if a machine has written a non-empty message on **decision**, the overall output is uniquely determined.

3.4 Probability Space and Relevant Random Variables

We consider the standard probability space over $\{0, 1\}^\omega$ (the set of infinite bit strings), where the probability of a cone $\bar{\beta}$ of $\beta \in \{0, 1\}^*$, which is the set of all infinite bit strings that have β as a prefix, is defined to be $2^{-|\beta|}$. Now, the probability space for **Rand** is defined in a standard way as the probability space on the infinite product of the probability space for $\{0, 1\}^\omega$. For example, given $\beta_1, \beta_2, \beta_3 \in \{0, 1\}^*$, the probability for the event $E = \{\alpha \in \text{Rand} \mid \alpha(i) \text{ is prefixed with } \beta_i \text{ for all } i \in \{1, 2, 3\}\}$ is $2^{-|\beta_1|} \cdot 2^{-|\beta_2|} \cdot 2^{-|\beta_3|}$. This probability space over **Rand** guarantees that all random variables defined next are measurable.

By $\mathcal{S}(1^\eta, a): \text{Rand} \rightarrow \{0, 1\}^* \cup \{\perp\}$ we denote the random variable that describes the overall output (i.e., output on **decision**) of runs of the system $\mathcal{S}(1^\eta, a)$. More precisely, for $\alpha \in \text{Rand}$ we define $\mathcal{S}(1^\eta, a)(\alpha)$ to be the overall output of the run of $\mathcal{S}^{(\alpha)}(1^\eta, a)$, where \perp denotes undefined output.

Given this random variable, the probability that a run has overall output 1 is

$$\text{Prob}[\mathcal{S}(1^\eta, a) = 1] .$$

By $\text{Time}(\mathcal{S}(1^\eta, a)) : \text{Rand} \rightarrow \mathbf{N} \cup \{\infty\}$ we denote the random variable that describes the overall number of transitions that have been taken by IITMs in mode **Compute** in runs of $\mathcal{S}(1^\eta, a)$. More precisely, for $\alpha \in \text{Rand}$ we define $\text{Time}(\mathcal{S}(1^\eta, a))(\alpha)$ to be the overall number of transitions that have been taken by IITMs in mode **Compute** in a run of $\mathcal{S}^{(\alpha)}(1^\eta, a)$. This number is ∞ if the run is infinite or some IITM did not terminate. Note that transitions taken in mode **CheckAddress** are not counted and also the emptying of input or output tapes before IITMs are activated is not counted.

Similarly, given a subsystem \mathcal{Q} of \mathcal{S} , by $\text{Time}_{\mathcal{Q}}(\mathcal{S}(1^\eta, a))$ we denote the random variable that describes the overall number of transitions in mode **Compute** that have been taken by IITMs in \mathcal{Q} in runs of $\mathcal{S}(1^\eta, a)$. Clearly, $\text{Time}_{\mathcal{S}}(\mathcal{S}(1^\eta, a)) = \text{Time}(\mathcal{S}(1^\eta, a))$ for all systems \mathcal{S} .

3.5 Equivalence/Indistinguishability of Systems

We first introduce negligible functions following [3].

Definition 3. A function $f : \mathbf{N} \times \{0, 1\}^* \rightarrow \mathbf{R}_{\geq 0}$ is called *negligible* if for all $c, d \in \mathbf{N}$ there exists $\eta_0 \in \mathbf{N}$ such that for all $\eta > \eta_0$ and all $a \in \bigcup_{\eta' \leq \eta^a} \{0, 1\}^{\eta'}$: $f(\eta, a) < \eta^{-c}$.⁴

A function $f : \mathbf{N} \times \{0, 1\}^* \rightarrow [0, 1]$ is called *overwhelming* if $1 - f$ is negligible.⁵

Two systems that produce overall output 1 with almost the same probability are called equivalent or indistinguishable:

Definition 4. Let $f : \mathbf{N} \times \{0, 1\}^* \rightarrow \mathbf{R}_{\geq 0}$ be a function. Two systems \mathcal{P} and \mathcal{Q} are called *f-equivalent* or *f-indistinguishable* ($\mathcal{P} \equiv_f \mathcal{Q}$) if and only if for every security parameter $\eta \in \mathbf{N}$ and external input $a \in \{0, 1\}^*$:

$$|\text{Prob}[\mathcal{P}(1^\eta, a) = 1] - \text{Prob}[\mathcal{Q}(1^\eta, a) = 1]| \leq f(\eta, a) .$$

Two systems \mathcal{P} and \mathcal{Q} are called *equivalent* or *indistinguishable* ($\mathcal{P} \equiv \mathcal{Q}$) if and only if there exists a negligible function f (Definition 3) such that $\mathcal{P} \equiv_f \mathcal{Q}$.

It is easy to see that for every two functions f, f' as in Definition 4 the relation \equiv_f is reflexive and that $\mathcal{P} \equiv_f \mathcal{Q}$ and $\mathcal{Q} \equiv_{f'} \mathcal{S}$ implies $\mathcal{P} \equiv_{f+f'} \mathcal{S}$. In particular, \equiv is reflexive and transitive.

4 Polynomial Time and Properties of Systems

In this section, we introduce notions of polynomial runtime for arbitrary systems, environmental systems, and protocol systems. We also state basic properties about such systems. We begin with further notation and terminology.

4.1 Further Notation and Terminology

Let \mathcal{S} be a system and M be an IITM. Recall that $\mathcal{T}(M)$, $\mathcal{T}_{in}(M)$, and $\mathcal{T}_{out}(M)$ denote the set of (names of) input and output tapes, the set (of names) of input tapes, and the set (of names) of output tapes of M , respectively.

A tape c in $\mathcal{T}(\mathcal{S})$ is called *internal* if there exist two IITMs M and M' in \mathcal{S} such that $c \in \mathcal{T}_{out}(M) \cap \mathcal{T}_{in}(M')$. Otherwise, c is called *external*. The set of *internal tapes* of \mathcal{S} is denoted by $\mathcal{T}_{int}(\mathcal{S})$ and the set of *external*

⁴We note that this definition of negligible is equivalent to the following: f is negligible if and only if for all positive polynomials p and q (i.e., $p(n) > 0$ and $q(n) > 0$ for all $n \in \mathbf{N}$) there exists $\eta_0 \in \mathbf{N}$ such that for all $\eta > \eta_0$ and all $a \in \bigcup_{\eta' \leq q(\eta)} \{0, 1\}^{\eta'}$: $f(\eta, a) < \frac{1}{p(\eta)}$. We further note that such negligible functions have the following properties: (i) If f and g are negligible, then $f + g$ is negligible. (ii) If f is negligible and p is a positive polynomial, then $g(\eta, a) := p(\eta + |a|) \cdot f(\eta, a)$ for all η, a is negligible.

⁵By $[0, 1]$ we denote the interval of all real numbers x such that $0 \leq x \leq 1$.

tapes of \mathcal{S} by $\mathcal{T}_{ext}(\mathcal{S})$. We call c an (*external*) *input tape* of \mathcal{S} if $c \in \mathcal{T}_{ext}(\mathcal{S})$ and $c \in \mathcal{T}_{in}(M)$ for some IITM M in \mathcal{S} . Analogously, c is called an (*external*) *output tape* of \mathcal{S} if $c \in \mathcal{T}_{ext}(\mathcal{S})$ and $c \in \mathcal{T}_{out}(M)$ for some IITM M in \mathcal{S} . The set of (*external*) input and output tapes of \mathcal{S} is denoted by $\mathcal{T}_{in}(\mathcal{S})$ and $\mathcal{T}_{out}(\mathcal{S})$, respectively.

Note that for every \mathcal{S} we have that $\mathbf{start} \in \mathcal{T}(\mathcal{S})$ implies $\mathbf{start} \in \mathcal{T}_{in}(\mathcal{S})$ and $\mathbf{decision} \in \mathcal{T}(\mathcal{S})$ implies $\mathbf{decision} \in \mathcal{T}_{out}(\mathcal{S})$.

We call two systems compatible if they provide the same external interface:

Definition 5. Two systems \mathcal{P} and \mathcal{Q} are *compatible* iff $\mathcal{T}_{in}(\mathcal{P}) = \mathcal{T}_{in}(\mathcal{Q})$ and $\mathcal{T}_{out}(\mathcal{P}) = \mathcal{T}_{out}(\mathcal{Q})$, i.e., \mathcal{P} and \mathcal{Q} coincide on their input and output tapes.

Given two systems \mathcal{P} and \mathcal{Q} , by

$$\mathcal{P} | \mathcal{Q}$$

we denote the parallel composition $\mathcal{P}' || \mathcal{Q}'$ where \mathcal{P}' and \mathcal{Q}' are obtained from \mathcal{P} and \mathcal{Q} by renaming the internal tapes of \mathcal{P} and \mathcal{Q} , respectively, such that $\mathcal{T}(\mathcal{P}') \cap \mathcal{T}_{int}(\mathcal{Q}') = \emptyset$ and $\mathcal{T}_{int}(\mathcal{P}') \cap \mathcal{T}(\mathcal{Q}') = \emptyset$. The intuition is that \mathcal{P} and \mathcal{Q} are different systems (e.g., a protocol and its environment) which communicate via their external tapes; they should not interfere on their internal tapes.

Definition 6. The systems \mathcal{Q} and \mathcal{P} are *connectable* if each common external tape of \mathcal{P} and \mathcal{Q} has complementary directions (input or output), i.e., for all $c \in \mathcal{T}_{ext}(\mathcal{P}) \cap \mathcal{T}_{ext}(\mathcal{Q})$, we have that $c \in \mathcal{T}_{in}(\mathcal{P}) \cap \mathcal{T}_{out}(\mathcal{Q})$ or $c \in \mathcal{T}_{out}(\mathcal{P}) \cap \mathcal{T}_{in}(\mathcal{Q})$. If \mathcal{Q} and \mathcal{P} are *connectable*, then we also say that \mathcal{P} *can be connected to* \mathcal{Q} . We denote by $\text{Con}(\mathcal{Q})$ the set of all systems \mathcal{P} that can be connected to \mathcal{Q} .

The systems $\mathcal{S}_1, \dots, \mathcal{S}_n$ are *connectable* if they are pairwise connectable (i.e., \mathcal{S}_i and \mathcal{S}_j are connectable for every $i, j \leq n$ such that $i \neq j$).

We note that if $\mathcal{S}_1, \dots, \mathcal{S}_n$ are connectable, then every common external tape of the systems is the external input tape of exactly one system \mathcal{S}_i and the external output tape of exactly one other system \mathcal{S}_j . So, there is no ambiguity about how these systems connect to each other in a parallel composition. In particular, the order in which these systems are composed does not matter. In fact, the composition operator ‘ $|$ ’ is associative (i.e., $\mathcal{P} | (\mathcal{Q} | \mathcal{S}) \equiv_0 (\mathcal{P} | \mathcal{Q}) | \mathcal{S}$) and commutative (i.e., $\mathcal{P} | \mathcal{Q} \equiv_0 \mathcal{Q} | \mathcal{P}$).

4.2 Polynomially Bounded Systems

The following notion captures that a system runs in polynomial time (except maybe with negligible probability), i.e., the overall number of transitions taken by the IITMs is bounded from above by a polynomial (in the security parameter plus the length of the external input).

Definition 7. A system \mathcal{S} is *almost bounded* if there exists a polynomial p such that:

$$f(\eta, a) := \text{Prob}[\text{Time}(\mathcal{S}(1^\eta, a)) > p(\eta + |a|)] \text{ for all } \eta \in \mathbf{N} \text{ and } a \in \{0, 1\}^*$$

is negligible (as a function in η and a).

We say that a system \mathcal{S} is *strictly bounded* if there exists a polynomial p such that for every security parameter $\eta \in \mathbf{N}$ and external input $a \in \{0, 1\}^*$: $\text{Time}(\mathcal{S}(1^\eta, a))(\alpha) \leq p(\eta + |a|)$ for all $\alpha \in \text{Rand}$.

Clearly, every strictly bounded system is almost bounded. The next lemma states that every almost bounded system can be simulated by a single strictly bounded IITM M except for a negligible error. In particular, it can be simulated (with negligible error) by a probabilistic polynomial time Turing machine.

Lemma 1. *For every almost bounded system \mathcal{S} there exists an IITM M such that M (as a system) is strictly bounded and $\mathcal{S} \equiv M$.*

Proof. This lemma is a special case of Lemma 7. □

4.3 Environments and Environmental Indistinguishability

We first introduce environmental systems and then define environmental indistinguishability, i.e., indistinguishability where the distinguisher is an environmental system. In the context of universal composability, environmental systems will play the role of an environment.

Intuitively, an environmental system (also called a universally bounded system) is a system that runs in (strict) polynomial time no matter to which system it is connected.

Definition 8. A system \mathcal{E} is called an *environmental system* or *universally strictly bounded*, or simply *universally bounded* if there exists a polynomial p such that for every security parameter $\eta \in \mathbf{N}$, external input $a \in \{0, 1\}^*$, and system \mathcal{S} that can be connected to \mathcal{E} (i.e., $\mathcal{S} \in \text{Con}(\mathcal{E})$) it holds true that $\text{Time}_{\mathcal{E}}((\mathcal{E} | \mathcal{S})(1^\eta, a))(\alpha) \leq p(\eta + |a|)$ for all $\alpha \in \text{Rand}$.⁶

Given a system \mathcal{S} , by $\text{Env}(\mathcal{S})$ we denote the set of all environmental systems \mathcal{E} such that \mathcal{E} and \mathcal{S} are connectable, i.e., $\mathcal{E} \in \text{Con}(\mathcal{S})$, and \mathcal{E} is universally bounded.

We call two systems environmentally indistinguishable if they cannot be distinguished by any environmental system:

Definition 9. Two systems \mathcal{P} and \mathcal{Q} are called *environmentally equivalent* or *environmentally indistinguishable* ($\mathcal{P} \cong \mathcal{Q}$) if and only if

1. \mathcal{P} and \mathcal{Q} are compatible and
2. $\mathcal{E} | \mathcal{P} \equiv \mathcal{E} | \mathcal{Q}$ for all $\mathcal{E} \in \text{Env}(\mathcal{P})$.

Environmental indistinguishability talks about arbitrary systems \mathcal{P} and \mathcal{Q} . In particular, $\mathcal{E} | \mathcal{P}$ or $\mathcal{E} | \mathcal{Q}$ for some $\mathcal{E} \in \text{Env}(\mathcal{P})$ are not necessarily almost bounded. In typical applications and in the context of universal composability, we are interested in systems \mathcal{P} and \mathcal{Q} such that $\mathcal{E} | \mathcal{P}$ and $\mathcal{E} | \mathcal{Q}$ are almost bounded for all environmental systems \mathcal{E} . This class of systems—environmentally bounded systems—will be defined in the next subsection. But first we make some more remarks about the notion of environmental indistinguishability.

Using the fact that the relationship \equiv is reflexive and transitive, the following lemma is easy to see.

Lemma 2. *The relationship \cong is reflexive and transitive.*

Remark 1. Analogously to Definition 8, one could define the notion of a *universally almost bounded system* which may take more than a polynomial number of steps in a negligible set of runs, i.e., there exists a polynomial p such that for every system $\mathcal{S} \in \text{Con}(\mathcal{E})$ it holds that

$$\text{Prob}[\text{Time}_{\mathcal{E}}((\mathcal{E} | \mathcal{S})(1^\eta, a)) > p(\eta + |a|)]$$

is negligible (as a function in η and a). But this does not make environmental systems more powerful. In particular, it would not change the notion of environmental indistinguishability because for every universally almost bounded system \mathcal{E} one could easily construct a universally strictly bounded system \mathcal{E}' such that $\mathcal{E} | \mathcal{S} \equiv \mathcal{E}' | \mathcal{S}$ for every system \mathcal{S} that can be connected to \mathcal{E} .

Remark 2. We note that the notion of a universally bounded system \mathcal{E} is equivalent to the following notion (stated informally): There exists a polynomial p such that for every sequence of incoming messages, \mathcal{E} “interacting” with this sequence of messages takes at most $p(\eta + |a|)$ transitions in mode **Compute**.

Remark 3. Environmental indistinguishability is defined with respect to one-bit output. That is, the probabilities of the environment outputs 1 (on tape decision) are compared. An alternative definition would be to require that the output (which may be more than a bit) produced by the environment when interacting with the systems is computationally indistinguishable, i.e., it cannot (with more than negligible probability) be distinguished by any polynomial time distinguisher. It is easy to show that this alternative definition is equivalent to the one-bit definition we use here, provided that $\text{start}, \text{decision} \notin \mathcal{T}(\mathcal{P}) = \mathcal{T}(\mathcal{Q})$: The idea is that in our version with one-bit output, the environment also plays the role of the distinguisher (see also remarks in [3]).

⁶Recall that $\text{Time}_{\mathcal{E}}((\mathcal{E} | \mathcal{S})(1^\eta, a))$ is the random variable that denotes the number of transitions taken by machines in \mathcal{E} in mode **Compute** in a run of $(\mathcal{E} | \mathcal{S})(1^\eta, a)$.

4.4 Protocols and Environmentally Bounded Systems

We now define the class of systems that run in polynomial time (except maybe with negligible probability) when combined with any environmental system. This is a very expressive class of systems and the one we will always consider in the context of universal composability. The runtime notion conceptually follows the one introduced in [16]. We also introduce protocol systems, which, as the name suggests, are systems which describe protocols; in the context of universal composability, these systems model ideal and real protocols.

Definition 10. A system \mathcal{S} is *environmentally (almost) bounded* if for every $\mathcal{E} \in \text{Env}(\mathcal{S})$: $\mathcal{E} | \mathcal{S}$ is almost bounded.

A system \mathcal{S} is *environmentally strictly bounded* if for every $\mathcal{E} \in \text{Env}(\mathcal{S})$: $\mathcal{E} | \mathcal{S}$ is strictly bounded.

Clearly, every environmentally strictly bounded system is environmentally bounded. The property of being environmentally bounded is typically easy to check (see also the discussion in Section 8). We also note that quantifying over all universally *almost* bounded systems \mathcal{E} would not have strengthened the notion of environmentally bounded systems (see Remark 1). Moreover, while the notion “environmentally bounded” talks about the overall runtime of the composed system $\mathcal{E} | \mathcal{S}$, we could equivalently have defined it by only restricting the runtime of \mathcal{S} , since \mathcal{E} is universally bounded anyway:

Lemma 3. *A system \mathcal{S} is environmentally bounded if and only if for every $\mathcal{E} \in \text{Env}(\mathcal{S})$ there exists a polynomial p such that $\text{Prob}[\text{Time}_{\mathcal{S}}((\mathcal{E} | \mathcal{S})(1^\eta, a)) > p(\eta + |a|)]$ is negligible (as a function in η and a); similarly for environmentally strictly bounded systems.*

We now define protocol systems.

Definition 11. A *protocol system* \mathcal{P} is a system such that (i) no tape in \mathcal{P} is named **start** or **decision** and (ii) for every IITM M occurring in \mathcal{P} such that M is not in the scope of a bang, we require that M accepts every incoming message in mode **CheckAddress**.

The motivation behind condition (ii) is that if M does not occur in the scope of a bang, then in every run of \mathcal{P} (in some context) there will be at most one copy of M . Hence, there is no reason to address different copies of M , and therefore, in mode **CheckAddress**, M should accept every incoming message. This condition will be used in the proofs of the composition theorems for unbounded self-composition.

In the above definition of protocol systems, we do not explicitly require that such systems are environmentally bounded—although this will typically be the case—in order to obtain more general results. For example, in the composition theorems we typically do not require the ideal protocol system to be environmentally bounded.

4.5 Properties of Systems

In this section, we summarize some fundamental and useful properties of systems. Some of these properties are not satisfied in some models for universal composability, causing partly severe problems (see also the discussion in Section 11).

The following lemma, which easily follows from the definition of systems, says that consistently changing the names of tapes in a system does not change the behavior of the system.

Lemma 4. *Let $\mathcal{S}_1, \dots, \mathcal{S}_k$ be connectable systems. Furthermore, for all $i \leq k$, let \mathcal{S}'_i be derived from \mathcal{S}_i by consistently (w.r.t. the other \mathcal{S}'_j) renaming external tapes, where, however, **start** and **decision** may not be renamed. Then,*

$$\mathcal{S}_1 | \dots | \mathcal{S}_k \equiv_0 \mathcal{S}'_1 | \dots | \mathcal{S}'_k$$

and

$$\mathcal{S}_1 | \dots | \mathcal{S}_k \text{ is almost/strictly bounded iff } \mathcal{S}'_1 | \dots | \mathcal{S}'_k \text{ is almost/strictly bounded.}$$

We now consider what we call a *dummy IITM* \mathcal{D} which simply forwards messages between entities: The dummy IITM has for all of its input tapes a corresponding output tape. The concrete set of input and output tapes that \mathcal{D} has depends on the entities between which \mathcal{D} is put. The dummy IITM accepts all messages on input tapes in mode **CheckAddress** and in mode **Compute** it simply copies a message received on an input tape to the corresponding output tape. We note that, except for the set of input and output tapes, \mathcal{D} does not depend on the entities between which it is put.

More precisely, let \mathcal{T}_{in} and \mathcal{T}_{out} be disjoint finite sets of tapes. Moreover, let $\mathcal{T}'_{in} = \{c' \mid c \in \mathcal{T}_{in}\}$ and $\mathcal{T}'_{out} = \{c' \mid c \in \mathcal{T}_{out}\}$ where c' is a new copy of c , i.e., a new tape with a new name.

We define

$$\mathcal{D} = \mathcal{D}(\mathcal{T}_{in}, \mathcal{T}_{out})$$

to be an IITM with input tapes $\mathcal{T}_{out} \cup \mathcal{T}'_{in}$ and output tapes $\mathcal{T}_{in} \cup \mathcal{T}'_{out}$. In mode **CheckAddress**, \mathcal{D} always accepts. In mode **Compute**, \mathcal{D} copies every message received on $c \in \mathcal{T}_{out}$ onto $c' \in \mathcal{T}'_{out}$ and every message received on $c' \in \mathcal{T}'_{in}$ onto $c \in \mathcal{T}_{in}$.

The following lemma says that the dummy IITM can be plugged between two systems without changing the behavior of the overall system. The proof of the lemma is straightforward.

Lemma 5. *Let \mathcal{P} and \mathcal{Q} be two connectable systems. Let $\mathcal{T}_{ext} = \mathcal{T}_{ext}(\mathcal{P}) \cap \mathcal{T}_{ext}(\mathcal{Q})$, $\mathcal{D} = \mathcal{D}(\mathcal{T}_{ext} \cap \mathcal{T}_{in}(\mathcal{P}), \mathcal{T}_{ext} \cap \mathcal{T}_{out}(\mathcal{P}))$, and \mathcal{Q}' be obtained from \mathcal{Q} by renaming all tapes c in \mathcal{T}_{ext} by c' . Then,*

$$\mathcal{P} \mid \mathcal{Q} \equiv_0 \mathcal{P} \mid \mathcal{D} \mid \mathcal{Q}'$$

and

$$\mathcal{P} \mid \mathcal{Q} \text{ is almost/strictly bounded iff } \mathcal{P} \mid \mathcal{D} \mid \mathcal{Q}' \text{ is almost/strictly bounded.}$$

While the above lemma is obvious and expected, it does not hold in all other models for universal composability, which is often very problematic. For instance, it does not hold in general in Canetti's UC model. However, the lemma is, for example, needed in order to prove that UC and dummy UC are equivalent security notions (see also Section 6.3), which in turn is needed to prove the composition theorem in the UC model (see also Section 11.1).

The following three lemmas also state fundamental properties of our general computational model and the runtime notions that we use. These properties are again essential for many general results, such as composition and joint state theorems. They are all about replacing a subsystem by a single IITM. Unfortunately, they are not satisfied in all other models for universal composability, causing severe problems in these models (see below and Section 11 for a discussion).

Lemma 6. *For every system \mathcal{S} there exists an IITM M such that the following conditions are satisfied:*

1. M and \mathcal{S} are compatible.
2. M accepts every message in mode **CheckAddress**.
3. For every system \mathcal{Q} that can be connected to \mathcal{S} :
 - (a) $\mathcal{S} \mid \mathcal{Q} \equiv_0 M \mid \mathcal{Q}$ and
 - (b) $\mathcal{S} \mid \mathcal{Q}$ is almost/strictly bounded if and only if $M \mid \mathcal{Q}$ is almost/strictly bounded.

Proof. We define an IITM M which simulates \mathcal{S} . The input and output tapes of M are the tapes in $\mathcal{T}_{in}(\mathcal{S})$ and $\mathcal{T}_{out}(\mathcal{S})$, respectively. The machine M accepts every message in mode **CheckAddress**. Hence, by construction, M is compatible with \mathcal{S} and accepts every message in mode **CheckAddress**. In mode **Compute**, M internally simulates the system \mathcal{S} according to the definition of running a system. In particular, M stores the configurations of all (previously) activated machines in \mathcal{S} . Upon input, M simulates the according machines in mode **CheckAddress** and in mode **Compute**. If, in mode **CheckAddress**, a message is rejected, M stops the current activation and does not produce output. Otherwise, the computation is continued in mode **Compute** (as usual, with the configuration of the simulated machine set to the one before the simulation of

the `CheckAddress` mode started.) If necessary, M creates new activated machines following the definition of running a system. If simulated machines produce internal output (i.e., output on internal tapes of \mathcal{S}), M continues the internal simulation. If some simulated machine produces external output (i.e., output on an external tape of \mathcal{S}), M simply produces this external output.

Clearly, M is an IITM because the runtime of M in mode `CheckAddress` is bounded by some polynomial (in fact, the runtime of M in mode `CheckAddress` is constant because M accepts every message). We note that the runtime of M in mode `Compute` might not be bounded but this only becomes relevant later; it is not a requirement for IITMs per se.

Since M perfectly simulates \mathcal{S} , we have

$$\mathcal{S} \mid \mathcal{Q} \equiv_0 M \mid \mathcal{Q}$$

for every system \mathcal{Q} that can be connected to \mathcal{S} .

Now, we prove 3. (b) and here we have to deal with the runtime of M in mode `Compute`. First, we note that \mathcal{Q} has exactly the same view in both systems because M perfectly simulates \mathcal{S} . In particular, we have that

$$\text{Time}_{\mathcal{Q}}((\mathcal{S} \mid \mathcal{Q})(1^\eta, a)) = \text{Time}_{\mathcal{Q}}((M \mid \mathcal{Q})(1^\eta, a)) \text{ , } ^7 \quad (1)$$

for all $\eta \in \mathbf{N}$ and $a \in \{0, 1\}^*$.

The implication from right to left in 3. (b) is easy to prove: Since M perfectly simulates in particular the mode `Compute` of all machines in \mathcal{S} (recall that $\text{Time}((\mathcal{S} \mid \mathcal{Q})(1^\eta, a))$ only counts the runtime in mode `Compute`), we have that $\text{Time}_{\mathcal{S}}((\mathcal{S} \mid \mathcal{Q})(1^\eta, a)) \leq \text{Time}_M((M \mid \mathcal{Q})(1^\eta, a))$, for all $\eta \in \mathbf{N}$ and $a \in \{0, 1\}^*$, and we obtain:

$$\begin{aligned} \text{Time}((\mathcal{S} \mid \mathcal{Q})(1^\eta, a)) &= \text{Time}_{\mathcal{S}}((\mathcal{S} \mid \mathcal{Q})(1^\eta, a)) + \text{Time}_{\mathcal{Q}}((\mathcal{S} \mid \mathcal{Q})(1^\eta, a)) \\ &\leq \text{Time}_M((M \mid \mathcal{Q})(1^\eta, a)) + \text{Time}_{\mathcal{Q}}((M \mid \mathcal{Q})(1^\eta, a)) \\ &= \text{Time}((M \mid \mathcal{Q})(1^\eta, a)) \end{aligned}$$

We conclude that $\mathcal{S} \mid \mathcal{Q}$ is almost/strictly bounded if $M \mid \mathcal{Q}$ is almost/strictly bounded.

To prove the direction from left to right, we assume that $\mathcal{S} \mid \mathcal{Q}$ is almost bounded (the case of strict boundedness is similar). We note that M , in mode `Compute`, also has to do all the maintenance work such as emptying all input and output tapes before simulated machines are activated in mode `CheckAddress` or `Compute`. So, to show that $M \mid \mathcal{Q}$ is almost bounded, we have to show that M can perform the simulation of the computations in mode `CheckAddress` and `Compute` and the maintenance work in polynomial time.

Since $\mathcal{S} \mid \mathcal{Q}$ is almost bounded, there exists a polynomial p such that

$$f(\eta, a) := \text{Prob}[\text{Time}((\mathcal{S} \mid \mathcal{Q})(1^\eta, a)) > p(\eta + |a|)] \quad (2)$$

is negligible (as a function in η and a). Hence, we have that the following holds in almost every run of $(\mathcal{S} \mid \mathcal{Q})(1^\eta, a)$ (i.e., except for a negligible set of runs): (i) the overall length of all messages that are sent (internally in \mathcal{S} and \mathcal{Q} and externally between \mathcal{S} and \mathcal{Q}) is bounded by $p(\eta + |a|)$, (ii) the number of activated instances of machines in \mathcal{S} is bounded by $p(\eta + |a|)$, and (iii) the size of all configurations (where only used random coins are counted) of activated instances of machines in \mathcal{S} is bounded by $p(\eta + |a|)$.

The runtime in mode `CheckAddress` of every IITM is polynomially bounded in the length of the configuration plus the security parameter plus the input message. Hence, because of (i), (ii), and (iii) and because M perfectly simulates \mathcal{S} , we conclude that the runtime of M in mode `Compute` is polynomially bounded. More precisely, there exists a polynomial p' such that

$$f'(\eta, a) := \text{Prob}[\text{Time}_M((M \mid \mathcal{Q})(1^\eta, a)) > p'(\eta + |a|)]$$

is negligible (as a function in η and a). From this, using (1) and (2), we conclude that for all $\eta \in \mathbf{N}$ and $a \in \{0, 1\}^*$:

$$\text{Prob}[\text{Time}((M \mid \mathcal{Q})(1^\eta, a)) > p(\eta + |a|) + p'(\eta + |a|)] \leq f(\eta, a) + f'(\eta, a) \text{ .}$$

Hence, $M \mid \mathcal{Q}$ is almost bounded. □

⁷That is, $\text{Time}_{\mathcal{Q}}((\mathcal{S} \mid \mathcal{Q})(1^\eta, a))(\alpha) = \text{Time}_{\mathcal{Q}}((M \mid \mathcal{Q})(1^\eta, a))(\alpha)$ for all random coins $\alpha \in \text{Rand}$.

The above lemma, which is also used in the proofs of the following two lemmas, is not satisfied in the UC model, which, as shown in [20] (see also the discussion in Section 11.1), causes the general joint state theorem to fail in the UC model.

The following lemma shows how one subsystem of an almost bounded system can be replaced by one universally bounded IITM (i.e., an environmental system which consists of a single IITM).

Lemma 7. *Let \mathcal{S} and \mathcal{Q} be connectable systems such that $\mathcal{S} \mid \mathcal{Q}$ is almost bounded and $\mathbf{start} \in \mathcal{T}(\mathcal{S})$ (i.e., \mathcal{S} contains a master IITM, and hence, \mathcal{Q} does not). Then there exists an IITM M such that the following conditions are satisfied:*

1. M and \mathcal{S} are compatible.
2. M accepts every message in mode `CheckAddress`.
3. M is universally bounded.
4. $M \mid \mathcal{Q}$ is almost bounded.
5. $M \mid \mathcal{Q} \equiv \mathcal{S} \mid \mathcal{Q}$.

Proof. By Lemma 6, we may assume that $\mathcal{S} = M'$ is a single IITM that accepts every message in mode `CheckAddress`. We define an IITM M which simulates \mathcal{S} . The input and output tapes of M are the tapes in $\mathcal{T}_{in}(\mathcal{S})$ and $\mathcal{T}_{out}(\mathcal{S})$, respectively.

Before we specify how M works in mode `Compute`, observe that, because $\mathcal{S} \mid \mathcal{Q}$ is almost bounded, there exists a polynomial p such that

$$f(\eta, a) := \text{Prob}[\text{Time}((\mathcal{S} \mid \mathcal{Q})(1^\eta, a)) > p(\eta + |a|)] \quad (3)$$

is negligible (as a function in η and a).

We now define M to simulate the system \mathcal{S} as follows: If invoked in mode `CheckAddress`, M will, just as \mathcal{S} , accept every message. In mode `Compute`, M will simulate \mathcal{S} in mode `Compute` where, however, not more than $p(\eta + |a|)$ transitions of the IITM \mathcal{S} are simulated overall (note that M knows $|a|$ because $\mathbf{start} \in \mathcal{T}(\mathcal{S})$, and hence, $\mathbf{start} \in \mathcal{T}_{in}(M)$). If this bound is reached, M goes into some final state from which M cannot take any transition. In particular, from that point on, in mode `Compute`, M ignores all incoming messages, does not make any further step, and does not produce output.

By construction, M is compatible with \mathcal{S} and it is easy to see that M is universally bounded. It remains to show that $M \mid \mathcal{Q}$ is almost bounded and that $M \mid \mathcal{Q} \equiv \mathcal{S} \mid \mathcal{Q}$. For this purpose, let $B(1^\eta, a)$ denote the event that in a run of $(\mathcal{S} \mid \mathcal{Q})(1^\eta, a)$ the system \mathcal{S} takes more than $p(\eta + |a|)$ transitions (in mode `Compute`), i.e., $B(1^\eta, a) = \{\alpha \in \text{Rand} \mid \text{Time}_{\mathcal{S}}((\mathcal{S} \mid \mathcal{Q})(1^\eta, a))(\alpha) > p(\eta + |a|)\}$.⁸ By (3), we have that $\text{Prob}[B(1^\eta, a)] \leq f(\eta, a)$. Let $B'(1^\eta, a)$ denote the event that in a run of $(M \mid \mathcal{Q})(1^\eta, a)$ the bound $p(\eta + |a|)$ is reached. By $\overline{B(1^\eta, a)}$ and $\overline{B'(1^\eta, a)}$, we denote the complement of event $B(1^\eta, a)$ and $B'(1^\eta, a)$, respectively.

Now, since M perfectly simulates the machine \mathcal{S} in mode `CheckAddress` and in mode `Compute` until $B'(1^\eta, a)$ occurs, we conclude that:

$$\text{Prob}[B'(1^\eta, a)] = \text{Prob}[B(1^\eta, a)] \quad (4)$$

$$\text{Prob}[\text{Time}_{\mathcal{Q}}((M \mid \mathcal{Q})(1^\eta, a)) > p(\eta + |a|) \wedge \overline{B'(1^\eta, a)}] = \text{Prob}[\text{Time}_{\mathcal{Q}}((\mathcal{S} \mid \mathcal{Q})(1^\eta, a)) > p(\eta + |a|) \wedge \overline{B(1^\eta, a)}] \quad (5)$$

$$\text{Prob}[(M \mid \mathcal{Q})(1^\eta, a) = 1 \wedge \overline{B'(1^\eta, a)}] = \text{Prob}[(\mathcal{S} \mid \mathcal{Q})(1^\eta, a) = 1 \wedge \overline{B(1^\eta, a)}] . \quad (6)$$

⁸Recall that we can identify a run ρ with security parameter η and external input a with the random coins $\alpha \in \text{Rand}$ used in this run since a run is uniquely determined by the random coins.

From this, we obtain:

$$\begin{aligned}
& \text{Prob}[\text{Time}_{\mathcal{Q}}((M \mid \mathcal{Q})(1^\eta, a)) > p(\eta + |a|)] \\
& \leq \text{Prob}[B'(1^\eta, a)] + \text{Prob}[\text{Time}_{\mathcal{Q}}((M \mid \mathcal{Q})(1^\eta, a)) > p(\eta + |a|) \wedge \overline{B'(1^\eta, a)}] \\
& \stackrel{(4),(5)}{=} \text{Prob}[B(1^\eta, a)] + \text{Prob}[\text{Time}_{\mathcal{Q}}((\mathcal{S} \mid \mathcal{Q})(1^\eta, a)) > p(\eta + |a|) \wedge \overline{B(1^\eta, a)}] \\
& \leq \text{Prob}[B(1^\eta, a)] + \text{Prob}[\text{Time}_{\mathcal{Q}}((\mathcal{S} \mid \mathcal{Q})(1^\eta, a)) > p(\eta + |a|)] \\
& \stackrel{(3)}{\leq} 2 \cdot f(\eta, a) .
\end{aligned}$$

Hence, since M is universally bounded, $M \mid \mathcal{Q}$ is almost bounded. Furthermore, we obtain:

$$\begin{aligned}
& |\text{Prob}[(M \mid \mathcal{Q})(1^\eta, a) = 1] - \text{Prob}[(\mathcal{S} \mid \mathcal{Q})(1^\eta, a) = 1]| \\
& \stackrel{(6)}{=} |\text{Prob}[(M \mid \mathcal{Q})(1^\eta, a) = 1 \wedge B'(1^\eta, a)] - \text{Prob}[(\mathcal{S} \mid \mathcal{Q})(1^\eta, a) = 1 \wedge B(1^\eta, a)]| \\
& \stackrel{(4)}{\leq} \text{Prob}[B(1^\eta, a)] \\
& \stackrel{(3)}{\leq} f(\eta, a) .
\end{aligned}$$

Hence, $M \mid \mathcal{Q} \equiv \mathcal{S} \mid \mathcal{Q}$. □

Definition 12. We denote the IITM M as constructed in the proof of Lemma 7 by $[\mathcal{S}]_{\mathcal{Q}}$.

The following lemma will allow us to “open” $[\mathcal{S}]_{\mathcal{Q}}$, i.e., replace $[\mathcal{S}]_{\mathcal{Q}}$ by \mathcal{S} , in a context different from \mathcal{Q} .

Lemma 8. *Let \mathcal{S} and \mathcal{Q} be connectable systems such that $\mathcal{S} \mid \mathcal{Q}$ is almost bounded, $\text{start} \in \mathcal{T}(\mathcal{S})$ (i.e., \mathcal{S} contains a master IITM), and $\text{decision} \notin \mathcal{T}(\mathcal{Q})$. Furthermore, let \mathcal{Q}' be a system which is compatible with \mathcal{Q} and satisfies the following condition: $\mathcal{E} \mid \mathcal{Q} \equiv \mathcal{E} \mid \mathcal{Q}'$ for every $\mathcal{E} \in \text{Env}(\mathcal{Q})$ such that $\mathcal{E} \mid \mathcal{Q}$ is almost bounded. Then,*

$$[\mathcal{S}]_{\mathcal{Q}} \mid \mathcal{Q}' \equiv \mathcal{S} \mid \mathcal{Q}' .$$

Moreover, if $[\mathcal{S}]_{\mathcal{Q}} \mid \mathcal{Q}'$ is almost bounded, then $\mathcal{S} \mid \mathcal{Q}'$ is almost bounded too.

Proof. Just as in the proof of Lemma 7, by Lemma 6 we may assume that \mathcal{S} is a single IITM that accepts every message in mode `CheckAddress`. Now, recall that by definition, $[\mathcal{S}]_{\mathcal{Q}}$ exactly simulates all transitions of \mathcal{S} up to a certain polynomial bound and that when running $[\mathcal{S}]_{\mathcal{Q}} \mid \mathcal{Q}$ this bound is reached with only negligible probability. It follows that the probability that this bound is reached when running the system $[\mathcal{S}]_{\mathcal{Q}} \mid \mathcal{Q}'$ is negligible as well. Otherwise, one could easily construct a universally bounded system \mathcal{E} such that \mathcal{E} can be connected to \mathcal{Q} (i.e., $\mathcal{E} \in \text{Env}(\mathcal{Q})$) and $\mathcal{E} \mid \mathcal{Q} \not\equiv \mathcal{E} \mid \mathcal{Q}'$, in contradiction to the assumption in Lemma 8: The system \mathcal{E} is defined to simulate $[\mathcal{S}]_{\mathcal{Q}}$ and output 1 on `decision` if and only if the bound is reached. (If in the simulation of $[\mathcal{S}]_{\mathcal{Q}}$ output shall be written on `decision` before the bound is reached, \mathcal{E} writes 0 on `decision` and terminates.) Since $\text{decision} \notin \mathcal{T}(\mathcal{Q})$, 1 is output on `decision` in runs of $\mathcal{E} \mid \mathcal{Q}$ and $\mathcal{E} \mid \mathcal{Q}'$, respectively, if and only if the bound is reached.

It follows that with overwhelming probability $[\mathcal{S}]_{\mathcal{Q}}$ exactly simulates \mathcal{S} in the system $[\mathcal{S}]_{\mathcal{Q}} \mid \mathcal{Q}'$. Thus, we obtain $[\mathcal{S}]_{\mathcal{Q}} \mid \mathcal{Q}' \equiv \mathcal{S} \mid \mathcal{Q}'$. Similarly, it is easy to see that if $[\mathcal{S}]_{\mathcal{Q}} \mid \mathcal{Q}'$ is almost bounded, then $\mathcal{S} \mid \mathcal{Q}'$ is almost bounded too. □

We note that in the above lemma the condition that \mathcal{Q} and \mathcal{Q}' are indistinguishable for all environments such that $\mathcal{E} \mid \mathcal{Q}$ are almost bounded is, although sufficient for our purposes, much stronger than what is needed for this lemma. As can be seen from the proof, it suffices to require that $[\mathcal{S}]_{\mathcal{Q}}^{\text{alert}} \mid \mathcal{Q} \equiv [\mathcal{S}]_{\mathcal{Q}}^{\text{alert}} \mid \mathcal{Q}'$ where $[\mathcal{S}]_{\mathcal{Q}}^{\text{alert}}$ is the environment \mathcal{E} constructed in the proof, i.e., it simulates $[\mathcal{S}]_{\mathcal{Q}}$ and outputs 1 on `decision` if and only if the bound is reached.

5 Composition Theorems for Environmental Indistinguishability

We now prove general composition theorems for environmental indistinguishability (\cong). They are the core of the composition theorems for universal composability, which are presented in Section 7. In fact, the composition theorems for universal composability security are merely corollaries of the theorems presented in this section.

In a nutshell, these theorems say that the concurrent composition of environmentally indistinguishable systems is environmentally indistinguishable. In particular, these theorems can be used to establish environmental indistinguishability in a modular way: it suffices to show environmental indistinguishability for subsystems in order to conclude that the concurrent composition of the subsystems are environmentally indistinguishable as well.

We first, in Section 5.1, present a composition theorem for the composition of a constant number of, possibly different, systems. We then state a theorem which says that if two systems are environmentally indistinguishable, then they are environmentally indistinguishable even if an unbounded number of concurrent sessions/copies of these systems may be executed (unbounded self-composition). We actually prove two versions of this theorem, one in which the systems are not aware of the session identifiers that are used to address the different sessions (see Section 5.2) and one in which they are aware of the session identifiers (see Section 5.3), where the latter theorem is a corollary of the former one. The theorems for a constant number of protocols and for unbounded self-composition can freely be combined to establish, in a modular way, environmental indistinguishability of more and more complex systems.

5.1 Composition Theorem for a Constant Number of Systems

We now present the composition theorem for the composition of a constant number of (possibly different) systems. The proof of this theorem uses merely the equational principles (Lemma 7 and 8) established in Section 4.5.

Theorem 4. *Let $k \geq 1$ and $\mathcal{S}, \mathcal{P}_1, \dots, \mathcal{P}_k, \mathcal{Q}_1, \dots, \mathcal{Q}_k$ be protocol systems⁹ such that the following conditions are satisfied:*

1. *For all $j \leq k$: \mathcal{P}_j and \mathcal{Q}_j are environmentally bounded and $\mathcal{P}_j \cong \mathcal{Q}_j$.*
2. *$\mathcal{S}, \mathcal{P}_1, \dots, \mathcal{P}_k$ are connectable (hence, $\mathcal{S}, \mathcal{Q}_1, \dots, \mathcal{Q}_k$ are connectable) and $\mathcal{S} | \mathcal{P}_1 | \dots | \mathcal{P}_k$ is environmentally bounded.*

Then, $\mathcal{S} | \mathcal{P}_1 | \dots | \mathcal{P}_k \cong \mathcal{S} | \mathcal{Q}_1 | \dots | \mathcal{Q}_k$ and $\mathcal{S} | \mathcal{Q}_1 | \dots | \mathcal{Q}_k$ is environmentally bounded.

Before we prove this theorem, we make some remarks.

Remark 4. We note that the system \mathcal{S} in the above theorem could be a simple forwarder or even a system without external tapes, so the theorem also holds if \mathcal{S} is omitted. That is, if $\mathcal{P}_1 | \dots | \mathcal{P}_k$ is environmentally bounded and the assumptions of the theorem hold, then $\mathcal{P}_1 | \dots | \mathcal{P}_k \cong \mathcal{Q}_1 | \dots | \mathcal{Q}_k$ and $\mathcal{Q}_1 | \dots | \mathcal{Q}_k$ is environmentally bounded.

Remark 5. We note that Condition 2. in the above theorem is easy to check. Connectability is a simple syntactic condition. Moreover, for typical applications, it is easy to check whether a system is environmentally bounded (see also the discussion in Section 8, in particular Lemma 15).

Proof of Theorem 4. We first show the theorem for the case $k = 1$. The proof of this case is depicted in Figure 2. For any $k \geq 1$, the theorem follows by applying the same argument iteratively, see below.

Assume that $k = 1$ and let $\mathcal{E} \in \text{Env}(\mathcal{S} | \mathcal{P}_1) = \text{Env}(\mathcal{S} | \mathcal{Q}_1)$. We may assume that $\text{start} \in \mathcal{T}(\mathcal{E} | \mathcal{S} | \mathcal{P}_1) = \mathcal{T}(\mathcal{E} | \mathcal{S} | \mathcal{Q}_1)$. Otherwise, there would not exist a master IITM, and hence, a run would always directly

⁹We note that a stronger variant of this theorem holds where $\mathcal{S}, \mathcal{P}_1, \dots, \mathcal{P}_k, \mathcal{Q}_1, \dots, \mathcal{Q}_k$ are arbitrary systems (not necessarily protocol systems) such that $\text{start}, \text{decision} \notin \mathcal{T}(\mathcal{P}_j)$ and $\text{start}, \text{decision} \notin \mathcal{T}(\mathcal{Q}_j)$ for all $j \leq k$.

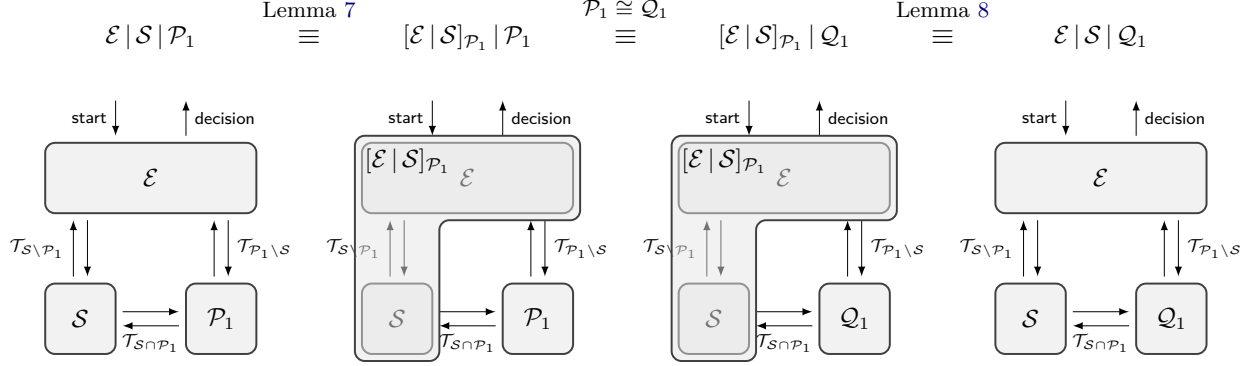


Figure 2: The proof of Theorem 4 for the case $k = 1$. We use the following abbreviations: $\mathcal{T}_{S \setminus P_1} := \mathcal{T}_{ext}(S) \setminus \mathcal{T}_{ext}(P_1)$, $\mathcal{T}_{P_1 \setminus S} := \mathcal{T}_{ext}(P_1) \setminus \mathcal{T}_{ext}(S)$, and $\mathcal{T}_{S \cap P_1} := \mathcal{T}_{ext}(S) \cap \mathcal{T}_{ext}(P_1)$.

terminate with empty overall output. Since $S | P_1$ is environmentally bounded, $\mathcal{E} | S | P_1$ is almost bounded and we have that:

$$\begin{aligned}
\mathcal{E} | S | P_1 &\equiv [\mathcal{E} | S]_{P_1} | P_1 && \text{(Lemma 7)} \\
&\equiv [\mathcal{E} | S]_{P_1} | Q_1 && (P_1 \cong Q_1) \\
&\equiv \mathcal{E} | S | Q_1 && (P_1 \cong Q_1, \text{ Lemma 8}).
\end{aligned}$$

We note that, in order to apply Lemmas 7 and 8, we need that $\text{start} \in \mathcal{T}(\mathcal{E} | S)$ and $\text{decision} \notin \mathcal{T}(P_1) = \mathcal{T}(Q_1)$, which is true because P_1 and Q_1 are protocol systems. Moreover, since Q_1 is environmentally bounded and $[\mathcal{E} | S]_{P_1}$ is universally bounded, $[\mathcal{E} | S]_{P_1} | Q_1$ is almost bounded and, by Lemma 8, we conclude that $\mathcal{E} | S | Q_1$ is almost bounded. Since this holds for all $\mathcal{E} \in \text{Env}(S | P_1) = \text{Env}(S | Q_1)$, we conclude that $S | P_1 \cong S | Q_1$ and $S | Q_1$ is environmentally bounded, which proves the theorem for $k = 1$.

We now prove the theorem for any $k \geq 1$. For every $r \leq k$, we define the r -th hybrid system:

$$\mathcal{H}_r := S | Q_1 | \cdots | Q_{r-1} | P_{r+1} | \cdots | P_k ,$$

which can be connected to \mathcal{P}_r or \mathcal{Q}_r . By applying the case “ $k = 1$ ” k times, we obtain that:

$$\begin{aligned}
S | P_1 | \cdots | P_k &= \mathcal{H}_1 | P_1 && \text{(syntactic reordering of subsystems)} \\
&\cong \mathcal{H}_1 | Q_1 && \text{(case “} k = 1 \text{”, with } \mathcal{H}_1 \text{ playing the role of } S) \\
&= \mathcal{H}_2 | P_2 && \text{(syntactic reordering of subsystems)} \\
&\cong \mathcal{H}_2 | Q_2 && \text{(case “} k = 1 \text{”)} \\
&\vdots \\
&= \mathcal{H}_k | P_k && \text{(syntactic reordering of subsystems)} \\
&\cong \mathcal{H}_k | Q_k && \text{(case “} k = 1 \text{”)} \\
&= S | Q_1 | \cdots | Q_k .
\end{aligned}$$

We note that in the above argument, the equalities ($=$) hold true up to syntactic reordering of subsystems. Furthermore, all these systems are environmentally bounded: By assumption $S | P_1 | \cdots | P_k = \mathcal{H}_1 | P_1$ is environmentally bounded. For all $\mathcal{H}_r | Q_r$, $r \geq 1$, this follows from the case “ $k = 1$ ” and that $\mathcal{H}_r | P_r = \mathcal{H}_{r-1} | Q_{r-1}$ is environmentally bounded. This concludes the proof of the theorem. \square

5.2 Composition Theorem for Unbounded Self-Composition

We now present the composition theorem for the unbounded self-composition of systems that are not aware of the session identifiers used to address the different sessions. As mentioned before, below we will present a composition theorem, as a corollary of this composition theorem, where systems are aware of the session identifiers.

We first need to introduce the notion of a session version of a system.

5.2.1 Session Versions of Systems

To model multiple copies/sessions of a system one can use the mode `CheckAddress` of IITMs combined with session identifiers (SIDs) as follows.

Given a system \mathcal{Q} , we define its session version $\underline{\mathcal{Q}}$, which allows us to address different copies of (IITMs occurring in) \mathcal{Q} by a particular SID. We first define the session version of a single IITM.

The *session version* \underline{M} of an IITM M is an IITM that exactly simulates M except that all messages received have to be prefixed by a particular SID (i.e., in mode `CheckAddress` the IITM \underline{M} will reject all messages not prefixed by the particular SID) and all messages sent out are prefixed by this SID. The SID \underline{M} will use is the one with which \underline{M} is first activated (hence, in the first activation, \underline{M} will accept the incoming message in mode `CheckAddress` and then store the SID). More precisely, \underline{M} behaves as follows in mode `CheckAddress` and `Compute`, respectively:

- When activated in mode `CheckAddress`, \underline{M} does the following: If \underline{M} has never been activated before, it accepts an incoming message m' only if the following is satisfied: (i) m' is of the form (s, m) ¹⁰ where s is interpreted as an SID, and (ii) the simulated M accepts m in mode `CheckAddress`. (In mode `Compute`, s will be stored by \underline{M} .) If \underline{M} was activated before, then \underline{M} will accept an incoming message m' only if the following is satisfied: (i) m' is of the form (s, m) where s is the SID that was stored in the first activation (in mode `Compute`), and (ii) m is accepted by the simulated M in mode `CheckAddress`.
- When activated in mode `Compute`, \underline{M} does the following: If \underline{M} has never been activated before (in mode `Compute`), then by the definition of \underline{M} in mode `CheckAddress` it follows that the incoming message is of the form (s, m) . Now, \underline{M} first stores s and then simulates M on input m in mode `Compute`. If M produces output, say m' , then \underline{M} sends the output (s, m') , i.e., prefixes m' with s . If \underline{M} was activated before (in mode `Compute`), then by definition of \underline{M} in mode `CheckAddress` it follows that the incoming message is of the form (s, m) where s is the SID that was stored in the first activation. Now, as before, \underline{M} simulates M on input m in mode `Compute` and prefixes the output produced (if any) with s .

Note that the IITM M , simulated within \underline{M} , is not aware of the SID that is used to address (a copy of) M . In particular, it cannot use the SID as part of the messages it produces.

Now, the *session version* $\underline{\mathcal{Q}}$ of a system \mathcal{Q} is obtained from \mathcal{Q} by replacing every IITM occurring in \mathcal{Q} by its session version. For instance, if $\mathcal{Q} = M_1 \mid M_2$, then $\underline{\mathcal{Q}} = \underline{M}_1 \mid \underline{M}_2$.

While in session versions as defined above, SIDs are always prefixed to messages, it clearly does not matter where exactly the SIDs occur in a message. The `CheckAddress` mode, being an arbitrary deterministic polynomial computation, is flexible enough to allow for many variants; for example, SIDs could be appended instead of prefixed to messages. The results on session versions do not depend on specific details.

It is sometimes useful to define a session version of a single IITM or an entire system with respect to a domain of SIDs (parameterized by the security parameter). In mode `CheckAddress` only those SIDs would be accepted that belong to the specified domain. (Clearly, we need to require that it can be decided efficiently in mode `CheckAddress` whether an SID belongs to the domain.) With this, we could, for example, define a session version \underline{M} of an IITM M which only accepts SIDs of the form (sid, pid) , where pid denotes a party and sid identifies the session in which this party runs. Hence, in a run of the system \underline{M} (in some

¹⁰We assume that (s, m) is a bit string which encodes the concatenation of the bit strings s and m in such a way that given s and m , the bit string (s, m) is computable in polynomial time (in $|s| + |m|$), and that given (s, m) the components s and m can be retrieved in polynomial time.

environment) all instances of \underline{M} would have SIDs of this form. In statements involving session versions, such as composition theorems, details of how the domains of SIDs are chosen are typically not important, as long as they are chosen consistently. We therefore omit such details in the statements.

We now state a fundamental property about the runtime of session versions of two indistinguishable systems. While this property might be expected, the proof of this theorem, which is presented in Appendix A and mainly follows ideas presented in [16], is non-trivial.

Lemma 9. *Let \mathcal{P} and \mathcal{Q} be protocol systems such that $!\underline{\mathcal{P}}$ and \mathcal{Q} are environmentally bounded and $\mathcal{P} \cong \mathcal{Q}$. Then, $!\underline{\mathcal{Q}}$ is environmentally bounded.*

5.2.2 The Composition Theorem for Session Versions

The following theorem says that if two systems are environmentally indistinguishable, then an unbounded number of copies of one system are environmentally indistinguishable from an unbounded number of copies of the other system, where to address different copies session versions are considered. In particular, SIDs are merely used as a means to address (IITMs belonging to) copies of systems. The systems (more precisely, the IITMs these systems consist of) are not and do not need to be aware of the SIDs that are used to address their copies, and the specific addressing mechanism used. As mentioned before, in Section 5.3 we will present a composition theorem, as a corollary of Theorem 5, where systems are aware of their own SID.

Theorem 5. *Let \mathcal{P} and \mathcal{Q} be protocol systems such that $!\underline{\mathcal{P}}$ and \mathcal{Q} are environmentally bounded and $\mathcal{P} \cong \mathcal{Q}$. Then, $!\underline{\mathcal{P}} \cong !\underline{\mathcal{Q}}$ and $!\underline{\mathcal{Q}}$ is environmentally bounded*

Before we prove this theorem, let us mention that typically \mathcal{P} and \mathcal{Q} will be environmentally *strictly* bounded and that this is easy to verify. By Lemma 17 it then follows immediately that $!\underline{\mathcal{P}}$ is environmentally strictly bounded as well, and hence, environmentally bounded. So Theorem 5 can be applied directly to typical protocol systems without additional effort.

Proof of Theorem 5. By Lemma 9, we know that $!\underline{\mathcal{Q}}$ is environmentally bounded.

Now, let $\mathcal{E} \in \text{Env}(!\underline{\mathcal{P}}) = \text{Env}(!\underline{\mathcal{Q}})$. To prove the theorem, it remains to show that $\mathcal{E} | !\underline{\mathcal{P}} \equiv \mathcal{E} | !\underline{\mathcal{Q}}$. By Lemma 7, we may assume that \mathcal{E} is a single IITM which, in mode `CheckAddress`, accepts all messages. Moreover, we may assume, without loss of generality, that `start`, `decision` $\in \mathcal{T}(\mathcal{E})$ and that \mathcal{E} is such that every message m that \mathcal{E} outputs (except if m is output on tape `decision`) is prefixed by some SID, i.e., $m = (s, m')$ for some bit strings s and m' :¹¹ since \mathcal{E} will only interact with session versions, message not of the form (s, m') would be rejected by these session versions anyway.

Since \mathcal{E} is universally bounded, it follows that there exists a polynomial $p_{\mathcal{E}}$ such that the number of different sessions (i.e., messages with distinct SIDs output by \mathcal{E}) is bounded from above by $p_{\mathcal{E}}(\eta + |a|)$ (where η is the security parameter and a is the external input). In particular, for every η, a and every run of $(\mathcal{E} | !\underline{\mathcal{P}})(1^\eta, a)$ or $(\mathcal{E} | !\underline{\mathcal{Q}})(1^\eta, a)$, there exist at most $p_{\mathcal{E}}(\eta + |a|)$ copies of $\underline{\mathcal{Q}}$ and $\underline{\mathcal{P}}$ in such a run.

The proof proceeds by a hybrid argument. In what follows, let $\underline{\mathcal{P}}'$ be the variant of $\underline{\mathcal{P}}$ obtained from $\underline{\mathcal{P}}$ by renaming every tape c occurring in $\underline{\mathcal{P}}$ as c' . Analogously, let $\underline{\mathcal{P}}''$ be obtained from $\underline{\mathcal{P}}$ by renaming every tape c occurring in $\underline{\mathcal{P}}$ as c'' . Similarly for $\underline{\mathcal{Q}}'$ and $\underline{\mathcal{Q}}''$.

We define an IITM \mathcal{E}_r (for every $r \in \mathbf{N}$) which basically simulates \mathcal{E} and which will run in the system $\mathcal{E}_r | !\underline{\mathcal{P}}'' | !\underline{\mathcal{Q}}' | \mathcal{P}$ or $\mathcal{E}_r | !\underline{\mathcal{P}}'' | !\underline{\mathcal{Q}}' | \mathcal{Q}$, respectively. The first $r - 1$ copies of the protocol invoked by \mathcal{E} will be copies of $\underline{\mathcal{Q}}'$, the r -th copy will be the external system \mathcal{P} or \mathcal{Q} , respectively, and the remaining copies will be copies of $\underline{\mathcal{P}}''$.

Formally, \mathcal{E}_r is obtained from \mathcal{E} as follows. (Recall that we assume that \mathcal{E} is a single IITM which accepts every message in mode `CheckAddress`). The IITM \mathcal{E}_r will always accept in mode `CheckAddress`. The behavior of \mathcal{E}_r in mode `Compute` is specified next.

First, we need to make sure that \mathcal{E}_r has the appropriate tapes to connect to the different entities. The IITM \mathcal{E} already has tapes to connect to the external tapes of \mathcal{P} and \mathcal{Q} . For each such tape c , we add to \mathcal{E}_r a tape c' and c'' to connect to the external tapes of $\underline{\mathcal{Q}}'$ and $\underline{\mathcal{P}}''$, respectively.

¹¹More formally, for every system $\mathcal{S} \in \text{Con}(\mathcal{E})$ and all parameters η, a in every run of $(\mathcal{E} | \mathcal{S})(1^\eta, a)$ the system \mathcal{E} should only output messages of the described form.

Next, we need to specify how \mathcal{E}_r redirects protocol invocations of \mathcal{E} in the way described above: \mathcal{E}_r keeps a list L of SIDs, which initially is empty, and the length l of the list, which initially is 0. By definition of $p_{\mathcal{E}}$, it will always hold that $l \leq p_{\mathcal{E}}(\eta + |a|)$. In the first activation with security parameter $\eta \in \mathbf{N}$ and external input $a \in \{0, 1\}^*$, \mathcal{E}_r starts a simulation of \mathcal{E} with security parameter η and external input a . In particular, if \mathcal{E} produces output, then so does \mathcal{E}_r , and if \mathcal{E}_r receives input, then \mathcal{E} is simulated with this input. However, as explained next, the behavior of \mathcal{E}_r deviates from that of \mathcal{E} when it comes to sending and receiving messages to the different copies of protocols.

1. If \mathcal{E} produces output m on some external output tape c to $\underline{\mathcal{P}}$ (and hence, $\underline{\mathcal{Q}}$) prefixed with s , then \mathcal{E}_r checks whether s occurs in L . If s does not occur in L , s is appended at the end of L and l is increased by 1. Let $j \in \{1, \dots, l\}$ be the position where s occurs in L .
 - (a) If $j < r$, then \mathcal{E}_r writes m on tape c' .
 - (b) If $j = r$, then \mathcal{E}_r outputs m' on c where m' is a message such that $m = (s, m')$, i.e., s is removed from m .
 - (c) If $j > r$, then \mathcal{E}_r writes m on tape c'' .
2. If \mathcal{E}_r receives input on tape c'' where c'' is an external tape of $\underline{\mathcal{P}}''$ corresponding to an external tape c of $\underline{\mathcal{P}}$, then \mathcal{E}_r behaves as \mathcal{E} in case input was received on tape c .
3. If \mathcal{E}_r receives input on tape c' where c' is an external tape of $\underline{\mathcal{Q}}'$ corresponding to an external tape c of $\underline{\mathcal{Q}}$, then \mathcal{E}_r behaves as \mathcal{E} in case input was received on tape c .
4. If \mathcal{E}_r receives input m on tape c where c is an external tape of \mathcal{P} (and hence, \mathcal{Q}), then \mathcal{E}_r behaves as \mathcal{E} in case input $(L[r], m)$ was received on tape c where $L[r]$ denotes the r -th entry of L . By construction, this entry exists in L (i.e., $r \leq l$) since \mathcal{E} must have invoked the r -th copy.

We also consider a variant \mathcal{E}_{\S} of \mathcal{E}_r which in the first activation chooses $r \in \{1, \dots, p_{\mathcal{E}}(\eta + |a|)\}$ uniformly at random and then behaves exactly like \mathcal{E}_r . Moreover, we consider the variant $\mathcal{E}_{p_{\mathcal{E}}}$ of \mathcal{E}_{\S} which uses $r = p_{\mathcal{E}}(\eta + |a|)$.

We define the following hybrid systems, for every $r \in \mathbf{N}$:

$$\begin{aligned} \mathcal{H}_r &:= \mathcal{E}_r \mid \underline{\mathcal{P}}'' \mid \underline{\mathcal{Q}}' \text{ and} \\ \mathcal{H}_{p_{\mathcal{E}}} &:= \mathcal{E}_{p_{\mathcal{E}}} \mid \underline{\mathcal{P}}'' \mid \underline{\mathcal{Q}}' \text{ and} \\ \mathcal{H}_{\S} &:= \mathcal{E}_{\S} \mid \underline{\mathcal{P}}'' \mid \underline{\mathcal{Q}}' , \end{aligned}$$

which can be connected to \mathcal{P} (and hence \mathcal{Q}). The system \mathcal{H}_r is illustrated in Figure 3.

By construction, for every $r \in \mathbf{N}$, the systems $\mathcal{E} \mid \underline{\mathcal{P}}$ and $\mathcal{H}_1 \mid \mathcal{P}$, $\mathcal{H}_r \mid \mathcal{Q}$ and $\mathcal{H}_{r+1} \mid \mathcal{P}$, and $\mathcal{E} \mid \underline{\mathcal{Q}}$ and $\mathcal{H}_{p_{\mathcal{E}}} \mid \mathcal{Q}$, respectively, behave exactly the same. In particular, for all $r \in \mathbf{N}$, we have that:

$$\mathcal{E} \mid \underline{\mathcal{P}} \equiv_0 \mathcal{H}_1 \mid \mathcal{P} , \tag{7}$$

$$\mathcal{H}_r \mid \mathcal{Q} \equiv_0 \mathcal{H}_{r+1} \mid \mathcal{P} , \text{ and} \tag{8}$$

$$\mathcal{E} \mid \underline{\mathcal{Q}} \equiv_0 \mathcal{H}_{p_{\mathcal{E}}} \mid \mathcal{Q} . \tag{9}$$

For (7) we use that \mathcal{P} is a protocol system. In particular, we use property (ii) of protocol systems (see Definition 11). If this property were not satisfied, i.e., \mathcal{P} contains an IITM M which is not in the scope of a bang but which could reject a message in mode `CheckAddress`, the following could happen. In a run of $(\mathcal{H}_1 \mid \mathcal{P})(1^\eta, a)$ a message is sent to M , but it is rejected by M (in mode `CheckAddress`). Then, since M is not in the scope of a bang, no new copy of M will be generated. Conversely, if in a run of $\mathcal{E} \mid \underline{\mathcal{P}}$ a message is sent to a copy of the session version \underline{M} of M prefixed with a SID generated by \mathcal{E} and the simulated M in \underline{M} would reject the message, then it could happen that a new copy of \underline{M} is generated (since \underline{M} is in the scope of a bang in $\mathcal{E} \mid \underline{\mathcal{P}}$) which then would not have a corresponding entity in a run of the system $(\mathcal{H}_1 \mid \mathcal{P})(1^\eta, a)$. In short, by property (ii) of protocol systems it is guaranteed that for IITMs that do not occur in the scope of a




Figure 3: The hybrid system $\mathcal{H}_r = \mathcal{E}_r | !\underline{Q}' | !\underline{P}''$ composed with the external session \underline{Q} or \underline{P} , respectively.

bang in \underline{P} only at most one copy is generated per SID in the run of $\mathcal{E} | !\underline{P}$. A similar argument is used to prove (8) and (9).

We now show that $!\underline{P}'' | !\underline{Q}'$ is environmentally bounded: Since $!\underline{P}'$ is environmentally bounded, it is easy to see that $!\underline{P}'' | !\underline{P}'$ is environmentally bounded too, and in particular $\underline{P}'' | \underline{P}'$ is environmentally bounded. By Theorem 4 (with $\mathcal{S} = \underline{P}''$ and $\underline{P}' \cong \underline{Q}'$), we obtain that $\underline{P}'' | \underline{P}' \cong \underline{P}'' | \underline{Q}'$ and that $\underline{P}'' | \underline{Q}'$ is environmentally bounded. Then, by Lemma 9, we conclude that $!\underline{P}'' | !\underline{Q}'$ is environmentally bounded.

From this it easily follows that $!\underline{P}'' | !\underline{Q}' | \underline{P}$ and $!\underline{P}'' | !\underline{Q}' | \underline{Q}$ are environmentally bounded as well. Since $\underline{P} \cong \underline{Q}$, by Theorem 4 (with $\mathcal{S} = !\underline{P}'' | !\underline{Q}'$), we now obtain that $!\underline{P}'' | !\underline{Q}' | \underline{P} \cong !\underline{P}'' | !\underline{Q}' | \underline{Q}$.

It is easy to see that $\mathcal{E}_\mathbb{S}$, $\mathcal{E}_{p_{\mathcal{E}}(\eta+|a|)}$, and \mathcal{E}_r , for every $r \in \mathbf{N}$, are universally bounded. So, we obtain that $\mathcal{H}_\mathbb{S} | \underline{P} \equiv \mathcal{H}_\mathbb{S} | \underline{Q}$. Hence, there exists a negligible function f such that for all $\eta \in \mathbf{N}$ and $a \in \{0, 1\}^*$:

$$\begin{aligned}
f(\eta, a) &\geq |\text{Prob}[(\mathcal{H}_\mathbb{S} | \underline{P})(1^\eta, a) = 1] - \text{Prob}[(\mathcal{H}_\mathbb{S} | \underline{Q})(1^\eta, a) = 1]| \\
&= \frac{1}{p_{\mathcal{E}}(\eta + |a|)} \cdot \left| \sum_{r=1}^{p_{\mathcal{E}}(\eta+|a|)} \text{Prob}[(\mathcal{H}_r | \underline{P})(1^\eta, a) = 1] - \text{Prob}[(\mathcal{H}_r | \underline{Q})(1^\eta, a) = 1] \right| \\
&\stackrel{(8)}{=} \frac{1}{p_{\mathcal{E}}(\eta + |a|)} \cdot |\text{Prob}[(\mathcal{H}_1 | \underline{P})(1^\eta, a) = 1] - \text{Prob}[(\mathcal{H}_{p_{\mathcal{E}}(\eta+|a|)} | \underline{Q})(1^\eta, a) = 1]| \\
&\stackrel{(7),(9)}{=} \frac{1}{p_{\mathcal{E}}(\eta + |a|)} \cdot |\text{Prob}[(\mathcal{E} | !\underline{P})(1^\eta, a) = 1] - \text{Prob}[(\mathcal{E} | !\underline{Q})(1^\eta, a) = 1]| .
\end{aligned}$$

We conclude that $\mathcal{E} | !\underline{P} \equiv \mathcal{E} | !\underline{Q}$. □

We note that by iteratively applying this composition theorem and Theorem 4, one can establish environmental indistinguishability for more and more complex systems in a modular way.

5.3 Composition Theorem for Unbounded Self-Composition of SID Dependent Systems

As explained before, in Theorem 5 SIDs are used merely as a means to address (IITMs belonging to) copies/sessions of systems. In particular, a session of a system is not aware of the SID that is used to address it, and the specific addressing mechanism employed. We now present a composition theorem for the unbounded self-composition of a system where sessions *are* aware of their SIDs. This theorem is basically a corollary of Theorem 5.

We first need to generalize the notion of a session version of an IITM and a system and introduce further notions in the context of session versions.

5.3.1 Generalized Session Versions

A function $\sigma: \{0,1\}^* \times \mathcal{T} \rightarrow \{0,1\}^* \cup \{\perp\}$ (where \mathcal{T} is a set of tape names) is called a *session identifier (SID) function* if it is computable in polynomial time (in the length of its input). Intuitively, an SID function assigns an SID to a message w.r.t. some tape (or \perp , if the message does not have an SID). For example, $\sigma_0(m, c) := 0$, for all m, c , assigns the SID 0 to every message. The following function takes the prefix of a message as its SID: $\sigma_{\text{prefix}}(m, c) := s$ if $m = (s, m')$ for some s, m' and $\sigma_{\text{prefix}}(m, c) := \perp$ otherwise, for all m, c . Clearly, many more examples are conceivable. We note that an SID function, besides the message, also gets a tape name as input because the way SIDs are extracted from messages may depend on the tape on which a message is received; this gives even more flexibility, although in most cases an SID of a message will be determined independently of the tape on which the message is received.

We now define the notion of a σ -session version, which generalizes the notion of session versions introduced in Section 5.2.1. Intuitively, a machine M is a σ -session version, if M in mode `CheckAddress` accepts only messages that have the same SID w.r.t. σ and M outputs only messages with the same SID. For a system to be a σ -session version, we require that all machines in that system are σ -session versions. More precisely, σ -session versions are defined as follows:

Definition 13. Let σ be an SID function and let M be an IITM such that $\mathcal{T}(M) \subseteq \mathcal{T}$ where \mathcal{T} is the set of tape names σ is defined on (i.e., σ is defined on all names of input and output tapes of M). Then, M is a σ -session machine (also called a σ -session version) if for every system \mathcal{S} such that \mathcal{S} and M are connectable the following conditions are satisfied for every η, a and every run ρ of $(\mathcal{S} \mid M)(1^\eta, a)$:

1. Whenever M is activated in ρ in mode `CheckAddress` with an input message m on tape c , then M rejects m if $\sigma(m, c) = \perp$.
2. If the first input message that M accepted in ρ in mode `CheckAddress` is m_0 on tape c_0 and (later) M is activated in mode `CheckAddress` in ρ with an input message m on tape c , then M rejects m if $\sigma(m, c) \neq \sigma(m_0, c_0)$.
3. Whenever M outputs a message m on tape c in ρ in mode `Compute`, then $\sigma(m, c) = \sigma(m_0, c_0)$ (where the first accepted message was m_0 on tape c_0 , see above).

A system \mathcal{Q} is a σ -session system/version if every IITM occurring in \mathcal{Q} is a σ -session version.

We note that if \mathcal{Q} is a protocol system and unless σ is trivial, i.e., allows a machine to accept all messages on all tapes (such as σ_0), every IITM occurring in \mathcal{Q} , by definition of protocol systems (see Definition 11), has to be in the scope of a bang because it does not accept all messages in mode `CheckAddress`.

We emphasize that by 2. above, the fact that $\sigma(m, c) = \sigma(m_0, c_0) \neq \perp$ does not mean that a machine has to accept m on c . This is only a necessary condition for the machine to be able to accept m on c in mode `CheckAddress`. In other words, every system which is a σ -session version can accept “less” messages than “suggested” by σ ; in particular, every system \mathcal{S} is a σ_0 -session version, since σ_0 allows a machine to accept every message.

The following two lemmas capture some more basic properties of σ -session versions. The first lemma says that if a system is a σ -session version, then it is also a σ' -session version for a function σ' that is at least as permissive as σ .

Lemma 10. *Let σ and σ' be SID functions such that for all $m, m' \in \{0,1\}^*, c, c' \in \mathcal{T}$: $\sigma'(m, c) = \sigma'(m', c') \neq \perp$ if $\sigma(m, c) = \sigma(m', c') \neq \perp$. Then, every σ -session system is a σ' -session system.*

The following lemma shows how standard session versions, as introduced in Section 5.2.1, and σ -session versions can be combined.

Lemma 11. *Let σ be an SID function and \mathcal{S} be a σ -session system. Then, $!\underline{\mathcal{S}}$ is a σ' -session system where for all $m \in \{0,1\}^*, c \in \mathcal{T}$:*

$$\sigma'(m, c) := \begin{cases} (s, \sigma(m', c)) & \exists s, m' : m = (s, m') \wedge \sigma(m', c) \neq \perp, \\ \perp & \text{otherwise} . \end{cases}$$

Combining the two lemmas above and using that every system is a σ_0 -session version, we obtain that the multi-session version $!\underline{\mathcal{S}}$ of every system \mathcal{S} is a σ_{prefix} -session version. Moreover, applying Lemma 11 to $!\underline{\mathcal{S}}$ yields that the multi-session version of the multi-session version of every system \mathcal{S} , i.e., $!\underline{\underline{\mathcal{S}}}$, is a $\sigma_{\text{prefix}}^{(2)}$ -session version where $\sigma_{\text{prefix}}^{(2)}(m, c) := (s, s')$ if $m = (s, (s', m'))$ for some s, s', m' and $\sigma_{\text{prefix}}^{(2)}(m, c) := \perp$ otherwise, for all m, c . Note that by Lemma 10 every $\sigma_{\text{prefix}}^{(2)}$ -session version (e.g., $!\underline{\underline{\mathcal{S}}}$) is a σ_{prefix} -session version. This is analogous to the (standard) session versions from Section 5.2.1: $!\underline{\underline{\mathcal{S}}}$ is a session version itself.

Now, given two systems \mathcal{P} and \mathcal{Q} which both are σ -session versions, the composition theorem basically states that if \mathcal{P} and \mathcal{Q} are single-session indistinguishable, i.e., \mathcal{P} and \mathcal{Q} are environmentally indistinguishable for environments that call a single session of \mathcal{P}/\mathcal{Q} only, then \mathcal{P} and \mathcal{Q} are environmentally indistinguishable ($\mathcal{P} \cong \mathcal{Q}$). We now formalize “single-session indistinguishability”. We first need to formalize systems that invoke only a single session of another system. Such systems are called σ -single session and they output messages with the same SID only. More precisely:

Definition 14. Let σ be an SID function. A system \mathcal{Q} is σ -single session if for every system \mathcal{S} such that \mathcal{S} and \mathcal{Q} are connectable the following is true for every η, a and in every run ρ of $(\mathcal{S} | \mathcal{Q})(1^\eta, a)$: Let $m_0 \neq \varepsilon$ (where ε is the empty bit string) be the first message output by \mathcal{Q} on some external tape c_0 (except the decision tape) in ρ . Then $\sigma(m_0, c_0) \neq \perp$ and every message $m \neq \varepsilon$ output by \mathcal{Q} on an external tape c (except the decision tape) in ρ satisfies $\sigma(m, c) = \sigma(m_0, c_0)$.

Let σ be an SID function. Given a system \mathcal{S} , by $\text{Env}_{\sigma\text{-single}}(\mathcal{S})$ we denote the set of all systems $\mathcal{E} \in \text{Env}(\mathcal{S})$ such that \mathcal{E} is σ -single session, i.e., $\text{Env}_{\sigma\text{-single}}(\mathcal{S})$ is the set of all σ -single session environmental systems that can be connected to \mathcal{S} .

Definition 15. Let σ be an SID function. Two systems \mathcal{P} and \mathcal{Q} are called σ -environmentally indistinguishable or *indistinguishable w.r.t. σ -single session environments*, denoted by $\mathcal{P} \cong_{\sigma\text{-single}} \mathcal{Q}$, if and only if

1. \mathcal{P} and \mathcal{Q} are compatible and
2. $\mathcal{E} | \mathcal{P} \equiv \mathcal{E} | \mathcal{Q}$ for all $\mathcal{E} \in \text{Env}_{\sigma\text{-single}}(\mathcal{P})$.

We will also need the following definition.

Definition 16. Let σ be an SID function. A system \mathcal{S} is σ -environmentally (almost) bounded or *environmentally bounded w.r.t. σ -single session environments* if for every $\mathcal{E} \in \text{Env}_{\sigma\text{-single}}(\mathcal{S})$: $\mathcal{E} | \mathcal{S}$ is almost bounded.

Clearly, every environmentally bounded system is σ -environmentally bounded for every SID function σ .

5.3.2 A Composition Theorem for σ -Session Versions

We are now able to formulate the composition theorem for the unbounded self-composition of systems that may depend on their SID:

Theorem 6. *Let σ be an SID function. Let \mathcal{P} and \mathcal{Q} be two protocol systems such that \mathcal{P} and \mathcal{Q} are σ -session versions, $!\underline{\mathcal{P}}$ is environmentally bounded, \mathcal{Q} is σ -environmentally bounded, and $\mathcal{P} \cong_{\sigma\text{-single}} \mathcal{Q}$. Then, $\mathcal{P} \cong \mathcal{Q}$ and \mathcal{Q} is environmentally bounded.*

We note that a stronger variant of the above theorem, where it is assumed that only \mathcal{P} (instead of $!\underline{\mathcal{P}}$) is environmentally bounded, also holds (in this case, the claim still is that $\mathcal{P} \cong \mathcal{Q}$ and \mathcal{Q} is environmentally bounded). Such a theorem can be proven by redoing the proofs of Theorem 5 and Lemma 9.¹² Here we want to prove Theorem 6 as a corollary of Theorem 5 and this requires the additional assumption.¹³ However, as discussed below Theorem 5, typically \mathcal{P} is environmentally strictly bounded and in this case, by Lemma 17, it follows immediately that $!\underline{\mathcal{P}}$ is environmentally bounded.

¹²As for the lemma, one assumes that (i) \mathcal{P} and \mathcal{Q} are σ -session versions, (ii) \mathcal{P} is environmentally bounded, (iii) \mathcal{Q} is σ -environmentally bounded, and (iv) $\mathcal{P} \cong_{\sigma\text{-single}} \mathcal{Q}$. For Lemma 9 the statement then is that \mathcal{Q} is environmentally bounded.

¹³In general, assuming that $!\underline{\mathcal{P}}$ is environmentally bounded is a stronger condition than just assuming that \mathcal{P} is environmentally bounded (see Lemma 19).

Proof of Theorem 6. The basic idea of the proof is as follows: We first define single IITMs $[\mathcal{P}]_\sigma$ and $[\mathcal{Q}]_\sigma$ which, in every environment, simulate a single session of \mathcal{P} and \mathcal{Q} (w.r.t. σ), respectively. Since $\mathcal{P} \cong_{\sigma\text{-single}} \mathcal{Q}$ we obtain $[\mathcal{P}]_\sigma \cong [\mathcal{Q}]_\sigma$. Then, by the composition theorem (Theorem 5), we obtain that $![\mathcal{P}]_\sigma \cong ![\mathcal{Q}]_\sigma$. Now, from $![\mathcal{P}]_\sigma$ and $![\mathcal{Q}]_\sigma$ we can conclude that $\mathcal{P} \cong \mathcal{Q}$: An environment \mathcal{E} for \mathcal{P} and \mathcal{Q} can be simulated by an environment \mathcal{E}' which simply prefixes every message m output on tape c with $\sigma(m, c)$ if $\sigma(m, c) \neq \perp$.

Following this idea, we first define $[\mathcal{P}]_\sigma$. As in the proof of Lemma 6, we construct $[\mathcal{P}]_\sigma$ as an IITM which is compatible with \mathcal{P} and accepts every message in mode **CheckAddress**. However, it only simulates a single session of \mathcal{P} w.r.t. σ . More specifically, $[\mathcal{P}]_\sigma$ (in any run with any system) does the following: If $[\mathcal{P}]_\sigma$ receives a message m on tape c with $\sigma(m, c) = \perp$, then $[\mathcal{P}]_\sigma$ ends its current activation with empty output. Let m_0 be the first message that $[\mathcal{P}]_\sigma$ receives on some external tape c_0 with $\sigma(m_0, c_0) \neq \perp$, then $[\mathcal{P}]_\sigma$ simulates \mathcal{P} with input m_0 on c_0 and if \mathcal{P} produces output m' on some external tape c' , $[\mathcal{P}]_\sigma$ produces output m' on c' . Now, whenever later $[\mathcal{P}]_\sigma$ receives a message m on tape c with $\sigma(m, c) \neq \sigma(m_0, c_0)$, then $[\mathcal{P}]_\sigma$ ends its current activation with empty output. Otherwise, i.e., if $\sigma(m, c) = \sigma(m_0, c_0)$, then $[\mathcal{P}]_\sigma$ simulates \mathcal{P} with input m on c as above.

By construction of $[\mathcal{P}]_\sigma$ and since \mathcal{P} is a σ -session version, it is easy to see that for every $\mathcal{E} \in \text{Env}_{\sigma\text{-single}}(\mathcal{P})$ we have that $\mathcal{E} \mid \mathcal{P} \equiv_0 \mathcal{E} \mid [\mathcal{P}]_\sigma$. (This does not necessarily hold for all $\mathcal{E} \in \text{Env}(\mathcal{P})$.) Because $!\mathcal{P}$ is environmentally bounded by assumption, and hence, in particular \mathcal{P} is environmentally bounded, it is not hard to see that $[\mathcal{P}]_\sigma$ is environmentally bounded. In fact, to conclude that $[\mathcal{P}]_\sigma$ is environmentally bounded it suffices to assume that \mathcal{P} is σ -environmentally bounded, since $[\mathcal{P}]_\sigma$ simulates only a single session of \mathcal{P} .

Just as $[\mathcal{P}]_\sigma$, we define $[\mathcal{Q}]_\sigma$. By assumption, \mathcal{Q} is a σ -session version and it is σ -environmentally bounded. Now, just as in the case of $[\mathcal{P}]_\sigma$, we obtain that (i) $\mathcal{E} \mid \mathcal{Q} \equiv_0 \mathcal{E} \mid [\mathcal{Q}]_\sigma$ for every $\mathcal{E} \in \text{Env}_{\sigma\text{-single}}(\mathcal{Q}) = \text{Env}_{\sigma\text{-single}}(\mathcal{P})$ and (ii) $[\mathcal{Q}]_\sigma$ is environmentally bounded.

Because $\mathcal{P} \cong_{\sigma\text{-single}} \mathcal{Q}$, it now easily follows that $\mathcal{E} \mid [\mathcal{P}]_\sigma \equiv \mathcal{E} \mid [\mathcal{Q}]_\sigma$ for all $\mathcal{E} \in \text{Env}(\mathcal{P})$, and thus,

$$[\mathcal{P}]_\sigma \cong [\mathcal{Q}]_\sigma . \quad (10)$$

By assumption, $!\mathcal{P}$ is environmentally bounded. It is not hard to see that this implies that $![\mathcal{P}]_\sigma$ is environmentally bounded too.¹⁴ By Lemma 9 it follows that $![\mathcal{Q}]_\sigma$ is environmentally bounded. Moreover, by Theorem 5 we obtain:

$$![\mathcal{P}]_\sigma \cong ![\mathcal{Q}]_\sigma . \quad (11)$$

We use (11) to show that

$$\mathcal{P} \cong \mathcal{Q} . \quad (12)$$

Let $\mathcal{E} \in \text{Env}(\mathcal{P})$. By Lemma 7, we may assume that \mathcal{E} is a single IITM which accepts every message in mode **CheckAddress**. Furthermore, without loss of generality, we may assume that $\text{start}, \text{decision} \in \mathcal{T}(\mathcal{E})$ and that \mathcal{E} only outputs messages m on tape c with $\sigma(m, c) \neq \perp$, for every tape $c \neq \text{decision}$: since \mathcal{P} and \mathcal{Q} are σ -session versions, they would reject any message m on tape c with $\sigma(m, c) = \perp$ anyway. We may also assume that on tape decision , \mathcal{E} only outputs 0 or 1.

Now, to show that $\mathcal{E} \mid \mathcal{P} \equiv \mathcal{E} \mid \mathcal{Q}$, we define an IITM $\mathcal{E}' \in \text{Env}(![\mathcal{P}]_\sigma)$. The external tapes of \mathcal{E}' are the external tapes of \mathcal{E} . In mode **CheckAddress**, \mathcal{E}' accepts all messages. In mode **Compute**, \mathcal{E}' simulates \mathcal{E} as follows:

- If \mathcal{E}' receives some input m (possibly empty input) on the tape **start**, then \mathcal{E}' simulates \mathcal{E} on input m on **start**. Otherwise, if \mathcal{E}' receives some input of shape (s, m) on some tape c with $\sigma(m, c) = s \neq \perp$, then \mathcal{E}' simulates \mathcal{E} on input m on c .
- If the simulated \mathcal{E} outputs some message m on the tape **decision**, then \mathcal{E}' outputs m on **decision**. Otherwise, if the simulated \mathcal{E} outputs some message m on some tape $c \neq \text{decision}$ with $\sigma(m, c) \neq \perp$, then \mathcal{E}' outputs $(\sigma(m, c), m)$ on c .

¹⁴To conclude that $![\mathcal{P}]_\sigma$ is environmentally bounded, the assumption that $!\mathcal{P}$ is environmentally bounded is required. Just assuming that \mathcal{P} is environmentally bounded would not be enough.

- Otherwise, i.e., \mathcal{E}' receives some other input or the simulated \mathcal{E} produces different output, \mathcal{E}' produces output “error” $\notin \{0, 1\}$ on decision.

Since \mathcal{E} never outputs messages with $\sigma(m, c) = \perp$ (expect maybe on tape decision) and by definition of \mathcal{E}' , $![\underline{\mathcal{P}}]_\sigma$, and $![\underline{\mathcal{Q}}]_\sigma$, it is easy to see that for all η, a :

$$\text{Prob} \left[(\mathcal{E}' \mid ![\underline{\mathcal{P}}]_\sigma)(1^\eta, a) = \text{“error”} \right] = \text{Prob} \left[(\mathcal{E}' \mid ![\underline{\mathcal{Q}}]_\sigma)(1^\eta, a) = \text{“error”} \right] = 0 .$$

Therefore it is easy to define a bijection between the runs of $(\mathcal{E}' \mid ![\underline{\mathcal{P}}]_\sigma)(1^\eta, a)$ and $(\mathcal{E} \mid \mathcal{P})(1^\eta, a)$ as well as between the runs of $(\mathcal{E}' \mid ![\underline{\mathcal{Q}}]_\sigma)(1^\eta, a)$ and $(\mathcal{E} \mid \mathcal{Q})(1^\eta, a)$. Hence, we obtain:

$$\mathcal{E}' \mid ![\underline{\mathcal{P}}]_\sigma \equiv_0 \mathcal{E} \mid \mathcal{P} \quad \text{and} \quad \mathcal{E}' \mid ![\underline{\mathcal{Q}}]_\sigma \equiv_0 \mathcal{E} \mid \mathcal{Q} . \quad (13)$$

Since \mathcal{E} is universally bounded, it follows easily that \mathcal{E}' is universally bounded too. Thus, $\mathcal{E}' \in \text{Env}(![\underline{\mathcal{P}}]_\sigma)$. By (11), we can conclude that $\mathcal{E}' \mid ![\underline{\mathcal{P}}]_\sigma \equiv \mathcal{E}' \mid ![\underline{\mathcal{Q}}]_\sigma$. It now follows that

$$\mathcal{E} \mid \mathcal{P} \stackrel{(13)}{\equiv_0} \mathcal{E}' \mid ![\underline{\mathcal{P}}]_\sigma \equiv \mathcal{E}' \mid ![\underline{\mathcal{Q}}]_\sigma \stackrel{(13)}{\equiv_0} \mathcal{E} \mid \mathcal{Q} .$$

With the bijection between runs of $\mathcal{E}' \mid ![\underline{\mathcal{Q}}]_\sigma(1^\eta, a)$ and $(\mathcal{E} \mid \mathcal{Q})(1^\eta, a)$ and the fact that $![\underline{\mathcal{Q}}]_\sigma$ is environmentally bounded, it also immediately follows that \mathcal{Q} is environmentally bounded. \square

We note that under the assumptions of Theorem 6, one can even show that $!\underline{\mathcal{Q}}$ is environmentally bounded, using that $![\underline{\mathcal{Q}}]_\sigma$ is environmentally bounded.

Clearly, by iteratively applying Theorems 4, 5, and 6, one obtains, in a modular way, environmental indistinguishability for more and more complex systems.

6 Notions of Universal Composability

In the literature, several notions of universal composability have been proposed: universal composability (UC) [2, 27], dummy UC [2], black-box simulatability [27], strong simulatability [19], and reactive simulatability [27].

While intuitively one would expect these notions to be equivalent (where for reactive simulatability this requires non-uniform environments/environments with external input), this is not the case in all models. In particular, it is not true in Canetti’s UC model, due to the incompleteness of the dummy adversary (see Section 4.5 and Remark 11).

In this section, the mentioned notions are defined and it is shown that these notions are equivalent in the IITM model. We also state further basic properties. But first we need to introduce some notation and terminology.

6.1 Further Notation and Terminology

To define the notions of universal composability we need the following kinds of systems: environmental, adversarial systems, and protocol systems. While we have already defined environmental and protocol systems (see Definitions 8 and 11), we still need to introduce adversarial systems and the way these systems are connected.

6.1.1 Network and I/O Tapes

To define how the different entities (environments, protocols, adversaries/simulators) connect to each other, we partition the set of external tapes of every system \mathcal{S} into *network* and *I/O tapes*. We denote the set of (external) network tapes of \mathcal{S} by $\mathcal{T}_{ext}^{net}(\mathcal{S})$ and the set of (external) I/O tapes of \mathcal{S} by $\mathcal{T}_{ext}^{io}(\mathcal{S})$. Each of the sets $\mathcal{T}_{ext}^{net}(\mathcal{S})$ and $\mathcal{T}_{ext}^{io}(\mathcal{S})$ is further partitioned into input and output tapes. We denote by $\mathcal{T}_{in}^{net}(\mathcal{S})$, $\mathcal{T}_{out}^{net}(\mathcal{S})$,

$\mathcal{T}_{in}^{io}(\mathcal{S})$, and $\mathcal{T}_{out}^{io}(\mathcal{S})$ the set of network input and output tapes and the set of I/O input and output tapes, respectively.

We assume that network tapes are named differently from I/O tapes also across different systems. This can easily be enforced, for instance, by using different prefixes for the different kind of tapes, such as *net* for network tapes and *io* for I/O tapes. By this, we obtain the following:

- (i) If \mathcal{P} and \mathcal{Q} are compatible, then they also agree on the type (network or I/O) of their external tapes, i.e., $\mathcal{T}_{in}^{net}(\mathcal{P}) = \mathcal{T}_{in}^{net}(\mathcal{Q})$, $\mathcal{T}_{out}^{net}(\mathcal{P}) = \mathcal{T}_{out}^{net}(\mathcal{Q})$, $\mathcal{T}_{in}^{io}(\mathcal{P}) = \mathcal{T}_{in}^{io}(\mathcal{Q})$, and $\mathcal{T}_{out}^{io}(\mathcal{P}) = \mathcal{T}_{out}^{io}(\mathcal{Q})$.
- (ii) If \mathcal{P} and \mathcal{Q} are connectable, then they agree on the type (network or I/O) of their common external tapes, i.e., for all $c \in \mathcal{T}_{ext}(\mathcal{P}) \cap \mathcal{T}_{ext}(\mathcal{Q})$, we have that $c \in \mathcal{T}_{ext}^{net}(\mathcal{P}) \cap \mathcal{T}_{ext}^{net}(\mathcal{Q})$ or $c \in \mathcal{T}_{ext}^{io}(\mathcal{P}) \cap \mathcal{T}_{ext}^{io}(\mathcal{Q})$.

Definition 17. The systems \mathcal{Q} and \mathcal{P} are *I/O-connectable* if \mathcal{Q} and \mathcal{P} are connectable and $\mathcal{T}_{ext}^{net}(\mathcal{Q}) \cap \mathcal{T}_{ext}^{net}(\mathcal{P}) = \emptyset$. In other words, \mathcal{Q} and \mathcal{P} only connect through their I/O tapes.

The systems $\mathcal{S}_1, \dots, \mathcal{S}_n$ are *I/O-connectable* if \mathcal{S}_i and \mathcal{S}_j are I/O-connectable for every $i, j \leq n$ such that $i \neq j$.

As already explained in Section 4.1 for connectability, if $\mathcal{S}_1, \dots, \mathcal{S}_n$ are I/O-connectable, then every common external tape of the systems is the external input tape of exactly one system \mathcal{S}_i and the external output tape of exactly one other system \mathcal{S}_j . So, there is no ambiguity in how these systems connect to each other in the parallel composition $\mathcal{S}_1 \mid \dots \mid \mathcal{S}_n$. Note that for I/O-connectability every common tape between these systems is an I/O tape, and hence, these systems only connect through their I/O tapes.

6.1.2 Adversarial Systems

We are now ready to define adversarial systems and how these connect to protocol systems.

Definition 18. An *adversarial system* \mathcal{A} is a system such that no tape in \mathcal{A} is named *start* or *decision*.

For a system \mathcal{S} , we denote by $\text{Adv}(\mathcal{S})$ the set of all adversarial systems \mathcal{A} such that \mathcal{A} can be connected to \mathcal{S} , $\mathcal{A} \mid \mathcal{S}$ is environmentally bounded, the set of external tapes of \mathcal{A} is disjoint from the set of I/O tapes of \mathcal{S} , and $\mathcal{A} \mid \mathcal{S}$ does not have any (external) network tapes (i.e., $\mathcal{A} \in \text{Con}(\mathcal{S})$, $\mathcal{T}_{ext}(\mathcal{A}) \cap \mathcal{T}_{ext}^{io}(\mathcal{S}) = \emptyset$, and $\mathcal{T}_{ext}^{net}(\mathcal{A} \mid \mathcal{S}) = \emptyset$). Thus, an adversary \mathcal{A} can connect to the network tapes of \mathcal{S} only.

For two systems \mathcal{S} and \mathcal{S}' , by $\text{Sim}^{\mathcal{S}'}(\mathcal{S})$ we denote the set of all adversarial systems \mathcal{A} such that \mathcal{A} can be connected to \mathcal{S} , $\mathcal{A} \mid \mathcal{S}$ is environmentally bounded, the set of external tapes of \mathcal{A} is disjoint from the set of I/O tapes of \mathcal{S} , and $\mathcal{A} \mid \mathcal{S}$ is compatible with \mathcal{S}' .

We note that environmental systems, which will run concurrently with a protocol and possibly an adversarial system, may contain tapes named *start* or *decision*. In particular, they may contain a master IITM (while protocol and adversarial systems may not). This choice is justified and motivated by results shown in [19, 11].

6.2 Defining the Notions of Universal Composability

The various security notions for universal composability proposed in the literature—(dummy) UC, black-box simulatability, strong simulatability, reactive simulatability—can be defined in a concise and simple way in the IITM model.

The basic idea behind these security notions is that security properties are specified as an ideal protocol/functionality \mathcal{F} . For example, \mathcal{F} might specify a secure channel (see also Section 10.2 for examples). Using such an ideal protocol, parties can carry out their security critical tasks, such as communicating over a secure channel, in a secure way. This functionality is secure by construction, it defines what security means. Now a real protocol \mathcal{P} realizes an ideal protocol \mathcal{F} if for every (real) adversary \mathcal{A} for \mathcal{P} there exists an ideal adversary/simulator \mathcal{I} for \mathcal{F} such that no environment can distinguish whether it interacts with $\mathcal{A} \mid \mathcal{P}$ or with $\mathcal{I} \mid \mathcal{F}$. Since \mathcal{F} is secure by definition and no environment can distinguish whether it interacts with the real or the ideal system, \mathcal{P} is as secure as \mathcal{F} , and hence, secure.

The different security notions proposed in the literature differ in the kind of real adversaries considered and in the order of quantification. For dummy UC, the real adversary is just the dummy adversary, i.e., the adversary which forwards all messages between the protocol and the environment. For strong simulatability, the real adversary is dropped altogether and the environment directly connects to the network interface of the real protocol. For black-box simulatability, one requires that there exists a simulator \mathcal{S} such that for every adversary \mathcal{A} the composition $\mathcal{A}|\mathcal{S}$ is a good ideal adversary. For reactive simulatability, the ideal adversary may depend on the environment that tries to distinguish the real from the ideal protocol.

In order to formally define the notion of dummy UC, we first need to define the dummy adversary for a protocol. Given a system \mathcal{P} , by $\mathcal{D}_{\mathcal{P}}$ we denote the *dummy adversary for \mathcal{P}* which is the dummy IITM $\mathcal{D}_{\mathcal{P}} = \mathcal{D}(\mathcal{T}_{in}^{net}(\mathcal{P}), \mathcal{T}_{out}^{net}(\mathcal{P}))$ (see Section 4.5) where all tapes $c' \in \mathcal{T}_{ext}(\mathcal{D}) \setminus \mathcal{T}_{ext}^{net}(\mathcal{P})$ are declared to be I/O tapes. The dummy adversary simply forwards all messages received from the protocol on a network tape to the environment on the corresponding I/O tape and all messages it receives from the environment on an I/O tape to the protocol on the corresponding network tape. We note that if \mathcal{P} is environmentally bounded, then $\mathcal{D}_{\mathcal{P}}|\mathcal{P}$ is environmentally bounded too, and hence, $\mathcal{D}_{\mathcal{P}} \in \text{Adv}(\mathcal{P})$.

The security notions defined in what follows are illustrated in Figures 4 to 7. In terms of runtime, the notions UC and dummyUC introduced below conceptually follow the definitions in [16].

Definition 19. Let \mathcal{P} and \mathcal{F} be protocol systems, the *real* and *ideal protocol*, respectively, such that \mathcal{P} is environmentally bounded.

1. *Strong Simulatability (SS)*: $\mathcal{P} \leq^{SS} \mathcal{F}$ iff $\exists \mathcal{S} \in \text{Sim}^{\mathcal{P}}(\mathcal{F})$:

$$\mathcal{P} \cong \mathcal{S}|\mathcal{F} .$$

2. *Universal Simulatability/Composability (UC)*: $\mathcal{P} \leq^{UC} \mathcal{F}$ iff $\forall \mathcal{A} \in \text{Adv}(\mathcal{P}) \exists \mathcal{I} \in \text{Sim}^{\mathcal{A}|\mathcal{P}}(\mathcal{F})$:

$$\mathcal{A}|\mathcal{P} \cong \mathcal{I}|\mathcal{F} .$$

3. *Dummy Version of UC (dummyUC)*: $\mathcal{P} \leq^{dumUC} \mathcal{F}$ iff $\exists \mathcal{I} \in \text{Sim}^{\mathcal{D}_{\mathcal{P}}|\mathcal{P}}(\mathcal{F})$:

$$\mathcal{D}_{\mathcal{P}}|\mathcal{P} \cong \mathcal{I}|\mathcal{F} .$$

4. *Black-box Simulatability (BB)*: $\mathcal{P} \leq^{BB} \mathcal{F}$ iff $\exists \mathcal{S} \in \text{Sim}^{\mathcal{P}}(\mathcal{F}) \forall \mathcal{A} \in \text{Adv}(\mathcal{P})$:

$$\mathcal{A}|\mathcal{P} \cong \mathcal{A}|\mathcal{S}|\mathcal{F} .$$

5. *Reactive Simulatability (RS)*: $\mathcal{P} \leq^{RS} \mathcal{F}$ iff $\forall \mathcal{A} \in \text{Adv}(\mathcal{P}) \forall \mathcal{E} \in \text{Env}(\mathcal{A}|\mathcal{P}) \exists \mathcal{I} \in \text{Sim}^{\mathcal{A}|\mathcal{P}}(\mathcal{F})$:

$$\mathcal{E}|\mathcal{A}|\mathcal{P} \cong \mathcal{E}|\mathcal{I}|\mathcal{F} .$$

We say that \mathcal{P} *SS-realizes* \mathcal{F} if $\mathcal{P} \leq^{SS} \mathcal{F}$; similarly for the other security notions. We often simply say that \mathcal{P} *realizes* \mathcal{F} without explicitly indicating the underlying security notion. (This is justified by Theorem 7, which states that all security notions are equivalent.)

For strong and black-box simulatability, if \mathcal{P} and \mathcal{F} do not have disjoint network tapes, there does not exist a simulator with $\mathcal{S} \in \text{Sim}^{\mathcal{P}}(\mathcal{F})$. We therefore always (implicitly) assume that the network tapes of \mathcal{F} are renamed first so that the set of network tapes of \mathcal{P} and \mathcal{F} are disjoint.

Remark 6. We emphasize that details such as addressing of machines by party/session IDs, corruption, and the structure of protocols are not, and do not need to be fixed in order to define the security notions. This makes the definitions and the underlying model particularly simple and expressive.

Remark 7. Recall that the notion of environmental indistinguishability that we use in Definition 19 is defined with respect to environments that output one-bit. According to Remark 3, using an alternative definition with a distinguisher and an environment that may output more than one bit would yield equivalent security notions.

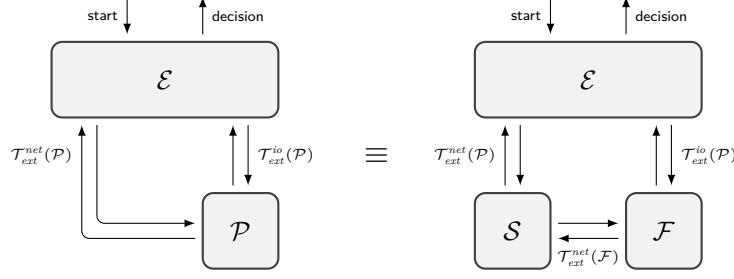


Figure 4: Strong simulatability (SS).




Figure 5: Universal simulatability/composability (UC) and reactive simulatability (RS).

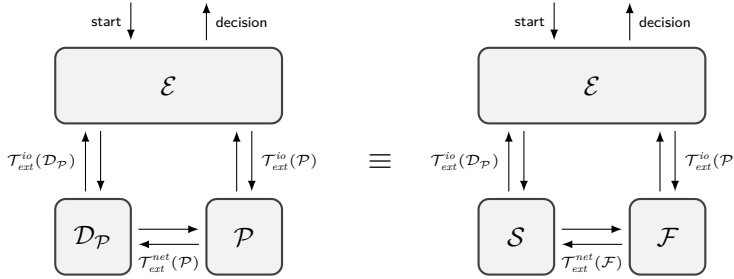


Figure 6: Dummy universal simulatability/composability (dummyUC).

Remark 8. Even though in the above definitions, \mathcal{F} is not required to be environmentally bounded, in applications it will always be environmentally bounded.

Remark 9. Note that all systems considered in Definition 19, 1. to 4. are environmentally bounded and the ones in 5. are almost bounded. For all systems, except for $\mathcal{A}|\mathcal{S}|\mathcal{F}$ in the definition of black-box simulatability, this follows immediately from the definitions of the systems. For the system $\mathcal{A}|\mathcal{S}|\mathcal{F}$ this follows using Theorem 4 as follows: Since $\mathcal{D}_{\mathcal{P}} \in \text{Adv}(\mathcal{P})$ (where $\mathcal{D}_{\mathcal{P}}$ is the dummy adversary defined above), we have that $\mathcal{D}_{\mathcal{P}}|\mathcal{P} \cong \mathcal{D}_{\mathcal{P}}|\mathcal{S}|\mathcal{F}$ and, hence, $\mathcal{P} \cong \mathcal{S}|\mathcal{F}$. By Theorem 4 (\mathcal{P} , $\mathcal{S}|\mathcal{F}$, and $\mathcal{A}|\mathcal{P}$ are environmentally bounded and $\mathcal{P} \cong \mathcal{S}|\mathcal{F}$), we obtain that $\mathcal{A}|\mathcal{S}|\mathcal{F}$ is environmentally bounded.

Remark 10. We note that all security notions imply that \mathcal{P} and \mathcal{F} have the same set of (external) I/O tapes (i.e., $\mathcal{T}_{ext}^{io}(\mathcal{P}) = \mathcal{T}_{ext}^{io}(\mathcal{F})$).

6.3 Relationships Between the Notions of Universal Composability

In this section, we show that all security notions introduced in the previous section are equivalent. We note that for the equivalence with reactive simulatability we use that the environment gets external input, and




Figure 7: Black-box simulatability (BB).

hence, is non-uniform. As shown in [15], reactive simulatability is not equivalent to universal simulatability in the uniform case; this is also true if, in the case of reactive simulatability, the external input provided to the environment is chosen before the ideal adversary. All other equivalences hold true even for uniform environments, that is, environments that do not receive external input.

Theorem 7. *Let \mathcal{P} and \mathcal{F} be protocol systems such that \mathcal{P} is environmentally bounded. Then:*

$$\mathcal{P} \leq^{SS} \mathcal{F} \text{ iff } \mathcal{P} \leq^{UC} \mathcal{F} \text{ iff } \mathcal{P} \leq^{dumUC} \mathcal{F} \text{ iff } \mathcal{P} \leq^{BB} \mathcal{F} \text{ iff } \mathcal{P} \leq^{RS} \mathcal{F}.$$

Proof. Most of the above equivalences can be proven by equational reasoning using the equational principles established in Section 4.5 and Theorem 4. Below, we prove the equivalence of strong and universal simulatability. We refer the reader to Appendix C for the other equivalences.

The basic idea of proving the equivalence of strong and universal simulatability is as follows: If $\mathcal{P} \leq^{SS} \mathcal{F}$, then there exists a simulator \mathcal{S} such that $\mathcal{P} \cong \mathcal{S} | \mathcal{F}$. Now for every adversary \mathcal{A} for \mathcal{P} , we define $\mathcal{I} := \mathcal{A} | \mathcal{S}$ to be the ideal adversary (simulator). We obtain that $\mathcal{A} | \mathcal{P} \cong \mathcal{I} | \mathcal{F}$. Conversely, if $\mathcal{P} \leq^{UC} \mathcal{F}$, then we can use the simulator obtained for the dummy adversary to prove $\mathcal{P} \leq^{SS} \mathcal{F}$. A detailed proof of the equivalence follows.

First, we show that $\mathcal{P} \leq^{SS} \mathcal{F}$ implies $\mathcal{P} \leq^{UC} \mathcal{F}$. By definition of $\mathcal{P} \leq^{SS} \mathcal{F}$, we obtain that there exists $\mathcal{S} \in \text{Sim}^{\mathcal{P}}(\mathcal{F})$ such that $\mathcal{P} \cong \mathcal{S} | \mathcal{F}$. Let $\mathcal{A} \in \text{Adv}(\mathcal{P})$. By Lemma 6, we may assume that both \mathcal{S} and \mathcal{A} are single IITMs that accept every message in mode CheckAddress.¹⁵ We define $\mathcal{I} := \mathcal{A} | \mathcal{S}$. Since $\mathcal{P} \cong \mathcal{S} | \mathcal{F}$ and $\mathcal{P}, \mathcal{S} | \mathcal{F}$, and $\mathcal{A} | \mathcal{P}$ are environmentally bounded, by Theorem 4 (note that \mathcal{A} and $\mathcal{S} | \mathcal{F}$ are protocol systems because \mathcal{A} and \mathcal{S} accept every message in mode CheckAddress) we obtain that $\mathcal{A} | \mathcal{P} \cong \mathcal{A} | \mathcal{S} | \mathcal{F} = \mathcal{I} | \mathcal{F}$ and $\mathcal{A} | \mathcal{S} | \mathcal{F} = \mathcal{I} | \mathcal{F}$ is environmentally bounded. Hence, we can also conclude that $\mathcal{I} \in \text{Sim}^{\mathcal{A} | \mathcal{P}}(\mathcal{F})$, which, altogether, proves that $\mathcal{P} \leq^{UC} \mathcal{F}$.

We now show that $\mathcal{P} \leq^{UC} \mathcal{F}$ implies $\mathcal{P} \leq^{SS} \mathcal{F}$. Since \mathcal{P} is environmentally bounded, the system $\mathcal{D}_{\mathcal{P}} | \mathcal{P}$ is environmentally bounded as well. Hence, $\mathcal{D}_{\mathcal{P}} \in \text{Adv}(\mathcal{P})$ and, by the definition of $\mathcal{P} \leq^{UC} \mathcal{F}$, there exists $\mathcal{I} \in \text{Sim}^{\mathcal{D}_{\mathcal{P}} | \mathcal{P}}(\mathcal{F})$ such that $\mathcal{D}_{\mathcal{P}} | \mathcal{P} \cong \mathcal{I} | \mathcal{F}$. We define $\mathcal{S} := \mathcal{I}'$ where \mathcal{I}' is obtained from \mathcal{I} by renaming the (external) I/O tapes c' of \mathcal{I} (i.e., the tapes that do not connect to \mathcal{F} but to an environment) to c and declaring them to be network tapes. Hence, $\mathcal{S} | \mathcal{F}$ is compatible with \mathcal{P} and $\mathcal{S} \in \text{Sim}^{\mathcal{P}}(\mathcal{F})$. Let $\mathcal{E} \in \text{Env}(\mathcal{P})$. Let \mathcal{E}' be the system obtained from \mathcal{E} by renaming the (external) network tapes c of \mathcal{E} (i.e., the tapes connecting to the network tapes of \mathcal{P}) to c' and declaring them to be I/O tapes. Hence, $\mathcal{E}' \in \text{Env}(\mathcal{D}_{\mathcal{P}} | \mathcal{P})$. Then, we obtain that:

$$\begin{aligned} \mathcal{E} | \mathcal{P} &\equiv \mathcal{E}' | \mathcal{D}_{\mathcal{P}} | \mathcal{P} && \text{(Lemma 5 and 4)} \\ &\equiv \mathcal{E}' | \mathcal{I} | \mathcal{F} && (\mathcal{D}_{\mathcal{P}} | \mathcal{P} \cong \mathcal{I} | \mathcal{F}) \\ &\equiv \mathcal{E} | \mathcal{S} | \mathcal{F} && \text{(Lemma 4)} \end{aligned}$$

We conclude that $\mathcal{P} \leq^{SS} \mathcal{F}$. □

¹⁵We note that this assumption is actually not required. Below, where we use it, we could have used the mentioned stronger variant of Theorem 4.

Since all security notions are equivalent, it does not matter which notion we use in the following. Since strong simulatability (\leq^{SS}) is the conceptually simplest notion, we typically use this notion.

Remark 11. As already mentioned at the beginning of this section, while intuitively one would expect that all security notions are equivalent, except for reactive simulatability (with uniform environments), equivalence does not hold true in all models, in particular this is the case for Canetti’s UC model. In the UC model, due to what is sometimes called the *incompleteness of the dummy adversary* UC and dummyUC are not equivalent in the UC model (see, e.g., [16]). However, completeness is needed in the proof of the composition theorem. As mentioned in [3, page 47], dummyUC and SS are not equivalent in the UC model either.

6.4 Reflexivity and Transitivity

The following two lemmas state that strong simulatability is a reflexive and transitive relationship. By Theorem 7, this is true for all security notions.

While again these properties are expected, they are not satisfied in all models. One trivial reason is that the real and ideal protocol are required to be syntactically different in some models: The ideal protocol often has to be a single machine, while this is not the case for the real protocol. More importantly, inadequate definitions of runtime notions can also cause problems (see, for instance, the discussion in [18]).

Lemma 12. *Let \mathcal{P} be an environmentally bounded protocol system. Then, $\mathcal{P} \leq^{SS} \mathcal{P}$.*

Proof. Reflexivity holds for \leq^{UC} because every valid adversary $\mathcal{A} \in \text{Adv}(\mathcal{P})$ is a valid simulator $\mathcal{A} \in \text{Sim}^{\mathcal{A}|\mathcal{P}}(\mathcal{P})$ and $\mathcal{A}|\mathcal{P} \cong \mathcal{A}|\mathcal{P}$. By Theorem 7, reflexivity also holds for \leq^{SS} . \square

Lemma 13. *Let \mathcal{P} , \mathcal{Q} , and \mathcal{F} be protocol systems such that \mathcal{P} and \mathcal{Q} are environmentally bounded, $\mathcal{P} \leq^{SS} \mathcal{Q}$, and $\mathcal{Q} \leq^{SS} \mathcal{F}$. Then, $\mathcal{P} \leq^{SS} \mathcal{F}$.*

Proof. We will show transitivity for \leq^{UC} and obtain transitivity for \leq^{SS} by Theorem 7. By Theorem 7, we obtain that $\mathcal{P} \leq^{UC} \mathcal{Q}$ and $\mathcal{Q} \leq^{UC} \mathcal{F}$. Let $\mathcal{A} \in \text{Adv}(\mathcal{P})$. By definition of $\mathcal{P} \leq^{UC} \mathcal{Q}$, there exists $\mathcal{I}_1 \in \text{Sim}^{\mathcal{A}|\mathcal{P}}(\mathcal{Q})$ such that $\mathcal{A}|\mathcal{P} \cong \mathcal{I}_1|\mathcal{Q}$. Since $\mathcal{I}_1 \in \text{Adv}(\mathcal{Q})$, by definition of $\mathcal{Q} \leq^{UC} \mathcal{F}$, there exists $\mathcal{I}_2 \in \text{Sim}^{\mathcal{I}_1|\mathcal{Q}}(\mathcal{F})$ such that $\mathcal{I}_1|\mathcal{Q} \cong \mathcal{I}_2|\mathcal{F}$. Clearly, $\mathcal{I}_2 \in \text{Sim}^{\mathcal{A}|\mathcal{P}}(\mathcal{F})$. By transitivity of \cong (see Lemma 2), we have that $\mathcal{A}|\mathcal{P} \cong \mathcal{I}_2|\mathcal{F}$. We conclude that $\mathcal{P} \leq^{UC} \mathcal{F}$. Finally, by Theorem 7, we conclude that $\mathcal{P} \leq^{SS} \mathcal{F}$. \square

7 Composition Theorems for the Realization Relations

We now prove general composition theorems for the realization relations. These theorems are at the heart of the universal composability paradigm and are the main motivation for this paradigm.

In a nutshell, these theorems say that if real protocols are secure individually, then their concurrent composition is secure as well. More precisely, if real protocols each realize some ideal protocol, then the concurrent composition of the real protocols realizes the concurrent composition of the corresponding ideal protocols, even if an unbounded number of sessions of the real/ideal protocols run concurrently. Therefore, these theorems can be used to analyze and design systems in a modular way: it suffices to show security of every individual real protocol in a single session, i.e., show that a single session of a real protocol realizes a single session of the corresponding ideal protocol, in order to conclude security of the concurrent composition of multiple sessions of these protocols.

The composition theorems presented in this section follow quite easily from the composition theorems for environmental indistinguishability presented in Section 5. Analogously to Section 5, we first, in Section 7.1, present a composition theorem for the composition of a constant number of, possibly different, protocols (more precisely, protocol systems). We then state a theorem which captures the security of a protocol when run in an unbounded number of sessions/copies (unbounded self-composition). We, similarly to Section 5, prove two versions of this theorem: one in which protocol participants are not aware of the session identifiers that are used to address the different protocol sessions (see Section 7.2) and one in which they are aware of the session identifiers (see Section 7.3). The theorems for a constant number of protocols and for unbounded

self-composition can freely be combined to establish, in a modular way, the security of more and more complex systems, as illustrated in Section 7.4.

All composition theorems are stated for strong simulatability, but by Theorem 7 they also hold for every other security notion introduced in Section 6.2, e.g., universal simulatability (UC).

Our composition theorems are, in many ways, more expressive than those proven in other models (see also the discussion in Sections 10 and 11). The fact that the IITM model does not a priori fix a specific addressing mechanism or a specific form of corruption and does not impose a specific structure on protocols allows us to prove general composition theorems which hold true no matter how these details are chosen. Moreover, unlike other models, our composition theorems do not restrict the environment to only access top-level protocols. This flexibility and generality of the theorems is also reflected in the fact that the general joint state theorem is an immediate consequence of our composition theorem. In other models, even stating the joint state theorem requires to change the model and/or introduce new concepts. Composition theorems with global setup can also be formalized without changing the model, as discussed in Section 10.1.

7.1 Composition Theorem for a Constant Number of Protocol Systems

We now present the composition theorem for the composition of a constant number of (possibly different) protocol systems. The theorem directly follows from Theorem 4.

Theorem 8. *Let $k \geq 1$ and $\mathcal{Q}, \mathcal{P}_1, \dots, \mathcal{P}_k, \mathcal{F}_1, \dots, \mathcal{F}_k$ be protocol systems with pairwise disjoint sets of network tapes such that the following conditions are satisfied:*

1. *For all $j \leq k$: \mathcal{P}_j is environmentally bounded and $\mathcal{P}_j \leq^{SS} \mathcal{F}_j$*
2. *$\mathcal{Q}, \mathcal{P}_1, \dots, \mathcal{P}_k$ are I/O-connectable and $\mathcal{Q} | \mathcal{P}_1 | \dots | \mathcal{P}_k$ is environmentally bounded.*

Then, $\mathcal{Q} | \mathcal{P}_1 | \dots | \mathcal{P}_k \leq^{SS} \mathcal{Q} | \mathcal{F}_1 | \dots | \mathcal{F}_k$.

Clearly, the theorem also hold true if the system \mathcal{Q} is dropped. We note that in the above theorem it is not required that, for all $j \leq k$, \mathcal{F}_j and $\mathcal{Q} | \mathcal{F}_1 | \dots | \mathcal{F}_k$ are environmentally bounded (although in applications this will typically be the case). Just as for Theorem 4, we finally remark that Condition 2. in the above theorem is easy to check. I/O-connectability is a simple syntactic condition. Moreover, for typical applications, it is easy to check whether a system is environmentally bounded (see also the discussion in Section 8).

Proof of Theorem 8. Since $\mathcal{P}_j \leq^{SS} \mathcal{F}_j$, for all $j \leq k$, there exists $\mathcal{S}_j \in \text{Sim}^{\mathcal{P}_j}(\mathcal{F}_j)$ such that $\mathcal{P}_j \cong \mathcal{S}_j | \mathcal{F}_j$. By Lemma 6, we may assume that all \mathcal{S}_j are single IITMs that accept all messages in mode `CheckAddress`.¹⁶ By definition of $\text{Sim}^{\mathcal{P}_j}(\mathcal{F}_j)$, we know that $\mathcal{S}_j | \mathcal{F}_j$ is environmentally bounded and \mathcal{P}_j and $\mathcal{S}_j | \mathcal{F}_j$ are compatible. Since $\mathcal{Q}, \mathcal{P}_1, \dots, \mathcal{P}_k$ are I/O-connectable, it follows that $\mathcal{S}_1 | \mathcal{F}_1, \dots, \mathcal{S}_k | \mathcal{F}_k$ are connectable. With this, by Theorem 4 (with $\mathcal{S} = \mathcal{Q}$ and $\mathcal{Q}_j = \mathcal{S}_j | \mathcal{F}_j$ for all $j \leq k$; note that \mathcal{Q}_j is a protocol system because \mathcal{S}_j accepts every message in mode `CheckAddress`), we obtain that $\mathcal{Q} | \mathcal{P}_1 | \dots | \mathcal{P}_k \cong \mathcal{Q} | \mathcal{S}_1 | \mathcal{F}_1 | \dots | \mathcal{S}_k | \mathcal{F}_k$ and that $\mathcal{Q} | \mathcal{S}_1 | \mathcal{F}_1 | \dots | \mathcal{S}_k | \mathcal{F}_k$ is environmentally bounded. In particular, we obtain that $\mathcal{S}_1 | \dots | \mathcal{S}_k \in \text{Sim}^{\mathcal{Q} | \mathcal{P}_1 | \dots | \mathcal{P}_k}(\mathcal{Q} | \mathcal{F}_1 | \dots | \mathcal{F}_k)$. Altogether, this proves that $\mathcal{Q} | \mathcal{P}_1 | \dots | \mathcal{P}_k \leq^{SS} \mathcal{Q} | \mathcal{F}_1 | \dots | \mathcal{F}_k$. \square

7.2 Composition Theorem for Unbounded Self-Composition

We now present a composition theorem for the unbounded self-composition of a protocol system. It basically says that if a real protocol securely realizes an ideal protocol, then an unbounded number of sessions (copies) of the real protocol securely realize an unbounded number of sessions of the ideal protocol, where both for the real and the ideal protocol we consider the session versions of these protocols to address different sessions. In particular, SIDs are merely used as a means to address certain (IITMs belonging to) sessions of protocols.

¹⁶We note that this assumption is not actually required. We could use the mentioned stronger variant of Theorem 4 when we use this theorem in the proof.

Protocol participants are not and do not need to be aware of the SID used to address their protocol sessions, and the specific addressing mechanism used. The theorem follows from Theorem 5.

As mentioned before, below we will present another composition theorem, which is a corollary of Theorem 6, where protocol participants are aware of their own SID.

Theorem 9. *Let \mathcal{P} and \mathcal{F} be protocol systems with disjoint sets of network tapes such that $!\underline{\mathcal{P}}$ is environmentally bounded and $\mathcal{P} \leq^{SS} \mathcal{F}$. Then, $!\underline{\mathcal{P}} \leq^{SS} !\underline{\mathcal{F}}$.*

We mention, similarly to the remark following Theorem 5, that \mathcal{P} will typically be environmentally *strictly* bounded and that this is easy to check. By Lemma 17 it then follows immediately that $!\underline{\mathcal{P}}$ is environmentally strictly bounded as well, and hence, environmentally bounded. So Theorem 9 can be applied directly to typical protocol systems without additional effort.

Proof of Theorem 9. Since $\mathcal{P} \leq^{SS} \mathcal{F}$, there exists $\mathcal{S} \in \text{Sim}^{\mathcal{P}}(\mathcal{F})$ such that $\mathcal{P} \cong \mathcal{S}|\mathcal{F}$. By Lemma 6, we may assume that \mathcal{S} is a single IITM which, in mode `CheckAddress`, accepts all messages. Hence, $\mathcal{S}|\mathcal{F}$ is a protocol system (in particular, Condition (ii) of Definition 11 is satisfied). By definition of $\text{Sim}^{\mathcal{P}}(\mathcal{F})$, $\mathcal{S}|\mathcal{F}$ is environmentally bounded and \mathcal{P} and $\mathcal{S}|\mathcal{F}$ are compatible. With this, by Theorem 5 (with $\mathcal{Q} = \mathcal{S}|\mathcal{F}$), we obtain that $!\underline{\mathcal{P}} \cong !\underline{\mathcal{S}}|\underline{\mathcal{F}}$ and $!\underline{\mathcal{S}}|\underline{\mathcal{F}}$ is environmentally bounded. The latter implies that $!\underline{\mathcal{S}} \in \text{Sim}^{!\underline{\mathcal{P}}}(!\underline{\mathcal{F}})$. This proves $!\underline{\mathcal{P}} \leq^{SS} !\underline{\mathcal{F}}$. \square

7.3 Composition Theorem for Unbounded Self-Composition of SID Dependent Protocols

As explained before, in Theorem 9, similarly to Theorem 5, SIDs are merely used as a means to address certain (IITMs belonging to) sessions of protocols. In particular, protocol participants are not aware of the SIDs used to address their protocol sessions, and the specific addressing mechanism that is employed. We now present a composition theorem for the unbounded self-composition of a protocol where the protocol participants *are* aware of their SIDs. This theorem is a corollary of Theorem 6.

Before we can state the composition theorem, we need to introduce the notion of single-session realizability, using the notion of σ -environmental indistinguishability ($\cong_{\sigma\text{-single}}$) introduced in Section 5.3.1: For two systems \mathcal{P} and \mathcal{F} and an SID function σ , we denote by $\text{Sim}_{\sigma\text{-single}}^{\mathcal{P}}(\mathcal{F})$ the set of all adversarial systems \mathcal{S} such that \mathcal{S} can be connected to \mathcal{F} , the set of external tapes of \mathcal{S} is disjoint from the set of I/O tapes of \mathcal{F} (i.e., $\mathcal{T}_{ext}(\mathcal{S}) \cap \mathcal{T}_{ext}^{io}(\mathcal{F}) = \emptyset$), $\mathcal{S}|\mathcal{F}$ and \mathcal{P} are compatible, and $\mathcal{S}|\mathcal{F}$ is σ -environmentally bounded. We note that $\text{Sim}^{\mathcal{P}}(\mathcal{F}) \subseteq \text{Sim}_{\sigma\text{-single}}^{\mathcal{P}}(\mathcal{F})$; the only difference between these two sets is that $\mathcal{S}|\mathcal{F}$ has to be environmentally bounded in one case and σ -environmentally bounded in the other case.

Definition 20. Let σ be an SID function and let \mathcal{P} and \mathcal{F} be protocol systems (the real and ideal protocol, respectively) such that \mathcal{P} is σ -environmentally bounded. We say that \mathcal{P} *single-session SS-realizes* \mathcal{F} w.r.t. σ or \mathcal{P} *SS-realizes* \mathcal{F} w.r.t. σ -single session environments, denoted by $\mathcal{P} \leq_{\sigma\text{-single}}^{SS} \mathcal{F}$, iff there exists $\mathcal{S} \in \text{Sim}_{\sigma\text{-single}}^{\mathcal{P}}(\mathcal{F})$ such that $\mathcal{P} \cong_{\sigma\text{-single}} \mathcal{S}|\mathcal{F}$.¹⁷

Now, we are able to formulate the composition theorem for unbounded self-composition of protocol systems that may depend on their SID:

Theorem 10. *Let σ be an SID function and let \mathcal{P} and \mathcal{F} be protocol systems such that \mathcal{P} and \mathcal{F} are σ -session versions with disjoint sets of network tapes, $!\underline{\mathcal{P}}$ is environmentally bounded, and $\mathcal{P} \leq_{\sigma\text{-single}}^{SS} \mathcal{F}$. Then, $\mathcal{P} \leq^{SS} \mathcal{F}$.*

We note that a stronger variant of the above theorem, where it is assumed that only \mathcal{P} (instead of $!\underline{\mathcal{P}}$) is environmentally bounded, also holds. Such a theorem can be obtained from the stronger variant of Theorem 6 which is sketched following Theorem 6 in Section 5.3.2. However, as discussed above, typically \mathcal{P} is environmentally strictly bounded and in this case, by Lemma 17, it follows immediately that $!\underline{\mathcal{P}}$ is

¹⁷As usual, if the sets of network tapes of \mathcal{P} and \mathcal{F} are not disjoint, we first rename the network tapes of \mathcal{F} accordingly.

environmentally bounded. So in concrete applications, Theorem 10 is just as easy to apply as the stronger variant.

Before we prove the above theorem, we show the following lemma, which basically states that if there exists a good single-session simulator w.r.t. some SID function σ , then there also exists a good single-session simulator which is a σ -session version. To state the lemma, we use the following terminology. We say that an IITM M is σ -complete if it satisfies all conditions stated in Definition 13 (σ -session version) but with Condition 2. replaced by the following stronger condition: If the first input message that M accepted in ρ in mode **CheckAddress** is m_0 on tape c_0 and (later) M is activated in mode **CheckAddress** in ρ with an input message m on tape c , then M accepts m iff $\sigma(m, c) = \sigma(m_0, c_0)$. In other words, σ determines exactly those messages accepted by M in mode **CheckAddress**.

Lemma 14. *Let σ be an SID function and let \mathcal{P} and \mathcal{F} be protocol systems such that \mathcal{P} and \mathcal{F} are σ -session versions, \mathcal{P} and \mathcal{F} have disjoint sets of network tapes, \mathcal{P} is σ -environmentally bounded, and $\mathcal{P} \leq_{\sigma\text{-single}}^{SS} \mathcal{F}$. Then, there exists $\mathcal{S} \in \text{Sim}_{\sigma\text{-single}}^{\mathcal{P}}(\mathcal{F})$ such that $\mathcal{P} \cong_{\sigma\text{-single}} \mathcal{S} | \mathcal{F}$ and \mathcal{S} is a single IITM which is σ -complete.*

Proof. Since $\mathcal{P} \leq_{\sigma\text{-single}}^{SS} \mathcal{F}$, there exists $\mathcal{S} \in \text{Sim}_{\sigma\text{-single}}^{\mathcal{P}}(\mathcal{F})$ such that:

$$\forall \mathcal{E} \in \text{Env}_{\sigma\text{-single}}(\mathcal{P}): \mathcal{E} | \mathcal{P} \equiv \mathcal{E} | \mathcal{S} | \mathcal{F} . \quad (14)$$

In what follows, we construct a simulator \mathcal{S}' which is a single IITM that basically simulates \mathcal{S} and is σ -complete.

By Lemma 6, we may assume that $\mathcal{S} = M$ is a single IITM which, in mode **CheckAddress**, accepts all messages. We now define the system $\mathcal{S}' := M'$ where M' is σ -complete. The basic idea behind the construction of M' is as follows: Let us consider a run of $\mathcal{E} | M | \mathcal{F}$ for some $\mathcal{E} \in \text{Env}_{\sigma\text{-single}}(\mathcal{P})$. Since \mathcal{E} is σ -single session, there exists an SID sid such that every message m output by \mathcal{E} in this run has SID sid , i.e., for every such message m and tape c we have that $\sigma(m, c) = sid$. Nevertheless, M might output messages with $\sigma(m, c) \neq sid$.

If M sent such a message to \mathcal{E} , then \mathcal{E} could easily distinguish \mathcal{P} from $M | \mathcal{F}$: Because \mathcal{E} is σ -single session and \mathcal{P} is an σ -session version, \mathcal{P} will only output messages to \mathcal{E} with SID sid . So, by (14), it can only happen with negligible probability that M sends a message m on tape c to \mathcal{E} with $\sigma(m, c) \neq sid$.

However, M might send a message m on tape c to \mathcal{F} with $\sigma(m, c) \neq sid$ (without \mathcal{E} noticing this). Since \mathcal{F} is a σ -session version, m will not be accepted by (a copy of) a machine in \mathcal{F} which has accepted messages with SID sid . Also, whenever (a copy of) a machine of \mathcal{F} receives a message, it outputs a message with the same SID. So, the copy of the machine of \mathcal{F} which accepts m must not (except with negligible probability) produce output to \mathcal{E} . In addition, with \mathcal{E} being σ -single session, \mathcal{E} cannot send messages to such a machine. Consequently, since such a machine cannot interact with \mathcal{E} , M' can simply simulate this machine internally.

From the above discussion the definition of M' suggests itself: M' is compatible with M and its **CheckAddress** mode is defined in such a way that M' is σ -complete, i.e., M' rejects all messages m on any tape c with $\sigma(m, c) = \perp$ and it accepts the first message m_0 on some tape c_0 with $\sigma(m_0, c_0) \neq \perp$. From then on, M' accepts a message m on tape c iff $\sigma(m, c) = \sigma(m_0, c_0)$. In mode **Compute**, M' internally simulates a copy of M and possibly machines of \mathcal{F} as follows.

- Whenever M' receives input m on some of its external tapes, M' internally simulates M with input m (on the same tape).
- Whenever (the internally simulated) M produces empty output, M' produces empty output.
- Whenever M outputs m on tape c with $\sigma(m, c) = sid$, M' outputs m (on the same tape).
- Whenever M outputs m on tape c to \mathcal{E} with $\sigma(m, c) \neq sid$, M' produces empty output; as discussed above, this will happen with only negligible probability.
- Whenever M outputs m on tape c to (a copy of) a machine of \mathcal{F} with $\sigma(m, c) \neq sid$, then M' internally simulates this copy with input m .

- Whenever an internally simulated machine of \mathcal{F} produces empty output, M' produces empty output.
- Whenever an internally simulated machine of \mathcal{F} produces output on its I/O interface to \mathcal{E} , M' produces empty output; as discussed above, this will happen with only negligible probability.
- Whenever an internally simulated machine of \mathcal{F} produces output m on an internal tape to another machine in \mathcal{F} , M' internally simulates the receiving machine with input m .
- Whenever an internally simulated machine of \mathcal{F} outputs m on its network interface (to M'), M' internally simulates M with input m .

From the above discussion and by (14), it easily follows that

$$\forall \mathcal{E} \in \text{Env}_{\sigma\text{-single}}(\mathcal{P}): \mathcal{E} | \mathcal{P} \equiv \mathcal{E} | \mathcal{S} | \mathcal{F} \equiv \mathcal{E} | \mathcal{S}' | \mathcal{F} . \quad (15)$$

Now, since $\mathcal{S} | \mathcal{F}$ is σ -environmentally bounded, it is easy to see that $\mathcal{S}' | \mathcal{F}$ is σ -environmentally bounded too. With $\mathcal{S} \in \text{Sim}_{\sigma\text{-single}}^{\mathcal{P}}(\mathcal{F})$ we thus obtain $\mathcal{S}' \in \text{Sim}_{\sigma\text{-single}}^{\mathcal{P}}(\mathcal{F})$. Also, by construction we have that \mathcal{S}' is σ -complete. \square

Theorem 10 now follows directly from the above lemma and Theorem 6:

Proof of Theorem 10. By Lemma 14, there exists $\mathcal{S} \in \text{Sim}_{\sigma\text{-single}}^{\mathcal{P}}(\mathcal{F})$ such that \mathcal{S} is a σ -complete IITM and $\mathcal{P} \cong_{\sigma\text{-single}} \mathcal{S} | \mathcal{F}$. Since \mathcal{S} is σ -complete, we can conclude that $\mathcal{P} \cong_{\sigma\text{-single}} !\mathcal{S} | \mathcal{F}$: since $\mathcal{E} \in \text{Env}_{\sigma\text{-single}}(\mathcal{P})$ invokes only a single session, i.e., sends only messages to \mathcal{S} and \mathcal{F} with the same SID (w.r.t. σ) and \mathcal{F} is a σ -session version, \mathcal{S} receives only messages with the same SID (w.r.t. σ) from \mathcal{F} and \mathcal{E} . So, even with $!\mathcal{S}$ only one instance of \mathcal{S} will be invoked in every run of $\mathcal{E} | !\mathcal{S} | \mathcal{F}$ and we have $\mathcal{P} \cong_{\sigma\text{-single}} !\mathcal{S} | \mathcal{F}$. Now, it is easy to see that $!\mathcal{S} | \mathcal{F}$ is a protocol system (in particular, Condition (ii) of Definition 11 is satisfied) and a σ -session version. Furthermore, $!\mathcal{S} | \mathcal{F}$ is σ -environmentally bounded because $\mathcal{S} | \mathcal{F}$ is σ -environmentally bounded by definition of $\text{Sim}_{\sigma\text{-single}}^{\mathcal{P}}(\mathcal{F})$. From this, by Theorem 6 (with $\mathcal{Q} = !\mathcal{S} | \mathcal{F}$), we obtain that $\mathcal{P} \cong !\mathcal{S} | \mathcal{F}$ and that $!\mathcal{S} | \mathcal{F}$ is environmentally bounded. The latter implies $!\mathcal{S} \in \text{Sim}^{\mathcal{P}}(\mathcal{F})$. Hence, we conclude that $\mathcal{P} \leq^{SS} \mathcal{F}$. \square

7.4 Composition Theorem for More Complex Systems

By iteratively applying Theorems 8, 9, and 10, one can construct more and more complex systems. For example, as an immediate consequence of Theorem 8 and 9 we obtain that if (an unbounded number of sessions of) an ideal protocol \mathcal{F} is used as a component in a more complex system \mathcal{Q} , then it can be replaced by its realization \mathcal{P} :

Corollary 1. *Let \mathcal{Q} , \mathcal{P} , and \mathcal{F} be protocol systems such that the following conditions are satisfied:*

1. $!\mathcal{P}$ is environmentally bounded and $\mathcal{P} \leq^{SS} \mathcal{F}$.
2. \mathcal{Q} and \mathcal{P} are I/O-connectable and $\mathcal{Q} | !\mathcal{P}$ is environmentally bounded.

Then, $\mathcal{Q} | !\mathcal{P} \leq^{SS} \mathcal{Q} | !\mathcal{F}$.

If we want \mathcal{P} and \mathcal{F} to be aware of their SIDs, we can use Theorem 10 instead of Theorem 9 and obtain the following corollary.

Corollary 2. *Let \mathcal{Q} , \mathcal{P} , and \mathcal{F} be protocol systems such that \mathcal{P} and \mathcal{F} are σ -session versions for some SID function σ and the following conditions are satisfied:*

1. $!\mathcal{P}$ is environmentally bounded and $\mathcal{P} \leq_{\sigma\text{-single}}^{SS} \mathcal{F}$.
2. \mathcal{Q} and \mathcal{P} are I/O-connectable and $\mathcal{Q} | \mathcal{P}$ is environmentally bounded.

Then, $\mathcal{Q} | \mathcal{P} \leq^{SS} \mathcal{Q} | \mathcal{F}$.

Clearly, by iteratively applying Theorems 8, 9, and 10 further, even more complex system can be constructed.

8 On the Composability of Runtime Notions

In this section, we discuss the composability of environmentally (almost/strictly) bounded systems. The first simple observation is that, in general, the composition of two environmentally (almost/strictly) bounded systems does not need to be environmentally bounded: For example, consider two environmentally (almost/strictly) bounded protocol systems \mathcal{S}_1 and \mathcal{S}_2 which connect via some external tapes. (The simplest example is that both \mathcal{S}_1 and \mathcal{S}_2 are single IITMs.) Then, \mathcal{S}_1 and \mathcal{S}_2 could “play ping-pong” with each other, i.e., they could send messages back and forth forever, and hence, in such a case they are not environmentally bounded.

However, in applications the composition of environmentally almost/strictly bounded systems is basically always environmentally almost/strictly bounded. Clearly, one can construct examples, such as the above, where this is not the case. However, these examples typically do not occur in applications. Moreover, in applications it is typically easy to see whether a system, including the composition of two environmentally almost/strictly bounded systems, is environmentally almost/strictly bounded.

We also observe that in applications protocol systems are typically *strictly* bounded, and for such systems we obtain useful general composability statements, as presented in the following subsection. In Section 8.2 we show that some of these general statements do not hold true for environmentally almost bounded systems.

8.1 On the Composability of Environmentally Strictly Bounded Systems

The following lemma shows that the composition of two environmentally strictly bounded systems that have disjoint tapes, and hence, do not communicate directly, is environmentally strictly bounded.

Lemma 15. *Let \mathcal{P} and \mathcal{Q} be two environmentally strictly bounded protocol systems such that the sets of external tapes of \mathcal{P} and \mathcal{Q} are disjoint. Then, $\mathcal{P} \mid \mathcal{Q}$ is environmentally strictly bounded.*

Proof. Let \mathcal{E} be a universally bounded system which can be connected to $\mathcal{P} \mid \mathcal{Q}$. We need to show that $\mathcal{E} \mid \mathcal{P} \mid \mathcal{Q}$ is strictly bounded.

First, we observe that because \mathcal{E} is universally bounded there exists a polynomial p_0 such that for all $\eta \in \mathbf{N}$ and $a \in \{0, 1\}^*$:

$$\text{Time}_{\mathcal{E}}((\mathcal{E} \mid \mathcal{P} \mid \mathcal{Q})(1^\eta, a)) \leq p_0(\eta + |a|) \quad .^{18} \quad (16)$$

In particular, the length of the overall output produced by \mathcal{E} in a run of $(\mathcal{E} \mid \mathcal{P} \mid \mathcal{Q})(1^\eta, a)$ is bounded from above by $p_0(\eta + |a|)$. We thus can construct an environmental system $\mathcal{E}' \in \text{Env}(\mathcal{P})$ that upon every activation produces random output such that for every sequence of messages sent from \mathcal{E} to \mathcal{P} in a run of $(\mathcal{E} \mid \mathcal{P} \mid \mathcal{Q})(1^\eta, a)$ the probability that this sequence of messages is sent from \mathcal{E}' to \mathcal{P} in a run of $(\mathcal{E}' \mid \mathcal{P})(1^\eta, a)$ is non-zero. (Here we need that \mathcal{E}' knows the length of a . This is true since \mathcal{P} and \mathcal{Q} are protocol systems, and hence, $\text{start} \notin \mathcal{T}(\mathcal{P} \mid \mathcal{Q})$ and \mathcal{E}' may use start .)

Clearly, \mathcal{E}' is universally bounded. Now, since \mathcal{P} is environmentally strictly bounded, there exists a polynomial p_1 such that for all $\eta \in \mathbf{N}$ and $a \in \{0, 1\}^*$:

$$\text{Time}_{\mathcal{P}}((\mathcal{E}' \mid \mathcal{P})(1^\eta, a)) \leq p_1(\eta + |a|) \quad .$$

Since with non-zero probability \mathcal{E}' (running with \mathcal{P}) sends the same sequence of messages to \mathcal{P} as \mathcal{E} (running with $\mathcal{P} \mid \mathcal{Q}$), we deduce that for all $\eta \in \mathbf{N}$ and $a \in \{0, 1\}^*$:

$$\text{Time}_{\mathcal{P}}((\mathcal{E} \mid \mathcal{P} \mid \mathcal{Q})(1^\eta, a)) \leq p_1(\eta + |a|) \quad . \quad (17)$$

Analogously, there exists a polynomial p_2 such that for all $\eta \in \mathbf{N}$ and $a \in \{0, 1\}^*$:

$$\text{Time}_{\mathcal{Q}}((\mathcal{E} \mid \mathcal{P} \mid \mathcal{Q})(1^\eta, a)) \leq p_2(\eta + |a|) \quad . \quad (18)$$

Combining (16), (17), and (18), we obtain that for all $\eta \in \mathbf{N}$ and $a \in \{0, 1\}^*$:

$$\text{Time}((\mathcal{E} \mid \mathcal{P} \mid \mathcal{Q})(1^\eta, a)) \leq p_0(\eta + |a|) + p_1(\eta + |a|) + p_2(\eta + |a|) \quad .$$

Thus, $\mathcal{E} \mid \mathcal{P} \mid \mathcal{Q}$ is strictly bounded. □

¹⁸That is, $\text{Time}_{\mathcal{E}}((\mathcal{E} \mid \mathcal{P} \mid \mathcal{Q})(1^\eta, a))(\alpha) \leq p_0(\eta + |a|)$ for all random coins $\alpha \in \text{Rand}$.

The above lemma can be generalized to protocol systems \mathcal{P} and \mathcal{Q} whose sets of external tapes are not disjoint, provided that the length of the messages that, say, \mathcal{P} sends to \mathcal{Q} is bounded by a polynomial in the security parameter plus the external input (of the environment) plus the input that \mathcal{P} gets from the environment. To state this precisely, let \mathcal{S} , \mathcal{S}' , and \mathcal{S}'' be connectable systems. Then, we define $\text{Flow}_{\mathcal{S} \rightarrow \mathcal{S}'}((\mathcal{S} | \mathcal{S}' | \mathcal{S}'')(1^\eta, a))$ to be the random variable over Rand which assigns to every $\alpha \in \text{Rand}$ the overall length of all messages that were sent from \mathcal{S} to \mathcal{S}' (i.e., the output by \mathcal{S} on input tapes of \mathcal{S}') in the run of $\mathcal{S} | \mathcal{S}' | \mathcal{S}''$ with security parameter η , external input a , and random coins α . Here we count all messages even messages that are rejected in mode `CheckAddress`. (Note that if the set of external tapes of \mathcal{S} and \mathcal{S}' are disjoint, then the flow from \mathcal{S} and \mathcal{S}' is empty.) Now we can prove the following generalization of Lemma 15.

Lemma 16. *Let \mathcal{P} and \mathcal{Q} be two environmentally strictly bounded protocol systems that are connectable and such that there exists a polynomial p such that for all $\mathcal{E} \in \text{Env}(\mathcal{P} | \mathcal{Q})$:*

$$\text{Prob} [\text{Flow}_{\mathcal{P} \rightarrow \mathcal{Q}}((\mathcal{E} | \mathcal{P} | \mathcal{Q})(1^\eta, a)) \leq p(\eta + |a| + \text{Flow}_{\mathcal{E} \rightarrow \mathcal{P}}((\mathcal{E} | \mathcal{P} | \mathcal{Q})(1^\eta, a)))] = 1$$

(i.e., the output from \mathcal{P} to \mathcal{Q} is bounded in length by a polynomial in the security parameter plus the length of the external input plus the length of the input \mathcal{P} received from \mathcal{E}). Then, $\mathcal{P} | \mathcal{Q}$ is environmentally strictly bounded.

Proof. Let \mathcal{E} and p be as in the assumption of the lemma. We may assume that p has only non-negative coefficients. Hence, $p(n) \leq p(n')$ for all $n, n' \in \mathbf{N}$ such that $n \leq n'$. Since \mathcal{E} is universally bounded, we know that there exists a polynomial p' such that $\text{Flow}_{\mathcal{E} \rightarrow \mathcal{P}}((\mathcal{E} | \mathcal{P} | \mathcal{Q})(1^\eta, a)) \leq p'(\eta + |a|)$ (i.e., the probability for this is 1) and such that $\text{Flow}_{\mathcal{E} \rightarrow \mathcal{Q}}((\mathcal{E} | \mathcal{P} | \mathcal{Q})(1^\eta, a)) \leq p'(\eta + |a|)$. Hence, by assumption, $\text{Flow}_{\mathcal{P} \rightarrow \mathcal{Q}}((\mathcal{E} | \mathcal{P} | \mathcal{Q})(1^\eta, a)) \leq p(\eta + |a| + p'(\eta + |a|))$. So, $\text{Flow}_{\mathcal{E} | \mathcal{P} \rightarrow \mathcal{Q}}((\mathcal{E} | \mathcal{P} | \mathcal{Q})(1^\eta, a)) \leq p'(\eta + |a|) + p(\eta + |a| + p'(\eta + |a|))$. As in the proof of Lemma 15 and since \mathcal{Q} is environmentally strictly bounded, we can deduce that there exists a polynomial q such that $\text{Time}_{\mathcal{Q}}((\mathcal{E} | \mathcal{P} | \mathcal{Q})(1^\eta, a)) \leq q(\eta + |a|)$. Hence, $\text{Flow}_{\mathcal{E} | \mathcal{Q} \rightarrow \mathcal{P}}((\mathcal{E} | \mathcal{P} | \mathcal{Q})(1^\eta, a)) \leq p'(\eta + |a|) + q(\eta + |a|)$. Now, again since \mathcal{P} is environmentally strictly bounded and as in the proof of Lemma 15, we can conclude that there exists a polynomial q' such that $\text{Time}_{\mathcal{P}}((\mathcal{E} | \mathcal{P} | \mathcal{Q})(1^\eta, a)) \leq q'(\eta + |a|)$. So, $\text{Time}_{\mathcal{Q} | \mathcal{P}}((\mathcal{E} | \mathcal{P} | \mathcal{Q})(1^\eta, a)) \leq q(\eta + |a|) + q'(\eta + |a|)$. It follows that $\mathcal{P} | \mathcal{Q}$ is environmentally strictly bounded. \square

The condition stated in Lemma 16 is often satisfied. For instance, if \mathcal{P} and \mathcal{Q} do not have common external tapes or in many settings where \mathcal{Q} is a subprotocol of \mathcal{P} . However, we emphasize that even if this condition is not satisfied, it is generally easy to see whether the composition of two environmentally strictly bounded systems is environmentally strictly bounded. In fact, as already mentioned, we would argue that for natural protocol systems it basically never happens that the composition is not environmentally strictly bounded.

The next lemma shows that the multi-session version of an environmentally strictly bounded system is environmentally strictly bounded.

Lemma 17. *Let \mathcal{S} be an environmentally strictly bounded protocol system. Then, $!\underline{\mathcal{S}}$ is environmentally strictly bounded.*

Proof. First, let $\mathcal{E} \in \text{Env}(!\underline{\mathcal{S}})$ be σ_{prefix} -single session, i.e., the environment \mathcal{E} invokes only a single session of $\underline{\mathcal{S}}$. Let $\mathcal{E}' \in \text{Env}(\mathcal{S})$ simulate \mathcal{E} except that it strips off the SID with which \mathcal{E} prefixes messages. Since \mathcal{E} is universally bounded, \mathcal{E}' is universally bounded too. Then it is easy to construct a bijection between runs of $\mathcal{E} | !\underline{\mathcal{S}}$ and $\mathcal{E}' | \mathcal{S}$ such that corresponding runs produce the same output and have the same probability of occurring.¹⁹ Since, by assumption, \mathcal{S} is environmentally strictly bounded, there exists a polynomial p in the security parameter plus the length of the external input such that for all runs of $\mathcal{E}' | \mathcal{S}$ the number of steps taken by \mathcal{S} is bounded by p . It follows that there exists a polynomial p' such that for every run of $\mathcal{E} | !\underline{\mathcal{S}}$ the

¹⁹Here we use that \mathcal{S} is a protocol system, in particular Condition (ii) of Definition 11. Without this condition, there could be several copies of a machine in $!\underline{\mathcal{S}}$, whereas in \mathcal{S} there is only one such copy. We also use that \mathcal{S} does not contain the tape `start` since the content of this tape would be interpreted differently by $\underline{\mathcal{S}}$ and \mathcal{S} .

number of steps taken by the session invoked by \mathcal{E} (i.e., the overall number of steps taken by machines in mode `Compute` belonging to that session) in this run is bounded p' .

Now, let $\mathcal{E} \in \text{Env}(!\underline{\mathcal{S}})$, where we do not require \mathcal{E} to be σ_{prefix} -single session. Since \mathcal{E} is universally bounded, there exists a polynomial q (in the security parameter plus the length of the external input) such that the overall number of steps taken by \mathcal{E} (in any run with any system) is bounded by q . In particular, the length of messages sent by \mathcal{E} to $!\underline{\mathcal{S}}$ is bounded by q . Hence, similarly to the proof of Lemma 15, we can construct a universally bounded environment $\mathcal{E}' \in \text{Env}(!\underline{\mathcal{S}})$ which randomly picks an SID sid (one bounded in length by q) and randomly picks at most q messages of length at most q and sends these messages prefixed by sid to $!\underline{\mathcal{S}}$. By construction, \mathcal{E}' is a σ_{prefix} -single session environment. Hence, by the above, we know that there exists a polynomial p such that for every run of $\mathcal{E}' | !\underline{\mathcal{S}}$ the number of steps taken by the session invoked by \mathcal{E}' in this run is bounded by p . Now, consider a run of $\mathcal{E} | !\underline{\mathcal{S}}$ and, for some SID, all machines with this SID in this run. The probability that this SID is picked by \mathcal{E}' and that \mathcal{E}' produces exactly the same input for the machines with that SID is non-zero. Hence, it follows that the number of steps taken by these machines in the considered run of $\mathcal{E} | !\underline{\mathcal{S}}$ is bounded by p .

Overall, we obtain that for all $\eta \in \mathbf{N}$ and $a \in \{0, 1\}^*$:

$$\text{Time}((\mathcal{E} | !\underline{\mathcal{S}})(1^\eta, a)) \leq q(\eta + |a|) + q(\eta + |a|) \cdot p(\eta + |a|) .$$

Hence, $\mathcal{E} | !\underline{\mathcal{S}}$ is strictly bounded. □

Clearly, the converse of Lemma 17 is also true, as stated in the next lemma. Furthermore, the multi-session version of the multi-session version of \mathcal{S} , i.e., $!\underline{\underline{\mathcal{S}}}$ —in this system all messages are prefixed by two SIDs, i.e., they are of the form $m = (s_1, (s_2, m'))$ —is environmentally bounded if $!\underline{\mathcal{S}}$ is:

Lemma 18. *Let \mathcal{S} be a protocol system such that $!\underline{\mathcal{S}}$ is environmentally strictly bounded. Then, \mathcal{S} and $!\underline{\underline{\mathcal{S}}}$ are environmentally strictly bounded.*

The proof of this lemma is straightforward and therefore omitted.

8.2 On the Composability of Environmentally Almost Bounded Systems

We now show that the Lemmas 15 and 17 (and hence, Lemma 16 because it is a generalization of Lemma 15) do not hold for environmentally almost bounded protocol systems in general. However, as mentioned, for applications this is not really relevant since protocol systems are typically strictly bounded and the composition of such systems typically is strictly bounded as well.

For the following lemma, we need the existence of so-called time-lock puzzles, a complexity assumption introduced in [28]. We define time-lock puzzles following [16]:²⁰

Definition 21. A *time-lock puzzle* consists of an ITM V (the *verifier*) and an ITM P (the *prover*) such that the following conditions are satisfied, where by $\langle P, V \rangle$ we denote the distribution of the output of V after an interaction with P :

1. Given an argument of the form $(1^\eta, s)$, V runs in polynomial time in η . Given an argument of the form $(1^\eta, s)$, P runs in polynomial time in $\eta + s$.
2. *Easiness.* For every polynomial p we have that

$$\min_{s \leq p(\eta)} \text{Prob}[\langle P(1^\eta, s), V(1^\eta, s) \rangle = 1]$$

is overwhelming (as a function in η). (We call s the *hardness* of the puzzle.)

²⁰We note that in [16], for the hardness assumption, the runtime of B does not depend on the length of the external input a , but only on the security parameter η . However, since here the runtime of environments depends on the length of a and, in the proof of the following lemma, we need that environments play the role of B , we adapted this notion. For uniform environments one could drop the external input altogether.

3. *Hardness.* For any ITM B running in polynomial time in the length of its first two arguments (i.e., in $\eta + |a|$) there exists a polynomial p such that

$$\sup_{s \geq p(\eta + |a|)} \text{Prob} [\langle B(1^\eta, a, s), V(1^\eta, s) \rangle = 1]$$

is negligible (as a function in η and a).

Lemma 19. *If time-lock puzzles exist, then the following statements are true.*

1. *There exist protocol systems \mathcal{P} and \mathcal{Q} with disjoint sets of external tapes such that \mathcal{P} and \mathcal{Q} are environmentally bounded but their composition $\mathcal{P} \mid \mathcal{Q}$ is not environmentally bounded.*
2. *There exists a protocol system \mathcal{S} such that \mathcal{S} is environmentally bounded but $!\mathcal{S}$ is not environmentally bounded.*

Proof. Hofheinz et al. [16, footnote 21 on page 34] sketched that if time-lock puzzles exist, then there exist environmentally bounded protocol systems \mathcal{P} and \mathcal{Q} with disjoint sets of external tapes such that \mathcal{P} and \mathcal{Q} are environmentally bounded but $!\mathcal{P} \mid !\mathcal{Q}$ is not environmentally bounded. Their idea can be used as follows to show that there even exist environmentally bounded systems \mathcal{P} and \mathcal{Q} such that $\mathcal{P} \mid \mathcal{Q}$ is not environmentally bounded.

Let (V_A, P_A) and (V_B, P_B) be two time-lock puzzles (with verifiers V_A/V_B and provers P_A/P_B). We define \mathcal{P} and \mathcal{Q} to have disjoint tapes. Also both systems are defined to have one input and one output tape and they accept every message in mode `CheckAddress`. Upon input s (this is the hardness of the puzzle), \mathcal{P} simulates $V_A(1^\eta, s)$ (if V_A sends a message to P_A , \mathcal{P} outputs this message on its output tape and if V_A wants to receive a message, \mathcal{P} waits for input on its input tape and then gives this input to V_A) until V_A halts and outputs a bit b . If $b = 0$ (i.e., the verifier does not accept), then \mathcal{P} terminates. Otherwise, \mathcal{P} continues and simulates $P_B(1^\eta, 2s)$ (here $2s$ means “2 times s ” where s is interpreted as a non-negative integer). When P_B halts, \mathcal{P} goes back to the start and again waits for input s , to simulate $V_A(1^\eta, s)$ and so on. \mathcal{Q} is defined as \mathcal{P} except that the puzzles A and B are swapped, i.e., \mathcal{Q} simulates V_B and P_A instead of V_A and P_B . So, \mathcal{P} (\mathcal{Q}) verifies a puzzle of type A (B) of hardness s and then solves a puzzle of type B (A) of hardness $2s$.

We now show that \mathcal{P} and \mathcal{Q} (in separation) are environmentally bounded. For this purpose, let \mathcal{E} be an environment of \mathcal{P} . Since \mathcal{E} is universally bounded, by the hardness assumption of the puzzles, there exists a polynomial p_1 such that \mathcal{E} can only solve puzzles of hardness $s \leq p_1(\eta + |a|)$ (except with negligible probability), where η is the security parameter and a is the external input to \mathcal{E} . Let s_1, \dots, s_n be the hardnesses that \mathcal{E} sends to \mathcal{P} in a run of $(\mathcal{E} \mid \mathcal{P})(1^\eta, a)$. Note that in every run $n \leq p_2(\eta + |a|)$, where p_2 is a polynomial that bounds the number of steps taken by \mathcal{E} . By definition of time-lock puzzles, there exist polynomials p_3 and p_4 such that, for all $i \leq n$, the runtime of (the simulated) $V_A(1^\eta, s_i)$ is bounded by $p_3(\eta)$ and the runtime of $P_B(1^\eta, s_i)$ is bounded by $p_4(\eta + s_i)$. We may assume that p_4 has only non-negative coefficients. Hence, $p_4(l) \leq p_4(l')$ for all $l, l' \in \mathbf{N}$ with $l \leq l'$. As argued above, the simulated $V_A(1^\eta, s_i)$, $i \leq n$, accepts (i.e., outputs 1) only if $s_i \leq p_1(\eta + |a|)$ (except with negligible probability). Hence, $P_B(1^\eta, 2s_i)$, $i \leq n$, is only simulated if $s_i \leq p_1(\eta + |a|)$. We obtain that the overall runtime of \mathcal{P} is bounded by

$$\sum_{i=1}^n p_3(\eta) + \sum_{i=1}^n p_4(\eta + 2s_i) \leq p_2(\eta + |a|) \cdot p_3(\eta) + p_2(\eta + |a|) \cdot p_4(\eta + 2p_1(\eta + |a|))$$

(except with negligible probability). Hence, \mathcal{P} is environmentally bounded. Analogously, \mathcal{Q} is environmentally bounded.

Now, we show that the system $\mathcal{P} \mid \mathcal{Q}$ is not environmentally bounded. Therefore, we construct an environment \mathcal{E} as follows. Basically \mathcal{E} will start with a puzzle of hardness 1 for \mathcal{P} , solves it itself and then forwards all messages between \mathcal{P} and \mathcal{Q} for η many runs of the puzzles. At the end, a puzzle of hardness 2^η is solved. Now we describe \mathcal{E} in more detail. First, \mathcal{E} sends hardness 1 to \mathcal{P} and simulates $P_A(1^\eta, 1)$ with \mathcal{P} . So, $V_A(1^\eta, 1)$, simulated by \mathcal{P} , accepts (except with negligible probability), and hence, \mathcal{P} starts the simulation of $P_B(1^\eta, 2)$. Then, \mathcal{E} sends hardness 2 to \mathcal{Q} , which then simulates $V_B(1^\eta, 2)$, where \mathcal{E} forwards all

messages between \mathcal{P} (and hence, $P_B(1^\eta, 2)$) and \mathcal{Q} (and hence, $V_B(1^\eta, 2)$). As a result, $V_B(1^\eta, 2)$, simulated by \mathcal{Q} , accepts, and hence, \mathcal{Q} starts the simulation of $P_A(1^\eta, 4)$. \mathcal{E} sends hardness 4 to \mathcal{P} and again forwards all messages between \mathcal{P} and \mathcal{Q} . This is iterated by \mathcal{E} η times, always doubling the hardness. At the end, a puzzle of hardness 2^η is solved, so, except with negligible probability, the overall runtime of $\mathcal{E}|\mathcal{P}|\mathcal{Q}$ can not be polynomially bounded in $\eta + |a|$ (by the hardness assumption of the puzzles). Thus, $\mathcal{E}|\mathcal{P}|\mathcal{Q}$ is not almost bounded. Since clearly \mathcal{E} is universally bounded, this means that $\mathcal{P}|\mathcal{Q}$ is not environmentally bounded. This proves the first statement of the lemma.

We now prove the second statement using the first one. Let \mathcal{P} and \mathcal{Q} be protocol systems as above. That is, \mathcal{P} and \mathcal{Q} are environmentally bounded but there exists an environment \mathcal{E} of $\mathcal{P}|\mathcal{Q}$ such that $\mathcal{E}|\mathcal{P}|\mathcal{Q}$ is not almost bounded. We define the protocol system \mathcal{S} as follows: \mathcal{S} chooses a bit b uniformly at random. If $b = 0$, \mathcal{S} behaves exactly like \mathcal{P} and otherwise it behaves exactly like \mathcal{Q} . (The tape names have to be adjusted appropriately.) Now, \mathcal{S} is environmentally bounded because \mathcal{P} and \mathcal{Q} are environmentally bounded, and \mathcal{S} runs either only \mathcal{P} or only \mathcal{Q} . However, $!\underline{\mathcal{S}}$ is not environmentally bounded: We construct an environment \mathcal{E}' of $!\underline{\mathcal{S}}$ as follows: \mathcal{E}' uses two sessions of $\underline{\mathcal{S}}$ and simulates \mathcal{E} where messages exchanged with \mathcal{P}/\mathcal{Q} are now exchanged with the first/second session of $\underline{\mathcal{S}}$, where now SIDs are used to address these sessions. Clearly, \mathcal{E}' is universally bounded because \mathcal{E} is universally bounded. Now, with probability $\frac{1}{4}$, the system $\mathcal{E}'|!\underline{\mathcal{S}}$ behaves exactly as $\mathcal{E}|\mathcal{P}|\mathcal{Q}$, and hence, $\mathcal{E}'|!\underline{\mathcal{S}}$ is not almost bounded. \square

We remark that the lemma holds both for uniform and non-uniform environments. Note that the environments constructed in the above proof are uniform, i.e., they ignore the external input.

The following lemma states that Lemma 18 also holds for environmentally almost bounded systems. Again, this lemma is easy to prove.

Lemma 20. *Let \mathcal{S} be a protocol system such that $!\underline{\mathcal{S}}$ is environmentally almost bounded. Then, \mathcal{S} and $!\underline{\mathcal{S}}$ are environmentally almost bounded.*

9 On Basing Universal Composability on Environmentally Strictly Bounded Systems

As already mentioned before, in applications the systems one has to deal with are typically environmentally strictly bounded. Also, as shown in Section 8, environmentally strictly bounded systems enjoy useful composability properties. Therefore, a question that suggests itself is whether we obtain useful notions of universal composability if we formulate these notions based on environmentally strictly bounded systems, instead of environmentally almost bounded systems: we would assume the protocol system \mathcal{P} to be environmentally strictly bounded and, for strong simulatability, we would require $\mathcal{S}|\mathcal{F}$ to be environmentally strictly bounded; similarly, the other security notions would be adapted.

In this section, we show that basing the security notions on environmentally *strictly* bounded systems, rather than environmentally almost bounded systems, yields unsuitable security notions. More specifically, we show the following properties:

1. The composition theorem for a constant number of protocol systems (the analog of Theorem 8) would not hold for any of the new security notions.
2. Strong simulatability and dummy UC would not be transitive relationships anymore.
3. Strong simulatability and dummy UC would not imply UC anymore.

We note, however, that the composition theorem for the unbounded self-composition of systems (the analog of Theorem 9) would still hold. The proof would even be simpler: due to Lemma 17 it would follow immediately that $!(\mathcal{S}|\mathcal{F})$ and all hybrid systems considered in the proof of this theorem are environmentally strictly bounded. In the case of environmentally almost bounded systems this required a more tailored and involved proof, namely, the proof of Lemma 9.

We further remark that Hofheinz et al. [16] considered a security definition conceptually similar to our strict variant of UC (they do not consider strict variants of SS or dummy UC) and that they showed that, for their notion, the universal composition theorem in the UC model does not hold. The counterexample used in their proof could also be used to show that the composition theorem does not hold for our strict variant of UC. However, the counterexample we use below is simpler and more general. In particular, it can be used to show that strict dummy UC and strict strong simulatability are unsuitable security notions too, a fact that could not have been proven with their example.

In what follows, we first introduce the security notions based on environmentally strictly bounded systems and then prove the negative results mentioned above.

9.1 Strict Simulatability

In order to define the security notions based on environmentally strictly bounded systems, we first need to introduce some notation.

In Section 6.2, we required that the system composed of an adversary/simulator and a protocol system is environmentally almost bounded. Now, we require that this composition is environmentally *strictly* bounded. Therefore, we introduce the sets $\text{Adv}_{\text{strict}}(\mathcal{S})$ and $\text{Sim}_{\text{strict}}^{\mathcal{S}'}(\mathcal{S})$, which are defined just as $\text{Adv}(\mathcal{S})$ and $\text{Sim}^{\mathcal{S}'}(\mathcal{S})$ (see Definition 18), respectively, except that for $\mathcal{A} \in \text{Adv}_{\text{strict}}(\mathcal{S})$ or $\mathcal{A} \in \text{Sim}_{\text{strict}}^{\mathcal{S}'}(\mathcal{S})$ we now require that $\mathcal{A}|\mathcal{S}$ is environmentally strictly bounded. We emphasize that, as before, \mathcal{A} itself is not required to be environmentally (strictly) bounded.

Definition 22. Let \mathcal{P} and \mathcal{F} be protocol systems such that \mathcal{P} is environmentally strictly bounded.

1. *Strict Strong Simulatability (strict SS):* $\mathcal{P} \leq_{\text{strict}}^{\text{SS}} \mathcal{F}$ iff $\exists \mathcal{S} \in \text{Sim}_{\text{strict}}^{\mathcal{P}}(\mathcal{F})$: $\mathcal{P} \cong \mathcal{S}|\mathcal{F}$.
2. *Strict Universal Simulatability/Composability (strict UC):* $\mathcal{P} \leq_{\text{strict}}^{\text{UC}} \mathcal{F}$ iff $\forall \mathcal{A} \in \text{Adv}_{\text{strict}}(\mathcal{P}) \exists \mathcal{I} \in \text{Sim}_{\text{strict}}^{\mathcal{A}|\mathcal{P}}(\mathcal{F})$: $\mathcal{A}|\mathcal{P} \cong \mathcal{I}|\mathcal{F}$.
3. *Dummy Version of strict UC (strict dummyUC):* $\mathcal{P} \leq_{\text{strict}}^{\text{dumUC}} \mathcal{F}$ iff $\exists \mathcal{I} \in \text{Sim}_{\text{strict}}^{\mathcal{D}|\mathcal{P}}(\mathcal{F})$: $\mathcal{D}|\mathcal{P} \cong \mathcal{I}|\mathcal{F}$ where $\mathcal{D} = \mathcal{D}(\mathcal{T}_{\text{in}}^{\text{net}}(\mathcal{P}), \mathcal{T}_{\text{out}}^{\text{net}}(\mathcal{P}))$.

Analogously, we could define strict variants of black-box and reactive simulatability. However, from the results presented for above notions it can be easily seen that the strict variants of black-box and reactive simulatability yield unsuitable security notions as well.

Using Lemmas 4 and 5, it is easy to see that strict SS is equivalent to strict dummyUC and that strict UC implies strict SS (and hence, strict dummyUC):

Lemma 21. *Let \mathcal{P} and \mathcal{F} be protocol systems such that \mathcal{P} is environmentally strictly bounded. Then, $\mathcal{P} \leq_{\text{strict}}^{\text{SS}} \mathcal{F}$ if and only if $\mathcal{P} \leq_{\text{strict}}^{\text{dumUC}} \mathcal{F}$ and $\mathcal{P} \leq_{\text{strict}}^{\text{SS}} \mathcal{F}$ if $\mathcal{P} \leq_{\text{strict}}^{\text{UC}} \mathcal{F}$.*

Since strict dummyUC and strict SS are equivalent, in what follows, we will discuss only strict SS and strict UC.

9.2 No Universal Composability for a Constant Number of Protocol Systems

We now show that the analog of Theorem 8, the composition theorem for a constant number of protocol systems, does not hold for strict SS and strict UC:

Lemma 22. *There exist protocol systems \mathcal{P}_1 , \mathcal{P}_2 , \mathcal{F}_1 , and \mathcal{F}_2 that satisfy the following conditions:*

1. \mathcal{P}_1 and \mathcal{P}_2 as well as \mathcal{F}_1 and \mathcal{F}_2 are I/O-connectable, respectively.
2. \mathcal{P}_1 , \mathcal{P}_2 , $\mathcal{P}_1|\mathcal{P}_2$, \mathcal{F}_1 , \mathcal{F}_2 , and $\mathcal{F}_1|\mathcal{F}_2$ are environmentally strictly bounded.
3. For all $i \in \{1, 2\}$: $\mathcal{P}_i \leq_{\text{strict}}^{\text{SS}} \mathcal{F}_i$ and $\mathcal{P}_i \leq_{\text{strict}}^{\text{UC}} \mathcal{F}_i$.




Figure 8: IITMs constructed in the proof of Lemma 22.

4. $\mathcal{P}_1 | \mathcal{P}_2 \not\leq_{\text{strict}}^{SS} \mathcal{F}_1 | \mathcal{F}_2$ and $\mathcal{P}_1 | \mathcal{P}_2 \not\leq_{\text{strict}}^{UC} \mathcal{F}_1 | \mathcal{F}_2$.

Proof. We define the systems \mathcal{P}_1 , \mathcal{P}_2 , \mathcal{F}_1 , and \mathcal{F}_2 as follows; see Figure 8 for an illustration of the systems. Every system is a single IITM. The systems \mathcal{P}_2 and \mathcal{F}_2 have the same I/O interface which consists of an I/O input tape c and an I/O output tape \bar{c} . The system \mathcal{F}_2 additionally has a network input tape and a network output tape but \mathcal{P}_2 does not have any network tapes. The systems \mathcal{P}_1 and \mathcal{F}_1 have no network tapes and they have the same I/O interface which consists of two I/O input tapes \bar{c} and c' and two I/O output tapes c and \bar{c}' . That is, with the tapes c and \bar{c} , they connect to $\mathcal{P}_2/\mathcal{F}_2$. Clearly, \mathcal{P}_1 and \mathcal{P}_2 as well as \mathcal{F}_1 and \mathcal{F}_2 are I/O-connectable, respectively. These systems are defined to accept all messages in mode **CheckAddress**.

In mode **Compute** the machines act as follows: When receiving a message m on c' (resp. \bar{c}) the machine \mathcal{P}_1 outputs m on c (resp. \bar{c}'). When receiving a message m on c the machine \mathcal{P}_2 outputs m on \bar{c} . When receiving a message m on the I/O tape c , the machine \mathcal{F}_2 outputs m on its network tape. When receiving a message m on the network tape, m is output on \bar{c} . Note that \mathcal{P}_1 , \mathcal{P}_2 , and \mathcal{F}_2 are simple dummy IITMs which merely forward all messages they receive. The machine \mathcal{F}_1 also only forwards messages but to which tape a message is forwarded depends on a random choice. More specifically, \mathcal{F}_1 , upon its first activation in mode **Compute**, first chooses $r \in \{1, \dots, 2^\eta\}$ uniformly at random (where η is the security parameter). If $r \neq 1$, then \mathcal{F}_1 forwards messages from c' to c and from \bar{c} to \bar{c}' just like \mathcal{P}_1 . Otherwise (i.e., $r = 1$), \mathcal{F}_1 still forwards messages from c' to c but it forwards messages from \bar{c} to c . That is, with the overwhelming probability $1 - 2^{-\eta}$, \mathcal{F}_1 behaves exactly as \mathcal{P}_1 . But with the negligible probability $2^{-\eta}$, \mathcal{F}_1 forwards all messages to c .

It is easy to see that \mathcal{P}_1 , \mathcal{P}_2 , $\mathcal{P}_1 | \mathcal{P}_2$, \mathcal{F}_1 , \mathcal{F}_2 , and $\mathcal{F}_1 | \mathcal{F}_2$ are environmentally strictly bounded (both for uniform and non-uniform environments). Furthermore, it is easy to see that $\mathcal{P}_1 \leq_{\text{strict}}^{SS} \mathcal{F}_1$ and $\mathcal{P}_1 \leq_{\text{strict}}^{UC} \mathcal{F}_1$, because, except with negligible probability, \mathcal{F}_1 behaves exactly like \mathcal{P}_1 (no simulator is needed). It is also easy to see that $\mathcal{P}_2 \leq_{\text{strict}}^{SS} \mathcal{F}_2$ and $\mathcal{P}_2 \leq_{\text{strict}}^{UC} \mathcal{F}_2$: the simulator can simply replay messages output by \mathcal{F}_2 back to \mathcal{F}_2 .

Next, we show that $\mathcal{P}_1 | \mathcal{P}_2 \not\leq_{\text{strict}}^{SS} \mathcal{F}_1 | \mathcal{F}_2$. Note that this implies $\mathcal{P}_1 | \mathcal{P}_2 \not\leq_{\text{strict}}^{UC} \mathcal{F}_1 | \mathcal{F}_2$ because strict UC implies strict SS (Lemma 21). As a warming up, we note that the simulator \mathcal{S} that simply replays messages output by \mathcal{F}_2 back to \mathcal{F}_2 (as it can be used to prove $\mathcal{P}_2 \leq_{\text{strict}}^{SS} \mathcal{F}_2$) is not a “good” simulator to prove $\mathcal{P}_1 | \mathcal{P}_2 \leq_{\text{strict}}^{SS} \mathcal{F}_1 | \mathcal{F}_2$ because the system $\mathcal{S} | \mathcal{F}_1 | \mathcal{F}_2$ is not environmentally *strictly* bounded: If $r = 1$ (in \mathcal{F}_1), which happens with probability $2^{-\eta}$, and an environment sends a message to $\mathcal{S} | \mathcal{F}_1 | \mathcal{F}_2$ via c' , then this message circles between \mathcal{F}_1 , \mathcal{F}_2 , and \mathcal{S} forever, and hence, the run does not terminate. (Note, however, that $\mathcal{S} | \mathcal{F}_1 | \mathcal{F}_2$ is environmentally almost bounded and with the simulator \mathcal{S} we obtain $\mathcal{P}_1 | \mathcal{P}_2 \leq_{\text{strict}}^{SS} \mathcal{F}_1 | \mathcal{F}_2$).

We now show that $\mathcal{P}_1 | \mathcal{P}_2 \not\leq_{\text{strict}}^{SS} \mathcal{F}_1 | \mathcal{F}_2$ by contradiction. Assume that $\mathcal{P}_1 | \mathcal{P}_2 \leq_{\text{strict}}^{SS} \mathcal{F}_1 | \mathcal{F}_2$. Then, there exists $\mathcal{S} \in \text{Sim}_{\text{strict}}^{\mathcal{P}_1 | \mathcal{P}_2}(\mathcal{F}_1 | \mathcal{F}_2)$ such that $\mathcal{P}_1 | \mathcal{P}_2 \cong \mathcal{S} | \mathcal{F}_1 | \mathcal{F}_2$.

First, we show that \mathcal{S} has basically no choice but to forward all messages it receives back to \mathcal{F}_2 . Note that \mathcal{S} only has two network tapes (one input and one output tape) which connect to \mathcal{F}_2 . For all $\eta, i \in \mathbb{N}$, let $p_{\eta, i}$ denote the probability that \mathcal{S} (with security parameter η) outputs 1 in the first i consecutive activations with input 1. (Here, the probability is taken over the random coins of \mathcal{S} .) We now show

that the probability $p_{\eta,i}$ is overwhelming for all polynomially bounded i (in η). Let q be a polynomial in η . We construct $\mathcal{E} \in \text{Env}(\mathcal{P}_1 | \mathcal{P}_2)$ as follows: \mathcal{E} sends the message 1 on tape c' (i.e., to \mathcal{P}_1 or \mathcal{F}_1) and waits to receive 1 on \bar{c}' . If \mathcal{E} does not receive 1 (either \mathcal{E} is activated with empty input on tape **start** or it receives some other message on \bar{c}'), \mathcal{E} outputs 0 on **decision**. If \mathcal{E} receives 1 on \bar{c}' , it again sends the message 1 on tape c' and behaves just as before. The system \mathcal{E} does this $q(\eta)$ times, unless it outputs 0 on tape **decision**. If in all these iterations, \mathcal{E} received 1 back from \mathcal{F}_2 , \mathcal{E} outputs 1 on **decision**. Clearly, if \mathcal{E} interacts with $\mathcal{P}_1 | \mathcal{P}_2$, it will always output 1 on **decision**, i.e., $\text{Prob}[(\mathcal{E} | \mathcal{P}_1 | \mathcal{P}_2)(1^\eta, a) = 1] = 1$ for all η, a . If \mathcal{E} interacts with $\mathcal{S} | \mathcal{F}_1 | \mathcal{F}_2$, then \mathcal{E} will output 1 only if \mathcal{S} forwarded 1 $q(\eta)$ times (and $r \neq 1$; where r is the number chosen by \mathcal{F}_1). That is, $\text{Prob}[(\mathcal{E} | \mathcal{S} | \mathcal{F}_1 | \mathcal{F}_2)(1^\eta, a) = 1] \leq p_{\eta,q(\eta)}$, for all η, a . Since $|\text{Prob}[(\mathcal{E} | \mathcal{P}_1 | \mathcal{P}_2)(1^\eta, a) = 1] - \text{Prob}[(\mathcal{E} | \mathcal{S} | \mathcal{F}_1 | \mathcal{F}_2)(1^\eta, a) = 1]|$ is negligible (as a function in η and a), $1 - p_{\eta,q(\eta)}$ is negligible, i.e., $p_{\eta,q(\eta)}$ is overwhelming. In particular, there exists $\eta_0 \in \mathbf{N}$ such that $p_{\eta,q(\eta)} > 0$ for all $\eta \geq \eta_0$.

Now, we use the above to show that $\mathcal{S} | \mathcal{F}_1 | \mathcal{F}_2$ is not environmentally strictly bounded, which contradicts the assumption that $\mathcal{S} \in \text{Sim}_{\text{strict}}^{\mathcal{P}_1 | \mathcal{P}_2}(\mathcal{F}_1 | \mathcal{F}_2)$. For this purpose, let $\mathcal{E} \in \text{Env}(\mathcal{S} | \mathcal{F}_1 | \mathcal{F}_2)$ be the following system: \mathcal{E} sends the message 1 on tape c' (i.e., to \mathcal{P}_1 or \mathcal{F}_1) and then halts. Since, by assumption, $\mathcal{S} | \mathcal{F}_1 | \mathcal{F}_2$ is environmentally strictly bounded, there exists a polynomial p such that for every security parameter $\eta \in \mathbf{N}$ and external input a the runtime of the system $\mathcal{Q} := \mathcal{E} | \mathcal{S} | \mathcal{F}_1 | \mathcal{F}_2$ is bounded by $p(\eta)$ for all runs of this system, i.e., $\text{Time}(\mathcal{Q}(1^\eta, a)) \leq p(\eta)$ (note that \mathcal{E} ignores the external input a). Let $q(\eta) := p(\eta) + 1$ for all $\eta \in \mathbf{N}$. If $r = 1$ in \mathcal{F}_2 , which happens with probability $2^{-\eta}$, then in such a run \mathcal{S} is activated (at least) $q(\eta)$ times with 1 and we know from what we have shown above that there exists $\eta_0 \in \mathbf{N}$ such that $p_{\eta,q(\eta)} > 0$ for all $\eta \geq \eta_0$, i.e., the probably that \mathcal{S} will return 1 at least $q(\eta)$ times is non-zero. In particular, there exists a run of \mathcal{Q} with security parameter η_0 (the external input a is ignore by \mathcal{E} , and hence, does not matter) that has non-zero probability such that the runtime of this run is bigger than $p(\eta_0)$, a contradiction. This shows that $\mathcal{S} | \mathcal{F}_1 | \mathcal{F}_2$ is not environmentally strictly bounded. We conclude that $\mathcal{P}_1 | \mathcal{P}_2 \not\leq_{\text{strict}}^{SS} \mathcal{F}_1 | \mathcal{F}_2$. \square

We note that the above proof works for both uniform and non-uniform environments. Hence, the statement of the lemma also holds in a uniform setting.

We emphasize that the protocol systems \mathcal{P}_1 , \mathcal{P}_2 , and \mathcal{F}_2 defined in the above proof are simple dummy IITMs which only forward/replay messages. The IITM \mathcal{F}_1 is simple too: It only forwards/replays messages and the “switch” r determines where to forward/replay the messages. Moreover, the simulators that are used in the above proof to show that \mathcal{P}_1 realizes \mathcal{F}_1 and \mathcal{P}_2 realizes \mathcal{F}_2 are also very simple: In fact, for \mathcal{P}_1 and \mathcal{F}_1 there is no simulator because these systems do not have any network tapes. The simulator \mathcal{S} which is used to show that \mathcal{P}_2 realizes \mathcal{F}_2 is a simple dummy IITM which replays all messages back to \mathcal{F}_2 . In particular, \mathcal{S} is environmentally strictly bounded. Altogether, since the protocols and simulators considered are all very simple, it is quite hard to restrict the class of protocols and/or simulators in order to obtain reasonable security notions based on environmental strict boundedness. Possible approaches are to enforce some acyclicity conditions on the flow of messages between protocols and/or simulators such that the cycles vanish. Such approaches have been taken in related work but they are unsatisfactory as discussed in Section 11.

9.3 No Transitivity

The proof of the following lemma, which shows that $\leq_{\text{strict}}^{SS}$ is not transitive, is very similar to the proof of Lemma 22.

Lemma 23. *There exist protocol systems \mathcal{Q}_1 , \mathcal{Q}_2 , and \mathcal{Q}_3 with pairwise disjoint sets of network tapes such that the following conditions are satisfied:*

1. \mathcal{Q}_1 , \mathcal{Q}_2 , and \mathcal{Q}_3 are environmentally strictly bounded.
2. $\mathcal{Q}_1 \leq_{\text{strict}}^{SS} \mathcal{Q}_2$.
3. $\mathcal{Q}_2 \leq_{\text{strict}}^{SS} \mathcal{Q}_3$.

4. $\mathcal{Q}_1 \not\leq_{\text{strict}}^{SS} \mathcal{Q}_3$.

Proof. Let $\mathcal{P}_1, \mathcal{P}_2, \mathcal{F}_1$, and \mathcal{F}_2 be the protocol systems defined in the proof of Lemma 22. Furthermore, let \mathcal{F}'_2 be obtained from \mathcal{F}_2 by renaming the network tapes such that the set of network tapes of \mathcal{F}_2 and \mathcal{F}'_2 are disjoint. We define:

$$\begin{aligned}\mathcal{Q}_1 &:= \mathcal{P}_1 | \mathcal{P}_2 \text{ ,} \\ \mathcal{Q}_2 &:= \mathcal{P}_1 | \mathcal{F}_2 \text{ ,} \\ \mathcal{Q}_3 &:= \mathcal{F}_1 | \mathcal{F}'_2 \text{ .}\end{aligned}$$

From the proof of Lemma 22, we already know that \mathcal{Q}_1 and \mathcal{Q}_3 are environmentally strictly bounded and that $\mathcal{Q}_1 \not\leq_{\text{strict}}^{SS} \mathcal{Q}_3$. Moreover, it is easy to see that $\mathcal{Q}_2 = \mathcal{P}_1 | \mathcal{F}_2$ is environmentally strictly bounded. It remains to show that $\mathcal{Q}_1 \leq_{\text{strict}}^{SS} \mathcal{Q}_2$ and $\mathcal{Q}_2 \leq_{\text{strict}}^{SS} \mathcal{Q}_3$.

In order to show $\mathcal{Q}_1 \leq_{\text{strict}}^{SS} \mathcal{Q}_2$ it suffices to consider a simulator \mathcal{S}_1 which simply replays all messages from \mathcal{Q}_2 back to \mathcal{Q}_1 . It is easy to see that $\mathcal{S}_1 | \mathcal{Q}_2$ is environmentally strictly bounded. With this, $\mathcal{Q}_1 \leq_{\text{strict}}^{SS} \mathcal{Q}_2$ follows immediately.

In order to show $\mathcal{Q}_2 \leq_{\text{strict}}^{SS} \mathcal{Q}_3$ it clearly suffices to consider a simulator \mathcal{S}_2 which simply forwards all messages from \mathcal{Q}_3 to the environment and vice versa. Again, it is easy to see that $\mathcal{S}_2 | \mathcal{Q}_3$ is environmentally strictly bounded. \square

We note that the above lemma does not carry over to strict UC (with the notation in the proof of Lemma 22 we have that $\mathcal{Q}_2 \not\leq_{\text{strict}}^{UC} \mathcal{Q}_3$; this can be shown analogously to the proof of Lemma 22). In fact, strict UC can easily be seen to be transitive:

Lemma 24. *Let $\mathcal{Q}_1, \mathcal{Q}_2$, and \mathcal{Q}_3 be environmentally strictly bounded protocol systems. If $\mathcal{Q}_1 \leq_{\text{strict}}^{UC} \mathcal{Q}_2$ and $\mathcal{Q}_2 \leq_{\text{strict}}^{UC} \mathcal{Q}_3$, then $\mathcal{Q}_1 \leq_{\text{strict}}^{UC} \mathcal{Q}_3$.*

Proof. The proof is analogous to the proof of Lemma 13. \square

9.4 Strict SS does not Imply Strict UC (Incompleteness of the Dummy Adversary)

The following lemma states that strict SS does not imply strict UC.

Lemma 25. *There exist environmentally strictly bounded protocol systems \mathcal{P} and \mathcal{F} such that $\mathcal{P} \leq_{\text{strict}}^{SS} \mathcal{F}$ and $\mathcal{P} \not\leq_{\text{strict}}^{UC} \mathcal{F}$.*

Proof. Let \mathcal{Q}_2 and \mathcal{Q}_3 be defined as in the proof of Lemma 23. With $\mathcal{P} = \mathcal{Q}_2$ and $\mathcal{F} = \mathcal{Q}_3$, the lemma follows. \square

Note that this lemma means that the dummy adversary is not complete (see Remark 11).

10 Instantiation of the IITM Model

The IITM model is a very flexible and expressive model which allows us to model cryptographic protocols in various ways. In this section, we suggest one possible way of modeling (multi-party) protocols in the IITM model, including the addressing of multiple sessions as well as the modeling of subprotocols and corruption; depending on the kind of setting, other approaches might be favorable. A general description of the suggested modeling of protocols is provided in Section 10.1, with a more concrete example presented in Section 10.2.

10.1 Modeling Protocols

A protocol with n roles can be modeled in the IITM model as the following protocol system:

$$\mathcal{P} := !M_1 \mid \cdots \mid !M_n \quad (19)$$

where M_i , for $i \in \{1, \dots, n\}$, is an IITM which models the i th role. Every machine is in the scope of a bang operator to model multiple sessions of the protocol (see below). Moreover, every machine M_i has I/O and network tapes. The network tapes are used to communicate with other machines over the (untrusted) network. As usual, the network is controlled by the adversary, and hence, all messages sent on a network tape go directly to the adversary and all message received on a network tape come from the adversary. Network tapes are also used to model corruption (see below). The I/O tapes are not controlled by the adversary. They are used by a machine M_i to communicate with the environment, such as (honest) users of the protocol or higher level protocols. For example, if \mathcal{P} is a key exchange protocol, then an I/O tape would be used to output a successfully established session key (to some user or, for example, to an instance of a secure channel protocol). Potentially, the I/O tapes can also be used by a machine M_i to communicate with other roles directly, rather than through the adversary or environment. In this case, M_i would directly be connected with another machine, say M_j , via I/O tapes. However, this would constitute an ideal channel of communication between two parties, which would typically be modeled as an ideal functionality serving as an (ideal) subprotocol of \mathcal{P} (see below). Nevertheless, the IITM model is flexible enough to model such ideal channels.

Addressing of multiple sessions. Due to the general concept of the mode `CheckAddress` there are many possible ways of how multiple sessions of a protocol can be addressed. We now describe one such approach, which is based on globally unique session identifiers (SIDs); we refer the reader to [22] for an addressing mechanism formulated in the IITM model with *locally* chosen and managed session identifiers.

To address multiple instances of a machine (role) M_i , and hence, multiple sessions of \mathcal{P} , by using (global) SIDs, M_i can be defined as follows: In mode `CheckAddress`, the machine M_i accepts an incoming message only if it is of the form (sid, m) for some SID sid and some message m . In mode `Compute`, the machine M_i records the SID contained in the first message it accepted. Later it will only accept messages in mode `CheckAddress` which are prefixed with the recorded SID, say sid . Moreover, in mode `Compute` the machine M_i will only output messages that are prefixed with sid .

This guarantees that in a run of \mathcal{P} (with some environment) there is at most one copy of every M_i for every SID sid ; we denote such a copy by $M_i[sid]$ and say that this copy is *addressed by* sid . The instances $M_1[sid], \dots, M_n[sid]$ (or a subset thereof if not all instances are present in a run) form a session of the protocol \mathcal{P} , the *session with SID* sid . We say that $M_i[sid]$ *belongs to session* sid . Note that all instances within one session share the same SID. This SID is globally unique, is given to an instance from outside the protocol, and is pre-established by the parties participating in one session, in the sense that the SID is established before the actual protocol starts to run.

We note that, with the notion from Section 5.3.1, \mathcal{P} is an σ_{prefix} -session version.²¹ In particular, the composition theorems can be applied, and hence, it suffices to reason about \mathcal{P} in the single-session setting (see below). Note that, by modeling M_i to be a σ_{prefix} -session version, an instance $M_i[sid]$ of M_i addressed by sid is aware of its SID sid , i.e., it can use sid in its computation, and, for example, include it in messages to be signed/encrypted. If this is not necessary, one could model M_i as a session version of some IITM M'_i , i.e., $M_i = \overline{M'_i}$ (see Section 5.2.1). In this case, M_i would be completely oblivious to its SID.

Typically, an SID is structured and contains, in addition to the actual SID, the names of the parties involved in the session. For example, to model that in one session s the i th role is played by party pid_i the SID would be of the form $sid = (s, pid_1, \dots, pid_n)$ and a machine M_i would be defined in such a way that an instance $M_i[sid]$ would run the i th role as party pid_i . (Clearly, this can be generalized to let several parties run one role in one session of the protocol.)

²¹Recall that, for all messages m and tapes c , $\sigma_{\text{prefix}}(m, c) := sid$ if $m = (sid, m')$ for some sid, m' and $\sigma_{\text{prefix}}(m, c) := \perp$ otherwise.

The modeling of globally unique and pre-established SIDs as described above is used in all other models for universal composability and it is hard-wired into these models (see Section 11). The IITM model is flexible enough to allow for other forms of addressing multiple protocol sessions. In particular, in [22] an alternative way of addressing multiple protocol session without pre-established and globally unique SIDs is presented within the IITM model. In this formulation, parties merely use locally chosen and managed SIDs. As further discussed in [22], this allows for a *faithful* analysis of protocols that do not use pre-established SIDs, which is the case for most real-world authentication, key exchange, and secure channel protocols.

Security proofs using composition theorems. In the universal composability paradigm, security of a protocol \mathcal{P} , as modeled above, typically means that it realizes some appropriate ideal protocol \mathcal{F} (e.g., in case of a key exchange protocol, an ideal key exchange functionality), i.e., \mathcal{P} is considered secure (w.r.t. \mathcal{F}) if $\mathcal{P} \leq^{SS} \mathcal{F}$. Of course, one can attempt to prove $\mathcal{P} \leq^{SS} \mathcal{F}$ directly. But this would require a proof which has to consider multiple concurrent sessions. Using the composition theorems for unbounded self-composition (Theorem 9, if the protocol does not depend on the SIDs, i.e., $M_i = \underline{M}'_i$ for some M'_i ; or Theorem 10 if the protocol depends on the SIDs) simplifies this proof because one has to consider only a single session of the protocol: For example, by Theorem 10, except for some (typically simple) checks concerning the runtime of the system, one has to show only that $\mathcal{P} \leq_{\sigma_{\text{prefix-single}}}^{SS} \mathcal{F}$ to obtain that $\mathcal{P} \leq^{SS} \mathcal{F}$. Hence, roughly speaking, it suffices to show that $M_1[side] \mid \dots \mid M_n[side]$ realizes $\mathcal{F}[side]$ for one session *side*.

Subprotocols and ideal functionalities. Complex protocols can often be/are often structured in a hierarchy of higher- and lower-level protocols. For example, a secure channel protocol might use a key exchange protocol or an authenticated channel as a subprotocol, and cryptographic primitives (such as encryption or digital signatures) could be modeled as subprotocols (see Section 10.2 and [23, 22]).

For the sake of the discussion here, let $\mathcal{P}' = !M'_1 \mid \dots \mid !M'_n$ be a subprotocol of \mathcal{P} with the same structure as \mathcal{P} and where the addressing of machines is defined just as for \mathcal{P} . Since \mathcal{P}' is supposed to be a subprotocol of \mathcal{P} , the machines in \mathcal{P}' typically connect via I/O tapes to the corresponding machines in \mathcal{P} , i.e., M'_i and M_i are connected via I/O tapes. By the addressing defined above for \mathcal{P} and \mathcal{P}' , every instance $M_i[side]$ will, via the I/O tapes, only interact with $M'_i[side]$ (since these instances output messages only of the form $(side, m)$ for some message m and such messages are not accepted by other instances). Instead of the subprotocol \mathcal{P}' , \mathcal{P} might be connected to an ideal protocol (or ideal functionality) \mathcal{F}' that provides the same I/O interface as \mathcal{P}' but provides the functionality of \mathcal{P}' in an ideal way (e.g., an ideal key exchange or an ideal cryptographic primitive).

Structuring a protocol like this again simplifies the proof of security of a (complex) protocol because the subprotocol can be analyzed in separation and then \mathcal{P} can be analyzed based on the ideal protocol as follows: To prove that $\mathcal{P} \mid \mathcal{P}'$ is secure, i.e., $\mathcal{P} \mid \mathcal{P}' \leq^{SS} \mathcal{F}$ for some ideal protocol/functionality \mathcal{F} , it suffice to show that

$$\mathcal{P}' \leq_{\sigma_{\text{prefix-single}}}^{SS} \mathcal{F}'$$

for some appropriate ideal protocol/functionality \mathcal{F}' and that

$$\mathcal{P} \mid \mathcal{F}' \leq_{\sigma_{\text{prefix-single}}}^{SS} \mathcal{F} .$$

From this, using the composition theorems (Theorem 8 and 10) and transitivity of \leq^{SS} (Lemma 13), it follows immediately that $\mathcal{P} \mid \mathcal{P}' \leq^{SS} \mathcal{F}$, which means that multiple sessions of \mathcal{P} , where every such session may use a session of \mathcal{P}' , realizes multiple sessions of \mathcal{F} . We emphasize that both proof steps require merely single-session reasoning and that the second proof step is further simplified because the subprotocol/functionality used by \mathcal{P} is idealized.

We note that the IITM model and the composition theorems are flexible enough to deal with much more complex scenarios than the one described above. For instance:

1. The reasoning can be iterated: \mathcal{P}' itself could use subprotocols and the composition theorems can be used to simplify the proof of $\mathcal{P}' \leq_{\sigma_{\text{prefix-single}}}^{SS} \mathcal{F}'$, just as in the case of \mathcal{P} above.

2. \mathcal{P} might use more than just one subprotocol in parallel.
3. \mathcal{P} could use multiple sessions of \mathcal{P}' per session. In this case, the sessions of \mathcal{P}' could, for instance, be addressed by a hierarchical SID (sid, sid', m) where sid is the SID that \mathcal{P} uses and sid' is an extra SID that is used to address the sessions of \mathcal{P}' within a session of \mathcal{P} . However, again single-session reasoning would suffice for both \mathcal{P} and \mathcal{P}' to establish security properties for \mathcal{P} (in composition with \mathcal{P}') for multiple sessions.
4. Obviously, \mathcal{P}' is not restricted to have the same structure as \mathcal{P} . For example, \mathcal{P}' could be a two-party/two-role protocol and \mathcal{P} could be an n -party/ n -role protocol where every two parties of one session in \mathcal{P} use one session of \mathcal{P}' . This way, for example, an n -party key exchange protocol could be built from a two-party key exchange protocol.

Modeling corruption. In the IITM model, the way corruption is modeled is not fixed and hard-wired into the model but is part of the specification of protocols, with the advantage that (i) the IITM model is simple, (ii) general theorems proven in the IITM model, such as composition theorems, hold true independently of how corruption is modeled, and (iii) corruption can be modeled in a very flexible way. We now describe one possible way of modeling corruption in the IITM model. Clearly, other ways are possible and in some cases might be desirable.

The adversary (or environment) who connects to the network interface of a protocol \mathcal{P} may send a special *corrupt* message to a network tape of (some instance of) a machine M_i in \mathcal{P} . When M_i receives such a message it considers itself corrupted and outputs its complete configuration to the adversary (clearly other options are conceivable and useful as well, see below). From then on M_i forwards all messages between the I/O and network interface, i.e., the adversary is in full control of the corrupted instance.²² If \mathcal{P} uses subprotocols/functionalities, as described above, by this the adversary would also gain full access to the I/O interface of the subprotocols the corrupted instance of M_i has access to. This models fully adaptive, active corruption of single instances. We note that, as always in universal composability settings, the distinguishing environment should have the possibility to know which instances are corrupted because, otherwise, a simulator could always corrupt instances in the ideal world and then perfectly simulate the real world, i.e., every protocol system would realize every other protocol system (with the same I/O interface). Therefore, a machine M_i is defined in such a way that, on the I/O interface, it accepts requests of the form *corrupted?* and answers *yes* if it has been corrupted, i.e., if it has received a *corrupt* message on the network interface before. As a result, an environment can ask whether an instance is corrupted.

If \mathcal{P} uses subprotocols/functionalities, then, typically, the corruption status of a machine in \mathcal{P} also depends on the corruption status of the subprotocols it uses. That is, an instance of a machine M_i might consider itself corrupted also if one of its subprotocol instances is corrupted (a fact that M_i can check by sending a *corrupted?* request). Note, however, that even if M_i returns to the environment that it is corrupted (because some part of the subprotocol is corrupted), then this does not necessarily mean that M_i has to consider itself completely controlled by the adversary. (Clearly, if desired, M_i could be modeled in such a way that in this case it considers itself to be fully controlled by the adversary.)

For example, if M_i models an instance of a key exchange protocol and uses a functionality for public-key encryption, then M_i , if asked whether it is corrupted, would return *yes* if it has been corrupted directly or if its public-key functionality has been corrupted, because in this case it could not provide security guarantees. It makes sense to model M_i in such a way that, even though the public-key functionality that M_i uses is corrupted, it still follows its prescribed protocol: the fact that the private key has been stolen by the adversary does not necessarily mean that every instance that uses the key is completely controlled by the adversary. However, if desired, one could also model M_i in such a way that if one of its subprotocols is corrupted, then M_i considers itself controlled by the adversary. This depends on the kind of corruption one would like to consider. Conversely, the adversary could corrupt only M_i but not the public-key functionality that M_i uses,

²²However, this instance of M_i would still make sure addressing conventions are followed. For example, if it is a σ_{prefix} -session version, it would still behave like a σ_{prefix} -session version with the SID it already has.

which would model that the private key of M_i is still not known by the adversary (e.g., because it is stored on a smart card), but the process (the instance of M_i) that uses the private key is corrupted.

It should be clear that the way of corruption sketched above allows for a very fine-grained and flexible modeling of corruption, ranging from the corruption of single instances to the corruption of complete parties. In order to corrupt a complete party, the adversary can corrupt every instance of M_i (and subprotocol instances used by this instance) that belong to the party an adversary wants to control. For such a form of corruption, instances would typically check whether one of their subprotocols are corrupted and then consider themselves completely controlled by the adversary as well. If one wants to make sure that if an instance is corrupted, then also all its subinstances are, an instance could check that if it has been explicitly corrupted by the adversary that then all its subinstances have been corrupted as well (and if this is not the case it could wait until the adversary has explicitly corrupted all subinstances).

So far we have not put a restriction on when corruption can occur, and hence, we modeled adaptive corruption. Clearly, static corruption can be modeled as well: For this purpose, upon its first activation (an instance of a) machine could first ask the adversary whether he wants to corrupt the machine. Subsequent corrupt messages would then be ignored by the machine.

While, as introduced above, explicit corruption meant that a machine provides its complete configuration to the adversary, other forms of corruption where a machine gives away, for example, merely its long-term or ephemeral keys are conceivable as well.

Composition with joint state. The composition theorems for unbounded self-composition allow us to conclude that a protocol is secure in a setting with multiple sessions, given that the protocol is secure in a single-session setting. However, these composition theorems assume that protocol sessions are completely disjoint: there does not exist a joint state between different sessions, i.e., between instances of the M_i 's with different SIDs. Such a joint state is, however, often desirable/required. For example, one would like to use the same long-term keys, such as public/private key pairs or shared long-term symmetric keys, in every protocol session. Such long-term keys constitute a joint state between different protocol sessions. Using fresh long-term keys in every session would be completely impractical.

We emphasize that it is, of course, straightforward to model a (practical) protocol in such a way that different instances make use of joint state. In particular, a protocol \mathcal{P} of the form $!M_1 \mid \cdots \mid !M_n$ as introduced above could, for example, be extended by a subprotocol, an ideal functionality \mathcal{F} (or its realization), say, such that in every run all instances of M_i access the same instance of \mathcal{F} .²³ The point is that with such a joint state the composition theorems for unbounded self-composition would not be applicable anymore, and hence, the security of such a protocol would have to be proven directly in a multi-session setting.

To be able to analyze protocols that use joint state in the single-session setting and still obtain security in the multi-session setting, Canetti and Rabin [10] were the first to propose composition theorems that allow for joint state—so-called composition theorems with joint state, or simply joint state (composition) theorems; see Section 11 and [20] for a discussion of [10].

In [10], Canetti and Rabin first proposed a general joint state theorem (in the UC framework). In [20], this theorem was stated in the IITM model and it was shown that, unlike in the UC framework, in the IITM model it is a direct consequence of the composition theorem for a constant number of systems (Theorem 8) and the formulation of this theorem does not require to introduce a joint state operator. Formulated in the IITM model, the general joint state theorem states that if $\mathcal{P}' \leq^{SS} \mathcal{F}'$, then $\mathcal{P} \mid \mathcal{P}' \leq^{SS} \mathcal{P} \mid \mathcal{F}'$, where \mathcal{F}' models some multi-session version of some ideal functionality and \mathcal{P}' is supposed to be a realization of \mathcal{F}' that utilizes some joint state across the sessions, e.g., \mathcal{P}' uses only a single copy of \mathcal{F}' for all sessions. Clearly, in the IITM model, this general joint state theorem is a direct consequence of Theorem 8. In the UC framework, it requires to extend the model and the notation and it requires a proof.

²³For this purpose, the `CheckAddress` mode of \mathcal{F} could be defined in such a way that all message are accepted. One could also consider a system $! \mathcal{F}$ where there is one instance of \mathcal{F} for every party, say. In this case, \mathcal{F} could be defined to be a σ_{prefix} -session version with messages of the form (pid, m) . Hence, there would be (at most) one instance of \mathcal{F} for every pid in a run. Note that in any case, to access (an instance of) \mathcal{F} , a machine M_i cannot be defined to be a σ_{prefix} -session version anymore, because while it is invoked with some SID sid , it outputs messages with a different prefix in order to access (an instance of) \mathcal{F} .

The general joint state theorem by itself does not say how a joint state realization looks like. The main challenge is always to find suitable joint state realizations for concrete ideal functionalities. As an example, we consider the joint state realization for public-key encryption, following [20]. In [20], an ideal functionality \mathcal{F}_{pke} for public-key encryption is defined; the details of this functionality are not important for this discussion.²⁴ This functionality is the “encryption/decryption-box” of one party. In particular, it encapsulates the public/private key pair of that party, where the private key stays in the functionality (except if the functionality is corrupted) and the public key is given out, and hence, can be distributed. The system $!\mathcal{F}_{\text{pke}}$ describes the multi-party version of \mathcal{F}_{pke} , i.e., in every run there is at most one instance $\mathcal{F}_{\text{pke}}[pid]$ of \mathcal{F}_{pke} per party pid . The system $!\underline{\mathcal{F}_{\text{pke}}}$ is the multi-session and multi-party version of \mathcal{F}_{pke} , i.e., there may be multiple sessions per party and instances of \mathcal{F}_{pke} are addressed by identifiers of the form (sid, pid) , denoting the session sid of party pid . In [20], a joint state realization $\mathcal{P}_{\text{pke}}^{\text{js}} | !\mathcal{F}_{\text{pke}}$ is proposed, where $\mathcal{P}_{\text{pke}}^{\text{js}}$ is a kind of multiplexer which handles multiple sessions per party but where encryption and decryption of all sessions of one party pid are handled by the instance $\mathcal{F}_{\text{pke}}[pid]$ of that party, i.e., if $\mathcal{P}_{\text{pke}}^{\text{js}}$ is asked to encrypt/decrypt a message for pid in some session (sid, pid') , $\mathcal{P}_{\text{pke}}^{\text{js}}$ uses the instance $\mathcal{F}_{\text{pke}}[pid]$ for that purpose. The basic idea is that SIDs sid are added to messages to be encrypted. Upon decryption in session sid , it is checked whether a message contains sid . By this, it is prevented that ciphertexts created in one session can be used in other sessions. In [20], it has been proven that the proposed joint state realization in fact realizes the multi-session and multi-party version of \mathcal{F}_{pke} , i.e.,

$$\mathcal{P}_{\text{pke}}^{\text{js}} | !\mathcal{F}_{\text{pke}} \leq^{SS} !\underline{\mathcal{F}_{\text{pke}}} . \quad (20)$$

With such a joint state realization it is possible to prove that a protocol \mathcal{P} that uses public-key encryption is secure in a multi-session setting where a party uses the same public/private key across all sessions by reasoning about just a single session of \mathcal{P} . To illustrate this, let \mathcal{P} be a (multi-session and multi-party) protocol of the form (19) that uses $\mathcal{P}_{\text{pke}}^{\text{js}} | !\mathcal{F}_{\text{pke}}$ for encryption. Recall that an instance of M_i is addressed by an SID of the form $sid = (s, pid_1, \dots, pid_n)$. If such an instance wants to encrypt a message m for party pid_j it would send an encryption request containing the message m to $\mathcal{P}_{\text{pke}}^{\text{js}} | !\mathcal{F}_{\text{pke}}$. For addressing purposes, the request also contains the SID (s, pid_1, \dots, pid_n) and the PID pid_j . With this, the message m is encrypted using the instance $\mathcal{F}_{\text{pke}}[pid_j]$. (More precisely, by the definition of $\mathcal{P}_{\text{pke}}^{\text{js}}$, $\mathcal{P}_{\text{pke}}^{\text{js}}$ asks $\mathcal{F}_{\text{pke}}[pid_j]$ to encrypt the message $((s, pid_1, \dots, pid_n), m)$.)

Now, assume that we want to prove that $\mathcal{P} | \mathcal{P}_{\text{pke}}^{\text{js}} | !\mathcal{F}_{\text{pke}}$ realizes \mathcal{F} , where \mathcal{F} is a multi-session and multi-party formulation of some ideal functionality. Then it suffices to show that:

$$\mathcal{P} | !\underline{\mathcal{F}_{\text{pke}}} \leq_{\sigma_{\text{prefix-single}}}^{SS} \mathcal{F} . \quad (21)$$

Note that to prove (21) only a single session (s, pid_1, \dots, pid_n) of \mathcal{P}/\mathcal{F} needs to be analyzed. Such a session might contain n instances of \mathcal{F}_{pke} , one for each pid_i . The analysis of such a session is further simplified due to the use of \mathcal{F}_{pke} , i.e., *ideal* public-key encryption.

From (21), using the joint state composition theorem for public-key encryption (20), the composition theorems (Theorem 8 and 10), and transitivity of \leq^{SS} (Lemma 13), it immediately follows that

$$\mathcal{P} | \mathcal{P}_{\text{pke}}^{\text{js}} | !\mathcal{F}_{\text{pke}} \leq^{SS} \mathcal{F} . \quad (22)$$

Moreover, if \mathcal{P}_{pke} is a realization of \mathcal{F}_{pke} , i.e., $\mathcal{P}_{\text{pke}} \leq^{SS} \mathcal{F}_{\text{pke}}$ (\mathcal{P}_{pke} could, for example, be an IND-CCA2-secure public-key encryption scheme), we obtain by Theorem 8 and 9 that

$$\mathcal{P} | \mathcal{P}_{\text{pke}}^{\text{js}} | !\underline{\mathcal{P}_{\text{pke}}} \leq^{SS} \mathcal{F} . \quad (23)$$

This says that \mathcal{P} , modeling a multi-session, multi-party version of some protocol which uses public-key encryption, realizes \mathcal{F} , modeling a multi-session, multi-party version of an ideal functionality. By the joint

²⁴The main idea of this functionality is that when given a message m to be encrypted, it encrypts a random message of the length of m instead of encrypting m . The resulting ciphertext c is stored in \mathcal{F}_{pke} along with m . This ciphertext, by construction, does not contain any information about m , except for the length of m . If later, \mathcal{F}_{pke} is asked to decrypt c , it looks up the corresponding plaintext in the table and returns this plaintext. We refer the reader to [20] for details.

state realization $\mathcal{P}_{\text{pke}}^{\text{js}} \mid \!\! \mid \mathcal{P}_{\text{pke}}$, \mathcal{P} uses only one public/private key pair for every party across all sessions of \mathcal{P} . We emphasize that to prove (23), we needed to show only (21). (Note that proving (20) is a once and for all proof, which does not depend on the context in which public-key encryption is used.)

However, as already mentioned before, in the realization $\mathcal{P} \mid \mathcal{P}_{\text{pke}}^{\text{js}} \mid \!\! \mid \mathcal{P}_{\text{pke}}$ the SIDs ($s, \text{pid}_1, \dots, \text{pid}_n$) are added to all encrypted messages (this is what $\mathcal{P}_{\text{pke}}^{\text{js}}$ does). While this is a good design principle, existing, in particular real-world protocols, typically do not follow this pattern. Hence, such protocols cannot be analyzed with the joint state theorem sketched above without severely modifying the protocols: adding SIDs to plaintexts is a severe modification of a protocol, which can turn an insecure protocol into a secure one. In order to analyze a protocol without modifying it, one could resort to multi-session analysis, with the disadvantage that the analysis is more complex and one does not make full use of modular analysis. Therefore, in [22] composition and joint state theorems were proposed (within the IITM model) which allow for establishing the security of a protocol w.r.t. multiple sessions by analyzing only a single session of the protocol, but without requiring to change the protocol by adding SIDs in messages or in any other way, as illustrated for several real-world protocols in [22].

Composition with global setup. Canetti et al. [6] extended the UC framework by a concept called *global setup*, which allows for modeling protocols that use a globally available trusted setup assumption such as a public-key infrastructure (PKI) or a common reference string (CRS). “Globally available” means that *every* component, including the attacker and other protocols, have access to the setup assumption. This is in contrast to composition with joint state (see above) where not everybody but only the analyzed protocol has access to the joint state. For example, all sessions of the protocol \mathcal{P} described above have access to the same copy of the public-key encryption functionality but no other protocol (or the environment) has (direct) access to this copy of the public-key encryption functionality. While a global setup for public-key encryption would not make sense, for other tasks, such as a PKI or a CRS, modeling a global setup functionality can be desirable. If, as described, all components have permanently access to a global setup functionality, the composition theorems need to be adjusted accordingly.

To enable global setup in the UC model, Canetti et al. [6] had to extend the UC model to a generalized UC (GUC) framework and define a new security notion, which they called *GUC-emulation/realizability*. The extension lets the environment invoke and interact directly with the global setup functionality (i.e., the globally available functionality); this is not possible in the UC model without extension because the environment can only interact directly with the highest-level protocol and the adversary. GUC-emulation is defined similarly to UC-emulation except that the environment is now allowed to access the global setup functionality. To formulate the composition theorem with global setup, Canetti et al. also had to introduce a new notion, called *EUC-emulation/realizability*, which is defined like GUC-emulation but restricts the environment to invoke only a single session of the protocol. The composition theorem with global setup that Canetti et al. prove is basically the following:

Theorem 11 (global setup composition theorem in the GUC model [6]; informal). *Let \mathcal{P} be a protocol that uses a global setup functionality \mathcal{G} . If $\mathcal{P} \mid \mathcal{G}$ EUC-emulates $\mathcal{F} \mid \mathcal{G}$ for some ideal functionality \mathcal{F} , then $\mathcal{Q} \mid \mathcal{P} \mid \mathcal{G}$ GUC-emulates $\mathcal{Q} \mid \mathcal{F} \mid \mathcal{G}$, where \mathcal{Q} is a protocol that uses (possibly multiple copies of) \mathcal{P} (or \mathcal{F}) as a subprotocol.*

We note that, in the above theorem, \mathcal{F} might or might not use \mathcal{G} but it still has to be present on the right-hand side because the environment has to be able to interact with it.

In the IITM model, to enable global setup, we do not need to extend the model or define a new security notion because the model is already flexible enough to express global setup. To formulate the composition theorems with global setup in the IITM model, we only need to define a *global setup (ideal) functionality* \mathcal{G} that can be any ideal functionality but it is parametric in its I/O and network tapes. When we combine \mathcal{G} with environmental, protocol, or adversarial systems (e.g., $\mathcal{E} \mid \mathcal{S} \mid \mathcal{F} \mid \mathcal{G}$ for an environment \mathcal{E} , a simulator \mathcal{S} , and an ideal functionality \mathcal{F}), we choose the parameters of \mathcal{G} such that every IITM has an input and an output tape to connect to \mathcal{G} . The tapes for environmental and protocol systems are chosen to be I/O tapes and the tapes for adversarial systems are network tapes (e.g., in the above system, \mathcal{E} and \mathcal{F} connect to \mathcal{G} via

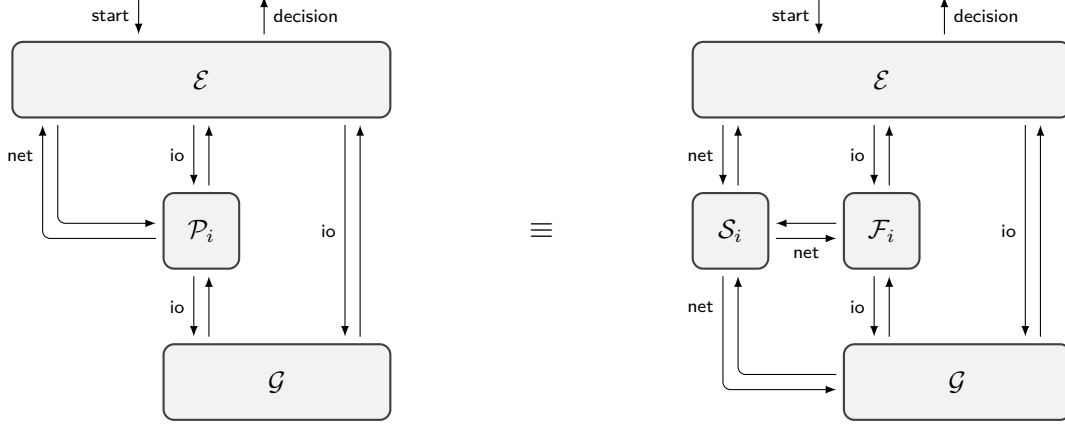


Figure 9: The systems in the assumptions of the composition theorem with global setup. Network tapes are labeled by *net* and I/O tapes are labeled by *io*. We have that $\mathcal{P}_i | \mathcal{G} \leq^{SS} \mathcal{F}_i | \mathcal{G}$, i.e., there exists a simulator \mathcal{S}_i such that $\mathcal{E} | \mathcal{P}_i | \mathcal{G} \equiv \mathcal{E} | \mathcal{S}_i | \mathcal{F}_i | \mathcal{G}$. Notice that \mathcal{G} provides the set of tapes needed in the specific system it runs. In particular, on the right-hand side, \mathcal{G} provides tapes to the simulator \mathcal{S}_i , which are not needed on the left-hand side.

I/O tapes and \mathcal{S} connects to \mathcal{G} via network tapes). This way, \mathcal{G} is globally available. Being a global setup functionality, \mathcal{G} should handle a request independently of the tape the request has been received, except that the response is sent on the corresponding output tape. In the following, we will keep the set of tapes \mathcal{G} has implicit. We just assume that \mathcal{G} has sufficiently many tapes such that all machines in the system in which \mathcal{G} runs can connect to \mathcal{G} .

Within the IITM model, the composition theorems with global setup are now stated as follows. We start with the theorem for composing a constant number of protocols.

Theorem 12 (informal). *Let \mathcal{P}_1 and \mathcal{P}_2 be protocol systems that use a global setup functionality \mathcal{G} . If $\mathcal{P}_1 | \mathcal{G} \leq^{SS} \mathcal{F}_1 | \mathcal{G}$ and $\mathcal{P}_2 | \mathcal{G} \leq^{SS} \mathcal{F}_2 | \mathcal{G}$ for some ideal functionalities \mathcal{F}_1 and \mathcal{F}_2 (see Figure 9), then $\mathcal{P}_1 | \mathcal{P}_2 | \mathcal{G} \leq^{SS} \mathcal{F}_1 | \mathcal{F}_2 | \mathcal{G}$.*

We emphasize that unlike in the UC model, this theorem (and the following one) is stated in the IITM model without changing the model or defining new security notions. Also, the proof of this theorem is completely analogous to the one for Theorem 8: one simply adds \mathcal{G} to all systems considered.

The global setup composition theorem for the unbounded self-composition of a protocol system is stated as follows. We formulate this theorem for protocols that are aware of their SID; a formulation analogous to Theorem 9 is of course also possible.

Theorem 13 (informal). *Let \mathcal{P} be a protocol system that uses a global setup functionality \mathcal{G} and is a σ -session version. If $\mathcal{P} | \mathcal{G} \leq_{\sigma\text{-single}}^{SS} \mathcal{F} | \mathcal{G}$ for some ideal functionality \mathcal{F} , then $\mathcal{P} | \mathcal{G} \leq^{SS} \mathcal{F} | \mathcal{G}$.*

Again, the proof is analogous to the corresponding theorem without global setup (Theorem 10). We also note that Theorems 12 and 13 can be combined to obtain corollaries analogous to those in Section 7.4.

10.2 A Concrete Example

We now present a concrete example of a protocol in order to illustrate the modeling sketched in the previous subsection. More specifically, we show how to model, according to the general description presented above, a simple key exchange protocol based on public-key encryption that is similar to the Needham-Schroeder Public-Key protocol [26] but that uses pre-established SIDs. We also explain how the security of this protocol can be proven in the IITM model along the lines sketched above using the (joint state) composition theorems

1. $A \rightarrow B: \{sid, N_A\}_{k_B}$
2. $B \rightarrow A: \{sid, N_A, N_B\}_{k_A}$
3. $A \rightarrow B: \{sid, N_B\}_{k_B}$

| | | |
|------------|---|---|
| sid | – | unique pre-established SID that contains the party names of A and B |
| k_A, k_B | – | public key of A and B , respectively |
| N_A, N_B | – | nonces, generated by A and B , respectively |
| $\{m\}_k$ | – | encryption of the message m with the public key k using an IND-CCA2 secure public-key encryption scheme |

Figure 10: An example key exchange protocol based on public-key encryption.

and an ideal functionality for public-key encryption, where security for this protocol means that it realizes a standard ideal key exchange functionality that we describe below. Finally, we illustrate how to build a secure channel from an ideal key exchange functionality (or its realization).

Description of our example protocol. Our example protocol is informally defined in Alice-Bob-Notation in Figure 10. There are two roles A and B with public keys k_A and k_B . We assume that the parties know the public keys of each other (i.e., we assume some form of a public-key infrastructure). We also assume that the two parties that play role A and B have already established a unique SID $sid = (sid', pid_A, pid_B)$ for a session of the protocol where pid_A and pid_B are the party names of the parties in role A and B , respectively. The SID sid is supposed to be unique in the sense that it is used in no other session.²⁵

As shown in Figure 10, role A performs the following actions. First, A generates a nonce N_A , encrypts it together with the SID sid under B 's public key k_B , and sends the obtained ciphertext to B . Then, A waits to receive a message that is encrypted under A 's public key k_A and contains sid , A 's nonce N_A , and some nonce N_B . Finally, A encrypts sid and the received nonce N_B under k_B and sends the obtained ciphertext to B . If something goes wrong in one of these steps, A aborts immediately, and otherwise, if the protocol runs through successfully, A outputs N_A as the session key.

Role B performs the following actions according to Figure 10. First, B waits to receive a message that is encrypted under k_B and contains sid and some nonce N_A . Then, B generates a nonce N_B , encrypts it together with sid and the received nonce N_A under k_A , and sends the obtained ciphertext to A . Finally, B waits to receive a message that is encrypted under k_B and contains sid and B 's nonce N_B . If the protocol runs through successfully, B outputs N_A as the session key, and otherwise aborts immediately as soon as an error occurs.

The above protocol differs from the original Needham-Schroeder Public-Key protocol in two aspects: First, our protocol uses pre-established and unique SIDs. Second, since the party names are already included in the SID, we do not need to include them in the messages otherwise.

Ideal key exchange functionality. We use a standard ideal functionality \mathcal{F}_{ke} for (authenticated) key exchange (see, e.g., [9, 7, 22]). This functionality describes one session of an ideal key exchange between two parties/roles. It has two sets of I/O tapes, one for each role. It first waits to receive key exchange requests from a role. The simulator (ideal adversary) is informed about such requests. If the simulator sends a message for one role to finish and both roles have sent their key exchange requests to \mathcal{F}_{ke} before,²⁶ \mathcal{F}_{ke} outputs the session key to that role, where the session key is chosen uniformly at random from $\{0, 1\}^\eta$ by \mathcal{F}_{ke} (η is the security parameter). (Of course, other distributions for the session key could be used.) The simulator has the ability to corrupt \mathcal{F}_{ke} , i.e., send a corrupt message to \mathcal{F}_{ke} , before any of the two roles have output a

²⁵This SID can, for example, be established by exchanging nonces as explained in [1]. Both parties send nonces N'_A, N'_B in plain to each other. The SID then is the concatenation of the nonces and the party names: $sid = (N'_A, N'_B, pid_A, pid_B)$.

²⁶The condition that both roles have sent their key exchange requests models *authenticated* key exchange and could be dropped, but is convenient.

key. More precisely, in one request the simulator can corrupt one role. In this case, no matter which role was corrupted, if the simulator instructs a role to output a key, the simulator can determine the key that is output; the simulator can choose different keys for each of the two roles. Altogether, an uncorrupted \mathcal{F}_{ke} guarantees that the key a role receives after having sent a key exchange request to \mathcal{F}_{ke} is a freshly generated key that is only given to the roles involved in the key exchange. The key is indistinguishable from random for an adversary even if the key is output by one role before the end of the protocol. Also, if both roles receive a key, the two keys are guaranteed to coincide. Conversely, a corrupted \mathcal{F}_{ke} , i.e., one or both roles are corrupted, does not provide security guarantees; the keys the roles obtain (if any) are determined by the simulator and they do not need to coincide. As usual, the environment may ask whether a role of \mathcal{F}_{ke} has been corrupted; such a request is only allowed on the tapes of the respective role.

As mentioned, \mathcal{F}_{ke} captures only a single key exchange between two roles (played by arbitrary parties). Key exchange for an unbounded number of sessions and between an unbounded number of pairs of parties can be described by the multi-session version $!\mathcal{F}_{\text{ke}}$ of \mathcal{F}_{ke} .²⁷ We define the domain of SIDs that $!\mathcal{F}_{\text{ke}}$ accepts to be SIDs of the form $sid = (sid', pid_A, pid_B)$.²⁸ Intuitively, an instance of $!\mathcal{F}_{\text{ke}}$ addressed by such an SID is an ideal functionality for the key exchange between the parties pid_A and pid_B in session sid' . (Note that by this, there can be multiple sessions for the key exchange between the same two parties.)

Modeling our example protocol in the IITM model. We now specify the above example protocol as a protocol system $\mathcal{P} = !M_A | !M_B$ following the general approach outlined in Section 10.1. We define M_A and M_B to be IITMs that are addressed by structured SIDs of the form $sid = (sid', pid_A, pid_B)$. Every message received and output by M_A and M_B is prefixed by the SID. In other words, M_A and M_B are σ_{prefix} -session versions with SIDs of the described form.

Ultimately, we are interested in the protocol system $\mathcal{P} | \mathcal{P}_{\text{pke}}^{\text{js}} | !\mathcal{P}_{\text{pke}}$, where \mathcal{P} uses (the joint state realization for) public-key encryption, and for this system we want to show that it realizes the ideal multi-session key exchange functionality $!\mathcal{F}_{\text{ke}}$ (see above). The system $\mathcal{P} | \mathcal{P}_{\text{pke}}^{\text{js}} | !\mathcal{P}_{\text{pke}}$ will precisely model our example protocol. Note that in a run of this system (in the context of some environment) there can be multiple sessions (with multiple parties) of our example protocol running concurrently and a party will use the same public/private key pair across all sessions, due to the joint state realization.

However, we do not want to show $\mathcal{P} | \mathcal{P}_{\text{pke}}^{\text{js}} | !\mathcal{P}_{\text{pke}} \leq^{SS} !\mathcal{F}_{\text{ke}}$ directly, but rather make full use of the composition and joint state theorems as described in Section 10.1. Therefore, we first replace $\mathcal{P}_{\text{pke}}^{\text{js}} | !\mathcal{P}_{\text{pke}}$ by the ideal functionality $!\mathcal{F}_{\text{pke}}$ (the multi-session version of the multi-party version of \mathcal{F}_{pke}) and show that $\mathcal{P} | !\mathcal{F}_{\text{pke}} \leq^{SS} !\mathcal{F}_{\text{ke}}$. By the composition theorems, this even only needs to be shown for a single session of $\mathcal{P} / \mathcal{F}_{\text{ke}}$. By the composition theorem, we can later replace $!\mathcal{F}_{\text{pke}}$ by its (joint-state) realization $\mathcal{P}_{\text{pke}}^{\text{js}} | !\mathcal{P}_{\text{pke}}$.

Now, let us describe \mathcal{P} in more detail. We define M_A and M_B such that they use $!\mathcal{F}_{\text{pke}}$, where every copy of \mathcal{F}_{pke} , denoted by $\mathcal{F}_{\text{pke}}[sid, pid]$, is addressed by an SID sid of the form (sid', pid_A, pid_B) as discussed above and a party name pid (see also Section 10.1). We assume that parties can ask \mathcal{F}_{pke} for the public key of the party that owns \mathcal{F}_{pke} . More precisely, given an instance $\mathcal{F}_{\text{pke}}[sid, pid]$ of \mathcal{F}_{pke} , which is the “encryption/decryption-box” of party pid , we assume \mathcal{F}_{pke} to be defined in such a way that parties can ask this instance to receive the public key of pid (in session sid). Clearly, the realization of \mathcal{F}_{pke} would require a public-key infrastructure.

Both machines M_A and M_B have I/O tapes to connect to the environment and to connect to $!\mathcal{F}_{\text{pke}}$. Furthermore, they have network tapes to connect to the adversary. The system $\mathcal{P} | !\mathcal{F}_{\text{pke}}$ is depicted in Figure 11 and a run of this system is depicted in Figure 12 (w.r.t. an environment that created several sessions). Let $sid = (sid', pid_A, pid_B)$ be an SID. Next, we describe the actions performed by $M_A[sid]$ and $M_B[sid]$, i.e., the copies of M_A and M_B with SID sid , in more detail. Since $\mathcal{P} | !\mathcal{F}_{\text{pke}}$ has to provide the same

²⁷Note that an instance of \mathcal{F}_{ke} does not need to be aware of its SID, and hence, it is not necessary to provide a multi-session and multi-party version \mathcal{F}'_{ke} of \mathcal{F}_{ke} as a σ_{prefix} -session version.

²⁸See Section 5.2.1 for a discussion on domains of SIDs. Domains were also briefly mentioned in Section 2.




Figure 11: $\mathcal{P} \mid !\underline{\mathcal{F}}_{\text{pke}}$ single-session realizes $!\underline{\mathcal{F}}_{\text{ke}}$ (see Claim 1). In this figure, horizontal arrows denote network tapes and vertical arrows denote I/O tapes. Every pair of input/output tapes is labeled by a name. The actual tape names are decorations of this name, e.g., \mathcal{F}_{ke} has the input tapes io_A^{in} , io_B^{in} , and net^{in} and the output tapes io_A^{out} , io_B^{out} , and net^{out} .

I/O interface as $!\underline{\mathcal{F}}_{\text{ke}}$, $M_A[\text{sid}]$ and $M_B[\text{sid}]$ first wait to receive key exchange requests. If, say, $M_A[\text{sid}]$ receives such a request, $M_A[\text{sid}]$ asks its instance $\underline{\mathcal{F}}_{\text{pke}}[\text{sid}, \text{pid}_A]$ to create a public/private key pair and asks for the public key of the instance $\underline{\mathcal{F}}_{\text{pke}}[\text{sid}, \text{pid}_A]$,²⁹ analogously for $M_B[\text{sid}]$. Both $M_A[\text{sid}]$ and $M_B[\text{sid}]$ also inform the adversary about the key exchange requests that they got. In the rest of the protocol, for encryption under k_A and decryption with the corresponding private key, the copy $\underline{\mathcal{F}}_{\text{pke}}[\text{sid}, \text{pid}_A]$ is used by $M_A[\text{sid}]$ and $M_B[\text{sid}]$, where $M_B[\text{sid}]$ of course uses $\underline{\mathcal{F}}_{\text{pke}}[\text{sid}, \text{pid}_A]$ only for encryption; analogously for k_B . Since the joint state realization prefixes all messages by the corresponding SID, we define $M_A[\text{sid}]$ and $M_B[\text{sid}]$ such that sid is not added to the plaintexts. For example, the first message that $M_A[\text{sid}]$ encrypts (using $\underline{\mathcal{F}}_{\text{pke}}[\text{sid}, \text{pid}_A]$) is the nonce N_A rather than the concatenation of sid and N_A . The protocol messages (i.e., the ciphertexts) are sent and received via the network tapes (i.e., to/from the adversary who represents the network). If $M_A[\text{sid}]$ (and analogously for $M_B[\text{sid}]$) successfully finishes its run (according to the protocol specification), it outputs the session key (i.e., N_A) on the I/O tape to the environment.

We define static corruption of M_A as follows, following the description in Section 10.1; corruption of M_B is defined analogously. The adversary may corrupt $M_A[\text{sid}]$ at the beginning (when the adversary is informed about the key exchange request that $M_A[\text{sid}]$ got) by sending a corrupt message to this instance. If corrupted, $M_A[\text{sid}]$ forwards all input to the adversary and lets the adversary determine its output. Additionally, $M_A[\text{sid}]$ allows the adversary to decrypt messages using $\underline{\mathcal{F}}_{\text{pke}}[\text{sid}, \text{pid}_A]$. (Encryption requests could also be allowed, but they do not make the adversary stronger, because he knows the public keys of all parties and can encrypt messages by himself.)

Just as in the case of \mathcal{F}_{pke} , the environment may ask the instance $M_A[\text{sid}]$ (analogously for $M_B[\text{sid}]$) about its corruption status. It returns *yes* if and only if it was directly corrupted or $\underline{\mathcal{F}}_{\text{pke}}[\text{sid}, \text{pid}_A]$ is corrupted. (We note that the adversary has the ability to corrupt a copy of \mathcal{F}_{pke} without explicitly corrupting the respective copy of M_A or M_B , and hence, these copies will still follow the prescribed protocol. However, M_A and M_B report to the environment that they are corrupted if their instance of \mathcal{F}_{pke} that they use is corrupted, because the key exchange protocol cannot guarantee a secure key exchange in this situation.)

It is easy to see that the protocol system $\mathcal{P} \mid !\underline{\mathcal{F}}_{\text{pke}}$ is environmentally strictly bounded. Hence, all preconditions of the composition theorems about runtime properties are trivially satisfied. Furthermore, it is an σ_{prefix} -session version (because all messages are prefixed by an SID) and we can apply the reasoning as described in Section 10.1.

First, we state that the protocol as modeled above is secure in a single-session setting (see also Figure 11).

Claim 1. $\mathcal{P} \mid !\underline{\mathcal{F}}_{\text{pke}} \leq_{\sigma_{\text{prefix}}\text{-single}}^{\text{SS}} !\underline{\mathcal{F}}_{\text{ke}}$.

The proof of this claim, which we omit for the sake of brevity, is relatively simple because one needs to

²⁹If that instance is invoked for the first time, and hence, has not generated a key pair yet, it first generates a key pair.

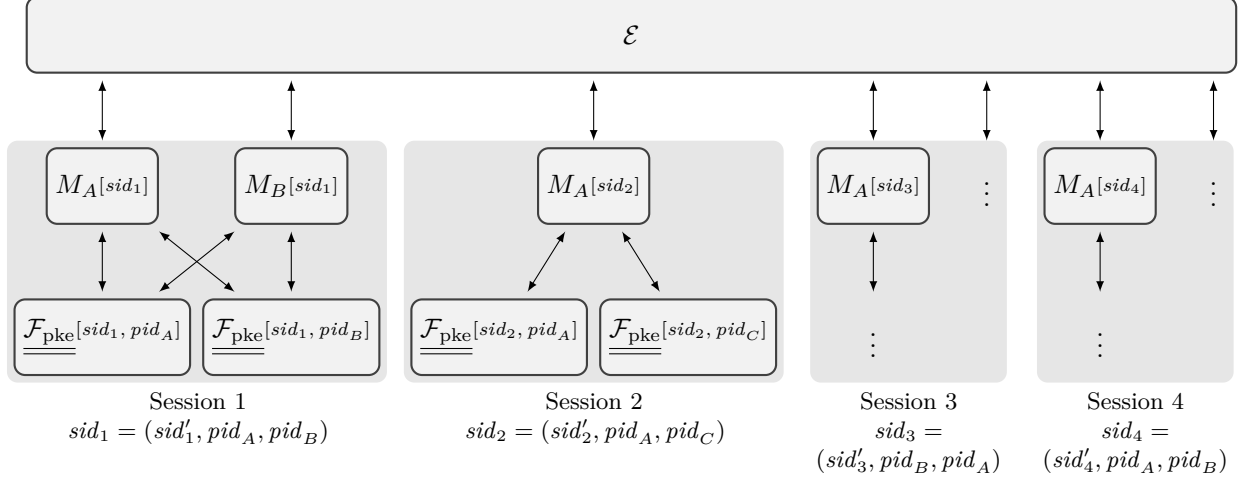


Figure 12: A run of $\mathcal{E} | \mathcal{P} | !\mathcal{F}_{\text{pke}}$ for some environment \mathcal{E} that created four sessions: two sessions with party pid_A playing role A and pid_B playing role B (Sessions 1 and 4), one session with pid_A playing role A and pid_C playing role B (Session 2), and one session with pid_B playing role A and pid_A playing role B (Session 3). Every box denotes an instance of a machine that has been created in this run. For example, $M_A[sid_1]$ is the instance of M_A that is addressed by sid_1 (i.e., party pid_A in role A talking to pid_B in session sid'_1) and $!\mathcal{F}_{\text{pke}}[sid_1, pid_A]$ is the instance of \mathcal{F}_{pke} that is addressed by sid_1 and pid_A (modeling the encryption/decryption-box of party pid_A in session sid_1). The arrows denote the connections between the instances via I/O tapes and addressing with SIDs. In addition, all instances connect to \mathcal{E} via network tapes. These connections are not displayed.

consider only a single session of the protocol and encryption is idealized by \mathcal{F}_{pke} . In particular, the proof can be done by information-theoretic arguments alone, without reductions to the security of the encryption scheme. We note that the above claim even holds in an information-theoretic setting with unbounded environments.

We further note that Claim 1 does not hold if the nonce N_B (instead of N_A) is used as the session key. Canetti and Herzog [8] showed that the Needham-Schroeder-Lowe (NSL) protocol is insecure when N_B is used as the session key. The reason here is the same: When A has output the session key, N_B in this case, but B has not yet received the last protocol message (a situation that an environment can easily create), then the environment can use B to test whether it is in the real world (i.e., interacting with $\mathcal{P} | !\mathcal{F}_{\text{pke}}$) or in the ideal world (i.e., interacting with $!\mathcal{F}_{\text{ke}}$); see [8] for details. Interestingly, even if N_B is used as a session key, the protocol would still satisfy the key usability property, i.e., it would realize $!\mathcal{F}_{\text{keyuse}}$, where $\mathcal{F}_{\text{keyuse}}$ is an ideal functionality for key usability introduced in [22]. This functionality works just like \mathcal{F}_{ke} but instead of outputting the session key it allows the party to use the session key in cryptographic operations, e.g., for (ideal) symmetric encryption. This functionality is often more useful than \mathcal{F}_{ke} in applications (see below for an example), and being weaker than key indistinguishability, it can be realized by more protocols.

As described in Section 10.1, from Claim 1 we can directly deduce, by applying the composition theorem and the joint state theorem for \mathcal{F}_{pke} , that \mathcal{P} is secure when multiple sessions run concurrently and a party uses the same public/private key pair across all sessions. Moreover, by the composition theorem, we can replace \mathcal{F}_{pke} by its realization \mathcal{P}_{pke} using an IND-CCA2 secure public-key encryption scheme (in conjunction with a public-key infrastructure). This deduction is depicted in Figure 13.

Claim 2. $\mathcal{P} | \mathcal{P}_{\text{pke}}^{\text{js}} | !\mathcal{P}_{\text{pke}} \leq^{SS} !\mathcal{F}_{\text{ke}}$.

We note that the protocol system $\mathcal{P} | \mathcal{P}_{\text{pke}}^{\text{js}} | !\mathcal{P}_{\text{pke}}$ precisely models our example protocol. A run of this system is depicted in Figure 14. As already mentioned, every party uses its public/private key pair across all protocol sessions. This is due to the use of the joint state realization. Furthermore, the joint state realization yields that the SID is added to every plaintext before encryption just as in the example protocol. Finally,

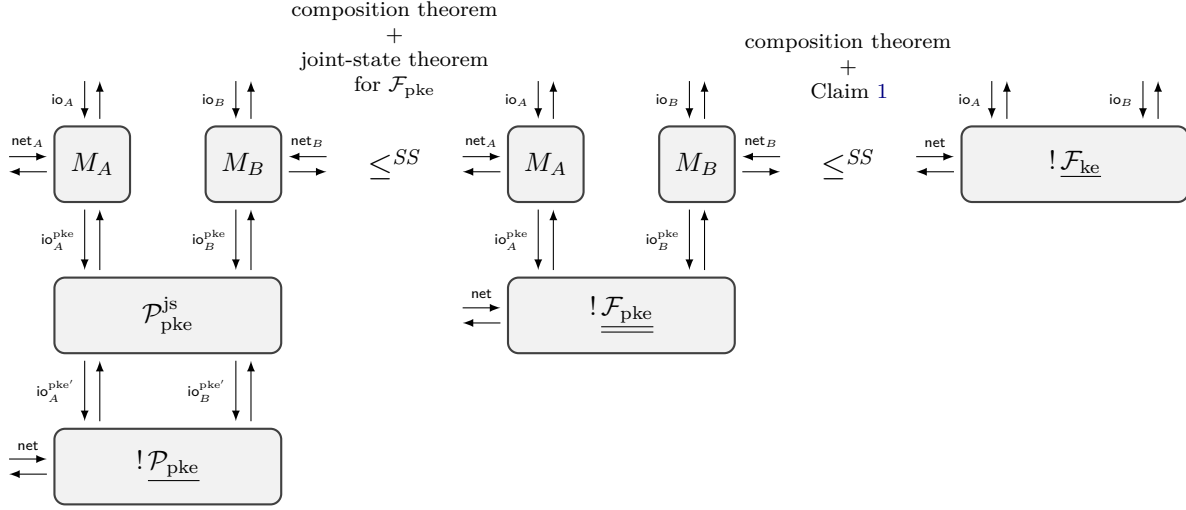


Figure 13: $\mathcal{P} \mid \mathcal{P}_{\text{pke}}^{\text{js}} \mid !\mathcal{P}_{\text{pke}}$ realizes $!\mathcal{F}_{\text{ke}}$ (see Claim 2). Tapes are labeled as in Figure 11.




Figure 14: The run of $\mathcal{E} \mid \mathcal{P} \mid \mathcal{P}_{\text{pke}}^{\text{js}} \mid !\mathcal{P}_{\text{pke}}$ that corresponds to the run depicted in Figure 12. In this run, the joint state between the sessions consists of the instances of \mathcal{F}_{pke} . For example, the instance $\mathcal{F}_{\text{pke}}[\text{pid}_A]$ (that models the encryption/decryption-box of party pid_A) is used by all four sessions.

the realization \mathcal{P}_{pke} uses a standard IND-CCA2 secure public-key encryption scheme, no modifications are required, again just like in our example protocol.

Of course, $\mathcal{P} | \mathcal{P}_{\text{pke}}^{\text{js}} | !\mathcal{P}_{\text{pke}}$ precisely models the example protocol only because the protocol uses pre-established SIDs in the way they are used by the joint state realization. As already mentioned in Section 10.1, one can also precisely model and analyze protocols in the IITM model even if they do not use pre-established SIDs (or not in the way stipulated by the joint state realization) [22].

We further emphasize that static corruption is modeled in a strong sense: The adversary can corrupt public/private keys for every party individually (upon first use) and he can corrupt individual roles of sessions (upon session start). Motivation for this kind of fine-grained corruption was already given in Section 10.1. Clearly, this includes that the adversary can take complete control of a party by corrupting its public/private key and its roles in all sessions.

Building secure channels from key exchange protocols. To illustrate how more complex protocols can be built from the ideal key exchange functionality (or a realization thereof) using the composition theorems, we consider a secure channel protocol.

An ideal secure channel protocol \mathcal{F}_{sc} (for two-parties/two-roles) works similarly to \mathcal{F}_{ke} described above: It expects to receive a request to establish a secure channel from both roles and upon every such request informs the simulator about it, who can then decide to corrupt a role of \mathcal{F}_{sc} , at which point the simulator obtains full control of \mathcal{F}_{sc} . At some point, the simulator can send a “session established” message to \mathcal{F}_{sc} for one of the two roles. If both roles have sent their requests to establish a secure channel to \mathcal{F}_{sc} before (this models authentication, similarly to \mathcal{F}_{ke}), \mathcal{F}_{sc} will inform the role that the session is established. From this point on, this role can send and receive messages to/from \mathcal{F}_{sc} . If a role sends a message via \mathcal{F}_{sc} , this message is delivered to the other role if the adversary instructs \mathcal{F}_{sc} to deliver the message; the adversary is informed about every “message send” request and learns the length of the message that is to be delivered. The functionality \mathcal{F}_{sc} guarantees confidentiality and integrity of the messages, that messages are not dropped, and delivered in order. As usual, the environment can ask about the corruption status of the roles in \mathcal{F}_{sc} (on the respective tapes). For more details, see, e.g., [3, 22]. Just like \mathcal{F}_{ke} , \mathcal{F}_{sc} captures only a single secure channel between two roles. Secure channels for an unbounded number of sessions and between an unbounded number of pairs of parties can be described by the multi-session version $!\mathcal{F}_{\text{sc}}$ of \mathcal{F}_{sc} where we define the domain of SIDs that \mathcal{F}_{sc} accepts to be SIDs of the form $sid = (sid', pid_A, pid_B)$.

An example of a simple real secure channel protocol is the following: The roles first establish a session key (using a key exchange protocol). Now, messages are encrypted under the session key using an authenticated symmetric encryption scheme, i.e., an IND-CPA and INT-CTXT secure scheme, along with a counter that is added to every message (in order to prevent messages to be dropped or reordered).

Formally, let \mathcal{P}_{ke} be a secure (multi-session) key exchange protocol, i.e., $\mathcal{P}_{\text{ke}} \leq^{SS} !\mathcal{F}_{\text{ke}}$. For example, \mathcal{P}_{ke} could be the example protocol $\mathcal{P} | \mathcal{P}_{\text{pke}}^{\text{js}} | !\mathcal{P}_{\text{pke}}$ discussed above. Furthermore, let $\mathcal{P}_{\text{sc}} = !M_A^{\text{sc}} | !M_B^{\text{sc}}$ be a secure channel protocol that uses the key exchange protocol \mathcal{P}_{ke} . The protocol system \mathcal{P}_{sc} is defined similarly to the key exchange protocol \mathcal{P} above except that it does not use \mathcal{F}_{pke} but \mathcal{P}_{ke} : Instances of M_A^{sc} and M_B^{sc} are again addressed by SIDs of the form (sid', pid_A, pid_B) . They first, at the I/O interface, wait to be instructed to establish a secure channel; the adversary is informed about such a request and can decide to corrupt that instance (see below). The instance then, if not corrupted, sends a key exchange request to \mathcal{P}_{ke} using its SID (sid', pid_A, pid_B) . When it receives a session key, it outputs to the party that the secure channel is established. The instance is now ready to receive and send messages. To send messages it uses a counter and an authenticated symmetric encryption scheme as sketched above.

As mentioned, an instance M_A^{sc} (or M_B^{sc}) can be corrupted by the adversary when the adversary is informed about the request to establish a secure channel by M_A^{sc} . When corrupted, M_A^{sc} provides full control to the adversary, including the interface of M_A^{sc} to \mathcal{P}_{ke} . (We do not require \mathcal{P}_{ke} to be corrupted as well.) If an instance is asked whether it is corrupted by the environment, it returns *yes* iff it was explicitly corrupted by the adversary or the corresponding role in the instance of \mathcal{P}_{ke} it interacts with was corrupted. We note that, by this, (static) corruption is again modeled in a quite general and fine-grained way.

Now, we want to show security of the protocol system $\mathcal{P}_{\text{sc}} | \mathcal{P}_{\text{ke}}$, i.e., the composition of the secure channel




Figure 15: $\mathcal{P}_{sc} | !\mathcal{F}_{ke}$ single-session realizes $!\mathcal{F}_{sc}$ (see Equation (24)). Tapes are labeled as in Figure 11.

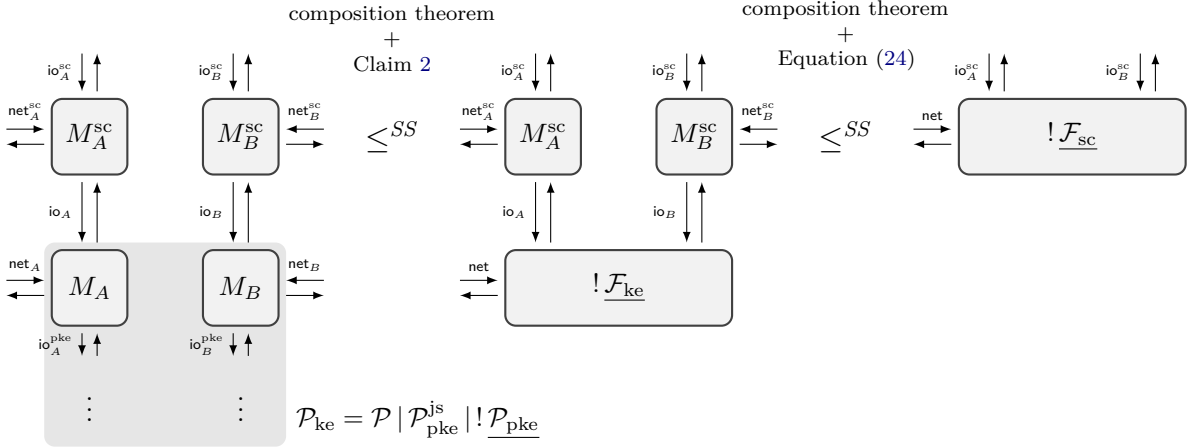


Figure 16: $\mathcal{P}_{sc} | \mathcal{P}_{ke}$ realizes $!\mathcal{F}_{sc}$ (see Equation (25)). Tapes are labeled as in Figure 11. There can be multiple instances of the machines M_A^{sc} and M_B^{sc} as well as of the machines M_A and M_B .

protocol \mathcal{P}_{sc} and the key exchange protocol \mathcal{P}_{ke} . More precisely, we want to show that $\mathcal{P}_{sc} | \mathcal{P}_{ke}$ realizes $!\mathcal{F}_{sc}$, i.e., $\mathcal{P}_{sc} | \mathcal{P}_{ke} \leq^{SS} !\mathcal{F}_{sc}$. For this purpose, it suffices to show that:

$$\mathcal{P}_{sc} | !\mathcal{F}_{ke} \leq_{\sigma_{\text{prefix-single}}}^{SS} !\mathcal{F}_{sc} . \quad (24)$$

This statement is depicted in Figure 15. From (24), $\mathcal{P}_{ke} \leq^{SS} !\mathcal{F}_{ke}$, and by Theorem 8 and 10, we immediately obtain the following (see also Figure 16):

$$\mathcal{P}_{sc} | \mathcal{P}_{ke} \leq^{SS} !\mathcal{F}_{sc} . \quad (25)$$

This statement says that multiple sessions of the secure channel protocol realize multiple sessions of the ideal secure channel functionality.

Note that the proof of (24) is relatively simple because it involves reasoning about only a single session of the protocol and because the key exchange is idealized. But it still requires a reduction proof to the security of the authenticated encryption scheme. This can be avoided if we do not use \mathcal{F}_{ke} , which outputs the session key, but instead an ideal key usability functionality \mathcal{F}_{keyuse} (as already mentioned above), which does not output the session key but instead allows \mathcal{P}_{sc} to use the session key in an ideal way (similar to encryption and decryption with \mathcal{F}_{pke}). Such a key usability functionality has been proposed in [22].

11 Related Work

As already mentioned in the introduction, the IITM model proposed here coincides with the (old) IITM model proposed in [18], except that we now use more general runtime notions, which conceptually build on those proposed by Hofheinz et al. in [16].

We concentrate our comparison of the IITM model with other models on the latest version of Canetti’s UC model [3] and the recent GNUC model [14] by Hofheinz and Shoup. Other models, including [27, 19, 5], are not considered as they either fix the runtime of a machine by a fixed polynomial in the security parameter, which has many problems as discussed in [18, 13], or do not consider systems with a polynomial number of copies of machines, which, however, is needed to even state general composition theorems. The recent Abstract Cryptographic model by Maurer and Renner [24] does not define a full model on the machine level, and hence, cannot directly be compared with the IITM model.

11.1 UC Model

The model by Canetti [3], often referred to as the UC model or the UC framework, has been evolving over time. In the first model, [2], the runtime of interactive Turing machines (ITMs) was bounded by a fixed polynomial in the security parameter, which as mentioned causes several problems. Therefore, Canetti proposed other variants of the model where the runtime of the machines may depend on the length of the input that they received. The latest variant of the UC model is presented in [3]. We will discuss this model in more detail in what follows. In-depth discussions of this model can also be found in [18, 20, 16, 14].

The UC model has several severe technical flaws, for example, concerning the validity of the composition and joint state theorems. While, formally speaking, the gaps in the model often invalidate cryptographic proofs carried out in the UC model in the literature, conceptually on a higher level of abstraction the cryptographic results and proofs are not invalidated. However, it is clearly highly unsatisfying if statements and proofs do not rest on solid ground.

In [14], a concrete example is given showing that the composition theorem does not hold true in the UC model. The main reason is that interactive machines do not necessarily know who they are interacting with. In the composition theorem, this leads to the problem that a simulator might not know what kind of simulation needs to be performed. As a result, one can construct, (natural) protocols \mathcal{Q} , \mathcal{F} , and \mathcal{P} such that \mathcal{P} realizes \mathcal{F} but the composition of \mathcal{Q} and \mathcal{P} does not realize the composition of \mathcal{Q} and \mathcal{F} , in contradiction to what the composition theorem says.

Another reason why the composition theorem is problematic is the fact that there does not exist a so-called dummy adversary in the UC model, i.e., an adversary which forwards all messages back and forth between the environment and the protocol. This in turn is due to the way the runtime of machines is defined in the UC model. In the UC model the *total* runtime of the ITMs is bounded by a polynomial in the security parameter and the length of the input received on I/O tapes (minus the runtime provided to “subroutine ITMs”); input received on communication tapes (the network interface) does not increase the runtime resources available to a machine. Now, in the UC model the adversary has I/O tapes to the environment but only network tapes to the protocol. Consequently, such an adversary cannot forward arbitrarily many and arbitrarily long messages from the protocol to the environment since it might not have the runtime resources to do that. In the UC model, the dummy adversary has to ask for resources from the environment in order to be able to forward arbitrary messages from the protocol to the environment. With a full-fledged dummy adversary, i.e., one that can forward all messages without asking for resources, one can prove that a real protocol realizes an ideal one by simply constructing a simulator for the dummy adversary (instead of constructing simulators for all real adversaries). However, as shown in [16], the restricted dummy adversary available in the UC model—the one that has to ask the environment for resources—is incomplete in this respect: considering only this adversary does not guarantee security for all adversaries. Now, the composition theorem in the UC model has only been proven with respect to the (restricted) dummy adversary. But since in the UC model dummy adversaries are incomplete in the sense explained above the proof of the composition theorem is incomplete—in addition to the fact that the composition theorem does not hold true due to the above mentioned orthogonal problem. It is not obvious how to fix this problem since dummy adversaries

are a central tool in the proof of the composition theorems, also in other models. In the IITM model, an unrestricted dummy adversary is obviously available as shown in Lemma 5. Completeness of the dummy adversary follows from Theorem 7 (see also Remark 11).

It was shown in [20] that the general joint state theorem (such theorems have been discussed in Section 10.1) does not hold true in the UC model either; Canetti and Rabin proposed this theorem in [10]. This is mainly due to the way the runtime of ITMs is defined in the UC model (see above). As a consequence of this definition, by sending many messages to an ITM (on the network interface), the ITM can be forced to stop. Also, in general, a single ITM M cannot simulate a concurrent composition of a fixed finite number of ITMs or an unbounded number of (copies of) ITMs: Assume that the ITM M internally simulates (a fixed number of) other machines. Now, by sending many messages to M intended for some internally simulated machine, M will eventually stop, and hence, cannot simulate the other machines internally anymore, even though, in the actual composition these machines could still perform actions. Due to the way the runtime of machines and systems of machines is defined in the IITM model these problems do not occur. In particular, in the IITM model one machine can simulate any system of machines in any polynomial context (see Lemmas 6 and 7) and an IITM cannot be exhausted, it can perform computations in every activation. In particular, it can always read its input in both modes, `CheckAddress` and `Compute`.

Besides the technical flaws of the UC model sketched above, the UC model is also quite complex and many things are hard-wired into the model. For example:

- The UC model uses control functions, dummy parties, and subroutines.
- The runtime a machine has depends on how much data it has sent to its subroutines, a quite complex and not very intuitive runtime notion.
- The way machines are addressed (by SIDs and PIDs) and the way corruption is handled is fixed in the UC model.

In the IITM model, control functions, dummy parties, and the notion of a subroutine are not needed. The runtime notions employed in the IITM model are quite intuitive and straightforward. The `CheckAddress` mode allows for a generic and very flexible addressing mechanism where we do not have to talk about SIDs and PIDs. This is important in order to model protocols in a faithful way [22]. Also, corruption is not hard-wired into the model. The way corruption is handled is entirely left to the specification of the protocols and functionalities, and hence, is very flexible (as illustrated in Section 10).

As pointed out in [14], there are in fact two problems with the way corruption is handled in the UC model. First, in order for the adversary to be able to corrupt a party, the adversary needs to receive an authorization message from the environment. Such a message contains a machine name which the adversary is allowed to corrupt. However, there are different sets of machines in the real and the ideal world, and hence, it is unclear what such a message means in the different worlds. Second, in the UC model an adversary can create machines and determine their program and IDs. By this he can “impersonate” or “hijack” honest machines without actually corrupting them. But then it is, in most cases, impossible for the simulator to achieve its goals because it is not allowed to corrupt the ideal functionality. So, as already explained in [14], formally most security claims in the literature are simply false.

The UC model is also less expressive than the IITM model. For example:

- In order to deal with joint state or global setup the UC model needed to be extended [10, 6]. This is not necessary in the IITM model as discussed in Section 10.1. The joint state theorem is even a direct consequence of the composition theorem in the IITM model. Since the UC model fixes things like the addressing of machines by SIDs and PIDs as well as the way machines can be corrupted, the composition and joint state theorems are less general. For example, if the way machines can be corrupted is changed, these theorems would have to be reproven.
- The way the runtime is defined in the UC model makes it hard to formulate protocols and functionalities. In fact, many, maybe most protocols/functionalities found in the literature cannot be expressed in the UC model. For example, it is very common that a protocol machine/functionality drops messages

received from a communication tape (the network interface) if they do not have the correct format. However, as explained before, in the UC model machines can easily be “exhausted” by sending them many useless messages on the communication tapes and there is nothing a machine can do about it. Thus, a machine will not have sufficient resources to even read incoming messages at some point and it will be forced to stop, which typically is not the intended behavior of functionalities/protocols. Dealing with this issue in a formally correct way in the UC model is, if at all possible, very tricky and has, as far as we know, never been done in the literature. Note that in the IITM model one can easily formulate protocols and functionalities within the class of environmentally bounded protocol systems that simply reject messages that have the wrong format.

11.2 GNUC model

The GNUC model was proposed by Hofheinz and Shoup in [14]. Compared to the IITM model, many things are hard-wired into the model and fixed in a specific way, as explained next.

Features of the GNUC model compared to the IITM model. In the GNUC model, a strict hierarchical structure is assumed for protocols: In a run of a system, every machine has a unique caller, hence machines form a call tree. Also, an acyclic (static) subroutine graph is assumed. In the IITM model, we do not put such restrictions on systems. A (protocol) system is simply a system of IITMs which are connected via (named) I/O tapes, with network tapes to the adversary; the notion of a subroutine is not even explicitly defined in the IITM model. For instance, it is possible to specify a system of the form $\mathcal{P}_1 | \mathcal{P}_2 | \mathcal{P}_3$ where all three protocols are mutually connected via I/O tapes or where \mathcal{P}_1 is connect to \mathcal{P}_2 , \mathcal{P}_2 to \mathcal{P}_3 , and \mathcal{P}_3 to \mathcal{P}_1 . In particular, no acyclic subroutine graph is imposed. Of course, if desired, acyclic subroutine graphs and call trees can also be modeled in the IITM model. For instance, if \mathcal{P}_2 is the only subroutine that the protocols \mathcal{P}_1 and \mathcal{P}_3 should be able to use, one would have I/O tapes only between \mathcal{P}_1 and \mathcal{P}_2 as well as between \mathcal{P}_3 and \mathcal{P}_2 , but not between \mathcal{P}_1 and \mathcal{P}_3 .³⁰ If one requires call trees, one can easily define the machines in \mathcal{P}_1 and \mathcal{P}_2 in such a way that they generate different IDs to address (different) copies of \mathcal{P}_2 . (As usual, the `CheckAddress` mode of machines in \mathcal{P}_2 could be defined as (σ -)session versions.) However, it is also possible to specify the protocols in such a way that, say, both an instance of \mathcal{P}_1 and an instance of \mathcal{P}_3 access the same instance of \mathcal{P}_2 , as required for example in joint state realizations. It is also not a problem to specify that all (instances of) machines in a system, including (the machines of) the protocol, the environment, and the adversary access the same instance of another machine, as required for example for settings with global state.

The GNUC model also fixes the way instances of machines are addressed by SIDs and PIDs in a very specific way. PIDs have to contain, besides the actual PID, a label whether a machine is a so-called real or ideal machine. Also, SIDs need to contain, besides the actual SID, the protocol name. It is required that if a machine calls a subroutine, then it extends its SID in a unique way. In this way, a call tree is enforced where every machine has a unique caller. This, however, is problematic in settings with joint and/or global state (see below). In contrast, the IITM model does not fix a specific addressing mechanism, but allows for a very flexible way of addressing machines by a general addressing mechanism, namely using the `CheckAddress` mode of IITMs. In particular, it would be easy to model exactly the kind of addressing enforced in the GNUC model. However, other ways of addressing machines in the IITM are possible and often desirable or necessary (e.g., for joint/global state), as illustrated in Section 10.

Corruption is also fixed in a specific way in the GNUC model. To corrupt a machine, a machine expects to receive a corrupt message at its I/O interface. This implies that in a call tree first the top-level machines have to be corrupted by the environment (the environment in the GNUC model can only access machines on the top-level). If a machine is corrupted, the adversary can corrupt subroutines of that machine so that altogether corruption spreads from the top of a call tree to the bottom. If a machine is corrupted, it forwards all messages to the adversary. However, the adversary is not allowed to send messages to the I/O interface of subroutines of a corrupted machine, except for special corrupt messages. This form of corruption is quite specific and restricted, as discussed below. In contrast, in the IITM model corruption is not hard-wired into

³⁰We note that even if there are I/O tapes between say \mathcal{P}_1 and \mathcal{P}_3 , this does not mean that the protocols have to use them. The protocols could be specified in such a way that they ignore such tapes and don’t communicate over such tapes.

the model. We do not fix any form of corruption at all. Corruption can be modeled in a very flexible way as part of the specification of protocols, as illustrated in Section 10. In particular, it would be easy to model the kind of corruption fixed in the GNUC model in the IITM model, i.e., to specify protocol systems in such a way that corruption is modeled as in the GNUC model. However, this is just one possible way of modeling corruption in the IITM model, other sometimes more desirable formulations, see the explanation below, are illustrated in Section 10.

As mentioned before, in the GNUC model, the environment can only access top-level machines. It cannot directly access lower-level machines, which, for example, in settings with global state would be necessary. Also, unlike the IITM model, the adversary cannot create machines. This can be useful, for instance, in a secure channel protocol where a machine should be created on the receiver side upon receipt of the first message.

In order to guarantee that a system consisting of an environment, an adversary, and a protocol runs in polynomial time in the security parameter, in the GNUC model the runtime of the adversary is restricted in that the number of bits he is allowed to send to the protocol is bounded by a polynomial in the number of bits the adversary received from the environment. This is a kind of acyclicity condition in terms of the length of the messages that may flow between the different system components (environment, adversary, protocol). However, by this the adversary/simulator is quite restricted. For example, consider the situation where an environment invokes many instances of a protocol/functionality. Now, even if every such instance contacts the adversary, the adversary could not respond to all instances because he is restricted in the number of bits he may send to the protocol. To slightly mitigate this problem, the concept of so-called “invited messages” is introduced in the GNUC model. These messages can be sent by the protocol to invite the adversary to send a message, even though he otherwise would not be allowed to send a message. However, the adversary is only allowed to send exactly the invited messages received from the protocol before. So if the protocol cannot know in advance which message the adversary will send, this mechanism is still insufficient. For example, due to the restriction put on the adversary/simulator, it is impossible to realize certain common functionalities in the GNUC model, e.g., for public-key and symmetric encryption, digital signatures, and MACs (see [3, 23, 21, 20]). In the IITM model, the adversary/simulator is not restricted in this way. In particular, the mentioned functionalities can easily be expressed and realized in the IITM model.

Due to the design choices and restrictions of the GNUC model sketched above, the GNUC model has several disadvantages in terms of simplicity, expressivity, and generality as discussed next.

Simplicity. The GNUC model introduces many concepts and fixes many details, as sketched above: distinction between regular and ideal machine, dummy parties, hierarchical protocol structure, very specific addressing mechanism with SID and PID having specific forms, specific form of corruption, invited and uninvited messages. This makes the model much more complex than the IITM model, where these concepts are not required and these things are not fixed. Also, when formulating protocols and functionalities in the GNUC model one has to think more about runtime issues. In particular, due to the need for invited messages the specifications tend to be artificial. In the IITM one does not have to think about runtime issues too much when specifying protocols and functionalities because typical formulations will simply be environmentally (strictly) bounded. A slight advantage of the runtime notion used in the GNUC model is that it guarantees that if a protocol running with an ideal functionality is polynomially bounded in the GNUC sense, then this is also the case when the ideal functionality is replaced by its realization. However, since it is typically easy to see whether a protocol is polynomially bounded in the IITM sense, i.e., whether it is environmentally (strictly) bounded, and since in most applications protocol systems are environmentally (strictly) bounded anyway, we do not consider this a big advantage; even more so considering the disadvantages of the runtime notion used in the GNUC model, which makes the model harder to use and less expressive (necessity of invited messages, certain natural functionality cannot be realized at all due to the restrictions put on simulators).

Expressivity and generality. The IITM model is much more expressive than the GNUC model in many respects. As already mentioned, the IITM model does not impose a hierarchical structure on protocol systems. This, in conjunction with the fact that in the IITM model the environment is not restricted to only access top-level protocols, not only allows for modeling a bigger class of protocols (see above), but also allows

us to handle joint and global state without changing the model. The joint state theorem is even a direct consequence of the composition theorem. In the GNUC model, dealing with joint state and global state required non-trivial extensions of the model and to reprove theorems, such as the composition theorem, where global state is not even dealt with in its full generality. The IITM model is also more expressive and flexible in terms of the kind of corruption it can handle, as already explained above and illustrated in Section 10. In fact, the way corruption is handled (and fixed) in the GNUC model is quite restricted and sometimes insufficient. As explained, machines can only be corrupted top-down. However, it might make sense to corrupt a subroutine without corrupting the top-level protocol. For example, a protocol might use some kind of double encryption to guarantee that if one encryption fails, confidentiality is still guaranteed. So the higher level protocol can potentially still achieve its task even though parts of the subroutines are corrupted. Another example is that in e-voting a fraction of mixnet servers can be corrupted while the overall e-voting system is still secure. Such settings cannot be handled in the GNUC model. Moreover, the fact that in the GNUC model an adversary cannot send messages via a corrupted machine to uncorrupted subroutines is also a restriction. Consider for example a system where several processes use the same secure channel (which might be established using some hardware tokens). If one of the processes is corrupted, the adversary should be allowed to send and receive messages over the secure channel via the corrupted process. In the current formulation of the GNUC model this cannot be modeled. Another aspect where the IITM model is much more expressive and flexible is the addressing of machines. This, as explained above, is fixed in the GNUC model by a specific use and form of SIDs and PIDs. Since, as explained in [22], such IDs are used in an essential way in protocols (not just for the purpose of addressing instances), the kind of protocols that can be designed in the GNUC model is restricted and the faithful analysis of existing protocols is often impossible. The notion of runtime used in the GNUC model is more restricted than the one used in the IITM model as well. In particular, certain (natural) functionalities cannot be realized due to the notion of runtime that is used, as already mentioned above. Finally, we remark that in the GNUC model, unlike the IITM model, the adversary may not create machines. However, this is inconvenient. It, for example, does not allow one to model a secure channel in such a way that only one party indicates that it wants to establish a secure channel and an instance of the other party is created when that instance receives the first message from the network.

The fact that the IITM model does not a priori fix certain details (addressing of machines, corruption, protocol structure, etc.) has also the big advantage that the theorems proved, such as the composition and joint state theorems, are much more general than in the GNUC model, as they hold true independently of specific choices; put otherwise, they hold true for all specific choices. The generality also often makes the proofs simpler and more elegant in that they are not cluttered with unessential details. In the GNUC model, the proof of the joint state theorem, for example, used in an essential way the concept of invited messages and the way corruption is defined.

Compared to the GNUC model, one could take the simplicity, expressivity, and generality of the IITM model against the model: For the design and analysis of concrete protocols several things have to be fixed, as part of the protocol specification, because they are not a priori fixed in the IITM model. However, we consider the high flexibility a big advantage because it makes the model much simpler as well as more expressive and general compared to the GNUC model and other models, as explained above. Also, one can easily define conventions (as illustrated in Section 10.1), e.g., for corruption, and later refer to them for concrete design and analysis tasks.

A Proof of Lemma 9

Let \mathcal{P}, \mathcal{Q} be protocol systems such that $!\mathcal{P}$ is environmentally bounded, \mathcal{Q} is environmentally bounded, and $\mathcal{P} \cong \mathcal{Q}$. Let $\mathcal{E} \in \text{Env}(!\mathcal{P}) = \text{Env}(!\mathcal{Q})$. In the following, by a hybrid argument where we replace the copies of \mathcal{P} by copies of \mathcal{Q} , we show that $\mathcal{E} | !\mathcal{Q}$ is almost bounded, and hence, that $!\mathcal{Q}$ is environmentally bounded.

By Lemma 7, we may assume that \mathcal{E} is a single IITM which, in mode `CheckAddress`, accepts all messages. Moreover, we may assume, without loss of generality, that \mathcal{E} is such that every message m that \mathcal{E} outputs (except if m is output on tape `decision`) is prefixed by some SID, i.e., $m = (s, m')$ for some bit strings s and

m' :³¹ since \mathcal{E} will only interact with session versions, message not of the form (s, m') would be rejected by these session versions anyway. Since \mathcal{E} is universally bounded, it follows that there exists a polynomial $p_{\mathcal{E}}$ such that the number of different sessions (i.e., messages with distinct SIDs output by \mathcal{E}) is bounded from above by $p_{\mathcal{E}}(\eta + |a|)$ (where η is the security parameter and a is the external input).

In what follows, let \underline{Q}' be the variant of \underline{Q} obtained from \underline{Q} by renaming every tape c occurring in \underline{Q} to c' . Analogously, let \underline{P}'' be obtained from \underline{P} by renaming every tape c occurring in \underline{P} to c'' . By this, we have that \underline{P} , \underline{Q}' , and \underline{P}'' have pairwise disjoint sets of external tapes, and hence, these systems are pairwise connectable.

We now define an IITM \mathcal{E}_r (for every $r \in \mathbf{N}$) which basically simulates \mathcal{E} and which will run in the system $\mathcal{E}_r | \underline{P}'' | \underline{Q}' | \underline{P}$ or $\mathcal{E}_r | \underline{P}'' | \underline{Q}' | \underline{Q}$, respectively. The IITM \mathcal{E}_r randomly shuffles the copies of the protocols invoked by \mathcal{E} by choosing a random permutation on $\{1, \dots, p_{\mathcal{E}}(\eta + |a|)\}$. Intuitively, this permutation is needed because \mathcal{E} must not “know” whether the copy invoked in the system $! \underline{P}'' | \underline{Q}' | \underline{P}$ and $! \underline{P}'' | \underline{Q}' | \underline{Q}$, respectively, is a copy of $\underline{P}/\underline{P}''$ or $\underline{Q}/\underline{Q}'$. This is needed in the proof of Lemma 28. The first $r - 1$ copies (after shuffling) of the protocol invoked by \mathcal{E} will be copies of \underline{Q}' , the r -th copy will be the external system \underline{P} or \underline{Q} , respectively, and the remaining copies will be copies of \underline{P}'' .

Formally, \mathcal{E}_r is obtained from \mathcal{E} as follows (recall that we assume that \mathcal{E} is a single IITM which accepts every message in mode `CheckAddress`). The IITM \mathcal{E}_r will also always accept every message in mode `CheckAddress`. The behavior of \mathcal{E}_r in mode `Compute` is specified next.

First, we need to make sure that \mathcal{E}_r has the appropriate tapes to connect to the different entities. The IITM \mathcal{E} already has tapes to connect to the external tapes of \underline{P} and \underline{Q} . For each such tape c , we add to \mathcal{E}_r a tape c' and c'' to connect to the external tapes of \underline{Q}' and \underline{P}'' , respectively.

Next, we need to specify how \mathcal{E}_r redirects protocol invocations of \mathcal{E} in the way sketched above: \mathcal{E}_r keeps a list L of SIDs, which initially is empty, and the length l of the list, which initially is 0. By definition of $p_{\mathcal{E}}$, it will always hold that $l \leq p_{\mathcal{E}}(\eta + |a|)$. Furthermore, in the first activation with security parameter $\eta \in \mathbf{N}$ and external input $a \in \{0, 1\}^*$, \mathcal{E}_r chooses a permutation π of $\{1, \dots, p_{\mathcal{E}}(\eta + |a|)\}$ uniformly at random. From now on, \mathcal{E}_r simulates \mathcal{E} with security parameter η and external input a . In particular, if \mathcal{E} produces output, then so does \mathcal{E}_r , and if \mathcal{E}_r receives input, then \mathcal{E} is simulated with this input. However, as explained next, the behavior of \mathcal{E}_r deviates from that of \mathcal{E} when it comes to sending and receiving messages to the different copies of protocols.

1. If \mathcal{E} produces output $m = (s, m')$, for some SID s and some message m' , on some external tape c of \underline{P} (and hence, \underline{Q}), then \mathcal{E}_r checks whether s occurs in L . If s does not occur in L , s is first appended at the end of L and l is increased by 1. Let $j \in \{1, \dots, l\}$ be the position where s occurs in L .
 - (a) If $\pi(j) < r$, then \mathcal{E}_r writes m on tape c' .
 - (b) If $\pi(j) = r$, then \mathcal{E}_r outputs m' on c .
 - (c) If $\pi(j) > r$, then \mathcal{E}_r writes m on tape c'' .
2. If \mathcal{E}_r receives input on tape c'' (where c'' is an external tape of \underline{P}'' corresponding to an external tape c of \underline{P}), then \mathcal{E}_r behaves as \mathcal{E} in case input was received on tape c .
3. If \mathcal{E}_r receives input on tape c' (where c' is an external tape of \underline{Q}' corresponding to an external tape c of \underline{Q}), then \mathcal{E}_r behaves as \mathcal{E} in case input was received on tape c .
4. If \mathcal{E}_r receives input m on tape c where c is an external tape of \underline{P} (and hence, \underline{Q}), then \mathcal{E}_r behaves as \mathcal{E} in case input $(L[\pi^{-1}(r)], m)$ was received on tape c where $L[\pi^{-1}(r)]$ denotes the $\pi^{-1}(r)$ -st entry of L . By construction, this entry exists in L (i.e., $\pi^{-1}(r) \leq l$) since \mathcal{E} must have invoked the $\pi^{-1}(r)$ -th copy.

It is easy to see that \mathcal{E}_r is universally bounded for every $r \in \mathbf{N}$ since \mathcal{E} is universally bounded.

We also consider a variant \mathcal{E}_{\S} of \mathcal{E}_r which in the first activation additionally chooses $r \in \{1, \dots, p_{\mathcal{E}}(\eta + |a|)\}$ uniformly at random and then behaves exactly like \mathcal{E}_r .

³¹More formally, for every system $\mathcal{S} \in \text{Con}(\mathcal{E})$ and all parameters η, a in every run of $(\mathcal{E} | \mathcal{S})(1^\eta, a)$ the system \mathcal{E} should only output messages of the described form.

We define the following hybrid systems, for every $r \in \mathbf{N}$:

$$\mathcal{H}_r := \mathcal{E}_r \mid !\underline{\mathcal{P}}'' \mid !\underline{\mathcal{Q}}' ,$$

which can be connected to \mathcal{P} (and hence \mathcal{Q}), i.e., $\mathcal{H}_r \in \text{Con}(\mathcal{P}) = \text{Con}(\mathcal{Q})$.

By $\mathcal{E}_{p_{\mathcal{E}}}$ and $\mathcal{H}_{p_{\mathcal{E}}}$ we denote the systems \mathcal{E}_r and \mathcal{H}_r , respectively, which use $r = p_{\mathcal{E}}(\eta + |a|)$. By construction, for every $r \in \mathbf{N}$, the systems $\mathcal{E} \mid !\underline{\mathcal{P}}$ and $\mathcal{H}_1 \mid \mathcal{P}$, $\mathcal{H}_r \mid \mathcal{Q}$ and $\mathcal{H}_{r+1} \mid \mathcal{P}$, and $\mathcal{E} \mid !\underline{\mathcal{Q}}$ and $\mathcal{H}_{p_{\mathcal{E}}} \mid \mathcal{Q}$, respectively, behave exactly the same (see below). In particular, for all $r \in \mathbf{N}$, we have that:

$$\mathcal{E} \mid !\underline{\mathcal{P}} \equiv_0 \mathcal{H}_1 \mid \mathcal{P} , \quad (26)$$

$$\mathcal{H}_r \mid \mathcal{Q} \equiv_0 \mathcal{H}_{r+1} \mid \mathcal{P} , \text{ and} \quad (27)$$

$$\mathcal{E} \mid !\underline{\mathcal{Q}} \equiv_0 \mathcal{H}_{p_{\mathcal{E}}} \mid \mathcal{Q} . \quad (28)$$

For (26) we use that \mathcal{P} is a protocol system. In particular, we use property (ii) of protocol systems (see Definition 11). If this property were not satisfied, i.e., \mathcal{P} contains an IITM M which is not in the scope of a bang but which could reject a message in mode **CheckAddress**, the following could happen. In a run of $(\mathcal{H}_1 \mid \mathcal{P})(1^\eta, a)$ a message is sent to M , but it is rejected by M (in mode **CheckAddress**). Then, since M is not in the scope of a bang, no new copy of M will be generated. Conversely, if in a run of $\mathcal{E} \mid !\underline{\mathcal{P}}$ a message is sent to a copy of the session version \underline{M} of M prefixed with the first SID generated by \mathcal{E} and the simulated M in \underline{M} would reject the message, then it could happen that a new copy of \underline{M} is generated (since \underline{M} is in the scope of a bang in $\mathcal{E} \mid !\underline{\mathcal{P}}$) which then would not have a corresponding entity in a run of the system $(\mathcal{H}_1 \mid \mathcal{P})(1^\eta, a)$. In short, by property (ii) of protocol systems it is guaranteed that for IITMs that do not occur in the scope of a bang in \mathcal{P} only at most one copy is generated per SID in the run of $\mathcal{E} \mid !\underline{\mathcal{P}}$. A similar argument is used to prove (27) and (28).

Since $\mathcal{E} \mid !\underline{\mathcal{P}}$ is almost bounded, it is easy to see that $\mathcal{H}_1 \mid \mathcal{P}$ is almost bounded too. Moreover, it is easy to see that $\mathcal{E} \mid !\underline{\mathcal{Q}}$ is almost bounded if and only if $\mathcal{H}_{p_{\mathcal{E}}} \mid \mathcal{Q}$ is almost bounded.

Next, we define the event (i.e., a set of runs or equivalently a set of random coins) that the j -th copy of the protocol takes more than $q(\eta + |a|)$ steps. More specifically, for every polynomial q , natural numbers $r, i, j, \eta \in \mathbf{N}$, and $a \in \{0, 1\}^*$ we define:

1. $B_{\mathcal{H}_r \mid \mathcal{P}}^{q,j} = B_{\mathcal{H}_r \mid \mathcal{P}}^{q,j}(1^\eta, a)$ is the following set of runs of $(\mathcal{H}_r \mid \mathcal{P})(1^\eta, a)$. A run ρ belongs to $B_{\mathcal{H}_r \mid \mathcal{P}}^{q,j}$ iff one of the following conditions is satisfied:
 - (a) $\pi(j) < r$ and the copies of machines in $\underline{\mathcal{Q}}'$ with SID $L[j]$ took more than $q(\eta + |a|)$ steps in ρ , where according to Section 3.4 only the steps in mode **Compute** are counted and where in addition the overhead created by the session version is not counted either; in other words, only the computation steps carried out by $\underline{\mathcal{Q}}'$ (and hence, $\underline{\mathcal{Q}}$) are counted.
 - (b) $\pi(j) = r$ and the machines in \mathcal{P} took more than $q(\eta + |a|)$ steps in ρ , with the steps counted as above.
 - (c) $\pi(j) > r$ and the copies of machines in $\underline{\mathcal{P}}''$ with SID $L[j]$ took more than $q(\eta + |a|)$ steps in ρ , with the steps counted as above.
2. $B_{\mathcal{H}_r \mid \mathcal{P}}^q := \bigcup_{i \in \mathbf{N}} B_{\mathcal{H}_r \mid \mathcal{P}}^{q,i}$.
3. $B_{\mathcal{H}_r \mid \mathcal{P}}^{q, \neq j} := \bigcup_{i \in \mathbf{N} \setminus \{j\}} B_{\mathcal{H}_r \mid \mathcal{P}}^{q,i}$.
4. A run ρ belongs to $B_{\mathcal{H}_r \mid \mathcal{P}}^{q, \pi^{-1}(i)}$ iff for the permutation, say π' , chosen in ρ , it hold that $\rho \in B_{\mathcal{H}_r \mid \mathcal{P}}^{q,j}$ with $j = \pi'^{-1}(i)$. The event $B_{\mathcal{H}_r \mid \mathcal{P}}^{q, \neq \pi^{-1}(i)}$ is defined analogously.
5. Analogously to the above events, we define the following events where the external system \mathcal{P} is replaced by \mathcal{Q} : $B_{\mathcal{H}_r \mid \mathcal{Q}}^{q,j}$, $B_{\mathcal{H}_r \mid \mathcal{Q}}^q$, $B_{\mathcal{H}_r \mid \mathcal{Q}}^{q, \neq j}$, $B_{\mathcal{H}_r \mid \mathcal{Q}}^{q, \pi^{-1}(i)}$, $B_{\mathcal{H}_r \mid \mathcal{Q}}^{q, \neq \pi^{-1}(i)}$.

We note that for all r , the systems $\mathcal{H}_r | \mathcal{P}$ (resp., $\mathcal{H}_r | \mathcal{Q}$) is almost bounded if and only if there exists a negligible function f and a polynomial q such that $\text{Prob}[B_{\mathcal{H}_r | \mathcal{P}}^q(1^\eta, a)] \leq f(\eta, a)$ (resp., $\text{Prob}[B_{\mathcal{H}_r | \mathcal{Q}}^q(1^\eta, a)] \leq f(\eta, a)$) for all $\eta \in \mathbf{N}$ and $a \in \{0, 1\}^*$. The direction from left to right is trivial. For the other direction first recall that (the simulated) \mathcal{E} is universally bounded, and hence, only a polynomial number of copies of the protocol is created (the length l of the list L in \mathcal{H}_r is bounded by $p_{\mathcal{E}}(\eta + |a|)$). Now since, by assumption, the number of steps taken by every copy of the protocol is polynomially bounded (except with negligible probability) and the number of steps taken by (the simulated) \mathcal{E} is polynomially bounded, the number of steps taken by $\mathcal{H}_r | \mathcal{P}$ (resp., $\mathcal{H}_r | \mathcal{Q}$) is polynomially bounded. (The overhead for the session versions is only polynomial.)

The next simple lemma shows that there exists a polynomial q_1 such that the event $B_{\mathcal{H}_1 | \mathcal{Q}}^{q_1}$ occurs with negligible probability. As argued above, this is equivalent to $\mathcal{H}_1 | \mathcal{Q}$ being almost bounded. By induction on r , using Theorem 4 and $\mathcal{P} \cong \mathcal{Q}$, one can even easily show that $\mathcal{H}_r | \mathcal{Q}$ is almost bounded for every $r \in \mathbf{N}$. This however does not suffice to prove Lemma 9 because we need to show that $\mathcal{H}_{p_{\mathcal{E}}} | \mathcal{Q}$ is almost bounded. That is, we need a uniform bound that is independent of r (as will be established in Lemma 28). This is in fact the only real difficulty in proving this lemma and the only reason the permutation π is required.

Lemma 26. *There exists a polynomial q_1 and a negligible function f_1 such that for all $\eta \in \mathbf{N}$ and $a \in \{0, 1\}^*$:*

$$\text{Prob}[B_{\mathcal{H}_1 | \mathcal{Q}}^{q_1}(1^\eta, a)] \leq f_1(\eta, a) .$$

Proof. By construction, $\mathcal{H}_1 | \mathcal{P}$ behaves exactly like $\mathcal{E} | \mathcal{P}$, which is almost bounded by assumption. By definition of \mathcal{H}_1 , the subsystem $! \mathcal{Q}'$ of \mathcal{H}_1 is never invoked, and hence, can be removed from \mathcal{H}_1 . Since $! \mathcal{P}$ is environmentally bounded by assumption, it is easy to see that the system $! \mathcal{P}'' | \mathcal{P}$ is environmentally bounded as well. Now, by Theorem 4 (with $\mathcal{S} = ! \mathcal{P}''$) and $\mathcal{P} \cong \mathcal{Q}$, we obtain that $\mathcal{E} | \mathcal{P}'' | \mathcal{Q}$ is almost bounded, and hence, $\mathcal{H}_1 | \mathcal{Q}$ is almost bounded. As argued above, this implies the existence of a polynomial q_1 and a negligible function f_1 with $\text{Prob}[B_{\mathcal{H}_1 | \mathcal{Q}}^{q_1}(1^\eta, a)] \leq f_1(\eta, a)$ for all $\eta \in \mathbf{N}$ and $a \in \{0, 1\}^*$. \square

Throughout the rest of the proof we fix the polynomial q_1 and the negligible function f_1 from the above lemma.

The next lemma shows that for all $r \leq p_{\mathcal{E}}(\eta + |a|)$ the difference of the probabilities that the events $B_{\mathcal{H}_r | \mathcal{P}}^{q_1, \neq \pi^{-1}(r)}$ and $B_{\mathcal{H}_r | \mathcal{Q}}^{q_1, \neq \pi^{-1}(r)}$ occur is negligible. Note that this statement does not consider the number of steps taken by the external protocol copy \mathcal{P} and \mathcal{Q} , respectively, because this copy, which corresponds to $\pi^{-1}(r)$, is excluded in these events.

Lemma 27. *There exists a negligible function f_2 such that for all $\eta \in \mathbf{N}$, $a \in \{0, 1\}^*$, and $r \leq p_{\mathcal{E}}(\eta + |a|)$:*

$$\left| \text{Prob} \left[B_{\mathcal{H}_r | \mathcal{P}}^{q_1, \neq \pi^{-1}(r)}(1^\eta, a) \right] - \text{Prob} \left[B_{\mathcal{H}_r | \mathcal{Q}}^{q_1, \neq \pi^{-1}(r)}(1^\eta, a) \right] \right| \leq f_2(\eta, a) .$$

Proof. We first prove the lemma for non-uniform environments, i.e., environments that use the external input a , because this proof is simpler and better highlights the main idea. In Appendix B we present a proof for this lemma in the uniform setting.

We define an IITM $D \in \text{Env}(\mathcal{P})$ which expects to receive (r, a) as external input and then simulates \mathcal{H}_r with external input a until the runtime bound $q_1(\eta + |a|)$ is hit by any of the internally simulated sessions. If the runtime bound is hit, D outputs 1 on decision. If the run stops and the runtime bound has not been hit, D outputs 0 on decision.

More formally, D is defined as follows. In mode **CheckAddress**, D accepts every message. In mode **Compute**, D behaves as follows: First, D parses the external input on **start** as (r, a) with $r \in \{1, \dots, p_{\mathcal{E}}(\eta + |a|)\}$ and $a \in \{0, 1\}^*$. If the external input is not of this form, D outputs 0 on decision. Otherwise, D simulates the system $\mathcal{H}_r = \mathcal{E}_r | \mathcal{P}'' | \mathcal{Q}'$ with external input a , i.e., it first activates \mathcal{E}_r with input a on **start** and then simulates all the machines in \mathcal{H}_r . If \mathcal{H}_r produces output, then so does D and if D receives input, then D forwards this input to the simulated \mathcal{H}_r . Additionally, D counts the number of transitions taken by all the simulated machines and checks if the conditions of the event $B_{\mathcal{H}_r | \mathcal{P}}^{q_1, \neq \pi^{-1}(r)}(1^\eta, a)$ and $B_{\mathcal{H}_r | \mathcal{Q}}^{q_1, \neq \pi^{-1}(r)}(1^\eta, a)$, respectively, are satisfied. (Note that D can do this even though it cannot inspect the number of transitions

in the external systems, \mathcal{P} and \mathcal{Q} , respectively.) If such a condition is satisfied (i.e., some internal session takes more than $q_1(\eta + |a|)$ steps), then D halts with output 1 on decision. If the run terminates (i.e., \mathcal{E}_r outputs something on decision or \mathcal{E}_r produces empty output, which terminates the run because \mathcal{E}_r is a master IITM) but these conditions are not satisfied, then D halts with output 0 on decision.

It is easy to see that D is universally bounded (i.e., $D \in \text{Env}(\mathcal{P})$) and that for all $\eta \in \mathbf{N}$, $a \in \{0, 1\}^*$, and $r \leq p_{\mathcal{E}}(\eta + |a|)$:

$$\begin{aligned} \text{Prob}[(D | \mathcal{P})(1^\eta, (r, a)) = 1] &= \text{Prob} \left[B_{\mathcal{H}_r | \mathcal{P}}^{q_1, \neq \pi^{-1}(r)}(1^\eta, a) \right] \quad \text{and} \\ \text{Prob}[(D | \mathcal{Q})(1^\eta, (r, a)) = 1] &= \text{Prob} \left[B_{\mathcal{H}_r | \mathcal{Q}}^{q_1, \neq \pi^{-1}(r)}(1^\eta, a) \right] . \end{aligned} \tag{29}$$

Since $\mathcal{P} \cong \mathcal{Q}$ and $D \in \text{Env}(\mathcal{P})$, we have that $D | \mathcal{P} \equiv D | \mathcal{Q}$. In particular, there exists a negligible function f'_2 such that for all $\eta \in \mathbf{N}$, $a \in \{0, 1\}^*$, and $r \leq p_{\mathcal{E}}(\eta + |a|)$:

$$\begin{aligned} f'_2(\eta, (r, a)) &\geq |\text{Prob}[(D | \mathcal{P})(1^\eta, (r, a)) = 1] - \text{Prob}[(D | \mathcal{Q})(1^\eta, (r, a)) = 1]| \\ &\stackrel{(29)}{=} \left| \text{Prob} \left[B_{\mathcal{H}_r | \mathcal{P}}^{q_1, \neq \pi^{-1}(r)}(1^\eta, a) \right] - \text{Prob} \left[B_{\mathcal{H}_r | \mathcal{Q}}^{q_1, \neq \pi^{-1}(r)}(1^\eta, a) \right] \right| . \end{aligned}$$

Let $f_2(\eta, a) := \max_{r \leq p_{\mathcal{E}}(\eta + |a|)} f'_2(\eta, (r, a))$. It is easy to see that f_2 is a negligible function. Now, obviously we have that

$$\left| \text{Prob} \left[B_{\mathcal{H}_r | \mathcal{P}}^{q_1, \neq \pi^{-1}(r)}(1^\eta, a) \right] - \text{Prob} \left[B_{\mathcal{H}_r | \mathcal{Q}}^{q_1, \neq \pi^{-1}(r)}(1^\eta, a) \right] \right| \leq f_2(\eta, a)$$

for all $\eta \in \mathbf{N}$, $a \in \{0, 1\}^*$, and $r \leq p_{\mathcal{E}}(\eta + |a|)$. This concludes the proof. \square

Throughout the rest of the proof, we fix the negligible function f_2 from the above lemma. Finally, we show that there exists a negligible function such that for all $r \leq p_{\mathcal{E}}(\eta + |a|)$ the probability that the event $B_{\mathcal{H}_r | \mathcal{Q}}^{q_1}$ occurs is bounded by this negligible function. This will conclude the proof.

Lemma 28. *There exists a negligible function f_3 such that for all $\eta \in \mathbf{N}$, $a \in \{0, 1\}^*$, and $r \leq p_{\mathcal{E}}(\eta + |a|)$:*

$$\text{Prob} \left[B_{\mathcal{H}_r | \mathcal{Q}}^{q_1}(1^\eta, a) \right] \leq f_3(\eta, a) .$$

Proof. Let $\eta \in \mathbf{N}$, and $a \in \{0, 1\}^*$. For all $r \leq p_{\mathcal{E}}(\eta + |a|)$ we define

$$t_r := t_r(1^\eta, a) := \text{Prob} \left[B_{\mathcal{H}_r | \mathcal{Q}}^{q_1}(1^\eta, a) \right] .$$

We need to show that there exists a negligible function that bounds t_r from above for every $r \leq p_{\mathcal{E}}(\eta + |a|)$.

Let $r \leq p_{\mathcal{E}}(\eta + |a|)$. If $r = 1$, then $t_r = t_1 \leq f_1(\eta, a)$ by Lemma 26. Next, we consider the case $r > 1$. Since the permutation π is chosen uniformly at random, we obtain for all $j \leq r$ the following equality:

$$\text{Prob} \left[B_{\mathcal{H}_r | \mathcal{Q}}^{q_1, \pi^{-1}(r)}(1^\eta, a) \setminus B_{\mathcal{H}_r | \mathcal{Q}}^{q_1, \neq \pi^{-1}(r)}(1^\eta, a) \right] = \text{Prob} \left[B_{\mathcal{H}_r | \mathcal{Q}}^{q_1, \pi^{-1}(j)}(1^\eta, a) \setminus B_{\mathcal{H}_r | \mathcal{Q}}^{q_1, \neq \pi^{-1}(j)}(1^\eta, a) \right] . \tag{30}$$

Intuitively, the event $B_1 := B_{\mathcal{H}_r | \mathcal{Q}}^{q_1, \pi^{-1}(r)}(1^\eta, a) \setminus B_{\mathcal{H}_r | \mathcal{Q}}^{q_1, \neq \pi^{-1}(r)}(1^\eta, a)$ says that the number of steps taken in the external protocol system \mathcal{Q} exceeded q_1 , but not the number of steps in the internal systems. The event $B_2 := B_{\mathcal{H}_r | \mathcal{Q}}^{q_1, \pi^{-1}(j)}(1^\eta, a) \setminus B_{\mathcal{H}_r | \mathcal{Q}}^{q_1, \neq \pi^{-1}(j)}(1^\eta, a)$ says that the number of steps taken in the j th copy of the internally simulated protocol system \mathcal{Q} exceeded q_1 , but not the number of steps taken in the other protocol systems, i.e., in the other internal copies and the external system. Since π is chosen uniformly at random, it does not make a difference whether q_1 is exceeded by the external system or one of the internal copies of \mathcal{Q} .

Formally, the equality (30) can easily be established by defining a bijection between the runs in B_1 and those in B_2 : a run $\rho \in B_1$ in which the permutation π was chosen is mapped to the corresponding run in B_2 where a permutation π' is chosen which coincides with π except that $\pi^{-1}(r)$ and $\pi^{-1}(j)$ are swapped, i.e., $\pi'^{-1}(r) = \pi^{-1}(j)$ and $\pi'^{-1}(j) = \pi^{-1}(r)$.

Now, with (30) we have that:

$$\begin{aligned}
\text{Prob}[B_{\mathcal{H}_r|\mathcal{Q}}^{q_1, \neq \pi^{-1}(r)}(1^\eta, a)] &\geq \text{Prob}\left[\bigcup_{j=1}^{r-1} B_{\mathcal{H}_r|\mathcal{Q}}^{q_1, \pi^{-1}(j)}(1^\eta, a)\right] \\
&\geq \text{Prob}\left[\bigcup_{j=1}^{r-1} B_{\mathcal{H}_r|\mathcal{Q}}^{q_1, \pi^{-1}(j)}(1^\eta, a) \setminus B_{\mathcal{H}_r|\mathcal{Q}}^{q_1, \neq \pi^{-1}(r)}(1^\eta, a)\right] \\
&= \sum_{j=1}^{r-1} \text{Prob}\left[B_{\mathcal{H}_r|\mathcal{Q}}^{q_1, \pi^{-1}(j)}(1^\eta, a) \setminus B_{\mathcal{H}_r|\mathcal{Q}}^{q_1, \neq \pi^{-1}(r)}(1^\eta, a)\right] \\
&\stackrel{(30)}{=} (r-1) \cdot \text{Prob}\left[B_{\mathcal{H}_r|\mathcal{Q}}^{q_1, \pi^{-1}(r)}(1^\eta, a) \setminus B_{\mathcal{H}_r|\mathcal{Q}}^{q_1, \neq \pi^{-1}(r)}(1^\eta, a)\right].
\end{aligned} \tag{31}$$

We conclude that:

$$\begin{aligned}
t_r &= \text{Prob}\left[B_{\mathcal{H}_r|\mathcal{Q}}^{q_1}(1^\eta, a)\right] \\
&= \text{Prob}\left[B_{\mathcal{H}_r|\mathcal{Q}}^{q_1, \neq \pi^{-1}(r)}(1^\eta, a)\right] + \text{Prob}\left[B_{\mathcal{H}_r|\mathcal{Q}}^{q_1, \pi^{-1}(r)}(1^\eta, a) \setminus B_{\mathcal{H}_r|\mathcal{Q}}^{q_1, \neq \pi^{-1}(r)}(1^\eta, a)\right] \\
&\leq \frac{r}{r-1} \cdot \text{Prob}\left[B_{\mathcal{H}_r|\mathcal{Q}}^{q_1, \neq \pi^{-1}(r)}(1^\eta, a)\right] && \text{due to (31)} \\
&\leq \frac{r}{r-1} \cdot \left(\text{Prob}\left[B_{\mathcal{H}_r|\mathcal{P}}^{q_1, \neq \pi^{-1}(r)}(1^\eta, a)\right] + f_2(\eta, a)\right) && \text{due to Lemma 27} \\
&\leq \frac{r}{r-1} \cdot \left(\text{Prob}\left[B_{\mathcal{H}_r|\mathcal{P}}^{q_1}(1^\eta, a)\right] + f_2(\eta, a)\right).
\end{aligned}$$

Since the systems $\mathcal{H}_r|\mathcal{P}$ and $\mathcal{H}_{r-1}|\mathcal{Q}$ behave exactly the same, we have that

$$\text{Prob}[B_{\mathcal{H}_r|\mathcal{P}}^{q_1}(1^\eta, a)] = \text{Prob}[B_{\mathcal{H}_{r-1}|\mathcal{Q}}^{q_1}(1^\eta, a)] = t_{r-1}$$

and we obtain the following simple recurrence relation:

$$t_r \leq \frac{r}{r-1} \cdot (t_{r-1} + f_2(\eta, a)).$$

By induction on r it can be shown that:

$$t_r \leq \left(\prod_{j=1}^{r-1} \frac{j+1}{j}\right) \cdot t_1 + \left(\sum_{j=1}^{r-1} \prod_{i=j}^{r-1} \frac{i+1}{i}\right) \cdot f_2(\eta, a). \tag{32}$$

From this we obtain:

$$\begin{aligned}
t_r &\leq \left(\prod_{j=1}^{r-1} \frac{j+1}{j}\right) \cdot t_1 + \left(\sum_{j=1}^{r-1} \prod_{i=j}^{r-1} \frac{i+1}{i}\right) \cdot f_2(\eta, a) && \text{due to (32)} \\
&= r \cdot t_1 + \left(\sum_{j=1}^{r-1} \frac{r}{j}\right) \cdot f_2(\eta, a) \\
&\leq r \cdot t_1 + (r-1) \cdot r \cdot f_2(\eta, a) \\
&\leq p_{\mathcal{E}}(\eta + |a|) \cdot f_1(\eta, a) + p_{\mathcal{E}}^2(\eta + |a|) \cdot f_2(\eta, a) && \text{due to Lemma 26.}
\end{aligned}$$

Hence, with

$$f_3(\eta, a) := p_{\mathcal{E}}(\eta + |a|) \cdot f_1(\eta, a) + p_{\mathcal{E}}^2(\eta + |a|) \cdot f_2(\eta, a)$$

it holds that

$$\text{Prob}[B_{\mathcal{H}_r|\mathcal{Q}}^{q_1}(1^\eta, a)] = t_r \leq f_3(\eta, a),$$

for all $r \leq p_{\mathcal{E}}(\eta + |a|)$. Note that f_3 is negligible and does not depend on r . \square

Lemma 28 immediately implies that $\text{Prob}[B_{\mathcal{H}_{p_{\mathcal{E}}(\eta+|a|)}}^{q_1} | \mathcal{Q}(1^\eta, a)] \leq f_3(\eta, a)$ for all $\eta \in \mathbf{N}$ and $a \in \{0, 1\}^*$. As argued above, it follows that $\mathcal{H}_{p_{\mathcal{E}}} | \mathcal{Q}$, and hence, $\mathcal{E} | !\underline{\mathcal{Q}}$ are almost bounded. Since this is true for any $\mathcal{E} \in \text{Env}(!\underline{\mathcal{Q}})$, we obtain that $!\underline{\mathcal{Q}}$ is environmentally bounded, which concludes the proof of Lemma 9.

B Proof of Lemma 27 for Uniform Environments

We present the proof of Lemma 27 for uniform environments, i.e., environments which do not get any external input. Let $\eta \in \mathbf{N}$. We define:

$$\delta_r := \delta_r(1^\eta) := \text{Prob} \left[B_{\mathcal{H}_r | \mathcal{P}}^{q_1, \neq \pi^{-1}(r)}(1^\eta) \right] - \text{Prob} \left[B_{\mathcal{H}_r | \mathcal{Q}}^{q_1, \neq \pi^{-1}(r)}(1^\eta) \right]$$

for all $r \leq p_{\mathcal{E}}(\eta)$. (Recall that $p_{\mathcal{E}}$ is the polynomial that bounds the number of different sessions created by \mathcal{E} . Since \mathcal{E} does not get any external input this polynomial is now in the security parameter alone.) Let $r^* \leq p_{\mathcal{E}}(\eta)$ such that it maximizes $|\delta_r|$, i.e., $|\delta_{r^*}| = \max\{|\delta_r| \mid r \leq p_{\mathcal{E}}(\eta)\}$. The basic idea of this proof is due to [17] and goes as follows: D , instead of reading r from the external input (recall that in this setting D does not obtain external input) first samples runs of $(\mathcal{H}_r | \mathcal{P})(1^\eta)$ and $(\mathcal{H}_r | \mathcal{Q})(1^\eta)$ and tries to compute r^* . Then D simulates \mathcal{H}_r as in the non-uniform case (see Appendix A) with $r = r^*$. Of course, D might only find an r such that $|\delta_r|$ approximates $|\delta_{r^*}|$, but this will be good enough.

Let $r^+, r^- \leq p_{\mathcal{E}}(\eta)$ such that it maximizes/minimizes δ_r , i.e., $\delta_{r^+} = \max\{\delta_r \mid r \leq p_{\mathcal{E}}(\eta)\}$ and $\delta_{r^-} = \min\{\delta_r \mid r \leq p_{\mathcal{E}}(\eta)\}$. Note that $|\delta_{r^*}| = \max\{|\delta_{r^+}|, |\delta_{r^-}|\}$. To simplify the proof, we show that both $|\delta_{r^+}|$ and $|\delta_{r^-}|$ are negligible (as functions in η), which implies that $|\delta_{r^*}|$ is negligible. Since $|\delta_{r^*}| \geq |\delta_r|$ for all r , this concludes the proof of the lemma.

First, we show that $|\delta_{r^+}|$ is negligible. Let q be a polynomial. We define D_q^+ , which is parameterized by q , as follows:

1. For every $r \leq p_{\mathcal{E}}(\eta)$ and $i \leq q(\eta)$, D_q^+ simulates a run of $(\mathcal{H}_r | \mathcal{P})(1^\eta)$ and checks if the event $B_{\mathcal{H}_r | \mathcal{P}}^{q_1, \neq \pi^{-1}(r)}(1^\eta)$ occurs. Note that with negligible probability this run might not be polynomial time, in which case D_q^+ aborts the run. Without loss of generality and for simplicity of presentation, we ignore such runs in the following. Let $P_i^r \in \{0, 1\}$ be 1 if the event occurred and 0 otherwise. Similarly, D_q^+ simulates runs of $(\mathcal{H}_r | \mathcal{Q})(1^\eta)$ and obtains the results $Q_i^r \in \{0, 1\}$.
2. For every $r \leq p_{\mathcal{E}}(\eta)$, D_q^+ computes:

$$\tilde{\delta}_r := \frac{1}{q(\eta)} \sum_{i=1}^{q(\eta)} P_i^r - Q_i^r .$$

3. D_q^+ computes \tilde{r}^+ which maximizes $\tilde{\delta}_r$, i.e., $\tilde{\delta}_{\tilde{r}^+} = \max\{\tilde{\delta}_r \mid r \leq p_{\mathcal{E}}(\eta)\}$. (If there exist multiple such \tilde{r}^+ , D_q^+ chooses one arbitrarily.)
4. D_q^+ simulates \mathcal{H}_r as in the non-uniform case (see the proof of Lemma 27 in Appendix A) with $r = \tilde{r}^+$.

Note that for all $r \leq p_{\mathcal{E}}(\eta)$ (considering \tilde{r}^+ as a random variable):

$$\delta_r = \text{Prob} \left[(D_q^+ | \mathcal{P})(1^\eta) = 1 \mid \tilde{r}^+ = r \right] - \text{Prob} \left[(D_q^+ | \mathcal{Q})(1^\eta) = 1 \mid \tilde{r}^+ = r \right] .$$

Furthermore, it is easy to see that D_q^+ is universally bounded (i.e., $D_q^+ \in \text{Env}(\mathcal{P})$). Since, $\mathcal{P} \cong \mathcal{Q}$, there exists a negligible function f_q such that:

$$f_q(\eta) \geq \left| \text{Prob}[(D_q^+ | \mathcal{P})(1^\eta) = 1] - \text{Prob}[(D_q^+ | \mathcal{Q})(1^\eta) = 1] \right| . \quad (33)$$

For every r , the random variables $P_1^r, \dots, P_{q(\eta)}^r$ (resp., $Q_1^r, \dots, Q_{q(\eta)}^r$) are independent and identically distributed, all Bernoulli distributed with success probability $p_{\mathcal{P}}^r := \text{Prob}[B_{\mathcal{H}_r | \mathcal{P}}^{q_1, \neq \pi^{-1}(r)}(1^\eta)]$ (resp., $p_{\mathcal{Q}}^r :=$

$\text{Prob}[B_{\mathcal{H}_r | \mathcal{Q}}^{q_1, \neq \pi^{-1}(r)}(1^\eta)]$. So, the expected value of $\tilde{\delta}_r$ (considered as a random variable where the probability space is over runs of D_q^+) is:

$$\mathbb{E}[\tilde{\delta}_r] = p_{\mathcal{P}}^r - p_{\mathcal{Q}}^r = \delta_r$$

and its variance is:

$$\text{Var}(\tilde{\delta}_r) = \frac{p_{\mathcal{P}}^r(1 - p_{\mathcal{P}}^r) + p_{\mathcal{Q}}^r(1 - p_{\mathcal{Q}}^r)}{q(\eta)} \leq \frac{1}{2q(\eta)} .$$

Using Chebyshev's inequality, we obtain that for every $r \leq p_{\mathcal{E}}(\eta)$ and $\varepsilon > 0$:

$$\text{Prob} [|\tilde{\delta}_r - \delta_r| \geq \varepsilon] \leq \frac{\text{Var}(\tilde{\delta}_r)}{\varepsilon^2} \leq \frac{1}{2q(\eta)\varepsilon^2} .$$

Let p be a polynomial. Then, choosing $\varepsilon := \frac{1}{p(\eta)}$ and $q(\eta) := \frac{p^3(\eta)}{2}$ yields for all r :

$$\text{Prob} \left[|\tilde{\delta}_r - \delta_r| \geq \frac{1}{p(\eta)} \right] \leq \frac{1}{p(\eta)} . \quad (34)$$

Let B_p be the event that in a run of D_q^+ (with some system) there exists an $r \leq p_{\mathcal{E}}(\eta)$ such that $|\tilde{\delta}_r - \delta_r| \geq \frac{1}{p(\eta)}$ (i.e., the sampling is off by too much). By $\overline{B_p}$ we denote the complement of B_p , i.e., that for all $r \leq p_{\mathcal{E}}(\eta)$: $|\tilde{\delta}_r - \delta_r| < \frac{1}{p(\eta)}$. With this, we obtain that:

$$\text{Prob} [B_p] \leq \sum_{r=1}^{p_{\mathcal{E}}(\eta)} \text{Prob} \left[|\tilde{\delta}_r - \delta_r| \geq \frac{1}{p(\eta)} \right] \stackrel{(34)}{\leq} \frac{p_{\mathcal{E}}(\eta)}{p(\eta)} . \quad (35)$$

It is easy to see that if B_p does not hold, then $|\delta_{\tilde{r}^+} - \delta_{r^+}| < \frac{2}{p(\eta)}$ where \tilde{r}^+ is the value computed by D_q^+ in step 3. That is, it holds that

$$\text{Prob} \left[|\delta_{\tilde{r}^+} - \delta_{r^+}| \geq \frac{2}{p(\eta)} \wedge \overline{B_p} \right] = 0 . \quad (36)$$

where, as usual, \tilde{r}^+ is a random variable and the probability space is taken over runs of D_q^+ .

By (36), we conclude that $\text{Prob} [\tilde{r}^+ = r \mid \overline{B_p}] = 0$ for all $r \leq p_{\mathcal{E}}(\eta)$ such that $|\delta_r - \delta_{r^+}| \geq \frac{2}{p(\eta)}$. Hence, for all $r \leq p_{\mathcal{E}}(\eta)$:

$$\text{Prob} [\tilde{r}^+ = r \mid \overline{B_p}] \cdot |\delta_r - \delta_{r^+}| < \frac{2}{p(\eta)} . \quad (37)$$

For all $r \leq p_{\mathcal{E}}(\eta)$, we define:

$$\delta'_r := \text{Prob} [(D_q^+ \mid \mathcal{P})(1^\eta) = 1 \mid \overline{B_p} \wedge \tilde{r}^+ = r] - \text{Prob} [(D_q^+ \mid \mathcal{Q})(1^\eta) = 1 \mid \overline{B_p} \wedge \tilde{r}^+ = r] . \quad (38)$$

That is, δ'_r is defined like δ_r but under the condition that B_p does not occur. It is easy to show that for all $r \leq p_{\mathcal{E}}(\eta)$:³²

$$|\delta'_r - \delta_r| \leq 2 \cdot \text{Prob} [B_p] . \quad (39)$$

³²We use the obvious fact that for $\text{Prob}[C] > 0$ it holds that $|(\text{Prob}[A] - \text{Prob}[B]) - (\text{Prob}[A \mid C] - \text{Prob}[B \mid C])| \leq 2 \cdot \text{Prob}[\overline{C}]$.

We conclude that:

$$\begin{aligned}
f_q(\eta) &\stackrel{(33)}{\geq} \left| \text{Prob}[(D_q^+ | \mathcal{P})(1^\eta) = 1] - \text{Prob}[(D_q^+ | \mathcal{Q})(1^\eta) = 1] \right| \\
&\geq \left| \text{Prob}[(D_q^+ | \mathcal{P})(1^\eta) = 1 | \overline{B_p}] - \text{Prob}[(D_q^+ | \mathcal{Q})(1^\eta) = 1 | \overline{B_p}] \right| - 2 \cdot \text{Prob}[B_p] \\
&\stackrel{(38)}{=} \left| \sum_{r=1}^{p_{\mathcal{E}}(\eta)} \text{Prob}[\tilde{r}^+ = r | \overline{B_p}] \cdot \delta'_r \right| - 2 \cdot \text{Prob}[B_p] \\
&\stackrel{(39)}{\geq} \left| \sum_{r=1}^{p_{\mathcal{E}}(\eta)} \text{Prob}[\tilde{r}^+ = r | \overline{B_p}] \cdot \delta_r \right| - 2p_{\mathcal{E}}(\eta) \cdot \text{Prob}[B_p] - 2 \cdot \text{Prob}[B_p] \\
&= \left| \delta_{r^+} + \sum_{r=1}^{p_{\mathcal{E}}(\eta)} \text{Prob}[\tilde{r}^+ = r | \overline{B_p}] \cdot (\delta_r - \delta_{r^+}) \right| - 2p_{\mathcal{E}}(\eta) \cdot \text{Prob}[B_p] - 2 \cdot \text{Prob}[B_p] \\
&\geq |\delta_{r^+}| - \left(\sum_{r=1}^{p_{\mathcal{E}}(\eta)} \text{Prob}[\tilde{r}^+ = r | \overline{B_p}] \cdot |\delta_r - \delta_{r^+}| \right) - 2p_{\mathcal{E}}(\eta) \cdot \text{Prob}[B_p] - 2 \cdot \text{Prob}[B_p] \\
&\stackrel{(37)}{>} |\delta_{r^+}| - p_{\mathcal{E}}(\eta) \cdot \frac{2}{p(\eta)} - 2p_{\mathcal{E}}(\eta) \cdot \text{Prob}[B_p] - 2 \cdot \text{Prob}[B_p] \\
&\stackrel{(35)}{\geq} |\delta_{r^+}| - \frac{2p_{\mathcal{E}}^2(\eta) + 4p_{\mathcal{E}}(\eta)}{p(\eta)}.
\end{aligned}$$

Since this holds for every polynomial p and f_q is negligible (recall that $q = \frac{p^3}{2}$), we conclude that $|\delta_{r^+}|$ is negligible.

To prove that $|\delta_{r^-}|$ is negligible, in the argument above, we simply have to replace D_q^+ by D_q^- , r^+ by r^- , and \tilde{r}^+ by \tilde{r}^- . D_q^- only differs from D_q^+ in step 3: D_q^- computes \tilde{r}^- such that it minimizes $\tilde{\delta}_r$. \square

C Proof of Theorem 7

In Section 6.3, we already proved that SS and UC are equivalent. To complete the proof of Theorem 7, we first show that UC implies dummyUC implies SS implies BB implies SS (Appendix C.1, C.2, C.3., and C.4, respectively), which altogether proves equivalence of SS, UC, dummyUC, and BB. As mentioned in Section 6, these equivalences also hold true in the case where the environment does not get external input, i.e., is uniform. In Appendix C.5, we then prove that SS is equivalent to RS, which concludes the proof of Theorem 7. As mentioned in Section 6, for this equivalence, we (have to) use that the environment gets external input, i.e., is non-uniform.

In what follows, let \mathcal{P} and \mathcal{F} be protocol systems. Furthermore, let \mathcal{P} be environmentally bounded. We assume that the set of network tapes of \mathcal{P} and \mathcal{F} are disjoint; otherwise the tapes can be renamed to satisfy this assumption.

C.1 UC \Rightarrow dummyUC

It is easy to see that $\mathcal{P} \leq^{UC} \mathcal{F}$ implies $\mathcal{P} \leq^{dumUC} \mathcal{F}$: It holds that $\mathcal{D}_{\mathcal{P}} = \mathcal{D}(\mathcal{T}_{in}^{net}(\mathcal{P}), \mathcal{T}_{out}^{net}(\mathcal{P})) \in \text{Adv}(\mathcal{P})$ (i.e., the dummy adversary is a valid adversary for \mathcal{P}). Hence, by definition of $\mathcal{P} \leq^{UC} \mathcal{F}$, there exists $\mathcal{I} \in \text{Sim}^{\mathcal{D}_{\mathcal{P}} | \mathcal{P}}(\mathcal{F})$ such that $\mathcal{D}_{\mathcal{P}} | \mathcal{P} \cong \mathcal{I} | \mathcal{F}$. We conclude that $\mathcal{P} \leq^{dumUC} \mathcal{F}$. \square

C.2 dummyUC \Rightarrow SS

The proof that dummyUC implies SS is exactly as the proof that UC implies SS, which is given in Section 6.3.

C.3 SS \Rightarrow BB

It follows directly from Theorem 4 that $\mathcal{P} \leq^{SS} \mathcal{F}$ implies $\mathcal{P} \leq^{BB} \mathcal{F}$: By definition of $\mathcal{P} \leq^{SS} \mathcal{F}$, there exists $\mathcal{S} \in \text{Sim}^{\mathcal{P}}(\mathcal{F})$ such that $\mathcal{P} \cong \mathcal{S} | \mathcal{F}$. Let $\mathcal{A} \in \text{Adv}(\mathcal{P})$. By Lemma 6, we may assume that both \mathcal{S} and \mathcal{A} are single IITMs that accept every message in mode `CheckAddress`.³³ Since $\mathcal{A} | \mathcal{P}$, \mathcal{P} , and $\mathcal{S} | \mathcal{F}$ are environmentally bounded and $\mathcal{P} \cong \mathcal{S} | \mathcal{F}$, by Theorem 4 (note that \mathcal{A} and $\mathcal{S} | \mathcal{F}$ are protocol systems because \mathcal{A} and \mathcal{S} accept every message in mode `CheckAddress`), we have that $\mathcal{A} | \mathcal{P} \cong \mathcal{A} | \mathcal{S} | \mathcal{F}$. We conclude that $\mathcal{P} \leq^{BB} \mathcal{F}$. \square

C.4 BB \Rightarrow SS

We show that $\mathcal{P} \leq^{BB} \mathcal{F}$ implies $\mathcal{P} \leq^{SS} \mathcal{F}$: By definition of $\mathcal{P} \leq^{BB} \mathcal{F}$ there exists $\mathcal{S} \in \text{Sim}^{\mathcal{P}}(\mathcal{F})$ such that $\mathcal{A} | \mathcal{P} \cong \mathcal{A} | \mathcal{S} | \mathcal{F}$ for all $\mathcal{A} \in \text{Adv}(\mathcal{P})$. Since the dummy adversary is a valid adversary for \mathcal{P} , i.e., $\mathcal{D}_{\mathcal{P}} \in \text{Adv}(\mathcal{P})$, we immediately obtain that $\mathcal{D}_{\mathcal{P}} | \mathcal{P} \cong \mathcal{D}_{\mathcal{P}} | \mathcal{S} | \mathcal{F}$. By Lemma 5, we conclude $\mathcal{P} \cong \mathcal{S} | \mathcal{F}$, and hence, $\mathcal{P} \leq^{SS} \mathcal{F}$. \square

C.5 SS \Leftrightarrow RS

It is trivially seen that $\mathcal{P} \leq^{UC} \mathcal{F}$ implies $\mathcal{P} \leq^{RS} \mathcal{F}$. Since, as shown before, UC and SS are equivalent, it follows that $\mathcal{P} \leq^{SS} \mathcal{F}$ implies $\mathcal{P} \leq^{RS} \mathcal{F}$. We now prove that $\mathcal{P} \leq^{RS} \mathcal{F}$ implies $\mathcal{P} \leq^{SS} \mathcal{F}$. The main argument is similar to the one presented in [3]. We therefore only present the proof sketch.

By definition of $\mathcal{P} \leq^{RS} \mathcal{F}$, for all $\mathcal{A} \in \text{Adv}(\mathcal{P})$ and all $\mathcal{E} \in \text{Env}(\mathcal{A} | \mathcal{P})$ there exists $\mathcal{I} \in \text{Sim}^{\mathcal{A} | \mathcal{P}}(\mathcal{F})$ such that:

$$\mathcal{E} | \mathcal{A} | \mathcal{P} \equiv \mathcal{E} | \mathcal{I} | \mathcal{F} . \quad (40)$$

We choose $\mathcal{A} = \mathcal{D}_{\mathcal{P}} = \mathcal{D}(\mathcal{T}_{in}^{net}(\mathcal{P}), \mathcal{T}_{out}^{net}(\mathcal{P}))$ to be the dummy adversary for \mathcal{P} . We also choose $\mathcal{E} \in \text{Env}(\mathcal{D}_{\mathcal{P}} | \mathcal{P})$ to be a “universal” Turing machine (more precisely, a universal IITM) which takes as external input (i.e., input on `start`) a tuple of the form $(a, e, 1^t)$ where e is an encoding of some IITM (representing an environmental system \mathcal{E}'), a is interpreted as an external input to \mathcal{E}' , and t is interpreted as a runtime. (By Lemma 7, we may assume that e encodes a single IITM which accepts every message in mode `CheckAddress`.) The universal IITM \mathcal{E} simulates \mathcal{E}' with external input a up to t steps. Clearly, \mathcal{E} is universally bounded because its runtime is polynomial in the security parameter plus the length of the external input. Now, given a security parameter η , external input a , and an environmental system \mathcal{E}' , there exists a tuple $(a, e, 1^t)$ of length polynomial in $\eta + |a|$ (it suffices to choose t polynomial in $\eta + |a|$ because \mathcal{E}' is universally bounded) such that \mathcal{E} with external input $(a, e, 1^t)$ precisely simulates \mathcal{E}' . Hence, (40) implies $\mathcal{E}' | \mathcal{D}_{\mathcal{P}} | \mathcal{P} \equiv \mathcal{E}' | \mathcal{I} | \mathcal{F}$, i.e., $\mathcal{P} \leq^{dumUC} \mathcal{F}$. Since we already know that $\mathcal{P} \leq^{dumUC} \mathcal{F}$ implies $\mathcal{P} \leq^{SS} \mathcal{F}$, we obtain $\mathcal{P} \leq^{SS} \mathcal{F}$. \square

References

- [1] B. Barak, Y. Lindell, and T. Rabin. Protocol Initialization for the Framework of Universal Composability. Technical Report 2004/006, Cryptology ePrint Archive, 2004. Available at <http://eprint.iacr.org/2004/006>.
- [2] R. Canetti. Universally Composable Security: A New Paradigm for Cryptographic Protocols. In *Proceedings of the 42nd Annual Symposium on Foundations of Computer Science (FOCS 2001)*, pages 136–145. IEEE Computer Society, 2001.
- [3] R. Canetti. Universally Composable Security: A New Paradigm for Cryptographic Protocols. Technical Report 2000/067, Cryptology ePrint Archive, December 2005. Available at <http://eprint.iacr.org/2000/067>.

³³We note that this assumption is actually not required. We could have used the mentioned stronger variant of Theorem 4 instead.

- [4] R. Canetti. Security and Composition of Cryptographic Protocols: A Tutorial. Technical Report 2006/465, Cryptology ePrint Archive, 2006. Available at <http://eprint.iacr.org/2006/465>.
- [5] R. Canetti, L. Cheung, D. Kaynar, M. Liskov, N. Lynch, O. Pereira, and R. Segala. Time-Bounded Task-PIOAs: A Framework for Analyzing Security Protocols. In S. Dolev, editor, *20th International Symposium on Distributed Computing (DISC 2006)*, pages 238–253. Springer, 2006.
- [6] R. Canetti, Y. Dodis, R. Pass, and S. Walfish. Universally Composable Security with Global Setup. In S. P. Vadhan, editor, *Theory of Cryptography, Proceedings of TCC 2007*, volume 4392 of *Lecture Notes in Computer Science*, pages 61–85. Springer, 2007.
- [7] R. Canetti, S. Halevi, J. Katz, Y. Lindell, and P. D. MacKenzie. Universally Composable Password-Based Key Exchange. In R. Cramer, editor, *Advances in Cryptology - EUROCRYPT 2005, 24th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, volume 3494 of *Lecture Notes in Computer Science*, pages 404–421. Springer, 2005.
- [8] R. Canetti and J. Herzog. Universally Composable Symbolic Analysis of Mutual Authentication and Key-Exchange Protocols. In S. Halevi and T. Rabin, editors, *Theory of Cryptography, Third Theory of Cryptography Conference, TCC 2006*, volume 3876 of *Lecture Notes in Computer Science*, pages 380–403. Springer, 2006.
- [9] R. Canetti and H. Krawczyk. Universally Composable Notions of Key Exchange and Secure Channels. In *Advances in Cryptology - EUROCRYPT 2002, International Conference on the Theory and Applications of Cryptographic Techniques, Proceedings*, volume 2332 of *Lecture Notes in Computer Science*, pages 337–351. Springer, 2002.
- [10] R. Canetti and T. Rabin. Universal Composition with Joint State. In *Advances in Cryptology, 23rd Annual International Cryptology Conference (CRYPTO 2003), Proceedings*, volume 2729 of *Lecture Notes in Computer Science*, pages 265–281. Springer, 2003.
- [11] A. Datta, R. Küsters, J. Mitchell, and A. Ramanathan. On the Relationships Between Notions of Simulation-Based Security. In J. Kilian, editor, *Proceedings of the 2nd Theory of Cryptography Conference (TCC 2005)*, volume 3378 of *Lecture Notes in Computer Science*, pages 476–494. Springer, 2005.
- [12] C. Hoare. *Communicating Sequential Processes*. Prentice-Hall, 1985.
- [13] D. Hofheinz, J. Müller-Quade, and D. Unruh. Polynomial Runtime in Simulatability Definitions. In *18th IEEE Computer Security Foundations Workshop (CSFW-18 2005)*, pages 156–169. IEEE Computer Society, 2005.
- [14] D. Hofheinz and V. Shoup. GNUC: A New Universal Composability Framework. Technical Report 2011/303, Cryptology ePrint Archive, 2011. Available at <http://eprint.iacr.org/2011/303>.
- [15] D. Hofheinz and D. Unruh. Comparing two Notions of Simulatability. In J. Kilian, editor, *Theory of Cryptography, Proceedings of TCC 2005*, volume 3378 of *Lecture Notes in Computer Science*, pages 86–103. Springer, 2005.
- [16] D. Hofheinz, D. Unruh, and J. Müller-Quade. Polynomial Runtime and Composability. Technical Report 2009/023, Cryptology ePrint Archive, 2009. Available at <http://eprint.iacr.org/2009/023>.
- [17] D. Hofheinz, D. Unruh, and J. Müller-Quade, 2011. [16, footnote 26 on page 40] and personal communication.
- [18] R. Küsters. Simulation-Based Security with Inexhaustible Interactive Turing Machines. In *Proceedings of the 19th IEEE Computer Security Foundations Workshop (CSFW-19 2006)*, pages 309–320. IEEE Computer Society, 2006.

- [19] R. Küsters, A. Datta, J. C. Mitchell, and A. Ramanathan. On the Relationships Between Notions of Simulation-Based Security. *Journal of Cryptology*, 21(4):492–546, 2008.
- [20] R. Küsters and M. Tuengerthal. Joint State Theorems for Public-Key Encryption and Digital Signature Functionalities with Local Computation. In *Proceedings of the 21st IEEE Computer Security Foundations Symposium (CSF 2008)*, pages 270–284. IEEE Computer Society, 2008.
- [21] R. Küsters and M. Tuengerthal. Universally Composable Symmetric Encryption. In *Proceedings of the 22nd IEEE Computer Security Foundations Symposium (CSF 2009)*, pages 293–307. IEEE Computer Society, 2009.
- [22] R. Küsters and M. Tuengerthal. Composition Theorems Without Pre-Established Session Identifiers. In Y. Chen, G. Danezis, and V. Shmatikov, editors, *Proceedings of the 18th ACM Conference on Computer and Communications Security (CCS 2011)*, pages 41–50. ACM, 2011.
- [23] R. Küsters and M. Tuengerthal. Ideal Key Derivation and Encryption in Simulation-based Security. In A. Kiayias, editor, *Topics in Cryptology – CT-RSA 2011, The Cryptographers’ Track at the RSA Conference 2011, Proceedings*, volume 6558 of *Lecture Notes in Computer Science*, pages 161–179. Springer, 2011.
- [24] U. Maurer and R. Renner. Abstract Cryptography. In B. Chazelle, editor, *Innovations in Computer Science - ICS 2010. Proceedings*, pages 1–21. Tsinghua University Press, 2011.
- [25] R. Milner. *A Calculus of Communicating Systems*. Springer, 1980.
- [26] R. Needham and M. Schroeder. Using Encryption for Authentication in Large Networks of Computers. *Communications of the ACM*, 21(12):993–999, 1978.
- [27] B. Pfitzmann and M. Waidner. A Model for Asynchronous Reactive Systems and its Application to Secure Message Transmission. In *IEEE Symposium on Security and Privacy*, pages 184–201. IEEE Computer Society, 2001.
- [28] R. L. Rivest, A. Shamir, and D. A. Wagner. Time-Lock Puzzles and Timed-Release Crypto. Technical Report MIT/LCS/TR-684, MIT, March 1996. Available at <http://theory.lcs.mit.edu/~rivest/RivestShamirWagner-timelock.ps>.