

Polynomial Runtime in Simulatability Definitions

Dennis Hofheinz, Jörn Müller-Quade, and Dominique Unruh
IAKS, Arbeitsgruppe Systemsicherheit
Fakultät für Informatik, Universität Karlsruhe, Am Fasanengarten 5
76131 Karlsruhe, Germany
{hofheinz,muellerq,unruh}@ira.uka.de

Abstract

We elaborate on the problem of polynomial runtime in simulatability definitions for multi-party computation. First, the need for a new definition is demonstrated by showing which problems occur with common definitions of polynomial runtime. Then, we give a definition which captures in an intuitive manner what it means for a protocol or an adversary to have polynomial runtime.

We show that this notion is suitable for simulatability definitions for multi-party computation. In particular, a composition theorem is shown for this notion.

Keywords: multi-party computation, reactive simulatability, universal composability.

1. Introduction

Recently, simulatability-based notions of security for multi-party protocols received a lot of attention. In particular, in the “reactive simulatability” modelling of Backes, Pfitzmann and Waidner (cf. [16, 6]) and the “universal composability” framework of Canetti (cf. [7]), both structural and constructive results could be formulated.

A simulatability-based notion of security considers a protocol π secure only relative to an idealisation τ of the respective protocol task. More concretely, π is considered as secure as τ , iff every attack on π can be simulated by a suitable attack on τ . Intuitively, this means that τ already exhibits every weakness of π . For the purpose of comparing π - and τ -attacks, a protocol user H is introduced,¹ which gives inputs to the parties, reads their outputs, and may talk to the respective adversary.

For modelling computational security, the mentioned models bound the computational complexity of all ma-

chines that participate in a protocol run (i.e., parties, adversary, and user) to strict polynomial-time in the security parameter k .² That is, every machine M halts after running $p_M(k)$ steps for a polynomial p_M which depends only on the machine M .

Next, we will describe shortcomings of such a definition in two common simulatability frameworks to motivate that this notion of polynomial-time is not intuitive. In Section 3, we attempt to give a more natural definition, for which we prove composition features in Section 4. In Section 5, we relate our definition to the existing one just described.

1.1. The UC framework

In the framework [7] of universal composability (UC), such a strict polynomial complexity bound on each machine can cause difficulties. First, it gets hard to formulate a cryptographic task like public key encryption (see, e.g., [7, 14, 11]) without fixing explicit runtime bounds (which might seem unnatural).

Moreover, the protocol environment may “kill” a polynomially bounded ideal functionality by activating it sufficiently often with nonsense inputs. A real implementation must now recognise that one party got “too many” inputs and stop service; this again may, depending on the protocol and the network model, not be possible.³

Furthermore, in the simulatability definition, the protocol environment \mathcal{Z} may depend on the ideal-model adversary \mathcal{S} .⁴ In particular, \mathcal{Z} may first of all activate \mathcal{S} with nonsense inputs until \mathcal{S} must have halted. In such a situation, τ

¹ In the framework of [7], this entity is called the (protocol) environment \mathcal{Z} .

² This excludes a very recent update on the UC framework, to be found at [10]; see below.

³ It is not helpful to explicitly bound the “dummy parties,” which relay in- and outputs to and from the ideal functionality: In a larger protocol in which the ideal functionality may be used, these dummy parties are omitted. Consequently, they cannot protect the ideal functionality from being “overwhelmed” by inputs. This would contradict secure composition of protocols.

⁴ \mathcal{S} attacks τ and thereby simulates an attack on π .

must still “look like π ” even if the ideal-model adversary \mathcal{S} has halted.

The preceding argument shows that in the original formulation from [7], the ideal functionalities \mathcal{F}_{PKE} and \mathcal{F}_{SIG} are unrealizable by any real protocol (i.e., by any protocol which works without other “helping” ideal functionalities).

In the special case of the functionality \mathcal{F}_{PKE} for public-key encryption, a solution was proposed in [14]: keep all machines polynomial *per activation*, and quantify only over environments \mathcal{Z} that guarantee a polynomial total running time of the complete protocol run (with both π and τ). However, for other functionalities which may play “ping-pong”⁵ with the ideal-model adversary, this notion would disallow any environment.

A similar approach was later used in [8]: here, all machines are polynomial per activation in the maximum of the security parameter and the input length; however, environment and adversary are strictly polynomially bounded. In this situation, an environment is no longer able to flood (and thereby disable) the dummy parties with wrong inputs; yet, \mathcal{Z} may still “kill” the ideal-model adversary.

On January 27, after completion of this manuscript, the paper [10] was updated to contain an alternative approach to solve the problem of polynomial runtime in an intuitive manner. This new model seems to be more restrictive than the model presented here, as a distinction of the real and the ideal model is possible based on the number of activations a machine allows before halting [9]. The model here abstracts from such “denial-of-service-attacks”. Therefore Canetti’s and our model consider different classes of secure protocols.

1.2. The Model of Reactive Simulatability

The original formulation of [16] is very similar to [7] with respect to the computational complexity of protocol machines. Concretely, [16] demand for computational security, that all machines are strictly polynomially bounded. For the notion of “standard simulatability” ([16]’s default notion of simulatability, in which an ideal-model adversary may depend on the protocol user⁶), this is not as difficult as for “universal simulatability.” The latter notion allows the protocol user H to depend on the ideal-model adversary⁷ A_2 , and thus to “kill” A_2 by sending lots of nonsensical input. Furthermore, the issue that an ideal functionality may be “killed” as described in Section 1.1 is also present.

This problem was addressed in [1, 2, 6] by allowing every machine to “block” selected connections. (To do so, a

⁵ I.e., both adversary and functionality immediately respond to any message from each other, thereby creating an infinite loop.

⁶ The protocol user H is the equivalent of the protocol environment \mathcal{Z} from the universal composability framework of [7].

⁷ The notation in [16] differs slightly from that in [7].

machine could set its so-called “length function” for that connection to zero.) So for example, the ideal-model adversary A_2 may—from a certain point in time on—block all connections from the user H , when the corresponding real-model adversary would have halted or blocked this connection. Thus, H is not able to “kill” A_2 anymore. Similarly, an ideal functionality is now able to “block” selected ports.

However, there are still reasons why one might consider this solution not satisfying. First, artificial polynomial bounds have to be stated for an idealisation of, say, public-key encryption. That is, concrete polynomial limitations on the message lengths and the number of encryptions have to be fixed to keep ideal and real protocol strictly polynomially bounded. So to achieve full generality, protocols have to be parameterised over, e.g., message lengths and the number of encryptions (as an example, cf. [5]).

Furthermore, notions like “polynomial fairness” of an adversary (which means that this adversary schedules messages between parties after a polynomial number of activations, cf. [4]) are not compatible with an a priori polynomially bounded adversary. This is so since the adversary is not able to schedule messages after it has halted, and thus no scheduling guarantees can be given.

Finally, the technical tool of length functions (which is the tool used to “block” a connection) might be considered artificial. There might be situations in which it is unrealistic to assume that a machine may block selected communication channels, but is still able to “listen” on other channels—consider a dial-in Internet connection, for example.

1.3. Other related work

In [12] the above problems with polynomial runtime have also been noticed. Their solution consists of introducing so-called guards, a generalisation of length functions. These guards may reject or modify incoming messages without wasting any of the total runtime of the concerned machine. This solves the problem of “killing” a machine by sending nonsensical inputs (these may be removed by the guard), but still requires that the amount of actual work a machine does is a priori bounded (e.g., a secure message transmission functionality would have an a priori limit of the number and length of messages transmitted).

1.4. Our Contributions

Motivated by the discussion above we give a new definition of polynomial runtime for simulatability and prove several desirable properties of our definition. The definition is stated in the model of reactive simulatability, but the concept is model-independent and should carry over to the UC framework.

The protocol user H will be chosen to be *weakly polynomial* (cf. [6]), i.e., it will in each activation be polynomially limited in the security parameter and the overall size of the input it gets on incoming ports. The adversary will be limited in the runtime of H . To guarantee this, two specific connections between the adversary and the user will be used to limit the adversary in the message volume communicated over these lines. A protocol user together with an adversary limited in this sense will be called *continuously polynomial*.

We stress that this definition allows users and adversaries that do not terminate at all. Specifically, they may run long enough to break every complexity-based cryptographic system. However, the definition guarantees that they may not do so in polynomial prefixes of H 's view. In fact, the definition guarantees that in polynomial prefixes of H 's view, both A and H take only a polynomial number of steps, *and* both of them send only messages of at most polynomial size to the protocol. This captures a very intuitive notion of polynomial runtime for protocol users and adversaries. This security notion is presented in Section 3.

Polynomial limitations of a protocol will be captured by the notion of *polynomially shaped* collections. Roughly, a set of machines is polynomially shaped if the total length of all messages sent by these machines is polynomial in the security parameter k plus the overall length of inputs which machines from this set got from machines outside this set. If additionally, all machines in the set are weakly polynomial we call this set *polynomially shaped weakly polynomial* (*ps-wp* for short). The notion of *ps-wp* is a natural definition of a protocol being “polynomially bounded in input length and security parameter” without having to give explicit a priori bounds for the lifetime of machines.

In Section 4, we prove a generalised composition theorem for *ps-wp* protocols. Specifically, in any *ps-wp* collection of machines, a functionality may be replaced by a secure implementation if the resulting collection of machines remains *ps-wp*.

We note that the set of *ps-wp* protocols is *not* closed under composition (i.e., there are *ps-wp* protocols which yield a non-*ps-wp* protocol if composed). We argue that this is not a flaw of our notion, but a “necessary evil” if one wants to catch the intuitive notion of a polynomially bounded protocol. Therefore, we construct an example of two protocols which are “intuitively polynomial” (and *ps-wp*), but which compose to a protocol that is non-polynomial in every intuitive way.

Additionally, we give a sub-notion of *ps-wp* protocols that is closed under composition. As a simple consequence, the mentioned *ps-wp* composition theorem shows that this notion allows for a *secure* composition of protocols (without any additional conditions on the complexity of the composed protocols).

In Section 5, we relate our new notion of security to the existing notion of polynomial security from [6]. More specifically, we prove that our notion is at least as strict as the one from [6].

In Section 7, we sketch how to apply our ideas to the UC framework.

Finally, in Appendix B we show that the generalisation of simulatable security to machines which are intuitively polynomial as defined in this work, but not strictly polynomial, will allow us to omit the formal concept of length functions, which was introduced in [1] to solve problems arising with strictly polynomial functionalities. More specifically, we show that removing length functions from protocol machines does not change the notion of security.

2. Review of Reactive Simulatability

In this section, we present the notion of reactive simulatability. This introduction only very roughly sketches the definitions, and the reader is encouraged to read [6] for more detailed information and formal definitions. A reader familiar with the model may skip this section and proceed to Section 3. Additionally, a glossary of important terms in the reactive simulatability framework can be found in Appendix A.

Reactive Simulatability is a definition of security which defines a protocol \hat{M}_1 (the *real protocol*) to be *as secure as* another protocol \hat{M}_2 (the *ideal protocol*, the *trusted host*), if for any adversary A_1 (also called the *real adversary*), and any *honest user* H , there is a *simulator* A_2 (also called the *ideal adversary*), s.t. the view of H is indistinguishable in the following two scenarios:

- The honest user H runs together with the real adversary A_1 and the real protocol \hat{M}_1
- The honest user H runs together with the simulator A_2 and the ideal protocol \hat{M}_2 .

Note that there is a security parameter k common to all machines, so that the notion of indistinguishability makes sense.

This definition allows to specify some trusted host—which is defined to be a secure implementation of some cryptographic task—as the ideal protocol, and then to consider the question, whether a real protocol is as secure as the trusted host (and thus also a secure implementation of that task). In order to understand the above definitions in more detail, we have to specify what is meant by machines “running together”. Consider a set of machines (called a *collection*). Each machine has so-called *simple in-ports* (written $p^?$), *simple out-ports* (written $p!$), and *clock out-ports* (written p^{cl}). Ports with the same name (p in our example) are considered to belong together and are associated




Figure 1. A connection

with a *buffer* \tilde{p} . These are then interconnected as in Figure 1 (note that some or all ports may originate from the same machine). Now, when a collection runs, the following happens: At every time, exactly one machine is activated. It may now read its simple in-ports (representing incoming network connections), do some work, and then write output to its simple out-ports. After such an activation the contents of the simple out-ports $p!$ are appended to the queue of messages stored in the associated buffer \tilde{p} . However, since now all messages are stored in buffers and will not be delivered by themselves, machines additionally have after each activation the possibility to write a number $n \geq 1$ to at most one clock out-port $p^<!$. Then the n -th undelivered message of buffer \tilde{p} will be written to the simple in-port $p?$ and deleted from the buffer's queue. The machine that has the simple in-port $p?$ will be activated next. So the clock out-ports control the scheduling. Usually, a connection is clocked by (i.e., the corresponding clock out-port is part of) the sender, or by the adversary. Since the most important use of a clock out-port is to write a 1 onto it (deliver the oldest message in the buffer), we say a machine clocks a connection or a message when a machine writes a 1 onto the clock port of that connection.

At the start of a run, or when no machine is activated at some point, a designated machine called the *master scheduler* is activated. For this, the master scheduler has a special port, called the *master clock port* $\text{clk}^<?$.

Note that not all collections can be executed, only so-called *closed* collections, where all connections have their simple in-, simple out-, and clock out-port. If a collection is not closed, we call the ports having no counterpart *free ports*.

In order to understand how this idea of networks relates to the above sketch of reactive simulatability, one has to get an idea of what is meant by a protocol. A protocol is represented by a so-called *structure* (\hat{M}, S) , consisting of a collection \hat{M} of the protocol participants (parties, trusted hosts, etc.), and a subset of the free ports of \hat{M} , the so-called *service ports* S . The service ports represent the protocol's interface (the connections to the protocol's users). The honest user can then only connect to the service ports (and to the adversary), all other free ports of the protocol are intended for the communication with the adversary

(they may e.g. represent side channels, possibilities of attack, etc.). Since usually a protocol does not explicitly communicate with an adversary, such free non-service ports are more commonly found with trusted hosts, explicitly modelling their imperfections.

With this information we can review the above “definition” of security. Namely, the honest user H , the adversary, and the simulator are nothing else but machines, and the protocols are structures. The view of H is then the restriction of the run (the transcripts of all states and in-/output of all machines during the protocols execution, also called trace) to the ports and state of H .

The definition, as presented so far, still has one drawback. We have not introduced the concept of a corruption. This can be accommodated by defining so-called systems. A *system* is a set of structures, where to each “corruption situation” (set of machines, which are corrupted) one structure corresponds. That is, when a machine is corrupted, it is not present anymore in the corresponding structure, and the adversary takes its place. For a trusted host, the corresponding system usually consists of structures for each corruption situation, too, where those connections of the trusted host, that are associated with a corrupted party, are under the control of the adversary.

We can now refine the definition of security as follows: A *real system* Sys_1 is as secure as an *ideal system* Sys_2 , if every structure in Sys_1 is as secure as the corresponding structure in Sys_2 .

A major advantage of a security definition by simulatability is the possibility of *composition*. The notion of composition can be sketched as follows: If we have on structure or system A (usually a protocol) implementing some other structure or system B (usually some primitive), and we have some protocol X^B (having B as a sub-protocol, i.e. using the primitive), then by replacing B by A in X^B , we get a protocol X^A which is as secure as X^B . This allows to modularly design protocols: first we design a protocol X^B , and then we find an implementation for B .

3. Continuously Polynomial Security

In this section, a new notion of polynomial runtime for the adversary and the protocol user H , *continuously polynomial*, is defined. For users and adversaries subject to our definition, terms like “guaranteed delivery after polynomial time” can be defined in a meaningful way. The definition of protocols which are *polynomially shaped* of Section 4 together with the restriction to weakly polynomial machines (ps-wp protocols) will ensure without explicit lifetime bounds that only polynomial-time computations are performed within polynomial time as seen by the protocol user H .

First, we demand from the protocol user H that it is *weakly polynomial*, as defined in [6]. There it is required that there is a polynomial p , such that in each activation, H runs at most $p(k + |I|)$ steps, where k is the security parameter, and $|I|$ is the length of all inputs H has received so far.⁸ We explicitly stress that this allows H s that do not halt, i.e., run infinitely long. It also does not forbid H to send messages to itself (possibly doubling the size of this “loop-back” message every time to get twice the computational power for the next activation), or to receive large messages.

To make sure that the induced security notion stays sensible, we will restrict to only polynomial prefixes of H 's view. That is, we consider only things that happen during polynomially-sized prefixes of H 's view.⁹ Here, the size of a view-prefix is the concatenated size of all inputs and outputs on H 's ports.

Second, an adversary A is required to be polynomial in H 's view. There are two obvious ways to do this: keeping A polynomial in the messages it *receives* from H , or keeping A polynomial in the messages it *sends* to H . We decided for a combination of both: in our definition, A must be polynomial in the size of the A - H -communication in *both* directions. We did so to give A more freedom: with the first notion, it would not be possible, in some cases, for A to simply forward protocol messages to H . Conversely, the second notion may forbid A to forward messages from H to the protocol. Thus, only our combined notion allows for a “dummy adversary” (an adversary that only acts as a relay between internal protocol lines and H). The concept of such a dummy adversary is useful, e.g., for proving concurrent composition properties.

However, this preliminary definition gives rise to a subtle problem with the proof of the composition theorem. In this proof, surrounding protocol machines are, for certain steps of the proof, simulated by the protocol user H . So an adversary considered in the proof of composability may have communication lines which are sometimes connected to protocol machines and sometimes connected to a protocol user mimicking these machines. Hence an adversary which is polynomial as described above could lose this property by the “regrouping” of machines during the composition proof, and the proof would fail.

Therefore, we introduce two specific communication lines which are *guaranteed* to connect the adversary and H .

⁸ This is similar to [13], where this approach is taken for the special case of secure function evaluations.

⁹ Alternatively, one could fix such a prefix with H and “hardwire” that bound into H to make it strictly polynomial in the traditional sense. However, in the case of standard security, the simulator is then chosen after H and thus knows the runtime bound of H . When trying to define notions like fairness (i.e., the property that the adversary eventually delivers messages), the simulator could then simply deliver all messages *after* the termination of the honest user H . This would circumvent the idea of a fair delivery.

The ports for these two lines will have names of the form $cpoly_{\dots}$, and such ports will not be allowed in any protocol. Now the total length of messages exchanged over these two specific lines is used as a lower bound for the “time” which has passed for H , and the adversary must be polynomial in this “volume” plus k . This volume includes the messages which is sent from the adversary to H in the same activation.

Counting a message, that is sent to H in the same activation, to the volume in which the adversary must be polynomial allows the adversary to receive (and, e.g., forward) arbitrarily long messages from the protocol. However, an adversary computing for a long time *must* send a long message to H to ensure that a long “elapse in time” is observed in the view of H . There is one important detail here: every prefix of the view of H is a sequence of results from whole activations. That is, if an adversary took a superpolynomial “debt” (e.g., by factoring a large integer), then the superpolynomial message which he is forced to send to H in the same activation will not be contained in any polynomial prefix of H 's view. So whenever the adversary is performing a superpolynomial number of computation steps, it is ensured that the result will not, not even in parts, be considered in the definition of security.

A further condition we impose on the adversary is the following: The adversary is required to read all incoming messages completely. This seemingly unnecessary condition has important consequences: Assume a protocol (e.g., for secure message transmission) in which a ciphertext is transmitted. Assume further that for generating a realistic first bit of the ciphertext, a runtime linear in the length of the message is required.¹⁰ Then a real adversary A could do the following: It intercepts the ciphertext, but reads only the first bit and forwards that bit to the honest user H . Since A only reads one bit, its running time is independent of the length of the transmitted message and it does not need to output anything on the $cpoly_{\dots}$ connection. However, the simulator now has the task to generate a realistic first bit, which takes a runtime linear in the length of the message. In the case of universal security, since the simulator is chosen before the honest user, this length may be larger than the number of steps the simulator may run without output on the $cpoly_{\dots}$ connection. So the simulator must output something there and the honest user can distinguish. By introducing the condition that the adversary reads all its inputs, this problem is fixed, since A now has to read the whole message, too, and hence also outputs on the $cpoly_{\dots}$ connection.

¹⁰ An example would be if the protocol prepended the bit $H^l(0)$ to the ciphertext, where l is the length of the message, and H a suitable function so that computing $H^l(0)$ cannot be done faster than in $\Omega(l)$. Clearly, an IND-CCA2 secure cryptosystem would not lose its security by such an addition.

As a technicality, messages sent from the adversary A to H over the specific line which influences A 's runtime must be delivered immediately to ensure the direct correspondence between runtime and messages received by H .¹¹

We turn to the actual definition:

Definition 3.1 (Continuously polynomial honest users and adversaries). *We call an honest user H continuously polynomial, if it is weakly polynomial, has ports $\text{cpoly_ha!}, \text{cpoly_ah?} \in \text{ports}(H)$, and the length function for cpoly_ah? is ∞ in every non-final state (i.e., all inputs on cpoly_ah? are written in full length to H 's view).*

We call an adversary A continuously polynomial, if

- *it has ports $\text{cpoly_ha?}, \text{cpoly_ah!}, \text{cpoly_ah}^{\text{q!}}$, and*
- *there is a polynomial p , s.t. for any closed collection \hat{C} of machines with $A \in \hat{C}$, and any possible view of A in \hat{C} (on security parameter 1^k), the following holds:*
 - *Let t_μ be the total number of Turing steps of A up to its μ -th activation (inclusive). Let c_μ be the total length of outputs on cpoly_ah! and inputs on cpoly_ha? up to A 's μ -th activation (inclusive). Then for all $\mu \in \mathbb{N}$ it is*

$$t_\mu \leq p(c_\mu + k).$$

- *Whenever A sends a message on cpoly_ah! , it is delivered immediately.*
- *A never sets its length functions to anything other than ∞ , and A always completely reads all incoming messages.¹²*

We can now define continuously polynomial security by simply restricting honest user and adversary to continuously polynomial ones:

Definition 3.2 (Continuously polynomial security). *Let (\hat{M}_1, S) and (\hat{M}_2, S) be structures (i.e., protocols), s.t. \hat{M}_1 and \hat{M}_2 have no port named cpoly_ah or cpoly_ha . Define¹³*

$$\begin{aligned} \text{Conf}_{\text{cpoly}}(\hat{M}_2, S) &:= \{(\hat{M}_2, S, H, A) \in \text{Conf}(\hat{M}_2, S) : \\ &\quad A \text{ and } H \text{ are continuously polynomial}\}, \\ \text{Conf}_{\text{cpoly}}^{\hat{M}_2}(\hat{M}_1, S) &:= \text{Conf}_{\text{cpoly}}(\hat{M}_1, S) \cap \text{Conf}^{\hat{M}_2}(\hat{M}_1, S). \end{aligned}$$

¹¹ To facilitate the presentation, we say that a message m from a machine M is delivered immediately over a port $p!$ if the receiving machine is activated with this message directly after M has entered a waiting state or a final state. In the model of [16, 6], this happens if the buffer \bar{p} is empty and M performs the commands $p! := m; p^{\text{q!}} := 1$.

¹² That is, in each activation, A takes at least $|I|$ steps, where $|I|$ is the length of A 's input in that activation.

¹³ Remember that in [6] $\text{Conf}^{\hat{M}_2}(\hat{M}_1, S)$ and $\text{Conf}(\hat{M}_2, S)$ are the sets of configurations (\hat{M}, S, H, A) so that H, A are valid honest user and adversary for the given protocol in the real and ideal model, respectively, and S is the set of service ports of the protocol $\hat{M} = \hat{M}_1, \hat{M}_2$, resp. Essentially, H and A are called valid if there are no open connections, and H only connects to service ports.

Less formally, the class of admissible honest users, adversaries and simulators is restricted to continuously polynomial ones.

If view is a view of some machine, then by $\text{pfx}_t(\text{view})$ we denote the longest prefix, s.t. the total length of all inputs and outputs in that prefix is bounded by $t \in \mathbb{N}$ (we will call such a prefix a t -prefix).

We call (\hat{M}_1, S) continuously polynomially as secure as (\hat{M}_2, S) (written: $\geq_{\text{sec}}^{\text{cpoly}}$), if for every configuration $\text{conf}_1 = (\hat{M}_1, S, H, A_1) \in \text{Conf}_{\text{cpoly}}^{\hat{M}_2}(\hat{M}_1, S)$, there exists a configuration $\text{conf}_2 = (\hat{M}_2, S, H, A_2) \in \text{Conf}_{\text{cpoly}}(\hat{M}_2, S)$ (essentially, this means that for continuously polynomial H, A_1 there is a continuously polynomial simulator A_2) s.t. for all polynomials l

$$\text{pfx}_{l(k)}(\text{view}_{\text{conf}_1, k}(H)) \approx_{\text{poly}} \text{pfx}_{l(k)}(\text{view}_{\text{conf}_2, k}(H)).$$

That is, for every adversary A_1 and user H that run with \hat{M}_1 , we require the existence of an adversary A_2 that runs with H and \hat{M}_2 , such that all polynomial prefixes of H 's view are indistinguishable in both protocols.

For universal security, (written: $\geq_{\text{sec}}^{\text{cpoly, uni}}$) we additionally require that A_2 does not depend on H .

4. A Generalised Composition Theorem

This section gives a generalised composition theorem for not necessarily terminating protocols. To this end, a new notion of polynomial runtime for *protocols* is introduced. For describing polynomial complexity, it is not only necessary to limit the computation time of a machine in each activation. It should also hold that superpolynomial “events” within the protocol yield a view for the user H having a superpolynomial representation. It should not pass unnoticed by H if a protocol machine gains superpolynomial computing power through a superpolynomial number of activations (which intuitively means that superpolynomial time must have passed) or by playing ping-pong with messages of growing size.

The definition of a *polynomially shaped* protocol ensures that each protocol machine can produce only messages of a total length which is polynomial in the length of the messages coming from outside the protocol, e.g. from the protocol user H or the adversary. The outside of the protocol is represented by a machine T in the definition below. If additionally, each protocol machine is weakly polynomial, then the number of Turing steps a protocol runs between two activations of H or the adversary is polynomially limited in the security parameter and the length of the overall protocol input.

Definition 4.1 (Polynomially shaped). *A collection \hat{C} of machines containing no master scheduler is called*

p -shaped for a function $p : \mathbb{N} \rightarrow \mathbb{N}$, if for all machines T s.t. $\hat{C} \cup T$ is closed (i.e., there are no open connections) the following property holds with overwhelming probability in the security parameter k :

Let o_μ denote the total length of the output of all machines in \hat{C} at position μ in the run of $\hat{C} \cup T$. Similarly, i_μ denotes the total length of the input of machines in \hat{C} on ports coming from T (i.e., ports $p?$ s.t. $p! \in \text{ports}(T)$). Further a_μ denotes the total number of activations of machines in \hat{C} at that point. Then

$$o_\mu + a_\mu \leq p(i_\mu + k).$$

The adversary or the user could try to gain superpolynomial computing power by playing “ping-pong” with a protocol which has no lifetime bound. However, this does not affect the security definition and computational assumptions can still be used, because security is defined by comparing only polynomial prefixes of the view of the user H . It is easy to see that results of a superpolynomial ping-pong cannot be contained in such a polynomial prefix if all machines are weakly polynomial, the protocol is polynomially shaped, and the user and the adversary are continuously polynomial. A superpolynomial number of invocations of the protocol either directly implies a superpolynomial view of the using machine H or it implies a superpolynomial view of the adversary. A result of such a superpolynomial computation can only appear in a superpolynomial view of the adversary. For a continuously polynomial adversary A and user H an event not visible in any polynomially view of the adversary cannot be visible in a polynomial prefix of the view of H . Even though the weakly polynomial machines could, in the long run, break any cryptosystem this does not imply distinguishability and computational assumptions can be used.

Next we generalise the composition theorem to continuously polynomial users and adversaries interacting with polynomially shaped protocols.

Note that the notion of polynomially shaped protocols is itself *not* closed under composition. A simple counterexample can be obtained from the two machines M_1, M_2 as follows. The machine M_1 has two input lines and one output line. It forwards each input to the output line and clocks the output line. The machine M_2 has one input line and one output line and acts as a repeater. It forwards each input to the output line and clocks the output line in the same activation. Both machines are polynomially shaped (as collections), but if we connect the two machines leaving one input line of M_1 open we obtain a collection which can generate infinite internal communication on one single input. This is a very bad effect as such a machine could run until it has solved some “hard” problem thereby invalidating computational assumptions.

So the generalised composition theorem states that a composed protocol is secure if it *remains* polynomially shaped. It is in the responsibility of the protocol designer to avoid “loops” when designing a protocol.

However, one can restrict the security definition to a subclass of polynomially shaped protocols which is closed under composition. Then the composition theorem still holds and e.g. loops cannot arise from composition.

A subclass of polynomially shaped protocols which is closed under composition can be obtained by restricting to protocols which give a shorter output than the total length of inputs given so far. This subclass contains a lot of natural protocols. It seems very difficult to find a subclass which is closed under composition and contains all natural protocols: for instance, a broadcast protocol has a larger output than the length of the input.

Intuitively, the generalised composition theorem says: Let a weakly polynomial protocol \hat{M}_1 use a sub-protocol \hat{M}'_0 such that the composition of \hat{M}_1 and \hat{M}'_0 is polynomially shaped. Let further \hat{M}_0 be a protocol which can connect to the protocol \hat{M}_1 in the same way as \hat{M}'_0 and for which the composition of \hat{M}_1 and \hat{M}_0 is polynomially shaped, too. Then the following holds: If \hat{M}_0 is at least as secure as \hat{M}'_0 according to Definition 3.2, then \hat{M}'_0 can be replaced by \hat{M}_0 without loss of security.

Theorem 4.2. *Let $(\hat{M}_0, S_0), (\hat{M}'_0, S_0), (\hat{M}_1, S_1)$ be structures (i.e., protocols), s.t. no port in \hat{M}_1, \hat{M}_0 , or \hat{M}'_0 is named `cpoly_ah` or `cpoly_ha`. Let then $(\hat{M}^\#, S) := (\hat{M}_1, S_1) \parallel (\hat{M}_0, S_0), (\hat{M}^*, S) := (\hat{M}_1, S_1) \parallel (\hat{M}'_0, S_0)$ (i.e., $\hat{M}^\#$ is the composition of \hat{M}_1 and \hat{M}_0 , while \hat{M}^* is the composition of \hat{M}_1 and \hat{M}'_0). Assume that*

- *The collections of machines $\hat{M}^\#$ and \hat{M}^* are polynomially shaped.*
- *The collection of machines \hat{M}_1 is weakly polynomial.*
- *It is $(\hat{M}_0, S_0) \geq_{\text{sec}}^{\text{cpoly}} (\hat{M}'_0, S_0)$.*
- *It is $\text{ports}(\hat{M}'_0) \cap S_1^c = \text{ports}(\hat{M}_0) \cap S_1^c$.¹⁴*

Then we have

$$(\hat{M}^\#, S) \geq_{\text{sec}}^{\text{cpoly}} (\hat{M}^*, S),$$

i.e., $\hat{M}^\#$ is continuously polynomially as secure as \hat{M}^* .

The same holds for universal security.

Proof. In the following proof, we assume all polynomials to be monotone. Furthermore, k always denotes the security parameter.

¹⁴ This is a formally necessary structural condition on the available ports, which also appear in the original version of the composition theorem, cf. [6] for details.

Let $conf_1 := (\hat{M}^\#, S, H, A_1) \in \text{Conf}_{\text{cpoly}}^{\hat{M}^\#}(\hat{M}^\#, S)$ be given (i.e., let some suitable continuously polynomial honest user H and adversary A_1 be given). To prove the theorem, we have to find a continuously polynomial simulator A_2 , s.t. $conf_2 := (\hat{M}^*, S, H, A_2) \in \text{Conf}_{\text{cpoly}}(\hat{M}^*, S)$ and

$$\text{pfx}_l(\text{view}_{conf_1}(H)) \approx_{\text{poly}} \text{pfx}_l(\text{view}_{conf_2}(H)) \quad (1)$$

for all polynomials l .

To prove universal security, we additionally need, that A_2 does not depend on H .

W.l.o.g. we can restrict our attention to honest users which do not terminate. Other honest users can be transformed into an honest user H' which 1. does not terminate, 2. is continuously polynomial, and for which 3. the view of the original H is a prefix of the new H' .

Consider the combination H' of $H \cup \hat{M}_1$. Since H and \hat{M}_1 are weakly polynomial, so is their combination H' . Since H does not terminate, the length function for cpoly_ah of H' is always ∞ , therefore H' is continuously polynomial.

Since $(\hat{M}_0, S_0) \geq_{\text{sec}}^{\text{cpoly}} (\hat{M}'_0, S_0)$ there is a continuously polynomial simulator A_2 , s.t.

$$\text{pfx}_L(\text{view}_{\hat{M}_0 \cup H' \cup A_1}(H')) \approx_{\text{poly}} \text{pfx}_L(\text{view}_{\hat{M}'_0 \cup H' \cup A_2}(H'))$$

for all polynomials L .

To show (1) from this, it is sufficient to show that for any polynomial l there is a polynomial L , s.t. the l -prefix of H is (with overwhelming probability) contained in the L -prefix in H' (intuitively, this means that the view of H does not grow superpolynomially by inclusion of \hat{M}_1).

First, consider the view of H in the real model (i.e. in the collection $H \cup A_1 \cup \hat{M}^\#$). Fix a polynomial l . Let then the random variable μ_k be the index in the run of the last element of the l -prefix of H 's view (more formally, the minimal μ_k , s.t. $\text{pfx}_{l(k)}(\text{view}(H))$ is contained in the first μ_k elements of the run).

Since A_1 is continuously polynomial, there exists a polynomial r (dependent on l) s.t. up to the μ_k -th step in the run the total length of A_1 's output is bounded by $r(k)$.

Since the total length of the output of H up to the μ_k -th step is bounded by $l(k)$ (by definition of l), we conclude that the total input of $\hat{M}^\#$ coming from H and A_1 is bounded by $l(k) + r(k)$. Since $\hat{M}^\#$ is polynomially shaped, it follows (by Definition 4.1) that the total output of $\hat{M}^\#$ is bounded by some polynomial $p(k)$ (dependent on l, r) with overwhelming probability.

So the length of the inputs and outputs of H' (being the combination of H and $\hat{M}_1 \subseteq \hat{M}^\#$) is bounded by $L_1(k) := l(k) + r(k) + l(k) + p(k) + p(k)$ (the summands being upper bounds for: in-/output of H ; output of A_1 ; output of H ; output of \hat{M}_1 ; output of $\hat{M}^\#$ (the latter appearing as input to H')). Therefore the l -prefix of H 's view appear with overwhelming probability in an L -prefix of the view of H' (in the real model).

Using the fact that \hat{M}^* is polynomially shaped, too, we get by analogous discussion that the l -prefix of H 's view appear with overwhelming probability in an L_2 -prefix of the view of H' . By choosing L as a polynomial bounding both L_1, L_2 , the remaining goal is shown, so (1) follows. \square

5. Relations to Polynomial Security

Continuously polynomial security allows for users and adversaries which are not strictly polynomial. On the other hand, every strictly polynomial pair of user and adversary can be interpreted as continuously polynomial ones—only the formally necessary cpoly_ah and cpoly_ha connections have to be added (but they need not be used).

However, this inclusion does not immediately imply that continuously polynomial security can be related in any way to the well-known concept of strictly polynomial security (for which only strictly polynomially bounded users and adversaries are considered). Namely, in case of continuously polynomial security, not only real adversaries, but also simulators may be drawn from a larger pool of possible adversaries. So in principle, continuously polynomial security of a system could mean that even for strictly polynomially bounded real attacks, a simulator might be necessary which is *not* polynomially bounded; strictly polynomial security might not follow from continuously polynomial one.

Fortunately, we can still show the following, not immediately obvious relation between continuously polynomial and strictly polynomial security:

Theorem 5.1. *Let (\hat{M}_1, S) and (\hat{M}_2, S) be polynomially shaped structures (i.e., protocols) satisfying $(\hat{M}_1, S) \geq_{\text{sec}}^{\text{cpoly}} (\hat{M}_2, S)$. Then $(\hat{M}_1, S) \geq_{\text{sec}}^{\text{poly}} (\hat{M}_2, S)$, i.e. continuously polynomial security implies strictly polynomial security for polynomially shaped protocols.*

Proof. Assume $(\hat{M}_1, S) \geq_{\text{sec}}^{\text{cpoly}} (\hat{M}_2, S)$. To prove that $(\hat{M}_1, S) \geq_{\text{sec}}^{\text{poly}} (\hat{M}_2, S)$ we have to show that for every $conf_1 := (H, A_1, \hat{M}_1, S) \in \text{Conf}_{\text{poly}}^{\hat{M}_1}(\hat{M}_1, S)$ (i.e., for any strictly polynomial honest user H and real adversary A_1), there is a simulator A_2 with $conf_2 := (H, A_2, \hat{M}_2, S) \in \text{Conf}_{\text{poly}}(\hat{M}_2, S)$ (i.e., a strictly polynomial adversary), s.t.

$$\text{view}_{conf_1}(H) \approx_{\text{poly}} \text{view}_{conf_2}(H). \quad (2)$$

Without loss of generality we can assume that no port of H and A_1 is named cpoly_ah or cpoly_ha .

First, since H and A_1 are strictly polynomial, and \hat{M}_1 is polynomially shaped, there is a polynomial p , s.t. $p(k)$ is with overwhelming probability an upper bound for the total length of all messages sent in a run of $\{H, A_1\} \cup \hat{M}_1$.

Therefore, we can construct a new real adversary A_1^p from A_1 as follows: We add new ports $\text{cpoly_ha}^?$, $\text{cpoly_ah}!$, and $\text{cpoly_ah}^<!$. A_1^p completely reads all its inputs and behaves as A_1 would (and ignores cpoly_ha -messages).

Only if the total length of the incoming messages received throughout the run exceeds $p(k)$, all messages are forwarded to H through `cpoly_ah` instead of simulating A_1 . Clearly, since A_1 was strictly polynomial, A_1^p is continuously polynomial.

Similarly, we construct a new honest user H' from H: We add new ports `cpoly_ah?`, `cpoly_ah!`, `cpoly_ah^!`. The length function on `cpoly_ah?` is set to ∞ , but any input on this port is ignored. No output is ever sent on the new ports. Clearly, since H was strictly polynomial, H' is continuously polynomial.

Intuitively, we have added a new connection between H and A_1 which is not used at all, but needed to fulfil the formal requirements of continuously polynomial honest users and adversaries. Since the new connection is not used, and A_1^p 's communication limit $p(k)$ is reached only with negligible probability, it immediately follows that

$$view_{conf_1}(H) \approx view_{H' \cup A_1^p \cup \hat{M}_1}(H'). \quad (3)$$

Since the machines H' and A_1^p are continuously polynomial, by $(\hat{M}_1, S) \geq_{sec}^{cpoly} (\hat{M}_2, S)$ there is a continuously polynomial simulator A_2^p s.t. for all polynomials l

$$pfx_l(view_{H' \cup A_1^p \cup \hat{M}_1}(H')) \approx_{poly} pfx_l(view_{H' \cup A_2^p \cup \hat{M}_2}(H')) \quad (4)$$

Since in runs of $H' \cup A_1^p \cup \hat{M}_1$ the adversary A_1^p sends anything on `cpoly_ah` only with negligible probability, A_2^p only sends with negligible probability on that port, too.

Therefore it is possible to construct a new simulator A_2 from A_2^p by removing the ports `cpoly_ah?`, `cpoly_ah!`, `cpoly_ah^!` (here A_2 simply terminates when A_2^p would have sent on `cpoly_ah`). Since only with negligible probability data is ever transmitted over these ports, it is immediate that

$$view_{H' \cup A_2^p \cup \hat{M}_2}(H') \approx view_{H \cup A_2 \cup \hat{M}_2}(H) \quad (5)$$

using the same identification of views as in (3).

Further, since A_2^p is continuously polynomial, and thus can only make a polynomial number of Turing steps while not receiving on `cpoly_ah` or sending on `cpoly_ah`, it follows that A_2 is strictly polynomial.

Setting $conf_2 := (H, A_2, \hat{M}_2, S)$, and combining (3), (4) and (5), we get

$$pfx_l(view_{conf_1}(H)) \approx_{poly} pfx_l(view_{conf_2}(H)) \quad (6)$$

for all polynomials l .

And since H and A_1 are strictly polynomial, and \hat{M}_1 is polynomially shaped, it follows from Definition 4.1 that there is a polynomial l s.t.

$$pfx_l(view_{conf_1}(H)) = view_{conf_1}(H)$$

with overwhelming probability (i.e., that the view is almost always of length at most $l(k)$).

The analogue holds for H, A_2 and \hat{M}_2 , so from (6) follows (2), which concludes the proof.

Note that the above proof does not work for universal security, since A_2 depends on p which again depends on H. \square

This theorem has several applications: first, it shows that continuously polynomial security is not “too weak” a security notion. In fact, anyone who would accept strictly polynomial security as a sufficiently strong security assumption should also find continuously polynomial security sufficiently strong.

Second, established results which need strictly polynomial security of a given system as a prerequisite can also be used with continuously polynomially secure systems. Consider the following example: You have proven continuously polynomial security for each of the many components of a large e-commerce protocol. The protocol and each of its components are—to avoid fixing a priori runtimes—formulated as a ps-wp protocol. Of course you use Theorem 4.2 to derive the security of the composed protocol. (Note that already this step would not have been possible with the strictly polynomial version of the composition theorem from [16], since for its application, the large protocol must be strictly polynomial-time.) Using [1, Theorem 5.1]¹⁵ and Theorem 5.1, you can now show that, e.g., integrity properties—as defined in [1]—the ideal version of the large protocol has are inherited by the composed (completely real) protocol. Since these steps involve composition of ps-wp systems, showing the same integrity properties of the composed real system is non-trivial when using only results which deal with strictly polynomial security.

6. A Simple Example

We will show the applicability of our definition using the very simple example of secure message transmission (SMT) over an authenticated channel using a one-time-pad. Note that despite its simplicity, such a functionality could not have been modelled in earlier approaches without bounding number and length of the messages (e.g., the SMT-functionality in [16] is parametrised by explicit bounds s and L for number and length of the messages).

To keep the presentation of this example simple, we assume a key exchange functionality KE that is has the following specification: When receiving a message of the form 1^L from party P_{Alice} , a random $K \in \{0, 1\}^L$ is sent to the parties P_{Alice} and P_{Bob} and a message 1^L is sent (with im-

¹⁵ This theorem states the preservation of integrity properties and is applicable even to protocols which are not polynomial-time.

mediate scheduling) to the adversary (informing him that a key exchange took place).¹⁶

We now want to implement the following functionality SMT: Whenever a message m is received from P_{Alice} , a message $1^{|m|}$ is sent (and immediately scheduled) to the adversary, and the message m is sent to P_{Bob} . (Note that here the adversary can reorder the messages, since he may choose when to schedule the delivery of m from SMT to P_{Bob} .)

The protocol we propose for SMT is fairly straightforward. When receiving a message m , P_{Alice} first requests a key of length $L := |m| + k$ from the functionality KE where k is the security parameter. Upon receipt of the key K it sends $c := (m0^k) \oplus K$ to P_{Bob} over an authenticated channel. P_{Alice} repeats this protocol for each new message.

Then, upon reception of a key K from KE and a ciphertext c from P_{Alice} , P_{Bob} calculates $\tilde{m} := c \oplus K$. If \tilde{m} has the form $m0^k$, P_{Bob} outputs m .

Obviously this protocol is ps-wp, for each input of length L it generates a communication volume of $5L + 4k$.

We now give a proof sketch that this protocol indeed realises SMT: First, consider the case that no party is corrupted. Then, for each adversary A_1 we construct a simulator A_2 as follows: A_2 simulates the adversary, as well as P_{Alice} and P_{Bob} . When the simulator A_2 receives a message 1^L from SMT (informing it that a message of length L is being sent), a random message $\tilde{m} \in \{0, 1\}^L$ is given to P_{Alice} as input, thus creating a fake view for the adversary. When P_{Bob} finally outputs the message $\tilde{m} \in \{0, 1\}^L$ (and the adversary schedules that output), the simulator schedules the delivery of the corresponding message m from SMT to the environment.

Since the adversary (and the honest user) does not learn the key K generated by KE, they may not distinguish whether the ciphertexts intercepted by the adversary correspond to the messages generated by the honest user, or to random messages of the same length generated by the simulator. However, one fine point must be taken care of: If several messages are in the process of being sent, the adversary may reorder the keys from KE differently on P_{Alice} 's and P_{Bob} 's side. Then it is possible that wrong messages get decoded. However, in order for this to happen, two generated keys have to match on the last k bits. Since the honest user H is continuously polynomial, for each prefix of length p of H 's view at most $O(p(k))$ messages are sent, thus at most $O(p(k))$ keys generated, so the probability of such a collusion of keys is bounded by $O(p(k)^2 2^{-k})$.

We add a short remark here: If instead of the one-time-pad an only computationally secure cipher had been used, we would additionally have to note that since the proto-

col is polynomially shaped, and the honest user and adversary are continuously polynomial, the adversary and honest user together can run at most a polynomial number of steps. Hence, they cannot break the cipher with more than a negligible probability.

The last thing left to check for the uncorrupted case is that our simulator is indeed continuously polynomial. Whenever the simulator gets a message 1^L from SMT, a simulation of P_{Alice} and P_{Bob} runs. The runtime needed for this simulation is polynomial in L . However, in the simulation P_{Alice} immediately sends a message of length $L + k$ which is passed to the simulated adversary. So the runtime needed for the simulation is polynomial in the length of the messages the simulated adversary gets. And since the simulated adversary is continuously polynomial, its runtime (which is also an upper bound for its incoming communication) is polynomial in its communication on the $\text{cpoly} \dots$ ports. So the total runtime of the simulator is polynomial in its communication on the $\text{cpoly} \dots$ ports (since all the communication of the simulated adversary on these ports is passed to H), and thus the simulator is continuously polynomial.

So at least in the uncorrupted case, our protocol is a continuously polynomially secure implementation of SMT.

The cases where P_{Alice} or P_{Bob} are corrupted are even easier, since here the simulator can learn the transmitted message. Checking that the simulator in these cases is also continuously polynomial is done very similarly to the uncorrupted case. We omit the details of these cases.

7. Applying our idea to the UC framework

We have shown how to allow for a more general class of polynomial-time protocols in the framework of reactive simulatability. Our approach can be adapted to the UC framework [7]. Several differences between the UC and the reactive simulatability framework that induce minor changes in our definitions are worth mentioning here:

- In the UC model, there is no concept of ports, the recipient of a message is dynamically specified by the sending machine. Therefore in Definition 3.1 we cannot consider the messages sent only over the $\text{cpoly} \dots$ ports. Instead, the messages intended to be sent over this connection must be marked in a special way, e.g., by a special prefix which is not allowed in messages sent to the protocol.
- In the UC model, indistinguishability of real and ideal protocols is not formulated in terms of the view, but in terms of the final output of the environment. Instead of quantifying over polynomial prefixes of the views in Def. 3.1 we would simply quantify only over environments that must terminate after a polynomial length of input and output.

¹⁶ This key exchange functionality could then easily be implemented by doing an L -bit Diffie-Hellman-style key exchange.

- In the UC model, it is possibly that additional machines appear during the execution of the protocol (these can model e.g., new participants, newly invoked subroutine threads, multiple instances of a functionality). The definition of a polynomially shaped protocol (Def. 4.1) should therefore require, that the outputs of *all* machines (including submachines that are created only during the execution of the protocol) are bounded polynomially in the external input of *all* machines. Only considering the machines present at the beginning of the protocol execution would not be sufficient, of course.

8. Conclusions

We have motivated and introduced a novel formulation of the intuitive requirement of simulatable security with respect to polynomially bounded attacks and protocol runs. We have shown that the induced security notion allows for composition and is at least as strong as the established notion of strictly polynomial security.

We have presented our approach in the modelling of reactive simulatability [6]. The ideas presented here should be applicable to the UC model [7], too.

Many of the oddities that arise with a combination of simulatable security and a strict polynomial bounding (as with strictly polynomial security) of all entities in a protocol are settled by our approach. Nonetheless, more radical techniques are possible: e.g., message scheduling and scheduling of activations could be separately managed by distinguished entities. In such a setting, machines can send messages which are scheduled *while* the sending machine remains activated. Then, a very intuitive formulation of “polynomial runtime,” which can even more closely model realistic protocol situations, would seem possible.

Acknowledgements

We thank Michael Backes and Ran Canetti for helpful discussions and comments.

References

- [1] M. Backes. *Cryptographically Sound Analysis of Security Protocols*. PhD thesis, Universität des Saarlandes, 2002. Online available at <http://www.zurich.ibm.com/~mbc/papers/PhDthesis.ps.gz>.
- [2] M. Backes. Unifying simulatability definitions in cryptographic systems under different timing assumptions. In R. Amadio and D. Lugiez, editors, *Concurrency Theory, Proceedings of CONCUR 2003*, number 2761 in Lecture Notes in Computer Science, pages 350–365. Springer-Verlag, 2003. Full version online available at <http://eprint.iacr.org/2003/114.ps>.
- [3] M. Backes. E-mail communication with the authors, June 2004.
- [4] M. Backes, D. Hofheinz, J. Müller-Quade, and D. Unruh. Fair and reliable networks in the context of simulatable security. Unpublished, 2004.
- [5] M. Backes, B. Pfizmann, and M. Waidner. A composable cryptographic library with nested operations. In *10th ACM Conference on Computer and Communications Security, Proceedings of CCS 2003*, pages 220–230. ACM Press, 2003. Extended abstract, extended version online available at <http://eprint.iacr.org/2003/015.ps>.
- [6] M. Backes, B. Pfizmann, and M. Waidner. Secure asynchronous reactive systems. IACR ePrint Archive, Mar. 2004. Online available at <http://eprint.iacr.org/2004/082.ps>.
- [7] R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42th Annual Symposium on Foundations of Computer Science, Proceedings of FOCS 2001*, pages 136–145. IEEE Computer Society, 2001. Full version online available at <http://www.eccc.uni-trier.de/eccc-reports/2001/TR01-016/revn01.ps>.
- [8] R. Canetti. Universally composable signature, certification and authentication. IACR ePrint Archive, Aug. 2004.
- [9] R. Canetti. E-mail communication with the authors, Jan. 2005.
- [10] R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. IACR ePrint Archive, Jan. 2005. Online available at <http://eprint.iacr.org/2000/067.ps>.
- [11] R. Canetti, H. Krawczyk, and J. B. Nielsen. Relaxing chosen-ciphertext security. In D. Boneh, editor, *Advances in Cryptology, Proceedings of CRYPTO 2003*, number 2729 in Lecture Notes in Computer Science, pages 565–582. Springer-Verlag, 2003. Full version online available at <http://eprint.iacr.org/2003/174.ps>.
- [12] A. Datta, R. Küsters, J. C. Mitchell, and A. Ramanathan. On the relationships between notions of simulation-based security. In J. Kilian, editor, *Theory of Cryptography, Proceedings of TCC 2005*, number 3378 in Lecture Notes in Computer Science, pages 476–494. Springer-Verlag, 2005. Online available at http://www.ti.informatik.uni-kiel.de/~kuesters/publications_html/DattaKuestersMitchellRamanathan-TCC-2005.ps.gz.
- [13] O. Goldreich. Secure multi-party computation. Unpublished, online available at <http://www.wisdom.weizmann.ac.il/~oded/PS/prot.ps>, Oct. 2002.
- [14] D. Hofheinz, J. Müller-Quade, and R. Steinwandt. On modeling IND-CCA security in cryptographic protocols. IACR ePrint Archive, Feb. 2003. Full version of [15].
- [15] D. Hofheinz, J. Müller-Quade, and R. Steinwandt. On modeling IND-CCA security in cryptographic protocols. In *4th Central European Conference on Cryptology, Proceedings of WARTACRYPT 2004*, pages 47–49, 2004. Extended abstract, full version is [14].

- [16] 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, Proceedings of SSP 2001*, pages 184–200. IEEE Computer Society, 2001. Full version online available at <http://eprint.iacr.org/2000/066.ps>.

A. Glossary

In this section we explain the technical terms of the reactive simulatability framework used in this paper. Longer and formal definitions can be found in [6].

$[\hat{C}]$: The completion of the collection \hat{C} . Results from adding all missing buffers to \hat{C} . **Conf_x(\hat{M}_2, S)**: Set of ideal configurations that are possible for structure (\hat{M}_2, S) . **Conf_x \hat{M}_2 (\hat{M}_1, S)**: Set of real configurations possible for structure (\hat{M}_1, S) . **ports(M)**: The set of all ports, a machine or collection M has. **to clock**: To write 1 onto a clock out-port. **EXPSMALL**: The set of exponentially small functions. **NEGL**: The set of negligible functions (asymptotically smaller than the inverse of any polynomial). **buffer**: Stores message sent from a simple out- to a simple in-port. Needs an input from a clock port to deliver. **clock out-port p^{cl}!**: A port used to schedule connection. **closed collection**: A collection is closed, if all ports have all their necessary counterparts. **collection**: A set of machines. **combination**: The combination of a set of machines is a new machine simulating the other machines. A set of machines can be replaced by its combination without changing the view of any machine. **composition**: Replacing sub-protocols by other sub-protocols. **computational security**: When in the security definition, honest user and adversary are restricted to machines running in polynomial time, and the views are computationally indistinguishable. **configuration**: A structure together with an honest user and an adversary. **free ports**: The free ports of a collection are those missing their counterpart. **honest user**: Represents the setting in which the protocol runs. Also called environment. **intended structure**: A structure from which a system is derived making a structure for every corruption situation. **master clock port clk^{cl}?**: A special port by which the master scheduler is activated. **master scheduler**: The machine that gets activated when no machine would get activated. **perfect security**: When in the security definition, the real and ideal run have to be identical, not only indistinguishable. Further the machines are completely unrestricted.¹⁷ **run**: The transcript of everything

that happens while a collection is run. Formally a random variable over sequences. $run_{conf,k,l}$ is the random variable of the run when running the configuration $conf$ upon security parameter k , restricted to its first l elements. If k is omitted, a family of random variables is meant. If l is omitted, we mean the full run. **service ports**: The ports of a structure to which the honest user may connect. They represent the interface of the protocol. As service ports are most often ports of a buffer, they are sometimes specified through the set S^c of their complementary ports; S^c consists of all ports which directly connect to a service port. **simple in-port p?**: A port of a machine, where it can receive messages from other machines. **simple out-port p!**: As simple in-port, but for sending. **statistical security**: When in the security definition the statistical distance of polynomial prefixes of the views have a statistical distance which lies in a set of small functions *SMALL* (in the security parameter k). Usually *SMALL* = *NEGL*. Further the machines are completely unrestricted.¹⁷ **structure**: A collection together with a set of service ports, represents a protocol. **view**: A subsequence of the run. The $view(M)$ of some collection or machine M consists of the run restricted to the ports and states of M . Possible indices are as with runs.

B. Length Functions

In Section 1, we mentioned that in security definitions which handle only strictly polynomial protocols it is often necessary to restrict the amount of data (lengths of inputs, number of invocations) a protocol can handle by some polynomial in the security parameter. We saw in Section 4 that the notion of continuously polynomial security allows to consider a much larger class of protocols, namely protocols which are ps-wp. This frees protocols from the necessity of terminating after some amount of input; rather protocols are only required to be polynomial in the “input from outside”.

In earlier versions of the reactive simulatability definitions and the modelling of universal composability (e.g., in [7]), the following problem arose: consider e.g. the seemingly trivial functionality/trusted host, that has two in-ports and two out-ports (representing two parties) and on each pair of in-/out-port would just echo every input. In order to make this functionality strictly polynomial, it is now necessary to restrict the amount of echoed data to some polynomial p . Then the functionality has to terminate after receiving p messages on the first port, otherwise it might have to spend superpolynomial time by ignoring the incoming messages on that port. Then of course the functionality would not echo anything on the second port, even if no message

¹⁷ In [6] a machine can in every activation for a given input and current state only reach one of a finite number of states (this convention has been chosen for simplicity [3]). However, this cannot even model the simple Turing machine that tosses (within one activation) coins until a 1 appears, and then stores the number of coin tosses. Therefore we will here adopt the convention that each state can have a countable number of potential successor states, from which one is chosen fol-

lowing some distribution depending on the input and the current state.

has been echoed there yet. This introduces a flow of information between the two echo ports which certainly was not the intention of the original functionality.

To handle this artefact and allow functionalities to “switch off” selected ports, [1] introduces so-called *length functions*. These allow a machine to set the maximal length of messages it can receive through a given port at a given time. In particular, by setting the length function on a port to 0, the port is blocked and will not be activated by messages on that port, so that ignored messages do not consume runtime.

Since with continuously polynomial security, we do not need strictly polynomial protocols, one might wonder whether it is still necessary to consider and use length functions in this modelling, since these are an answer to a problem which is actually solved by our modelling in another manner. This question will be addressed in the present section, where we will show that we can in fact assume all protocol machines to have no length functions.¹⁸

The question is therefore whether a ps-wp protocol/functionality with length functions can be modified into another ps-wp protocol/functionality without length functions so that the security is not affected. Fortunately the following straightforward modification already has the desired property: we say a machine M' results from another machine M by removing length functions if M' has no length functions, but otherwise behaves as M does. That is, when receiving a message, the content of the message after the prefix the length function of M indicates is ignored, and only that prefix is used for the simulation of M (or the message is ignored, if the length function is 0). In other words, M' simulates the length functions of M without actually having them. When \hat{M} is a collection, removing length functions means removing them from every machine in \hat{M} .¹⁹

Since obviously the difference between M and M' is only a formal property, not a difference in behaviour, we would expect M' to be a suitable replacement for M . This is confirmed by the following

Lemma B.1. *Let \hat{M}_i ($i = 1, 2$) be collections without master schedulers, and let \hat{M}'_i result from \hat{M}_i by removing length functions. Then it holds that*

- \hat{M}_i is polynomially shaped iff \hat{M}'_i is.
- If \hat{M}_i is weakly polynomial, so is \hat{M}'_i .

¹⁸ Formally, by a machine *without length functions* we mean a machine, whose length functions are ∞ in every non-final state.

¹⁹ A careful study of the definition of machines in [6] shows, that formally we can define the machine resulting from removing length functions from a machine $M = (\text{name}, \text{Ports}, \text{States}, \delta, l, \text{Ini}, \text{Fin})$ simply as $M' = (\text{name}, \text{Ports}, \text{States}, \delta, \infty, \text{Ini}, \text{Fin})$, where ∞ denotes the length function yielding ∞ for all ports and non-final states.

- *The following are equivalent:*

$$\begin{aligned} (\hat{M}_1, S) \geq_{\text{sec}} (\hat{M}_2, S), & \quad (\hat{M}_1, S) \geq_{\text{sec}} (\hat{M}'_2, S), \\ (\hat{M}'_1, S) \geq_{\text{sec}} (\hat{M}_2, S), & \quad (\hat{M}'_1, S) \geq_{\text{sec}} (\hat{M}'_2, S) \end{aligned}$$

Here \geq_{sec} denotes one of the following security notions: *perfect / statistical / strictly polynomial / continuously polynomial in the flavours of standard or universal security*.

The main idea of the proof is that the removal of length functions does not change the behaviour of the protocol, therefore the equivalences of the three security relations. Then it remains to be seen that the machines do not need superpolynomial runtime in the input to ignore the inputs (this shows the modified machines to be weakly polynomial), and that the amount of output does not change (this shows the resulting structures to be polynomially shaped). Note that such a property would not hold for strictly polynomial structures, since by removing a length function from a blocked port the resulting machine would have to ignore but accept an unbounded number of messages on that port, which is not allowed for strictly polynomial machines. The full proof goes as follows:

Proof. Let \hat{C}_i be some collection, s.t. $\hat{M}_i \cup \hat{C}_i$ is a closed collection (i.e., no port is unconnected). Then removing the length functions from \hat{M}_i yields a collection $\hat{M}'_i \cup \hat{C}_i$, so that the run of $\hat{M}'_i \cup \hat{C}_i$ differs from that of $\hat{M}_i \cup \hat{C}_i$ only in the following points: 1. the inputs of machines in \hat{M}'_i are changed (i.e., they are longer since with unmodified \hat{M}_i they were added to the run in truncated form), 2. there are additional activations of machines in \hat{M}'_1 with empty outputs.

Now let any machine T without length functions be given, s.t. $\hat{M}_i \cup T$ is closed. Consider then a run *run* of $T \cup \hat{M}_i$ with security parameter k and the corresponding run *run'* of $T \cup \hat{M}'_i$ (i.e., the runs result from the same random choices). Let $\mu \in \mathbb{N}$. Then let t_μ denote the total length of the output of T , a_μ the number of activations of machines in \hat{M}_i , and o_μ the total length of the output of machines in \hat{M}_i , all up to the μ -th activation of T in *run* (cf. Definition 4.1). Let t'_μ , a'_μ , and o'_μ be defined analogously for *run'*. By setting $\hat{C}_i := \{T\}$ the considerations at the beginning of the proof tell us that

$$a_\mu \leq a'_\mu, \quad t_\mu = t'_\mu, \quad o_\mu = o'_\mu.$$

Note further that whenever a simple machine (no master scheduler) is activated, some other machine necessarily sent a nonempty message to that effect. This allows to conclude $a'_\mu \leq t'_\mu + o'_\mu$.

If then \hat{M}_i is p -shaped then from these inequalities we get with overwhelming probability for all μ

$$\begin{aligned} a'_\mu + o'_\mu &\leq t'_\mu + 2o'_\mu = t'_\mu + 2o_\mu \\ &\leq t'_\mu + 2p(t_\mu + k) \leq (2p + \text{id})(t'_\mu + k), \end{aligned}$$

so \hat{M}'_i is $(2p + \text{id})$ -shaped.

If on the other hand \hat{M}'_i is p -shaped, it is

$$a_\mu + o_\mu \leq a'_\mu + o'_\mu \leq p(t'_\mu + k) = p(t_\mu + k),$$

so \hat{M}_i is p -shaped. So the claim follows that \hat{M}_i is polynomially shaped iff \hat{M}'_i is.

Now assume some weakly polynomial machine M is given, and M' results by removing length functions. Let some input sequence for M resp. M' be given. Then for activation μ we distinguish two cases: First, the length function of M is not zero on the port containing input. Then M' only has to ignore any trailing input, which can be done with an overhead polynomial in the running time of M in that activation. Second, if the length function of M is zero, the overhead of M' is constant, i.e., in particular polynomially bounded in the size of the non-empty input. So summarising we see that the overhead of M' is polynomial in the running time of M and the length of the input, so M' is weakly polynomial, too. Therefore this shows the claim \hat{M}'_1 is weakly polynomial if \hat{M}_1 is.

Considering again the results from the beginning of the proof, and letting \hat{C}_1 be the honest user together with the real adversary, we see that the view of the honest user is not changed by removing the length functions from the machines in \hat{M}_1 , so $(\hat{M}_1, S) \geq_{\text{sec}} (\hat{M}_2, S)$ is equivalent to $(\hat{M}'_1, S) \geq_{\text{sec}} (\hat{M}_2, S)$ and $(\hat{M}_1, S) \geq_{\text{sec}} (\hat{M}'_2, S)$ is equivalent with $(\hat{M}'_1, S) \geq_{\text{sec}} (\hat{M}'_2, S)$. Similarly with \hat{C}_2 being the honest user together with the simulator, we see that $(\hat{M}_1, S) \geq_{\text{sec}} (\hat{M}_2, S)$ is equivalent with $(\hat{M}_1, S) \geq_{\text{sec}} (\hat{M}'_2, S)$. This shows the third claim. \square