

Information-Theoretically Secure Protocols and Security Under Composition*

Eyal Kushilevitz[†]
Department of Computer Science
Technion Israel Institute of Technology
Haifa 32000, Israel.
eyalk@cs.technion.ac.il

Yehuda Lindell
Department of Computer Science
Bar-Ilan University
Ramat Gan, 52900, Israel.
lindell@cs.biu.ac.il

Tal Rabin
IBM T.J.Watson Research
New York, USA.
talr@us.ibm.com

December 20, 2009

Abstract

We investigate the question of whether security of protocols in the information-theoretic setting (where the adversary is computationally unbounded) implies the security of these protocols under concurrent composition. This question is motivated by the folklore that all known protocols that are secure in the information-theoretic setting are indeed secure under concurrent composition. We provide answers to this question for a number of different settings (i.e., considering perfect versus statistical security, and concurrent composition with adaptive versus fixed inputs). Our results enhance the understanding of what is necessary for obtaining security under composition, as well as providing tools (i.e., composition theorems) that can be used for proving the security of protocols under composition while considering only the standard stand-alone definitions of security.

Keywords: Theory of cryptography, secure multiparty computation, information-theoretic security, security under composition

*An extended abstract of this work appeared in the 38th *STOC*, 2006.

[†]Research supported by grant 36/03 from the Israel Science Foundation.

1 Introduction

1.1 Background

In the setting of secure multiparty computation, a set of parties with private inputs wishes to jointly compute some function of their inputs. Loosely speaking, the security requirements of such a computation are that nothing is learned from the protocol other than the output (privacy), and that the output is distributed according to the prescribed functionality (correctness). More exactly, the result of an execution of a secure protocol must be like the result of an “ideal execution” with a trusted party who honestly computes the function for the parties (cf. [Can00] or [Gol04, Section 7.1]). These security requirements must hold in the face of a malicious adversary who controls some subset of the parties and can arbitrarily deviate from the protocol instructions. We focus on secure function evaluation here (although we do not know of any difference between this setting and that of reactive functionalities). Our results hold for both *static* and *adaptive* corruptions.

Secure multiparty computation has been studied in a number of different scenarios. One important distinction relates to the power of the adversary, yielding two main settings:

1. **The information-theoretic setting:** In this setting, the adversary is computationally unbounded. Thus, security here does not rely on any unproven complexity assumptions; rather, it is “information theoretic”. There are two levels of security that have been considered here:
 - (a) *Perfect security:* Very informally, here the result of a real execution of the protocol with a real adversary must be exactly the same as the result of an ideal execution with a trusted party and an ideal-world adversary/simulator.
 - (b) *Statistical security:* Here, the result of the real protocol execution need “only” be statistically close to the result of an ideal execution.

It has been shown that, assuming that more than $2/3$ of the parties are honest (or assuming a simple majority if the parties are additionally given a broadcast channel), it is possible to securely compute any functionality [BGW88, CCD88, RB89, Bea89]. We note that these results assume that the parties can communicate via perfectly-secure communication channels.

2. **The computational setting:** In this setting, the adversary is assumed to run in probabilistic polynomial-time, and the security of protocols typically relies on the assumed hardness of some problem (like factoring a large composite into its prime factors). Under appropriate cryptographic assumptions, it has been shown that any functionality can be securely computed, even if an overwhelming majority of the parties are corrupted [Yao86, GMW87].¹

Another important distinction relates to the setting in which the protocol is executed. We relate here to two possible scenarios:

1. **The stand-alone model:** In this setting, the secure protocol is executed only once, and it is assumed that this is the only protocol being executed. This was the standard setting for analyzing protocols initially; e.g., the results of [Yao86, GMW87, BGW88, CCD88, RB89, Bea89], cited above, were all obtained in this setting.
2. **Security under composition:** In this setting, a protocol is executed many times, possibly alongside other (secure and insecure) protocols. Security in such cases is referred to as *security under composition*. There are many different types of composition, and we will mention two here:

¹When no honest majority is assumed, the definition of security is slightly relaxed so that fairness is not required.

- (a) *Concurrent general composition (a.k.a. universal composition)*: In this setting, a protocol is run many times in an arbitrary network, possibly alongside other (secure and insecure) protocols. *Concurrent* composition implies that the scheduling of the delivery of all messages sent (even messages sent between honest parties) is fully controlled by the adversary. See, e.g., [DM00, PW00, Can01].
- (b) *Concurrent self composition*: In this setting, a single protocol is run many times concurrently (with itself) in a network. See, e.g., [Fei90, DDN00, DNS04]. Again, the scheduling of the delivery of all messages sent is under the control of the adversary.

The question of security of protocols under concurrent composition has received a lot of attention recently (see, e.g. [DNS04, DM00, PW00, Can01, Lin03a, Lin03b, Lin04] and much more).

It was believed that all protocols (or at least, all known protocols) that are secure in the information-theoretic setting, even when proven secure only in the stand-alone model, are secure under concurrent composition. In this paper, we study this belief and more generally the interesting connection between information-theoretic security and security under composition. Our aim in initiating this study is twofold. First, as our results demonstrate, understanding this connection deepens our understanding of what is required for obtaining security under composition. Second, due to its complex nature, the task of proving the security of protocols under composition is a difficult one and we obtain a number of results that simplify this task.

1.2 Our Results

As we have mentioned, the folklore prior to our work seems to have been that all protocols that are secure in the information-theoretic stand-alone setting are also secure under concurrent general composition. The reason for this folklore appears to be based on the fact that all known protocols in the information-theoretic setting are proven secure using a **straight-line black-box simulator**,² and the existence of such a simulator was believed to suffice for proving the security of a protocol under concurrent general composition. Indeed it is common to hear the claim that universal composability (that implies security under concurrent general composition) is essentially the same as stand-alone security with the exception that the simulator is not allowed to rewind the adversary. We begin by showing that this folklore belief is false. That is, we obtain the following, informally stated, proposition:

Theorem 1.1 (counter-example – informal): *There exist protocols that are statistically secure in the stand-alone information theoretical model and are proven secure using straight-line black-box simulation, and yet are not secure under concurrent general composition.*

The idea behind the proof of Theorem 1.1 is quite simple and is due to the issue of “delayed inputs”. Namely, in the setting of concurrent general composition, the parties run completely asynchronously and so some parties may begin a protocol execution before others begin. In fact, some parties may begin a protocol execution before others have even determined their input to that execution. This results in a possible scenario whereby the first message that one party sends in a protocol execution can be set as the input of another party to that execution. If the first message of the parties has high entropy, then in the stand-alone setting one party’s first message will equal another party’s input with only negligible probability. Thus, one can construct a protocol that will fail if this happens (and so

²A **black-box simulator** is one that works while being given only *oracle access* to the adversary; see [Gol01, Sec. 4.5] for a detailed definition. Such a simulator is **straight-line** if it interacts with the adversary in the same way as real parties. That is, it first sends the adversary all of the messages it expects to receive from the honest parties in the first round. It then proceeds round by round, without ever going back.

will not be secure under concurrent general composition), but is otherwise statistically secure with straight-line simulation.

Although Theorem 1.1 demonstrates that stand-alone information-theoretically secure protocols are not necessarily secure under concurrent general composition, the fact that the source of the problem seems to be “delayed inputs” (as described above) implies that there may be a simple way to circumvent the problem. In addition, the proof of Theorem 1.1 relies inherently on the ability of the stand-alone protocol to fail (albeit with negligible probability). This raises the question as to whether or not delayed inputs are problematic in the setting of *perfect* security. We consider these questions and more.

Perfect security. We prove the following theorem regarding perfectly secure protocols.

Theorem 1.2 (perfect security – informal): *Every protocol that is perfectly secure in the stand-alone model, and has a straight-line black-box simulator, is secure under concurrent general composition.*

The intuition behind the proof of Theorem 1.2 is that if a protocol is not secure under concurrent general composition, then there must be some setting of inputs and random tapes for the honest parties in the concurrent setting for which the protocol “fails” (i.e., does not behave as it should). Although such a “fail event” may happen with only negligible probability in the stand-alone setting, this is enough to contradict perfect security that requires that protocols *never* fail.

Contrasting Theorem 1.2 and Theorem 1.1 we uncover a fundamental gap between the statistical and perfect notions of security; namely, perfectly secure protocols have a *real advantage* over statistically secure protocols, and so perfect security is not just an issue of aesthetics. To the best of our knowledge, this is the first example of such a separation.

Beyond its theoretical interest, Theorem 1.2 provides a useful tool for proving the security of protocols under concurrent composition. Specifically, it suffices to consider the standard stand-alone definitions and security in the far more complex setting of concurrent composition can be derived as long as a perfect, black-box straight-line simulator is constructed. Thus, a corollary of Theorem 1.2 is that the protocol of [BGW88] is secure under concurrent general composition.³

We remark that, under strictly more stringent conditions, a theorem similar to Theorem 1.2 was proven in [DM00], and that the theorem of [DM00] is also sufficient to derive the security of [BGW88] under concurrent general composition. See Section 1.3, for a more detailed comparison of their work with ours.

Fixed inputs. As we have seen, *statistical* security under straight-line black-box simulation is insufficient for obtaining security under concurrent general composition. Furthermore, our proof of this relies on the fact that one party’s input to be derived from another party’s behavior in the protocol. This raises the question as to whether straight-line black-box simulation suffices when inputs cannot be dynamically chosen, and in particular in the setting of *fixed inputs* where each honest party receives all of its inputs before any execution begins. That is, at the onset, each party receives a vector of inputs $\mathbf{x} = (x_1, x_2, \dots)$ and uses x_i as its input to the i^{th} protocol execution. (We note that despite the fact that this looks like a very relaxed notion of security, until now there was no known separation between security under general composition with fixed or with adaptively-chosen inputs.) We show that security under straight-line black-box simulation *does suffice* for obtaining security under concurrent general composition with “fixed inputs”. That is, we prove:

³The protocol of [BGW88] was designed for the synchronous setting whereas, in the setting of concurrency, the network is essentially asynchronous [B84, F89, BCG93]. Thus, what we really refer to here is an appropriately modified version of [BGW88]; see Section 2, for a detailed discussion on the issue of synchronicity.

Theorem 1.3 (fixed inputs – informal): *Every protocol that is secure in the stand-alone model, and has a straight-line black-box simulator, is secure under concurrent general composition with fixed inputs.*

Theorem 1.3 holds for computational, statistical and perfect security. Therefore, as a corollary to Theorem 1.3, we obtain that the protocols of [CCD88, RB89, CDD⁺99] – which have straight-line black-box simulators – are all secure under concurrent general composition with fixed inputs. Also, by combining Theorem 1.3 with Theorem 1.1, we obtain the first separation between security under general composition with fixed inputs and with adaptively-chosen inputs. That is, we have the following corollary:

Corollary 1.4 *There exist protocols that are secure under concurrent general composition with fixed inputs, and are not secure under concurrent general composition with adaptively-chosen inputs.*

Since general composition implies self composition, we have that every protocol that is secure in the stand-alone model, and has a straight-line black-box simulator, is secure under *concurrent self composition with fixed inputs*. This is interesting for historical reasons as the classic problem of concurrent zero-knowledge [DNS04] is exactly in the setting of concurrent self composition with fixed inputs.

Theorem 1.3 states that any protocol that is proven secure with a straight-line black-box simulator is secure under concurrent general composition with fixed inputs (and therefore also secure under concurrent self composition with fixed inputs). Since self composition is a more restricted setting, this begs the question as to whether straight-line black-box simulation – without fixed inputs – suffices for achieving concurrent *self* composition. However, by the equivalence shown in [Lin04] (stating that concurrent self and general composition are almost equivalent for adaptively chosen inputs), this is clearly not possible because, as stated in Theorem 1.1, black-box straight-line simulation does *not* suffice for concurrent general composition.

Input availability/start synchronization. Finally, we ask if there is a simple property of protocols that can be required, in addition to a protocol having a straight-line black-box simulator, so that it will be secure under concurrent general composition. We have shown that requiring fixed inputs, i.e. that all inputs to *all* invocations of the executed protocols be set prior to the start of *any* protocol, yields security under concurrent general composition. However, this is a very restrictive requirement that applies to the invocations and not to the protocol, and it severely limits the applicability of the protocols. We therefore ask whether it is possible to achieve concurrent general composition by imposing a less restrictive (and more realistic) property on the protocol. Our aim is to define a minimal property that circumvents the problem of “delayed inputs”, with the hope that this suffices for achieving security under composition.

We show that it is sufficient to require that the inputs of all parties participating in each *specific* invocation of a protocol be fixed before this execution begins. We call this property *input availability*. We stress that the inputs need not be fixed before *all* protocol executions begin (as in the assumption of fixed inputs); rather, we require that the inputs of all parties to *each specific* execution be fixed before that execution begins. Formally, this holds if all honest parties have read (and fixed) their input to an execution before any of the honest parties sends the first message of the protocol. In order to see why this seemingly small modification makes a difference, recall Theorem 1.3 that states that the existence of a straight-line black-box simulator is sufficient to demonstrate security under composition as long as the parties’ inputs are fixed before any execution begins. Now, if the parties’ inputs are not fixed at the onset, but are fixed before a specific execution begins, then it is still possible to view everything that happened before the specific execution as *auxiliary input* that determines the parties’ inputs to that execution. This enables us to reduce the security of the specific execution to the

stand-alone setting, while the auxiliary input is used to emulate all the other concurrent executions. Stated differently, when input availability holds, the setting is essentially the same as that of fixed inputs (cf. Theorem 1.3) and so security under concurrent general composition holds.

We stress that when input availability does *not* hold, it is not possible to emulate all concurrent executions outside of a specific execution using an auxiliary input. This is because, by definition, auxiliary input is fixed before the protocol execution begins and thus must be *independent* of the random tapes of the parties. However, if a party’s input is determined only after other parties have already begun the execution, then there is no guarantee of independence. Indeed, this is exactly how we prove the counter-example of Theorem 1.1.

Importantly, it is possible to ensure that the above property holds using a simple protocol preamble that we call **start synchronization**. Specifically, when a party has received its input to an execution and is ready to begin executing, it broadcasts a message to all other parties that it is ready to begin. The parties then only commence the actual execution *after* they have heard from all other parties that they are ready. This has the effect of ensuring that the actual protocol execution only begins after all parties’ inputs have already been fixed (because no party begins until it hears from all other parties that they are ready to start, meaning that they already have input). This modification captures the essence of what is needed for the composition which is that the inputs for each single execution of a protocol be known prior to the start of the execution. Thus, we have the following theorem:

Theorem 1.5 (black-box straight-line and start synchronization – informal): *Every protocol that is secure in the stand-alone model and has start synchronization and a straight-line black-box simulator is secure under concurrent general composition.*

We remark that Theorem 1.5 holds for computational, statistical and perfect security. Since any protocol can easily be modified so that it has start synchronization, we achieve the following corollary (informally stated):

Corollary 1.6 *Assuming a broadcast channel and an honest majority, every (standard, non-reactive) functionality can be computed by a statistically secure protocol under concurrent general composition.*

Note that Corollary 1.6 is stated in terms of *functionalities* while our other results are stated as results about *protocols*; the corollary is obtained by applying Theorem 1.5 to the protocols of [RB89, Bea89], once modified by adding start synchronization. We remark that this is the first proof of the existence of such protocols. The main novelty of this corollary is that it is the first proof that there exist protocols that achieve fairness in the setting of concurrent general composition, assuming an honest majority (previously, this was only known for the case that more than two thirds of the parties are honest). The fact that the protocol of [RB89] is universally composable and thus secure under concurrent general composition was claimed with a high-level proof sketch in an early version of [Can01])

Summary. Table 1 below summarizes our results; the security level “all” means that the result holds for computational, statistical *and* perfect security. We remark that the proof of Theorem 1.3 is actually a simple special case of the proof of Theorem 1.5. We therefore first prove Theorem 1.5 (in Section 6) and only then Theorem 1.3 (in Section 7).

A definitional discussion. A fundamental difference between the definitions of security in the stand-alone setting and security under composition via *universal composability* [Can01] is the question of whether the distinguisher is *passive* and views the result of the execution after it terminates (as in the stand-alone setting) or whether it is *interactive* and views the intermediate stages of the computation as well (as in universal composability). (Note that although we formally use the definition of concurrent

Theorem	Simulator	Security Level	Composition
Theorem 1.1	straight-line, black-box	statistical or computational	<i>not necessarily</i> secure under concurrent general composition
Theorem 1.2	straight-line, black-box	perfect	secure under concurrent general composition
Theorem 1.3	straight-line, black-box	all	secure under concurrent general composition with fixed inputs
Theorem 1.5	straight-line, black-box	all + start synchronization	secure under concurrent general composition

Table 1: A summary of our results

general composition, and not universal composability, these are known to be equivalent [Lin03b].) In this light, it is possible to technically interpret our results as a study of the relation between security under passive and interactive distinguishers, and not as a study of the relation between stand-alone security and security under composition. Nevertheless, the equivalence of universal composability (a stand-alone definition with an interactive distinguisher) and concurrent general composition (a definition that considers composition and a passive distinguisher) demonstrates that these perspectives are essentially the same.

1.3 Related and Subsequent Work

As we have mentioned, the question of composition of information-theoretic protocols was previously studied in [DM00]. They showed that any protocol that is secure under their definition (which is a slight modification of the definition of [MR91]) is secure under concurrent general composition. There are three main requirements in the definition considered in [DM00]: **(a)** the protocol must be perfectly secure, **(b)** a straight-line black-box simulator must be used, and **(c)** there must be a fixed “committal round” at which point all of the parties’ inputs are fully defined by the protocol traffic, but no parties have yet received output. In contrast, our proof of Theorem 1.2 requires (a) and (b), but not (c). (We note that the proof of [DM00] relies explicitly on this property, as stated in [DM00, Section 4.3].) Apart from Theorem 1.2, there is no overlap between our work and the work of [DM00]. Indeed, they leave the question of statistical security open, and we provide both positive and negative answers (depending on the setting).

An interesting question that arises from our work relates to the necessity (or lack thereof) of straight-line simulation in our results. Specifically, both Theorems 1.2 and 1.3 require that the original black-box simulator be “straight-line”. This raises the question as to whether the requirement that the simulator be straight-line is inherent or whether similar theorems could be proven also if the simulator is not straight-line. In the conference version of this paper [KLR06] we claimed that the theorems could be extended to any black-box simulator by showing a method to convert a rewinding polynomial-time simulator into a straight-line exponential-time simulator (even though the resulting simulator is not efficient, this can still be of some interest in the information-theoretic setting). However, our proof was erroneous and it was shown in [BMU07] that there exist perfectly secure protocols (for the stand-alone model) for which *any* black-box simulator must use rewinding and which are *not* secure under concurrent composition even with an inefficient simulator.

We remark that the proof of a similar result to Theorem 1.1 was obtained independently and concurrently in [HU06]. In [HU06] it is also proven that any protocol that is perfectly secure under bounded concurrent general composition is also secure under (unbounded) concurrent general composition. We prove a much stronger result here; namely, that it suffices to prove perfect security in the *stand-alone* model.

2 On Concurrent Composition and the Synchronous Model

In this paper, we consider protocols that were proven secure in the stand-alone model, and analyze their security in the setting of concurrent composition. There are two “standard” modes of message scheduling in the stand-alone setting: *synchronous scheduling* and *asynchronous scheduling*. In the first, all parties send the messages of a given round at the same time and these messages arrive at their destination instantaneously.⁴ In the asynchronous setting [B84, F89, BCG93], the adversary is given full control over the scheduling of the sending and receiving of messages, with the requirement that *all messages be eventually delivered*. Now, in the standard model of concurrent computation [Can01], the adversary is given full control over the delivery of messages, as in the asynchronous model, but need not ever deliver messages sent by the honest parties. That is, unlike in the asynchronous setting where all messages must eventually be delivered, in the standard concurrent setting the adversary may completely block some of the honest parties and *never* deliver their messages. As such, it is not possible to guarantee output delivery or fairness (at least when concurrency is modeled in this way). Thus, the best that we can hope for in the concurrent setting is *security with abort* (meaning that the adversary may receive output while the honest parties do not), as in the stand-alone setting with no honest majority.⁵

As we have mentioned, in this paper we analyze the security under concurrent composition of protocols that were designed for the stand-alone model. In order to maintain compatibility between these models, we consider the security under concurrent composition of protocols that are *secure with abort in the stand-alone model when the adversary has full control over message delivery*. We stress that this includes the capability of the adversary to completely block messages between honest parties, unlike the stand-alone asynchronous model where all messages must eventually be delivered. Since this is not the standard model for stand-alone computation (neither for the synchronous nor asynchronous model), we need to discuss how to execute these protocols when the adversary has full control including blocking. For the synchronous model, we show below how to convert any secure protocol designed as a stand-alone protocol (with or without abort) into a protocol that is secure *with abort* in the model where the adversary has full control over message delivery (including blocking). The case of the asynchronous setting is more straightforward. The adversary’s actions in the concurrent setting (where message blocking is allowed) is actually a prefix of an adversary’s actions in the asynchronous setting where all messages must eventually be delivered (this can be seen by viewing a blocked message simply as a message that was not yet delivered). This immediately implies that the security of any protocol designed in the asynchronous setting is not violated by the scheduling of an adversary who can also block messages. We stress, however, that if the adversary does not complete the delivery of messages as required in the asynchronous setting then output delivery and fairness may not be achieved. This implies that the level of security obtained is still only security with abort. We conclude that asynchronous stand-alone protocols can be executed unmodified in the presence of an adversary who is also allowed to block messages, with the result that security with abort is achieved. We now proceed to describe the transformation for the synchronous protocols.

⁴It is typically assumed that the adversary may receive the messages from the honest parties in any given round before it sends its own. This is called “rushing”.

⁵In the asynchronous stand-alone model of secure computation, output delivery and fairness are guaranteed to some extent. Loosely speaking, the definition of [BCG93] allows the adversary to exclude up to t honest parties inputs from the computation (where t denotes the maximum number of corrupted parties), but guarantees fairness and output delivery for all others. This can be achieved because the adversary must eventually deliver all messages; see [BCG93] for details. (We remark that the exclusion of t parties is possible because the honest parties cannot know if those t parties are corrupted and will never send their messages or they are honest and the adversary has not yet delivered their messages. Thus, they cannot wait for all messages because they may be from corrupted parties and so may never arrive.) In contrast, when no parties’ messages need to be eventually delivered, it is not possible to do better in general than security with abort.

Modifying synchronous protocols for the concurrent setting. Protocols that were designed for the stand-alone synchronous model are proven secure based on the assumption that the protocol advances from round to round based on some global clock (or round synchronizer). In particular, it is assumed that all parties send and receive their round i messages before proceeding to round $i + 1$. When the adversary has the ability to schedule the messages as it wishes then some parties may receive all round i messages and therefore proceed to send their round $i + 1$ messages before others have concluded the previous round. In such a case, the security properties of the protocol may be violated.

In order to solve this problem, we modify the stand-alone synchronous protocol so that no party will proceed to round $i + 1$ before all parties conclude round i , irrespective of when and if messages are delivered. We achieve this by adding a “synchronization step” between each round of the protocol. Specifically, after an honest party has sent its round i messages, it is instructed to wait for all of its incoming round i messages. Then, upon receiving all round i messages, the party broadcasts a notification to all other parties that it has completed round i . It then continues to wait to receive notifications that *all* parties have completed round i . Once all notifications are received, the parties proceed to round $i + 1$. (We can also assume that honest parties concatenate the round index to every message they send so that the adversary cannot replay messages from previous rounds.) Such modifications can be trivially made to every protocol that is secure *with or without abort* in the synchronous setting, and the result is a protocol that remains *secure with abort* when the adversary has full control over message delivery. (We stress that when applying the transformation to a protocol that is secure without abort, meaning that fairness and output delivery are guaranteed, the result is still a protocol that is only secure with abort. Thus, in all cases fairness and output delivery are no longer guaranteed since the adversary can mount a simple denial-of-service attack by not forwarding a parties’ messages.)

3 Definitions and Notation

In this paper we study the connection between a number of different notions of security for multiparty computation. Specifically, we refer to information-theoretic security (perfect and statistical) and to computational security. In addition, we consider stand-alone computation (where only a single protocol execution takes place) versus concurrent general composition. Finally, we consider the restriction to the case that all inputs are fixed ahead of time.

In this section, we present the definitions and terminology that we use in this paper. We provide only brief outlines of the definitions and refer the reader to [Gol04, Chapter 7] for definitions of secure multiparty computation in the stand-alone setting, and [Lin03b] for the definition of security under concurrent general composition.

Preliminaries. A function μ is negligible if for every polynomial p there exists an integer N such that for every $k > N$ it holds that $\mu(k) < 1/p(k)$. Two distribution ensembles $\{X(k, a)\}_{k \in \mathbb{N}, a \in \{0,1\}^*}$ and $\{Y(k, a)\}_{k \in \mathbb{N}, a \in \{0,1\}^*}$ are computationally indistinguishable, denoted $\{X(k, a)\} \stackrel{c}{\equiv} \{Y(k, a)\}$, if for every non-uniform polynomial-time distinguisher D there exists a negligible function μ such that for all $a \in \{0, 1\}^*$ and $k \in \mathbb{N}$

$$|\Pr[D(X(k, a)) = 1] - \Pr[D(Y(k, a)) = 1]| \leq \mu(k)$$

Two distribution ensembles as above are statistically close, denoted $\{X(k, a)\} \stackrel{s}{\equiv} \{Y(k, a)\}$, if for every non-uniform distinguisher (not necessarily polynomial-time) there exists a negligible function μ such

that for all a and k ,

$$|\Pr[D(X(k, a)) = 1] - \Pr[D(Y(k, a)) = 1]| \leq \mu(k)$$

We write $\{X(k, a)\} \equiv \{Y(k, a)\}$ if the distributions are identical.

Secure multiparty computation. In the setting of secure multiparty computation, n parties, denoted P_1, \dots, P_n , wish to jointly compute some function f of their inputs in a secure fashion. The standard way of defining security in this scenario is to require that a real protocol “behaves” just like an ideal execution involving a trusted third party who computes the function for the parties [GL90, Bea91, MR91, Can00]. More specifically, an ideal execution is defined in which the parties simply hand their inputs to a trusted third party, who then computes the function and returns the designated outputs to each party. When considering security with abort, the trusted party first hands the output to the adversary, who may then decide whether or not the honest parties also receive output [Gol04, Chapter 7]. A protocol is said to be **secure with abort** if for every adversary attacking a real execution of the protocol, there exists an ideal-world adversary (running in an ideal execution) such that the output of the honest parties and the ideal-world adversary in the ideal execution is “close” to the output of the honest parties and the real adversary in the real protocol execution. As we have mentioned, the real model that we consider in this paper is one where the adversary has full control over the delivery of messages between parties, including the ability to never deliver messages.

There are a number of different settings for this definition. In the **information-theoretic setting**, the real-world adversary and ideal-world adversary are computationally unbounded. Here, there are two notions of security: **perfect security**, where the outputs of the ideal model are required to be *identically distributed* to the outputs of the real model, and **statistical security**, where the outputs of the ideal model are required to be *statistically close* to the outputs of the real model. We say that a protocol **computes a function f with perfect security** (respectively, **statistical security**) if it meets the requirements in the information-theoretic setting (with the appropriate notion of “closeness”; identical distribution for perfect security and statistical closeness for statistical security). In the **computational setting**, the real-world and ideal-world adversaries run in probabilistic polynomial-time and the outputs of the ideal model are required to be *computationally indistinguishable* from the outputs of the real model. In this setting, we say that a protocol **computes a function f with computational security**.

Security is defined as follows. Let k be the security parameter. We denote a real execution of a protocol ρ with adversary \mathcal{A} by $\text{REAL}_{\rho, \mathcal{A}(z)}(k, \mathbf{x})$, where z is the auxiliary input of \mathcal{A} and \mathbf{x} is the vector of the parties’ inputs. Likewise, we denote by $\text{IDEAL}_{\mathcal{S}(z)}^f(k, \mathbf{x})$ an execution in the ideal model where the trusted party computes the function f , \mathcal{S} is the ideal-world adversary (or simulator) and k , z and \mathbf{x} are as above. For a full specification of the ideal and real executions, see [Can00] and [Gol04, Chapter 7]. We are now ready to present the definitions.

Definition 1 (secure computation):

1. **Information-theoretic security:** A protocol ρ computes f with statistical security and abort if for every real-model adversary \mathcal{A} there exists an ideal-model adversary \mathcal{S} running in time that is polynomial in the running-time of \mathcal{A} such that,

$$\left\{ \text{IDEAL}_{\mathcal{S}(z)}^f(k, \mathbf{x}) \right\}_{k \in \mathbb{N}; z, \mathbf{x} \in \{0,1\}^*} \stackrel{s}{\equiv} \left\{ \text{REAL}_{\rho, \mathcal{A}(z)}(k, \mathbf{x}) \right\}_{k \in \mathbb{N}; z, \mathbf{x} \in \{0,1\}^*}$$

Protocol ρ computes f with perfect security and abort if $\left\{ \text{IDEAL}_{\mathcal{S}(z)}^f(k, \mathbf{x}) \right\} \equiv \left\{ \text{REAL}_{\rho, \mathcal{A}(z)}(k, \mathbf{x}) \right\}$.

2. **Computational security:** A protocol ρ computes f with computational security and abort if for every non-uniform probabilistic polynomial-time real-model adversary \mathcal{A} there exists a non-uniform probabilistic polynomial-time ideal-model adversary \mathcal{S} such that,

$$\left\{ \text{IDEAL}_{\mathcal{S}(z)}^f(k, \mathbf{x}) \right\}_{k \in \mathbb{N}; z, \mathbf{x} \in \{0,1\}^*} \stackrel{c}{=} \left\{ \text{REAL}_{\rho, \mathcal{A}(z)}(k, \mathbf{x}) \right\}_{k \in \mathbb{N}; z, \mathbf{x} \in \{0,1\}^*}$$

All of our results in this paper hold for both static and adaptive adversaries. An adaptive adversary can “adaptively” choose which honest parties to corrupt throughout the execution, based on its view so far. In the case of a static adversary, the set of corrupted parties is fixed at the onset.

Black-box straight-line simulation. In our presentation, we will refer to black-box straight-line simulation. Informally speaking, a **black-box simulator** is a universal ideal-world adversary (i.e., it is a single simulator that works for all real adversaries) that interacts with the real adversary in a black-box way only (i.e., it uses only oracle access to the real adversary). Furthermore, a black-box simulator is **straight-line** if the real adversary maintains state between oracle calls, in the same way as in a real protocol execution. Stated differently, a black-box straight-line simulator interacts with the real adversary in essentially the same way as real parties in a real protocol interaction (the only difference is that the simulator may simultaneously play the role of many parties). When \mathcal{S} is a black-box simulator, we denote the output of an ideal-model execution by $\text{IDEAL}_{\mathcal{S}(z)}^f(k, \mathbf{x})$; note that \mathcal{S} is given oracle access only to \mathcal{A} and in particular is not given the auxiliary input z .⁶ This is the standard way of defining black-box simulation; see [Gol01, Sec. 4.5].

Concurrent general composition. In the setting of concurrent general composition, a secure protocol ρ that computes a function f is run concurrently to an arbitrary other protocol π (where π represents many arbitrary protocols running concurrently in the network). In order to define security, a real execution of π with secure protocol ρ is compared to an idealized setting where protocol π is run, but the parties also have access to a trusted party that computes f . Furthermore, instead of running ρ , the parties send their inputs to the trusted party (like in the regular ideal model). This setting is called the **hybrid model**, because a trusted party is used for computing f , but real messages are also sent for computing π . The security requirement here is that for every real adversary \mathcal{A} and every arbitrary protocol π , there exists an ideal-world adversary \mathcal{S} such that for every set of inputs to π , the output of the ideal-world adversary and the honest parties in an ideal/hybrid execution of π with a trusted party computing f is “close” to the output of the real adversary and the honest parties in a real execution of π with ρ . We stress that the inputs to ρ (or equivalently to f) are determined by π . This models the setting that ρ is run in an arbitrary network (modeled by π) and the inputs to ρ are influenced by other protocol executions that take place in this network. We denote by $\text{REAL}_{\pi, \mathcal{A}(z)}(\mathbf{x})$ a real execution of ρ with π , with adversary \mathcal{A} , inputs \mathbf{x} and auxiliary input z for \mathcal{A} . Furthermore, we denote by $\text{HYBRID}_{\pi, \mathcal{S}(z)}^f(\mathbf{x})$ a hybrid execution of protocol π with ideal calls to f , and inputs \mathbf{x} and z as above. Detailed definitions of security under concurrent general composition can be found in [Lin03b]. (We remark that although the above refers to a single execution of a protocol ρ , security in this case is actually equivalent to security in the case where a polynomial number of protocols ρ_1, ρ_2, \dots are concurrently with each other and with π .)

Concurrent general composition with fixed inputs. The only difference between this case and the above definition is that the input to ρ is fixed before the execution of π begins. Formally, this

⁶Observe that if \mathcal{S} is given the auxiliary input, then it is no longer really a black-box simulator. In particular, one can define a “universal” adversary \mathcal{A} that receives for auxiliary input the description of an adversary and then runs that adversary. If \mathcal{S} is given the auxiliary input, then it has the strategy of \mathcal{A} .

is modeled by handing each party a pair of inputs (x_i, y_i) where x_i is party P_i 's input to π and y_i is its input to ρ . (Likewise, when considering the more general setting that π is run concurrently with protocols ρ_1, ρ_2, \dots then the inputs to all execution of ρ_1, \dots are fixed before π begins. In this case, each party receives a vector of inputs $(x_i, y_i^1, y_i^2, \dots)$ such that x_i is party P_i 's input to π and y_i^j is its input to ρ_j .) Security is required to hold for all possible vectors of inputs of all possible lengths (that is, we quantify over all adversaries, all polynomials $p(\cdot)$, and all possible vectors of inputs of length $p(n)$). As above, we will use the simpler formulation where a single execution of a protocol ρ is run together with an arbitrary protocol π .

Universal composability [Can01]. Universal composability is a definition of security that considers a stand-alone execution of a protocol in a special setting involving an environment machine \mathcal{Z} , in addition to the honest parties and adversary. As with the standard definition, ideal and real models are considered where a trusted party carries out the computation in the ideal model and the real protocol is run in the real model. The environment adaptively chooses the inputs for the honest parties, interacts with the adversary throughout the computation, and receives the honest parties' outputs. Security is formulated by requiring the existence of an ideal-model simulator \mathcal{S} so that no environment \mathcal{Z} can distinguish between the case that it runs with the real adversary \mathcal{A} in the real model and the case that it runs with the ideal-model simulator \mathcal{S} in the ideal model. The importance of this definition is a composition theorem that states that any protocol that is universally composable is secure under concurrent general composition [Can01]. It has also been shown that full equivalence holds; that is, any protocol that is secure under concurrent general composition (as formulated above) is also universally composable [Lin03b]. Thus, all of our results here for concurrent general composition hold equivalently for universal composability.

4 A Counter-Example to the Folklore

It seems to be well-accepted folklore that any protocol that is secure with a *straight-line black-box simulator* is secure under concurrent composition. This folklore probably stems from the fact that all known protocols for the information-theoretic setting have black-box straight-line simulators and are assumed (though most of them do not have proofs) to be secure under concurrent composition. In this section we show that this folklore is false. Specifically, we present a protocol that is statistically secure *and* has a straight-line black-box simulator, yet is not secure under concurrent general composition. (Recall that in the setting of concurrent general composition, a secure protocol runs concurrently to arbitrary other protocols.) We note that the simulator that we present for this protocol also runs in polynomial-time. Thus, the counter-example holds also for the computational setting.

The main idea behind our counter-example is as follows. In the stand-alone model, the parties' inputs are all fixed before the execution begins. In contrast, in the setting of concurrent general composition, a party's input to the secure protocol may depend on messages that it receives in another protocol. Furthermore, since the scheduling is concurrent (and controlled by the adversary), it is possible that a party's input to a secure protocol ρ is defined *after* other parties have already begun running ρ . Such an event cannot happen in the stand-alone model, where all the parties' inputs are fixed before the execution begins.

Universal composability – discussion. We note that our example highlights an important issue regarding the definition of universal composability [Can01]. Namely, this definition has two stringencies over other definitions. First, it requires that the simulator be straight-line and black-box. Second, it requires that the inputs may be chosen adaptively (by an external environment). It seems to have been

generally believed (although not stated in [Can01]) that the first stringency is the main one needed for achieving composition. Our example shows that both are actually necessary.

The counter-example. Our counter-example consists of a certain protocol that computes the three-party constant function f defined by $f(x_1, x_2, x_3) = (0, 0, 0)$, where $|x_1| = |x_2| = |x_3| = k$. (I.e., all parties receive 0, irrespective of the values of their inputs; obviously, this function can also be computed by other protocols that do guarantee any type of security.) Specifically, we prove the following proposition:

Theorem 2 *Assuming an honest majority of participants, there exists a function f and a protocol that computes f with statistical/computational security in the stand-alone model and with black-box straight-line simulation, that is not secure under concurrent general composition. This holds for both static and adaptive corruptions.*

Proof: We begin by presenting a three-party protocol that computes the function $f(x_1, x_2, x_3) = (0, 0, 0)$ with statistical security in the stand-alone model and with black-box straight-line simulation, in the case of an honest majority.

Protocol 1

1. Party P_2 chooses a random string r_2 of length $k/2$ and sends it to P_1 .
2. Party P_1 chooses a random string r_1 of length $k/2$ and sends it to P_2 .
3. Parties P_1 and P_2 define $r = (r_1, r_2)$, where (a, b) denotes the concatenation of a and b . If one of the parties does not receive the r_i value from the other, then it just waits. If one of the parties receives an invalid value (i.e., one not of length $k/2$), then it sets r to a uniformly distributed string of length k .
4. Parties P_1 and P_2 both send r to P_3 .
5. If P_3 received the same string from both P_1 and P_2 , and the string equals its input x_3 , then it outputs 1. Otherwise, it outputs 0.
6. P_1 and P_2 always output 0.

We begin by showing that Protocol 1 is secure in the stand-alone model. Intuitively, the protocol is secure because P_3 only outputs 1 in the case that r equals its input. However, since at least one of P_1 or P_2 is honest, and thus chooses r_i at random ($i \in \{1, 2\}$), the probability that r equals P_3 's input is at most $2^{-k/2}$.

Claim 4.1 *Protocol 1 computes f with statistical/computational security in the stand-alone model with black-box straight-line simulation and an honest majority.*

Proof: In order to prove Claim 4.1, we construct a simulator \mathcal{S} as follows.

1. If no parties are corrupted, then \mathcal{S} internally runs the code of all parties for the adversary \mathcal{A} and outputs whatever \mathcal{A} outputs.
2. If P_3 is corrupted, then \mathcal{S} internally runs the code of the honest parties and hands the adversary \mathcal{A} controlling P_3 the messages that it expects to receive from the honest parties. At the end, \mathcal{S} outputs whatever \mathcal{A} outputs.

3. If P_1 is corrupted, then \mathcal{S} interacts with the adversary \mathcal{A} controlling P_1 just as the honest P_2 would (sending it a random r_2 and receiving back r_1). Upon receiving r_1 from \mathcal{A} , simulator \mathcal{S} simulates sending $r = (r_1, r_2)$ to P_3 . At the end, \mathcal{S} outputs whatever \mathcal{A} outputs.
4. If P_2 is corrupted, then \mathcal{S} behaves in a symmetrical way to the case that P_1 is corrupted.

Since we assume an honest majority, these are all the possible corruption scenarios. We claim that the above simulation results in outputs that are statistically close to a real execution. Notice that the only possible difference between a real execution with an adversary \mathcal{A} and an ideal execution with \mathcal{S} is in the case that the messages that P_3 receives from P_1 and P_2 equal its input x_3 (because in this case P_3 outputs 1, whereas in an ideal execution it always outputs 0). However, we claim that this can only happen with negligible probability. In the case that one of P_1 or P_2 receive an invalid string r_i , this is clear (because the aborted party sends a truly random string of length k to P_3). In all other cases, this follows from the fact that at least one of P_1 or P_2 are honest (recall that we assume an honest majority here). Therefore, the entropy of the string received by P_3 is at least $k/2$. Since x_3 is fixed before the execution began, the probability that an honest P_3 outputs 1 is at most $2^{-k/2}$. We conclude that Protocol 1 computes f with statistical (and hence computational) security,⁷ completing the proof of Claim 4.1. ■

We now prove that Protocol 1 is *not* secure under concurrent general composition.

Claim 4.2 *Protocol 1 does not compute f with statistical or computational security, in the setting of concurrent general composition.*

Proof: The idea behind the proof of this claim is that the input of P_3 is not necessarily fixed before P_1 and P_2 start running their protocol. Since we are working in a concurrent setting, this is inevitable. We begin by describing a protocol π that contains an ideal call to f . (The protocol π is the “arbitrary protocol” that runs concurrent to the secure Protocol 1.)

Arbitrary protocol π :

1. Party P_1 sends a random string $s \in_R \{0, 1\}^n$ to P_3 .
2. Parties P_1 and P_2 send the trusted party computing f the input 0^n .
3. Party P_3 sends the trusted party computing f the input s (as received from P_1) and outputs whatever output it receives back from the trusted party.
4. Parties P_1 and P_2 record their outputs whenever they receive them from the trusted party and they output these recorded values.

Consider now a real execution of π with Protocol 1 replacing the trusted computation of f (Steps 2-4), and an adversary \mathcal{A} who controls a corrupted P_1 . In such an execution, parties P_1 and P_2 begin running Protocol 1. (Note that P_3 does not start running Protocol 1 until it receives its input s from P_1 in π .) P_1 receives a string r_2 from the honest P_2 and chooses its own random r_1 . The adversarial strategy of \mathcal{A} controlling P_1 here is to send $s = r = (r_1, r_2)$ to P_3 *as part of the protocol π* (i.e., Step 1 in Protocol π), and to send r_1 to P_2 *as part of Protocol 1* (i.e., Step 2 in Protocol 1). \mathcal{A} then proceeds

⁷One minor issue here relates to the fact that by our description, it is possible that P_3 does not receive output. This occurs if it waits forever for P_1 and P_2 to send their messages. The simplest way of dealing with this is just to state that our definition of security does not require the parties to generate output. This makes sense when we move to the full concurrent setting, where the adversary controls all scheduling. If we want to assume some synchrony or timing, this is also fine. Then, P_3 will just wait the appropriate amount of time and will output 0 if it didn't receive both messages.

with Protocol 1 following the instructions of P_1 honestly. (In particular, it sends $r = (r_1, r_2)$ to P_3 and outputs 0.)

Notice that in a real execution of π with Protocol 1 and adversary \mathcal{A} , in Protocol 1 party P_3 always receives messages from P_1 and P_2 that equal its input x_3 . This is because P_3 sets its input x_3 to equal s , and P_1 is able to make $s = r$. Therefore, P_3 always outputs 1 from such an execution of Protocol 1 with π and \mathcal{A} . By the definition of π , this means that P_3 also always outputs 1 from π in such an execution with \mathcal{A} . (Recall that in a real execution of π with Protocol 1, the output of f is replaced with the output of Protocol 1.)

It remains to note that for all simulators working in a hybrid/ideal model where π is run together with an ideal call to a trusted party computing f , the output of the honest parties P_2 and P_3 from π is always 0, because this is the output *always* sent by the trusted party. Therefore, the outputs of the real and ideal (hybrid) executions are easily distinguishable. This completes the proof of Claim 4.2. ■

Theorem 2 follows by combining Claims 4.1 and 4.2. ■

5 Perfect Security and Concurrent General Composition

Theorem 2 states that black-box straight-line simulation and *statistical security* do not suffice for achieving security under concurrent general composition. In this section, we show that if the protocol has a black-box straight-line simulator and is *perfectly secure*, then it *is* secure under concurrent general composition.

Theorem 3 *Let f be a function and let ρ be a protocol that computes f with perfect security and abort in the stand-alone model, with a black-box straight-line simulator. Then, ρ computes f with perfect security under concurrent general composition. Furthermore, if the black-box simulator for ρ in the stand-alone model runs in polynomial-time, then so does the simulator for ρ in the setting of concurrent general composition. This holds for both static and adaptive adversaries.*

Proof: The idea behind the proof of this theorem is as follows. Assume, by contradiction, that ρ is not secure under concurrent general composition. Loosely speaking, this means that there exists a protocol π and an adversary \mathcal{A} such that a real execution of π and ρ with \mathcal{A} cannot be simulated in the hybrid world where π is run together with a trusted party who computes f .⁸ In particular, it must be that the output distribution of the adversary and honest parties in ρ , when running it together with π , is not the same as their output distribution when they just send their inputs to a trusted party computing f . The main idea here is to use this fact to attack a stand-alone execution of ρ . Specifically, we construct a stand-alone adversary \mathcal{A}_ρ who attacks ρ by internally simulating the entire π execution for \mathcal{A} , while running the protocol ρ externally with the honest parties. If \mathcal{A}_ρ 's simulation of π is “good”, then it follows that the *stand-alone* execution of ρ will be the same as when it is run together with π . Thus, the output of ρ in this stand-alone execution will not be the same as the output of an ideal execution with a trusted party computing f . The problem with this approach is that in order for \mathcal{A}_ρ 's simulation of π to be “good”, it must somehow guess the honest party's inputs and random coins in a way that will make the combination of the internal π simulation and the external ρ execution look the same as a full external execution of π with ρ . This is especially problematic because in the setting of concurrent composition, the honest parties' inputs to ρ may be *determined* by π (whereas in the stand-alone model, the inputs of the honest parties are fixed ahead

⁸We remark that we prove the theorem directly for security under concurrent general composition, without going through the definition of universal composability [Can01]. Those who feel more comfortable using universal composability can just think of the protocol π as the environment \mathcal{Z} and everything remains the same.

of time). Nevertheless, in the case of *perfect security*, this is not a problem because we only need \mathcal{A}_ρ 's attack on ρ to “succeed” with nonzero probability. In particular, even a statistical distance of $2^{-n^{O(1)}}$ between the real and ideal executions is not allowed. Now, with very small probability, \mathcal{A}_ρ 's guesses are “correct”, and in this case, the output distribution and execution of ρ with \mathcal{A}_ρ is the same as after an execution of π and ρ with \mathcal{A} . By the contradicting assumption, this output distribution is not simulatable in the ideal model, and so we have that the real and ideal distributions are not identical, in contradiction to the perfect security of ρ in the stand-alone model. We note that the above contains the main idea behind the proof. However, the structure of the actual proof is a little different. We now proceed to the formal proof. For the sake of simplicity, we provide the proof for the case of static adversaries. In the case of adaptive corruptions, the only difference is that the \mathcal{A}_ρ needs to corrupt a party whenever \mathcal{A} does and then provide \mathcal{A} with the entire state of that party. This is not a problem because the state of all parties in π is known completely by \mathcal{A} (who runs the entire π execution internally) and the state of a newly corrupted party in ρ is learned by \mathcal{A}_ρ when it carries out the corruption. (The simulation strategy in the case of adaptive corruptions is derived in the same way.)

In this proof, without loss of generality, we only consider the case that the arbitrary protocol π has a *single* ideal call to f . We note that in the case that a single universal simulator is provided for all protocols π (as will be the case here), this is equivalent to the case that π has any polynomial number of ideal calls to f (see the full version of [Lin03b] for a proof). Recall that we denote by $\text{REAL}_{\pi^\rho, \mathcal{A}(z)}(\mathbf{x})$ the distribution of the outputs of the adversary \mathcal{A} with auxiliary input z and the honest parties in a real execution of π with ρ when given input vector \mathbf{x} . Likewise, we denote by $\text{HYBRID}_{\pi, \mathcal{S}^{\mathcal{A}(z)}}^f(\mathbf{x})$ the distribution of the outputs of the adversary \mathcal{S} and the honest parties in a hybrid execution of π with a trusted party computing f when given input vector \mathbf{x} . (We don't use a security parameter k here in the notation because it is not needed when considering perfect security.)

Let f be a function, let ρ be a protocol that computes f with perfect security, and let \mathcal{S}_ρ be the black-box straight-line ideal-model simulator for ρ that is assumed to exist. Then, we construct a simulator \mathcal{S} for the setting of concurrent general composition as follows. Let \mathcal{A} be a real-model adversary and let π be an arbitrary protocol that contains a single ideal call to f . Then, \mathcal{S} invokes the real adversary \mathcal{A} and forwards all π -messages untouched between \mathcal{A} and the honest parties (recall that \mathcal{S} runs a real execution of π with external parties). In contrast, when \mathcal{A} begins the ρ execution, \mathcal{S} invokes a copy of \mathcal{S}_ρ and forwards all ρ -messages between \mathcal{A} and \mathcal{S}_ρ (in contrast to the π -messages that are sent externally, ρ is internally simulated by \mathcal{S} using \mathcal{S}_ρ). More specifically, when \mathcal{A} outputs a ρ -message to be sent to an honest party, \mathcal{S} hands this message to \mathcal{S}_ρ . Likewise, \mathcal{S}_ρ 's reply is handed back to \mathcal{A} as if coming from an honest party. In addition, whenever \mathcal{S}_ρ sends an input to the trusted party computing f , simulator \mathcal{S} sends the same input. Likewise, outputs sent to \mathcal{S} from the trusted party are forwarded to \mathcal{S}_ρ . (We note that the π execution continues by just forwarding π -messages between \mathcal{A} and the honest parties, even while the simulation of ρ takes place. Notice that the real execution of π can continue concurrently with the simulation of ρ because \mathcal{S}_ρ is *straight-line*, and so it never needs to rewind.⁹) Whenever \mathcal{A} halts, simulator \mathcal{S} copies \mathcal{A} 's output to its own output tape and halts.

We now prove that for every arbitrary protocol π (with a single ideal call to f) and every real adversary \mathcal{A} ,

$$\left\{ \text{HYBRID}_{\pi, \mathcal{S}^{\mathcal{A}(z)}}^f(\mathbf{x}) \right\}_{z, \mathbf{x} \in \{0,1\}^*} \equiv \left\{ \text{REAL}_{\pi^\rho, \mathcal{A}(z)}(\mathbf{x}) \right\}_{z, \mathbf{x} \in \{0,1\}^*}. \quad (1)$$

⁹If \mathcal{S}_ρ were to rewind \mathcal{A} , then a problem would arise in the case that a π -message that \mathcal{A} sent before rewinding is changed by \mathcal{A} after rewinding. In particular, since these messages are sent *externally* to honest parties, they cannot be changed at a later stage.

Assume, by contradiction, that this is not the case. This implies that there exists a protocol π , an adversary \mathcal{A} and a vector of inputs \mathbf{x} such that Eq. (1) does not hold.

First, assume that the output of each honest party in π includes its local input to and output from the ideal call to f (or likewise, its local input to and output from the execution of ρ). Clearly Eq. (1) does not hold here as well, because we only provided more output. Next, modify the protocol π to π' so that the honest parties output their local input to and output from f (or ρ) only, and not their original π -output. This is the only difference between π and π' . In particular, the honest parties run π to the end as before, following the same instructions. We now claim that Eq. (1) also does not hold with respect to π' . That is,

$$\left\{ \text{HYBRID}_{\pi', \mathcal{S}^{\mathcal{A}(z)}}^f(\mathbf{x}) \right\} \not\equiv \left\{ \text{REAL}_{\pi', \mathcal{A}(z)}(\mathbf{x}) \right\} . \quad (2)$$

This requires justification because in π' the output distribution is different (and contains less “information”). In order to see that Eq. (2) holds, we show that if this were not the case, then Eq. (1) would hold. This follows from the fact that the only difference between $\text{HYBRID}_{\pi, \mathcal{S}}^f(\mathbf{x})$ and $\text{REAL}_{\pi, \mathcal{A}}(\mathbf{x})$ for the honest parties is that the output from f is obtained from the trusted party in HYBRID and from the execution of ρ in REAL. Furthermore, protocol π only refers to the input to ρ and output from ρ and thus only the input to and output from ρ influence the honest parties’ output in π . This implies that if the *joint input/output distribution* to ρ are identically distributed in HYBRID and REAL, then the honest parties’ outputs from π are identically distributed in HYBRID and REAL.¹⁰ One subtle point that needs to be clarified is that the output of the honest parties in π is also influenced by the messages that they receive from the adversary. However, by the definition of \mathcal{S} , the messages received by the honest parties *within the execution of π* are identical in the HYBRID and REAL executions; indeed, \mathcal{S} merely forwards these messages unmodified between \mathcal{A} and the honest parties. This contradicts the assumption that Eq. (1) does not hold, and so we conclude that Eq. (2) holds.

Now, let \mathbf{x}_ρ be a vector of inputs to ρ such that in the event that the inputs determined by π' to f (or equivalently to ρ) equal \mathbf{x}_ρ , the $\text{HYBRID}_{\pi', \mathcal{S}}^f(\mathbf{x})$ and $\text{REAL}_{\pi', \mathcal{A}}(\mathbf{x})$ distributions are not identical. Such a vector of inputs \mathbf{x}_ρ to ρ must exist, because otherwise the HYBRID and REAL distributions would be identical. (More formally, we can break up the distributions of $\text{HYBRID}_{\pi', \mathcal{S}}^f(\mathbf{x})$ and $\text{REAL}_{\pi', \mathcal{A}}(\mathbf{x})$ according to all possible sets of inputs to f and ρ in an execution of π' with \mathbf{x} . Then, it must be that for at least one of these “sub-distributions”, the HYBRID and REAL distributions are not identical.)

We are now ready to construct an adversary \mathcal{A}_ρ that attacks a *stand-alone execution of ρ* , and succeeds when the input vector to ρ equals \mathbf{x}_ρ . Adversary \mathcal{A}_ρ works by internally emulating all the honest parties in the execution of π' , and externally running the ρ execution. Specifically, \mathcal{A}_ρ sets the inputs of the (internally-emulated) parties running π' to \mathbf{x} , and also chooses uniformly distributed random-tapes for these parties. Then, \mathcal{A}_ρ invokes \mathcal{A} and perfectly emulates an execution of $\text{REAL}_{\pi', \mathcal{A}}(\mathbf{x})$, by passing all π' -messages between \mathcal{A} and the internally-emulated honest parties. In contrast, when the execution of ρ is reached, \mathcal{A}_ρ sends \mathcal{A} ’s ρ -messages externally to the real honest parties running ρ . Likewise, ρ -messages received externally by \mathcal{A}_ρ are internally handed back to \mathcal{A} . We note that the internal emulation of π continues concurrently to the external execution of ρ (according to the scheduling determined by \mathcal{A}). At the end of the external execution of ρ , adversary \mathcal{A}_ρ guesses the outputs of the honest parties (from the set of all possible outputs) and uses them in the continuation of the internal emulation of π .¹¹ \mathcal{A}_ρ continues this internal emulation of π' until \mathcal{A} halts. At this point

¹⁰This is not necessarily true for the adversary’s output. However, the adversary’s output is unchanged in π and π' .

¹¹This guess can be made efficient by guessing the honest parties’ inputs and random tapes. In the case that the messages sent by the real honest parties are consistent with the guessed inputs and random tapes, \mathcal{A}_ρ sets their outputs according to these values and the messages sent by \mathcal{A} to the honest parties (if the guess of \mathcal{A}_ρ is correct here, it follows that \mathcal{A}_ρ holds the view of the honest parties and so can compute their exact outputs). Otherwise, \mathcal{A}_ρ sets their outputs

\mathcal{A}_ρ outputs \mathcal{A} 's output, as well as its guesses for the honest parties' inputs to ρ and their outputs. \mathcal{A}_ρ then halts.

First, observe that with nonzero probability, the inputs of the internally-emulated honest parties to ρ equal \mathbf{x}_ρ (and therefore match the real inputs of the external parties). This must hold because otherwise \mathbf{x}_ρ would not have been chosen above (an \mathbf{x}_ρ which appears with zero probability cannot contribute to a statistical difference between the REAL and HYBRID distributions). Furthermore, the guess of \mathcal{A}_ρ with respect to the honest parties' outputs from ρ is also correct with nonzero probability. Now, conditioned on these guesses being correct, the output of the honest parties from ρ along with the portion of \mathcal{A}_ρ 's output that is copied from \mathcal{A} 's output is distributed exactly according to $\text{REAL}_{\pi', \mathcal{A}(z)}(\mathbf{x})$. Furthermore, in an ideal-model simulation by \mathcal{S}_ρ , once again conditioned on the guesses being correct, the output of the (ideal) honest parties along with the appropriate portion of \mathcal{S}_ρ 's output is distributed exactly according to $\text{HYBRID}_{\pi', \mathcal{S}\mathcal{A}(z)}^f(\mathbf{x})$.

Notice finally that because \mathcal{A}_ρ outputs its guesses for the inputs and outputs of the honest parties (and whether or not they are correct can be derived from comparing \mathcal{A}_ρ 's output to the real inputs and outputs of the honest parties), the “sub-distribution” of the guesses of \mathcal{A}_ρ being correct is disjoint from the case that the guesses of \mathcal{A}_ρ are incorrect. Since it follows from the contradicting assumption that this sub-distribution is not identically distributed in the real and ideal executions, we conclude that

$$\left\{ \text{IDEAL}_{\mathcal{S}_\rho^{\mathcal{A}_\rho(z)}}^f(\mathbf{x}_\rho) \right\} \not\equiv \left\{ \text{REAL}_{\rho, \mathcal{A}_\rho(z)}(\mathbf{x}_\rho) \right\}$$

in contradiction to the perfect security of ρ with the black-box simulator \mathcal{S}_ρ .

The “furthermore” in the theorem statement follows from the fact that the simulator \mathcal{S} that we constructed merely invokes \mathcal{A} and the simulator \mathcal{S}_ρ for ρ . Thus, if \mathcal{S}_ρ runs in time that is polynomial in the running-time of \mathcal{A} , so does \mathcal{S} . This completes the proof. ■

The computational setting. In the information-theoretic setting, ideal secure channels are assumed to exist. However, in the computational model only authenticated channels are used. Thus, information-theoretic protocols cannot be used directly in the computational setting. Fortunately, this is easy to fix; it suffices to use computationally secure channels as formulated in [CK02] (these can be constructed from any CCA2-secure public-key encryption scheme, or any public-key encryption scheme when assuming authenticated channels. We denote by ρ' the protocol which is identical to ρ except that all messages are sent encrypted, according to the scheme shown in [CK02]. This yields the following corollary:

Corollary 4 *Assume the existence of public-key encryption schemes. Let f be a function and let ρ be a protocol that computes f with perfect security and abort in the stand-alone model and has a black-box straight-line simulator. Then, protocol ρ' defined above securely computes f under concurrent general composition in the computational setting.*

6 Straight-Line Black-Box Simulation and Start Synchronization

Theorem 2 states that statistical security and black-box straight-line simulation do not suffice for achieving security under concurrent general composition. In this section, we show that a mild additional requirement can be made on the protocol so that it does suffice. Namely, we prove that if a secure protocol has a black-box straight-line simulator, and no party sends any message that depends

to some default values. Notice, that if \mathcal{A}_ρ 's input and random-tape guesses are correct, then so are the output guesses. Furthermore, the input and random-tape guesses are correct with nonzero probability.

on its input or random-tape until all parties have announced that they have started, then the protocol *is* secure under concurrent general composition. This result holds for all levels of security, even computational security.

Start synchronization. We say that a protocol ρ has start synchronization if it begins with the following steps for *all* parties:

1. Send “begin ρ ” to all parties
2. Wait until “begin ρ ” messages are received from *all* parties. Then, continue the execution of ρ .

We stress that the inputs to the protocol are read by each party before it sends its begin ρ message, and are therefore fixed before the actual execution of ρ begins (or, more exactly, before any honest party sends a message that depends on its input or random tape). We prove the following theorem.

Theorem 5 *Let f be a function and let ρ be a protocol that computes f with computational (resp. statistical or perfect) security and abort in the stand-alone model, and with a black-box straight-line simulator. Furthermore, assume that ρ has start synchronization. Then, ρ computes f with computational (resp. statistical or perfect) security under concurrent general composition. This holds for both static and adaptive adversaries.*

Proof: The case of perfect security follows from Theorem 3, even without start synchronization. It therefore remains to prove the theorem for the computational and statistical cases. We present the proof for the computational setting; only minor modifications are needed for the statistical case. Also, as in the proof of Theorem 3, we provide the proof only for the case of static adversaries; dealing with adaptive corruptions here is straightforward. The intuition behind this theorem comes from the counter-example of Section 4. Specifically, notice that the security of the protocol there is compromised by having one party’s input depend on messages sent by the other parties in ρ itself. When there is start synchronization, this cannot happen. Technically, the proof utilizes start synchronization by the following observation. Assume that all of the parties’ inputs to ρ in a concurrent execution with some protocol π , depend only on π (as is indeed the case when there is start synchronization). Then, all of the computation of π can be thrown into an auxiliary input for an adversary attacking a stand-alone execution of ρ . Thus, we can use the stand-alone security of ρ to derive security under concurrent general composition. (Note that an auxiliary input can be correlated to the parties’ inputs to ρ but not to their random tapes in ρ . Thus, this strategy cannot work if there is no start synchronization, as is demonstrated in the counter-example of Section 4.)

We begin the proof in a very similar fashion to the proof of Theorem 3. That is, we use the simulator \mathcal{S}_ρ that exists for ρ (by the assumption that it is secure in the stand-alone model) in order to construct a simulator \mathcal{S} for the setting of concurrent general composition. As in the proof of Theorem 3, the simulator \mathcal{S} works by forwarding all of the π -messages untouched (where π is an arbitrary protocol that contains an ideal call to f), and using \mathcal{S}_ρ to deal with all of the ρ -messages. More specifically, let f be a function, let ρ be a protocol that computes f with computational security, and let \mathcal{S}_ρ be the black-box straight-line ideal-model simulator for ρ that is assumed to exist. Then, we construct a simulator \mathcal{S} for the setting of concurrent general composition as follows. Let \mathcal{A} be a real-model adversary and let π be an arbitrary protocol that contains a single ideal call to f (as discussed in the proof of Theorem 3 considering a single call to f suffices). Then, \mathcal{S} invokes the real adversary \mathcal{A} and forwards all π -messages untouched between \mathcal{A} and the honest parties (recall that \mathcal{S} runs a real execution of π with external parties). In contrast, when \mathcal{A} begins the ρ execution, \mathcal{S} invokes a copy of \mathcal{S}_ρ and forwards all ρ -messages between \mathcal{A} and \mathcal{S}_ρ (in contrast to the π -messages that are sent externally, ρ is internally simulated by \mathcal{S} using \mathcal{S}_ρ). More specifically, when \mathcal{A} outputs a ρ -message

to be sent to an honest party, \mathcal{S} hands this message to \mathcal{S}_ρ . Likewise, \mathcal{S}_ρ 's reply is handed back to \mathcal{A} as if coming from an honest party. In addition, whenever \mathcal{S}_ρ sends an input to the trusted party computing f , simulator \mathcal{S} sends the same input. Likewise, outputs sent to \mathcal{S} from the trusted party are forwarded to \mathcal{S}_ρ . We note that the π execution continues by just forwarding π -messages between \mathcal{A} and the honest parties, even while the simulation of ρ takes place. Notice that the real execution of π can continue concurrently with the simulation of ρ because \mathcal{S}_ρ is *straight-line*, and so never needs to rewind (this is where the straight-line property of \mathcal{S}_ρ is used). Whenever \mathcal{A} halts, simulator \mathcal{S} copies \mathcal{A} 's output to its own output tape and halts.

Let k be the security parameter (i.e., all parties are assumed to run in time that is polynomial in k). Recall that we denote by $\text{REAL}_{\pi^\rho, \mathcal{A}(z)}(k, \mathbf{x})$ a real execution of the protocol π with real protocol ρ , where \mathcal{A} has auxiliary input z and the vector of the parties' inputs equals \mathbf{x} . Likewise, we denote by $\text{HYBRID}_{\pi, \mathcal{S}^{\mathcal{A}(z)}}^f(k, \mathbf{x})$ the hybrid execution of π with ideal calls to f , where the parties' inputs are \mathbf{x} and the ideal-world adversary/simulator is given oracle (i.e., black-box) access to $\mathcal{A}(z)$.

We now prove that for every arbitrary protocol π (with a single ideal call to f) and every real adversary \mathcal{A} ,

$$\left\{ \text{HYBRID}_{\pi, \mathcal{S}^{\mathcal{A}(z)}}^f(k, \mathbf{x}) \right\}_{\mathbf{x}, z \in \{0,1\}^*; k \in \mathbb{N}} \stackrel{c}{=} \left\{ \text{REAL}_{\pi^\rho, \mathcal{A}(z)}(k, \mathbf{x}) \right\}_{\mathbf{x}, z \in \{0,1\}^*; k \in \mathbb{N}} \quad (3)$$

Assume, by contradiction, that this is not the case. This implies that there exists a protocol π , a distinguisher D , an adversary \mathcal{A} , a polynomial p and a vector of inputs \mathbf{x} and auxiliary input z such that for infinitely many k 's,

$$\left| \Pr \left[D(\text{HYBRID}_{\pi, \mathcal{S}^{\mathcal{A}(z)}}^f(k, \mathbf{x})) = 1 \right] - \Pr \left[D(\text{REAL}_{\pi^\rho, \mathcal{A}(z)}(k, \mathbf{x})) \right] \right| > \frac{1}{p(k)}$$

We use this to contradict the security of ρ . That is, we show that there exists a distinguisher D_ρ , real-world adversary \mathcal{A}_ρ , polynomial q and an auxiliary input z_ρ and vector of inputs \mathbf{x}_ρ such that for infinitely many k 's,

$$\left| \Pr \left[D_\rho(\text{IDEAL}_{\mathcal{S}_\rho^{\mathcal{A}_\rho(z_\rho)}}^f(k, \mathbf{x}_\rho)) = 1 \right] - \Pr \left[D(\text{REAL}_{\rho, \mathcal{A}_\rho(z_\rho)}(k, \mathbf{x}_\rho)) \right] \right| > \frac{1}{q(k)} \quad (4)$$

in contradiction to the assumption that \mathcal{S}_ρ is a "good simulator".

As we have mentioned above, the key point in the proof is to use the auxiliary input z_ρ for \mathcal{A}_ρ in a stand-alone execution of ρ in place of the execution of π together with ρ . Specifically, a secure protocol must remain secure when the adversary receives any auxiliary input, even one that is correlated with the honest parties' inputs. In particular, for every vector of inputs \mathbf{x}_ρ to ρ , security must hold when the auxiliary input z_ρ consists of a vector of inputs \mathbf{x} and random coins \mathbf{r} for the honest parties (and the adversary) such that when π is run with the adversary \mathcal{A} , inputs \mathbf{x} and random tapes \mathbf{r} , the honest parties' inputs to f (or equivalently to ρ) are those given in \mathbf{x}_ρ . (Note that the random tapes \mathbf{r} in z_ρ are actually chosen uniformly from all those meeting the requirement.) Before proceeding, we stress that defining such an auxiliary input is only possible if the inputs of the parties in ρ are *independent* of the ρ -execution itself. This is ensured because ρ begins with start synchronization.

Next, we describe an adversary \mathcal{A}_ρ attacking ρ . Adversary \mathcal{A}_ρ internally invokes \mathcal{A} and perfectly simulates the honest parties' actions in π by running the π -instructions of the honest parties on the inputs and random-tapes specified in its auxiliary input z_ρ . In contrast to the π -messages that are internally emulated by \mathcal{A}_ρ for \mathcal{A} , all of the messages belonging to ρ are forwarded by \mathcal{A}_ρ from \mathcal{A} to the external honest parties and back. At the end of the execution of ρ , adversary \mathcal{A}_ρ halts, outputting \mathcal{A} 's current view.

Given the above, we show that by the contradicting assumption, there exists at least one vector of inputs \mathbf{x}_ρ to ρ with associated auxiliary input z_ρ such that a real execution of ρ with adversary \mathcal{A}_ρ , auxiliary input z_ρ and inputs \mathbf{x}_ρ , can be distinguished from an ideal execution of f with adversary/simulator $\mathcal{S}_\rho^{\mathcal{A}_\rho(z_\rho)}$ and inputs \mathbf{x}_ρ . We prove this by showing that if this is not the case (i.e., if Eq. (4) does *not* hold), then this would imply that Eq. (3) holds (in contradiction to the assumption). To see this, notice that an efficient distinguisher D_ρ who receives the adversary's view after the execution of ρ is completed (including its outputs from ρ or f), the auxiliary input z_ρ of the adversary, and the honest parties' outputs from ρ or f , can complete the execution of π itself and obtain the outputs of \mathcal{A} and all the honest parties. Now, if D_ρ receives outputs from a real execution of ρ , the output distribution that it will internally compute will be exactly $\text{REAL}_{\pi, \mathcal{A}(z)}(k, \mathbf{x})$. In contrast, if D_ρ receives outputs from an ideal execution of f , the output distribution that it will internally compute will be exactly $\text{HYBRID}_{\pi, \mathcal{S}^{\mathcal{A}(z)}}^f(k, \mathbf{x})$. (This holds due to the construction of \mathcal{S} that runs \mathcal{S}_ρ for the ρ part of the execution, and everything else is run according to \mathcal{A} .) We conclude that if Eq. (3) does not hold, then there exist \mathbf{x}_ρ and z_ρ such that D_ρ can distinguish between IDEAL and REAL as in Eq. (4). This contradicts the assumed security of ρ . ■

7 Black-Box Straight-Line Simulation and Fixed Inputs

Theorem 2 states that black-box straight-line simulation does not suffice for achieving security under concurrent general composition. In this section, we show that it does however suffice for achieving security under concurrent general composition with *fixed inputs*. Recall that in the setting of *fixed inputs*, each party P_i receives a pair of inputs (x_i, y_i) before the execution begins. Then, party P_i uses x_i as its input to π and it uses y_i as its input to ρ .

Theorem 6 *Let f be a probabilistic polynomial-time function and let ρ be a protocol that computes f with computational (resp., statistical/perfect) security and abort in the stand-alone model under black-box straight-line simulation. Then, ρ computes f with computational (resp., statistical/perfect) security under concurrent general-composition with fixed inputs.*

The proof of this theorem is a simple, special case of the proof of Theorem 5 where instead of using the auxiliary input to fix the inputs to ρ , they are already fixed ahead of time.

Acknowledgements

We would like to thank the anonymous referees for their very helpful comments.

References

- [BMU07] M. Backes, J. Müller-Quade and D. Unruh. On the Necessity of Rewinding in Secure Multiparty Computation. In the *4th TCC*, Springer-Verlag (LNCS 4392), pages 157–173, 2007.
- [Bea89] D. Beaver. Multiparty Protocols Tolerating Half Faulty Processors. In *CRYPTO'89*, Springer-Verlag (LNCS 435), pages 560–572, 1990.
- [Bea91] D. Beaver. Foundations of Secure Interactive Computing. In *CRYPTO'91*, Springer-Verlag (LNCS 576), pages 377–391, 1991.

- [BCG93] M. Ben-Or, R. Canetti and O. Goldreich. Asynchronous Secure Computation. In the *25th STOC*, pages 52–61, 1993.
- [BGW88] M. Ben-Or, S. Goldwasser, and A. Wigderson. Completeness Theorems for Noncryptographic Fault-Tolerant Distributed Computations. In the *20th STOC*, pages 1–10, 1988.
- [BKR94] M. Ben-Or, B. Kelmer and T. Rabin. Asynchronous Secure Computations with Optimal Resilience. In *13th PODC*, pages 183–192, 1994.
- [B84] G. Bracha. An Asynchronous $(n - 1)/3$ -Resilient Consensus Protocol. In the *3rd PODC*, pages 154–162, 1984.
- [Can00] R. Canetti. Security and Composition of Multiparty Cryptographic Protocols. *Journal of Cryptology*, 13(1):143–202, 2000.
- [Can01] R. Canetti. Universally Composable Security: A New Paradigm for Cryptographic Protocols. In the *42nd FOCS*, pages 136–145, 2001. The full version can be found at <http://eprint.iacr.org/2000/067>.
- [CK02] R. Canetti and H. Krawczyk. Universally Composable Notions of Key Exchange and Secure Channels. In *EUROCRYPT 2002*, Springer-Verlag (LNCS 2332), pages 337–351, 2002.
- [CCD88] D. Chaum, C. Crépeau, and I. Damgård. Multiparty Unconditionally Secure Protocols. In the *20th STOC*, pages 11–19, 1988.
- [CDD⁺99] R. Cramer, I. Damgård, S. Dziembowski, M. Hirt, and T. Rabin. Efficient Multiparty Computations with Dishonest Minority. In *EUROCRYPT'99*, Springer-Verlag (LNCS 1592), pages 311–326, 1999.
- [DDN00] D. Dolev, C. Dwork, and M. Naor. Non-malleable Cryptography. *SIAM Journal on Computing*, 30(2):391–437, 2000.
- [DM00] Y. Dodis and S. Micali. Parallel Reducibility for Information-Theoretically Secure Computation. In *CRYPTO 2000*, Springer-Verlag (LNCS 1880), pages 74–92, 2000.
- [DNS04] C. Dwork, M. Naor, and A. Sahai. Concurrent zero-knowledge. *Journal of the ACM*, 51(6):851–898, 2004.
- [Fei90] U. Feige. Alternative Models for Zero Knowledge Interactive Proofs. PhD thesis, Weizmann Institute of Science, 1990. Available from <http://www.wisdom.weizmann.ac.il/~feige>.
- [F89] P. Feldman. Asynchronous Byzantine Agreement in Constant Expected Time. Unpublished manuscript, 1989.
- [GMW87] O. Goldreich, S. Micali, and A. Wigderson. How to Play Any Mental Game. Information-Theoretically Secure Protocols and Security under Composition. In *Proc. 36th STOC*, pages 109–118. ACM, 2006.
- [Gol01] O. Goldreich. *Foundations of Cryptography – Volume 1*. Cambridge University Press, 2001.
- [Gol04] O. Goldreich. *Foundations of Cryptography – Volume 2*. Cambridge University Press, 2004.
- [GL90] S. Goldwasser and L. Levin. Fair Computation of General Functions in Presence of Immoral Majority. In *CRYPTO'90*, Springer-Verlag (LNCS 537), pages 77–93, 1990.

- [KLR06] E. Kushilevitz, Y. Lindell, and T. Rabin. Information-Theoretically Secure Protocols and Security under Composition. In the *36th STOC*, pages 109–118, 2006.
- [HU06] D. Hofheinz and D. Unruh. Simulatable Security and Polynomially Bounded Concurrent Composability. In the *2006 IEEE Symposium on Security and Privacy*, pages 169–183, 2006.
- [Lin03a] Y. Lindell. Bounded-Concurrent Secure Two-Party Computation Without Setup Assumptions. In the *35th STOC*, pages 683–692, 2003.
- [Lin03b] Y. Lindell. General Composition and Universal Composability in Secure Multi-Party Computation. In the *Journal of Cryptology*, 22(3):395-428, 2009.
- [Lin04] Y. Lindell. Lower Bounds for Concurrent Self Composition. In the *Journal of Cryptology*, 21(2):200–249, 2008.
- [MR91] S. Micali and P. Rogaway. Secure Computation. In *CRYPTO'91*, Springer-Verlag (LNCS 576), pages 392–404, 1991.
- [PW00] B. Pfitzmann and M. Waidner. Composition and Integrity Preservation of Secure Reactive Systems. In *7th ACM Conference on Computer and Communication Security*, pages 245–254, 2000.
- [RB89] T. Rabin and M. Ben-Or. Verifiable Secret Sharing and Multiparty Protocols with Honest Majority. In the *21st STOC*, pages 73–85, 1989.
- [Yao86] A. C-C. Yao. How to Generate and Exchange Secrets. In the *27th FOCS*, pages 162–167, 1986.