Concurrent Composition of Secure Protocols in the Timing Model

Yael Tauman Kalai[†][*]
M.I.T
yael@csail.mit.edu

Yehuda Lindell[†]
Dept. of Computer Science
Bar-Ilan University, ISRAEL
lindell@cs.biu.ac.il

Manoj Prabhakaran[†]
University of Illinois at Urbana-Champaign
mmp@cs.uiuc.edu

July 16, 2006

Abstract

In the setting of secure multiparty computation, a set of mutually distrustful parties wish to securely compute some joint function of their inputs. In the stand-alone case, it has been shown that *every* efficient function can be securely computed. However, in the setting of concurrent composition, broad impossibility results have been proven for the case of no honest majority and no trusted setup phase. These results hold both for the case of general composition (where a secure protocol is run many times concurrently with arbitrary other protocols) and self composition (where a single secure protocol is run many times concurrently).

In this paper, we investigate the feasibility of obtaining security in the concurrent setting, assuming that each party has a local clock and that these clocks proceed at approximately the same rate. We show that under this mild timing assumption, it is possible to securely compute *any* multiparty functionality under concurrent *self* composition. Loosely speaking, we also show that it is possible to securely compute *any* multiparty functionality under concurrent *general* composition, as long as the secure protocol is run only with protocols whose messages are delayed by a specified amount of time. On the negative side, we show that it is impossible to achieve security under concurrent general composition with no restrictions whatsoever on the network (like the aforementioned delays), even in the timing model.

Keywords: theory of cryptography, secure multiparty computation, concurrent composition, timing assumptions.

[*]Supported in part by NSF CyberTrust grant CNS-0430450
[†]Part of this work was carried out while the authors were all at IBM T.J.Watson Research, New York.

1 Introduction

1.1 Background

In the setting of secure multiparty computation, a set of parties with private inputs wish to jointly compute some functionality 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 an incorruptible trusted party who honestly computes the function for the parties (cf. [11] or [28, 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. Powerful feasibility results have been shown for this problem in both the information-theoretic and computational settings [48, 30, 8, 17]. In the computational setting, it has been shown that *any* multiparty probabilistic polynomial-time functionality can be securely computed for any number of corrupted parties, assuming the existence of enhanced trapdoor permutations [48, 30, 28].

Security under concurrent composition. The above-described feasibility results relate only to the stand-alone setting, where a single protocol is run in isolation. However, in modern network settings, protocols must remain secure even when many protocol executions take place concurrently and are being attacked in a coordinated manner. Unfortunately, the security of a protocol in the stand-alone setting does not necessarily imply its security under concurrent composition. Therefore, an important research goal is to re-establish the feasibility results of the stand-alone setting for the setting of concurrent composition. There are two main types of concurrent composition that have been considered:

1. *Concurrent self composition:* In this setting, a single protocol is executed many times concurrently in a network. Formally, "concurrency" means that the adversary has full control over the scheduling of all messages in all the executions.

2. *Concurrent general composition:* In this setting, a protocol is run many times in an arbitrary network. That is, the protocol is run many times concurrently, alongside other secure and insecure protocols, again with the scheduling being fully controlled by the adversary.

On the positive side, it has been shown that in the case of an honest majority, or a trusted setup phase (e.g., for generating a common reference string or for generating a secure public-key infrastructure), any functionality can be securely computed under concurrent general composition [12, 16, 3]. Thus, in these cases, we obtain the same broad feasibility results of the stand-alone model (except that in the stand-alone model, neither an honest majority nor a trusted setup phase is needed).

When considering the case of *no honest majority* and *no trusted setup* in the setting of concurrent composition, the situation is completely different. Recent impossibility results have demonstrated that in such a setting, large classes of functionalities cannot be securely computed [14, 12, 15, 37, 38]. These results hold for both concurrent general composition *and* concurrent self composition. In fact, these two types of composition have been shown to be (almost) equivalent [38]. Therefore, in the natural setting of no trusted setup and no honest majority (including the important *two-party case*), it is impossible to construct protocols that remain secure in the setting of full concurrency.

There are a number of possible ways to overcome these impossibility results. One direction is to weaken the security requirements; this approach was taken by [42, 46]. Another direction, and the one taken in this paper, is to introduce realistic assumptions on the adversary or network, while providing the same strong security guarantees as for the stand-alone setting. Needless to say, it is best to not assume any restriction whatsoever. However, as we have mentioned, this is not possible. We therefore consider a very reasonable network restriction that holds in real networks today.

Timing assumptions. The network restriction that we consider is a *timing* assumption on the network. Timing assumptions were first used in the context of secure protocol composition by [22] who used them to achieve (efficient) zero-knowledge protocols that remain secure under concurrent self composition. (An equivalent formulation of these assumptions was given in [29], and our presentation is more according to this latter formulation.) There are two specific assumptions involved here:

- *Assumption 1 – bounded clock drift:* It is assumed that the parties' local clocks proceed at approximately the same rate. Specifically, there exists a *global* bound $\epsilon \geq 1$ such that when one local clock advances t time units, every other local clock advances t' time units where $t/\epsilon \leq t' \leq t\epsilon$. We stress that there is *no assumption* regarding the synchronization of the parties' local clocks with respect to each other (and, in particular, they may read completely different times).
- *Assumption 2 – maximum latency:* It is assumed that an upper bound Δ is known on the time it takes for a message to be computed, sent and delivered from one party to another. In other words, Δ is the *maximum latency* over the network (plus the time it takes to carry out the local computation for generating the message that is sent). For simplicity, we assume that all local computation is instantaneous, and that Δ measures the latency only (or, in other words, the time that it takes for the adversary to deliver messages).

The second of these two assumptions is far more problematic than the first. This is due to the fact that in real settings, the variance of network latency can be very large. Thus, a global upper bound would have to be very large. As we will see, taking such a high upper bound would greatly hinder performance. In addition, any reasonable bound is unlikely to always hold, thus potentially compromising the security of the protocol. In contrast, local clocks are usually very accurate, at least with respect to the drift.

Motivated by this observation, we relate to these assumptions differently. More specifically, our definition of *security* for the timing model relies *only* on the first assumption regarding the clock drift. Therefore, security holds as long as the drifts of the clocks are not too far apart, and *irrespective of the network latency* (which may, however, cause the execution to terminate unsuccessfully). The latency assumption is only used to ensure *liveness* (or non-triviality of protocols). Namely, we only require that the protocol terminates successfully if it does not come under attack and the latency is indeed lower than Δ.

The use of timing assumptions. As in other works, the timing assumptions are used for introducing time-out and delay operations in the protocol instructions. A time-out command is of the form: "if more than $f(\Delta, \epsilon)$ time units have passed since message x was sent (or received), and message y has not yet been received, then output time-out and halt the execution" (where f is a function specified by the protocol). A delay command is of the form: "wait $g(\Delta, \epsilon)$ time units before sending message y". Typically, the use of these operations is to limit the interleaving of different protocol executions. Specifically, delay and time-out commands are used to ensure properties of the form: if part A of execution i begins after part B of execution j begins, then part B of execution j is completed before part A of execution i is completed. This is achieved by timing-out if B takes too long and delaying to makes sure that A takes long enough, as depicted in Figure 1.

The differences in the lengths of part A and part B in the different executions shown in Figure 1 are due to the control that the adversary has over message delivery. We stress that the time-out and delay instructions depend on the parties' local clocks only, and so do not rely on any global synchronization.

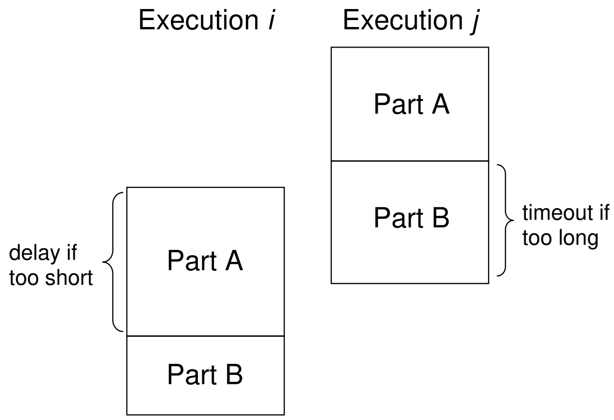

Figure 1: limiting the interleaving (notice that part A must take longer than part B).

Limiting time-out damage. As we have described, `time-out` instructions are used in order to limit the adversary's power in scheduling the executions. However, if the network latency during a protocol execution is higher than usual (say, due to high network traffic), then a `time-out` can occur even when no adversary is present.[1] This therefore raises the question of what actions honest parties should take after a `time-out` occurs. In particular, can a timed-out protocol execution be safely restarted? Fortunately, our protocols have the property that they remain secure if they are restarted from scratch after a time-out.

In order to further clarify this point, we distinguish between two types of failures: `abort` failures that occur due to foul play by participating parties (and are present even in the stand-alone case when there is no honest majority), and `time-out` failures that occur due to high network latency (or by the adversary stalling or blocking messages that are sent). In the case of an *abort failure* (again, even in the stand-alone case), security is *not* guaranteed if the honest parties restart the protocol execution. This is due to the fact that the adversary may have received its own output, and based on this output has decided to cause the honest parties to abort.[2] In contrast, we argue that it should be possible to restart protocol executions that are halted due to *time-out failures*. This is due to the fact that such a failure can occur even if there is no adversarial interference, and just due to the network latency being high at the time of the execution. In order to ensure that timed-out executions can be restarted without any damage to security, our definition of security requires that time-outs (by honest participants) are only allowed to occur in an early stage of the protocol, before any information about the output is revealed. We note that since timed-out protocols can be restarted safely, a relatively optimistic estimate on the network latency can be taken with the cost being that timed-out protocol executions are simply restarted. (There is a tradeoff here between choosing a large Δ that slows down all protocol executions and choosing a

[1] Of course, one could set Δ to be an upper bound that includes latencies that are far higher than the average. However, as we will see, this would have the effect of significantly delaying all executions.

[2] For the sake of concreteness, consider the case that parties run a coin-tossing protocol. Then, the first party to receive output can cause an abort if the first bit of the output is 1. By re-executing upon abort, this party can bias the outcome so that the resulting string always has the first bit set to 0.

small Δ that results in more executions being timed-out and restarted.) We remark that previous works that used timing assumptions considered only the problem of concurrent zero-knowledge, where essentially no output is generated. The above discussion is therefore a "non-issue" in that case, and protocol executions can always be restarted.

1.2 Our Main Results

We investigate the feasibility of constructing protocols that are secure under concurrent composition in the timing model. We consider both self and general composition, with the following main results:

1. Concurrent self composition: We show that in the timing model, every multiparty functionality $\mathcal{F}$ can be securely computed under concurrent self composition. We note that the model of concurrent self composition that we consider here is one *without* fixed roles. Thus, for example, in the setting of zero-knowledge, parties can play both the prover and verifier roles simultaneously. This is the first zero-knowledge protocol (for the setting of unbounded concurrency) that has this property.

2. Concurrent general composition: For this setting, we have both positive and negative results:

   (a) *Positive result:* Loosely speaking, we show that in the timing model, every multiparty functionality $\mathcal{F}$ can be securely computed under concurrent general composition, as long as the arbitrary protocols that are running in the network all have the property that their messages are delayed by some specified amount of time (the exact delay required is specified in Theorem 6 and is independent of $\mathcal{F}$). We stress that our protocol is only proven secure in the case that the arbitrary protocols running together with it do not use time (beyond the delays that we impose).

   (b) *Negative result:* We show that it is impossible to achieve security under concurrent general composition if no restrictions (like delays) are imposed on the arbitrary protocols running in the network.

Positive results. We now elaborate on our main positive results. We note that all of our results relate to the setting of no honest majority. Therefore, our definition of security does not guarantee fairness. That is, the adversary may receive output while the honest parties do not. This is standard for the case of no honest majority, even in the stand-alone model. (Note that our protocols do not guarantee output delivery even in the event that an honest majority does happen to be present. Again, this is standard for protocols that achieve security against any number of corruptions.) Our result regarding concurrent self composition is *informally* stated as follows.

Theorem 1.1 *Assume that there exist enhanced trapdoor permutations and dense cryptosystems.*[3] *Then, in the timing model, any multiparty functionality $\mathcal{F}$ can be securely computed under concurrent self composition, in the presence of static adversaries.*

The proof of Theorem 1.1 gives the first construction of a protocol that achieves security (in the standard sense) under concurrent self composition, without relying on a trusted setup phase, an

[3]An *enhanced* trapdoor permutation has the property that it is hard to invert even given the coins used to sample the value from the range; see [28, Appendix C.1]. A dense cryptosystem is one for which uniform strings are valid public-keys with noticeable probability; see [19]. We inherit these assumptions from [16].

assumed honest majority, or an assumed a-priori bound on the number of executions taking place concurrently. As we have mentioned above, our definition for concurrent self composition makes no limitation on the roles played by the parties. Thus, Theorem 1.1 implies the existence of a concurrent zero-knowledge protocol (in the timing model) that remains secure even when the adversary can carry out a concurrent man-in-the-middle attack (i.e., where the adversary may simultaneously verify and prove many concurrent proofs). This is the first zero-knowledge protocol with this property. (We note that an analogous result is known in a setting where the number of concurrent executions is a priori bounded [43].) We also note that in fact, we achieve a stronger result than that stated in Theorem 1.1. Namely, any set of protocols that are constructed according to the methodology presented in this paper remain secure when concurrently composed together.

In order to state our positive result for concurrent general composition in more detail, we introduce the following terminology. A protocol that does not use any timing instruction (and in particular never refers to a clock) is called timing-free. We say that a (timing-free) protocol π is δ-delayed if every message in π is delayed by at least δ time units before it is sent. We stress that the contents of the messages specified by π are untouched. Furthermore, there is no assumption regarding the delaying of messages by corrupted parties. Rather, security is guaranteed as long (*and only as long*) as the *honest parties* delay all π-messages, as instructed. We now informally state our main result:

Theorem 1.2 (main theorem – informal): *Assume that there exist enhanced trapdoor permutations and dense cryptosystems, and let $\mathcal{F}_1, \ldots, \mathcal{F}_t$ be any set of multiparty functionalities. Then, there exist protocols $\rho_1, \ldots, \rho_t$ in the timing model that securely compute $\mathcal{F}_1, \ldots, \mathcal{F}_t$ under concurrent general composition* (*in the presence of static adversaries*), *as long as the arbitrary protocols running concurrently in the network together with $\rho_1, \ldots, \rho_t$ are δ-delayed, for some parameter δ.*

Another equivalent way of stating the security guarantee achieved by Theorem 1.2 is that it is possible to securely compute any multiparty functionality $\mathcal{F}$ under concurrent general composition, as long as all the protocols running concurrently in the network are either δ-delayed or are constructed "according to our methodology". In the theorem, the protocols realizing $\mathcal{F}_1, \ldots, \mathcal{F}_t$ are those that are constructed "according to our methodology". We also stress that the protocol ρ_i that securely realizes $\mathcal{F}_i$ is the same, irrespective of the other functionalities $\mathcal{F}_j$ and protocols ρ_j that are being run together with it. Furthermore, the parameter δ is also independent of the functionalities.

We note that Theorem 1.1 follows immediately from Theorem 1.2. This is due to the fact that in Theorem 1.2 security is guaranteed as long as the secure protocols are run alongside arbitrary protocols that are δ-delayed. In Theorem 1.1, on the other hand, security is guaranteed only under concurrent self composition, where the set of arbitrary protocols running alongside is empty.

We prove Theorems 1.1 and 1.2 by first constructing a protocol that securely realizes the common random string (CRS) functionality under concurrent general composition (as long as the arbitrary protocols running together with it are delayed). The CRS functionality simply hands each party a uniformly distributed string, and as such is essentially a multiparty coin-tossing functionality. We then rely on the fact that any efficient functionality can be securely computed under concurrent general composition in the common random string model [16]. Thus combining our protocol for securely realizing the CRS functionality together with a delayed version of the protocol of [16], instantiated separately for each functionality $\mathcal{F}_i$, we obtain that $\mathcal{F}_1, \ldots, \mathcal{F}_t$ can be securely computed under concurrent general composition with any δ-delayed protocol.

Discussion. The proof of Theorem 1.2 gives the first construction of a protocol that achieves security (in the standard sense) in such a general setting of composition, without relying on a trusted setup phase or an assumed honest majority. Of course, as we have mentioned, this comes at the price of the honest parties delaying all messages of protocols running concurrently to our protocol. On the one hand, these delay instructions can easily be carried out by all honest parties. However, on the other hand, they impose a severe slow-down that is unlikely to be tolerated in real settings. Thus, reducing the number and length of the delays imposed by our protocol is an important question for future research. A more important shortcoming of our result is that we do not achieve security if the arbitrary protocol π that is run together with our protocols uses time in its own instructions. This question is also left open. Despite the above shortcomings, our protocol provides the first feasibility result for this setting, and it demonstrates that timing assumptions *can* be used to bypass the broad impossibility results for achieving security under concurrent general composition.

Negative results. As we have mentioned, the timing assumptions are used for introducing time-out and delay instructions in the protocol. Furthermore, our use of delays is extensive, since we do not only insert delay instructions into our secure protocol ρ, but we also require that *every* message of every arbitrary (timing-free) protocol π that runs concurrently to ρ is also delayed. This is clearly a drawback of our result. However, we show that some sort of "time-based" modification of π is *essential* for achieving security. Recall that a protocol π is timing-free if it never looks at its clock and so contains no time-based instruction (and in particular no delay or time-out instructions). We say that a protocol ρ is secure under concurrent general composition with timing-free protocols, if it is secure when run concurrently with any timing-free protocol π. We prove the following theorem (stated informally here):

Theorem 1.3 *There exist large classes of efficient functionalities that cannot be securely computed under concurrent general composition with timing-free protocols, even in the timing model, unless an honest majority or a trusted setup phase are assumed.*

We conclude that some timing-based modification must be introduced into π. The question of how many delays must be introduced into π, and of what length, is left open by this work.

1.3 Related Work

Secure computation was first studied in the stand-alone model, where it was shown that any multi-party functionality can be securely computed [48, 30, 8, 17]. The study of concurrent composition of protocols was initiated by [26] in the context of witness indistinguishability, and was next considered by [21] in the context of non-malleability. Until recently, the majority of work on concurrent composition was in the context of concurrent zero-knowledge [22, 47]. The concurrent composition of protocols for *general secure computation* was only considered much later. Specifically, the first definition and composition theorem for security under concurrent general composition was presented by [20] for the case of perfect security in the information-theoretic setting. Next, [45] considered the computational setting and the case that a secure protocol is executed once in an arbitrary network. The general case for the computational setting, where many secure protocol executions may take place (again, in an arbitrary network) was then considered in the definition (and composition theorem) of universal composability [12]. It was also shown that any functionality can be securely realized in this setting assuming an honest majority [12], or assuming a trusted setup phase in the form of a common random string [16], or in the form of a key registration functionality [3]. However, in the case of no honest majority or trusted setup, broad impossibility results have been demonstrated for universal composability, concurrent general composition and concurrent self composition [15, 37, 38].

These impossibility results justify and provide motivation for considering restricted network settings and weaker notions of security. One type of restriction that has been considered for concurrent self composition is that of m-bounded concurrency, where an upper bound m on the global number of concurrent executions is assumed [1]. In this model, both positive results [36, 44, 43] and lower bounds [36, 38] have been demonstrated. In our opinion, the timing model is a more realistic assumption than that of bounded concurrency. A different way of bypassing the aforementioned impossibility results (and one not taken in this paper) is to consider weaker notions of security. This approach was taken by the works [42, 46, 5, 39] who all provide "additional power" to the ideal adversary (i.e., they allow the simulator to run longer than the real-model adversary). We remark that such solutions provide weaker security guarantees.

As we have mentioned, timing assumptions were introduced in the cryptographic context by [22]. Subsequently, they have been used in a number or works, including [24, 23, 34, 29]. However, all of these works considered the security of specific cryptographic tasks (namely, zero-knowledge and authentication-type protocols). Furthermore, they all considered security under a limited form of concurrent self composition. This paper is the first to use timing assumptions in order to construct a secure protocol for *any* multiparty functionality, that remains secure under concurrent self composition, and under concurrent general composition with (timing-free) delayed protocols.

2 Definitions and Tools

2.1 Preliminaries

We denote the security parameter by n. A function $\mu(\cdot)$ is negligible in n (or just negligible) if for every polynomial $p(\cdot)$ there exists a value N such that for all $n > N$ it holds that $\mu(n) < 1/p(n)$. A machine is said to run in polynomial-time if its number of steps is polynomial in the *security parameter*, irrespective of the length of its input. Formally, each machine has a security-parameter tape upon which 1^n is written. The machine is then polynomial in the contents of this tape.

Let $X = \{X(n,a)\}_{n\in\mathsf{N},a\in\{0,1\}^*}$ and $Y = \{Y(n,a)\}_{n\in\mathsf{N},a\in\{0,1\}^*}$ be distribution ensembles. Then, we say that X and Y are computationally indistinguishable, denoted $X \stackrel{c}{\equiv} Y$, if for every non-uniform polynomial-time distinguisher D there exists a function $\mu(\cdot)$ that is negligible in n, such that for *every* $a \in \{0,1\}^*$,

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

Typically, the distributions X and Y will denote the output vectors of the parties in real and ideal executions, respectively. In this case, a denotes the parties' inputs.

2.2 Security Under Concurrent General Composition in the Timing Model

In this section, we present the definition of concurrent general composition in the timing model. This is a direct extension of the definition of concurrent general composition in the standard (non-timed) model, as defined for example in [12, 37]. Informally speaking, concurrent general composition considers the case that a protocol ρ for securely computing some ideal functionality $\mathcal{F}$, is run concurrently (many times) with arbitrary protocols running in the network. This arbitrary network is modelled as a "calling protocol" π with respect to the functionality $\mathcal{F}$. That is, π is a protocol that contains, among other things, "ideal calls" to a trusted party that computes the functionality $\mathcal{F}$. This means that in addition to standard messages sent between the parties, protocol π's specification contains instructions of the type "send the value x to the trusted party and receive back output y". Then, the real-world scenario is obtained by replacing the ideal calls to

$\mathcal{F}$ in protocol π with real executions of protocol ρ. (When we say that an ideal call to $\mathcal{F}$ is replaced by an execution of ρ, this means that the parties run ρ upon the same inputs that π instructs them to send to the trusted party computing $\mathcal{F}$.) The composed protocol is denoted by π^ρ and it takes place without any trusted help.

We note that in this composed protocol, messages of π may be sent concurrently to the executions of ρ (even though π "calls" ρ). In addition, the inputs are determined by π and may therefore be influenced by previous ρ-outputs and the party's overall view in the arbitrary network. We stress that although the above seems very similar to "modular sequential composition" of [11], it is fundamentally different. Most importantly, here there is no restriction on the scheduling of the executions of ρ (or equivalently calls to $\mathcal{F}$) with respect to π. For example, π may include many calls to ρ and these executions may be run concurrently with each other and with other messages of π. (This can be achieved by having π be multi-threaded and thus contain instructions like "invoke n executions of ρ simultaneously".) To take this a step further, π may not even fix the scheduling ahead of time; rather it may contain instructions like "upon receiving a request from some party to run ρ proceed using the input x". We also remark that π models a large network and may therefore involve many more parties than in any single execution of ρ. Furthermore, there is no limitation on the different sets of parties running ρ; they may sometimes be distinct, partially intersecting or the same.

Now, security is defined for ρ by requiring that for every protocol π that contains ideal calls to $\mathcal{F}$, an adversary interacting with the composed real protocol π^ρ (where there is no trusted help) can do no more harm than in an execution of π where a trusted party computes all the calls to $\mathcal{F}$. This therefore means that ρ behaves just like an ideal call to $\mathcal{F}$, even when it is run concurrently with any arbitrary protocol π. The above informal description is the same in the standard (non-timed) model and in the timing model. The only difference is that in the timing model, the parties have access to local clocks. This will be described below.

We stress one more issue regarding the above formulation of security. As we have mentioned, π is supposed to represent a real dynamic network like the Internet. However, it is fixed ahead of time, unlike real networks. The reason why this suffices is because we quantify over *all* protocols π in the definition of security. Thus, if there exists a real-world scenario in which the protocol ρ does not behave like an ideal call to $\mathcal{F}$, it is possible to retroactively take the real-world execution and use it to define a protocol π. It then follows that ρ does not behave like an ideal call to $\mathcal{F}$ with respect to this π, contradicting the definition of security. Note that this holds even if the scheduling of the protocols in the real Internet-like setting depends dynamically on the messages sent. This is due to the fact that, as mentioned above, π does not necessarily fix the scheduling of protocols ahead of time, but may leave this to the adversary and participating parties.

Secure multiparty computation. A multiparty computation task for a set of parties $P_1, \ldots, P_m$ is cast by specifying a (probabilistic polynomial-time) multiparty ideal functionality $\mathcal{F}$ that receives inputs from these parties and provides outputs. The aim of the computation is for the parties to jointly compute the functionality $\mathcal{F}$. According to the standard ideal/real model paradigm [31, 6, 40, 11, 28], a real protocol execution is compared to an ideal execution where a *trusted third party* computes $\mathcal{F}$ for the parties. Instead of explicitly considering such a trusted party, we sometimes talk about the parties (and adversary) communicating directly with the ideal functionality. This is just shorthand for saying that the parties communicate with the trusted party computing the functionality.

Basic network model. As is typical for secure multiparty computation, we assume that the parties are all linked by authenticated (but not necessarily private) channels. Thus, the adversary cannot modify a message sent by one honest parties to another honest party, without being detected. This assumption can be realized using a public-key infrastructure for secure digital signatures. We also assume that each party has a *unique identity.* Under the assumption that a public-key infrastructure is in place, a party's identity could be taken to be its public-key. (This assumes, however, that the adversary cannot copy a party's public-key.) We remark that although the aim of this work is to remove setup assumptions, the existence of authenticated channels is assumed by almost all work on secure multiparty computation, even in the stand-alone model. Indeed the standard security guarantees are not achievable without assuming authenticated channels; see [2] for discussion and work on secure computation without authenticated channels. In any case, we stress that the setup assumption required for achieving authenticated channels (see [13]), is far weaker than both the common-reference string model used in [16] and the key-registration model used in [3].

Adversarial behavior. In this work we consider malicious, static adversaries. That is, the adversary controls an a priori fixed subset of the parties who are said to be corrupted. The corrupted parties follow the instructions of the adversary in their interaction with the honest parties, and may arbitrarily deviate from the protocol specification. The adversary also receives the view of the corrupted parties at every stage of the computation. In our model, the adversary also has full control over the scheduling of the delivery of all messages. Thus, the network is asynchronous. Finally, as we will see below, the adversary has *some* control over the clocks of the honest parties.

The $\mathcal{F}$-hybrid model. Let π be an arbitrary protocol that utilizes ideal interaction with a trusted party computing the multiparty functionality $\mathcal{F}$ (recall that π actually models arbitrary network activity). There may be many copies of the functionality, and so these copies are differentiated by unique session identifiers or *sid*s. A protocol π that runs in the $\mathcal{F}$-hybrid model contains two types of messages: standard messages and ideal messages. A standard message is one that is sent between two parties that are participating in the execution of π, using the point-to-point network (or broadcast channel, if assumed). An ideal message is one that is sent by a participating party (or the adversary) to the ideal functionality $\mathcal{F}$, or from the ideal functionality to a participating party (or the adversary). Ideal messages are typically *inputs* and *outputs* for the functionality being computed by the trusted party. However, in order to model execution failures, there are two "special" ideal messages (or instructions) that the adversary can send to the trusted party. The first is an *abort instruction* which is due to the fact that in our setting (of no honest majority) there is no guaranteed output delivery. The second is a *time-out* instruction that is unique to the timing model. In more detail:

1. Abort instructions: These instructions play the same role as in stand-alone secure computation for the case of no honest majority. That is, in the case that an honest party receives an invalid message in a real protocol execution, it would halt and output abort (meaning that malicious behavior has been detected). The ideal adversary must therefore also be able to cause an honest party to output abort in the ideal (or hybrid) model.

   Thus, the adversary can issue instructions of the type $(\mathsf{abort}, sid, P_i)$ to the trusted party. Upon receiving such a message, the trusted party forwards (abort, sid) to P_i, who in turn sets its output from execution sid to abort.[4] We stress that an abort instruction can be issued at any time and for any party. Furthermore, once an honest party receives abort, it halts the execution (and should refuse to restart it).

[4]We note that in known protocols for stand-alone secure computation without an honest majority in the synchronous model, abort is also output when a party doesn't receive all of its messages in a given round. In our model, a party will either just continue waiting (possibly forever), or will output time-out if so instructed by the protocol.

2. Time-out instructions: These instructions are included in order to model the case that an honest party is instructed to output time-out in a protocol execution (in the timing model). The purpose of a time-out output is to indicate that the execution can be safely repeated (unlike when the output is abort). As with abort, the ideal adversary must also be able to cause an honest party to output time-out in an ideal/hybrid execution. However, the mechanism for time-outs is different than for aborts. In particular, a time-out can only be issued *before* any output was generated.

   Formally, the adversary may issue an instruction of the type $(\text{time-out}, sid)$ to the trusted party. Upon receiving such a message, the trusted party checks if it previously sent any output in execution sid.[5] If no outputs have been sent, then the trusted party sends $(\text{time-out}, sid)$ to all parties and halts the execution. Otherwise (if outputs have been sent), the trusted party just ignores the time-out instruction.

   Note that as used here, the time-out mechanism has no relation to *time*. Indeed, in the hybrid model there are no clocks.

We remark that if one party receives time-out, then no parties receive output. This is due to the fact that the functionality sends time-out to all parties (and so if the adversary delivers an output to an honest party, it can only be time-out). This is in contrast to aborts, where some parties may receive abort and some may receive their prescribed output. (The issue of whether all parties receive abort together or not is discussed in [32].)

Notice that the computation of π is a "hybrid" between the ideal model (where a trusted party carries out the entire computation) and the real model (where the parties interact with each other only). Specifically, the messages of π are sent directly between the parties, and the trusted party is only used in the ideal calls to $\mathcal{F}$.

As is standard for concurrent settings, the adversary controls the scheduling of all messages, including both standard and ideal messages. This means that even if the trusted party sends the same output to all parties at the same time, the honest parties only receive their output if and when the adversary decides to deliver it. As usual, we assume that the parties are connected via authenticated channels. Therefore, the adversary can read all standard messages, and may use this knowledge to decide when, if ever, to deliver a message. (We remark that the adversary cannot, however, modify messages or insert messages of its own.) In contrast, the channels connecting the participating parties and the trusted party are both authenticated *and* private. More precisely, ideal messages are comprised of a public header and a private body. The contents of a message that belong in the header or body is specified by the functionality definition. In general, the public header contains information like the name and session identifier of the functionality for which the message is intended. We stress that although the adversary delivers the entire message, it can only read the public header, and cannot read the private body. However, we adopt the convention that the *length* of this private body is given to the adversary. (This models the fact that the lengths of inputs and outputs cannot be fully hidden from the adversary.)[6]

Computation in the $\mathcal{F}$-hybrid model proceeds as follows. The computation begins with the adversary receiving the inputs and random tapes of the corrupted parties. Throughout the execution, the adversary controls these parties and can instruct them to send any standard and ideal

[5] In the case that the functionality is a simple function, all outputs are generated and sent at the same time. Thus, this reduces to the trusted party just checking if it has computed the function output yet. In the case of *reactive functionalities* where computation takes place over a number of phases, outputs are generated at different times. Here, the trusted party checks that *no outputs* were generated until this time.

[6] For the majority of this paper, the ideal functionality that we consider generates a public common random string. Therefore, all communication between the parties and functionality can be made part of the public header.

messages that it wishes. In addition to controlling the corrupted parties, the adversary delivers all the standard and ideal messages (when and if it wishes to do so) by copying them from outgoing communication tapes to incoming communication tapes. The series of activations is sequential. That is, the adversary is activated first, at which time it can carry out any arbitrary computation. It concludes its activation by writing a message to the incoming communication tape of either a party or an ideal functionality. A party (or an ideal functionality) that receives a message on its incoming communication tape is immediately activated. When it halts, the adversary is activated once again.[7] Upon being activated, the honest parties always follow the specification of protocol π. Specifically, upon receiving a message (delivered by the adversary), the party reads the message, carries out a local computation as instructed by π, and writes standard and/or ideal messages to its outgoing communication tape, as instructed by π. Likewise, the ideal functionality follows its prescribed instructions (and is never corrupted). At the end of the computation, the honest parties write the output value prescribed by π on their output tapes, the corrupted parties output a special corrupted symbol and the adversary outputs an arbitrary function of its view. Let n be the security parameter, let $\mathcal{S}$ be an adversary for the $\mathcal{F}$-hybrid model with auxiliary input $z \in \{0,1\}^*$, let $I \subseteq [m]$ be the set of corrupted parties, and let $\overline{x} = (x_1, \ldots, x_m) \in (\{0,1\}^*)^m$ be the vector of the parties' inputs to π. Then, the hybrid execution of π with ideal functionality $\mathcal{F}$, denoted $\text{HYBRID}^{\mathcal{F}}_{\pi,\mathcal{S},I}(n, \overline{x}, z)$, is defined as the output vector of all parties and $\mathcal{S}$ from the above hybrid execution.

The real model. Let ρ be a multiparty protocol. Intuitively, the composition of protocol π with ρ is such that a real execution of protocol ρ takes the place of an ideal call to $\mathcal{F}$.[8]

In the real model, each party holds the code of a probabilistic interactive Turing machine (ITM) that works according to the specification of the protocol ρ.[9]When π instructs a party to send an ideal message (i.e., input) α to the ideal functionality $\mathcal{F}$ with session identifier sid, the party creates a new instantiation of the ITM for ρ, associates the identifier sid with this machine, and invokes it with input α. Any message that it later receives that is earmarked for ρ with identifier sid, it forwards to this ITM. All other messages (that are not earmarked for ρ) are answered according to π. Finally, when the execution of ρ with identifier sid concludes and a value β is written on the output tape of the ITM, the party copies β to the incoming communication tape for π, as if β is an ideal message (i.e., output) received from the copy of the ideal functionality $\mathcal{F}$ with identifier sid. This composition of π with ρ is denoted π^ρ and takes place without any trusted help. Thus, the computation proceeds in the same way as in the hybrid model, except that all messages are standard. (Like in the hybrid model, the adversary controls message delivery and can also read messages sent, but cannot modify or insert messages.) Let n be the security parameter, let $\mathcal{A}$ be an adversary for the real model with auxiliary input z, let $I \subseteq [m]$ be the set of corrupted parties, and let $\overline{x}$ be the vector of the parties' inputs to π. Then, the real execution of π with ρ, denoted $\text{REAL}_{\pi^\rho,\mathcal{A},I}(n, \overline{x}, z)$, is defined as the output vector of all the parties and $\mathcal{A}$ from the above real execution.

So far the description of the real model is consistent with the standard non-timed model. In our timing model, in addition to the above, each party has a local clock. In order to model parties

[7]The adversary can activate parties at the beginning of the execution, before there are messages to deliver, by sending them a special "begin computation" message.

[8]Recall that though we refer to π as a protocol, it could in fact be an arbitrary asynchronous environment, consisting of multiple protocol executions.

[9]Note that each party receives the same machine and thus the same set of instructions for ρ. This means that separate, fixed roles are not defined for the different parties. Rather, assigning the roles (if different roles exist, like for example in zero-knowledge proof where there are distinct prover and verifier roles) is assumed to be a part of the functionality.

with clocks, we add a clock tape to the interactive Turing machines that model the parties in the network; we call such a modified machine an ITMC (interactive Turing machine with a clock). As we will see below, the adversary is the only machine to update the clock tapes of the parties. The leeway given to the adversary in its control over these tapes determines the model being considered. For example, if the adversary has full control and can write any values that it wishes to the clock tapes, then this is equivalent to a non-timed, fully asynchronous model. On the other extreme, if the adversary initializes all clocks to 0 and adds 1 to each clock at the same time, then this is equivalent to the fully synchronous model.[10] In the timing model, as introduced by [22], the adversary is somewhat limited in its power over the clock tapes. Specifically, the adversary can initialize the values of the clock tapes to any values that it wishes (this initialization takes place at the onset of the computation and models the fact that we do *not* require synchronized clocks). Following this initialization step, the adversary may update the clock of any party that it wishes, under the constraint that a bound on the *clock drift* is preserved. Loosely speaking, this restriction states that the clocks of all machines proceed at approximately the same rate (within a factor of ϵ).

More formally, let $M_1, \ldots, M_\ell$ be the ITMC's in the network and let $a_1, a_2, \ldots$ be the series of global states of all machines in the network, where a_j denotes the global state after the j^{th} activation of a machine by the real-model adversary. (Note that we do not include activations of the adversary, but just of the participating parties.) Denote the contents of the clock tape of machine M_i in activation a_j by $\mathsf{clock}_i(a_j)$, and let $\mathsf{clock}_i(a_0)$ be the initial value of its clock tape. Then, adversarial control over the clocks is modeled as follows:

1. Before the computation begins, the adversary is allowed to write any values that it wishes to the parties' clock tapes (if a value is not written, then the default is 0). These are the *initial* clock values.

2. Every time that the adversary is activated, it is given write access to the clock tapes of all the parties. This write access is limited in a natural way in that the adversary is only allowed to increase the current value. We stress that writing to a party's clock tape does *not* activate it (in this way, it is different than writing to a party's incoming communication tape).

The above describes *how* the adversary updates the clock tapes; it does not specify any limitations over these updates. In the timing model, it is assumed that the clocks all proceed within ϵ units of each other. That is, let $\epsilon \geq 1$ be a constant. We say that an adversary is ϵ-drift preserving if for every pair of parties P_i and P_j and for every $k \geq 1$,

$$\frac{1}{\epsilon} \cdot (\mathsf{clock}_j(a_k) - \mathsf{clock}_j(a_{k-1})) \leq \mathsf{clock}_i(a_k) - \mathsf{clock}_i(a_{k-1}) \leq \epsilon \cdot (\mathsf{clock}_j(a_k) - \mathsf{clock}_j(a_{k-1})) \quad (1)$$

In other words, whenever a party's clock is increased by some value δ, then all other clocks must be increased by some value between δ/ϵ and $\delta\epsilon$. An *equivalent* and more explicit way of stating this requirement is as follows.

Let $\epsilon \geq 1$ be a constant. Then, we say that an adversary is ϵ-drift preserving if there exist a series of values $\delta_1, \delta_2, \ldots$ so that for every i and every $k \geq 1$,

$$\delta_k \leq \mathsf{clock}_i(a_k) - \mathsf{clock}_i(a_{k-1}) \leq \delta_k \cdot \epsilon$$

[10] Of course, just updating the clocks together does not necessarily force the adversary to activate all the parties in parallel (or essentially in parallel, by activating them sequentially in a round robin fashion). Nevertheless, a protocol can force this by having a party abort if it does not receive its round i messages when its clock reads i.

This means that between activation a_{k-1} and activation a_k, the clocks of all parties have increased by a value which is between δ_k and $\delta_k\epsilon$. Intuitively, one can think of δ_k as being the objective real time (although there may be a number of values δ_k that fulfill this condition).[11]

The rest of the execution is the same as in the (non-timed) real execution described above. Let n be the security parameter, let $\mathcal{A}$ be an ϵ-drift preserving adversary for the real model with auxiliary input z, let $I \subseteq [m]$ be the set of corrupted parties, and let $\overline{x}$ be the vector of the parties' inputs to π. Then, the real execution of π with ρ, denoted $\text{REAL}^{\epsilon}_{\pi^\rho,\mathcal{A},I}(n,\overline{x},z)$, is defined as the output vector of all the parties and $\mathcal{A}$ from the above real execution.

Security under concurrent composition in the timing model. Having defined the hybrid and real models, we can now define security of protocols under concurrent composition. We first define security under concurrent general composition. Loosely speaking, the definition asserts that for any context, or calling protocol π, the real execution of π^ρ emulates the hybrid execution of π which utilizes ideal calls to $\mathcal{F}$. This is formulated by saying that for every real-model adversary there exists a hybrid model adversary for which the output distributions are computationally indistinguishable. The fact that the above emulation must hold for *every* protocol π that utilizes ideal calls to $\mathcal{F}$, means that *general composition* is being considered (recall that π represents arbitrary network activity).

As we have described, we don't actually achieve concurrent general composition in the most general sense. Rather, we only obtain security when the arbitrary protocols running in the network are all delayed by some parameter δ. Furthermore, these protocols (before the δ-delays are added) are timing-free; i.e., the protocols never instruct the honest parties to read their clock tapes. Formally,

Definition 1 (delayed protocols): *Let π be any timing-free protocol (in the real model or in the $\mathcal{F}$-hybrid model for some $\mathcal{F}$), and let $\delta = \delta(n)$ be a function. Then, π_δ is the protocol obtained from π by having every honest party* delay *sending every (standard or ideal) message by δ local time units.*[12] *We call π_δ a δ-*delayed protocol.

We stress that the above notion of δ-delayed protocols is only defined for protocols π that are timing free. Thus, any time we refer to a δ-delayed protocol π_δ, it is implicit that the original protocol π was timing-free. We are now ready to define our notion of security:

Definition 2 (security under concurrent general composition in the timing model): *Let ρ be a probabilistic polynomial-time protocol, let $\mathcal{F}$ be an ideal functionality, and let ϵ be a constant. We say that ρ* securely realizes $\mathcal{F}$ under concurrent general composition in the timing model with ϵ-drift *if for every probabilistic polynomial-time m-party protocol π in the $\mathcal{F}$-hybrid model and every ϵ-drift preserving non-uniform polynomial-time real-model adversary $\mathcal{A}$ for π^ρ, there exists a probabilistic non-uniform polynomial-time hybrid-model adversary $\mathcal{S}$ such that for every $I \subseteq [m]$:*

$$\left\{\text{HYBRID}^{\mathcal{F}}_{\pi,\mathcal{S},I}(n,\overline{x},z)\right\}_{n\in\mathsf{N};\overline{x}\in(\{0,1\}^*)^m;z\in\{0,1\}^*} \stackrel{\text{c}}{\equiv} \left\{\text{REAL}^{\epsilon}_{\pi^\rho,\mathcal{A},I}(n,\overline{x},z)\right\}_{n\in\mathsf{N};\overline{x}\in(\{0,1\}^*)^m;z\in\{0,1\}^*} \quad (2)$$

[11] Clearly this alternate condition implies Eq. (1). Conversely, taking $\delta_k = \min_j\{\mathsf{clock}_j(a_k) - \mathsf{clock}_j(a_{k-1})\}$, Eq. (1) implies that $\max_j\{\mathsf{clock}_j(a_k) - \mathsf{clock}_j(a_{k-1})\} \leq \delta_k\epsilon$, which in turn implies that for every i, $\delta_k \leq \mathsf{clock}_i(a_k) - \mathsf{clock}_i(a_{k-1}) \leq \delta_k\epsilon$.

[12] That is, a message generated is sent out when the party is active again and δ time units have passed by the local clock. Depending on the exact local clock reading when the party is activated, the delay introduced may be more than δ.

If there exists a δ such that $\left\{\text{HYBRID}^{\mathcal{F}}_{\pi,\mathcal{S},I}(n,\overline{x},z)\right\} \stackrel{\text{c}}{\equiv} \left\{\text{REAL}^{\epsilon}_{(\pi_\delta)^\rho,\mathcal{A},I}(n,\overline{x},z)\right\}$ *holds (note that here π is δ-delayed), then we say that ρ* securely realizes $\mathcal{F}$ under concurrent general composition with δ-delays in the timing model with ϵ-drift.

We stress that the real execution with $\mathcal{A}$ takes place in a model with time and clocks, whereas the hybrid execution with $\mathcal{S}$ takes place in a model with no timing at all.

The above definition can be extended to deal with a number of protocols $\rho_1,\ldots,\rho_t$ such that each ρ_i securely realizes some functionality $\mathcal{F}_i$. We will need this for stating our main result:

Definition 3 *We say that $\rho_1,\ldots,\rho_t$* securely realize $\mathcal{F}_1,\ldots,\mathcal{F}_t$ under concurrent general composition with δ-delays in the timing model with ϵ-drift *if for every (non-timed) probabilistic polynomial-time m-party protocol π in the $(\mathcal{F}_1,\ldots,\mathcal{F}_t)$-hybrid model and every ϵ-drift preserving non-uniform polynomial-time real-model adversary $\mathcal{A}$ for $(\pi_\delta)^{\rho_1,\ldots,\rho_t}$, there exists a probabilistic non-uniform polynomial-time hybrid-model adversary $\mathcal{S}$ such that for every $I \subseteq [m]$:*

$$\left\{\text{HYBRID}^{\mathcal{F}_1,\ldots,\mathcal{F}_t}_{\pi,\mathcal{S},I}(n,\overline{x},z)\right\}_{n,\overline{x},z} \stackrel{\text{c}}{\equiv} \left\{\text{REAL}^{\epsilon}_{(\pi_\delta)^{\rho_1,\ldots,\rho_t},\mathcal{A},I}(n,\overline{x},z)\right\}_{n,\overline{x},z}$$

where ρ_i is called in the place of any ideal call to $\mathcal{F}_i$.

In the setting of concurrent self composition, a secure protocol is run many times concurrently, but it is the only protocol running in the network. In this setting, which is a special case of the setting of concurrent general composition, the "arbitrary" (non-timed) protocol π contains calls to the ideal functionality, and nothing else (and thus is no longer arbitrary). We let $\lambda_{\mathcal{F}}$ denote the set of (non-timed) protocols that contain only calls to ideal functionality $\mathcal{F}$ (and in particular have no standard messages). We have the following definition, that basically states that a protocol is secure under concurrent self composition if it is secure under concurrent general composition with any protocol from the class of protocols $\lambda_{\mathcal{F}}$:

Definition 4 (security under concurrent self composition in the timing model): *Let ρ be a probabilistic polynomial-time protocol and let $\mathcal{F}$ be an ideal functionality. Then, ρ* securely realizes $\mathcal{F}$ under concurrent self composition in the timing model with ϵ-drift *if for every $\pi \in \lambda_{\mathcal{F}}$ and for every ϵ-drift preserving non-uniform polynomial-time real-model adversary $\mathcal{A}$, there exists a probabilistic non-uniform polynomial-time hybrid-model adversary $\mathcal{S}$ such that for every $I \subseteq [m]$:*

$$\left\{\text{HYBRID}^{\mathcal{F}}_{\pi,\mathcal{S},I}(n,\overline{x},z)\right\}_{n,\overline{x},z} \stackrel{\text{c}}{\equiv} \left\{\text{REAL}^{\epsilon}_{\pi^\rho,\mathcal{A},I}(n,\overline{x},z)\right\}_{n,\overline{x},z}$$

where $\overline{x} \in (\{0,1\}^*)^m$ and $z \in \{0,1\}^*$.

We remark that if Definition 4 had considered protocols in $(\lambda_{\mathcal{F}})_\delta$ rather than protocols in $\lambda_{\mathcal{F}}$, where $(\lambda_{\mathcal{F}})_\delta$ is the set of protocols obtained from $\lambda_{\mathcal{F}}$ by adding δ-delays before the sending of each ideal message, then Definition 4 would clearly be a special case of Definition 2 when considering concurrent general composition with δ-delays; the only difference is that now the set of δ-delayed protocols is much more restricted. Moreover, if a protocol ρ securely realizes $\mathcal{F}$ according to Definition 4 when only considering protocols in $(\lambda_{\mathcal{F}})_\delta$, then there exists a protocol ρ' that securely realizes $\mathcal{F}$ according to Definition 4 (which considers non-timed protocols in $\lambda_{\mathcal{F}}$); protocol ρ' simply runs ρ while delaying the sending of the first message by δ time units. This implies that our result on concurrent self composition follows as an immediate corollary from our result on concurrent general composition with δ-delays.

Non-trivial protocols in the timing model. As we discussed in the introduction, the timing model relies on two assumptions: the clock-drift ϵ and the maximum network latency Δ. However, the security of a protocol should rely solely on the more realistic assumption regarding clock drift. Therefore, our above definition refers to the clock drift, but makes no mention of the network latency; the latency assumption is only used in order to guarantee *non-triviality*. Loosely speaking, a protocol is non-trivial if the honest parties are guaranteed to receive their outputs (according to the functionality definition) in executions where the adversary is "well-behaved". More specifically, in the context of the timing model, a protocol is non-trivial for Δ and ϵ if in each execution in which the adversary is ϵ-drift preserving, delivers all messages in time at most Δ, and does not corrupt any party, none of the parties output time-out or abort.

Definition 5 (non-triviality): *We say that a protocol ρ is* non-trivial under timing assumptions (Δ, ϵ) *if in any execution of ρ where:*

1. *The real-model adversary $\mathcal{A}$ has not corrupted any of the participating parties, and*
2. *The real-model adversary $\mathcal{A}$ is ϵ-drift preserving and delivers all the messages of ρ within Δ time units (according to the clocks of all the parties – see explanation below),*

it holds that all parties receive an output that does not equal time-out *or* abort.

Notice that item (2) in Definition 5 refers to delivery within Δ time units according to the clocks of *all parties*. More specifically, this means that if a message is sent by one honest party P to another honest party P' when the vector of clocks of all parties reads $(t_1, \ldots, t_m)$, then the message is received by P' when the vector of clocks is such that for *every* i, party P_i's clock reads at most $t_i + \Delta$. Thus, Δ is an upper bound on the latency with respect to all local clocks (and not with respect to some specific clock).

Modeling delays and time-outs. As we have discussed in the Introduction, our secure protocols utilize the clocks by introducing delay and time-out instructions. Such instructions can be carried out in our model as follows:

1. *Delay instructions:* If a party P_i is instructed to delay sending a message x by c time units, then it chooses a random identifier *delay-id* and writes $(x, \textit{delay-id}, c, \textit{time})$ on its work tape, where *time* is the current contents of its clock tape. It then writes (delay,*delay-id*,c) on its outgoing communication tape concluding the activation. Upon receiving a message (send,*delay-id*) from the adversary in a future activation, party P_i first checks that c units have passed according to its clock (i.e., that the current contents of its clock is at least $\textit{time} + c$, where c and *time* are the values in the tuple indexed by *delay-id*). If not, then it halts this activation. If yes, then it writes the delayed message x on its outgoing communication tape, concluding the activation. (We note that our decision to write the length c of the delay on the outgoing communication tape is arbitrary and makes no difference to our result.)
2. *Time-out instructions:* If a party P_i (or an ITMC that it runs as a subprotocol) has an instruction to time-out if it doesn't receive a specific message within c time units from the present time, then P_i writes the current contents of its clock tape on its work tape. Then, when it receives the specific message, it outputs time-out if the current contents of its clock tape is greater than the previously recorded value plus c.

Discussion – local computation time. In our definitions, we have included a local clock on machines and use this to measure the time that it takes for messages to be sent and received over the network. A more general model would also include issues such as the time that it takes for local computation. The focus of this paper is a secure protocol that utilizes timing assumptions, and not the issue of modelling time in its most general fashion. Our model therefore assumes that local computation is immediate (this can be seen because the adversary is not activated while local computations take place and so cannot update the clocks). One approach for generalizing the model is to have the adversary be activated after every single step of the transition function of an ITMC. We leave these questions of modelling for future work.

2.3 Tools

Our protocol uses a number of different tools and primitives. In this section, we briefly describe these tools and provide references to full definitions.

Witness indistinguishable and witness hiding proofs [26]. We consider the interactive proof system between a probabilistic polynomial-time verifier and a probabilistic polynomial-time prover who is given an auxiliary input (typically, an NP-witness). Such an interactive proof is witness indistinguishable if interactions in which the prover uses different "legitimate" auxiliary-inputs are computationally indistinguishable from each other [26]. Recall that any zero-knowledge proof system is also witness indistinguishable. Furthermore, witness indistinguishable proofs remain witness indistinguishable under concurrent composition. Witness hiding proofs have the property that a verifier cannot obtain a witness from its interaction with the prover. For example, if a prover proves that it knows the preimage of some one-way function using a witness-hiding proof, then the interaction will not help any probabilistic polynomial-time verifier to compute a preimage. Witness hiding proofs can be constructed from witness indistinguishable proofs by considering "double statements" with independent witnesses, of the form "I know the preimage of one of v_1 and v_2" [26]. See [27, Section 4.6] for a full treatment of witness indistinguishable and witness hiding proofs.

Strong proofs of knowledge [27]. A proof of knowledge [33, 7] is an interactive proof which convinces a verifier that the prover "knows" a witness to a certain statement. This is in contrast to a regular interactive proof, where the verifier is just convinced of the validity of the statement. The concept of "knowledge" for machines is formalized by saying that if a prover can convince the verifier, then there exists an efficient procedure that can "extract" a witness from this prover (thus the prover knows a witness because it can run the extraction procedure on itself). More formally, a proof of knowledge has the property that for every machine P^* there exists a *knowledge extractor* K, such that if P^* convinces V with probability p, then K "extracts" a valid witness from the prover P^* with probability that is negligibly close to p. A strong proof of knowledge, as defined by Goldreich [27, Sec. 4.7.6], is a proof of knowledge where the knowledge extractor runs in *strict* polynomial-time and fulfills the following more stringent requirement: There exists a negligible function $\mu(n)$ such that if a given prover convinces the honest verifier to accept with probability greater than $\mu(n)$, then the knowledge extractor succeeds in obtaining a witness with probability at least $1 - \mu(n)$. See [27, Sec. 4.7.6] for a full treatment.

We remark that there exist witness indistinguishable strong proofs of knowledge with any *super-constant* number of rounds. (The construction of [27] uses a super-logarithmic number of sequential executions of the 3-round zero-knowledge proof for Hamiltonicity [10]. However, using the same ideas, it can be shown that by running $\log n$ parallel executions of the proof of Hamiltonicity, any super-constant number of sequential repetitions is actually enough. We can therefore reduce this to any super-constant number of rounds $\alpha(n) = \omega(1)$.) We also remark that it has been shown that under exponential hardness assumptions, there *do not exist* witness indistinguishable *strong* proofs of knowledge with a constant number of rounds, even using non-black-box techniques [4].

3 Constructing Secure Protocols in the Timing Model

In this section we prove our main positive results, which consists of proving Theorem 1.2 (and obtaining Theorem 1.1 as a corollary). We begin by formally restating Theorem 1.2.

Theorem 6 (Theorem 1.2 – restated): *Assume that there exist enhanced trapdoor permutations and dense cryptosystems, and let* Δ *and* ϵ *be constants where* $1 \leq \epsilon < \sqrt[3]{1.5}$.[13] *Then, there exists a function* $\delta(n) \stackrel{\rm def}{=} \alpha(n) \cdot \Delta \cdot \epsilon$ *such that for any set of probabilistic polynomial-time functionalities* $\mathcal{F}_1, \ldots, \mathcal{F}_t$ *there exist probabilistic polynomial-time protocols* $\rho_1, \ldots, \rho_t$ *that securely realize* $\mathcal{F}_1, \ldots, \mathcal{F}_t$ *under concurrent general composition with* δ*-delays in the timing model with* ϵ*-drift. Furthermore, each* ρ_i *is non-trivial under timing assumptions* (Δ, ϵ)*, and does not depend on the other functionalities* $\{\mathcal{F}_j\}_{j \neq i}$.

The majority of the proof of Theorem 6 involves showing how to securely realize the "common random string" (CRS) functionality under concurrent general composition with any δ-delayed protocol π_δ. We begin by formally defining the CRS functionality in Section 3.1. Next, in Section 3.2, we show that in order to prove Theorem 6, it suffices to securely realize the CRS functionality under concurrent general composition with δ-delays. Finally, Sections 3.3 to 3.6 are devoted to showing how to indeed securely realize the CRS functionality under concurrent composition with δ-delays.

3.1 The CRS Functionality

We now formally define the common random string functionality, denoted $\mathcal{F}_{\rm CRS}$. Intuitively, the functionality simply chooses a random string and sends it to all parties. Any party can send the functionality a crsgen request. Once the functionality receives such a request, it generates a random string $R_{\rm CRS}$ and sends it to the adversary and all the parties. Recall that the adversary controls the delivery of messages between $\mathcal{F}_{\rm CRS}$ and the parties; therefore, the fact that $\mathcal{F}_{\rm CRS}$ sends the output does not mean that the parties receive it immediately (or even that they will ever receive it). Recall also that parties may receive time-out and abort outputs, as defined in Section 2.2. A formal description of $\mathcal{F}_{\rm CRS}$ appears in Figure 2.

Functionality $\mathcal{F}_{\rm CRS}$

Let n be the security parameter and let $p(\cdot)$ be a fixed polynomial.[14] Let $P_1, \ldots, P_m$ be the set of all parties, and let $\mathcal{S}$ be the adversary. The functionality $\mathcal{F}_{\rm CRS}$ proceeds as follows:

> Upon receiving a message $(\mathsf{crsgen}, sid, \{P_{i_1}, \ldots, P_{i_k}\})$, choose a uniformly distributed string $R_{\rm CRS} \in_R \{0,1\}^{p(n)}$, send $(\mathsf{crsgen}, sid, \{P_{i_1}, \ldots, P_{i_k}\}, R_{\rm CRS})$ to $\mathcal{S}$ and to all parties $P_{i_1}, \ldots, P_{i_k}$, and halt.[15]

Figure 2: The ideal multiparty $\mathcal{F}_{\rm CRS}$ functionality

[13] This limitation on ϵ is needed in our proof of security.

[14] $\mathcal{F}_{\rm CRS}$ is parameterized by a polynomial $p(\cdot)$ that determines the length of the common random string generated. If desired, this can be included as input with almost no difference to the protocol (the only necessary addition is for the parties to negotiate the value of $p(\cdot)$ at the onset).

[15] When the set of parties $\{P_{i_1}, \ldots, P_{i_k}\}$ is clear from the context, we sometimes let the message $(sid, \mathsf{compute})$ be a shorthand for the message $(\mathsf{crsgen}, sid, \{P_{i_1}, \ldots, P_{i_k}\})$. Similarly, we let $(sid, R_{\rm CRS})$ be a shorthand for

We note that the $\mathcal{F}_{\text{CRS}}$ functionality sends only *uniformly distributed* strings (in contrast to some prior definitions which allowed any efficiently samplable distribution). This is crucial for our implementation since we use a coin-tossing protocol.

3.2 Reducing the Problem to Realizing the CRS Functionality

In the $\mathcal{F}_{\text{CRS}}$-hybrid model, all parties are given access to the $\mathcal{F}_{\text{CRS}}$ functionality. As we have mentioned, it follows from [16] and from the composition theorem given in [12], that if the $\mathcal{F}_{\text{CRS}}$ functionality can be securely realized under concurrent general composition then *any* functionality $\mathcal{F}$ can be securely realized under concurrent general composition (assuming the existence of enhanced trapdoor permutations and dense cryptosystems).[16]

Unfortunately, it is known that the $\mathcal{F}_{\text{CRS}}$ functionality *cannot* be securely realized under concurrent general composition, unless an honest majority or a trusted setup phase is assumed [12, 14, 15]. Moreover, as we show in Section 5, *even in the timing model*, the $\mathcal{F}_{\text{CRS}}$ functionality *cannot* be securely realized under concurrent general composition. Instead, we show how to securely realize the $\mathcal{F}_{\text{CRS}}$ functionality in the timing model under concurrent general composition *with δ-delays*, as defined in Definition 2. In this section, we show that this implies that any series of multiparty functionalities can be securely realized in the timing model under concurrent general composition with δ-delays, as defined in Definition 3. Our proof follows the following series of steps. Let $\mathcal{F}_1, \ldots, \mathcal{F}_t$ be a series of multiparty functionalities and let ρ be a protocol that securely realizes the $\mathcal{F}_{\text{CRS}}$ functionality in the timing model under concurrent general composition with δ-delays, for δ as defined in the theorem. Then,

1. We first use [16] to obtain timing-free protocols $\sigma_1, \ldots, \sigma_t$ such that each σ_i securely realizes $\mathcal{F}_i$ under concurrent general composition in the $\mathcal{F}_{\text{CRS}}$-hybrid model. (Note that each σ_i is secure in a model without time.)

2. Next, we wish to replace the ideal calls to $\mathcal{F}_{\text{CRS}}$ in each σ_i by invocations of the real protocol ρ. However, since ρ is only secure when all other protocols that are run concurrently in the network are δ-delayed, we first insert delays into all the protocols $\sigma_1, \ldots, \sigma_t$. We denote by $\hat{\sigma}_i$ the δ-delayed version of σ_i, as defined in Definition 1. Note that inserting delays into σ_i does not affect its security because the security of σ_i holds under any scheduling.

3. Finally, we define the protocol ρ_i to be the composition of $\hat{\sigma}_i$ with ρ (i.e., $\hat{\sigma}_i$ uses real calls to ρ instead of ideal calls to $\mathcal{F}_{\text{CRS}}$).

We now explain intuitively why $\rho_1, \ldots, \rho_t$ securely realize $\mathcal{F}_1, \ldots, \mathcal{F}_t$. Let π be any arbitrary protocol. We argue that the result of a real execution of π with $\rho_1, \ldots, \rho_t$ is indistinguishable from the result of an execution of π with ideal calls to $\mathcal{F}_1, \ldots, \mathcal{F}_t$,[17] as follows. Define a new protocol π' that includes the timing-free protocol π and the timing-free protocols $\sigma_1, \ldots, \sigma_t$, where each σ_i replaces the ideal calls to $\mathcal{F}_i$. The protocol π' is timing-free and it runs in the $\mathcal{F}_{\text{CRS}}$-hybrid model, because each σ_i runs in the $\mathcal{F}_{\text{CRS}}$-hybrid model. The fact that an execution of π' with $\mathcal{F}_{\text{CRS}}$ is indistinguishable from a hybrid execution of π with $\mathcal{F}_1, \ldots, \mathcal{F}_t$ follows from the security of

$(\mathsf{crsgen}, sid, \{P_{i_1}, \ldots, P_{i_k}\}, R_{\text{CRS}})$.

[16] Actually, the result in [16] holds only for the class of "well-formed" functionalities. However, in the case of static adversaries, this only limits the functionalities to those that are "unaware" of which parties are corrupted and which are honest. Since in our definition of the computational model the ideal functionality is not given this information, it follows that *all efficient functionalities* can be securely realized.

[17] Of course, this is completely informal and what we mean is that the conditions of Definition 3 are met.

$\sigma_1, \ldots, \sigma_t$ under concurrent general composition in the $\mathcal{F}_{\text{CRS}}$-hybrid model. Next, we use the fact that ρ securely realizes $\mathcal{F}_{\text{CRS}}$ when run concurrently with any δ-delayed protocol. In particular, it securely realizes $\mathcal{F}_{\text{CRS}}$ when run concurrently with π'_δ (the δ-delayed version of π'). Thus, it follows that the real protocol π'_δ with ρ, denoted $(\pi'_\delta)^\rho$, is indistinguishable from π' with ideal calls to $\mathcal{F}_{\text{CRS}}$, which is in turn indistinguishable from π with ideal calls to $\mathcal{F}_1, \ldots, \mathcal{F}_t$. Finally, note that $(\pi'_\delta)^\rho$ is exactly the same as the composition of π with the protocols $\hat{\sigma}_1, \ldots, \hat{\sigma}_t$. Thus, the theorem follows.

We proceed to provide a formal proof of the above; i.e., of the fact that the protocols of [16] (with delays of length δ inserted before the sending of each message), composed with a protocol for securely realizing the $\mathcal{F}_{\text{CRS}}$ functionality in the timing model under concurrent general composition with δ-delays, results with protocols that securely realize $\mathcal{F}_1, \ldots, \mathcal{F}_t$ under concurrent general composition with δ-delays. We remark that the proofs are straightforward, and just involve justifying a few technical points related to the above informal reasoning. First, the hybrid model of [16] is not quite the same as the one defined here, since in our definition the adversary can send time-out instructions to the trusted party. This is in contrast to the hybrid model of [16] which does not include time-out instructions (this technicality was not expressed in the above informal discussion but is proven below). Second, we do not securely realize the $\mathcal{F}_{\text{CRS}}$ functionality under standard concurrent general composition (even though the security of the σ_i protocols was proven in this setting); rather, we securely realize it under concurrent general composition *with δ-delays.* Despite the above differences, we show that our construction is secure.

We note one technicality that must be dealt with. According to our definition, a time-out can only be issued before any outputs are obtained. Thus, a policy must be determined regarding what the delayed protocol $\hat{\sigma}_i$ should do in case $\mathcal{F}_{\text{CRS}}$ returns time-out (note that in the model of [16] where σ was constructed, the functionality $\mathcal{F}_{\text{CRS}}$ never returns time-out; this issue only arises in the timing model). There are two possible approaches here. First, we could define that the output of σ in this case is abort (and not time-out). The drawback of this approach is that we lose the advantage of a time-out that enables parties to restart the execution. The second approach is to rely on the fact that the protocol σ from [16] can be written so that there is only a single call to $\mathcal{F}_{\text{CRS}}$, and this takes place at the very onset of the execution (before any outputs are produced). We prefer this latter approach as it results in a protocol that can be restarted in the case that high network latency is the only reason that the protocol did not terminate successfully.

We first claim that time-out instructions can be added to a secure timing-free protocol σ without making any real difference. As discussed above, we assume that σ contains only a single call to $\mathcal{F}_{\text{CRS}}$, and this takes place at the onset of the execution. In the proof below, we will refer both to the $\mathcal{F}_{\text{CRS}}$-hybrid model defined in this paper where time-out instructions are allowed and to the $\mathcal{F}_{\text{CRS}}$-hybrid model of [16] where there are no time-out instructions. We will call these the $\mathcal{F}_{\text{CRS}}$-hybrid model with and without time-outs, respectively.

Claim 3.1 *Let $\mathcal{F}$ be a functionality and let σ be a timing-free protocol that contains a single call to $\mathcal{F}_{\text{CRS}}$ at the onset of the protocol and securely realizes $\mathcal{F}$ under concurrent general composition in the $\mathcal{F}_{\text{CRS}}$-hybrid model without* time-outs. *Define the protocol σ' to be the same as σ except that if an honest party receives* time-out *as output from the ideal $\mathcal{F}_{\text{CRS}}$ functionality, then it outputs* time-out *in σ' and halts. Then, protocol σ' securely realizes $\mathcal{F}$ under concurrent general composition in the $\mathcal{F}_{\text{CRS}}$-hybrid model with* time-outs, *as defined in Section 2.2.*

Proof Sketch: The only difference between σ and σ' is with respect to the possibility of obtaining time-out. There are two cases here:

1. *Case 1 – the $\mathcal{F}_{\text{CRS}}$ functionality sends* time-out *to all honest parties:* In this case, the ideal-model adversary/simulator for σ' sends a time-out instruction to the trusted party computing

the functionality $\mathcal{F}$. Since $\mathcal{F}_{\text{CRS}}$ is called at the onset of the execution, and thus before any outputs are generated, this implies that the time-out instruction is "accepted" by the trusted party, and thus it will also send time-out to all honest parties.

2. *Case 2 – the $\mathcal{F}_{\text{CRS}}$ functionality does not send* time-out *to the honest parties:* In this case, the execution of σ' is exactly the same as σ. Therefore, the simulator for σ' follows exactly the same strategy as the simulator for σ.

Note that the simulator for σ' knows whether it is in case 1 or case 2, and so can efficiently implement the above strategy. Furthermore, by the definition of the hybrid model, all parties receive time-out or none do; thus these are the only two cases. This completes the proof sketch. ■

So far, we have shown that adding time-out instructions does not make any difference with respect to a protocol's security in the $\mathcal{F}_{\text{CRS}}$-hybrid model. The next claim shows that in order to securely realize any t functionalities $\mathcal{F}_1, \ldots, \mathcal{F}_t$, it suffices to prove that the $\mathcal{F}_{\text{CRS}}$ functionality can be securely realized under concurrent general composition with δ-delayed protocols.

Claim 3.2 *For any set of probabilistic polynomial-time multiparty functionalities $\mathcal{F}_1, \ldots, \mathcal{F}_t$, let $\sigma'_1, \ldots, \sigma'_t$ be timing-free protocols such that for every i, σ'_i securely realizes $\mathcal{F}_i$ under concurrent general composition in the $\mathcal{F}_{\text{CRS}}$-hybrid model with* time-outs*. Furthermore, let ρ be a real-world protocol that securely realizes $\mathcal{F}_{\text{CRS}}$ (with* time-outs*) under concurrent general composition with δ-delays in the timing model with ϵ-drift. For every i, define the real-world protocol $\rho_i = (\hat{\sigma}_i)^\rho$, where $\hat{\sigma}_i = (\sigma'_i)_\delta$ is the δ-delayed version of σ'_i. Then, the real protocols $\rho_1, \ldots, \rho_t$ securely realize $\mathcal{F}_1, \ldots, \mathcal{F}_t$ under concurrent general composition with δ-delays in the timing model with ϵ-drift.*

Proof Sketch: Fix any set of probabilistic polynomial-time multiparty functionalities $\mathcal{F}_1, \ldots, \mathcal{F}_t$, and let $\sigma'_1, \ldots, \sigma'_t$ be as in the claim statement. Let π be an arbitrary protocol that contains ideal calls to the functionalities $\mathcal{F}_1, \ldots, \mathcal{F}_t$. Consider now the real-world protocol $(\pi_\delta)^{\rho_1,\ldots,\rho_t}$, where each ρ_i is as defined in the claim statement, where ρ_i replaces every ideal call to $\mathcal{F}_i$ in π, and where the messages of π are delayed by δ units of time. It holds that

$$(\pi_\delta)^{\rho_1,\ldots,\rho_t} = ((\pi_\delta)^{\hat{\sigma}_1,\ldots,\hat{\sigma}_t})^\rho = ((\pi^{\sigma'_1,\ldots,\sigma'_t})_\delta)^\rho,$$

where the first equality from the fact that $\rho_i = (\hat{\sigma}_i)^\rho$, and the second equality follows from the fact that $\hat{\sigma}_i = (\sigma'_i)_\delta$. Now, by the assumption in the claim, ρ securely realizes $\mathcal{F}_{\text{CRS}}$ under concurrent general composition with δ-delays in the timing model with ϵ-drift. Thus, for every real-model adversary $\mathcal{A}$ there exists an adversary $\mathcal{H}$ in the $\mathcal{F}_{\text{CRS}}$-hybrid model with time-outs such that

$$\begin{aligned}\left\{\text{REAL}^\epsilon_{(\pi_\delta)^{\rho_1,\ldots,\rho_t},\mathcal{A},I}(n,\overline{x},z)\right\}_{n,\overline{x},z} &\equiv \left\{\text{REAL}^\epsilon_{((\pi^{\sigma'_1,\ldots,\sigma'_t})_\delta)^\rho,\mathcal{A},I}(n,\overline{x},z)\right\}_{n,\overline{x},z} \\ &\stackrel{\text{c}}{\equiv} \left\{\text{HYBRID}^{\mathcal{F}_{\text{CRS}}}_{\pi^{\sigma'_1,\ldots,\sigma'_t},\mathcal{H},I}(n,\overline{x},z)\right\}_{n,\overline{x},z}\end{aligned}$$

Next, we use the fact that σ'_i securely realizes $\mathcal{F}_i$ under concurrent general composition in the $\mathcal{F}_{\text{CRS}}$-hybrid model with time-outs (note that there is actually no time in this model). Specifically, this implies that for every adversary $\mathcal{H}$ in the $\mathcal{F}_{\text{CRS}}$-hybrid model with time-outs, there exists an adversary $\mathcal{S}$ in the $\mathcal{F}_1, \ldots, \mathcal{F}_t$-hybrid model with time-outs such that

$$\left\{\text{HYBRID}^{\mathcal{F}_{\text{CRS}}}_{\pi^{\sigma'_1,\ldots,\sigma'_t},\mathcal{H},I}(n,\overline{x},z)\right\}_{n,\overline{x},z} \stackrel{\text{c}}{\equiv} \left\{\text{HYBRID}^{\mathcal{F}_1,\ldots,\mathcal{F}_t}_{\pi,\mathcal{S},I}(n,\overline{x},z)\right\}_{n,\overline{x},z}$$

Combining the above, we have that $\rho_1, \ldots, \rho_t$ securely realize $\mathcal{F}_1, \ldots, \mathcal{F}_t$ under concurrent general composition with δ-delays in the timing model with ϵ-drift, as required. ■

Under the assumption that enhanced trapdoor permutations and dense cryptosystems exist, the result of [16] provides us with protocols σ as required in Claim 3.1. Thus, using Claim 3.2 we conclude that in order to prove Theorem 6 it suffices to prove that there exists a protocol ρ that securely realizes $\mathcal{F}_{\text{CRS}}$ under concurrent general composition with δ-delays in the timing model with ϵ-drift. The rest of this section is devoted to this task of securely realizing $\mathcal{F}_{\text{CRS}}$.

Remark: *Above we introduced the* "hybrid model without time-outs" *only in order to refer to the functionalities in [16] which do not accept* time-out *instructions from the adversary. For the rest of this paper, whenever we refer to a hybrid model, we mean the model defined in Section 2.2 that allows* time-outs*.*

3.3 Overview of the Protocol for $\mathcal{F}_{\text{CRS}}$ and its Security Analysis

Before proceeding to describe the actual protocol for securely realizing the $\mathcal{F}_{\text{CRS}}$ functionality, we provide a high-level overview of the construction. The basic structure of the protocol is an extension of the two-party coin-tossing protocol of [35] (which is in turn an extension of Blum's protocol [9]). In this protocol, each party first commits to a randomly chosen value and provides a zero-knowledge proof of knowledge of the committed value. In the next phase of the protocol, each party reveals its committed value, without actually decommitting, and provides a zero-knowledge proof that the revealed value is indeed the one that was committed to. The idea behind this construction is that due to the soundness of the proofs, a corrupted party has no choice but to reveal the value that it committed to in the first phase. Thus, the binding property of the commitment scheme is preserved. Intuitively, this means that the adversary cannot bias the outcome of the coin-tossing protocol, because it is bound to the corrupted parties' committed values before it sees the honest parties' committed values. However, in order to prove the security of the coin-tossing protocol according to the simulation paradigm, it is necessary to construct a simulator that can force the outcome of the protocol to be the exact string R_{CRS} that is generated by the ideal functionality $\mathcal{F}_{\text{CRS}}$. The zero-knowledge proofs of knowledge are included in order to facilitate such a simulation. Specifically, it is true that the adversary must reveal the correctly committed value due to the soundness of the proofs. However, the simulator can run the simulator for the zero-knowledge protocol, and can effectively cheat. Thus, the simulator can reveal any value that it wishes and is not bound by the commitment scheme (note that decommitment never actually takes place; rather the committed value is revealed and the zero-knowledge proof is used to determine that it is correct). This observation is used in the following way. In the first phase of the protocol, the simulator commits to random values for the honest parties, and extracts all of the values committed to by the corrupted parties (it does this by running the knowledge extractor on the proofs of knowledge of the committed values). Next, given the corrupted parties' values, it chooses *new* random strings for the honest parties so that the XOR of the extracted corrupted parties' values and the new honest parties values equals R_{CRS} exactly. Finally, the simulator reveals the new (fake) honest party values and simulates the zero-knowledge proofs claiming that the revealed values are indeed the committed ones (which they are not). The adversary's view in this simulation is indistinguishable from in a real protocol execution due to the hiding property of the commitments and the zero-knowledge property of the proofs.

A crucial point in the above security argument is that the proofs of knowledge must be run *independently of each other* (in order to ensure that the adversary does not "copy" a proof from an honest party). The same holds also for the zero-knowledge proofs of consistency in phase 2. (Here the reason is slightly different. During simulation, the simulator actually "cheats" by proving an incorrect theorem. We need to ensure that the adversary cannot use the cheating of the simulator in order to cheat itself.) In the stand-alone case, this independence is achieved by simply running the proofs *sequentially*. Technically, this enables the rewinding of the proofs of knowledge provided by the adversary (for extraction in the first phase) and the rewinding of the zero-knowledge proofs verified by the adversary (for simulation in the second phase) without overlapping and therefore without interfering with any of the other proofs. In our case, however, we must achieve security under *concurrent composition*. Therefore, it is not possible to enforce any specific scheduling that will ensure independence between the proofs (or that they don't overlap during rewinding).

As a first step towards solving this problem (and as a solution to another problem), we limit the rewinding stages to be "early" on in the protocol. In particular, rewinding takes place only before the decommitment values are revealed and so before the common reference string can be learned by the adversary. This is a necessary step because *time-out* instructions are crucial for enabling "proper rewinding," and as we have discussed in Subsection 2.2, a time-out must only occur before the adversary can learn the output. We achieve this by using the specific zero-knowledge arguments of knowledge of Feige and Shamir [25] (an argument is a proof system where soundness is only guaranteed for polynomial-time cheating provers). Loosely speaking, the Feige-Shamir argument system consists of two witness-indistinguishable proofs of knowledge (WIPOKs, for short); first the verifier proves that it knows one of two independent secrets; next, the prover proves either that it knows one of the verifier's secrets or that it knows the real witness. The soundness of this protocol follows from the fact that a WIPOK for statements with multiple independent witnesses is *witness hiding*. Therefore, the prover could not have obtained the secret from the first WIPOK and must use the real witness in the second WIPOK. The zero-knowledge property is demonstrated by first extracting a secret from the verifier in the first stage, and then proving the second WIPOK using knowledge of this secret. Note that the second stage of the simulation requires no rewinding, and that this is the only part of the proof that depends on the statement being proved.

To be more precise, our protocol consists of three phases. In Phase 1, each player runs a WIPOK that it knows one of two independent secrets (this is the first WIPOK of the Feige-Shamir argument system). Then, in Phase 2, each player commits to a random value, and runs a single WIPOK that it either knows the value that it committed to or that it knows one of the secrets of the verifier (completing the Feige-Shamir argument that was initiated in Phase 1). Thus, by the end of Phase 2, each player has committed to some value and has proved in zero-knowledge to each of the other players that it knows the value that it committed to. Notice that phases 1 and 2 correspond to the first part of the coin-tossing protocol of [35]. The coin-tossing protocol is then completed in Phase 3 where each player reveals the value that it committed to in Phase 2 (without decommitting), and proves that it is the correct value. This proof is a single WIPOK that it either knows the decommitment information that corresponds to this value or that it knows one of the secrets of the verifier. Once again, combining this WIPOK with that of phase 1, we obtain a Feige-Shamir argument. Thus, both the proofs of Phase 2 and of Phase 3 (which consist of only a single WIPOK) are actually zero-knowledge, as required by the coin-tossing protocol. An important property of this protocol is that the only rewinding needed is **(a)** to extract the secrets from the first Feige-Shamir WIPOK in Phase 1 (enabling simulation later), and **(b)** to extract the committed value from the adversary in Phase 2. This implies that all rewinding takes place before Phase 3, which is where the committed values are revealed. Furthermore, all rewinding is actually for the purpose of *extraction* only.[18]

[18]This strategy simplifies the proof of security, because it turns out to be "much easier" to extract than simulate. This is especially true because we use *strong* proofs of knowledge, rather than ordinary ones; see below.

Until now, we have focused on how to limit the rewinding to the early stage of the protocol, and to witness extraction only. However, a far more crucial issue is how we carry out this extraction (i.e., rewinding) in the concurrent setting. It is here that we use the timing assumptions, via *time-out* and *delay* instructions, in an inherent way. Informally speaking, there are two issues that must be dealt with when considering concurrent composition here: **(a)** the WIPOK protocols must self-compose (i.e., we should be able to extract and enforce independence when many WIPOK executions take place concurrently), and **(b)** the WIPOK executions should remain secure (again, enabling extraction and independence) when run concurrently with an arbitrary δ-delayed protocol π. We separately explain, at an intuitive level, the security of the WIPOKs under these two types of composition.[19]

Composition with arbitrary δ-delayed protocols. The main problem that arises when running a secure protocol ρ concurrently to an arbitrary other protocol π, is that the adversary may be able to generate some dependence between π and the secure protocol ρ. (For example, π messages may have the same format as ρ messages and so an adversary can just forward messages from one protocol to another). On a more technical level, the proof of security works by constructing a hybrid-model simulator who runs π externally, while internally simulating ρ. Now, if the simulator needs to rewind ρ, it cannot proceed with π because the π-messages are sent to external parties and so cannot be retracted. Thus, it is crucial that while rewinding the WIPOKs in order to extract, the simulator does not need to send any π-messages externally. By setting δ to be the amount of time that it takes to complete a WIPOK, we have that the rewinding spans only over this amount of time. Thus, if π is δ-delayed, we have that no new π messages need to be dealt with during rewinding. We note that the length of the WIPOK is forced to be at most δ by timing-out if it takes too long. Thus, as described in the introduction, the needed effect is obtained by combining time-out and delay instructions together.

Concurrent self-composition. The main concern that arises here is that of *independence*. That is, when many WIPOK executions are run concurrently, the adversary can carry out a man-in-the-middle (or mauling) attack, in which it takes messages received in one execution and forwards (or modifies) them in another execution. Such a strategy enables it to "copy" a proof provided by an honest party, and contradicts the requirement of independence.

In order to prevent such an attack, it suffices to ensure that no (relevant) WIPOK in one session occurs concurrently with any (relevant) WIPOK in another session. However, in a setting where we cannot coordinate between multiple sessions of the protocol, this is impossible. We therefore have the parties prove *many* WIPOKs in every session, according to a carefully designed *scheduling strategy*. Our scheduling is based on the Chor-Rabin scheduling [18], with modifications necessary due to the fact that we work in the concurrent setting with timing (whereas they worked in the fully synchronous model). Our scheduling has the property that for every two sessions, there exists at least one WIPOK in the first session that does not overlap with *any* of the WIPOKs of the second session. We call a scheduling that has this property a *pairwise-disjoint scheduling*, and discuss it further in Sections 3.4 and 4.[20] We note that we make essential use of the timing assumptions in order to construct this scheduling.

Use of strong proofs of knowledge. We actually use *strong* proofs of knowledge in our protocol, rather than ordinary ones. (Recall that such a proof has the property that if the prover convinces

[19] We caution the reader that the formal proof of security does not separate out in this fashion.

[20] We remark that the Chor-Rabin scheduling was also used by [21] in a concurrent-type setting in order to achieve non-malleable commitments (without timing assumptions). Our setting differs in that we have many executions (and in this way it is "harder"), but we also utilize timing assumptions (and in this way it is "easier").

the verifier with non-negligible probability, then the extractor obtains a witness with overwhelming probability. Furthermore, the running-time of the extractor is independent of the probability that the prover convinces the verifier.) We do not know if this is essential, but we also do not know how to prove the security of our protocol otherwise.[21] Loosely speaking, we use strong proofs of knowledge in order to obtain the following effect. Our simulation strategy works by running in a "straight-line simulation mode" until a WIPOK is reached. When the beginning of such a proof is reached, we leave this mode and enter an "extraction mode," where rewinding takes place. We then run the extractor, while internally simulating the *future messages* (that is, the strategy is actually one of look-ahead, rather than rewinding back). Now, if a strong proof of knowledge is used, then after the extractor terminates, we are guaranteed that the following holds: either the extractor succeeded in obtaining a witness, or if it did not, we know that the prover will only succeed in convincing the verifier with negligible probability (in which case, we will not need the witness because the session will be aborted with all but negligible probability). Thus, there is no uncertainty (of course, beyond the negligible probability that the above will not hold). In contrast, in a regular proof of knowledge, such a look-ahead would fail because even if the extractor did not obtain a witness, it may still happen that the prover will convince the verifier. Thus, we would need to use a "rewind back" strategy where after the prover convinces the verifier, we would go back and obtain the witness. This type of strategy seems to be more difficult when dealing with the external π-messages (although, as mentioned above, we do not know whether or not the difficulties are inherent).

3.4 Scheduling

Our goal is to construct a protocol that securely realizes the $\mathcal{F}_{\text{CRS}}$ functionality in the timing model, in a general multiparty network where sessions are being executed concurrently. One of the major risks in this concurrent setting is related to the notion of *malleability*. Loosely speaking, this refers to an adversary who interleaves different executions of the protocol, and chooses its messages in one execution based on messages that it receives in the other executions. Consider, for example, many interleaved executions of a (regular, stand-alone) zero-knowledge proof of knowledge. In this setting, even if an adversary succeeds in convincing a verifier that it knows some secret s, it does not necessarily mean that the adversary actually knows s. Rather, it may be the case that there is some other party that is concurrently proving to the adversary that it knows the same secret s, and the adversary is simply relaying the messages between these two executions. Such a strategy is known as a "*man-in-the-middle*" attack. In order to construct secure protocols, it is necessary to prevent such attacks.

Our idea for preventing such mauling attacks is based on [18], who introduce a method for concurrently alternating and interleaving protocol executions in the *fully synchronous model*, while preserving independence. Loosely speaking, they construct an $O(\log n)$-round n-party protocol, in which each party (concurrently) carries out several zero-knowledge proofs sequentially, so that at least one of its proofs is "independent" from the proofs of the other parties.

More specifically, they associate with each party P_i a unique identifier $id^i \in \{0,1\}^{2m}$ that contains exactly m ones and m zeros (since the number of parties is polynomial in n, the value m can be set to be $O(\log n)$). The protocol consists of $2m$ phases, where in each phase some of the parties play the role of prover (and all parties verify). A party plays the prover in a zero knowledge proof in phase k if and only if the k^{th} bit of its identifier is 1 (i.e., party P_i will play the prover in

[21] This is the first work that we are aware of that utilizes strong proofs of knowledge in an essential way, rather than just in order to simplify the construction and proof.

phase k if and only if $(id^i)_k = 1$). In total, every party plays the prover's role during half of the phases, and for every two parties P_i and P_j, there is at least one phase in which P_i acts as a prover while P_j acts only as a verifier, and vice versa. This follows from the fact that for every $i \neq j$, id^i and id^j are distinct and they both have the same number of ones and zeros. Therefore, there exist two distinct indices k and k' such that: **(a)** $(id^i)_k = 1$ and $(id^j)_k = 0$, and **(b)** $(id^i)_{k'} = 0$ and $(id^j)_{k'} = 1$. Thus, in phase k party P_i proves and party P_j only verifies, and in phase k' party P_j proves and party P_i only verifies. Intuitively, this prevents P_i from using P_j as an oracle for supplying its proofs. As we explain below, even though this method seems to guarantee only pairwise independence, it actually achieves overall independence. We show that a similar idea can be used to achieve independence in a concurrent setting, in the timing model.

To this end we define the notion of a *pairwise disjoint* scheduling. In Section 4 we show how to construct a pairwise disjoint scheduling in the timing model. In Section 3.5, we show how such a scheduling can be used to design a protocol that securely realizes the $\mathcal{F}_{\text{CRS}}$ functionality under concurrent general composition with delays, in the timing model.

Pairwise-disjoint scheduling. Consider one pre-specified protocol σ, which needs to be executed concurrently in many different sessions, where each session has a unique identifier. The aim of a pairwise-disjoint scheduling is to ensure that different concurrent executions of σ are somewhat "independent". Intuitively, the idea is to achieve independence by requiring the parties to act as follows: Instead of running a single execution of Protocol σ in a given session, the parties execute σ several times in that session according to some pre-specified "*pairwise-disjoint*" scheduling S. Loosely speaking, this scheduling ensures that when looking at any two distinct sessions (each containing at least one honest party), there exists at least one execution in *each* of the sessions that does not intersect (i.e., overlap) with *any* execution in the other session.

We define the notion of a *pairwise-disjoint scheduling algorithm* S that receives for input a protocol σ, a unique session identifier sid, and the network timing assumptions Δ and ϵ. The algorithm $S(\sigma, sid, \Delta, \epsilon)$ then outputs a schedule consisting of many executions of σ with the property that for every two distinct sessions sid and sid' there exists at least one execution in $S(\sigma, sid, \Delta, \epsilon)$ that does not overlap with any of the executions of $S(\sigma, sid', \Delta, \epsilon,)$ and vice versa. We stress a crucial point here. When considering many different sessions, it may be the case that *every* execution of σ in a session sid overlaps with some other execution of σ in some other session. However, it is guaranteed that for every session sid', there exists at least *one* execution of σ in session sid that does not overlap with *any* of the executions in session sid'. This type of pairwise disjointness suffices since in our proof the simulator simulates all the honest provers except for one chosen prover which will be an "external prover." It is only this "external prover" that cannot be rewound. Thus, it suffices to ensure that for each session there exists one execution which does not overlap with the proofs of the "external prover." This is exactly what a pairwise disjoint scheduling ensures.

In what follows, we formally define the syntax of a scheduling algorithm. We are only interested in schedules which are polynomial-time[22] and non-trivial (where the parties output time-out only if the network delay is too long). We therefore incorporate these requirements into the definition.

Definition 7 (non-trivial scheduling algorithm): *A* non-trivial scheduling algorithm *is an algorithm S that receives for input a protocol σ, a session identifier sid, and a pair (Δ, ϵ), and outputs a schedule Σ consisting of polynomially many executions of σ together with* delay *and* time-out *instructions that are polynomial in Δ and ϵ. Furthermore Σ is non-trivial (as defined in Definition* 5*).*

[22]This means that the number of executions of σ in the schedule is polynomial in Δ, ϵ and n, and also that all honest parties will finish the schedule (either terminate normally or abort or time-out) within time which is polynomial. In particular, due to the non-triviality condition, this requires that all the delays are for polynomial durations.

Before proceeding further, we define what it means for an execution of a protocol σ to *overlap* with another execution. Let σ_1 and σ_2 be two executions of Protocol σ, and let P_1 and P_2 be any two honest participants in σ_1 and σ_2 respectively. Then, σ_1 and σ_2 overlap according to P_1 and P_2, if P_1 sends a σ_1 message *after* P_2 has sent its first σ_2 message but *before* P_2 sends its last σ_2 message, or if P_2 sends a σ_2 message *after* P_1 has sent its first σ_1 message but *before* P_1 sends its last σ_1 message. This is shown in Figure 3 (a) and in Figure 3 (b), respectively. Therefore, if σ_1 and σ_2 overlap according to P_1 and P_2 then there is a message (of σ_1 or of σ_2) that was sent while P_1 was executing σ_1 and while P_2 was executing σ_2. Notice that the notion of overlapping is defined with respect to a pair of parties. This is due to the fact that parties do not necessarily begin and conclude executions at the same time in an asynchronous network (and so σ_1 and σ_2 may not overlap according to some pairs, and may overlap according to others). We therefore always refer to overlapping *according to a specified pair of parties.* We are now ready to define what it means for a schedule to be *pairwise-disjoint.*

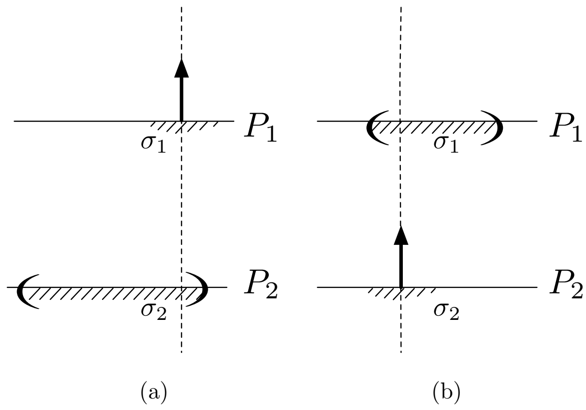

Figure 3: A schematic representation of the two situations when two WIPOKs σ_1 and σ_2 are said to overlap with respect to two parties P_1 and P_2, respectively. The horizontal timelines indicate the progress of the execution. In situation (a), the upward pointing bold arrow represents a message in σ_1 sent out by P_1. In the lower timeline for P_2, the hashed region between the parenthesis indicates when P_2 is executing σ_2. The vertical dotted line indicates that while a message in σ_1 was sent out by P_1, P_2 was executing σ_2. In situation (b), the roles of P_1 and P_2 are reversed.

Definition 8 (non-trivial pairwise-disjoint scheduling): *The schedule output by a non-trivial scheduling algorithm S is said to be* pairwise-disjoint *if it has the following property: Let $sid_1 \neq sid_2$ be any identifiers of the same length and assume that $\Sigma_1 = S(\sigma, sid_1, \Delta, \epsilon)$ and $\Sigma_2 = S(\sigma, sid_2, \Delta, \epsilon)$ are run in a network with an ϵ-drift preserving adversary, such that both sid_1 and sid_2 have at least one honest participant each. Then for any two honest parties P_1 and P_2 in sessions sid_1 and sid_2 respectively, there exists an execution of σ in Σ_2, denoted by σ_2, such that for every execution σ_1 in Σ_1, it holds that σ_1 and σ_2 do not overlap according to P_1 and P_2.*

Note that if P_2 times-out session sid_2 before some execution σ^* in $S(\sigma, sid_1, \Delta, \epsilon)$ was initiated, then in particular σ^* does not overlap with any execution in $S(\sigma, sid_2, \Delta, \epsilon)$, according to P_1 and P_2. This fact will be used in the proof of Theorem 12 in Section 4.

In Section 4, we prove the following theorem which will be used in order to construct our protocol for securely realizing the $\mathcal{F}_{\text{CRS}}$ functionality.

Theorem 9 *There exists a non-trivial pairwise-disjoint scheduling algorithm for any protocol σ, any network delay Δ, any clock-drift ϵ such that $1 \leq \epsilon < \sqrt[3]{1.5}$, and any set of identifiers $sid \in \{0,1\}^{\text{poly}(n)}$. Furthermore, in each execution of σ in the scheduling, honest parties output* time-out *if the execution runs for longer than $\Delta \cdot \mathsf{rounds}(\sigma)$ time units according to the party's local clock, where $\mathsf{rounds}(\sigma)$ denotes the number of rounds of communication in an execution of σ.*

Before proceeding, we explain again (but in more detail) why *pairwise* disjointness suffices. In our protocol, we use a pairwise disjoint scheduling for WIPOKs. Then, at some stage in our proof of security of the protocol, we focus on a single session *sid*, and argue that the simulation (i.e., the extraction from WIPOK proofs) in all sessions $sid' \neq sid$ can be carried out without rewinding during any WIPOK of session *sid*. This can be achieved because the pairwise disjointness property of the schedule guarantees that for every session *sid'*, there exists at least one WIPOK in *sid'* that does not overlap with any WIPOK in *sid*. We can therefore extract from the non-overlapping WIPOK in *sid'* without rewinding any of the WIPOK proofs in *sid*. Since this is true for all sessions $sid' \neq sid$, we are able to simulate without rewinding any WIPOK proof in *sid*, as required.

3.5 The Protocol for $\mathcal{F}_{\text{CRS}}$

The protocol below refers to a one-way function f and a commitment scheme C. We denote by $C(r; s)$ a commitment to r using random coins s. For simplicity, we use a non-interactive commitment scheme. Such schemes are known to exist assuming the existence of 1–1 one-way functions. Our protocol uses a broadcast primitive, and we assume that *all* messages are sent over this channel. We remark that in the case that output delivery is not guaranteed (as in our model here), broadcast that is secure under concurrent general composition can be easily implemented in a standard point-to-point network [32].

As was mentioned in Subsection 3.3, the protocol is based on a natural extension of the coin-tossing protocol of [35] to the multiparty setting, with the following high-level differences. First, instead of using just any zero-knowledge proof of knowledge, we use the zero-knowledge arguments of knowledge of [25] that are constructed from two witness-indistinguishable proofs of knowledge.[23] Second, we use *strong* proofs of knowledge, rather than "ordinary" ones, so that if the prover convinces an honest verifier with non-negligible probability, a witness can be extracted with overwhelming probability in polynomial time. Third, some of these strong proofs of knowledge are given according to a pairwise disjoint scheduling. We now present the protocol.

Protocol $\boldsymbol{\rho}$ (protocol for realizing the $\mathcal{F}_{\text{CRS}}$ functionality in a general multiparty network, assuming time bounds Δ and ϵ):

- **Participating Parties:** $P_1, \ldots, P_k$ (some subset of the parties in the entire network).
- **Common Input:** the security parameter n, a session identifier $sid \in \{0,1\}^m$, and global constants Δ and ϵ.
- **The Protocol:** The protocol proceeds in three phases.
  - Phase One:
    1. Each party P_i chooses a pair of values $w_1^i, w_2^i \in_R \{0,1\}^n$, and computes $v_1^i = f(w_1^i), v_2^i = f(w_2^i)$, where f is a one-way function.

[23] We note that looking at our protocol it is not clear that we use the zero-knowledge proof of knowledge of [25], since the two witness-indistinguishable proofs of knowledge appear in different phases of the protocol, and moreover, we use the first witness-indistinguishable proof of knowledge for two different zero-knowledge proofs. Thus, our protocol does not exactly follow the syntax of [25] though the concept is similar.

2. Each party P_i proves independently to all other parties that it knows $f^{-1}(v_1^i)$ or $f^{-1}(v_2^i)$. Formally, P_i proves that it knows a witness for the relation

$$R_1^i \stackrel{\text{def}}{=} \{((v_1^i, v_2^i), w) \mid v_1^i = f(w) \text{ or } v_2^i = f(w)\}.$$

The proofs are given according to some arbitrary order; say the party with the smallest ID proves first, then the party with the second to smallest ID, and so on.[24] Each P_i carries out a proof that has the following properties:

(a) The proof is an $\alpha(n)$-round witness-indistinguishable strong proof of knowledge, for some pre-specified super-constant function $\alpha(\cdot)$.[25] (Henceforth, we denote this proof by WISPOK, for short).

(b) The proof is carried out in a parallel manner. That is, P_i sends the first message of the proof to all other parties. It then waits for the responses from all the parties, and only then sends the second message to all the parties, and so on.

(c) The first and the last messages of the proof are sent by the verifier. (This convention is for convenience.)

We let σ denote such a proof system (i.e., σ is a protocol consisting of $k-1$ WISPOKs in which a single party P_i gives a proof in parallel to all other parties P_j). Each party P_i repeats this proof σ several times, according to *any* non-trivial pairwise-disjoint scheduling $\Sigma = S(\sigma, sid, \Delta, \epsilon)$ in which each execution of σ is timed-out by honest parties if more than $\tau \stackrel{\text{def}}{=} \alpha(n) \cdot \Delta$ units of time pass on their local clocks. (The existence of such a scheduling is guaranteed in Theorem 9.) Denote by $\Sigma(i)$ the schedule in which P_i plays the prover. Then, the protocol works by first running $\Sigma(1)$, then $\Sigma(2)$ and so on until $\Sigma(k)$ (assuming party P_1 has the smallest ID, party P_2 has the second to smallest ID, and so on).

If a party P_i receives a time-out message in any execution of σ in any schedule $\Sigma(j)$, then it broadcasts time-out to all the parties, outputs $(\textsf{time-out}, sid)$ and halts. Any party receiving such a time-out message also outputs $(\textsf{time-out}, sid)$ and halts.

- Phase Two: Each party P_i operates as follows.

1. Party P_i chooses $r_i \in_R \{0,1\}^{p(n)}$ and broadcasts a commitment $c_i = C(r_i; s_i)$ to all the parties, where C is a perfectly binding commitment scheme and s_i is a random string. P_i waits for the commitments from all other parties to arrive before proceeding.

2. Party P_i proves in parallel to every other party P_j that it knows either $f^{-1}(v_1^j)$ or $f^{-1}(v_2^j)$ or a pair (r_i, s_i) such that $c_i = C(r_i; s_i)$, using an $\alpha(n)$-round witness-indistinguishable strong proof of knowledge. Formally, P_i proves that it knows a witness for the relation

$$R_2^{i,j} \stackrel{\text{def}}{=} \{((v_1^j, v_2^j, c_i), (w, r, s)) \mid v_1^j = f(w) \text{ or } v_2^j = f(w) \text{ or } c_i = C(r; s)\}.$$

Time-Out Mechanism: For every proof that party P_i participated in (either as a prover or as a verifier), it checks that no more than $\tau = \alpha(n) \cdot \Delta$ local time units have passed from the time that the proof began until the time that it ended. If more time passed, then P_i broadcasts time-out to all the parties, outputs $(\textsf{time-out}, sid)$ and halts the execution. Any party receiving such a time-out message also outputs $(\textsf{time-out}, sid)$ and halts the execution.

[24] Note that by requiring the proofs to be given sequentially we automatically obtain "independence" between proofs that belong to the *same* session.

[25] Recall that such proofs are known to exist for *any* super-constant function $\alpha(\cdot)$.

3. Once Party P_i finished its proof and verified the proofs of all other parties, it broadcasts a $\mathsf{Phase2over}$ message to all other parties. It then waits for the same message to arrive from all other parties before proceeding. After this it will never output $(\mathsf{time\text{-}out}, sid)$.

We remark that the proofs in this phase can be scheduled in any way.

- Delay Mechanism: Before continuing to Phase 3, each party P_i waits $\tau\epsilon$ local time units.
- Phase Three: Party P_i broadcasts r_i to all other parties (without decommitting) and, using a (regular, not necessarily a strong) 3-round witness indistinguishable proof of knowledge, proves in parallel to every other party P_j that it either knows a preimage for one of v_1^j, v_2^j or that it knows s such that $c_i = C(r_i; s)$. Formally, P_i proves in parallel that it knows a witness for the relation

$$R_3^{i,j} \stackrel{\text{def}}{=} \{((v_1^j, v_2^j, c_i, r_i), (w, s)) \mid v_1^j = f(w) \text{ or } v_2^j = f(w) \text{ or } c_i = C(r_i; s)\}.$$

We remark that the proofs in this phase can be scheduled in any way.

- Each party P_i defines $R = r_1 \oplus r_2 \oplus \ldots \oplus r_k$, where r_j is the string it received in the previous step from party P_j, and r_i is the string that it broadcasted to all other parties.[26]
- **Output:** Each party outputs (sid, R).

This completes the description of the protocol.

Conventions. If an honest party receives a message that does not have a valid format or if it rejects a proof that it verifies, then the party broadcasts an abort message to all other parties and halts the execution.[27] Any party receiving such an abort message also halts the execution. We also assume that all messages are sent together with the session identifier sid, which is part of the common input. This enables the correct assignment of messages to their intended sessions. We stress that the security of the protocol does not rely on this assignment being correct. Rather, this mechanism just ensures successful termination when honest parties interact.

3.6 Proof of Security

We now show that Protocol ρ securely realizes the $\mathcal{F}_{\text{CRS}}$ functionality in the timing model, even when run many times concurrently with an arbitrary other protocol π, *as long as* all the messages in π are delayed by $\tau\epsilon$ local time units, where $\tau = \alpha(n) \cdot \Delta$.[28] In other words, Protocol ρ securely realizes $\mathcal{F}_{\text{CRS}}$ under concurrent general composition with $\tau\epsilon$-delays in the timing model with ϵ-drift. As we have seen in Section 3.2, this (along with the non-triviality condition) suffices for proving Theorem 6.

Theorem 10 *Let Δ and ϵ be fixed constants, such that $1 \leq \epsilon < \sqrt[3]{1.5}$, and let $\tau = \alpha(n) \cdot \Delta$. Then, assuming the existence of 1–1 one-way functions, Protocol ρ securely realizes the $\mathcal{F}_{\text{CRS}}$ functionality under concurrent general composition with $\tau\epsilon$-delays (as defined in Definition 1) in the timing model with ϵ-drift, and in the presence of static malicious adversaries. Furthermore, Protocol ρ is non-trivial under timing assumptions (Δ, ϵ).*

[26] Note that since all the r_i's were broadcast it must be the case that all the honest parties have the same R.

[27] Recall that when a party times-out it behaves differently. Namely, it does not send an abort message, but rather sends a $\mathsf{time\text{-}out}$ message.

[28] Recall that $\alpha(n)$ is the number of rounds in the WISPOKs of Protocol ρ.

Proof: We begin by proving that Protocol ρ is non-trivial. In order to see this, notice that in ρ an honest party outputs a time-out message only if a WISPOK takes more than $\tau = \alpha(n) \cdot \Delta$ local time units or if the (pairwise-disjoint) schedule instructs it to time-out. Since the WISPOKs consist of $\alpha(n)$ rounds, if the latency of the network is at most Δ (according to *all* local clocks) then each WISPOK will conclude within at most $\tau = \alpha(n) \cdot \Delta$ local time units (recall that we assume that local computation is instantaneous). This together with the fact that the schedule used in ρ is non-trivial, implies that ρ is non-trivial.

We now proceed to prove the security of Protocol ρ. Let Δ and ϵ be any fixed constants such that $1 \leq \epsilon < \sqrt[3]{1.5}$. Let $\pi_{\tau\epsilon}$ be an arbitrary $\tau\epsilon$-delayed multiparty protocol that may contain ideal calls to the $\mathcal{F}_{\text{CRS}}$ functionality. Let $\mathcal{A}$ be any static non-uniform probabilistic polynomial-time ϵ-drift preserving adversary that runs protocol π^ρ in the timing model. We begin by describing the hybrid-model simulator $\mathcal{S}$ that runs π in the $\mathcal{F}_{\text{CRS}}$-hybrid model (in a model without time).

The simulator $\mathcal{S}$ simulates the real-model adversary $\mathcal{A}$ internally. The aim of $\mathcal{S}$ is to force the output of the coin tossing protocol ρ in any given session to equal the common random string obtained from the $\mathcal{F}_{\text{CRS}}$ functionality.

In order to force a coin-tossing session to output some given random string R_{CRS}, in each session $\mathcal{S}$ will do the following: for every corrupted party P_j, it will *extract* from $\mathcal{A}$ both a value w_j such that $f(w_j) = v_1^j$ or $f(w_j) = v_2^j$, and a value r_j, which is the decommitment of c_j (sent by P_j in the beginning of Phase 2). These values will be extracted before entering Phase 3 of this session. Then, Phase 3 will be simulated in a "straight-line" manner: $\mathcal{S}$ will simulate each honest party P_i sending a random r_i such that $R_{\text{CRS}} = \oplus_{l=1}^{k} r_l$, and proving to each party P_j that r_i was committed to (even though it was not) using the previously extracted witness w_j.

Thus, the simulation by $\mathcal{S}$ consists of an extraction mode and a straight-line simulation mode. The "rewinding" takes place only when $\mathcal{S}$ is in the extraction mode (the rest of the simulation is "straight-line"). In the extraction mode $\mathcal{S}$ rewinds $\mathcal{A}$ internally, without rewinding the simulated protocol. That is, $\mathcal{S}$ pauses the simulation, and internally creates a copy of its simulated world. Then the extraction is carried out in a look-ahead manner. That is, $\mathcal{S}$ (forward) simulates the messages that the honest parties in ρ will send to $\mathcal{A}$ after the paused point, and then rewinds this simulation. The time constraints ensure that messages from π (that are sent externally by $\mathcal{S}$) never have to be sent while $\mathcal{S}$ is in the extraction mode, where rewinding takes place. We now formally describe $\mathcal{S}$.

The simulator $\mathcal{S}$: $\mathcal{S}$ is a "hybrid-world" adversary, that interacts with the parties running protocol π in the $\mathcal{F}_{\text{CRS}}$-hybrid model (without time). The aim of $\mathcal{S}$ is to create the same effect in the $\mathcal{F}_{\text{CRS}}$-hybrid model, as the real-model adversary $\mathcal{A}$ does in a real execution of π^ρ (in a model with time).

As was previously mentioned, $\mathcal{S}$'s operations consist of two modes of operation: straight-line simulation mode and extraction mode. $\mathcal{S}$ starts and ends in the straight-line simulation mode, but frequently leaves it and enters the extraction mode. In the straight-line mode, $\mathcal{S}$ interacts with the honest parties (in the $\mathcal{F}_{\text{CRS}}$-hybrid model), while updating internally simulated states of the adversary $\mathcal{A}$ and of the honest parties running the program for the protocol ρ. In the extraction-mode, these simulated states are frozen, while $\mathcal{S}$ applies an extraction subroutine. The output of the extraction subroutine will be needed for continuing the straight-line mode.

We shall denote by $\mathcal{P}_i^{sid}$ the simulated ρ program of an honest party P_i for a session sid, which $\mathcal{S}$ uses in the straight-line simulation mode. The simulation enters the extraction mode every time that $\mathcal{P}_i^{sid}$ is about to take part as a verifier in one of the WISPOKs given by a corrupted player in the protocol. In this extraction mode S calls the extraction subroutine. This subroutine will try to find the witness used by the corrupted prover in that WISPOK. The straight-line simulation continues when the extraction subroutine returns.

We proceed to define the simulator by first describing the straight-line simulation mode and then describing the extraction subroutine.

STRAIGHT-LINE MODE: $\mathcal{S}$ internally runs $\mathcal{A}$, and for each honest party, $\mathcal{S}$ simulates the various tapes that $\mathcal{A}$ expects to have access to (namely, the communication tapes and the clock-tape). It also maintains simulated states (i.e., the work-tape) of the ρ programs of the honest parties. As was mentioned above, we denote by $\mathcal{P}_i^{sid}$ the program simulated by $\mathcal{S}$, corresponding to the ρ program of an honest party P_i in session sid. The simulated programs $\mathcal{P}_i^{sid}$ communicate directly with $\mathcal{A}$.

In addition, $\mathcal{S}$ needs to let the π program of the external honest parties communicate with $\mathcal{A}$ (here we mean the real parties with whom $\mathcal{S}$ interacts in the hybrid model). For π messages from $\mathcal{A}$ to a party P_i, this is done simply by sending the messages out to P_i (i.e., they are copied onto P_i's incoming message tape). However, upon receiving a π message from an external honest party P_i, simulator $\mathcal{S}$ needs to simulate the delay of P_i before forwarding it to $\mathcal{A}$ (because in the hybrid world π messages are sent out without any delay, in contrast to the real world). Therefore, $\mathcal{S}$ waits $\tau\epsilon$ time units according to P_i's simulated local clock before sending the received π-message to $\mathcal{A}$. Finally, $\mathcal{S}$ generates the same input-output as $\mathcal{A}$. More formally:

- Whenever a session sid with parties $P_{i_1}, \ldots, P_{i_k}$ is begun, $\mathcal{S}$ sends $(\mathsf{crsgen}, sid, \{P_{i_1}, \ldots, P_{i_k}\})$ to the $\mathcal{F}_{\text{CRS}}$ functionality.[29] We assume that at least one party in $\{P_{i_1}, \ldots, P_{i_k}\}$ is honest, since the case that all parties are corrupted is trivial.
- $\mathcal{S}$ initiates the program $\mathcal{P}_i^{sid}$ corresponding to each honest participant P_i.
- If at any point $\mathcal{P}_i^{sid}$ outputs $(\mathsf{time\text{-}out}, sid)$, $\mathcal{S}$ sends $(\mathsf{time\text{-}out}, sid)$ to the $\mathcal{F}_{\text{CRS}}$ functionality and delivers the message $(\mathsf{time\text{-}out}, sid)$ from the $\mathcal{F}_{\text{CRS}}$ functionality to P_i.
- Whenever $\mathcal{A}$ sends a π message to some party P_i, $\mathcal{S}$ sends the π message externally to party P_i.
- Whenever $\mathcal{S}$ (externally) receives a π message from some honest party P_i, it stores the message in an internal delay buffer. Then, after $\tau\epsilon$ time units according to P_i's internally simulated local clock, it forwards the π message to $\mathcal{A}$.
- For all except one honest party, $\mathcal{P}_i^{sid}$ runs exactly the program specified by the protocol ρ. We denote the index of this one chosen honest party by $\text{H}(sid)$; when sid is clear from the context we shall write H instead of $\text{H}(sid)$.[30] The program $\mathcal{P}_{\text{H}(sid)}^{sid}$ is identical to ρ in Phases 1 and 2, and differs from ρ only in Phase 3. We shall describe the differences shortly.
- In Phases 1 and 2, when $\mathcal{P}_{\text{H}}^{sid}$ receives the first message of a WISPOK in which it plays *verifier* and a corrupted party plays *prover*, $\mathcal{S}$ applies the extraction subroutine (to be defined later) to that WISPOK. The output of the extraction subroutine is recorded for later reference. (Note that extraction is only carried out when $\mathcal{P}_{\text{H}}^{sid}$ plays the verifier.)
- At the point that $\mathcal{P}_{\text{H}}^{sid}$ enters Phase 3 of the protocol, $\mathcal{S}$ carries out the following two checks:

[29] Actually, this crsgen message to the functionality may have been sent by one of the honest parties participating in session sid. This is inconsequential.

[30] This honest party can be arbitrarily chosen – say, the one with the "smallest" identity among all honest participants.

1. $\mathcal{S}$ checks the output of the extraction subroutine applied to each of the Phase 1 and Phase 2 WISPOKs given to $\mathcal{P}_{\mathrm{H}}^{sid}$ by a corrupted party. If in any of them, the extraction subroutine failed to extract a valid witness for the statement of that WISPOK, then $\mathcal{S}$ outputs fail_1 and halts.
2. Let $(v_1^{\mathrm{H}}, v_2^{\mathrm{H}})$ be the first message that $\mathcal{P}_{\mathrm{H}}^{sid}$ sends in session sid. Recall that w such that $f(w) \in \{v_1^{\mathrm{H}}, v_2^{\mathrm{H}}\}$ is a valid witness in all of the Phase 2 WISPOKs that $\mathcal{P}_{\mathrm{H}}^{sid}$ verifies. If the extraction subroutine, applied to any of the Phase 2 WISPOKs given to $\mathcal{P}_{\mathrm{H}}^{sid}$ by a corrupted party outputs such a w as a witness, then $\mathcal{S}$ outputs fail_2 and halts.

Note that if $\mathcal{S}$ did not output fail_1 then for every corrupted party P_j, the extraction subroutine applied to each of the Phase 1 WISPOKs given by P_j must have returned w^j such that $f(w^j) \in \{v_1^j, v_2^j\}$. Furthermore, for every honest party P_i, $\mathcal{S}$ can look up such a w^i from $\mathcal{P}_i^{sid}$ (because $\mathcal{S}$ runs the code of P_i internally).

Similarly, if $\mathcal{S}$ did not output fail_2 either, then for every corrupted party P_j, the extraction subroutine, applied to each of the Phase 2 WISPOKs given by P_j, must have produced the witness (r_j, s_j) such that $c_j = C(r_j; s_j)$, where c_j is the commitment sent by P_j in Phase 2 of session sid (recall that the only valid witnesses for this WISPOK are either the above mentioned witness (r_j, s_j) or w such that $f(w) \in \{v_1^{\mathrm{H}}, v_2^{\mathrm{H}}\}$, where extraction of the latter witness results in fail_2). In addition, for every honest party P_i, $\mathcal{S}$ can look up r_i from $\mathcal{P}_i^{sid}$ (again, because $\mathcal{S}$ runs $\mathcal{P}_i^{sid}$).

Thus, if the above two checks passed (namely, $\mathcal{S}$ did not output fail_1 or fail_2) then $\mathcal{S}$ has obtained values w^i and r_i, *for all participants* P_i. (As we will see, for all $i \neq \mathrm{H}$, the values w^i and r_i will be needed by $\mathcal{S}$ to continue the simulation.)

- If the above two checks passed then $\mathcal{S}$ acts as follows:
  1. $\mathcal{S}$ sends $(sid, \mathsf{compute})$ to the $\mathcal{F}_{\mathrm{CRS}}$ functionality, and receives (sid, R_{CRS}) in response.[31] Then using the r_i values as given above, $\mathcal{S}$ computes
  $$r = R_{\mathrm{CRS}} \oplus \left(\bigoplus_{i \neq \mathrm{H}} r_i \right). \tag{3}$$
  2. $\mathcal{S}$ hands r (from Eq. (3)) and $\{w^i\}_{i \neq \mathrm{H}}$ to $\mathcal{P}_{\mathrm{H}}^{sid}$.
- $\mathcal{P}_{\mathrm{H}}^{sid}$ proceeds with the simulation of Phase 3. (Notice that its instructions here differ from the program specified by ρ for the honest parties.)
  1. In the beginning of Phase 3, $\mathcal{P}_{\mathrm{H}}^{sid}$ does not send the value r_{H} that it committed to in Phase 2, as instructed by protocol ρ. Rather, it sends the value r given to it by $\mathcal{S}$.
  2. After sending r, $\mathcal{P}_{\mathrm{H}}^{sid}$ proves to each party P_j (in the WIPOK of Phase 3) that this "fake" value r is the value that it committed to in Phase 2. This is done using the alternative witness w^j given to it by $\mathcal{S}$.[32]

[31] At this point of the protocol it is guaranteed that no honest party has or will output $(\mathsf{time\text{-}out}, sid)$, because they must all have sent $\mathsf{Phase2over}$ messages (since $\mathcal{P}_{\mathrm{H}}^{sid}$ entered Phase 3). Hence it is possible for $\mathcal{S}$ to send a $(sid, \mathsf{compute})$ message to $\mathcal{F}_{\mathrm{CRS}}$, thereby receiving back (sid, R_{CRS}).

[32] Its ability to use the "fake" value $r = R_{\mathrm{CRS}} \oplus \left(\bigoplus_{i \neq \mathrm{H}} r_i\right)$, rather than the value that it committed to, is exactly what allows the output of this session to equal R_{CRS}. Note that in order to use this "fake" value r it must know all

- If there exists a (corrupted) party P_j that broadcasted $r'_j \neq r_j$ in the beginning of Phase 3 and $\mathcal{P}_{\mathrm{H}}^{sid}$ accepts its Phase 3 WIPOK, then $\mathcal{S}$ outputs fail_3.

- For every honest party P_i, if $\mathcal{P}_i^{sid}$ outputs (sid, R) then $\mathcal{S}$ delivers the message (sid, R_{CRS}) from $\mathcal{F}_{\mathrm{CRS}}$ to P_i.[33]

This completes the description of the simulator except for the extraction subroutine.

THE EXTRACTION SUBROUTINE: Recall that in Phases 1 and 2, when $\mathcal{P}_{\mathrm{H}}^{sid}$ receives the first message of a WISPOK in which it plays verifier and a corrupted party P_j plays prover, $\mathcal{S}$ calls the extraction subroutine. (We shall denote such a WISPOK by $\textsc{wispok}_j^{sid}$.) The extraction subroutine will try to extract a witness for the statement of $\textsc{wispok}_j^{sid}$ by constructing a stand-alone prover $\mathcal{Q}_j^{sid}$ from $\mathcal{A}$, and then applying the strong proof of knowledge extractor to $\mathcal{Q}_j^{sid}$. The stand-alone prover $\mathcal{Q}_j^{sid}$ is defined as follows.

$\mathcal{Q}_j^{sid}$ is a stand-alone (cheating) prover who proves a single strong proof of knowledge to an external verifier. $\mathcal{Q}_j^{sid}$ simulates $\mathcal{A}$ exactly like $\mathcal{S}$ does, continuing from the point after the extraction subroutine is invoked, except for the following differences:

- In $\mathcal{S}$, the program $\mathcal{P}_{\mathrm{H}}^{sid}$ plays the verifier of the WISPOK: i.e., it receives the WISPOK messages from a prover P_j, and responds to them as a verifier. Instead, in $\mathcal{Q}_j^{sid}$, the program $\mathcal{P}_{\mathrm{H}}^{sid}$ relays out the incoming WISPOK messages from P_j to an external verifier. When it receives a response from the external verifier, it forwards it internally to P_j as its own response.

- Since $\mathcal{Q}_j^{sid}$ is a stand-alone prover, unlike $\mathcal{S}$, it cannot interact with the honest parties running the protocol π in the hybrid world. So all messages generated by $\mathcal{S}$ for these parties are ignored. Furthermore, there are no incoming messages from the π protocol. However the messages that arrived earlier and were stored internally in the delaying buffers of $\mathcal{S}$ will be used just like $\mathcal{S}$ did originally. (As we will see later, this suffices and no "new" π messages are needed.)

  Note also that since $\mathcal{Q}_j^{sid}$ is a stand-alone prover, it cannot interact with the different instances of $\mathcal{F}_{\mathrm{CRS}}$. However, these can all be perfectly simulated internally by $\mathcal{Q}_j^{sid}$.

- $\mathcal{Q}_j^{sid}$ does not invoke the extraction subroutine that $\mathcal{S}$ invokes. Instead, when the extraction subroutine needs to be called, it is just assumed to return $\perp$ (this ensures that $\mathcal{Q}_j^{sid}$ is well-defined).

- $\mathcal{Q}_j^{sid}$ halts as soon as it receives an accept or reject message from the outside verifier. Also, if $\mathcal{P}_{\mathrm{H}}^{sid}$'s local clock reaches a time where the original $\mathcal{P}_{\mathrm{H}}^{sid}$ would have timed-out, then $\mathcal{Q}_j^{sid}$ halts.

The key point to notice is that $\mathcal{Q}_j^{sid}$ is a stand-alone adversary who proves a single strong proof of knowledge to an external verifier. The extraction subroutine applies the strong knowledge

the alternative witnesses $\{w^i\}_{i \neq \mathrm{H}}$, which is why $\mathcal{S}$ must apply the extraction subroutine to the WISPOKs of Phase 1. The reason $\mathcal{S}$ must apply the extraction subroutine to the WISPOKs of Phase 2 is in order to obtain all the values $\{r_i\}_{i \neq \mathrm{H}}$, which are needed in order to determine the "fake" value $r = R_{\mathrm{CRS}} \oplus \left(\bigoplus_{i \neq \mathrm{H}} r_i\right)$.

[33]Notice that if $\mathcal{P}_i^{sid}$ produced an output (sid, R) then it must be the case that $R = R_{\mathrm{CRS}}$. If $R \neq R_{\mathrm{CRS}}$ then there exists a j such that $r_j \neq r'_j$ (follows from the fact that $R = r'_1 \oplus \cdots \oplus r'_k$ and $R_{\mathrm{CRS}} = r_1 \oplus \cdots \oplus r_k$). In this case, either $\mathcal{S}$ outputs fail_3 or $\mathcal{P}_{\mathrm{H}}^{sid}$ does not accept the Phase 3 WIPOK of P_j and thus will halt the execution. In both cases $\mathcal{P}_i^{sid}$ would not produce an output.

extractor K to the prover $\mathcal{Q}_j^{sid}$ (recall that if $\mathcal{Q}_j^{sid}$ convinces an honest verifier V in the proof with probability greater than $\mu(n)$ for some negligible function μ, then K obtains a witness with probability at least $1-\mu(n)$).

This completes the description of the extraction subroutine. Note that the extraction subroutine is invoked on all the Phase 1 and Phase 2 WISPOKs given to $\mathcal{P}_{\mathrm{H}}^{sid}$ by any corrupted party P_j in any session *sid* (with at least one honest player). This ensures that if the WISPOKs convince $\mathcal{P}_{\mathrm{H}}^{sid}$ with non-negligible probability, then the simulator will obtain the corresponding witnesses with overwhelming probability, by applying the extraction subroutine. (Of course, this is the case assuming that $\mathcal{Q}_j^{sid}$ convinces the verifier with essentially the same probability that $\mathcal{P}_{\mathrm{H}}^{sid}$ is convinced. This will be proven below.)

Proof of the simulation. First note that $\mathcal{S}$ runs in strict polynomial-time if $\mathcal{A}$ runs in strict polynomial-time (because the knowledge extractor of a strong proof of knowledge runs in strict polynomial-time, and the only rewinding carried out by $\mathcal{S}$ is in applying the knowledge extractor). We now prove that the output distribution of $\mathcal{S}$ and the honest parties running π in the $\mathcal{F}_{\mathrm{CRS}}$-hybrid model is computationally indistinguishable from the output distribution of an ϵ-drift preserving adversary $\mathcal{A}$ and the honest parties in a real execution of Protocol π^ρ in the timing model. In order to prove this, we first show that $\mathcal{S}$ outputs a fail message with negligible probability. Given this, we then introduce hybrid experiments which bridge the difference between the $\mathcal{F}_{\mathrm{CRS}}$-hybrid execution and the real execution, to prove the claimed indistinguishability.

We now prove that $\mathcal{S}$ outputs a fail message with at most negligible probability. Recall that there are three types of failures: fail_1, fail_2 and fail_3. Intuitively, fail_1 occurs if there exists a WISPOK for which the extraction subroutine fails to output a corresponding witness, and yet $\mathcal{P}_{\mathrm{H}}^{sid}$ accepts the WISPOK. fail_2 occurs if the extraction subroutine, applied to any of the Phase 2 WISPOKs, outputs the "wrong witness;" i.e., instead of extracting the committed value r_j together with the corresponding randomness s_j (such that $c_j = C(r_j, s_j)$), it somehow extracts a witness w such that $f(w) \in \{v_1^{\mathrm{H}}, v_2^{\mathrm{H}}\}$. fail_3 occurs if there exists a (corrupted) party P_j that in the beginning of Phase 3, sends a value r_j' which is different from the value r_j extracted in the extraction subroutine, and yet $\mathcal{P}_{\mathrm{H}}^{sid}$ accepts the WIPOK in Phase 3.

We show that each of these failures occurs with negligible probability.

$\mathcal{S}$ OUTPUTS fail_1 WITH NEGLIGIBLE PROBABILITY: Recall that $\mathcal{S}$ outputs fail_1 if there exists a session *sid* such that $\mathcal{P}_{\mathrm{H}}^{sid}$ enters Phase 3, and there exists a corrupted party P_j such that for one of its (Phase 1 or Phase 2) WISPOKs given to $\mathcal{P}_{\mathrm{H}}^{sid}$ in this session, the extraction subroutine failed to extract a witness. Note that it must be the case that $\mathcal{P}_{\mathrm{H}}^{sid}$ has accepted this WISPOK, since otherwise it would have never reached Phase 3. Thus, the occurrence of fail_1 implies that there exists a session *sid* and a corrupted party P_j such that the extraction subroutine failed to extract a witness, and yet $\mathcal{P}_{\mathrm{H}}^{sid}$ accepted the WISPOK. In other words, the strong knowledge extractor K failed to obtain a witness from the stand-alone prover $\mathcal{Q}_j^{sid}$, yet later in the simulation, $\mathcal{S}$ accepts that proof from $\mathcal{A}$. Intuitively, this should not happen because K has the property that if a prover convinces the honest verifier with non-negligible probability, then it successfully extracts with overwhelming probability. However, this is not immediate because K attempts to extract from the stand-alone adversary $\mathcal{Q}_j^{sid}$, whereas $\mathcal{P}_{\mathrm{H}}^{sid}$ verifies the proof from the original adversary $\mathcal{A}$. Thus, the essence here is to show that $\mathcal{Q}_j^{sid}$ convinces an honest verifier with the same probability that $\mathcal{A}$ convinces $\mathcal{P}_{\mathrm{H}}^{sid}$.

Claim 11 *For any corrupted party P_j participating in session sid, let* $\text{WISPOK}_j^{sid,\ell}$ *denote the ℓ^{th}* WISPOK *of P_j in this session. Then, the stand-alone prover $\mathcal{Q}_j^{sid}$, constructed by the extractor in* *the beginning of* $\text{WISPOK}_j^{sid,\ell}$*, convinces an honest verifier with exactly the same probability as* $\mathcal{P}_{\text{H}}^{sid}$ *accepts* $\text{WISPOK}_j^{sid,\ell}$ *in the straight-line simulation by* $\mathcal{S}$*.*

Proof: The main observation involved is that after $\text{WISPOK}_j^{sid,\ell}$ begins, the fact that no further extraction procedures are run and no new π-messages are received, makes no difference in the straight-line mode, *until after the* WISPOK *is finished.* This is ensured by the time-out for the WISPOK, by the fact that the output of the extraction subroutines in a session are not used until the session enters Phase 3 (together with the fact that a delay occurs before entering Phase 3), and by the assumption that π is a delayed protocol. We elaborate below.

First, we construct a simulator $\mathcal{S}'$ which is the same as $\mathcal{S}$ except that it does not invoke the extraction subroutine after the point at which $\text{WISPOK}_j^{sid,\ell}$ has begun. Thus, if a WISPOK, denoted $\text{WISPOK}_{j'}^{sid',\ell'}$, of Phase 1 or Phase 2 in a session *sid'* starts after the point at which $\text{WISPOK}_j^{sid,\ell}$ has begun, the simulator $\mathcal{S}'$ will not run the extraction subroutine for $\text{WISPOK}_{j'}^{sid',\ell'}$, whereas $\mathcal{S}$ would. Now, recall that $\mathcal{S}$ does not use the output that this extraction subroutine returns until session *sid'* enters Phase 3. We claim that the delay between Phase 2 and Phase 3 in the protocol ensures that $\mathcal{S}$ will enter Phase 3 in session *sid'* only after $\text{WISPOK}_j^{sid,\ell}$ has already concluded. This follows from the following facts:

1. When $\text{WISPOK}_j^{sid,\ell}$ began, session *sid'* did not yet finish Phase 2 (because session *sid'* must still at least run $\text{WISPOK}_{j'}^{sid',\ell'}$).

2. $\text{WISPOK}_j^{sid,\ell}$ is timed-out by $\mathcal{P}_{\text{H}}^{sid}$ if it does not conclude within τ local time units. By the assumption on the bounded clock drifts, this is at most $\tau\epsilon$ local time units according to $\mathcal{P}_{\text{H}'}^{sid'}$'s clock.

3. $\mathcal{P}_{\text{H}'}^{sid'}$ waits at least $\tau\epsilon$ local time units between Phase 2 and Phase 3.

Thus $\mathcal{S}$ and $\mathcal{S}'$ identically simulate the interaction between $\mathcal{P}_{\text{H}}^{sid}$ and $\mathcal{A}$, until $\text{WISPOK}_j^{sid,\ell}$ concludes. Therefore, the probability that $\text{WISPOK}_j^{sid,\ell}$ is accepted by $\mathcal{P}_{\text{H}}^{sid}$ is equal in both cases.

Next we modify $\mathcal{S}'$ to obtain a stand-alone machine $\mathcal{S}''$ which ignores all communication with the honest parties (in the π protocol) after the point at which $\text{WISPOK}_j^{sid,\ell}$ has begun. Note that if $\mathcal{S}'$ receives a π-message from a party P_i, it will be delivered to $\mathcal{A}$ only after a delay of $\tau\epsilon$ time units according to P_i's local clock. The restriction on the drifts of the clocks ensures that this delay is at least τ time units according to the local clock of $\mathcal{P}_{\text{H}}^{sid}$. So, if $\mathcal{S}'$ received this message *after* $\text{WISPOK}_j^{sid,\ell}$ has begun, it will not be used until $\mathcal{P}_{\text{H}}^{sid}$ concludes $\text{WISPOK}_j^{sid,\ell}$. This is because $\mathcal{P}_{\text{H}}^{sid}$ will conclude the WISPOK (by timing-out if necessary) within τ local time units after $\text{WISPOK}_j^{sid,\ell}$ has begun (which is at most $\tau\epsilon$ on P_i's local clock).[34] Hence the probability that $\mathcal{P}_{\text{H}}^{sid}$ accepts $\text{WISPOK}_j^{sid,\ell}$ in $\mathcal{S}''$ is equal to that in $\mathcal{S}'$.

Finally, we note that the system consisting of the stand-alone prover $\mathcal{Q}_j^{sid}$ interacting with an external honest verifier, is the same system as emulated by the stand-alone machine $\mathcal{S}''$. The role of the external verifier is played honestly by $\mathcal{P}_{\text{H}}^{sid}$ in $\mathcal{S}''$. Thus the probability that $\mathcal{Q}_j^{sid}$ can convince an honest verifier is exactly equal to the probability that $\mathcal{P}_{\text{H}}^{sid}$ will accept $\text{WISPOK}_j^{sid,\ell}$ in the execution of $\mathcal{S}''$ or $\mathcal{S}$. ■

[34] Notice that we are using here the fact that the the first and last message of $\text{WISPOK}_j^{sid,\ell}$ are sent by the verifier, since we assume that once $\text{WISPOK}_j^{sid,\ell}$ begins, it also begins according to the verifier $\mathcal{P}_{\text{H}}^{sid}$.

Now, let $\mu(n)$ be the negligible error function of the strong proof of knowledge. That is, if a prover convinces an honest verifier with probability greater than $\mu(n)$, then K successfully extracts with probability greater than $1-\mu(n)$. We define three events: "K-fail" if K fails to extract a witness from $\mathcal{Q}_j^{sid}$, "$\mathcal{S}$-pass" if $\mathcal{P}_{\mathrm{H}}^{sid}$ accepts $\text{WISPOK}_j^{sid,\ell}$, and "$\mathsf{good\text{-}proof}$" if the probability that an honest verifier accepts the proof given by the stand-alone prover $\mathcal{Q}_j^{sid}$ is at least $\mu(n)$. Then, the probability that $\mathcal{S}$ outputs fail_1 corresponding to $\text{WISPOK}_j^{sid,\ell}$ is bounded by

$$\begin{aligned}
\Pr[K\text{-}\mathsf{fail} \wedge \mathcal{S}\text{-}\mathsf{pass}] &= \Pr[K\text{-}\mathsf{fail} \wedge \mathcal{S}\text{-}\mathsf{pass} \wedge \mathsf{good\text{-}proof}] + \Pr[K\text{-}\mathsf{fail} \wedge \mathcal{S}\text{-}\mathsf{pass} \wedge \neg\mathsf{good\text{-}proof}] \\
&\leq \Pr[K\text{-}\mathsf{fail}|\mathsf{good\text{-}proof}]\Pr[\mathsf{good\text{-}proof}] + \Pr[\mathcal{S}\text{-}\mathsf{pass}|\neg\mathsf{good\text{-}proof}]\Pr[\neg\mathsf{good\text{-}proof}] \\
&\leq \mu(n)\Pr[\mathsf{good\text{-}proof}] + \mu(n)\Pr[\neg\mathsf{good\text{-}proof}] \\
&= \mu(n). \qquad (4)
\end{aligned}$$

$\mathcal{S}$ OUTPUTS fail_2 OR fail_3 WITH NEGLIGIBLE PROBABILITY: Recall that $\mathcal{S}$ outputs fail_2 if the extraction subroutine applied to a Phase 2 WISPOK of some session sid outputs w such that $f(w) \in \{v_1^{\mathrm{H}(sid)}, v_2^{\mathrm{H}(sid)}\}$. It outputs fail_3 if in Phase 3 of some session sid there exists a corrupted party P_j that does the following: **(a)** it sends a value r_j' different from the value r_j extracted from the extraction subroutine (applied to the Phase 2 WISPOK given by P_j in session sid), and **(b)** it succeeds in proving that it either knows w such that $f(w) \in \{v_1^{\mathrm{H}(sid)}, v_2^{\mathrm{H}(sid)}\}$ or that r_j' is indeed the value it committed to in Phase 2. However, since the second half of **(b)** is false, the soundness of the WIPOK would require that the first half of **(b)** be true, namely that it knows w.

Thus the cause for either of these failures (fail_2 or fail_3) is essentially that the adversary knows w such that $f(w) \in \{v_1^{\mathrm{H}(sid)}, v_2^{\mathrm{H}(sid)}\}$. (Note that these $(v_1^{\mathrm{H}(sid)}, v_2^{\mathrm{H}(sid)})$ values are chosen by an honest party.) Our proof that fail_2 or fail_3 is unlikely will use the argument that it is unlikely that the adversary can obtain such a w. Intuitively, this is due to the fact that w is only used in proving witness-indistinguishable proofs, which are also witness hiding. However, the actual proof is more complicated due to the fact that the adversary does not have to explicitly guess such a w, but merely succeed in giving a proof of knowledge of w, when concurrently interacting with the honest parties in multiple sessions. In order to prove that this is not feasible, we shall show how to construct a stand-alone machine M which interacts with an external machine $\mathcal{E}$. The machine $\mathcal{E}$ sends a pair (v_1, v_2), like in Phase 1 of our protocol, followed by many WISPOKs to M, to prove that it knows w such that $f(w) \in \{v_1, v_2\}$. Our construction of M will be such that if $\mathcal{S}$ outputs fail_2 or fail_3 with non-negligible probability, then M can also output w at the end of this interaction with non-negligible probability. Since f is a one-way function and the proofs are witness indistinguishable (and hence witness hiding), this will lead to a contradiction. We note that the formal proof relies heavily on the fact that the scheduling is pairwise disjoint.

M is constructed in two steps. First we describe a modified simulator $\mathcal{T}$, and then, depending on whether it is fail_2 or fail_3 that occurs with non-negligible probability, we show how to build M from $\mathcal{T}$.

The main feature of $\mathcal{T}$ is that, in a randomly chosen session sid^*, it interacts with the above mentioned external prover $\mathcal{E}$ (instead of with the internally simulated honest protocol program $\mathcal{P}_{\mathrm{H}}^{sid^*}$). We shall ensure that if $\mathcal{S}$ outputs fail_2 or fail_3 with non-negligible probability, then so does $\mathcal{T}$.

OVERVIEW OF $\mathcal{T}$: $\mathcal{T}$ emulates part of the hybrid system consisting of $\mathcal{F}_{\text{CRS}}$ and $\mathcal{S}$, but with the emulated $\mathcal{S}$ modified as follows: for a randomly chosen session sid^*, the simulated program $\mathcal{P}_{\mathrm{H}}^{sid^*}$ is not entirely run internally; instead part of the Phase 1 protocol is carried out by an external program $\mathcal{E}$, with which $\mathcal{T}$ interacts. The extraction subroutines in $\mathcal{S}$ are modified in such a way that they do not use the internal state of $\mathcal{E}$ (and in particular they do not "rewind" $\mathcal{E}$). These modifications will be such that $\mathcal{T}$ outputs fail_2 or fail_3 with non-negligible probability if $\mathcal{S}$ did so in the original hybrid system. The proof of this fact crucially depends on the way Phase 1 WISPOKs are scheduled; we will use the fact that the scheduling is pairwise disjoint to argue that even without rewinding $\mathcal{E}$, the extraction procedure can still be carried out in $\mathcal{T}$.

Overview of M: M will run $\mathcal{T}$ described above, as well as the rest of the hybrid system (namely, the honest parties running the protocol π). M *does not* include $\mathcal{E}$ mentioned above. Instead, it *interacts with* $\mathcal{E}$. Furthermore, M attempts to compute the witness w while interacting with $\mathcal{E}$, as mentioned earlier. If $\mathcal{T}$ outputs fail_2, the witness should have been extracted by the extractor in $\mathcal{T}$. Thus, M can output this witness. If $\mathcal{T}$ outputs fail_3, then M will construct a stand-alone prover for the Phase 3 WIPOK (corresponding to which $\mathcal{T}$ outputs fail_3) and use an extractor on this prover to obtain w (because, as mentioned earlier, in this case w will be the only valid witness for the WIPOK). In either case M will be able to output w with non-negligible probability.

Contradiction given M: Notice that M, which interacts with $\mathcal{E}$ as above, can output w with at most negligible probability. This is due to the following two observations:

1. Given the pair (v_1, v_2) which is computed by $\mathcal{E}$ (by choosing w_1, w_2 at random and setting $v_i = f(w_i)$), it is infeasible for M to find w such that $f(w) \in \{v_1, v_2\}$ (this follows from the fact that f is a one-way function).

2. The WISPOKs that $\mathcal{E}$ provides to M are *witness hiding* [26] (this follows from the fact that the proofs are witness indistinguishable with independent witnesses; see [27] for further details), and thus do not give M any non-negligible advantage in guessing w.

Thus, in order to prove that $\mathcal{S}$ has negligible probability of outputting fail_2 or fail_3, it suffices to show that if $\mathcal{S}$ outputs fail_2 or fail_3 with non-negligible probability, then M, which interacts with $\mathcal{E}$ as above, outputs w with non-negligible probability.

It remains only to construct M as claimed, which in turn is built from $\mathcal{T}$.

Construction of $\mathcal{T}$: First we present the details of the construction of $\mathcal{T}$, as well as the proof that it outputs fail_2 or fail_3 with non-negligible probability if $\mathcal{S}$ does so. The construction is carried out through a series of modifications to $\mathcal{S}$. The goal is to bring the simulator to a state where it does not need to rewind the Phase 1 WISPOKs of $\mathcal{P}_{\mathrm{H}}^{sid^*}$ (step 6). This will enable us to safely replace this part of $\mathcal{P}_{\mathrm{H}}^{sid^*}$ by the external machine $\mathcal{E}$ (step 7). In order to eliminate the rewinding of the Phase 1 WISPOKs of $\mathcal{P}_{\mathrm{H}}^{sid^*}$, we modify $\mathcal{S}$ so that rather than running the extraction subroutine on all Phase 2 WISPOKs, it runs it only on the Phase 2 WISPOKs of session sid^* (step 5). Then we further modify $\mathcal{S}$ so that, rather than running the extraction subroutine on all the Phase 1 WISPOKs, it runs it on a single Phase 1 WISPOK in each session; namely, the one which is pairwise disjoint to (i.e., does not overlap with) any of the Phase 1 WISPOKs of $\mathcal{P}_{\mathrm{H}}^{sid^*}$. (This point of the proof is exactly where the pairwise disjointness comes in.)

Formally, the construction of $\mathcal{T}$ is carried out through a series of seven modifications to $\mathcal{S}$. After each modification we show that if the probability of outputting fail_2 or fail_3 is non-negligible in the previous step, it continues to be so in this step too. The simulator in step 7 corresponds to $\mathcal{T}$. We now begin with the modifications:

1. First modify $\mathcal{S}$ so that it never outputs fail_1, and does not check if the fail_1 condition holds. We denote the modified simulator by $\mathcal{S}_1$. Since $\mathcal{S}$ outputs fail_1 with negligible probability,

it follows that $\mathcal{S}_1$ and $\mathcal{S}$ are statistically close, and in particular, the probability with which they output fail_2 and fail_3 is the same up to a negligible factor.

2. Modify $\mathcal{S}_1$ to obtain a new simulator $\mathcal{S}_2$ that behaves similarly to $\mathcal{S}_1$ with the following differences: Instead of accessing an external $\mathcal{F}_{\text{CRS}}$ functionality, it internally implements it. (Thus the honest parties obtain their outputs from $\mathcal{F}_{\text{CRS}}$ implemented by $\mathcal{S}_2$.) Furthermore, in Phase 3 of each session sid, instead of first drawing a random R_{CRS} (on behalf of $\mathcal{F}_{\text{CRS}}$) and then defining $r = \bigoplus_{i \neq \text{H}} r_i \oplus R_{\text{CRS}}$, it first draws a random r and defines $R_{\text{CRS}} = \bigoplus_{i \neq \text{H}} r_i \oplus r$. (See Eq. (3); recall that r_i is the value that party P_i committed to in the beginning of Phase 2, and if P_i is corrupted then r_i is obtained by applying the extraction subroutine to the Phase 2 WISPOK given by P_i.) Note that the output distributions of $\mathcal{S}_1$ and $\mathcal{S}_2$ are identical, and in particular the probability with which $\mathcal{S}_1$ and $\mathcal{S}_2$ output fail_2 and fail_3 is the same.

3. Next we observe that the r_j values extracted from the Phase 2 WISPOKs are used twice by the simulator $\mathcal{S}_2$:

   (a) To check the fail_3 condition.

   (b) To compute R_{CRS}, which is needed when some $\mathcal{P}_i^{sid}$ produces an output (sid, R). In this case, $\mathcal{F}_{\text{CRS}}$ (implemented by the simulator) sends R_{CRS} to P_i.

   We claim that the second usage of the r_j values is not essential. In order to see this, we modify $\mathcal{S}_2$ so that instead of computing $R_{\text{CRS}} = r_1 \oplus \cdots \oplus r_k$ and sending it to P_i (thereby using the r_j values), it computes $\mathsf{R}' = r'_1 \oplus \cdots \oplus r'_k$ and sends R' to P_i.[35] As was pointed out in footnote 33, if $R_{\text{CRS}} \neq \mathsf{R}'$ then it must be the case that either $\mathcal{S}$ outputs fail_3 or $\mathcal{P}_{\text{H}}^{sid}$ rejects one of the Phase 3 WIPOKs that it verifies, both of which result in P_i not receiving any output. Thus if P_i does receive an output it must be the case that $R_{\text{CRS}} = \mathsf{R}'$. Therefore this modification does not change anything in the system, except to make it explicit that the extracted values r_j are used only for determining if fail_3 occurs. We denote the new simulator by $\mathcal{S}_3$.

4. We next define $\mathcal{S}_4$ which behaves identically to $\mathcal{S}_3$ except for the following: $\mathcal{S}_4$ chooses a random session and outputs fail_2 or fail_3 only if it happens in the chosen session. (In other sessions if $\mathcal{S}_3$ would have output fail_2 or fail_3 and halted, $\mathcal{S}_4$ does not even check for the failure condition and so might continue executing.) Note that there are only polynomially sessions possible (as the adversary and the polynomially many parties are all assumed to be strict probabilistic polynomial-time machines). Hence if $\mathcal{S}_3$ outputs fail_2 or fail_3 with non-negligible probability, so does $\mathcal{S}_4$. (The reason that this holds is that with probability 1/poly, the *first* session in which fail_2 or fail_3 occurs will be chosen, and the simulation until that point is identical.)

   We shall denote by $\mathcal{S}_4^{sid^*}$ the resulting simulator when $\mathcal{S}_4$ picks a session with session identifier sid^* as its random choice. All the simulators defined below also choose a random session in the beginning. We use similar notation to denote them.

5. $\mathcal{S}_5^{sid^*}$ is the same as $\mathcal{S}_4^{sid^*}$ with the following difference. $\mathcal{S}_5^{sid^*}$ does not run the Phase 2 extraction subroutines for any session except sid^*. (The Phase 1 extraction subroutines are run for all sessions.) Note that $\mathcal{S}_4^{sid^*}$ does not use the extracted values from Phase 2 in

[35] Recall that r'_j is the (supposedly committed) value sent by party P_j at the beginning of Phase 3 of this session, and note that (sid, R') is the output of $\mathcal{P}_{\text{H}}^{sid}$ in this session.

any other session except sid^*. This is because it neither calculates R_{CRS} nor checks the fail_2 and fail_3 conditions in those sessions. Thus $\mathcal{S}_5^{sid^*}$ and $\mathcal{S}_4^{sid^*}$ output $\mathsf{fail}_2/\mathsf{fail}_3$ with the same probability.

6. $\mathcal{S}_6^{sid^*}$ is designed not to rewind $\mathcal{P}_{\text{H}}^{sid^*}$ (so that in the next step we can replace the internally simulated $\mathcal{P}_{\text{H}}^{sid^*}$ by an external machine $\mathcal{E}$).

   Recall that in $\mathcal{S}_5^{sid^*}$, for each session sid and for each Phase 1 WISPOK $\text{WISPOK}_j^{sid,\ell}$ ($1 \leq \ell \leq m+2$) proven by a corrupt party P_j to the honest verifier $\mathcal{P}_{\text{H}}^{sid}$, the extraction subroutine builds a stand-alone prover $\mathcal{Q}_j^{sid,\ell}$ (which we earlier abbreviated as $\mathcal{Q}_j^{sid}$). In $\mathcal{S}_6^{sid^*}$ this stand-alone prover is modified so that it aborts the proof if, during its internally simulated execution, $\mathcal{P}_{\text{H}}^{sid^*}$ sends a message (as the prover) in a Phase 1 WISPOK in session sid^*.[36] (This depends on the adversarial scheduling of the sessions sid and sid^*.) We shall denote this modified stand-alone prover by $\mathcal{Q}_{j\backslash\text{H}(sid^*)}^{sid,\ell}$. Other than this replacement $\mathcal{S}_6^{sid^*}$ is identical to $\mathcal{S}_5^{sid^*}$.

   We need to argue that even with this replacement the probability of $\mathcal{S}_6^{sid^*}$ outputting fail_2 or fail_3 does not significantly change. For this, it suffices to prove that for every session sid, if Phase 1 of session sid is successfully concluded (in the straight-line mode), then for every corrupted party P_j in session sid, the extractor (for $\mathcal{S}_6^{sid^*}$) will succeed in obtaining a witness from one of the Phase 1 WISPOKs proven by P_j in session sid with overwhelming probability. Intuitively, this follows from the *pairwise disjoint* scheduling of the Phase 1 WISPOKs, as explained below.

   If Phase 1 of a session sid is successfully concluded (in the straight-line mode), then the fact that the scheduling of the Phase 1 WISPOKs is pairwise disjoint guarantees that there exists of a WISPOK proven by P_j in session *sid that does not overlap, according to* $\mathcal{P}_{\text{H}}^{sid}$ *and* $\mathcal{P}_{\text{H}}^{sid^*}$, *with any of the Phase 1* WISPOK*s given by* $\mathcal{P}_{\text{H}}^{sid^*}$ *in session* sid^*.[37] Intuitively, the extractor (for $\mathcal{S}_6^{sid^*}$) should succeed in extracting a witness from this "disjoint" WISPOK with overwhelming probability. However, this is not immediate because, similarly to the proof that fail_1 occurs only with negligible probability, the extractor does not extract from the real (cheating) prover, but rather from the corresponding stand-alone prover. The fact that in the straight-line mode, say $\text{WISPOK}_j^{sid,\ell}$, was accepting and was non-overlapping with the proofs given by $\mathcal{P}_{\text{H}}^{sid^*}$, does not necessarily imply that this will remain so in the look-ahead run in the extraction mode (since it may be that in the straight-line mode this happened by chance). In other words, it is still possible that, with significant probability, an honest verifier will reject the proof given by $\mathcal{Q}_{j\backslash\text{H}(sid^*)}^{sid,\ell}$, or $\mathcal{Q}_{j\backslash\text{H}(sid^*)}^{sid,\ell}$ will abort because of $\mathcal{P}_{\text{H}}^{sid^*}$ sending a message. Nevertheless, if $\text{WISPOK}_j^{sid,\ell}$ was accepting and non-overlapping, we can expect that, with non-negligible probability, this will remain the case when the real (cheating) prover is replaced with the stand-alone prover. Thus, the knowledge extractor will be able to extract the witness with overwhelming probability, even though it is not allowed to rewind the (now external) prover $\mathcal{P}_{\text{H}}^{sid^*}$. A more formal argument follows.

   Consider any session sid and any corrupted party P_j in session sid. Similar to the proof that fail_1 occurs with only negligible probability, we define the following events:

[36]Note that in $\mathcal{S}_5^{sid^*}$ the only Phase 2 WISPOKs which are extracted from are those of session sid^*, which occur strictly after the Phase 1 of sid^*. Hence these extractions never rewind the Phase 1 proofs given by $\mathcal{P}_{\text{H}}^{sid^*}$, and need not be modified.

[37]Recall that non-overlapping between two sessions sid and sid^* is with respect to two honest parties, one from each session. Here we take $\mathcal{P}_{\text{H}}^{sid}$ to be the honest party in sid, and $\mathcal{P}_{\text{H}}^{sid^*}$ to be the honest party in sid^*.

- K-fail_ℓ is the event that K fails to extract a witness from $\mathcal{Q}^{sid,\ell}_{j\setminus\mathrm{H}(sid^*)}$
- $\mathcal{S}$-pass_ℓ is the event that $\mathcal{P}^{sid}_{\mathrm{H}}$ accepts $\text{WISPOK}^{sid,\ell}_j$
- $\mathsf{disjoint}_\ell$ is the event that $\text{WISPOK}^{sid,\ell}_j$ does not overlap with any WISPOK given by $\mathcal{P}^{sid^*}_{\mathrm{H}}$, according to $\mathcal{P}^{sid}_{\mathrm{H}}$ and $\mathcal{P}^{sid^*}_{\mathrm{H}}$.

  Remark: Notice that if event $\mathsf{disjoint}_\ell$ holds then while P_j is proving $\text{WISPOK}^{sid,\ell}_j$, party $\mathcal{P}^{sid^*}_{\mathrm{H}}$ does not send any message (as a prover) in any phase 1 WISPOK.[38]
- $\mathsf{good\text{-}proof}_\ell$ is the event that the probability that an honest verifier accepts the proof given by the stand-alone prover $\mathcal{Q}^{sid,\ell}_{j\setminus\mathrm{H}(sid^*)}$ is at least $\mu(n)$ (where $\mu(n)$ is a negligible function such that if a prover convinces an honest verifier with probability greater than $\mu(n)$, then K successfully extracts with probability greater than $1-\mu(n)$).

Notice that

$$\Pr\left[K\text{-}\mathsf{fail}_\ell \mid \mathsf{good\text{-}proof}_\ell\right] \le \mu(n). \tag{5}$$

Moreover, using similar arguments to the ones given in the proof of Claim 11, and using the remark above, one can show that

$$\Pr\left[\mathcal{S}\text{-}\mathsf{pass}_\ell \wedge \mathsf{disjoint}_\ell \mid \neg\mathsf{good\text{-}proof}_\ell\right] \le \mu(n). \tag{6}$$

Now we can bound the probability of $\mathcal{S}^{sid^*}_6$ not being able to extract a witness in the extraction mode from any of the WISPOKs given by P_j in session sid, while in the straight-line mode all those WISPOKs are accepted by $\mathcal{P}^{sid}_{\mathrm{H}}$.

$$\begin{aligned}
\Pr\left[\bigwedge_\ell (K\text{-}\mathsf{fail}_\ell \wedge \mathcal{S}\text{-}\mathsf{pass}_\ell)\right] &= \Pr\left[\left(\bigwedge_\ell (K\text{-}\mathsf{fail}_\ell \wedge \mathcal{S}\text{-}\mathsf{pass}_\ell)\right) \wedge \left(\bigvee_\ell \mathsf{good\text{-}proof}_\ell\right)\right] \\
&\quad + \Pr\left[\bigwedge_\ell (K\text{-}\mathsf{fail}_\ell \wedge \mathcal{S}\text{-}\mathsf{pass}_\ell \wedge \neg\mathsf{good\text{-}proof}_\ell)\right] \\
&\le \Pr\left[\bigvee_\ell (K\text{-}\mathsf{fail}_\ell \wedge \mathsf{good\text{-}proof}_\ell))\right] \\
&\quad + \Pr\left[\bigwedge_\ell (\mathcal{S}\text{-}\mathsf{pass}_\ell \wedge \neg\mathsf{good\text{-}proof}_\ell)\right]
\end{aligned}$$

By Eq. (5) and by the union bound, the first term is bounded by $(m+2)\mu(n)$. To bound the second term we use the guarantee from pairwise disjoint scheduling, namely: $\bigwedge_\ell \mathcal{S}\text{-}\mathsf{pass}_\ell \implies \bigvee_\ell \mathsf{disjoint}_\ell$. Then we have

$$\begin{aligned}
\Pr\left[\bigwedge_\ell (\mathcal{S}\text{-}\mathsf{pass}_\ell \wedge \neg\mathsf{good\text{-}proof}_\ell)\right] &= \Pr\left[\left(\bigwedge_\ell (\mathcal{S}\text{-}\mathsf{pass}_\ell \wedge \neg\mathsf{good\text{-}proof}_\ell)\right) \wedge \left(\bigvee_\ell \mathsf{disjoint}_\ell\right)\right] \\
&\le \Pr\left[\bigvee_\ell (\mathcal{S}\text{-}\mathsf{pass}_\ell \wedge \neg\mathsf{good\text{-}proof}_\ell \wedge \mathsf{disjoint}_\ell)\right] \\
&\le \sum_\ell \Pr\left[\mathcal{S}\text{-}\mathsf{pass}_\ell \wedge \mathsf{disjoint}_\ell \mid \neg\mathsf{good\text{-}proof}_\ell\right]
\end{aligned}$$

[38] This follows from our assumption that the first and last message in every WISPOK is sent by the verifier.

Using Eq. (6), this is bounded by $(m+2)\mu(n)$. Thus, $\Pr\left[\bigwedge_\ell (K\text{-}\mathsf{fail}_\ell \wedge \mathcal{S}\text{-}\mathsf{pass}_\ell)\right]$ is negligible. Hence we can conclude that, except with negligible probability, the extractor of $\mathcal{S}_6^{sid^*}$ will also be able to extract the witnesses in all the (polynomially many) sessions which reach Phase 3 in the straight-line mode.

Note that the only difference between $\mathcal{S}_5^{sid^*}$ and $\mathcal{S}_6^{sid^*}$ is in the way the Phase 1 witnesses of corrupted parties are extracted. Since both $\mathcal{S}_5^{sid^*}$ and $\mathcal{S}_6^{sid^*}$ succeed in extracting these witnesses (with overwhelming probability) for every session that reaches Phase 3, and since these witnesses are used only in Phase 3, we would like to conclude and say that the output distributions of $\mathcal{S}_5^{sid^*}$ and $\mathcal{S}_6^{sid^*}$ are *statistically* indistinguishable. However, there is a subtle point here: The witnesses w^j obtained by $\mathcal{S}_5^{sid^*}$ and $\mathcal{S}_6^{sid^*}$ may be distributed differently. But, since the simulation uses these witnesses only in WIPOKs we conclude that the output distributions of $\mathcal{S}_5^{sid^*}$ and $\mathcal{S}_6^{sid^*}$ are *computationally* indistinguishable, which in particular implies that the probability with which they output fail_2 and fail_3 is the same (up to a negligible factor). The formal reduction (reducing any algorithm that distinguishes between the outputs of $\mathcal{S}_5^{sid^*}$ and $\mathcal{S}_6^{sid^*}$ to an algorithm that breaks the witness indistinguishability property of the WIPOK) is straightforward, and omitted here.

7. Finally we define $\mathcal{S}_7^{sid^*}$ which replaces the internal simulation of the first message (namely, $(v_1^{\mathrm{H}(sid^*)}, v_2^{\mathrm{H}(sid^*)})$ and the WISPOKs given by $P_{\mathrm{H}(sid^*)}$ in session sid^* by externally received messages. That is, $\mathcal{S}_7^{sid^*}$ interacts with an external machine $\mathcal{E}$ that picks (w_1, w_2), sets $v_i = f(w_i)$, sends them to $\mathcal{S}_7^{sid^*}$, and then engages in multiple WISPOKs to prove knowledge of w such that $f(w) \in \{v_1, v_2\}$. Internally, $\mathcal{S}_7^{sid^*}$ uses this to replace (part of) the computation carried out by $\mathcal{P}_{\mathrm{H}}^{sid^*}$. In other words, the program of $\mathcal{P}_{\mathrm{H}}^{sid^*}$ will be considered split into an external machine $\mathcal{E}$ (which sends $(v_1^{\mathrm{H}(sid^*)}, v_2^{\mathrm{H}(sid^*)})$ and carries out the proofs of Phase 1) and an internal machine (which carries out the rest of the protocol execution). The extractors will not have access to the state of the external machine.

   Recall that the stand-alone provers used by the extractors in $\mathcal{S}_6^{sid^*}$ abort if, during their internally simulated execution, $\mathcal{P}_{\mathrm{H}}^{sid^*}$ sends a message (as a prover) in a Phase 1 WISPOK. Thus, the extractor used by $\mathcal{S}_6^{sid^*}$ does not need access to the internal state of $\mathcal{E}$. Therefore, $\mathcal{S}_7^{sid^*}$ can use the same extraction procedure, which implies that the probability that $\mathcal{S}_7^{sid^*}$ outputs fail_2 or fail_3 is the same as the probability that $\mathcal{S}_6^{sid^*}$ does so.

$\mathcal{T}$ is the same as $\mathcal{S}_7^{sid^*}$, with sid^* chosen randomly. The above series of steps shows that if $\mathcal{S}$ outputs fail_2 or fail_3 with non-negligible probability, then $\mathcal{T}$ also outputs fail_2 or fail_3 with non-negligible probability.

Construction of M: We seek to construct a machine M such that, if $\mathcal{T}$ has a non-negligible probability of outputting fail_2 or fail_3 in its interaction with $\mathcal{E}$, then with non-negligible probability, when M interacts with $\mathcal{E}$ it will succeed in computing w such that $f(w) \in \{v_1, v_2\}$, where (v_1, v_2) is the pair sent to it by $\mathcal{E}$.

The machine M emulates the entire system of honest parties running π and the simulator $\mathcal{T}$. However, it does not simulate $\mathcal{E}$. Instead M itself interacts with $\mathcal{E}$. We construct M separately for the following two cases.

$\mathcal{T}$ outputs fail_2 with non-negligible probability: While emulating the system, if $\mathcal{T}$ outputs fail_2, then M can output the witness w that caused $\mathcal{T}$ to fail. This witness, by definition of fail_2, equals w such that $f(w) \in \{v_1, v_2\}$, where (v_1, v_2) is the first message sent by $\mathcal{E}$. This is in contradiction to the fact that the WISPOKs are witness hiding.

$\mathcal{T}$ **outputs** fail_3 **with non-negligible probability:** Recall that $\mathcal{T}$ outputs fail_3 if some P_j sent in the beginning of Phase 3 a value $r'_j \neq r_j$, where r_j was the value extracted in Phase 2. Let sid^* be the random session chosen by $\mathcal{T}$ (i.e., $\mathcal{T}$ is identical to $\mathcal{S}_7^{sid^*}$). In $\mathcal{T}$, when $\mathcal{P}_{\rm H}^{sid^*}$ enters Phase 3, M will randomly pick a corrupt party P_j and construct a stand-alone prover corresponding to P_j's Phase 3 WIPOK to $\mathcal{P}_{\rm H}^{sid^*}$. The stand-alone prover is constructed by modifying $\mathcal{P}_{\rm H}^{sid^*}$ to simply relay messages between P_j and an external verifier. This construction is similar to, but simpler than that of $\mathcal{Q}_j^{sid^*}$ described earlier. Recall that there, $\mathcal{Q}_j^{sid^*}$ worked exactly like the simulator, continuing from the point where the extraction subroutine was invoked, *except* that $\mathcal{Q}_j^{sid^*}$ (unlike the simulator) did not interact with the honest parties running π and did not invoke the extraction subroutines that $\mathcal{S}$ invokes. But now, the stand-alone prover *includes* the honest parties running the π protocol, and also runs all extraction subroutines. Note that by running the extraction subroutines there is no danger in rewinding $\mathcal{E}$ since $\mathcal{E}$ is not active any more when M reaches Phase 3 of the session sid^*.

Now, M applies a knowledge extractor to this stand alone prover, and if it extracts a witness w such that $f(w) \in \{v_1, v_2\}$, then M outputs w. Note that $\mathcal{T}$ outputs fail_3 when for some party $P_{j'}$, the value $r_{j'}$ extracted in Phase 2 is different from the value $r'_{j'}$ that it sent out in Phase 3, and yet its Phase 3 WIPOK is accepted. Note that the only valid witnesses for this WIPOK are values w such that $f(w) \in \{v_1, v_2\}$. Now, since the probability of $\mathcal{T}$ outputting fail_3 is non-negligible, and since there are only polynomially many (corrupt) parties from which M picked P_j, with non-negligible probability M picked party $P_{j'}$, and thus convinces the external verifier of a statement with the only witnesses being w such that $f(w) \in \{v_1, v_2\}$. This implies that the knowledge extractor when run on M, will succeed in outputting such a w with non-negligible probability. (This knowledge extractor runs in expected, and not strict, polynomial-time. Nevertheless, using standard arguments, we can obtain a strict polynomial-time machine that obtains w with non-negligible probability.)

Completing the proof for fail_2 **and** fail_3**.** This completes the construction of M, and also the proof that $\mathcal{S}$ outputs fail_2 or fail_3 with only negligible probability.

THE HYBRIDS: Above we have shown that $\mathcal{S}$ outputs fail_1, fail_2 or fail_3 only with negligible probability. We now prove that the output distributions of $\mathcal{S}$ and the honest parties running π in the $\mathcal{F}_{\rm CRS}$-hybrid model are indistinguishable from that of $\mathcal{A}$ and the honest parties running π^ρ in the real world with timing. For this we note that $\mathcal{S}$ in the hybrid world almost perfectly emulates the real world interaction, but with a few differences. The main difference is that in the simulated world in every session sid there is one party $\mathcal{P}_{\rm H}^{sid}$ that deviates from the protocol. This is the case since the simulator gets a random string $R_{\rm CRS}$ from the functionality and needs to simulate the protocol so that its output will be equal to $R_{\rm CRS}$.

We shall build some hybrid simulators to bridge the gap between the real and hybrid worlds.

- HYBRID SIMULATOR $\mathcal{H}_1$: This is similar to $\mathcal{S}_2$ as defined earlier: It implements $\mathcal{F}_{\rm CRS}$ internally and defines $R_{\rm CRS}$ by randomly picking r and setting $R_{\rm CRS} = \bigoplus_{i \neq {\rm H}} r_i \oplus r$ (however it outputs fail_1 just like $\mathcal{S}$ does). As argued above, this does not change anything in the system, and in particular the output distributions remain unchanged.
- HYBRID SIMULATOR $\mathcal{H}_2$: Recall that when $\mathcal{P}_i^{sid}$ produces an output (sid, R), $\mathcal{H}_1$ delivers the output $(sid, R_{\rm CRS})$ from $\mathcal{F}_{\rm CRS}$ to P_i (after $\mathcal{P}_{\rm H}^{sid}$ produces an output). In contrast, the simulator $\mathcal{H}_2$ will hand P_i the output R generated by $\mathcal{P}_i^{sid}$ in the simulation. Note that if $\mathcal{P}_i^{sid}$ outputs (sid, R) and $\mathcal{H}_1$ did not output fail_3 (and $\mathcal{P}_{\rm H}^{sid}$ produced an output) then it must be the case that with overwhelming probability $R = R_{\rm CRS}$, since the fact that $\mathcal{H}_1$ did not output fail_3 and that $\mathcal{P}_{\rm H}^{sid}$ produced an output implies that all parties must have sent in Phase 3 the decommitment

value which was extracted by the extracted subroutine. Therefore, the output distributions of $\mathcal{H}_1$ and $\mathcal{H}_2$ are statistically close.

- Hybrid Simulator $\mathcal{H}_3$: $\mathcal{H}_3$ is defined exactly as $\mathcal{H}_2$ is, except with the following difference: instead of running $\mathcal{P}_{\mathrm{H}}^{sid}$ in every session sid (with at least one honest player H), $\mathcal{H}_3$ runs another program ${\mathcal{P}'}_{\mathrm{H}}^{sid}$. This program is exactly like $\mathcal{P}_{\mathrm{H}}^{sid}$, except that in Phase 3, instead of sending r received from $\mathcal{S}$, it sends out r_{H} as instructed by the honest program of ρ.

  The hiding property of the Phase 2 commitment scheme, the hiding property of the Phase 2 WISPOK, and the fact that r and the committed value r_{H} are identically distributed (both are uniformly distributed) imply that the output distributions of $\mathcal{H}_2$ and $\mathcal{H}_3$ are computationally indistinguishable.

- Hybrid Simulator $\mathcal{H}_4$: $\mathcal{H}_4$ uses exactly the program specified by ρ for $\mathcal{P}_{\mathrm{H}}^{sid}$. Note that the only difference between ${\mathcal{P}'}_{\mathrm{H}}^{sid}$ used by $\mathcal{H}_3$ and the program specified by ρ is that while giving Phase 3 WIPOK to a party P_j, ${\mathcal{P}'}_{\mathrm{H}}^{sid}$ uses the alternate witness provided by $\mathcal{S}$ (namely w^j such that $f(w^j) \in \{v_1^j, v_2^j\}$) instead of what is specified by the protocol ρ. The witness indistinguishable property of this WIPOK implies that the output distributions of $\mathcal{H}_3$ and $\mathcal{H}_4$ are computationally indistinguishable.

Now note that the system run by $\mathcal{H}_4$ and the real world system are identical, except that $\mathcal{H}_4$ also runs the extraction subroutines and might output fail depending on the outputs from these subroutines. Other than that, the extraction subroutines are not used in the system (because we replaced the $\mathcal{P}_{\mathrm{H}}^{sid}$ programs by the original programs specified by ρ). Now since $\mathcal{S}$ outputs fail with negligible probability and the output of $\mathcal{H}_4$ is indistinguishable from that of $\mathcal{S}$, we see that $\mathcal{H}_4$ also outputs fail only with negligible probability. Thus, it follows that the output of the system with $\mathcal{H}_4$ is indistinguishable from that of the real world system. From the line of reasoning above, we conclude that the distribution of the output of the system consisting of $\mathcal{S}$ and the honest parties running π in the $\mathcal{F}_{\mathrm{CRS}}$-hybrid world is indistinguishable from the output of the system consisting of $\mathcal{A}$ and the honest parties running π^ρ in the real world (with time). ■

4 Pairwise-Disjoint Scheduling

In this section, we construct a pairwise-disjoint scheduling algorithm, thereby proving Theorem 9 of Section 3.4. On a very high level, the idea is that for each session $sid \in \{0,1\}^m$, the schedule output by $S(\sigma, sid, \Delta, \epsilon)$ is such that protocol σ is executed $m+2$ times, with delays between each execution (here we make use of the timing model). The crux of the idea is that the delays depend on the bits of sid, so that for any $sid \neq sid'$ the executions of $S(\sigma, sid, \Delta, \epsilon)$ and $S(\sigma, sid', \Delta, \epsilon)$ will not be aligned. The schedule is enforced by requiring the parties to "time-out" if the execution is too long, say if it takes more that τ local time units, where τ is a function of σ and Δ (and the delays depend on this parameter τ). In our specific protocol, σ is a strong proof-of-knowledge with $\alpha(n)$ rounds, and we set $\tau \stackrel{\mathrm{def}}{=} \alpha(n)\cdot\Delta$.[39]

Motivation to the schedule. Due to the technical nature of the schedule and its proof, we first provide a lengthy discussion explaining the idea behind the construction. Recall that our aim is to obtain pairwise disjointness, meaning that for every two sessions sid and sid', there exists at least

[39]Note that since σ consists of $\alpha(n)$ rounds and Δ is an upper bound on the latency according to *all* clocks, we have that $\tau \stackrel{\mathrm{def}}{=} \alpha(n)\cdot\Delta$ is an upper bound on σ's overall running time, assuming that all messages are delivered with Δ time units (and assuming local computation is instantaneous).

one execution of σ in sid that does not overlap with *any* execution of σ in sid'. As a first try, suppose that the schedule consists of running σ twice, with a delay between each execution that is "large" and directly proportional to the session ID sid. For example, interpret the value $sid \in \{0,1\}^m$ as an integer in the range $[1, \ldots, 2^m]$ and delay $2sid \cdot \tau$ time units between the executions, where τ is an upper bound on how long σ should run. Furthermore, time-out an execution of σ if it runs longer than τ time units. Now, let $sid' \neq sid$ be two different sessions. Denote by σ_1, σ_2 the two executions of σ in session sid, and denote by σ'_1, σ'_2 the two executions of σ in session sid'. Without taking the clock drift ϵ into account for now, we have the following cases:

1. *Execution σ_1 overlaps with execution σ'_1:* Notice that σ_2 is delayed by $2sid \cdot \tau$ time units, whereas σ'_2 is delayed by $2sid' \cdot \tau$ time units. Since $sid' \neq sid$, there is a difference of at least 2τ time units between the delay before σ_2 and the delay before σ'_2. The fact that each execution of σ takes at most τ time units ensures that the σ'_2 execution does not overlap with σ_2. Also, the fact that the delay before σ'_2 is longer than τ time units implies that σ'_2 does not overlap with σ_1.

2. *Execution σ_2 overlaps with execution σ'_2:* The same analysis as above yields that σ'_1 does not overlap with σ_1 or σ_2.

3. *Execution σ_1 overlaps with execution σ'_2:* In this case, it follows immediately that σ'_1 concluded before σ_1 began (because there is a delay of more than τ time units between σ'_1 and σ'_2). Thus, σ'_1 does not overlap with σ_1 or σ_2.

4. *Execution σ_2 overlaps with execution σ'_1:* As above, it follows that σ'_2 does not overlap with σ_1 or σ_2.

We therefore obtain that the above is a pairwise-disjoint schedule. However, this schedule is problematic because the length of the delays are *exponential* in the length of sid. Thus, unless there is an *a priori* polynomial bound on the number of sessions (in which case, sid can be of length $O(\log n)$), we obtain that the schedule is not polynomial in the security parameter.

We solve this problem by using a more involved scheduling strategy, adapted from the strategy of Chor and Rabin [18]. We now recall this strategy (already described in Section 3.4). It was observed in [18] that if the identifiers sid and sid' are encoded (one-to-one) into $2m$-bit strings containing m zeros and m ones, then for any two different identifiers $sid \neq sid'$, there is at least one bit position where the encoding of sid has a zero and that of sid' has a one. Suppose now that the time is divided into $2m$ distinct slots (each slot corresponding to a bit of the encoding of the identifier), and executions of σ in the session sid are run only in the slots where the encoding of sid has a one in that slot. Then there is a slot in which an execution of σ is run in sid', but not in sid. The improvement over the previous scheme is that this encoding is compact (i.e., linear), rather than exponential, in the length of sid.

However, there are numerous complications in adapting this strategy to our setting. Firstly, unlike the setting considered in [18], we consider executions of ρ occurring in different sessions at different times. Therefore, two encodings which are different may be shifted with respect to each other in a way that all the positions with ones align with each other (e.g. the ones in 0110 and 1100 can be aligned with each other by shifting one of the two strings by one position). This problem is solved simply by prepending a one to the encoding (for convenience in later analysis, we shall actually add a one to both ends of the encoding). We therefore have that the above encodings become 101101 and 111001, respectively, and shifting in either direction will result in independence.

Another problem that arises is due to the fact that in our setting, it is not possible to define distinct time-slots (because the parties' clocks are not synchronized). Therefore, one execution of

σ in session sid can partially overlap with two executions of σ in session sid'. We solve this by introducing delays between the time slots in each session. We note that it suffices to delay for at least the maximum time that it takes to conclude an execution of σ. (It is possible to limit the maximum time for any execution of σ by using a time-out instruction.) We thereby obtain that any execution of σ in session sid can overlap with *at most one* execution of σ in session sid'.

The final complication that arises is due to the fact that the parties' local clocks do not proceed at exactly the same rate, but rather can drift. Since the rates at which the local clocks of the different parties proceed may vary adversarially (up to a factor ϵ), it is possible that two different schedules from different sessions may perfectly overlap. For example, suppose that the schedule for session sid is $10^i10^j1^k$ and the schedule in sid' is $10^j10^i1^k$ (with say $i > j$). Furthermore, suppose that an honest party P is participating in session sid, and another honest party P' is participating in session sid'. Then, the adversary can cause the executions of P and P' to overlap by first running the clock of P faster than that of P' by a factor of i/j (starting after the first execution of σ, up to the second execution of σ), and then running it slower by a factor of j/i (after finishing the second execution of σ and until reaching the third execution of σ).[40] Now, note that although P and P' use the prescribed distinct schedules, the adversary can make every execution of σ in sid fully coincide with every execution of σ in sid'. However, for this to work, it must hold that i/j is less than ϵ. Intuitively, if i and j are small (and ϵ is not too large, and in particular $\epsilon << i/j$) then the adversary does not have enough leeway to align these two execution. This holds because even if the adversary runs P's clock faster than P''s clock by a factor of ϵ, the delay of P between these two executions will still be much longer than the delay of P'. Thus, if we make sure that there are no long runs of zeros in the encoding used, we can use our scheduling for values of ϵ that are reasonably larger than one (but not too large). The particular encoding we use (which is sometimes called the "Manchester encoding") ensures that there will be at most two consecutive zeros. For this encoding, we prove that if two consecutive executions of session sid overlap with two consecutive executions of session sid', then it must be the case that the number of zeros between these executions is the same in both sessions, assuming ϵ is less than $\sqrt[3]{1.5}$. The reason for this specific bound on ϵ is quite technical, and is needed for our proof to go through. Our complete description of the schedule, and the formal proof, take all of the above discussed factors into account.

Convention. We assume for simplicity (and without loss of generality) that in protocol σ there exists one party that sends the first message which is of the form "start" and the last message which is of the form "end" to all of the parties that participate in the protocol. This ensures that (when the adversary does not corrupt parties and delivers all messages within time Δ) the duration of the protocol is roughly the same for all parties participating in σ.

The construction. We now present our construction of a pairwise disjoint scheduling. We associate with each session sid a unique session identifier u^{sid} which is a vector of zeros and ones, so that the number of ones is the same for each identifier. Loosely speaking, each 1 entry will correspond to an execution of σ.

Formally, our scheduling algorithm, on input a protocol σ, a session identifier sid, and time-bounds Δ and ϵ, operates as follows. We specify the delay and time-out mechanisms in terms of some parameters d, τ, $\tau_{\text{MIN}}(\cdot)$ and $\tau_{\text{MAX}}(\cdot)$. We shall fix these parameters later, as functions of Δ and ϵ.

1. Associate with session $sid \in \{0,1\}^m$ a vector $u^{sid} = (u_1^{sid}, \ldots, u_{2m+2}^{sid}) \in \{0,1\}^{2m+2}$, defined as follows:

[40] We ignore the "delaying slots" between the time slots for this discussion.

(a) $u_1^{sid} = 1$ and $u_{2m+2}^{sid} = 1$.

(b) For every $j \in \{1, \ldots, m\}$, if $sid_j = 1$ then $(u_{2j}^{sid}, u_{2j+1}^{sid}) = (1, 0)$, and if $sid_j = 0$ then $(u_{2j}^{sid}, u_{2j+1}^{sid}) = (0, 1)$.

Notice that u^{sid} has exactly $m+2$ ones and m zeros. Moreover, it has at most two consecutive zeros. $S(\sigma, sid, \Delta, \epsilon)$ will consist of $m+2$ executions of σ, one execution corresponding to each 1 entry of the u^{sid} vector.

2. Carry out $m+2$ executions of σ according to the following scheduling.

   (a) Set $j = 1$.

   (b) If $u_j^{sid} = 1$ then carry out an execution of σ and then continue to step 2c. Otherwise, continue immediately to step 2c

   (c) Wait d local time units (d will be specified later).

   (d) Set $j \stackrel{\text{def}}{=} j + 1$.

   (e) If $j \leq 2m + 2$ then goto step 2b.

3. Time-Out Mechanism: In each of the above executions of σ, each participant checks that no more than τ local time units passed from the time that it received its first message of the execution ("start"), to the time that it received its last message of the execution ("end"). If more time passes before the execution is over, then it outputs $(\mathsf{time\text{-}out}, sid)$ on its output tape and halts the execution.

4. Delay Mechanism: For any $x \in \{0, 1, 2\}$ and for any two consecutive executions of σ in $S(\sigma, sid, \Delta, \epsilon)$ that correspond to two 1's with x zeros in between, each party P participating in session sid, checks that the delay (according to P's local clock) between these two executions (i.e., the time between receiving its last message in one execution and receiving its first message in the next execution), is between $\tau_{\mathrm{MIN}}(x)$ and $\tau_{\mathrm{MAX}}(x)$. Here $\tau_{\mathrm{MIN}}(\cdot)$ and $\tau_{\mathrm{MAX}}(\cdot)$ are increasing functions, to be specified later. For each honest participant, if the delay is too short or too long then it outputs $(\mathsf{time\text{-}out}, sid)$ on its output tape and halts the execution.

Theorem 12 *Assume that $1 \leq \epsilon <$. Then the above scheduling is a non-trivial pairwise-disjoint scheduling, for the following parameters:*

$$\begin{aligned} \tau &\geq \alpha(n) \cdot \Delta \\ d &> (2\tau\epsilon^2 + \Delta(1+\epsilon)\epsilon)/(3 - 2\epsilon^3) \\ \tau_{\mathrm{MIN}}(x) &= (x+1)d/\epsilon - \Delta \\ \tau_{\mathrm{MAX}}(x) &= (x+1)d\epsilon + \Delta \end{aligned}$$

Note that the efficiency of the scheduling depends on ϵ. The closer ϵ is to $\sqrt[3]{1.5}$, the greater the delay is, and the less efficient the scheduling is. (This is due to the $(3 - 2\epsilon^3)$ factor in the denominator of d.)

Proof: First, we collect a few inequalities, which we shall refer to throughout the proof.

$$\tau_{\mathrm{MIN}}(0) > \tau\epsilon \tag{7}$$

$$\tau_{\mathrm{MIN}}(1) > (2\tau + \tau_{\mathrm{MAX}}(0))\epsilon \tag{8}$$

$$\tau_{\mathrm{MIN}}(2) > (2\tau + \tau_{\mathrm{MAX}}(0))\epsilon \tag{9}$$

$$\tau_{\mathrm{MIN}}(2) > (2\tau + \tau_{\mathrm{MAX}}(1))\epsilon \tag{10}$$

We note that these inequalities easily follow from the inequalities listed in the hypothesis.[41]

Assume for the sake of contradiction that S is not a pairwise-disjoint scheduling for some protocol σ, and timing parameters (Δ, ϵ) such that $1 \leq \epsilon < \sqrt[3]{1.5}$. Thus, there exists an ϵ-drift preserving adversary, and two distinct sessions sid and sid', such that the following holds. There exist honest parties P and P' participating in sessions sid and sid' respectively, such that according to P and P', every execution of σ in $S(\sigma, sid', \Delta, \epsilon)$ overlaps with at least one of the executions of σ in $S(\sigma, sid, \Delta, \epsilon)$. For simplicity of notation, throughout this proof we denote $S(\sigma, sid, \Delta, \epsilon)$ by Σ, and $S(\sigma, sid', \Delta, \epsilon)$ by Σ'. Further, we shall use "overlaps" as a short hand for "overlaps according to P and P'".

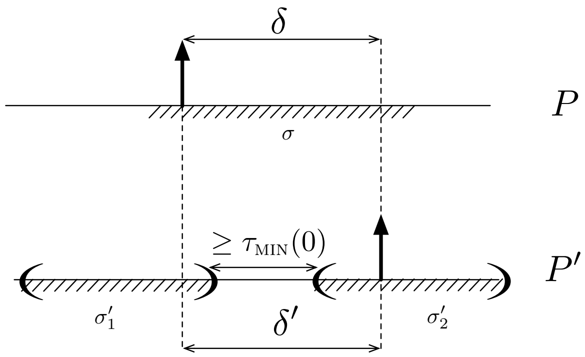

Figure 4: An illustration of the argument that a single execution σ in one schedule can overlap at most one execution in another schedule, with respect to two parties P and P'. In the scenario shown, the first overlap is according to situation (a) of Figure 3 and the second overlap is according to situation (b). δ is measured using P's local clock and δ' using P''s local clock.

We first show that any execution in Σ can overlap with at most one execution in Σ'. This is due to the delay inserted between each execution. More specifically, assume that there is one execution σ in Σ which overlaps with two executions σ'_1 and σ'_2 in Σ'. The fact that σ and σ'_1 overlap implies that there is a message (of σ or of σ'_1) that was sent while P was executing σ and while P' was executing σ'_1. Similarly, the fact that σ and σ'_2 overlap implies that there is a message (of σ or of σ'_2) that was sent while P was executing σ and while P' was executing σ'_2. Let the time between sending these two messages be δ as measured by the clock of P, and δ' as measured by the clock of P'. Since the clock drift factor is at most ϵ, we have $\delta' \leq \delta\epsilon$. Note that the executions σ'_1 and σ'_2 are separated by at least $\tau_{\text{MIN}}(0)$ local time units, according to P''s clock. This is the case since otherwise P' would timeout the execution before σ'_2 really started, which would imply that σ does not overlap σ'_2 according to P and P', contradicting our assumption. Thus, the above mentioned messages must also be separated by at least that much time, i.e., $\delta' \geq \tau_{\text{MIN}}(0)$. Finally, we note that since both the messages were sent out while the honest party P was executing σ, then $\delta \leq \tau$. Combining the above inequalities we get $\tau_{\text{MIN}}(0) \leq \delta' \leq \delta\epsilon \leq \tau\epsilon$. This contradicts Eq. (7). See

[41] This can be seen as follows. The denominator of the delay d is at most 1 (assuming $1 \leq \epsilon < \sqrt[3]{1.5}$), which implies that $d > 2\tau\epsilon^2 + \Delta(1+\epsilon)\epsilon$. Thus, $\tau_{\text{MIN}}(0) = d/\epsilon - \Delta > (2\tau\epsilon + \Delta(1+\epsilon)) - \Delta = 2\tau\epsilon + \Delta\epsilon > \tau\epsilon$, implying Eq. (7). Next, in order to prove Eq. (8) and Eq. (10) it suffices to prove that $\tau_{\text{MIN}}(x) - \tau_{\text{MAX}}(x-1)\epsilon > 2\tau\epsilon$ (this can be seen by simply manipulating the equations). In order to prove that $\tau_{\text{MIN}}(x) - \tau_{\text{MAX}}(x-1)\epsilon > 2\tau\epsilon$, note that $\tau_{\text{MIN}}(x) - \tau_{\text{MAX}}(x-1)\epsilon = (x+1)d/\epsilon - \Delta - (xd\epsilon + \Delta)\epsilon = d(x/\epsilon - x\epsilon^2 + 1/\epsilon) - \Delta(1+\epsilon)$. Since $1/\epsilon - \epsilon^2 \leq 0$, the latter equality is smallest when $x = 2$. Thus, $\tau_{\text{MIN}}(x) - \tau_{\text{MAX}}(x-1)\epsilon \geq d(2/\epsilon - 2\epsilon^2 + 1/\epsilon) - \Delta(1+\epsilon) = d(3-2\epsilon^3)/\epsilon - \Delta(1+\epsilon) > (2\tau\epsilon + \Delta(1+\epsilon) - \Delta(1+\epsilon) = 2\tau\epsilon$, as desired. Finally, note that Eq. (9) follows immediately from Eq. (10).

Figure 4 For an illustration of this argument.

We thus have that any execution in Σ can overlap with at most one execution in Σ'. This, together with our assumption that S is not a pairwise-disjoint scheduling, implies that every execution in Σ overlaps with exactly one execution in Σ'. Moreover, it must be the case that for every $l \in [m+2]$, the l^{th} execution in Σ overlaps only with the l^{th} execution in Σ'.

Fix any $l \in [m+1]$. Let x' be the number of zeros between the l^{th} one and the $l+1$'st one in $u^{sid'}$. Note that the encoding guarantees that $x' \in \{0,1,2\}$. We prove that the number of zeros between the l^{th} one and the $l+1$'st one in u^{sid} is also x'. This will imply that $u^{sid} = u^{sid'}$, which in turn will imply that $sid = sid'$, contradicting our assumption that sid and sid' are distinct.

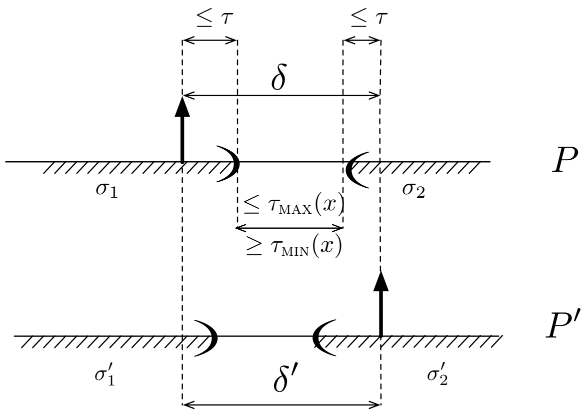

Figure 5: An illustration of inequality 11. All time intervals shown (except δ') are measured according to P's local clock.

Take any two (consecutive) executions σ_1 and σ_2 in Σ that overlap with two consecutive executions σ_1' and σ_2' in Σ' respectively. Let x be the number of zeros between the ones corresponding to σ_1 and σ_2 in u^{sid}. Similarly let x' be the number of zeros between the ones corresponding to σ_1' and σ_2' in $u^{sid'}$. We need to show that $x = x'$.

Since σ_1 and σ_1' overlap, a message (of σ_1 or of σ_1') was sent while P was executing σ_1 and while P' was executing σ_1'. Call the sending of this message the "first event." Similarly, since σ_2 and σ_2' overlap, a message (of σ_2 or of σ_2') was sent while P was executing σ_2 and while P' was executing σ_2'. We call the sending of this message the "second event." Let δ denote the duration between these two events according to the clock of P, and δ' the duration between them according to the clock of P'. Then,

$$\delta/\epsilon \leq \delta' \leq \delta\epsilon.$$

Now, since σ_1 and σ_2 are separated by x zeros, and P is an honest party, we are assured that

$$\tau_{\text{MIN}}(x) \leq \delta \leq \tau_{\text{MAX}}(x) + 2\tau. \tag{11}$$

This is the case since P, being honest, checks that each of these executions run for at most τ time units. It also checks that the delay between the last message of σ_1 and the first message of σ_2 is in the range $[\tau_{\text{MIN}}(x), \tau_{\text{MAX}}(x)]$. Note that these checks must be satisfied since otherwise P would timeout, and thus would not participate in σ_2. Therefore, σ_2 and σ_2' would not overlap according to P and P', contradicting our assumption. Since the first and second events occur in the middle of σ_1 and σ_2 respectively, we are assured that Inequality 11 holds. This is illustrated in Figure 5. Similarly, we are assured that

$$\tau_{\text{MIN}}(x') \le \delta' \le \tau_{\text{MAX}}(x') + 2\tau.$$

The above three displayed inequalities imply

$$\tau_{\text{MIN}}(x) \le \delta \le \delta'\epsilon \le (2\tau + \tau_{\text{MAX}}(x'))\epsilon$$
$$\tau_{\text{MIN}}(x') \le \delta' \le \delta\epsilon \le (2\tau + \tau_{\text{MAX}}(x))\epsilon$$

From these two inequalities we can easily derive contradictions for all the combinations $(x, x') = (1, 0)$, $(x, x') = (2, 0)$, $(x, x') = (0, 1)$, $(x, x') = (2, 1)$, $(x, x') = (0, 2)$ and $(x, x') = (1, 2)$. For instance, setting $(x, x') = (1, 0)$ or $(x, x') = (0, 1)$, we obtain

$$2\tau + \tau_{\text{MAX}}(0) \ge \delta' \ge \delta/\epsilon \ge \tau_{\text{MIN}}(1)/\epsilon$$

which contradicts Eq. (8). Similarly, setting $(x, x') = (2, 0)$ or $(x, x') = (0, 2)$ contradicts Eq. (9), and setting $(x, x') = (2, 1)$ or $(x, x') = (1, 2)$ contradicts Eq. (10). Hence we conclude that $x' = x$, as required. This shows that the scheduling is indeed pairwise disjoint.

It remains to show that it is non-trivial. For this, consider a scheduling Σ being executed in the presence of an adversary who does not corrupt any party and delivers all messages within time Δ by the clocks of *all* the parties. Firstly, since the protocol has $\alpha(n)$ rounds, setting the time-out for an individual execution to be $\tau = \alpha(n) \cdot \Delta$ ensures that no party times-out an execution. We need to also ensure that for every party, the checks on the delays between the executions are also satisfied. Recall our convention that a designated party sends out "start" and "end" messages to every party in the protocol; call this party P. For any two executions σ_1 and σ_2, corresponding to two ones with x zeros in between, party P delays $\delta \stackrel{\text{def}}{=} (x+1)d$ local time units between the "end" message of σ_1 and the "start" message of σ_2. By the clock of another party P' this duration will be measured as δ', where $\delta/\epsilon \le \delta' \le \delta\epsilon$. However P' considers the time at which these two messages reach it (rather than when they were sent). At one extreme, the "end" message may be delivered instantaneously and the subsequent "start" message delivered with a delay of Δ (according to P''s clock), in which case the time between the arrival of the two messages will be $\delta' + \Delta$. At the other extreme, "end" is delayed by Δ, while "start" reaches instantaneously, making the time between the two arrivals $\delta' - \Delta$. Thus, the delay between the two messages will be in the range $[\delta' - \Delta, \delta' + \Delta]$ which is in turn in the range $[\delta/\epsilon - \Delta, \delta\epsilon + \Delta]$. Since $\delta = (x + 1)d$, this range is the same as $[\tau_{\text{MIN}}(x), \tau_{\text{MAX}}(x)]$. Thus no party will time-out in the schedule. Also note that m and $O(\alpha(n))$ are bounded by a polynomial. Hence the schedule will be completed in polynomial number of steps and within polynomial number of time units according to any party. Thus the scheduling algorithm is polynomial and non-trivial. ■

5 Impossibility for Non-Delayed General Composition

In this section, we prove that in the timing model, in order to construct a protocol ρ that securely realizes $\mathcal{F}$ when running concurrently to some arbitrary protocol π, it holds that π must consider time in some way. (Recall that we assumed that π is a $\tau\epsilon$-delayed protocol.) In order to state this result, we first define the notion of a timing-free protocol. Intuitively, such a protocol does not use timing in its instructions. Formally, in our model, a timing-free protocol does not read the clock tape. (The "plain model" in the theorem refers to the model as defined in this paper, without for example, any trusted setup phase.)

Theorem 13 *In the plain model and without an assumed honest majority, there exist probabilistic polynomial-time functionalities that cannot be securely computed (by a non-trivial protocol) under concurrent general composition with timing-free protocols, even in the (Δ, ϵ)-timing model, for* any *Δ and* any $\epsilon \geq 1$.

We prove this theorem by showing that for every protocol ρ in the timing model, if ρ is secure under concurrent general composition with timing-free protocols, then it can be modified to become secure under *1-bounded parallel general composition* in a model with no timing. (In the setting of 1-bounded parallel general composition, a secure protocol ρ is executed *once* in parallel with an arbitrary protocol π.) This suffices for proving Theorem 13 because impossibility of this case is proven explicitly in [37]. As in [37], we also limit ourselves to 2-party protocols.

The intuition behind the proof of Theorem 13 is as follows. If a secure protocol ρ is run together with a timing-free protocol π, then this means that the adversary has *full control* over the scheduling of the messages of π. Now, consider a single execution of ρ together with π. Since the adversary can schedule π-messages as it wishes, it can force π to run perfectly in parallel with ρ. Notice that this holds irrespective of the timing instructions used in ρ. We conclude that ρ must remain secure when run in parallel with an arbitrary protocol π, in contradiction to the impossibility results of [37]. We now proceed to the formal proof.

Proof: Let $\Delta \geq 1$ and $\epsilon \geq 1$ be any values, and let ρ be a 2-party protocol that securely realizes a functionality $\mathcal{F}$ under concurrent general composition with timing-free protocols, in the (Δ, ϵ)-timing model.[42] Denote the participating parties by P_1 and P_2.

We construct a modified protocol ρ' that is timing-free. In ρ', instead of using the clock, the parties simulate the clock themselves by incrementing a counter *on each activation* that is is initialized to 0. (This is equivalent to setting the counter to equal the round number in the protocol, if it is defined.) This simulated clock is then made available to Protocol ρ (or more precisely, to the computation specified by Protocol ρ). Note that ρ' consists of two components: a clock simulation protocol and the original protocol ρ in the timing-model.

We now show that if ρ is non-trivial and secure under concurrent general composition in the timing model, then ρ' is non-trivial and secure under 1-bounded *parallel* general composition (in the timing-free model).[43] We note that if the adversary in the timing-free model activates the same party multiple times before activating the other party, then in ρ' the simulated clocks would have an unavoidable drift. This is problematic because in this case ρ does not give any guarantee of security. However, we consider *parallel* general composition for ρ'. In this setting, the adversary strictly alternates between activating P_1 and P_2. Furthermore, in the $i+1^{\text{th}}$ activation of a party, the adversary delivers it the i^{th}-round message from ρ' and the i^{th}-round message from π (where π is the arbitrary protocol running concurrently with ρ'). We call such an adversary for the parallel setting a *round-robin adversary.* The formal arguments are given in the proof of the following claim.

Claim 14 *Let ρ be a two-party protocol and let $\Delta, \epsilon \geq 1$ be any values. If ρ is non-trivial and securely realizes a functionality $\mathcal{F}$ under concurrent general composition in the (Δ, ϵ)-timing model (even when run concurrently with timing-free protocols), then ρ' as described above is a non-trivial*

[42] We stress that a contradiction will be derived for *any* choice of $\Delta, \epsilon \geq 1$. Note that $\epsilon \geq 1$ by definition, and that $\Delta \geq 1$ is the smallest increment possible.

[43] The notion of non-trivial protocols has also been considered in the timing-free model since the trivial protocol that just hangs and never generates output securely realizes all functionalities. Therefore, as in the timing model, only non-trivial protocols are of interest. In the timing-free model, a protocol is called non-trivial if output is guaranteed in the event that the adversary corrupts no parties and (eventually) delivers all messages. As expected, the impossibility results of [37] for parallel general composition hold only for non-trivial protocols.

protocol that securely realizes $\mathcal{F}$ under 1-bounded parallel general composition in the timing-free model.

Proof: Let π be an arbitrary timing-free two-party protocol. In order to prove the security claim on ρ', we need to show that for any given *round-robin adversary*, there exists a simulator $\mathcal{S}$ such that the output distribution of $\mathcal{A}$ and the honest parties running π and ρ' in the real model is computationally indistinguishable from the output distribution of $\mathcal{S}$ and the honest parties running π with ideal access to $\mathcal{F}$ in the $\mathcal{F}$-hybrid model. In order to construct $\mathcal{S}$, we first we show an intermediate adversary $\mathcal{H}$ (who interacts with the parties running ρ in the timing model) such that the output distributions of the adversary and honest parties in the following two scenarios are identical:

- *Scenario A:* The honest parties and the adversary $\mathcal{A}$ run π and ρ' in the timing-free (real) model.
- *Scenario B:* The honest parties and the adversary $\mathcal{H}$ run π and ρ in the real model with time.

We now describe the adversary $\mathcal{H}$ in the timing model. $\mathcal{H}$ internally invokes $\mathcal{A}$ and perfectly emulates all of $\mathcal{A}$'s actions. This means that $\mathcal{H}$ delivers messages whenever $\mathcal{A}$ does (thereby activating the recipients) and passes $\mathcal{A}$ the messages that it receives. In addition to this emulation, $\mathcal{H}$ needs to increment the clocks of the honest parties (because $\mathcal{H}$ works in the timing model, unlike $\mathcal{A}$). This is carried out simply by having $\mathcal{H}$ increment the clocks of all honest parties by 1 at the beginning of each round-robin round.

Before proceeding, we show that the outputs of the honest parties and adversaries are identical in scenarios A and B, described above. This follows from the fact that in both scenarios, the clock of each party is incremented by 1 between every activation. Furthermore, $\mathcal{H}$ carries out exactly the same actions as $\mathcal{A}$. (The only difference is that in scenario A, the clocks are updated in sequence upon each activation, whereas in scenario B, they are all updated together. However, since parties only read their clocks upon activation, this is exactly the same.) We therefore have that for every round-robin adversary $\mathcal{A}$ in the timing-free real model with π and ρ' there exists an adversary $\mathcal{H}$ in the timing model with π and ρ such that the output distributions in both cases are identical.

Next, notice that as long as $\mathcal{H}$ is ϵ-drift preserving, the assumed security of ρ implies that there exists a simulator $\mathcal{S}$ such that the output distribution of an execution with $\mathcal{S}$ and the honest parties running π in the $\mathcal{F}$-hybrid model is indistinguishable from an execution with $\mathcal{H}$ and the honest parties running π and ρ in the real timing model. This suffices because $\mathcal{H}$ satisfies the drift condition for any ϵ (notice that the clocks of all the honest parties are always the same). Combining the above two steps, we obtain that ρ' securely realizes $\mathcal{F}$ under one-bounded parallel general composition.

To complete the proof of the claim, we shall show that if ρ is non-trivial then so is ρ'. Recall that ρ' is non-trivial (in the timing-free model) if in the case that $\mathcal{A}$ corrupts no parties and delivers all messages, then all parties receive output. In order to see that this holds, first recall that $\mathcal{H}$ essentially just emulates $\mathcal{A}$. Therefore, if $\mathcal{A}$ corrupts no parties, then so does $\mathcal{H}$. Furthermore, by the assumption that $\mathcal{A}$ is a round-robin adversary, we know that it *always* delivers all messages immediately (i.e., all round i messages are received in round $i+1$). Therefore, $\mathcal{H}$ delivers all messages within time $\Delta = 1$. Finally, as we have shown above, $\mathcal{H}$ is always ϵ-drift preserving (for any $\epsilon \geq 1$). We conclude that in an execution of ρ' in which $\mathcal{A}$ does not corrupt any parties, the analogous execution of ρ with $\mathcal{H}$ is such that $\mathcal{H}$ corrupts no parties, is ϵ-drift preserving and delivers all messages within time $\Delta = 1$. Therefore, by Definition 5 and the assumption that ρ is non-trivial, we have that in this execution of ρ with $\mathcal{H}$, the honest parties all obtain their output (and this output does not equal time-out). By the equivalence between scenarios A and B above, we obtain that in the execution of ρ' with $\mathcal{A}$, the parties also all receive output. That is, ρ' is non-trivial. This completes the proof of non-triviality and of the claim. ■

As we have mentioned above, the proof of the theorem follows immediately from the above claim and the impossibility results for 1-bounded parallel general composition (in the timing-free model) as proven in [37]. ■

Remark. Theorem 13 states that there *exist* functionalities that cannot be securely computed under concurrent general composition with timing-free protocols. However, the proof actually shows that this setting inherits all of the impossibility results of [37], which are in turn inherited from [15]. Thus, we actually obtain very broad impossibility results that hold for large classes of functionalities.

Acknowledgements

We would like to thank Ran Canetti, Oded Goldreich, Shafi Goldwasser and Shai Halevi for helpful discussions and comments. We would also like to thank the anonymous referees for their many very helpful comments.

References

[1] B. Barak. How to Go Beyond the Black-Box Simulation Barrier. In *42nd FOCS*, pages 106–115, 2001.

[2] B. Barak, R. Canetti, Y. Lindell, R. Pass and T. Rabin. Secure Computation Without Authentication. In *CRYPTO 2005*, Springer-Verlag (LNCS 3621), pages 361–377, 2005.

[3] B. Barak, R. Canetti, J. B. Nielsen and R. Pass. Universally Composable Protocols with Relaxed Set-Up Assumptions. In *45th FOCS*, pages 186–195, 2004.

[4] B. Barak, Y. Lindell and S. Vadhan. Lower Bounds for Non-Black-Box Zero-Knowledge. In *44th FOCS,* pages 384–393, 2003.

[5] B. Barak and A. Sahai How To Play Almost Any Mental Game Over The Net – Concurrent Composition via Super-Polynomial Simulation. In *46th FOCS*, pages 543–552, 2005.

[6] D. Beaver. Foundations of Secure Interactive Computing. In *CRYPTO'91*, Springer-Verlag (LNCS 576), pages 377–391, 1991.

[7] M. Bellare and O. Goldreich. On Defining Proofs of Knowledge. In *CRYPTO'92*, Springer-Verlag (LNCS 740), pages 390–420, 1992.

[8] M. Ben-Or, S. Goldwasser and A. Wigderson. Completeness Theorems for Non-Cryptographic Fault-Tolerant Distributed Computation. In *20th STOC,* pages 1–10, 1988.

[9] M. Blum. Coin Flipping by Phone. *IEEE Spring COMPCOM*, pages 133–137, 1982.

[10] M. Blum. How to Prove a Theorem So No One Else Can Claim It. *Proceedings of the International Congress of Mathematicians*, pages 1444–1451, USA, 1987.

[11] R. Canetti. Security and Composition of Multiparty Cryptographic Protocols. *Journal of Cryptology*, 13(1):143–202, 2000.

[12] R. Canetti. Universally Composable Security: A New Paradigm for Cryptographic Protocols. In *42nd FOCS*, pages 136–145, 2001.

[13] R. Canetti. Universally Composable Signature, Certification, and Authentication. In the *17th Computer Security Foundations Workshop*, pages 219–235, 2004.

[14] R. Canetti and M. Fischlin. Universally Composable Commitments. In *CRYPTO 2001*, Springer-Verlag (LNCS 2139), pages 19–40, 2001.

[15] R. Canetti, E. Kushilevitz and Y. Lindell. On the Limitations of Universal Composable Two-Party Computation Without Set-Up Assumptions. *Journal of Cryptology,* 19(2):135-167, 2006.

[16] R. Canetti, Y. Lindell, R. Ostrovsky and A. Sahai. Universally Composable Two-Party and Multi-Party Computation. In *34th STOC*, pages 494–503, 2002.

[17] D. Chaum, C. Crepeau and I. Damgard. Multiparty Unconditionally Secure Protocols. In *20th STOC*, pages 11–19, 1988.

[18] B. Chor and M. Rabin. Achieving Independence in Logarithmic Number of Rounds. In *6th PODC*, pages 260–268, 1987.

[19] A. De Santis and G. Persiano. Zero-Knowledge Proofs of Knowledge Without Interaction. In *33rd FOCS,* pages 427–436, 1992.

[20] Y. Dodis and S. Micali. Parallel Reducibility for Information-Theoretically Secure Computation. In *CRYPTO 2000*, Springer-Verlag (LNCS 1880), pages 74–92, 2000.

[21] D. Dolev, C. Dwork and M. Naor. Non-Malleable Cryptography. *SIAM Journal on Computing,* 30(2):391–437, 2000.

[22] C. Dwork, M. Naor, and A. Sahai. Concurrent Zero-Knowledge. *Journal of the ACM*, 51(6):851–898, 2004.

[23] C. Dwork and M. Naor. Zaps and Their Applications. In *41st FOCS*, pages 283–293, 2000.

[24] C. Dwork and A. Sahai. Concurrent Zero-Knowledge: Reducing the Need for Timing Constraints. In *CRYPTO'98*, Springer-Verlag (LNCS 1462), pages 442–457, 1998.

[25] U. Feige and A. Shamir. Zero-Knowledge Proofs of Knowledge in Two Rounds. In *CRYPTO'89*, Springer-Verlag (LNCS 435), pages 526–544, 1989.

[26] U. Feige and A. Shamir. Witness Indistinguishability and Witness Hiding Protocols. In *22nd STOC*, pages 416–426, 1990.

[27] O. Goldreich. *Foundations of Cryptography: Volume 1 – Basic Tools.* Cambridge University Press, 2001.

[28] O. Goldreich. *Foundations of Cryptography: Volume 2 – Basic Applications.* Cambridge University Press, 2004.

[29] O. Goldreich. Concurrent Zero-Knowledge With Timing Revisited. In 34*th STOC*, pages 332–340, 2002.

[30] O. Goldreich, S. Micali and A. Wigderson. How to Play any Mental Game – A Completeness Theorem for Protocols with Honest Majority. In 19*th STOC,* pages 218–229, 1987.

[31] 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.

[32] S. Goldwasser and Y. Lindell. Secure Computation Without Agreement. *Journal of Cryptology,* 18(3):247–287, 2005.

[33] S. Goldwasser, S. Micali and C. Rackoff The Knowledge Complexity of Interactive Proof Systems. *SIAM Journal on Computing,* 18(1):186–208, 1989.

[34] J. Katz. Efficient and Non-malleable Proofs of Plaintext Knowledge and Applications. In *EUROCRYPT 2003*, Springer-Verlag (LNCS 2656), pages 211–228, 2003.

[35] Y. Lindell. Parallel Coin-Tossing and Constant-Round Secure Two-Party Computation. *Journal of Cryptology*, 16(3):143–184, 2003.

[36] Y. Lindell. Bounded-Concurrent Secure Two-Party Computation Without Setup Assumptions. In 35*th STOC*, pages 683–692, 2003.

[37] Y. Lindell. General Composition and Universal Composability in Secure multiparty Computation. In 44*th FOCS*, pages 394–403, 2003.

[38] Y. Lindell. Lower Bounds for Concurrent Self Composition. In the *1st Theory of Cryptography Conference* (TCC), Springer-Verlag (LNCS 2951), pages 203–222, 2004.

[39] T. Malkin, R. Moriarty and N. Yakovenko. Generalized Environmental Security from Number Theoretic Assumptions. In the *3rd Theory of Cryptography Conference* (TCC), Springer-Verlag (LNCS 3876), pages 343–359, 2006.

[40] S. Micali and P. Rogaway. Secure Computation. Unpublished manuscript, 1992. Preliminary version in *CRYPTO'91*, Springer-Verlag (LNCS 576), pages 392–404, 1991.

[41] M. Naor. Bit Commitment using Pseudorandom Generators. *Journal of Cryptology*, 4(2):151–158, 1991.

[42] R. Pass. Simulation in Quasi-Polynomial Time, and Its Application to Protocol Composition. In *Eurocrypt 2003*, Springer-Verlag (LNCS 2656), pages 160–176, 2003.

[43] R. Pass. Bounded-Concurrent Secure Multiparty Computation with a Dishonest Majority. In the 36*th STOC,* pages 232–241, 2004.

[44] R. Pass and A. Rosen Bounded-Concurrent Secure Two-Party Computation in a Constant Number of Rounds. In 44*th FOCS*, pages 404–413, 2003.

[45] 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.

[46] M. Prabhakaran and A. Sahai. New Notions of Security: Universal Composability Without Trusted Setup. In 36*th STOC*, pages 242–251, 2004.

[47] R. Richardson and J. Kilian. On the Concurrent Composition of Zero-Knowledge Proofs. In *EUROCRYPT'99*, Springer-Verlag (LNCS 1592), pages 415–431, 1999.

[48] A. Yao. How to Generate and Exchange Secrets. In 27*th FOCS*, pages 162–167, 1986.