Universally Composable Synchronous Computation

Jonathan Katz
University of Maryland
jkatz@cs.umd.edu

Ueli Maurer
ETH Zurich
maurer@inf.ethz.ch

Björn Tackmann
ETH Zurich
bjoern.tackmann@inf.ethz.ch

Vassilis Zikas
University of Maryland
vzikas@cs.umd.edu

Abstract

In synchronous networks, protocols can achieve security guarantees that are not possible in an asynchronous world: i.e., they can simultaneously achieve *input completeness* (all honest parties' inputs are included in the computation) and *guaranteed termination* (honest parties do not "hang" indefinitely). In practice truly synchronous networks rarely exist, but synchrony can be emulated if channels have (known) latency and parties have loosely synchronized clocks.

The framework of universal composability (UC) is inherently asynchronous, but several approaches for adding synchrony to the framework have been proposed. However, we show that the existing proposals do *not* provide the expected guarantees. Given this, we propose a "clean slate" approach to defining synchrony in the UC framework by introducing functionalities exactly meant to model, respectively, bounded-delay networks and loosely synchronized clocks. We show that the expected guarantees of synchronous computation can be realized given these functionalities, and that previous models can all be expressed within our new framework.

1 Introduction

In synchronous networks, protocols can achieve both *input completeness* (all honest parties' inputs are included in the computation) and *guaranteed termination* (honest parties do not "hang" indefinitely). On the other hand, it is known that these properties cannot simultaneously be ensured in an asynchronous world [CM89, BCG93].

The classical synchronous communication model assumes that protocols proceed in rounds: the current round is known to all parties, and messages sent in some round are delivered by the beginning of the next round. While this is a strong model that rarely corresponds to real-world networks, the model is still useful since it can be *emulated* under the relatively mild assumptions of a known bound on the network latency and loose synchronization of the (honest) parties' clocks. In fact, it is fair to say that these two assumptions are exactly what is meant when speaking of "synchrony" in real-world networks.

The framework of *universal composability* (UC) [Can01] assumes, by default, completely asynchronous communication; in fact, even eventual message delivery is not guaranteed. (We assume familiarity with the UC framework; see Section 2 for a high-level overview.) Protocol designers working in the UC setting are thus faced with two choices: either work in an asynchronous network and give up on input completeness [BCG93] or guaranteed termination [KLR06], or else modify the UC framework so as to incorporate synchronous communication somehow.

Several ideas for adding synchrony to the UC framework have been proposed. Canetti [Can05] introduced an ideal functionality $\mathcal{F}_{\text{SYN}}$ that was intended exactly to model synchronous communication in a general-purpose fashion. We prove later, however, that $\mathcal{F}_{\text{SYN}}$ does *not* provide the guarantees expected of a synchronous network (cf. Section 5.1). Nielsen [Nie03] and Hofheinz and Müller-Quade [HM04] also propose ways of modeling synchrony with composition guarantees, but their approaches are *ad hoc* (in that they are specifically tailored to their desired applications) and are not sufficiently general to model, e.g., synchrony in an incomplete network, or the case when synchrony holds only in part of a network (say, because certain links do not have bounded delay while others do). In all these cases it appears fair to say that the proposed modifications to the UC framework are complex, to the point where it is unclear that they are adequately capturing some intuitive notion of synchrony. The timing model considered in, e.g., [DNS98, Gol02, KLP05] comes closer to capturing this intuition, but it turns out (as we show later in this work) that they also do not provide the guarantees expected from a synchronous network.

1.1 Our Results

As just discussed, existing approaches to modeling synchrony in the UC framework are either ad hoc, or fail to capture the properties desired from synchronous networks. This motivates us to explore a "clean slate" approach to modeling synchronous communication in the UC framework. In doing so, we also aim for an intuitively appealing model that more closely corresponds to the actual synchrony assumptions being made in the real world. For this reason, the approach we take is to introduce functionalities specifically intended to (independently) model the two assumptions of bounded network delay and loose clock synchronization. An additional benefit of separating the assumptions in this way is that we can also study the case when only one of the assumptions holds.

We begin by formally defining a functionality corresponding to (authenticated) communication channels with *bounded delay*. Unfortunately, this alone is not sufficient for achieving guaranteed termination. (Throughout, we will always want input completeness to hold.) Intuitively, this is because bounded-delay channels alone—without any global clock—only provide the same "eventual message delivery" guarantee of classical asynchronous networks [BCG93, Can96]. It thus becomes clear that what is missing when only bounded-delay channels are available is some notion of *time*. To rectify this, we further introduce a functionality $\mathcal{F}_{\text{CLOCK}}$ that directly corresponds to the presence of loosely synchronized clocks among some set of parties. After defining an appropriate notion of terminating functionalities (exactly meant to capture the security requirement of guaranteed termination), we then show that standard protocols can be used to securely realize any *terminating* functionality in a hybrid world where $\mathcal{F}_{\text{CLOCK}}$ and bounded-delay channels are available. Note that previous approaches for addressing guaranteed termination in UC restrict the class of considered adversaries, namely to those that *explicitly allow* the protocol to terminate; as a result of this restriction, the models do no longer provide universal composition guarantees. To the contrary, our approach is to model termination as a property of functionalities, which allows us to maintain composition. Furthermore, it bridges the gap between the theoretical model and the realistic scenario, where the synchronized clocks of the parties ensure that the adversary *cannot* stall their local computation *even if he tries to*. Overall, our results show that the two functionalities we propose—meant to model, independently, bounded network delay and loosely synchronized clocks—enable us to capture exactly the security guarantees provided by traditional synchronous networks.

We then revisit prior approaches to modeling synchrony in light of our results. In particular, we show that Canetti's $\mathcal{F}_{\text{SYN}}$ functionality can be realized (essentially) from bounded-delay channels alone! Given our results showing that bounded-delay channels are insufficient for guaranteed termination, this implies that $\mathcal{F}_{\text{SYN}}$—which, recall, was specifically introduced as a way of modeling the guarantees provided by fully synchronous networks—cannot provide guaranteed termination. This serves as a good indication of the subtle nature of the problem. On the positive side, we show a general way to translate results from the existing models into our framework while additionally guaranteeing termination.

1.2 Organization of the Paper

In Section 2 we include a brief description of the UC model [Can01] and introduce the necessary notation and terminology. In Section 3, we review the model of completely asynchronous networks, describe its limitations, and introduce a functionality modeling bounded-delay channels. In Section 4 we introduce a functionality $\mathcal{F}_{\text{CLOCK}}$ meant to model loose clock synchronization; there, we also explore the guarantees provided by this functionality. Further, we define the notion of computation with guaranteed termination within the UC framework, and show how this notion can be achieved using $\mathcal{F}_{\text{CLOCK}}$ and bounded-delay channels. In Section 5, we revisit previous models for synchronous computation.

2 Preliminaries

Simulation-based Security. Simulation-based security definitions follow the real-world/ideal-world paradigm: In the real world, the players execute the protocol and can communicate over channels as defined by the model. In the ideal world, the players securely access an ideal functionality $\mathcal{F}$ that obtains inputs from the players, runs the program that specifies the task to be achieved by the protocol, and returns the resulting outputs to the players. Intuitively, a protocol *securely realizes* the functionality $\mathcal{F}$ if, for any real-world adversary $\mathcal{A}$ attacking the protocol execution, there is an ideal-world adversary $\mathcal{S}$, also called the *simulator*, that emulates $\mathcal{A}$'s attack in the ideal evaluation of $\mathcal{F}$. The "quality" of the simulation is specified by considering a distinguisher $\mathcal{Z}$, called the *environment*, which interacts, in a well defined manner, with the parties and the adversary or the simulator and tries to distinguish between the two worlds.

The advantage of this type of security definition is that it satisfies strong composability properties: let π_1 be a protocol that securely realizes a functionality $\mathcal{F}_1$. If a protocol π_2, using the functionality $\mathcal{F}_1$ as a subroutine, securely realizes a functionality $\mathcal{F}_2$, then the protocol $\pi_2^{\pi_1/\mathcal{F}_1}$, which results when replacing the calls to $\mathcal{F}_1$ by invocations of π_1, securely realizes $\mathcal{F}_2$ (without calls to $\mathcal{F}_1$). Therefore, it is sufficient to analyze the security of the simpler protocol π_2 in the $\mathcal{F}_1$-*hybrid* model, where the players run π_2 with access to the ideal functionality $\mathcal{F}_1$. For more details on composability of protocols and a formal handling of composition, the reader is referred to, e.g., [Can00, DM00, Can01, BPW07, MR11].

Security Models Considered in this Work. All security models discussed in this work are similar to the simulation-based UC framework [Can01], and the definitions are based on the same paradigm. The entities taking part in the execution (protocol machines, functionalities, adversary, and environment) are described as *interactive Turing machines* (ITMs) that have specific commands to achieve communication by writing to each other's tapes. The execution is initiated by the environment that provides input to and obtains output from the protocol machines, and also communicates with the adversary. The adversary has access to the ideal functionalities in the hybrid models and also serves as a network among the protocol machines. During the execution, the ITMs are activated one-by-one, where the exact order of the activations depends on the considered model.

Functionalities and Protocols. The functionalities considered in this work are (reactive) n-party functionalities, specified as interactive Turing machines that communicate with the (honest) parties taking part in the actual protocol execution and the adversary. The specification of the functionality also models the adversary's ability to obtain leaked information and/or to influence the computation, also in case the adversary corrupts parties (which models misbehavior). A protocol is specified as an ITM π executed by the parties and potentially ideal functionalities (if needed to execute π).

Notation and Convention. We consider protocol executions among a certain set of players $\mathcal{P}$, often referred to as the *player set*, where every $p_i \in \mathcal{P}$ is associated to a unique party ID. We refer to a party with party ID i as p_i or simply as party i.[1] A protocol execution involves the following types of ITMs: the environment $\mathcal{Z}$, the adversary $\mathcal{A}$, the protocol machines $\pi_1, \ldots, \pi_n$, and (possibly) ideal functionalities $\mathcal{F}_1, \ldots, \mathcal{F}_m$. The contents of the environment's output tape after the execution is described by the random variable $\text{EXEC}_{\pi,\{\mathcal{F}_1,\ldots,\mathcal{F}_m\},\mathcal{A},\mathcal{Z}}(k, z)$, where $k \in \mathbb{N}$ is the *security parameter* and $z \in \{0,1\}^*$ is the input to the environment $\mathcal{Z}$.[2] We use the following convention: when no protocol π is specified in the above notation, we mean that we take as π the "dummy" protocol φ which forwards all its inputs to the ideal functionality and directly outputs every message it receives from it (see [Can05]). We say that a protocol π securely realizes $\mathcal{F}$ in the $\mathcal{F}'$-hybrid model if

$$\forall \mathcal{A}\ \exists \mathcal{S}\ \forall \mathcal{Z} : \text{EXEC}_{\pi,\mathcal{F}',\mathcal{A},\mathcal{Z}} \approx \text{EXEC}_{\mathcal{F},\mathcal{S},\mathcal{Z}},$$

where "$\approx$" denotes indistinguishability of the respective distribution ensembles.

Delayed outputs. As in [Can05], the statement "the functionality sends a *(private) delayed output* y to party i" describes the following process: the functionality sends a message to the adversary requesting his permission to output y to party i; as soon as the adversary agrees, the output y is delivered. The statement "the functionality sends a *public delayed output* y to party i" corresponds to the same process, where the permission request also includes the full message y.

[1] In order to simplify the description, at times we use $\mathcal{P}$ as being the set of the corresponding party ID's.

[2] For the details of the execution, we follow the description in [Can05]. The slight change of notation has been made to better represent the assumed hybrids with a view towards the reductions.

All of our functionalities $\mathcal{F}$ use *standard corruption* as defined in [Can05], which is also described in Section A.1 of the appendix. At any point in the execution, we denote by $\mathcal{H}$ the set of "honest" parties that have not (yet) been corrupted. Finally, all of our functionalities use a fixed player set $\mathcal{P}$, and formally we assume that the session ID of a functionality $\mathcal{F}$ is of the form $sid = (\mathcal{P}, sid')$ with $sid' \in \{0,1\}^*$.

3 Synchronous Protocols in an Asynchronous Environment

Protocols in asynchronous networks cannot achieve input completeness and guaranteed termination simultaneously [CM89, BCG93]. Intuitively, the reason is that honest parties cannot distinguish the case where an input has been delayed—and to satisfy input completeness they should wait for this input—from the case where the sender is corrupted and did not send a message. In particular, when modeling secure computation within the UC framework, one typically relaxes (at least one of) the above properties. More precisely, there are two main network models for asynchronous protocols: On the one hand, there are fully asynchronous channels that only provide authenticity and/or privacy guarantees [CLOS02, Can05]; on the other hand, there is the model of channels with bounded delay which might be a publicly known constant or unknown [BCG93]. In the following, we formalize the channels assumed in each of the two settings and discuss how the resulting networks can be used by synchronous protocols. The results presented in this section formally confirm—in the language of the UC framework—facts about the power of synchrony assumptions that are known or folklore in the distributed systems literature.

3.1 Fully Asynchronous Network

We model the communication in a fully asynchronous network by the functionality $\mathcal{F}_{\text{SMT}}$ from [Can05]. This functionality corresponds to a fully asynchronous channel. As soon as p_s inputs the message m he wants to send, $\mathcal{F}_{\text{SMT}}$ informs the adversary. The asynchrony of the channel is modeled by the fact that the channel delivers the received message to p_r only when the adversary instructs it to. To model different privacy-guarantees, $\mathcal{F}_{\text{SMT}}$ is parametrized by a leakage function $\ell(\cdot)$ that determines the information about m that is leaked to the adversary in case both p_s and p_r are honest. In particular, the authenticated channel, denoted as $\mathcal{F}_{\text{AUTH}}$, is modeled by $\mathcal{F}_{\text{SMT}}$ parametrized by the identity function $\ell(m) = m$, and the secure channel, denoted as $\mathcal{F}_{\text{SEC}}$, is modeled by $\mathcal{F}_{\text{SMT}}$ with the constant function $\ell(m) = \perp$.[3] An important feature of $\mathcal{F}_{\text{SMT}}$ is that it allows for *adaptive message replacement*, i.e., for an honest sender sending a message m, the adversary can, depending on the leakage $\ell(m)$, corrupt the sender and replace the sent message. A detailed description of $\mathcal{F}_{\text{SMT}}$ is given in Section A.2 of the appendix.

In this model, Canetti et al. [CLOS02] showed that, assuming a common reference string, any (well-formed) functionality can be implemented, without guaranteed termination, by an $\{\mathcal{F}_{\text{SMT}}, \mathcal{F}_{\text{CRS}}\}$-hybrid protocol. Moreover, the combination of [KLR06, AL11] shows that appropriate modifications of the protocols of [GMW87, BGW88, CCD88] (which are designed for the synchronous setting) are sufficient to achieve general secure computation without termination in this asynchronous setting.

The following lemma formalizes the intuition that a fully asynchronous network cannot be used for implementing terminating computation, i.e., computation which cannot be stalled by the adversary. More precisely, for a functionality $\mathcal{F}$, denote by $[\mathcal{F}]^{\text{NT}}$ the *non-terminating relaxation* of $\mathcal{F}$ which is defined as follows: $[\mathcal{F}]^{\text{NT}}$ is the functionality which behaves as $\mathcal{F}$, but whenever $\mathcal{F}$ outputs a value to some honest party, $[\mathcal{F}]^{\text{NT}}$ provides this output in a delayed manner.[4] The lemma states that there are functionalities $\mathcal{F}$ that are not realizable in the $\mathcal{F}_{\text{SMT}}$-hybrid model, but their delayed relaxations $[\mathcal{F}]^{\text{NT}}$ are. We point out that the statement of the lemma holds even for stand-alone security, i.e., for environments that do not interact with the adversary during the protocol execution. Additionally, the impossibility applies to all functionalities that cannot be computed locally[5] (see [KMR09]) with guaranteed termination as defined in Section 4 (which is an implication of Lemma 6), and even it we assume *bidirectional* channels, i.e., for each copy of $\mathcal{F}_{\text{SMT}}$, an independent copy of $\mathcal{F}_{\text{SMT}}$ is assumed with the roles of sender and receiver inverted.

Lemma 1. *There are functionalities $\mathcal{F}$ such that $[\mathcal{F}]^{NT}$ can be implemented in the $\mathcal{F}_{\text{SMT}}$-hybrid model, but $\mathcal{F}$ cannot be implemented, even if we assume bidirectional $\mathcal{F}_{\text{SMT}}$-channels.*

[3] This correspond to *ideally* secure channels as defined in [BGW88, CCD88]. Yet, as such a channel cannot be implemented without further assumptions, one typically resorts to the length function $\ell(m) = |m|$ [CK02].

[4] For a discussion about delayed outputs we refer to Section 2.

[5] In the case of SFE, we refer to locally computable functions as *trivial*.

Proof (idea). Consider the functionality $\mathcal{F}$ which behaves as $\mathcal{F}_{\text{SMT}}$, but with the following add-on: upon receiving a special "fetch" message from the receiver p_r, outputs y to p_r, where $y = m$ if the sender has input the message m, and $y = \perp$ (i.e., a default value), otherwise.[6] It is straightforward to verify that $\mathcal{F}$ cannot be realized from $\mathcal{F}_{\text{SMT}}$ channels, but $[\mathcal{F}]^{\text{NT}}$ can be realized from $\mathcal{F}_{\text{SMT}}$ by simply invoking the dummy protocol. For a detailed proof, we refer to Appendix A.2. □

3.2 Eventual-Delivery Channels

The functionality $\mathcal{F}_{\text{SMT}}$ does not capture the guarantee that a message *will* be delivered *eventually*, independently of the adversary's strategy; a property often attributed to asynchronous communication [BCG93]. We show how to modify $\mathcal{F}_{\text{SMT}}$ to capture this guarantee (similar to the functionality $\mathcal{F}$ from the proof of Lemma 1): The receiver can enforce delivery of the message using "fetch" requests to the channel. The potential delay of the channel is modeled by ignoring a certain number D of such requests before delivering the actual message to p_r; to model the fact that the delay might be arbitrary, we allow the adversary to set the value of D. Yet, $\mathcal{A}$ is allowed to set it only once and cannot choose a delay which is longer than its running time.[7] We denote this modified functionality as $\mathcal{F}_{\text{ED-SMT}}$. Similar to $\mathcal{F}_{\text{SMT}}$, $\mathcal{F}_{\text{ED-SMT}}$ is parametrized by a leakage function $\ell(\cdot)$.

Functionality $\mathcal{F}_{\text{ED-SMT}}(p_s, p_r, \ell(\cdot))$

Initialize $M := \perp$ and $D := 0$.

- Upon receiving a message m from p_s, set $D := 1$ and $M := m$ and send $\ell(M)$ to the adversary.
- Upon receiving a message $(\texttt{fetch})$ from p_r
  1. Set $D := D - 1$.
  2. If $D = 0$ then send M to p_r (otherwise no message is sent and, as it is defined in the UC framework, the activation is given back to the environment).
- Upon receiving a message $(\texttt{delay}, T)$ from the adversary, if T is a valid delay (i.e., encodes a natural number in unary notation) and no $\texttt{delay}$ message was received so far, then set $D := T$; otherwise ignore the message.
- *(adaptive message replacement)*: Upon receiving a $(\texttt{corrupt}, p_s, m', T')$ from $\mathcal{A}$: if $D > 0$ and T' is a valid delay, then set $D = T'$ and set $M := m'$.

Unlike fully asynchronous channels, eventual-delivery channels allow for protocols which are guaranteed to (eventually) terminate, at the cost of violating input completeness.[8] For instance, the protocol of [BCG93] securely implements any functionality where the inputs of up to $\frac{n}{4}$ parties might be ignored, i.e., replaced by a default value. Alternatively, one can achieve secure computation of a functionality without termination, similarly to the fully asynchronous channels setting. Yet, the eventual-delivery channels, by themselves, do not allow to compute functionalities with strong termination guarantees. In fact, the result of Lemma 1 holds even if we replace $\mathcal{F}_{\text{SMT}}$ by $\mathcal{F}_{\text{ED-SMT}}$. This is stated in the following lemma, which again translates to both stand-alone security and to arbitrary functionalities that are not locally computable. The proof is similar to the one of Lemma 1 and appears in Appendix A.3.

Lemma 2. *There are functionalities $\mathcal{F}$ such that $[\mathcal{F}]^{NT}$ can be implemented in the $\mathcal{F}_{\text{ED-SMT}}$-hybrid model, but $\mathcal{F}$ cannot be implemented, even if we assume bidirectional $\mathcal{F}_{\text{ED-SMT}}$-channels.*

3.3 Bounded-Delay Channels with a Known Upper-Bound

One might be tempted to think that the impossibility of Lemma 2 is based on the fact that no upper bound on the network's latency is known. We show that this is not the case. In particular, we describe a functionality that models

[6] As we shall show, this functionality corresponds to a channel with zero-delay.

[7] Technically, this is enforced by accepting the delay-number only when given in unary notation.

[8] In particular, Lemma 1 extends to the case where $\mathcal{F}$ is this eventual-delivery channel $\mathcal{F}_{\text{ED-SMT}}$: The simulator can extend the delivery of the message only by a polynomial number of steps, and the environment can issue sufficiently many queries at the receiver's interface.

a channel with a known upper bound δ on its latency; note that this is one of the two standard assumptions in the synchronous model. This functionality works as $\mathcal{F}_{\text{ED-SMT}}$, but it is parametrized by a (strictly) positive number δ which is an upper bound on the delay of the channels, i.e, delays $T > \delta$ chosen by the adversary are ignored. Furthermore, the sender/receiver can query the functionality for learning the value of δ. The δ-delayed channels functionality $\mathcal{F}^{\delta}_{\text{BD-SMT}}$ is described in detail in Appendix A.4.

In reality a channel with latency δ' can be used to substitute a channel with latency $\delta > \delta'$. Using our formulation of bounded-delay channels, this intuition translates into the following UC statement: "for any $0 < \delta' < \delta$, $\mathcal{F}^{\delta}_{\text{BD-SMT}}$ can be implemented in the $\mathcal{F}^{\delta'}_{\text{BD-SMT}}$-hybrid model". Indeed, the simple $\mathcal{F}^{\delta'}_{\text{BD-SMT}}$-hybrid protocol that drops $\delta - \delta'$ $(\texttt{fetch})$-queries implements $\mathcal{F}^{\delta}_{\text{BD-SMT}}$, where the simulator only follows the adversary's instructions. However, the oposite is not in general true. The following lemma confirms, and formalizes in UC, this separation, i.e., that channels with smaller upper bound on the delay are *strictly* stronger when termination is required. The proof is similar to the proof of Lemma 2 and can be found in Appendix A.4. Again, the impossibility extends to both stand-alone security and, as shown in Lemma 6, (not locally computable) functionalities which satisfy the guaranteed termination notion defined in the following section.

Lemma 3. *For any $0 < \delta' < \delta$ the functionality $\left[\mathcal{F}^{\delta'}_{\text{BD-SMT}}\right]^{NT}$ can be implemented in the $\mathcal{F}^{\delta}_{\text{BD-SMT}}$-hybrid model, but $\mathcal{F}^{\delta'}_{\text{BD-SMT}}$ cannot be implemented, even if we assume bidirectional $\mathcal{F}^{\delta}_{\text{BD-SMT}}$-channels.*

The above lemma already suggests that bounded-delay channels, without additional synchronization assumptions such as synchronized clocks, are not sufficient for terminating computation. While Lemma 3 only handles the case where the assumed channel has a strictly positive upper-bound on the delay, the (more general) impossibility in Lemma 6 holds even for *instant-delivery* channels, i.e., bounded-delay channels which become ready to deliver as soon as they get input from the sender.

4 Synchronous Computation with Guaranteed Termination

The assumption of bounded-delay channels (without clocks, these degrade to channels with eventual delivery), is not by itself sufficient for achieving both input completeness and termination for several natural functionalities. In this section, we introduce the functionality $\mathcal{F}_{\text{CLOCK}}$ that, together with the bounded-delay channels $\mathcal{F}^{\delta}_{\text{BD-SMT}}$, allows synchronous protocols to satisfy both properties simultaneously (in a universally composable way). In particular, we define what it means for a protocol to UC-implement a given multi-party function *with guaranteed termination*, and show how to construct an $\{\mathcal{F}_{\text{CLOCK}}, \mathcal{F}^{\delta}_{\text{BD-SMT}}\}$-protocol that satisfies this definition.

4.1 The Functionality $\mathcal{F}_{\text{CLOCK}}$

To motivate the functionality $\mathcal{F}_{\text{CLOCK}}$, we examine how synchronous protocols in reality use the assumptions of bounded-delay (with a known upper bound) channels and synchronized clocks to satisfy the input-completeness and the termination properties simultaneously. In particular, this is achieved by having the players assign to each round a time-slot that is long enough to incorporate the time for computing and sending all next-round messages, plus the network delay. Note that it is sufficient, at the cost of having longer rounds, to assume that the clocks are not advancing in a fully synchronized manner but there is an known upper bound on the maximum clock-drift [Gol02, KLP05].

The above discussion demonstrates that one of the guarantees used by synchronous protocols in reality is that, at certain points in their execution, the parties know that all honest parties have performed all their operations for some round. Furthermore, these points do not depend on the adversary's behavior. The goal of the functionality $\mathcal{F}_{\text{CLOCK}}$ is to provide this guarantee to UC protocols. But as $\mathcal{F}_{\text{CLOCK}}$ is an ordinary UC functionality, it has no means of knowing whether a player has finished with its intended operations for a certain round. This problem is resolved by having each player $p \in \mathcal{P}$ notify $\mathcal{F}_{\text{CLOCK}}$ when it reaches this state. In particular, $\mathcal{F}_{\text{CLOCK}}$ works as follows: It keeps track of the set of players who have completed their actions in a vector of indicator bits $(d_i, \ldots, d_n)$, where $d_i = 1$ if p_i has reported that it has performed all its actions for the current round and $d_i = 0$, otherwise. As soon as every honest party has reported that it has performed all its actions, $\mathcal{F}_{\text{CLOCK}}$ resets $d_i = 0$ for every $p_i \in \mathcal{P}$.[9] In

[9]As usual for UC functionalities, whenever some party is corrupted, $\mathcal{F}_{\text{CLOCK}}$ is notified and updates the set $\mathcal{H}$ accordingly.

addition to notifications, any player p_i might send a synchronization request to $\mathcal{F}_{\text{CLOCK}}$, which is answered with the actual value of d_i. If p_i observes that d_i has switched, then it can conclude that all honest parties have completed their respective duties. Note that $\mathcal{F}_{\text{CLOCK}}$ does *not* wait for the corrupted parties to report that they are done with their actions, which also means that $\mathcal{F}_{\text{CLOCK}}$ cannot be realized based on well-formed functionalities. Nevertheless, as discussed above, it can be implemented by synchronous protocol in reality by using time. In order to allow the adversary enough time to perform its attack between the synchronization points, whenever $\mathcal{F}_{\text{CLOCK}}$ receives a notification from an honest player, it informs the adversary about this fact.

Functionality $\mathcal{F}_{\text{CLOCK}}(\mathcal{P})$

Initialize for each $i \in \mathcal{P}$ a bit $d_i = 0$.

- Upon receiving message (`RoundOK`) from Party i set $d_i := 1$. If for all $j \in \mathcal{H} : d_j = 1$, then reset $d_j := 0$ for all $j \in \mathcal{P}$. Send $(\texttt{switch}, i)$ to $\mathcal{A}$.
- Upon receiving (`RequestRound`) from p_i, send d_i to p_i.

In the remainder of this paper, we use instant-delivery channels, i.e., $\mathcal{F}^{\delta}_{\text{BD-SMT}}$ with $\delta = 1$. This simplifies the notation and is not a restriction of the model and can be achieved in reality by extending the delays appropriately. Consequently, we completely omit the delay parameter, i.e., we write $\mathcal{F}_{\text{BD-SMT}}$ instead of $\mathcal{F}^{1}_{\text{BD-SMT}}$. Furthermore, we use $\mathcal{F}_{\text{BD-SEC}}$ and $\mathcal{F}_{\text{BD-AUTH}}$ to denote the corresponding authenticated and secure bounded-delay channel with $\delta = 1$, respectively.

Synchronous protocols in the $\{\mathcal{F}_{\text{CLOCK}},\mathcal{F}_{\text{BD-SMT}}\}$**-hybrid model.** We next show how a synchronous protocol can be described in the UC setting as an $\{\mathcal{F}_{\text{CLOCK}}, \mathcal{F}_{\text{BD-SMT}}\}$-hybrid protocol. The program of every player in such a protocol is described by a sequence of send, receive, and compute operations. Each such operation is annotated by the index of some round in which it should be executed. Without loss of generality, we assume that in each round r, first all parties receive their messages from round $r - 1$, then they compute their messages for round r, and finally they send these messages. The functionalities $\mathcal{F}_{\text{CLOCK}}$ and $\mathcal{F}_{\text{BD-SMT}}$ are used by each party p_i in a straightforward manner: On the onset of the protocol execution, each p_i sets its local round index to 1. Whenever p_i receives a message from some party other than $\mathcal{F}_{\text{CLOCK}}$ (i.e., from $\mathcal{Z}$, $\mathcal{A}$, or some other functionality) and is done with its actions for the current round, p_i sends a (`RequestRound`) message to $\mathcal{F}_{\text{CLOCK}}$, which answers back with an indicator bit d_i. Upon receiving this bit d_i, p_i uses it to detect whether or not every party is done with the current round and either proceeds to the next round or waits for further activations. More precisely, p_i executes the following steps: If $d_i = 1$ then ignore it and halt, else (i.e., if $d_i = 0$) do the following:

- If p_i has performed all its operations for the current round, then p_i informs the clock about this by sending it a (`RoundOK`) message.
- Otherwise, p_i executes the next operation for the current round.

Perfect vs. Imperfect Clock Synchronization. The above described functionality $\mathcal{F}_{\text{CLOCK}}$ models the following guarantee: Once a single party observes that a round is completed, every party will immediately be able to observe this event. As a "real world" assumption, this means that all parties perform the round switch at exactly the same time, which is only achievable if the parties' clocks are in *perfect* synchronization. One can relax the $\mathcal{F}_{\text{CLOCK}}$ functionality to model more realistic synchrony assumptions by incorporating "delays" as for the bounded-delay channel $\mathcal{F}_{\text{BD-SMT}}$. Such a "relaxed" version $\mathcal{F}^{-}_{\text{CLOCK}}$ of the clock is described in Appendix A.5.

We point out that if one uses the imperfect clock directly in the above described synchronous protocol compiler, the resulting protocol would not necessarily be secure. Indeed, depending on the delay, some party might start sending its round $r + 1$ messages before some other party has even received its round r messages. This can be fixed by introducing a "dummy" re-synchronization round between every two rounds (where every party sends a $\perp$ message to all parties).

4.2 Defining Guaranteed Termination

To show that the above described model allows to examine (synchronous) protocols with respect to both input completeness and guaranteed termination, we have to formalize what it means for a UC protocol to actually compute some specification with guaranteed termination. For simplicity, we restrict ourselves to non-reactive functionalities (secure function evaluation, or SFE). The case of reactive functionalities can be obtained as a straightforward extension and is explained in Appendix A.7.

Let $f : (\{0,1\}^*)^n \times R \longrightarrow (\{0,1\}^*)^n$ denote an n party randomized function, where the i-th component of f's input (resp. output) corresponds to the input (resp. output) of p_i, and the $(n+1)$-th input $r \in R$ corresponds to the randomness used by f. In order to define what it means for a UC protocol π among the parties in $\mathcal{P}$ to UC-securely evaluate f *with guaranteed termination* we use an approach inspired by constructive cryptography [MR11, Mau11]. The basic idea is to model the guarantee as a property of the ideal functionality that is to be achieved by the protocol, and to use the real world/ideal world paradigm to formalize the implied requirement for the protocol. Remember that the security definition states that, for any real-world adversary, there must exist an ideal world adversary, the simulator, such that whatever the adversary can do in the real-world the simulator can also do in the ideal world. Intuitively, the guaranteed termination property puts the following restriction on the influence of the (real-world) adversary on the protocol execution: the adversary cannot stall the computation indefinitely. This is modeled by requiring that the corresponding functionality does not allow the simulator to stop the (ideal) evaluation, i.e., in order to be terminating, the SFE functionality should not wait for the simulator *indefinitely*. In other words, the environment must be able to make the functionality proceed and (eventually) produce its outputs, independent of the simulator's strategy. Note that we do *not* require that, in the actual execution with an arbitrary environment $\mathcal{Z}$, the functionality will indeed provide output. The reason is that $\mathcal{Z}$ models the "world" in which the distinguishing experiment (between real and ideal world) takes place, and if $\mathcal{Z}$ aborts the execution at an early stage (and provides output), this corresponds to an abortion of the random experiment (and not of the protocol execution).[10] In contrast, we *do* require that the computation proceeds (and finishes) if $\mathcal{Z}$ makes queries to only the honest parties' interfaces.

In [Can05], the functionality $\mathcal{F}_{\text{SFE}}$ was introduced to capture secure evaluation of a multi-party function f. On a high level, $\mathcal{F}_{\text{SFE}}$ works as follows: it is parametrized by the function f; any honest party might either submit input to $\mathcal{F}_{\text{SFE}}^f$ or request output. Upon input x_i from some party p_i, $\mathcal{F}_{\text{SFE}}^f$ records x_i and notifies $\mathcal{A}$. When some party requests its output, $\mathcal{F}_{\text{SFE}}^f$ checks if all honest parties have submitted inputs; if so, $\mathcal{F}_{\text{SFE}}^f$ evaluates f on the received inputs (missing inputs of corrupted parties are replaced by default values), stops accepting further inputs, and outputs to p_i its output of the evaluation. A detailed description of $\mathcal{F}_{\text{SFE}}^f$ appears in Appendix A.6.

Although functionality $\mathcal{F}_{\text{SFE}}^f$ has the properties outlined in the discussion above, Lemma 4 shows that, for most choices of the function f, there exists no synchronous protocol implementing it from any reasonable network functionality. In order to specify "reasonable" we introduce the following terminology: we say that a network-functionality $\mathcal{F}_{\text{NET}}$ provides *separable rounds* if for any synchronous $\mathcal{F}_{\text{NET}}$-hybrid protocol which communicates exclusively through $\mathcal{F}_{\text{NET}}$, $\mathcal{F}_{\text{NET}}$ activates the adversary at least once in every round. The following lemma states that for any function f which requires more than one synchronous round in order to be evaluated, the corresponding SFE functionality $\mathcal{F}_{\text{SFE}}^f$ cannot be securely realized by any synchronous protocol in the $\mathcal{F}_{\text{NET}}$-hybrid model. Note that this includes many interesting functionalities such as broadcast, coin-tossing, etc.

Lemma 4. *For any function f and any network functionality $\mathcal{F}_{\text{NET}}$ with separable rounds, there exists an $\mathcal{F}_{\text{NET}}$-hybrid protocol π which securely realizes $\mathcal{F}_{\text{SFE}}^f$ only if π computes its output in one round.*

Proof (sketch). Assume, towards contradiction, that π is a two-round protocol securely computing $\mathcal{F}_{\text{SFE}}^f$, and consider the environment $\mathcal{Z}$ that provides input to all parties and, subsequently, directly requests the output from some honest party. Because, $\mathcal{F}_{\text{NET}}$ provides separable rounds, after all the inputs have been submitted, the adversary will be activated at least twice before the protocols can generate outputs. However this is not the case for the simulator in the ideal evaluation of $\mathcal{F}_{\text{SFE}}^f$. Hence, the dummy adversary cannot be simulated, which contradicts the security of π. $\square$

[10] In fact, any environment $\mathcal{Z}$ can be transformed into an environment $\mathcal{Z}'$ that completes the execution and achieves the same advantage as $\mathcal{Z}$.

In order to obtain an SFE functionality that matches the intuition of guaranteed termination, we need to circumvent the above impossibility. For this, we relax $\mathcal{F}_{\text{SFE}}$ as follows: we parametrize it with a function $Rnd(k)$ in the security parameter which corresponds to the minimum number of rounds required for evaluating f; the functionality $\mathcal{F}_{\text{SFE}}^{f,Rnd}$ is included in Appendix A.6. One can easily verify that for any (polynomial) round function $Rnd(\cdot)$ the functionality $\mathcal{F}_{\text{SFE}}^{f,Rnd}$ will eventually terminate (if there are sufficiently many queries at the honest parties' interfaces) independently of the simulator's strategy.

Definition 1 (Guaranteed Termination)**.** We say that a protocol UC-securely evaluates a function f *with guaranteed termination*, if it UC-realizes a functionality $\mathcal{F}_{\text{SFE}}^{f,Rnd}$ for some round function $Rnd(\cdot)$.

Remark (Lower Bounds). Note that the above formulation offers a language for making UC-statements about (lower bounds on) the round complexity of certain problems in the synchronous setting. In particular, the question whether $\mathcal{F}_{\text{SFE}}^{f,Rnd}$ can be realized by a synchronous protocol, corresponds to the question: "Does there exist a synchronous protocol π which securely evaluates f in $Rnd(k)$ rounds?", where k is the security parameter. As an example, the statement: "a function f needs at least r rounds to be evaluated" is translated to "there exists no synchronous protocol which UC securely realizes the functionality $\mathcal{F}_{\text{SFE}}^{f,Rnd(k)=r'}$, where $r' < r$."

The following theorem allows us to translate known results on feasibility of secure multi-party computation, e.g., [GMW87, BGW88, CCD88, RB89] into our synchronous universally composable setting with termination. The only modification is that the protocols start with a void synchronization round where no honest party sends or receives any message. For a synchronous protocol ρ, we denote by $\hat{\rho}$ the protocol which is obtained by extending ρ with such a start-synchronization round. The proof is based on ideas from [KLR06] and can be found in Appendix A.8.

Theorem 5. *Let f be a function and let ρ be a protocol that, according to the notion of [Can00], realizes f with computational (or statistical or perfect) security in the stand-alone model, with an efficient straight-line black-box simulator. Then $\hat{\rho}$ UC-realizes f with computational (or statistical or perfect) security and guaranteed termination in the $\{\mathcal{F}_{\text{CLOCK}}, \mathcal{F}_{\text{BD-SEC}}\}$-hybrid model with static adversaries.*

Using the above theorem, one can easily translate the statements of [GMW87, BGW88, CCD88] in our model, as the corresponding protocols are proved to be secure in the stand-alone model, with an efficient straight-line black-box simulator.

4.3 The Need for Both Synchronization and Bounded-Delay

In this section, we formalize the intuition that each one of the two "standard" synchrony assumptions, i.e., bounded-delay channels or synchronized clocks, is *alone* not sufficient for computation with guaranteed termination. In particular, we first show in Lemma 6 that bounded-delay channels (even with instant delivery) are, by themselves, not sufficient; subsequently, we show in Lemma 7 that (even perfectly) synchronized clocks are also not sufficient, even in combination with eventual delivery channels (with no known bound on the delay).

Lemma 6. *There are functions f such that for any (efficient) round-function Rnd and any $\delta > 0$: $\left[\mathcal{F}_{\text{SFE}}^{f,Rnd}\right]^{NT}$ can be realized in the $\mathcal{F}_{\text{BD-SMT}}^{\delta}$-hybrid model, but $\mathcal{F}_{\text{SFE}}^{f,Rnd}$ cannot, even when we assume bidirectional $\mathcal{F}_{\text{BD-SMT}}^{\delta}$-channels.*

Proof idea. Consider the two-party function f which, on input a bit $x_1 \in \{0,1\}$ from party p_1 (and nothing from p_2), outputs x_1 to p_2 (and nothing to p_1). The SFE functionality with guaranteed termination allows p_2 to distinguish between the cases (1) p_1 is honest but did not send something *yet* (because he was not yet activated by $\mathcal{Z}$), and (2) p_1 is corrupted and did not send. However, in the real world, the behavior of the bounded-delay channel in the above two cases is identical.

The functionality $\left[\mathcal{F}_{\text{SFE}}^{f,Rnd}\right]^{\text{NT}}$ can be implemented from $\mathcal{F}_{\text{BD-SMT}}$: p_1 simply has to send the input to p_2 via the $\mathcal{F}_{\text{BD-SMT}}$-channel. The simulator makes sure that the output in the ideal model is delivered to the p_2 only after $\mathcal{Z}$ acknowledges the delivery. A detailed proof can be found in Appendix A.9. □

In reality, synchronous clocks alone are not sufficient for synchronous computation if there is no *known* upper bound on the channels (even if they have guaranteed *eventual* delivery). This statement can be formalized in our model with the clock functionality $\mathcal{F}_{\text{CLOCK}}$ and the eventual-delivery channels $\mathcal{F}_{\text{ED-SMT}}$. The proof is similar to Lemma 3 and can be found in Appendix A.9.

Lemma 7. *There are functions f such that for any (efficient) round-function Rnd: $\left[\mathcal{F}_{\text{SFE}}^{f,Rnd}\right]^{NT}$ can be realized in the $\{\mathcal{F}_{\text{ED-SMT}}, \mathcal{F}_{\text{CLOCK}}\}$-hybrid model, but $\mathcal{F}_{\text{SFE}}^{f,Rnd}$ cannot, even when we assume bidirectional $\mathcal{F}_{\text{ED-SMT}}$-channels.*

4.4 Atomicity of Send/Receive Operations and Rushing

Hirt and Zikas [HZ10], pointed out that the standard formulation of a "rushing" adversary [Can00] in the synchronous setting puts a restriction on the order of the send/receive operations within one synchronous round. Demonstrating the flexibility of our framework, we show a simple way to pinpoint this inconsistency by showing that the rushing assumption corresponds to a functionality which cannot even be computed using $\mathcal{F}_{\text{CLOCK}}$ and $\mathcal{F}_{\text{BD-SMT}}$.

Intuitively, a rushing adversary is allowed to deliver a message to some party i *only* after some party j has finished sending *all* its messages. This is explicitly stated in [Can00], where the notion of "synchronous computation with *rushing*" is defined as follows:

"The computation proceeds in rounds; each round proceeds in mini-rounds, as follows. Each mini-round starts by allowing $\mathcal{A}$ to corrupt parties one by one in an adaptive way, as long as at most t parties are corrupted altogether. (The behavior of the system upon corruption of a party is described below.) Next $\mathcal{A}$ chooses an uncorrupted party, p, that was not yet activated in this round and activates it. Upon activation, p receives the messages sent to it in the previous round, generates its messages for this round*, and the next mini-round begins."*

In reality, it is arguable whether we can obtain the above guarantee by simply assuming bilateral bounded-delay channels and synchronized clocks. Indeed, when a computer sends multiple messages, this is typically not implemented as an atomic operation, but the messages are buffered on the network interface and sent one-by-one. Hence, to achieve the above guarantee, one would need to assume that the total time it takes for the sender to put all the messages on the network minus the *minimum* latency of the network is not sufficient for a party to become corrupted. In the following, we describe a UC-functionality which captures the above guarantee. We refer to this functionality as a *simultaneous multi-send* channel, and denote it as $\mathcal{F}_{\text{MS}}$.

On a high level, the functionality $\mathcal{F}_{\text{MS}}$ can be described as a channel that allows a sender p_i to send a vector of messages $(x_1, \ldots, x_n)$ to the respective receivers $p_1, \ldots, p_n$ as an atomic operation. In particular, $\mathcal{F}_{\text{MS}}$ can be derived from $\mathcal{F}_{\text{BD-SMT}}$ with the following modifications: first, instead of a single receiver, in $\mathcal{F}_{\text{MS}}$ there is a set $\mathcal{P}$ of receivers, and instead of a single message, the sender inputs a vector of $|\mathcal{P}|$ messages, one for each party in $\mathcal{P}$. As soon as one of the parties in $\mathcal{P}$ receives his message, the adversary can not replace any of the remaining messages that correspond to honest receivers, not even by corrupting the sender. The detailed description of the functionality is included in Appendix A.10. As with the case of the bounded-delay channels, we denote by $\mathcal{F}_{\text{MS-AUTH}}$ the multisend channel which leaks the transmitted vector to the adversary.

The following lemma states that the delayed relaxation of $\mathcal{F}_{\text{MS-AUTH}}$ cannot be implemented from $\mathcal{F}_{\text{BD-SEC}}$ and $\mathcal{F}_{\text{CLOCK}}$ when arbitrary many players can be corrupted. Note that this implies that $\mathcal{F}_{\text{MS-AUTH}}$ can also not be implemented from $\mathcal{F}_{\text{BD-SEC}}$ and $\mathcal{F}_{\text{CLOCK}}$.

Lemma 8. *Let $\mathcal{P}$ be a player set with $|\mathcal{P}| > 3$ and let $p \in \mathcal{P}$. Then there exists no protocol which UC-securely realizes $[\mathcal{F}_{\text{MS-AUTH}}(p, \mathcal{P})]^{NT}$ in the $\{\mathcal{F}_{\text{CLOCK}}, \mathcal{F}_{\text{BD-AUTH}}\}$-hybrid model and tolerates a corrupted majority.*

Proof (sketch). Garay et al. [GKKZ11] showed that if atomic multi-send (along with a setup for digital signatures) is assumed, then the broadcast protocol from Dolev and Strong [DS82] UC-securely realizes broadcast (without guaranteed termination) in the presence of an adaptive adversary who corrupts any number of parties. Hence, if there exist a protocol for realizing $[\mathcal{F}_{\text{MS-AUTH}}(p, \mathcal{P})]^{\text{NT}}$ in the synchronous model, i.e., in the $\{\mathcal{F}_{\text{CLOCK}}, \mathcal{F}_{\text{BD-AUTH}}\}$-hybrid world, with corrupted majority and adaptive adversary, then one could also realize broadcast in this model, contradicting the impossibility result from [HZ10]. $\square$

The above lemma implies that the traditional notion of "synchronous computation with *rushing*" cannot be, in general, achieved in the UC model unless some non-trivial property is assumed on the communication channel. Yet, $[\mathcal{F}_{\text{MS-AUTH}}]^{\text{NT}}$ *can* be UC-realized from $\{[\mathcal{F}_{\text{BD-AUTH}}]^{\text{NT}}, [\mathcal{F}_{\text{COM}}]^{\text{NT}}\}$, where $\mathcal{F}_{\text{COM}}$ denotes the standard UC-commitment functionality [CF01]. The idea is the following: in order to simultaneously multi-send a vector $(x_1, \ldots, x_n)$ to the players $p_1, \ldots, p_n$, the sender, sends an independent commitment on x_i to every recipient p_i, who acknowledges the receipt (using the channel $[\mathcal{F}_{\text{BD-SEC}}]^{\text{NT}}$). After receiving all such acknowledgments, the sender, in a second round, opens all commitments. The hiding property of $\mathcal{F}_{\text{COM}}$ ensures that the adversary, unless he corrupts the sender in the first round, learns the committed messages x_i only after every player has received his commitment; but, from that point on, $\mathcal{A}$ can no longer change the sent message, due to the binding property. For completeness we state the above in the following lemma. The proof is as described above and can be found in Appendix A.10.

Lemma 9. *For each player set $\mathcal{P}$ and $p \in \mathcal{P}$, there exists a synchronous $\{[\mathcal{F}_{\text{BD-AUTH}}]^{NT}, [\mathcal{F}_{\text{COM}}]^{NT}\}$-hybrid protocol which UC securely realizes $[\mathcal{F}_{\text{MS-AUTH}}(p, \mathcal{P})]^{NT}$.*

Combining the above Lemmas 8 and 9, and using the trivial fact that the delayed relaxation of any functionality $\mathcal{F}$ can be implemented in the $\mathcal{F}$-hybrid model, we can extend the result of Canetti and Fischlin [CF01] on impossibility of UC commitments to our synchronous setting.

Corollary 10. *There exists no protocol which UC securely realizes the commitment functionality $\mathcal{F}_{\text{COM}}$ in the $\{\mathcal{F}_{\text{BD-AUTH}}, \mathcal{F}_{\text{CLOCK}}\}$-hybrid model.*

5 Existing Synchronous Models as Special Cases

In this section, we revisit existing models for synchronous computation. We show that either they are sufficient only for non-terminating computation (which can be achieved even in a fully asynchronous environment), or they are somewhat *ad hoc* and cannot easily be adapted to model, e.g., synchrony in an incomplete network, or the case when synchrony holds only in part of a network.

5.1 The $\mathcal{F}_{\text{SYN}}$-Hybrid Model

In the latest revision of the UC-framework [Can05], synchronous computation is captured by defining a *synchronous network* functionality $\mathcal{F}_{\text{SYN}}$. On a high-level, $\mathcal{F}_{\text{SYN}}$ corresponds to an authenticated network with storage, which proceeds in a round-based fashion; in each round r, every party which is associated to $\mathcal{F}_{\text{SYN}}$ sends a vector of messages to $\mathcal{F}_{\text{SYN}}$, where it is guaranteed that (1) the adversary cannot change the message sent by an honest party without corrupting this party, and (2) the round index is only increased after every honest party as well as the adversary have submitted their messages for that round. Furthermore, $\mathcal{F}_{\text{SYN}}$ allows the parties to query the current value of r along with the messages of that round r. A detailed description of $\mathcal{F}_{\text{SYN}}$ is included in Appendix B.

By inspection of $\mathcal{F}_{\text{SYN}}$, one can realize that it allows the adversary to stall the protocol execution (by not sending the `Advance-Round`-message). Hence, $\mathcal{F}_{\text{SYN}}$ cannot be used for evaluating a function f with guaranteed termination, unless f can be computed locally (cf. Sections 3.2 and 3.3). This is stated in the following lemma. The idea of the proof is the following: because we require termination, for every protocol which securely realizes $\mathcal{F}_{\text{SFE}}$ and for every adversary, the environment $\mathcal{Z}$ which gives inputs to all honest parties and issues sufficiently many `fetch` requests should be able to make π generate its output from the evaluation. This must in particular hold when the adversary never sends an `Advance-Round`-message to $\mathcal{F}_{\text{SYN}}$, which leads to contradiction. A more detailed proof can be found in Appendix B.

Lemma 11. *For every non-trivial n-party function f and for every round function Rnd there exists no $\mathcal{F}_{\text{SYN}}$-hybrid protocol which securely realizes $\mathcal{F}_{\text{SFE}}^{f,Rnd}$ with guaranteed termination.*

$\mathcal{F}_{\text{SYN}}$ defines a point in the computation (chosen by $\mathcal{A}$) in which the round index advances for every player simultaneously, and the following lemma demonstrates that the existence of such a point is essentially the only advantage that $\mathcal{F}_{\text{SYN}}$ offers on top of what can be achieved from bounded-delay channels. More precisely, denote by $\mathcal{F}_{\text{SYN}}^{-}$ the functionality that behaves as $\mathcal{F}_{\text{SYN}}$, except for the (`Advance-Round`)-message: $\mathcal{F}_{\text{SYN}}^{-}$ advances the round, but, for each player, allows the adversary to specify a delay $t_d^{(i)}$ (polynomial in the security parameter) such that the first $t_d^{(i)}$ (receive)-requests of p_i are answered with the previous round messages (which is the same result as if the adversary does not send (Advance-Round) to $\mathcal{F}_{\text{SYN}}$). Arguably, this is only a mild relaxation of the functionality: The only slackness introduced is that the adversary can delay the message delivery, but this inconsistency can only be detected by the parties if they have access to a channel which is guaranteed to deliver faster than the adversary's delay. The functionality $\mathcal{F}_{\text{SYN}}^{-}$ can be implemented by a protocol over asynchronous channels with eventual delivery. The protocol follows the ideas of [Awe85, KLP05]: In each round r, each party p_i sends one message to the parties p_j with $j \in \mathcal{P} \setminus \{i\}$. After receiving messages from all p_j, p_i sends an acknowledgment to all p_j. Once all such acknowledgments have been received, p_i prepares the messages received in that round for local output upon request and starts the next round (as soon as messages have been provided as local input). The implemented functionality $\mathcal{F}_{\text{SYN}}^{-}$ has a delay of $2n$ for each round. This proves the following lemma:

Lemma 12. *There exist a protocol that UC-realizes the functionality $\mathcal{F}_{\text{SYN}}^{-}$ in the $\mathcal{F}_{\text{BD-SMT}}$-hybrid model.*

5.2 The Timing Model

In this section, we relate the assumptions from the "Timing Model" [Gol02, KLP05] to our model. As we show, similarly to the $\mathcal{F}_{\text{SYN}}$-hybrid model, the assumptions also do not capture the essence of guaranteed termination, i.e, the adversary is able to stall the computation. In particular, we show that protocols in the Timing Model can be formulated as $\mathcal{F}_{\text{AUTH}}$-hybrid protocols (in UC) such that security statements are preserved. However, in contrast to $\mathcal{F}_{\text{SYN}}$, which "lives" in the UC framework, security statements in the Timing Model cannot be automatically transferred to the UC setting. Indeed, there is a type-mismatch between functionalities/protocols in the two frameworks. We resolve this by introducing a way of translating statements in the Timing Model into statements in UC so that security properties are preserved. This technique is inspired by [Bac03, MPR10, MR11] and is of independent interest.

Roughly speaking, the "Timing" assumption is implemented in [KLP05] by extending the model of computation, so that each party, in addition to its communication and computation tapes, also has a *clock tape* that is writable for the adversary in a monotone and "bounded-drift"-preserving manner: The adversary can only increase the value of the clocks, and, for any two parties, the distance ϵ of their clocks' speed (drift) at any point in time is bounded by a known constant. The value of each party's clock-tape defines the local time of this party. Depending on their local time, protocols delay sending messages or time-out if a message has not arrived as expected. We formalize this assumption in UC in the following straightforward manner: We introduce a functionality $\mathcal{F}_{\text{TIME}}$ (see Appendix C) that maintains a clock value for each party and allows the adversary to advance this clock, in the same monotone and "bounded-drift"-preserving way. Instead of reading the local clock tape, $\mathcal{F}_{\text{TIME}}$-hybrid protocols obtain their local time value by querying $\mathcal{F}_{\text{TIME}}$. The following lemma states that $\mathcal{F}_{\text{TIME}}$ can be implemented from asynchronous authenticated communication. The idea of the proof is the following: The protocol τ that implements the functionality $\mathcal{F}_{\text{TIME}}$ from pairwise authenticated channels maintains, for each party p_i, a local integer variable t_i that corresponds to p_i's local time. Using the authenticated network, the parties ensure that the local time values increase with bounded drift. The protocol and the security proof are given in Appendix C.

Lemmas 13 and 14 demonstrate that timed protocols cannot UC-realize more functionalities than non-timed protocols, which is consistent with [KLP05, Theorem 2]. In contrast, timed protocols *can* securely realize arbitrary functionalities under composition with δ-delayed protocols [KLP05, Theorem 1].

Lemma 13. *Let $\mathcal{P}$ be a player set and $\epsilon \geq 1$. The functionality $\mathcal{F}_{\text{TIME}}(\mathcal{P}, \epsilon)$ can be UC-realized from pairwise authenticated channels $\mathcal{F}_{\text{AUTH}}(i, j)$ for all pairs of $i, j \in \mathcal{P}$, $i \neq j$.*

We use the following idea for resolving the type mismatch between security statements in UC and in the Timing Model: A functionality $\mathcal{F}$ in the timing model is compiled to a functionality in UC which behaves exactly as $\mathcal{F}$ but ensures that the interfaces are compatible with the UC model of computation. On a high-level, the functionality compiler $T(\cdot)$ works as follows: $T(\mathcal{F})$ behaves as $\mathcal{F}$, but for any input it receives from an honest party, it informs the adversary about this event (without leaking information on the contents). Whenever $\mathcal{F}$ outputs a value y to some party, $T(\mathcal{F})$ issues a (private) delayed output y instead.[11] A detailed description of the functionality compiler $T(\cdot)$ appears in Appendix C. The following Lemma 14 states that any security statement about a functionality $\mathcal{F}$ in the Timing Model can be translated into a statement about $T(\mathcal{F})$ in the $\{\mathcal{F}_{\text{TIME}}, \mathcal{F}_{\text{AUTH}}\}$-hybrid model (in UC). In fact, the translation is both constructive and uniform, i.e., we describe a protocol compiler $C(\cdot)$ that translates a protocol in the Timing Model into a corresponding protocol in the $\{\mathcal{F}_{\text{TIME}}, \mathcal{F}_{\text{AUTH}}\}$-hybrid model. The compiler $C(\cdot)$ and the proof of the lemma can also be found in Appendix C.

Lemma 14. *For an arbitrary functionality $\mathcal{F}$ and a protocol π in the Timing Model, π securely realizes $\mathcal{F}$ (in the timing model) if and only if the compiled protocol $C(\pi)$ UC realizes $T(\mathcal{F})$ in the $\{\mathcal{F}_{\text{TIME}}, \mathcal{F}_{\text{AUTH}}\}$-hybrid model in the presence of a static adversary.*

5.3 Models with Guaranteed Termination

Nielsen's Framework [Nie03]. The framework described in [Nie03] is an adaptation of the asynchronous framework of [Can01] to authenticated synchronous networks. While the general structure of the security definition is adopted, the definition of protocols and their executions differs considerably. For instance, the "subroutine" composition of two protocols is defined in a "lock-step" way: the round switches occur at the same time. Similarly to our bounded-delay channels, messages in transfer can be replaced if the sender becomes corrupted. Lemma 15 states that we can translate, along the lines of Section 5.2, any security statement in the model of [Nie03] into a security statement about a synchronous protocol in the $\{\mathcal{F}_{\text{CLOCK}}, \mathcal{F}_{\text{BD-AUTH}}\}$-hybrid model. As in the previous section, the translation is done by a functionality compiler $T(\cdot)$ that resolves the type mismatch between the functionalities in UC and in [Nie03], and a corresponding protocol compiler $C(\cdot)$. The descriptions of these compilers along with the proof of Lemma 15 are provided in Section D of the appendix, where we also briefly describe the model of [Nie03] and point out a small flaw in the specification. We emphasize that the converse statement of Lemma 15 does *not* hold, i.e., there are UC statements about synchronous protocols that cannot be modeled in the [Nie03] framework. For instance, our model of UC synchronous computation allows protocols to use further functionalities that run mutually asynchronously with the synchronous network, which cannot be modeled in [Nie03].

Lemma 15. *For an arbitrary functionality $\mathcal{F}$ and a protocol π in [Nie03], π securely realizes $\mathcal{F}$ (in [Nie03]) if and only if the compiled protocol $C(\pi)$ UC realizes $T(\mathcal{F})$ in the $\{\mathcal{F}_{\text{CLOCK}}, \mathcal{F}_{\text{BD-AUTH}}\}$-hybrid model.*

Hofheinz and Müller-Quade's Framework [HM04]. The framework of [HM04] also models authenticated synchronous networks based on the framework of [Can01], but the rules of the protocol execution are differ considerably: The computation proceeds in rounds, and each round is split into three phases. In each phase, only a subset of the involved ITMs are activated according to a specific scheme. The adversary has a relaxed *rushing* property: while being the last to specify the messages for a round, he cannot corrupt parties within a round. This corresponds to a network with guarantees that are stronger than simultaneous multi-send: Once the first message of an honest party is provided to the adversary, all messages of honest parties are guaranteed to be delivered correctly.[12] We model this strong guarantee in UC by the functionality $\mathcal{F}_{\text{NET+}}$, which is a modified version of $\mathcal{F}_{\text{MS}}$, in Appendix E. As before, we translate the security statements of [HM04] to our model (where $\mathcal{F}_{\text{NET+}}$ is used instead of $\mathcal{F}_{\text{AUTH}}$) through a pair of compilers $(T(\cdot), C(\cdot))$, which take care of the syntactic mismatch.

Lemma 16. *For an arbitrary functionality $\mathcal{F}$ and a protocol π in [HM04], π securely realizes $\mathcal{F}$ (in [HM04]) if and only if the compiled protocol $C(\pi)$ UC realizes $T(\mathcal{F})$ in the $\{\mathcal{F}_{\text{CLOCK}}, \mathcal{F}_{\text{NET+}}\}$-hybrid model.*

6 Conclusions

We provided a universally composable framework for synchronous secure computation. The standard assumptions of a synchronous network are that there is a known bound on the latency of the network and that the players have (partially) synchronized clocks. We described functionalities modeling each of these two assumptions in the UC framework of Canetti [Can01], and introduced a "clean slate" definition of secure computation with guaranteed

[11] Note that the delay is inherent in the functionalities, as they send their outputs through asynchronous secure channels.

[12] This is an advantage in the context of [HM04], as the goal is an impossibility result.

termination within the UC framework. Furthermore, we showed how results from the synchronous protocols literature can be expressed within this framework while preserving all their security properties, including guaranteed termination and composition with arbitrary protocols (synchronous or not). Finally, we revisited existing security frameworks for synchronous secure computation and showed how to translate them in our model. Interestingly, we observed that some of them, including the $\mathcal{F}_{\mathrm{SYN}}$-hybrid model [Can01]—which, was specifically introduced as a way of modeling the guarantees provided by fully synchronous networks—are insufficient for capturing guaranteed termination.

References

[AL11] Gilad Asharov and Yehuda Lindell. A full proof of the BGW protocol for perfectly-secure multiparty computation. In *Advances in Cryptology — CRYPTO 2011*, 2011.

[Awe85] Baruch Awerbuch. Complexity of network synchronization. *Journal of the ACM*, 32(4):804–823, 1985.

[Bac03] Michael Backes. Unifying simulatability definitions in cryptographic systems under different timing assumptions. In Roberto Amadio and Denis Lugiez, editors, *Proceedings of CONCUR 2003*, volume 2761 of *Lecture Notes in Computer Science*, pages 350–365. Springer, 2003.

[BCG93] Michael Ben-Or, Ran Canetti, and Oded Goldreich. Asynchronous secure computation. In *Proceedings of the twenty-fifth annual ACM Symposium on Theory of Computing*, pages 52–61. ACM, 1993.

[BGW88] Michael Ben-Or, Shafi Goldwasser, and Avi Widgerson. Completeness theorems for non-cryptographic fault-tolerant distributed computation. In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing*, pages 1–10. ACM, 1988.

[BPW07] Michael Backes, Birgit Pfitzmann, and Michael Waidner. The reactive simulatability (RSIM) framework for asynchronous systems. *Information and Computation*, 205(12):1685–1720, December 2007.

[Can96] Ran Canetti. *Studies in Secure Multiparty Computation and Applications*. PhD thesis, The Weizmann Institute of Science, 1996.

[Can00] Ran Canetti. Security and composition of multiparty cryptographic protocols. *Journal of Cryptology*, 13:143–202, April 2000.

[Can01] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science*, pages 136–145. IEEE, 2001. Extended version in [Can05].

[Can05] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. Cryptology ePrint Archive, Report 2000/067, December 2005.

[CCD88] David Chaum, Claude Crépeau, and Ivan Damgård. Multiparty unconditionally secure protocols. In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing*, pages 11–19. ACM, 1988.

[CF01] Ran Canetti and Marc Fischlin. Universally composable commitments. In Joe Kilian, editor, *Advances in Cryptology — CRYPTO 2001*, volume 2139 of *Lecture Notes in Computer Science*, pages 19–40. IACR, Springer-Verlag, 2001.

[CK02] Ran Canetti and Hugo Krawczyk. Universally composable notions of key exchange and secure channels. In Lars R. Knudsen, editor, *Advances in Cryptology — EUROCRYPT 2002*, volume 3027 of *Lecture Notes in Computer Science*, pages 337–351. IACR, Springer-Verlag, 2002.

[CLOS02] Ran Canetti, Yehuda Lindell, Rafail Ostrovsky, and Amit Sahai. Universally composable two-party and multi-party secure computation. In *Proceedings of the 34th ACM Symposium on Theory of Computing*. ACM, 2002.

[CM89] Benny Chor and Lior Moscovici. Solvability in asynchronous environments. In *30th Annual Symposium on Foundations of Computer Science*, pages 422–427. IEEE, 1989.

[DM00] Yevgeniy Dodis and Silvio Micali. Parallel reducibility for information-theoretically secure computation. In Mihir Bellare, editor, *Advances in Cryptology — CRYPTO 2000*, volume 1880 of *Lecture Notes in Computer Science*, pages 74–92. IACR, Springer-Verlag, 2000.

[DNS98] Cyntha Dwork, Moni Naor, and Amit Sahai. Concurrent zero-knowledge. In *Proceedings of the 30th annual ACM Symposium on Theory of Computing*. ACM, 1998.

[DS82] Danny Dolev and H. R. Strong. Polynomial algorithms for multiple processor agreement. In *Proceedings of the 14th Annual ACM Symposium on Theory of Computing*, pages 401–407. ACM, 1982.

[GKKZ11] Juan A. Garay, Jonathan Katz, Ranjit Kumersan, and Hong-Sheng Zhou. Adaptively secure broadcast, revisited. In *Proceedings of the 30th Annual ACM Symposium on Principles of Distributed Computing*. ACM, 2011.

[GMW87] Oded Goldreich, Silvio Micali, and Avi Widgerson. How to play any mental game—A completeness theorem for protocols with honest majority. In *Proceedings of the 19th Annual ACM Symposium on Theory of Computing*, pages 218–229. ACM, 1987.

[Gol02] Oded Goldreich. Concurrent zero-knowledge with timing, revisited. In *Proceedings of the 34th annual ACM Symposium on Theory of Computing*, pages 332–340. ACM, 2002.

[HM04] Dennis Hofheinz and Jörn Müller-Quade. A synchronous model for multi-party computation and the incompleteness of oblivious transfer. In *Proceedings of FCS*, 2004.

[HZ10] Martin Hirt and Vassilis Zikas. Adaptively secure broadcast. In Henri Gilbert, editor, *Advances in Cryptology — EUROCRYPT 2010*, volume 6110 of *Lecture Notes in Computer Science*, pages 466–485. IACR, Springer-Verlag, 2010.

[KLP05] Yael Taumann Kalai, Yehuda Lindell, and Manoj Prabhakaran. Concurrent general composition of secure protocols in the timing model. In *Proceedings of the thirty-seventh annual ACM Symposium on Theory of Computing*, pages 644–653. ACM, 2005.

[KLR06] Eyal Kushilevitz, Yehuda Lindell, and Tal Rabin. Information-theoretically secure protocols and security under composition. In *Proceedings of the thirty-eighth annual ACM Symposium on Theory of Computing*, pages 109–118. ACM, 2006.

[KMR09] Robin Künzler, Jörn Müller-Quade, and Dominik Raub. Secure computability of functions in the IT setting with dishonest majority and applications to long-term security. In Omer Reingold, editor, *Theory of Cryptography*, volume 5444 of *Lecture Notes in Computer Science*. IACR, Springer, 2009.

[Mau02] Ueli Maurer. Indistinguishability of random systems. In Lars R. Knudsen, editor, *Advances in Cryptology — EUROCRYPT 2002*, volume 2332 of *Lecture Notes in Computer Science*, pages 110–132. IACR, Springer-Verlag, 2002.

[Mau11] Ueli Maurer. Constructive cryptography: A new paradigm for security definitions and proofs. In Sebastian Mödersheim and Catuscia Palamidessi, editors, *TOSCA 2011—Theory of Security and Applications*, Lecture Notes in Computer Science. Springer-Verlag, 2011. Available online: `http://www.crypto.ethz.ch/research/ac/TOSCA11.pdf`.

[MPR10] Hemanta Maji, Manoj Prabhakaran, and Mike Rosulek. Cryptographic complexity classes and computational intractability assumptions. In *Innovations in Computer Science*. Tsinghua University Press, 2010.

[MR11] Ueli Maurer and Renato Renner. Abstract cryptography. In *Innovations in Computer Science*. Tsinghua University Press, 2011.

[Nie03] Jesper Buus Nielsen. *On Protocol Security in the Cryptographic Model*. PhD thesis, University of Aarhus, June 2003.

[RB89] Tal Rabin and Michael Ben-Or. Verifiable secret sharing and multiparty protocols with honest majority. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing*, pages 73–85. ACM, 1989.

A Supplementary Material

This section contains supplementary material that has been deferred from the main paper due to lack of space, but which is useful to consult while reading the paper. Some of the functionalities presented here are material imported from other papers and restated using our conventions, and some of the functionalities are only described informally in the body of the paper.

A.1 Conventions Concerning the UC Framework

In the UC framework [Can01], the potential dishonesty of players is modeled by the fact that the adversary $\mathcal{A}$ may corrupt protocol machines (adaptively) during the execution by sending them a special $(\texttt{corrupt})$ message. We only consider Byzantine party-corruption, which means that the party sends its entire current local state to $\mathcal{A}$ and, in all future activations, follows $\mathcal{A}$'s instructions. We generally assume that parties do not delete state, and hence the entire execution history can be derived from the party's current state. For more details on party corruption, see [Can05, Section 4.1].

Our functionalities $\mathcal{F}$ use *standard corruption* as described in [Can05]: At any point in time, $\mathcal{F}$ will accept messages of the type $(\texttt{corrupt}, i)$ for $i \in \mathcal{P}$ from the adversary $\mathcal{A}$, in which case $\mathcal{F}$ marks the player p_i as corrupted and replies with $(\texttt{ok})$. Whenever a corrupted party p_i inputs a message x, the functionality $\mathcal{F}$ sends a message $(\texttt{input}, i, x)$ to $\mathcal{A}$ and accepts as a reply $(\texttt{replace}, \tilde{x})$, before continuing the execution with input $\tilde{x}$. Whenever $\mathcal{F}$ would send an output message y to a corrupted party p_i, it sends a message $(\texttt{output}, i, y)$ to $\mathcal{A}$ and accepts as a reply $(\texttt{replace}, \tilde{y})$. $\mathcal{F}$ then outputs $\tilde{y}$ to p_i. We use the symbol $\mathcal{H}$ to denote (at the current point of the execution) the subset of $\mathcal{P}$ for which no message $(\texttt{corrupt}, i)$ has been sent.[13]

A.2 The Asynchronous Communication Channel

In [Can05], Canetti introduces the generic communication functionality $\mathcal{F}_{\text{SMT}}$ that can be parametrized by a leakage function $\ell : \{0,1\}^* \to \{0,1\}^*$. This functionality can be used to model, e.g., authenticated or secure communication channels by instantiating the function ℓ appropriately. Below, the functionality is restated using our conventions.

Functionality $\mathcal{F}_{\text{SMT}}(\ell, p_s, p_r)$ as in [Can05]

The functionality $\mathcal{F}^{\ell}_{\text{SMT}}$ is parametrized by a leakage function $\ell : \{0,1\}^* \to \{0,1\}^*$.

- Upon receiving $(\texttt{send}, m)$ from p_s, send $(\texttt{sent}, \ell(m))$ to the adversary and provide private delayed output $(\texttt{sent}, m)$ to p_r.

Lemma 1. *There are functionalities $\mathcal{F}$ such that $[\mathcal{F}]^{NT}$ can be implemented in the $\mathcal{F}_{\text{SMT}}$-hybrid model, but $\mathcal{F}$ cannot be implemented, even when we assume bidirectional $\mathcal{F}_{\text{SMT}}$-channels.*

Proof. We describe a functionality $\mathcal{F}$ that satisfies the conditions of the lemma: $\mathcal{F}$ is a two party functionality that corresponds to a channel between a sender p_s and a receiver p_r. $\mathcal{F}$ behaves as $\mathcal{F}_{\text{SMT}}$, but does not necessarily wait for the adversary's approval for delivering a message. Instead, upon receiving a special $(\texttt{fetch})$-message from the receiver p_r, $\mathcal{F}$ outputs y to p_r, where $y = m$ if the sender has input the message m, and $y = \perp$ (i.e., a default value), otherwise. The functionality $\mathcal{F}$ is described in detail in the following:

Functionality $\mathcal{F}(\ell, p_s, p_r)$

The functionality $\mathcal{F}$ is parametrized by a leakage function $\ell : \{0,1\}^* \to \{0,1\}^*$.

- Upon receiving $(\texttt{send}, m)$ from p_s, send $(\texttt{sent}, \ell(m))$ to the adversary and provide a private delayed output $(\texttt{sent}, m)$ to p_r.
- Upon receiving $(\texttt{fetch})$ from p_r: if a message m was received from p_s and was not yet delivered to p_r, then send m to p_r, otherwise send $\perp$ to p_r.

[13] Note that the corruption in [Can05] starts by $\mathcal{Z}$ issuing such a message to $\mathcal{A}$. From this point in time, $\mathcal{A}$ is allowed by the control function to send corruption requests $(\texttt{corrupt}, i)$ to the ideal functionalities or $(\texttt{corrupt})$ to the protocol machine π_i.

We show that $\mathcal{F}$ cannot be realized from $\mathcal{F}_{\text{SMT}}$. Indeed, consider the following environment $\mathcal{Z}$ with the dummy adversary (which acts as a forwarder between the environment and the functionality): both p_s and p_r are honest, and $\mathcal{Z}$ gives a uniformly chosen input-message $m \in_R \{0,1\}$ to p_s; as soon as $\mathcal{Z}$ receives from the adversary the leaked value $\ell(m)$ from the channel, it activates p_r for fetching a message from the channel (formally, $\mathcal{Z}$ provides input $(\texttt{fetch})$ to p_r). The (honest) party p_r is expected to answer with a message m' as received from the channel.[14] Note that any action other than reporting a message, i.e., producing no output or sending a message to p_s (through the $\mathcal{F}_{\text{SMT}}$-channel in the other direction) is detected by $\mathcal{Z}$, as in both cases $\mathcal{Z}$ is activated before receiving output from p_r, and can be used for distinguishing. Using the information whether or not m' equals m, the environment can decide whether it witnessed an ideal-world or a real-world execution. In the ideal-world case (where the parties interact via $\mathcal{F}$) p_r's output will always be $m' = m$. In the real world, as the adversary has not instructed the channel $\mathcal{F}_{\text{SMT}}$ to deliver the message, the (honest) p_r has no information on m. Because m is chosen uniformly from $\{0,1\}$, the probability that $m' = m$ is $1/2$, which provides a noticeable distinguishing advantage to $\mathcal{Z}$.

For the reduction of $[\mathcal{F}]^{\text{NT}}$ to $\mathcal{F}_{\text{SMT}}$, we observe that $[\mathcal{F}]^{\text{NT}}$ can be realized from $\mathcal{F}_{\text{SMT}}$ using the dummy protocol φ. The simulator $\mathcal{S}$ uses the adversary $\mathcal{A}$ attacking the execution of φ in a black-box manner and behaves as follows: Throughout the simulation, $\mathcal{S}$ forwards all messages sent between $\mathcal{A}$ and $\mathcal{Z}$; whenever $\mathcal{A}$ requests to corrupt a party, $\mathcal{S}$ requests to corrupt this party in the ideal world (simulating the internal state is easy); furthermore, the simulator forwards to $\mathcal{A}$ all the messages that are sent from $[\mathcal{F}]^{\text{NT}}$, except for the notification about the $(\texttt{fetch})$-message sent (to $[\mathcal{F}]^{\text{NT}}$) by p_r. ($\mathcal{S}$ is informed about each $(\texttt{fetch})$-message, because $[\mathcal{F}]^{\text{NT}}$ issues a delayed output as a response to each such message). It is straightforward to verify that the above $\mathcal{S}$ is a good simulator for the adversary $\mathcal{A}$. ☐

A.3 Eventual Delivery Channels

Lemma 2. *There are functionalities $\mathcal{F}$ such that $[\mathcal{F}]^{NT}$ can be implemented in the $\mathcal{F}_{\text{ED-SMT}}$-hybrid model, but $\mathcal{F}$ cannot be implemented, even when we assume bidirectional $\mathcal{F}_{\text{ED-SMT}}$-channels.*

Proof. For the functionality $\mathcal{F}$ introduced in the proof of Lemma 1 (see Page 17), we show that $\mathcal{F}$ cannot be realized from $\mathcal{F}_{\text{ED-SMT}}$, but $[\mathcal{F}]^{\text{NT}}$ can. Indeed, the dummy protocol φ implements $[\mathcal{F}]^{\text{NT}}$ based on $\mathcal{F}_{\text{ED-SMT}}$: The simulator $\mathcal{S}$ uses the hybrid model adversary $\mathcal{A}$ in a black-box manner and behaves as follows. Throughout the simulation, $\mathcal{S}$ forwards all the messages sent between $\mathcal{A}$ and $\mathcal{Z}$; furthermore, whenever $\mathcal{A}$ requests to corrupt a party, $\mathcal{S}$ requests to corrupt this party in the ideal world (the internal state of the protocol machines can be simulated easily). From the messages $\mathcal{A}$ sends to $\mathcal{F}_{\text{ED-SMT}}$, $\mathcal{S}$ can learn the actual delay T that the $\mathcal{A}$ wishes to put on the message transmission. In order to apply the same delay on the $[\mathcal{F}]^{\text{NT}}$-hybrid world, $\mathcal{S}$ stalls the output of any message sent through $[\mathcal{F}]^{\text{NT}}$ (recall that $[\mathcal{F}]^{\text{NT}}$ allows the simulator to deliver output-messages if and when he wishes, as they are issued in a delayed manner) until T $(\texttt{fetch})$-messages have been sent from p_r to $[\mathcal{F}]^{\text{NT}}$ ($\mathcal{S}$ is informed about each $(\texttt{fetch})$-message, because $[\mathcal{F}]^{\text{NT}}$ issues a delayed output as a response to each such message). It is straightforward to verify that the above $\mathcal{S}$ is a good simulator for the adversary $\mathcal{A}$.

We next turn into proving the impossibly of implementing $\mathcal{F}$ from $\mathcal{F}_{\text{ED-SMT}}$. The idea is similar to the impossibility proof of Lemma 1: In particular, consider the dummy adversary $\mathcal{A}$ and the environment $\mathcal{Z}$ that corrupts no party and has the following distinguishing strategy: $\mathcal{Z}$ provides as input to p_s a uniformly random bit m. As soon as it receives the leaked value $\ell(m)$ from the functionality (forwarded by $\mathcal{A}$), $\mathcal{Z}$ instructs $\mathcal{A}$ to give delay $T = 2$ to $\mathcal{F}_{\text{ED-SMT}}$. Subsequently, $\mathcal{Z}$ activates p_r with input $(\texttt{fetch})$; as a result, if $\mathcal{Z}$ is witnessing the ideal-world execution, he will receive the message m from p_r (note that $\mathcal{F}$ does not allow the simulator to stall this reply); hence, in the real world, the protocol of p_r has to also output some m' to $\mathcal{Z}$. Indeed, the other choices of p_r are to not produce any output, or to send $(\texttt{fetch})$ to $\mathcal{F}_{\text{ED-SMT}}$, or to invoke the $\mathcal{F}_{\text{ED-SMT}}$ in the other direction (i.e., as a sender); in all three cases the environment is activated before receiving output which allows it to detect that it is witnessing the real-world execution. However, as p_r has no information on m, the probability that $m' = m$ is $1/2$, which allows $\mathcal{Z}$ to distinguish with noticeable advantage. ☐

[14]Recall that in UC a party cannot become corrupted without the environments permission, hence the simulator cannot stop p_r from reporting the messages it receives in the ideal world.

A.4 Bounded Delay Channels with a Known Upper-Bound

The bounded-delay channel functionality $\mathcal{F}^{\delta}_{\text{BD-SMT}}$ formalizes a channel for which the receiver knows the maximum delay of the message. This is modeled by the query $(\texttt{LearnBound})$, upon which the channel tells the maximum number of activations needed to obtain the message, once it has been submitted by the sender.

Functionality $\mathcal{F}^{\delta}_{\text{BD-SMT}}(p_s, p_r, \ell(\cdot))$

Initialize $M := \perp$ and $D := 0$.

- Upon receiving a message m from p_s, set $D := 1$ and $M := m$ and send $\ell(M)$ to the adversary.
- Upon receiving a message $(\texttt{LearnBound})$ from any party $p_i \in \{p_s, p_r\}$ (resp. from $\mathcal{A}$), send δ to p_i (resp. to $\mathcal{A}$).
- Upon receiving a message $(\texttt{fetch})$ from p_r:
  1. Set $D := D - 1$
  2. If $D = 0$, then send M to p_r
- Upon receiving $(\texttt{delay}, T)$ for the adversary, if $T \leq \delta$ and no $\texttt{delay}$ message was received so far, then set $D := T$; otherwise ignore the message.
- Upon receiving $(\texttt{corrupt}, p_s, m', T')$ from $\mathcal{A}$: if $D > 0$ and T' is a valid delay, then set $D := T'$ and $M := m'$.

Lemma 3. *For any $0 < \delta' < \delta$ the functionality $\left[\mathcal{F}^{\delta'}_{\text{BD-SMT}}\right]^{NT}$ can be implemented in the $\mathcal{F}^{\delta}_{\text{BD-SMT}}$-hybrid model, but $\mathcal{F}^{\delta'}_{\text{BD-SMT}}$ cannot be implemented, even if we assume bidirectional $\mathcal{F}^{\delta}_{\text{BD-SMT}}$-channels.*

Proof. We first prove the impossibility result, namely that $\mathcal{F}^{\delta'}_{\text{BD-SMT}}$ cannot be implemented in the $\mathcal{F}^{\delta}_{\text{BD-SMT}}$-hybrid model. The idea is similar to the proof of Lemma 2: In particular, consider the dummy adversary $\mathcal{A}$ and the environment $\mathcal{Z}$ that corrupts no party and provides as input to p_s a uniformly random bit m. After receiving the leaked value $\ell(m)$ from the functionality (forwarded by $\mathcal{A}$), $\mathcal{Z}$ instructs $\mathcal{A}$ to choose the delay $T = \delta$ at the functionality. Note that in the real ($\mathcal{F}^{\delta}_{\text{BD-SMT}}$-hybrid) world, this means that the channel will ignore the next δ $(\texttt{fetch})$-requests from p_r; however, in the ideal world, the functionality $\mathcal{F}^{\delta'}_{\text{BD-SMT}}$ will ignore at most $\delta' < \delta$ $(\texttt{fetch})$-requests from p_r. Subsequently, $\mathcal{Z}$ activates p_r with input $(\texttt{fetch})$ δ' times; as a result, if $\mathcal{Z}$ is witnessing the ideal-world execution, it will receive the message m from p_r (note that $\mathcal{F}$ does not allow the simulator to further delay this reply); hence, in the real world, the protocol of p_r has to also output some m' to $\mathcal{Z}$. Indeed, the other choices of p_r are to not produce any output, or to send $(\texttt{fetch})$ to $\mathcal{F}^{\delta}_{\text{BD-SMT}}$, or to use the channel in the opposite direction (i.e., as a sender); in all three cases, $\mathcal{Z}$ is activated before p_r receives the next output, which allows $\mathcal{Z}$ to detect that it is witnessing the real-world execution. However, as p_r has no information on m, the probability that $m' = m$ is $1/2$, which results in a noticeable distinguishing advantage for $\mathcal{Z}$.

The reduction of $\left[\mathcal{F}^{\delta'}_{\text{BD-SMT}}\right]^{\text{NT}}$ to $\mathcal{F}^{\delta}_{\text{BD-SMT}}$ is done along the lines of the proof of Lemma 2, but we have to take care of some details occurring due to the different bounds δ and δ'. We show that $\left[\mathcal{F}^{\delta'}_{\text{BD-SMT}}\right]^{\text{NT}}$ can be realized from $\mathcal{F}^{\delta}_{\text{BD-SMT}}$ using the simple protocol $\pi_{\delta'}$ that, on every $(\texttt{fetch})$-query, activates the adversary. On the other hand, upon an activation from the adversary, the protocol issues a $(\texttt{fetch})$-query to $\mathcal{F}^{\delta}_{\text{BD-SMT}}$ and outputs the result. Also, on a $(\texttt{LearnBound})$-query, the protocol answers δ' instead of δ. The simulator $\mathcal{S}$ works as follows: It simulates copies of $\mathcal{A}$, $\mathcal{F}^{\delta}_{\text{BD-SMT}}$, and the protocol machines internally and instructs the (ideal) channel to proceed without additional delay (note that $\left[\mathcal{F}^{\delta'}_{\text{BD-SMT}}\right]^{\text{NT}}$ is defined to make only delayed outputs). $\mathcal{S}$ analyzes the communication among the simulated $\mathcal{A}$ and $\mathcal{F}^{\delta}_{\text{BD-SMT}}$ to determine the chosen delay, and examines the interaction between the protocol machines, $\mathcal{A}$, and the simulated channel. From this communication, $\mathcal{S}$ can easily deduce when the receiving protocol machine would provide output and can instruct the ideal channel to deliver the message. $\square$

A.5 The Imperfectly Synchronized Clock

The "relaxed" clock $\mathcal{F}^{-}_{\text{CLOCK}}$ models imperfectly synchronized clocks available to the parties. The functionality is modeled by modifying $\mathcal{F}_{\text{CLOCK}}$ as follows: For each party p_i, $\mathcal{F}^{-}_{\text{CLOCK}}$ maintains a value t_i^d that corresponds to the number of queries needed by p_i before learning that the round has switched. In particular, as soon as every honest party has sent $\mathcal{F}^{-}_{\text{CLOCK}}$ a $(\texttt{RoundOK})$ message, the clock resets the indicators $d_i := 0$ for $i \in \mathcal{P}$, but, for each party p_i, the first $t_i^d - 1$ $(\texttt{RequestRound})$ messages after the switch are answered with 1, and only the following queries are answered with 0 (which corresponds to the round switch).

Functionality $\mathcal{F}^{-}_{\text{CLOCK}}(\mathcal{P})$

For each $i \in \mathcal{P}$, initialize a bit $d_i := 0$, an initial delay $t_i^d := 1$, and a current delay $t_i := 1$.

- Upon receiving message $(\texttt{RoundOK})$ from Party i set $d_i := 1$. If, for all $j \in \mathcal{H} : d_j = 1$, then reset $d_j := 0$ for all $j \in \mathcal{P}$ and set $t_i := t_i^d$. Send $(\texttt{switch}, i)$ to $\mathcal{A}$.
- Upon receiving $(\texttt{RequestRound})$ from p_i:
  - If $d_i = 1$ or $t_i > 1$, then set $t_i := t_i - 1$ and return 1 to p_i.
  - If $d_i = 0$ and $t_i = 1$, then return 0 to p_i.
- Upon receiving a message $(\texttt{delay}, i, T)$ for the adversary, if T is a valid delay (i.e., encodes a natural number in unary notation), then set $t_i^d := T$; otherwise ignore the message.

A.6 Secure Function Evaluation

An (albeit hard to implement) version of the secure function evaluation (SFE) functionality $\mathcal{F}_{\text{SFE}}$ appears in [Can05]. The functionality models the strictest version of termination one can hope for: Once all honest parties have provided input, the functionality will produce output in the next activation. This functionality cannot be implemented if the network leaks any information about transferred messages (e.g., the fact that a message is transferred). Also, depending on the set-up used by protocols, the functionality must potentially be weakened to drop guarantees if too many parties become corrupted.

Functionality $\mathcal{F}^{f}_{\text{SFE}}(\mathcal{P})$

$\mathcal{F}^{f}_{\text{SFE}}$ proceeds as follows, given a function $f : (\{0,1\}^* \cup \{\perp\})^n \times R \to (\{0,1\}^*)^n$ and a player set $\mathcal{P}$. Initialize the variables $x_1, \ldots, x_n, y_1, \ldots, y_n$ to a default value $\perp$. Next:

- Upon receiving input $(\texttt{input}, v)$ from some party p_i with $i \in \mathcal{P}$, set $x_i := v$ and send a message $(\texttt{input}, i)$ to the adversary.
- Upon receiving input $(\texttt{output})$ from some party p_i with $i \in \mathcal{P}$, do:
  - If x_j has been set for all $j \in \mathcal{H}$, and $y_1, \ldots, y_n$ have not yet been set, then choose $r \xleftarrow{R} R$ and set $(y_1, \ldots, y_n) := f(x_1, \ldots, x_n, r)$.
  - Output y_i to p_i.

To be able to implement the functionality $\mathcal{F}_{\text{SFE}}$ on realistic networks, we augment it by commands that allow to "lose activations" to the adversary to allow the simulator to emulate the network traffic. This functionality still provides guaranteed termination, but it can be implemented, e.g., in the $\{\mathcal{F}_{\text{CLOCK}}, \mathcal{F}_{\text{BD-SMT}}\}$-hybrid model by protocols adapted from, e.g., [GMW87, BGW88, CLOS02]. Again, depending on the set-up and the protocol used, the functionality must potentially drop guarantees if too many players are corrupted.

Functionality $\mathcal{F}_{\text{SFE}}^{f,Rnd}(\mathcal{P})$

$\mathcal{F}_{\text{SFE}}^{f,Rnd}$ proceeds as follows, given a function $f : (\{0,1\}^* \cup \{\perp\})^n \times R \to (\{0,1\}^*)^n$, a round function Rnd, and a player set $\mathcal{P}$. For each $i \in \mathcal{P}$, initialize variables x_i and y_i to a default value $\perp$ and a current delay $t_i := |\mathcal{P}| + 1$. Moreover, initialize a global round counter $\ell := 1$. Next:

- Upon receiving input $(\mathtt{input}, v)$ from some party p_i with $i \in \mathcal{P}$, set $x_i := v$ and send a message $(\mathtt{input}, i)$ to the adversary.
- Upon receiving input $(\mathtt{output})$ from some party p_i with $i \in \mathcal{P}$, if $p_i \in \mathcal{H}$ and x_i has not yet been set then ignore p_i's message, else do:
  - If $t_i > 1$, then set $t_i := t_i - 1$. If (now) $t_j = 1$ for all $j \in \mathcal{H}$, then set $\ell := \ell + 1$ and $t_j := |\mathcal{P}| + 1$ for all $j \in \mathcal{P}$. Send $(\mathtt{activated}, i)$ to the adversary.
  - else, if $t_i = 1$ but $\ell < Rnd$, then send $(\mathtt{early})$ to p_i.
  - else:
    - if x_j has been set for all $j \in \mathcal{H}$, and $y_1, \ldots, y_n$ have not yet been set, then choose $r \xleftarrow{R} R$ and set $(y_1, \ldots, y_n) := f(x_1, \ldots, x_n, r)$.
    - Output y_i to p_i.

The $|\mathcal{P}| + 1$ activations in each round are necessary to simulate the fact that for the parties the protocol needs enough activations to send their messages within a round and notify $\mathcal{F}_{\text{CLOCK}}$.

A.7 Secure (Reactive) Multi-Party Computation

We next describe the functionality $\mathcal{F}_{\text{MPC}}$ for reactive multi-party computation. For simplicity, we first describe the extension of $\mathcal{F}_{\text{SFE}}$ from [Can05] to the reactive case. As with the non-reactive function evaluation (SFE), this MPC functionality is too strong to be implemented in "reasonable" synchronous networks. However, introducing it will allow us to deal with reactiveness separately from the problems that occur due to delays and synchronization issues.

A reactive computation can be specified as an ordered sequence of secure function evaluations which can maintain a joint state. The state used to evaluate any function in this sequence is passed on to the subsequent functions. For simplicity we assume, as in the case of $\mathcal{F}_{\text{SFE}}$, that each time a party is expected to give input to (resp. receive output from) the computation, *every* party gives –a potentially void– input (resp. receives output). More precisely, the computation is described as a vector of functions $\vec{f} = (f_1, \ldots, f_m)$, where each $f_\lambda \in \vec{f}$ takes as input a vector of values from $\{0,1\}^* \cup \{\perp\}$, a uniformly random value r from a known domain R, as well as a *state-vector* $\vec{\mathcal{S}_\lambda} \in ((\{0,1\}^* \cup \{\perp\})^n \times R)^{(\lambda-1)}$. The state-vector $\vec{\mathcal{S}_\lambda}$ contains the inputs and randomness used for evaluating the functions $f_1, \ldots, f_{\lambda-1}$; each $f_\lambda \in \vec{f}$ outputs a vector of strings $\vec{y}_\lambda \in \{0,1\}^n$.

The functionality $\mathcal{F}_{\text{MPC}}^{\vec{f}}$ is parametrized by the vector of functions $\vec{f} = (f_1, \ldots, f_m)$. For each $f_\lambda \in \vec{f}$, $\mathcal{F}_{\text{MPC}}$ might receive an input $x_{i,\lambda}$ from party i at any point before f_λ is evaluated (as soon as f_λ is evaluated and the output has been given to some party, $\mathcal{F}_{\text{MPC}}^{\vec{f}}$ stops accepting inputs for f_λ). In that case, $\mathcal{F}_{\text{MPC}}$ records $x_{i,\lambda}$ as p_i's input to the evaluation of f_λ. We denote the vector of the parties' inputs to the function f_λ as $\vec{x}_\lambda = (x_{1,\lambda}, \ldots, x_{n,\lambda})$. Similarly, we denote the vector of the parties' outputs from f_λ as $\vec{y}_\lambda = (y_{1,\lambda}, \ldots, y_{n,\lambda})$. We say that $\vec{x}_\lambda$ (resp. $\vec{y}_\lambda$) *has been set* if all the *honest* parties have handed $\mathcal{F}_{\text{MPC}}$ their input for f_λ (resp. the output of f_λ has already been computed). When a request for producing output for some f_λ is received from some honest p_j, $\mathcal{F}_{\text{MPC}}$ checks that the input $\vec{x}_\lambda$ and all outputs $\vec{y}_1, \ldots, \vec{y}_{\lambda-1}$ have been set; if so, $\mathcal{F}_{\text{MPC}}$ computed f_λ's outcome $\vec{y}_\lambda = (y_{1,\lambda}, \ldots, y_{n,\lambda})$ using the inputs and the state-vector $\vec{\mathcal{S}_\lambda}$, records the output, and hands to p_j his part of the output, i.e., $y_{i,\lambda}$. Recall that the state consist of the inputs to all previous functions and the corresponding randomness, i.e. $\vec{\mathcal{S}_\lambda} = (\vec{x}_1, \ldots, \vec{x}_{\lambda-1}, r_1, \ldots, r_{\lambda-1})$. A detailed description of $\mathcal{F}_{\text{MPC}}^{\vec{f}}$ appears in the following:

Functionality $\mathcal{F}_{\text{MPC}}^{\vec{f}}(\mathcal{P})$

$\mathcal{F}_{\text{MPC}}^{\vec{f}}$ proceeds as follows, given a vector of functions $\vec{f} = (f_1, \dots, f_m)$ and a player set $\mathcal{P}$, where $|\mathcal{P}| = n$, and for $\lambda = 1, \dots, m, f_\lambda : (\{0,1\}^* \cup \{\perp\})^n \times ((\{0,1\}^* \cup \{\perp\})^n \times R)^{(\lambda-1)} \times R \to (\{0,1\}^*)^n$. Initialize the variables $x_{1,1}, \dots, x_{n,m}, y_{1,1}, \dots, y_{n,m}$ to a default value $\perp$ and $\vec{\mathcal{S}}_0 := (\perp, \dots, \perp)$. Next:

- Upon receiving input $(\texttt{input}, \lambda, v)$ from some party p_i with $i \in \mathcal{P}$ and $\lambda \in \{1, \dots, m\}$, if $\vec{y}_\lambda$ has not yet been set, then set $x_{i,\lambda} := v$ and send a message $(\texttt{input}, i, \lambda)$ to the adversary.
- Upon receiving input $(\texttt{output}, \lambda)$ from some party p_i with $i \in \mathcal{P}$ and $\lambda \in \{1, \dots, m\}$, do:
  - If $\vec{x}_\lambda$ and $\vec{y}_1, \dots, \vec{y}_{\lambda-1}$ have been set, and $\vec{y}_\lambda$ has not yet been set, then choose $r_\lambda \overset{R}{\leftarrow} R$ and set $(y_{1,\lambda}, \dots, y_{n,\lambda}) := f(\vec{x}_\lambda, \vec{\mathcal{S}}_{\lambda-1}, r_\lambda)$; also set $\vec{\mathcal{S}}_\lambda := (\vec{x}_1, \dots, \vec{x}_\lambda, r_1, \dots, r_\lambda)$.
  - Output $y_{i,\lambda}$ to p_i.

As already mentioned, the above functionality cannot be realized under "reasonable" synchronicity assumptions. Indeed, one can verify that this would contradict Lemma 4, as $\mathcal{F}_{\text{SFE}}^{f}$ is the same as $\mathcal{F}_{\text{MPC}}^{\vec{f}'}$, where $\vec{f}' = (f)$. For that reason, we modify it along the ways of the modification we did for the SFE functionality from [Can05]: the MPC functionality is parametrized, in addition to the vector of functions $\vec{f} = (f_1, \dots, f_m)$, by a vector of round-functions $\vec{Rnd} = (Rnd_1(\cdot), \dots, Rnd_m(\cdot))$, where each $Rnd_\lambda \in \vec{Rnd}$ is a function of the security parameter which specifies the number of rounds that are used for evaluating f_λ; as in the case of $\mathcal{F}_{\text{SFE}}$, in each round the simulator is given $|\mathcal{P}|+1$ activations for each party p_i, in order to simulate the messages sent from p_i in that round. Because Rnd_λ corresponds to the number of rounds needed for the evaluation of f_λ, and the output of f_λ should be generated only after the outputs of $f_1, \dots, f_{\lambda-1}$ have been set, if the MPC functionality receives a request for output (from some $p_i \in \mathcal{H}$) for f_λ in round $\ell < \sum_{\rho=1}^{\lambda} Rnd_\rho$, then it notifies p_i that it is too early for generating this output.

Functionality $\mathcal{F}_{\text{MPC}}^{\vec{f},\vec{Rnd}}(\mathcal{P})$

$\mathcal{F}_{\text{SFE}}^{\vec{f},\vec{Rnd}}$ proceeds as follows, given a vector of functions $\vec{f} = (f_1, \dots, f_m)$, a vector of round functions $\vec{Rnd} = (Rnd_1, \dots, Rnd_m)$, and a player set $\mathcal{P}$, where $|\mathcal{P}| = n$, and for $\lambda = 1, \dots, m, f_\lambda : (\{0,1\}^* \cup \{\perp\})^n \times ((\{0,1\}^* \cup \{\perp\})^n \times R)^{(\lambda-1)} \times R \to (\{0,1\}^*)^n$. Initialize the variables $x_{1,1}, \dots, x_{n,m}, y_{1,1}, \dots, y_{n,m}$ to a default value $\perp$, initialize $\vec{\mathcal{S}}_0 := (\perp, \dots, \perp)$, and for each pair $(i, \lambda) \in \{1, \dots, |\mathcal{P}|\} \times \{1, \dots, m\}$ initialize the delay $t_{i,\lambda} := |\mathcal{P}| + 1$. Moreover, initialize a global round counter $\ell := 1$. Next:

- Upon receiving input $(\texttt{input}, \lambda, v)$ from some party p_i with $i \in \mathcal{P}$ and $\lambda \in \{1, \dots, m\}$, if $\vec{y}_\lambda$ has not yet been set, then set $x_{i,\lambda} := v$ and send a message $(\texttt{input}, i, \lambda)$ to the adversary.
- Upon receiving input $(\texttt{output}, \lambda)$ from some party p_i with $i \in \mathcal{P}$ and $\lambda \in \{1, \dots, m\}$, if $p_i \in \mathcal{H}$ and for some $\rho \in \{1, \dots, \lambda\}$ the input $x_{i,\rho}$ has not yet been set then ignore p_i's message, else do:
  - If $t_{i,\lambda} > 1$, then set $t_{i,\lambda} := t_{i,\lambda} - 1$. If (now) $t_{j,\lambda} = 1$ for all $j \in \mathcal{H}$, then set $\ell := \ell + 1$ and $t_{j,\lambda} := |\mathcal{P}| + 1$ for all $j \in \mathcal{P}$. Send $(\texttt{activated}, i, \lambda)$ to the adversary
  - else, if $t_{i,\lambda} = 1$ but $\ell < \sum_{\rho=1}^{\lambda} Rnd_\rho$, then send $(\texttt{early})$ to p_i
  - else:
    - If $\vec{x}_\lambda$ and $\vec{y}_1, \dots, \vec{y}_{\lambda-1}$ have been set, and $\vec{y}_\lambda$ has not yet been set, then choose $r_\lambda \overset{R}{\leftarrow} R$ and set $(y_{1,\lambda}, \dots, y_{n,\lambda}) := f(\vec{x}_\lambda, \vec{\mathcal{S}}_{\lambda-1}, r_\lambda)$; also set $\vec{\mathcal{S}}_\lambda := (\vec{x}_1, \dots, \vec{x}_\lambda, r_1, \dots, r_\lambda)$.
    - Output $y_{i,\lambda}$ to p_i.

Definition 2 (Guaranteed Termination – Reactive Computation)**.** We say that a protocol UC-securely realizes a (possibly reactive) computation described by the function-vector $\vec{f}$ *with guaranteed termination*, if it UC-realizes a functionality $\mathcal{F}_{\text{MPC}}^{\vec{f},\vec{Rnd}}$ for some vector of round functions $\vec{Rnd}$.

It is straightforward to verify that all the results of Section 4.2 that are stated for $\mathcal{F}_{\text{SFE}}^{f}$ and for $\mathcal{F}_{\text{SFE}}^{f,Rnd}$, i.e., Lemma 4 and Theorem 5, apply also to $\mathcal{F}_{\text{MPC}}^{\vec{f}}$ and $\mathcal{F}_{\text{MPC}}^{\vec{f},\vec{Rnd}}$, respectively.

A.8 Universally Composable SFE with Guaranteed Termination

In this section, we translate a result from Kushilevitz et al. [KLR06] into our setting to obtain UC-secure protocols for secure function evaluation with guaranteed termination. The structure of the proof of Theorem 5 is similar to the proof of [KLR06, Theorem 7.1] with some adaptations.

Because of the considered model of communication channels, it is important that the simulator can faithfully emulate the "communication pattern" of the protocol. In contrast, most of the literature on secure function evaluation assumes that communication channels are perfectly secure: they do not even leak the information that a message transfer takes place. This is in contrast to our channel $\mathcal{F}_{\text{BD-SMT}}$, which activates the adversary upon message transmission and, hence, leaks this information. But as our proof considers "standard synchronous" protocols that proceed in rounds and, within each such round, send exactly one message to each one of the other parties, this assumption can be seen to be fulfilled. This also implies that the implemented functionality $\mathcal{F}_{\text{SFE}}^{f,c}$ will require c additional activations from each honest party p_i and issue a message $(\texttt{activated}, i)$ to $\mathcal{A}$ upon each such activation.

As, in our model, the parties have access to the clock functionality $\mathcal{F}_{\text{CLOCK}}$, we redefine the notion of start synchronization from [KLR06] in a simpler way: We say that a protocol ρ has *start synchronization* if it begins with the following steps for *all* parties:

1. Send $(\texttt{RoundOK})$ to $\mathcal{F}_{\text{CLOCK}}$,
2. In each further activation, query $(\texttt{RequestRound})$ to $\mathcal{F}_{\text{CLOCK}}$. Once $\mathcal{F}_{\text{CLOCK}}$ answers with 0, proceed with the execution of ρ.

Note that the goal of start synchronization is that the first message of the protocol that actually depends on the input of an honest party is sent only after the inputs to all honest parties are fixed. Technically, this is needed to "fix" the honest users' inputs from the environment $\mathcal{Z}$ in the simulated UC-execution. Recall that for a synchronous protocol ρ, we denote by $\hat{\rho}$ the protocol which is obtained by extending ρ with such a start-synchronization round.

Theorem 5. *Let f be a function and let ρ be a protocol that, according to the notion of [Can00], realizes f with computational (or statistical or perfect) security in the stand-alone model, with an efficient straight-line black-box simulator. Then $\hat{\rho}$ UC-realizes f with computational (or statistical or perfect) security and guaranteed termination in the $\{\mathcal{F}_{\text{CLOCK}}, \mathcal{F}_{\text{BD-SEC}}\}$-hybrid model with static adversaries.*

Note that we do not use the fully non-adaptive version of [Can00] to start from, but the "initially-adaptive" version where the adversary is allowed to adaptively choose the players to be corrupted until the first protocol message is sent. This is equivalent to the non-adaptive version by [Can00, Section 4.1, Remark 5]. The same argument is used implicitly in the proof of [KLR06]. As in [KLR06], we prove the statement only for the computational setting, the extension to the statistical case is straightforward.

Also note that the network implied by [Can00] provides strong synchrony guarantees. In particular, all honest parties compute on the messages that are in the channels at the point in time when the first honest party starts the computation of a round. (In reality, it would be conceivable that the first party p_i has finished computation and sending for round r before another party p_j's clock tells p_j to start computing. In this case, the adversary could inject round $r-1$ messages to p_j that actually depend on the round r messages of p_i.) This strong synchrony assumption becomes explicit in our model by doubling the round number of ρ and implement alternating "receive" and "compute/send" rounds.

The synchronous protocol ρ has to be detailed slightly to implement the interface of $\mathcal{F}_{\text{SFE}}^{f,Rnd}$. For each pair of parties and round of the protocol, an independent channel $\mathcal{F}_{\text{BD-SMT}}$ is used. After the local computation of the messages for round r, the following $|\mathcal{P}|+1$ activations are consumed to send the messages $m_{i,j,r}$ via the channels to the protocol machines ρ_j and send $(\texttt{RoundOK})$ to $\mathcal{F}_{\text{CLOCK}}$. Upon each further activation, send $(\texttt{RequestRound})$ to $\mathcal{F}_{\text{CLOCK}}$ and output $(\texttt{early})$ locally if $d = 1$, and proceed with setting $r := r + 1$ and sending $(\texttt{fetch})$ to the channels $\mathcal{F}_{\text{BD-SMT}}$ to obtain all messages.

Proof of Theorem 5. By the assumption on the security of ρ, we know that there is a stand-alone simulator $\mathcal{S}_\rho$ with the stated additional properties. Without loss of generality, we assume that ρ has start synchronization, i.e., $\rho = \rho'$

(if this is not the case, then we cas trivially extend ρ with the two "start-synchronizing" steps sketched above). We describe a UC-simulator $\mathcal{S}$ for the protocol ρ executed in the $\{\mathcal{F}_{\text{CLOCK}}, \mathcal{F}_{\text{BD-SMT}}\}$-hybrid model with the dummy adversary $\mathcal{D}$, which makes black-box use of $\mathcal{S}_\rho$.

The UC-Simulator $\mathcal{S}$.

Create instances of the assumed stand-alone simulator $\mathcal{S}_\rho$, the ideal functionalities $\mathcal{F}_{\text{CLOCK}}$ and $\mathcal{F}_{\text{BD-SMT}}$, and the dummy adversary $\mathcal{D}$.

- Forward all messages between the environment $\mathcal{Z}$ and the emulated adversary $\mathcal{D}$ (via input and subroutine output tapes), as well as between $\mathcal{D}$ and the simulated ideal functionalities.
- Forward the corruption-related communication between $\mathcal{S}_\rho$ and $\mathcal{F}_{\text{SFE}}^{f,Rnd}$ (via the respective communication tapes).
- Receiving $(\texttt{activated}, i)$ from $\mathcal{F}_{\text{SFE}}^{f,Rnd}$:
  - If this is the first such message in a round r, provide the messages sent via $\mathcal{F}_{\text{BD-SMT}}(j,i)$ with $j \in \mathcal{P} \setminus \mathcal{H}$ for round $r-1$ and obtain the round r messages $m_{i,j,r}$ from ρ_i to ρ_j with $j \in \mathcal{P} \setminus \mathcal{H}$ as well as the leakages $\ell(m_{i,j,r})$ for $i, j \in \mathcal{H}$ as a reply.
  - Using the communication pattern of the protocol, simulate either the leakage $\ell(m_{i,j,r})$ from $\mathcal{F}_{\text{BD-SMT}}(i,j)$ if $j \in \mathcal{H}$, the message $m_{i,j,r}$ if $j \in \mathcal{P} \setminus \mathcal{H}$, or the notification $(\texttt{switch}, i)$ to $\mathcal{D}$.

In the following, we have to show that if there exists an environment $\mathcal{Z}$ that distinguishes between the real execution of protocol ρ and the ideal execution of $\mathcal{F}_{\text{SFE}}^{f,Rnd,n+1}$ with simulator $\mathcal{S}$ with non-negligible advantage, then we can construct a stand-alone adversary $\mathcal{A}_\mathcal{Z}$ that attacks ρ in the stand-alone model with non-negligible probability.

The Stand-alone Adversary $\mathcal{A}_\mathcal{Z}$.

Split the auxiliary input into a random tape and an auxiliary input for $\mathcal{Z}$. Create and run an instance of the environment $\mathcal{Z}$ with the random tape and input as above, as well as the (dummy) adversary $\mathcal{D}$, the clock $\mathcal{F}_{\text{CLOCK}}$ and the channels $\mathcal{F}_{\text{BD-SMT}}(i,j)$ for $i, j \in \mathcal{P}$, $i \neq j$. In this simulated computation, forward the communication between the simulated parties $\mathcal{D}$, $\mathcal{Z}$, $\mathcal{F}_{\text{CLOCK}}$, and $\mathcal{F}_{\text{BD-SMT}}$.

- Run the simulation until $\mathcal{Z}$ provides the first input to a protocol, answer corruption requests from $\mathcal{D}$ with simulated "vanilla" ρ-states, and corrupt the corresponding players in ρ.
- The UC-interaction is simulated almost entirely, with the exception for the computation of the honest parties ρ_i. Instead, on every other round switch at $\mathcal{F}_{\text{CLOCK}}$, provide all messages injected by $\mathcal{D}$ into $\mathcal{F}_{\text{BD-SMT}}(j,i)$ in the name of the corrupted players $j \in \mathcal{P} \setminus \mathcal{H}$ to the real ρ_i.
- Being activated with the messages $m_{i,j,r}$ with $i \in \mathcal{H}$, $j \in \mathcal{P} \setminus \mathcal{H}$, and $\ell(m_{i,j,r})$ with $i, j \in \mathcal{H}$, use these values in the further simulation until the next but one round switch of $\mathcal{F}_{\text{CLOCK}}$ signals that the round r messages of the corrupted parties are also fixed.
- After the execution of a protocol machine ρ_i in the UC execution would terminate, output the transcript of the execution (including the auxiliary input) and halt. Note that this transcript is complete, as by the time that ρ_i *does* produce output, it is guaranteed that all messages of the final round must be in the channels.

The main idea of the proof is the following. Since the protocol ρ has start synchronization, the input and the random tape of $\mathcal{Z}$ determine the inputs that $\mathcal{Z}$ gives to the honest protocol machines (because the rest of the execution within $\mathcal{A}_\mathcal{Z}$ is deterministic. This allows us to restrict our attention to this particular set of inputs in the stand-alone execution. If the environment distinguishes the real and the ideal UC-executions, then the distribution of the environment's "view" of the execution of ρ (as a random variable consisting of all messages sent and received by the environment) together with the outputs of the honest ρ_i must be distinguishable. But this corresponds to the fact that the transcripts in the stand-alone executions of ρ with $\mathcal{A}_\mathcal{Z}$ and $\mathcal{S}_\rho$ (can be extended to transcripts which) are distinguishable.

The protocol execution begins with a corruption phase: According to the definition of static corruption in [Can05], only parties p_i for which the environment sent a $(\texttt{corrupt})$ message prior to the first input to a protocol machine may be corrupted during the protocol execution (by the definition of the control function). This is compatible with the "initially-adaptive" corruption allowed for $\mathcal{A}_{\mathcal{Z}}$ by the description in [Can00]. In the real stand-alone execution of ρ with $\mathcal{A}_{\mathcal{Z}}$, the inputs provided by the environment $\mathcal{Z}$ to the honest parties p_i are determined by the auxiliary input and the random tape of the (simulated) environment $\mathcal{Z}$: By the start synchronization, the first message of the protocol ρ (not equal to "begin ρ") is sent only after *all* parties p_i have sent their synchronization messages, acknowledging their input. All computations that $\mathcal{A}_{\mathcal{Z}}$, $\mathcal{D}$, the honest ρ_i, and the ideal functionalities have performed up to this point are deterministic. In the following description, we restrict our attention to executions in the stand-alone model where the inputs to p_i and ρ_i indeed coincide (both in the real execution with $\mathcal{A}_{\mathcal{Z}}$ and the black-box ideal execution with $\mathcal{A}_{\mathcal{Z}}$ and $\mathcal{S}_\rho$).

The environment's "view" in the real UC-execution is the same as the simulated "view" in the stand-alone execution within $\mathcal{A}_{\mathcal{Z}}$ restricted to the cases where the inputs and outputs of the honest ρ_i indeed coincide.[15] For all inputs except for those that provoke a message from the channel $\mathcal{F}_{\text{BD-SMT}}(i,j)$ to $\mathcal{D}$, this is straightforward: $\mathcal{A}_{\mathcal{Z}}$ only performs a straightforward simulation of the other ITMs. But for the messages from $\mathcal{F}_{\text{BD-SMT}}(i,j)$, the inputs and messages sent to the protocol machines ρ_i are equal (resp. have the same distribution) by assumption, and the computation of the stand-alone and the UC-version of ρ are exactly the same.

The environment's "view" in the ideal UC-execution is the same as the simulated "view" in the stand-alone execution in case the inputs and outputs of the honest ρ_i coincide. Again, responses that do not contain (simulated) messages from the honest ρ_i are the same by the fact that both $\mathcal{A}_{\mathcal{Z}}$ and $\mathcal{S}$ perform a straightforward simulation of the real setting. But responses that involve messages are also the same in both cases: The simulator $\mathcal{S}$ in the UC-execution uses the messages supplied by $\mathcal{S}_\rho$ as message intended from an honest ρ_i to a corrupted ρ_j, and the same messages are used by $\mathcal{A}_{\mathcal{Z}}$ in the ideal (black-box straight-line) execution with $\mathcal{S}$.

Our goal is to conclude that, if an environment $\mathcal{Z}$ distinguishes with non-negligible advantage between the real and ideal UC-executions, the stand-alone transcripts of the real execution with $\mathcal{A}_{\mathcal{Z}}$ and the ideal execution with $\mathcal{A}_{\mathcal{Z}}$ and $\mathcal{S}_\rho$ can also be distinguished with non-negligible advantage. We construct a distinguisher D as follows: By the above construction, the output of $\mathcal{A}_{\mathcal{Z}}$ contains the auxiliary input to $\mathcal{A}_{\mathcal{Z}}$, which determines the input and random tape to $\mathcal{Z}$. Hence, D can replay the complete execution and obtain the state of $\mathcal{Z}$ at the time the execution was aborted. D continues the simulation of the real model and provides $\mathcal{Z}$ with the output of the honest protocol machines ρ_i as obtained from the transcript. Once the environment $\mathcal{Z}$ halts with output bit b, D uses the same bit b as its output.

The auxiliary input to $\mathcal{A}_{\mathcal{Z}}$ is interpreted as a tuple $x = (r, z)$ where $r \in \{0,1\}^*$ is used as a random tape and $z \in \{0,1\}^*$ as an auxiliary output to $\mathcal{Z}$. Hence, if we consider the ensembles corresponding to the stand-alone settings where the variable R corresponding to $\mathcal{Z}$'s random tape is chosen according to the correct distribution, we obtain

$$\Big\{D\big(\text{REAL}^{\mathsf{SA}}_{\rho,\mathcal{A}_{\mathcal{Z}}}(k,(R,z),\mathbf{x}(R,z))\big)\Big\}_{k\in\mathbb{N},z\in\{0,1\}^*} \equiv \Big\{\text{REAL}^{\mathsf{UC}}_{\rho,\mathcal{D},\{\mathcal{F}_{\text{CLOCK}},\mathcal{F}_{\text{BD-SMT}}\},\mathcal{Z}}(k,z)\Big\}_{k\in\mathbb{N},z\in\{0,1\}^*}$$

and

$$\Big\{\text{IDEAL}^{\mathsf{UC}}_{\mathcal{F}^{f,c}_{\text{SFE}},\mathcal{S},\mathcal{Z}}(k,z)\Big\}_{k\in\mathbb{N},z\in\{0,1\}^*} \equiv \Big\{D\big(\text{IDEAL}^{\mathsf{SA}}_{f,\mathcal{S}^{\mathcal{A}}}(k,(R,z),\mathbf{x}(R,z))\big)\Big\}_{k\in\mathbb{N},z\in\{0,1\}^*},$$

where $\mathbf{x}(R,z)$ denotes the inputs that $\mathcal{Z}$ computes for the honest parties p_i, given random tape R and auxiliary input z. By the above arguments, these are the only parameters that determine the inputs. The statement holds, as all values in the simulated UC-executions are determined according to exactly the same computations as in the corresponding "real" UC-executions. Hence, by the assumption that $\mathcal{Z}$ is a successful distinguisher, we conclude that

$$\Big\{D\big(\text{REAL}^{\mathsf{SA}}_{\rho,\mathcal{A}_{\mathcal{Z}}}(k,x)\big)\Big\}_{k\in\mathbb{N},x\in\{0,1\}^*} \not\approx \Big\{D\big(\text{IDEAL}^{\mathsf{SA}}_{f,\mathcal{S}^{\mathcal{A}}}(k,x)\big)\Big\}_{k\in\mathbb{N},x\in\{0,1\}^*},$$

which concludes the main argument.

As the functionality $\mathcal{F}^{f,c}_{\text{SFE}}$ models guaranteed termination, this proof implies that the protocol ρ achieves guaranteed termination if it is executed using channels of the type $\mathcal{F}_{\text{BD-SMT}}$. The guaranteed-delivery property of the

[15] To make this argument formal, one could use the theory of random systems from [Mau02].

channels is used in the proof to conclude that the messages provided to the honest parties ρ_i are indeed available in the beginning of each "receive" sub-round by the fact that all honest parties inject them into the channels before they agree to switch rounds, and the channel $\mathcal{F}_{\text{BD-SMT}}$ immediately makes the messages available to the intended receiver. $\square$

A.9 The Need for Both Synchronization and Bounded-Delay

Lemma 6. *There are functions f such that for any (efficient) round-function Rnd and any $\delta > 0$: $\left[\mathcal{F}_{\text{SFE}}^{f,Rnd}\right]^{NT}$ can be realized in the $\mathcal{F}_{\text{BD-SMT}}^{\delta}$-hybrid model, but $\mathcal{F}_{\text{SFE}}^{f,Rnd}$ cannot, even when we assume bidirectional $\mathcal{F}_{\text{BD-SMT}}^{\delta}$-channels.*

Proof. Consider the deterministic two-party function f that on input a bit $x_1 \in \{0,1\}$ from p_1 (and no input from p_2) outputs x_1 to p_2 (f outputs $y_2 := 0$ if $x_1 = \perp$) and nothing to p_2. We consider the cases (I) $\delta = 1$ and (II) $\delta > 1$ separately.

Case I ($\delta = 1$): For the impossibility, assume, towards a contradiction, that there exists a $\mathcal{F}_{\text{BD-SMT}}^{1}$-hybrid protocol π which securely realizes $\mathcal{F}_{\text{SFE}}^{f,Rnd}$. Consider the following two scenarios which involve a dummy adversary and an environment that never corrupts p_2:

- Scenario 1: p_1 is corrupted and the adversary is instructed to force p_1 to never send any message. As a result, in protocol π, the $(\texttt{fetch})$-queries issued by p_2 will result in an activation of $\mathcal{Z}$.
- Scenario 2: p_1 is honest, but never activated by the environment. Again, in protocol π, the $(\texttt{fetch})$-queries issued by p_2 will result in an activation of $\mathcal{Z}$. Hence p_2 cannot distinguish this scenario from Scenario 1.

In Scenario 1, after receiving $(|\mathcal{P}|+1)Rnd$ queries of the type $(\texttt{output})$ from $\mathcal{Z}$, the (real world) program of p_2 should output 0 (this is the default input of p_1) to $\mathcal{Z}$, as this is the result in the ideal world. However, in Scenario 2, p_2 should output $(\texttt{early})$ to the environment, independently of the number of $(\texttt{output})$-messages received. Indeed, as both parties are honest, $\mathcal{F}_{\text{SFE}}$ will wait for p_1's input before producing output. The environment chooses to implement one of the two above scenarios at random; as the two scenarios are indistinguishable for p_2, the action taken by p_2 will be different from the "correct action" with probability at least $1/2$, which allows $\mathcal{Z}$ to distinguish.

$\left[\mathcal{F}_{\text{SFE}}^{f,Rnd}\right]^{\text{NT}}$ can be implemented from $\mathcal{F}_{\text{BD-SMT}}^{1}$ by the simple protocol in which p_1 sends its input via the channel $\mathcal{F}_{\text{BD-SMT}}^{1}$. Indeed, the simulator is informed about who is corrupted and about when messages are sent to $\left[\mathcal{F}_{\text{SFE}}^{f,Rnd}\right]^{\text{NT}}$ and can delay the outputs of the functionality arbitrarily, therefore, it can ensure that the exchanged outputs are only handed to the (dummy) parties when the environments expects to see them.

Case II ($\delta > 1$): The impossibility of implementing $\mathcal{F}_{\text{SFE}}^{f,Rnd}$ from $\mathcal{F}_{\text{BD-SMT}}^{\delta}$ is argued as follows: Assume, towards a contradiction, that there exists a protocol π which UC realizes $\mathcal{F}_{\text{SFE}}^{f,Rnd}$ in the $\mathcal{F}_{\text{BD-SMT}}^{\delta}$-hybrid world. Because faster channels are stronger (see Section 3.3), this implies that there exists a $\mathcal{F}_{\text{BD-SMT}}^{1}$-hybrid protocol π' (in fact, $\pi' = \pi$) which UC-securely implements $\mathcal{F}_{\text{SFE}}^{f,Rnd}$, contradicting the impossibility statements from Case I.

To complete the proof, we show that $\left[\mathcal{F}_{\text{SFE}}^{f,Rnd}\right]^{\text{NT}}$ can be implemented from $\mathcal{F}_{\text{BD-SMT}}^{\delta}$: Lemma 3 implies that $\left[\mathcal{F}_{\text{BD-SMT}}^{1}\right]^{\text{NT}}$ can be implemented in the $\mathcal{F}_{\text{BD-SMT}}^{\delta}$-hybrid model. Furthermore, in Case I we showed that $\mathcal{F}_{\text{BD-SMT}}^{1}$ is sufficient for $\left[\mathcal{F}_{\text{SFE}}^{f,Rnd}\right]^{\text{NT}}$, which extents to $\left[\mathcal{F}_{\text{BD-SMT}}^{1}\right]^{\text{NT}}$. Indeed, relaxing the hybrid $\mathcal{F}_{\text{BD-SMT}}^{1}$ to its non-terminating relaxation does not have any effect on the statement as the functionality we want to achieve is already of this type. In particular, we use the same protocol as in Case I, the only difference in the simulation is that whenever the adversary delays some output in the $\left[\mathcal{F}_{\text{BD-SMT}}^{1}\right]^{\text{NT}}$, the simulator also delays the outputs in ideal model accordingly. $\square$

Lemma 7. *There are functions f such that for any (efficient) round-function Rnd: $\left[\mathcal{F}_{\text{SFE}}^{f,Rnd}\right]^{NT}$ can be realized in the $\{\mathcal{F}_{\text{ED-SMT}}, \mathcal{F}_{\text{CLOCK}}\}$-hybrid model, but $\mathcal{F}_{\text{SFE}}^{f,Rnd}$ cannot, even when we assume bidirectional $\mathcal{F}_{\text{ED-SMT}}$-channels.*

Proof. We prove the statement for f being the two-party function used in the proof of Lemma 6, i.e., on input a bit $x_1 \in \{0,1\}$ from p_1 (and no input from p_2) f outputs x_1 to p_2 (f outputs $y_2 := 0$ if $x_1 = \perp$) and nothing to p_2.

The sufficiency of $\{\mathcal{F}_{\text{ED-SMT}}, \mathcal{F}_{\text{CLOCK}}\}$ for $\left[\mathcal{F}_{\text{SFE}}^{f,Rnd}\right]^{\text{NT}}$ follows along the lines of the sufficiency condition in Case II in the proof of Lemma 6. Indeed, the channels $\mathcal{F}_{\text{ED-SMT}}$ are sufficient even without $\mathcal{F}_{\text{CLOCK}}$.

We show that $\mathcal{F}_{\text{SFE}}^{f,Rnd}$ cannot be implemented by an $\{\mathcal{F}_{\text{ED-SMT}}, \mathcal{F}_{\text{CLOCK}}\}$-hybrid protocol. Consider the dummy adversary $\mathcal{A}$ and the environment $\mathcal{Z}$ which corrupts no party and has the following distinguishing strategy: $\mathcal{Z}$ starts by giving a uniformly random two-bit input $m \in \{0,1\}^2$ to p_1. As soon as it receives the leak function $\ell(m)$ from the functionality (forwarded by $\mathcal{A}$), $\mathcal{Z}$ instructs $\mathcal{A}$ to give delay $T = Rnd \cdot (|\mathcal{P}|+1)+2$ to $\mathcal{F}_{\text{ED-SMT}}$. Subsequently, $\mathcal{Z}$ activates p_2 with input $(\texttt{fetch})$ $Rnd \cdot (|\mathcal{P}|+1)+1$ times in a row; as a result, if $\mathcal{Z}$ is witnessing the ideal-world execution, he will receive the message m from p_2 (note that $\mathcal{F}_{\text{SFE}}^{f,Rnd}$ does not allow the simulator to stall this reply). Hence, p_2's protocol (in the real world) also has to output the message m to $\mathcal{Z}$ after $(Rnd \cdot (|\mathcal{P}|+1)+1)$ activations. However, the protocol of p_2 cannot get any information on m from $\mathcal{F}_{\text{ED-SMT}}$, because by setting the delay to $T = Rnd \cdot (|\mathcal{P}|+1)+2$, the environment makes sure that the $Rnd \cdot (|\mathcal{P}|+1)+1$ are not sufficient for making $\mathcal{F}_{\text{ED-SMT}}$ deliver m to p_2. The only possible way out for p_2 would be to interact with the functionality $\mathcal{F}_{\text{CLOCK}}$. However, because p_1 is activated only once $\mathcal{F}_{\text{CLOCK}}$ allows for at most one bit of communication (i.e., the bit indicating whether or not p_1 has switch his indicator bit d_1 in this one activation). Because m is a two-bit message, the probability that p_2 outputs m is at most 1/2, which gives $\mathcal{Z}$ a noticeable distinguishing advantage. ☐

A.10 Simultaneous Multi-Send

The various models for synchronous computations considered in the literature define adversarial models with different levels of power. One particular model has been coined *rushing* by [Can00]. In this model, the adversary is not allowed to corrupt parties *while they are sending their messages for a certain round*. Turned around, this corresponds to the guarantee that a party that begins to send messages is able to transmit all of the messages for the current round without interference of the adversary. This is exactly modeled by the *simultaneous multi-send* functionality $\mathcal{F}_{\text{MS}}$ defined below. For simplicity, we describe $\mathcal{F}_{\text{MS}}$ in the strict (zero-delay) eventual delivery formulation. The relaxed version of the functionality is defined analogously to $\mathcal{F}_{\text{BD-SMT}}$ and $\mathcal{F}_{\text{CLOCK}}^{-}$. We point out, that in order to model the fact that the adversary cannot replace the sender's message *after* it has been seen by some party, we slightly modify the way which the functionality deals with `replace` requests.

Functionality $\mathcal{F}_{\text{MS}}(\ell, i, \mathcal{P})$

- Upon receiving a vector of messages $\vec{m} = (m_1, \ldots, m_n)$ from p_i, record $\vec{m}$ and send a message $(\texttt{sent}, \ell(\vec{m}))$ to the adversary.
- Upon receiving $(\texttt{fetch})$ from $p_j \in \mathcal{P}$, output m_j to p_j ($m_j = \perp$ is $\vec{m}$ has not been recorded).
- *(restricted response to* `replace`*)* Upon receiving a $(\texttt{replace}, \vec{m}')$ request from the adversary for replacing p_i's input (after issuing a request for corrupting p_i), if no (honest or corrupted) p_j received m_j *before* p_i got corrupted, then replace $\vec{m}$ by $\vec{m}'$.

Lemma 9. *For each player set $\mathcal{P}$ and $p \in \mathcal{P}$, there exists a synchronous $\{[\mathcal{F}_{\text{BD-AUTH}}]^{NT}, [\mathcal{F}_{\text{COM}}]^{NT}\}$-hybrid protocol which UC securely realizes $[\mathcal{F}_{\text{MS-AUTH}}(p, \mathcal{P})]^{NT}$.*

Proof. Let π be the protocol that proceeds as follows (we assume that the designated sender $p_i \in \mathcal{P}$ is specified in the session ID, and denote the corresponding protocol machine by π_i). On input a message vector $\vec{m}$, the machine π_i initializes one copy $[\mathcal{F}_{\text{COM}}]_j^{\text{NT}}$ and issues $(\texttt{commit}, m_j)$ to $[\mathcal{F}_{\text{COM}}]_j^{\text{NT}}$ for each receiver $p_j \in \mathcal{P} \setminus \{p_i\}$. In the following activations, the machine π_i repeatedly polls the channels $[\mathcal{F}_{\text{BD-SEC}}]^{\text{NT}}$ and waits to receive a message $(\texttt{ack})$ from each p_j. Afterward, π_i proceeds by sending $(\texttt{open})$ messages to all $[\mathcal{F}_{\text{COM}}]_j^{\text{NT}}$. The protocol machines of the receiving parties p_j wait for the notification $(\texttt{receipt})$ from the functionality $[\mathcal{F}_{\text{COM}}]_j^{\text{NT}}$ and react by sending a message $(\texttt{ack})$ via the channel $[\mathcal{F}_{\text{BD-SEC}}]^{\text{NT}}$ to π_i. After receiving both the message $(\texttt{open}, m_j)$ from $[\mathcal{F}_{\text{COM}}]_j^{\text{NT}}$ and the query $(\texttt{fetch})$ on the input tape, π_j outputs m_j locally.

We describe the simulator $\mathcal{S}$ that is used to prove that π achieves the desired goals. Upon receiving $(\texttt{sent}, \vec{m})$ from $[\mathcal{F}_{\text{MS-AUTH}}]^{\text{NT}}$, $\mathcal{S}$ simulates (in each of the subsequent activations of π_i) a public delayed message $(\texttt{receipt})$

from $[\mathcal{F}_{\text{COM}}]_j^{\text{NT}}$ for one of the $p_j \neq p_i$. Whenever the environment acknowledges the delivery of such a message for $[\mathcal{F}_{\text{COM}}]_j^{\text{NT}}$, $\mathcal{S}$ simulates the leakage of the message $(\texttt{ack})$ on the channel $[\mathcal{F}_{\text{BD-AUTH}}]^{\text{NT}}$ from p_j to p_i. On every subsequent activation of p_i, the machine $\mathcal{S}$ simulates the behavior of the channels $[\mathcal{F}_{\text{BD-AUTH}}]^{\text{NT}}$ that π_i is polling, until all acknowledgments (would) have arrived—this is by simply counting the respective activations. On the following activations of p_i, $\mathcal{S}$ simulates the notifications $(\texttt{open}, m_j)$ from the functionalities $[\mathcal{F}_{\text{COM}}]_j^{\text{NT}}$, and after the environment acknowledged this opening and $\mathcal{S}$ obtained the notification for the output from $[\mathcal{F}_{\text{MS-AUTH}}]^{\text{NT}}$, $\mathcal{S}$ acknowledges this output. If the environment requests to corrupt a party, the simulator can easily construct the expected internal state since all data that does not solely represent the state of the protocol is the message vector $\vec{m}$ which is provided to $\mathcal{S}$ in the first activation.

Note that the two executions are indeed indistinguishable: The messages produced by the simulator have exactly the correct distribution (except for the opening message, all messages only represent state transitions of the protocol, and the messages m_j in the opening messages are exactly the ones obtained from $[\mathcal{F}_{\text{MS-AUTH}}]^{\text{NT}}$). Also, the restriction that $\mathcal{S}$ cannot replace the message vector if p_i was honest while sending $\vec{m}$ *and* the first message was already delivered to a p_j is consistent with the restrictions in the real execution: If p_i is corrupted *after* the first commitment is opened, all other messages to the honest receivers are already immutable in the respective commitment functionalities (by the strict ordering of events enforced by π_i and the acknowledgements sent upon receiving a $(\texttt{receipt})$) and can not be changed unless the respective receivers are corrupted. Hence, the same restrictions apply in the real model. $\square$

B Canetti's Universally Composable Synchronous Network (The $\mathcal{F}_{\text{SYN}}$-Hybrid Model)

In the latest version of the UC framework [Can05], Canetti describes a functionality that models a synchronous network. We describe this functionality $\mathcal{F}_{\text{SYN}}$ adapted to our conventions. We assume that all the parties in $\mathcal{P}$ are aware that the corresponding synchronous session has started. (If we do not want to make this assumption, we can have $\mathcal{F}_{\text{SYN}}$ send an initialization message to every player in $\mathcal{P}$.)

Functionality $\mathcal{F}_{\text{SYN}}$ $(\mathcal{P})$

Initialize a round counter $r := 1$

- Upon receiving input $(\texttt{send}, \vec{M})$ from a party $p \in \mathcal{P}$, where $\vec{M}$ is a vector of n messages (one for each party in $\mathcal{P}$), record $(p, \vec{M}, r)$ and output $(p, \vec{M}, r)$ to the adversary. (If p later becomes corrupted then the record $(p, \vec{M}, r)$ is deleted.)
- Upon receiving a message $(\texttt{Advance-Round}, \vec{N})$ from the adversary, do: If there exists $p \in \mathcal{H}$ for which no record $(p, \vec{M}, r)$ exists then ignore the message. Else:
  1. Interpret $\vec{N}$ as the list of message sent by corrupted parties in the current round r.
  2. Prepare for each $p \in \mathcal{P}$ the list L_p^r of messages that were sent to it in round r.
  3. Update $r := r + 1$.
- Upon receiving input $\texttt{receive}$ from party p, output (L_p^{r-1}, r) to p. (Let $L_p^0 = \perp$).

Lemma 11. *For every non-trivial n-party function f and for every round function Rnd there exists no $\mathcal{F}_{\text{SYN}}$-hybrid protocol which securely realizes $\mathcal{F}_{\text{SFE}}^{f,Rnd}$ with guaranteed termination.*

Proof. Let f be a non-trivial (i.e., not locally computable) function, and assume, towards contradiction that there exists a protocol π which securely realizes $\mathcal{F}_{\text{SFE}}^{f,Rnd}$ in the $\mathcal{F}_{\text{SYN}}$-hybrid model. As f is not locally computable, there are input vectors $\vec{x}$ and $\vec{x'}$ and a party $p_i \in \mathcal{P}$ such that $x_i = x'_i$ but $(f(\vec{x}, R))_i \not\approx (f(\vec{x'}, R))_i$. Consider the following environment: $\mathcal{Z}$ corrupts no party, tosses a coin and hands inputs either $\vec{x}$ or $\vec{x'}$ to the parties. Subsequently, $\mathcal{Z}$ activates the parties $Rnd \cdot (|\mathcal{P}| + 1)$ times in a round-robin fashion using $(\texttt{output})$-queries. Throughout this computation, $\mathcal{Z}$ ignores all activations from the adversary and never activates the adversary, but $\mathcal{Z}$ records the values y_i provided as output by the honest parties in the last round of the computation.

In the ideal world execution (i.e., in the $\mathcal{F}_{\text{SFE}}^{f,Rnd}$-hybrid model with some simulator $\mathcal{S}$), by the definition of $\mathcal{F}_{\text{SFE}}^{f,Rnd}$, this behavior will result in an output vector $\vec{y}$ that is distributed as either $f(\vec{x}, R)$ or $f(\vec{x'}, R)$, depending on the inputs provided by the environment. In contrast, in the $\mathcal{F}_{\text{SYN}}$-hybrid world with the dummy adversary $\mathcal{A}$, the network $\mathcal{F}_{\text{SYN}}$ will not allow the parties to exchange any message (as the round does not advance). As a result, the distribution of the output y_i is the same independent of whether the input provided by $\mathcal{Z}$ is $\vec{x}$ or $\vec{x'}$ (as all messages received by the protocol machine of party p_i are exactly the same). $\mathcal{Z}$ uses the distinguisher guaranteed for $(f(\vec{x}, R))_i$ and $(f(\vec{x'}, R))_i$ and obtains half the distinguishing advantage (which is still non-negligible). $\square$

C Transferring Statements from the Timing Model

The goal of the "Timing Model" [KLP05] is to examine the guarantees that can be used by cryptographic protocols if certain guaranteed concerning time are assumed. The respective guarantees considered by the model are two-fold:

Bounded clock drift: The parties have clocks that are, for some global parameter $\epsilon \geq 1$, *ϵ-drift preserving*: When some local clock advanced by time δ, all other local clocks must have advanced by time δ' with $\delta/\epsilon < \delta' < \epsilon \cdot \delta$.

Maximum latency: There is an upper bound Δ on the time it takes to compute, send, and deliver messages between parties via the assumed channels.

To be able to compare our security model to the Timing Model from [KLP05], we provide an analogous extension of the protocol machines. The model of [KLP05] defines the protocol machines to be ITMCs—interactive Turing machines with clocks. These protocol machines have a specific *clock tape* that can be written to by the adversary, under the restriction that the values are advanced according to the *ϵ-drift preserving* property. Unlike for other tapes, an ITM is not activated when the adversary writes to the clock tape. We model this extension of the ITMs as an ideal functionality $\mathcal{F}_{\text{TIME}}$ that is available to the protocol machines, where instead of accessing the local clock tape, the protocol machines access the ideal functionality $\mathcal{F}_{\text{TIME}}$. As in [KLP05], the adversary fulfills the task of actually advancing the value of the clock, where the functionality $\mathcal{F}_{\text{TIME}}$ ensures that the ϵ-drift preserving property is preserved.

Functionality $\mathcal{F}_{\text{TIME}}(\mathcal{P}, \epsilon)$

For every player $p_i \in \mathcal{P}$, the functionality $\mathcal{F}_{\text{TIME}}$ maintains an integer t_i (p_i's current time) and two bits $a_i, s_i \in \{0, 1\}$. For initialization, $\mathcal{F}_{\text{TIME}}$ sets $a_i := 0$ and $s_i := 1$ and expects to receive a vector $\vec{T}$ of $|\mathcal{P}|$ integers from $\mathcal{A}$ and $t_i := \vec{T}_i$ (otherwise $t_i = 0$ for all $p_i \in \mathcal{P}$).

- Upon receiving $(\texttt{deactivate}, i)$ from $\mathcal{A}$, set $a_i := 0$ and $s_i := 0$, and send $(\texttt{ok})$ to $\mathcal{A}$.
- Upon receiving $(\texttt{set}, (\rho_i)_{i \in \mathcal{P}})$ from $\mathcal{A}$, check whether $a_i = 0$ for all $i \in \mathcal{P}$ and whether setting t_i to ρ_i for all $i \in \mathcal{P}$ preserves the property of being ϵ-drift preserving. If this is the case, set $t_i := \rho_i$ and $s_i := 1$ for all $i \in \mathcal{P}$. Send $(\texttt{ok})$ to $\mathcal{A}$.
- Upon receiving $(\texttt{activate}, i)$ from $\mathcal{A}$, if $s_i = 1$ then set $a_i := 1$. Send $(\texttt{ok})$ to $\mathcal{A}$.
- Upon receiving $(\texttt{time})$ from p_i, if $a_i = 1$ then send $(\texttt{time}, t_i)$ to p_i, otherwise send $(\texttt{inactive})$ to p_i.

The different scheduling schemes used in [Can05] and [KLP05] impose some further technical details on $\mathcal{F}_{\text{TIME}}$: While in [Can05], the "master scheduler" activated by default is defined to be the environment $\mathcal{Z}$, this task is fulfilled by the adversary $\mathcal{A}$ in [KLP05]. To be able to show the relation between security statements in the two models, we have to define the functionality $\mathcal{F}_{\text{TIME}}$ such that it "protects" parties from being activated without the acknowledgment of the adversary.

Translating functionalities. As for the different definitions of protocol executions, there are several differences between the definitions of ITMs in [KLP05] and those in [Can05]. In particular, ITMs in [KLP05] have an explicit outgoing communication tape, on which all messages (both "communication" messages intended for communication over the network and "ideal" messages directed to an ideal functionality) are written (albeit the adversary cannot access the contents of the ideal messages, but merely sees their length). ITMs can write an (a priori) arbitrary number of messages for arbitrary receivers to their outgoing communication tape in a single activation. In [Can05], in contrast, communication is defined as writing directly to the respective input tape of the receiving ITI. This implies that ideal messages are not communicated via $\mathcal{A}$, and that, in a single activation, messages can only be sent to a single receiver. For ideal functionalities, the differences that must be leveraged by the transformation T are: For inputs to and outputs from parties, the lengths of the respective messages are leaked to $\mathcal{A}$. Also, $\mathcal{A}$ may schedule the delivery of any such "ideal" message.

We define the transformation mapping T as follows. The functionality $T(\mathcal{F})$, upon receiving an input message x from any honest party p_i, sends a notification $(\texttt{input}, p_i, |x|)$ to the adversary. After the adversary has acknowledged this input, $T(\mathcal{F})$ simulates x as an input to $\mathcal{F}$. All outputs y of $\mathcal{F}$ to a player p_i are sent by $T(\mathcal{F})$ as (private) delayed outputs,[16] where additionally the length $|y|$ is leaked.

Translating protocols. The protocol compiler C translates the ITMCs (interactive Turing machines with clocks) from [KLP05] to $\mathcal{F}_{\text{TIME}}$-hybrid ITMs. At the beginning of each activation (except for replies from $\mathcal{F}_{\text{TIME}}$), $C(\pi)$ sends a message $(\texttt{time})$ to the functionality $\mathcal{F}_{\text{TIME}}$ and writes the returned time value to the simulated clock tape of π. If $\mathcal{F}_{\text{TIME}}$ replies with $(\texttt{inactive})$, then $C(\pi)$ buffers the input (if necessary) and halts. The protocol compiler also leverages for the different semantics of communication. This means that the $\mathcal{F}_{\text{AUTH}}$-hybrid protocol $C(\pi)$, instead of writing network messages to the outgoing communication tape, inserts them into the authenticated channel $\mathcal{F}_{\text{AUTH}}(p_i, p_j)$ that p_i shares with the receiver p_j. After the initial input is obtained, $C(\pi)$ ignores all further activations via the input tape.[17]

Proof of Lemma 14. Following [Can05, Claim 12], it is sufficient to show the that a protocol is secure in the Timing Model if and only if it is secure in UC with respect to specialized simulators. We describe how the adversarial interfaces in Timing Model and in the $\{\mathcal{F}_{\text{TIME}}, \mathcal{F}_{\text{AUTH}}\}$-hybrid model of UC can be translated. Adversaries or simulators can be translated between the Timing Model and UC by applying the suitable conversion strategies at both the interface to the "execution" and the interface to the environment.

The interface to the real execution. In the execution in the Timing Model, the adversary can communicate with the protocol machines π_i (by sending a message "earmarked" for π_i to p_i, respectively receiving a message from π_i via p_i), and can write (valid) time-values to the clock tapes of the p_i. In UC, the adversary can communicate with the protocol machines $C(\pi)_i$, the authenticated channels $\mathcal{F}_{\text{AUTH}}(p_i, p_j)$, and the time functionality $\mathcal{F}_{\text{TIME}}$ via the respective communication tapes. The actions in the Timing Model translate as follows.

Updating the clock tapes of $(\pi_i)_{i\in\mathcal{P}}$ to $(\rho_i)_{i\in\mathcal{P}}$: sending $(\texttt{deactivate}, i)$ to $\mathcal{F}_{\text{TIME}}$ for all $i \in \mathcal{P}$, sending a message $(\texttt{set}, (\rho_i)_{i\in\mathcal{P}})$ to $\mathcal{F}_{\text{TIME}}$, and activating all p_i by sending $(\texttt{activate}, i)$ to $\mathcal{F}_{\text{TIME}}$.
Obtaining messages sent by π_i: obtaining the messages one-by-one by activating $C(\pi)_i$ until no further message is sent. Messages for π_j correspond to $(\texttt{leak}, m)$ from $\mathcal{F}_{\text{AUTH}}(p_i, p_j)$.
Delivering message from π_i to π_j: issuing $(\texttt{deliver}, m)$ to $\mathcal{F}_{\text{AUTH}}(p_i, p_j)$.
Sending a message for a corrupt π_i: inserting the message into the corresponding $\mathcal{F}_{\text{AUTH}}(p_i, p_j)$.

Note that, in between two activations of players, the Timing Model-adversary must update either the clock tapes of all parties or none at all. This stems from the definition of the ϵ-drift preserving property. In UC, this is captured by the functionality $\mathcal{F}_{\text{TIME}}$ that takes a vector of values for all parties.

The interface to the ideal execution. In the execution in the Timing Model, the adversary can communicate with the ideal functionality $\mathcal{F}$ using the communication tapes of $\mathcal{F}$. Also, the adversary can forward ideal messages from $\mathcal{F}$ to p_i and vice versa by copying them to the corresponding tape. In UC, the adversary can communicate with $T(\mathcal{F})$ via communication tapes.

Obtaining ideal messages from $\mathcal{F}$ for the ψ_i: obtaining notifications (including the message length) for the delayed output from $T(\mathcal{F})$.
Obtaining ideal messages from ψ_i for $\mathcal{F}$: obtaining notifications $(\texttt{input}, p_i, |x|)$ from $T(\mathcal{F})$.
Delivering output from $\mathcal{F}$: acknowledging the delivery of the delayed output at $T(\mathcal{F})$.
Delivering input to $\mathcal{F}$: acknowledging the delivery of the delayed input at $T(\mathcal{F})$.
Communication with $\mathcal{F}$: communication with $T(\mathcal{F})$.

[16] A notification is sent to $\mathcal{A}$, who may later decide to deliver the message. For details, see [Can05].

[17] In contrast to the Timing Model, protocols in the model of [Can05] can only send a single message within an activation. On the other hand, this property of the Timing Model leaks some additional information to the adversary, namely the information that the messages on the outgoing communication tape are all messages to be sent prior to the next input. In order to equip the UC-adversary with the same power, we define that the protocol π sends an empty message to the adversary if it is activated for sending but there are no pending messages.

The interface to the environment. In UC, the adversary communicates with the environment directly via its local input and the environment's subroutine output tape. In the Timing Model, the role of the environment is taken by the malicious protocol ψ, which communicates with the adversary using the generic communication tapes. The translation is straightforward: A message from or to $\mathcal{Z}$ corresponds directly to a message from or to some protocol machine running ψ. Note that, in UC, the environment is the master scheduler, whereas this task is fulfilled by the adversary in the Timing Model. This difference must be accounted for by carefully designing input/output behavior of all involved systems.

Security in the Timing Model implies security in UC. To show security in UC with respect to specialized simulators, we have to show that for each adversary $\mathcal{A}$ and environment $\mathcal{Z}$, there is a simulator $\mathcal{S}$ such that the outcomes of the real and the ideal execution are indistinguishable. We convert $\mathcal{A}$ into a Timing Model-adversary $\mathcal{A}'$ and $\mathcal{Z}$ into a malicious protocol $\psi^{\mathcal{Z}}$ and conclude that there is a good Timing Model-simulator $\mathcal{S}'$. From this simulator, we construct a UC-simulator $\mathcal{S}$ and show that the output of $\mathcal{Z}$ and the Timing Model-distinguisher D are the same for both the real and the ideal executions.

The adversary $\mathcal{A}'$ is constructed from the adversary $\mathcal{A}$ by the above described interface transformations. We detail the translation of $\mathcal{A}$'s clock tape handling to $\mathcal{A}'$: The definition of the adversarial interface of $\mathcal{F}_{\text{TIME}}$ guarantees that the values of the clock tapes of all parties advance at the same time, so $\mathcal{A}'$ can write the values $(\rho_i)_{i\in\mathcal{P}}$ to the clock tapes of $(\pi_i)_{i\in\mathcal{P}}$ after seeing the $(\texttt{set}, (\rho_i)_{i\in\mathcal{P}})$ message from $\mathcal{A}$ between two activations of parties (this is necessary for $\mathcal{A}'$ to be ϵ-drift preserving). Also, the fact that $\mathcal{A}'$ controls the activations of the π_i (by means of the $\psi_i^{\mathcal{Z}}$, see below) allows $\mathcal{A}'$ to only activate p_i if reading the time tape would succeed for the protocol $C(\pi_i)$ and can otherwise activate protocol machine $\psi_0^{\mathcal{Z}}$.

The distinguisher $\mathcal{Z}$ is converted into a malicious protocol $\psi^{\mathcal{Z}}$ as follows. We define $\psi_0^{\mathcal{Z}}$ to run the ITM $\mathcal{Z}$ and relay all messages intended for $\mathcal{A}$ to $\mathcal{A}'$, and mask the messages from $\mathcal{A}'$ as messages from $\mathcal{A}$. Inputs to and outputs from $C(\pi)_0$ can be processed locally by $\psi_0^{\mathcal{Z}}$, inputs for $C(\pi)_i$ with $i \neq 0$ are sent via $\mathcal{A}'$ to the respective $\psi_i^{\mathcal{Z}}$, which only acts as a forwarder between $C(\pi)_i$ and $\psi_0^{\mathcal{Z}}$. The protocol machine $\psi_0^{\mathcal{Z}}$ allows the adversary to schedule the input to π_0. When $\mathcal{Z}$ generates local output, $\psi_0^{\mathcal{Z}}$ outputs the same message (the local output of $\psi_i^{\mathcal{Z}}$ for $i \neq 0$ is constant).

By the fact that π is secure in the Timing Model, we know that there is a good simulator $\mathcal{S}'$. From this simulator, we construct a simulator $\mathcal{S}$ in UC using the translations described above.

Finally, we can track the messages sent (by $\mathcal{Z}$) in both executions and verify that the messages returning to $\mathcal{Z}$ are computed using the same functions on identically distributed inputs in the execution in UC and the execution within $\psi_0^{\mathcal{Z}}$ in the Timing Model, in both the real and the ideal cases. Note that, in UC, the environment $\mathcal{Z}$ is the master scheduler that is also activated in the beginning of the execution, whereas this task is fulfilled by the adversary $\mathcal{A}'$ in the Timing Model.

In more detail, we have to argue that, for each activation, the inputs and the state of the ITMs π_i, $\mathcal{A}$, and $\mathcal{Z}$ have the same distribution in the executions $\text{EXEC}_{C(\pi),\{\mathcal{F}_{\text{TIME}},\mathcal{F}_{\text{AUTH}}\},\mathcal{A},\mathcal{Z}}$ in UC and $\text{EXEC}_{\pi,\mathcal{A}',\psi^{\mathcal{Z}}}$ in the Timing Model. We can then conclude the same for the outputs by an inductive argument. Note that we also have to argue that the clock tape of π_i has the same contents. The following types of activations occur in the real execution:

Local input at π_i: π_i computes based on the input from $\mathcal{Z}$ and the clock tape, which is guaranteed to have consistent contents by potentially buffering the input (either within $C(\pi)$ or within $\psi^{\mathcal{Z}}$ by the fact that $\psi_i^{\mathcal{Z}}$ awaits the acknowledgment of $\mathcal{A}'$). Moreover, it is easy for $\mathcal{A}'$ to track which π_i can be activated and what the contents of their clock tapes must be like.

Receive message at π_i: The (new) values considered by π_i are the contents of the message and the clock tape. Again, the consistency of the clock tapes is guaranteed by providing the messages to π_i only if p_i is marked as active in $\mathcal{F}_{\text{TIME}}$.

Leaked message from $\mathcal{F}_{\text{AUTH}}$ at $\mathcal{A}$: The contents of the messages is generated by π_i and the "envelopes" can be easily simulated by $\mathcal{A}'$.

Reply from $\mathcal{F}_{\text{TIME}}$ to $\mathcal{A}$: The messages are constant (the purpose is to return the activation).

Communication between $\mathcal{Z}$ at $\mathcal{A}$: The correct "forwarding" of these messages is guaranteed by the "protocol" used among $\psi^{\mathcal{Z}}$ and $\mathcal{A}'$.

Empty activation from $C(\pi_i)$ at $\mathcal{A}$: This can easily be simulated as $\mathcal{A}'$ keeps track of the messages that are otherwise buffered by $C(\pi_i)$.

Output of π_i at $\mathcal{Z}$: The output is simply forwarded by $\psi_i^{\mathcal{Z}}$.

Of course, we must also provide the same arguments for the ideal executions:

Inputs to $\mathcal{F}$ (from ψ_i/$\mathcal{Z}$): The input notifications of $T(\mathcal{F})$ are wrapped as messages among the $\psi_i^{\mathcal{Z}}$ by $\mathcal{S}$, and the input is acknowledged by $\mathcal{S}$ as soon as $\mathcal{S}'$ decides to deliver the messages (among the $\psi_i^{\mathcal{Z}}$ and the ideal input message to $\mathcal{F}$). This guarantees that $\mathcal{F}$ will obtain the same inputs in the same order in both cases.

Outputs of $\mathcal{F}$ to ψ_i/$\mathcal{Z}$: The notifications for delayed outputs from $T(\mathcal{F})$ are collected by $\mathcal{S}$ and simulated as ideal messages from $\mathcal{F}$ to ψ_i. $\mathcal{S}$ asks $T(\mathcal{F})$ to deliver the outputs once $\mathcal{S}'$ delivers the ideal messages to $\psi_i^{\mathcal{Z}}$ and the (forwarded) message to $\psi_0^{\mathcal{Z}}$. This guarantees that $\mathcal{Z}$ obtains the outputs in the correct order.

Communication between $\mathcal{F}$ and $\mathcal{S}'$: The communication is simply forwarded by the definition of $T(\mathcal{F})$ and $\mathcal{S}$.

Messages among $\mathcal{S}'$ and ψ_i/$\mathcal{Z}$: Messages from $\mathcal{Z}$ are "wrapped" as messages from ψ_i by $\mathcal{S}$ before handing them to $\mathcal{S}'$. This means that messages from $\mathcal{Z}$ to $\mathcal{A}$ are wrapped as special messages from ψ_0 (as done by $\psi_0^{\mathcal{Z}}$), and the input notifications from $T(\mathcal{F})$ are wrapped as described above. On the other hand, $\mathcal{S}$ also "unwraps" the messages originating from $\mathcal{S}'$ before handing them to $\mathcal{Z}$. The consistent behavior of $\mathcal{S}$ and $\psi_0^{\mathcal{Z}}$ guarantees that communication between $\mathcal{S}'$ and $\mathcal{Z}$ works identically in both settings.

The input to the UC-execution is one auxiliary input string for $\mathcal{Z}$, whereas in the Timing Model, the input is one such string for each party p_i, and one for the adversary $\mathcal{A}'$. Hence, the auxiliary input for $\mathcal{Z}$ is provided to p_0, while the inputs to all other parties as well as the adversary are empty. Note that all parties p_i with $i \neq 0$ provide empty output, while p_0 forwards the output of $\mathcal{Z}$. As $\mathcal{Z}$'s "view" is identical in both the ideal executions in the Timing Model and UC and in the real executions in the Timing Model and UC, the distinguishing advantage of $\mathcal{D}'$ simply forwarding $\mathcal{Z}$'s decision in the Timing Model is at least as large as the advantage of $\mathcal{Z}$ in UC.

Note that security in the Timing Model is defined as

$$\forall \psi, \mathcal{A}' \; \exists \mathcal{S}' : \left\{ \mathrm{EXEC}_{\pi,\mathcal{A}',\psi}(k,x,z) \right\}_{k\in\mathbb{N},x\in(\{0,1\}^*)^m,z\in\{0,1\}^*} \overset{c}{\approx} \left\{ \mathrm{EXEC}_{\mathcal{F},\mathcal{S}',\psi}(k,x,z) \right\}_{k\in\mathbb{N},x\in(\{0,1\}^*)^m,z\in\{0,1\}^*}.$$

By the above constructions, we can transform any pair $\mathcal{Z}, \mathcal{A}$ of a UC-environment and a UC-adversary into such a pair $\psi^{\mathcal{Z}}, \mathcal{A}'$. The simulator $\mathcal{S}'$ guaranteed by the above definition can be converted into a UC-simulator $\mathcal{S}$, such that the (computational) distance of the output vectors of the two Timing Model experiments is as large as the (statistical) distance of the outcomes of the two UC experiments, meaning that we have

$$\forall \mathcal{Z}, \mathcal{A} \; \exists \mathcal{S} : \left\{ \mathrm{EXEC}_{C(\pi),\{\mathcal{F}_{\text{TIME}},\mathcal{F}_{\text{AUTH}}\},\mathcal{A},\mathcal{Z}}(k,x) \right\}_{k\in\mathbb{N},x\in\{0,1\}^*} \approx \left\{ \mathrm{EXEC}_{T(\mathcal{F}),\mathcal{S},\mathcal{Z}}(k,x) \right\}_{k\in\mathbb{N},x\in\{0,1\}^*},$$

which is equivalent to the original definition of UC by [Can05, Claim 12].

Security in UC implies security in the Timing Model. We employ the above described interface conversion strategy to convert the Timing Model-adversary $\mathcal{A}'$, malicious protocol ψ, and distinguisher $\mathcal{D}'$ into a UC-adversary $\mathcal{A}$ and environment $\mathcal{Z}^{\mathcal{D}',\psi}$. The adversary $\mathcal{A}$ notifies the environment $\mathcal{Z}^{\mathcal{D}',\psi}$ upon the first activation originating from $\mathcal{A}'$ for each party ψ_i such that $\mathcal{Z}^{\mathcal{D}',\psi}$ can take care of providing ψ_i with the correct input. The environment $\mathcal{Z}^{\mathcal{D}',\psi}$ emulates ψ_i for all p_i, forwards the local input to the $C(\pi)_i$ and the messages of the ψ_i to the adversary. Finally, $\mathcal{Z}$ tailors a transcript of the corresponding Timing Model-execution and applies $\mathcal{D}'$ to this transcript to determine the final output. Also, we employ the inverse conversion to transform the guaranteed UC-simulator $\mathcal{S}$ into a Timing Model-simulator $\mathcal{S}'$.

Translating the handling of the clock tapes is straightforward: Once $\mathcal{A}'$ has finished the writing operation (i.e., proceeds with the next activation), $\mathcal{A}$ deactivates all parties at $\mathcal{F}_{\text{TIME}}$, submits a vector of new clock values, and re-activates all parties.

The remainder of the argument is as above: The distribution of the inputs obtained and outputs generated by the ITMs π_i, ψ_i, and $\mathcal{A}'$ is the same in the execution in the Timing Model and the one in UC. The transcript for $\mathcal{D}'$ is obtained in a deterministic way (and also has the same distribution). As the "decision" taken by $\mathcal{Z}$ is the same that is done by $\mathcal{D}'$, the distance is the same as well.

More detailed, what are the actions that can happen within the system:

Local input to π_i: The ITM π_i computes based on this local input (which is easily forwarded from the simulated ψ_i by $\mathcal{Z}$) and the contents of the clock tape (which is kept in a consistent state by the above described strategy).

Local input to ψ_i: This is derived from $\mathcal{Z}^{\mathcal{D}',\psi}$'s auxiliary input and is provided by $\mathcal{Z}^{\mathcal{D}',\psi}$ at the correct point in time (according to the clock tape of ψ_i), which is ensured by the collaboration of $\mathcal{A}$ and $\mathcal{Z}^{\mathcal{D}',\psi}$, once $\mathcal{A}'$ schedules p_i for the first time.

Message delivered for π_i, ψ_i, or local output from π_i to ψ_i: The consistency of both these messages and the contents of p_i's clock tape is argued as above. In the Timing Model, the messages among the ψ_i are scheduled and delivered via $\mathcal{A}'$. Hence, $\mathcal{Z}^{\mathcal{D},\psi}$ (and $\mathcal{A}$) forward all messages among the simulated ψ_i via $\mathcal{A}'$ to guarantee the correct scheduling.

Local output of ψ_i: This output is faithfully included in the transcript by $\mathcal{Z}^{\mathcal{D}',\psi}$.

Messages obtained by $\mathcal{A}'$: The messages generated by the (simulated) ψ_i are easily obtained and forwarded to $\mathcal{A}$ by $\mathcal{Z}^{\mathcal{D}',\psi}$, the messages generated by the π_i can be extracted from $C(\pi_i)$ by the above described scheme.

Of course, we have to provide a similar analysis for the ideal setting (the involved protagonists are $\mathcal{F}$, ψ, $\mathcal{S}$, and $\mathcal{D}'$):

Input from ψ_i to $\mathcal{F}$: The protocol machine ψ_i sends its input to $\mathcal{F}$ by means of an ideal message. In $\mathcal{Z}^{\mathcal{D}',\psi}$, this is translated into an input to $T(\mathcal{F})$ at interface i, which yields an input notification to $\mathcal{S}$. Once $\mathcal{S}$ acknowledges this input to $\mathcal{F}$, $\mathcal{S}'$ will deliver the ideal message to $\mathcal{F}$.

Communication between $\mathcal{F}$ and $\mathcal{S}$: The conversion between the two different communication modes is done identically by $T(\mathcal{F})$ and $\mathcal{S}'$.

Local input to ψ_i: As in the real case.

Message delivered for ψ_i: In the execution in the Timing Model, the communication among the ψ_i is scheduled by $\mathcal{S}'$. In UC, $\mathcal{Z}^{\mathcal{D}',\psi}$ simulates the protocol machines ψ_i, providing the messages to $\mathcal{A}$ to obtain the scheduling via $\mathcal{A}'$. The simulator $\mathcal{S}'$ behaves consistently: It provides $\mathcal{S}$ with the notifications that would have been provided by $\mathcal{Z}^{\mathcal{D}',\psi}$ and also acts according to $\mathcal{S}$'s replies.

Output from $\mathcal{F}$ to ψ_i: The messages from $\mathcal{F}$ to ψ_i are ideal messages that are scheduled by $\mathcal{S}'$. In UC, the functionality $T(\mathcal{F})$ first sends the notification for the delayed output to $\mathcal{S}$ (which has to acknowledge the delivery). Hence, $\mathcal{S}'$ can deliver the message once $\mathcal{S}$ sends the acknowledgment.

As the outputs of the simulated ψ are used to construct the transcript, the input to $\mathcal{D}'$ has the same distribution in both cases.

The input to the Timing Model-execution consists of one auxiliary string for each p_i along with one such string for the adversary $\mathcal{A}'$. These inputs can all be encoded into one single auxiliary string provided to $\mathcal{Z}^{\mathcal{D}',\psi}$ in UC, which recovers the original contents and provides it as an input to the ψ_i and (via $\mathcal{A}$) to $\mathcal{A}'$.

Note that security in UC is defined as

$$\forall \mathcal{A}\ \exists \mathcal{S}\ \forall \mathcal{Z} : \left\{ \text{EXEC}_{C(\pi),\{\mathcal{F}_{\text{TIME}},\mathcal{F}_{\text{AUTH}}\},\mathcal{A},\mathcal{Z}}(k,x) \right\}_{k\in\mathbb{N},x\in\{0,1\}^*} \approx \left\{ \text{EXEC}_{T(\mathcal{F}),\mathcal{S},\mathcal{Z}}(k,x) \right\}_{k\in\mathbb{N},x\in\{0,1\}^*}.$$

Our goal is to conclude security in the Timing Model, which is

$$\forall \psi, \mathcal{A}'\ \exists \mathcal{S}' : \left\{ \text{EXEC}_{\pi,\mathcal{A}',\psi}(k,x,z) \right\}_{k\in\mathbb{N},x\in(\{0,1\}^*)^m,z\in\{0,1\}^*} \stackrel{c}{\approx} \left\{ \text{EXEC}_{\mathcal{F},\mathcal{S}',\psi}(k,x,z) \right\}_{k\in\mathbb{N},x\in(\{0,1\}^*)^m,z\in\{0,1\}^*}.$$

Starting from a Timing Model-adversary $\mathcal{A}'$ (and a protocol ψ), we construct a UC-adversary $\mathcal{A}$, and transform the guaranteed UC-simulator $\mathcal{S}$ into a Timing Model-simulator $\mathcal{S}'$. From ψ and the best distinguisher $\mathcal{D}'$ for Timing Model-transcripts, we construct a UC-environment that has (by construction) distinguishing advantage at least as good as $\mathcal{D}'$. Hence, by the fact that the protocol is UC-secure, we can conclude that it is also secure in the Timing Model. □

After having proven that security statements (regarding statements about universal or general composability) from the Timing Model can be faithfully translated into the $\{\mathcal{F}_{\text{TIME}}, \mathcal{F}_{\text{AUTH}}\}$-hybrid model of UC, we can show that the clocks provided to the parties do not provide them with further guarantees. More formally, we show that the functionality $\mathcal{F}_{\text{TIME}}$ can be implemented only from the channels $\mathcal{F}_{\text{AUTH}}$. This is consistent with [KLP05, Theorem 13], which shows that it is impossible to implement arbitrary functionalities in the Timing model, if one considers general composition.

Proof of Lemma 13. The protocol τ implements $\mathcal{F}_{\text{TIME}}$ in the $\mathcal{F}_{\text{AUTH}}$-hybrid model.

Protocol $\tau(\mathcal{P})$

For (each) player p_i, initialize variables $t_j^{(i)} := 0$ and $a_j^{(i)} := 0$ for all $j \in \mathcal{P}$. The protocol assumes that p_i shares a channel $\mathcal{F}_{\text{AUTH}}(p_i, p_j)$ with each p_j for $j \in \mathcal{P}$. If several messages are to be sent in a single activation, they are stored in a buffer and sent during the next activations.

- Upon a $(\texttt{set-inactive}, \rho_j)$ message from p_j, set $a_j^{(i)} := 0$.
  - If $a_i^{(i)} = 1$ and $t_i^{(i)} \leq t_\ell^{(i)}$ for all $\ell \in \mathcal{P}$, then become inactive.
  - If for all $\ell \in \mathcal{P} : ((a_\ell^{(i)} = 0 \ \wedge \ t_\ell^{(i)} \geq t_i^{(i)}) \ \vee \ t_\ell^{(i)} > t_i^{(i)})$, then set $t_i^{(i)} := t_i^{(i)} + 1$ and $a_i^{(i)} := 1$, and send $(\texttt{set-active}, t_i^{(i)})$ to all $\ell \in \mathcal{P} \setminus \{i\}$.
- Upon a $(\texttt{set-active}, \rho_j)$ message from p_j with $\rho_j > t_j^{(i)}$, set $a_j^{(i)} := 1$ and $t_j^{(i)} := \rho_j$.
- Upon local input $(\texttt{time})$, if $a_i^{(i)} = 0$ then output $(\texttt{inactive})$, otherwise output $(\texttt{time}, t_i^{(i)})$.
- Upon any activation via the communication tape while no outgoing messages are pending and $a_j^{(i)} = 1$ for all $j \in \mathcal{P}$, become inactive.
- Becoming inactive: Set $a_i^{(i)} := 0$ and send a message $(\texttt{set-inactive}, t_i^{(i)})$ to all $j \in \mathcal{P} \setminus \{i\}$.

Messages that do not match any of the rules are ignored.

The simulator $\mathcal{S}$ for the protocol τ and the dummy adversary $\mathcal{D}$ proceeds as follows. All messages between $\mathcal{Z}$ and (the simulated) $\mathcal{D}$ are simply forwarded. $\mathcal{S}$ also simulates protocol machines τ_i running protocol τ for each player p_i, and maintains buffers for the (simulated) functionalities $\mathcal{F}_{\text{AUTH}}(p_i, p_j)$. During the execution, $\mathcal{S}$ behaves as follows:

- For a message m sent among (the simulated) τ_i and τ_j, simulate $(\texttt{leak}, m)$ from $\mathcal{F}_{\text{AUTH}}(p_i, p_j)$ to $\mathcal{D}$.
- If $\mathcal{D}$ activates τ_j or issues a message $(\texttt{deliver})$ to a (non-empty) channel $\mathcal{F}_{\text{AUTH}}(p_i, p_j)$, run the simulated τ_j with the respective input, and:
  - if τ_j changes $a_j^{(j)}$ from 1 to 0, then send $(\texttt{deactivate}, j)$ to $\mathcal{F}_{\text{TIME}}$. If τ_j was the last (honest) protocol machine to make this transition *for the current round value*, then increase the clock counter for all $i \in \mathcal{P}$ using the $(\texttt{set}, (\rho_i)_{i\in\mathcal{P}})$ message to $\mathcal{F}_{\text{TIME}}$.
  - if τ_j changes $a_j^{(j)}$ from 0 to 1, then send $(\texttt{activate}, j)$ to $\mathcal{F}_{\text{TIME}}$.

The described simulator is perfect. In particular, we show that any query by $\mathcal{Z}$ will lead to the same state transitions and replies in the two executions.

There are three man arguments that build the proof: First, the transitions of the variables $a_j^{(i)}$ and $t_j^{(i)}$ within the players in the real execution and the simulated protocol machines in the ideal execution is consistent for every query of $\mathcal{Z}$. Second, throughout the complete ideal execution, the invariant $a_i = a_i^{(i)}$, $a_i = 1 \Rightarrow t_i = t_i^{(i)}$, $(a_i = 0 \ \wedge s_i = 0) \Rightarrow t_i = t_i^{(i)}$, $(a_i = 0 \ \wedge s_i = 1) \Rightarrow t_i = t_i^{(i)} + 1$, and $t_\ell^{(j)} \leq t_\ell$ is always preserved, where a_i, t_i are the values in $\mathcal{F}_{\text{TIME}}$ and $a_i^{(j)}$, $t_i^{(j)}$ are the values within the simulated protocol machine τ_j. Third, the responses of the real and ideal model on the same queries are always identical, given that the state is consistent. (The fact that the protocol is deterministic simplifies some arguments.)

The environment can initiate the following queries at its interfaces (and those of $\mathcal{D}$):

Query $(\texttt{time})$ **at** τ_j**:** In both cases, the response is $(\texttt{inactive})$ if $a_j = 0$ and $(\texttt{time}, t_j)$ if $a_j = 1$ (consistently for the values $a_j^{(j)}$ and $t_j^{(j)}$).

Query $(\texttt{deliver})$ **at** $\mathcal{F}_{\text{AUTH}}(p_i, p_j)$**:** If the buffer has the same contents in both cases prior to this query, then the contents is the same after this query. Also, if the buffer is empty, then $\mathcal{Z}$ is activated in both cases.

If the first message in the buffer is a $(\texttt{set-active}, \rho_i)$ message, then both the real and the simulated τ_j will update their internal view of τ_i's state, but will not change their own state.

If the first message in the buffer is a $(\texttt{set-inactive}, \rho_j)$ message, then the state transition of the real and the ideal τ_j will also be the same. We have to argue that the effect on the behavior of $\mathcal{F}_{\text{TIME}}$ coincides with the behavior of τ_j for $(\texttt{time})$-queries.

$a_j^{(j)} = 1 \rightsquigarrow a_j^{(j)} = 0$**:** $\mathcal{S}$ sends $(\texttt{deactivate}, j)$ to $\mathcal{F}_{\text{TIME}}$, which sets in $a_j = 0$ and $s_j = 0$. Since neither t_j nor $t_j^{(j)}$ changed, the invariant is preserved.

If $a_\ell = 0$ and $s_\ell = 0$ for all $\ell \in \mathcal{H}$, then the invariant implies that $t_\ell = t_\ell^{(\ell)}$. Also, $\mathcal{S}$ (which can easily keep track of this) will advance $t_\ell := t_\ell + 1$ for all $\ell \in \mathcal{P}$ by sending $\big(\texttt{set}, (t_\ell + 1)_{\ell \in \mathcal{P}}\big)$, which will also set $s_\ell := 1$ for all $\ell \in \mathcal{P}$. Hence, after the switch, we have $a_\ell = 0$, $s_\ell = 1$, and $t_\ell = t_\ell^{(\ell)} + 1$ for all $\ell \in \mathcal{H}$, so the invariant is preserved.

$a_j^{(j)} = 0 \rightsquigarrow a_j^{(j)} = 1$**:** $\mathcal{S}$ will send $(\texttt{activate}, j)$ to $\mathcal{F}_{\text{TIME}}$, which induces $a_j := 1$.

Note that τ_j only switches if $a_\ell^{(j)} = 0$ or $t_\ell^{(j)} > t_j^{(j)}$ for all $\ell \in \mathcal{P}$. If there is an $\ell \in \mathcal{H}$ with $a_\ell^{(j)} = 1$, then $t_j = t_\ell \geq t_\ell^{(j)} > t_j^{(j)}$ and hence $s_j = 1$ and $t_j = t_j^{(j)} + 1$. Otherwise, if $a_\ell^{(j)} = 0$ for all $\ell \in \mathcal{H}$, then we must have $t_\ell \geq t_\ell^{(j)} \geq t_j^{(j)}$, so all parties are inactive and (at least) in the "current" round, so we also have $s_j = 1$ and $t_j = t_j^{(j)} + 1$. Hence, after this transition, we have $a_j = a_j^{(j)} = 1$ and $t_j^{(j)} = t_j$.

Empty activation at τ_j**:** This is a special case of the above.

The initial state of the system also fulfills the invariant. Note that while corrupted parties can easily stall the computation (by simply withholding the respective messages), they cannot undermine the consistency guarantees of the protocol. Also, note that the protocol is non-trivial: In presence of an adversary that forwards messages among the protocol machines and occasionally activates them via the network, the protocol will indeed make progress. ☐

Having seen that the availability of partially synchronized clocks cannot be used to UC-securely implement protocols that cannot be implemented without these clocks, one might be tempted to ask whether it is possible to achieve more under relaxed definitions of security. Indeed, [KLP05] provide techniques to construct protocols secure under a relaxed notion of composition. A protocol π is said to achieve security under *composition with δ-delayed protocols*, if all messages sent by protocols executed concurrently with π are delayed by a time-span of at least δ.

D Nielsen's Model

Nielsen [Nie03] adapts the asynchronous UC framework of Canetti [Can01] to the setting of authenticated synchronous networks. While the general structure of the security definition is adopted, the frameworks differ considerably in their concrete definitions. For instance, Nielsen defines the composition of two synchronous protocols in a "lock-step" way: the round switches of all protocols occur at the same time. In Section D.1, we sketch the basic structure of Nielsen's framework [Nie03] and point out an ambiguity in the original formulation. In Section D.2, we show the relation of this framework to our $\{\mathcal{F}_{\text{CLOCK}}, \mathcal{F}_{\text{BD-AUTH}}\}$-hybrid model (in the framework of [Can01]). In particular, we show that the framework also allows to capture guaranteed termination. For a more detailed introduction to the model, refer to Chapter 3 of [Nie03].

D.1 Description of the Model

The fundamental structure of the security definition in [Nie03] is adopted from the framework of *universally composable security* in [Can01]: The security of a protocol π is defined by an ideal functionality $\mathcal{F}$, and π is deemed secure if there is a simulator $\mathcal{S}$ such that an execution of the protocol π in the synchronous network is indistinguishable from an execution of $\mathcal{F}$ with the simulator $\mathcal{S}$.

The synchronous protocol π among parties $\mathcal{P} = [n]$ is executed in rounds, and the model guarantees that each party p_i is activated exactly once per round. In each round ℓ, the protocol machine π_i of party p_i obtains local input as well as the local output generated by sub-protocols and, for each party p_j with $j \in \mathcal{P}$, one message that π_j sent in round $\ell - 1$. During the activation, π_i produces local output, local input for each sub-protocol, and, for each $j \in \mathcal{P}$, one message that will be delivered to π_j in round $\ell + 1$.

The course of the execution is directed by the environment $\mathcal{Z}$. In each round ℓ, $\mathcal{Z}$ may choose the order in which the (honest) parties p_i with $i \in \mathcal{H}$ are activated. $\mathcal{Z}$ chooses the messages sent to p_i by the corrupted parties immediately before the activation, which corresponds to the *strongly rushing* property. Also, the messages generated by p_i for other parties are given to $\mathcal{Z}$ immediately after the activation. After all activations of a round are completed, $\mathcal{Z}$ may switch the computation to the next round. At any point of the computation, $\mathcal{Z}$ may *corrupt* honest parties p_j, obtaining the internal state[18] of π_j. The fact that $\mathcal{Z}$ controls p_j in the further execution is modeled by having $\mathcal{Z}$ specify all future messages sent by p_j. Ideal functionalities must explicitly specify the gained capabilities of $\mathcal{Z}$ in terms of messages leaked to and received from the simulator.

The composition of two ideal functionalities is defined as executing both functionalities in a "lock-step" way. This means that the rounds are defined globally for the complete execution, and that all ideal functionalities switch rounds synchronously. Hence, this also holds for protocols executed both in parallel or as sub-protocols. Note that the communication between a protocol and a sub-protocol also adheres to these rounds of execution, which implies that if a protocol is composed of several layers of sub-protocols, each such layer introduces one round of delay for messages propagated from the functionalities to the player's interfaces.

Ambiguity of Corruption. Upon corruption of a party p_i round r, the model sets $C := C \cup \{i\}$ to store the information that p_i has been corrupted. At a later activation $(\texttt{activate}, j, x_{j,r}, (m_{r,j,r-1})_{r \in C})$ in the same round r, the adversary is—technically—also allowed to specify the message $m_{i,j,r-1}$ which was sent by p_i *in round* $\ell - 1$. We assume that this behavior is not intended, and propose fixing the ambiguity by explicitly specifying whether p_i was corrupted in round $r - 1$, for instance by keeping a sequence of corruption sets $(C^{r'})_{r' \in \mathbb{N}}$ with $C^r \subseteq C^{r+1}$.

D.2 Relation to our Model

In this section, we prove that, for a protocol π and a functionality $\mathcal{F}$ defined in Nielsen's model, there is a compiled protocol $C(\pi)$ that implements the translated functionality $T(\mathcal{F})$ in the $\{\mathcal{F}_{\text{CLOCK}}, \mathcal{F}_{\text{BD-AUTH}}\}$-hybrid model of [Can01] if and only if π implements $\mathcal{F}$ in the model of [Nie03]. Moreover, the translation $T(\cdot)$ preserves the termination guarantees (if a functionality achieves guaranteed termination in [Nie03], then the transformed functionality $T(\mathcal{F})$ also achieves this in our model).

[18] The original version of [Nie03] was based on an earlier version of [Can05] and only leaked the random tape. Leaking the complete internal state is consistent with the "standard corruption" of [Can05, Page 68], possibly including also the local inputs and messages received in previous executions. Also, corruption in [Nie03] is PID-wise.

Translating the Synchronous Functionalities from [Nie03].

For a player set $\mathcal{P}$, the functionality initializes $r := 1$, $\mathcal{H} = \mathcal{P}$, as well as $r_i := 1$, $c_i := 0$, and $r_i := 0$ for all $i \in \mathcal{P}$.

- On each activation while $r_i > r$ and $c_i \geq |\mathcal{P}|$ for all $i \in \mathcal{H}$:
  - If the received message is $(\texttt{end round}, v'_{\mathcal{F}})$ from $\mathcal{A}$, then produce outputs $y_{i,r}$ for all $i \in \mathcal{P}$ as defined by the functionality $\mathcal{F}$. Set $r := r + 1$, $c_i := 0$ for all $i \in \mathcal{P}$, and send $(y_{j,r})_{j \in \mathcal{P} \setminus \mathcal{H}}$ to $\mathcal{A}$.
  - Any other activation is interpreted as $(\texttt{end round}, \perp)$ message[a] from $\mathcal{A}$ (except that the functionality does not send the corrupted parties' outputs to $\mathcal{A}$).
- On input $(\texttt{input}, x_{i,r_i})$ with $r = r_i = r_i$ from party p_i with $i \in \mathcal{H}$, compute the value $v_{\mathcal{F}}$ as specified by the functionality $\mathcal{F}$ and set $r_i := r_i + 1$ and $c_i := 0$. Send $(\texttt{leaked}, i, v_{\mathcal{F}})$ to $\mathcal{A}$.
- On an empty activation from p_i with $r_i > r$, set $c_i := c_i + 1$. If $r_i = r > r_i$, set $r_i := r_i$. Send $(\texttt{activated}, i)$ to $\mathcal{A}$.
- On input $(\texttt{output})$ from party p_i or $(\texttt{output}, j)$ with $j \in \mathcal{P} \setminus \mathcal{H}$ from $\mathcal{A}$, reply with $(\texttt{output}, y_{i,r})$.
- On input $(\texttt{corrupt}, p_i)$ from $\mathcal{A}$ with $i \in \mathcal{H}$, set $\mathcal{H} := \mathcal{H} \setminus \{i\}$, obtain the value $v_{\mathcal{F}}$ from $\mathcal{F}$ by simulating the corresponding corruption message for $\mathcal{F}$ and output $v_{\mathcal{F}}$ to $\mathcal{A}$.

[a] Assumed to be the default input.

Translating Protocols. A protocol machine π_i in the model of [Nie03], in each round r, takes a local input $x_{i,r}$ and a set of messages $(m_{j,i,r-1})_{j \in \mathcal{P}}$ (where the message $m_{i,i,r-1}$ models keeping state), and produces local output $y_{i,r}$ as well as messages $(m_{i,j,r})_{j \in \mathcal{P}}$. The protocol π is transformed into an ITM π' in the $\{\mathcal{F}_{\text{CLOCK}}, \mathcal{F}_{\text{BD-AUTH}}\}$-hybrid model as follows. The ITM π' keeps an (initially empty) internal buffer `buf` with messages that have been computed but not yet output or sent via a channel. Set the round counter to $r_i = 1$ and define $m_{j,i,0} = \perp$ for $j \in \mathcal{P}$.

- On local input $(\texttt{output})$, output the latest value $y_{i,r'}$ that is marked as ready for output.
- On local input $(\texttt{input}, x_{i,r_i})$, use π to compute $\left((m_{i,j,r_i})_{j \in \mathcal{P}}, y_{i,r_i}\right) := \pi\left((m_{j,i,r_i-1})_{j \in \mathcal{P}}, x_{i,r_i}\right)$ and store the messages $m_{i,j,r}$ in `buf`. Deliver the first message.
- Upon an empty activation, deliver one message from `buf`. That is, send $(\texttt{send}, m_{i,j,r})$ to the channel $\mathcal{F}_{\text{BD-AUTH}}(i,j)$. Once `buf` is empty, send $(\texttt{RoundOK})$ to $\mathcal{F}_{\text{CLOCK}}$.
- Upon an empty activation, check via $(\texttt{RequestRound})$ with $\mathcal{F}_{\text{CLOCK}}$ whether the round switch occurred (i.e., $d = 0$). In this case, mark $y_{i,r}$ as ready for output and send $(\texttt{RoundOK})$ to $\mathcal{F}_{\text{CLOCK}}$.
- Upon each activation, check via $(\texttt{RequestRound})$ with $\mathcal{F}_{\text{CLOCK}}$ whether the next round switch occurred (i.e., $d = 0$). In this case, set $r_i := r_i + 1$ and obtain the messages m_{j,i,r_i-1} by sending $(\texttt{fetch})$ to $\mathcal{F}_{\text{BD-AUTH}}(j,i)$ for all $j \in \mathcal{P} \setminus \{i\}$ and proceed with the first step.

Lemma 15 is an immediate consequence of the following lemma.

Lemma 17. *Let π be a protocol and $\mathcal{F}$ be an ideal functionality in the model of [Nie03]. Then, the protocol π implements the functionality $\mathcal{F}$ in the model of [Nie03] if and only if $C(\pi)$ UC-implements the functionality $T(\mathcal{F})$ in the $\{\mathcal{F}_{\text{CLOCK}}, \mathcal{F}_{\text{BD-AUTH}}\}$-hybrid model.*

Proof. Intuitively, we first show that any environment that is a good synchronous distinguisher for the model of [Nie03] is also a good distinguisher for the $\{\mathcal{F}_{\text{CLOCK}}, \mathcal{F}_{\text{BD-AUTH}}\}$-hybrid model (i.e., security in UC implies security in [Nie03]). Hence, we describe a transformation that converts a synchronous environment from into a UC-environment. This resulting environment $\mathcal{C}^{\mathcal{Z}}$ proceeds as follows.

- A query $(\texttt{activate}, i, x_{i,r}, (m_{j,i,r-1})_{j \in \mathcal{C}})$ from $\mathcal{Z}$ is handled by injecting the messages $m_{j,i,r-1}$ into the channels $\mathcal{F}_{\text{BD-AUTH}}(j,i)$ and providing local input $(\texttt{input}, x_{i,r})$ to π'_i. Then, repeatedly activate π'_i until all messages for round r are in the channels $\mathcal{F}_{\text{BD-AUTH}}(i,j)$. Provide the vector $(m_{i,j,r})_{j \in \mathcal{P} \setminus \{i\}}$ of messages leaked on the channels $\mathcal{F}_{\text{BD-AUTH}}(i,j)$ to $\mathcal{Z}$.

- On query $(\texttt{end round})$ from $\mathcal{Z}$, request the local outputs $y_{j,r}$ by first activating π'_j and receiving $(\texttt{switch}, j)$ from $\mathcal{F}_{\text{CLOCK}}$ for each $j \in \mathcal{H}$ (this is done twice in rounds), and then querying $(\texttt{output})$ again at each π'_j for $j \in \mathcal{H}$. Return the vector $(y_{j,\ell})_{j\in\mathcal{H}}$ to $\mathcal{Z}$.
- On query $(\texttt{corrupt}, i)$ from $\mathcal{Z}$, issue the same request to $\mathcal{F}_{\text{CLOCK}}$, all $\mathcal{F}_{\text{BD-AUTH}}(i, j)$ and $\mathcal{F}_{\text{BD-AUTH}}(j, i)$ with $j \in \mathcal{P}$, and instruct the adversary to corrupt π'_i. The internal state of π'_i contains the internal state of π_i. Extract this information and return it to $\mathcal{Z}$.
- If $\mathcal{Z}$ issues a query that would not be allowed in [Nie03], then output 0 and halt. Otherwise, once $\mathcal{Z}$ provides local output, write the same value on the output tape and halt.

Claim: The described "adaptor" perfectly emulates the environment's view. Formally,

$$\forall\, \mathcal{Z}: \ \text{REAL}_{\pi,\mathcal{Z}} \equiv \text{HYB}^{\{\mathcal{F}_{\text{CLOCK}},\mathcal{F}_{\text{BD-AUTH}}\}}_{\pi',\mathcal{D},\mathcal{C}^{\mathcal{Z}}},$$

with the real model of [Nie03] and the UC-hybrid model.

We use an inductive argument: Prior to each query, all random variables in the executions have the same distribution (all messages, inputs, and outputs can be interpreted as random variables defined on the random tapes) in the two cases. During each query, the exact same transformations are applied to these random variables in either case; so the equivalence is extended to the outputs. By induction on the number of queries, we conclude that the output of the two execution also has the same distribution. Again, we differentiate between the three different types of queries.

- For each $i \in \mathcal{H}$, one query $(\texttt{activate}, i, x_{i,\ell}, (m_{j,i,r-1})_{j\in\mathcal{C}})$ is allowed for each round r. In $\text{REAL}_{\pi,\mathcal{Z}}$, such a query results in the output $(m_{i,j,r})_{j\in\mathcal{P}\setminus\{i\}}$. In the converted execution, injecting the messages $m_{j,i,r-1}$ into the corresponding channels is allowed for $j \in \mathcal{C}$, as the channels allow for replacing messages (this is important in case p_j is corrupted after the activation in the current round.) Providing input $(\texttt{input}, x_{i,r})$ to π'_i results in the same distribution for inputs to π_i as in the execution $\text{REAL}_{\pi,\mathcal{Z}}$, so the values computed by π_i also have the same distribution. By making π'_i send all messages $m_{i,j,r}$ to the channels, the converter obtains the tuple $(m_{i,j,r})_{j\in\mathcal{P}\setminus\{i\}}$ with the same distribution as in the "real" execution.
- After all honest parties have been activated in round r, the $(\texttt{end round})$-query is answered with $(y_{i,r})_{i\in\mathcal{H}}$. The converter activates all honest parties and queries their outputs, and the equivalence follows by the same argument as above.
- Upon corruption, the converter obtains the internal state of π'_i. By the definition of π', this includes the state of π_i as expected by the environment $\mathcal{Z}$.

Altogether, this proves the above claim for all environments $\mathcal{Z}$.

Claim: The described "adaptor" indistinguishably emulates the environment's view in the ideal case, if the synchronous simulator $\mathcal{T}$ is constructed from the UC-simulator as described below. Formally,

$$\forall\, \mathcal{S}\, \exists\, \mathcal{T}\, \forall\, \mathcal{Z}: \ \text{IDEAL}_{\mathcal{F},\mathcal{T},\mathcal{Z}} \approx \text{EXEC}_{\mathcal{F}',\mathcal{S},\mathcal{C}^{\mathcal{Z}}},$$

where the left random variable is defined by an execution in [Nie03] and the right random variable is defined by a UC-execution.

Converting Simulators from UC to the Synchronous Model.

The simulator $\mathcal{T}$ is given oracle access to $\mathcal{S}$ and proceeds as follows.

- On query $(\texttt{activate}, i, v_{\mathcal{F}}, (m_{j,i,r-1})_{j\in\mathcal{C}})$, first inject the messages $m_{j,i,r-1}$ into the channels simulated by $\mathcal{S}$ (expect $\mathcal{S}$ to answer consistently with the behavior of real channels—in particular, if a some p_j was corrupted after inserting a message into the channel, the message will be replaced). Then, provide the message $(\texttt{leaked}, i, v_{\mathcal{F}})$ to $\mathcal{S}$ as a message from $T(\mathcal{F})$ and repeatedly send $(\texttt{activated}, i)$ until all messages $m_{i,j,r}$ are supposed to be in (and leaked on) the channels. Provide the tuple $(m_{i,j,r})_{j\in\mathcal{P}\setminus\{i\}}$ as output.
- On query $(\texttt{end round})$, send one message $(\texttt{activated}, i)$ for each $i \in \mathcal{H}$, and expect the message $(\texttt{switch}, i)$ from $\mathcal{F}_{\text{CLOCK}}$ to be simulated. Before $\mathcal{S}$ simulates the last such notification, it may pose the query $(\texttt{end round}, v_{\mathcal{F}})$ to $T(\mathcal{F})$ (otherwise $\mathcal{T}$ sets $v_{\mathcal{F}} := \perp$). Provide $v_{\mathcal{F}}$ as output.
- On input $(y_{i,r})_{i\in\mathcal{C}}$, record this tuple. If $\mathcal{S}$ provided a value $v_{\mathcal{F}}$, provide the tuple $(y_{i,r})_{i\in\mathcal{C}}$ to $\mathcal{S}$ and expect the last missing $(\texttt{switch}, i)$-message. Upon future $(\texttt{output}, j)$ queries from $\mathcal{S}$ to $T(\mathcal{F})$ with $j \in \mathcal{P} \setminus \mathcal{H}$, reply with $(\texttt{output}, y_{j,r})$. Simulate one further message $(\texttt{activated}, i)$ for each $i \in \mathcal{H}$ to $\mathcal{S}$ and expect to receive the corresponding $(\texttt{switch}, i)$-messages as a response.
- On query $(\texttt{corrupt}, i)$, generate the corruption requests for $\mathcal{F}_{\text{CLOCK}}$, all $\mathcal{F}_{\text{BD-AUTH}}(i,j)$ and $\mathcal{F}_{\text{BD-AUTH}}(j,i)$, and π'_i for $\mathcal{S}$. For each of these requests, expect $\mathcal{S}$ to answer according to the behavior of the respective systems. If $\mathcal{S}$, at any future point in time, makes a $(\texttt{corrupt}, i)$ request to $T(\mathcal{F})$, query $(\texttt{corrupt}, i)$ to obtain a value $v_{\mathcal{F}}$ and provide this value to $\mathcal{S}$. Expect a message from $\mathcal{S}$ to the environment that describes the internal state of π'_i, and extract and output the state of π_i.
- If any message from $\mathcal{S}$ is unexpected, halt immediately.

- On queries $(\texttt{activate}, i, x_{i,r}, (m_{j,i,r-1})_{j\in\mathcal{C}})$ for $i \in \mathcal{H}$, the messages returned to the environment are determined equivalently. The synchronous execution first provides $(p_i, x_{i,r})$ to $\mathcal{F}$ to obtain the value $v_{\mathcal{F}}$, before $\mathcal{T}$ injects the messages $m_{j,i,r-1}$ for $j \in \mathcal{C}$ into the channels simulated by $\mathcal{S}$, provides a message $(\texttt{leaked}, i, v_{\mathcal{F}})$ to $\mathcal{S}$, and provokes the computation for the simulated π'_i by repeatedly issuing $(\texttt{activated}, i)$ to $\mathcal{S}$. In the UC-execution, the same messages are injected into the channels simulated by $\mathcal{S}$, and the $(\texttt{input}, x_{i,r})$ given to $T(\mathcal{F})$ also triggers the output of the value $v_{\mathcal{F}}$ to $\mathcal{S}$ by $\mathcal{F}$. This results in an equally distributed message $(\texttt{leaked}, i, v_{\mathcal{F}})$ to $\mathcal{S}$. In both cases, $\mathcal{S}$ repeatedly obtains $(\texttt{activated}, i)$ messages until all messages $m_{i,j,r}$ are leaked on the channels. Since the activations of the systems π_j, $\mathcal{S}$, and $\mathcal{F}$ occur in the same order with equally distributed inputs, all variables have the same distribution after the $\texttt{activate}$-query.
- Upon the $(\texttt{end round})$-query, $\mathcal{S}$ is allowed to produce a $(\texttt{end round}, v_{\mathcal{F}})$-message to $\mathcal{F}$ (or otherwise accept that $v_{\mathcal{F}} = \perp$). In the synchronous execution, the $(\texttt{end round})$ message is given to $\mathcal{T}$, which sends one $(\texttt{activate}, i)$ message for each $i \in \mathcal{H}$ to $\mathcal{S}$ to potentially obtain $v_{\mathcal{F}}$ as a response to the last message. If $\mathcal{S}$ provides such a value $v_{\mathcal{F}}$ it is provided as an output and given to $\mathcal{F}$, and the resulting messages $(y_{i,r})_{i\in\mathcal{P}\setminus\mathcal{H}}$ are returned to $\mathcal{S}$ via $\mathcal{T}$. In the UC-execution, the converter activates the π'_i (in rounds), which yields messages $(\texttt{activated}, i)$ to $\mathcal{S}$, and (by the assumption on $\mathcal{S}$, otherwise we can easily build a good distinguisher) the simulator $\mathcal{S}$ translates these into messages $(\texttt{switch}, i)$. Before simulating the last such message, $\mathcal{S}$ may issue a $(\texttt{end round}, v_{\mathcal{F}})$-query to $T(\mathcal{F})$, which will be given to $\mathcal{F}$, and the messages $(y_{i,r})_{i\in\mathcal{P}\setminus\mathcal{H}}$ are returned to $\mathcal{S}$. If $\mathcal{S}$ does not send such a message, $v_{\mathcal{F}} := \perp$ provided to $\mathcal{F}$ as the simulator's input by $T(\cdot)$, as $\mathcal{C}'$ issues a further round of activations to the π'_i with $i \in \mathcal{H}$. This is consistent with the synchronous execution. In both cases, $\mathcal{S}$ can now access the values $y_{i,r}$ for $i \in \mathcal{C}$ generated by $\mathcal{F}$ on input $v_{\mathcal{F}}$. In [Nie03], the tuple $(y_{i,r})_{i\in\mathcal{H}}$ is directly given to $\mathcal{Z}$; in the converted execution, the converter extracts the same messages by requesting the outputs via $(\texttt{output})$ at all π'_i with $i \in \mathcal{H}$.
- For a $(\texttt{corrupt}, i)$, the simulator $\mathcal{T}$ is notified and converts it into corruption queries for $\mathcal{F}_{\text{CLOCK}}$, the channels, and π'_i. The same messages are generated and given to the simulator $\mathcal{S}$ by the converter $\mathcal{C}$. If (at any future point in time) $\mathcal{S}$ sends a $(\texttt{corrupt}, i)$ query to $\mathcal{F}'$, then $\mathcal{T}$ asks $\mathcal{F}$ to corrupt p_i and provides the obtained value $v_{\mathcal{F}}$ to $\mathcal{S}$. In our model, the $(\texttt{corrupt}, i)$ message is given to $\mathcal{F}'$, which obtains the value $v_{\mathcal{F}}$ in the same way and provides it to $\mathcal{S}$. Hence, the distribution of the representation of π'_i's internal state generated by $\mathcal{S}$ is the same, hence the state of π_i extracted by $\mathcal{T}$ and $\mathcal{C}$ also has the same distribution.

Several above arguments make the assumption that the queries or answers generated by $\mathcal{S}$ are of a certain format. These points are highlighted by explicitly making the simulator $\mathcal{T}$ fail if the answers are different. Yet, in all of these cases, the ideal execution with $\mathcal{S}$ would clearly be distinguishable from the real execution, and since $\mathcal{S}$ is assumed to be a good simulator, such a condition is violated with at most negligible probability.

Combining the above two claims, we conclude that

$$\exists \mathcal{S}\, \forall \mathcal{Z} : \ \text{REAL}_{\pi,\mathcal{Z}} \approx \text{HYB}^{\{\mathcal{F}_{\text{CLOCK}},\mathcal{F}_{\text{BD-AUTH}}\}}_{\pi',\mathcal{D},\mathcal{C}^{\mathcal{Z}}} \approx \text{IDEAL}_{\mathcal{F}',\mathcal{S},\mathcal{C}^{\mathcal{Z}}} \approx \text{IDEAL}_{\mathcal{F},\mathcal{S},\mathcal{Z}},$$

where the first and the third indistinguishability are proven above and the middle indistinguishability follows from the security statement in the hybrid model.

Security in [Nie03] **Implies Security in UC.** We describe a converter $\mathcal{C}'$ that transforms a UC-environment into an synchronous environment. This converter $\mathcal{C}'$ behaves as follows.

- On input $(\texttt{input}, x_{i,r})$ at π'_i, if this is not the first such input in the current round, it is dropped. Otherwise, in any round except for the first one, assemble the vector $(m_{j,i,r-1})_{j\in\mathcal{C}}$ from the messages obtained at the channels (set all messages to $\perp$ in the first round). Call $(\texttt{activate}, i, x_{i,r}, (m_{j,i,r-1})_{j\in\mathcal{C}})$ to obtain $(m_{i,j,\ell})_{j\in\mathcal{P}\setminus\{i\}}$, and store these messages in a buffer $\texttt{buf}_i$. Simulate the first message on a channel $\mathcal{F}_{\text{AUTH}}(i,j)$.
- On further activations of π'_i, if $\texttt{buf}_i$ is not empty, remove the first message from $\texttt{buf}_i$ and simulate it on the channel $\mathcal{F}_{\text{AUTH}}(i,j)$. For the first activation after $\texttt{buf}_i$ is empty, simulate a message $(\texttt{switch}, i)$ from $\mathcal{F}_{\text{CLOCK}}$, and ignore all further such activations until this message has been simulated for all parties p_i with $i \in \mathcal{H}$.
- On the next empty activation, call $(\texttt{end round})$ and record the tuple $(y_{i,r})_{i\in\mathcal{H}}$. For each empty activation at π'_j for $j \in \mathcal{H}$, mark the output $y_{j,r}$ as ready for output and simulate a message $(\texttt{switch}, j)$ from $\mathcal{F}_{\text{CLOCK}}$.
- On input $(\texttt{output})$ at π'_i, return $(\texttt{output}, y_{i,r'})$ with the latest output $y_{i,r'}$ marked as ready.
- On a corruption message $(\texttt{corrupt}, i)$ to the functionality $\mathcal{F}_{\text{CLOCK}}$, issue the call $(\texttt{corrupt}, i)$ to obtain the internal state of π_i and mark π_i as corrupted in $\mathcal{F}_{\text{CLOCK}}$. In the further activations, simulate the execution of π_i as if it were honest. Yet, in future rounds, if $\mathcal{Z}$ did not provide messages for the channels $\mathcal{F}_{\text{AUTH}}(i,j)$, set the corresponding messages $m_{i,j,r}$ to a default value $\perp$.
- On a corruption message $(\texttt{corrupt}, i)$ to a functionality $\mathcal{F}_{\text{BD-AUTH}}(i,j)$ or $\mathcal{F}_{\text{BD-AUTH}}(j,i)$, issue the call $(\texttt{corrupt}, i)$ to obtain the internal state of π_i and mark p_i as corrupted for the corresponding channel. In the further activations, simulate the execution of π'_i as if it were honest. (In particular, if p_i is not corrupted at $\mathcal{F}_{\text{CLOCK}}$, then proceed only once the simulated π'_i would have sent $(\texttt{RoundOK})$.) From here on, allow $\mathcal{Z}$ to replace messages in the corresponding channel. If p_i is corrupted in several channels or in channels and $\mathcal{F}_{\text{CLOCK}}$, it looses further guarantees in the straightforward way.
- Upon a corruption message $(\texttt{corrupt})$ to the protocol machine π'_i, issue the call $(\texttt{corrupt}, i)$ to obtain the internal state of π_i. Since π'_i only uses π_i in a straightforward way, the state of π'_i can be easily simulated. If p_i is not corrupted on all channels and the $\mathcal{F}_{\text{CLOCK}}$, force $\mathcal{Z}$ to use the corresponding honest interfaces.

Claim: The described "adaptor" perfectly emulates the environment's view. Formally,

$$\forall \mathcal{Z} : \ \text{HYB}^{\{\mathcal{F}_{\text{CLOCK}},\mathcal{F}_{\text{BD-AUTH}}\}}_{\pi',\mathcal{D},\mathcal{Z}} \equiv \text{REAL}_{\pi,\mathcal{C}'^{\mathcal{Z}}},$$

where the left random variable is defined by a UC-execution and the right random variable is defined by a synchronous execution.

We use the same inductive argument as above to handle the queries of $\mathcal{Z}$ one-by-one.

- The $(\texttt{input}, x_{i,r})$-query is only allowed once for each $i \in \mathcal{H}$ and r, and ignored otherwise in both cases. The values $y_{i,r}$ and $m_{i,j,r}$ for $j \in \mathcal{P}$ are computed by an invocation of the protocol π_i on equivalently distributed inputs: For $x_{i,r}$, this follows by the induction, so we only have to show that the messages $m_{j,i,r-1}$ are consistent. For $j \in \mathcal{H}$, this is clear, and parties p_j that are corrupted only at $\mathcal{F}_{\text{CLOCK}}$ still perform their computations honestly, so $\mathcal{C}'$ can ensure the consistency. The same argument holds if the corrupted p_j is the receiver at some channel. If p_j is the sender or the protocol machine π'_i is corrupted, then $\mathcal{Z}$ can inject messages $m_{j,i,r-1}$ only before π_i switches to round r, so $\mathcal{C}'$ can include these messages in the $\texttt{activate}$-query.
- Activating π'_i after the $(\texttt{input}, x_{i,r})$-query results in a message $m_{i,j,r}$ being leaked to $\mathcal{Z}$, before a single message $(\texttt{switch}, i)$ from $\mathcal{F}_{\text{CLOCK}}$ is sent. The same output is produced by $\mathcal{C}'$.

- For each $i \in \mathcal{H}$, activating π'_i after all parties have sent their messages in the current round causes π'_i to mark the freshly recorded output value $y_{i,r}$ as ready and output a message $(\texttt{switch}, i)$. This is done consistently by $\mathcal{C}'$. All further activations are ignored until all parties have been activated.
- The queries $(\texttt{output})$ and $(\texttt{output}, j)$ with $j \in \mathcal{P} \setminus \mathcal{H}$ are answered consistently because the same values are marked as ready in both cases, and the distribution of $y_{i,r}$ is the same by induction.
- A corruption message $(\texttt{corrupt}, i)$ to either the functionality $\mathcal{F}_{\text{CLOCK}}$ or a functionality $\mathcal{F}_{\text{BD-AUTH}}(i,j)$ or $\mathcal{F}_{\text{BD-AUTH}}(j,i)$ has no immediate effect on the computation; the indirect effects are described in the input-step.
- Upon a corruption message $(\texttt{corrupt})$ to the protocol machine π'_i, the (simulated) internal state of π'_i is leaked to $\mathcal{Z}$ in both cases. By the construction of the protocol machine π'_i, the state can be easily derived from the state of π_i.

Claim: The described "adaptor" indistinguishably emulates the environment's view in the ideal case, if the UC-simulator $\mathcal{S}$ is constructed from the synchronous simulator as described below. Formally,

$$\forall \mathcal{T}\, \exists \mathcal{S}\, \forall \mathcal{Z} :\ \text{EXEC}_{\mathcal{F}',\mathcal{S},\mathcal{Z}} \approx \text{IDEAL}_{\mathcal{F},\mathcal{T},\mathcal{C}'^{\mathcal{Z}}},$$

where the left random variable is defined by a UC-execution and the right random variable is defined by a synchronous execution.

Converting Synchronous Simulators to UC.

The simulator $\mathcal{S}$ is given oracle access to $\mathcal{T}$ and proceeds as follows.

- On input $(\texttt{leaked}, i, v_{\mathcal{F}})$ from $T(\mathcal{F})$, compute $(m_{i,j,r})_{j\in\mathcal{P}\setminus\{i\}} := \mathcal{T}((m_{i,j,r-1})_{i\in\mathcal{C}}, v_{\mathcal{F}})$ and store the messages in a buffer $\texttt{buf}_i$. Simulate the first message on a channel $\mathcal{F}_{\text{BD-AUTH}}(i,j)$.
- On input $(\texttt{activated}, i)$ where $\texttt{buf}_i$ is not empty, remove the first message m from $\texttt{buf}_i$ and simulate m as a message on the respective channel $\mathcal{F}_{\text{BD-AUTH}}(i,j)$. On the first such activation after $\texttt{buf}_i$ is empty, simulate a message $(\texttt{switch}, i)$ from $\mathcal{F}_{\text{CLOCK}}$. Once this message has been simulated for all $i \in \mathcal{H}$, proceed to the next step.
- On further input $(\texttt{activated}, i)$, if this is the first such message, send $(\texttt{end round})$ to $\mathcal{T}$ to obtain the value $v_{\mathcal{F}}$ and Send $(\texttt{end round}, v_{\mathcal{F}})$ to $T(\mathcal{F})$ to obtain the outputs $(y_{i,r})_{i\in\mathcal{C}}$, and provide these values to $\mathcal{T}$. For each such message, simulate one further message $(\texttt{switch}, i)$ for each $i \in \mathcal{H}$.
- Upon $(\texttt{corrupt}, i)$ for $\mathcal{F}_{\text{CLOCK}}$, issue $(\texttt{corrupt}, i)$ to $T(\mathcal{F})$ to obtain $v_{\mathcal{F}}$ and use $\mathcal{T}$ to obtain the state of π_i, from which the state of π'_i can be generated. Unless explicitly corrupted, simulate π_i as honest (but note that π'_i looses the guarantees of $\mathcal{F}_{\text{CLOCK}}$). Yet, if in a round not all messages expected by π'_i are in the channels, set the undefined $m_{i,j,r}$ to $\perp$.
- On a message $(\texttt{corrupt}, i)$ to $\mathcal{F}_{\text{BD-AUTH}}(i,j)$ or $\mathcal{F}_{\text{BD-AUTH}}(j,i)$, obtain the internal state of π'_i as above, mark p_i as corrupted for the channel. If, for some $j \in \mathcal{H}$, $\mathcal{Z}$ provides an input message m to $\mathcal{F}_{\text{AUTH}}(i,j)$, record this as $m_{i,j,r-1} := m$ for the activation. (If $j \in \mathcal{H}$ is activated without such an input, set $m_{i,j,r-1} := \perp$.) Unless explicitly corrupted, simulate π'_i as honest.
- Upon $(\texttt{corrupt})$ to π'_i, obtain the state of π'_i as above and leak it to $\mathcal{Z}$. If p_i is not corrupted on all channels and the $\mathcal{F}_{\text{CLOCK}}$, make $\mathcal{Z}$ use the corresponding honest interfaces.

- The $(\texttt{input}, x_{i,r})$-query is only allowed once per honest p_i and round r, and ignored otherwise in both cases. In the adapted execution, $\mathcal{C}'$ executes the $\texttt{activate}$-query. As a result, the input is provided to $\mathcal{F}$, and the value $v_{\mathcal{F}}$ produced by $\mathcal{F}$ is given to the simulator $\mathcal{T}$ together with the messages $(m_{j,i,r-1})_{j\in\mathcal{C}}$. $\mathcal{T}$ produces messages $(m_{i,j,r})_{j\in\mathcal{P}\setminus\{i\}}$, which are buffered by $\mathcal{C}'$. The first message is simulated on $\mathcal{F}_{\text{BD-AUTH}}(i,j)$. In the UC-execution, the input is provided to $T(\mathcal{F})$, which also provides it to $\mathcal{F}$ and leaks the resulting output $v_{\mathcal{F}}$ to $\mathcal{S}$. The simulator $\mathcal{S}$ uses $\mathcal{T}$ to compute $(m_{i,j,r})_{j\in\mathcal{P}\setminus\{i\}}$, buffers these messages, and simulates the first message on $\mathcal{F}_{\text{BD-AUTH}}(i,j)$. It remains to show that the messages $(m_{j,i,r-1})_{j\in\mathcal{C}}$ used by $\mathcal{S}$ have the correct distribution. But this holds since $\mathcal{S}$ and $\mathcal{C}'$ handle the injection of messages into the channels $\mathcal{F}_{\text{BD-AUTH}}(j,i)$ for $j \in \mathcal{P} \setminus \mathcal{H}$ in the same way.
- Upon further activations of π'_i, the buffered messages $(m_{i,j,r})_{j\in\mathcal{P}\setminus\{i\}}$ are simulated as sent over the channels $\mathcal{F}_{\text{AUTH}}(i,j)$, and by simulating one message $(\texttt{switch}, i)$ afterwards (and ignoring all further activations until

the messages have been simulated for all honest parties). This is done by $\mathcal{C}'$ in the adapted execution, and $\mathcal{S}$ produces these messages when notified by $T(\mathcal{F})$ via $(\texttt{activated}, i)$ in the UC-execution.

- Upon the first further activation of π'_i, the simulator $\mathcal{T}$ is invoked on input $(\texttt{end round})$ in both cases, and the resulting value $v_{\mathcal{F}}$ is given to $\mathcal{F}$ (via $\mathcal{C}'$ and the model in one case, and via $\mathcal{S}$ and $T(\cdot)$ in the other case). The values $(y_{i,r})_{j\in\mathcal{P}\setminus\mathcal{H}}$ are provided to $\mathcal{T}$ either directly by the model, or by $T(\cdot)$ and $\mathcal{S}$. The values $(y_{i,r})_{j\in\mathcal{H}}$ are stored by $\mathcal{C}'$ and $T(\mathcal{F})$, respectively. Since $\mathcal{F}$ computes on equally distributed values, the values $y_{i,r}$ also have the same distribution in both cases. For the first such activation of π'_i, the value $y_{i,r}$ is made available for output in both cases and the message $(\texttt{switch}, i)$ is simulated, either by $\mathcal{C}'$ or by $T(\mathcal{F})$ and $\mathcal{S}$, respectively.
- By the above arguments, the returned value $(\texttt{output}, y_{i,r'})$ upon input $(\texttt{output})$ at π'_i is always consistent.
- Corruption messages $(\texttt{corrupt}, i)$ to the functionality $\mathcal{F}_{\text{CLOCK}}$ or the channels $\mathcal{F}_{\text{BD-AUTH}}(i,j)$ or $\mathcal{F}_{\text{BD-AUTH}}(j,i)$ have no immediate effect on the computation. The indirect effects are described in the previous steps.
- Upon a corruption message $(\texttt{corrupt})$ to the protocol machine π'_i, the simulated internal state of π'_i is leaked to $\mathcal{Z}$ in both cases. As above, the state is obtained in the same way. By the construction of the protocol machine π'_i, the state can be easily derived from the state of π_i.

By the same arguments as above, this shows that if no environment can distinguish the real execution of π and the ideal execution of $\mathcal{F}$ in the model of [Nie03], then no environment can distinguish the $\{\mathcal{F}_{\text{CLOCK}}, \mathcal{F}_{\text{BD-AUTH}}\}$-hybrid execution of π' and the ideal execution of $\mathcal{F}'$ in UC. This concludes the proof. □

Both the model of [Nie03] and UC allow for a "universal" composition operation (which is the terminology from [Can01] for using sub-protocols). In the words of [MR11], this concept is captured by the statement that the reductions implied by the security models are serially composable with respect to a serial protocol composition. For the full proof of Lemma 15, we have to generalize the statement of Lemma 17 to protocols defined in hybrid models.

We first generalize the mapping $T(\cdot)$ to $T^c(\cdot)$, where the parameter $c \in \mathbb{N}$ denotes the number of empty activations required by the functionality before providing output to the parties. This is necessary because the $T(\mathcal{G})$-hybrid protocol $C(\pi)$ must, besides sending messages and synchronizing using the clock, also provide the activations to the functionality $T(\mathcal{G})$. In more detail, the statement we are going to prove is that π implements the functionality $\mathcal{F}$ in the $\mathcal{G}$-hybrid model of [Nie03] if and only if $C(\pi)$ UC-implements $T^{c+n+1}(\mathcal{F})$ in the $\{T^c(\mathcal{G}), \mathcal{F}_{\text{CLOCK}}, \mathcal{F}_{\text{BD-AUTH}}\}$-hybrid model. The protocols $C(\pi)$, after computing the messages $(m_{i,j,r})_{i\in\mathcal{P}}$, the output $y_{i,r}$, and the input $x^{\mathcal{G}}_{i,r}$ for $\mathcal{G}$, first provides input $x^{\mathcal{G}}_{i,r}$ as well as sufficiently many activations to $T(\mathcal{G})$. Then, $C(\pi)$ sends the messages $(m_{i,j,r})_{j\in\mathcal{P}\setminus\{i\}}$ via the channels $\mathcal{F}_{\text{BD-AUTH}}(i,j)$. In the beginning of the next round, $C(\pi)$ also requests the output $y_{i,r}$ from $T(\mathcal{G})$, which is also taken as an input to the computation of the messages for the next round.

The converters and the simulators in the proof of Lemma 17 must be adapted only slightly: The increased number of empty activations for honest parties has to be taken into account, and the values communicated at the adversarial interface of $\mathcal{G}$ must be forwarded to and from $\mathcal{Z}$.

E Hofheinz and Müller-Quade's Model

Hofheinz and Müller-Quade [HM04] devise another model for synchronous computation based on the paradigms of the framework of universally composable security [Can01]. The definition of a protocol execution differs strongly from both the models of [Can01] and [Nie03]. The model assumes a network of pairwise authenticated communication channels between the honest parties.

E.1 Description of the Model

The execution is defined as an interaction between ITMs. The definition of these machines is similar to the one defined in [Can05], but differs slightly with respect to the communication and the definition of efficiency. As in [Can01], the entities involved in the execution are the environment $\mathcal{Z}$, the adversary $\mathcal{A}$, the protocol machines π, and the ideal functionalities $\mathcal{F}_j$.

The execution proceeds in rounds, each of which is further divided into three phases: the *attack phase*, the *party computation phase*, and the *ideal functionality computation phase*. In each of these phases, only a subset of the ITMs is activated, and the interaction within each of these phases follows a specific set of rules. The *attack phase* models the activities of $\mathcal{Z}$ and $\mathcal{A}$. In particular, the allowed interaction corresponds to the actions of corrupted parties; in this phase, the adversary is allowed to corrupt further parties, interact with the functionalities $\mathcal{F}_j$ in the name of the corrupted parties, and generate messages to the honest parties in the name of the corrupted ones. In the *party computation phase*, all honest parties are activated in parallel, obtain their inputs, the outputs of the $\mathcal{F}_j$, and the messages from all other parties, and generate the respective responses. After this phase, only the messages to the $\mathcal{F}_j$ and the adversary are delivered immediately. In the *ideal functionality computation phase*, all $\mathcal{F}_j$ are activated with the inputs generated by the honest parties. After this phase, the computation proceeds to the *attack phase* of the next round.

Upon corruption of a party p_i, $\mathcal{A}$ immediately obtains p_i's internal state containing the complete history with tapes and head positions. From there on, $\mathcal{A}$ may write arbitrary messages on the outgoing communication tape in the name of p_i and obtains all messages that are sent to p_i.

E.2 Relation to our Model

It turns out that the network implicitly assumed by [HM04] is very strong. Yet, the model *does* allow for guaranteed termination. Overall, one can show that it is embedded into the $\{\mathcal{F}_{\text{CLOCK}}, \mathcal{F}_{\text{NET+}}\}$-hybrid world of our model. The network $\mathcal{F}_{\text{NET+}}$ is sketched below, but the proof is left for the full version of this paper. Note that, in particular, the message sent by an honest party can no longer be changed after they are input to $\mathcal{F}_{\text{NET+}}$.

Functionality $\mathcal{F}_{\text{NET+}}(\mathcal{P})$

Initialize a round counter $r := 1$ and a bit $d_i := 0$ for each $i \in \mathcal{P}$. For each round, proceed as follows.

- Upon receiving input $(\texttt{send}, \vec{M})$ from a party $i \in \mathcal{P}$, where $d_i = 0$ and $\vec{M}$ is a vector of $|\mathcal{P}|$ messages, record $(i, \vec{M}, r)$, set $d_i := 1$ and output $(\texttt{input}, i)$ to the adversary.
- Once $d_i = 1$ for all $i \in \mathcal{H}$, provide the recorded tuples $(i, \vec{M}, r)$ to $\mathcal{A}$ for all $i \in \mathcal{H}$, and allow the adversary to specify the messages to be sent in the name of the players p_j with $j \in \mathcal{P} \setminus \mathcal{H}$.
- Upon receiving input $(\texttt{receive})$ from party p_j with $j \in \mathcal{P}$: if $d_i = 1$ for all $i \in \mathcal{H}$, set $d_i := 0$ for all $i \in \mathcal{H}$ and set $r := r + 1$ (if the adversary did not specify any messages, they are defined to be $\perp$). Output all messages that have been sent to p_i in round $r - 1$.

The fact that the model of [HM04] assumes this network can be seen by analyzing the phases that describe the computation: messages are generated by all honest parties in parallel, without giving the adversary the possibility to interfere. Moreover, the message delivery mechanism is defined in such a way that the adversary cannot prevent messages that have been sent while a party was still honest from being delivered (to other honest parties).