



Round-Preserving Parallel Composition of Probabilistic-Termination Cryptographic Protocols*

Ran Cohen

Northeastern University, Boston, USA
rancohen@ccs.neu.edu

Sandro Coretti

IOHK, Zurich, Switzerland
sandro.coretti@iohk.io

Juan Garay

Department of Computer Science and Engineering, Texas A&M University, College Station, USA
garay@cse.tamu.edu

Vassilis Zikas

Purdue University, West Lafayette, USA
vzikas@cs.purdue.edu

Communicated by Jonathan Katz.

Received 9 September 2017 / Revised 5 February 2021 / Accepted 5 February 2021
Online publication 9 March 2021

Abstract. An important benchmark for multi-party computation protocols (MPC) is their *round complexity*. For several important MPC tasks, such as broadcast, (tight) lower bounds on the round complexity are known. However, some of these lower bounds can be circumvented when the termination round of every party is not a priori known, and simultaneous termination is not guaranteed. Protocols with this property are called *probabilistic-termination (PT)* protocols. Running PT protocols in parallel affects the round complexity of the resulting protocol in somewhat unexpected ways. For instance, an execution of m protocols with constant expected round complexity might take $O(\log m)$ rounds to complete. In a seminal work, Ben-Or and El-Yaniv (Distributed Computing '03) developed a technique for a parallel execution of arbitrarily many broadcast protocols, while preserving expected round complexity. More recently, Cohen *et al.* (CRYPTO '16) devised a framework for universal composition of PT protocols, and provided the first composable parallel-broadcast protocol with a simulation-based proof. These constructions crucially rely on the fact that broadcast is “privacy-free,” and do not generalize to arbitrary protocols in a straightforward way. This raises the question of whether it is possible to execute *arbitrary* PT protocols in parallel, without increasing the round complexity. In this paper we tackle this question

*A preliminary version of this work appeared at *ICALP 2017 (Track A)* [18]. Most of the work was carried out while R. C. was at Tel Aviv University, S. C. at New York University, J. G. at Yahoo Research, and V. Z. at RPI.

and provide both feasibility and infeasibility results. We construct a round-preserving protocol compiler, tolerating any dishonest minority of actively corrupted parties, that compiles arbitrary protocols into a protocol realizing their parallel composition, while having a black-box access to the underlying *protocols*. Furthermore, we prove that the same cannot be achieved, using known techniques, given only black-box access to the *functionalities* realized by the protocols, unless merely security against semi-honest corruptions is required, for which case we provide a protocol. To prove our results, we utilize the language and results by Cohen *et al.*, which we extend to capture parallel composition and reactive functionalities, and to handle the case of an honest majority.

1. Introduction

Secure multi-party computation (MPC) [33,61] allows a set of parties to jointly perform a computation on their inputs, in such a way that no coalition of cheating parties can learn any information beyond what is revealed by their outputs (privacy) or affect the outputs of the computation in any way other than by choosing their own inputs (correctness). Since the first seminal works on MPC [6,13,33,56,61], it has been studied in a variety of different settings and for numerous security notions: There exist protocols secure against passively corrupted (aka semi-honest) parties and against actively corrupted (aka malicious) parties; the underlying network can be synchronous or asynchronous; and the required security guarantees can be information-theoretic or computational—to name but a few of the axes along which the MPC task can be evaluated.

The prevalent model for the design of MPC protocols is the *synchronous* model, where the protocol proceeds in rounds. In this setting, the round complexity, i.e., the number of rounds it takes for a protocol to deliver outputs, is arguably the most important efficiency metric. Tight lower bounds are known on the round complexity of several MPC tasks. For example, for the well-known problems of Byzantine agreement (BA) and broadcast [48,53], it is known that any deterministic protocol against an active attacker corrupting a linear fraction of the parties has linear round complexity [25,28]. This result has quite far-reaching consequences as, starting with the seminal MPC works mentioned above, a common assumption in the design of secure protocols has been that the parties have access to a broadcast channel, which they potentially invoke in every round. In reality, such a broadcast channel might not be available and would have to be implemented by a broadcast protocol designed for a point-to-point network. It follows that even though the round complexity of many MPC protocols is linear in the multiplicative depth of the circuit being computed, their actual running time depends on the number of parties, when executed over point-to-point channels.

The above lower bound on the number rounds for deterministic BA holds even if less than t parties are cheating, as long as all honest parties are required to complete the protocol together, at the same round [26]. Indeed, *randomized* BA protocols that circumvent this lower bound and run in *expected constant* number of rounds (e.g., [4,27,29,44,50,55]) do not provide simultaneous termination, i.e., once a party completes the protocol's execution, it cannot know whether all honest parties have also terminated or if some honest parties are still running the protocol; in particular, the termination round of each party is not a priori known. A protocol with this property is said to have *probabilistic termination (PT)*.

As pointed out by Ben-Or and El-Yaniv [5], when several such PT protocols are executed in parallel, the expected round complexity of the combined execution might no longer be constant (specifically, might not be equal to the maximum of the expected running times of the individual protocols). Indeed, when m protocols, whose termination round is geometrically distributed (and so, have constant expected round complexity), are run in parallel, the expected number of rounds that elapse before all of them terminate is $\Theta(\log m)$ [16]. While an elegant mechanism was proposed in [5] for implementing parallel calls to broadcast such that the total expected number of rounds remains constant, it did not provide any guarantees to remain secure under composition, raising questions about its usability in a higher-level protocol (such as the MPC setting described above). Such a shortcoming was recently addressed by Cohen et al. [16] who provided a framework for universal composition of PT protocols (building upon the framework of universal composability [8]). An application of their result was the first composable protocol for parallel broadcast (with a simulation-based proof) that can be used for securely replacing broadcast channels in arbitrary protocols, and whose round complexity is constant in expectation.

Indeed, an immediate application of the composable parallel-broadcast protocol from [16] is plugging it into broadcast-model MPC protocols in order to obtain point-to-point protocols with a round complexity that is independent of the number of parties. In the information-theoretic setting, this approach yields protocols whose round complexity depends on the depth of the circuit computing the function [6, 13, 21, 56], whereas in the computational setting, assuming standard cryptographic assumptions, this approach yields expected-constant-round protocols [2, 3, 23, 32, 34, 41, 47, 51]. However, the resulting point-to-point protocols have probabilistic-termination on their own. The techniques used for composing PT broadcast protocols in parallel crucially rely on the fact that broadcast is a *privacy-free* functionality, and a naïve generalization of this approach to arbitrary PT protocols fails to be secure. This raises the question of whether it is possible to execute *arbitrary* PT protocols in parallel, without increasing the round complexity.

We remark that circumventing lower bounds on round complexity is just one of the areas where such PT protocols have been successfully used. Indeed, randomizing the termination round has been proven to be a very useful technique in circumventing impossibilities and improving efficiency for many cryptographic protocols. Notable examples include *non-committing encryption* [24], cryptographic protocols designed for rational parties [1, 30, 31, 35, 37, 52], concurrent zero-knowledge protocols [10, 14], and parallel repetition of interactive arguments [36, 38]. The rich literature on such protocols motivates a thorough investigation of their security and composability. As mentioned above, in [16] the initial foundations were laid out for such an investigation, but what was proven for arbitrary PT protocols was a round-preserving *sequential* composition theorem, leaving parallel composition as an open question.

1.1. Our Contributions

Protocol-black-box composition In this work, we investigate the issue of parallel composition for *arbitrary* protocols with probabilistic termination. In particular, we develop a compiler such that given functionalities $\mathcal{F}_1, \dots, \mathcal{F}_M$ and protocols π_1, \dots, π_M , where for every $i \in [M]$, protocol π_i realizes \mathcal{F}_i (possibly using correlated randomness as

setup¹), then the compiled protocol realizes the parallel composition of the functionalities, denoted as $(\mathcal{F}_1 \parallel \dots \parallel \mathcal{F}_M)$.

Our compiler uses the underlying protocols in a black-box manner,² guarantees output delivery (i.e., secure without abort), and is resilient against a computationally unbounded active adversary, adaptively corrupting up to $t < n/2$ parties (which is optimal [56]). Moreover, our compiler is round-preserving, meaning that if the maximal (expected) round complexity of each protocol is μ , then the expected round complexity of the compiled protocol is $O(\mu)$. For example, if each protocol π_i has constant expected round complexity, then so does the compiled protocol. Recall that this task is quite complicated even for the simple case of BA ([5, 16]). For arbitrary functionalities it is even more involved, since as we show, the approach from [5] cannot be applied in a *functionally black-box* way in this case.³ Thus, effectively, our result is the first round-preserving parallel composition result for arbitrary multi-party protocols/tasks with probabilistic termination.

We now describe the ideas underlying our compiler. In [5] (see also [16]), a round-preserving parallel-broadcast protocol was constructed by iteratively running, for a constant number of rounds, *multiple* instances of BA protocols (each instance is executed multiple times in parallel, in a batch), hoping that at least one execution of every BA instance will complete. By choosing the multiplicity suitably, this would occur with constant probability, and therefore, the process is only needed to be repeated a constant expected number of times.

At first sight it might seem that this idea can be applied to arbitrary tasks, but this is not the case. Intuitively, the reason is that if the tasks that we want to compose in parallel have privacy requirements, then making the parties run them in (parallel) “batches” with the same input might compromise privacy, since the adversary will be able to use different inputs and learn multiple outputs of the function(s). This issue is not relevant for broadcast, because it is a “privacy-free” functionality; the adversary may learn the result of multiple computations using the same inputs for honest parties, without compromising security.

To cope with the above issue, our parallel-composition compiler generalizes the approach of [5] in a privacy-preserving manner. At a high level, it wraps the batching technique by an MPC protocol which restricts the parties to use the same input in all protocols for the same function. In particular, the compiler is defined in the *Setup-Commit-then-Prove* hybrid model [11, 42], where the parties receive private correlated randomness that allows each party to commit to its input values and later execute multiple instances of every protocol, each time proving that the same input value is used in all executions.

¹A trusted setup phase is needed for implementing broadcast in the honest-majority setting. Note that as shown in [15, 17], some interesting functions can be computed without such a setup phase.

²Following [40], by a black-box access to a protocol we mean a black-box usage of a semi-honest MPC protocol computing its next-message function.

³Loosely speaking, a functionally black-box protocol, as defined in [57], is a protocol that can compute a function f without knowing the code of f , i.e., given only an oracle access to the function f . Note that in this model, each ideal functionality \mathcal{F}_i has an oracle access to the function f_i it computes. We note that most MPC protocols, e.g., [6, 33, 61], are *not* functionally black-box, as they need an explicit representation of the function in the form of a circuit.

The constructions in [11,42] for realizing the Setup-Commit-then-Prove functionality are designed for the dishonest-majority setting and therefore allow for a premature abort. Since we assume an honest majority, we require security with *guaranteed output delivery*. A possible way around would be, as is common in the MPC literature, to restart the protocol upon discovering some cheating or add for each abort a recovery round; this, however, would induce a linear overhead (in the number of parties) on the round complexity of the protocol.

Instead, in order to recover from a misbehavior by corrupted parties, we modify the Setup-Commit-then-Prove functionality and secret-share every committed random string between all the parties, using an error-correcting secret-sharing scheme (aka robust secret sharing [12,22,56]).⁴ In case a party is identified as cheating, every party broadcasts the share of the committed randomness corresponding to that party, reconstructs the correlated randomness for that party, and locally computes the messages corresponding to this party in every instance of every protocol. We also prove that the modified Setup-Commit-then-Prove functionality can be realized in a constant number of rounds, thus yielding no (asymptotic) overhead on the round complexity of the compiler. We refer the reader to Sect. 6.1 for more details.

Functionally black-box composition Our feasibility result described above shows how to compute the parallel composition $(\mathcal{F}_1 \parallel \dots \parallel \mathcal{F}_M)$ of the functionalities $\mathcal{F}_1, \dots, \mathcal{F}_M$, given a black-box access to the protocols π_1, \dots, π_M . That is, the compiler crucially relies on knowing the *messages* (i.e., the *transcript*) of the protocols' executions, but it does not need to know *how* these messages are generated. Indeed, given the messages of the underlying protocols, it is possible to exploit MPC techniques to enforce consistency between multiple executions of the protocols.

We proceed to ask whether a similar compiler can be constructed *without* having access to the transcripts of the protocols' executions. Specifically, we investigate the question of whether there exists a protocol that securely realizes $(\mathcal{F}_1 \parallel \dots \parallel \mathcal{F}_M)$ given only black-box access to the *functionalities* $\mathcal{F}_1, \dots, \mathcal{F}_M$, but not to protocols realizing them. Each \mathcal{F}_i is modeled as an ideal functionality that receives the inputs from the parties and returns the output in a probabilistic-termination fashion. We note that this question only makes sense if the functionalities $\mathcal{F}_1, \dots, \mathcal{F}_M$ are not fixed in advance, but belong to some *class* of functionalities (see [57]), since otherwise a protocol may always ignore the functionalities $\mathcal{F}_1, \dots, \mathcal{F}_M$ and implement $(\mathcal{F}_1 \parallel \dots \parallel \mathcal{F}_M)$ from scratch. We thus consider M classes of functions $\mathcal{C}_1, \dots, \mathcal{C}_M$, where the i 'th ideal functionality \mathcal{F}_i has an oracle access to one of the functions in \mathcal{C}_i , denoted as $\mathcal{F}_i^{\mathcal{C}_i}$.

When considering semi-honest corruptions, we show that there indeed exists a round-preserving protocol for parallel composition of arbitrary deterministic functionalities $\mathcal{F}_1^{\mathcal{C}_1}, \dots, \mathcal{F}_M^{\mathcal{C}_M}$ in a functionally black-box manner. The composition follows the “batching-technique” from [5,16] of invoking each functionality $\mathcal{F}_i^{\mathcal{C}_i}$ multiple times in parallel for a constant number of rounds, such that with constant probability at least one computation in a batch will successfully produce an output. The reason that this

⁴A (t, n) secret-sharing scheme is error correcting, if the reconstruction algorithm outputs the correct secret even when up to t shares are arbitrarily modified.

approach works is that if the adversary follows the protocol, it cannot provide different inputs to different invocations of $\mathcal{F}_i^{\mathcal{C}}$. This works for deterministic computations since it is crucial to ensure that the adversary obtains one output value per batch even if more than one invocation produces outputs. Note that using standard techniques a randomized functionality can be reduced to a deterministic one in which each party provides an additional random string r_i , and the functionality uses $r_1 \oplus \dots \oplus r_n$ as its random coins. We refer the reader to Sect. 5 for more details.

When considering malicious corruptions, we prove that there exist function-classes for which no “natural” protocol (that is based on any of the known techniques from the literature) can compute their parallel composition in a round-preserving manner, while accessing the functions in a black-box way, tolerating even a single adversarial party. More precisely, we consider the function class \mathcal{C} of distributed point functions, consisting of functions $f_\alpha(x_1, x_2, \lambda, \dots, \lambda)$ (parametrized by $\alpha \in \{0, 1\}^\kappa$, where κ is the security parameter and λ is the empty string) that return α if $x_1 \oplus x_2 = \alpha$ and 0 otherwise. We show that for this class it holds that:

1. Calling to each of the ideal functionalities $\mathcal{F}_i^{\mathcal{C}}$ until termination (i.e., until *all* parties receive the output) will not be round-preserving.
2. It is impossible to compute the parallel composition without calling every ideal functionality (until *some* parties receive the output).
3. Using the same input value in more than one call to any of the ideal functionalities will break privacy.

We refer the reader to Sect. 6.2 for the details of the lower bound.

We phrase our results using the framework for composition of protocols with probabilistic termination [16], and extend it (a side result of independent interest) to include parallel composition, reactive functionalities (in order to capture the Setup-Commit-then-Prove functionality), and to the higher corruption threshold of $t < n/2$.

Related work Most relevant to our work are results on compositional aspects of broadcast protocols with probabilistic termination. Lindell et al. [49] studied sequential composition of probabilistic-termination Byzantine agreement protocols, and Ben-Or and El-Yaniv [5] a round-preserving parallel composition of such protocols. A simplification for parallel composition of probabilistic-termination BA protocols that are based on leader election was presented by Fitzi and Garay [29], and was used by Katz and Koo [44] for analyzing the exact round complexity of probabilistic-termination broadcast in the context of secure computation. Cohen et al. [16] studied probabilistic-termination protocols in the UC framework and constructed a composable parallel-broadcast protocol with a simulation-based proof.

Organization of the paper The rest of the paper is organized as follows. In Sect. 2 we describe the network model, the basics of the probabilistic-termination framework by Cohen et al. [16], and other tools that are used throughout the paper. In Sect. 3 we extend the framework to the honest-majority setting, and in Sect. 4 define parallel composition of PT protocols. Section 5 presents the protocol that achieves round-preserving parallel composition for arbitrary functionalities in a functionally black-box manner against semi-honest adversaries. Section 6 is dedicated to active corruptions; first, the round-

preserving protocol-black-box construction is presented, followed by the negative result on round-preserving functionally black-box composition in the case of active corruptions. For ease of exposition, finer details of the model and PT framework are presented in “Appendix.”

2. Model and Preliminaries

In the following, we introduce some necessary notation and terminology. We denote by κ the security parameter. For $n \in \mathbb{N}$, let $[n] = \{1, \dots, n\}$. Let poly denote the set of all positive polynomials, and let PPT denote a probabilistic algorithm that runs in *strictly* polynomial time. A function $\nu: \mathbb{N} \rightarrow [0, 1]$ is *negligible* if $\nu(\kappa) < 1/p(\kappa)$ for every $p \in \text{poly}$ and large enough κ . Given a random variable X , we write $x \leftarrow X$ to indicate that x is selected according to X .

The statistical distance between two random variables X and Y over a finite set \mathcal{U} , denoted $\text{SD}(X, Y)$, is defined as $\frac{1}{2} \cdot \sum_{u \in \mathcal{U}} |\Pr[X = u] - \Pr[Y = u]|$. Two distribution ensembles $X = \{X(\kappa)\}_{\kappa \in \mathbb{N}}$ and $Y = \{Y(\kappa)\}_{\kappa \in \mathbb{N}}$ are *statistically close* (denoted $X \stackrel{s}{\equiv} Y$) if $\text{SD}(X, Y) \leq \nu(\kappa)$ for a negligible function ν .

2.1. Synchronous Protocols in UC

We consider synchronous protocols in the model of Katz et al. [46], which is designed on top of the universal composability framework of Canetti [8]. More specifically, we consider n parties P_1, \dots, P_n and a computationally unbounded, adaptive t -adversary that can dynamically corrupt up to t parties during the protocol execution. Synchronous protocols in [46] are protocols that run in a hybrid model where parties have access to a simple “clock” functionality $\mathcal{F}_{\text{CLOCK}}$. This functionality keeps an indicator bit, which is switched once *all honest parties* request the functionality to do so, i.e., once all honest parties have completed their operations for the current round. In addition, all communication is done over bounded-delay secure channels, where each party requests the channel to fetch messages that are sent to him, such that the adversary is allowed to delay the message delivery by a bounded and a priori known number of fetch requests. Stated differently, once the sender has sent some message, it is guaranteed that the message will be delivered within a known number of activations of the receiver. For simplicity, we assume that every message is delivered within a single fetch request. A more detailed overview of [46] can be found in “Appendix B.”

2.2. The Probabilistic-Termination Framework

Cohen et al. [16] extended the UC framework to capture protocols with probabilistic termination, i.e., protocols without a fixed output round and without simultaneous termination. This section outlines their techniques; additional details can be found in “Appendix C.”

Canonical synchronous functionalities The main idea behind modeling probabilistic termination is to separate the functionality to be computed from the round complexity that is required for the computation. The atomic building block in [16] is a functionality template called a *canonical synchronous functionality (CSF)*, which is a simple two-round functionality with explicit (one-round) input and (one-round) output phases. The functionality \mathcal{F}_{CSF} has two parameters: (1) a (possibly) randomized function f that receives $n + 1$ inputs (n inputs from the parties and one additional input from the adversary) and (2) a leakage function l that determines what information about the input values is leaked to the adversary.

The functionality \mathcal{F}_{CSF} proceeds in two rounds: In the first (input) round, all the parties hand \mathcal{F}_{CSF} their input values, and in the second (output) round, each party receives its output. Whenever some input is submitted to \mathcal{F}_{CSF} , the adversary is handed some leakage function of this input; the adversary can use this leakage for deciding which parties to corrupt and which input values to use for corrupted parties. Additionally, the adversary is allowed to input an extra message, which—depending on the function f —might affect the output(s). The detailed description of \mathcal{F}_{CSF} is given in Fig. 7 in “Appendix C.1.” As a side contribution, in Definition C.1, we extend the definition of CSF to the reactive setting.

Wrappers and traces Computation with probabilistic termination is captured by defining *output-round randomizing wrappers*. Such wrappers address the issue that while an ideal functionality abstractly describes a protocol’s task, it does not describe its round complexity. Each wrapper is parametrized by a distribution (more precisely, an efficient probabilistic sampling algorithm) D that may depend on a specific protocol implementing the functionality. The wrapper samples a round $\rho_{\text{term}} \leftarrow D$, by which all parties are guaranteed to receive their outputs. Two wrappers are considered: The first, denoted $\mathcal{W}_{\text{strict}}$, ensures in a strict manner that all (honest) parties terminate together in round ρ_{term} ; the second, denoted $\mathcal{W}_{\text{flex}}$, is more flexible and allows the adversary to deliver outputs to individual parties at any time before round ρ_{term} . The detailed descriptions of the two wrappers can be found in “Appendix C.4.”

As pointed out in [16], it is not sufficient to inform the simulator \mathcal{S} about the round ρ_{term} . In many cases, the wrapper should explain to \mathcal{S} *how* this round was sampled; concretely, the wrapper provides \mathcal{S} with the random coins that are used to sample ρ_{term} . In particular, \mathcal{S} learns the entire *trace* of calls to ideal functionalities that are made by the protocol in order to complete by round ρ_{term} . A trace basically records which hybrids were called by a protocol’s execution, and in a recursive way, for each hybrid, which hybrids would have been called by a protocol realizing that hybrid. The recursion ends when the base case is reached, i.e., when the protocol is defined using the *atomic* functionalities that are “assumed” by the model.⁵ We refer the reader to “Appendix C.3” for further intuition and an illustrating example, and formally define a trace as follows:

Definition 2.1. (Traces) A *trace* is a rooted tree of depth at least 1, in which all nodes are labeled by functionalities and where every node’s children are *ordered*. The root and

⁵The atomic functionalities considered in this work are the CSFs for the point-to-point communication functionality \mathcal{F}_{SMT} and the correlated-randomness functionality for broadcast $\mathcal{F}_{\text{CORR-BC}}$.

all internal nodes are labeled by wrapped CSFs (by either of the two wrappers), and the leaves are labeled by unwrapped CSFs. The *trace complexity* of a trace T , denoted $c_{\text{tr}}(T)$, is the number of leaves in T . Moreover, denote by $\text{flex}_{\text{tr}}(T)$ the number nodes labeled by flexibly wrapped CSFs in T .

In this work, we consider an augmented definition of traces, which allows parallel calls to ideal functionalities at the same round. Specifically, a trace is augmented with another (potentially empty) layer of nodes, such that each leaf, in the original definition of a trace, may have a list of (unordered) children. The trace complexity is defined as in Definition 2.1, as the number of original leaves (before augmenting with unsorted node-lists). We note that the composition theorems (below) trivially extend to use this augmented definition of a trace. To simplify notations, we denote by $[\mathcal{F}_1, \dots, \mathcal{F}_m]$ the node with unordered list of children $\mathcal{F}_1, \dots, \mathcal{F}_m$ (modeling a parallel call to these functionalities). We will also denote by

$$\left(\left([\mathcal{F}_1]^k, \dots, [\mathcal{F}_m]^k \right) \right)^l$$

a trace consisting of l (ordered) sequences of leaves, each augmented with an (unordered) list of k nodes of each \mathcal{F}_i . This corresponds to l sequential calls to k parallel instances of each \mathcal{F}_i .

Sequential composition of probabilistic-termination protocols When a set of parties execute a probabilistic-termination protocol, or equivalently, invoke a flexibly wrapped CSF, they might get out-of-sync and start the next protocol in different rounds. The approach in [16] for dealing with sequential composition is to start by designing simpler protocols, that are in a so-called *synchronous normal form*, where the parties remain in-sync throughout the execution, and next, compile these protocols into slack-tolerant protocols.

Definition 2.2. (*Synchronous normal form*) Let $\mathcal{F}_1, \dots, \mathcal{F}_m$ be canonical synchronous functionalities. A synchronous protocol π in the $(\mathcal{F}_1, \dots, \mathcal{F}_m)$ -hybrid model is in *synchronous normal form (SNF)* if in every round exactly one ideal functionality \mathcal{F}_i is invoked by all honest parties, and in addition, no honest party hands inputs to other CSFs before this instance halts.

SNF protocols are designed as an intermediate step only, since the hybrid functionalities $\mathcal{F}_1, \dots, \mathcal{F}_m$ are two-round CSFs and, in general, cannot be realized by real-world protocols. In order to obtain protocols that can be realized in the real world, [16] introduced *slack-tolerant* variants of both the strict and the flexible wrappers, denoted $\mathcal{W}_{\text{sl-strict}}$ and $\mathcal{W}_{\text{sl-flex}}$. These wrappers are parametrized by a slack parameter $c \geq 0$ and can be used even if parties provide inputs within $c + 1$ consecutive rounds (i.e., they tolerate input slack of c rounds); furthermore, the wrappers ensure that all honest parties obtain output within two consecutive rounds (i.e., they reduce the slack to $c = 1$). The detailed definitions of the slack-tolerant wrappers are given in ‘‘Appendix C.5.’’ In order to convert SNF protocols into protocols that realize functionalities with slack tolerance, [16] con-

structured a deterministic-termination compiler Comp_{DT} , a probabilistic-termination compiler Comp_{PT} , and a probabilistic-termination with slack-reduction compiler Comp_{PTR} . Loosely speaking, the composition theorems provide the following guarantees:

1. If an SNF protocol π realizes a wrapped CSF $\mathcal{W}_{\text{strict}}^D(\mathcal{F})$ in the $(\mathcal{F}_1, \dots, \mathcal{F}_m)$ -hybrid model, then $\text{Comp}_{\text{DT}}^c(\pi)$ realizes $\mathcal{W}_{\text{sl-strict}}^{D',c}(\mathcal{F})$ in the $(\mathcal{W}_{\text{sl-strict}}^{D_1,c}(\mathcal{F}_1), \dots, \mathcal{W}_{\text{sl-strict}}^{D_m,c}(\mathcal{F}_m))$ -hybrid model (where D' is defined using D, D_1, \dots, D_m).
2. If an SNF protocol π realizes a wrapped CSF $\mathcal{W}_{\text{strict}}^D(\mathcal{F})$ in the $(\mathcal{F}_1, \dots, \mathcal{F}_m)$ -hybrid model, then $\text{Comp}_{\text{PT}}^c(\pi)$ realizes $\mathcal{W}_{\text{sl-flex}}^{D',c}(\mathcal{F})$ in the $(\mathcal{W}_{\text{sl}}^{D_1,c}(\mathcal{F}_1), \dots, \mathcal{W}_{\text{sl}}^{D_m,c}(\mathcal{F}_m))$ -hybrid model (where $\mathcal{W}_{\text{sl}}^{D_i,c}(\mathcal{F}_i)$ is $\mathcal{W}_{\text{sl-flex}}^{D_i,c}(\mathcal{F}_i)$ if $i \in I$ and $\mathcal{W}_{\text{sl-strict}}^{D_i,c}(\mathcal{F}_i)$ if $i \notin I$, given a subset $I \subseteq [m]$ (of indices) of functionalities to be wrapped using the flexible wrapper).
3. If an SNF protocol π realizes a wrapped CSF $\mathcal{W}_{\text{flex}}^D(\mathcal{F})$ in the $(\mathcal{F}_1, \dots, \mathcal{F}_m)$ -hybrid model, then $\text{Comp}_{\text{PTR}}^c(\pi)$ realizes $\mathcal{W}_{\text{sl-flex}}^{D',c}(\mathcal{F})$ in the $(\mathcal{W}_{\text{sl}}^{D_1,c}(\mathcal{F}_1), \dots, \mathcal{W}_{\text{sl}}^{D_m,c}(\mathcal{F}_m))$ -hybrid model (where $\mathcal{W}_{\text{sl}}^{D_i,c}(\mathcal{F}_i)$ is $\mathcal{W}_{\text{sl-flex}}^{D_i,c}(\mathcal{F}_i)$ if $i \in I$ and $\mathcal{W}_{\text{sl-strict}}^{D_i,c}(\mathcal{F}_i)$ if $i \notin I$, for a subset I as above).

The compilers maintain the security and the asymptotic (expected) round complexity of the original SNF protocols. At the same time, the compilers take care of any potential slack that is introduced by the protocol and ensure that the resulting protocol can be safely executed even if the parties do not start the protocol simultaneously. More precise descriptions of the compilers can be found in ‘‘Appendix C.6.’’ As a side contribution, we extend this framework to the honest-majority setting in ‘‘Appendix C.’’

Finally, in [16], the authors also provided protocols for realizing wrapped variants of the atomic CSF functionality for secure point-to-point communication. This suggested the following design paradigm for realizing a wrapped functionality $\mathcal{W}_{\text{sl-strict}}(\mathcal{F})$ (resp., $\mathcal{W}_{\text{sl-flex}}(\mathcal{F})$): First, construct an SNF protocol for realizing $\mathcal{W}_{\text{strict}}(\mathcal{F})$ (resp., $\mathcal{W}_{\text{flex}}(\mathcal{F})$) using CSF hybrids $\mathcal{F}_1, \dots, \mathcal{F}_m$. Next, for each of the non-atomic hybrids \mathcal{F}_i , show how to realize $\mathcal{W}_{\text{strict}}(\mathcal{F}_i)$ (resp., $\mathcal{W}_{\text{flex}}(\mathcal{F}_i)$) using CSF hybrids $\mathcal{F}'_1, \dots, \mathcal{F}'_{m'}$. Proceed in this manner until all CSF hybrids are atomic functionalities. Finally, repeated applications of the composition theorems above yield a protocol for $\mathcal{W}_{\text{sl-strict}}(\mathcal{F})$ (resp., $\mathcal{W}_{\text{sl-flex}}(\mathcal{F})$) using only atomic functionalities as hybrids.

2.3. A Lemma on Termination Probabilities

The following lemma, which will be used in our positive results, provides a constant lower bound on the probability that when running simultaneously (i.e., in parallel) N copies of M probabilistic-termination protocols π_1, \dots, π_M , at least one copy of each π_i will complete after R rounds, for suitable choices of N and R .

Lemma 2.3. *Let $M, N, R \in \mathbb{N}$. For $i \in [M]$ and $j \in [N]$, let X_{ij} be independent random variables over the natural numbers, such that X_{i1}, \dots, X_{iN} are identically distributed with expectation μ_i . Denote $Y_i = \min\{X_{i1}, \dots, X_{iN}\}$ and $\mu = \max\{\mu_1, \dots, \mu_M\}$. Then, for any constant $0 < \epsilon < 1$, if $R > \mu$ and $N > \frac{\log(M/\epsilon)}{\log(R/\mu)}$, it holds that $\Pr[\forall i : Y_i < R] \geq 1 - \epsilon$.*

Proof First, notice that

$$\begin{aligned}
\Pr[\exists i : Y_i \geq R] &\leq \sum_{i \in [M]} \Pr[Y_i \geq R] \\
&\leq \sum_{i \in [M]} \left(\prod_{j \in [N]} \Pr[X_{ij} \geq R] \right) \\
&\stackrel{(*)}{\leq} \sum_{i \in [M]} \left(\prod_{j \in [N]} \left(\frac{\mu_j}{R} \right) \right) \\
&\leq \sum_{i \in [M]} \left(\prod_{j \in [N]} \left(\frac{\mu}{R} \right) \right) \\
&= M \cdot \left(\frac{\mu}{R} \right)^N,
\end{aligned}$$

where $(*)$ follows from Markov's inequality. Therefore,

$$\Pr[\forall i : Y_i < R] = 1 - \Pr[\exists i : Y_i \geq R] \geq 1 - M \cdot \left(\frac{\mu}{R} \right)^N.$$

Finally, since $R > \mu$ it holds that $1 > \mu/R$. Therefore, for constant $0 < \epsilon < 1$, by setting

$$N > \frac{\log(M/\epsilon)}{\log(R/\mu)},$$

it holds that

$$\Pr[\forall i : Y_i < R] \geq 1 - M \cdot \left(\frac{\mu}{R} \right)^N \geq 1 - \epsilon.$$

□

3. Probabilistic Termination with an Honest Majority

In this section, we extend the probabilistic-termination framework [16] (that was defined for the 2/3-majority setting) to the honest-majority regime.

3.1. Fast Sequential Composition

The composition theorems from [16] are defined for $t < n/3$ (as they focused on perfect security). When moving to the honest-majority setting, i.e., $t < n/2$, the compilers and composition theorems follow in a straightforward way. The main difference lies in the

usage of the Bracha-termination technique (specifically, in Theorem 3.2), where the “termination messages” in the compiled protocol $\pi' = \text{Comp}_{\text{PTR}}^c(\pi, D_1, \dots, D_m, I)$ must be authenticated. Therefore, there is an additional hybrid functionality that is required for generating correlated randomness to be used for authenticating messages. If information-theoretic security is required, correlated randomness for information-theoretic signatures [54] can be used, whereas if computational security suffices, a public-key infrastructure (PKI) can be used.

- **CORRELATED RANDOMNESS.** The *correlated-randomness* functionality, parametrized by a distribution D , is defined as follows. The function to compute is $f_{\text{CORR}}(\lambda, \dots, \lambda, a) = (R_1, \dots, R_n)$, where $(R_1, \dots, R_n) \leftarrow D$, and the leakage function is $l_{\text{CORR}}(\lambda, \dots, \lambda) = \perp$. We denote by $\mathcal{F}_{\text{CORR}}^D$ the functionality \mathcal{F}_{CSF} when parametrized with the above functions f_{CORR} and l_{CORR} . We denote by $\mathcal{F}_{\text{CORR-BC}}$ the functionality $\mathcal{F}_{\text{CORR}}^{D_{\text{BC}}}$, where D_{BC} is the distribution for correlated randomness needed for information-theoretic broadcast [54], and by \mathcal{F}_{PKI} functionality $\mathcal{F}_{\text{CORR}}^{D_{\text{PKI}}}$, where D_{PKI} is the distribution for correlated randomness needed for a public-key infrastructure using standard digital signatures.

We state without proof the composition theorems for the honest-majority setting. The proofs follow in similar lines to [16]. See “Appendix C.6” for additional details including the definition of full-trace.

Theorem 3.1. *Let $\mathcal{F}, \mathcal{F}_1, \dots, \mathcal{F}_m$ be canonical synchronous functionalities, let $t < n/2$, and let π be an SNF protocol that UC-realizes $\mathcal{W}_{\text{strict}}^D(\mathcal{F})$, with information-theoretic (resp., computational) security, in the $(\mathcal{F}_1, \dots, \mathcal{F}_m)$ -hybrid model, for some depth-1 distribution D , in the presence of an adaptive, malicious t -adversary, and assuming that all honest parties receive their inputs at the same round. Let D_1, \dots, D_m be arbitrary distributions over traces, let $D^{\text{full}} = \text{full-trace}(D, D_1, \dots, D_m)$, and let $c \geq 0$.*

Then, protocol $\pi' = \text{Comp}_{\text{DT}}^c(\pi, D_1, \dots, D_m)$ UC-realizes $\mathcal{W}_{\text{sl-strict}}^{D^{\text{full},c}}(\mathcal{F})$, with information-theoretic (resp., computational) security, in the $(\mathcal{W}_{\text{sl-strict}}^{D_1,c}(\mathcal{F}_1), \dots, \mathcal{W}_{\text{sl-strict}}^{D_m,c}(\mathcal{F}_m))$ -hybrid model, in the presence of an adaptive, malicious t -adversary, assuming that all honest parties receive their inputs within $c + 1$ consecutive rounds.

Furthermore, the expected round complexity of the compiled protocol π' is

$$B_c \cdot \sum_{i \in [m]} d_i \cdot E[\text{ctr}(T_i)],$$

where d_i is the expected number of calls in π to hybrid \mathcal{F}_i , T_i is a trace sampled from D_i , and $B_c = 3c + 1$ is the blow-up factor.

Theorem 3.2. *Let $\mathcal{F}, \mathcal{F}_1, \dots, \mathcal{F}_m$ be canonical synchronous functionalities, let $t < n/2$, and let π be an SNF protocol that UC-realizes $\mathcal{W}_{\text{flex}}^D(\mathcal{F})$, with information-theoretic (resp., computational) security, in the $(\mathcal{F}_1, \dots, \mathcal{F}_m)$ -hybrid model, for some depth-1 distribution D , in the presence of an adaptive, malicious t -adversary, and assuming that all honest parties receive their inputs at the same round. Let $I \subseteq [m]$ be the subset (of indices) of functionalities to be wrapped using the flexible wrapper, let D_1, \dots, D_m be*

arbitrary distributions over traces, denote $D^{\text{full}} = \text{full-trace}(D, D_1, \dots, D_m)$, and let $c \geq 0$. Assume that \mathcal{F} and \mathcal{F}_i , for every $i \in I$, are public-output functionalities.

Then, the compiled protocol $\pi' = \text{Comp}_{\text{PTR}}^c(\pi, D_1, \dots, D_m, I)$ UC-realizes $\mathcal{W}_{\text{sl-flex}}^{D^{\text{full}},c}(\mathcal{F})$, with information-theoretic (resp., computational) security, in the $(\mathcal{F}_{\text{CORR-BC}}, \mathcal{W}(\mathcal{F}_1), \dots, \mathcal{W}(\mathcal{F}_m))$ -hybrid model, where $\mathcal{W}(\mathcal{F}_i) = \mathcal{W}_{\text{sl-flex}}^{D_i,c}(\mathcal{F}_i)$ if $i \in I$ and $\mathcal{W}(\mathcal{F}_i) = \mathcal{W}_{\text{sl-strict}}^{D_i,c}(\mathcal{F}_i)$ if $i \notin I$, in the presence of an adaptive, malicious t -adversary, assuming that all honest parties receive their inputs within $c + 1$ consecutive rounds.

Furthermore, the expected round complexity of the compiled protocol π' is

$$B_c \cdot \sum_{i \in [m]} d_i \cdot E[c_{\text{tr}}(T_i)] + 2 \cdot \sum_{i \in [m]} d_i \cdot E[\text{flex}_{\text{tr}}(T_i)] + 2,$$

where d_i is the expected number of calls in π to hybrid \mathcal{F}_i , T_i is a trace sampled from D_i , and $B_c = 3c + 1$ is the blow-up factor.

Theorem 3.3. Let $\mathcal{F}, \mathcal{F}_1, \dots, \mathcal{F}_m$ be canonical synchronous functionalities, let $t < n/2$, and let π be an SNF protocol that UC-realizes $\mathcal{W}_{\text{strict}}^D(\mathcal{F})$, with information-theoretic (resp., computational) security, in the $(\mathcal{F}_1, \dots, \mathcal{F}_m)$ -hybrid model, for some depth-1 distribution D , in the presence of an adaptive, malicious t -adversary, and assuming that all honest parties receive their inputs at the same round. Let $I \subseteq [m]$ be the subset (of indices) of functionalities to be wrapped using the flexible wrapper, let D_1, \dots, D_m be arbitrary distributions over traces, denote $D^{\text{full}} = \text{full-trace}(D, D_1, \dots, D_m)$, and let $c \geq 0$. Assume that \mathcal{F} and \mathcal{F}_i , for every $i \in I$, are public-output functionalities.

Then, the compiled protocol $\pi' = \text{Comp}_{\text{PT}}^c(\pi, D_1, \dots, D_m, I)$ UC-realizes $\mathcal{W}_{\text{sl-flex}}^{D^{\text{full}},c}(\mathcal{F})$, with information-theoretic (resp., computational) security, in the $(\mathcal{W}(\mathcal{F}_1), \dots, \mathcal{W}(\mathcal{F}_m))$ -hybrid model, where $\mathcal{W}(\mathcal{F}_i) = \mathcal{W}_{\text{sl-flex}}^{D_i,c}(\mathcal{F}_i)$ if $i \in I$ and $\mathcal{W}(\mathcal{F}_i) = \mathcal{W}_{\text{sl-strict}}^{D_i,c}(\mathcal{F}_i)$ if $i \notin I$, in the presence of an adaptive, malicious t -adversary, assuming that all honest parties receive their inputs within $c + 1$ consecutive rounds.

Furthermore, the expected round complexity of the compiled protocol π' is

$$B_c \cdot \sum_{i \in [m]} d_i \cdot E[c_{\text{tr}}(T_i)] + 2 \cdot \sum_{i \in [m]} d_i \cdot E[\text{flex}_{\text{tr}}(T_i)],$$

where d_i is the expected number of calls in π to hybrid \mathcal{F}_i , T_i is a trace sampled from D_i , and $B_c = 3c + 1$ is the blow-up factor.

3.2. Fast Parallel Broadcast

Cohen et al. [16], based on Hirt and Zikas [39], defined the unfair parallel-broadcast functionality, in which the functionality informs the adversary which messages it received, and allows the adversary, based on this information, to corrupt senders and replace their input messages.

- UNFAIR PARALLEL BROADCAST. In the *unfair parallel-broadcast* functionality, each party P_i with input x_i distributes its input to all the parties. The adversary is

Protocol π_{PBC}

1. In the first round, upon receiving $(\text{input}, \text{sid}, x_i)$ with $x_i \in V$ from the environment, P_i secret shares x_i using a (t, n) error-correcting secret sharing scheme, $(x_i^1, \dots, x_i^n) \leftarrow \text{Share}(x_i)$. Next, P_i sends for every party P_j its share (sid, x_i^j) . Denote by \tilde{x}_j^i the value received from P_j (replace invalid/missing values by zero).
2. In the second round, P_i broadcasts the values $\mathbf{x}_i = (\tilde{x}_1^i, \dots, \tilde{x}_n^i)$ using the unfair parallel-broadcast functionality, i.e., P_i sends $(\text{input}, \text{sid}, \mathbf{x}_i)$ to $\mathcal{F}_{\text{UPBC}}$. Denote by $\mathbf{y}_j = (y_1^j, \dots, y_n^j)$ the value received from P_j (replace invalid/missing values by zero). Next, P_i reconstructs all the input values, i.e., for every $j \in [n]$ computes $y_j = \text{Recon}(y_j^1, \dots, y_j^n)$ (in case $y_j = \perp$ set $y_j \leftarrow 0$), and outputs $(\text{output}, \text{sid}, (y_1, \dots, y_n))$.

Fig. 1. The parallel-broadcast protocol, in the $(\mathcal{F}_{\text{PSMT}}, \mathcal{F}_{\text{UPBC}})$ -hybrid model.

allowed to learn the content of each input value from the leakage function (and so it can corrupt parties and change their messages prior to their distribution, based on this information). The function to compute is $f_{\text{UPBC}}(x_1, \dots, x_n, a) = ((x_1, \dots, x_n), \dots, (x_1, \dots, x_n))$ and the leakage function is $l_{\text{UPBC}}(x_1, \dots, x_n) = (x_1, \dots, x_n)$. We denote by $\mathcal{F}_{\text{UPBC}}$ the functionality \mathcal{F}_{CSF} when parametrized with the above functions f_{UPBC} and l_{UPBC} .

The protocol of Katz and Koo [45] realizes (a wrapped version of) $\mathcal{F}_{\text{UPBC}}$ when the parties have correlated-randomness setup. The following result follows.

Theorem 3.4. *Let $c \geq 0$ and $t < n/2$. There exists an efficiently sampleable distribution D such that the functionality $\mathcal{W}_{\text{sl-flex}}^{D,c}(\mathcal{F}_{\text{UPBC}})$ has an expected-constant-round complexity, and can be UC-realized in the $(\mathcal{F}_{\text{SMT}}, \mathcal{F}_{\text{CORR-BC}})$ -hybrid model, with information-theoretic security, in the presence of an adaptive, malicious t -adversary, assuming that all honest parties receive their inputs within $c + 1$ consecutive rounds.*

The parallel-broadcast functionality is similar to the unfair version, except that the adversary cannot corrupt parties based on the messages they send.

- **PARALLEL BROADCAST.** In the *parallel-broadcast* functionality, each party P_i with input x_i distributes its input to all the parties. Unlike the unfair version, the adversary only learns the length of the honest parties' messages before their distribution, i.e., the leakage function is $l_{\text{PBC}}(x_1, \dots, x_n) = (|x_1|, \dots, |x_n|)$. It follows that the adversary cannot use the leaked information in a meaningful way when deciding which parties to corrupt. The function to compute is identical to the unfair version, i.e., $f_{\text{PBC}}(x_1, \dots, x_n, a) = ((x_1, \dots, x_n), \dots, (x_1, \dots, x_n))$. We denote by \mathcal{F}_{PBC} the functionality \mathcal{F}_{CSF} when parametrized with the above functions f_{PBC} and l_{PBC} .

We next show how to realize the parallel-broadcast functionality \mathcal{F}_{PBC} in the $\mathcal{F}_{\text{UPBC}}$ -hybrid model, in the honest-majority setting. The construction follows [16], where the only difference is that for $t < n/2$, perfectly correct error-correcting secret sharing (see Definition A.1) cannot be achieved, and a negligible error probability is introduced. We describe this protocol, denoted π_{PBC} , in Fig. 1.

Theorem 3.5. *Let $c \geq 0$ and $t < n/2$. There exists an efficiently sampleable distribution D such that the functionality $\mathcal{W}_{\text{sl-flex}}^{D,c}(\mathcal{F}_{\text{PBC}})$ has an expected-constant-round complexity, and can be UC-realized in the $(\mathcal{F}_{\text{SMT}}, \mathcal{F}_{\text{CORR-BC}})$ -hybrid model, with information-theoretic security, in the presence of an adaptive malicious t -adversary, assuming that all honest parties receive their inputs within $c + 1$ consecutive rounds.*

The proof of the theorem follows in the same lines of the proof of [16, Theorem 5.6].

3.3. Fast SFE in the Point-to-Point Model

We conclude this section by showing how to construct a UC-secure SFE protocol which computes a given circuit in expected $O(d)$ rounds, independently of the number of parties, in the point-to-point channels model. The protocol is obtained by taking the protocol of Cramer et al. [21], denoted π_{SFE} . This protocol relies on (parallel) broadcast and (parallel) point-to-point channels, and therefore it can be described in the $(\mathcal{F}_{\text{PSMT}}, \mathcal{F}_{\text{PBC}})$ -hybrid model.

Theorem 3.6. *Let f be an n -party function, let C an arithmetic circuit computing f with multiplicative depth d , let $c \geq 0$, and let $t < n/2$. Then, there exists an efficiently sampleable distribution D such that the functionality $\mathcal{W}_{\text{sl-flex}}^{D,c}(\mathcal{F}_{\text{SFE}}^f)$ has round complexity $O(d)$ in expectation, and can be UC-realized in the $(\mathcal{F}_{\text{SMT}}, \mathcal{F}_{\text{CORR-BC}})$ -hybrid model, with information-theoretic security, in the presence of an adaptive, malicious t -adversary, assuming that all honest parties receive their inputs within $c + 1$ consecutive rounds.*

Using an adaptively secure constant-round MPC in the broadcast model, such as the protocol of Damgård and Ishai [23] (or the two-round protocol that exists under stronger assumptions [20]), we obtain expected-constant-round MPC over point-to-point channels.

Theorem 3.7. *Let f be an n -party function, let $c \geq 0$, let $t < n/2$, and assume that one-way functions exist. Then, there exists an efficiently sampleable distribution D such that the functionality $\mathcal{W}_{\text{sl-flex}}^{D,c}(\mathcal{F}_{\text{SFE}}^f)$ has round complexity $O(1)$ in expectation, and can be UC-realized in the $(\mathcal{F}_{\text{SMT}}, \mathcal{F}_{\text{PKI}})$ -hybrid model, with computational security, in the presence of an adaptive, malicious t -adversary, assuming that all honest parties receive their inputs within $c + 1$ consecutive rounds.*

4. Functionally Black-Box Protocols and Parallel Composition

In this section, we extend the probabilistic-termination framework [16] to capture the notions of functionally black-box protocols [57] and of parallel composition of canonical synchronous functionalities.

Functionally black-box protocols Most MPC protocols in the literature explicitly require a representation of the function to be computed, usually in the form a Boolean circuit, an arithmetic circuit, or a RAM program. This representation is a common param-

eter to the protocol and is known to all of the parties. Rosulek [57] asked whether this is inherent or if it possible to securely compute a function without “knowing its code.” Specifically, Rosulek [57] defined functionally black-box (FBB) protocols, in which the parties compute a function from within some function class, where each party is given only a local oracle access to the function (i.e., the party can query the oracle to learn the evaluation of the function on inputs of its choice). A feasibility result for semi-honest two-party FBB protocols was given in [57] for a class of functions related to blind signatures, as well as an impossibility result (that was strengthened in [43]) ruling out generic FBB protocols, also in the semi-honest case.

Looking ahead, in Sects. 5 and 6.2 we consider parallel composition of FBB computations. That is, given public function-classes $\mathcal{C}_1, \dots, \mathcal{C}_M$ and (hidden) functions $g_i \in \mathcal{C}_i$, each party is given a local oracle access to g_1, \dots, g_M (as in [57]), but in addition the parties can globally use M ideal functionalities to jointly compute each g_i ; every ideal functionality is parametrized by a function class \mathcal{C}_i and is given an oracle access to the corresponding g_i . Because the parties can use an ideal functionality to compute each g_i , the impossibility results from [43,57] do not apply, and indeed, in Sect. 5 we show that in the semi-honest setting it is possible to compose in parallel FBB computations in a round-preserving manner. In Sect. 6.2 we extend the technique from [43] to rule out round-preserving FBB parallel composition in the malicious setting.

We formalize the notion of functionally black-box protocols of Rosulek [57] in the language of canonical synchronous functionalities. As in [57], we focus on secure function evaluation. The SFE functionality $\mathcal{F}_{\text{SFE}}^g$ (see Sect. C.1), parametrized by an n -party function g , is defined as the CSF $\mathcal{F}_{\text{CSF}}^{f_{\text{SFE}}, l_{\text{SFE}}}$, where $f_{\text{SFE}}(x_1, \dots, x_n, a) = g(x_1, \dots, x_n)$ (i.e., computes the function g while ignoring the adversary’s input a) and the leakage function is $l_{\text{SFE}}(x_1, \dots, x_n) = (|x_1|, \dots, |x_n|)$. The following definition explains what we mean by a protocol that realizes the secure function evaluation functionality in a black-box way with respect to the function g .

Definition 4.1. Let $\mathcal{C} = \{g: (\{0, 1\}^*)^n \rightarrow (\{0, 1\}^*)^n\}$ be a class of n -party functions. Denote by $\mathcal{F}_{\text{SFE}}^{\mathcal{C}}$ the CSF, implemented as an (uninstantiated) oracle machine that in order to compute $f_{\text{SFE}}^{\mathcal{C}}(x_1, \dots, x_n, a)$, queries the oracle with (x_1, \dots, x_n) and stores the response (y_1, \dots, y_n) . The leakage function $l_{\text{SFE}}(x_1, \dots, x_n) = (|x_1|, \dots, |x_n|)$ is unchanged.

Then, a protocol $\pi = (\pi_1, \dots, \pi_n)$ is a functionally black-box (FBB) protocol for (a wrapped version of) $\mathcal{F}_{\text{SFE}}^{\mathcal{C}}$, if for every $f \in \mathcal{C}$, the protocol $\pi^f = (\pi_1^f, \dots, \pi_n^f)$ UC-realizes $\mathcal{F}_{\text{SFE}}^f$.

Parallel composition of CSFs The parallel composition of CSFs is defined in a natural way as the CSF that evaluates the corresponding functions in parallel.

Definition 4.2. Let f_1, \dots, f_M be n -input functions. We define the $(n \cdot M)$ -input function $(f_1 \parallel \dots \parallel f_M)$ as follows. Upon input (x_1, \dots, x_n) , where each x_i is an M -tuple (x_i^1, \dots, x_i^M) , the output is the M -tuple defined as

$$(f_1 \parallel \cdots \parallel f_M)(\mathbf{x}_1, \dots, \mathbf{x}_n) = \left((y_1^1, \dots, y_1^M), \dots, (y_n^1, \dots, y_n^M) \right),$$

where $(y_1^j, \dots, y_n^j) = f_j(x_1^j, \dots, x_n^j)$.

Let $\mathcal{F}_{\text{CSF}}^{f_1, l_1}, \dots, \mathcal{F}_{\text{CSF}}^{f_M, l_M}$ be CSFs and denote $\mathcal{F}_i = \mathcal{F}_{\text{CSF}}^{f_i, l_i}$. The parallel composition of $\mathcal{F}_1, \dots, \mathcal{F}_M$, denoted as $(\mathcal{F}_1 \parallel \cdots \parallel \mathcal{F}_M)$, is the CSF defined by the function $(f_1 \parallel \cdots \parallel f_M)$ and the leakage function $(l_1 \parallel \cdots \parallel l_M)$.

5. Round-Preserving Parallel Composition: Passive Security

In this section, we show that round-preserving parallel composition is feasible, in a functionally black-box manner, facing semi-honest adversaries. The underlying idea of our protocol π_{PFBB} (standing for **parallel functionally black-box**), formally presented in Fig. 2, is based on a simplified form of the parallel-broadcast protocol of Ben-Or and El-Yaniv [5]. The protocol proceeds in iterations, where in each iteration, the parties invoke, in parallel and using the same input values, sufficiently many instances of each (oracle-aided) ideal functionality, but only for a constant number of rounds. If some party received an output in at least one invocation of every ideal functionality, it distributes all output values and the protocol completes; otherwise, the protocol resumes with another iteration. This protocol retains privacy for deterministic functions with public output,⁶ since the adversary is semi-honest, and so corrupted parties will provide the same input values to all instances of each ideal functionality.

Intuitively, during the simulation of the protocol, the simulator should imitate every call for every ideal functionality toward the adversary. A subtle issue is that in order to do so, the simulator must know the exact trace that is sampled by each instance of each ideal functionality during the execution of the real protocol. Therefore, it is indeed essential for the simulator to receive the *random coins* used to sample the trace for the entire protocol, by the ideal functionality computing the parallel composition (see Sect. 2.2). By defining the trace-distribution sampler in a way that consists of all (potential) sub-traces for every instance of every ideal functionality, the simulator can induce the exact random coins used to sample the correct sub-trace for every ideal functionality that is invoked.

Theorem 5.1. *Let $\mathcal{C}_1, \dots, \mathcal{C}_M$ be deterministic function-classes with public output, let $\mathcal{F}_{\text{SFE}}^{\mathcal{C}_1}, \dots, \mathcal{F}_{\text{SFE}}^{\mathcal{C}_M}$ be oracle-aided secure function evaluation functionalities, and let $t < n/2$. Let D_1, \dots, D_M be distributions, such that for every $j \in [M]$, the round complexity of $\mathcal{W}_{\text{flex}}^{D_j}(\mathcal{F}_{\text{SFE}}^{\mathcal{C}_j})$ has expectation μ_j . Denote $\mu = \max\{\mu_1, \dots, \mu_M\}$.*

Then, there exists a distribution D with expectation $\mu' = O(\mu)$ such that $\mathcal{W}_{\text{flex}}^D(\mathcal{F}_{\text{SFE}}^{\mathcal{C}_1} \parallel \cdots \parallel \mathcal{F}_{\text{SFE}}^{\mathcal{C}_M})$ can be UC-realized by an FBB protocol in the $(\mathcal{F}_{\text{SMT}}, \mathcal{W}_{\text{flex}}^{D_1}(\mathcal{F}_{\text{SFE}}^{\mathcal{C}_1}), \dots, \mathcal{W}_{\text{flex}}^{D_M}(\mathcal{F}_{\text{SFE}}^{\mathcal{C}_M}))$ -hybrid model, with information-theoretic security, in the presence of an

⁶Although the result holds for deterministic functionalities, we note that using standard techniques every functionality can be transformed to an equivalent deterministic functionality in a black-box way.

Protocol $\pi_{\text{PFBB}}(\mathbb{N}, \mathbb{R}, \mathbb{L})$

1. Each party P_i , upon receiving $(\text{input}, \text{sid}, \mathbf{x}_i)$ with $\mathbf{x}_i = (x_i^1, \dots, x_i^M)$, initializes the iteration index $\alpha \leftarrow 0$.
2. Initiate (in parallel) \mathbb{N} instances of every ideal functionality $\mathcal{W}_{\text{flex}}^{D_j}(\mathcal{F}_{\text{SFE}}^{C_j})$ as follows:
 - (a) For every $j \in [\mathbb{M}]$ and $k \in [\mathbb{N}]$, send $(\text{input}, \text{sid}_{j, \tilde{k}}, \mathbf{x}_i^j)$ to $\mathcal{W}_{\text{flex}}^{D_j}(\mathcal{F}_{\text{SFE}}^{C_j})$ (where $\tilde{k} = \alpha \cdot \mathbb{N} + k$ and $\text{sid}_{j, \tilde{k}} = \text{sid} \circ j \circ \tilde{k}$).
 - (b) Execute for rounds $\rho = 1, \dots, \mathbb{R}$:
 - i. Every party P_i sends $(\text{fetch-output}, \text{sid}_{j, \tilde{k}})$ to $\mathcal{W}_{\text{flex}}^{D_j}(\mathcal{F}_{\text{SFE}}^{C_j})$.
 - ii. If received $(\text{output}, \text{sid}_{j, \tilde{k}}, y_j)$ with $y_j \neq \perp$, add y_j to the (initially empty) set S_i^j .
 - (c) If for every $j \in [\mathbb{M}]$, $S_i^j \neq \emptyset$, then set $\mathbf{y} = (y_1, \dots, y_M)$, where $y_j \in S_i^j$ is arbitrarily chosen, and send (sid, \mathbf{y}) to all parties.
 - (d) If P_i received a value (sid, \mathbf{y}) from some party, it outputs $(\text{output}, \text{sid}, \mathbf{y})$ and halts. Else, set $\alpha \leftarrow \alpha + 1$. If $\alpha < \mathbb{L}$, goto Step 2; else, output \perp and halt.

Fig. 2. FBB parallel composition in the $(\mathcal{F}_{\text{SMT}}, \mathcal{W}_{\text{flex}}^{D_1}(\mathcal{F}_{\text{SFE}}^{C_1}), \dots, \mathcal{W}_{\text{flex}}^{D_M}(\mathcal{F}_{\text{SFE}}^{C_M}))$ -hybrid model.

adaptive, semi-honest t -adversary, assuming that all honest parties receive their inputs at the same round.

In particular, if for every $j \in [\mathbb{M}]$, the expectation μ_j is constant, then μ' is constant.

The proof of Theorem 5.1 follows immediately from the following lemma.

Lemma 5.2. *Consider the notations in Theorem 5.1 and let $0 < \epsilon < 1$.*

Then, for any $R > \mu$, $N > \frac{\log(M/\epsilon)}{\log(R/\mu)}$ and $L = \text{poly}(\kappa)$, protocol $\pi_{\text{PFBB}}(\mathbb{N}, \mathbb{R}, \mathbb{L})$ is an FBB protocol for $\mathcal{W}_{\text{flex}}^{D_{\text{pfbb}}}(\mathcal{F}_{\text{SFE}}^{C_1} \parallel \dots \parallel \mathcal{F}_{\text{SFE}}^{C_M})$, for a distribution D_{pfbb} with expectation $\mu_{\text{pfbb}} = O(R)$, in the $(\mathcal{F}_{\text{SMT}}, \mathcal{W}_{\text{flex}}^{D_1}(\mathcal{F}_{\text{SFE}}^{C_1}), \dots, \mathcal{W}_{\text{flex}}^{D_M}(\mathcal{F}_{\text{SFE}}^{C_M}))$ -hybrid model, with information-theoretic security, in the presence of an adaptive, semi-honest t -adversary, assuming that all honest parties receive their inputs at the same round.

Proof We start by defining the sampling algorithm for the distribution D_{pfbb} , parametrized by \mathbb{N} , \mathbb{R} , \mathbb{L} , and distributions D_1, \dots, D_M . The sampler initially sets $\alpha \leftarrow 0$ and a trace T with a root labeled by $\mathcal{W}_{\text{flex}}^{D_{\text{pfbb}}}(\mathcal{F}_{\text{SFE}}^{C_1} \parallel \dots \parallel \mathcal{F}_{\text{SFE}}^{C_M})$ and no children. Next, independently sample traces $T_j^{\tilde{k}} \leftarrow D_j$, for $j \in [\mathbb{M}]$ and $k \in [\mathbb{N}]$ (where $\tilde{k} = \alpha \cdot \mathbb{N} + k$), and append

$$\left(\left(\left[\mathcal{W}_{\text{flex}}^{D_1}(\mathcal{F}_{\text{SFE}}^{C_1}) \right]^{\mathbb{N}}, \dots, \left[\mathcal{W}_{\text{flex}}^{D_M}(\mathcal{F}_{\text{SFE}}^{C_M}) \right]^{\mathbb{N}} \right) \right)^{\mathbb{R}}, \mathcal{F}_{\text{PSMT}} \right)$$

to the (initially empty) ordered set of leaves of the trace T (i.e., calling sequentially \mathbb{R} times, to \mathbb{N} parallel instances of each of the functionalities $\mathcal{W}_{\text{flex}}^{D_j}(\mathcal{F}_{\text{SFE}}^{C_j})$ (for every $j \in [\mathbb{M}]$), followed by a call to $\mathcal{F}_{\text{PSMT}}$). If for every $j \in [\mathbb{M}]$, there exists $k \in [\mathbb{N}]$, such

that $c_{\text{tr}}(T_j^{\tilde{k}}) < \mathbb{R}$, then output T and halt. Else, set $\alpha \leftarrow \alpha + 1$. If $\alpha < \mathbb{L}$, repeat the sampling process; otherwise output T and halt.

Following Lemma 2.3, for $\mathbb{R} > \mu$ and $\mathbb{N} > \frac{\log(M/c)}{\log(\mathbb{R}/\mu)}$, it holds that in every iteration, at least one invocation of $\mathcal{W}_{\text{flex}}^{D_j}(\mathcal{F}_{\text{SFE}}^{C_j})$ will produce output, for every $j \in [M]$, with a constant probability. It follows that the expected number of iterations until all honest parties receive output and the protocol terminates is constant, and since each iteration consists of $O(\mathbb{R})$ rounds, the entire execution completes within $O(\mathbb{R})$ rounds in expectation, as required. The failure probability that the protocol will not terminate within $\mathbb{L} = \text{poly}(\kappa)$ iterations is negligible.

Let \mathcal{A} be a semi-honest adversary, we now construct a simulator \mathcal{S} for \mathcal{A} . Initially, \mathcal{S} sets the values $\alpha \leftarrow 0$ and $\mathbf{y} \leftarrow \perp$, and starts by receiving leakage messages (leakage, sid , P_i , (l_1, \dots, l_M)) and a trace message (trace, sid , T) from $\mathcal{W}_{\text{flex}}^{D_{\text{ptfbb}}}(\mathcal{F}_{\text{SFE}}^{C_1} \parallel \dots \parallel \mathcal{F}_{\text{SFE}}^{C_M})$, where T is a depth-1 trace of the form

$$\left(\left(\left(\left[\mathcal{W}_{\text{flex}}^{D_1}(\mathcal{F}_{\text{SFE}}^{C_1}) \right]^{\mathbb{N}}, \dots, \left[\mathcal{W}_{\text{flex}}^{D_M}(\mathcal{F}_{\text{SFE}}^{C_M}) \right]^{\mathbb{N}} \right) \right)^{\mathbb{R}}, \mathcal{F}_{\text{PSMT}} \right)^q$$

(i.e., q iterations of calling sequentially \mathbb{R} times, to \mathbb{N} parallel instances of each $\mathcal{W}_{\text{flex}}^{D_j}(\mathcal{F}_{\text{SFE}}^{C_j})$, for every $j \in [M]$, followed by a call to $\mathcal{F}_{\text{PSMT}}$). More precisely, \mathcal{S} receives the coins that were used by the functionality $\mathcal{W}_{\text{flex}}^{D_{\text{ptfbb}}}(\mathcal{F}_{\text{SFE}}^{C_1} \parallel \dots \parallel \mathcal{F}_{\text{SFE}}^{C_M})$ to sample the trace T and using these coins, \mathcal{S} samples the same traces $T_j^{\tilde{k}}$ that were used to define T .

In order to simulate the α 'th iteration, \mathcal{S} sends the message (leakage, $\text{sid}_{j,\tilde{k}}$, P_i , l_j) to \mathcal{A} , for every $j \in [M]$, every $k \in [\mathbb{N}]$, and every honest P_i (where $\tilde{k} = \alpha \cdot \mathbb{N} + k$), and receives (input, $\text{sid}_{j,\tilde{k}}$, x_i^j) from \mathcal{A} on behalf of every corrupted party P_i . Since \mathcal{A} is semi-honest, it holds that the same x_i^j is used for each corrupted party P_i in all instances of the functionality $\mathcal{W}_{\text{flex}}^{D_j}(\mathcal{F}_{\text{SFE}}^{C_j})$. In the first iteration, \mathcal{S} sends to $\mathcal{W}_{\text{flex}}^{D_{\text{ptfbb}}}(\mathcal{F}_{\text{SFE}}^{C_1} \parallel \dots \parallel \mathcal{F}_{\text{SFE}}^{C_M})$, on behalf of every corrupted party P_i , the message (input, sid , \mathbf{x}_i), with $\mathbf{x}_i = (x_i^1, \dots, x_i^M)$. When \mathcal{A} sends (fetch-output, $\text{sid}_{j,\tilde{k}}$) requests in round ρ , \mathcal{S} answers with \perp if $c_{\text{tr}}(T_j^{\tilde{k}}) < \rho$, and with (output, $\text{sid}_{j,\tilde{k}}$, y_j) otherwise, where $\mathbf{y} = (y_1, \dots, y_M)$; if $\mathbf{y} = \perp$ (i.e., on the first time), send (early-output, sid , P_i) to $\mathcal{W}_{\text{flex}}^{D_{\text{ptfbb}}}(\mathcal{F}_{\text{SFE}}^{C_1} \parallel \dots \parallel \mathcal{F}_{\text{SFE}}^{C_M})$, receive (output, sid , \mathbf{y}) and store $\mathbf{y} = (y_1, \dots, y_M)$. In case \mathcal{A} sends (early-output, $\text{sid}_{j,\tilde{k}}$, P_i), for some corrupted P_i , send (output, $\text{sid}_{j,\tilde{k}}$, y_j) to \mathcal{A} . If the simulated protocol completes during the α 'th iteration, \mathcal{S} sends (early-output, sid , P_i) to the functionality $\mathcal{W}_{\text{flex}}^{D_{\text{ptfbb}}}(\mathcal{F}_{\text{SFE}}^{C_1} \parallel \dots \parallel \mathcal{F}_{\text{SFE}}^{C_M})$ on behalf of every party.

Proving that no environment can distinguish between its view in an execution of π_{PFBB} in the $(\mathcal{F}_{\text{SMT}}, \mathcal{W}_{\text{flex}}^{D_1}(\mathcal{F}_{\text{SFE}}^{C_1}), \dots, \mathcal{W}_{\text{flex}}^{D_M}(\mathcal{F}_{\text{SFE}}^{C_M}))$ -hybrid model, and its view when interacting with \mathcal{S} in the ideal computation of $\mathcal{W}_{\text{flex}}^{D_{\text{ptfbb}}}(\mathcal{F}_{\text{SFE}}^{C_1} \parallel \dots \parallel \mathcal{F}_{\text{SFE}}^{C_M})$, follows via a standard hybrid argument. Starting with the execution of π_{PFBB} , the invocations of the ideal functionalities $\mathcal{W}_{\text{flex}}^{D_1}(\mathcal{F}_{\text{SFE}}^{C_1}), \dots, \mathcal{W}_{\text{flex}}^{D_M}(\mathcal{F}_{\text{SFE}}^{C_M})$ are replaced, one-by-one, with the answers of

the simulator \mathcal{S} . Since \mathcal{S} perfectly emulates each such call using the trace it received from the ideal functionality $\mathcal{W}_{\text{flex}}^{D_{\text{pfb}}}$ ($\mathcal{F}_{\text{SFE}}^{\mathcal{C}_1} \parallel \dots \parallel \mathcal{F}_{\text{SFE}}^{\mathcal{C}_M}$), it follows immediately that the views of the environment in two neighboring hybrids are indistinguishable; therefore, the view of the environment in the simulation is indistinguishable from its view in the execution of π_{PFB} . \square

6. Round-Preserving Parallel Composition: Active Security

In this section, we consider security against active adversaries. First, in Sect. 6.1, we show how to compute the parallel composition of probabilistic-termination functionalities, in a round-preserving manner, using a black-box access to *protocols* realizing the individual functionalities. In Sect. 6.2, we investigate the question of whether there exists a *functionally* black-box round-preserving malicious protocol for the parallel composition of probabilistic-termination functionalities, and show that for a natural extension of protocols, following the techniques from [5], this is not the case—i.e., there exist functions such that no such protocol with black-box access to them can compute their parallel composition, in a round-preserving manner, tolerating even a single adversarial party.

6.1. Feasibility of Round-Preserving Parallel Composition

In this section, we show how to compile multiple protocols, realizing probabilistic-termination functionalities, into a single protocol that realizes the parallel composition of the functionalities, in a round-preserving manner, and while only using black-box access to the underlying protocols. We start by providing a high-level description of the compiler.

The compiler receives as input protocols π_1, \dots, π_M , where each protocol π_j is defined in the point-to-point model, in which the parties are given correlated randomness in a secure setup phase, i.e., in the $(\mathcal{F}_{\text{SMT}}, \mathcal{F}_{\text{CORR}}^{D_{\text{corr}}})$ -hybrid model.⁷ It follows that the next-message function for each party in each protocol is a deterministic function that receives the input value, correlated randomness, private randomness, and history of incoming messages, and outputs a vector of n messages to be sent in the following round (one message for each party); we denote by $f_{\text{next-msg}}^{\pi_j, i}$ the next-message function for party P_i in protocol π_j . In particular, we note that the entire transcript of the protocol is fixed once the input value, correlated randomness, and private randomness of each party are determined.

The underlying ideas of the compiler are inspired by the constructions in [5, 16], where a round-preserving parallel-broadcast protocol was constructed by iteratively running, for a constant number of rounds, multiple instances of BA protocols (each instance is executed multiple times in parallel), until at least one execution of every BA instance is completed. This approach is indeed suitable for computing “privacy-free” functionalities

⁷This captures, for example, broadcast-model protocols, where the broadcast channel is realized using an expected-constant-round protocol, see Theorems 3.6 and 3.7.

(such as broadcast), where the adversary may learn the result of multiple computations using the same inputs for honest parties, without compromising security. However, when considering the parallel composition of arbitrary functions, running two instances of a protocol *using the same input values* will violate privacy, since the adversary can use different inputs to learn multiple outputs of the function.

The parallel-composition compiler generalizes the above approach in a privacy-preserving manner. The compiler follows the GMW paradigm [33] and is defined in the Setup-Commit-then-Prove hybrid model [11, 42], which generates committed correlated randomness for the parties and ensures that all parties follow the protocol specification. This mechanism allows each party to commit to its input values and later execute multiple instances of each protocol, while proving that the same input value is used in all executions. For simplicity and without loss of generality, we assume that each function is deterministic and has a public output. In this case, it is ensured that if two parties receive output values in two executions of π_j , then they receive the same output value. The private random coins that are used in each execution only affect the termination round, but not the output value. Using this simplification, we can remove the leader-election phase from the output-agreement technique in [5, 16] and directly use the termination technique from Bracha [7].

Another obstacle is to recover from corruptions without increasing the round complexity. Indeed, in case some party misbehaves, e.g., by using different input values in different instances of the same protocol π_j , then the Setup-Commit-then-Prove functionality ensures that all honest parties will identify the cheating party. In this case, the parties cannot recover by, for example, backtracking and simulating the cheating party, as this will yield a round complexity that is linear in the number of parties. Furthermore, the protocol must resume in a way such that all instances of a specific protocol π_j will use the same input value that the identified corrupted party used throughout the protocol's execution until it misbehaved (since the cheating party might have learned an output value in one of the executed protocols).

To this end, we slightly adjust the Setup-Commit-then-Prove functionality and secret-share every committed random string r_i (the correlated randomness for party P_i) among all the parties, using an error-correcting secret-sharing scheme (see Sect. A.1). Note that this can be done information theoretically as we assume an honest majority [12, 22, 56]. In case a party is identified as cheating, every party broadcasts the share of the committed randomness corresponding to that party, reconstructs the correlated randomness for that party, and from that point onwards, locally computes the messages corresponding to this party in every instance of every protocol. Using this approach, every round in the original protocols π_1, \dots, π_M is expanded by a constant number of rounds, and the overall round complexity is preserved.

We prove the following theorem.

Theorem 6.1. *Let $\mathcal{F}_1, \dots, \mathcal{F}_M$ be CSFs, let $t < n/2$, and let $c \geq 1$. Let π_1, \dots, π_M be SNF protocols such that for every $j \in [M]$, protocol π_j UC-realizes $\mathcal{W}_{\text{flex}}^{D_j}(\mathcal{F}_j)$ with expected round complexity μ_j and information-theoretic security, in the $(\mathcal{F}_{\text{SMT}}, \mathcal{F}_{\text{CORR}}^{D_j^{\text{corr}}})$ -hybrid model (for a distribution D_j and a distribution D_j^{corr} in NC^0), in the presence of*

an adaptive, malicious t -adversary, assuming that all honest parties receive their inputs at the same round. Denote $\mu = \max\{\mu_1, \dots, \mu_M\}$.

Then, $\mathcal{W}_{\text{sl-flex}}^{D,c}(\mathcal{F}_1 \parallel \dots \parallel \mathcal{F}_M)$, for some distribution D with expectation $\mu' = O(\mu)$, can be UC-realized with information-theoretic security by a protocol π in the $(\mathcal{F}_{\text{SMT}}, \mathcal{F}_{\text{CORR-BC}})$ -hybrid model, in the same adversarial setting, assuming that all honest parties receive their inputs within $c + 1$ consecutive rounds. In addition, protocol π requires only black-box access to the protocols π_1, \dots, π_M .

In particular, if for every $j \in [M]$, the expectation μ_j is constant, then μ' is constant.

Proof (sketch) Without loss of generality, we assume that every CSF \mathcal{F}_j is deterministic and has public output. The proof for randomized functionalities with private output follows using standard techniques. In Lemma 6.3, we prove that the compiled protocol $\pi_{\text{PBB}} = \text{Comp}(\pi_1, \dots, \pi_M, N, R, L)$, for $R > \mu$, $N > \frac{\log(M/\epsilon)}{\log(R/\mu)}$ and $L = \text{poly}(\kappa)$, UC-realizes $\mathcal{W}_{\text{flex}}^{D_{\text{pbb}}}(\mathcal{F}_1 \parallel \dots \parallel \mathcal{F}_M)$ with expected round complexity $O(R)$ and information-theoretic security, in the $(\mathcal{F}_{\text{PBC}}, \mathcal{F}_{\text{SCP}})$ -hybrid model. In Lemma 6.2, we show that a wrapped version of $\mathcal{F}_{\text{SCP}}(\mathcal{P}, \mathcal{D}_{\text{parallel}}(\pi_1, \dots, \pi_M, N \cdot L, R, \ell), \mathcal{R}_{\text{parallel}}, \Pi)$ (explained below) can be implemented, such that every call can be UC-realized in the $(\mathcal{F}_{\text{PSMT}}, \mathcal{F}_{\text{PBC}})$ -hybrid model with constant round complexity and information-theoretic security. Following Theorem 3.5, \mathcal{F}_{PBC} can be UC-realized in the \mathcal{F}_{SMT} -hybrid model with expected-constant-round complexity and information-theoretic security. The proof follows from the sequential composition theorems, Theorems 3.1, 3.2 and 3.3. \square

We now proceed to define the Setup-Commit-Then-Prove Functionality \mathcal{F}_{SCP} and prove Lemma 6.2 in Sect. 6.1.1, and to prove Lemma 6.3 in Sect. 6.1.2.

6.1.1. The Setup-Commit-Then-Prove Functionality

An important building block in our parallel-composition compiler is the Setup-Commit-then-Prove functionality. This functionality is used in order to allow parties to execute multiple instances of a protocol, using the same inputs, while ensuring input consistency. In addition, in case a party misbehaves and tries to deviate from the protocol or to use different inputs in different executions, the functionality allows all parties to identify this misbehavior and recover, while increasing the round complexity only by a constant factor. This functionality was defined by Ishai et al. [42], based on the Commit-then-Prove functionality of Canetti et al. [11], and was used in order to compile any semi-honest protocol into a protocol that is secure with identifiable abort, facing malicious adversaries, by generating committed setup for the parties and allow them to prove NP-statements in zero knowledge.

The Setup-Commit-then-Prove functionality \mathcal{F}_{SCP} is a reactive functionality. (The notion of CSF is extended to the reactive setting in ‘‘Appendix C.2.’’) The functionality is formally defined in Fig. 3 and is parametrized by a party-set \mathcal{P} , a distribution \mathcal{D} , a vector of n NP-relations \mathcal{R} , and a (t, n) error-correcting secret-sharing scheme Π . In the first call to the functionality, the parties do not send inputs (more precisely, send the empty input λ). The functionality samples correlated randomness $(r_1, \dots, r_n) \leftarrow \mathcal{D}$, hands r_i to P_i and stores r_i as the committed witness for P_i . In addition, the functionality secret shares r_i between all parties. All subsequent calls are used to prove NP-statements on

Functionality $\mathcal{F}_{\text{SCP}}(\mathcal{P}, \mathcal{D}, \vec{\mathcal{R}} = (\mathcal{R}_1, \dots, \mathcal{R}_n), \Pi)$

The functionality \mathcal{F}_{SCP} is a reactive CSF, with vectors $\mathbf{f}_{\text{SCP}} = (f_{\text{SCP}}^1, \dots, f_{\text{SCP}}^q)$ and $\mathbf{l}_{\text{SCP}} = (l_{\text{SCP}}^1, \dots, l_{\text{SCP}}^q)$ of functions (defined below). \mathcal{F}_{SCP} is parametrized by party-set $\mathcal{P} = \{P_1, \dots, P_n\}$, a distribution \mathcal{D} , a vector $\vec{\mathcal{R}}$ of NP relations, and a (t, n) error-correcting secret-sharing scheme $\Pi = (\text{Share}, \text{Recon})$.

Setup-Commit Phase: The function $f_{\text{SCP}}^1(x_{1,1}, \dots, x_{n,1}, a_1)$ is defined as follows. Initially, sample $(r_1, \dots, r_n) \leftarrow \mathcal{D}$ and for every $i \in [n]$, compute $(v_i^1, \dots, v_i^n) \leftarrow \text{Share}(r_i)$. The output for party P_i is $y_{i,1} = (r_i, v_i^1, \dots, v_i^n)$. The leakage function is $l_{\text{SCP}}^1(x_{1,1}, \dots, x_{n,1}) = \lambda$.

Prove Phase: For $k > 1$, the function $f_{\text{SCP}}^k(x_{1,k}, \dots, x_{n,k}, a_k)$ is defined as follows. For every $i \in [n]$, check if the relation $\mathcal{R}_i(x_{i,k}, r_i) = 1$. If so, set $b_i = 1$; else, set $b_i = 0$. The output for every party P_i is $y_{i,k} = ((x_{1,k}, b_1), \dots, (x_{n,k}, b_n))$. The leakage function is $l_{\text{SCP}}^k(x_{1,1}, \dots, x_{n,k}) = (x_{1,k}, \dots, x_{n,k})$.

Fig. 3. The Setup-Commit-then-Prove functionality.

the committed witnesses, i.e., on the k 'th call, for $k > 1$, P_i sends as its input a statement $x_{i,k}$; the functionality verifies whether $\mathcal{R}_i(x_{i,k}, r_i) = 1$ and sends $x_{i,k}$ and the result to all parties.

Ishai et al. [42] showed how to realize a slightly reduced version of the Setup-Commit-then-Prove functionality (in which the functionality does not secret share the witnesses), unconditionally, with identifiable abort, facing an arbitrary number of corrupted parties. Their protocol is based on the ‘‘MPC-in-the-head’’ approach that was put forth by Ishai et al. [40] with the goal of using information-theoretic MPC to construct a zero-knowledge protocol. In a nutshell, given a statement x and a witness w , the prover secret shares the witness as $\omega = \omega_1 \oplus \dots \oplus \omega_m$ (for some m) and emulates in its head an m -party t -secure MPC protocol with perfect correctness (e.g., BGW [6]) where the i 'th party has input (ω_i, x) . The MPC task is to reconstruct the witness $\omega = \omega_1 \oplus \dots \oplus \omega_m$, evaluate the relation $\mathcal{R}(x, \omega)$ and provide each party with the resulting bit. Next, the prover sends to the verifier a commitment to the local views of each of the virtual m parties (consisting of the input ω_i , the random coins, and the incoming messages). In turn, the verifier asks the prover to open t of these commitments, and verifies that these views are consistent and that the output is 1. Intuitively, completeness holds by the correctness of the MPC, soundness holds by the perfect security, and zero knowledge since the views of t parties leak no information about other inputs; so the overall security of the zero-knowledge protocol reduces to that of the commitment scheme.

In the following lemma, we adjust and simplify the protocol from [42] for the honest-majority setting, and show how to realize \mathcal{F}_{SCP} with guaranteed output delivery.

Lemma 6.2. *Let $\mathcal{R} = (\mathcal{R}_1, \dots, \mathcal{R}_n)$ be a vector of NP-relations, let \mathcal{D} be an efficiently sampleable distribution in NC^0 , and let $t < n/2$. Consider the representation of the reactive CSF, $\mathcal{F}_{\text{SCP}}(\mathcal{P}, \mathcal{D}, \mathcal{R}, \Pi)$, as a sequence of (non-reactive) CSFs $(\mathcal{F}_{\text{CSF}}^{\tilde{f}_{\text{SCP}}^1, l_{\text{SCP}}^1}, \dots, \mathcal{F}_{\text{CSF}}^{\tilde{f}_{\text{SCP}}^q, l_{\text{SCP}}^q})$, where \tilde{f}_{SCP}^j is defined as in Definition C.1. Denote $\mathcal{F}_{\text{SCP}}^j = \mathcal{F}_{\text{CSF}}^{\tilde{f}_{\text{SCP}}^j, l_{\text{SCP}}^j}$.*

Then, there exists a vector of distributions $\mathbf{D} = (D_1, \dots, D_q)$ such that for every $j \in [q]$, the wrapped functionality $\mathcal{W}_{\text{strict}}^{D_j}(\mathcal{F}_{\text{SCP}}^j)$ can be UC-realized in the $(\mathcal{F}_{\text{PSMT}}, \mathcal{F}_{\text{PBC}})$ -

hybrid model with constant-round complexity and information-theoretic security, in the presence of an adaptive, malicious t -adversary.

Proof (sketch) We start with a high-level description of the protocol in [42], for a single prover that proves a single statement (the extension to the multi-instance version of many provers that prove many statements follows via the JUC theorem [9]). As this protocol is designed for the dishonest-majority setting, it does not guarantee output delivery, but achieves identifiable abort.

As mentioned above, the protocol follows the “MPC-in-the-head” approach [40,41], where the prover emulates in its head a protocol, where m servers, each has as input a share of the witness ω , compute for a public statement x the function $b = \mathcal{R}(x, \omega)$, and output (x, b) . The prover publicly commits to the view of each server, and the verifiers challenge the prover to open the views of some of the servers. The verifiers accept the statement x , if and only if all opened views are consistent.

For simplicity, consider an ideal functionality $\mathcal{F}_{\text{CORR}}^{\text{ISCP}}$ for generating the correlated randomness to the parties. As we argue later, this functionality can be realized in constant rounds in the broadcast model via generic MPC. (Looking ahead, the distribution \mathcal{D} will consist of n independent copies of the following values, one for each party that acts as the prover.) The parties receive the following correlated randomness from $\mathcal{F}_{\text{CORR}}^{\text{ISCP}}$:

- The prover receives signed secret shares of the witness $(\omega_i, \sigma(\omega_i))$, for $i \in [n]$, where $(\omega_1, \dots, \omega_n)$ are shares of the prover’s witness ω , and $\sigma(\omega_i)$ is an information-theoretic signature of ω_i .
- The prover also receives random strings v_1, \dots, v_m along with corresponding signatures $\sigma(v_1), \dots, \sigma(v_m)$, that will be used for committing to the server’s views in the m -party protocol.
- Every P_i receives a challenge string c_i along with an information-theoretic signature $\sigma(c_i)$.
- Every party P_i receives a verification key for each of the signature values (note that all of the signing keys are hidden from the parties).

The prove phase, consists of three rounds:

1. The prover emulates in its head the protocol for computing $b = \mathcal{R}(x, \omega_1 \oplus \dots \oplus \omega_m)$ and broadcasts, for every $j \in [m]$, a commitment to the view of the j ’th server as $\text{view}_j \oplus v_j$ along with $\sigma(v_j)$.
2. Every party P_i broadcasts the committed random string c_i along with its commitment $\sigma(c_i)$. All parties locally compute $c = \sum c_j$ and use it to choose a subset $\mathcal{J} \subseteq [m]$. In case some party P_i sends invalid values, all parties identify P_i as corrupted and abort.
3. The prover broadcasts view_j and $\sigma(\omega_j)$, for every $j \in \mathcal{J}$, and all parties validate consistency.

Since in this work we consider an honest majority, we can simplify the protocol and achieve security with guaranteed output delivery. In the second round, instead of having each P_i broadcasting the committed randomness $(c_i, \sigma(c_i))$, we have the parties jointly compute the XOR function to agree on a subset $\mathcal{J} \subseteq [m]$; that is, each party enters a (locally chosen) random string as its input and receives the XOR of all strings

(if a party does not input a value it is simply ignored and treated as 0). Since this functionality can be represented using a constant-depth circuit, we can use the protocol of Cramer et al. [21], whose round complexity is $O(d)$, where d is the depth of the circuit, and provides information-theoretic security in the broadcast-hybrid model against an adaptive malicious t -adversary. A second modification we require is that $\mathcal{F}_{\text{CORR}}^{\text{ISCP}}$ secret shares the witness ω between all parties using an error-correcting secret-sharing scheme, which can be easily achieved in the honest-majority setting.

It is left to show that the sampling algorithm for the correlated randomness that is used in [42] can be represented using a constant-depth circuit; the proof will then follow using again the protocol from [21]. By assumption, sampling a value from the distribution \mathcal{D} can be done using a constant-depth circuit. In addition, the information-theoretic commitments in [42] are computed using information-theoretic signatures (see ‘‘Appendix A.2’’), such that *no party knows the signing key*. This means that the adversary cannot generate signatures on its own. In addition, each signature will only need to be verified once (by each party). Using the information-theoretic signatures construction from [60] (see Theorem A.3), the degree of the polynomial, that is used to generate the signature, is bounded by the number of signatures the adversary is allowed to see from each honest party, which in our case is constant. We conclude that the randomness-sampling algorithm can be represented using a constant-depth circuit. \square

In the following, we will consider the distribution $\mathcal{D}_{\text{parallel}}(\pi_1, \dots, \pi_M, q, R, \ell)$ for protocols π_1, \dots, π_M , where each π_j is defined in the $(\mathcal{F}_{\text{PSMT}}, \mathcal{F}_{\text{CORR}}^{D_j^{\text{corr}}})$ -hybrid model, that prepares (at most) q executions of each protocol, for exactly R rounds. For each of the q instances of every protocol π_j , the correlated randomness consists of three parts: first, sample correlated randomness for π_j from the distribution D_j^{corr} ; next, sample independent random coins (local for each party) from the corresponding distribution D_{π_j} , suitable for R rounds; finally, sample random coins that are used to mask all the communication (explained below). The parameter $\ell = \text{poly}(\kappa)$ represents the maximum between the input length and the maximal message length in the protocols π_1, \dots, π_M . The distribution is formally defined in Fig. 4.

The masking of the communication in the ρ 'th round of the k 'th instance of protocol π_j is performed as follows:

- When P_i wants to send messages (m_i^1, \dots, m_i^n) (where m_i^u is sent privately to P_u), party P_i sets for every $u \in [n]$, $\tilde{m}_i^u = m_i^u \oplus r_{i,j,k,\rho,u}^{\text{mask}}$ and broadcasts $\tilde{\mathbf{m}}_i = (\tilde{m}_i^1, \dots, \tilde{m}_i^n)$.
- When P_i receives messages $\tilde{\mathbf{m}}_u = (\tilde{m}_u^1, \dots, \tilde{m}_u^n)$ from P_u , P_i computes $m_u^i = \tilde{m}_u^i \oplus r_{u,j,k,\rho,i}^{\text{mask}}$, and uses m_u^i as the message P_u sent him.

The vector of relations $\mathcal{R}_{\text{parallel}} = (\mathcal{R}_{\text{parallel}}^1, \dots, \mathcal{R}_{\text{parallel}}^n)$, described below, will be used to verify that every party sends its messages in multiple executions of a protocol according to the specification of the protocol, while using the same input value in all executions. Formally, for every $i \in [n]$, the relation $\mathcal{R}_{\text{parallel}}^i$ consists of pairs $((\alpha, \rho, \tilde{\mathbf{m}}, \tilde{\mathbf{h}}, \mathbf{r}_i)$, satisfying:

- α is an integer representing the iteration number.

Distribution $\mathcal{D}_{\text{parallel}}(\pi_1, \dots, \pi_M, q, \mathbb{R}, \ell)$

The distribution $\mathcal{D}_{\text{parallel}}$ is parametrized by protocols π_1, \dots, π_M , where each π_j is defined in the $(\mathcal{F}_{\text{PSMT}}, \mathcal{F}_{\text{CORR}}^{D_j^{\text{corr}}})$ -hybrid model, and integers q, \mathbb{R}, ℓ . Denote by $D_{\pi_j}(\mathbb{R})$ the distribution that samples independent random coins for each party for \mathbb{R} rounds of protocol π_j .

1. For every $j \in [M]$:

(a) For every $k \in [q]$:

i. Sample correlated randomness for the k 'th instance of protocol π_j ,

$$(r_{1,j,k}^{\text{corr}}, \dots, r_{n,j,k}^{\text{corr}}) \leftarrow D_j^{\text{corr}}.$$

ii. Sample independent random coins for \mathbb{R} rounds in the k 'th instance of protocol π_j ,

$$(r_{1,j,k}^{\text{prot}}, \dots, r_{n,j,k}^{\text{prot}}) \leftarrow D_{\pi_j}(\mathbb{R}).$$

iii. Sample randomness $(r_{1,j,k}^{\text{mask}}, \dots, r_{n,j,k}^{\text{mask}})$ to mask the communication in the k 'th instance of π_j , where for every $i \in [n]$, $r_{i,j,k}^{\text{mask}}$ is set as follows:

A. For every $\rho \in [\mathbb{R}]$ and $u \in [n]$, sample $r_{i,j,k,\rho,u}^{\text{mask}} \leftarrow \{0,1\}^\ell$ (used to mask the message from P_i to P_u in the ρ 'th round of the k 'th execution of π_j),

B. For every $\rho \in [\mathbb{R}]$, set $r_{i,j,k,\rho}^{\text{mask}}$ as $(r_{i,j,k,\rho,1}^{\text{mask}}, \dots, r_{i,j,k,\rho,n}^{\text{mask}})$ and $(r_{1,j,k,\rho,i}^{\text{mask}}, \dots, r_{n,j,k,\rho,i}^{\text{mask}})$.

C. Set $r_{i,j,k}^{\text{mask}} = (r_{i,j,k,1}^{\text{mask}}, \dots, r_{i,j,k,\mathbb{R}}^{\text{mask}})$.

(b) For every $i \in [n]$, denote the random coins as $r_{i,j}^{\text{corr}} = (r_{i,j,1}^{\text{corr}}, \dots, r_{i,j,q}^{\text{corr}})$, $r_{i,j}^{\text{prot}} = (r_{i,j,1}^{\text{prot}}, \dots, r_{i,j,q}^{\text{prot}})$ and $r_{i,j}^{\text{mask}} = (r_{i,j,1}^{\text{mask}}, \dots, r_{i,j,q}^{\text{mask}})$; in addition, denote $\mathbf{r}_i^j = (r_{i,j}^{\text{corr}}, r_{i,j}^{\text{prot}}, r_{i,j}^{\text{mask}})$.

2. For every $i \in [n]$, sample randomness $\mathbf{r}_i^{\text{input}} \leftarrow (\{0,1\}^\ell)^M$ to mask the input value and denote $\mathbf{r}_i = (r_i^{\text{input}}, \mathbf{r}_i^1, \dots, \mathbf{r}_i^M)$.

3. Return $(\mathbf{r}_1, \dots, \mathbf{r}_n)$.

Fig. 4. The correlated-randomness distribution for parallel composition.

- ρ is an integer representing the round number.
- The vector of messages $\tilde{\mathbf{m}} = (\tilde{\mathbf{m}}_{i,1,\rho}, \dots, \tilde{\mathbf{m}}_{i,M,\rho})$ is structured such that for every $j \in [M]$, $\tilde{\mathbf{m}}_{i,j,\rho} = (\tilde{\mathbf{m}}_{i,j,1,\rho}, \dots, \tilde{\mathbf{m}}_{i,j,N,\rho})$, and for every $k \in [N]$, $\tilde{\mathbf{m}}_{i,j,k,\rho} = (\tilde{m}_{i,j,k,\rho,1}, \dots, \tilde{m}_{i,j,k,\rho,n})$ represents the message P_i broadcasts in the ρ 'th round of the k 'th execution of protocol π_j , in the α 'th iteration.
- The vector of history-messages $\tilde{\mathbf{h}} = (\tilde{\mathbf{h}}^{\text{input}}, \tilde{\mathbf{h}}_1, \dots, \tilde{\mathbf{h}}_M)$ is structured such that $\tilde{\mathbf{h}}^{\text{input}} = (\tilde{\mathbf{m}}_1^{\text{input}}, \dots, \tilde{\mathbf{m}}_n^{\text{input}})$ and for $j \in [M]$, $\tilde{\mathbf{h}}_j = (\tilde{h}_{j,1}, \dots, \tilde{h}_{j,N})$, and each $\tilde{h}_{j,k}$ represents the history-transcript until the ρ 'th round of the k 'th execution of protocol π_j , in the α 'th iteration.
- The random coins are $\mathbf{r}_i = (r_i^{\text{input}}, \mathbf{r}_i^1, \dots, \mathbf{r}_i^M)$, with $\mathbf{r}_i^j = (r_{i,j}^{\text{corr}}, r_{i,j}^{\text{prot}}, r_{i,j}^{\text{mask}})$, as defined in $\mathcal{D}_{\text{parallel}}$.
- For every $j \in [M]$ and $k \in [N]$, denote by $\mathbf{m}_{i,j,k,\rho} = (m_{i,j,k,\rho,1}, \dots, m_{i,j,k,\rho,n})$ the unmasked messages, where for every $u \in [n]$, $m_{i,j,k,\rho,u} = \tilde{m}_{i,j,k,\rho,u} \oplus r_{i,\tilde{k},j,\rho,u}^{\text{mask}}$, with $\tilde{k} = \alpha \cdot N + k$. Similarly, denote by $h_{i,j,k}$ the unmasked history obtained from $\tilde{h}_{i,j,k}$ using the corresponding randomness in $r_{i,j,\tilde{k}}^{\text{mask}}$. Denote $(x_i^1, \dots, x_i^M) = \tilde{\mathbf{m}}_i^{\text{input}} \oplus \mathbf{r}_i^{\text{input}}$. Then, for every j and k , it holds that

$$\mathbf{m}_{i,j,k,\rho} = f_{\text{nxt-msg}}^{\pi_j,i}(x_i^j, r_{i,j,\bar{k}}^{\text{corr}}, r_{i,j,\bar{k}}^{\text{prot}}, h_{i,j,k}).$$

That is, $\mathbf{m}_{i,j,k,\rho}$ is the output of the next-message function of P_i in protocol π_j on input x_i^j , correlated randomness $r_{i,j,\bar{k}}^{\text{corr}}$, private randomness $r_{i,j,\bar{k}}^{\text{prot}}$ and history $h_{i,j,k}$.

In the sequel, for simplicity, we will denote by \mathcal{F}_{SCP} the functionality

$$\mathcal{F}_{\text{SCP}}(\mathcal{P}, \mathcal{D}_{\text{parallel}}(\pi_1, \dots, \pi_M, N \cdot L, R, \ell), \mathcal{R}_{\text{parallel}}, \Pi).$$

6.1.2. Round-Preserving Parallel-Composition Compiler

We are now ready to present our protocol-black-box (PBB) parallel-composition compiler, formally described in Fig. 5.

Lemma 6.3. *Consider the notation in Theorem 6.1 and let $0 < \epsilon < 1$.*

Then, for parameters $R > \mu$, $N > \frac{\log(M/\epsilon)}{\log(R/\mu)}$, and $L = \text{poly}(\kappa)$, the compiled protocol $\pi_{\text{PBB}} = \text{Comp}(\pi_1, \dots, \pi_M, N, R, L)$ UC-realizes $\mathcal{W}_{\text{flex}}^{D_{\text{pbb}}}(\mathcal{F}_1 \parallel \dots \parallel \mathcal{F}_M)$ with expected round complexity $\mu_{\text{pbb}} = O(R)$, in the $(\mathcal{F}_{\text{PBC}}, \mathcal{F}_{\text{SCP}})$ -hybrid model and same adversarial setting, assuming that all honest parties receive their inputs at the same round. In addition, the compiler requires only black-box access to the protocols π_1, \dots, π_M .

To prove security of the construction, we construct a simulator for the dummy adversary, which simulates the functionality \mathcal{F}_{SCP} and all honest parties. At a high level, the simulation proceeds as follows. Since every protocol π_j realizes \mathcal{F}_j , there exists a simulator \mathcal{S}_j for the dummy adversary. In order to simulate the k 'th instance of each protocol π_j , the simulator \mathcal{S} invokes an instance of \mathcal{S}_j , denoted \mathcal{S}_j^k , and receives correlated randomness $\tilde{r}_{i,j,k}^{\text{corr}}$ for every corrupted party P_i . The simulator \mathcal{S} samples randomness from the distribution $\mathcal{D}_{\text{parallel}}$, adjusts the correlated randomness for every corrupted party accordingly and hands the adversary $(\mathbf{r}_i, v_1^i, \dots, v_n^i)$ as the answer from the first call to \mathcal{F}_{SCP} , where (v_1^i, \dots, v_n^i) are shares of zero, for every $i \in [n]$. For the input-commitment message, the simulator broadcasts commitments of zero for the honest parties (i.e., random messages). The k 'th instance of π_j is simulated now using \mathcal{S}_j^k , where \mathcal{S} masks/unmasks the messages between \mathcal{A} and \mathcal{S}_j^k , appropriately. Every message sent by \mathcal{A} on behalf of a corrupted party P_i is validated by \mathcal{S} according to the relation $\mathcal{R}_{\text{parallel}}^i$, and in case it is invalid, \mathcal{S} locally computes the messages for P_i , using its input and correlated randomness. In case party P_i gets corrupted, the simulator corrupts the dummy party in the ideal computation and learns its input; next, \mathcal{S} hands the input to each simulator \mathcal{S}_j^k , receives the random coins for P_i in each instance of π_j , updates the random coins \mathbf{r}_i accordingly and hands it to \mathcal{A} . This is a valid simulation for the dummy adversary, following the security guarantees of each simulator \mathcal{S}_j and of the secret-sharing scheme Π .

Proof For simplicity of notation, denote $\mathcal{F}_{\text{PBB}} = \mathcal{W}_{\text{flex}}^{D_{\text{pbb}}}(\mathcal{F}_1 \parallel \dots \parallel \mathcal{F}_M)$. We start by defining the sampling algorithm for the distribution D_{pbb} , parametrized by N, R, L and distributions D_1, \dots, D_M . The sampler initially sets $\alpha \leftarrow 0$ and a trace T with root

labeled by \mathcal{F}_{PBB} and children $(\mathcal{F}_{\text{SCP}}^1, \mathcal{F}_{\text{PBC}})$. Next, independently sample traces $T_j^{\tilde{k}} \leftarrow D_j$, for $j \in [M]$ and $k \in [N]$ (where $\tilde{k} = \alpha \cdot N + k$), and append

$$\left(\left(\mathcal{F}_{\text{SCP}}^i, \mathcal{F}_{\text{PBC}} \right)^R, \mathcal{F}_{\text{PBC}}, \mathcal{F}_{\text{PBC}}, \mathcal{F}_{\text{PBC}} \right)$$

to the leaves of the trace T (i.e., calling R times, sequentially, to $(\mathcal{F}_{\text{SCP}}^i, \mathcal{F}_{\text{PBC}})$, followed by three sequential calls to \mathcal{F}_{PBC}). If for every $j \in [M]$, there exists $k \in [N]$ such that $c_{\text{tr}}(T_j^{\tilde{k}}) < R$, then output T and halt. Else, set $\alpha \leftarrow \alpha + 1$. If $\alpha < L$, repeat the sampling process; otherwise output T and halt.

Following Lemma 2.3, for $R > \mu$ and $N > \frac{\log(M/\epsilon)}{\log(R/\mu)}$, it holds that in every iteration, at least one execution of π_j will produce output, for every $j \in [M]$, with a constant probability. It follows that the expected number of iterations until all honest parties receive output and the protocol terminates is constant, and since each iteration consists of R rounds, the entire execution completes within expected $O(R)$ rounds, as required. The failure probability that the protocol will not terminate within $L = \text{poly}(\kappa)$ iterations is negligible.

We construct a simulator \mathcal{S} for the dummy adversary \mathcal{A} . Let \mathcal{Z} be an environment. The simulator \mathcal{S} uses, in a black-box way, the simulators \mathcal{S}_j , for $j \in [M]$, where every \mathcal{S}_j simulates the dummy adversary for protocol π_j . The simulators \mathcal{S}_j are guaranteed to exist since every π_j UC-realizes \mathcal{F}_j . Each simulator \mathcal{S}_j is invoked (at most) $q = L \cdot N$ times; denote by \mathcal{S}_j^k the k 'th invocation of \mathcal{S}_j .

We consider the representation of the reactive CSF \mathcal{F}_{SCP} as a sequence of (non-reactive) CSFs $(\mathcal{F}_{\text{CSF}}^{\tilde{j}_{\text{SCP}}^1, l_{\text{SCP}}^1}, \dots, \mathcal{F}_{\text{CSF}}^{\tilde{j}_{\text{SCP}}^{q'}, l_{\text{SCP}}^{q'}})$, where \tilde{j}_{SCP}^j is defined as in Definition C.1 and $q' = L \cdot R + 1$. Denote $\mathcal{F}_{\text{SCP}}^j = \mathcal{F}_{\text{CSF}}^{\tilde{j}_{\text{SCP}}^j, l_{\text{SCP}}^j}$.

The simulator \mathcal{S} proceeds as follows:

- Initially, set $\mathbf{x}_1 = \dots = \mathbf{x}_n = \perp$ and $\mathbf{y} = \perp$.
- Simulating the first call to functionality \mathcal{F}_{SCP} :
 1. Sample correlated randomness $(\mathbf{r}_1, \dots, \mathbf{r}_n) \leftarrow \mathcal{D}_{\text{parallel}}(\pi_1, \dots, \pi_M, q, R, \ell)$, where $q = N \cdot L$ and $\mathbf{r}_i = (\mathbf{r}_i^{\text{input}}, \mathbf{r}_i^1, \dots, \mathbf{r}_i^M)$ with $\mathbf{r}_i^j = (r_{i,j}^{\text{corr}}, r_{i,j}^{\text{prot}}, r_{i,j}^{\text{mask}})$, for $i \in [n]$ and $j \in [M]$.
 2. For every $j \in [M]$ and $k \in [q]$, invoke \mathcal{S}_j^k , request the correlated randomness for every corrupted party P_i and get from \mathcal{S}_j^k the values $\tilde{r}_{i,j,k}^{\text{corr}}$. Set $r_{i,j,k}^{\text{corr}} \leftarrow \tilde{r}_{i,j,k}^{\text{corr}}$ in \mathbf{r}_i .
 3. Receive from \mathcal{A} the message $(\text{input}, \text{sid}, \lambda)$, for a corrupted P_i (that \mathcal{A} sends to \mathcal{F}_{SCP}).
 4. For every party P_i , compute $(\tilde{v}_i^1, \dots, \tilde{v}_i^n) \leftarrow \text{Share}(0^{|r_i|})$.
 5. Send the message $(\text{output}, \text{sid}, (\mathbf{r}_i, \tilde{v}_i^1, \dots, \tilde{v}_i^n))$ to \mathcal{A} as the response from \mathcal{F}_{SCP} , for every corrupted P_i .
- Simulating the first (input-commitment) message:
 1. For every honest party P_i , send a random string $\tilde{\mathbf{m}}_i^{\text{input}}$ to \mathcal{A} .

Protocol $\text{Comp}(\pi_1, \dots, \pi_M, N, R, L)$

1. Each party P_i , upon receiving $(\text{input}, \text{sid}, \mathbf{x}_i)$ with $\mathbf{x}_i = (x_i^1, \dots, x_i^M)$, sends $(\text{input}, \text{sid}, \lambda)$ to \mathcal{F}_{SCP} and receives $(\text{output}, \text{sid}, (r_i, v_i^1, \dots, v_i^n))$, where $r_i = (r_i^{\text{input}}, r_i^1, \dots, r_i^M)$ with $r_i^j = (r_{i,j}^{\text{corr}}, r_{i,j}^{\text{prot}}, r_{i,j}^{\text{mask}})$, for $j \in [M]$, as defined in $\mathcal{D}_{\text{parallel}}$, and where v_i^u is the i 'th share of the committed randomness r_u of party P_u . In addition, set the iteration index $\alpha \leftarrow 0$.
2. Every party P_i broadcasts $\tilde{m}_i^{\text{input}} = \mathbf{x}_i \oplus r_i^{\text{input}}$.
3. Denote $\tilde{\mathbf{m}}^{\text{input}} = (\tilde{m}_1^{\text{input}}, \dots, \tilde{m}_n^{\text{input}})$. In case some party P_i did not broadcast a valid value, each party locally sets for P_i a default input value $\tilde{\mathbf{x}}_i = (\tilde{x}_1, \dots, \tilde{x}_n)$ and default random coins \tilde{r}_i , and locally computes $\tilde{m}_i^{\text{input}}$ for that party.
4. initiate (in parallel) N instances of every protocol π_j as follows:

- (a) For every $j \in [M]$ and $k \in [N]$, initialize the history of each execution, $h_{i,j,k}, \tilde{h}_{i,j,k} \leftarrow \lambda$. Denote the private history of all instances of π_j as $\tilde{h}_i^j = (h_{i,j,1}, \dots, h_{i,j,N})$, and the public history as $\tilde{h}_i^j = (\tilde{h}_{i,j,1}, \dots, \tilde{h}_{i,j,N})$; finally, denote the public history of the entire protocol as $\tilde{\mathbf{h}}_i = (\tilde{m}_i^{\text{input}}, \tilde{h}_i^1, \dots, \tilde{h}_i^M)$.
- (b) Execute for rounds $\rho = 1, \dots, R$:
 - i. Every party P_i computes its ρ -round messages for the k 'th execution of π_j , for every $j \in [M]$ and $k \in [N]$:

$$m_{i,j,k,\rho} = (m_{i,j,k,\rho,1}, \dots, m_{i,j,k,\rho,n}) = \mathcal{F}_{\text{next-msg}}^{\pi_j, i}(x_i^j, r_{i,j,k}^{\text{corr}}, r_{i,j,k}^{\text{prot}}, h_{i,j,k}),$$

where $\tilde{k} = \alpha \cdot N + k$. Next, mask the messages as $\tilde{m}_{i,j,k,\rho} = (\tilde{m}_{i,j,k,\rho,1}, \dots, \tilde{m}_{i,j,k,\rho,n})$, where for $u \in [n]$, set $\tilde{m}_{i,j,k,\rho,u} = m_{i,j,k,\rho,u} \oplus r_{i,j,k,\rho,u}^{\text{mask}}$. Finally, prepare the message $\tilde{\mathbf{m}}_{i,\rho} = (\tilde{m}_{i,1,\rho}, \dots, \tilde{m}_{i,M,\rho})$, where for every $j \in [M]$, $\tilde{\mathbf{m}}_{i,j,\rho} = (\tilde{m}_{i,j,1,\rho}, \dots, \tilde{m}_{i,j,N,\rho})$.

- ii. Every party P_i sends $(\text{input}, \text{sid}, (\alpha, \rho, \tilde{\mathbf{m}}_{i,\rho}, \tilde{\mathbf{h}}_1))$ to \mathcal{F}_{SCP} , and receives the message $(\text{output}, \text{sid}, (((\alpha, \rho, \tilde{\mathbf{m}}_{1,\rho}, \tilde{\mathbf{h}}_1), b_1), \dots, ((\alpha, \rho, \tilde{\mathbf{m}}_{n,\rho}, \tilde{\mathbf{h}}_n), b_n)))$.
 - iii. For every $u \in [n]$ with $b_u = 0$, every P_i broadcasts v_u^i and all parties reconstruct the committed randomness of P_u as $r_u = \text{Recon}(v_u^1, \dots, v_u^n)$.
 - iv. For every $u \in [n]$ with $b_u = 1$, party P_i unmaskes the message $\tilde{\mathbf{m}}_{u,\rho}$ sent by P_u , by computing $m_{u,j,k,\rho,i} = \tilde{m}_{u,j,k,\rho,i} \oplus r_{u,j,k,\rho,i}^{\text{mask}}$, for $j \in [M]$ and $k \in [N]$.
 - v. Every party P_i locally computes the ρ -round messages $\tilde{\mathbf{m}}_{u,\rho}$ on behalf of all parties P_u that have been identified as corrupted thus far, using the committed input $\tilde{m}_u^{\text{input}}$ and reconstructed randomness r_u .
 - vi. Every party P_i appends all unmasked ρ -round messages $(m_{1,j,k,\rho,i}, \dots, m_{n,j,k,\rho,i})$ to $h_{i,j,k}$, and all masked ρ -round messages $(\tilde{m}_{1,j,k,\rho}, \dots, \tilde{m}_{n,j,k,\rho})$ to $\tilde{h}_{i,j,k}$.
- (c) For every $j \in [M]$ and $k \in [N]$, if party P_i completed the k 'th execution of π_j with output, let $y_{i,j,k}$ be the output value ($y_{i,j,k} = \perp$ otherwise). Denote $S_i^j = \{y_{i,j,k} \mid y_{i,j,k} \neq \perp\}$. If for every $j \in [M]$, $S_i^j \neq \emptyset$, then set $\mathbf{y} = (y_1, \dots, y_M)$, where $y_j \in S_i^j$ is arbitrarily chosen, and broadcasts (sid, \mathbf{y}) .
 - (d) If P_i received identical values (sid, \mathbf{y}) from (at least) $t + 1$ parties, it broadcasts (sid, \mathbf{y}) .
 - (e) If party P_i received identical values (sid, \mathbf{y}) from (at least) $n - t$ parties, it outputs $(\text{output}, \text{sid}, \mathbf{y})$ and halts. Else, set $\alpha \leftarrow \alpha + 1$. If $\alpha < L$, goto Step 4; else, output \perp and halt.

Fig. 5. The parallel-composition compiler in the $(\mathcal{F}_{\text{PBC}}, \mathcal{F}_{\text{SCP}})$ -hybrid model.

2. For every corrupted party P_i , receive the message $\tilde{m}_i^{\text{input}}$ from \mathcal{A} and extract the input value $x_i = \tilde{m}_i^{\text{input}} \oplus r_i^{\text{input}}$. In case \mathcal{A} does not send a message for P_i , set x_i and r_i to predetermined default values and locally compute $\tilde{m}_i^{\text{input}} = x_i \oplus r_i$.
 3. Set $\tilde{m}^{\text{input}} = (\tilde{m}_1^{\text{input}}, \dots, \tilde{m}_n^{\text{input}})$.
- Send inputs to \mathcal{F}_{PBB} :
 1. For every corrupted party P_i , send the message $(\text{input}, \text{sid}, x_i)$ to \mathcal{F}_{PBB} .
 2. Upon receiving message $(\text{leakage}, \text{sid}, P_i, (l_1, \dots, l_M))$ from \mathcal{F}_{PBB} , for an honest party P_i , store the leakage information.
 3. In addition, \mathcal{S} receives $(\text{trace}, \text{sid}, T)$, from \mathcal{F}_{PBB} , where T is a depth-1 trace of the form

$$\left(\mathcal{F}_{\text{SCP}}^1, \mathcal{F}_{\text{PBC}}, \left(\left(\mathcal{F}_{\text{SCP}}^i, \mathcal{F}_{\text{PBC}} \right)^{\text{R}}, \mathcal{F}_{\text{PBC}}, \mathcal{F}_{\text{PBC}}, \mathcal{F}_{\text{PBC}} \right)^{q''} \right),$$

(i.e., initially, calls to $\mathcal{F}_{\text{SCP}}^1$ and \mathcal{F}_{PBC} , followed by q'' iterations of calling R times $(\mathcal{F}_{\text{SCP}}^i, \mathcal{F}_{\text{PBC}})$ and 3 rounds of \mathcal{F}_{PBC} in order to agree on termination). More precisely, \mathcal{S} receives the coins that were used by the functionality \mathcal{F}_{PBB} to sample the trace T from D_{pbb} . Using these coins, \mathcal{S} samples the same traces $T_j^{\tilde{k}}$ that were used to define T .

- Simulating the α 'th iteration:
 1. For every $i \in [n]$, $j \in [M]$ and $k \in [N]$, set $h_{i,j,k}, \tilde{h}_{j,k} \leftarrow \lambda$. Denote $\mathbf{h}_i^j = (h_{i,j,1}, \dots, h_{i,j,N})$, $\tilde{\mathbf{h}}^j = (\tilde{h}_{j,1}, \dots, \tilde{h}_{j,N})$ and $\tilde{\mathbf{h}} = (\tilde{h}^1, \dots, \tilde{h}^M)$. For $k \in [N]$, let $\tilde{k} = \alpha \cdot N + k$.
 2. For $j \in [M]$ and $k \in [N]$, send to $\mathcal{S}_j^{\tilde{k}}$, on behalf of every honest party, the leakage messages $(\text{leakage}, \text{sid}, P_i, l_j)$, as well as $(\text{trace}, \text{sid}, T_j^{\tilde{k}})$. In addition, send to $\mathcal{S}_j^{\tilde{k}}$ the message $(\text{input}, \text{sid}, x_i^j)$, for every corrupted party P_i , and receive from $\mathcal{S}_j^{\tilde{k}}$ the messages $(\text{input}, \text{sid}, x_i^j)$ that are sent to $\mathcal{W}_{\text{flex}}^{D_j}(\mathcal{F}_j)$, for the corrupted parties.
 3. Execute for rounds $\rho = 1, \dots, \text{R}$:
 - (a) For every honest P_i , send $(\alpha, \rho, \tilde{m}_{i,\rho}, \tilde{\mathbf{h}})$ to \mathcal{A} as the leakage from \mathcal{F}_{SCP} for P_i , where $\tilde{m}_{i,\rho}$ is prepared as follows. For every corrupted party P_u , receive from $\mathcal{S}_j^{\tilde{k}}$ the message $m_{i,j,k,\rho,u}$ and set $\tilde{m}_{i,j,k,\rho,u} = m_{i,j,k,\rho,u} \oplus r_{i,j,k,\rho,u}^{\text{mask}}$. For every honest P_u , sample a random message $\tilde{m}_{i,j,k,\rho,u}$. Set $\tilde{\mathbf{m}}_{i,\rho} = (\tilde{m}_{i,1,\rho}, \dots, \tilde{m}_{i,M,\rho})$, where $\tilde{m}_{i,j,\rho} = (\tilde{m}_{i,j,1,\rho}, \dots, \tilde{m}_{i,j,N,\rho})$ with $\tilde{m}_{i,j,k,\rho} = (\tilde{m}_{i,j,k,\rho,1}, \dots, \tilde{m}_{i,j,k,\rho,n})$.
 - (b) For every corrupted party P_i , receive from \mathcal{A} the message $(\text{input}, \text{sid}, (\alpha, \rho, \tilde{m}_{i,\rho}, \tilde{\mathbf{h}}_i))$ (intended to \mathcal{F}_{SCP}), and check whether $\tilde{\mathbf{h}}_i = \tilde{\mathbf{h}}$ and $\mathcal{R}_{\text{parallel}}^i((\alpha, \rho, \tilde{m}_{i,\rho}, \tilde{\mathbf{h}}_i), \mathbf{r}_i) = 1$. If so, set $b_i = 1$ else $b_i = 0$.
 - (c) For every corrupted party P_i with $b_i = 1$, it holds that $\tilde{m}_{i,\rho} = (\tilde{m}_{i,1,\rho}, \dots, \tilde{m}_{i,M,\rho})$, where $\tilde{m}_{i,j,\rho} = (\tilde{m}_{i,j,1,\rho}, \dots, \tilde{m}_{i,j,N,\rho})$ and $\tilde{m}_{i,j,k,\rho} =$

- ($\tilde{m}_{i,j,k,\rho,1}, \dots, \tilde{m}_{i,j,k,\rho,n}$). Denote, for every honest P_u , $m_{i,j,k,\rho,u} = \tilde{m}_{i,j,k,\rho,u} \oplus r_{i,j,k,\rho,u}^{\text{mask}}$.
- (d) For every honest party P_i , denote $\tilde{h}_i = \tilde{h}$ and $b_i = 1$, and send to the adversary \mathcal{A} (`output, sid, ((($\alpha, \rho, \tilde{m}_{1,\rho}, \tilde{h}_1$), b_1), \dots , ($\alpha, \rho, \tilde{m}_{n,\rho}, \tilde{h}_n$), b_n))`) as the response from \mathcal{F}_{SCP} , on behalf of every corrupted party.
 - (e) For every corrupted party P_i with $b_i = 1$ and every honest P_u , forward to \mathcal{S}_j^k the message $m_{i,j,k,\rho,u}$.
 - (f) For every corrupted party P_i with $b_i = 0$, compute $(v_i^1, \dots, v_i^n) \leftarrow \text{Share}(r_i)$ and send v_i^u to \mathcal{A} , on behalf of every honest party P_u . (Note that r_i may change as the simulation proceeds, in case P_i gets corrupted dynamically.)
 - (g) For every corrupted party that has been previously identified, locally compute the messages $m_{i,j,k,\rho,u}$, for every honest party P_u , using (x_i, r_i) and $h_{i,j,k}$, and send them to \mathcal{S}_j^k .
 - (h) Append the messages $(\tilde{m}_{1,\rho}, \dots, \tilde{m}_{n,\rho})$ to the global history \tilde{h} . For every corrupted party P_i , append the messages $(m_{1,j,k,\rho,i}, \dots, m_{n,j,k,\rho,i})$ to the local history $h_{i,j,k}$.
 - (i) Upon receiving (`early-output, sid, P_i`) from \mathcal{S}_j^k for an honest party P_i , simulate the honest party receiving the output message.
 - (j) Upon receiving (`early-output, sid, P_i`) from \mathcal{S}_j^k for a corrupted party P_i , if $y = \perp$, send (`early-output, sid, P_i`) to \mathcal{F}_{PBB} and receive back (`output, sid, y`), with $y = (y_1, \dots, y_M)$. Send (`output, sid, y_j`) to \mathcal{S}_j^k .
 - (k) For every $j \in [M]$ and $k \in [N]$, check whether $\rho = c_{\text{tr}}(T_j^k)$. If so and $y = \perp$, send (`early-output, sid, P_i`) to \mathcal{F}_{PBB} and receive back (`output, sid, y`), with $y = (y_1, \dots, y_M)$. Send (`output, sid, y_j`) to \mathcal{S}_j^k for every corrupted party.
- Explaining corruption requests: upon a corruption request for P_i , proceed as follows.
 1. Corrupt the dummy party P_i and learn its input $x_i = (x_i^1, \dots, x_i^M)$.
 2. For every $j \in [M]$ and $k \in [L \cdot N]$:
 - (a) Instruct \mathcal{S}_j^k to corrupt P_i and provide it with input x_i^j .
 - (b) Receive from \mathcal{S}_j^k the view of P_i , i.e., its internal randomness $\tilde{r}_{i,j,k}^{\text{prot}}$, correlated randomness $\tilde{r}_{i,j,k}^{\text{corr}}$ and history $h_{i,j,k}$.
 - (c) For every $u \in [n]$ and every message $m_{u,j,k,\rho,i}$ in $h_{i,j,k}$ (message from P_u to P_i in the ρ 'th round of the k 'th instance of π_j , for $\rho \in [R]$), compute $\tilde{r}_{u,j,k,\rho,i}^{\text{mask}} = m_{u,j,k,\rho,i} \oplus \tilde{m}_{u,j,k,\rho,i}$. Similarly, compute $\tilde{r}_{i,j,k,\rho,u}^{\text{mask}} = m_{i,j,k,\rho,u} \oplus \tilde{m}_{i,j,k,\rho,u}$.
 - (d) Compute $\tilde{r}_i^{\text{input}} = \tilde{m}_i^{\text{input}} \oplus x_i$.
 - (e) Adjust the random coins $r_{i,j,k}^{\text{corr}} \leftarrow \tilde{r}_{i,j,k}^{\text{corr}}$, $r_{i,j,k}^{\text{prot}} \leftarrow \tilde{r}_{i,j,k}^{\text{prot}}$, $r_{i,j,k}^{\text{mask}} \leftarrow \tilde{r}_{i,j,k}^{\text{mask}}$ and $r_i^{\text{input}} \leftarrow \tilde{r}_i^{\text{input}}$ in r_i .

3. Provide the input x_i and the adjusted correlated randomness r_i to \mathcal{Z} as the internal state of party P_i .

We next prove that no environment can distinguish between interacting with the dummy adversary and the honest parties running protocol π' in the $(\mathcal{F}_{\text{PBC}}, \mathcal{F}_{\text{SCP}})$ -hybrid model, from interacting with the simulator \mathcal{S} and the dummy honest parties computing \mathcal{F}_{PBB} , except for a negligible probability. We prove this using a series of hybrid games. The output of each game is the output of the environment.

The game $\text{HYB}_{\pi_{\text{PBB}}, \mathcal{A}, \mathcal{Z}}^1$. This is exactly the execution of the protocol π_{PBB} in the $(\mathcal{F}_{\text{PBC}}, \mathcal{F}_{\text{SCP}})$ -hybrid model with environment \mathcal{Z} and dummy adversary \mathcal{A} .

The games $\text{HYB}_{\pi_{\text{PBB}}, \mathcal{A}, \mathcal{Z}}^{2, i^*}$, for $0 \leq i^* \leq n$. In these games, we modify $\text{HYB}_{\pi_{\text{PBB}}, \mathcal{A}, \mathcal{Z}}^1$ as follows. During the first call to \mathcal{F}_{SCP} , for $i^* < i \leq n$, compute $(v_i^1, \dots, v_n^1) \leftarrow \text{Share}(r_i)$ (as before) and send (v_i^1, \dots, v_n^1) as the shares for every party P_i . For every $i \leq i^*$, compute $(v_i^1, \dots, v_n^1) \leftarrow \text{Share}(r_i)$ and $(\tilde{v}_i^1, \dots, \tilde{v}_n^1) \leftarrow \text{Share}(0^{|\mathcal{r}_i|})$. Next, send to every corrupted P_k the shares $(\tilde{v}_1^k, \dots, \tilde{v}_n^k)$ and to every honest party the shares (v_1^i, \dots, v_n^i) .

Claim 6.4. $\text{HYB}_{\pi_{\text{PBB}}, \mathcal{A}, \mathcal{Z}}^1 \stackrel{s}{\equiv} \text{HYB}_{\pi_{\text{PBB}}, \mathcal{A}, \mathcal{Z}}^{2, n}$.

Proof Note that $\text{HYB}_{\pi_{\text{PBB}}, \mathcal{A}, \mathcal{Z}}^1 \stackrel{s}{\equiv} \text{HYB}_{\pi_{\text{PBB}}, \mathcal{A}, \mathcal{Z}}^{2, 0}$. For every $0 \leq \tilde{i} < n$, it holds that $\text{HYB}_{\pi_{\text{PBB}}, \mathcal{A}, \mathcal{Z}}^{2, \tilde{i}} \stackrel{s}{\equiv} \text{HYB}_{\pi_{\text{PBB}}, \mathcal{A}, \mathcal{Z}}^{2, \tilde{i}+1}$. This follows from the security of the ECSS scheme and the honest-majority assumption. In particular, the shares $\{v_1^u, \dots, v_n^u\}_u$ held by the adversary (for all corrupted P_u), completely hide for that random coins r_i of every honest party P_i , and are compatible even if P_i gets adaptively corrupted and r_i is revealed. This holds since all honest parties receive valid shares of r_i , therefore, r_i will be correctly reconstructed even if all corrupted parties have incorrect shares. The claim follows using a standard hybrid argument. \square

The game $\text{HYB}_{\pi_{\text{PBB}}, \mathcal{A}, \mathcal{Z}}^3$. In this game, we modify $\text{HYB}_{\pi_{\text{PBB}}, \mathcal{A}, \mathcal{Z}}^{2, n}$ as follows. The input-commitment message $\tilde{m}_i^{\text{input}}$, sent by an honest party is uniformly chosen. In addition, in every call to \mathcal{F}_{SCP} the returned value b_i for every honest party P_i is always set to 1. Finally, the random coins for honest parties that are corrupted adaptively are set as $r_i^{\text{input}} = \tilde{m}_i^{\text{input}} \oplus x_i$.

Claim 6.5. $\text{HYB}_{\pi_{\text{PBB}}, \mathcal{A}, \mathcal{Z}}^{2, n} \stackrel{s}{\equiv} \text{HYB}_{\pi_{\text{PBB}}, \mathcal{A}, \mathcal{Z}}^3$.

Proof The claim immediately follows since honest parties always send correct messages. \square

The games $\text{HYB}_{\pi_{\text{PBB}}, \mathcal{A}, \mathcal{Z}}^{4, j^*, k^*}$, for $0 \leq j^* \leq M$ and $0 \leq k^* \leq q = N \cdot L$. In these games, we modify $\text{HYB}_{\pi_{\text{PBB}}, \mathcal{A}, \mathcal{Z}}^3$ as follows. For $(j, k) < (j^*, k^*)$ (i.e., for $j < j^*$, or $j = j^*$

and $k < k^*$), the k 'th execution of protocol π_j is replaced with the simulated transcript generated by \mathcal{S}_j^k .

More specifically, the experiment interacts with the ideal functionality \mathcal{F}_{PBB} . Initially, it receives the leakage messages (leakage , sid , P_i , (l_1, \dots, l_M)) and the coins used to sample a trace T from D_{pbb} (via the message (trace , sid , T)); using these random coins, sample the traces T_j^k from D_j .

Each simulator \mathcal{S}_j^k (for $(j, k) < (j^*, k^*)$) is invoked and is given the leakage information l_j and the trace T_j^k . The simulator \mathcal{S}_j^k provides correlated randomness $\tilde{r}_{i,j,k}^{\text{corr}}$ for every corrupted party; the random coins r_i for corrupted P_i are adjusted accordingly. Once \mathcal{A} sends the input-commitment message $\tilde{m}_i^{\text{input}}$, for a corrupted P_i , extract the input value $x_i = (x_i^1, \dots, x_i^M)$ and hand \mathcal{S}_j^k the value x_i^j as the input value for P_i . The interaction with \mathcal{S}_j^k is done as in the simulation, i.e., validate the messages from \mathcal{A} and unmask valid messages before forwarding to \mathcal{S}_j^k , and mask messages from \mathcal{S}_j^k using $r_{i,j,k}^{\text{corr}}$. Messages from identified corrupted parties are locally computed and sent to \mathcal{S}_j^k . When \mathcal{S}_j^k sends (early-output , sid , \cdot) requests, respond with the output value y_j from $y = (y_1, \dots, y_j)$, where if $y = \perp$, then forward the message to \mathcal{F}_{PBB} .

Claim 6.6. $\text{HYB}_{\pi_{\text{PBB}}, \mathcal{A}, \mathcal{Z}}^3 \stackrel{s}{\equiv} \text{HYB}_{\pi_{\text{PBB}}, \mathcal{A}, \mathcal{Z}}^{4, M, q}$.

Proof Note that $\text{HYB}_{\pi_{\text{PBB}}, \mathcal{A}, \mathcal{Z}}^3 \stackrel{s}{\equiv} \text{HYB}_{\pi_{\text{PBB}}, \mathcal{A}, \mathcal{Z}}^{4, 0, 0}$. For every $0 \leq j < M$ and $0 \leq k < q$, it holds that $\text{HYB}_{\pi_{\text{PBB}}, \mathcal{A}, \mathcal{Z}}^{4, j, k} \stackrel{s}{\equiv} \text{HYB}_{\pi_{\text{PBB}}, \mathcal{A}, \mathcal{Z}}^{2, j, k+1}$ and similarly, $\text{HYB}_{\pi_{\text{PBB}}, \mathcal{A}, \mathcal{Z}}^{4, j, q} \stackrel{s}{\equiv} \text{HYB}_{\pi_{\text{PBB}}, \mathcal{A}, \mathcal{Z}}^{2, j+1, 1}$; otherwise, there exists an environment that violates the simulation of \mathcal{S}_j (resp., \mathcal{S}_{j+1}). The claim follows using a standard hybrid argument. \square

The proof of Lemma 6.3 follows since $\text{HYB}_{\pi_{\text{PBB}}, \mathcal{A}, \mathcal{Z}}^{4, M, q}$ (for $q = N \cdot L$) exactly describes the simulation done by \mathcal{S} . \square

6.2. An Impossibility of FBB Round-Preserving Parallel Composition

In this section, we prove that for a natural class of protocols, following and/or extending in various ways the techniques from Ben-Or and El-Yaniv [5],⁸ there exist functions such that no protocol can compute their parallel composition in a round-preserving manner, while accessing the functions in a black-box way, tolerating even a single adversarial party. Although this is not a general impossibility result, it indicates that the batching approach of [5] is limited to semi-honest security (Sect. 5) and/or functionally white-box transformations.

We observe that this impossibility serves as an additional justification for the optimality of our protocol-black-box parallel composition (Sect. 6.1). Indeed, on the one hand, it formally confirms the generic observation that the naïve parallel composition of a set of PT functionalities does not preserve their round complexity. On the other hand, and

⁸To our knowledge, the only known techniques for round-preserving parallel composition are those of Ben-Or and El-Yaniv [5] and are only for the specific case of Byzantine agreement.

most importantly, it proves that all existing techniques for composing PT functionalities in parallel in the natural (FBB) manner fail in preserving the round complexity. Hence, the only known existing round-preserving composition for such functionalities is the protocol-black-box compiler presented in Sect. 6.1 or more inefficient non-black-box techniques. The wideness of the class of excluded protocols by our impossibility result justifies our conjecture that there exists no round-preserving FBB protocol for parallel composition of PT functionalities. Proving this conjecture is in our opinion a very interesting research direction.

We first argue informally why the approach of [5], cannot be directly extended to privacy-sensitive functions. The idea in [5] for allowing each of the n parties to broadcast its value is to have each of the n parties participate in $m = O(\log n)$ parallel invocations (hereafter called *batches*, to avoid confusion with the goal of parallel broadcast for different messages) of broadcast as sender with the same input. Each of those batches is executed in parallel for a fixed (constant) number of rounds (for the same broadcast message); this increases the probability that sufficiently many parties receive output from each batch. At the end of each batch execution, the parties check whether they jointly hold the output, and if not, they repeat the computation of the batches. It might seem that this idea can be applied to arbitrary tasks, but this is not the case. The reason is that this idea fails if the functionality has any privacy requirements, is that the adversary can input different values on different calls of the functionality within a batch and learn more information on the input.

Batched-parallel composition The above issue with privacy appears whenever a function is invoked twice in the same round on the same inputs from honest parties. Indeed, in this case the adversary can use different inputs to each invocation and learn information as sketched above. The same attack can be extended to composition protocols which invoke the function in two different rounds ρ and ρ' ; as long as the adversary knows these rounds, he can still launch the above attack on privacy. Generalizing the claim even further, for specific classes of functions, it suffices that there are two (possibly different) functions which are evaluated on the same inputs in rounds ρ and ρ' . This excludes protocols that might attempt to avoid using some functionality $\mathcal{W}_{\text{flex}}^{D_j}(\mathcal{F}_j)$ by invoking some other $\mathcal{W}_{\text{flex}}^{D_{j'}}(\mathcal{F}_{j'})$.

To capture the above generalization, we define the class of *batched-parallel composition* protocols: A protocol π implementing the PT parallel composition $\mathcal{W}_{\text{flex}}^D(\mathcal{F}_1 \parallel \dots \parallel \mathcal{F}_M)$ in the $(\mathcal{W}_{\text{flex}}^{D_1}(\mathcal{F}_1), \dots, \mathcal{W}_{\text{flex}}^{D_M}(\mathcal{F}_M))$ -hybrid model (for some distributions D, D_1, \dots, D_M) is a *batched-parallel composition* protocol if it has the following structure. It proceeds in rounds, where in each round the protocol might initiate (possibly multiple) calls to any number of the hybrid functionalities $\mathcal{W}_{\text{flex}}^{D_j}(\mathcal{F}_j)$ and/or continue calls that were initiated in previous rounds. Furthermore, there exist two publicly known protocol rounds ρ and ρ' , and indices $j, j', \ell \in [M]$, such that for the input vector $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$ that π gives to $\mathcal{W}_{\text{flex}}^D(\mathcal{F}_1 \parallel \dots \parallel \mathcal{F}_M)$ (where $\mathbf{x}_i = (x_i^1, \dots, x_i^M)$) the following properties are satisfied:

1. In round ρ the functionality $\mathcal{W}_{\text{flex}}^{D_j}(\mathcal{F}_j)$ is called on input $\mathbf{x}^\ell = (x_1^\ell, \dots, x_n^\ell)$ and at least two of its rounds are executed.⁹
2. In round ρ' the functionality $\mathcal{W}_{\text{flex}}^{D_{j'}}(\mathcal{F}_{j'})$ is also called on input \mathbf{x}^ℓ and at least two of its rounds are executed.

Note that the protocol in [5] (as well as our semi-honest protocol from Sect. 5) is an example of a batched-parallel composition protocol for $\rho = \rho' = 1$ and for $j = j' = \ell$ being the index of any one of the hybrid functionalities. Indeed, in the first round of these protocols, each functionality is invoked at least twice on the same inputs. In particular, protocols that follow this structure, e.g., even ones that do not call all functionalities in every phase, or those that have variable batch sizes, can be described as such batched-parallel composition protocols.

We next show that there are classes of functions $\mathcal{C}_1, \dots, \mathcal{C}_M$ such and for any protocol π that securely computes the parallel composition $\mathcal{W}_{\text{flex}}^D(\mathcal{F}_{\text{SFE}}^{\mathcal{C}_1} \parallel \dots \parallel \mathcal{F}_{\text{SFE}}^{\mathcal{C}_M})$ while given hybrid access to PT functionalities $\mathcal{W}_{\text{flex}}^{D_i}(\mathcal{F}_{\text{SFE}}^{\mathcal{C}_i})$ the following properties hold *simultaneously*:

1. π has to call each of the hybrids $\mathcal{W}_{\text{flex}}^{D_i}(\mathcal{F}_{\text{SFE}}^{\mathcal{C}_i})$ (for at least 2 rounds each).¹⁰
2. The naïve solution of π calling each of the $\mathcal{W}_{\text{flex}}^{D_i}(\mathcal{F}_{\text{SFE}}^{\mathcal{C}_i})$'s in parallel until they terminate is not round-preserving (for an appropriate choice of D_i 's.)
3. π cannot be a batched-parallel composition protocol.

The above shows that the classes $\mathcal{C}_1, \dots, \mathcal{C}_M$ not only exclude the existence of a batched-parallel composition protocol, but they also exclude all other known solutions. This implies that for this classes of functions, every known approach—and generalizations thereof—fail to compute the parallel composition of the corresponding functionality in an FBB and round-preserving manner.

Theorem 6.7. *Let $M = O(\kappa)$. There exist n -party function-classes $\mathcal{C}_1, \dots, \mathcal{C}_M$ and distributions D_1, \dots, D_M , such that the following properties hold in the presence of a malicious adversary corrupting any one of the parties:*

1. *The protocol that calls each $\mathcal{W}_{\text{flex}}^{D_i}(\mathcal{F}_{\text{SFE}}^{\mathcal{C}_i})$ in parallel (once) until it terminates is not round-preserving (i.e., its expected round complexity is asymptotically higher than that of the distributions D_i).*
2. *Any $(\mathcal{F}_{\text{SMT}}, \mathcal{W}_{\text{flex}}^{D_1}(\mathcal{F}_{\text{SFE}}^{\mathcal{C}_1}), \dots, \mathcal{W}_{\text{flex}}^{D_M}(\mathcal{F}_{\text{SFE}}^{\mathcal{C}_M}))$ -hybrid protocol for computing $\mathcal{W}_{\text{flex}}^D(\mathcal{F}_{\text{SFE}}^{\mathcal{C}_1} \parallel \dots \parallel \mathcal{F}_{\text{SFE}}^{\mathcal{C}_M})$ has to make a meaningful call (i.e., a call that executes at least two rounds) to each PT hybrid $\mathcal{W}_{\text{flex}}^{D_i}(\mathcal{F}_{\text{SFE}}^{\mathcal{C}_i})$.*
3. *There exists no functionally black-box batched-parallel composition protocol for computing $\mathcal{W}_{\text{flex}}^D(\mathcal{F}_{\text{SFE}}^{\mathcal{C}_1} \parallel \dots \parallel \mathcal{F}_{\text{SFE}}^{\mathcal{C}_M})$ in the $(\mathcal{F}_{\text{SMT}}, \mathcal{W}_{\text{flex}}^{D_1}(\mathcal{F}_{\text{SFE}}^{\mathcal{C}_1}), \dots, \mathcal{W}_{\text{flex}}^{D_M}(\mathcal{F}_{\text{SFE}}^{\mathcal{C}_M}))$ -hybrid model, where D has (asymptotically) the same expectation as D_1, \dots, D_M .*

⁹Note that any call to a PT (i.e., wrapped) functionality that executes less than two rounds is useless as it can be simulated by the protocol that does nothing (without of course increasing the round complexity). The reason is that, by definition of PT functionalities, any such call gives no output (the first round is an input-only round).

¹⁰Note that this does not mean that π is not round preserving as the calls might be in parallel.

Proof Let $D_1 = \dots = D_M$ be the geometric distribution with parameter $1/2$. (This means that the expected round complexity of each $\mathcal{W}_{\text{flex}}^{D_i}(\mathcal{F}_{\text{SFE}}^{\mathcal{C}_i})$ is constant.) Then, Property 1 follows immediately from the observation in [5] (see also [16]), which implies that the expectation of the round complexity of the naïve protocol that executes each functionality in parallel until its completion will be $\Theta(\log M)$, which is super-constant.

We next turn to Property 2. Toward proving it we first prove the following useful lemma.

Lemma 6.8. *There exists a family of functions $\mathcal{C} = \{f_\alpha\}_{\alpha \in \{0,1\}^\kappa}$ such that there exists no FBB protocol for computing the family of (oracle-aided) n -party SFE-functionalities $\{\mathcal{F}^{f_\alpha}\}_{f_\alpha \in \mathcal{C}}$, which is secure against a semi-honest adversary corrupting any one of the n parties.*

Proof Let f_α be the function that takes inputs $x_1 \in \{0,1\}^\kappa$ and $x_2 \in \{0,1\}^\kappa$ from P_1 and P_2 , respectively (and a default empty string λ from every P_j with $j \in \{3, \dots, n\}$) and outputs y_i to each P_i as follows:

- $y_1 = \begin{cases} x_2, & \text{if } x_1 \oplus x_2 = \alpha \\ 0^\kappa, & \text{otherwise.} \end{cases}$
- $y_2 = \begin{cases} x_1, & \text{if } x_1 \oplus x_2 = \alpha \\ 0^\kappa, & \text{otherwise.} \end{cases}$
- For every $j \in \{3, \dots, n\}$, $y_j = \begin{cases} x_1 \oplus x_2, & \text{if } x_1 \oplus x_2 = \alpha \\ 0^\kappa, & \text{otherwise.} \end{cases}$

The argument that there exists no FBB protocol for the above function family is inspired by [43, Theorem 1]. Concretely, assume toward contradiction that such a protocol π exists¹¹ and consider the following experiment. Pick x_1, x_2, α uniformly and independently at random and run π^{f_α} with inputs x_1 and x_2 for P_1 and P_2 , respectively, and input λ for all other parties. Then, we argue that the following events have negligible probability of occurring (where the probability is taken over the choice of x_1, x_2, α and the random coins $\mathbf{r} = (r_1, \dots, r_n)$ of the parties):

- (A) Any of the parties (i.e., any of the π_i 's) queries its oracle f_α with $(p, q, \lambda, \dots, \lambda)$ such that $p \oplus q = \alpha$.
- (B) Any of the parties queries its oracle f_α with $(p, q, \lambda, \dots, \lambda)$ such that $p \oplus q = x_1 \oplus x_2$.

The fact that (A) occurs with negligible probability is due to the fact that α is uniformly random.

The fact that (B) occurs with negligible probability follows from the protocol's privacy (and is, in fact, independent of the distribution of α). Indeed, suppose that the probability that P_i makes a query $(p, q, \lambda, \dots, \lambda)$ such that $p \oplus q = x_1 \oplus x_2$ is noticeable (i.e., not negligible). Consider an adversary corrupting P_i and outputting the list of values $p \oplus q$, for each $(p, q, \lambda, \dots, \lambda)$ that P_i makes to its oracle. By assumption, this list will include $x_1 \oplus x_2$ with noticeable probability. However, in the ideal-world execution,

¹¹Note that by definition of FBB [57] protocols, $\pi = (\pi_1, \dots, \pi_n)$ and each of the π_i 's is an oracle machine which can query the function f_α is trying to compute (see Sect. 4).

the simulator, even if it knows α , will be unable to produce a list of values containing $x_1 \oplus x_2$ with noticeable probability, since x_1 and x_2 are chosen uniformly at random, and a simulator corrupting a single party cannot learn both x_1 and x_2 . This implies that the above adversary cannot be simulated, which contradicts the protocol's privacy.

Let now R denote the set of all $(\mathbf{r}, x_1, x_2, \alpha)$ such that in the above experiment, neither event (A) nor (B) occurs. From the above argument we know that $\Pr[(\mathbf{r}, x_1, x_2, \alpha) \in R] > 1 - \nu(\kappa)$, for a negligible function ν .

Next, as in [43, Theorem 1], we consider the coupled experiment in which we use the same \mathbf{r}, x_1, x_2 as above, but run the protocol $\pi^{f_{\alpha^*}}$ where $\alpha^* = x_1 \oplus x_2$. As in [43, Theorem 1], we can prove that this experiment proceeds identically as the original one (which, recall differs only on the oracle calls); in particular, all oracle queries will be answered by 0^κ to all parties. The reason for this is that an oracle query $(p, q, \lambda, \dots, \lambda)$ is answered by a value other than 0^κ in the first experiment only if $p \oplus q = \alpha$ and in the second experiment only if $p \oplus q = x_1 \oplus x_2$, which for the combinations of $(\mathbf{r}, x_1, x_2, \alpha) \in R$ does not occur. This in particular implies that the output vector $\mathbf{y} = (y_1, \dots, y_n)$ that the parties P_1, \dots, P_n receive from the protocol will be identical in both experiments for $(\mathbf{r}, x_1, x_2, \alpha) \in R$.

It follows that, since $\Pr[(\mathbf{r}, x_1, x_2, \alpha) \in R] > 1 - \nu(\kappa)$, the distribution of \mathbf{y} in both experiments can be at most negligible apart. However, since the protocol is required to output the correct value with overwhelming probability, in the first experiment (where x_1, x_2, α are uniformly distributed) $\Pr[\mathbf{y} = (0^\kappa)^n] = 1 - \text{negl}(\kappa)$ and in the second experiment (where $\alpha^* = x_1 \oplus x_2$) $\Pr[\mathbf{y} \neq (0^\kappa)^n] = 1 - \text{negl}(\kappa)$. This creates a noticeable advantage in distinguishing the outputs of the two experiments leading to contradiction. \square

Lemma 6.9. *For the above function class \mathcal{C} , let $\mathcal{C}_i = \mathcal{C}$ for each $i \in [M]$. Then, any FBB protocol for computing $\mathcal{W}_{\text{flex}}^D(\mathcal{F}_{\text{SFE}}^{\mathcal{C}_1} \parallel \dots \parallel \mathcal{F}_{\text{SFE}}^{\mathcal{C}_M})$ with hybrid access to all $\mathcal{W}_{\text{flex}}^{D_i}(\mathcal{F}_{\text{SFE}}^{\mathcal{C}_i})$'s needs to invoke at least two rounds of every hybrid $\mathcal{W}_{\text{flex}}^{D_i}(\mathcal{F}_{\text{SFE}}^{\mathcal{C}_i})$.*

Proof Assume toward contradiction that there exists a protocol π computing $\mathcal{W}_{\text{flex}}^D(\mathcal{F}_{\text{SFE}}^{\mathcal{C}_1} \parallel \dots \parallel \mathcal{F}_{\text{SFE}}^{\mathcal{C}_M})$ such that for some i , protocol π might execute at most one round of $\mathcal{W}_{\text{flex}}^{D_i}(\mathcal{F}_{\text{SFE}}^{\mathcal{C}_i})$. We will prove that no such protocol can exist for a uniformly random choice of $\alpha = (\alpha_1, \dots, \alpha_M)$.¹² This is sufficient, since by an averaging argument, it implies that there exists no such protocol that is secure for all choices of α , as required by the definition in [57]. (Indeed, if there exists a protocol that is secure for all choices of α (independent of α) there exists one that is secure for a randomly chosen α .)

First, we observe that any call to a functionality $\mathcal{W}_{\text{flex}}^{D_j}(\mathcal{F}_{\text{SFE}}^{\mathcal{C}_j})$ that executes less than two rounds can be simulated by the protocol that does nothing during this execution (without of course increasing the round complexity). The reason is that, by definition of flexibly wrapped CSFs (Sect. 2.2), any such call gives no output (the first round is an

¹²Consistently with [57], we assume that although α is chosen uniformly at random, it is known to the environment, the adversary and the simulator in the proof (in particular, we can assume that the environment chooses α to be uniformly random in each of the experiments and hands it to the adversary/simulator).

input-only round). Thus, we can assume without loss of generality that π makes no call to $\mathcal{W}_{\text{flex}}^{D_i}(\mathcal{F}_{\text{SFE}}^{C_i})$.

Next, observe that any protocol that computes $\mathcal{W}_{\text{flex}}^D(\mathcal{F}_{\text{SFE}}^{C_1} \parallel \dots \parallel \mathcal{F}_{\text{SFE}}^{C_M})$ can be adjusted to compute $\mathcal{W}_{\text{flex}}^{D_i}(\mathcal{F}_{\text{SFE}}^{C_i})$ for any $i \in [M]$. Indeed, one simply needs to invoke the protocol for $\mathcal{W}_{\text{flex}}^D(\mathcal{F}_{\text{SFE}}^{C_1} \parallel \dots \parallel \mathcal{F}_{\text{SFE}}^{C_M})$ and take its i 'th output as the output of $\mathcal{W}_{\text{flex}}^{D_i}(\mathcal{F}_{\text{SFE}}^{C_i})$. Hence, π can be trivially turned to a protocol for computing $\mathcal{W}_{\text{flex}}^{D_i}(\mathcal{F}_{\text{SFE}}^{C_i})$ (without ever accessing $\mathcal{W}_{\text{flex}}^{D_i}(\mathcal{F}_{\text{SFE}}^{C_i})$). Note that since we assume that π is FBB, the parties are given access to all other functionalities $\mathcal{W}_{\text{flex}}^{D_1}(\mathcal{F}_{\text{CSF}}^{f_{\alpha_1}}), \dots, \mathcal{W}_{\text{flex}}^{D_{i-1}}(\mathcal{F}_{\text{CSF}}^{f_{\alpha_{i-1}}}), \mathcal{W}_{\text{flex}}^{D_{i+1}}(\mathcal{F}_{\text{CSF}}^{f_{\alpha_{i+1}}}), \dots, \mathcal{W}_{\text{flex}}^{D_M}(\mathcal{F}_{\text{CSF}}^{f_{\alpha_M}})$, and can make oracle queries to all underlying functions $\{f_{\alpha_1}, \dots, f_{\alpha_M}\} \setminus \{f_{\alpha_i}\}$.

Finally, we observe the argument of Lemma 6.8 can be easily extended to the above scenario, where the aim is to compute $\{\mathcal{F}^{f_{\alpha_i}}\}_{f_{\alpha_i} \in \mathcal{C}}$ in an FBB manner, but where the parties have oracle access to the function f_{α_i} (as required by the definition of FBB protocols [57]) and, in addition, have oracle access to $f_{\alpha_1}, \dots, f_{\alpha_{i-1}}, f_{\alpha_{i+1}}, \dots, f_{\alpha_M}$ and ideal access to $\mathcal{W}_{\text{flex}}^{D_1}(\mathcal{F}_{\text{CSF}}^{f_{\alpha_1}}), \dots, \mathcal{W}_{\text{flex}}^{D_{i-1}}(\mathcal{F}_{\text{CSF}}^{f_{\alpha_{i-1}}}), \mathcal{W}_{\text{flex}}^{D_{i+1}}(\mathcal{F}_{\text{CSF}}^{f_{\alpha_{i+1}}}), \dots, \mathcal{W}_{\text{flex}}^{D_M}(\mathcal{F}_{\text{CSF}}^{f_{\alpha_M}})$. The reason is that the behavior of functions $f_{\alpha_1}, \dots, f_{\alpha_{i-1}}, f_{\alpha_{i+1}}, \dots, f_{\alpha_M}$ is independent of f_{α_i} and can be therefore trivially simulated by means of an information-theoretic MPC protocol (recall we only have one corruption) that implements a globally accessible oracle to $f_{\alpha_1}, \dots, f_{\alpha_{i-1}}, f_{\alpha_{i+1}}, \dots, f_{\alpha_M}$ and ideal functionalities $\mathcal{W}_{\text{flex}}^{D_1}(\mathcal{F}_{\text{CSF}}^{f_{\alpha_1}}), \dots, \mathcal{W}_{\text{flex}}^{D_{i-1}}(\mathcal{F}_{\text{CSF}}^{f_{\alpha_{i-1}}}), \mathcal{W}_{\text{flex}}^{D_{i+1}}(\mathcal{F}_{\text{CSF}}^{f_{\alpha_{i+1}}}), \dots, \mathcal{W}_{\text{flex}}^{D_M}(\mathcal{F}_{\text{CSF}}^{f_{\alpha_M}})$. (Since α is chosen uniformly at random their role in computing f_{α_i} can be emulated by choosing independent α_j 's for $j \in [M] \setminus \{i\}$.) Thus, if there were a protocol which would compute $\{\mathcal{F}^{f_{\alpha_i}}\}_{f_{\alpha_i} \in \mathcal{C}}$ in the above hybrid world, then it can be trivially converted to a protocol which does not access the hybrids or the oracle calls to the functions other than f_{α_i} , which would contradict Lemma 6.8. Observe that the above extension is independent of the round complexity, and adding a PT structure to the hybrids does not affect the impossibility. \square

We complete the proof of the theorem by proving Property 3 for the above choice of C_1, \dots, C_M and D_1, \dots, D_M .

Lemma 6.10. *There exists no functionally black-box batched-parallel composition protocol for computing $\mathcal{W}_{\text{flex}}^D(\mathcal{F}_{\text{SFE}}^{C_1} \parallel \dots \parallel \mathcal{F}_{\text{SFE}}^{C_M})$ in the $(\mathcal{F}_{\text{SMT}}, \mathcal{W}_{\text{flex}}^{D_1}(\mathcal{F}_{\text{SFE}}^{C_1}), \dots, \mathcal{W}_{\text{flex}}^{D_M}(\mathcal{F}_{\text{SFE}}^{C_M}))$ -hybrid model, tolerating a static adversary actively corrupting any one of the parties.*

Proof As before, it suffices to prove that there exists no batched-parallel composition protocol π that is secure for computing $\mathcal{W}_{\text{flex}}^D(\mathcal{F}_{\text{SFE}}^{f_{\alpha_1}} \parallel \dots \parallel \mathcal{F}_{\text{SFE}}^{f_{\alpha_M}})$ in the $(\mathcal{F}_{\text{SMT}}, \mathcal{W}_{\text{flex}}^{D_1}(\mathcal{F}_{\text{SFE}}^{f_{\alpha_1}}), \dots, \mathcal{W}_{\text{flex}}^{D_M}(\mathcal{F}_{\text{SFE}}^{f_{\alpha_M}}))$ -hybrid model, for a uniformly random choice of $\alpha = (\alpha_1, \dots, \alpha_M)$.

Assume toward contradiction that such a protocol π exists, which is secure against a malicious (i.e., active) adversary corrupting any one party, and assume without loss of generality that party P_1 is corrupted. Let $(\rho, j), (\rho', j')$, and ℓ denote the values that

are assumed to exist by the fact that π is a batched-parallel composition protocol and denote by $\mathcal{W}_{\text{flex}}^{D_j}(\mathcal{F}_{\text{SFE}}^{f_{\alpha_j}})$ and $\mathcal{W}_{\text{flex}}^{D_{j'}}(\mathcal{F}_{\text{SFE}}^{f_{\alpha_{j'}}})$ the corresponding functionalities indexed by j and j' . We will denote by $\mathbf{x}_1 = (x_1^1, \dots, x_1^M) \in (\{0, 1\}^\kappa)^M$ the input of P_1 and by $\mathbf{x}_2 = (x_2^1, \dots, x_2^M) \in (\{0, 1\}^\kappa)^M$ the input of P_2 to $\mathcal{W}_{\text{flex}}^D(\mathcal{F}_{\text{SFE}}^{f_{\alpha_1}} \parallel \dots \parallel \mathcal{F}_{\text{SFE}}^{f_{\alpha_M}})$. Consider an environment that chooses all inputs to the parties uniformly at random but hands its adversary the first $\kappa - 2$ bits of the input x_2^ℓ . (Recall that the batched-parallel composition requires that P_2 inputs x_2^ℓ to $\mathcal{W}_{\text{flex}}^{D_j}(\mathcal{F}_{\text{SFE}}^{f_{\alpha_j}})$ in round ρ and to $\mathcal{W}_{\text{flex}}^{D_{j'}}(\mathcal{F}_{\text{SFE}}^{f_{\alpha_{j'}}})$ in round ρ' .)

Denote the output from the evaluation of $\mathcal{W}_{\text{flex}}^D(\mathcal{F}_{\text{SFE}}^{f_{\alpha_1}} \parallel \dots \parallel \mathcal{F}_{\text{SFE}}^{f_{\alpha_M}})$ as $(\mathbf{y}_1, \dots, \mathbf{y}_n)$, where each $\mathbf{y}_i = (y_i^1, \dots, y_i^M)$ is the output of P_i . Because all inputs are independently and uniformly distributed, the simulator gains no information on the missing (i.e., last) two bits of x_2^ℓ neither by using its knowledge of the α_ℓ , nor by the inputs and outputs of any of the parallelly composed $\mathcal{W}_{\text{flex}}^D(\mathcal{F}_{\text{SFE}}^{f_{\alpha_1}} \parallel \dots \parallel \mathcal{F}_{\text{SFE}}^{f_{\alpha_M}})$ other than its ℓ' th output. In other words, the only way that the simulator might learn additional information on the missing two last bits of the input x_2^ℓ of the honest party P_2 is from the corrupted P_1 's ℓ' th output y_1^ℓ of the ideal functionality $\mathcal{W}_{\text{flex}}^D(\mathcal{F}_{\text{SFE}}^{f_{\alpha_1}} \parallel \dots \parallel \mathcal{F}_{\text{SFE}}^{f_{\alpha_M}})$.

In the analysis below we use the following notation. Given a string $\mathbf{x} = (x_1, \dots, x_m) \in \{0, 1\}^m$, denote by $\mathbf{x}[i, \dots, j]$, for $i < j$, the substring (x_i, \dots, x_j) . Consider the following cases for the input x_1^ℓ that \mathcal{S} hands to the ℓ' th functionality in $\mathcal{W}_{\text{flex}}^D(\mathcal{F}_{\text{SFE}}^{f_{\alpha_1}} \parallel \dots \parallel \mathcal{F}_{\text{SFE}}^{f_{\alpha_M}})$:

1. $x_1^\ell[1, \dots, \kappa - 2] \neq x_2^\ell[1, \dots, \kappa - 2] \oplus \alpha_\ell[1, \dots, \kappa - 2]$. In this case, by correctness of π , the ℓ' th output y_1^ℓ equals 0^κ independently of the last two bits of x_1^ℓ . Hence, in this case the simulator is able to output the last two bits of x_2^ℓ with probability $1/4$ (i.e., the best he can do is guess).
2. $x_1^\ell[1, \dots, \kappa - 2] = x_2^\ell[1, \dots, \kappa - 2] \oplus \alpha_\ell[1, \dots, \kappa - 2]$. In this case, we consider the following event

$$E_1 : x_1^\ell[\kappa - 1, \kappa] = x_2^\ell[\kappa - 1, \kappa] \oplus \alpha_\ell[\kappa - 1, \kappa].$$

- If E_1 occurs, then the simulator will see that the output $y_1^\ell \neq 0^\kappa$, so can output $x_2^\ell[\kappa - 1, \kappa] = x_1^\ell[\kappa - 1, \kappa] \oplus \alpha_\ell[\kappa - 1, \kappa]$ as his guess for the last two bits of x_2^ℓ , which will always (with probability 1) be the correct guess (by definition of the function). Note that $\Pr[E_1] = 1/4$ since the simulator knows α_ℓ (by the definition of FBB [57]) and has no information on $x_2^\ell[\kappa - 1, \kappa]$.
- If E_1 does not occur, then \mathcal{S} will see that $y_1^\ell = 0^\kappa$, from which it can deduce that $x_2^\ell[\kappa - 1, \kappa] \neq x_1^\ell[\kappa - 1, \kappa] \oplus \alpha_\ell[\kappa - 1, \kappa]$, but gets no more information on $x_2^\ell[\kappa - 1, \kappa]$. Hence, the probability of outputting $x_2^\ell[\kappa - 1, \kappa]$ in this case is at most $1/3$ (i.e., the probability of guessing among the 2-bit strings that are not equal to $x_1^\ell[\kappa - 1, \kappa] \oplus \alpha_\ell[\kappa - 1, \kappa]$).

Hence, the total probability that a simulator outputs a correct guess for $x_2^\ell[\kappa - 1, \kappa]$ is

$$P_S \leq 1/4 + (3/4)(1/3) = 1/2.$$

To complete the proof, we will describe an adversary who outputs the two last bits of x_2^ℓ , i.e., $x_2^\ell[\kappa - 1, \kappa]$, with probability noticeably higher than $1/2$; this implies a noticeable distinguishing advantage between the real world and the ideal world.

The adversary chooses two different random two-bit strings $\mathbf{b}, \mathbf{b}' \in \{0, 1\}^2$ and inputs $(x_2^\ell[1, \dots, \kappa - 2], \mathbf{b}) \oplus \alpha_j$ on the execution of $\mathcal{W}_{\text{flex}}^{D_j}(\mathcal{F}_{\text{SFE}}^{f_{\alpha_j}})$ in round ρ and $(x_2^\ell[1, \dots, \kappa - 2], \mathbf{b}') \oplus \alpha_{j'}$ on the $\mathcal{W}_{\text{flex}}^{D_{j'}}(\mathcal{F}_{\text{SFE}}^{f_{\alpha_{j'}}})$ in round ρ' . Once the adversary receives P_1 's outputs from the above two functionalities (denote them as $\hat{y}_{1,j}$ and $\hat{y}_{1,j'}$)¹³ he does the following: If $\hat{y}_{1,j} \neq 0^\kappa$ or $\hat{y}_{1,j'} \neq 0^\kappa$ (an event that happens with probability $1/2$ since there are four possible two-bit strings and one of them makes the output $\neq 0^\kappa$) then the adversary outputs $\mathbf{b} \oplus \alpha_j[\kappa - 1, \kappa]$ (or $\mathbf{b}' \oplus \alpha_{j'}[\kappa - 1, \kappa]$, respectively) as his guess of the last two bits of x_2^ℓ . By correctness of the protocol, except for negligible probability the guess will be correct in this case. Otherwise, the adversary outputs a random string from $T = \{00, 01, 10, 00\} \setminus \{\mathbf{b}, \mathbf{b}'\}$; the probability of outputting a correct guess in this case is $1/2$ since it has to be one of the strings in T . Hence, the overall probability that this adversary outputs the right guess for the last two bits of x_2 is $3/4 - \nu$, where ν is a negligible function implied by the error probabilities in the above protocol. Hence, the output of the adversary is distinguishable from the output of the best simulator which contradicts the assumed security of π . \square

This completes the proof of Theorem 6.7. \square

A. Preliminaries (Cont'd)

A.1. Error-Correcting Secret Sharing

Definition A.1. A (t, n) error-correcting secret sharing (ECSS) scheme over a message space \mathcal{M} consists of a pair of algorithms (**Share**, **Recon**) satisfying the following properties:

1. *t*-privacy For every $m \in \mathcal{M}$, and every subset $\mathcal{I} \subseteq [n]$ of size $|\mathcal{I}| \leq t$, the distribution of $\{s_i\}_{i \in \mathcal{I}}$ is independent of m , where $(s_1, \dots, s_n) \leftarrow \text{Share}(m)$.
2. *Reconstruction from up to t erroneous shares* For every $m \in \mathcal{M}$, every $s = (s_1, \dots, s_n)$, and every $s' = (s'_1, \dots, s'_n)$ such that $\Pr_{S \leftarrow \text{Share}(m)} [S = s] > 0$ and $|\{i \mid s_i = s'_i\}| \geq n - t$, it holds that $m = \text{Recon}(s')$ (except for a negligible probability).

ECSS can be constructed information-theoretically, with a negligible positive error probability, when $t < n/2$ [12,22,56].

A.2. Information-Theoretic Signatures

Parts of the following section are taken almost verbatim from [42].

¹³Recall that, by definition of probabilistic-termination SFE, the adversary is always able to learn the output in the second round.

Key Generation: The algorithm for key generation $\text{Gen}(1^\kappa, n, \ell_S)$ is as follows:

1. For $(j, k) \in \{0, \dots, n-1\} \times \{0, \dots, \ell_S\}$, choose $a_{ij} \in_R F$ uniformly at random and set the signing key to be (the description of) the multi-variate polynomial

$$\text{sk} := f(y_1, \dots, y_{n-1}, x) = \sum_{k=0}^{\ell_S} a_{0,k} x^k + \sum_{j=1}^{n-1} \sum_{k=0}^{\ell_S} a_{j,k} y_j x^k.$$

2. For $i \in [n]$, choose vector $\vec{v}_i = (v_{i,1}, \dots, v_{i,n-1}) \in_R F^{n-1}$ uniformly at random and set the i 'th verification key to be

$$\text{vk}_i = (\vec{v}_i, f_{\vec{v}_i}(x)),$$

where $f_{\vec{v}_i}(x) = f(v_{i,1}, \dots, v_{i,n-1}, x)$.

Signature Generation: The algorithm for signing a message $m \in F$, given the above signing key, is (a description of) the following polynomial

$$\text{Sign}(m, \text{sk}) := g(y_1, \dots, y_{n-1}) := f(y_1, \dots, y_{n-1}, m)$$

Signature Verification: The algorithm for verifying a signature $\sigma = g(y_1, \dots, y_n)$ on a given message m using the i 'th verification key is

$$\text{Ver}(m, \sigma, \text{vk}_i) = \begin{cases} 1, & \text{if } g(\vec{v}_i) = f_{\vec{v}_i}(m) \\ 0, & \text{otherwise} \end{cases}$$

Fig. 6. Construction of information-theoretic signatures [60].

\mathcal{P} -verifiable Information-Theoretic Signatures We recall the definition and construction of information-theoretic signatures [58, 59] but slightly modify the terminology to what we consider to be more intuitive. The signature scheme (in particular the key-generation algorithm) needs to know the total number of verifiers or alternatively the list \mathcal{P} of their identities. Furthermore, as usually for information-theoretic primitives, the key-length needs to be proportional to the number of times that the key is used. Therefore, the scheme is parameterized by two natural numbers ℓ_S and ℓ_V which will be upper bounds on the number of signatures that can be generated and verified, respectively, without violating the security.

A \mathcal{P} -verifiable signature scheme consists of a triple of randomized algorithms $(\text{Gen}, \text{Sign}, \text{Ver})$, where:

1. $\text{Gen}(1^\kappa, n, \ell_S, \ell_V)$ outputs a pair (sk, vk) , where $\text{sk} \in \{0, 1\}^\kappa$ is a signing key, $\text{vk} = (\text{vk}_1, \dots, \text{vk}_n) \in (\{0, 1\}^\kappa)^n$ is a verification-key-vector consisting of (private) verification sub-keys, and $\ell_S, \ell_V \in \mathbb{N}$.
2. $\text{Sign}(m, \text{sk})$ on input a message m and the signing-key sk outputs a signature $\sigma \in \{0, 1\}^{\text{poly}(\kappa)}$.
3. $\text{Ver}(m, \sigma, \text{vk}_i)$ on input a message m , a signature σ and a verification sub-key vk_i , outputs a decision bit $d \in \{0, 1\}$.

Definition A.2. A \mathcal{P} -verifiable signature scheme $(\text{Gen}, \text{Sign}, \text{Ver})$ is said to be *information-theoretically* (ℓ_S, ℓ_V) -secure if it satisfies the following properties:

- (completeness) A valid signature is accepted from any honest receiver:

$$\Pr[\text{Gen} \rightarrow (\text{sk}, (\text{vk}_1, \dots, \text{vk}_n)); \text{ for } i \in [n] : (\text{Ver}(m, \text{Sign}(m, \text{sk}), \text{vk}_i) = 1)] = 1.$$

– Let $\mathcal{O}_{\text{sk}}^S$ denote a signing oracle (on input m , $\mathcal{O}_{\text{sk}}^S$ outputs $\sigma = \text{Sign}(m, \text{sk})$) and $\mathcal{O}_{\text{vk}}^V$ denote a verification oracle (on input (m, σ, i) , $\mathcal{O}_{\text{vk}}^V$ outputs $\text{Ver}(m, \sigma, \text{vk}_i)$). Also, let $\mathcal{A}^{\mathcal{O}_{\text{sk}}^S, \mathcal{O}_{\text{vk}}^V}$ denote a computationally unbounded adversary that makes at most ℓ_S calls to $\mathcal{O}_{\text{sk}}^S$ and at most ℓ_V calls to $\mathcal{O}_{\text{vk}}^V$, and gets to see the verification keys indexed by some subset $\mathcal{I} \subsetneq [n]$. The following properties hold:

- (unforgeability) $\mathcal{A}^{\mathcal{O}_{\text{sk}}^S, \mathcal{O}_{\text{vk}}^V}$ cannot generate a valid signature on message m' of his choice, other than the one he queries $\mathcal{O}_{\text{sk}}^S$ on (except with negligible probability). Formally,

$$\Pr \left[\begin{array}{l} \text{Gen} \rightarrow (\text{sk}, \text{vk}); \text{ for some } \mathcal{I} \subsetneq [n] \text{ chosen by } \mathcal{A}^{\mathcal{O}_{\text{sk}}^S, \mathcal{O}_{\text{vk}}^V} : \\ \left(A^{\mathcal{O}_{\text{sk}}^S, \mathcal{O}_{\text{vk}}^V}(\text{vk}|_{\mathcal{I}}) \rightarrow (m, \sigma, j) \right) \wedge (m \text{ was not queried to } \mathcal{O}_{\text{sk}}^S) \wedge \\ (j \in [n] \setminus \mathcal{I}) \wedge (\text{Ver}(m, \sigma, \text{vk}_j) = 1) \end{array} \right] \\ = \text{negl.}$$

- (consistency)¹⁴ $\mathcal{A}^{\mathcal{O}_{\text{sk}}^S, \mathcal{O}_{\text{vk}}^V}$ cannot create a signature that is accepted by some (honest) party and rejected by some other even after seeing ℓ_S valid signatures and verifying ℓ_V signatures (except with negligible probability). Formally,

$$\Pr \left[\begin{array}{l} \text{Gen} \rightarrow (\text{sk}, \text{vk}); \text{ for some } \mathcal{I} \subsetneq [n] \text{ chosen by } \mathcal{A}^{\mathcal{O}_{\text{sk}}^S, \mathcal{O}_{\text{vk}}^V}(\text{sk}) : \\ A^{\mathcal{O}_{\text{sk}}^S, \mathcal{O}_{\text{vk}}^V}(\text{sk}, \text{vk}|_{\mathcal{I}}) \rightarrow (m, \sigma) \\ \exists i, j \in [n] \setminus \mathcal{I} \text{ s.t. } \text{Ver}(m, \sigma, \text{vk}_i) \neq \text{Ver}(m, \sigma, \text{vk}_j) \end{array} \right] \\ = \text{negl.}$$

In [59,60] a signature scheme satisfying the above notion of security was constructed. These signatures have a deterministic signature generation algorithm **Sign**. In the following (Fig. 6) we describe the construction from [59] (as described by [60] but for a single signer). We point out that the keys and signatures in the described scheme are elements of a sufficiently large finite field F (i.e., $|F| = O(2^{\text{poly}(\kappa)})$); one can easily derive a scheme for strings of length $\ell = \text{poly}(\kappa)$ by applying an appropriate encoding: for example, map the i 'th element of F to the i 'th string (in the lexicographic order) and vice versa. We say that a value σ is a *valid signature* on message m (with respect to a given key setup (sk, vk)), if for every honest P_i it holds that $\text{Ver}(m, \sigma, \text{vk}_i) = 1$.

Theorem A.3. ([60]) *Assuming $|F| = \Omega(2^\kappa)$ and $\ell_S = \text{poly}(\kappa)$ the above signature scheme (Fig. 6) is an information-theoretically $(\ell_S, \text{poly}(\kappa))$ -secure \mathcal{P} -verifiable signature scheme.*

¹⁴This property is often referred to as transferability.

B. Synchronous Protocols in UC (Cont'd)

In this section, we give complementary material to Sect. 2.1 and in particular we include a high-level overview of the formulation of synchronous UC from [46]. More concretely, Katz et al. [46] introduced a framework for universally composable synchronous computation. For self containment we describe here the basics of the model and introduce some terminology that simplifies the description of corresponding functionalities.

Synchronous protocols can be cast as UC protocols which have access to a special clock functionality $\mathcal{F}_{\text{CLOCK}}$, which allows them to coordinate round switches as described below, and communicate over bounded-delay channels.¹⁵ In a nutshell, the clock functionality works as follows: It stores a bit b which is initially set to 0 and it accepts from each party two types of messages: `CLOCK-UPDATE` and `CLOCK-READ`. The response to `CLOCK-READ` is the value of the bit b to the requester. Each `CLOCK-UPDATE` is forwarded to the adversary, but it is also recorded, and upon receiving such a `CLOCK-UPDATE` message from all honest parties, the clock functionality updates b to $b \oplus 1$. It then keeps working as above, until it receives again a `CLOCK-UPDATE` message from all honest parties, in which case it resets b to $b \oplus 1$ and so on.

Such a clock can be used as follows to ensure that honest parties remain synchronized, i.e., no honest party proceeds to the next round before all (honest) parties have finished the current round: Every party stores a local variable where it keeps (its view of) the current value of the clock indicator b . At the beginning of the protocol execution this variable is 0 for all parties. In every round, every party uses all its activations (i.e., messages it receives) to complete all its current-round instructions, and only then sends `CLOCK-UPDATE` to the clock signaling to the clock that it has completed its round; following `CLOCK-UPDATE`, all future activations result to the party sending `CLOCK-READ` to the clock until its bit b is flipped; once the party observes that the bit b has flipped, it starts its next round. For the sake of clarity, we do not explicitly mention $\mathcal{F}_{\text{CLOCK}}$ in our constructions.

In [46], for each message that is to be sent in the protocol, the sender and the receiver are given access to an independent single-use channel.¹⁶ We point out, that instead of the bounded-delay channels, in this work we will assume very simple CSFs that take as input from the sender the message he wishes to send (and a default input from other parties) and deliver the output to the receiver in a fetch mode. Such a simple secure-channel SFE can be realized in a straightforward manner from bounded-delay channels and a clock $\mathcal{F}_{\text{CLOCK}}$.

As is common in the synchronous protocols literature, throughout this work we will assume that protocols have the following structure: In each round every party sends/receives a (potentially empty) message to all parties and hybrid functionalities. Such a protocol can be described in UC in a regular form using the methodology from [46] as follows: Let $\mu \in \mathbb{N}$ denote the maximum number of messages that any party

¹⁵As argued in [46], bounded-delay channels are essential as they allow parties to detect whether or not a message was sent within a round.

¹⁶As pointed out in [46], an alternative approach would be to have a multi-use communication channel; as modeling the actual communication network is out of the scope of the current work, we will use the more standard and formally treated model of single-use channels from [46].

P_i might send to all its hybrids during some round.¹⁷ Every party in the protocol uses exactly μ activations in each round. That is, once a party P_i observes that the round has changed, i.e., the indicator-bit b of the clock has being flipped, P_i starts its next round as described above. However, this round finishes only after P_i receives μ additional activations. Note that P_i uses these activations to execute his current round instructions; since μ is a bound to the number of hybrids used in any round by any party, μ activations are enough for the party to complete its round (If P_i finishes the round early, i.e., in less than μ activations, it simply does nothing until the μ activations are received.) Once μ activations are received in the current round, P_i sends `CLOCK-UPDATE` to the clock and then keeps sending `CLOCK-READ` messages, as described above, until it observes a flip of b indicating that P_i can go to the next round.

In addition to the regular form of protocol execution, Katz et al. [46] described a way of capturing in UC the property that a protocol is guaranteed to terminate in a given number of rounds. The idea is that a synchronous protocol in regular form, which terminates after r rounds, realizes the following functionality \mathcal{F} . The functionality \mathcal{F} keeps track of the number of times every honest party sends μ activations/messages and delivers output as soon as this has happened r times. More concretely, imitating an r -round synchronous protocol with μ activations per party per round, upon being instantiated, \mathcal{F} initiates a global round-counter $\tau = 0$ and an indicator variable $\tau_i := 0$ for each $P_i \in \mathcal{P}$; as soon as some party P_i sends μ messages to \mathcal{F} , while the round-counter τ is the same, \mathcal{F} sets $\tau_i := 1$ and performs the following check:¹⁸ if $\tau_i = 1$ for every honest P_i then increase $\tau := \tau + 1$ and reset $\tau_i = 0$ for all $P_i \in \mathcal{P}$. As soon as $\tau = r$, the functionality \mathcal{F} enters a “delivery” mode. In this mode, whenever a message `fetch-output` is received from some party P_i , \mathcal{F} outputs to P_i its output. (If \mathcal{F} has no output to P_i is outputs \perp .)

We refer to a functionality that has the above structure, i.e., which keeps track of the current round τ by counting how many times every honest party has sent a certain number μ of messages, as a *synchronous functionality*. To simplify the description of our functionalities, we introduce the following terminology. We say that a *synchronous functionality* \mathcal{F} is in round ρ if the current value of the above internal counter in \mathcal{F} is $\tau = \rho$.

We note that protocols in the synchronous model of [46] enjoy the strong composition properties of the UC framework. However, in order to have protocols being executed in a lock-step mode, i.e., where all protocols complete their round within the same clock-tick, Katz et al. [46] make use of the composition with joint-state (JUC) [9]. The idea is the parties use an $\mathcal{F}_{\text{CLOCK}}$ -hybrid protocol $\hat{\pi}$ that emulates toward each of the protocols, sub-clocks and assigns to each sub-clock a unique sub-session ID (`ssid`). Each of these sub-clocks is local to its calling protocol, but $\hat{\pi}$ makes sure that it gives a `CLOCK-UPDATE` to the actual (joint) clock functionality $\mathcal{F}_{\text{CLOCK}}$, only when all sub-clocks have received such a `CLOCK-UPDATE` message. This ensures that all clocks will switch their internal bits at the same time with the bigger clock, which means that the protocols using them will

¹⁷In the simple case where the parties only use point-to-point channels, $\mu = 2(n - 1)$, since each party uses $n - 1$ channels as sender and $n - 1$ as receiver to exchange his messages for each round with all other n parties.

¹⁸To make sure that the simulator can keep track of the round index, \mathcal{F} notifies \mathcal{S} about each received input, unless it has reached its delivery state defined below.

Functionality $\mathcal{F}_{\text{CSF}}^{f,l}(\mathcal{P})$

\mathcal{F}_{CSF} proceeds as follows, parametrized by a function $f: (\{0,1\}^* \cup \{\perp\})^{n+1} \times R \rightarrow (\{0,1\}^*)^n$ and a leakage function $l: (\{0,1\}^* \cup \{\perp\})^n \rightarrow \{0,1\}^*$, and running with parties $\mathcal{P} = \{P_1, \dots, P_n\}$ and \mathcal{S} .

- Initially, set the input values x_1, \dots, x_n , output values y_1, \dots, y_n , and adversary's value a to \perp .
- In round $\rho = 1$:
 - Upon receiving (**adv-input**, sid, v) from the adversary, set $a \leftarrow v$.
 - Upon receiving a message (**input**, sid, v) from some party $P_i \in \mathcal{P}$, set $x_i \leftarrow v$ and send (**leakage**, $\text{sid}, P_i, l(x_1, \dots, x_n)$) to the adversary.
- In round $\rho = 2$:
 - Upon receiving (**adv-input**, sid, v) from the adversary, if $y_1 = \dots = y_n = \perp$, set $a \leftarrow v$. Otherwise, discard the message.
 - Upon receiving (**fetch-output**, sid) from some party $P_i \in \mathcal{P}$, if $y_1 = \dots = y_n = \perp$, then choose $r \leftarrow R$ and compute $(y_1, \dots, y_n) = f(x_1, \dots, x_n, a, r)$. Next, send (**output**, sid, y_i) to P_i and (**fetch-output**, sid, P_i) to the adversary.

Fig. 7. The canonical synchronous functionality.

be mutually synchronized. This property can be formally proved by a direct application of the JUC theorem. For further details the interested reader is referred to [9, 46].

C. The Probabilistic-Termination Framework (Cont'd)

In this section, we provide supplementary material for Sect. 2.2.

C.1. Canonical Synchronous Functionalities

The description of the *canonical synchronous functionality (CSF)* is given in Fig. 7. As a generalization of the SFE functionality, CSFs are powerful enough to capture any deterministic well-formed functionality. In fact, all the basic (unwrapped) functionalities considered in this work will be CSFs.

The functionality \mathcal{F}_{CSF} is parametrized by a (possibly) randomized function f that receives $n + 1$ inputs (n inputs from the parties and one additional input from the adversary) and a leakage function l that determines what information about the input values is leaked to the adversary. In the first (input) round, all the parties hand \mathcal{F}_{CSF} their input values, and in the second (output) round, each party receives its output. Whenever some input is submitted to \mathcal{F}_{CSF} , the adversary is given some leakage function of this input; the adversary can use this leakage for deciding which parties to corrupt and which input values to use for corrupted parties. For example, in a broadcast protocol such as [25] the adversary may decide to adaptively corrupt the broadcaster and replace its message based on the leakage received. Additionally, the adversary is allowed to input an extra message that for some functionalities (e.g., Byzantine agreement) might affect the output.

We now describe a few standard functionalities from the MPC literature as CSFs, we refer the reader to [16] for additional examples.

- SECURE MESSAGE TRANSMISSION (AKA SECURE CHANNEL). In the *secure message transmission* (SMT) functionality, a sender P_i with input x_i sends its input to a receiver P_j . The function to compute is $f_{\text{SMT}}^{i,j}(x_1, \dots, x_n, a) = (\lambda, \dots, x_i, \dots, \lambda)$ (where x_i is the value of the j 'th coordinate) and the leakage function is $l_{\text{SMT}}^{i,j}(x_1, \dots, x_n) = y$, where $y = |x_i|$ in case P_j is honest and $y = x_i$ in case P_j is corrupted. We denote by $\mathcal{F}_{\text{SMT}}^{i,j}$ the functionality \mathcal{F}_{CSF} when parametrized with the above functions $f_{\text{SMT}}^{i,j}$ and $l_{\text{SMT}}^{i,j}$, for sender P_i and receiver P_j .
- BROADCAST. In the (standard) *broadcast* functionality, a sender P_i with input x_i distributes its input to all the parties, i.e., the function to compute is $f_{\text{BC}}^i(x_1, \dots, x_n, a) = (x_i, \dots, x_i)$. The adversary only learns the length of the message x_i before its distribution, i.e., the leakage function is $l_{\text{BC}}^i(x_1, \dots, x_n) = |x_i|$. This means that the adversary does not gain new information about the input of an honest sender before the output value for all the parties is determined, and in particular, the adversary *cannot* corrupt an honest sender and change its input *after* learning the input message. We denote by $\mathcal{F}_{\text{BC}}^i$ the functionality \mathcal{F}_{CSF} when parametrized with the above functions f_{BC}^i and l_{BC}^i , for sender P_i .
- SECURE FUNCTION EVALUATION. In the *secure function evaluation* functionality, the parties compute a randomized function $g(x_1, \dots, x_n)$, i.e., the function to compute is $f_{\text{SFE}}^g(x_1, \dots, x_n, a) = g(x_1, \dots, x_n)$. The adversary learns the length of the input values via the leakage function, i.e., the leakage function is $l_{\text{SFE}}(x_1, \dots, x_n) = (|x_1|, \dots, |x_n|)$. We denote by $\mathcal{F}_{\text{SFE}}^g$ the functionality \mathcal{F}_{CSF} when parametrized with the above functions f_{SFE}^g and l_{SFE} , for computing the n -party function g .

C.2. Reactive CSFs

We proceed to extend the notion of CSF to *reactive* CSFs (RCSFs), i.e., CSFs with multiple input/output phases. Correspondingly, a reactive CSF is parametrized by two vectors of functions $\mathbf{f} = (f_1, \dots, f_q)$ and $\mathbf{l} = (l_1, \dots, l_q)$. The description of reactive CSFs can be found in Fig. 8.

Definition C.1. Let $\mathcal{F}_{\text{RCSF}}^{\mathbf{f}, \mathbf{l}}$ be a reactive CSF, with $\mathbf{f} = (f_1, \dots, f_q)$ and $\mathbf{l} = (l_1, \dots, l_q)$, let $t < n/2$, and let $\Pi = (\text{Share}, \text{Recon})$ be a (t, n) error-correcting secret-sharing scheme. For every $k \in [q]$, denote by \tilde{f}_k the function that on inputs $(\tilde{x}_1, \dots, \tilde{x}_n, a)$, with $\tilde{x}_i = (x_i, s_i)$, first reconstructs the state $s = \text{Recon}(s_1, \dots, s_n)$, next samples random coins r and computes $(y_1, \dots, y_n) = f_k(s, x_1, \dots, x_n, a, r)$, and finally shares the new state $s' = (s, x_1, \dots, x_n, a, r)$ as $(s'_1, \dots, s'_n) \leftarrow \text{Share}(s')$ and outputs $\tilde{y}_i = (y_i, s'_i)$ to P_i .

A protocol π UC-realizes $\mathcal{W}_{\text{strict}}^{\mathbf{D}}(\mathcal{F}_{\text{RCSF}}^{\mathbf{f}, \mathbf{l}})$, for a vector of distributions $\mathbf{D} = (D_1, \dots, D_q)$, if π consists of q sub-protocols (π_1, \dots, π_q) , such that for every $k \in [q]$, sub-protocol π_k UC-realizes $\mathcal{W}_{\text{strict}}^{\mathbf{D}_k}(\mathcal{F}_{\text{CSF}}^{\tilde{f}_k, l_k})$. In addition, each party P_i in π keeps a value s_i , initially set to \perp , that is used as the second input for each sub-protocol. Upon completing the execution of each sub-protocol, party P_i updates s_i to be the second output value received.

Functionality $\mathcal{F}_{\text{RCSF}}^{f,l}(\mathcal{P})$

$\mathcal{F}_{\text{RCSF}}$ proceeds as follows, parametrized by two vectors of functions $\mathbf{f} = (f_1, \dots, f_q)$ and $\mathbf{l} = (l_1, \dots, l_q)$, where $f_k: ((\{0,1\}^* \cup \{\perp\})^{n+1} \times R)^k \rightarrow (\{0,1\}^*)^n$ and $l_k: ((\{0,1\}^* \cup \{\perp\})^n)^k \rightarrow \{0,1\}^*$ for each $k \in [q]$, and running with parties $\mathcal{P} = \{P_1, \dots, P_n\}$ and an adversary \mathcal{S} .

- Initially, set the input values $x_{1,1}, \dots, x_{n,q}$, the output values $y_{1,1}, \dots, y_{n,q}$ and the adversary's values a_1, \dots, a_q to \perp ; in addition, set the state $\mathbf{s}_0 = \perp$.
- In round $\rho = 2k - 1$, with $k \in [q]$:
 - Upon receiving $(\text{adv-input}, \text{sid}, v)$ from the adversary, set $a_k \leftarrow v$.
 - Upon receiving a message $(\text{input}, \text{sid}, v)$ from some party $P_i \in \mathcal{P}$, set $x_{i,k} \leftarrow v$ and send $(\text{leakage}, \text{sid}, P_i, l_k(x_{1,1}, \dots, x_{n,k}))$ to the adversary.
- In round $\rho = 2k$, with $k \in [q]$:
 - Upon receiving $(\text{adv-input}, \text{sid}, v)$ from the adversary, if $y_{1,k} = \dots = y_{n,k} = \perp$, set $a_k \leftarrow v$. Otherwise, discard the message.
 - Upon receiving $(\text{fetch-output}, \text{sid})$ from some party $P_i \in \mathcal{P}$, if $y_{1,k} = \dots = y_{n,k} = \perp$, then choose $r_k \leftarrow R$ and compute $(y_{1,k}, \dots, y_{n,k}) = f(\mathbf{s}_{k-1}, x_{1,k}, \dots, x_{n,k}, a_k, r_k)$. Next, set, $\mathbf{s}_k = (\mathbf{s}_{k-1}, x_{1,k}, \dots, x_{n,k}, a_k, r_k)$ send $(\text{output}, \text{sid}, y_{i,k})$ to P_i and $(\text{fetch-output}, \text{sid}, P_i)$ to the adversary.

Fig. 8. The reactive canonical synchronous functionality.

C.3. Execution Traces

As discussed in Sect. 2.2, a computation with probabilistic termination is modeled in [16] by augmenting CSFs that capture the actual computational task, with *output-round randomizing wrappers* that capture the round structure of the protocol realizing the task. The underlying idea is to ensure that an environment, that can always observe how many rounds the execution of a protocol takes, will see an indistinguishable round structure in the corresponding ideal computation. Before formally describing the wrappers in “Appendix C.4”, we illustrate the notion of *execution traces*, which is central to enable the simulator emulating the round structure.

Recall that the randomizing wrappers are parametrized by a *round sampler* D , that may depend on a specific protocol implementing the functionality. The round sampler D samples a round number ρ_{term} by which all parties are guaranteed to receive their outputs no matter what the adversary strategy is. The strict wrapper $\mathcal{W}_{\text{strict}}$ ensures that all (honest) parties terminate together in the round ρ_{term} , whereas the flexible wrapper $\mathcal{W}_{\text{flex}}$ allows the adversary to deliver outputs to individual parties at any time before round ρ_{term} .

Consider an arbitrary functionality \mathcal{F} that is realized by some protocol π . If \mathcal{F} is to provide guaranteed termination (whether probabilistic or not), it must enforce an upper bound on the number of rounds that elapse until all parties receive their outputs. If the termination round of π is not fixed (but may depend on random choices made during its execution), this upper bound must be chosen according to the distribution induced by π .

Thus, in order to simulate correctly, the functionality \mathcal{F} and π 's simulator \mathcal{S} must coordinate the termination round, and therefore, \mathcal{F} must pass the upper bound it samples to \mathcal{S} . However, it is not sufficient to simply inform the simulator about the guaranteed-termination upper bound ρ_{term} . Intuitively, the reason is that protocol π may make

probabilistic choices as to the order in which it calls its hybrids (and, even worse, these hybrids may even have probabilistic termination themselves). Thus, \mathcal{F} needs to sample the upper bound based on π and the protocols realizing the hybrids called by π . As \mathcal{S} needs to emulate the entire protocol execution, it is now left with the task of trying to sample the protocol's choices conditioned on the upper bound it receives from \mathcal{F} . In general, however, it is unclear whether such a reverse sampling can be performed in (strict) polynomial time.

To avoid this issue and allow for an efficient simulation, we have \mathcal{F} output all the coins that were used for sampling the round ρ_{term} to \mathcal{S} . Because \mathcal{S} knows the round-sampler algorithm, it can reproduce the entire computation of the sampler and use it in its simulation. In fact, as we discuss below, it suffices for our proofs to have \mathcal{F} output a *trace* of its choices to the simulator instead of all the coins that were used to sample this trace.

Execution traces As mentioned above, in the synchronous communication model, the execution of the ideal functionality must take the same number of rounds as the protocol. For example, suppose that the functionality \mathcal{F} in our illustration above is used as a hybrid by a higher-level protocol π' . The functionality \mathcal{G} realized by π' must, similarly to \mathcal{F} , choose an upper bound on the number of rounds that elapse before parties obtain their outputs. However, this upper bound now depends not only on π' itself but also on π (in particular, when π is a probabilistic-termination protocol).

Given the above, the round sampler of a functionality needs to keep track of how the functionality was realized. This can be achieved via the notion of a *trace*. A trace basically records which hybrids were called by a protocol, and in a recursive way, for each hybrid, which hybrids would have been called by a protocol realizing that hybrid. The recursion ends with the hybrids that are “assumed” by the model, called *atomic* functionalities (in this paper atomic functionalities are the secure point-to-point communication functionality and the correlated-randomness functionality for broadcast).

Building on our running illustration above, suppose protocol π' (realizing \mathcal{G}) makes ideal hybrid calls to \mathcal{F} and to some atomic functionality \mathcal{H} . Assume that in an example execution, π' happens to make (sequential) calls to instances of \mathcal{H} and \mathcal{F} in the following order: \mathcal{F} , then \mathcal{H} , and finally \mathcal{F} again. Moreover, assume that \mathcal{F} is replaced by protocol π (realizing \mathcal{F}) and that π happens to make two (sequential) calls to \mathcal{H} upon the first invocation by π' , and three (sequential) calls to \mathcal{H} the second time. (We assume that both π and π' call exactly one hybrid in every round.) Then, this would result in the trace depicted in Fig. 9.

Assume that π is a probabilistic-termination protocol and π' a deterministic-termination protocol. Consequently, this means that \mathcal{F} is in fact a flexibly wrapped functionality of some CSF \mathcal{F}' , i.e., $\mathcal{F} = \mathcal{W}_{\text{flex}}^{D_{\mathcal{F}}}(\mathcal{F}')$, where the distribution $D_{\mathcal{F}}$ samples (from a distribution induced by π) depth-1 traces with root $\mathcal{W}_{\text{flex}}^{D_{\mathcal{F}}}(\mathcal{F}')$ and leaves \mathcal{H} .¹⁹ Similarly, \mathcal{G} is a strictly wrapped functionality of some CSF \mathcal{G}' , i.e., $\mathcal{G} = \mathcal{W}_{\text{strict}}^{D_{\mathcal{G}}}(\mathcal{G}')$, where the distribution $D_{\mathcal{G}}$ first samples (from a distribution induced by π') a depth-1

¹⁹Note that the root node of the trace sampled from $D_{\mathcal{F}}$ is merely *labeled* by $\mathcal{W}_{\text{flex}}^{D_{\mathcal{F}}}(\mathcal{F}')$, i.e., this is not a circular definition.




Fig. 9. Example of an execution trace.




Fig. 10. An execution trace with probabilistic-termination and deterministic-termination protocols.

trace with root $\mathcal{W}_{\text{strict}}^{D_G}(G')$ and leaves $\mathcal{W}_{\text{flex}}^{D_F}(F')$ as well as \mathcal{H} . Then, each leaf node $\mathcal{W}_{\text{flex}}^{D_F}(F')$ is replaced by a trace (independently) sampled from D_F . Thus, the example trace from Fig. 9 would look as in Fig. 10.

C.4. Strict and Flexible Wrappers

This section contains the definitions of the strict wrapper $\mathcal{W}_{\text{strict}}$ and of the flexible wrapper $\mathcal{W}_{\text{flex}}$.

Strict-wrapper functionality The *strict-wrapper functionality*, formally defined in Fig. 11, is parametrized by (a sampler that induces) a distribution D over traces, and internally runs a copy of a CSF functionality \mathcal{F} . Initially, a trace T is sampled from D ; this trace is given to the adversary once the first honest party provides its input. The trace T is used by the wrapper to define the termination round $\rho_{\text{term}} \leftarrow c_{\text{tr}}(T)$. In the first round, the wrapper forwards all the messages from the parties and the adversary to (and from) the functionality \mathcal{F} . Next, the wrapper essentially waits until round ρ_{term} , with the exception that the adversary is allowed to send (`adv-input`, `sid`, \cdot) messages and change its input to the function computed by the CSF. Finally, when round ρ_{term} arrives, the wrapper provides the output generated by \mathcal{F} to all parties.

Flexible-wrapper functionality The *flexible-wrapper functionality*, defined in Fig. 12, follows in similar lines as the strict wrapper. The difference is that the adversary is allowed to instruct the wrapper to deliver the output to each party at any round. In order to accomplish this, the wrapper assigns a termination indicator `termi`, initially set to 0, to each party. Once the wrapper receives an (`early-output`, `sid`, \cdot) request from the adver-

Wrapper Functionality $\mathcal{W}_{\text{strict}}^D(\mathcal{F})$

$\mathcal{W}_{\text{strict}}$, parametrized by an efficiently sampleable distribution D , internally runs a copy of \mathcal{F} and proceeds as follows:

- Initially, sample a trace $T \leftarrow D$ and compute the output round $\rho_{\text{term}} \leftarrow c_{\text{tr}}(T)$. Send $(\text{trace}, \text{sid}, T)$ to the adversary.^a
- At all times, forward $(\text{adv-input}, \text{sid}, \cdot)$ messages from the adversary to \mathcal{F} .
- In round $\rho = 1$: Forward $(\text{input}, \text{sid}, \cdot)$ messages from each party $P_i \in \mathcal{P}$ to \mathcal{F} . In addition, forward $(\text{leakage}, \text{sid}, \cdot)$ messages from \mathcal{F} to the adversary.
- In rounds $\rho > 1$: Upon receiving a message $(\text{fetch-output}, \text{sid})$ from some party $P_i \in \mathcal{P}$, proceed as follows:
 - If $\rho = \rho_{\text{term}}$, forward the message to \mathcal{F} , and the response $(\text{output}, \text{sid}, y_i)$ to P_i .
 - Else, send $(\text{fetch-output}, \text{sid}, P_i)$ to the adversary.

^aTechnically, the trace is sent to the adversary at the first activation of the functionality along with the first message.

Fig. 11. The strict-wrapper functionality.

Wrapper Functionality $\mathcal{W}_{\text{flex}}^D(\mathcal{F})$

$\mathcal{W}_{\text{flex}}$, parametrized by an efficiently sampleable distribution D , internally runs a copy of \mathcal{F} and proceeds as follows:

- Initially, sample a trace $T \leftarrow D$ and compute the output round $\rho_{\text{term}} \leftarrow c_{\text{tr}}(T)$. Send $(\text{trace}, \text{sid}, T)$ to the adversary.^a In addition, initialize termination indicators $\text{term}_1, \dots, \text{term}_n \leftarrow 0$.
- At all times, forward $(\text{adv-input}, \text{sid}, \cdot)$ messages from the adversary to \mathcal{F} .
- In round $\rho = 1$: Forward $(\text{input}, \text{sid}, \cdot)$ messages from each party $P_i \in \mathcal{P}$ to \mathcal{F} . In addition, forward $(\text{leakage}, \text{sid}, \cdot)$ messages from \mathcal{F} to the adversary.
- In rounds $\rho > 1$:
 - Upon receiving $(\text{fetch-output}, \text{sid})$ from some party $P_i \in \mathcal{P}$, proceed as follows:
 - * If $\text{term}_i = 1$ or $\rho = \rho_{\text{term}}$ (and P_i did not receive output yet), forward the message to \mathcal{F} , and the output $(\text{output}, \text{sid}, y_i)$ to P_i .
 - * Else, send $(\text{fetch-output}, \text{sid}, P_i)$ to the adversary.
 - Upon receiving $(\text{early-output}, \text{sid}, P_i)$ from the adversary, set $\text{term}_i \leftarrow 1$.

^aTechnically, the trace is sent to the adversary at the first activation of the functionality along with the first message.

Fig. 12. The flexible wrapper functionality.

sary for P_i , it sets $\text{term}_i \leftarrow 1$. Now, when a party P_i sends a $(\text{fetch-output}, \text{sid})$ request, the wrapper checks if $\text{term}_i = 1$, and lets the party receive its output in this case (by forwarding the $(\text{fetch-output}, \text{sid})$ request to \mathcal{F}). When the guaranteed-termination round ρ_{term} arrives, the wrapper provides the output to all parties that did not receive it yet.

Wrapper Functionality $\mathcal{W}_{\text{sl-strict}}^{D,c}(\mathcal{F})$

$\mathcal{W}_{\text{sl-strict}}$, parametrized by an efficiently sampleable distribution D and a non-negative integer c , internally runs a copy of \mathcal{F} and proceeds as follows:

- Initially, sample a trace $T \leftarrow D$ and compute the output round $\rho_{\text{term}} \leftarrow B_c \cdot c_{\text{tr}}(T)$, where $B_c = 3c+1$. Send $(\text{trace}, \text{sid}, T)$ to the adversary.^a Initialize slack counters $c_1, \dots, c_n \leftarrow 0$.
- At all times, forward $(\text{adv-input}, \text{sid}, \cdot)$ messages from the adversary to \mathcal{F} .
- In rounds $\rho = 1, \dots, 2c + 1$: Upon receiving a message from some party $P_i \in \mathcal{P}$, proceed as follows:
 - If the message is $(\text{input}, \text{sid}, \cdot)$, forward it to \mathcal{F} , forward the $(\text{leakage}, \text{sid}, \cdot)$ message \mathcal{F} subsequently outputs to the adversary, and set P_i 's local slack $c_i \leftarrow \rho - 1$.
 - Else, send $(\text{fetch-output}, \text{sid}, P_i)$ to the adversary.
- In rounds $\rho > 2c + 1$: Upon receiving a message $(\text{fetch-output}, \text{sid})$ from some party $P_i \in \mathcal{P}$, proceed as follows:
 - If $\rho = \rho_{\text{term}} + c_i$, send the message to \mathcal{F} , and the output $(\text{output}, \text{sid}, y_i)$ to P_i .
 - Else, send $(\text{fetch-output}, \text{sid}, P_i)$ to the adversary.

^aTechnically, the trace is sent to the adversary at the first activation of the functionality along with the first message.

Fig. 13. The slack-tolerant strict-wrapper functionality.

C.5. Slack-Tolerant Wrappers

Slack-tolerant strict wrapper The *slack-tolerant strict wrapper* $\mathcal{W}_{\text{sl-strict}}^{D,c}$, formally defined in Fig. 13, is parametrized by an integer $c \geq 0$, which denotes the amount of slack tolerance that is added, and a distribution D over traces. The wrapper $\mathcal{W}_{\text{sl-strict}}$ is similar to $\mathcal{W}_{\text{strict}}$ but allows parties to provide input within a window of $2c + 1$ rounds and ensures that they obtain output with the same slack they started with. The wrapper essentially increases the termination round by a factor of $B_c = 3c + 1$, which is due to the slack-tolerance technique used to implement the wrapped version of the atomic parallel SMT functionality (see [16]).

Slack-tolerant flexible wrapper The *slack-tolerant flexible wrapper* $\mathcal{W}_{\text{sl-flex}}^{D,c}$, formally defined in Fig. 14, is parametrized by an integer $c \geq 0$, which denotes the amount of slack tolerance that is added, and a distribution D over traces. The wrapper $\mathcal{W}_{\text{sl-flex}}$ is similar to $\mathcal{W}_{\text{flex}}$ but allows parties to provide input within a window of $2c + 1$ rounds and ensures that all honest parties will receive their output within two consecutive rounds. The wrapper essentially increases the termination round to

$$\rho_{\text{term}} = B_c \cdot c_{\text{tr}}(T) + 2 \cdot \text{flex}_{\text{tr}}(T) + c,$$

where the blow-up factor B_c is as explained above, and the additional factor of 2 results from the termination protocol of Bracha [7] used for every flexibly wrapped CSF, which increases the round complexity by at most two additional rounds (recall that $\text{flex}_{\text{tr}}(T)$ denotes the number of such CSFs), and c is due to the potential slack. $\mathcal{W}_{\text{sl-flex}}$ allows the adversary to deliver output at any round prior to ρ_{term} but ensures that all parties

Wrapper Functionality $\mathcal{W}_{\text{sl-flex}}^{D,c}(\mathcal{F})$

$\mathcal{W}_{\text{sl-flex}}$, parametrized by an efficiently sampleable distribution D and a non-negative integer c , internally runs a copy of (the public-output functionality) \mathcal{F} and proceeds as follows:

- Initially, sample a trace $T \leftarrow D$ and compute the output round $\rho_{\text{term}} \leftarrow B_c \cdot c_{\text{tr}}(T) + 2 \cdot \text{flex}_{\text{tr}}(T) + c$, where $B_c = 3c + 1$. Send $(\text{trace}, \text{sid}, T)$ to the adversary.^a Initialize termination indicators $\text{term}_1, \dots, \text{term}_n \leftarrow 0$.
- At all times, forward $(\text{adv-input}, \text{sid}, \cdot)$ messages from the adversary to \mathcal{F} .
- In rounds $\rho = 1, \dots, 2c + 1$: Upon receiving a message from some party $P_i \in \mathcal{P}$, proceed as follows:
 - If the message is $(\text{input}, \text{sid}, \cdot)$, send it to \mathcal{F} and forward the $(\text{leakage}, \text{sid}, \cdot)$ message \mathcal{F} subsequently outputs to the adversary.
 - Else, send $(\text{fetch-output}, \text{sid}, P_i)$ to the adversary.
- In rounds $\rho > 2c + 1$:
 - Upon receiving a message $(\text{fetch-output}, \text{sid})$ from some party $P_i \in \mathcal{P}$, proceed as follows:
 - * If $\text{term}_i = 1$ or $\rho = \rho_{\text{term}}$, forward the message to \mathcal{F} , and the output $(\text{output}, \text{sid}, y)$ to P_i . Record the output value y .
 - * Else, output $(\text{fetch-output}, \text{sid}, P_i)$ to the adversary.
 - Upon receiving $(\text{get-output}, \text{sid})$ from the adversary, if the output value y was not recorded yet, send $(\text{fetch-output}, \text{sid})$ to \mathcal{F} on behalf of some party P_i . Next, send $(\text{output}, \text{sid}, y)$ to the adversary.
 - Upon receiving $(\text{early-output}, \text{sid}, P_i)$ from the adversary, set $\text{term}_i \leftarrow 1$ and $\rho_{\text{term}} \leftarrow \min\{\rho_{\text{term}}, \rho + 1\}$.

^aTechnically, the trace is sent to the adversary at the first activation of the functionality along with the first message.

Fig. 14. The slack-tolerant flexible wrapper functionality.

obtain output with a slack of at most one round. Moreover, it allows the adversary to obtain the output using the $(\text{get-output}, \text{sid})$ command, which is necessary in order to simulate the termination protocol.

C.6. Compilers and Composition Theorems

Deterministic-termination compiler Let $\mathcal{F}, \mathcal{F}_1, \dots, \mathcal{F}_m$ be canonical synchronous functionalities, and let π be an SNF protocol that UC-realizes the strictly wrapped functionality $\mathcal{W}_{\text{strict}}^D(\mathcal{F})$, for some depth-1 distribution D , in the $(\mathcal{F}_1, \dots, \mathcal{F}_m)$ -hybrid model, assuming that all honest parties receive their inputs at the same round. The compiler $\text{Comp}_{\text{DT}}^c$, parametrized with a slack parameter $c \geq 0$, receives as input the protocol π and distributions D_1, \dots, D_m over traces and replaces every call to a CSF \mathcal{F}_i with a call to the wrapped CSF $\mathcal{W}_{\text{sl-strict}}^{D_i,c}(\mathcal{F}_i)$. We denote the output of the compiler by $\pi' = \text{Comp}_{\text{DT}}^c(\pi, D_1, \dots, D_m)$.

The compiled protocol π' realizes $\mathcal{W}_{\text{sl-strict}}^{D^{\text{full}},c}(\mathcal{F})$, for a suitably adapted distribution D^{full} , assuming all parties start within $c + 1$ consecutive rounds. Consequently, the compiled protocol π' can handle a slack of up to c rounds while using hybrids that are realizable themselves. Calling the wrapped CSFs instead of the original CSFs $\mathcal{F}_1, \dots, \mathcal{F}_m$

affects the trace corresponding to \mathcal{F} . The new trace $D^{\text{full}} = \text{full-trace}(D, D_1, \dots, D_m)$ is obtained as follows:

1. Sample a trace $T \leftarrow D$, which is a depth-1 tree with a root label $\mathcal{W}_{\text{strict}}^D(\mathcal{F})$ and leaves from the set $\{\mathcal{F}_1, \dots, \mathcal{F}_m\}$.
2. Construct a new trace T' with a root label $\mathcal{W}_{\text{sl-strict}}^{D^{\text{full},c}}(\mathcal{F})$.
3. For each leaf node $\mathcal{F}' = \mathcal{F}_i$, for some $i \in [m]$, sample a trace $T_i \leftarrow D_i$ and append the trace T_i to the first layer in T' (i.e., replace the node \mathcal{F}' with T_i).
4. Output the resulting trace T' .

To illustrate what a full trace is, consider a two-round SFE protocol π that uses a broadcast channel in both rounds (e.g., the protocol from [20] that is adaptively secure and guarantees output delivery). That is, the SNF protocol π is defined in the \mathcal{F}_{PBC} -hybrid model, where \mathcal{F}_{PBC} is an uninstantiable two-round CSF.²⁰ The distribution D_π outputs a trace consisting of a root labeled with $\mathcal{W}_{\text{strict}}^{D_\pi}(\mathcal{F}_{\text{SFE}})$, and two (ordered) children $(\mathcal{F}_{\text{PBC}}, \mathcal{F}_{\text{PBC}})$. Now, consider the compiled protocol $\pi' = \text{Comp}_{\text{DT}}^c(\pi, D_1, D_2)$, where D_1 and D_2 are the distributions corresponding to the protocol of Dolev and Strong [25], consisting of $t + 1$ point-to-point rounds, i.e., a root labeled with $\mathcal{W}_{\text{sl-strict}}^{D_1,c}(\mathcal{F}_{\text{PBC}})$, and $(t + 1)(3c + 1)$ (ordered) children $(\mathcal{F}_{\text{PSMT}}, \dots, \mathcal{F}_{\text{PSMT}})$ (recall that each round is replaced by $3c + 1$ rounds to support a slack of c rounds). Then, the full trace $D_\pi^{\text{full}} = \text{full-trace}(D_\pi, D_1, D_2)$ outputs a root labeled with $\mathcal{W}_{\text{sl-strict}}^{D_\pi^{\text{full},c}}(\mathcal{F}_{\text{SFE}})$ with two (ordered) children $\mathcal{W}_{\text{sl-strict}}^{D_1,c}(\mathcal{F}_{\text{PBC}})$ and $\mathcal{W}_{\text{sl-strict}}^{D_2,c}(\mathcal{F}_{\text{PBC}})$, each has $(t + 1)(3c + 1)$ (ordered) children $(\mathcal{F}_{\text{PSMT}}, \dots, \mathcal{F}_{\text{PSMT}})$. (To simplify the example and focus on the meaning of a full trace, we swipe some details under the rug, such as the correlated randomness functionality that is needed for the protocol from [20] and that the protocol from [25] achieves unfair broadcast rather than broadcast.)

Probabilistic-termination compiler Let $\mathcal{F}, \mathcal{F}_1, \dots, \mathcal{F}_m$ be canonical synchronous functionalities, and let π be an SNF protocol that UC-realizes the flexibly wrapped functionality $\mathcal{W}_{\text{flex}}^D(\mathcal{F})$ in the $(\mathcal{F}_1, \dots, \mathcal{F}_m)$ -hybrid model, for some depth-1 distribution D , assuming all parties start at the same round. Define the following compiler $\text{Comp}_{\text{PTR}}^c$, parametrized by a slack parameter $c \geq 0$. The compiler receives as input the protocol π , distributions D_1, \dots, D_m over traces, and a subset $I \subseteq [m]$ indexing which CSFs \mathcal{F}_i are to be wrapped with $\mathcal{W}_{\text{sl-flex}}$ and which with $\mathcal{W}_{\text{sl-strict}}$; every call in π to a CSF \mathcal{F}_i is replaced with a call to the wrapped CSF $\mathcal{W}_{\text{sl-flex}}^{D_i,c}(\mathcal{F}_i)$ if $i \in I$ or to $\mathcal{W}_{\text{sl-strict}}^{D_i,c}(\mathcal{F}_i)$ if $i \notin I$.

In addition, the compiler adds the termination procedure, based on an approach originally suggested by Bracha [7], which ensures all honest parties will terminate within two consecutive rounds:

- As soon as a party is ready to output a value y (according to the prescribed protocol) or upon receiving at least $t + 1$ messages (`end`, `sid`, y) for the same value y (whichever happens first), it sends (`end`, `sid`, y) to all parties.

²⁰A recent lower bound shows that Byzantine agreement is unlikely to be realized in two rounds, even in expectation [19].

- Upon receiving $n - t$ messages (end , sid , y) for the same value y , a party outputs y as the result of the computation and halts.

This termination technique only applies to public-output functionalities; therefore, only CSFs with public output can be wrapped by $\mathcal{W}_{\text{sl-flex}}$. We denote the output of the compiler by $\pi' = \text{Comp}_{\text{PTR}}^c(\pi, D_1, \dots, D_m, I)$.

The compiled protocol π' UC-realizes the wrapped functionality $\mathcal{W}_{\text{sl-flex}}^{D^{\text{full}},c}(\mathcal{F})$, for the adapted distribution $D^{\text{full}} = \text{full-trace}(D, D_1, \dots, D_m)$. Consequently, the compiled protocol π' can handle a slack of up to c rounds, while using hybrids that are realizable themselves, and ensuring that the output slack is at most one round (as opposed to π). Calling the wrapped hybrids instead of the CSFs affects the trace corresponding to \mathcal{F} in exactly the same way as in the case with deterministic termination.²¹

The probabilistic-termination compiler $\text{Comp}_{\text{PTR}}^c$ is suitable for SNF protocols that implement a flexibly wrapped functionality, e.g., the (adjusted) protocol of Feldman and Micali [27] that realizes randomized Byzantine agreement. Indeed, such protocols introduce new slack; hence, the slack-reduction technique described above is needed to control the new slack and reduce it to $c = 1$. As pointed out in [16], in some cases the SNF protocol may realize a strictly wrapped functionality; however, some of the hybrids are to be wrapped using the flexible wrapper. An example for the latter type of probabilistic-termination protocols is the BGW protocol [6] that has deterministic termination in the broadcast model, yet, once the broadcast channel is implemented using randomized protocols, the obtained protocol has probabilistic termination. For that reason, a second probabilistic-termination compiler $\text{Comp}_{\text{PT}}^c$, without the slack-reduction procedure, was introduced in [16].

To continue with our illustration, consider again a two-round SFE protocol π that uses a broadcast channel in both rounds, where the distribution D_π outputs a trace consisting of a root labeled with $\mathcal{W}_{\text{strict}}^{D_\pi}(\mathcal{F}_{\text{SFE}})$ and two (ordered) children (\mathcal{F}_{PBC} , \mathcal{F}_{PBC}). Now, instead of using the protocol of Dolev and Strong [25], consider the parallel-broadcast protocol from Theorem 3.5 that consists of r constant-round phases, followed by two rounds for slack-reduction; the number of rounds in a phase is denoted by const , and the number of phases r is sampled from a geometric distribution. The compiled protocol is $\pi' = \text{Comp}_{\text{PT}}^c(\pi, D_1, D_2)$, where D_1 and D_2 are the distributions that output a root labeled with $\mathcal{W}_{\text{sl-flex}}^{D_1,c}(\mathcal{F}_{\text{PBC}})$, and $r \cdot \text{const} \cdot (3c + 1) + 2$ (ordered) children ($\mathcal{F}_{\text{PSMT}}, \dots, \mathcal{F}_{\text{PSMT}}$), where r is sampled from the corresponding geometric distribution (recall that each round is replaced by $3c + 1$ rounds to support a slack of c rounds). Then, the full trace $D_\pi^{\text{full}} = \text{full-trace}(D_\pi, D_1, D_2)$ outputs a root labeled with $\mathcal{W}_{\text{sl-flex}}^{D_\pi^{\text{full}},c}(\mathcal{F}_{\text{SFE}})$ with two (ordered) children $\mathcal{W}_{\text{sl-flex}}^{D_1,c}(\mathcal{F}_{\text{PBC}})$ and $\mathcal{W}_{\text{sl-flex}}^{D_2,c}(\mathcal{F}_{\text{PBC}})$, where the first child has $r_1 \cdot \text{const} \cdot (3c + 1) + 2$ (ordered) children ($\mathcal{F}_{\text{PSMT}}, \dots, \mathcal{F}_{\text{PSMT}}$) and the second child has $r_2 \cdot \text{const} \cdot (3c + 1) + 2$ (ordered) children ($\mathcal{F}_{\text{PSMT}}, \dots, \mathcal{F}_{\text{PSMT}}$), where r_1 and r_2 are independently sampled from the corresponding geometric distribution.

²¹Of course, the root of the trace T sampled from D is a flexibly wrapped functionality $\mathcal{W}_{\text{sl-flex}}^{D,c}(\mathcal{F})$ in the probabilistic-termination case.

References

- [1] G. Asharov, Y. Lindell, Utility dependence in correct and fair rational secret sharing, in: S. Halevi (ed.) *CRYPTO 2009*, volume 5677 of LNCS (Springer, Heidelberg, 2009), pp. 559–576
- [2] G. Asharov, A. Jain, A. López-Alt, E. Tromer, V. Vaikuntanathan, D. Wichs, Multiparty computation with low communication, computation and interaction via threshold FHE, in D. Pointcheval, T. Johansson (eds.) *EUROCRYPT 2012*, volume 7237 of LNCS (Springer, Heidelberg, 2012), pp. 483–501
- [3] D. Beaver, S. Micali, P. Rogaway, The round complexity of secure protocols (extended abstract), in *22nd ACM STOC* (ACM Press, 1990), pp. 503–513
- [4] M. Ben-Or, Another advantage of free choice: Completely asynchronous agreement protocols (extended abstract), in R.L. Probert, N.A. Lynch, N. Santoro (eds.) *2nd ACM PODC* (ACM, 1983), pp. 27–30
- [5] M. Ben-Or, R. El-Yaniv, Resilient-optimal interactive consistency in constant time. *Distributed Computing*, 16(4):249–262 (2003)
- [6] M. Ben-Or, S. Goldwasser, A. Wigderson, Completeness theorems for non-cryptographic fault-tolerant distributed computation (extended abstract), in *20th ACM STOC* (ACM Press, 1988), pp. 1–10
- [7] G. Bracha, An asynchronous $[(n - 1)/3]$ -resilient consensus protocol, in R.L. Probert, N.A. Lynch, N. Santoro (eds.) *3rd ACM PODC* (ACM, 1984), pp. 154–162
- [8] R. Canetti, Universally composable security: a new paradigm for cryptographic protocols, in *42nd FOCS* (IEEE Computer Society Press, 2001), pp. 136–145
- [9] R. Canetti, T. Rabin, Universal composition with joint state, in: D. Boneh (ed.), *CRYPTO 2003*, volume 2729 of LNCS (Springer, Heidelberg, 2003), pp. 265–281
- [10] R. Canetti, J. Kilian, E. Petrank, A. Rosen, Black-box concurrent zero-knowledge requires $\omega(\log n)$ rounds, in *33rd ACM STOC* (ACM Press, 2001), pp. 570–579
- [11] R. Canetti, Y. Lindell, R. Ostrovsky, A. Sahai, Universally composable two-party and multi-party secure computation, in: *34th ACM STOC* (ACM Press, 2002), pp. 494–503
- [12] A. Cevallos, S. Fehr, R. Ostrovsky, Y. Rabani, Unconditionally-secure robust secret sharing with compact shares, in: D. Pointcheval, T. Johansson (eds.) *EUROCRYPT 2012*, volume 7237 of LNCS (Springer, Heidelberg, 2012), pp. 195–208
- [13] D. Chaum, C. Crépeau, I. Damgård, Multiparty unconditionally secure protocols (extended abstract), in: *20th ACM STOC* (ACM Press, 1988), pp. 11–19
- [14] K.-M. Chung, R. Pass, W.-L.D. Tseng, The knowledge tightness of parallel zero-knowledge, in: R. Cramer (ed.) *TCC 2012*, volume 7194 of LNCS (Springer, Heidelberg, 2012), pp. 512–529
- [15] R. Cohen, Y. Lindell, Fairness versus guaranteed output delivery in secure multiparty computation, in: P. Sarkar, T. Iwata (eds.) *ASIACRYPT 2014, Part II*, volume 8874 of LNCS (Springer, Heidelberg, 2014), pp. 466–485
- [16] R. Cohen, S. Coretti, J.A. Garay, V. Zikas, Probabilistic termination and composability of cryptographic protocols, in: M. Robshaw, J. Katz (eds.) *CRYPTO 2016, Part III*, volume 9816 of LNCS (Springer, Heidelberg, 2016), pp. 240–269
- [17] R. Cohen, I. Haitner, E. Omri, L. Rotem, Characterization of secure multiparty computation without broadcast, in: E. Kushilevitz, T. Malkin (eds.) *TCC 2016-A, Part I*, volume 9562 of LNCS (Springer, Heidelberg, 2016), pp. 596–616
- [18] R. Cohen, S. Coretti, J.A. Garay, V. Zikas, Round-preserving parallel composition of probabilistic-termination cryptographic protocols, in I. Chatzigiannakis, P. Indyk, F. Kuhn, A. Muscholl (eds.) *ICALP 2017*, volume 80 of LIPIcs (2017), pp. 37:1–37:15
- [19] R. Cohen, I. Haitner, N. Makriyannis, M. Orland, A. Samorodnitsky, On the round complexity of randomized Byzantine agreement, in: *DISC* (2019), pp. 12:1–12:17
- [20] R. Cohen, A. Shelat, D. Wichs, Adaptively secure MPC with sublinear communication complexity, in: *CRYPTO 2019, Part II*, volume 11693 of LNCS (Springer, Heidelberg, 2019), pp. 30–60
- [21] R. Cramer, I. Damgård, S. Dziembowski, M. Hirt, T. Rabin, Efficient multiparty computations secure against an adaptive adversary, in: J. Stern (ed.) *EUROCRYPT'99*, volume 1592 of LNCS (Springer, Heidelberg, 1999), pp. 311–326
- [22] R. Cramer, I. Damgård, S. Fehr, On the cost of reconstructing a secret, or VSS with optimal reconstruction phase, in: J. Kilian (ed.) *CRYPTO 2001*, volume 2139 of LNCS (Springer, Heidelberg, 2001), pp. 503–523

- [23] I. Damgård, Y. Ishai, Constant-round multiparty computation using a black-box pseudorandom generator, in: V. Shoup (ed.) *CRYPTO 2005*, volume 3621 of LNCS (Springer, Heidelberg, 2005), pp. 378–394
- [24] I. Damgård, J.B. Nielsen, Improved non-committing encryption schemes based on a general complexity assumption, in M. Bellare (ed.) *CRYPTO 2000*, volume 1880 of LNCS (Springer, Heidelberg, 2000), pp. 378–394
- [25] D. Dolev, H.R. Strong, Authenticated algorithms for Byzantine agreement, *SIAM Journal on Computing*, 12(4):656–666 (1983)
- [26] D. Dolev, R. Reischuk, H.R. Strong, Early stopping in Byzantine agreement, *Journal of the ACM*, 37(4):720–741 (1990)
- [27] P. Feldman, S. Micali, An optimal probabilistic protocol for synchronous Byzantine agreement, *SIAM Journal on Computing*, 26(4):873–933 (1997)
- [28] M.J. Fischer, N.A. Lynch, A lower bound for the time to assure interactive consistency, *Information Processing Letters*, 14(4):183–186 (1982)
- [29] M. Fitzi, J.A. Garay, Efficient player-optimal protocols for strong and differential consensus, in E. Borowsky, S. Rajsbaum (eds.) *22nd ACM PODC* (ACM, 2003), pp. 211–220
- [30] G. Fuchsbauer, J. Katz, D. Naccache, Efficient rational secret sharing in standard communication networks, in D. Micciancio (ed.) *TCC 2010*, volume 5978 of LNCS (Springer, Heidelberg, 2010), pp. 419–436
- [31] J.A. Garay, J. Katz, U. Maurer, B. Tackmann, V. Zikas, Rational protocol design: cryptography against incentive-driven adversaries, in: *54th FOCS* (IEEE Computer Society Press, 2013), pp. 648–657
- [32] S. Garg, C. Gentry, S. Halevi, M. Raykova, Two-round secure MPC from indistinguishability obfuscation, in: Y. Lindell (ed.) *TCC 2014*, volume 8349 of LNCS (Springer, Heidelberg, 2014), pp. 74–94. https://doi.org/10.1007/978-3-642-54242-8_4
- [33] O. Goldreich, S. Micali, A. Wigderson, How to play any mental game or A completeness theorem for protocols with honest majority, in: A. Aho (ed.) *19th ACM STOC* (ACM Press, 1987), pp. 218–229
- [34] S.D. Gordon, F. Liu, E. Shi, Constant-round MPC with fairness and guarantee of output delivery, in R. Gennaro, M.J.B. Robshaw (eds.) *CRYPTO 2015, Part II*, volume 9216 of LNCS (Springer, Heidelberg, 2015), pp. 63–82
- [35] A. Groce, J. Katz, Fair computation with rational players, in D. Pointcheval, T. Johansson (eds.) *EUROCRYPT 2012*, volume 7237 of LNCS (Springer, Heidelberg, 2012), pp. 81–98
- [36] I. Haitner, A parallel repetition theorem for any interactive argument, in: *50th FOCS* (IEEE Computer Society Press, 2009), pp. 241–250
- [37] J.Y. Halpern, V. Teague, Rational secret sharing and multiparty computation: extended abstract, in L. Babai (ed.) *36th ACM STOC* (ACM Press, 2004), pp. 623–632
- [38] J. Håstad, R. Pass, D. Wikström, K. Pietrzak, An efficient parallel repetition theorem, in D. Micciancio (ed.) *TCC 2010*, volume 5978 of LNCS (Springer, Heidelberg, 2010), pp. 1–18
- [39] M. Hirt, V. Zikas, Adaptively secure broadcast, in H. Gilbert (ed.) *EUROCRYPT 2010*, volume 6110 of LNCS (Springer, Heidelberg, 2010), pp. 466–485
- [40] Y. Ishai, E. Kushilevitz, R. Ostrovsky, A. Sahai, Zero-knowledge from secure multiparty computation, in: D.S. Johnson, U. Feige (eds.) *39th ACM STOC* (ACM Press, 2007), pp. 21–30
- [41] Y. Ishai, M. Prabhakaran, A. Sahai, Founding cryptography on oblivious transfer-efficiently, in: D. Wagner (ed.) *CRYPTO 2008*, volume 5157 of LNCS (Springer, Heidelberg, 2008), pp. 572–591
- [42] Y. Ishai, R. Ostrovsky, V. Zikas, Secure multi-party computation with identifiable abort, in: J.A. Garay, R. Gennaro (eds.) *CRYPTO 2014, Part II*, volume 8617 of LNCS (Springer, Heidelberg, 2014), pp. 369–386
- [43] Y. Ishai, E. Kushilevitz, M. Prabhakaran, A. Sahai, C. Yu, Secure protocol transformations, in: M. Robshaw, J. Katz (eds.) *CRYPTO 2016, Part II*, volume 9815 of LNCS (Springer, Heidelberg, 2016), pp. 430–458
- [44] J. Katz, C.-Y. Koo, On expected constant-round protocols for Byzantine agreement, in C. Dwork (ed.) *CRYPTO 2006*, volume 4117 of LNCS (Springer, Heidelberg, 2006), pp. 445–462
- [45] J. Katz, C.-Y. Koo, Round-efficient secure computation in point-to-point networks, in: M. Naor (ed.) *EUROCRYPT 2007*, volume 4515 of LNCS (Springer, Heidelberg, 2007), pp. 311–328
- [46] J. Katz, U. Maurer, B. Tackmann, V. Zikas, Universally composable synchronous computation, in: A. Sahai (ed.) *TCC 2013*, volume 7785 of LNCS (Springer, Heidelberg, 2013), pp. 477–498

- [47] J. Kilian, Founding cryptography on oblivious transfer, in: *20th ACM STOC* (ACM Press, 1988), pp. 20–31
- [48] L. Lamport, R.E. Shostak, M.C. Pease, The Byzantine generals problem, *ACM Trans. Program. Lang. Syst.*, 4(3):382–401 (1982)
- [49] Y. Lindell, A. Lysyanskaya, T. Rabin, Sequential composition of protocols without simultaneous termination, in: A. Ricciardi (ed.) *21st ACM PODC* (ACM, 2002), pp. 203–212
- [50] S. Micali, Fast and furious Byzantine agreement, in C.H. Papadimitriou (ed.) *ITCS 2017* (ACM, 2017)
- [51] P. Mukherjee, D. Wichs, Two round multiparty computation via multi-key FHE, in M. Fischlin, J.-S. Coron (eds.) *EUROCRYPT 2016, Part II*, volume 9666 of LNCS (Springer, Heidelberg, 2016), pp. 735–763
- [52] S.J. Ong, D.C. Parkes, A. Rosen, S.P. Vadhan, Fairness with an honest minority and a rational majority, in: O. Reingold (ed.) *TCC 2009*, volume 5444 of LNCS (Springer, Heidelberg, 2009), pp. 36–53
- [53] M.C. Pease, R.E. Shostak, L. Lamport, Reaching agreement in the presence of faults, *Journal of the ACM*, 27(2):228–234 (1980)
- [54] B. Pfitzmann, M. Waidner, Unconditional Byzantine agreement for any number of faulty processors, in: *STACS*, volume 577 of LNCS (Springer, 1992), pp. 339–350
- [55] M.O. Rabin, Randomized Byzantine generals, in: *24th FOCS* (1983), pp. 403–409
- [56] T. Rabin, M. Ben-Or, Verifiable secret sharing and multiparty protocols with honest majority (extended abstract), in: *21st ACM STOC* (ACM Press, 1989), pp. 73–85
- [57] M. Rosulek, Must you know the code of f to securely compute f ?, in: R. Safavi-Naini, R. Canetti (eds.) *CRYPTO 2012*, volume 7417 of LNCS (Springer, Heidelberg, 2012), pp. 87–104
- [58] T. Seito, T. Aikawa, J. Shikata, T. Matsumoto, Information-theoretically secure key-insulated multi-receiver authentication codes, in: D.J. Bernstein, T. Lange (eds.) *AFRICACRYPT 10*, volume 6055 of LNCS (Springer, Heidelberg, 2010), pp. 148–165
- [59] J. Shikata, G. Hanaoka, Y. Zheng, H. Imai, Security notions for unconditionally secure signature schemes, In L.R. Knudsen (ed.) *EUROCRYPT 2002*, volume 2332 of LNCS (Springer, Heidelberg, 2002), pp. 434–449
- [60] C. Swanson, D.R. Stinson, Unconditionally secure signature schemes revisited, in: S. Fehr (ed.) *ICITS 11*, volume 6673 of LNCS (Springer, Heidelberg, 2011), pp. 100–116
- [61] A.C.-C. Yao, Protocols for secure computations (extended abstract), in: *23rd FOCS* ((IEEE Computer Society Press, 1982), pp. 160–164