Fair and Robust Multi-Party Computation using a Global Transaction Ledger

Aggelos Kiayias
aggelos@di.uoa.gr

Hong-Sheng Zhou
hszhou@vcu.edu

Vassilis Zikas
vzikas@inf.ethz.edu

October 29, 2015

Abstract

Classical results on secure multi-party computation (MPC) imply that fully secure computation, including fairness (either all parties get output or none) and robustness (output delivery is guaranteed), is impossible unless a majority of the parties is honest. Recently, cryptocurrencies like Bitcoin where utilized to leverage the fairness loss in MPC against a dishonest majority. The idea is that when the protocol aborts in an unfair manner (i.e., after the adversary receives output) then honest parties get compensated by the adversarially controlled parties.

Our contribution is three-fold. First, we put forth a new formal model of secure MPC with compensation and we show how the introduction of suitable ledger and synchronization functionalities makes it possible to express completely such protocols using standard interactive Turing machines (ITM) circumventing the need for the use of extra features that are outside the standard model as in previous works. Second, our model, is expressed in the universal composition setting with global setup and is equipped with a composition theorem that enables the design of protocols that compose safely with each other and within larger environments where other protocols with compensation take place; a composition theorem for MPC protocols with compensation was not known before. Third, we introduce the first robust MPC protocol with compensation, i.e., an MPC protocol where not only fairness is guaranteed (via compensation) but additionally the protocol is guaranteed to deliver output to the parties that get engaged and therefore the adversary, after an initial round of deposits, is not even able to mount a denial of service attack without having to suffer a monetary penalty. Importantly, our robust MPC protocol requires only a *constant* number of (coin-transfer and communication) rounds.

1 Introduction

Secure multiparty computation (MPC) enables a set of parties to evaluate the output of a known function $f(\cdot)$ on inputs they privately contribute to the protocol execution. The design of secure MPC protocols, initiated with the seminal works of Yao [Yao82] and Goldreich et al. [GMW87] has evolved to a major effort in computer security engineering. Beyond privacy, a secure MPC protocol is highly desirable to be *fair* (either all parties learn the output or none) and *robust* (the delivery of the output is guaranteed and the adversary cannot mount a "denial of service" against the protocol). Achieving fairness and robustness in a setting where there is an arbitrary number of corruptions, as desirable as it may appear, is prohibited by strong impossibility results stemming from the work of Cleve [Cle86] who showed that coin-flipping is infeasible in any setting where there is no honest majority among parties that execute the protocol. These impossibility results, combined with the importance of the properties that they prevent, strongly motivate the exploration of alternate – yet still realistic – models that would enable fair and robust MPC protocols.

With the advent of Bitcoin [Nak08] and other decentralized cryptocurrencies, the works of [ADMM14a, ADMM14b, BK14, KB14] showed a new direction for circumvention of the impossibility results regarding the fairness property: enforcing fairness could be achieved through imposing monetary penalties. In this setting a breach of fairness by the adversary is still possible but it results in the honest parties collecting a compensation in a way that is determined by the protocol execution. At the same time, in case fairness is not breached, it is guaranteed that no party loses any money (despite the fact that currency transfers may have taken place between the parties). The rationale here is that a suitable monetary penalty suffices in most practical scenarios to force the adversary to operate in the protocol fairly.

While the main idea of fairness with penalties sounds simple enough, its implementation proves to be quite challenging. The main reason is that the way a crypto-currency operates does not readily provide a trusted party that will collect money from all participants and then either return it or redistribute it according to the pre-agreed penalty structure. This is because crypto-currencies are decentralized and hence no single party is ever in control of a money transfer beyond the owner of a set of coins. The mechanism used in [ADMM14a, ADMM14b, BK14, KB14] to circumvent the above problem is the capability[1] of the Bitcoin network to issue transactions that are "time-locked", i.e., become valid only after a specific time and prior to that time may be superseded by other transactions that are posted in the public ledger. Superseded time-locked transactions become invalid and remain in the ledger without ever being redeemed.

While the above works are an important step for the design of MPC protocols with properties that circumvent the classical impossibility results, several critical open questions remain to be tackled; those we address herein are as follows.

Our Results. Our contribution is three-fold. First, we put forth a new formal model of secure MPC with compensation and we show how the introduction of suitable ledger and synchronization functionalities makes it possible to express completely such protocols using standard interactive Turing machines (ITM) circumventing the need for the use of extra features that are outside the standard model (in comparison, the only previous model [BK14] resorted to specialized ITM's that utilize resources outside the computational model[2]). Second, our model is equipped with a composition theorem that enables the design of protocols that compose safely with each other and within larger environments where other protocols with compensation take place; a composition theorem for this class of protocols was not known before and requires a new framework for synchronization in the global UC setting that can be of independent interest. Third, we introduce the first robust MPC protocol with compensation, i.e., an MPC protocol where not only fairness is guaranteed (via compensation) but additionally the protocol is guaranteed to deliver output to the parties that get engaged and therefore the adversary is not even able to mount a denial of service attack without having to suffer a monetary penalty. In more details we have the following.

- We put forth a new model that utilizes two ideal functionalities and expresses the ledger of transactions and a clock in the sense of [KMTZ13] that is connected to the ledger and enables

[1] Note that this feature is currently not fully supported.

[2] An ITM with the special features of "wallet" and "safe" was introduced in [BK14] to express the ability of ITM's to store and transfer "coins." Such coins were treated as physical quantities that were moved between players but also locked in safes in a way that parties were then prevented to use them in certain ways (in other words such safes were not local but were affected from external events).

parties to synchronize their protocol interactions. Our ledger functionality enable us to abstract all the necessary features of the underlying cryptocurrency. Contrary to the only previous formalization approach [BK14, KB14], our modeling allows the entities that participate in an MPC execution to be regular interactive Turing machines (ITM) and there is no need to equip them with additional physical features such as “safes” and “locks.” Furthermore the explicit inclusion of the clock functionality (which is only alluded to in [BK14,KB14]) and a synchronous framework for protocol design given such clock reveal the exact dependencies between the ledger and the clock functionality that are necessary in order for MPC with compensation protocols to be properly described. We express our model within a general framework that we call Q-fairness and robustness and may be of independent interest as it can express meaningful relaxations of fairness and robustness in the presence of a global ideal functionality.

- We prove a composition theorem that establishes that protocols in our framework are secure in a universally composable fashion. Our composition proof treats the clock and ledger functionalities as global setups in the sense of [CDPW07, CJS14]. We emphasize that this is a critical design choice: the fact that the ledger is a global functionality ensures that any penalties that are incurred to the adversary that result to credits towards the honest parties will be globally recognized. This should be contrasted to an approach that utilizes regular ideal functionalities which may be only accessible within the scope of a single protocol instance and hence any penalty bookkeeping they account may vanish with the completion of the protocol. Providing a composition theorem for MPC protocols with compensation was left as an open question in [BK14].
- We finally present a new protocol for fair and robust secure MPC with compensation. The robustness property we prove guarantees that once the protocol passes an initial round of deposits, parties are guaranteed to obtain output or be compensated. This is in contrast to fair MPC with compensation [ADMM14a, ADMM14b, BK14, KB14] where the guarantee is that compensation takes place only in case the adversary obtains output while an honest party does not. To put it differently, it is feasible for the adversary to lead the protocol to a deadlock where no party receives output however the honest parties have wasted resources by introducing transactions in the ledger. We remark that it is in principle possible to upgrade the protocols of [ADMM14a, ADMM14b, BK14, KB14] to the robust MPC setting by having them perform an MPC with identifiable abort, cf. [GMW87, IOZ14], (in such protocol the party that causes the abort can be identified and excluded from future executions). However even using such protocol the resulting robust MPC with compensation will need in the worst case a *linear* number of deposit/communication rounds in the number of malicious parties. Contrary to that, our robust protocol can be instantiated so that it requires a constant number of deposit/communication rounds independently of the number of parties that are running the protocol. Our construction uses time-locked transactions in a novel way to ensure that parties do progress in the MPC protocol or otherwise transactions are suitably revertible to a compensation for the remaining parties. The structure of our transactions is quite more complex than what can be presently supported by bitcoin; we describe in high-level how our protocol can be implemented via Ethereum[3] contracts.

Related work. In addition to the previous works [ADMM14a, ADMM14b, BK14, KB14] in fair MPC with compensation, very recently, Ruffing et al. [RKS15] address equivocation issues via penalty mechanism, and design decentralized “non-equivocation” contracts.

[3] `http://www.ethereum.org`.

There are a number of other works that attempted to circumvent the impossibility results for fairness in the setting of dishonest majority by considering alternate models. Contrary to the approach based on cryptocurrencies these works give an advantage to the protocol designer with respect to the adversarial strategy for corruption. For instance, in [GKM$^+$13] a rational adversary is proposed and the protocol designer is privy to the utility function of the adversary. In [ALZ13] a reputation system is used and the protocol designer has the availability of the reputation information of the parties that will be engaged in the protocol. Finally in [GGJ$^+$15] a two tiered model is proposed where the protocol designer is capable of distinguishing two distinct sets of servers at the onset of the computation that differ in terms of their corruptibility.

Global setups were first put forth in [CDPW07] motivated by notion of deniability in cryptographic protocols. In our work we utilize global functionalities for universal composition (without the deniability aspect) as in [CJS14] where a similar approach was taken for the case of the use of the random oracle as a global setup functionality for MPC.

Fairness was considered from the resource perspective, cf. [BN00, Pin03, GMPY06], where it is guaranteed due to the investment of proportional resources between the parties running the protocol, and the optimistic perspective, cf. [ASW97, ASW98, CC00], where a trusted mediator can be invoked in the case of an abort. We finally note that without any additional assumptions, due to the impossibility results mentioned above, one can provide fairness only with certain high probability that will be affecting the complexity of the resulting protocol, see, e.g., [GK09] and references therein.

In concurrent and independent work, Kosba et al [KMS$^+$15] propose a framework for composable protocols based on a ledger. and explore a notion of fairness with compensation. Our work goes beyond fairness and provides a treatment of robustness. Furthermore we provide a synchronous framework with a global clock (of independent interest) that uses the ledger as a *global* setup to achieve fairness and robustness and we prove a composition theorem for our framework.

Organization. We start with preliminaries in Section 2. Then in Sections 3 and 4, we lay down a formal framework for designing composable fair protocols in the presence of globally available trusted resources. In Section 3, we introduce two shared functionalities $\bar{\mathcal{G}}_{\text{CLOCK}}$ and $\bar{\mathcal{G}}_{\text{LEDGER}}$ respectively to formulate the trust resources that are provided by Bitcoin-like systems. Subsequently, in Section 4, we put forth a new formal framework for secure MPC with compensation: we introduce the notions of Q-fairness, and Q-robustness via wrapper functionalities; we then consider the realization of such wrapper functionalities, and further provide a composition theorem. In Section 5, we present a protocol in our new framework to achieve our new notions of fairness and robustness. Implementing our protocol within Ethereum is discussed in the section 6. Proofs of our theorems are presented in Appendix A.

2 Preliminaries

Throughout the paper we assume an (often implicit) security parameter denoted as κ. For a number $n \in \mathbb{N}$ we denote by $[n]$ the set $[n] = \{1, \ldots, n\}$ and denote by 0^n (resp. 1^n) the all-zero (resp. all-ones) string of length n. For a randomized algorithm Alg we denote by $\mathrm{Alg}(x; r)$ the output of Alg on input x and random coins r. To avoid always explicitly writing the coins r, we shall denote by $y \overset{\$}{\leftarrow} \mathrm{Alg}(x)$ the operation of running Alg on input x (and uniformly random coins) and storing the output on variable y. We write $f : X \xrightarrow{\$} Y$ to denote a probabilistic function with domain X and range Y. We use the standard definition of *negligible* and *overwhelming* (e.g., see [Gol01]).

For a multiparty function $f : (\{0,1\}^* \cup \{\lambda\})^n \rightarrow (\{0,1\}^* \cup \{\perp\})^n$ for parties in $\mathcal{P} = \{p_1, \ldots, p_n\}$ and for a set $\mathcal{P} \subseteq \mathcal{P}$, we denote by $f|_{|\mathcal{P}'|}$ the restriction of f to the parties in $\mathcal{P}'$, namely, if each $p_i \in \mathcal{P}'$ has input x_i, then the output of $f|_{|\mathcal{P}'|}$ is the output of f evaluated on inputs x_i for each $p_i \in \mathcal{P}'$ and $x_j = \lambda$ for each $p_j \in \mathcal{P} \setminus \mathcal{P}'$.

We describe our results in the extension of Canetti's UC framework [Can01] to allow for global setups, known as GUC [CDPW07]. As argued above, this is the natural model to consider execution in the present of a globally synchronized clock and a ledger/bulletin board. Consistently with the (G)UC notation, we denote local (UC) functionalities by calligraphic letters, as in $\mathcal{F}$, and add a bar to denote global functionalities, as in $\bar{\mathcal{G}}$. Furthermore, we denote by ϕ, the dummy protocol. Note that in GUC ϕ might receive inputs for its (UC) hybrids and/or for the global setup, where an implicit mechanism is assumed to allow the environment to define the intended recipient of each submitted input to ϕ. For a protocol π, a (local) UC functionality $\mathcal{F}$ and a global setup $\bar{\mathcal{G}}$ we denote by $\textsc{Exec}^{\bar{\mathcal{G}},\mathcal{F}}_{\pi,\mathcal{A},\mathcal{Z}}$ the output of the environment $\mathcal{Z}$ in an execution of π having hybrid access to $\bar{\mathcal{G}}$ and $\mathcal{F}$ in the presence of adversary $\mathcal{A}$. We assume some familiarity with the UC and/or the GUC framework.

Correlated Randomness as a Sampling Functionality Our protocols are in the *correlated randomness* model, i.e., they assume that the parties initially, before receiving their inputs, receive appropriately correlated random strings. In particular, the parties jointly hold a vector $\vec{R} = (R_1, \ldots, R_n) \in (\{0,1\}^*)^n$, where P_i holds R_i, drawn from a given efficiently samplable distribution $\mathcal{D}$. This is, as usual, captured by giving the parties initial access to an ideal functionality $\mathcal{F}^{\mathcal{D}}_{\text{CORR}}$, known as a *sampling functionality*, which, upon receiving a default input from any party, samples $\vec{R}$ from $\mathcal{D}$ and distributes it to the parties (cf. Figure 1). Hence, a protocol in the correlated randomness model is formally an $\mathcal{F}^{\mathcal{D}}_{\text{CORR}}$-hybrid protocol. Formally, a sampling functionality $\mathcal{F}^{\mathcal{D}}_{\text{CORR}}$ is parameterized by an efficiently computable sampling distribution $\mathcal{D}$ and the (ID's of the parties in) the player set $\mathcal{P}$.

Functionality $\mathcal{F}^{\mathcal{D}}_{\text{CORR}}$

Functionality $\mathcal{F}^{\mathcal{D}}_{\text{CORR}}$ interacts with a set of parties $\mathcal{P} = \{P_1, \ldots, P_n\}$, the adversary $\mathcal{S}$ and the environment $\mathcal{Z}$. The functionality is parameterized with a distribution sampler $\mathcal{D}$.

- Upon receiving $(\text{REQUEST}, \mathsf{sid})$ from any party or the adversary, set $\vec{R} = (R_1, \ldots, R_n) \leftarrow \mathcal{D}$ and for each $P_i \in \mathcal{P}$ send $(\text{REQUEST}, \mathsf{sid}, R_i)$ to P_i (or to the adversary if P_i is corrupted).

Figure 1: The correlated randomness functionality.

3 Model

In this section and next section, we lay down a formal framework for designing composable fair protocols in the presence of globally available trusted resources. we introduce in the current section, shared (in the sense of the GUC model [CDPW07]) functionalities $\bar{\mathcal{G}}_{\text{CLOCK}}$ and $\bar{\mathcal{G}}_{\text{LEDGER}}$ respectively to formulate the trust resources that are provided by Bitcoin-like systems. We stress that these two functionalities can be thought of as a single global functionality and in our description are allowed to communicate. Nonetheless, we choose to describe then as two separate functionalities, because as we argue, the clock $\bar{\mathcal{G}}_{\text{CLOCK}}$ can also be used alone (without $\bar{\mathcal{G}}_{\text{LEDGER}}$) to naturally model synchronous computation with a global notion of time.

3.1 Global Clock Functionality and Synchronous Protocol Executions

In this section we describe how to model execution of synchronous protocols that can access a global-clock setup. This is an adaptation of the original idea by Katz et al. [KMTZ13], where a clock was modelled as UC functionality that is local to the calling protocol, and is of independent interest as a model for the design of synchronous protocols. In addition to being a more realistic model for capturing time in UC, the notion of the global clock allows for synchronous execution of any protocols that choose to use it.

Before defining our clock, we recall the reader the clock and model of synchronous execution from [KMTZ13] and then highlight the main differences. The clock in [KMTZ13] is a UC functionality that keeps an indicator bit b originally set to 0. The parties can send to the clock special "update" messages, and as soon as the clock sees that all honest parties agree to update the state is sets $b := b \oplus 1$. The clock then continues to receive "update" messages, and again, as soon as the clock sees that all honest parties have requested to update after the last switch of the bit b it switches it again. To make sure that the adversary is given enough activations, whenever the clock receives an "update" message from the honest party it notifies the adversary. In addition to "update" messages, the parties can send the clock a "read" message which the clock replies with the current value of b.

The use of such a clock to keep a round structure is as follows: Whenever a party observes a switch of the bit b, it interprets it as a round advance. Thus, a synchronous protocol with access to such a clock is executed as follows. In each round, every party performs all its protocol instructions for the current round, and at the end sends an "update" message to the clock; from the point where the party updates (its round has finished) it queries ("reads") the clock with each following activation to detect when all parties have also finished their rounds (i.e., when the value of b switches). Once this happens, the party starts its next protocol round.

An issue with the above clock is that in order to execute two protocols using the same clock we need to make use of the joint-state UC theorem [CR03]. Instead, in this work we take an alternative modelling approach and define a shared clock functionality $\bar{\mathcal{G}}_{\text{CLOCK}}$. This functionality can be viewed as a shared version of the clock functionality which was defined by Katz et al. [KMTZ13]. The main intuition behind our clock functionality is that all honest parties can use it to ensure that they proceed with their rounds at the same pace. On a hight level, the clock operates as follows: any party that wishes to be synchronised with the global clock can send $(\text{REGISTER}, \mathsf{sid})$ to the clock and subsequently it can send it $(\text{CLOCK-UPDATE}, \mathsf{sid})$ commands, where sid is $\bar{\mathcal{G}}_{\text{CLOCK}}$'s identifier. The clock stores a global-time counter τ (initially set to 0), and as soon it is instructed by all currently honest parties and by associated shared functionalities[4] to advance the time (i.e., receives $(\text{CLOCK-UPDATE}, \mathsf{sid})$ it increases its state-counter τ by 1.

The main difference between our formulation and that by Katz et al. [KMTZ13] is that in [KMTZ13] the clock is a UC functionality which is local to a single protocol and waits for an "update" message by every honest party to advance its state; however, here we intend to have the clock to be ac-

[4] Certain global functionalities, such as the ledger defined in the following section, might depend on time and, therefore, need to be synchronized with the clock.

cessed *globally* and used by arbitrary protocols. Therefore we give the power to the environment to define the clock's speed. Indeed, if there are no associated shared functionalities, the environment can instruct dummy parties to send inputs $(\textsc{clock-update}, \mathsf{sid})$ to $\bar{\mathcal{G}}_{\text{CLOCK}}$ and advance the clock whenever it wishes. An additional difference is that in [KMTZ13], the clock state is binary while here, in our formulation, the state τ is a positive integer which indicates the time that has passed from point zero (i.e., from the beginning of time).

Next, we elaborate and explain how to use the global clock to design synchronous protocols. We remark that the model of synchronous protocol execution of [KMTZ13] cannot be used in our setting as the environment can make the clock advance before honest parties have time to take actions in any round. Indeed, in the ideal setting the environment can keep sending $(\textsc{clock-update}, \mathsf{sid})$ to the dummy parties, which will forward it to the clock making its state to advance; to make sure that the protocol is indistinguishable, honest parties would have to do the same, thereby giving away the activations that they need for executing their protocol instructions such as send and receive operations.[5] This might, at first, seem like a bug but it is in fact a feature. It captures the fact that since time is a quantity that should be in the control of the environment, if the environment chooses to advance time too fast then some protocol might not have enough time to perform their operations for each round, and might therefore need to give up.

To make sure that the environment cannot exploit such fast-forwarding of the clock we use the following idea: We allow the clock to receive from honest parties or (non-shared) ideal functionalities a special (Clock-Fast) message, which makes it set an internal indicator from 0 to 1. This indicator will be added onto the response of the clock to clock-read queries, and will make any synchronous protocol or corresponding functionality that reads the clock and observes this indicator being set to one to immediately terminate with a default value. This way we ensure that an environment that tries such a fast-forward distinguishing attack will be forced to make any synchronous protocol behave in a default way, a behavior which, as we see, is easily imitated in the ideal world. The detailed description of the clock functionality can be found in Figure 2.

We stress that having a global $\bar{\mathcal{G}}_{\text{CLOCK}}$-hybrid model makes the mode of execution of synchronous protocols more intuitive compared to [KMTZ13]. Here is how synchronous protocols are executed in this setting. First, as is the case in real-life synchronous protocols, we assume that the protocol participants have agreed on the starting time τ_0 of their protocol and also on the duration of each round.[6] We abstract this knowledge by assuming the parties know a function $\texttt{Round2Time} : \mathbb{Z} \to \mathbb{Z}$ which maps protocol rounds to time (according to the global clock) in which the round should be completed. For $\rho \in \mathbb{Z}$, $\texttt{Round2Time}(\rho)$ is the time in which the ρth round of the protocol should be completed. To make sure that no party proceeds to round $\rho + 1$ of the protocol before all honest parties have completed round ρ , we require that any two protocol rounds are at least two clock-ticks apart (see [KMTZ13] for a discussion); formally, for all $\rho \geq 0$, it holds that $\texttt{Round2Time}(\rho + 1) \geq \texttt{Round2Time}(\rho) + 2$.

A synchronous protocol in the above setting proceeds as follows where the parties keep locally track of the current round ρ in the protocol they are in:

- Upon receiving a $(\textsc{clock-update}, \mathsf{sid})$ input (from its environment) where sid is the ID of $\bar{\mathcal{G}}_{\text{CLOCK}}$, party P_i forwards it to $\bar{\mathcal{G}}_{\text{CLOCK}}$.

[5] The communication channels we are using are fetch-type bounded delivery channels as in [KMTZ13]. In such channels, the receiver needs to issue "fetch"-requests which are answered only if a message is ready for delivery. We refer to [KMTZ13] for details.

[6] Different protocols might proceed at a different pace.

Functionality $\bar{\mathcal{G}}_{\textsc{clock}}$

Shared functionality $\bar{\mathcal{G}}_{\textsc{clock}}$ is globally available to all participants. The shared functionality is parameterized with variables τ, a bit $d_{\bar{\mathcal{G}}_{\textsc{ledger}}}$ a set $\mathcal{P}'$ and a bit $\texttt{fast}$ and is associated with a ledger shared functionality $\bar{\mathcal{G}}_{\textsc{ledger}}$.

Initially, $\tau := 0$, $d_{\bar{\mathcal{G}}_{\textsc{ledger}}} := 0$, $\texttt{fast} := 0$ and $\mathcal{P}' := \emptyset$.

- Upon receiving $(\textsc{register}, \mathsf{sid})$ from some party P, set $\mathcal{P}' := \mathcal{P}' \cup \{P\}$ and if P was not registered before, set $d_P := 0$; subsequently, forward $(\textsc{register}, \mathsf{sid}, P)$ to $\mathcal{A}$.
- Upon receiving $(\textsc{clock-update}, \mathsf{sid})$ from $\bar{\mathcal{G}}_{\textsc{ledger}}$ set $d_{\bar{\mathcal{G}}_{\textsc{ledger}}} := 1$ and forward $(\textsc{clock-update}, \mathsf{sid}, \bar{\mathcal{G}}_{\textsc{ledger}})$ to $\mathcal{A}$
- Upon receiving $(\textsc{clock-update}, \mathsf{sid})$ from some honest party $P \in \mathcal{P}'$ set $d_i := 1$; then if $d_{\bar{\mathcal{G}}_{\textsc{ledger}}} := 1$ and $d_P = 1$ for all honest parties in $\mathcal{P}'$, then set $\tau := \tau + 1$ and reset $d_{\bar{\mathcal{G}}_{\textsc{ledger}}} := 0$ and $d_P := 0$ for all parties in $\mathcal{P}'$. Forward $(\textsc{clock-update}, \mathsf{sid}, P)$ to $\mathcal{A}$.
- Upon receiving $(\textsc{clock-read}, \mathsf{sid})$ from any participant (including the environment, the adversary, or any ideal—shared or local—functionality) return $(\textsc{clock-read}, \mathsf{sid}, \tau, \texttt{fast})$ to the requestor.
- Upon receiving (Clock-Fast) from any honest party or ideal functionality, set $\texttt{fast} := 1$.

Figure 2: The clock functionality.

– Upon receiving a $(\textsc{clock-read}, \mathsf{sid})$ input (from its environment), party P_i forwards it to $\bar{\mathcal{G}}_{\textsc{clock}}$ and outputs the response to the environment.

– Upon receiving a (Clock-Fast) input (from its environment), party P_i forwards it to $\bar{\mathcal{G}}_{\textsc{clock}}$.

– Upon receiving any message $(\textsc{Input}, \mathsf{sid}')$ where sid' is the session ID of a protocol P_i is involved in, do the following: Send $(\textsc{clock-read}, \mathsf{sid})$ to $\bar{\mathcal{G}}_{\textsc{clock}}$ and denote the response by $(\textsc{clock-read}, \mathsf{sid}, \tau, \texttt{fast})$; if $\texttt{fast} = 1$ then output Clock-Fast to the environment. Otherwise do:

  - if $\tau \leq \texttt{Round2Time}(\rho - 1)$ halt;
  - else, if $\texttt{Round2Time}(\rho - 1) < \tau \leq \texttt{Round2Time}(\rho)$ execute the next pending round$-\rho$ instruction (if all the instructions for round ρ are finished halt.)
  - else, if $\tau > \texttt{Round2Time}(\rho)$ and there are still pending instructions for the current round, send (Clock-Fast) to $\bar{\mathcal{G}}_{\textsc{clock}}$.
  - else, i.e., if $\tau > \texttt{Round2Time}(\rho)$ and P_i has completed all round-ρ instruction, then set $\rho := \rho + 1$ and halt.

It is easy to verify that the above mode of operation will guarantee that the parties are never out-of-sync, since as soon as the first party issues a Clock-Fast message for the clock, all synchronous protocols will enter the mode of outputting Clock-Fast for every input that the environment hands them (that is not intended for the clock). However, there is one more thing that needs to be taken care of. Since in the real-world the parties go to a default mode (where they output Clock-Fast to every query) when the environment does not give them sufficient time, this should also be the case in the ideal world. To achieve this we use another idea inspired by the guaranteed termination functionality from [KMTZ13]: Let π be a synchronous protocol with round-to-time function $\texttt{Round2Time} : \mathbb{Z} \to \mathbb{Z}$, where in each round, each party needs exactly m activations to perform its instructions[7]. We introduce a wrapper $\tilde{\mathcal{W}}$ which, at a high level, forwards messages to and from its wrapped functionality but stores a round-index and checks, as the protocol would, that every party issues to the wrapped functionality, at least m activations for each round ρ in the intended interval. If this is not the case the wrapper sends (Clock-Fast) to $\bar{\mathcal{G}}_{\text{CLOCK}}$ and responds with Clock-Fast form that point on. The detailed description can be found in Figure 3.

Wrapper Functionality $\tilde{\mathcal{W}}(\mathcal{F})$

The wrapper $\tilde{\mathcal{W}}(\mathcal{F})$ interacts with a set of parties $\mathcal{P} = \{P_1, \ldots, P_n\}$, the adversary $\mathcal{S}$ and the environment $\mathcal{Z}$, as well as the shared clock functionality $\bar{\mathcal{G}}_{\text{CLOCK}}$. It is also parameterised with a round-to-time function $\texttt{Round2Time} : \mathbb{Z} \to \mathbb{Z}$, and the number m of activationss that each party should receive in each round. It locally stores a round index ρ initialized to $\rho := 1$.

- Upon receiving a message from $\mathcal{S}$ it forwards it to $\mathcal{F}$.
- Upon receiving a message from $\mathcal{F}$ for $\mathcal{S}$ it forwards it to $\mathcal{S}$.
- Upon receining a message from $\mathcal{F}$ for some party P_i it forwards it to P_i.
- Upon receiving a message from an honest party, P_i sends $(\text{CLOCK-READ}, \mathsf{sid})$ to $\bar{\mathcal{G}}_{\text{CLOCK}}$ and denotes the response by $(\text{CLOCK-READ}, \mathsf{sid}, \tau, \texttt{fast})$; if $\texttt{fast} = 1$ then it outputs Clock-Fast to the environment. Otherwise do:
  - if $\tau \leq \texttt{Round2Time}(\rho - 1)$ ignore the mesage;
  - else, if $\texttt{Round2Time}(\rho - 1) < \tau \leq \texttt{Round2Time}(\rho)$ and P_i has sent less than m messages to $\tilde{\mathcal{W}}(\mathcal{F})$ in the interval $[\texttt{Round2Time}(\rho - 1) + 1, \texttt{Round2Time}(\rho)]$ then relay the message to $\mathcal{F}$.
  - else, if $\tau > \texttt{Round2Time}(\rho)$ and P_i has sent less than m messages to $\tilde{\mathcal{W}}(\mathcal{F})$ in the interval $[\texttt{Round2Time}(\rho - 1) + 1, \texttt{Round2Time}(\rho)]$ then send (Clock-Fast) to $\bar{\mathcal{G}}_{\text{CLOCK}}$.
  - else, i.e., if $\tau > \texttt{Round2Time}(\rho)$ and P_i has sent less than m messages to $\tilde{\mathcal{W}}(\mathcal{F})$ in the interval $[\texttt{Round2Time}(\rho - 1) + 1, \texttt{Round2Time}(\rho)]$, then set $\rho := \rho + 1$ and halt.

Figure 3: The Syncronized wrapper functionality.

3.2 Global Ledger Functionality

Functionality $\bar{\mathcal{G}}_{\text{LEDGER}}$ provides the abstraction of a public ledger in Bitcoin-like systems (e.g., Bitcoin, Litecoin, Namecoin, Ethereum, etc). Intuitively, the public ledger could be accessed globally by protocol parties or other entities including the environment $\mathcal{Z}$. Protocol parties or the environment can generate transactions; and these valid transactions will be gathered by a set of ledger maintainers (e.g., miners in Bitcoin-like systems) in a certain order as the state of the ledger. More concretely, whenever the ledger maintainers receive a vector of transactions $\vec{\texttt{tx}}$, they first add the transactions in a buffer, assuming they are valid with respect to the existing transactions and the state of the ledger; thus, in this way a vector of transactions is formed in the buffer. After a certain amount of time, denoted by $\texttt{T}$, which will be also referred to as a *ledger round*, all transactions in the buffer will be "glued" into the ledger state in the form of a block. The adversary is allowed to permute the buffer prior to its addition to the ledger. In Bitcoin, $\texttt{T}$ is 10 minutes (approximately); thus in about every 10 minutes, a new block of transactions will be included into the ledger, and the ledger state will be updated correspondingly.

To enable the ledger to be aware of time, the ledger maintainers are allowed to "read" the state of another publicly available functionality $\bar{\mathcal{G}}_{\text{CLOCK}}$ defined above. Furthermore, to ensure that

[7] One can make any synchronous protocol have this form by introducing dummy instructions.

Functionality $\bar{\mathcal{G}}_{\text{LEDGER}}$

Shared functionality $\bar{\mathcal{G}}_{\text{LEDGER}}$ is globally available to all participants. The shared functionality is parameterized with a predicate $\mathsf{Validate}$, a constant $\mathtt{T}$, and variables state, buffer and $\mathsf{counter}$.

Initially, $\mathsf{state} := \varepsilon$, $\mathsf{buffer} := \varepsilon$, and $\mathsf{counter} := 0$.

- Upon receiving $(\textsc{submit}, \mathsf{sid}, \vec{\mathtt{tx}})$ from some participant, If $\mathsf{Validate}(\mathsf{state}, (\mathsf{buffer}, \vec{\mathtt{tx}})) = 1$, then set $\mathsf{buffer} := \mathsf{buffer}||\vec{\mathtt{tx}}$. Go to *State Extend.*
- Upon receiving $(\textsc{read}, \mathsf{sid})$ from a party P or $\mathcal{A}$, if P is honest set $b = \mathsf{state}$ else set $b = (\mathsf{state}, \mathsf{buffer})$.
  1. Execute *State Extend.*
  2. Return $(\textsc{read}, \mathsf{sid}, b)$ to the requestor.
- Upon receiving $(\textsc{permute}, \mathsf{sid}, \pi)$ from $\mathcal{A}$ apply permutation π on the elements of buffer.

State Extend: Send $(\textsc{clock-read}, \mathsf{sid})$ to $\bar{\mathcal{G}}_{\text{CLOCK}}$ and receive $(\textsc{clock-read}, \mathsf{sid}, \tau)$ from $\bar{\mathcal{G}}_{\text{CLOCK}}$. If $|\tau - \mathtt{T} \cdot \mathsf{counter}| > \mathtt{T}$, then set $\mathsf{state} := \mathsf{state}||\mathsf{Blockify}(\tau, \mathsf{buffer})$ and $\mathsf{buffer} := \varepsilon$ and $\mathsf{counter} := \mathsf{counter} + 1$. Subsequently, send $(\textsc{clock-update}, \mathsf{sid})$ to $\bar{\mathcal{G}}_{\text{CLOCK}}$ where sid is the ID of $\bar{\mathcal{G}}_{\text{CLOCK}}$.

Figure 4: The public ledger functionality.

the ledger is activated at least once in each time-tick[8] (i.e., each advance of the $\bar{\mathcal{G}}_{\text{CLOCK}}$ state) we have the ledger, with every message it gets from a party other than the adversary, send a $(\textsc{clock-update}, \mathsf{sid})$ message to $\bar{\mathcal{G}}_{\text{CLOCK}}$. (Recall that, as defined, $\bar{\mathcal{G}}_{\text{CLOCK}}$ always waits for at least one such message from the ledger before advancing its time counter.)

We remark that all gathered transactions should be "valid" which is defined by a predicate $\mathsf{Validate}$. In different systems, predicate $\mathsf{Validate}$ will take different forms. For example, in the Bitcoin system, the predicate $\mathsf{Validate}$ should make sure that for each newly received transaction that transfers v coins from the original wallet address $\mathtt{address}_o$ to the destination wallet address $\mathtt{address}_d$, the original wallet address $\mathtt{address}_o$ should have v or more than v coins, and the transaction should be generated by the original wallet holder (as shown by the issuance of a digital signature). Furthermore, prior to each vector of transactions becoming block, the vector is passed through a function $\mathsf{Blockify}(\cdot)$ that homogenizes the sequence of transactions in the form of a block. Moreover, in some systems like Bitcoin, it may add a special transaction called a "coinbase" transaction that implements a reward mechanism for the ledger maintainers.

In Figure 4 we provide the details of the ledger functionality.

4 Q-Fairness and Q-Robustness

In this section, we provide a formal framework for secure computation with fair and robust compensation. In the spirit of [GMPY06], our main tool is a wrapper functionality. Our wrapper functionality is equipped with a predicate $\mathsf{Q}_{\bar{\mathcal{G}}}$ which is used to make sure that the outcome of the protocol execution is consistent with appropriate conditions on the state of the global setup $\bar{\mathcal{G}}$. Intuitively, the predicate $\mathsf{Q}_{\bar{\mathcal{G}}}$ works as a filter, such that if certain "bad" event occurs (e.g., an abort), then the wrapped functionality will restrict the simulators influence. More concretely, the

[8] This is essential to ensure that updates are done in a time-consistent manner.

predicate $\mathsf{Q}_{\bar{\mathcal{G}}}$ has three modes $\mathsf{Q}_{\bar{\mathcal{G}}}^{\mathrm{Init}}$, $\mathsf{Q}_{\bar{\mathcal{G}}}^{\mathrm{Dlv}}$ and, $\mathsf{Q}_{\bar{\mathcal{G}}}^{\mathrm{Abt}}$, where $\mathsf{Q}_{\bar{\mathcal{G}}}^{\mathrm{Init}}$ specifies under which condition (on the global setup's state) the protocol should start executing; $\mathsf{Q}_{\bar{\mathcal{G}}}^{\mathrm{Dlv}}$ specifies under which condition parties should receive their output; and $\mathsf{Q}_{\bar{\mathcal{G}}}^{\mathrm{Abt}}$ specifies under which condition the simulator is allowed to force parties to abort. With foresight $\mathsf{Q}_{\bar{\mathcal{G}}}^{\mathrm{Init}}$ will ensure that the protocol is executed only if all honest participants have enough coins; $\mathsf{Q}_{\bar{\mathcal{G}}}^{\mathrm{Dlv}}$ will ensure that honest parties do not lose coins if they execute the protocol; and $\mathsf{Q}_{\bar{\mathcal{G}}}^{\mathrm{Abt}}$ will ensure that honest parties might be forced to an "unfair" abort (i.e, where the adversary has received his output) only if they are compensated by earning coins (from the corrupted parties). We will call an implementation of a wrapped version of $\mathcal{F}$ a Q-fair implementation of $\mathcal{F}$. [9]

Our definition of $\mathsf{Q}_{\bar{\mathcal{G}}}$-fairness can be instantiated with respect to any global setup that upon receiving a READ symbol (from any protocol participant or functionality) it returns its public state trans. Concretely, let $\bar{\mathcal{G}}$ be global ideal functionality and let $\mathsf{Q}_{\bar{\mathcal{G}}}$ a predicate, as above, with respect to such $\bar{\mathcal{G}}$. Let also $\mathcal{F}$ be a non-reactive functionality[10] which allows for fair evaluation of a given function (SFE) in the sense of [GMPY06], i.e., it has two modes of delivering output: (i) delayed delivery: $(\textsc{deliver}, \mathsf{sid}, m, P)$ signifying delayed output delivery[11] of m to party P, (ii) fair delivery: $(\textsc{fair-deliver}, \mathsf{sid}, (m, P_{i_1}), \ldots, (m, P_{i_k}), (m_{\mathcal{S}}, \mathcal{S}))$ that results in simultaneous delivery of outputs $m_{i_1}, \ldots m_{i_k}$ to parties $P_{i_1}, \ldots, P_{i_k}$ and output $m_{\mathcal{S}}$ to $\mathcal{S}$. We note that (G)UC does not have an explicit mechanism for simultaneous delivery of outputs. Thus, when we refer to simultaneous delivery of a vector $(m_{i_1}, \ldots, m_{i_k})$ to parties $P_{i_1}, \ldots, P_{i_k}$, respectively, we imply that the functionality prepares all the output to be delivered in a "fetch mode" as defined in [KMTZ13]; that is:

- The functionality registers the pairs $(m_{i_1}, P_{i_1}), \ldots, (m, P_{i_k})$ as "ready to fetch" and sends the set $\{(m_{i_j}, P_{i_j}) | P_{i_j}$ is corrupted $\}$ to $\mathcal{S}$.
- Upon receiving an input $(\textsc{fetch-output}, P_i)$ from party P_i, if a message (m_i, P_i) has been registered as "ready to fetch" then remove it from the "ready to fetch" set and output it to P_i (if more than one such messages are registered, deliver and remove from the "ready to fetch" set the first, chronologically, registered such pair); otherwise send $(\textsc{fetch-output}, P_i)$ to $\mathcal{S}$.

4.1 $\mathsf{Q}_{\bar{\mathcal{G}}}$-Fairness

The wrapper functionality $\mathcal{W}$ that will be used in the definition of Q-fair (secure) computation is given in Figure 5. The intuition is as follows: Prior to handing inputs to the (wrapped) functionality $\mathcal{F}$, the parties can request the wrapper to generate on their behalf a *resource-setup* (by executing an associated resource-setup generation algorithm Gen) which allows them to update the global setup $\bar{\mathcal{G}}$; this resource setup consists of a public component $RS_{P,\mathsf{sid}}^{\mathtt{pub}}$ and a private component $RS_{P,sid}^{\mathtt{priv}}$. [12] Both these values are given to the simulator, and the public component is handed to the party.

From the point when parties receive their inputs the Q predicate is used as a filter to specify the wrapper's behavior and add the fairness guarantees. More concretely, upon receiving an input from a party, the wrapper checks on the global setup to ensure that $\mathsf{Q}^{\mathrm{Init}}$ is true, and if it is not true it aborts (i.e., sets all honest parties' outputs to $\perp$ and blocks any communication between $\mathcal{F}$

[9] We note that whenever it is clear from the context we may drop the subscript $\bar{\mathcal{G}}$ in $\mathsf{Q}_{\bar{\mathcal{G}}}, \mathsf{Q}_{\bar{\mathcal{G}}}^{\mathrm{Init}}, \mathsf{Q}_{\bar{\mathcal{G}}}^{\mathrm{Dlv}}, \mathsf{Q}_{\bar{\mathcal{G}}}^{\mathrm{Abt}}$.

[10] A non-reactive functionality does not accept any input from honest parties after generating output.

[11] Delayed output delivery is a standard (G)UC mechanism where the adversary is allowed to schedule the output at a time of its choosing.

[12] In the case of bitcoin-like ledgers these will correspond to a wallet (public-key) and a corresponding secret key.

and the adversary). This means that if the environment has not set up the experiment properly,[13] then the experiment will not be executed and the wrapped functionality will become useless. This formally resolves the question "What happens if some party does not have sufficient coins to play the protocol?" which leads to some ambiguity in existing bitcoin-based definitions of computation with fair compensation [BK14].

The predicates $\mathsf{Q}^{\mathrm{Dlv}}$ and $\mathsf{Q}^{\mathrm{Abt}}$ are used to filter out attempts of the simulator to deliver outputs or abort when $\mathsf{Q}^{\mathrm{Dlv}}$ and $\mathsf{Q}^{\mathrm{Abt}}$ are violated.[14] Concretely, any such attempt will be ignored if the corresponding predicate is not satisfied.

Intuitively, by requiring the protocol to implement such a wrapped version of a functionality, we will ensure that the parties might only abort if $\mathsf{Q}^{\mathrm{Abt}}$ is true, and might output a valid (non-$\perp$) value if $\mathsf{Q}^{\mathrm{Dlv}}$. As we shall see in Section 4.2, by a trivial modification of the fairness wrapper, we can capture a stronger property which we will call Q-robustness; the latter, roughly, guarantees that honest parties which start the protocol will either receive their output (and $\mathsf{Q}^{\mathrm{Dlv}}$ being true) or will abort and increase their revenue. (I.e., there is no way for the adversary to make the protocol abort after the first honest party has sent its first input-dependent message).

Definition 4.1. *We say protocol π realizes functionality $\mathcal{F}$ with $\mathsf{Q}_{\bar{\mathcal{G}}}$-fairness with respect to global functionality $\bar{\mathcal{G}}$, provided the following statement is true. For all adversaries $\mathcal{A}$, there is a simulator $\mathcal{S}$ so that for all environments $\mathcal{Z}$ it holds:*

$$\textsc{Exec}^{\bar{\mathcal{G}}}_{\pi,\mathcal{A},\mathcal{Z}} \approx \textsc{Exec}^{\bar{\mathcal{G}},\mathcal{W}_{\mathsf{Q},\bar{\mathcal{G}}}(\mathcal{F})}_{\mathcal{S},\mathcal{Z}}$$

More generally, the protocol σ realizes $\mathcal{H}$ with $\mathsf{Q}'_{\bar{\mathcal{G}}}$ fairness using a functionality $\mathcal{F}$ with fairness $\mathsf{Q}_{\bar{\mathcal{G}}}$ provided that for all adversaries $\mathcal{A}$, there is a simulator $\mathcal{S}$ so that for all environments $\mathcal{Z}$, it holds:

$$\textsc{Exec}^{\bar{\mathcal{G}},\mathcal{W}_{\mathsf{Q},\bar{\mathcal{G}}}(\mathcal{F})}_{\pi,\mathcal{A},\mathcal{Z}} \approx \textsc{Exec}^{\bar{\mathcal{G}},\mathcal{W}_{\mathsf{Q}',\bar{\mathcal{G}}}(\mathcal{H})}_{\mathcal{S},\mathcal{Z}}$$

We note that, both protocol π and the functionality $(\mathcal{W}_{\mathsf{Q},\bar{\mathcal{G}}}(\mathcal{F}), \bar{\mathcal{G}})$ are with respect to the global functionality[15] $\bar{\mathcal{G}}$. By following the very similar proof idea in [CDPW07], we can prove the following lemma and theorem:

Lemma 4.2. *Let $\mathsf{Q}_{\bar{\mathcal{G}}}$ be a predicate with respect to global functionality $\bar{\mathcal{G}}$. Let π be a protocol that realizes the functionality $\mathcal{F}$ with $\mathsf{Q}_{\bar{\mathcal{G}}}$-fairness. Let σ be a protocol in $(\mathcal{W}_{\mathsf{Q},\bar{\mathcal{G}}}(\mathcal{F}), \bar{\mathcal{G}})$-hybrid world. Then for all adversaries $\mathcal{A}$, there is a simulator $\mathcal{S}$ so that for all environments $\mathcal{Z}$, it holds*

$$\textsc{Exec}^{\bar{\mathcal{G}}}_{\sigma^{\pi},\mathcal{A},\mathcal{Z}} \approx \textsc{Exec}^{\bar{\mathcal{G}},\mathcal{W}_{\mathsf{Q},\bar{\mathcal{G}}}(\mathcal{F})}_{\sigma,\mathcal{S},\mathcal{Z}}$$

Theorem 4.3. *Let $\mathsf{Q}_{\bar{\mathcal{G}}}$ and $\mathsf{Q}'_{\bar{\mathcal{G}}}$ be predicates with respect to global functionality $\bar{\mathcal{G}}$. Let π be a protocol that realizes the functionality $\mathcal{F}$ with $\mathsf{Q}_{\bar{\mathcal{G}}}$-fairness. Let σ be a protocol in $(\mathcal{W}_{\mathsf{Q},\bar{\mathcal{G}}}(\mathcal{F}), \bar{\mathcal{G}})$-hybrid world that realizes the functionality $\mathcal{H}$ with $\mathsf{Q}'_{\bar{\mathcal{G}}}$-fairness. Then for all adversaries $\mathcal{A}$, there is a simulator $\mathcal{S}$ so that for all environments $\mathcal{Z}$ it holds:*

$$\textsc{Exec}^{\bar{\mathcal{G}}}_{\sigma^{\pi},\mathcal{A},\mathcal{Z}} \approx \textsc{Exec}^{\bar{\mathcal{G}},\mathcal{W}_{\mathsf{Q}',\bar{\mathcal{G}}}(\mathcal{H})}_{\mathcal{S},\mathcal{Z}}$$

Please see appendix for the proof details.

[13] In the case of a bitcoin-ledger this corresponds to the environment not transferring to some protocol-related wallet sufficient funds to execute the protocol.

[14] As we will see, in bitcoin-like instantiations, $\mathsf{Q}^{\mathrm{Dlv}}$ will be satisfied when no honest party has a negative balance, and $\mathsf{Q}^{\mathrm{Abt}}$ will be satisfied when every honest party has a (strictly) positive balance.

[15] In GUC framework [CDPW07], this is also called, $\bar{\mathcal{G}}$-subroutine respecting.

Wrapper Functionality $\mathcal{W}_{\mathsf{Q},\bar{\mathcal{G}}}(\mathcal{F})$

The wrapper $\mathcal{W}_{\mathsf{Q},\bar{\mathcal{G}}}(\mathcal{F})$ interacts with a set of parties $\mathcal{P} = \{P_1, \ldots, P_n\}$, the adversary $\mathcal{S}$ and the environment $\mathcal{Z}$, as well as shared functionality $\bar{\mathcal{G}}$. It is parameterized with a predicate $\mathsf{Q} = (\mathsf{Q}^{\mathrm{Init}}, \mathsf{Q}^{\mathrm{Dlv}}, \mathsf{Q}^{\mathrm{Abt}})$ and a resource-setup generating algorithm $\mathrm{Gen} : 1^* \overset{\$}{\to} (\{0,1\}^*)^2$ and wraps any given non-reactive n-party functionality $\mathcal{F}$ with the two output-delivery modes (delayed and fair) described in Section 4.1. The functionality also keeps an indicator bit b, initially set to 0, indicating whether or not $\mathcal{S}$ is blocked from sending messages to $\mathcal{F}$.

- *Allocating Resources.* Upon receiving ($\textsc{alocate}, \mathsf{sid}$) from a party P, if a message ($\textsc{alocate}, \mathsf{sid}$) has already been received for P then ignore it; else send ($\textsc{Coins}, \mathsf{sid}, P$) to $\mathcal{S}$ and upon receiving ($\textsc{Coins}, \mathsf{sid}, P, r$) from $\mathcal{S}$ compute $(RS^{\mathsf{pub}}_{P,\mathsf{sid}}, RS^{\mathtt{priv}}_{P,\mathsf{sid}}) \leftarrow \mathrm{Gen}(1^{\kappa}; r)$ and sends a delayed output ($\textsc{deliver}, \mathsf{sid}, RS^{\mathsf{pub}}_{P,\mathsf{sid}}, P$) to P.
- Upon receiving any message M from $\mathcal{F}$ to be delivered to its simulator, if $b = 0$ forward M to $\mathcal{S}$.
- Upon receiving a message ($\textsc{Forward}, M$) from $\mathcal{S}$, if $b = 0$ then forward M to $\mathcal{F}$ as a message coming from its simulator.
- *Receiving input for $\mathcal{F}$.* Upon receiving ($\textsc{input}, \mathsf{sid}, x$) from a party P, send $\textsc{read}$ to $\bar{\mathcal{G}}$, denote the response by trans and if $\neg\mathsf{Q}^{\mathrm{Init}}(RS^{\mathsf{pub}}_{P,\mathsf{sid}}, \mathsf{trans})$ then set $b := 1$ and issue a message ($\textsc{fair-deliver}, \mathsf{sid}, (\bot, P_1), \ldots, (\bot, P_n), (\bot, \mathcal{S})$) (i.e., simultaneously deliver $\bot$ to all parties and ignore all future messages except ($\textsc{fetch-output}, \cdot$) messages. Otherwise, forward ($\textsc{input}, \mathsf{sid}, x$) to $\mathcal{F}$ as input for P.
- *Generating delayed output.* Upon receiving a message from $\mathcal{F}$ marked ($\textsc{deliver}, \mathsf{sid}, m, P$) forwards m to party P via delayed output.
- *Registering fair output.* Upon receiving a message from $\mathcal{F}$ that is marked for fair delivery ($\textsc{fair-deliver}, \mathsf{sid}, \mathsf{mid}, (m_1, P_{i_1}), \ldots, (m_k, P_{i_k}), (m_{\mathcal{S}}, \mathcal{S})$), it forwards ($\mathsf{mid}, P_{i_1}, \ldots, P_{i_k}, m_{\mathcal{S}}$) to $\mathcal{S}$.
- Q-*fair delivery.* Upon receiving (Q-$\textsc{deliver}, \mathsf{sid}, \mathsf{mid}$) from $\mathcal{S}$ then provided that a message ($\mathsf{mid}, \ldots$) has been delivered to $\mathcal{S}$ operate as follows. For each pair of the form (m, P) associated with mid: Let $\mathrm{L} = \{(m, P)|\ P \text{ is uncorrupted}\}$. Send $\{(m, P)|P \text{ is corrupted}\}$ to $\mathcal{S}$. (If some currently honest P becomes corrupted later on, remove (m, P) from sending and send (m, P) to $\mathcal{S}$.) Subsequently perform the following.
  - On input a message ($\textsc{deliver}, \mathsf{sid}, \mathsf{mid}, P$) from $\mathcal{S}$, provided that the record mid contains the pair $(m, P) \in \mathrm{L}$, send $\textsc{read}$ to $\bar{\mathcal{G}}$, denote the response by trans and if $\neg\mathsf{Q}^{\mathrm{Dlv}}(\mathsf{sid}, P, RS^{\mathsf{pub}}_{P,\mathsf{sid}}, \mathsf{trans})$ then ignore the message. Else, remove (m, P) from L and register (m, P) as "ready to fetch".
  - On input a message ($\textsc{abort}, \mathsf{sid}, \mathsf{mid}, P$) from $\mathcal{S}$, provided that the record mid contains the pair $(m, P) \in \mathrm{L}$, send $\textsc{read}$ to $\bar{\mathcal{G}}$, denote the response by trans and if $\neg\mathsf{Q}^{\mathrm{Abt}}(\mathsf{sid}, P, RS^{\mathsf{pub}}_{P,\mathsf{sid}}, \mathsf{trans})$ then ignore the message. Else, remove (m, P) from L and register $(\bot, P)$ as "ready to fetch".
- Upon receiving an input ($\textsc{fetch-output}, P$) from party P, if a message (m, P) has been registered as "ready to fetch" then remove it from the "ready to fetch" set and output it to P_i (if more than one such messages are registered, deliver and remove from the "ready to fetch" set the first, chronologically, registered such pair); otherwise send ($\textsc{fetch-output}, P_i$) to $\mathcal{S}$.

Figure 5: The Q-Fairness wrapper functionality.

Is the ledger functionality sufficient for Q fairness? We will construct secure computation protocols based on the ledger functionality $\bar{\mathcal{G}}_{\text{LEDGER}}$ together with other trusted setups. We may wonder if we can construct secure computation protocol from $\bar{\mathcal{G}}_{\text{LEDGER}}$ only. The answer if negative. Indeed, we prove the following statement

Theorem 4.4. *Let $\mathsf{Q}_{\bar{\mathcal{G}}}$ be a predicate with respect to global functionality $\bar{\mathcal{G}} = \bar{\mathcal{G}}_{\text{LEDGER}}$. There exists no protocol in the $\bar{\mathcal{G}}_{\text{LEDGER}}$ hybrid world which realizes the commitment functionality $\mathcal{F}_{\text{COM}}$ with $\mathsf{Q}_{\bar{\mathcal{G}}}$ fairness.*

The proof idea is very similar to the well-known Canetti-Fischlin [CF01] impossibility proof. Please refer to Appendix A.3 for the proof.

4.2 $\mathsf{Q}_{\bar{\mathcal{G}}}$-Robustness

The above wrapper $\mathcal{W}$ allows the simulator to delay delivery of messages arbitrarily. Thus, although the predicates do guarantee the promised notion of fairness, the resulting functionality lacks the other relevant property that we discussed in the introduction, namely robustness. In the following we define Q-robustness which will ensure that if any party starts executing the protocol on its input (i.e., the protocol does not abort due to lack of resources for some party), then every honest party is guaranteed to either receive its output without loosing revenue, or receive bottom and a compensation. This property can be obtained by modifying the wrapper $\mathcal{W}$ using an idea from [KMTZ13] so that in addition to the global-setup-related guarantees induced by predicate Q, it also preserves the guaranteed termination property of the wrapped functionality.[16]

More concretely, in [KMTZ13], a functionality was augmented to have guaranteed termination, by ensuring that given appropriately many activations (i.e., dummy inputs), from its honest interface, it computes its output.[17] In the same spirit, a wrapper which ensures Q-robustness is derived from $\mathcal{W}$ via the following modification: As soon as a fair-output is registered (i.e., upon the wrapper receiving $(\text{FAIR-DELIVER}, \mathsf{sid}, \mathsf{mid}, (m_1, P_{i_1}), \ldots, (m_k, P_{i_k}), (m_{\mathcal{S}}, \mathcal{S}))$ from its inner functionality) it initiates a counter $\lambda = 0$ and an indicator variable $\lambda_{i_j} := 0$ for each $P_{i_j} \in \{P_{i_1}, \ldots, P_{i_k}\}$; whenever a message is received from some $P_{i_j} \in \{P_{i_1}, \ldots, P_{i_k}\}$, the wrapper sets $\lambda_{i_j} := 1$ and does the following check: if $\lambda_{i_j} = 1$ for all $P_{i_j} \in \{P_{i_1}, \ldots, P_{i_k}\}$ then increase $\lambda := \lambda + 1$ and reset $\lambda_{i_j} = 0$ for all $P_{i_j} \in \{P_{i_1}, \ldots, P_{i_k}\}$. As soon as λ reaches a set threshold T, the wrapper simultaneously delivers each $((m_1, P_{i_k}), \ldots, (m_k, P_{i_k})$ (i.e., prepares them to be fetched) without waiting for the simulator and does not accept any inputs other than $(\text{FETCH-OUTPUT}, \cdot)$ from that point on. When this happens, we will say that the wrapper *reached its termination limit.* We denote by $\hat{\mathcal{W}}^T$ the wrapper from Figure 5 modified as described above. Note that the wrapper is parameterized by the termination threshold T.

The intuition why this modification ensures guaranteed termination is the same as in [KMTZ13]: if the environment wishes the experiment to terminate, the it can make it terminate irrespective of the simulator's strategy. Thus a protocol which realizes such a wrapper should also have such a guaranteed termination (the adversary cannot stall the computation indefinitely.)

[16] That is, we want to ensure that if the functionality $\mathcal{F}$ has guaranteed termination then the wrapped functionality will also have guaranteed termination.

[17] Of course, the simulator needs to be given sufficiently many activation so that he can provide its own inputs and perform the simulation (for details we refer the interested reader to [KMTZ13]).

Definition 4.5. *We say protocol π realizes functionality $\mathcal{F}$ with $\mathsf{Q}_{\bar{\mathcal{G}}}$-robustness with respect to global functionality $\bar{\mathcal{G}}$, provided the following statement is true. There exists a threshold T such that for all adversaries $\mathcal{A}$, there is a simulator $\mathcal{S}$ so that for all environments $\mathcal{Z}$ it holds:*

$$\textsc{Exec}^{\bar{\mathcal{G}}}_{\pi,\mathcal{A},\mathcal{Z}} \approx \textsc{Exec}^{\bar{\mathcal{G}},\hat{\mathcal{W}}^{T}_{\mathsf{Q},\bar{\mathcal{G}}}(\mathcal{F})}_{\mathcal{S},\mathcal{Z}}.$$

Moreover, whenever the wrapper reaches its termination limit, then for the state trans *of the global setup $\bar{\mathcal{G}}$ upon termination it holds that* $\mathsf{Q}^{Dlv}_{\bar{\mathcal{G}}}(\mathsf{sid}, P, RS^{\mathsf{pub}}_{P,\mathsf{sid}}, \mathsf{trans})$ *for every party $P \in \mathcal{P}$.*

The composition theorems for Q-fairness from Section 4.1 can be adapted in a straight-forward manner to Q-robustness. The statements and proofs are as in the previous section and are omitted. We note in passing that since the wrapper $\hat{\mathcal{W}}$ is in fact a wrapper which restricts the behavior of $\mathcal{S}$ on top of the restrictions which are applied by the Q-fairness wrapper $\mathcal{W}$, a protocol which is Q-robustness is also Q-fair with respect to the same predicate Q.

4.3 Computation with Fair/Robust Compensation

We are now ready to instantiate the notion of Q-fairness with a compensation mechanism. For the case when $\bar{\mathcal{G}}$ corresponds to a Bitcoin-like ledger, e.g., $\bar{\mathcal{G}} = \bar{\mathcal{G}}_{\text{LEDGER}}$, and $\mathsf{Q}_{\bar{\mathcal{G}}}$ provides compensation of c coins, where $c > 0$, in the case of an abort, the resource-setup generation algorithm Gen a pair of $(\mathtt{address}, \mathsf{sk})$ where $\mathtt{address}$ is a bitcoin address and sk is the corresponding secret-key and the predicate $\mathsf{Q}^{\mathsf{coin}}_{\bar{\mathcal{G}}} = (\mathsf{Q}^{\text{C-Init}}_{\bar{\mathcal{G}}}, \mathsf{Q}^{\text{C-Dlv}}_{\bar{\mathcal{G}}}, \mathsf{Q}^{\text{C-Abt}}_{\bar{\mathcal{G}}})$ operates as follows. On input a session ID sid, a party id P, a wallet address $RS^{\mathsf{pub}}_{P,\mathsf{sid}}$, and a string trans which is parsed as a bitcoin ledger that contains transactions:[18]

- $\mathsf{Q}^{\text{C-Init}}_{\bar{\mathcal{G}}}$ outputs true if and only if the balance of all transactions (both incoming and outgoing) that concern $RS^{\mathsf{pub}}_{P,\mathsf{sid}}$ in trans and carry the meta-data sid is higher than a fixed pre-agreed initialization amount.[19]
- $\mathsf{Q}^{\text{C-Dlv}}_{\bar{\mathcal{G}}}$ outputs true if and only if the balance of all transactions (both incoming and outgoing) that concern $RS^{\mathsf{pub}}_{P,\mathsf{sid}}$ in trans and carry the meta-data sid is greater or equal to 0.
- $\mathsf{Q}^{\text{C-Abt}}_{\bar{\mathcal{G}}}$ outputs true if and only if the balance of all transactions (both incoming and outgoing) that concern $RS^{\mathsf{pub}}_{P,\mathsf{sid}}$ in trans and carry the meta-data sid is greater or equal to a fixed pre-agreed compensation amount.

If a protocol π realizes a functionality $\mathcal{F}$ with $\mathsf{Q}^{\mathsf{coin}}_{\bar{\mathcal{G}}}$-fairness (resp. $\mathsf{Q}^{\mathsf{coin}}_{\bar{\mathcal{G}}}$-robustness), i.e., with respect to the global functionality $\bar{\mathcal{G}}_{\text{LEDGER}}$, we say that π realizes $\mathcal{F}$ with fair compensation (resp. with robust compensation). Because our results are proved for $\mathsf{Q}^{\mathsf{coin}}_{\bar{\mathcal{G}}}$, to keep the notation simple in the remainder of the paper we might drop the superscript from $\mathsf{Q}^{\mathsf{coin}}_{\bar{\mathcal{G}}}$, i.e., we write Q or $\mathsf{Q}_{\bar{\mathcal{G}}}$ instead of $\mathsf{Q}^{\mathsf{coin}}_{\bar{\mathcal{G}}}$.

[18] Transactions in trans can also be marked with metadata.

[19] In our construction $\mathsf{Q}^{\text{C-Init}}_{\bar{\mathcal{G}}}$ will check additional properties for the initial set of transactions that concern $RS^{\mathsf{pub}}_{P,\mathsf{sid}}$; specifically, not only that a fixed amount μ is present but also that it is distributed in a special way.

5 Our $\mathsf{Q}^{\mathsf{coin}}_{\bar{\mathcal{G}}}$-Robust Protocol Compiler

In this section we present our fair and robust protocol compiler. Our compiler compiles a synchronous protocol $\pi_{\mathtt{SH}}$ which is secure (i.e., private) against a corrupted majority in the semi-honest correlated randomness model (e.g, an OT-hybrid protocol where the OT's have been pre-computed) into a protocol π which is secure with fair-compensation in the malicious correlated randomness model. The high-level idea is the following: We first compile $\pi_{\mathtt{SH}}$ into a protocol in the malicious correlated randomness model, which is executed over a broadcast channel and is secure with publicly identifiable abort. (Roughly, this means that someone observing the protocol execution can decide, upon abort, which party is not executing its code.) This protocol is then transformed into a protocol with fair compensation as follows: Every party (after receiving his correlated randomness setup) posts to the ledger transactions that the other parties can claim only if they, later, post transactions that prove that they follow their protocol. Transactions that are not claimed this way are returned to the source address; thus, if some party does not post such a proof it will not be able to claim the corresponding transaction, and will therefore leave the honest parties with a positive balance as their transactions will be refunded. Observe that these are not standard Bitcoin transactions, but they have a special format which is described in the following.

Importantly, the protocol we describe is guaranteed to either produce output in as many (Bitcoin) rounds as the rounds of the original malicious protocol, or to compensate all honest parties. This *robustness* property is achieved by a novel technique which ensures that as soon as the honest parties make their initial transaction, the adversary has no way of preventing them from either computing their output or being compensated. Informally, our technique consists of splitting the parties into "islands" depending on the transactions they post (so that all honest parties are on the same island) and then allowing them to either compute the function within their island, or if they abort to get compensated. (The adversary has the option of being included or not in the honest parties' island.)

5.1 MPC with Publicly Identifiable Abort

As a first step in our compiler we invoke the semi-honest to malicious with identifiable abort compiler of Ishai, Ostrovsky, and Zikas [IOZ14] (hereafter referred to as the *IOZ compiler*). This compiler takes a semi-honest protocol $\pi_{\mathtt{SH}}$ in the correlated randomness model and transforms it to a protocol in the malicious correlated randomness model (for an appropriate setup) which is secure with identifiable abort, i.e., when it aborts, every party learns the identity of a corrupted party. The compiler in [IOZ14] follows the so called GMW paradigm [GMW87], which in a nutshell has every party commit to its input and randomness for executing the semi-honest protocol $\pi_{\mathtt{SH}}$ and then has every party run $\pi_{\mathtt{SH}}$ over a broadcast channel, where in each round ρ every party broadcasts his round ρ messages and proves in zero-knowledge that the broadcasted message is correct, i.e., that he knows the input and randomness that are consistent with the initial commitments and the (public) view of the protocol so far. The main difference of the IOZ compiler and the GMW compiler is that the parties are not only committed to their randomness, but they are also committed to their entire setup string, i.e., their private component of the correlated randomness. In the following, for the sake of completeness, we enumerate some key properties of the resulting maliciously secure protocol $\pi_{\mathtt{Mal}}$ (which is based on the compiler in [IOZ14]) that will be important for our construction:

- Every party is committed to his setup, i.e., the part of the correlated randomness it holds. That is, every party P_i receives from the setup his randomness (which we refer to as P_i's *private* *component* of the setup) along with one-to-many commitments[20] on the private components of all parties. Without loss of generality, we also assume that a common-reference string (CRS) and a public-key infrastructure (PKI) are included in every party's setup. We refer to the distribution of this correlated randomness as $\mathcal{D}_{\mathtt{Mal}}$.

- The protocol $\pi_{\mathtt{Mal}}$ uses *only* the broadcast channel for communication.
- Given the correlated randomness setup, the protocol $\pi_{\mathtt{Mal}}$ is completely deterministic. This is achieved in [IOZ14] by ensuring that all the randomness used in the protocol, even the one needed for the zero-knowledge proofs, is part of the private components that are distributed by the sampling functionality.[21]
- $\pi_{\mathtt{Mal}}$ starts off by having every party broadcast a one-time pad encryption of its input with its (committed) randomness and a NIZK that it knows the input and randomness corresponding to the broadcasted message.
- By convention, the next-message function of $\pi_{\mathtt{Mal}}$ is such that if in any round the transcript seen by a party is an aborting transcript, i.e., is not consistent with an accepting run of the semi-honest protocol, then the party outputs $\perp$. Recall that the identifiable abort property ensures that in this case every party will also output the identity of a malicious party (the same for all parties).
- There is a (known) upper bound on the number ρ_{c} of rounds of $\pi_{\mathtt{Mal}}$.

We remark that, given appropriate setup, the IOZ-compiler achieves information-theoretic security, and needs therefore to build information-theoretic commitments and zero-knowledge proofs. As in this work we are only after computational security, we modify the IOZ compiler so that we use (computationally) UC secure one-to-many commitments [CLOS02] and computationally UC secure non-interactive zero-knowledge proofs (NIZKs) instead if their information-theoretic instantiation suggested in [IOZ14]. Both the UC commitment and the NIZKs can be built in the CRS model. Moreover, the use of UC secure instantiations of zero-knowledge and commitments ensures that the resulting protocol will be (computationally) secure.

Using the setup within a subset of parties. A standard property of many protocols in the correlated-randomness model is that once the parties in $\mathcal{P}$ have received the setup, any subset $\mathcal{P}' \subset \mathcal{P}$ is able to use it to perform a computation of a $|\mathcal{P}'|$-party function amongst them while ignoring parties in $\mathcal{P} \setminus \mathcal{P}'$. More concretely, assume the parties in $\mathcal{P}$ have been handed a setup allowing them to execute some protocol π for computing any $|\mathcal{P}|$-party function f; then for any $\mathcal{P}' \subseteq \mathcal{P}$, the parties in $\mathcal{P}'$ can use their setup within a protocol $\pi|_{\mathcal{P}'}$ to compute any $|\mathcal{P}'|$-party function $f|_{|\mathcal{P}'|}$. This property which will prove very useful for obtaining computation with robustness or compensation, is also satisfied by the IOZ protocol, as the parties in $\mathcal{P}'$ can simply ignore the commitments (public setup component) corresponding to parties in $\mathcal{P} \setminus \mathcal{P}'$. It should be noted that this is not an inherent property of the correlated randomness model: e.g., protocols based on threshold encryption do not immediately satisfy this property (as players would have to readjust the threshold).

Making Identifiability Public. The general idea of our protocol is to have every party issue transactions by which he commits to transferring a certain amount of coins per party for each protocol round. All these transactions are issued at the beginning of the protocol execution. Every

[20] These are commitments that can be opened so that every party agrees on whether or not the opening succeeded.

[21] As an example, the challenge for the zero-knowledge proofs is generated by the parties opening appropriate parts of their committed random strings.

party can claim the "committed" coins transferred to him associated to some protocol round ρ only under the following conditions: (1) the claim is posted in the time-interval corresponding to round ρ; (2) the party has claimed all his transferred coins associated to the previous rounds; and (3) the party has posted a transaction which includes his valid protocol message for round ρ.

In order to ensure that a party cannot claim his coins unless he follows the protocol, the ledger (more concretely the validation predicate) should be able to check that the party is indeed posting its valid next message. In other words, in each round ρ, P_i's round-ρ message acts as a witness for P_i claiming all the coins committed to him associated with this round ρ. To this direction we make the following modification to the protocol: Let $f(x_1, \ldots, x_n) = (y_1, \ldots, y_n)$ denote the n-party function we wish to compute, and let f^{+1} be the $(n+1)$-party function which takes input x_i from each P_i, $i \in [n]$, and no input from P_{n+1} and outputs y_i to each P_i and a special symbol (e.g., 0) to P_{n+1}. Clearly, if $\pi_{\mathtt{SH}}$ is a semi-honest n-party protocol for computing f over broadcast, then the $n+1$ protocol $\pi_{\mathtt{SH}}^{+1}$ (in which every P_i with $i \in [n]$ executes $\pi_{\mathtt{SH}}$ and P_{n+1} simply listens to the broadcast channel and outputs 0) is a semi-honest secure protocol for f^{+1}.

Now if $\pi_{\mathtt{Mal}}^{+1}$ denotes the $(n+1)$-party malicious protocol which results by applying the above modified IOZ compiler on the $(n+1)$-party semi-honest protocol $\pi_{\mathtt{SH}}^{+1}$ for computing the function f^{+1}, then, by construction this protocol computes function f^{+1} with identifiable abort and has the following additional properties:

- Party P_{n+1} does not make any use of his private randomness whatsoever; this is true because he broadcasts no messages and simply verifies the broadcasted NIZKs.
- If some party P_i, $i \in [n]$ deviates from running $\pi_{\mathtt{SH}}$ with the correlated (committed) randomness as distributed from the sampling functionality, then this is detected by all parties, including P_{n+1} (and protocol $\pi_{\mathtt{Mal}}^{+1}$ aborts identifying P_i as the offender). This follows by the soundness of the NIZK which P_i needs to provide proving that he is executing $\pi_{\mathtt{SH}}$ in every round.

Due to P_{n+1}'s role as an observer who gets to decide if the protocol is successful (P_{n+1} outputs 0) or some party deviated (P_{n+1} observes that the corresponding NIZK verification failed) in the following we will refer to P_{n+1} in the above protocol as the *judge*. The code of the judge can be used by anyone who has the public setup and wants to follow the protocol execution and decide whether it should abort or not given the parties' messages. Looking ahead, the judge's code in the protocol will be used by the ledger to decide wether or not a transaction that claims some committed coins is valid.

5.2 Special Transactions supported by our Ledger

In this section we specify the Validate and the Blockify predicates that are used for achieving our protocol's properties. More specifically, our protocol uses the following type of transactions which transfer v coins from wallet $\mathtt{address}_i$ to wallet $\mathtt{address}_j$ conditioned on a statement Σ:

$$\mathbb{B}_{v,\mathtt{address}_i,\mathtt{address}_j,\Sigma,\mathtt{aux},\sigma_i,\tau} \tag{1}$$

where σ_i is a signature of the transaction, which can be verified under wallet $\mathtt{address}_i$; τ is the time-stamp, i.e., the current value of the clock when this transaction is added to the state by the ledger—note that this timestamp is added by the ledger and not by the users,—$\mathtt{aux} \in \{0,1\}^*$ is an arbitrary string[22]; and the statement Σ consists of three arguments, i.e., $\Sigma = (\mathtt{arg1}, \mathtt{arg2}, \mathtt{arg3})$,

[22] This string will be included to the Ledger's state as soon as the transaction is posted and can be, therefore, referred to by other spending statements.

which are processed by the $\mathsf{Validate}$ predicate in order to decide if the transaction is valid (i.e., if it will be included in the ledger's next block).

The $\mathsf{Validate}$ predicate. The validation happens by processing the arguments of Σ in a sequential order, where if while processing of some argument the validation rejects, algorithm $\mathsf{Validate}$ stops processing at that point and this transaction is dropped. The arguments are defined/processed as follows:

Time-Restrictions: The first argument is a pair $\texttt{arg1} = (\tau_-, \tau_+) \in \mathbb{Z} \times (\mathbb{Z}^+ \cup \{\infty\})$ of points in time. If $\tau_- > \tau_+$ then the transaction is invalid (i.e., it will be dropped by the ledger). Otherwise, before time τ_- the coins in the transaction "remain" blocked, i.e., no party can spend them; from time τ_- until time τ_+, the money can be spent by the owner of wallet $\texttt{address}_j$ provided that the spending statement satisfies also the rest of the requirements/arguments in statement Σ (listed below). After time τ_+ the money can be spent by the owner of wallet $\texttt{address}_i$ without any additional restrictions (i.e., the rest of the arguments in Σ are not parsed). As a special case, if $\tau_+ = \infty$ then the transferred coins can be spent from $\texttt{address}_j$ at any point (provided the spending statement is satisfied); we say then that the transaction is *time-unrestricted*,[23] otherwise we say that the transaction is *time restricted.*

Spending Link: Provided that the processing of the first argument, as above, was not rejecting, the $\mathsf{Validate}$ predicate proceeds to the second argument, which is a unique "anchor", $\texttt{arg2} = \alpha \in \{0,1\}^*$. Informally, this serves as a unique identifier for linked transactions[24]; that is, when $\alpha \neq \perp$, then the $\mathsf{Validate}$ algorithm of the ledger looks in the ledger's state and buffer to confirm that the balance of transactions to/from the wallet address $\texttt{address}_i$ with this anchor $\texttt{arg2}$ is at least $v' \geq v$ coins. That is, the sum of coins in the state or in the buffer with receiver address $\texttt{address}_i$ and anchor $\texttt{arg2}$ minus the sum of coins in the state or in the buffer with sender address $\texttt{address}_i$ and anchor $\texttt{arg2}$ is greater equal to v. If this is not the case then the transaction is rendered invalid; otherwise the validation of this argument succeeds and the algorithm proceeds to the next argument.

State-Dependent Condition: The last argument to be validated is $\texttt{arg3}$, which is a relation $\mathcal{R} : \mathcal{S} \times \mathcal{B} \times \mathcal{T} \to \{0,1\}$, where $\mathcal{S}$, $\mathcal{B}$, and $\mathcal{T}$ are the domains of possible ledger-states, ledger-buffers, and transactions, respectively (in a given encoding). This argument defines which type of transactions can spend the coins transferred in the current transaction. That is, in order to spend the coins, the receiver needs to submit a transaction $\texttt{tx} \in \mathcal{T}$ such that $\mathcal{R}(\mathsf{state}, \mathsf{buffer}, \texttt{tx}) = 1$ at the moment when $\texttt{tx}$ is to be validated and inserted in the buffer. In our construction this is the part of the transaction that we will take advantage to detect cheating (and thus $\mathcal{R}$ will encode a NIZK verifier etc.).

We point out that as with standard Bitcoin transactions, the validation predicate will always also check validity of the signature σ_i with respect to the wallet $\texttt{address}_i$. Moreover, the standard Bitcoin-like transactions can be trivially casted as transactions of the above type by setting $\alpha = \perp$ and $\Sigma = ((0,\infty), \perp, \mathcal{R}_\emptyset)$, where $\mathcal{R}_\emptyset$ denotes the relation which is always true.

To simplify the structure of our special transactions and ease their implementation, we impose the following additional constraints: whenever a time-restriction is given, i.e., $\text{arg}_1 = (\tau_-, \tau_+)$ then

[23] This is the case with standard Bitcoin transactions.

[24] Looking ahead $\texttt{arg2}$ will be used to point to specific transactions of a protocol instance. The mechanism may be simulated by generating multiple addresses however it is more convenient for the protocol description and for this reason we adopt it.

it must be that $\alpha \neq \bot$. Furthermore, if a time-restricted transaction is present with anchor α from $\mathtt{address}_1$ to $\mathtt{address}_2$, the only transactions that are permitted with anchor α in the ledger would be time-unrestricted transactions originating from either $\mathtt{address}_2$ within the specified time-window, or $\mathtt{address}_1$ after the specified time window.

The Blockify *algorithm.* This algorithm simply groups transactions in the current buffer and adds a timestamp from the current round. We choose to ignore any additional functionality (e.g., such as a reward mechanism for mining that is present in typical cryptocurrencies — however such mechanism can be easily added independently of our results).

5.3 The Protocol

Let $\pi_{\mathtt{Mal}}^{+1}$ denote the protocol described in section 5.1. Let $\mathtt{Round2Time}(1)$ denote the time in which the parties have agreed to start the protocol execution. Without loss of generality we assume that $\mathtt{Round2Time}(1) > \mathtt{T}+1$ where $\mathtt{T}$ is the number clock ticks for each block generation cf. Figure 4.[25] Furthermore, for simplicity, we assume that each party P_i receives its input x_i with its first activation from the environment at time $\mathtt{Round2Time}(1)$ (if some honest party does not have an input by that time it will execute the protocol with a default input, e.g., 0).

Informally, the protocol proceeds as follows: In a pre-processing step, before the parties receive input, the parties invoke the sampling functionality for $\pi_{\mathtt{Mal}}^{+1}$ to receive their correlated randomness.[26] The public component of this randomness includes their protocol-associated wallet $\mathtt{address}_i$ which they output (to the environment). This corresponds to the resources allocation step in the Q-robustness wrapper $\hat{\mathcal{W}}$. The environment is then expected to submit ρ_{c} special (as above) transactions for each pair of parties $P_i \in \mathcal{P}$ and $P_j \in \mathcal{P}$; the source wallet-address for each such transaction is P_i's, i.e., $\mathtt{address}_i$ and the target wallet-address for is P_j's, i.e., $\mathtt{address}_j$, and the corresponding anchors are as follows: $\alpha_{i,j,\rho} = (\mathsf{pid}, i, j, \rho)$, for $(i,j,\rho) \in [n]^2 \times [\rho_{\mathrm{c}}]$, where[27] pid is the (G)UC protocol ID for $\pi_{\mathtt{Mal}}^{+1}$. Since by assumption, $\mathtt{Round2Time}(1) > \mathtt{T}+1$, the environment has sufficient time to submit these transaction so that by the time the protocol starts they have been posted on the ledger.

At time $\mathtt{Round2Time}(1)$ the parties receive their inputs and initiate the protocol execution by first checking that sufficient funds are allocated to their wallets linked to the protocol executions by appropriate anchors, as above. If some party does not have sufficient funds then it broadcasts an aborting message and all parties abort.[28] This aborting in case of insufficient funds is consistent with the behavior of the wrapper $\hat{\mathcal{W}}$ when $\mathsf{Q}_{\mathcal{G}}^{\text{C-Init}}$ is false. Otherwise, parties make the special transactions that commit them (see below) into executing the protocol, and then proceed into claiming them one-by-one by executing their protocol in a round-by-round fashion.

Note that each protocol round lasts one ledger round so that the parties have enough time to claim their transactions. This means that $\mathtt{Round2Time}(i+1) - \mathtt{Round2Time}(i) \geq \mathtt{T}$, which guarantees that any transaction submitted for round ρ, $\rho = 1, \ldots, \rho_{\mathrm{c}} - 1$, of the protocol, has been posted on the ledger by the beginning of round $\rho + 1$. Observe that by using a constant round protocol $\pi_{\mathtt{Mal}}^{+1}$

[25] That is we assume that at least one ledger rounds plus one extra clock-ticks have passed from the beginning of the time.

[26] In an actual application, the parties will use an unfair protocol for computing the correlated randomness. As this protocol has no inputs, an abort will not be unfair (i.e., the simulator can always simulate the view of the adversary in an aborting execution.)

[27] Recall that we assume $|\mathcal{P}| = n$.

[28] Note that this is a fair abort and no party has spent any time into making transactions.

(e.g., the modified compiled protocol from [IOZ14] instantiated with a constant round semi-honest protocol) we can ensure that our protocol will terminate in a constant number of ledger rounds and every honest party will either receive its input, or will have a positive balance in its wallet.

Remark 5.1 (On availability of funds)**.** *Unlike existing works, we choose to explicitly treat the issue of how funds become available to the protocol by making the off-line transfers external to the protocol itself (i.e., the environment takes care of them). However, the fact that the environment is in charge of "pouring" money into the wallets that are used for the protocol does not exclude that the parties might be actually the ones having done so. Indeed, the environment's goal is to capture everything that is done on the side of, before, or after the protocol, including other protocols that the parties might have participated in. By giving the environment enough time to ensure these transactions are posted we ensure that some honest party not having enough funds corresponds to an environment that makes the computation abort (in a fair way and only in the pre-processing phase, before the parties have invested time into posting protocol transactions).*

Here is how we exploit the power of our special transactions in order to arrange that the balance of honest parties is positive in case of an abort. We require that the auxiliary string of a transaction of a party P_j which claims a committed transaction for some round ρ includes his ρ-round protocol message. We then have the relation of this transaction be such that it evaluates to 1 if only if this is indeed P_j's next message. Thus, effectively the validate predicate implements the judge in $\pi^{+1}_{\mathtt{Mal}}$ and can, therefore, decide if some party aborted: if some party broadcasts a message that would make the judge abort, then the validate predicate drops the corresponding transaction and all claims for committed transactions corresponding to future rounds, thus, all other parties are allowed to reclaim their committed coins starting from the next round.

Before we give the protocol description there is a last question: how is the ledger able to know which parties should participate in the protocol? Here is the problem: The adversary might post in the first round (as part of the committing transaction for the first round) a fake, maliciously generated setup. Since the ledger is not part of the correlated randomness sampling, it would be impossible to decide which is the good setup. We solve this issue by the following technique that is inspired by [BCL$^+$05]: The ledger[29] groups together parties that post the same setup; these parties form "islands", i.e, subsets of $\mathcal{P}$. For each such subset $\mathcal{P}' \subseteq \mathcal{P} \cup \{P_{n+1}\}$ which includes the judge P_{n+1}, the ledger acts as if the parties in $\mathcal{P}'$ are executing the protocol $\pi^{+1}_{\mathtt{Mal}}|_{\mathcal{P}'}$ (which, recall, is the restriction of $\pi^{+1}_{\mathtt{Mal}}$ to the parties in $\mathcal{P}'$) for computing the $|\mathcal{P}'|$-party function $f^{+1}|_{\mathcal{P}'}(\vec{x})$ defined as follows: let the function to be computed be $f(\vec{x})$, where $\vec{x} = (x_1, \ldots, x_n)$, and f^{+1} be as above, then $f^{+1}|_{\mathcal{P}'}(\vec{x}) = f^{+1}(\vec{x}_{\mathcal{P}'})$ where $\vec{x}_{\mathcal{P}'} = (x'_1, \ldots, x'_n)$ with $x'_i = x_i$ for $P_i \in \mathcal{P}'$ and x'_i being a default value for every $P_i \notin \mathcal{P}'$. This solves the problem as all honest parties will be in the same island $\mathcal{P}' \subset \mathcal{P}$ (as they will all post the same value for public randomness); thus if the adversary chooses not to post this value on behalf of some corrupted party, he is effectively setting this party's input to a default value, a strategy which is easily simulatable. (Of course, the above solution will allow the adversary to also have "islands" of only corrupted parties that might execute the protocol, but this is also a fully simulatable strategy and has no effect on fair-compensation whatsoever—corrupted parties are not required to have a positive balance upon abort).

The final protocol $\pi^{\mathbb{B}}_{\mathtt{Mal}}$ is detailed in the following. The protocol ID is sid. The function to be computed is $f(x_1, \ldots, x_n)$. The protocol parties are $\mathcal{P} = \{P_1, \ldots, P_n\}$. We assume all parties have

[29] Throughout the following description, we say that the ledger does some check to refer to the process of checking a corresponding relation, as part of validating a special transaction.

registered with the clock functionality in advance and are therefore synchronized once the following steps start.

Phase 1: Setup Generation

Time $\tau_{-2} = \texttt{Round2Time}(1) - \texttt{T} - 2$:

The parties invoke the sampling functionality (cf. Figure 1) for $\mathcal{D}_{\texttt{Mal}}$, i.e., every party $P_i \in \mathcal{P}$ starts off by sending the sampling functionality a message $(\textsc{request}, \mathsf{sid})$; the sampling functionality returns $(R_i^{\texttt{priv}}, R^{\texttt{pub}})$ to P_i where $R_i^{\texttt{priv}}$ is P_i's private component (including all random coins he needs to run the protocol, along with his signing key sk_i) of the setup and $R^{\texttt{pub}}$ is the public component (the same for every party P_j) which includes the vector of UC commitments $(\mathsf{Com}_1, \ldots, \mathsf{Com}_n)$, where for $j \in [n]$, Com_j is a commitment to $R_j^{\texttt{priv}}$, along with a vector of public (verification) keys $(\mathsf{vk}_1, \ldots, \mathsf{vk}_n)$ corresponding to the signing keys $(\mathsf{sk}_1, \ldots, \mathsf{sk}_n)$ and a common reference string CRS. Every party outputs its own public key, as its wallet address for the protocol, i.e., $\texttt{address}_i = \mathsf{vk}_i$.

Phase 2: Inputs and Protocol Execution

Time $\tau_{-1} = \texttt{Round2Time}(1) - 1$:

Every party $P_i \in \mathcal{P}$ receives its input x_i ($x_i = 0$ if no input is received in the first activation of P_i for time $\texttt{Round2Time}(1)$) and does the following to check that it has sufficient fund available: P_i reads the current state from the ledger. If the state does not include for each $(i, j, \rho) \in [n]^2 \times [\rho_c]$ a transaction $\mathbb{B}_{c,\texttt{address},\texttt{address}_i,\Sigma^0_{i,j,\rho},\texttt{aux}^0_{i,j,\rho},\sigma,\tau}$, for some arbitrary $\texttt{address}$ and where $\Sigma^0_{i,j,\rho} = ((0, \infty), (\mathsf{sid}, i, j, \rho), \mathcal{R}_\emptyset)$ then P_i broadcasts $\perp$ and every party aborts the protocol execution with output $\perp$ (i.e., no party does anything from that point on.Recall that ρ_c is the upper bound on the number of rounds of $\pi^{+1}_{\texttt{Mal}}$, cf. Section 5.1.

Time $\tau_0 = \texttt{Round2Time}(1)$:

Every P_i submits to the ledger the following "commitment" transactions:[30]

1. For each $P_j \in \mathcal{P} : \mathbb{B}_{c,\texttt{address}_i,\texttt{address}_j,\Sigma_{i,j,1},\texttt{aux}_{i,j,1},\sigma,\tau}$, where $\texttt{aux}_{i,j,1} = R^{\texttt{pub}}$ and $\Sigma_{i,j,1} = (\texttt{arg1}_{i,j,1}, \texttt{arg2}_{i,j,1}, \texttt{arg3}_{i,j,1})$ with
   - $\texttt{arg1}_{i,j,1} = (\texttt{Round2Time}(1) + \texttt{T}, \texttt{Round2Time}(1) + 2\texttt{T} - 1)$
   - $\texttt{arg2}_{i,j,1} = (\mathsf{sid}, i, j, 1)$
   - $\texttt{arg3}_{i,j,1} = \mathcal{R}_{i,j,1}$ defined as follows: Let $\mathcal{P}^{+1} = \mathcal{P} \cup \{P_{n+1}\}$, where P_{n+1} denotes the judge, be the player set implicit in $R^{\texttt{pub}}$, [31] and let $\mathcal{P}_i^{+1} \subseteq \mathcal{P}^{+1}$ denote the island of party i including the judge, i.e., the set of parties (wallets), such that in the first block posted after time $\texttt{Round2Time}(1)$ all parties $P_k \in \mathcal{P}_i^{+1}$ have exactly one transaction for every $P_j \in \mathcal{P}$ with $\texttt{arg1}_{k,j,1} = (\texttt{Round2Time}(1) + \texttt{T}, \texttt{Round2Time}(1) + 2\texttt{T} - 1)$, $\texttt{arg2}_{k,j,1} = (\mathsf{sid}, k, j, 1)$, and $\texttt{aux}^1_{k,j,1} = R^{\texttt{pub}}$. Furthermore, let $\pi^{+1}_{\texttt{Mal}}|_{\mathcal{P}_i^{+1}}$ be the protocol with public identifiability for computing $f^{+1}|_{\mathcal{P}_i^{+1}}$, described above and

[30] Recall that, by definition of the clock, every party has as much time as it needs to complete all the steps below before the clock advances time.

[31] Recall that $R^{\texttt{pub}}$ includes commitments to all parties' private randomness (including the judge's P_d) used for running the protocol, which is an implicit representation of the player set.

denote by $R^{\texttt{pub}}|_{\mathcal{P}_i^{+1}}$ the restriction of the public setup to the parties in $\mathcal{P}_i^{+1}$. Then $\mathcal{R}_{i,j,1}(\mathsf{state}, \mathsf{buffer}, \texttt{tx}) = 1$ if and only if the protocol of the judge with public setup $R^{\texttt{pub}}|_{\mathcal{P}_i^{+1}}$ accepts the auxiliary string $\texttt{aux}_{\texttt{tx}}$ in $\texttt{tx}$ as P_i's first message in $\pi^{+1}_{\texttt{Mal}}|_{\mathcal{P}_i^{+1}}$ (i.e., it does not abort in the first round).

2. For each protocol round $\rho = 2, \ldots, \rho_{\mathrm{c}}$ and each $P_j \in \mathcal{P}$: each party posts the transaction: $\mathbb{B}_{c,\texttt{address}_i,\texttt{address}_j,\Sigma_{i,j,\rho},\texttt{aux}^1_{i,j,\rho},\sigma,\tau}$, where $\texttt{aux}^1_{i,j,\rho} = R^{\texttt{pub}}$ and $\Sigma_{i,j,\rho} = (\texttt{arg1}, \texttt{arg2}, \texttt{arg3})$ with
   - $\texttt{arg1} = (\texttt{Round2Time}(\rho) + \texttt{T}, \texttt{Round2Time}(\rho + 1) + 2\texttt{T} - 1)$
   - $\texttt{arg2} = (\mathsf{sid}, i, j, \rho)$.
   - $\texttt{arg3} = \mathcal{R}_{i,j,\rho}$ defined as follows: Let $\mathcal{P}_i^{+1}, \pi^{+1}_{\texttt{Mal}}|_{\mathcal{P}_i^{+1}}$ be defined as above (and assume $\mathcal{P}_i^{+1} = \{P_{i_1}, \ldots, P_{i_m}\}$. Then $\mathcal{R}_{i,j,\rho}(\mathsf{state}, \mathsf{buffer}, \texttt{tx}) = 1$ if and only if, for each $r = 1, \ldots, \rho - 1$ and each party $P_{i_k} \in \mathcal{P}_i^{+1}$, the state state includes transactions in which the auxiliary input is $\texttt{aux}_{i_k,r}$ and the protocol of the judge with public setup $R^{\texttt{pub}}|_{\mathcal{P}_i^{+1}}$, and transcript $(\texttt{aux}_{i_1,1}, \ldots, \texttt{aux}_{i_m,1}), \ldots, (\texttt{aux}_{i_1,\rho-1}, \ldots, \texttt{aux}_{i_m,\rho-1})$, accepts the auxiliary string $\texttt{aux}$ in $\texttt{tx}$ as P_i's next (ρ-round) message in $\pi^{+1}_{\texttt{Mal}}|_{\mathcal{P}_i^{+1}}$ (i.e., it does not abort in the ρ-th round).

Phase 3: Claiming Committed Transactions/Executing the Protocol

Time $\tau \geq \texttt{Round2Time}(1)$:

For each $\rho = 1, \ldots, \rho_{\mathrm{c}} + 1$, every P_i does the following at time $\texttt{Round2Time}(\rho)$,:

1. If $\tau = \texttt{Round2Time}(\rho_{\mathrm{c}} + 1)$ then go to Step 4; otherwise do the following:
2. Read the ledger's state, and compute $\mathcal{P}_i^{+1}, \pi^{+1}_{\texttt{Mal}}|_{\mathcal{P}_i^{+1}}$ as above.
3. If the state state is not aborting for $\mathcal{P}_i^{+1} = \{P_{i_1}, \ldots, P_{i_m}\}$, i.e., it includes for each $r = 1, \ldots, \rho - 1$ and each party $P_{i_k} \in \mathcal{P}_i^{+1}$ a transaction in which the auxiliary input is $\texttt{aux}_{i_k,r}$ such that P_i executing $\pi^{+1}_{\texttt{Mal}}|_{\mathcal{P}_i^{+1}}$ with public setup $R^{\texttt{pub}}|_{\mathcal{P}_i^{+1}}$, private setup $R_i^{\texttt{priv}}$, and transcript $(\texttt{aux}_{i_1,1}, \ldots, \texttt{aux}_{i_m,1}), \ldots, (\texttt{aux}_{i_1,\rho-1}, \ldots, \texttt{aux}_{i_m,\rho-1})$ for the first $r - 1$ rounds does not abort, then compute P_i's message for round ρ, denoted as msg_ρ, and submit to the ledger for each $P_k \in \mathcal{P}_i^{+1}$ a transaction $\mathbb{B}_{c,\texttt{address}_i,\texttt{address},\Sigma'_{k,i,\rho},\texttt{aux}^{\rho}_{k,i,\rho},\sigma,\tau}$, where $\texttt{aux}^{\rho}_{k,i,\rho} = \mathrm{msg}_\rho$, $\texttt{address}$ is the address that was the input of the first transaction with link $(\mathsf{sid}, i, k, \rho)$ and $\Sigma'_{k,i,\rho} = (\texttt{arg1}, \texttt{arg2}, \texttt{arg3})$ with
   - $\texttt{arg1} = (0, \infty)$
   - $\texttt{arg2} = (\mathsf{sid}, k, i, \rho)$
   - $\texttt{arg3} = \mathcal{R}_{\emptyset}$. For each such transaction posted enter $(\mathsf{sid}, k, i, \rho)$ in a set of "claimed" transactions CLAIM_i.
4. Otherwise, i.e., if the state state is aborting, then prepare for each round $r = 1, \ldots, \rho - 1$, and each $P_k \in \mathcal{P}$ a transaction by which the committed transaction towards P_k corresponding to round r is claimed back to $\texttt{address}_i$, i.e., $\mathbb{B}_{c,\texttt{address}_k,\texttt{address}_i,\Sigma,\texttt{aux},\sigma,\tau}$, where $\texttt{aux} =\perp$ and $\Sigma = (\texttt{arg1}, \texttt{arg2}, \texttt{arg3})$ with
   - $\texttt{arg1} = (0, \infty)$
   - $\texttt{arg2} = (\mathsf{sid}, i, k, r)$

- $\mathtt{arg3} = \mathcal{R}_{\emptyset}$.

  The above transaction is posted as long as it is not claimed already, i.e., $(\mathsf{sid}, i, k, r) \in \mathrm{CLAIM}_i$ in a previous step.

This completes the description of the protocol. The protocol terminates in $O(\rho_{\mathrm{c}})$ ledger rounds. A depiction of the transactions that are associated with a protocol round is given in Figure 6

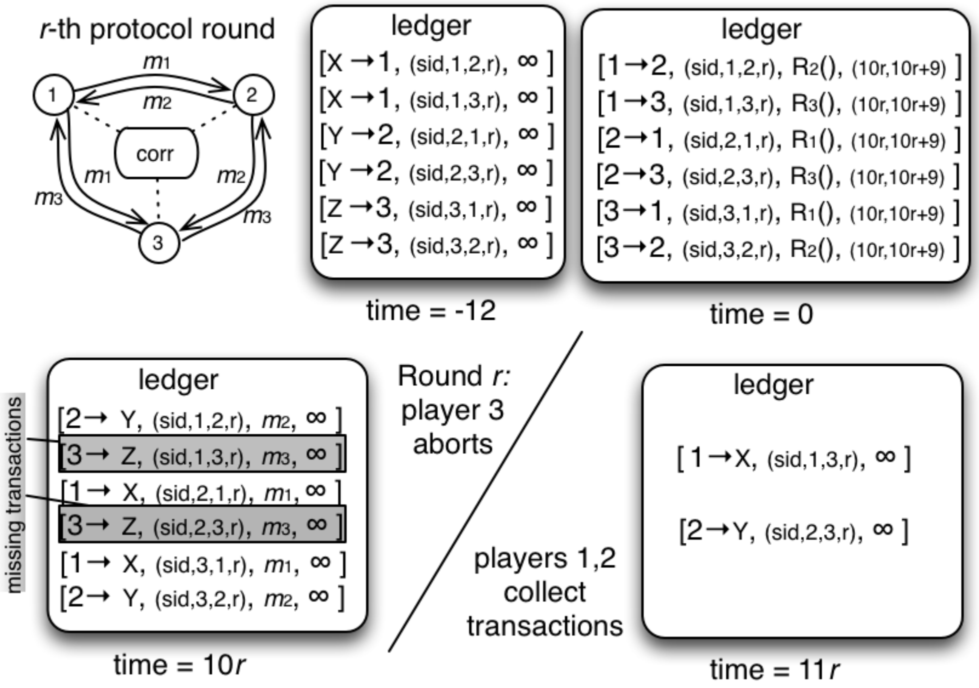
r-th protocol round
m1
m2
1
2
corr
m1
m2
m3
m3
3
ledger
[X →1, (sid,1,2,r), ∞]
[X →1, (sid,1,3,r), ∞]
[Y →2, (sid,2,1,r), ∞]
[Y →2, (sid,2,3,r), ∞]
[Z →3, (sid,3,1,r), ∞]
[Z →3, (sid,3,2,r), ∞]
time = -12
ledger
[1→2, (sid,1,2,r), R2(), (10r,10r+9)]
[1→3, (sid,1,3,r), R3(), (10r,10r+9)]
[2→1, (sid,2,1,r), R1(), (10r,10r+9)]
[2→3, (sid,2,3,r), R3(), (10r,10r+9)]
[3→1, (sid,3,1,r), R1(), (10r,10r+9)]
[3→2, (sid,3,2,r), R2(), (10r,10r+9)]
time = 0
missing transactions
ledger
[2→ Y, (sid,1,2,r), m2, ∞]
[3→ Z, (sid,1,3,r), m3, ∞]
[1→ X, (sid,2,1,r), m1, ∞]
[3→ Z, (sid,2,3,r), m3, ∞]
[1→ X, (sid,3,1,r), m1, ∞]
[2→ Y, (sid,3,2,r), m2, ∞]
time = 10r
Round r: player 3 aborts
players 1,2 collect transactions
ledger
[1→X, (sid,1,3,r), ∞]
[2→Y, (sid,2,3,r), ∞]
time = 11r

Figure 6: The transactions associated with the first round r of our protocol compiler. $R_i(\cdot)$ is a relation which is true given the r-th round message of P_i (for the given correlated randomness and previous messages); m_i is the message of player P_i for round r. Player 3 aborts in the r-th round of the protocol and players 1,2 collect their reward.

Observe that by using a constant-round protocol $\pi_{\mathtt{Mal}}$ [IOZ14], we obtain a protocol with constantly many ledger rounds. Furthermore, as soon as an honest party posts a protocol-related transaction, he is guaranteed to either receive his output or have a positive balance (of at least c coins) after $O(\rho_{\mathrm{c}})$ ledger rounds. The following theorem states the achieved security. We assume the protocol is executed in the synchronous model of Section 3.1.

Theorem 5.2. *Let $\bar{\mathcal{G}} = (\bar{\mathcal{G}}_{\mathrm{LEDGER}}, \bar{\mathcal{G}}_{\mathrm{CLOCK}})$, The above protocol in the $(\bar{\mathcal{G}}, \mathcal{F}_{\mathrm{CORR}}^{\mathcal{D}_{\mathtt{Mal}}})$-hybrid world realizes $\tilde{\mathcal{W}}(\mathcal{F})$ with robust compensation.*

Proof sketch. We first prove that the above protocol is simulatable, by sketching the corresponding simulator $\mathcal{S}$. If the protocol aborts already before the parties make their transactions, then the simulator can trivially simulate such an abort, as he needs to just receive the state of the ledger and see if all wallets corresponding to honest parties have sufficient funds to play the protocol. In the following we show that the rest of the protocol (including the ledger's contents) can be simulated so that if there is an abort, honest parties' wallets have a positive balance as required by Q fairness. First we observe that the simulator $\mathcal{S}$ can easily decide the islands in which the parties are split, as he internally simulates the sampling functionality. Any island other than the one of honest parties (all honest parties will be in the same island because they will post transactions including the same public setup-component) is trivially simulatable as it only consists of adversarial parties and no guarantee is given about their wallets by Q-fairness. Therefore, it suffices to provide a simulator for the honest parties' island. To this direction, the simulator uses the simulator $\mathcal{S}_{\pi_{\mathtt{Mal}}^{+1}}$ which is guaranteed to exist from the security of $\pi_{\mathtt{Mal}}^{+1}$ to decide which messages to embed in the transactions of honest parties (the messages corresponding to corrupted parties are provided by the adversary). If $\mathcal{S}_{\pi_{\mathtt{Mal}}^{+1}}$ would abort, then $\mathcal{S}$ interacts the ideal functionality to abort and continues by claiming back all the committed transactions to the honest parties' wallets, as the protocol would. The soundness of the simulation of $\mathcal{S}_{\pi_{\mathtt{Mal}}^{+1}}$ ensure that the output of the parties and the contents of the ledger in the real and the ideal world are indistinguishable.

The fact that the protocol will eventual terminate given sufficient rounds of activating every party (i.e., in the terminology of Definition 4.5, given a sufficiently high threshold T) follows by inspection of the protocol: in each round every party needs at most a (fixed) polynomial number of activations to post the transactions corresponding to his current-round message-vector. (In fact, the polynomial is only needed in the initial committing-transactions round and from that point on it is linear). To complete the proof, we argue that **(1)** when the protocol does not abort, every honest party has a non-negative balance, and **(2)** when the protocol aborts, then honest parties have a positive balance of at least c coins as required by predicate Q for the simulator to be able to complete its simulation and deliver the (possibly aborting) outputs. These properties are argued as follows:

Property **(1)**: The parties that are not in the honest parties' islands cannot claim any transaction that honest parties make towards them as the ledger will see they as not in the island and reject them. Thus by the last round every honest party will have re-claimed all transactions towards parties not in his island. As far as parties in the honest island are concerned, if no abort occurs then every party will claim all the transactions from parties in his island, and therefore his balance will be 0.

Property **(2)**: Assume that the protocol aborts because some (corrupted) P_i broadcasts an inconsistent message in some round ρ. By inspection of the protocol one can verify that honest parties will be able to claim all transaction-commitments done to them up to round ρ (as they honestly execute their protocol) plus all committed transactions that they made for rounds $\rho + 1 \ldots, \rho_{\mathrm{c}}$. Additionally, because P_i broadcasts an inconsistent message in round ρ, he will be unable to claim transactions of honest parties done from round ρ and on; these bitcoins will be reclaimed by the honest parties, thus giving their wallets a positive balance of at least c coins.

We refer to Appendix A.4 for more details of the proof. □

6 Using Ethereum Contracts

In this section we elaborate on the feasibility of implementing our construction using Ethereum contracts. Ethereum is a type of virtual machine that operates over a blockchain protocol and enables the execution of complex transactions, [Woo14]. The state of Ethereum is comprised of accounts that are determined by their balance and a transaction counter. Transactions in Ethereum contain the recipient account, a signature identifying the sender, the amount of *ether* (Ethereum's currency) and the data to send, as well as two values called $\mathsf{startgas}$ and $\mathsf{gasprice}$. These two values signify that in order for transactions to be processed, "gas" needs to be spent that will be collected by the miner running the transaction; $\mathsf{gasprice}$ represents the cost per computational step and $\mathsf{startgas}$ represents the initial gas value that the sender funds the transaction with. Gas can be funded with ether and corresponds to the product $\mathsf{startgas} \cdot \mathsf{gasprice}$.

Transaction recipients are regular accounts as well as "smart" contracts. A contract is a special account that has its own code that is executed whenever it receives a transaction or a message. Contracts are stateful in the sense that they can maintain data in a local (virtual) memory that has 2^{256} entries. A contract when executed can change its local state as well as generate new transactions.

The contract is executed by the miner that processes an incoming transaction for the contract as part of the state update function of the Ethereum blockchain. The decision to execute the contract depends also on the investment made by the transaction that is incoming; contract code may be expensive to run and thus a miner may refuse to execute the code of the contract if it is not sufficiently funded. When a contract is executed the code of the contract has access to various contextual information such as the current block timestamp, the current block number and so on. Using the current block timestamp, in particular, the contract is able to make time-sensitive decisions. For instance, in this way, a party may generate a contract that conditionally transfers some funds to someone that provides a specific type of data in a transaction. The funds may be locked in the contract while after a certain time the funds in the contract may be withdrawn back by the entity that initiated the contract.

Implementing $\bar{\mathcal{G}}_{\text{LEDGER}}$ as a smart contract. Next we describe how to implement our $\bar{\mathcal{G}}_{\text{LEDGER}}$ with the special transactions of section 5.2. First we note that Ethereum contracts are not able to inspect the blockchain when they are executed. It follows that a single conditional transaction as the one our protocol requires cannot be implemented as a contract. Nevertheless, the whole $\bar{\mathcal{G}}_{\text{LEDGER}}$ can be implemented as a single smart contract denoted as SC that operates as an ethereum application[32]. The SC will use contract storage to maintain account balances for each `address` that is associated with a protocol execution. Account balances will also have the feature that they can have a certain amount of funds from them put on hold. The parties, in order to use SC, will have to initialize their accounts. This will be accomplished by sending a special initialization transaction that indicates an internal SC address `address` to be initialized and the initial amount v in ether that will fund the account. This transaction will transfer v ether to SC from the sender and the smart contract will keep the funds in a reserve while it will introduce `address` in its local state. The account will be given an initial amount of internal currency that is proportional to v according to an exchange rate between SC's internal currency and ether; furthermore, a fee for processing the initialization transaction may be applied by SC.

Recall the type of transactions required by our ledger,

$$\mathsf{tx} = (\mathbb{B}_{v,\texttt{address}_i,\texttt{address}_j,\Sigma=(\tau_-,\tau_+),\alpha,\mathcal{R}),\texttt{aux},\sigma})$$

To issue such a transaction using SC, a corresponding Ethereum $\mathsf{tx}^{\mathsf{ether}}$ transaction will be generated from the sender directed to the account of the smart contract SC that passes as values $\mathsf{tx}^{\mathsf{ether}}.\mathsf{data}[i]$, $i = 0, 1, 2, \ldots$, the transaction elements $v, \texttt{address}_i, \texttt{address}_j, \tau_-, \tau_+, \alpha, \mathcal{R}, \texttt{aux}, \sigma$. The transaction

[32] For a high-level description of building applications within Ethereum see https://github.com/ethereum/wiki/wiki/White-Paper

$\mathsf{tx}^{\mathsf{ether}}$ may also transfer some amount of ether to the smart contract as transaction processing fee (but this is optional). Each incoming transaction tx will be validated (see below) by SC and placed in a temporary local storage buffer (the contract has access to local storage via $\mathsf{contract.storage}$[.]). The contract SC will also maintain a $\mathsf{counter}$ signifying the number of blocks in the internal ledger. Whenever SC is activated it will check the current time ct (observe that SC has access to current time via $\mathsf{block.timestamp}$) and if it is found that more than $\mathtt{T}$ clock ticks have passed since the previous block generation it will take all transactions in buffer and organize them in a block.

We describe next how transactions are validated and processed. We focus on how processing transactions takes place. This shows how to implement $\mathsf{Blockify}$ however validating transactions can be simply extracted from the description below as the precondition that is necessary to process the transaction successfully. Simple transactions, i.e. of the form $\Sigma = ((0, \infty), \perp, \mathcal{R}_\emptyset)$, are processed by SC in a straightforward manner: the internal balance of $\mathtt{address}_i$ is debited by v and the balance of $\mathtt{address}_j$ is credited provided that the signature σ is valid. Time-restricted transactions that are of the form $\Sigma = ((\tau_-, \tau_+), \alpha, \mathcal{R})$ are stored in the SC local state and SC puts a hold for an amount equal to v on account $\mathtt{address}_i$ (however no credit or debit is applied); the hold is marked with the anchor label α. When a transaction has no time-restriction but comes with a predicate $\mathcal{R} \neq \mathcal{R}_\emptyset$, processing for validity by SC requires that the predicate $\mathcal{R}$ is true. If a time-unrestricted transaction for value v from $\mathtt{address}'_1$ to $\mathtt{address}'_2$ is given such that $\alpha \neq \perp$ and a time-restricted transaction for time window (τ_-, τ_+) with the same anchor α and same value v has been previously issued from $\mathtt{address}_1$ to $\mathtt{address}_2$, SC processes it as follows: if $\mathtt{address}_2 = \mathtt{address}'_1$, it is checked that $\mathsf{ct} \in (\tau_-, \tau_+)$ and provided that otherwise the transaction is valid, the amount v is removed from hold, debited in $\mathtt{address}_1$, and credited to $\mathtt{address}'_2$. On the other hand, if it holds $\mathtt{address}_1 = \mathtt{address}'_1$, it is checked that $\mathsf{ct} > \tau_+$, and provided that otherwise the transaction is valid, the amount v is removed from hold and debited in $\mathtt{address}_1$, and credited to $\mathtt{address}'_2$.

The above implementation in conjunction with our protocol provides Q-fairness and robustness, as defined in Section 4.3, when measured in the internal currency of the smart contract SC. Running the protocol though mandates interacting with the smart contract that in itself requires issuing Ethereum transactions that cost ether for paying the miners that implement collectively the virtual machine (such ether is transformed to gas and is spent when the contract code is executed by the miners).

7 Acknowledgements

The first author was supported by ERC project CODAMODA. This research was partly conducted while the second and third authors were visiting the National and Kapodistrian U. of Athens. The authors thank Andrew Miller for helpful discussions.

References

[ADMM14a] Marcin Andrychowicz, Stefan Dziembowski, Daniel Malinowski, and Lukasz Mazurek. Fair two-party computations via the bitcoin deposits. In *1st Workshop on Bitcoin Research 2014 (in Assocation with Financial Crypto)*, 2014. http://eprint.iacr.org/2013/837.

[ADMM14b] Marcin Andrychowicz, Stefan Dziembowski, Daniel Malinowski, and Lukasz Mazurek. Secure multiparty computations on bitcoin. In *2014 IEEE Symposium on Security and Privacy*, pages 443–458. IEEE Computer Society Press, May 2014.

[ALZ13] Gilad Asharov, Yehuda Lindell, and Hila Zarosim. Fair and efficient secure multiparty computation with reputation systems. In Kazue Sako and Palash Sarkar, editors, *ASIACRYPT 2013, Part II*, volume 8270 of *LNCS*, pages 201–220. Springer, Heidelberg, December 2013.

[ASW97] N. Asokan, Matthias Schunter, and Michael Waidner. Optimistic protocols for fair exchange. In *ACM CCS 97*, pages 7–17. ACM Press, April 1997.

[ASW98] N. Asokan, Victor Shoup, and Michael Waidner. Optimistic fair exchange of digital signatures (extended abstract). In Kaisa Nyberg, editor, *EUROCRYPT'98*, volume 1403 of *LNCS*, pages 591–606. Springer, Heidelberg, May / June 1998.

[BCL$^+$05] Boaz Barak, Ran Canetti, Yehuda Lindell, Rafael Pass, and Tal Rabin. Secure computation without authentication. In Victor Shoup, editor, *CRYPTO 2005*, volume 3621 of *LNCS*, pages 361–377. Springer, Heidelberg, August 2005.

[BK14] Iddo Bentov and Ranjit Kumaresan. How to use bitcoin to design fair protocols. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part II*, volume 8617 of *LNCS*, pages 421–439. Springer, Heidelberg, August 2014.

[BN00] Dan Boneh and Moni Naor. Timed commitments. In Mihir Bellare, editor, *CRYPTO 2000*, volume 1880 of *LNCS*, pages 236–254. Springer, Heidelberg, August 2000.

[Can01] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd FOCS*, pages 136–145. IEEE Computer Society Press, October 2001.

[CC00] Christian Cachin and Jan Camenisch. Optimistic fair secure computation. In Mihir Bellare, editor, *CRYPTO 2000*, volume 1880 of *LNCS*, pages 93–111. Springer, Heidelberg, August 2000.

[CDPW07] Ran Canetti, Yevgeniy Dodis, Rafael Pass, and Shabsi Walfish. Universally composable security with global setup. In Salil P. Vadhan, editor, *TCC 2007*, volume 4392 of *LNCS*, pages 61–85. Springer, Heidelberg, February 2007.

[CF01] Ran Canetti and Marc Fischlin. Universally composable commitments. In Joe Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 19–40. Springer, Heidelberg, August 2001.

[CJS14] Ran Canetti, Abhishek Jain, and Alessandra Scafuro. Practical UC security with a global random oracle. In Gail-Joon Ahn, Moti Yung, and Ninghui Li, editors, *ACM CCS 14*, pages 597–608. ACM Press, November 2014.

[Cle86] Richard Cleve. Limits on the security of coin flips when half the processors are faulty (extended abstract). In Juris Hartmanis, editor, *STOC*, pages 364–369. ACM, 1986.

[CLOS02] Ran Canetti, Yehuda Lindell, Rafail Ostrovsky, and Amit Sahai. Universally composable two-party and multi-party secure computation. In *34th ACM STOC*, pages 494–503. ACM Press, May 2002.

[CR03] Ran Canetti and Tal Rabin. Universal composition with joint state. In Dan Boneh, editor, *CRYPTO 2003*, volume 2729 of *LNCS*, pages 265–281. Springer, Heidelberg, August 2003.

[GGJ$^+$15] Juan A. Garay, Ran Gelles, David S. Johnson, Aggelos Kiayias, and Moti Yung. A little honesty goes a long way - the two-tier model for secure multiparty computation. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *TCC 2015, Part I*, volume 9014 of *LNCS*, pages 134–158. Springer, Heidelberg, March 2015.

[GK09] S. Dov Gordon and Jonathan Katz. Complete fairness in multi-party computation without an honest majority. In Omer Reingold, editor, *TCC 2009*, volume 5444 of *LNCS*, pages 19–35. Springer, Heidelberg, March 2009.

[GKM+13] Juan A. Garay, Jonathan Katz, Ueli Maurer, Björn Tackmann, and Vassilis Zikas. Rational protocol design: Cryptography against incentive-driven adversaries. In *54th FOCS*, pages 648–657. IEEE Computer Society Press, October 2013.

[GMPY06] Juan A. Garay, Philip D. MacKenzie, Manoj Prabhakaran, and Ke Yang. Resource fairness and composability of cryptographic protocols. In Shai Halevi and Tal Rabin, editors, *TCC 2006*, volume 3876 of *LNCS*, pages 404–428. Springer, Heidelberg, March 2006.

[GMW87] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. In Alfred Aho, editor, *19th ACM STOC*, pages 218–229. ACM Press, May 1987.

[Gol01] Oded Goldreich. *Foundations of Cryptography: Basic Tools*, volume 1. Cambridge University Press, Cambridge, UK, 2001.

[IOZ14] Yuval Ishai, Rafail Ostrovsky, and Vassilis Zikas. Secure multi-party computation with identifiable abort. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part II*, volume 8617 of *LNCS*, pages 369–386. Springer, Heidelberg, August 2014.

[KB14] Ranjit Kumaresan and Iddo Bentov. How to use bitcoin to incentivize correct computations. In Gail-Joon Ahn, Moti Yung, and Ninghui Li, editors, *ACM CCS 14*, pages 30–41. ACM Press, November 2014.

[KMS+15] Ahmed Kosba, Andrew Miller, Elaine Shi, Zikai Wen, and Charalampos Papamanthou. Hawk: The blockchain model of cryptography and privacy-preserving smart contracts. Cryptology ePrint Archive, Report 2015/675, 2015. `http://eprint.iacr.org/2015/675`.

[KMTZ13] Jonathan Katz, Ueli Maurer, Björn Tackmann, and Vassilis Zikas. Universally composable synchronous computation. In Amit Sahai, editor, *TCC 2013*, volume 7785 of *LNCS*, pages 477–498. Springer, Heidelberg, March 2013.

[Nak08] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. http://bitcoin.org/bitcoin.pdf, 2008.

[Pin03] Benny Pinkas. Fair secure two-party computation. In Eli Biham, editor, *EUROCRYPT 2003*, volume 2656 of *LNCS*, pages 87–105. Springer, Heidelberg, May 2003.

[RKS15] Tim Ruffing, Aniket Kate, and Dominique Schröder. Liar, liar, coins on fire!: Penalizing equivocation by loss of bitcoins. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, Denver, CO, USA, October 12-6, 2015*, pages 219–230, 2015.

[Woo14] Gavin Wood. Ethereum: A secure decentralized transaction ledger. 2014. http://gavwood.com/paper.pdf.

[Yao82] Andrew Chi-Chih Yao. Protocols for secure computations (extended abstract). In *23rd FOCS*, pages 160–164. IEEE Computer Society Press, November 1982.

A Proofs

A.1 Proof of Lemma 4.2

Proof. The proof idea is similar to that for composition theorem in [CDPW07]. Here we need to show that for all PPT real world adversary $\mathcal{A}$ there exists PPT simulator $\mathcal{S}$ so that for all PPT environment $\mathcal{Z}$, the following holds:

$$\textsc{Exec}_{\sigma,\mathcal{S},\mathcal{Z}}^{\bar{\mathcal{G}},\mathcal{W}_{\mathsf{Q},\bar{\mathcal{G}}}(\mathcal{F})} \overset{\mathrm{c}}{\approx} \textsc{Exec}_{\sigma^{\pi},\mathcal{A},\mathcal{Z}}^{\bar{\mathcal{G}}} \tag{2}$$

By the condition that π realizes $\mathcal{F}$ with $\mathsf{Q}_{\bar{\mathcal{G}}}$-fairness with respect to global functionality $\bar{\mathcal{G}}$, we have: $\forall \mathcal{A}'$ $\exists \mathcal{S}'$ so that $\forall \mathcal{Z}'$

$$\textsc{Exec}^{\bar{\mathcal{G}}, \mathcal{W}_{\mathsf{Q},\bar{\mathcal{G}}}(\mathcal{F})}_{\mathcal{S}',\mathcal{Z}'} \overset{\text{c}}{\approx} \textsc{Exec}^{\bar{\mathcal{G}}}_{\pi,\mathcal{A}',\mathcal{Z}'} \tag{3}$$

We next prove the theorem.

We first describe the real execution $\textsc{Exec}^{\bar{\mathcal{G}}}_{\sigma^{\pi},\mathcal{A},\mathcal{Z}}$. Let K be a polynomial upper bound on the number of instances of π that are invoked by σ, and let $\pi[k]$ denote the k-th copy of protocol π. Here let the adversary $\mathcal{A} = \left(\mathcal{A}^{\sigma}, \mathcal{A}^{\pi[1]}, \mathcal{A}^{\pi[2]}, \ldots, \mathcal{A}^{\pi[K]}\right)$, where each $\mathcal{A}^{\pi[k]}$ is interacting with the k-th instance of π. Note that, the environment provides inputs to (and receives outputs from) the "father" protocol σ, and the protocol σ provides inputs to (and receives outputs from) its own subroutines $\pi[k]$'s. We remark that, here all protocol instances (including the father protocol and the subroutines) are allowed to access to the global functionality $\bar{\mathcal{G}}$.

We then describe the $\mathcal{W}_{\mathsf{Q},\bar{\mathcal{G}}}(\mathcal{F})$-hybrid execution $\textsc{Exec}^{\bar{\mathcal{G}}, \mathcal{W}_{\mathsf{Q},\bar{\mathcal{G}}}(\mathcal{F})}_{\sigma,\widetilde{\mathcal{A}},\mathcal{Z}}$. Now we let $\mathcal{W}_{\mathsf{Q},\bar{\mathcal{G}}}(\mathcal{F})[k]$ denote the k-th copy of functionality $\mathcal{W}_{\mathsf{Q},\bar{\mathcal{G}}}(\mathcal{F})$ that invoked by protocol σ. Similarly, we define the adversary $\widetilde{\mathcal{A}} = \left(\mathcal{A}^{\sigma}, \mathcal{S}^{\pi[1]}, \mathcal{S}^{\pi[2]}, \ldots, \mathcal{S}^{\pi[K]}\right)$, where each $\mathcal{S}^{\pi[k]}$ is interacting with the k-th instance of $\mathcal{W}_{\mathsf{Q},\bar{\mathcal{G}}}(\mathcal{F})$. As mentioned before, here the environment provides inputs to protocol σ, and protocol σ provides input to its own subroutines, functionality copies $\mathcal{W}_{\mathsf{Q},\bar{\mathcal{G}}}(\mathcal{F})[k]$'s; all protocol instances are allowed to access to the global functionality $\bar{\mathcal{G}}$.

Based on the above description, we next show the two worlds $\textsc{Exec}^{\bar{\mathcal{G}}}_{\sigma^{\pi},\mathcal{A},\mathcal{Z}}$ and $\textsc{Exec}^{\bar{\mathcal{G}}, \mathcal{W}_{\mathsf{Q},\bar{\mathcal{G}}}(\mathcal{F})}_{\sigma,\widetilde{\mathcal{A}},\mathcal{Z}}$ are indistinguishable through a hybrid argument. We define hybrids $\mathbf{Hyb}^k$ as follows:

- Let σ^k denote the following protocol instances:
  - an instance of σ;
  - $k-1$ instances of π, denoted $\pi[1], \ldots, \pi[k-1]$;
  - $K-k+1$ instances of $\mathcal{W}_{\mathsf{Q},\bar{\mathcal{G}}}(\mathcal{F})$, denoted $\mathcal{W}_{\mathsf{Q},\bar{\mathcal{G}}}(\mathcal{F})[k], \ldots, \mathcal{W}_{\mathsf{Q},\bar{\mathcal{G}}}(\mathcal{F})[K]$;
- Let $\mathcal{A}^k$ denote the following adversary copies:
  - $\mathcal{A}^{\sigma}$;
  - $k-1$ instances of $\mathcal{A}^{\pi}$, denoted $\mathcal{A}^{\pi[1]}, \ldots, \mathcal{A}^{\pi[k-1]}$;
  - $K-k+1$ instances of $\mathcal{S}^{\pi}$, denoted $\mathcal{S}^{\pi[k]}, \ldots, \mathcal{S}^{\pi[K]}$;

Define $\mathcal{B}$, which consists of K copies of adversaries, and K copies of protocol/functionality instances. Define $\mathcal{D}$, which consists of the k-th copy of adversary $\mathcal{A}$, and the k-th copy of protocol/functionality instance π. If $\mathcal{A} = \mathcal{S}^{\pi[k]}$ and $\pi = \mathcal{W}_{\mathsf{Q},\bar{\mathcal{G}}}(\mathcal{F})[k]$, then $\mathcal{B}$ is identical to $\mathbf{Hyb}^k$. If $\mathcal{A} = \mathcal{A}^{\pi[k]}$ and $\pi = \pi[k]$, then $\mathcal{B}$ is identical to $\mathbf{Hyb}^{k+1}$. We next show that adjacent hybrids are indistinguishable.

Claim A.1. *For $k \in \{1, \ldots, K\}$, the hybrids $\mathbf{Hyb}^k$ and $\mathbf{Hyb}^{k+1}$ are indistinguishable to any* PPT *$\mathcal{Z}$.*

Proof. By contradiction, assume there is a PPT $\mathcal{Z}$ who can tell the difference between $\mathbf{Hyb}^k$ and $\mathbf{Hyb}^{k+1}$. That means, $\textsc{Exec}^{\bar{\mathcal{G}}}_{\sigma^k,\mathcal{A}^k,\mathcal{Z}} \overset{\text{c}}{\not\approx} \textsc{Exec}^{\bar{\mathcal{G}}}_{\sigma^{k+1},\mathcal{A}^{k+1},\mathcal{Z}}$. Based on such $\mathcal{Z}$, we can define $\mathcal{Z}^k$ to simulate the interaction of all the rest of the network except the k-th place of the subroutine.

Based on the definition of coercion hybrid $\mathbf{Hyb}^k$ above, we can easily see that $\textsc{Exec}^{\bar{\mathcal{G}},\mathcal{W}_{\mathsf{Q},\bar{\mathcal{G}}}(\mathcal{F})}_{\mathcal{S}^{\pi[k]},\mathcal{Z}^k}$ is the representation of $\textsc{Exec}^{\bar{\mathcal{G}}}_{\sigma^k,\mathcal{A}^k,\mathcal{Z}}$. Similarly, we can easily see that $\textsc{Exec}^{\bar{\mathcal{G}}}_{\pi,\mathcal{A}^{\pi[k]},\mathcal{Z}^k}$ is the representation of $\textsc{Exec}^{\bar{\mathcal{G}}}_{\sigma^{k+1},\mathcal{A}^{k+1},\mathcal{Z}}$.

Based on the assumption that $\textsc{Exec}^{\bar{\mathcal{G}}}_{\sigma^k,\mathcal{A}^k,\mathcal{Z}} \overset{c}{\not\approx} \textsc{Exec}^{\bar{\mathcal{G}}}_{\sigma^{k+1},\mathcal{A}^{k+1},\mathcal{Z}}$, we immediately have $\textsc{Exec}^{\bar{\mathcal{G}},\mathcal{W}_{\mathsf{Q},\bar{\mathcal{G}}}(\mathcal{F})}_{\mathcal{S}^{\pi[k]},\mathcal{Z}^k} \overset{c}{\not\approx}$ $\textsc{Exec}^{\bar{\mathcal{G}}}_{\pi,\mathcal{A}^{\pi[k]},\mathcal{Z}^k}$. However, this contradicts to the premise in Equation 3. That means, our assumption that $\textsc{Exec}^{\bar{\mathcal{G}}}_{\sigma^k,\mathcal{A}^k,\mathcal{Z}} \overset{c}{\not\approx} \textsc{Exec}^{\bar{\mathcal{G}}}_{\sigma^{k+1},\mathcal{A}^{k+1},\mathcal{Z}}$ is not true. This completes our proof of the claim that $\mathbf{Hyb}^k$ and $\mathbf{Hyb}^{k+1}$ are indistinguishable for $k \in \{1, \ldots, K\}$. □ □

Finally, we note that hybrid $\mathbf{Hyb}^1$ is identical to $\textsc{Exec}^{\bar{\mathcal{G}},\mathcal{W}_{\mathsf{Q},\bar{\mathcal{G}}}(\mathcal{F})}_{\sigma,\widetilde{\mathcal{A}},\mathcal{Z}}$, and the hybrid $\mathbf{Hyb}^{K+1}$ is identical to $\textsc{Exec}^{\bar{\mathcal{G}}}_{\sigma^\pi,\mathcal{A},\mathcal{Z}}$. Based on the claim above, we can see that $\mathbf{Hyb}^1 \overset{c}{\approx} \mathbf{Hyb}^2 \overset{c}{\approx} \cdots \overset{c}{\approx} \mathbf{Hyb}^K \overset{c}{\approx} \mathbf{Hyb}^{K+1}$. This implies that $\textsc{Exec}^{\bar{\mathcal{G}}}_{\sigma^\pi,\mathcal{A},\mathcal{Z}} \overset{c}{\approx} \textsc{Exec}^{\bar{\mathcal{G}},\mathcal{W}_{\mathsf{Q},\bar{\mathcal{G}}}(\mathcal{F})}_{\sigma,\widetilde{\mathcal{A}},\mathcal{Z}}$, which completes the proof of the lemma. □ □

A.2 Proof of Theorem 4.3

Proof. By the condition that π realizes $\mathcal{F}$ with $\mathsf{Q}_{\bar{\mathcal{G}}}$-fairness with respect to global functionality $\bar{\mathcal{G}}$, we have: $\forall \mathcal{A}' \ \exists \mathcal{S}'$ so that $\forall \mathcal{Z}'$

$$\textsc{Exec}^{\bar{\mathcal{G}},\mathcal{W}_{\mathsf{Q},\bar{\mathcal{G}}}(\mathcal{F})}_{\mathcal{S}',\mathcal{Z}'} \overset{c}{\approx} \textsc{Exec}^{\bar{\mathcal{G}}}_{\pi,\mathcal{A}',\mathcal{Z}'} \tag{4}$$

By the condition that σ realizes $\mathcal{H}$ with $\mathsf{Q}'_{\bar{\mathcal{G}}}$-fairness with respect to global functionality $\bar{\mathcal{G}}$, in the $\mathcal{W}_{\mathsf{Q},\bar{\mathcal{G}}}(\mathcal{F})$-hybrid world, we have: $\forall \mathcal{A}' \ \exists \mathcal{S}'$ so that $\forall \mathcal{Z}'$,

$$\textsc{Exec}^{\bar{\mathcal{G}},\mathcal{W}_{\mathsf{Q}',\bar{\mathcal{G}}}(\mathcal{H})}_{\mathcal{S}',\mathcal{Z}'} \overset{c}{\approx} \textsc{Exec}^{\bar{\mathcal{G}},\mathcal{W}_{\mathsf{Q},\bar{\mathcal{G}}}(\mathcal{F})}_{\sigma,\mathcal{A}',\mathcal{Z}'} \tag{5}$$

Using the transitivity of indistinguishability, we have for all PPT real world adversary $\mathcal{A}$ there exists PPT simulator $\mathcal{S}$ so that for all PPT environment $\mathcal{Z}$, the following holds:

$$\textsc{Exec}^{\bar{\mathcal{G}},\mathcal{W}_{\mathsf{Q}',\bar{\mathcal{G}}}(\mathcal{H})}_{\mathcal{S},\mathcal{Z}} \overset{c}{\approx} \textsc{Exec}^{\bar{\mathcal{G}}}_{\sigma^\pi,\mathcal{A},\mathcal{Z}} \tag{6}$$

□

A.3 Proof of Theorem 4.4

Canetti and Fischlin [CF01] show the impossibility of realizing $\mathcal{F}_{\text{COM}}$ in the plain model. Roughly speaking, if a protocol π UC-realizes $\mathcal{F}_{\text{COM}}$, then an ideal world simulator $\mathcal{S}$ should be able to be constructed and satisfy the following properties:

- When the committer is corrupted, $\mathcal{S}$ must be able to "extract" the committed value once the commitment phase is done. That is, $\mathcal{S}$ has to come up with a value x such that the committer will almost never be able to successfully decommit to any $x' \neq x$. This is so since in the ideal process $\mathcal{S}$ has to explicitly provide $\mathcal{F}_{\text{COM}}$ with a committed value.

- When the receiver is corrupted, $\mathcal{S}$ has to be able to simulate a fake commitment phase and yet can be opened to any value at the time of opening. This is so since $\mathcal{S}$ has to provide adversary $\mathcal{A}$ and environment $\mathcal{Z}$ with the simulated commitment c before the value committed to is known. All this needs to be done without rewinding the environment $\mathcal{Z}$.

Intuitively, these requirements look impossible to meet: a simulator that has the above abilities can be used by a dishonest receiver to "extract" the committed value from an honest committer. This intuition can indeed be formalized to show that in the plain model it is impossible to UC-realize $\mathcal{F}_{\text{COM}}$ by any two-party protocol. This idea extends to the $\bar{\mathcal{G}}_{\text{LEDGER}}$ hybrid world for realizing commitment functionality with Q fairness.

Proof. Suppose, for contradiction, that there exists protocol π that realizes $\mathcal{F}_{\text{COM}}$ in the $\bar{\mathcal{G}}_{\text{LEDGER}}$ hybrid world with Q fairness. WLOG, we can consider a special Q predicate that it is never triggered. Assume at the end of the commitment phase, receiver acknowledges the committer by a $\mathsf{receipt}$ message. Consider an execution of π by an adversarial committer $\mathcal{A}_C$ and an honest receiver R, and WLOG we assume that the adversary merely forwards the communication messages between the environment $\mathcal{Z}_C$ and the honest receiver R (Note that this adversarial behavior is implementable by an adversary as the adversary does not need to apply any transformation on the state and merely forwards it). Here $\mathcal{Z}_C$ privately chooses a random bit b at the beginning and then runs the protocol of the honest committer based on input bit b and R's answers, and then in the name of the committer sends the generated messages to R. Once $\mathcal{Z}_C$ received a $\mathsf{receipt}$ message from R at the end of committing stage, it starts running the honest opening protocol in the name of the committer, and receives bit b' from R at the end of opening stage. Finally, $\mathcal{Z}_C$ outputs 1 iff $b' = b$. We know that if both committer and receiver are honest in an execution of π, then in the opening phase the receiver always outputs the bit committed to by the committer, i.e., $b' = b$ always holds. By assumption that π realizes $\mathcal{F}_{\text{COM}}$ in the $\bar{\mathcal{G}}_{\text{LEDGER}}$ hybrid world with Q fairness, there should exist an ideal world simulator $\mathcal{S}$ that interacts with $\mathcal{F}_{\text{COM}}$ as well as $\bar{\mathcal{G}}_{\text{LEDGER}}$ and generates a view for $\mathcal{Z}_C$ that is indistinguishable from a real execution with π in the $\bar{\mathcal{G}}_{\text{LEDGER}}$ hybrid world. We note that, $\mathcal{S}$ must make sure $b = b'$ almost always, where b' is the bit that $\mathcal{S}$ sends to $\mathcal{F}_{\text{COM}}$. This means that the simulator $\mathcal{S}$ must be able to generate the correct bit b before the opening phase.

Next based on this $\mathcal{S}$, we are able to construct another environment, $\mathcal{Z}_R$, and a corrupted receiver $\mathcal{A}_R$, such that $\mathcal{Z}_R$ successfully distinguishes between an execution of π and an interaction with $\mathcal{F}_{\text{COM}}$ for any simulator $\mathcal{S}_R$. $\mathcal{Z}_R$ and $\mathcal{A}_R$ proceed as follows: $\mathcal{Z}_R$ chooses a random bit b and hands b as input to the honest committer C; $\mathcal{A}_R$ simply runs $\mathcal{S}$ and forwards all interaction between the committer and $\mathcal{S}$, and between $\bar{\mathcal{G}}$ and $\mathcal{S}$ (again this strategy is implementable by an adversary as the adversary does not need to apply any transformation on the state); once $\mathcal{A}_R$ receives a bit b', it is passed to $\mathcal{Z}_R$ who then outputs 1 iff. $b = b'$.

Note that $\mathcal{S}$ can extract the committed bit b almost always, without rewinding or any additional information. In contrast, when $\mathcal{Z}_R$ interacts with $\mathcal{F}_{\text{COM}}$, the $\mathcal{S}_R$'s view is independent of b, and thus $b = b'$ with probability exactly one half. Therefore, $\mathcal{Z}_R$ can tell the difference between its interaction with the $\bar{\mathcal{G}}$ hybrid world or with $\mathcal{F}_{\text{COM}}$ and ideal world for any $\mathcal{S}_R$.

□

A.4 Proof of Theorem 5.2

To prove the theorem, we need to construct a simulator $\mathcal{S}$ so that for all PPT adversary $\mathcal{A}$ and PPT environment $\mathcal{Z}$, the execution in the $\bar{\mathcal{G}} = (\bar{\mathcal{G}}_{\text{LEDGER}}, \bar{\mathcal{G}}_{\text{CLOCK}})$ and $\mathcal{F}^{\mathcal{D}}_{\text{CORR}}$ hybrid world and the simulated execution in the ideal world are indistinguishable. We give the construction of the simulator $\mathcal{S}$ as follows.

The simulator emulates a copy of the adversary $\mathcal{A}$ internally, and also needs to provide an emulated view for $\mathcal{A}$. First, the simulator emulates $\mathcal{A}$'s communication with the environment $\mathcal{Z}$. Note that the simulator interacts with the external global functionalities $\bar{\mathcal{G}}_{\text{LEDGER}}$ and $\bar{\mathcal{G}}_{\text{CLOCK}}$, and also with the external wrapper functionality $\mathcal{W} = \mathcal{W}_{\bar{\mathcal{G}},\mathsf{Q}}(\mathcal{F})$.

Now the simulator internally emulates a copy of $\mathcal{F}^{\mathcal{D}}_{\text{CORR}}$; more concretely, the simulator chooses random strings $\{\gamma_i\}$, and then uses such $\{\gamma_i\}$ to compute $(\{R^{\texttt{priv}}_i\}_{i\in[n]}, R^{\mathsf{pub}})$. Note that in the ideal world, when the environment $\mathcal{Z}$ activates a dummy party P_i, expecting some resource-setup as response, the dummy party P_i sends $(\text{ALOCATE}, \mathsf{sid})$ to $\mathcal{W}$, and the wrapper functionality $\mathcal{W}$ sends $(\text{COINS}, \mathsf{sid}, P_i)$ to $\mathcal{S}$; and now $\mathcal{S}$ returns $\mathcal{W}$ with $(\text{COINS}, \mathsf{sid}, P_i, \gamma_i)$; note that the wrapper functionality then run $\text{Gen}(1^\kappa; \gamma_i)$ to compute the response for the dummy P_i.

Internally, the simulator $\mathcal{S}$ now emulates a copy of $\tilde{P}_i$ with $R^{\texttt{priv}}_i$ and R^{pub}, and instructs the wrapper functionality $\mathcal{W}$ to return $(\text{DELIVER}, \mathsf{sid}, R^{\mathsf{pub}}, P_i)$ to dummy P_i. We remark that, $R^{\texttt{priv}}_i$ will be P_i's private component (including all random coins he needs to run the protocol, along with his signing key sk_i) of the setup and R^{pub} is the public component (the same for every party P_j) which includes a vector of public (verification) keys $(\mathsf{vk}_1, \ldots, \mathsf{vk}_n)$ corresponding to the signing keys $(\mathsf{sk}_1, \ldots, \mathsf{sk}_n)$ and other information.

In the ideal world, when $\mathcal{Z}$ activates a dummy party P_i with $(\text{INPUT}, \mathsf{sid}, x)$, the dummy party forwards this to the wrapper functionality $\mathcal{W}$. Now the wrapper functionality sends READ to $\bar{\mathcal{G}}_{\text{LEDGER}}$, obtains the response trans from $\bar{\mathcal{G}}_{\text{LEDGER}}$; if $\neg\mathsf{Q}^{\text{Init}}(R^{\mathsf{pub}}, \mathsf{trans})$, i.e., some party has no sufficient amount funds for participating in the protocol, the wrapper functionality $\mathcal{W}$ will simultaneously deliver $\perp$ to all parties and the simulator $\mathcal{S}$. Now the simulator $\mathcal{S}$ instructs the internally simulated $\tilde{P}_i$ to broadcast $\perp$, and every internally simulated party aborts with output $\perp$.

If all parties have sufficient amount funds, then the wrapper functionality $\mathcal{W}$ forwards $(\text{INPUT}, \mathsf{sid}, x)$ to $\mathcal{F}$ as input for P_i. When $(\text{INPUT}, \mathsf{sid}, P_i)$ is sent to the simulator $\mathcal{S}$, now the simulator uses the simulator $\mathcal{S}_{\pi^{+1}_{\texttt{Mal}}}$ which is guaranteed to exist from the security of $\pi^{+1}_{\texttt{Mal}}$ to decide which messages to embed in the transactions of honest parties (the messages corresponding to corrupted parties are provided by the adversary). Note that the simulator $\mathcal{S}$ is aware of $R^{\texttt{priv}}_i$, and such $R^{\texttt{priv}}_i$ can be used for embedding $\pi^{+1}_{\texttt{Mal}}$ messages into the ledger transactions.

If $\mathcal{S}_{\pi^{+1}_{\texttt{Mal}}}$ does not abort, Q^{Dlv} will be triggered; now $\mathcal{S}$ interacts the wrapper functionality to fair-delivery the output messages to parties. The soundness of the simulation of $\mathcal{S}_{\pi^{+1}_{\texttt{Mal}}}$ ensures that the output of the parties and the contents of the ledger in the real and the ideal world are indistinguishable.

Now if $\mathcal{S}_{\pi^{+1}_{\texttt{Mal}}}$ would abort, Q^{Abt} will be triggered; now $\mathcal{S}$ interacts the wrapper functionality to abort and continues by claiming back all the committed transactions to the honest parties' wallets, as the protocol would. The soundness of the simulation of $\mathcal{S}_{\pi^{+1}_{\texttt{Mal}}}$ ensures that the output of the parties and the contents of the ledger in the real and the ideal world are indistinguishable.

To complete the proof, we need to argue that the protocol terminates given sufficient rounds of activating every party (i.e., in the terminology of Definition 4.5, given a sufficiently high threshold T) and that when it does the following properties are satisfied which will ensure that the simulator is able to complete its simulation and deliver the (possibly aborting) outputs: **(1)** when the protocol does not abort, every honest party has a non-negative balance, and **(2)** when the protocol aborts, then honest parties have a positive balance of at least c coins as required by predicate Q.

The fact that the protocol will eventual terminate given sufficient rounds of activating every party (i.e., in the terminology of Definition 4.5, given a sufficiently high threshold T) follows by inspection of the protocol: in each round every party needs at most a (fixed) polynomial number of activations to post the transactions corresponding to his current-round message-vector. (In fact, the polynomial is only needed in the initial committing-transactions round and from that point on it is linear).

The properties are argued as follows:

- Property **(1)**: The parties that are not in the honest parties' islands cannot claim any transaction that honest parties make towards them as the ledger will see they as not in the island and reject them. Thus by the last round every honest party will have re-claimed all transactions towards parties not in his island. As far as parties in the honest island are concerned, if no abort occurs then every party will claim all the transactions from parties in his island, and therefore his balance will be 0.
- Property **(2)**: Assume that the protocol aborts because some (corrupted) P_i broadcasts an inconsistent message in some round ρ. By inspection of the protocol one can verify that honest parties will be able to claim all transaction-commitments done to them up to round ρ (as they honestly execute their protocol) plus all committed transactions that they made for rounds $\rho+1\ldots,\rho_{\mathrm{c}}$. Additionally, because P_i broadcasts an inconsistent message in round ρ, he will be unable to claim transactions of honest parties done from round ρ and on; these bitcoins will be reclaimed by the honest parties, thus giving their wallets a positive balance of at least c coins.