Shielded Computations in Smart Contracts Overcoming Forks

Vincenzo Botta[2], Daniele Friolo[1], Daniele Venturi[1], and Ivan Visconti[2]

[1]*Department of Computer Science, Sapienza University of Rome, Italy*
[2]*DIEM, University of Salerno, Italy*

Abstract

In this work we consider executions of smart contracts on forking blockchains (e.g., Ethereum), and study security and delay issues due to forks. As security notion for modeling executions of smart contracts, we focus on secure multi-party computation (MPC). In particular we consider on-chain MPC executions with the aid of smart contracts. The classical double-spending problem tells us that messages of the MPC protocol should be confirmed on-chain before playing the next ones, thus slowing down the entire execution.

In this work we show how to design smart contracts on forking blockchains reducing the number of confirmations, still maintaining security and fairness. Our contributions are twofold:

- We design a compiler that takes any "digital and universally composable" MPC protocol (with or without honest majority) and transforms it into another one (for the same task and same setup) where all messages are played on-chain without delays and still security is maintained. The special requirements on the starting protocol mean that messages consists only of bits (e.g., no hardware token is sent) and security holds also in the presence of other protocols. Then we show that our compiler satisfies fairness with penalties as long as honest players only wait once.
- For the concrete case of fairly tossing multiple coins with penalties, we notice that the lottery protocol of Andrychowicz *et al.* (S&P '14) becomes insecure if players do not wait for the confirmations of several transactions. In addition, we present a smart contract that instead retains security even when all honest players immediately answer to transactions appearing on-chain. This second result improves what our generic compiler achieves for this specific functionality. We analyze the performance using Ethereum as testbed.

By reducing the number of confirmations, both of our protocols offer a significant speed-up in terms of completion time.

Keywords: blockchains, finality, forks, smart contracts, secure computation.

Contents

1 Introduction

The rise of blockchains[1] is progressively changing the way transactions are executed over the Internet. Indeed, the traditional client-server paradigm turns out to be insufficient when many parties want to perform a distributed computation, especially in cases where features like public verifiability and automatic punishment are desired. Blockchains through the execution of smart contracts naturally allow many players to perform a joint computation, even when they are not simultaneously online; moreover, they allow to publicly check the actions of all players[2] and enforce a proper behavior through financial punishments.

1.1 Forks, Finality and Double Spending

Typical blockchains experience some delays before a transaction can be considered confirmed. Indeed, a large part of the most used blockchains consists of a list of blocks that can temporary fork. In such cases, fork-resolution mechanisms decide which branch is eventually part of the list of blocks and which one is discarded, at the price of cutting off some transactions that for some time have appeared on the blockchain. This finality limitations generate delays and uncertainty and a significant effort has been made recently to obtain blockchains with better finality [MMNT19, BG17, PS17, PS18, GHM+17, CPS18].

The existence of transactions that appear and then disappear from a blockchain is the source of the (in)famous double-spending attack. In such attack, the adversary performs a payment thorough a transaction on the blockchain in order to receive a service off-chain. If subsequently, due to a fork, the transaction related to the payment disappears from the blockchain, then the attacker gets the money back and can spend it for something else. Therefore, at the end of the day, the off-chain service was received for free and the same coins can be successfully spent twice. The crucial point of the double spending attack is that, while the payment transaction disappears, the obtained service is not canceled since it is not linked to the payment transaction happening on chain.

The solution to the double spending problem is pretty harsh: the receiver of a payment will have to wait long time—i.e., until the transaction is confirmed and becomes irreversible—before taking future actions. Obviously, this can be problematic when an entire process consists of many sequential transactions and the confirmation time is long.

Interestingly, the double spending problem seems to disappear when instead the service consists of another on-chain transaction that is connected to the payment transaction. Indeed, in this case, if as consequence of a fork the payment transaction disappears, then the service transaction disappears too. This chaining of transactions related to the same process can be easily implemented through smart contracts. Indeed, a smart contract can have an initial state s_1 that is updated transaction by transaction obtaining s_2, s_3, and so on. Let t_i be the transaction that changes the state from s_i to s_{i+1}. If because of a fork the state goes back to s_i from a state s_j, only t_i is again applicable to s_i, while instead all transactions $t_{i+1}, \ldots, t_{j-1}$ are not applicable to state s_i. Therefore, by invalidating t_i (similarly to the double spending problem where money is used in a new transaction making t_i invalid), then $t_{i+1}, \ldots, t_{j-1}$ will be invalidated too.

This motivates the possibility of running smart contracts efficiently, without waiting that every single transaction is confirmed before broadcasting the next one.

---

[1] Throughout the paper, we use the terms "blockchain" and "distributed ledger" interchangeably.

[2] We will often use the two terms "party" and "player" as synonyms.

1.2 Attacks to Smart Contracts

Since transactions take long time to be confirmed in a forking blockchain, the full execution of a smart contract with multiple sequential transactions might take too long. It would thus be natural to speed up the execution of smart contracts by rushing and playing messages immediately. Indeed, as mentioned above, by appropriately chaining the transactions of a smart contract, attacks consisting in exploiting the cancellation of a transaction like the double-spending attack are not effective,[3] and therefore rushing could be a valid option.

However, we notice that forks can help an adversary to mount more subtle attacks. For example, an honest player could answer to some transaction A by sending another transaction B as soon as A appears on the blockchain. Obviously, in case of forks, the transaction A could appear on the blockchain in different branches, and then multiple copies of B would follow A. While at first sight this seems to be fine, an adversary computing A can exploit his view of B in a branch of a fork to play adaptively a message different than A in another branch, invalidating some expected security property of the smart contract. Indeed, different transactions A1 and A2 could be played by the adversary in the two branches of a fork, and (potentially different) transactions B1 and B2 sent by a honest player might be required and played as answers. Notice that the honest player could become aware of the fork only after the fact, i.e., after A1 and B1 have been played already. Indeed, because of a fork, transactions A1 and B1 could disappear, and only now that a transaction A2 appears instead of A1 the honest players realizes that existence of a fork. The honest player therefore will have to compute B2 to continue the execution of the smart contract. The fact that the adversary can play A2 adaptively after having seen B1 can produce a deviation from the expected behavior of the smart contract, therefore compromising the appealing transparency and robustness guarantees of this technology. The above scenario can be a serious threat for confidential data of honest players.

1.3 Why MPC on Blockchains?

Blockchains offer the public verifiability of an entire distributed computation so that in case of dispute everyone can verify what happened and when. Moreover smart contracts can automatically punish whoever violates some a-priori established rules during the execution of the process. Clearly the above advantages are useful also when players are interested in running a computation preserving privacy, therefore using MPC. A popular example of MPC that can benefit from a blockchain is e-voting since several schemes rely on a bulletin board (i.e., a blockchain) to get a public verifiability property named universal verifiability. In addition in [ADMM14, ADMM16], Andrychowicz *et al.* have shown how to use blockchains to add fairness through penalties to MPC protocols with dishonest majority, somehow circumventing the impossibility results of Cleve [Cle86] that holds without setup. Notice that playing on-chain adds the interesting benefit of running a protocol without requiring players to be online at the same time. Moreover, differently from an MPC protocol running on a WAN using TCP/IP, where players have to know each other's IP addresses beforehand, with the aid of a ledger any player can join an MPC execution by just reading a transaction containing useful informations needed to participate[4] (e.g. the functionality, the minimum number of parties, any other identification information).

The above interesting features and the dilemma about playing immediately risking security

---

[3]Recall that we are focusing on smart contract that during the intermediate state updates do not have off-chain impact.

[4]Note that blockchain identifiers are usually public pseudonyms not necessary correlated with the real user identities. This gives to each player another layer of privacy w.r.t. IP address identification.

or waiting for confirmation making the entire process very slow motivate our work aiming at obtaining smart contracts for fast/fair/secure/publicly-verifiable MPC on forking blockchains.

1.4 Our Contributions

Our main contributions are outlined below.

Insecurity of smart contracts with quick players. Consider a simple smart contract executed by two players, Alice and Bob, willing to establish jointly a random string:

1. Alice starts the protocol by sending to the smart contract a commitment to a random string r_1;
2. Bob sends a random string r_2 to the smart contract;
3. Alice then opens the commitment, and if the opening is valid the common string is defined to be $r = r_1 \oplus r_2$.

For concreteness, say that Alice is honest and Bob is corrupted, and assume that a fork happens after Alice already sent the commitment. If Bob runs the protocol honestly on the first branch, he gets to see Alice's opening, and thus he can completely bias the output on the other branch by just sending $r_2' = r' \oplus r_1$ to the smart contract, for any value r' of his choice. This motivating example clearly shows that, unless one has proven some kind of resilience to forks, it is certainly preferable to always wait that transactions are confirmed, at the price of having very slow executions of the smart contract. Such slowness could be unacceptable in some applications.

Defining on-chain MPC with quick players. The execution of a smart contract through transactions sent by different players is a computation involving multiple parties, and therefore when considering "security" of such computations we naturally refer to secure MPC. As our first conceptual contribution, we formalize different ways how to execute an MPC protocol in the presence of a blockchain. Our definition builds on the model of blockchain protocols, introduced in [PSS17, GG17]. Intuitively, a blockchain protocol allows the players to keep a consistent record of transactions satisfying: (i) *consistency* (i.e., the view of the blockchain obtained by different players is identical up to pruning k blocks from the chain); and (ii) *liveness* (i.e., if all honest parties attempt to broadcast a message, then after w rounds, an honest party will see that message at depth k in the ledger).

Hence, running an MPC protocol π with the aid of a blockchain protocol simply means that the players exchange messages using the blockchain. Intuitively, a player is called *non-quick* if she always waits that the previous messages are confirmed on the blockchain before sending the next one. On the other hand, a *quick player* sends its next message by just looking at its current view of the blockchain (without pruning blocks). Apart from these changes, security is defined similarly as in the standard real-ideal world paradigm.

General-purpose MPC with quick players. Having motivated the problem of running MPC protocols on forking blockchains, we show a general compiler to obtain smart contracts that implement ideal multi-party functionalities retaining security in the presence of forks and allowing players to rush.[5] Our compiler starts from the observation that a stand-alone MPC protocol could be insecure when executed on a blockchain. To be concrete, a rewinding simulator of the MPC protocol can not be used to prove the security of the on-chain MPC protocol, since

---

[5]In this work all our positive result consist of on-chain protocols for secure computation that are stand-alone secure, with security preserved under sequential composition. The reason why we do not try to obtain universal composability is that existing notions of universal composability with a ledger [CGJ19] rely on non-forking ledger functionalities and therefore on non-quick players.

rewinding would have the unclear meaning of rewinding the blockchain. Moreover, we do not want to give control of the blockchain to the simulator (i.e., no control of the majority of the stake, of the computational power, and so on) since our result aims at being generic w.r.t. the type of blockchain used. Essentially, the simulator is going to incarnate just the honest players of the MPC protocol during the simulation. In order perform a simulation in the presence of a concurrently played blockchain protocol, (i.e., rewinding is not possible and the blockchain is generic and therefore not controlled by the simulator), we therefore require the initial protocol π received in input by the compiler to be universally composable secure. This guarantees the existence of a straight-line simulator and allows us to avoid simulators that "control" the blockchain[6], therefore allowing our results to be applicable to generic blockchains. Additionally, we require π to have only "digital" communication. The reason is that players when running the protocol on chain must produce messages that consists of bits only. Therefore an exchange of hardware (e.g., tamper-proof tokens, PUFs) in π can not be accepted.

In order to preserve security with respect to fork attacks, our compiler makes sure that, whenever an execution of the MPC protocol is repeated in multiple branches, each honest player protects herself from a fork attack by refusing to play again a message of the same execution of the protocol in case the blockchain shows a different prefix in the transcript of the execution. Specifically, if on one branch $\mathcal{B}_2$ there is a player that changes the message already played in a different branch $\mathcal{B}_1$, then each honest player that played already in $\mathcal{B}_1$ and is asked to play again on input a different prefix in $\mathcal{B}_2$ will abort the execution in $\mathcal{B}_2$. Clearly, this strategy forces a unique execution regardless of forks, and therefore security holds even in the presence of fully quick players.

Finally, notice that the original protocol might require private and authenticated channels. Since the entire traffic of our protocol will be redirected to the blockchain, we will use public-key encryption and digital signatures to emulate the private and authenticated channels. The first message of each player in the compiled protocol will therefore consist of a pair of public keys, one to receive encrypted messages and one to allow others to verify signatures of messages.

Fairness through penalties. In Andrychowicz et al. [ADMM14, ADMM16] and in Kumaresan et al. [BK14, KB14] it was shown how to add fairness (i.e., the adversary should be discouraged from learning the output before others to then decide whether the honest players should receive the output) through penalties. The idea is that a player should deposit some coins of the underlying cryptocurrency and the smart contract should return the coins back only in case the player completes correctly the execution of the protocol defined by the smart contract. Fairness is a very useful property and we want to upgrade our previous result obtaining fairness through penalties. Recall that we are planning to do so still admitting that the blockchain could fork and trying to obtain fast executions avoiding as much as possible to wait for confirmations of transactions. Our first idea to obtain fairness consists of running the underlying UC protocol π for a different functionality. For the aim of simplicity we are considering a single output, but it can be easily generalized to multiple outputs. Let π be the input protocol and π_{bc} be the protocol obtained by applying our generic compiler to π. The easier way that one can think to obtain fairness with penalties is to add a deposit in the first round of π_{bc} and wait that this first round is confirmed. Each player that behaves honestly will take back the deposit, meanwhile when adversarial parties send an incorrect message or they abort the execution by sending no messages at all, they will be penalized. We call π_f this protocol that tries to achieve fairness with penalties. In π_f each party generating an abort in the execution is considered adversarial and will be penalized. π_f does not achieve fairness with penalties for the

---

[6]Typically a simulator that controls the blockchain requires some specific assumptions on the blockchain like in [GG17] where only some restricted proof-of-stake blockchains were compatible with the simulation.

following reason: if an honest party P_i is forced to abort the execution in the case the adversary send him an incorrect message, P_i will be penalized, differently from what is stated in fairness with penalties. To obtain a fair with penalties protocol π_{fair}, we use the following technique, borrowed from [KB14, BK14]. Let's consider a protocol π' running with parties $P_1, \ldots, P_n$ for a functionality f' that, where each P_i holds input x_i, given the output $y \leftarrow f(x_1, \ldots, x_n)$, secret shares y into $(\sigma_1, \ldots, \sigma_n)$ with full threshold, generates a set of commitments $C = (\gamma_1, \ldots, \gamma_n)$ such that γ_i is the commitment of σ_i. Each player P_i obtains as an output the pair (C, σ_i)[7]. (i) We compile π' with our generic compiler, obtaining π'_{bc}. (ii) In our protocol π_{fair}, parties $P_1, \ldots, P_n$ first engage in π'_{bc}. After π'_{bc} ends, each P_i obtains the output (C, σ_i). Now, each P_i has a limited time t_1 to send his tuple C to a smart contract together with a payment of some deposit. (iv) If everyone sent the correct tuple C, each player P_i has another time shift t_2 to send their share σ_i of γ_i to receive back their deposit. Else, if after t_2, $(\sigma_1, \ldots, \sigma_n)$ are posted to the smart contract, each P_i can reconstruct the output by using all the collected shares. Else, players that have not opened their share within t_2, will be penalized since their coins will remain deposited forever.

We prove that fairness can be achieved if honest parties playing π_{fair} wait for confirmation only of step (ii). The reason why the above construction requires the confirmation of phase (ii) is that otherwise the adversary can try to generate an abort during the execution of π'_{bc} after learning the output of the entire protocol π_{fair} on a different branch.

Now, let's say that t is the time needed for transaction confirmation in the blockchain, and r the number of rounds of π'_{bc}, π_{fair} requires around $r + 2t$ blocks to complete the on-chain execution. This is due to the fact that confirmation is required only for the phase in which commitments are posted in the blockchain and for confirming the final reconstructed output. To maintain security of Kumaresan et al. protocols by blindly posting messages on-chain, the overall execution requires around $r \cdot t$ blocks to be successfully terminated.

Fair lottery with penalties and fully quick players. We analyze a variant of the generic attack described earlier to the well-known smart contract[8] of Andrychowicz *et al.* [ADMM14, ADMM16], for securely realizing multi-party lotteries. The main difference with the toy example from above is that in their work each player commits to a random value r_i between 1 and n (where n is the total number of participants to the lottery), and then, after all the commitments have been opened, the winner of the lottery is defined to be the player $w = r_1 + \ldots + r_n \pmod{n} + 1$. An appealing feature of this protocol is that it achieves fairness with penalties: If a malicious player aborts the protocol (e.g., it does not open the commitment before a certain time bound), then a previously deposited amount of coins is automatically transferred to the honest players (i.e., to those that correctly opened the commitment on time). Such a feature is particularly important in light of the negative result by Cleve [Cle86] on achieving fairness without honest majority.

We note that in the protocol of Andrychowicz *et al.* it is vital that players are non-quick and therefore post new transactions only after the previous ones are already confirmed on the blockchain. Indeed, in the presence of quick players, a simple variant of the attack described above would allow a malicious party to commit to a value r_i such that $\sum_i r_i \pmod{n} + 1 = i$, assuming that all players already opened the commitments on a minor branch of a fork.

As our second contribution, we go beyond the limits of the protocol of [ADMM14, ADMM16], and present a smart contract that implements the lottery functionality efficiently, indeed it remains secure even in the presence of quick players. Fairness with penalties can be added without affecting the efficiency of the protocol. In fact, the smart contract we design is more

---

[7]P_i implicitly receives also any decommitment information of γ_i.
[8]The protocol of [ADMM14, ADMM16] is based on Bitcoin, but this makes no difference for our attack.

general, in that it allows the players to establish a common, uniformly random, string (which in turn allows to run a lottery).

The main idea in our construction consists of combining unique signatures [Lys02] and random oracles as follows: first of all, players compute unique signatures on input the concatenation of the ordered sequence of their public keys. Notice that as long as at least one player is honest, we have a long string that no PPT player could predict when selecting his public key. Then this long string is given in input to a random oracle that therefore gives in output a uniformly distributed string. The simulator will program the random oracle therefore forcing in the simulation the same random string obtained in the ideal-world execution.

There is still an attack that can be mounted. Assume that in the presence of a fork the entire protocol is executed in a branch. The adversary could take advantage of the output in one branch to decide to play the same first round or a different first round in the other branch biasing successfully the distribution of the output. To circumvent this problem, we make executions in different branches completely independent by also passing a branch id as input to the unique signature evaluation procedure. As branch id we take the hash of the block containing the last deposit. Therefore, when a protocol is entirely run in a branch, we have that the two branch ids are different and thus there is no point in adaptively choosing the same or a different message in another branch. Indeed, in any case, the outputs in different branches will be completely independent. In order to deal with multiple executions of the real-world protocol in different branches, we will also have a simulator that will play multiple times in the ideal world. Since the output of the protocol is a random string, it can be then used in many applications, not only to run a multi-party lottery. Note that our protocol is around 50% more efficient the lottery of Andrychowicz *et al.*. Let's say that t is the number of blocks needed for transaction confirmation, then our coin tossing protocol can be run by using only $t + 1$ blocks, whereas *Andrychowicz et al.* requires $2 \cdot t$ blocks to be completed.

Notice that this result makes no use of finality of transactions on a blockchain (i.e., no player needs to know after how many blocks a transaction can be considered permanent). The protocol therefore can be run in the presence of fully quick players, and is therefore very efficient.

We stress that we consider the adversary as a player that tries to exploit the existence of forks in order to bias the output of the smart contract. We are not modelling the adversary of the smart contract as a player that has control over forks, deciding which branch will eventually be discarded and which one will become permanently part of the blockchain. Obviously, a powerful adversary that has control over the forks can always play the protocol on each branch to then select the one that produced the output that she likes the most. This is unavoidable when there is no use of finality of transactions. Nevertheless, notice that in many cases this is not a problem. Indeed think of the need of establishing a random string to then use it as first round of a statistically hiding commitment scheme or as common reference string for a non-interactive zero-knowledge proof. In such scenarios the adversary can freely select a random string from any polynomially large set of randomly sampled strings without compromising any security. In other cases like playing bingo, the fact that the adversary can decide the string out of several candidates can be an issue.

1.5 Related Work

Following [ADMM14, ADMM16], several other works focus on achieving fairness with penalties for different applications of interest, including lotteries [BK14], decentralized poker [KMB15, BKM17], and general-purpose computation [BK14, KMS$^+$16, KB16, KVV16]. In particular, the line of works by Kumaresan *et al.* relies on an elegant paradigm working in two phases: During the first phase, the players run an MPC protocol to obtain the output in hidden form

(e.g., a secret sharing of the output); since the output is hidden, such a protocol can be executed off chain, as malicious aborts do not violate fairness. During the second phase, the output is then reconstructed in a fair manner on chain.

Unfortunately, the security of this paradigm in the presence of quick players is difficult to assess, as it relies on intermediate ideal functionalities (such as the "claim-or-refund" and "multi-lock" functionality [BK14, KB14]) that, while they can be implemented using Bitcoin or Ethereum, offer a-priori no security guarantee in the presence of blockchain forks. Moreover, known results about designing protocols in a hybrid model that allows to make calls to a functionality are applicable only to the classical setting where multiple executions of the same instance of the protocol due to forks are not possible. Also note that performing a large part of the computation off chain hinders one of the main advantages of blockchain-aided MPC (i.e., public verifiability of the entire process). Our results, in contrast, consider MPC protocols run completely on-chain through smart contracts.

A different line of works, shows how to perform MPC in the presence of an abstract transaction ledger [KZZ16, GG17, BMTZ17, SSV19, CGJ19], of which Bitcoin and Ethereum are possible implementations. However, such an idealized ledger does not account for the possibility of forks, thus (implicitly) meaning that the players using it are modeled as non-quick.

Our second contribution is a protocol to generate a randomness. it is known that there exist protocols suited for blockchains generating random values. A well known implementation is RANDAO [RT]. The smart contract introduced in RANDAO is similar to what we describe in Fig. 7 as a smart contract implementation of the Andrychowicz *et al.* lottery protocol [ADMM16]. The smart contract is divided in three phases: (i) the first phase collects valid hashes calculated as $sha3(s_i)$ and a fixed deposit d, where s_i is the secret value belonging to party P_i, (ii) the second phase collects valid s_i for each party P_i thus refunding back d, where an s_i is valid if only if $sha3(s_i)$ is equal to the value sent by P_i in the first phase, (iii) if each P_i provided the correct s_i the third phase computes a random number from all the values $(s_1, s_2, \ldots)$, otherwise the deposit d of each P_j who did not provided his s_j is confiscated and redistributed among the other players. The function used to compute the random number can be customized and can be chosen by the initializer of the smart contract. As we show in Sec. 5.2 even this smart contract is subject to attacks if the parties does not wait for the confirmation of at least the first phase of the protocol. On the contrary, our parallel coin tossing protocol described in Sec. 5.3 is secure even if the parties does not wait and block confirmation.

2 Preliminaries

2.1 Notation

Given an integer n, we let $[n] = \{1, \ldots, n\}$. If x is a string, we denote its length by $|x|$; if $\mathcal{X}$ is a set, $|\mathcal{X}|$ is the number of elements in $\mathcal{X}$. When x is chosen randomly in $\mathcal{X}$, we write $x \leftarrow_\$ \mathcal{X}$. When A is an algorithm, we write $y \leftarrow_\$ \mathsf{A}(x)$ to denote a run of A on input x and output y; if A is randomized, then y is a random variable and $\mathsf{A}(x; \omega)$ denotes a run of A on input x and random coins $\omega \in \{0, 1\}^*$.

Throughout the paper, we denote the security parameter by $\lambda \in \mathbb{N}$. A function $\nu(\lambda)$ is negligible in λ (or just negligible) if it decreases faster than the inverse of every polynomial in λ, i.e. $\nu(\lambda) \in O(1/p(\lambda))$ for every positive polynomial $p(\cdot)$. A machine is said to be probabilistic polynomial time (PPT) if it is randomized, and its number of steps is polynomial in the security parameter.

For a random variable $\mathbf{X}$, we write $\mathbb{P}[\mathbf{X} = x]$ for the probability that $\mathbf{X}$ takes a particular value x in its domain. A distribution ensemble $\mathbf{X} = \{\mathbf{X}(\lambda)\}_{\lambda \in \mathbb{N}}$ is an infinite sequence of random

variables indexed security parameter $\lambda \in \mathbb{N}$. Two distribution ensembles $\mathbf{X} = \{\mathbf{X}(\lambda)\}_{\lambda \in \mathbb{N}}$ and $\mathbf{Y} = \{\mathbf{Y}(\lambda)\}_{\lambda \in \mathbb{N}}$ are said to be *computationally indistinguishable*, denoted $\mathbf{X} \approx_c \mathbf{Y}$ if for every non-uniform PPT algorithm D there exists a negligible function $\nu(\cdot)$ such that:

$$\left| \mathbb{P}\left[\mathsf{D}(\mathbf{X}(\lambda)) = 1\right] - \mathbb{P}\left[\mathsf{D}(\mathbf{Y}(\lambda)) = 1\right] \right| \leq \nu(\lambda).$$

When the above equation holds for all (even unbounded) distinguishers D, we say that $\mathbf{X}$ and $\mathbf{Y}$ are statistically close, denoted $\mathbf{X} \approx_s \mathbf{Y}$.

2.2   Standard Primitives

Public-key encryption.   A public-key encryption (PKE) scheme is a tuple of polynomial-time algorithms $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ specified as follows. (i) The randomized algorithm Gen takes as input the security parameter, and outputs a pair of keys (pk, sk); (ii) The randomized algorithm Enc takes as input a public key pk and a message $m \in \mathcal{M}$, and outputs a ciphertext c; (iii) The deterministic algorithm Dec takes as input a secret key sk and a ciphertext c, and outputs a value in $\mathcal{M} \cup \{\perp\}$ (where $\perp$ denotes decryption error). Correctness says that for every key $\lambda \in \mathbb{N}$, every (pk, sk) in the support of $\mathsf{Gen}(1^\lambda)$, and every message $m \in \mathcal{M}$, it holds that $\mathsf{Dec}(sk, \mathsf{Enc}(pk, m)) = m$ with probability one over the randomness of Enc.

Definition 1 (Semantic security)**.** We say that $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ satisfies semantic security if for all PPT attackers $\mathsf{A} := (\mathsf{A}_0, \mathsf{A}_1)$ there exists a negligible function $\nu(\cdot)$ such that:

$$\left| \mathbb{P}\left[b' = b : \begin{array}{l} (pk, sk) \leftarrow_\$ \mathsf{Gen}(1^\lambda); (m_0, m_1, z) \leftarrow_\$ \mathsf{A}_0(pk) \\ b \leftarrow_\$ \{0,1\}; c \leftarrow_\$ \mathsf{Enc}(pk, m_b); b' \leftarrow_\$ \mathsf{A}_1(z, c) \end{array} \right] - \frac{1}{2} \right| \leq \nu(\lambda).$$

Signature schemes.   A signature scheme is a tuple of polynomial-time algorithms $(\mathsf{Gen}, \mathsf{Sign}, \mathsf{Verify})$ specified as follows. (i) The randomized algorithm Gen takes as input the security parameter and outputs a secret key sk together with a public verification key pk; (ii) The deterministic algorithm Sign takes as input the secret key sk and a message $x \in \{0,1\}^*$ and outputs a signature y; (iii) The randomized algorithm Verify takes as an input the verification key pk, a message/signature pair (x, y) and outputs a decision bit.

   Correctness says that for all $\lambda \in \mathbb{N}$, for all $(pk, sk) \in \mathsf{Gen}(1^\lambda)$, and for all $x \in \{0,1\}^*$ it holds that $\mathsf{Verify}(pk, x, \mathsf{Sign}(sk, x)) = 1$ (with probability one over the coin tosses of Verify).

   We will need so called *unique* signature schemes, which satisfy two properties known as uniqueness and unforgeability as defined below.

Definition 2 (Uniqueness)**.** For every pk, x, y_0, y_1 with $y_0 \neq y_1$ there exists a negligible function $\nu(\cdot)$ such that the following holds for either $i = 0$ or $i = 1$:

$$\mathbb{P}\left[\mathsf{Verify}(pk, x, y_i) = 1\right] \leq \nu(\lambda).$$

   In words, for every string pk and every x, there exists at most one value y that is a accepting signature of x.

Definition 3 (Unforgeability)**.** For all PPT *valid* attackers A there exists a negligible function $\nu(\cdot)$ such that:

$$\mathbb{P}\left[\mathsf{Sign}(sk, x) = y : \begin{array}{l} (pk, sk) \leftarrow_\$ \mathsf{Gen}(1^\lambda) \\ (x, y) \leftarrow_\$ \mathsf{A}^{\mathsf{Sign}(sk, \cdot)}(pk) \end{array} \right] \leq \nu(\lambda),$$

where A is called *valid* if it never queries m to its oracle.

   Unique signatures are sometimes also known under the name of verifiable unpredictable functions, and exist based on a variety of assumptions [BR96, MRV99, Lys02, DY05].

Commitment schemes. A non-interactive commitment Commit is a PPT algorithm taking as input a message $m \in \{0,1\}^\ell$, and outputting a commitment $\gamma = \mathsf{Commit}(m; \delta)$, where $\delta \in \{0,1\}^*$ is the randomness used to generate the commitment. The pair (m, δ) is called the opening.

Intuitively, a secure commitment satisfies two properties called binding and hiding. The first property says that it is hard to open a commitment in two different ways. The second property says that a commitment hides the underlying message.

Definition 4 (Binding). We say that a non-interactive commitment Commit is *perfectly binding* if pairs (m_0, δ_0), (m_1, δ_1) such that $m_0 \neq m_1$ and $\mathsf{Commit}(m_0; \delta_0) = \mathsf{Commit}(m_1; \delta_1)$ do not exist.

Definition 5 (Hiding). We say that a non-interactive commitment Commit is *computationally hiding* if for all non-uniform PPT adversaries A the following quantity is negligible

$$\left| \mathbb{P}\left[\mathsf{A}^{\mathsf{LR}(0,\cdot,\cdot)}(1^\lambda) = 1 \right] - \mathbb{P}\left[\mathsf{A}^{\mathsf{LR}(1,\cdot,\cdot)}(1^\lambda) = 1 \right] \right|,$$

where the oracle $\mathsf{LR}(b, \cdot, \cdot)$ with hard-wired $b \in \{0,1\}$ takes as input pairs of messages $m_0, m_1 \in \{0,1\}^\ell$, and outputs $\mathsf{Commit}(m_b)$.

Secret Sharing Schemes. An n-party secret sharing scheme $(\mathsf{Share}, \mathsf{Recon})$ is a pair of poly-time algorithms specified as follows. (i) The randomized algorithm Share takes as input a message $m \in \mathcal{M}$ and outputs n shares $\sigma = (\sigma_1, \ldots, \sigma_n) \in \mathcal{S}_1 \times \cdots \times \mathcal{S}_n$; (ii) The deterministic algorithm Recon takes as input a subset of the shares, say $\sigma_\mathcal{I}$ with $\mathcal{I} \subseteq [n]$, and outputs a value in $\mathcal{M} \cup \{\bot\}$.

Definition 6 (Threshold secret sharing). Let $n \in \mathbb{N}$. For any $t \leq n$, we say that $(\mathsf{Share}, \mathsf{Recon})$ is an (t, n)-secret sharing scheme if it satisfies the following properties.
- **Correctness**: For any message $m \in \mathcal{M}$, and for any $\mathcal{I} \subseteq [n]$ such that $|\mathcal{I}| \geq t$, we have that $\mathsf{Recon}(\mathsf{Share}(m)_\mathcal{I}) = m$ with probability one over the randomness of Share.
- **Privacy**: For any pair of messages $m_0, m_1 \in \mathcal{M}$, and for any $\mathcal{U} \subset [n]$ such that $|\mathcal{U}| < t$, we have that
$$\{\mathsf{Share}(1^\lambda, m_0)_\mathcal{U}\}_{\lambda \in \mathbb{N}} \approx_c \{\mathsf{Share}(1^\lambda, m_1)_\mathcal{U}\}_{\lambda \in \mathbb{N}}.$$

2.3 Multi-Party Computation

We recall standard notion of UC-security for multi-party computation (MPC). Let $f : (\{0,1\}^*)^n \rightarrow (\{0,1\}^*)^n$ be a function, and consider n players $\mathsf{P}_1, \ldots, \mathsf{P}_n$ executing a protocol π for computing f. Our default network model consists of the players interacting in synchronous rounds via private and authenticated point-to-point channels.

Intuitively, the security of π is formalized by comparing its execution in the real world (where an attacker may corrupt a subset of the players) with the ideal execution in which a trusted party computes the function f on behalf of the players.

The real model. In the real world, the protocol π is run in the presence of an adversary A coordinated by a non-uniform environment $\mathsf{Z} = \{\mathsf{Z}_\lambda\}_{\lambda \in \mathbb{N}}$. At the outset, Z chooses the inputs $(1^\lambda, x_i)$ for each player P_i, and gives $\mathcal{I}$, $\{x_i\}_{i \in \mathcal{I}}$ and z to A, where $\mathcal{I} \subseteq [n]$ represents the set of corrupted players and z is some auxiliary input. For simplicity, we only consider static corruptions (i.e., the environment decides who is corrupt at the beginning of the protocol). The parties then start running π, with the honest players P_i behaving as prescribed in the protocol (using input x_i), and with malicious parties behaving arbitrarily (directed by A). The attacker may delay sending the messages of the corrupted parties in any given round until after the

honest parties send their messages in that round; thus, for every r, the round-r messages of the corrupted parties may depend on the round-r messages of the honest parties.

At some point, A gives to Z an arbitrary function of its view, and Z additionally receives the outputs of the honest parties and must output a bit. We denote by $\mathbf{REAL}_{\pi,\mathsf{A},\mathsf{Z}}(\lambda)$ the random variable corresponding to Z's guess.

The ideal model. In the ideal world, a trusted third party evaluates the function f on behalf of a set of dummy players $(\mathsf{P}_i)_{i\in[n]}$. As in the real setting, Z chooses the inputs $(1^\lambda, x_i)$ for each honest player P_i, and gives $\mathcal{I}$, $\{x_i\}_{i\in\mathcal{I}}$ and z to the ideal adversary S, corrupting the dummy parties $(\mathsf{P}_i)_{i\in\mathcal{I}}$. Hence, honest parties send their input $x_i' = x_i$ to the trusted party, whereas the parties controlled by S might send an arbitrary input x_i'. The trusted party computes $(y_1, \ldots, y_n) = f(x_1', \ldots, x_n')$, and sends y_i to P_i. Finally, S gives to Z an arbitrary function of its view, and Z additionally receives the outputs of the honest parties and must output a bit. We denote by $\mathbf{IDEAL}_{f,\mathsf{S},\mathsf{Z}}(\lambda)$ the random variable corresponding to Z's guess.

The above specification of the ideal model automatically implies *fairness* (i.e., corrupted parties get the output if and only if honest parties do as well). Unfortunately, as shown by Cleve [Cle86], such a strong guarantee is impossible to achieve for some functionalities without assuming honest majority. For this reason, we also consider a weaker flavor of the ideal model yielding a middle-ground notion known as security with aborts, which is possible to achieve even in the presence of honest minority. Let $\mathcal{H} := [n] \setminus \mathcal{I}$. The only difference with the above specification is that the trusted party at first forwards only the outputs $\{y_i\}_{i\in\mathcal{I}}$ to the ideal adversary S. Hence, S might send either a message $(\texttt{continue}, \mathcal{H}')$ or $\texttt{abort}$ to the trusted party. In the former case, all the honest parties in $\mathcal{H}'$ are given their output y_i whereas the honest parties in $\mathcal{H} \setminus \mathcal{H}'$ receive an abort symbol $\bot$. In the latter case, all honest parties receive $\bot$. We denote by $\mathbf{IDEAL}_{f_\bot,\mathsf{S},\mathsf{Z}}(\lambda)$ the random variable corresponding to Z's final guess.

The definition. We are now ready to define security.

Definition 7 (UC-Secure MPC). Let π be an n-party protocol for computing a function $f : (\{0,1\}^*)^n \to (\{0,1\}^*)^n$. We say that π t-securely UC-realizes f in the presence of malicious adversaries if such that for every PPT adversary A there exists a PPT simulator S such that for every non-uniform PPT environment Z corrupting at most t parties the following holds:

$$\{\mathbf{REAL}_{\pi,\mathsf{A},\mathsf{Z}}(\lambda)\}_{\lambda\in\mathbb{N}} \approx_c \{\mathbf{IDEAL}_{f,\mathsf{S},\mathsf{Z}}(\lambda)\}_{\lambda\in\mathbb{N}}.$$

When replacing $\mathbf{IDEAL}_{f,\mathsf{S},\mathsf{Z}}(\lambda)$ with $\mathbf{IDEAL}_{f_\bot,\mathsf{S},\mathsf{Z}}(\lambda)$ we say that π t-securely computes f *with aborts* in the presence of malicious adversaries.

2.4 A Blockchain Model

Below, we describe verbatim the blockchain model of [GG17] (which in turn builds on [PSS17, GKL15]). A blockchain protocol Γ consists of the following algorithms:

- $\mathsf{UpdateState}(1^\lambda)$: It is a stateful algorithm that take as input a security parameter $\lambda \in \mathbb{N}$, and maintains a local state $st \in \{0,1\}^*$ which essentially consists of the entire blockchain (i.e., the sequence of minted blocks).
- $\mathsf{GetRecords}(1^\lambda, st)$: It takes as input the security parameter and a state $st \in \{0,1\}^*$. It outputs the longest *ordered* sequence of valid blocks (or simply blockchain) $\mathcal{B} = (\beta_1, \beta_2, \ldots)$ contained in the state variable, where each block β in the chain itself contains an *unordered* sequence of records/messages $(m_1, m_2, \ldots)$.

- Broadcast($1^\lambda, m$): It takes as input the security parameter and a message $m \in \{0,1\}^*$, and broadcasts the message over the network to all nodes executing the blockchain protocol. It does not give any output.

The blockchain protocol is also parameterized by a validity predicate V that captures the semantics of any particular blockchain application. The validity predicate takes as input a sequence of blocks $\mathcal{B}$ and outputs a bit, where the value 1 certifies the validity of the blockchain $\mathcal{B}$. Since V is immaterial for our purposes, in what follows we simply omit it.

Blockchain execution. Each participant in the protocol runs UpdateState to keep track of the latest blockchain state. This corresponds to listening on the broadcast network for messages from other nodes. GetRecords is used to extract an ordered sequence of blocks encoded in the blockchain state variable, which is considered as the common public ledger among all the nodes. Finally, Broadcast is used by a party when it wants to post a new message on the blockchain; such messages are accepted by the blockchain protocol only if they satisfy the validity predicate given the current state.

The execution of a blockchain protocol $\Gamma = (\mathsf{UpdateState}, \mathsf{GetRecords}, \mathsf{Broadcast})$ is directed by an environment $\mathsf{Z}(1^\lambda)$ which activates the parties as either honest or corrupt, and is also responsible for providing inputs/records to all parties in each round. All the corrupted parties are controlled by the adversary A, which is also responsible for delivery of all network messages. Honest parties start by executing UpdateState on input 1^λ, with an empty local state $st = \varepsilon$. Then, the protocol execution proceeds in rounds that model times steps, as detailed below.

- In round $r \in \mathbb{N}$, each honest player P_i potentially receives messages from Z, and incoming network messages (delivered by A). It may then perform any computation, broadcast a message to all other players (which will be delivered by the adversary as explained below), and update its local state st_i. It could also attempt to add a new block to its chain i.e., run the mining procedure.
- The attacker A is responsible to deliver all messages sent by parties (honest or corrupted) to all other parties. The adversary cannot modify the content of messages broadcast by honest players, but it may delay or reorder the delivery of a message as long as it eventually delivers all messages within a certain time limit.
- At any point, Z can communicate with A or access $\mathsf{GetRecords}(1^\lambda, st_i)$ where st_i is the local state of player P_i.

With the notation $\mathcal{B} \preceq \mathcal{B}'$, we denote that the blockchain $\mathcal{B}$ is a prefix of $\mathcal{B}'$. We also let $\mathcal{B}^{\lceil k}$ be the chain resulting from pruning the last k blocks in $\mathcal{B}$. Let $\mathbf{EXEC}_{\Gamma,\mathsf{A},\mathsf{Z}}(\lambda)$ be the random variable denoting the joint view of all parties in the execution of protocol Γ with adversary A, and environment Z. Note that this view fully determines the execution.

Blockchain properties. We now define two natural guarantees that are respected by an ideal ledger. The first property, called *consistency*, intuitively states that the view of the blockchain obtained by different players is identical up to pruning a certain number of blocks from the top of the chain. Let $\mathsf{Consistent}^k(\cdot)$ be the predicate that returns 1 iff for all rounds $r \leq \tilde{r}$, and all parties $\mathsf{P}_i, \mathsf{P}_j$ (potentially the same) such that P_i is honest at round r with blockchain $\mathcal{B}$ and player P_j is honest at round $\tilde{r}$ with blockchain $\tilde{\mathcal{B}}$, we have that $\mathcal{B}^{\lceil k} \preceq \tilde{\mathcal{B}}$.

Definition 8 (Chain consistency). *A blockchain protocol Γ satisfies $k(\cdot)$-consistency with adversary A and environment Z, if there exists a negligible function $\nu(\cdot)$ such that for every $\bar{k} > k(\lambda)$, the following holds:*

$$\mathbb{P}\left[\mathsf{Consistent}^{\bar{k}}(view) = 1 : \ view \leftarrow^{\$} \mathbf{EXEC}_{\Gamma,\mathsf{A},\mathsf{Z}}(\lambda)\right] \geq 1 - \nu(\lambda).$$

We note that previous work considered an even stronger property, called *persistence*, stipulating that if some honest player reports a message m at depth k in its local ledger, then m will be always reported in the same position and equal or more depth by all honest parties. We omit a formal definition, as this property is not required for our purposes.

The second property, called *liveness*, intuitively says that if all honest parties attempt to broadcast a message m, then after w rounds, an honest party will see m at depth k in the ledger. Let $\mathsf{Live}^k(\cdot, w)$ be the predicate that returns 1 iff for any w consecutive rounds $r, \ldots, r + w$ there exists some round $r' \in [r, r + w]$ and index $i \in [n]$ such that: (1) P_i is honest and received a message m at round r, and (2) for every player P_j that is honest at $r + w$ with blockchain $\mathcal{B}$, it holds that $m \in \mathcal{B}^{\lceil k}$.

Definition 9 (Liveness). A blockchain protocol Γ satisfies $(w(\cdot), k(\cdot))$-liveness with adversary A and environment Z, if there exists a negligible function $\nu(\cdot)$ such that for every $\bar{w} \geq w(\lambda)$ the following holds:

$$\mathbb{P}\left[\mathsf{Live}^k(view, \bar{w}) = 1 : \; view \leftarrow_\$ \mathbf{EXEC}_{\Gamma, \mathsf{A}, \mathsf{Z}}(\lambda)\right] \geq 1 - \nu(\lambda).$$

2.5 Threat Model

We assume that the adversary is computationally bounded, and when there is a fork in the blockchain, we pragmatically assume that the adversary has no power to choose which branch will be confirmed. We point out that if one would like to consider a very strong adversary with even 49% of the computational power of the network, then clearly our assumption does not hold. However, we stress that with such an adversary even the 6-block rule in Bitcoin does not make much sense. To guarantee that a delicate transaction (i.e., the coinbase transaction) is confirmed with a strong enough adversary, up to 144 blocks are necessary in Bitcoin [BW], meaning 1 day to communicate even a single protocol message. Therefore if one would like to consider such strong adversaries even a protocol requiring only one confirmation becomes totally impractical.

Our proposed protocols work in the dishonest majority setting.

Since the overall protocol execution is run on-chain, in our scenarios is not clear what means to fix the identity of the parties "before" the execution of the protocol. We say that a party P_i participates to a protocol execution only after he publishes the first protocol message.

Therefore, each protocol we are going to present in this paper falls in the following class: at the beginning each player P_i sends the first protocol message to "register" himself to the protocol, and waits until a fixed time t that the other players post the first message. If, after time t has passed, all parties sent the first message, the protocol continues as expected; Otherwise, we have two cases:

- (i) If the protocol is suited to run also with less the n parties, the registered parties continue a normal execution.
- (ii) If the protocol works only with n parties (or more), the execution is aborted by the registered parties.

We notice that our PCT protocol can be executed even if only two parties decide to participate in the protocol. This means that n does not need to be fixed beforehand. Moreover, registered parties are not incentivized to abort the protocol by not sending the second message in PCT due to financial compensation (parties must send a collateral deposit together with the first message). This makes DoS attacks in which the attacker aborts the protocol multiple times (making honest parties waste time and money) economically infeasible.

Our generic compiler, instead, depends on the class of protocols it is applied to. Firstly, the underlying protocol could need n or more parties to be correctly executed. This issue is not

directly solvable on the financial side, since our compiler outputs a protocol achieving *fair output delivery with penalties*, meaning that abort is still possible during the computation phase, but a financial penalty is given to the parties aborting in the output phase (i.e. when not every honest party receives the output). In fact, differently from the compiler, PCT achieves *fair computation with penalties*, meaning that an aborting player can be penalized at each step of computation.

Fair computation is not easily achievable for a generic compiler: as it can be seen in Section 4, even taking as an input a protocol achieving identifiable abort, it cannot be compiled to a protocol achieving fair computation with penalties that is secure in the presence of quick players. This is due to the fact that the compiler will abort when an attacker against quick players exploits fork to send different messages, thus provoking an abort during the computation phase. In this case, financial compensation can be performed only by waiting for the block confirmation of each message, thus boiling down to the abort identifiability of the underlying protocol. Unfortunately, this removes the main advantage of our compiler, that is, being quick.

3 Running MPC on Forking Blockchains

In this section, we formalize different ways how to run an MPC protocol with the aid of a blockchain. In §3.1 we specify what it means to run an MPC protocol on the blockchain both in the presence of quick and non-quick players. The security definition appears in §3.2.

3.1 Blockchain-Aided MPC

Next, we define what it means to run an n-party protocol π for securely computing some function $f : (\{0,1\}^*)^n \to (\{0,1\}^*)^n$ over a blockchain protocol Γ.

Intuitively, running π on Γ simply means that the players write the protocol's messages on the blockchain instead of using point-to-point connections. However, since the blockchain may fork, the protocol's participants have to choose how to manage possibly unconfirmed blocks that are part of the current chain. Looking ahead, this choice will have impact both on the efficiency and on the security of the protocol execution. In particular, we distinguish between *quick* and *non-quick* players as formalized below.

Non-quick execution. Roughly speaking, a player is said to be *non-quick* if it always decides its next message by looking at the transcript of the protocol that is obtained by pruning the last k blocks of the blockchain, where k is the parameter for the consistency property of the underlying blockchain.

Definition 10 (Non-quick player). Let $\Gamma = (\mathsf{UpdateState}, \mathsf{GetRecords}, \mathsf{Broadcast})$ be a blockchain protocol with k-consistency. A player P_i is said to be *non-quick* if it behaves as follows:
- Initialize $\tau_i^{(0)} := \varepsilon$, $st_i := \varepsilon$ and $r_i := 0$.
- Run the following loop:
    - Update the state st_i by running $\mathsf{UpdateState}(1^\lambda)$, and retrieve $\mathcal{B}_i \leftarrow_{\$} \mathsf{GetRecords}(st_i)$ until the partial transcript $\tau^{(r_i)}$ is contained in $\mathcal{B}_i^{\lceil k}$.
    - If the protocol is over (i.e., the transcript $\tau^{(r_i)}$ is sufficient for determining the output), output the value y_i as a function of $\tau_i^{(r_i)}$ and terminate.
    - Else, compute the next protocol message $m_i^{(r_i+1)}$, invoke $\mathsf{Broadcast}(m_i^{(r_i+1)})$, and set $r_i := r_i + 1$.

Quick execution. On the other hand, a player is *quick* if it decides and broadcasts its next message by looking at the latest version of the blockchain (i.e., without pruning blocks). Since the consistency property does not hold for the last k blocks, quick players may retrieve different protocol's transcripts as the protocol proceeds. In particular, it may happen that at a given time step party P_i reads from the blockchain a partial transcript $\tau^{(\tilde{r})}$, whereas at a later time step the same player reads $\tau^{(\tilde{r}')}$ for some $\tilde{r}' < \tilde{r}$. This is due to the fact that some of the messages contained in $\tau^{(\tilde{r})}$ may end up in unconfirmed blocks, and thus be discarded.

Definition 11 (Quick execution). Let $\Gamma = (\mathsf{UpdateState}, \mathsf{GetRecords}, \mathsf{Broadcast})$ be a blockchain protocol with k-consistency. A player P_i is said to be *quick* if it behaves as follows:
- Initialize $\tau_i^{(0)} := \varepsilon$ and $st_i := \varepsilon$.
- Run the following loop:
    - Update the state st_i by running $\mathsf{UpdateState}(1^\lambda)$, and let $\mathcal{B}_i \leftarrow_\$ \mathsf{GetRecords}(st_i)$.
    - Let $\tilde{r} \geq 0$ be the maximum value such that the partial transcript $\tau^{(\tilde{r})} \in \mathcal{B}_i$.
    - If the protocol is over (i.e., the transcript $\tau^{(\tilde{r})}$ is sufficient for determining the output), output the value y_i as a function of $\tau^{(\tilde{r})}$ and terminate.
    - Else, compute the next protocol message $m_i^{(\tilde{r}+1)}$ and invoke $\mathsf{Broadcast}(m_i^{(\tilde{r}+1)})$.

More generally, we call φ-quick a player that is non-quick until a partial transcript $\tau^{(\varphi)}$ is at least k blocks deep in the blockchain, and afterwards it starts being quick. We sometimes call φ the *finality parameter*. Note that a 0-quick player is identical to a quick player, whereas an ∞-quick player is identical to a non-quick player. We call (χ, φ)-quick a player that is quick for the first χ rounds, and then behaves like a φ-quick player.

3.2 Security in the Presence of Quick Players

We can now define security of MPC protocols running on the blockchain. As in the standard setting, the definition compares a protocol execution in the real world with one in the ideal setting where a trusted party is made available. The main difference with the standard definition is that the attacker A is given black-box access to the algorithms in Γ, which it can use arbitrarily. The simulator is not allowed to control the blockchain (i.e. it must simulate the view of the adversary while invoking the algorithms in Γ on behalf of the honest players).

The real model: This is the execution of π on Γ, where the honest players are φ-quick. As usual, the adversary A is coordinated by a non-uniform distinguisher D. At the outset, D chooses the inputs $(1^\lambda, x_i)$ for each player P_i, and gives $\mathcal{I}, \{x_i\}_{i \in \mathcal{I}}$ and z to A, where $\mathcal{I} \subseteq [n]$ represents the set of corrupted players and z is some auxiliary input. The parties then start running π on Γ, with the honest players P_i being φ-quick and behaving as prescribed in π (using input x_i), and with malicious parties behaving arbitrarily (directed by A). At some point, A gives to D an arbitrary function of its view; note that the latter includes the view generated via $\mathbf{EXEC}_{\Gamma,\mathsf{A},\mathsf{D}}(\lambda)$ in the blockchain protocol. Finally, D receives the outputs of the honest parties and must output a bit. We denote by $\mathbf{REAL}_{\pi,\mathsf{A},\mathsf{D}}^{\Gamma,\varphi}(\lambda)$ the random variable corresponding to D's guess.

The ideal model: This is identical to the ideal model for standard MPC (App. 2.3), with the only difference that the simulator S is also responsible for simulating the attacker's view corresponding to the interaction of the honest players with the blockchain. The latter is achieved using the algorithms of the underlying blockchain protocol Γ. We denote by $\mathbf{IDEAL}_{f,\mathsf{S},\mathsf{D}}^{\Gamma}(\lambda)$ and $\mathbf{IDEAL}_{f_\perp,\mathsf{S},\mathsf{D}}^{\Gamma}(\lambda)$ the random variable corresponding to D's guess in the ideal world, where the latter is for the case of security with aborts.

Definition 12 (Secure MPC in the presence of quick players). Let π be an n-party protocol run over a blockchain protocol Γ. We say that π t-securely computes f in the presence of φ-quick

players and malicious adversaries if for every PPT adversary A there exists a PPT simulator S such that for every non-uniform PPT distinguisher D corrupting at most t parties the following holds:

$$\left\{\mathbf{REAL}_{\pi,\mathsf{A},\mathsf{D}}^{\Gamma,\varphi}(\lambda)\right\}_{\lambda\in\mathbb{N}} \approx_c \left\{\mathbf{IDEAL}_{f,\mathsf{S},\mathsf{D}}^{\Gamma}(\lambda)\right\}_{\lambda\in\mathbb{N}}.$$

When replacing $\mathbf{IDEAL}_{f,\mathsf{S},\mathsf{D}}^{\Gamma}(\lambda)$ with $\mathbf{IDEAL}_{f_\perp,\mathsf{S},\mathsf{D}}^{\Gamma}(\lambda)$ we say that π t-securely computes f *with aborts* in the presence of φ-quick players and malicious adversaries.

Remark 1 (On $\varphi = \infty$). *One may think that every protocol π that t-securely computes f (with or without aborts) in the presence of malicious adversaries, must t-securely compute f (with or without aborts) in the presence of ∞-quick (i.e., non-quick) players and malicious adversaries.*

Remark 2 (On $\varphi = 0$). *Note that when the players are fully quick (i.e., $\varphi = 0$), the adversary's view in the real world may include multiple executions of the original protocol π (upon the same inputs chosen by the distinguisher). This view may not be possible to simulate in the ideal world, where the simulator can invoke the ideal functionality f only once.*

For this reason, whenever $\varphi = 0$, we implicitly assume that the simulator is allowed to query the ideal functionality f multiple times. Note that this yields a meaningful security guarantee only for certain functionalities f, similarly to the setting of resettably secure computation [GS09].

Remark 3 (On the power of the adversary). *We stress that we assume that the adversary of the MPC protocol has no impact on the execution of the consensus protocol of the underlying blockchain. Note that if we would instead assume that the adversary of the MPC protocol also creates new branches and/or contributes in deciding which branch of a fork is eventually confirmed on the blockchain then he can have an unfair advantage. Indeed the adversary can start more branches when he does not like the output computed in a branch, and/or can decide which output among the various outputs appearing in different branches should be confirmed on the blockchain. Obviously the above unfair advantages are unavoidable and our protocol is still secure by introducing the unavoidable real-world attack into the ideal world, similarly to the classical fairness issue resolved through aborts in the ideal world.*

Remark 4 (On public verifiability). *We notice that any on-chain MPC protocol with quick players admits the case where a honest player complete her execution computing an output that does not necessarily correspond to the transcript that others later on will see on the blockchain. In other words, the local output computed by players could not match the publicly verifiable execution that remains visible on the blockchain. The reason why public verifiability could fail is that an execution of the protocol could be entirely contained in a branch of a fork that will not become permanent in the blockchain. The above issue is intrinsic in all protocols played on-chain in the presence of forks and quick players. An obvious solution for a honest player consists of waiting that the last message of the protocol is confirmed on the blockchain and only after that the computation ends returning the computed output.*

Random oracle model. Our result in §5.3 are secure in the random oracle model (ROM). The definition remains the same, except that each player in the real world has now access to a truly random hash function Hash chosen at the beginning of the experiment. The simulator of the ideal world can program the random oracle.

4   A Simple Compiler

In this section, we propose and analyze a simple transformation that allows to run any MPC protocol safely on the blockchain, even when the players are quick. The description of our

compiler appears in §4.1, while in §4.2 we analyze its security. Finally, in §4.3, we discuss how to extend our generic transformation in order to achieve fairness with penalties, as long as the players start being quick after the confirmation of the first round.

4.1 Compiler Description

Intuitively our transformation proceeds as follows. Our starting point is any MPC protocol π UC-securely computing an n-party functionality $f : (\{0,1\}^*)^n \to (\{0,1\}^*)^n$ in the presence of malicious players (with aborts). Hence, the honest players fix the random tape for running π and simply execute protocol π by broadcasting their messages on the blockchain. Furthermore, each honest player P_i keeps track of the longest protocol transcript α_i generated so far and, in the presence of a fork, aborts the execution in case the view on a given branch is not consistent with α_i. This intuitively ensures that the underlying protocol π is run only once, even in the presence of forks.

Since the initial protocol π may require private channels between the players, we need to augment the above transformation in such a way that subsets of honest parties can exchange messages in a confidential and authenticated manner. Let $m_{i,j}^{(r)}$ be the message that P_i sends to P_j at the generic round r. The latter is achieved by having P_i encrypting $m_{i,j}^{(r)}$ using the public encryption key ek_j of P_j, and then signing the resulting ciphertext $c_{i,j}^{(r)}$ with its own private signing key sk_i, which is the standard way of building a secure channel.

Generic Compiler π^*

Let π be an n-party ρ-round protocol, and $\Gamma = (\mathsf{UpdateState}, \mathsf{GetRecords}, \mathsf{Broadcast})$ be a blockchain protocol. Further, let $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ be a PKE scheme and $(\mathsf{Gen}', \mathsf{Sign}, \mathsf{Verify})$ be a signature scheme, both with domain $\{0,1\}^*$ (see App. 2.2 for the formal definitions). The protocol π^* proceeds as follows:
- For $i \in [n]$, each player P_i initializes $st_i := \varepsilon$, samples $(ek_i, dk_i) \leftarrow\!\!\!{}_{\$}\, \mathsf{Gen}(1^\lambda)$, $(vk_i, sk_i) \leftarrow\!\!\!{}_{\$}\, \mathsf{Gen}'(1^\lambda)$, and $\omega_i \leftarrow\!\!\!{}_{\$} \{0,1\}^*$, and invokes $\mathsf{Broadcast}(ek_i\|vk_i\|i)$.
- For $i \in [n]$, each player P_i keeps running $st_i \leftarrow\!\!\!{}_{\$}\, \mathsf{UpdateState}(1^\lambda)$ and $\mathcal{B}_i \leftarrow\!\!\!{}_{\$}\, \mathsf{GetRecords}(st_i)$ until all the messages $(ek_j, vk_j)_{j \in [n]} \in \mathcal{B}_i$.
- For $i \in [n]$, each player P_i sets $\tau^{(0)} := (ek_j, vk_j)_{j \in [n]}$ and $\alpha_i := \tau^{(0)}$, and then runs the following loop:
  1. Update the state st_i by running $\mathsf{UpdateState}(1^\lambda)$, and let $\mathcal{B}_i \leftarrow\!\!\!{}_{\$}\, \mathsf{GetRecords}(st_i)$.
  2. Let $\tilde{r} \geq 0$ be the maximum value such that the partial transcript $\tau^{(\tilde{r})} \in \mathcal{B}_i$. Then:
     - If the ciphertexts in $\tau^{(\tilde{r})}$ are not consistent with those in α_i, output $\perp$ and terminate.
     - Else if $\tilde{r} = \rho$, output the value y_i as a function of $\tau^{(\rho)}$ and terminate.
     - Else, go to the next step and if α_i is a prefix of $\tau^{(\tilde{r})}$ let $\alpha_i := \tau^{(\tilde{r})}$.
  3. For each $j \in [n]$, with $j \neq i$, and for each $r \leq \tilde{r}$, decrypt the ciphertexts $c_{j,i}^{(r)}$ and use the corresponding values $m_{j,i}^{(r)}$ to compute the messages $m_{i,j}^{(\tilde{r}+1)}$ to be sent at round $\tilde{r}+1$ (using the corresponding portion of the random tape ω_i).
  4. Finally, let $c_{i,j}^{(\tilde{r}+1)} \leftarrow\!\!\!{}_{\$}\, \mathsf{Enc}(ek_j, m_{i,j}^{(\tilde{r}+1)})$ (using again random coins coming from ω_i) and $\sigma_{i,j}^{(\tilde{r}+1)} = \mathsf{Sign}(sk_i, c_{i,j}^{(\tilde{r}+1)})$, and invoke $\mathsf{Broadcast}((c_{i,j}^{(\tilde{r}+1)}\|\sigma_{i,j}^{(\tilde{r}+1)})_{j \in [n]\setminus\{i\}})$.

Figure 1: Generic compiler for obtaining blockchain-aided MPC with quick players.

We refer the reader to Fig. 1 for a formal description. Note that wlog. we assume that for each round of the underlying protocol π every P_i sends a single message to each $\mathsf{P}_{j \neq i}$ over a private and authenticated channel. Moreover, P_i picks a sufficiently long random tape ω_i that is then used to run π over Γ. Observe that ω_i includes both the randomness required to compute the messages in π and the random coins used to encrypt them. In particular, in the presence of forks, an honest P_i that does not abort broadcasts on the blockchain exactly the same ciphertexts on multiple branches.

4.2 Security Analysis

The theorem below establishes the security of our generic compiler.

Theorem 1. *Let* $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ *be a semantically secure PKE scheme, and* $(\mathsf{Gen}', \mathsf{Sign}, \mathsf{Verify})$ *be a (deterministic) unforgeable signature scheme. Furthermore, let* π *be an* n-*party* ρ-*round protocol that* t-*securely UC-realizes a functionality* f *with aborts in the presence of malicious adversaries. Then, the protocol* π^* *of Fig. 1* t-*securely computes* f *with aborts in the presence of quick players and malicious adversaries.*

UC security is needed due to the fact that the attacker in the real world may interact with the blockchain by posting messages and reading its state. As shown in [CGJ19], such blockchain-active adversaries render standard simulation techniques (e.g., black-box rewinding) moot. Note also that Remark 2 does not hold for our protocol. If the adversary tries to furnish two different inputs in two different branches it can be spotted by some honest player, leading to an abort. Therefore only one possible input can be given to the functionality.

We need to show that for every PPT adversary A^*, there exists a PPT simulator S^* such that no non-uniform PPT distinguisher D^* can tell apart the experiments $\mathbf{REAL}_{\pi,\mathsf{A}^*,\mathsf{D}^*}^{\Gamma,0}(\lambda)$ and $\mathbf{IDEAL}_{f_\perp,\mathsf{S}^*,\mathsf{D}^*}^{\Gamma}(\lambda)$. In particular, the simulator S^* needs to simulate the interaction of the honest players with the blockchain protocol Γ as it happens in the real experiment. Intuitively, S^* relies on the simulator S guaranteed by the underlying protocol π as follows. At the beginning, S^* samples the public/secret keys for encryption/signatures for the honest players. Then, S^* runs A^* reading its messages from the emulated execution of the blockchain protocol Γ, and simulates its view as follows: (i) The round-r messages $m_{j,i}^{(r)}$ sent by the honest players P_j to the malicious players P_i are obtained from the simulator S; (ii) The round-r messages $m_{j,j'}^{(r)}$ that are exchanged by the honest players $\mathsf{P}_j, \mathsf{P}_{j'}$ are replaced with the all-zero string. Of course, S^* does additional bookkeeping in order to simulate a real execution of the protocol using the blockchain; in particular, S^* needs to check that the attacker plays consistently on different branches of a fork, and simulate an abort whenever the latter does not happen. Moreover, when S extracts the inputs for the malicious parties, the simulator S^* forwards the same inputs to the trusted party, obtains the outputs for the malicious parties, and sends it to S. Finally, S^* completes the simulation consistently with the choice of S of aborting or not.

Very roughly, the security of the PKE scheme and of the signature scheme imply that the view of the attacker is identical to that in a real execution of protocol π, so that security of π^* follows by that of π.

Proof. We begin by describing the simulator S^*. Let S be the PPT simulator guaranteed by the malicious security of π. Upon input the set of corrupted parties $\mathcal{I}$, inputs $(x_i)_{i \in \mathcal{I}}$, and auxiliary input z, the simulator S^* proceeds as follows:

1. Initialize S upon input $(\mathcal{I}, (x_i)_{i \in \mathcal{I}}, z)$, with uniformly chosen random tape $\omega_{\mathsf{sim}} \leftarrow_\$ \{0,1\}^*$.
2. For each $j \notin \mathcal{I}$, sample $(ek_j, dk_j) \leftarrow_\$ \mathsf{Gen}(1^\lambda)$, $(vk_j, sk_j) \leftarrow_\$ \mathsf{Gen}'(1^\lambda)$, $\omega_j \leftarrow_\$ \{0,1\}^*$, and invoke $\mathsf{Broadcast}(ek_j||vk_j||j)$.
3. For each $j \notin \mathcal{I}$, keep running $st_j \leftarrow_\$ \mathsf{UpdateState}(1^\lambda)$ and $\mathcal{B}_j \leftarrow_\$ \mathsf{GetRecords}(st_j)$ until all the messages $(ek_i, vk_i)_{i \in [n]} \in \mathcal{B}_j$. Set $\tau_j^{(0)} := (ek_i, vk_i)_{i \in [n]}$ and $\alpha_j := \tau^{(0)}$.
4. For each $j \notin \mathcal{I}$, emulate the behavior of party P_j as follows:
   (a) Update the state st_j by running $\mathsf{UpdateState}(1^\lambda)$, and let $\mathcal{B}_j \leftarrow_\$ \mathsf{GetRecords}(st_j)$.
   (b) Let $\tilde{r} \geq 0$ be the maximum value such that the partial transcript $\tau^{(\tilde{r})} \in \mathcal{B}_j$. Then:
      - If the ciphertexts in $\tau^{(\tilde{r})}$ are not consistent with those in α_j, send `abort` to the trusted party, simulate A^* aborting in the real protocol, and terminate.
      - Else, go to the next step and if α_j is a prefix of $\tau^{(\tilde{r})}$ let $\alpha_j := \tau^{(\tilde{r})}$.

(c) Extract from $\tau_j^{(\tilde{r})}$ the ciphertexts $(c_{i,j}^{(\tilde{r})})_{i \in \mathcal{I}}$ and the signatures $(\sigma_{i,j}^{(\tilde{r})})_{i \in \mathcal{I}}$ that A^* (on behalf of each corrupted player P_i) forwards to P_j. If there exists $i \in \mathcal{I}$ such that $\mathsf{Verify}(vk_i, \sigma_{i,j}^{(\tilde{r})}) = 0$, send $\mathtt{abort}$ to the trusted party, simulate A^* aborting in the real protocol, and terminate. Else, for each $r \le \tilde{r}$, decrypt the ciphertexts $c_{i,j}^{(r)}$ using the decryption key dk_j, and pass the corresponding messages $((m_{i,j}^{(1)})_{i \in \mathcal{I}, j \in \mathcal{H}}, \ldots, (m_{i,j}^{(\tilde{r})})_{i \in \mathcal{I}, j \in \mathcal{H}})$ to S. Hence:

- Upon receiving $\mathtt{abort}$ from S, send $\mathtt{abort}$ to the trusted party, simulate A^* aborting in the real protocol, and terminate.
- Upon receiving $(x_i)_{i \in \mathcal{I}}$ from S, send $(x_i)_{i \in \mathcal{I}}$ to the trusted party, obtain the outputs $(y_i)_{i \in \mathcal{I}}$, and forward $(y_i)_{i \in \mathcal{I}}$ to S. In case S replies with $(\mathtt{continue}, \mathcal{H}')$, send $(\mathtt{continue}, \mathcal{H}')$ to the trusted party and terminate.
- Upon receiving a set of messages $(m_{j,i}^{(\tilde{r}+1)})_{j \in \mathcal{H}, i \in \mathcal{I}}$—corresponding to the simulated messages that each honest player P_j sends to the corrupted party P_i— for each $j \in \mathcal{H}$ and $i \in \mathcal{I}$ compute $c_{j,i}^{(\tilde{r}+1)} \leftarrow_\$ \mathsf{Enc}(ek_i, m_{j,i}^{(\tilde{r}+1)})$ (using coins from ω_j) and $\sigma_{j,i}^{(\tilde{r}+1)} = \mathsf{Sign}(sk_j, c_{j,i}^{(\tilde{r}+1)})$. Then, for each $j, j' \in \mathcal{H}$, let $c_{j,j'}^{(\tilde{r}+1)} \leftarrow_\$ \mathsf{Enc}(ek_{j'}, 0^{|m_{j,j'}^{(\tilde{r}+1)}|})$ (using coins from ω_j) and $\sigma_{j,j'}^{(\tilde{r}+1)} \leftarrow_\$ \mathsf{Sign}(sk_j, c_{j,j'}^{(\tilde{r}+1)})$, and finally invoke $\mathsf{Broadcast}((c_{j,i}^{(\tilde{r}+1)} || \sigma_{j,i}^{(\tilde{r}+1)})_{i \in [n] \setminus \{j\}})$.

To conclude the proof, we consider a sequence of hybrid experiments (ending with the real experiment) and argue that each pair of hybrids is computationally close thanks to the properties of the underlying cryptographic primitives.

Hybrid $\mathbf{H}_3(\lambda)$: This experiment is identical to $\mathbf{IDEAL}_{f_\perp, \mathsf{S}^*, \mathsf{D}^*}^\Gamma(\lambda)$.

Hybrid $\mathbf{H}_2(\lambda)$: Identical to $\mathbf{H}_3(\lambda)$ except that we replace the ciphertexts $(c_{j,j'}^{(r)})_{j \in \mathcal{H}, j' \in \mathcal{H} \setminus \{j\}}$ that each honest party P_j sends to the other honest players $\mathsf{P}_{j'}$ with an encryption of the real messages $(m_{j,j'}^{(r)})_{j \in \mathcal{H}, j' \in \mathcal{H} \setminus \{j\}}$ that the same parties would send in a real execution of π. Note that the other ciphertexts $(c_{j,i}^{(r)})_{j \in \mathcal{H}, i \in \mathcal{I}}$ are still emulated using the simulator, and the output of the experiment is determined by the trusted party.

The inputs for the honest parties are chosen to be the values $(x_i)_{i \in \mathcal{H}}$ chosen by the distinguisher D^* at the beginning of the experiment, and the random tape of each player is chosen uniformly once and for all as in the real world.

Hybrid $\mathbf{H}_1(\lambda)$: Identical to $\mathbf{H}_2(\lambda)$ except that we artificially abort if A^* modifies one of the ciphertexts $(c_{j,i}^{(r)})_{j \in \mathcal{H}, i \in [n] \setminus \{j\}}$ corresponding to the messages that each honest player sends in a given round. Note that these ciphertexts correspond to both the real messages $(m_{j,j'}^{(r)})_{j \in \mathcal{H}, j' \in \mathcal{H} \setminus \{j\}}$ and the simulated messages $(m_{j,i}^{(r)})_{j \in \mathcal{H}, i \in \mathcal{I}}$.

Hybrid $\mathbf{H}_0(\lambda)$: This experiment is identical to $\mathbf{REAL}_{\pi^*, \mathsf{A}^*, \mathsf{D}^*}^{\Gamma, 0}(\lambda)$.

Lemma 1. $\{\mathbf{H}_3(\lambda)\}_{\lambda \in \mathbb{N}} \approx_c \{\mathbf{H}_2(\lambda)\}_{\lambda \in \mathbb{N}}$.

Proof. We reduce to semantic security of $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$. Let $h = |\mathcal{H}|$. For $k \in [0, h]$, consider the hybrid experiment $\mathbf{H}_{3,k}(\lambda)$ in which the distribution of the ciphertexts $(c_{j,j'}^{(r+1)})_{j \in \mathcal{H}, j' \in \mathcal{H} \setminus \{j\}}$ is modified as in $\mathbf{H}_2(\lambda)$ only for the first h honest parties. Clearly, $\{\mathbf{H}_{3,0}(\lambda)\}_{\lambda \in \mathbb{N}} \equiv \{\mathbf{H}_3(\lambda)\}_{\lambda \in \mathbb{N}}$ and $\{\mathbf{H}_{3,h}(\lambda)\}_{\lambda \in \mathbb{N}} \equiv \{\mathbf{H}_2(\lambda)\}_{\lambda \in \mathbb{N}}$.

Next, we prove that for every $k \in [0, h]$ it holds that $\{\mathbf{H}_{3,k}(\lambda)\}_{\lambda \in \mathbb{N}} \approx_c \{\mathbf{H}_{3,k+1}(\lambda)\}_{\lambda \in \mathbb{N}}$ which concludes the proof of the lemma. By contradiction, assume that there exists an index $k \in [0, h]$, and a pair of PPT algorithms $(\mathsf{D}^*, \mathsf{A}^*)$ such that D^* can distinguish the two experiments $\mathbf{H}_{3,k}(\lambda)$ and $\mathbf{H}_{3,k+1}(\lambda)$ with non-negligible probability. We construct a PPT attacker B breaking semantic security of $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ as follows:

- Receive the target public encryption key ek^* from the challenger.
- Run D^*, receiving the set of corrupted parties $\mathcal{I}$, the inputs $(x_i)_{i \in [n]}$, and the auxiliary input z, then pass $(\mathcal{I}, (x_i)_{i \in \mathcal{I}}, z)$ to A^*.
- Interact with A^* as described in the ideal experiment, except that:
  - The public encryption key for player P_{k+1} is set to be the target public key ek^*.
  - For each $j \le k$, when it comes to simulating the ciphertexts $(c_{j,j'}^{(r)})_{j' \in \mathcal{H} \setminus \{j\}}$, use the real messages $(m_{j,j'}^{(r)})_{j' \in \mathcal{H} \setminus \{j\}}$, encrypt them using the public encryption key $ek_{j'}$ of $\mathsf{P}_{j'}$, and sign the ciphertexts with the secret key sk_j (which is known to the reduction).
  - When it comes to simulating the ciphertexts $(c_{k+1,j'}^{(r)})_{j' \in \mathcal{H} \setminus \{j\}}$, forward the pair of plaintexts $(m_{k+1,j'}^{(r)}, 0^{|m_{k+1,j'}^{(r)}|})_{j' \in \mathcal{H} \setminus \{k+1\}}$ to the left-or-right encryption oracle and sign the corresponding ciphertexts using the secret signing key sk_{k+1} of P_{k+1} (which is known to the reduction).
  - For each $j > k+1$, when it comes to simulating the ciphertexts $(c_{j,j'}^{(r)})_{j' \in \mathcal{H} \setminus \{j\}}$, use the dummy messages $(0^{|m_{j,j'}^{(r)}|})_{j' \in \mathcal{H} \setminus \{j\}}$, encrypt them using the public encryption key $ek_{j'}$ of $\mathsf{P}_{j'}$, and sign the ciphertexts with the secret key sk_j (which is known to the reduction).
- Finally, run D^* upon the final output generated by A^*, and return whatever D^* outputs.

Note that the reduction can indeed simulate the interaction with the blockchain protocol Γ as in the ideal experiment, and moreover it can generate the real messages $(m_{j,j'}^r)_{j \in \mathcal{H}, j' \in \mathcal{H} \setminus \{j\}}$ as it knows the parties' inputs $(x_i)_{i \in [n]}$. By inspection, the simulation performed by B is perfect in the sense that when the challenger encrypts the messages $m_{k+1,j'}^{(r)}$ the view of $(\mathsf{D}^*, \mathsf{A}^*)$ is identical to that in $\mathbf{H}_{3,k+1}(\lambda)$. Similarly, when the challenger encrypts the dummy messages $0^{|m_{k+1,j'}^{(r)}|}$ the view of $(\mathsf{D}^*, \mathsf{A}^*)$ is identical to that in $\mathbf{H}_{3,k}(\lambda)$. Hence, B breaks semantic security of $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ with non-negligible probability, concluding the proof. $\qquad\square$

Lemma 2. $\{\mathbf{H}_2(\lambda)\}_{\lambda \in \mathbb{N}} \approx_c \{\mathbf{H}_1(\lambda)\}_{\lambda \in \mathbb{N}}$.

Proof. Let $\mathbf{BAD}$ be the event that an artificial abort happens in $\mathbf{H}_1(\lambda)$. Note that this means that, for some $j \in \mathcal{H}$, the attacker A^* replaces one of the ciphertexts $c_{j,i}^{(r)}$ that P_j would send to P_i in the real protocol with a different ciphertext $\tilde{c}_{j,i}^{(r)}$, in such a way that the corresponding signature $\tilde{\sigma}_{j,i}^{(r)}$ is still accepting. Clearly, the experiments $\mathbf{H}_2(\lambda)$ and $\mathbf{H}_1(\lambda)$ are identical conditioning on $\mathbf{BAD}$ not happening, and does it suffices to show that $\mathbb{P}[\mathbf{BAD}]$ is negligible.

Given a PPT distinguisher D^* and a PPT attacker A^* such that A^* provokes event $\mathbf{BAD}$ in a run of $\mathbf{H}_2(\lambda)$ with non-negligible probability, we can construct a PPT attacker B breaking security of the signature scheme $(\mathsf{Gen}', \mathsf{Sign}, \mathsf{Verify})$. The reduction works as follows:
- Receive the target public verification key vk^* from the challenger.
- Choose a random j^* as a guess for the index corresponding to the honest party for which A^* provokes the bad event.
- Run D^*, receiving the set of corrupted parties $\mathcal{I}$, the inputs $(x_i)_{i \in [n]}$, and the auxiliary input z, then pass $(\mathcal{I}, (x_i)_{i \in \mathcal{I}}, z)$ to A^*.
- Interact with A^* as described in $\mathbf{H}_2(\lambda)$, except that:
  - The public verification key for player P_{j^*} is set to be the target public key vk^*.
  - When it comes to simulating the round-r messages from party P_{j^*}, generate the ciphertexts $(c_{j^*,i}^{(r)})_{i \in [n] \setminus \{j^*\}}$ as done in $\mathbf{H}_2(\lambda)$, and then forward each of $c_{j^*,i}^{(r)}$ to the challenger, obtaining the corresponding signature $\sigma_{j^*,i}^{(r)}$ that is needed in order to complete the simulation.

– Keep updating the local state of P_{j^*} until an index $i \in [n] \setminus \{j^*\}$ is found such that the partial transcript α_{j^*} contains a pair $(\tilde{c}_{j^*,i}^{(r)}, \tilde{\sigma}_{j^*,i}^{(r)})$ such that $\mathsf{Verify}(vk_{j^*}, \tilde{c}_{j^*,i}^{(r)}, \tilde{\sigma}_{j^*,i}^{(r)}) = 1$ and $\tilde{c}_{j^*,i}^{(r)}$ is different from the original ciphertext $c_{j^*,i}^{(r)}$ previously sent on behalf of P_{j^*}.

- If no such pair is found, abort the simulation. Else, return $(\tilde{c}_{j^*,i}^{(r)}, \tilde{\sigma}_{j^*,i}^{(r)})$.

Note that the simulation performed by B is perfect, in that the view of $(\mathsf{D}^*, \mathsf{A}^*)$ is identical to that in a run of $\mathbf{H}_2(\lambda)$. Moreover, conditioning on B guessing the index j^* correctly, the reduction is successful in breaking security of the signature scheme exactly with probability at least $\mathbb{P}[\mathbf{BAD}]$, which is non-negligible. Since the former event also happens with non-negligible probability, this concludes the proof. $\qquad\square$

Lemma 3. $\{\mathbf{H}_1(\lambda)\}_{\lambda \in \mathbb{N}} \approx_c \{\mathbf{H}_0(\lambda)\}_{\lambda \in \mathbb{N}}$.

Proof. The proof is by reduction to UC-security of the underlying protocol π. By contradiction, assume that there exists a PPT adversary A^* and a non-uniform PPT distinguisher D^* that can distinguish between $\mathbf{H}_1(\lambda)$ and $\mathbf{H}_0(\lambda)$ with non-negligible probability. Consider the following PPT attacker A, initialized with a set of corrupted parties $\mathcal{I}$, inputs $(x_i)_{i \in \mathcal{I}}$ for the malicious players, and auxiliary input $z = (z^*, (ek_i, dk_i)_{i \in [n]}, (vk_i, sk_i)_{i \in [n]})$ which will be specified later:

- Pass $(\mathcal{I}, (x_i)_{i \in \mathcal{I}}, z^*)$ to A^*.
- For each $i \in \mathcal{I}$, upon receiving the round-r messages $(m_{j,i}^{(r)})_{j \in \mathcal{H}}$ from the honest players to the malicious players, let $c_{j,i}^{(r)} \leftarrow_{\$} \mathsf{Enc}(ek_i, m_{j,i}^{(r)})$ and $\sigma_{j,i}^{(r)} = \mathsf{Sign}(sk_j, c_{j,i}^{(r)})$, and emulate broadcasting $(c_{j,i}^{(r)}, \sigma_{j,i}^{(r)})_{j \in \mathcal{H}, i \in \mathcal{I}}$ via the blockchain protocol.
- For each $j \in \mathcal{H}$, upon receiving the round-r messages $(m_{i,j}^{(r)})_{i \in \mathcal{I}}$ that A^* wants to send to the honest parties, let $c_{i,j}^{(r)} \leftarrow_{\$} \mathsf{Enc}(ek_j, m_{i,j}^{(r)})$ and $\sigma_{i,j}^{(r)} = \mathsf{Sign}(sk_i, c_{i,j}^{(r)})$, and emulate broadcasting $(c_{i,j}^{(r)}, \sigma_{i,j}^{(r)})_{i \in \mathcal{I}, j \in \mathcal{H}}$ via the blockchain protocol.
- For each $j \in \mathcal{H}$, compute the messages $(m_{j,j'}^{(r)})_{j' \in \mathcal{H} \setminus \{j\}}$ exchanged between honest parties as done in $\mathbf{H}_0$ (which is the same in $\mathbf{H}_1(\lambda)$), let $c_{j,j'}^{(r)} \leftarrow_{\$} \mathsf{Enc}(ek_{j'}, m_{j,j'}^{(r)})$ and $\sigma_{j,j'}^{(r)} = \mathsf{Sign}(sk_j, c_{j,j'}^{(r)})$, and emulate broadcasting $(c_{j,j'}^{(r)}, \sigma_{j,j'}^{(r)})_{j \in \mathcal{H}, j' \in \mathcal{H} \setminus \{j\}}$ via the blockchain protocol.
- In case a fork appears during the simulation of the underlying blockchain protocol, replicate the messages from the honest players as done in the other branches (using exactly the same randomness). On the other hand, if the messages from A^* differ from those sent on the simulation of a previous branch, simulate A^* aborting and terminate.
- Output whatever A^* outputs.

Additionally, let Z be the following PPT distinguisher:

- Run D^*, receiving the set of corrupted parties $\mathcal{I}$, the inputs $(x_i)_{i \in [n]}$, and the auxiliary input z^*, then sample (ek_i, dk_i) and (vk_i, sk_i) for all $i \in [n]$, and pass $(\mathcal{I}, (x_i)_{i \in \mathcal{I}}, z)$ to the above defined attacker A, where $z = (z^*, (ek_i, dk_i)_{i \in [n]}, (vk_i, sk_i)_{i \in [n]})$.
- Upon receiving the final output from A, pass it to D^* and output whatever D^* outputs.

By inspection, in case the attacker A is playing in a real execution of protocol π, the view of D^* is identical to that in an execution of $\mathbf{H}_0(\lambda)$ with A^* controlling the malicious parties. Similarly, in case the view of A is emulated using the simulator S (corrupting the dummy parties controlled by A) of protocol π, the view of D^* is identical to that in an execution of $\mathbf{H}_1(\lambda)$ with A^* controlling the malicious parties. It follows that Z can distinguish between $\mathbf{REAL}_{\pi, \mathsf{A}, \mathsf{Z}}(\lambda)$ and $\mathbf{IDEAL}_{f_\perp, \mathsf{S}, \mathsf{Z}}(\lambda)$ with non-negligible probability, a contradiction. $\qquad\square$

The theorem now follows directly by combining the above lemmas. $\qquad\square$

4.3 On Fairness with Penalties

In their work, Andrychowicz *et al.* [ADMM14, ADMM16] proposed a different notion of fairness for MPC protocols that run on blockchain: fairness with penalties. This notion states that if an adversary in an MPC protocol decides to abort the execution of the protocol it will be financially penalized. To obtain the penalization in the lottery protocol, Andrychowicz *et al.* added a deposit step in the protocol.

Our compiler described in §4 does not have suffer we discuss here how to obtain fairness with penalties following in part the outline of [KB14, BK14].

Let's now assume the existence of an (n, n)-secret sharing scheme $(\mathsf{Share}, \mathsf{Recon})$, non-interactive commitment schemes (see App. 2.2 for the formal definitions), and consider a functionality f' that first calculates $y \leftarrow f(x_1, \ldots, x_n)$, where each party P_i holds x_i, and then calculates the shares of the output $(\sigma_1, \ldots, \sigma_n) \leftarrow_\$ \mathsf{Share}(y)$, the commitments $C = (\gamma_1, \ldots, \gamma_n)$, where $\gamma_i \leftarrow_\$ \mathsf{Commit}(\sigma_i)$, and outputs (C, σ_i) to each player P_i. Let's call π' the protocol realizing f', we can apply our generic compiler to π' to obtain a protocol π'_{bc} that can be run in the blockchain. Our protocol π_{fair}, running with players $\mathsf{P}_1 \ldots, \mathsf{P}_n$ works as follows

(i) *Protocol execution*: All the players engage in the protocol π'_{bc}. A party P_i aborts the execution if π'_{bc} aborts. Otherwise, obtains (C, σ_i) in the last round.

(ii) *Smart contract*: P_1 publishes the smart contract depicted in Fig. 2.

(iii) *Commitment phase*: For each $i \in [n]$, P_i triggers `deposit`(C_i) together with d coins, where d is a fixed deposit. If some player does not publish his commitments with the deposit or there is a disagreement on the commitments within `time1` (i.e., a player P_j sends $C_j \neq C_i$ for some $P_{i \neq j}$, or deposits a value $d_i < d$, P_i abort the execution. Recall that abort in this phase is still fine, since no information about the output y is released. Otherwise, if `time1` has passed, go to the Opening Phase.

(iv) *Opening phase*: For each $i \in [n]$, P_i opens his commitment by sending `openCom`(i, σ_i), thus receiving back his d coins, wait that all the openings are published in the smart contract (within `time2`) and calculates $y \leftarrow \mathsf{Recon}(\sigma_1, \ldots, \sigma_n)$. If, after `time2`, some share is missing, P_i aborts the execution.

During the last phase, if some player did not open the commitment or sent an incorrect value, the smart contract will penalize him by freezing his deposit. Thus, the adversary is not incentivized to send an incorrect share.

This attempt to add fairness with penalties, however, introduces an attack. Given an n-party protocol $\pi_{f'}^\Gamma$ obtained by the compiler described in §4 applied to $\pi_{f'}$, with the addition of the smart contract, commit and opening phases described above, we have the following scenario:

- For all $i \in [n]$, party P_i runs π_{fair} obtaining (C, σ_i).
- For all $i \in [n]$, party P_i triggers `deposit`(C) together with d coins to the smart contract.
- For all $i \in [n-1]$, party P_i opens his commitment by triggering `openCom`(σ_i).
- Wlog., we say that P_n is an adversary. P_n compute the output y. If P_n does not like y in the current branch, P_n can try to exploit a fork happening during the execution of π_{fair} to change the in a different branch to obtain a new couple (C', σ_n').
- The honest parties $\mathsf{P}_1, \ldots, \mathsf{P}_{n-1}$ notice that there is a message published by P_n that differs from the value previously received always by P_n. Since the transcript obtained from the blockchain differs from the transcript stored in their local state, they will abort.

The protocol described is not fair, since we can construct a counterexample that prove that the unfair party P_n can obtain the output without being penalized.

$\mathsf{P}_1, \ldots, \mathsf{P}_n$ will play $\pi_{f'}^\Gamma$ to obtain (C, σ_i). As the smart contract is published, P_n will trigger `deposit`(C) together with d coins. At this point, P_n waits that all P_i, with $i \in [n]$ publish the

The **General Compiler Smart Contract** runs with players $P_1, \ldots, P_n$ and consists of two main functions `deposit` and `openCom` and two fixed timestamps `time1`,`time2`.

Commitment Phase: In round t_1, when `deposit`(C_i) together with d coins is triggered from a party P_i, store (i, C_i). Then, if $\forall i, j : C_i = C_j$ proceed to the Opening Phase. Otherwise, for all i, if the message (i, C_i) has been stored, send message d coins to P_i and terminate.

Opening Phase: In round t_2, when `openCom`(i, σ_i) is triggered from P_i, check if $\mathsf{Commit}(\sigma_i) = \gamma_i$, where γ_i is obtained from parsing $C_i = (\gamma_1, \ldots, \gamma_n)$ (recall that all the C_i are the same), and send d coins back to P_i.

Figure 2: General compiler smart contract.

opening σ_i.

When P_n sees $\sigma_1, \ldots, \sigma_{n-1}$ on the blockchain, computes the output. If P_n dislikes the output, then he tries to exploit a branch created before the end of the execution of π_{fair} to change messages in that branch to obtain an advantage. Since P_n publishes different messages on different branches of the blockchain, there exist some party P_i, with $i \in [n]$ that will notice it, causing an abort in the protocol.

Let's call b_1 the branch where P_n learned the output and b_2 the branch exploited to change the execution of π_{fair}. We have two cases:

- If b_1 is the branch that will be confirmed on the blockchain, P_n will be penalized.
- If b_2 is the branch that will be confirmed on the blockchain, P_n will cause an abort in the protocol before that the commitment phase starts. In this case he does not get penalized for learning the output.

With this counterexample we show that the proposed solution is not enough to obtain fairness with penalties, since P_n has the possibility to learn the output without incurring in any punishment. It is possible to obtain fairness with penalties with our general compiler, and waiting that the commitment phase is confirmed on the ledger. Since in this case the commitment phase is finalized, an adversary A cannot learn y unless she decides to lose the d coins deposited in the commitment phase. Since the commitment phase is confirmed on the ledger, A cannot find a fork to exploit the execution of the protocol on another branch. Yet, A can cause an abort in the protocol, but if it happens before the commitment phase she will not learn the output y. If the abort happens after the commitment phase, A will learn the output but will be penalized.

Theorem 2. *Let's assume the existence of non-interactive commitment schemes and (n, n)-secret sharing schemes. Let π'_{bc} be an n-party ρ-round protocol realizing f' in the presence of quick players. Then, the protocol π_{fair} described above securely realizes f satisfying fairness with penalties in the presence of $(\rho, 1)$-quick players.*

Proof Sketch. We can claim security of the the compiled protocol π'_{bc} obtained by applying the general compiler to π', by referring to the same proof of Theorem. 1. Now, we argue that the overall protocol π_{fair} achieves fairness with penalties. As mentioned before, aborts during the execution of π'_{bc} are acceptable, since the adversary cannot learn any information about the output. After the committing phase, that is finalized, the adversary could try to exploit different branches to send different openings of his commitments. We have the following timeline: The execution started in a branch b_1 and a forks happens right after the committing phase,

generating a branch b_2. Wlog. of generality we can extend this argument to multiple parallel executions in different branches. We have the following scenarios:

Scenario 1: A corrupted player abort in both branches. Since the commitments are finalized, fairness with penalties follows in a straightforward manner, since he did not open his commitments in each branch, and so also in the confirmed one.

Scenario 2: A corrupted player opens his share in b_1 and aborts the execution in b_2 (after the commitment phase). If b_1 gets confirmed, the honest parties will learn the output. If b_2 gets confirmed, it automatically boils down to Scenario 1.

Scenario 3: A corrupted player P_i opens his share in b_1 and tries to open on a different share in b_2. Since the commitment is always confirmed, the adversary cannot try to change his commitment by exploiting forks. If A is able to open the commitment by providing two different shares, then we can define an adversary $\mathsf{A}_{\mathsf{com}}$ breaking the binding property of the underlying commitment scheme with non-negligible probability. That means that at least in one of the two branches P_i gets penalized, and if he provides the correct opening in one of the branches and it gets confirmed, honest players will learn the output.

$\square$

5 Parallel Coin Tossing

Motivated by the shortcoming of our compiler when considering fairness with penalties in the presence of quick players (cf. §4.3), in this section we explore an alternative solution for the case of parallel coin tossing. Recall that coin tossing allows a set of players to agree on a uniformly random string, and has many important applications (as it, e.g., allows to easily implement a distributed lottery). Importantly, our protocol is compatible with the standard techniques to achieve fairness with penalties, but does not exploit finality (thus allowing the players to being fully quick).

In order to build some intuition on the design of our solution, we begin by recalling the protocol by Andrychowicz *et al.* [ADMM16] in §5.1. Hence, in §5.2, we show that their protocol becomes completely insecure in the presence of quick players. This naturally leads to our new protocol, which we describe and analyze in §5.3.

5.1 The Protocol of Andrychowicz *et al.*

Recall that in the Bitcoin ledger, each account is associated to a pair of keys (pk, sk), where pk is the verification key of a signature scheme—representing the address of an account—while sk is the corresponding secret key used to sign (the body of) the transactions. Each block on the ledger contains a list of transactions, and new blocks are issued by an entity called *miner*. The blockchain is maintained via a consensus mechanism based on the proof of work; users willing to add a transaction to the ledger forward it to the miners, which will try to include it in the next minted block.

In the description below, we say that a transaction is *valid* if it is computed correctly (i.e., the signature is valid, the coins have not been spent already, and so on), and that it is *confirmed* if it appears in the common-prefix of all the miners (i.e., it is at least k-blocks deep in the ledger). Each transaction Tx contains the following information:

- A set of input transactions $\mathsf{Tx}_1, \mathsf{Tx}_2, \cdots$ from which the coins needed for the actual transaction Tx are taken;
- A set of input scripts containing the input for the output scripts of $\mathsf{Tx}_1, \mathsf{Tx}_2, \cdots$;
- An output script defining in which condition Tx can be claimed;
- The number of coins taken from the redeemed transactions;

- A time lock t specifying when Tx becomes valid (i.e., a time-locked transaction won't be accepted by the miners before time t has passed).

The construction by [ADMM14, ADMM16] relies on a primitive called *time-locked commitment*. Let n denote the number of parties. Each party P_j creates $n-1$ $\mathsf{Commit}_{i\neq j}^j$ transactions containing a commitment to its lottery value. In particular, the output script of such a transaction ensures that it can be claimed either by P_j via an Open_i^j transaction exhibiting a valid opening for the commitment, or by another transaction that is signed by both P_j and P_i. Before posting these transactions on the ledger, P_j creates a time-locked transaction $\mathsf{PayDeposit}_i^j$ redeeming Commit_i^j, sends it off-chain to each $\mathsf{P}_{i\neq j}$, and finally posts all the Commit_i^j transactions on the ledger. In case P_j does not open the commitment before time τ, then each recipient of a $\mathsf{PayDeposit}_i^j$ transaction can sign it and post it on the ledger; since time τ has passed, the miners will now accept the transaction as a valid transaction redeeming Commit_i^j. More in details:

Deposit phase: Each player P_j computes a commitment $y_j = \mathsf{Hash}(x_j\|\delta_j)$, where δ_j is some randomness, sends off-chain the $\mathsf{PayDeposit}_i^j$ transactions (with time-lock τ) to each $\mathsf{P}_{i\neq j}$, and posts the Commit_i^j transactions.

Betting phase: P_j bets one coin in the form of a transaction $\mathsf{PutMoney}_j$ (redeeming a previous transaction held by P_j, and with P_j's signature as output script). All the players agree and sign off-chain a $\mathsf{Compute}$ transaction taking as input all the $(\mathsf{PutMoney}_j)_{j\in[n]}$ transactions, and then the last player that receives the $\mathsf{Compute}$ transaction posts it on the ledger. In order to claim this transaction, a player $\mathsf{P}_{w'}$ must exhibit the openings of the commitments of all participants: The script checks that the openings are valid, computes the index of the winner w (as a function of the values $x_1, \ldots, x_n$), and checks that $w' = w$ (i.e., the only participant that can claim the $\mathsf{Compute}$ transaction is the winner of the lottery).

Compensation phase: After time τ, in case some player P_j did not send all of its $\{\mathsf{Open}_i^j\}_{i\in[n],i\neq j}$ transactions, all the other players $\mathsf{P}_{i\neq j}$ can post the $\mathsf{PayDeposit}_i^j$ transaction, thus obtaining a compensation.

5.2 A Simple Attack

The main idea behind our attack is that, in the presence of quick players, the protocol's messages can end-up on unconfirmed blocks. By looking at different branches of a fork, an attacker can try to change an old unconfirmed transaction by re-posting it, with the hope that it will end-up on a different branch and become part of the common prefix. This essentially corresponds to a reset attack on the protocol.

The construction described in §5.1 relies on the (implicit) assumption that the players are non-quick. In particular, each player P_j should wait to post its $\mathsf{PutMoney}_i^j$ transaction only after all the Commit_i^j transactions are confirmed on the ledger, in such a way that all players are aligned on the same branch (and so the miners have the $\{\mathsf{Commit}_i^j\}_{i\in[n],j\neq i}$ transactions in their common prefix). In the case of quick players, when a fork occurs, an attacker can take advantage of the openings of the other parties played in a faster branch in order to bias the result of the lottery on a slower branch. If eventually the slower branch remains permanently in the blockchain, then clearly the attack is successful.

For concreteness, let us focus on Blum's coin tossing, in which the winner is defined to be $w = x_1 + \ldots + x_n \bmod n + 1$. Consider the following scenario:

- The (quick) players $\mathsf{P}_1, \ldots, \mathsf{P}_n$ run a full instance of the protocol; note that this requires at least 3 blocks.
- The attacker P_n hopes to see a fork containing all the $\{\mathsf{Commit}_j^i\}$ transactions of the other $n-1$ players.

- Since the attacker P_n now knows the openings $x_1, \ldots, x_{n-1}$, it can post a new set of $\{\mathsf{Commit}'^i_n\}_{i \in [n], i \neq n}$ transactions containing a commitment to a value x'_n such that $x_1 + \ldots + x_{n-1} + x'_n \mod n + 1 = n$.

In case the new set of transactions ends up on a different branch which is finally included in the common prefix, P_n wins the lottery. In the next section, we propose a new protocol that does not suffer from this problem.

5.3 Our New Lottery Protocol

We now present a parallel coin-tossing (PCT) protocol on blockchain that is secure in the presence of quick players. The main challenge that we face is that the protocol must prevent an adversary from choosing adaptively her contribution to the coin tossing in a branch of a fork, after possibly seeing the contributions of the other players in different branches.

We tackle this problem by requiring that each honest party computes his contribution by evaluating a unique signature (see App. 2.2 for the formal definition) upon input the public keys of all players. Notice that if the adversary A sees some signatures in a branch, and changes her public key in another branch, then A cannot predict the signatures of the honest players on this other branch by the unforgeability property of the signature scheme, and thus A will not manage to bias the final output. Hence, we hash the concatenation of all the signatures in order to determine the final output. Assuming that the hash function is modelled as a random oracle, we would like to argue that the output of the protocol looks uniform.

However, the following subtlety arises. Wlog., assume that only P_n is corrupt and that the protocol proceeds until the end on a given branch of the blockchain. Denote by pk_n the public key chosen by the attacker. Further, assume that A notices another branch where all honest players have already sent their public keys. Now, the adversary can either: (i) publish a different public key pk'_n, or (ii) publish the same public key pk_n as in the other branch. In case A "likes" the outcome of the protocol on the first branch, she will choose option (ii) and thus can bias the protocol output.

To avoid the above attack, we identify each branch with a string bid that is uniquely associated to it, and include bid as part of the message to sign. Intuitively, this solves the previous problem as, even if all the public keys stay the same on two different branches, the value bid will change thus ensuring that the protocol output will also be different (and uniformly random). We proceed with a more detailed description of our protocol (see also Fig. 3).[9]

- One of the players chooses a random value sid that represents the identifier of the current protocol execution, and publishes sid on the blockchain.
- Each player P_i willing to participate generates the public and private keys for the unique signature $(pk_i, sk_i) \leftarrow_\$ \mathsf{Gen}(1^\lambda)$, and publishes pk_i on the blockchain.
- Each player P_i lets $y_i = \mathsf{Sign}(sk_i, pk_1 || \cdots || pk_n || sid || bid)$, where bid is the hash of the blockchain[10] up to the block that contains the last public key, and publishes y_i on the blockchain.
- Each player P_i checks that $\mathsf{Verify}(pk_j, x, y_j) = 1$ for all $j \neq i$, where $x = pk_1 || \cdots || pk_n || sid || bid$, and outputs $\mathsf{Hash}(y_1 || \cdots || y_n)$.

We stress that thanks to the value bid, the protocol execution becomes branch dependent. In particular, the chances of success of a corrupted P_j to bias the output are not affected by

---

[9]Note that our protocol can be run on generic blockchains. In 6.2 , we provide an implementation using Ethereum smart contracts, but the protocol can also be implemented in Bitcoin using the opcode OP_RETURN in case players do not need to get fairness with penalties.

[10]For efficiency the hash can be more simply applied to the block containing pk_n. Nevertheless, for the sake of simplicity of the protocol description and of the security analysis we will stick with hashing the entire blockchain.

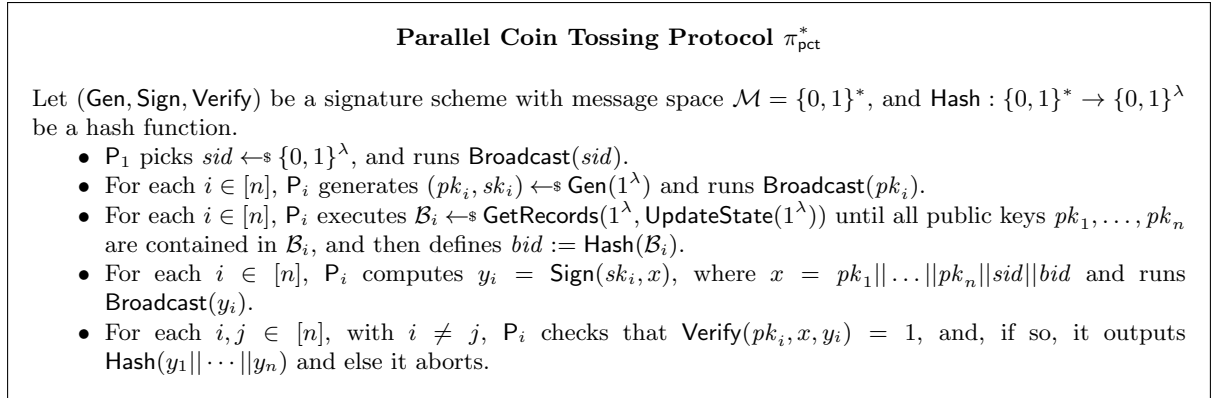the potential use of different public keys in branches of a fork corresponding to a protocol run with a given *sid*.

Parallel Coin Tossing Protocol π_{pct}^*

Let $(\mathsf{Gen}, \mathsf{Sign}, \mathsf{Verify})$ be a signature scheme with message space $\mathcal{M} = \{0,1\}^*$, and $\mathsf{Hash} : \{0,1\}^* \to \{0,1\}^\lambda$ be a hash function.

- P_1 picks $sid \leftarrow\!\!\$ \{0,1\}^\lambda$, and runs $\mathsf{Broadcast}(sid)$.
- For each $i \in [n]$, P_i generates $(pk_i, sk_i) \leftarrow\!\!\$ \mathsf{Gen}(1^\lambda)$ and runs $\mathsf{Broadcast}(pk_i)$.
- For each $i \in [n]$, P_i executes $\mathcal{B}_i \leftarrow\!\!\$ \mathsf{GetRecords}(1^\lambda, \mathsf{UpdateState}(1^\lambda))$ until all public keys $pk_1, \dots, pk_n$ are contained in $\mathcal{B}_i$, and then defines $bid := \mathsf{Hash}(\mathcal{B}_i)$.
- For each $i \in [n]$, P_i computes $y_i = \mathsf{Sign}(sk_i, x)$, where $x = pk_1||\dots||pk_n||sid||bid$ and runs $\mathsf{Broadcast}(y_i)$.
- For each $i, j \in [n]$, with $i \neq j$, P_i checks that $\mathsf{Verify}(pk_i, x, y_i) = 1$, and, if so, it outputs $\mathsf{Hash}(y_1||\cdots||y_n)$ and else it aborts.

Figure 3: Our new protocol for parallel coin tossing.

Security analysis. Let f^{pct} be the n-party functionality that picks a uniformly random string ω and sends it to all the n parties. The theorem below establishes the security of our coin-tossing protocol in the (programmable) random oracle model. We note that the security of the original protocol by Andrychowicz *et al.* [ADMM14, ADMM16] also relies on the random oracle heuristic, as do all currently known analysis of blockchain protocols [GKL15, PSS17].

Theorem 3. *Assuming that $(\mathsf{Gen}, \mathsf{Sign}, \mathsf{Verify})$ is a unique signature scheme, the protocol of Fig. 3 securely implements the functionality f^{pct} in the presence of quick players and malicious adversaries with aborts, in the programmable random oracle model.*

We need to show that for every PPT adversary A, there exists a PPT simulator S such that no non-uniform PPT distinguisher D can tell apart the experiments $\mathbf{REAL}_{\pi,\mathsf{A},\mathsf{D}}^{\Gamma,0}(\lambda)$ and $\mathbf{IDEAL}_{f_\perp^{\mathsf{pct}},\mathsf{S},\mathsf{D}}^{\Gamma}(\lambda)$. In particular, the simulator needs to simulate the interaction of the honest players with the blockchain protocol Γ as it happens in the real experiment.

Recall that the ideal coin-tossing functionality does not take any input, and returns a random string ω to all parties. Intuitively, the simulator S will just emulate a real execution of the protocol by reading the attacker's messages from the blockchain. Furthermore, after querying the ideal functionality, S will program the random oracle upon input ω. At the same time, S needs to simulate the answer to A's random oracle queries q, which is done as follows:

- If A already queried the random oracle upon q, obtaining answer $\tilde{\omega}$, then S returns the same value $\tilde{\omega}$;
- Else, if $pk_1, \dots, pk_n$ are not all available on the blockchain, S answers the query q with a fresh random value $\tilde{\omega}$;
- Otherwise, S parses q as $\tilde{y}_1||\dots||\tilde{y}_n$, and checks that $\mathsf{Verify}(pk_i, x, \tilde{y}_i) = 1$ for all $i \in [n]$, where $x = pk_1||\dots||pk_n||sid||bid$ and bid is derived by hashing the blockchain up to the block containing the last public key. If the check passes, S programs $\mathsf{Hash}(q)$ to the value ω obtained by $f_\perp^{\mathsf{pct}}$, else it answers with a fresh random value $\tilde{\omega}$.

The above strategy does not consider the fact that the real protocol may be run multiple times on different branches. However, since each execution has associated a different $bid' \neq bid$, the simulator can simply query again the functionality in order to obtain a new random value ω', and program the random oracle on ω' in the execution corresponding to bid'. Roughly

speaking, the attacker A can potentially take advantage from the following three actions: (i) A can refuse to publish her signed message; (ii) A can try to change the public keys of corrupted parties over different branches; (iii) A can try to choose the signature that produces the best possible output. Action (i) is equivalent to aborting the protocol, and can be easily handled by the simulator, since we achieve security with aborts. Action (ii) is tackled by the use effect of the of bid: regardless of A using the same or another public key, the outcome of the protocol in different branches are always independent. Finally, as for action (iii), note that this attack is prevented by the uniqueness property of the signature scheme that ensures that for every (possibly malicious) public key pk and every input x, there exists at most one valid signature y that is not rejected by the verification algorithm.

Proof of Theorem 3. We begin by describing the simulator. Upon input the set of corrupted parties $\mathcal{I}$, and auxiliary input z, the simulator S proceeds as follows:[11]

1. Initialize an empty array $\mathcal{L}$. Sample a random value $sid \leftarrow_\$ \{0,1\}^\lambda$, and run $\mathsf{Broadcast}(sid)$.
2. Generate $(pk_i, sk_i) \leftarrow_\$ \mathsf{Gen}(1^\lambda)$ and run $\mathsf{Broadcast}(pk_i)$ for each $i \notin \mathcal{I}$.
3. Keep running $\mathsf{GetRecords}(\mathsf{UpdateState}(1^\lambda))$ until all public keys $pk_1, \dots, pk_n$ are published on the blockchain; when that happens define bid to be the hash of the blockchain up to the block containing the last public key.
4. If $\mathcal{L}$ does not already contain the value bid, query the ideal functionality $f_\perp^{\mathsf{pct}}$, obtaining a random value ω, and store (bid, ω) in $\mathcal{L}$. Then complete the simulation for branch bid as follows:
   (a) For each $i \notin \mathcal{I}$, let $y_i = \mathsf{Sign}(sk_i, x)$ where $x = pk_1||\cdots||pk_n||sid||bid$. Then run $\mathsf{Broadcast}(y_i)$.
   (b) Keep running $\mathsf{GetRecords}(\mathsf{UpdateState}(1^\lambda))$ until all values $y_1, \dots, y_n$ are published on the blockchain; when that happens check that $\mathsf{Verify}(pk_i, x, y_i) = 1$ for all $i \in [n]$. If any of these checks fails, send $\texttt{abort}$ to $f_\perp^{\mathsf{pct}}$, simulate A aborting, and terminate. Else, in case $\mathsf{Hash}(y_1||\cdots||y_n)$ was already set to $\omega' \neq \omega$, abort the simulation and terminate. In this last case, we say that the simulator fails to program the random oracle.
5. Upon input a random oracle query q from A, answer as follows:
   (a) If there exists a pair $(bid, \omega) \in \mathcal{L}$ for which it is possible to parse $q := y_1||\cdots||y_n$ so that $\mathsf{Verify}(pk_i, x, y_i) = 1$ for all $i \in [n]$—where $x = pk_1||\cdots||pk_n||sid||bid$ for the values $sid, (pk_1, \dots, pk_n)$ that appear in the simulation of branch bid—program $\mathsf{Hash}(q) := \omega$ and answer query q with ω. If ω was already given as output for a different query then we say that the simulator fails creating a collision when programming the random oracle.
   (b) Else, return a random value (maintaining consistency among repeated queries).

Notice that S simulates perfectly the messages written on the blockchain by the honest players in a real execution of π_{pct}^*, including their interaction with the blockchain protocol Γ. Moreover, in each branch bid, the simulator perfectly emulates an abort of the protocol due to the fact that A sends invalid signatures. Hence, the only difference between the real and ideal experiment is that in the latter, for each branch bid, the simulator forces the protocol output to be a fresh random value received from the ideal coin-tossing functionality. Consider the following events:

Event BAD$_1$: The event becomes true in case the simulator fails to program the random oracle in step 4b of the simulation.

---

[11]For simplicity, we assume that the player P_1 initiating the protocol is honest; if not, it is easy to adapt the simulation by having S using the value sid written by A on the blockchain.

Event BAD$_2$: The event becomes true in case, the simulator fails creating a collision as described in step 5a. This is possible when for a given branch bid, there exists an index $i \in \mathcal{I}$ such that the attacker produces two outputs y_i and y_i' such that $y_i \neq y_i'$ and $\mathsf{Verify}(pk_i, x, y_i) = \mathsf{Verify}(pk_i, x, y_i') = 1$, and queries the random oracle first using y_i' therefore obtaining ω, and then querying y_i therefore obtain again ω.

Event BAD$_3$: The event becomes true in case there exist two branches with different public keys, but for which the value bid is the same.

Let $\mathbf{BAD} = \mathbf{BAD}_1 \cup \mathbf{BAD}_2 \cup \mathbf{BAD}_3$. We claim that conditioning on $\overline{\mathbf{BAD}}$, the experiments $\mathbf{REAL}^{\Gamma,0}_{\pi,\mathsf{A},\mathsf{D}}(\lambda)$ and $\mathbf{IDEAL}^{\Gamma}_{f^{\mathsf{pct}}_{\perp},\mathsf{S},\mathsf{D}}(\lambda)$ are identical. This is because conditioning on $\mathbf{BAD}_1$ not happening, the attacker never queries the random oracle on $pk_1||\cdots||pk_n||sid||bid$ before the protocol execution on branch bid terminates, and thus the final output ω on that branch is uniformly distributed from the point of view of A (as it happens to be in the ideal world). Furthermore, conditioning on $\mathbf{BAD}_2$ and $\mathbf{BAD}_3$ not happening, the simulator correctly assigns a single random value ω to each branch identified by bid.

Next, we show that each of the above events happens at most with negligible probability. By a standard argument, this concludes the proof as the computational distance between $\mathbf{REAL}^{\Gamma,0}_{\pi,\mathsf{A},\mathsf{D}}(\lambda)$ and $\mathbf{IDEAL}^{\Gamma}_{f^{\mathsf{pct}}_{\perp},\mathsf{S},\mathsf{D}}(\lambda)$ is at most equal to the probability of event $\mathbf{BAD}$.

Lemma 4. *For all PPT* A, *there exists a negligible function* $\nu_1(\cdot)$ *such that* $\mathbb{P}\left[\mathbf{BAD}_1\right] \leq \nu_1(\lambda)$.

Proof. Notice that event $\mathbf{BAD}_1$ happens if and only if there exists a protocol execution corresponding to a branch bid for which the attacker A queries the random oracle upon input $y_1||\cdots||y_n$ before these values appear on the blockchain. Intuitively, this requires that A forges a signature for one of the public keys corresponding to one of the honest players, and thus $\mathbb{P}\left[\mathbf{BAD}_1\right]$ must be negligible. The reduction is straightforward: Given a PPT attacker provoking event $\mathbf{BAD}_1$ with non-negligible probability, we can construct a PPT attacker A' as follows. Initially, A' tries to guess the index $i \notin \mathcal{I}$ and the branch index $j \in \mathtt{poly}(\lambda)$ corresponding to the protocol execution in which A will provoke event $\mathbf{BAD}_1$. Hence, A' simulates the execution of protocol π^*_{pct} with A as done in the real experiment, except that on the j-th branch it sets pk_i to be the public key pk^* received from the challenger.

Finally, A' waits that A makes a random oracle query $y_1||\cdots||y_n$ such that $\mathsf{Verify}(pk_i, x, y_i) = 1$, where $x = pk_1||\cdots||pk_n||sid||bid$; if the latter does not happen, A' aborts, else it forwards (x, y_i) to the challenger. The proof follows by observing that, with non-negligible probability, A' does not abort, and thus it breaks unforgeability with non-negligible probability. $\square$

Lemma 5. *For all PPT* A, *there exists a negligible function* $\nu_2(\cdot)$ *such that* $\mathbb{P}\left[\mathbf{BAD}_2\right] \leq \nu_2(\lambda)$.

Proof. Notice that event $\mathbf{BAD}_2$ directly contradicts uniqueness of the signature scheme (Gen, $\mathsf{Sign}, \mathsf{Verify}$). Hence, $\mathbb{P}\left[\mathbf{BAD}_2\right]$ must be negligible. The reduction is straightforward: Given a PPT attacker provoking event $\mathbf{BAD}_2$ with non-negligible probability, we can construct a PPT attacker A' as follows. Initially, A' tries to guess the index $i \in \mathcal{I}$ and the branch index $j \in \mathtt{poly}(\lambda)$ corresponding to the protocol execution in which A will provoke event $\mathbf{BAD}_2$. Hence, A' simulates the execution of protocol π^*_{pct} with A as done in the real experiment.

Finally, A' waits that A publishes on the j-th branch the two values y_i, y_i' which make the event $\mathbf{BAD}_2$ become true; if the latter does not happen, A' aborts, else it forwards (pk_i, y_i, y_i') to the challenger where the public key pk_i is the public key corresponding to the i-th player on the j-th branch. The proof follows by observing that, with non-negligible probability, A' does not abort, and thus it breaks uniqueness with non-negligible probability. $\square$

Lemma 6. *For all PPT* A, *there exists a negligible function* $\nu_3(\cdot)$ *such that* $\mathbb{P}\left[\mathbf{BAD}_3\right] \leq \nu_3(\lambda)$.

Proof. Notice that event $\mathbf{BAD}_3$ directly contradicts collision resistance of the hash function Hash. Hence, $\mathbb{P}[\mathbf{BAD}_3]$ must be negligible. The reduction is straightforward: Given a PPT attacker provoking event $\mathbf{BAD}_3$ with non-negligible probability, we can construct a PPT attacker A' that simply emulates a protocol execution with A as in the real experiment. The values bid, as well as the answers to A's random oracle queries, are obtained by querying the target random oracle. Hence, whenever A' finds a fork with two different branches $\mathcal{B}$ and $\mathcal{B}'$ such that $\mathsf{Hash}(\mathcal{B}) = \mathsf{Hash}(\mathcal{B}')$, it outputs $(\mathcal{B}, \mathcal{B}')$ and stops. Since A provokes event $\mathbf{BAD}_3$ with non-negligible probability, A' wins with the same probability. This concludes the proof. $\qquad\square$

Putting the above lemmas together, by a union bound, there exists a negligible function $\nu(\cdot)$ such that $\mathbb{P}[\mathbf{BAD}] \leq \nu(\lambda)$, as desired. $\qquad\square$

Fairness with penalties. We now discuss how to augment the protocol π^*_{pct} in order to achieve fairness with penalties. First of all, each party should publish also a deposit along with her public key on the blockchain. The deposit can be redeemed by showing a valid signature on the value $x = pk_1||\cdots||pk_n||sid||bid$.

Assume that P_n is corrupted. The adversary can wait that the honest parties publish their value $y_1, \ldots, y_{n-1}$ on a given branch, and thus locally compute the output $\mathsf{Hash}(y_1||\cdots||y_n)$, where y_n is P_n's signatures on x corresponding to public key pk_n. Now, if P_n does not like the output it can either: (i) publish y_n in any case , or (ii) decide not to publish y_n. In case (i), P_n plays honestly, takes back his deposit and every player obtains the output. In case (ii), P_n aborts the protocol, but loses her deposit.

Note that the penalties mechanism for our protocol is too sophisticated for the scripting language used in Bitcoin. Instead in Ethereum we can design a smart contract to define the PCT protocol, having fairness with penalties and without penalizing the efficiency. In Appendix 6.2 we give details about how the smart contract works.

We call $\tilde{\pi}^*_{\mathsf{pct}}$ the fair (with penalties) version of protocol π^*_{pct}.

The informal description of the smart contract used in $\tilde{\pi}^*_{\mathsf{pct}}$ is given in Fig. 4 and the protocol is described below:

(i) *Setup phase:* At the beginning, one of the players creates the smart contract. When the contract is posted on the blockchain, the constructor automatically generates a unique session identifier sid.

(ii) *Deposit phase:* For each $i \in [n]$, P_i can decide to participate to the PCT protocol by triggering the function `deposit` to send a safety deposit and his public key pk_i of an unique signature scheme. After `time1` blocks have passed, if $(pk_1, \ldots, pk_n)$ are collected by the smart contract, it computes bid as $\mathsf{Hash}(\mathcal{B})$, where $\mathcal{B}$ is the blockchain that contains $(pk_1, \ldots, pk_n)$. The deposit phase ends and parties can start to redeem their deposit.

(iii) *Claim phase:* For each $i \in [n]$, P_i can claim his deposit back by triggering the function `claim` of the smart contract and sending a value y_i such that $\mathsf{Verify}(pk_i, x, y_i) = 1$, where $x = pk_1||\cdots||pk_n||sid||bid$, and pk_i is the public key of P_i. After `time2` blocks have passed, the claim phase ends and the smart contract computes and publishes the output as $\mathsf{Hash}(y_1||\cdots||y_n)$.

Theorem 4. *Assuming that* $(\mathsf{Gen}, \mathsf{Sign}, \mathsf{Verify})$ *is a unique signature scheme, the protocol* $\tilde{\pi}^*_{\mathsf{pct}}$ *described in Fig. 4 securely realize* f^{pct} *and satisfies fairness with penalties in the presence of quick players and malicious adversaries, in the programmable random oracle model.*

Proof Sketch. The privacy of of the protocol is proven by following the same outline of Theorem. 3.

We have to prove that $\tilde{\pi}^*_{\mathsf{pct}}$ is fair (with penalties).

The **Parallel Coin Tossing Smart Contract** runs with players $\mathsf{P}_1, \ldots, \mathsf{P}_n$ and consists of two main functions `deposit` and `claim` and two fixed timestamps `time1`,`time2` and a session id sid.

Deposit Phase: In round t_1, when `deposit`(pk_i) together with d coins is triggered from a party P_i, store (i, pk_i). Then, if $(pk_1, \ldots, pk_n)$ are stored, compute and store $bid := \mathsf{Hash}(\mathcal{B})$ and proceed to the Claim Phase. Otherwise, for all i, if the message (i, pk_i) has been stored, send back d coins to P_i and terminate.

Claim Phase: In round t_2, when `claim`(i, y_i) is triggered from P_i, check if $\mathsf{Verify}(pk_i, x, y_i) = 1$, where $x = pk_1 || \cdots || pk_n || sid || bid$. If the check is correct send d coins back to P_i.

Compute Phase: If, after `time2`, all the y_i are correctly claimed, compute and publish $\mathsf{Hash}(y_1, \ldots, y_n)$.

Figure 4: Smart contract for parallel coin tossing.

There are four possible scenarios that can happen and in each of them either all parties learn the output or the adversary A loses her deposit. An output out of $\tilde{\pi}_{\mathsf{pct}}^*$ is considered valid if it is confirmed on the blockchain. The three scenarios are described below.

Scenario 1: A does not exploit branches to play different public keys on different execution of $\tilde{\pi}_{\mathsf{pct}}^*$. Fairness follows from the smart contract execution (i.e., if A does not provide the signature of x, A will lose her deposit).

Scenario 2: There is the following time-line: there is a fork with two branches b_1 and b_2 and the Setup Phase is published before the the fork, but the Deposit Phase is executed after the fork. A aborts (i.e., A does not provide the signature of x) the execution of $\tilde{\pi}_{\mathsf{pct}}^*$ in b_1 and exploits b_2 in the following way: in b_2 a corrupted player P_i double spends (for any kind of transaction) the coins deposited in b_1. In this case, either b_1 gets confirmed, thus, boiling down to Scenario 1, or b_2 get confirmed and all the transactions sent to the smart contract of $\tilde{\pi}_{\mathsf{pct}}^*$ in branch b_1 are not valid b_2 since the deposit of P_i is previously spent in another transaction. It guarantees fairness since in b_1 the adversary is punished, and in b_2 there's no execution, and so no valid output

Scenario 3: This scenario follows the same time-line of Scenario 2. A aborts (i.e., A does not provide the signature of x) the execution of $\tilde{\pi}_{\mathsf{pct}}^*$ in b_1 and exploits b_2 in the following way: a corrupted player P_i publishes a different public key pk_i in b_2 (wlog., we analyze the case with two executions but it can be extended to multiple executions). In this case, the execution of $\tilde{\pi}_{\mathsf{pct}}^*$ is restarted from the beginning of the Deposit Phase in b_2 and the output that A learned on b_1 is not valid anymore. If b_1 gets confirmed it boils down to Scenario 1, otherwise the honest players learn the valid output computed in b_2. It guarantees fairness because in b_1 the adversary is punished, while in b_2 all parties will learn the output.

Scenario 4: There is the following time-line: there is a fork with two branches b_1 and b_2 and the Deposit Phase is published before the the fork, but the Claim Phase is executed after the fork. A corrupted party P_i sends $y_i = \mathsf{Sign}(sk_i, x)$ in the Claim Phase in b_1 and disliking the output out of the execution of $\tilde{\pi}_{\mathsf{pct}}^*$ in b_1, P_i exploits b_2 to send $y_i' \neq y_i$. We note that if y_i' is a valid signature for x under secret key sk_i we can create an adversary A' breaking the uniqueness of the signature scheme. It means that y_i' cannot be a valid signature. If b_1 gets confirmed P_i is not penalized and every player will learn out, otherwise if b_2 gets confirmed P_i is penalized and no party will learn the output.

$\square$

Experimental Evaluation.   We also provide in App. 6 some experiments to show noticeable improvements of our PCT protocol with respect to the lottery protocol of Andrychowicz *et al.*. in terms of number of blocks needed for completion of the protocol and gas consumption of the smart contracts.

6   Experimental Evaluation

We evaluate the efficiency of a generic protocol π^* obtained compiling an UC-secure n-party ρ-round MPC protocol π both in the case of MPC with aborts and fairness with penalties. Moreover, we evaluate the efficiency of π^*_{pct} compared with the protocol from [ADMM14, ADMM16].

To evaluate the efficiency in the best case we consider the following assumptions:
- Transactions in the last k blocks are considered not confirmed yet.
- All parties send the message at round i of the protocol as soon as they read all messages form round $i - 1$ on $\mathcal{B}_i^{\lceil k}$, where k is 0 in case of quick executions.
- Whenever a player broadcasts a transaction, it appears in the next block.
- All messages in a round of the MPC protocol fit in a single block.

In case of non-quick executions if we have a ρ-round MPC protocol π running on the blockchain, the number of blocks needed to complete the execution with the previous assumption is $\rho \cdot k$.

6.1   Analysis

We now describe the cost of our naive compiler, and give a comparison between our coin tossing protocol and the one of Andrychowicz et al. In particular we will interpret the execution time as the number of blocks needed to complete the steps of each protocol. In the next section, we will describe our smart contract implementation. To allow for a fair comparison between our coin tossing protocol and the one presented in [ADMM14, ADMM16], we implemented both of protocols in Solidity using Ethereum smart contracts.[12] See Fig. 6 and Fig. 7 for the corresponding code.

Generic compiler.   Given an n-party ρ-round UC-secure MPC protocol π, we evaluate the efficiency of a protocol π^* obtained compiling π with the compiler described in Fig. 1. Our compiler adds only one round to the execution of the protocol, in which the parties publish their encryption keys of the underlying encryption scheme and signature keys of the signature scheme. We analyze the number of block needed to end π^* in case of standard MPC with aborts and in case of fairness with penalties. As we noted in §4.3, to obtain fairness with penalties in π^*, the players of the protocol have to wait that the first round (the deposit) became final on the blockchain before continuing to run the protocol with a quick execution.

Let us consider the case of MPC with aborts first. Since in this setting the protocol can be run a full-quick mode, if the underlying protocol π has ρ rounds, then π^* will have $\rho + 1$ rounds and the number of blocks needed to end the computation is $\rho + 1$.

In case of fairness with penalties, the execution time will be $\rho + w$ blocks, where $w \geq k$ is the liveness parameter. Since we assume that in the ideal conditions all the players broadcast the deposit message at the same time, their deposit will be posted and confirmed after at most w blocks.

---

[12]Notice that the average time for a new block to appear is around 15 seconds[eth].

Lottery and parallel coin tossing. For both the protocols, in the deposit phase, a timeout $\bar{t}$ must be provided by the contract creator, so that players have time to send their deposits together with the corresponding additional information required by the protocol. In the ideal conditions described above, we are allowed to stretch this timeout to only one block. The same argument applies to the opening phase of [ADMM14, ADMM16]. A comparison is described below:

- *Lottery:* Due to the expressiveness of smart contracts, our implementation of [ADMM14, ADMM16] requires one step less than the original implementation using in bitcoin. Specifically, we can embed the betting phase in the commit/deposit phase, by just requiring that the players deposit 1 more coin. Since in their setting block confirmation is required at each step, the overall execution takes exactly $3 \cdot w$ blocks (including one round for posting the smart contract).
- *Our PCT:* As proven in §5.3, our coin tossing protocol can be executed in full-quick mode. The entire execution lasts 3 blocks in total (including 1 block for posting the contract). We point out that in the most pessimistic setting, where all the messages will appear to the state of the honest player after w blocks for each step, the overall execution takes $3 \cdot w$, as much as [ADMM16].

6.2 Smart Contracts

We now describe our smart contract implementations in details, and comment on the relative gas consumption.[13]

Lottery protocol by [ADMM16]. The smart contract execution is described as follows:
- *Setup phase:* A player publishes the smart contract in Fig. 7 on page 37, indicating in the constructor the addresses of the players' wallet, and the committing phase timeout `time1` and opening phase timeout `time2`.
- *Committing phase:* The players trigger the `commit` function upon input the commitment to some value and a deposit of `minDep` $= n(n-1)+1$, where $n(n-1)$ coins are used for the penalty mechanism and 1 coin is used to put money for the lottery. After `time1` blocks, the n commitments are collected and the committing phase ends.
- *Opening phase:* All the participants open their commitment by triggering the `openCom` function, and the winner can then claim his bet by triggering the `claimWinner` function of the smart contract (if all the parties have opened).
- *Compensation phase:* If, after `time2` blocks, some player did not open his commitment, the function `payDeposit` can be triggered, so that all player that hasn't open before `time2` will be penalized and the players who had opened receive their deposit back together with a fraction of the adversaries' deposit.

Our coin tossing protocol. The description the smart contract execution of π^*_{pct} on Fig. 6 works as follows:
- *Setup phase:* At the beginning, one of the players creates the smart contract specifying a deposit amount `minDep` and timeout `time1`. When the contract is posted on the blockchain, the constructor automatically generates a unique session identifier *sid* by triggering `generateSid`.
- *Deposit phase:* For each $i \in [n]$, P_i can decide to participate to the PCT protocol by triggering the function `deposit` to send a safety deposit and his public key pk_i for an unique signature scheme. After that $(pk_1, \ldots, pk_n)$ are collected by the smart contract

---

[13]For our purposes we are not concentrating on optimizing the code.

and `time1` blocks are passed, the deposit phase ends parties can start to redeem their deposit. [14]

- *Claim phase:* During this phase, each player P_i can claim his deposit back by triggering the function `claim` of the smart contract and sending a value y_i such that $\mathsf{Verify}(pk_i, x, y_i) = 1$, where $x = pk_1 || \cdots || pk_n || sid || bid$, pk_i is the public key of P_i. For the signature verification, we can use a unique signature scheme with fast verification like BLS Signatures [BLS01], invoked with `BLSVerify` in the code, or RSA-FDH [BR96]. [15]

GAS consumption. As it can be seen in Fig. 5, PCT is more expensive in terms of GAS consumption than Lottery. It is well motivated by the fact that Lottery uses only hash function to compute the commitments and no other expensive computations. Our PCT protocol needs also unique signatures. It can be seen anyway as an affordable cost to pay to achieve efficiency still maintaining the same security guarantees.


Figure 5: GAS consumption comparison between our smart contract implementation of PCT (§5.3) protocol and Lottery (§5.1).

7 Acknowledgments

---

[14]Note that when some party P_i sends his deposit to the smart contract, the variable *bid* is updated with the hash of the last block of the contract state (uniquely identifying the branch in which the smart contract state is updated). This implies that *bid* is fixed after the last player sends his public key.

[15]We do not explicitly implement this signature, but solidity implementations are available for testing [sol].

```solidity
1   pragma solidity ^0.4.0;
2
3   contract ParallelCoinTossing {
4       struct Player {
5           bool isPlaying;
6           bool hasClaimed;
7           string pk;
8           uint d; //Player's deposit
9           uint c; //Player's claim
10          }
11      address[] playersAddr;
12      mapping(address => Player) players;
13      uint sid, time1, time2;
14      bytes32 bid;
15
16      //flags
17      bool claimPhase = false; //true if the claimPhase starts
18
19      //common message to be signed
20      uint x;
21
22      constructor(uint _time1, uint _time2) public   {
23          sid = generateSid(); //session id
24          time1 = _time1; //first timelock
25          time2 = _time2; //second timelock
26      }
27      function deposit(string pubKey) public payable {
28          require (msg.sender.balance >= minDep && msg.value >= minDep && players[msg.sender].d == 0 && now
                  < time1);
29          playersAddr.push(msg.sender); //add the public key of the current sender
30          Player p = players[msg.sender];
31          p.isPlaying = true;
32          p.pk = pubKey;
33          p.d = msg.value;   //msg.value is the deposit value of the player
34          bid = block.blockhash(now); //Every time he receives a public key, it updates the blockhash, so
                  that the correct bid is the blockchain state of the last public key deposited.
35          }
36      function claim(uint y) public {
37          require (claimPhase && now < time2 && players[msg.sender].isPlaying && !players[msg.sender].
                  hasClaimed && BLSVerify(players[msg.sender].pk,x,y));
38          Player p = players[msg.sender];
39          p.c = y;
40          p.hasClaimed = true;
41          msg.sender.transfer(p.d);
42  }
43
44      //automatic check functions run after a certain time
45      function checkDeposit() public {
46          require (!claimPhase && now >= time1);
47          uint n = playersAddr.length;
48          x = sha3(player[playersAddr[0]].pk||...||players[playersAddr[n-1]].pk||sid||bid);
49          claimPhase = true;
50          }
51  }
```

Figure 6: Pseudocode implementation of our smart contract for realizing parallel coin tossing. For simplicity, we omit an explicit definition of the concatenation function in the computation of x.

```solidity
1   pragma solidity >=0.4.21 <0.6.0;
2
3   contract FairLottery {
4     struct Player {
5       address addr;
6       bool hasCommitted, hasOpened, isPlaying;
7       uint balance, index;
8       bytes32 com;
9       int opn;
10      }
11    uint public n, time1, time2;
12    address[] addresses;
13    mapping (address => Player) players;
14
15  constructor(address[] _addresses, uint _time1, uint _time2) public { //creates a new instance of the
          lottery for a set of prescribed players
16      addresses = _addresses;
17      for (uint i = 0; i < addresses.length; i++) {
18        Player p = players[addresses[i]];
19        p.isPlaying = true;
20        p.index = i;
21        }
22      n = addresses.length;
23      time1 = _time1;
24      time2 = _time2;
25      }
26    function commit(bytes32 _com) public payable { //sha3 value commit (n*(n-1) coins + 1 coin for bet)
27      require(msg.value >= (n*(n-1)+1)) &&    players[msg.sender].isPlaying && !players[msg.sender].
            hasCommitted);
28      Player p = players[msg.sender];
29      p.com = _com;
30      p.hasCommitted = true;
31      p.balance = msg.value;
32      }
33    function openCom(int openVar) public   { //opening of the commitment
34      require(players[msg.sender].hasCommitted && now > time1 && now < time2 && !players[msg.sender].
            hasOpened && sha3(openVar) == players[msg.sender].com);
35      Player p = players[msg.sender];
36      p.hasOpened = true;
37      msg.sender.transfer(n*(n-1)); //pays the sender back
38      }
39    function payDeposit() public { //compensation function
40      require(players[msg.sender].isPlaying && now >= time2);
41      uint index = players[msg.sender].index;
42        for (uint i = 0; i < n; i++)
43          if (i != index && !players[addresses[i]].hasOpened) msg.sender.transfer(players[msg.sender].
              balance/n); //the player msg.sender get is compensation of n coins
44      }
45    function claimWinner(uint[] secrets) public { //function triggered by the winner
46      require (secrets.length == n && checkWinner(secrets,msg.sender));
47      msg.sender.transfer(n);   //redeem the won coins
48  }
49  //private local functions
50    function checkWinner(uint[] secrets, address _sender) private returns (bool) {
51      int sum = 0;
52        for (uint i = 0; i < secrets.length; i++) {
53          if (sha3(secrets[i] != players[addresses[i]].com) return false;
54          sum += secrets[i];
55        }
56      if ((sum%n != players[_sender].index))
57        return false;
58      return true;
59      }
60    }
```

Figure 7: Pseudocode implementation of the lottery protocol by Andrychowicz *et al.* [ADMM16], when using smart contracts.

References

[ADMM14] 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, 2014.

[ADMM16] Marcin Andrychowicz, Stefan Dziembowski, Daniel Malinowski, and Lukasz Mazurek. Secure multiparty computations on bitcoin. *Commun. ACM*, 59(4):76–84, 2016.

[BG17] Vitalik Buterin and Virgil Griffith. Casper the friendly finality gadget. *CoRR*, abs/1710.09437, 2017.

[BK14] Iddo Bentov and Ranjit Kumaresan. How to use bitcoin to design fair protocols. In *CRYPTO*, pages 421–439, 2014.

[BKM17] Iddo Bentov, Ranjit Kumaresan, and Andrew Miller. Instantaneous decentralized poker. In *ASIACRYPT*, pages 410–440, 2017.

[BLS01] Dan Boneh, Ben Lynn, and Hovav Shacham. Short signatures from the weil pairing. In *ASIACRYPT*, pages 514–532, 2001.

[BMTZ17] Christian Badertscher, Ueli Maurer, Daniel Tschudi, and Vassilis Zikas. Bitcoin as a transaction ledger: A composable treatment. In *CRYPTO*, pages 324–356, 2017.

[BR96] Mihir Bellare and Phillip Rogaway. The exact security of digital signatures - how to sign with RSA and Rabin. In *EUROCRYPT*, pages 399–416, 1996.

[BW] Bitcoin Wiki. Confirmation in bitcoin. `https://en.bitcoin.it/wiki/Confirmation`.

[CGJ19] Arka Rai Choudhuri, Vipul Goyal, and Abhishek Jain. Founding secure computation on blockchains. In *EUROCRYPT*, pages 351–380, 2019.

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

[CPS18] T.-H. Hubert Chan, Rafael Pass, and Elaine Shi. Pala: A simple partially synchronous blockchain. *IACR Cryptology ePrint Archive*, 2018.

[DY05] Yevgeniy Dodis and Aleksandr Yampolskiy. A verifiable random function with short proofs and keys. In *PKC*, pages 416–431, 2005.

[eth] The ethereum average block time chart. `https://etherscan.io/chart/blocktime`. Accessed: 2020-06-11.

[GG17] Rishab Goyal and Vipul Goyal. Overcoming cryptographic impossibility results using blockchains. In *TCC*, pages 529–561, 2017.

[GHM+17] Yossi Gilad, Rotem Hemo, Silvio Micali, Georgios Vlachos, and Nickolai Zeldovich. Algorand: Scaling byzantine agreements for cryptocurrencies. In *Proceedings of the 26th Symposium on Operating Systems Principles, Shanghai, China, October 28-31, 2017*, pages 51–68, 2017.

[GKL15] Juan A. Garay, Aggelos Kiayias, and Nikos Leonardos. The bitcoin backbone protocol: Analysis and applications. In *EUROCRYPT*, pages 281–310, 2015.

[GS09]   Vipul Goyal and Amit Sahai. Resettably secure computation. In *EUROCRYPT*, pages 54–71, 2009.

[KB14]   Ranjit Kumaresan and Iddo Bentov. How to use bitcoin to incentivize correct computations. In *ACM CCS*, pages 30–41, 2014.

[KB16]   Ranjit Kumaresan and Iddo Bentov. Amortizing secure computation with penalties. In *ACM CCS*, pages 418–429, 2016.

[KMB15]   Ranjit Kumaresan, Tal Moran, and Iddo Bentov. How to use bitcoin to play decentralized poker. In *ACM CCS*, pages 195–206, 2015.

[KMS+16]   Ahmed E. Kosba, Andrew Miller, Elaine Shi, Zikai Wen, and Charalampos Papamanthou. Hawk: The blockchain model of cryptography and privacy-preserving smart contracts. In *IEEE S&P*, pages 839–858, 2016.

[KVV16]   Ranjit Kumaresan, Vinod Vaikuntanathan, and Prashant Nalini Vasudevan. Improvements to secure computation with penalties. In *ACM CCS*, pages 406–417, 2016.

[KZZ16]   Aggelos Kiayias, Hong-Sheng Zhou, and Vassilis Zikas. Fair and robust multi-party computation using a global transaction ledger. In *EUROCRYPT*, pages 705–734, 2016.

[Lys02]   Anna Lysyanskaya. Unique signatures and verifiable random functions from the DH-DDH separation. In *CRYPTO*, pages 597–612, 2002.

[MMNT19]   Bernardo Magri, Christian Matt, Jesper Buus Nielsen, and Daniel Tschudi. Afgjort: A partially synchronous finality layer for blockchains. *IACR Cryptology ePrint Archive*, 2019.

[MRV99]   Silvio Micali, Michael O. Rabin, and Salil P. Vadhan. Verifiable random functions. In *IEEE FOCS*, pages 120–130, 1999.

[PS17]   Rafael Pass and Elaine Shi. Hybrid consensus: Efficient consensus in the permissionless model. In *DISC*, pages 39:1–39:16, 2017.

[PS18]   Rafael Pass and Elaine Shi. Thunderella: Blockchains with optimistic instant confirmation. In *EUROCRYPT*, pages 3–33, 2018.

[PSS17]   Rafael Pass, Lior Seeman, and Abhi Shelat. Analysis of the blockchain protocol in asynchronous networks. In *EUROCRYPT*, pages 643–673, 2017.

[RT]   Randao's Team. Randao: A dao working as rng of ethereum. `https://github.com/randao/randao`.

[sol]   Solidity bls signatures. `https://github.com/kfichter/solidity-bls`. Smart contract implementation of BLS Signatures.

[SSV19]   Alessandra Scafuro, Luisa Siniscalchi, and Ivan Visconti. Publicly verifiable proofs from blockchains. In *PKC*, pages 374–401, 2019.