

Secure Multiparty Computations on BitCoin

Marcin Andrychowicz, Stefan Dziembowski, Daniel Malinowski and Łukasz Mazurek
University of Warsaw

Abstract

BitCoin is a decentralized digital currency, introduced in 2008, that has recently gained noticeable popularity. Its main features are: (a) it lacks a central authority that controls the transactions, (b) the list of transactions is publicly available, and (c) its syntax allows more advanced transactions than simply transferring the money. The goal of this paper is to show how these properties of BitCoin can be used in the area of secure multiparty computation protocols (MPCs).

Firstly, we show that the BitCoin system provides an attractive way to construct a version of “timed commitments”, where the committer has to reveal his secret within a certain time frame, or to pay a fine. This, in turn, can be used to obtain fairness in some multiparty protocols. Secondly, we introduce a concept of multiparty protocols that work “directly on BitCoin”. Recall that the standard definition of the MPCs guarantees only that the protocol “emulates the trusted third party”. Hence ensuring that the inputs are correct, and the outcome is respected is beyond the scope of the definition. Our observation is that the BitCoin system can be used to go beyond the standard “emulation-based” definition, by constructing protocols that link their inputs and the outputs with the real BitCoin transactions.

As an instantiation of this idea we construct protocols for secure multiparty lotteries using the BitCoin currency, without relying on a trusted authority (one of these protocols uses the BitCoin-based timed commitments mentioned above). Our protocols guarantee fairness for the honest parties no matter how the looser behaves. For example: if one party interrupts the protocol then her money is transferred to the honest participants. Our protocols are practical (to demonstrate it we performed their transactions in the actual BitCoin system), and can be used in real life as a replacement for the online gambling sites. We think that this paradigm can have also other applications. We discuss some of them.

I. INTRODUCTION

Secure multiparty computation (MPC) protocols, originating from the seminal works of Yao [37] and Goldreich et al. [25], allow a group of mutually distrusting parties to compute a joint function f on their private inputs. Typically, the security of such protocols is defined with respect to the *ideal model* where f is computed by a trusted party T_f . More precisely: it is required that during the execution of a protocol the parties cannot learn more information about the inputs of the other participants than they would learn if f was computed by T_f who: (a) receives the inputs from the parties, (b) computes f , and (c) sends the output back to the parties. Moreover, even if some parties misbehave and do not follow the protocol, they should not be able to influence the output of the honest parties more than they could in the ideal model by modifying their own inputs.

As an illustration of the practical meaning of such security definition consider the case when there are only two participants, called Alice and Bob, and the function that they compute is a conjunction $f_\wedge(a, b) = a \wedge b$, where $a, b \in \{0, 1\}$ are Boolean variables denoting the inputs of Alice and Bob, respectively. This is sometimes called the *marriage proposal problem*, since one can interpret the input of each party as a declaration if she/he wants to marry the other one. More precisely: suppose $a = 1$ if and only if Alice wants to marry Bob, and $b = 1$ if and only if Bob wants to marry Alice. In this case $f_\wedge(a, b) = 1$ if and only if *both* parties want to marry each other,

and hence, if, e.g., $b = 0$ then Bob has no information about Alice’s input. Therefore the privacy of Alice is protected. One can also consider randomized functions f , the simplest example being the *coin tossing problem* [6] where the computed function $f_{\text{rnd}} : \{\perp\} \times \{\perp\} \rightarrow \{0, 1\}$ takes no inputs, and outputs a uniformly random bit. Yet another generalization are the so-called *reactive functionalities* where the trusted party T maintains a state and the parties can interact with T in several rounds. One example of such a functionality is the *mental poker* [35] where T simulates a card game, i.e. she first deals a deck of cards and then ensures that the players play the game according to the rules.

It was shown in [25] that for any efficiently-computable function f (or, more general, any reactive functionality) there exists an efficient protocol that securely computes it, assuming the existence of the trapdoor-permutations. If the minority of the parties is malicious (i.e. does not follow the protocol) then the protocol always terminates, and the output is known to each honest participant. If not, then the malicious parties can terminate the protocol after learning the output, preventing the honest parties from learning it. It turns out [16] that in general this problem, called the lack of *fairness*, is unavoidable, although there has been some effort to overcome this impossibility result by relaxing the security requirements [26], [11], [4], [31]. Note that in case of the two-player protocols it makes no sense to assume that the majority of the players is honest, as this would simply mean that none of the players is malicious. Hence, the two-party protocols in general do not provide complete fairness (unless the security definition is weakened).

Since the introduction of the MPCs there has been a significant effort to make these protocols efficient [28], [5], [17] and sometimes even to use them in the real-life applications such as, e.g., the online auctions [22]. On the other hand, perhaps surprisingly, the MPCs have not been used in many other areas where seemingly they would fit perfectly. One prominent example is the internet gambling: it may be intriguing that currently gambling over the internet is done almost entirely with the help of the web-sites that play the roles of the “trusted parties”, instead of using the coin flipping or the mental poker protocols. This situation is clearly unsatisfactory from the security point of view, especially since in the past there were cases when the operators of these sites abused their privileged position for their own financial gain (see e.g. [32]). Hence, it may look like the multiparty techniques that eliminate the need for a trusted party would be a perfect replacement for the traditional gambling sites (an additional benefit would be a reduced cost of gambling, since the gambling sites typically charge fees for their service).

In our opinion there are at least two main reasons why the MPCs are not used for online gambling. The first reason is that the multiparty protocols do not provide fairness in case there is no honest majority among the participants. Consider for example a simple two-party lottery based on the coin-tossing protocol: the parties first compute a random bit b , if $b = 0$ then Alice pays \$1 to Bob, if $b = 1$ then Bob pays \$1 to Alice, and if the protocol did not terminate correctly then the parties do not pay any money to each other. In this case a malicious party, say Alice, could prevent Bob from learning the output if it is equal to 0, making 1 the only possible output of a protocol. Since this easily generalizes to the multiparty case, hence it is clear that the gambling protocol would work only if the majority is honest, which is not a realistic assumption in the fully distributed internet environment (there are many reasons for this, one of them being the *sybil* attacks where one malicious party creates and controls several “fake” identities, easily obtaining the “majority” among the participants).

The second reason is even more fundamental, as it comes directly from the inherent limitations of the MPC security definition, namely: such protocols do not provide security beyond the trusted-party emulation. This drawback of the MPCs is rarely mentioned in the literature as it seems obvious that

in most of the real-life applications cryptography cannot be “responsible” for controlling that the users provide the “real” input to the protocol and that they respect the output. Consider for example the marriage proposal problem: it is clear that even in the ideal model there is no technological way to ensure that the users honestly provide their input to the trusted party, i.e. nothing prevents one party, say Bob, to lie about his feelings, and to set $b = 1$ in order to learn Alice’s input a . Similarly, forcing both parties to respect the outcome of the protocol and indeed marry cannot be guaranteed in a cryptographic way. This problem is especially important in the gambling applications: even in the simplest “two-party lottery” example described above, there exists no cryptographic method to force the loser to transfer the money to the winner.

One pragmatic solution to this problem, both in the digital and the non-digital world is to use the concept of “reputation”: a party caught on cheating (i.e. providing the wrong input or not respecting the outcome of the game) damages her reputation and next time may have trouble finding another party willing to gamble with her. Reputation systems have been constructed and analyzed in several papers (see, e.g. [33] for an overview), however they seem too cumbersome to use in many applications, one reason being that it is unclear how to define the reputation of the new users in the scenarios when the users are allowed to pick new names whenever they want [23].

Another option is to exploit the fact that the financial transactions are done electronically, and hence one could try to “incorporate” the final transaction (transferring \$1 from the loser to the winner) into the protocol, in such a way that the parties learn who won the game only when the transaction has already been performed. It is unfortunately not obvious how to do it within the framework of the existing electronic cash systems. Obviously, since the parties do not trust each other, we cannot accept solutions where the winning party learns e.g. the credit card number, or the account password of the loser. One possible solution would be to design a multiparty protocol that simulates, in a secure way, a simultaneous access to all the online accounts of the participants and executes a wire transfers in their name.¹ Even if theoretically possible, this solution is clearly very hard to implement in real life, especially since the protocol would need to be adapted to several banks used by the players (and would need to be updated whenever they change). The same problems occur obviously also if above we replace the “bank” with some other financial service (like PayPal). One could consider using Chaum’s Ecash [13], or one of its variants [14], [12]. Unfortunately, none of these systems got widely adopted in real-life. Moreover, they are also bank-dependent, meaning that even if they get popular, one would face a challenge of designing a protocol that simulates the interaction of a real user with a bank, and make it work for several different banks.

We therefore turn our attention to BitCoin, which is a decentralized digital currency introduced in 2008 by Satoshi Nakamoto² [30]. BitCoin has recently gained a noticeable popularity (its current market capitalization is over 5 billion USD) mostly due to its distributed nature and the lack of a central authority that controls the transactions. Because of that it is infeasible for anyone to take control over the system, create large amounts of coins (to generate inflation), or shut it down. The money is transferred directly between two parties — they do not have to trust anyone else and transaction fees are zero or very small. Another advantage is the pseudo-anonymity³ — the users

¹Note that this would require, in particular, “simulating” the web-browser and the SSL sessions, since each individual user should not learn the contents of the communication between the “protocol” and his bank, as otherwise he could interrupt the communication whenever he realizes that the “protocol” ordered a wire transfer from his account. Moreover, one would need to assume that the transactions cannot be cancelled once they were ordered.

²This name is widely believed to be a pseudonym.

³A very interesting modification of BitCoin that provides real cryptographic anonymity has been recently proposed in [29].

are only identified by their public keys that can be easily created, and hence it is hard to link the real person with the virtual party spending the money. In Section II we describe the main design principles of BitCoin, focusing only on the most relevant parts of this system. A more detailed description can be found in Nakamoto’s original paper [30], the BitCoin wiki webpage en.bitcoin.it (sections particularly relevant to our work are: “Transactions” and “Contracts”), or other papers on BitCoin [29], [15], [2], [34]. In the sequel “ ₿ ” denotes the BitCoin currency symbol.

A. Our contribution

We study how to do “MPCs on BitCoin”. First of all, we show that the BitCoin system provides an attractive way to construct a version of “timed commitments” [7], [24], where the committer has to reveal his secret within a certain time frame, or to pay a fine. This, in turn, can be used to obtain fairness in certain multiparty protocols. Hence it can be viewed as an “application of BitCoin to the MPCs”.

What is probably more interesting is our second idea, which in some inverts the previous one by showing an “application of the MPCs to BitCoin”, namely we introduce a concept of multiparty protocols that work directly on BitCoin. As explained above, the standard definition of the MPCs guarantees only that the protocol “emulates the trusted third party”. Hence ensuring that the inputs are correct, and the outcome is respected is beyond the scope of the definition. Our observation is that the BitCoin system can be used to go beyond the standard “emulation-based” definition, by constructing protocols that link the inputs and the outputs with the real BitCoin transactions. This is possible since the BitCoin lacks a central authority, the list of transactions is public, and its syntax allows more advanced transactions than simply transferring the money.

As an instantiation of this idea we construct protocols for secure multiparty lotteries using the BitCoin currency, without relying on a trusted authority. By “lottery” we mean a protocol in which a group of parties initially invests some money, and at the end one of them, chosen randomly, gets all the invested money (called the *pot*). Our protocols can work in purely peer-to-peer environment, and can be executed between players that are anonymous and do not trust each other. Our constructions come with a very strong security guarantee: no matter how the dishonest parties behave, the honest parties will never get cheated. More precisely, each honest party can be sure that, once the game starts, it will always terminate and will be fair.

Our two main constructions are as follows. The first protocol (Section IV) can be executed between any number of parties. Its security is obtained via the so-called *deposits*: each user is required to initially put aside a certain amount of money, which will be payed back to her once she completes the protocol honestly. Otherwise the deposit is given to the other parties and “compensates” them the fact that the game terminated prematurely. This protocol uses the timed commitment scheme described above. A certain drawback of this protocol is that the deposits need to be relatively large, especially if the protocol is executed among larger groups of players. More precisely to achieve security the deposit of each player needs to be $N(N - 1)$ times the size of the bet (observe that for the two-party case it simply means that the deposit is twice the size of the bet).

We also describe (in Section V) a protocol that does not require the use of deposits at all. This comes at a price: the protocol works only for two parties, and its security relies on an additional assumption (see Section V for more details).

The only cost that the participants need to pay in our protocols are the BitCoin transaction fees. The typical BitCoin transactions are currently free. However, the participants of our protocols need to make a small number of non-standard transactions (so-called “strange transactions”, see Section

II), for which there is usually some small fee (currently around $0.00005\text{฿} \approx 0.03\text{\$}$). To keep the exposition simple we initially present our results assuming that the fees are zero, and later, in Section VI, argue how to extend the definitions and security statements to take into account also the non-zero fees. For the sake of simplicity we also assume that the bets in the lotteries are equal to 1 ฿ . It should be straightforward to see how to generalize our protocols to other values of the bets.

Our constructions are based on the coin-tossing protocol of Blum [6]. We managed to adapt this protocol to our model, without the need to modify the current BitCoin system. We do not use any generic methods like the MPC or zero-knowledge compilers, and hence the protocols are very efficient. The only cryptographic primitives that we use are the commitment schemes, implemented using the hash functions (which are standard BitCoin primitives). Our protocols rely strongly on the advanced features of the BitCoin (in particular: the so-called “transaction scripts”, and “time-locks”). Because of the lack of space we only sketch the formal security definitions. The security proofs will appear in an extended version of this paper. We executed our transactions on the real BitCoin. We provide a description of these transactions and a reference to them in the BitCoin chain.

B. Applications and future work

Although, as argued in Section I-C below, it may actually make economic sense to use our protocols in practice, we view gambling mostly as a motivating example for introducing a concept that can be called “MPCs on BitCoin”, and which will hopefully have other applications. One (rather theoretical) example of such application is the “millionaires problem” where Alice and Bob want to establish who is richer.⁴ It is easy to see that Alice and Bob can (inefficiently) determine who has more coins by applying the generic MPC and zero-knowledge techniques. This is possible since the only inputs that are needed are (a) the contents of the BitCoin ledger (more precisely: its subset consisting of the non-redeemed transactions), which is public, and (b) Alice’s and Bob’s private keys used as their respective private inputs (see Section II for the definitions of these terms). Obviously, using this protocol makes sense only if, for some reason, each party is interested in proving that she is the richer one. This is because every participant can easily pretend to be poorer than she really is and “hide” his money by transferring it to some other address (that he also controls). Since we do not think that this protocol is particularly relevant to practical applications, we do not describe it in detail here. Let us only observe that, interestingly, this protocol is in some sense dual to the coin-tossing protocol, as it uses the BitCoin system to verify the correctness of the inputs, instead guaranteeing that the outcome is respected (as it is the case with the coin-tossing)⁵.

We think that analyzing what functionalities can be computed this way (taking into account the problem of the participants “pretending to be poorer than they really are”) may be an interesting research direction. Other possible future research directions are: constructing protocols secure against “malleability attacks” and “eavesdropping attacks” (see Sec. V for more details) that do not require the deposits, showing general completeness results, providing a more formal framework to analyze the deposit-based technique (this can probably be done using the tools from the “rational cryptography” literature [27], [1], [21]). An interesting research direction is also to construct

⁴The formal definition is as follows: let $a, b \in \mathbf{N}$ denote the amount of coins that Alice and Bob respectively own. In this case the parties compute the function $f_{\text{mill}} : \mathbf{N} \times \mathbf{N} \rightarrow \{A, B\}$ defined as: $f_{\text{mill}}(a, b) = A$ if and only if $a \geq b$ and $f_{\text{mill}}(a, b) = B$ otherwise.

⁵The reader may be tempted to think that a similar protocol could be used with the eCash [13]. This is not the case, as in eCash there is no method of proving in zero-knowledge that the money has not been spent (since the list of transactions is not public).

protocols for games other than simple lotteries (like card, or board games). We also think that our protocols can provide a good motivation for research in the formal analysis of the BitCoin schemes that involve transactions with non-standard scripts. Is it possible to formally specify and verify the properties of such protocols in the spirit of [18]? Such formal approach would probably need to involve some elements of the temporal logic (see e.g. [3]) to capture the properties provided by the time-locks.

C. Economic analysis

Besides of being conceptually interesting, we think that our protocols can have direct practical applications in the online gambling, which is a significant market: it is estimated that there are currently 1,700 gambling sites worldwide handling bets worth over \$4 billion per year [20]. Some of these sites are using BitCoin. The largest of them, *SatoshiDice*, has been recently sold for over \$12 million [10]. All of the popular sites charge a fee for their service, called the *house edge* (on top of the BitCoin transaction fees). Currently, the honesty of these sites can be verified only *ex post facto*: they commit to their randomness before the game starts and later prove to the public that they did not cheat. Hence, nothing prevents them from cheating and then disappearing from the market (using the MPC terminology: such protocols provide security only against the “covert adversaries” [36]). Of course, this means that the users need to continually monitor the behavior of the gambling sites in order to identify the honest ones. This system, called the “mathematically provable fairness” is described in a recent article [9], where it is advised to look on a particular page, called *Mem’s Bitcoin Gambling List*, to check the gambling sites’ reputation. This simple reputation system can of course be attacked in various ways. Moreover, one can expect that the sites with more established reputation will have a higher house edge, and indeed the *SatoshiDice* site charges more than the other, less well-known, sites. Currently *SatoshiDice* house edge is around 2% [9].

Compared to the gambling sites, our protocols have the following advantage. First of all, the security guarantee is stronger, as it does not depend on the honesty of any trusted party. Secondly, in our protocols there is obviously no “house edge”. On a negative side, the BitCoin transaction fees can be slightly larger in our case than in the case of the gambling sites (since we have more transactions, and some of them are “strange”). At the moment of writing this paper, using our solution is cheaper than using *SatoshiDice* for bets larger than, say, \$5, but of course whether our protocols become really widely used in practice depends on several factors that are hard to predict, like the value of the fees for the “strange transactions”.

We also note that, although our initial motivation was the peer-to-peer lottery, it can actually make a lot of sense for the online gambling services to use our solutions, especially the two-party protocol. Of course the business model of such services makes sense only if there is non-zero house edge. This is not a problem since our protocols can be easily used in lotteries where the expected payoff is positive for one party (in this case: the gambling service) and it is negative for the other one (the client). Such “provably guaranteed fairness” can be actually a good selling line for some of these services.

D. Related work

Most of the related work has been already described in the earlier sections. Previous papers on BitCoin analyze the BitCoin transaction graph [34], or suggest improvements of the current version of the system. This includes important work of Barber et al. [2] who study various security aspects of BitCoin and Miers et al. [29] who propose a BitCoin system with provable anonymity. Our

paper does not belong to this category, and in particular our solutions are fully compatible with the current version of BitCoin (except of the „malleability” and „eavesdropping” problem concerning the last protocol, Section V)

The work most relevant to ours in the BitCoin literature is probably Section 7.1 of [2] where the authors construct a secure “mixer”, that allows two parties to securely “mix” their coins in order to obtain unlinkability of the transactions. They also construct commitment schemes with time-locks, however some important details are different, in particular, in the normal execution of the scheme the money is at the end transferred to the receiver. Also, the main motivation of this work is different: the goal of [2] is to fix an existing problem in BitCoin (“linkability”), while our goal is to use BitCoin to perform tasks that are hard (or impossible) to perform by other methods.

Commitments in the context of the BitCoin were already considered in [15], however, the construction and its applications are different — the main idea of [15] is to use the BitCoin system as a replacement of a trusted third party in time-stamping. The notion of “deposits” has already been used in BitCoin (see en.bitcoin.it/wiki/Contracts, Section “Example 1”⁶), but the application described there is different: the “deposit” is a method for a party with no reputation to prove that she is not a spambot by temporarily sacrificing some of her money.

II. A SHORT DESCRIPTION OF BITCOIN

BitCoin works as a peer-to-peer network in which the participants jointly emulate the central server that controls the correctness of transactions. In this sense it is similar to the concept of the multiparty computation protocols. Recall that, as described above, a fundamental problem with the traditional MPCs is that they cannot provide fairness if there is no honest majority among the participants, which is particularly difficult to guarantee in the peer-to-peer networks where the sybil attacks are possible. The BitCoin system overcomes this problem in the following way: the honest majority is defined in terms of the “majority of computing power”. In other words: in order to break the system, the adversary needs to control machines whose total computing power is comparable with the combined computing power of all the other participants of the protocol. Hence, e.g., the sybil attack does not work, as creating a lot of fake identities in the network does not help the adversary. In a moment we will explain how this is implemented, but let us first discuss the functionality of the trusted party that is emulated by the users.

One of the main problems with the digital currency is the potential double spending: if coins are just strings of bits then the owner of a coin can spend it multiple times. Clearly this risk could be avoided if the users had access to a trusted ledger with the list of all the transactions. In this case a transaction would be considered valid only if it is posted on the board. For example suppose the transactions are of a form: “user X transfers to user Y the money that he got in some previous transaction T_y ”, signed by the user X. In this case each user can verify if money from transaction T_y has not been already spent by X. The functionality of the trusted party emulated by the BitCoin network does precisely this: it maintains a full list of transactions that happened in the system. The format of the BitCoin transactions is in fact more complex than in the example above. Since it is of a special interest for us, we describe it in more detail in Section II-A.

The BitCoin ledger is implemented using the concept of the Proofs of Work (PoWs) [19] in the following clever way. The users maintain a chain of *blocks*. The first block B_0 , called the *genesis block*, was generated by the designers of the system in January 2009. Each new block B_i contains a list T_i of new transactions, the cryptographic hash of the previous block $H(B_{i-1})$, and some random salt R . The key point is that not every R works for given T_i and $H(B_{i-1})$. In fact,

⁶Accessed on 13.11.2013.

the system is designed in such a way that it is moderately hard to find a valid R . Technically it is done by requiring that the binary representation of the hash of $(T_i || H(B_{i-1}) || R)$ starts with a certain number m of zeros (the procedure of extending the chain is called *mining*, and the machines performing it are called *miners*). The hardness of finding the right R depends of course on m , and this parameter is periodically adjusted to the current computing power of the participants in such a way that the extension happens on average every 10 minutes. The system contains an incentive to work on finding the new blocks. We will not go into the details of this, but let us only say that one of the side-effects of this incentive system is the creation of new coins⁷.

The idea of the block chain is that the longest chain C is accepted as the proper one. If some transaction is contained in a block B_i and there are several new blocks on top of it, then it is infeasible for an adversary with less than a half of the total computing power of the Bitcoin network to revert it — he would have to mine a new chain C' bifurcating from C at block B_{i-1} (or earlier), and C' would have to be longer than C . The difficulty of that grows exponentially with number of new blocks on top of B_i . In practice the transactions need 10 to 20 minutes (i.e. 1-2 new blocks) for reasonably strong confirmation and 60 minutes (6 blocks) for almost absolute certainty that they are irreversible.

To sum up, when a user wants to pay somebody in Bitcoins, he creates a transaction and broadcasts it to other nodes in the network. They validate this transaction, send it further and add it to the block they are mining. When some node solves the mining problem, it broadcasts its block to the network. Nodes obtain a new block, validate transactions in it and its hash and accept it by mining on top of it. Presence of the transaction in the block is a confirmation of this transaction, but some users may choose to wait for several blocks on top of it to get more assurance.

A. The Bitcoin transactions

The Bitcoin currency system consists of *addresses* and *transactions* between them. An address is simply a public key pk .⁸ Normally every such a key has a corresponding private key sk known only to one user. The private key is used for signing the transactions, and the public key is used for verifying the signatures. Each user of the system needs to know at least one private key of some address, but this is simple to achieve, since the pairs (sk, pk) can be easily generated offline. We will frequently denote key pairs using the capital letters (e.g. A), and refer to the private key and the public key of A by: $A.sk$ and $A.pk$, respectively (hence: $A = (A.sk, A.pk)$). We will also use the following convention: if $A = (A.sk, A.pk)$ then let $\text{sig}_A(m)$ denote a signature on a message m computed with $A.sk$ and let $\text{ver}_A(m, \sigma)$ denote the result (true or false) of the verification of a signature σ on message m with respect to the public key $A.pk$.

1) *Simplified version:* We first describe a simplified version of the system and then show how to extend it to obtain the description of the real Bitcoin. We do not describe how the coins are created as it is not relevant to this paper. Let $A = (A.sk, A.pk)$ be a key pair. In our simplified view a transaction describing the fact that an amount v (called the *value* of a transaction) is transferred from an address $A.pk$ to an address $B.pk$ has the following form

$$T_x = (y, B.pk, v, \text{sig}_A(y, B.pk, v)),$$

⁷The number of coins that are created in the system is however limited, and therefore Bitcoin is expected to have no inflation.

⁸Technically an address is a *hash* of pk . In our informal description we decided to assume that it is simply pk . This is done only to keep the exposition as simple as possible, as it improves the readability of the transaction scripts later in the paper.

where y is an index of a previous transaction T_y . We say that $B.pk$ is the recipient of T_x , and that the transaction T_y is an *input* of the transaction T_x , or that it is *redeemed* by this transaction (or redeemed by the address $B.pk$). More precisely, the meaning of T_x is that the amount v of money transferred to $A.pk$ in transaction T_y is transferred further to $B.pk$. The transaction is valid only if (1) $A.pk$ was a recipient of the transaction T_y , (2) the value of T_y was at least v (the difference between v and the value of T_y is called the *transaction fee*), (3) the transaction T_y has not been redeemed earlier, and (4) the signature of A is correct. Clearly all of these conditions can be verified publicly.

The first important generalization of this simplified system is that a transaction can have several “inputs” meaning that it can accumulate money from several past transactions $T_{y_1}, \dots, T_{y_\ell}$. Let A_1, \dots, A_ℓ be the respective key pairs of the recipients of those transactions. Then a multiple-input transaction has the following form:

$$T_x = (y_1, \dots, y_\ell, B.pk, v, \text{sig}_{A_1}(y_1, B.pk, v), \dots, \text{sig}_{A_\ell}(y_\ell, B.pk, v)),$$

and the result of it is that $B.pk$ gets the amount v , provided it is at most equal to the sum of the values of transactions $T_{y_1}, \dots, T_{y_\ell}$. This happens only if *none* of these transactions has been redeemed before, and *all* the signatures are valid. Moreover, each transaction can have a *lock-time* t that tells at what time the transaction becomes final (t can refer either to a block index or to the real physical time). In this case we have:

$$T_x = (y_1, \dots, y_\ell, B.pk, v, t, \text{sig}_{A_1}(y_1, B.pk, v, t), \dots, \text{sig}_{A_\ell}(y_\ell, B.pk, v, t)).$$

Such a transaction becomes valid only if time t is reached and if none of the transactions $T_{y_1}, \dots, T_{y_\ell}$ has been redeemed by that time (otherwise it is discarded). Each transaction can also have several outputs, which is a way to divide money between several users. We ignore this fact in our description since we will not use it in our protocols.

2) *A more detailed version:* The real BitCoin system is significantly more sophisticated than what is described above. First of all, there are some syntactic differences, the most important for us being that each transaction T_x is identified not by its index, but by its hash $H(T_x)$. Hence, from now on we will assume that $x = H(T_x)$.


The main difference is, however, that in the real BitCoin the users have much more flexibility in defining the condition on how the transaction T_x can be redeemed. Consider for a moment the simplest transactions where there is just one input and no time-locks. Recall that in the simplified system described above, in order to redeem a transaction the recipient $A.pk$ had to produce another transaction T_x signed with his private key $A.sk$. In the real BitCoin this is generalized as follows: each transaction T_y comes with a description of a function (*output-script*) π_y whose output is Boolean. The transaction T_x redeeming the transaction T_y is valid if π_y evaluates to true on input T_x . Of course, one example of π_y is a function that treats T_x as a pair (a message m_x , a signature σ_x), and checks if σ_x is a valid signature on m_x with respect to the public key $A.pk$. However, much more general functions π_y are possible. Going further into details, a transaction looks as follows: $T_x = (y, \pi_x, v, \sigma_x)$, where $[T_x] = (y, \pi_x, v)$ is called the *body*⁹ of T_x and σ_x is a “witness” that is used to make the script π_y evaluate to true on T_x (in the simplest case σ_x is a signature on $[T_x]$). The scripts are written in the BitCoin scripting language, which is a stack based, not Turing-complete language (there are no loops in it). It provides basic arithmetical operations on numbers, operations on stack, if-then-else statements and some cryptographic functions like calculating hash function

⁹In the original BitCoin documentation this is called “simplified T_x ”. We chosen to rename it to “body” since we find the original terminology slightly misleading.

or verifying a signature. The generalization to the multiple-input transactions with time-locks is straightforward: a transaction has a form:

$$T_x = (y_1, \dots, y_\ell, \pi_x, v, t, \sigma_1, \dots, \sigma_\ell),$$

where the body $[T_x]$ is equal to $(y_1, \dots, y_\ell, \pi_x, v, t)$, and it is valid if (1) time t is reached, (2) every $\pi_i([T_x], \sigma_i)$ evaluates to true, where each π_i is the output script of the transaction T_{y_i} , and (3) none of these transactions has been redeemed before. We will present the transactions as boxes. The redeeming of transactions will be indicated with arrows (the arrows will be labelled with the transaction values). For example a transaction $T_x = (y_1, y_2, \pi_x, v, t, \sigma_1, \sigma_2)$ will be represented as:



The transactions where the input script is a signature, and the output script is a verification algorithm are the most common type of transactions. We will call them *standard transactions*, and the address against which the verification is done will be called the *recipient* of a transaction. Currently some miners accept only such transactions. However, there exist other ones that do accept the non-standard (also called *strange*) transactions, one example being a big mining pool¹⁰ called *Eligius* (that mines a new block on average once per hour). We also believe that in the future accepting the general transaction will become standard, maybe at a cost of a slightly increased fee. This is important for our applications since our protocols rely heavily on the extended form of transactions.

B. Security Model

To reason formally about the security we need to describe the attack model that corresponds to the current BitCoin system. We assume that the parties are connected by an insecure channel and have access to the BitCoin chain. Let us discuss these two assumptions in detail. First, recall that our protocol should allow any pair of users on the internet to engage in a protocol. Hence, we cannot assume that there is any secure connection between the parties (as this would require that they can verify their identity, which obviously is impossible in general), and therefore any type of a man-in-the-middle attack is possible.

The only “trusted component” in the system is the BitCoin chain. For the sake of simplicity in our model we will ignore the implementation details of it, and simply assume that the parties have access to a trusted third party denoted Ledger, whose contents is publicly available. One very important aspect that needs to be addressed are the security properties of the communication channel between the parties and the Ledger. Firstly, it is completely reasonable to assume that the parties can verify Ledger’s authenticity. In other words: each party can access the current contents of the Ledger. In particular, the users can post transactions on the Ledger. After a transaction is posted it appears on the Ledger (provided it is valid), however it may happen not immediately, and some delay is possible. There is an upper bound \max_{Ledger} on this delay. This corresponds to an

¹⁰Mining pools are coalitions of miners that perform their work jointly and share the profits.

assumption that sooner or later every transaction will appear in some BitCoin block. We use this assumption very mildly and e.g. $\max_{\text{Ledger}} = 1$ day is also ok for us (the only price for this is that in such case we have to allow the adversary to delay the termination of the protocol for time $O(\max_{\text{Ledger}})$). Each transaction posted on the Ledger has a time stamp that refers to the moment when it appeared on the Ledger.

What is a bit less obvious is how to define privacy of the communication between the parties and the Ledger, especially the question of the privacy of the writing procedure. More precisely, the problem is that it is completely unreasonable to assume that a transaction is secret until it appears on the Ledger (since the transactions are broadcast between the nodes of the network). Hence we do not assume it. This actually poses an additional challenge in designing the protocols because of the problem of the *malleability*¹¹ of the transactions. Let us explain it now. Recall that the transactions are referred to by their hashes. Suppose a party P creates a transaction T and, before posting it on the Ledger, obtains from some other party P' a transaction T' that redeems T (e.g.: T' may be time-locked and serve P to redeem T if P' misbehaves). Obviously T' needs to contain (in the signed body) a hash $H(T)$ of T . However, if now P posts T then an adversary (allied with malicious P') can produce another transaction \hat{T} whose semantics is the same as T , but whose hash is different (this can be done, e.g., by adding some dummy instructions to the input scripts of T). The adversary can now post \hat{T} on the Ledger and, if he is lucky, \hat{T} will appear on the Ledger instead of T ! It is possible that in the future versions of the BitCoin system this issue will be addressed and the transactions will not be malleable. In Section V we propose a scheme that is secure under this assumption. We would like to stress that our main schemes (Section III and IV) do *not* assume non-malleability of the transactions, and are secure even if the adversary obtains full information on the transactions before they appear on the Ledger.

We do not need to assume any privacy of the reading procedure, i.e. each party accesses pattern to Ledger can be publicly known. We assume that the parties have access to a perfect clock and that their internal computation takes no time. The communication between the parties also takes no time, unless the adversary delays it. These assumptions are made to keep the model as simple as possible, and the security of our protocols does not depend on these assumptions. In particular we assume that the network is asynchronous and our protocols are also secure if the communication takes some small amount of time. For simplicity we also assume that the transaction fees are zero. The extension to non-zero transaction fees is discussed in Section. VI.

III. BITCOIN-BASED TIMED COMMITMENT SCHEME

We start with constructing a BitCoin-based timed-commitment scheme [7], [24]. Recall that a commitment scheme [6], [8] consists of two phases: the *commitment phase* *Commit* and the *opening phase* *Open*. Typically a commitment scheme is executed between two parties: a committer C and a recipient. To be more general we will assume that there are n recipients, denoted P_1, \dots, P_n . The committer starts the protocol with some secret value x . This value will become known to every recipient after the opening phase is executed. Informally, we require that, if the committer is honest, then before the opening phase started, the adversary has no information about x (this property is called “hiding”). On the other hand, every honest recipient can be sure that, no matter how a malicious sender behaves, the commitment can be open in exactly one way, i.e. it is impossible for the committer to “change his mind” and open with some $x' \neq x$. This property is called “binding”. Although incredibly useful in many applications, the standard commitment schemes suffer from the

¹¹See en.bitcoin.it/wiki/Transaction_Malleability.

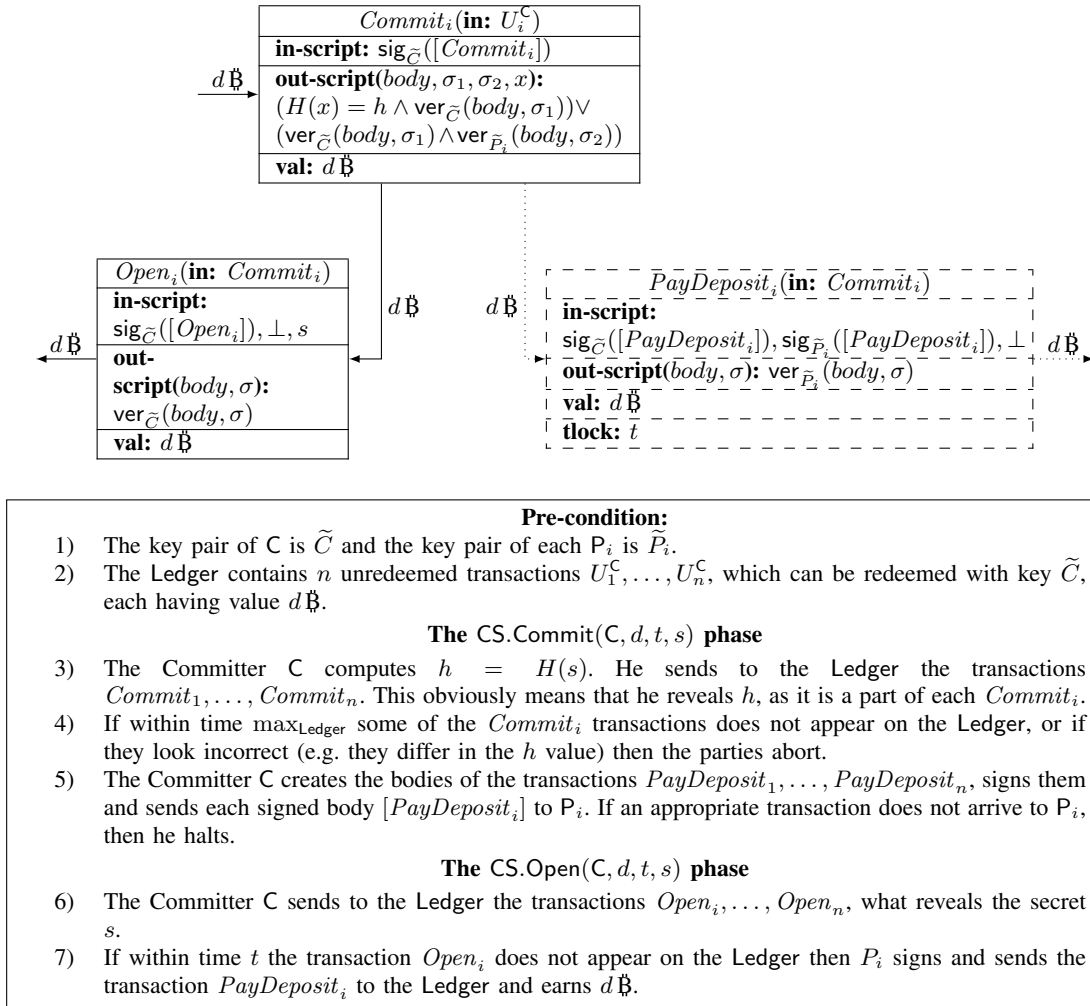


Fig. 1. The CS protocol. The scripts' arguments, which are omitted are denoted by \perp .

following problem: there is no way to force the committer to reveal his secret x , and, in particular, if he aborts before the *Open* phase starts then x remains secret.

BitCoin offers an attractive way to deal with this problem. Namely: using the BitCoin system one can force the committer to back his commitment with some money, called the *deposit*, that will be given to the other parties if he refuses to open the commitment within some time t .

We now sketch the definition of a *BitCoin-based commitment scheme*. First, assume that before the protocol starts the Ledger contains n unredeemed standard transactions U_1^C, \dots, U_n^C that can be redeemed with a key known only to C, each having value $d\mathbb{B}$ (for some parameter d). In fact, in real life it would be enough to have just one transaction, that would later be “splitted” inside of the protocol. This would, however, force us to use the multiple-output transactions which we want to avoid, in order not to additionally complicate the description of the system.

The protocol is denoted $CS(C, d, t, s)$ and it consists of two phases: the *commitment* phase, denoted $CS.Commit(C, d, t, s)$ (where s contains the message to which C commits and some randomness) and the *opening* phase $CS.Open(C, d, t, s)$. The honest committer always opens his commitment by time t . In this case he gets back his money, i.e. the Ledger consists of standard

transactions that can be redeemed with a key known only to him, whose total value¹² is $(d \cdot n) \text{฿}$.

The security definition in the standard commitment scheme: assuming that the committer is honest, the adversary does not learn any significant information about x before the opening phase, and each honest party can be sure that there is at most one value x that the committer can open in the opening phase. Each recipient can also abort the commitment phase (which happens if he discovers that the Committer is cheating, or if the adversary disturbs the communication). However, there is one additional security guarantee: if the committer did not open the commitment by time t then every other party earns $d \text{฿}$. More precisely: for every honest P_i the Ledger contains a transaction that can be redeemed with a key known only to P_i .

Let us also comment on the formal aspects. To satisfy the page limit, we do not provide the full formal model, however, from the discussion above it should be clear how such a model can be defined. We allow negligible error probabilities both in binding and in hiding. Also, the last security property (concerning the deposits) has to hold only with overwhelming probability. As these notions are asymptotic, this requires using a security parameter, denoted by k . Of course, in reality the parameter k is partially fixed by the BitCoin specification (e.g. we cannot modify the length of the outputs of the hash functions).

A. The implementation

Our implementation can be based on any standard commitment scheme as long as it is *hash-based*, by which we mean that it has the following structure. Let H be a hash function. During the commitment phase the committer sends to the recipient some value denoted h (which essentially constitutes his “commitment” to x), and in the opening phase the committer sends to the recipient a value s , such that $H(s) = h$. If $H(s) \neq h$ then the recipient does not accept the opening. Otherwise he computes x from s (there exists an algorithm that allows him to do it efficiently). One example of such a commitment scheme is as follows. Suppose $x \in \{0, 1\}^*$. In the commitment phase C computes $s := (x||r)$, where r is chosen uniformly at random from $\{0, 1\}^k$, and sends to every recipient $h = H(s)$. In the opening phase the committer sends to every recipient s , the recipient checks if indeed $h = H(s)$, and recovers x by stripping-off the last k bits from s . The binding property of this commitment follows from the collision-resistance of the hash function H , since to be able to open the commitment in two different ways a malicious sender would need to find collisions in H . For the hiding property we need to assume that H is a random oracle. We think that this is satisfactory since anyway the security of the BitCoin PoWs relies on the random oracle assumption. Clearly, if H is a random oracle then no adversary can obtain any information about x if he does not learn s (which an honest C keeps private until the opening phase).

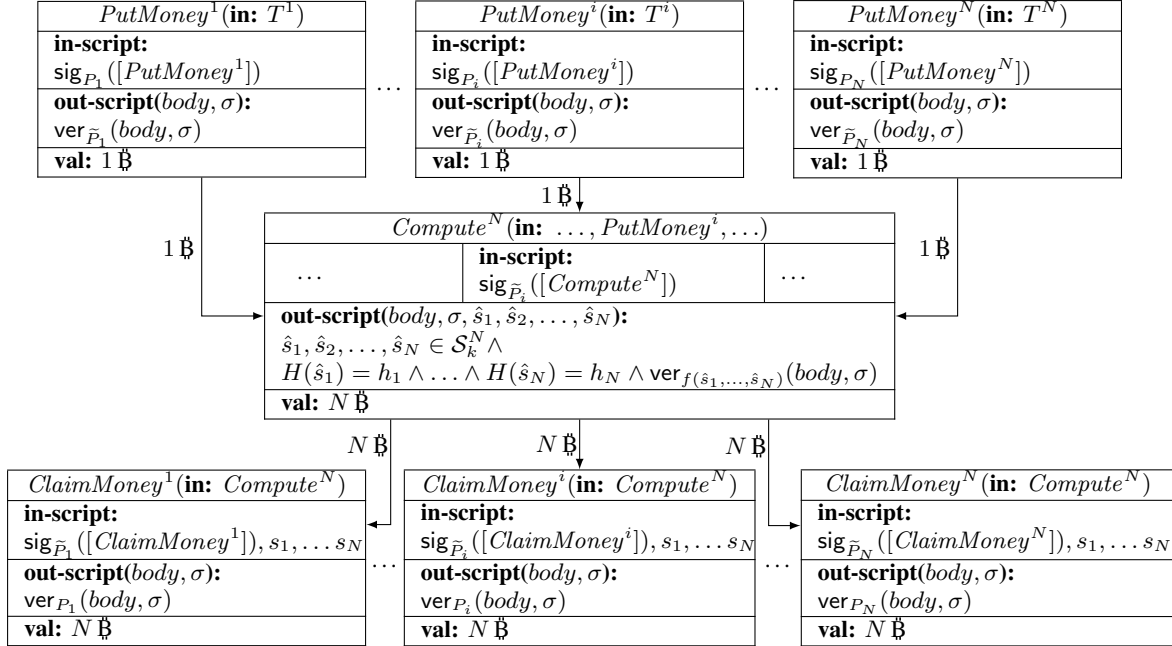
The basic idea of our protocol is as follows. The committer will talk independently to each recipient P_i . For each of them he will create in the commitment phase a transaction $Commit_i$ with value d that normally will be redeemed by him in the opening phase with a transaction $Open_i$. The transaction $Commit_i$ will be constructed in such a way that the $Open_i$ transaction has to automatically open the commitment. Technically it will be done by constructing the output script $Commit_i$ in such a way that the redeeming transaction has to provide s (which will therefore become publicly known as all transactions are publicly visible). Of course, this means that the money of the committer is “frozen” until he reveals s . However, to set a limit on the waiting time of the recipient, we also require the committer to send to P_i a transaction $PayDeposit_i$ that can redeem $Commit_i$ if time t passes. Of course, P_i , after receiving $PayDeposit_i$ needs to check if it

¹²In case of non-zero transaction fees this value can be decreased by these fees. This remark applies also to the amounts d redeemed by the recipients.

is correct. The commitment scheme and the transactions are depicted on Figure 1 (page 12). We now state the following lemma, whose proof appears in the full version of this paper.

Lemma 1: The CS scheme on Figure 1 is a BitCoin-based commitment scheme.

IV. THE LOTTERY PROTOCOL



- Pre-condition:**
- 1) For each i , player P_i holds a pair of keys $(P_i.sk, P_i.pk)$.
 - 2) For each i , the Ledger contains a standard transaction T^i that have value 1 ₿ each and whose recipient is P_i . The Ledger contains also transactions $\{U_j^i\}$, where $i, j \in \{1, \dots, N\}$ and $i \neq j$, such that each U_j^i can be redeemed by P^i and has value $d \text{ ₿}$ (for some parameter d)
- Initialization phase:**
- 3) For each i , player P_i generates a pair of keys $(\tilde{P}_i.sk, \tilde{P}_i.pk)$ and sends his public key $\tilde{P}_i.pk$ to all other players.
 - 4) For each i , player P_i chooses his secret s_i .
- Deposits phase:**
- 5) Let t be the current time. For each $i = 1, \dots, N$ the commitment phase $CS.Commit(P_i, d, t + 4 \cdot \max_{\text{Ledger}}, s_i)$ is executed using the transactions $\{U_j^i\}$ as inputs.
- Execution phase:**
- 6) For each i , player P_i puts transaction $PutMoney^i$ to the Ledger. The players halt if any of those transactions did not appear on the Ledger before time $t + 2 \cdot \max_{\text{Ledger}}$.
 - 7) For each $i \geq 2$, player P_i computes his signature on transaction $Compute^N$ and sends it to the player P_1 .
 - 8) Player P_1 puts all received signatures (and his own) into inputs of transaction $Compute^N$ and puts it to the Ledger. If $Compute^N$ did not appear on Ledger in time $t + 3 \cdot \max_{\text{Ledger}}$, then the players halt.
 - 9) For each i , the player P_i puts his *Open* transactions on Ledger what reveals his secret and sends back to him the deposits he made during the executions of CS protocol from Step. 5. If some player did not revealed his secret in time $t + 4 \cdot \max_{\text{Ledger}}$, then other players send the appropriate *PayDeposit* transactions from that player CS protocols to the Ledger to get $N \text{ ₿}$.
 - 10) The player, that is the winner (i.e. $P_{f(s_1, \dots, s_N)}$), gets the pot by sending the transaction $ClaimMoney^{f(s_1, \dots, s_N)}$ to the Ledger.

Fig. 2. The MultiPlayersLottery protocol.

As discussed in the introduction, as an example of an application of the “MPCs on BitCoin” concept we construct a protocol for a lottery executed among a group of parties P_1, \dots, P_N . We say that a protocol is a *fair lottery protocol* if it is *correct* and *secure*. To define correctness assume all the parties are following the protocol and the communication channels between them are secure (i.e. it reliably transmits the messages between the parties without delay).

We assume that before the protocol starts, the Ledger contains unredeemed standard transactions T^1, \dots, T^N known to all the parties, all of value 1 ₿ and each T^i can be redeemed with a key known only to P_i . Moreover, since we will use the commitment scheme from Section III, the parties need to have money to pay the “deposits”. This money will come from transactions $\{U_j^i\}$, where $i, j \in \{1, \dots, N\}$ and $i \neq j$, such that each U_j^i can be redeemed only by P_i and has value $d \text{ ₿}$ (for some parameter d whose value will be determined later). We assume that these transactions are on the Ledger before the protocol starts. The protocol has to terminate in time $O(\max_{\text{Ledger}})$ and at the moment of termination, the Ledger has to contain a standard transaction with value $N \text{ ₿}$ which can be redeemed with a key known only to P_w , where w is chosen uniformly at random from the set $\{1, \dots, N\}$. The Ledger also contains transactions for paying back the deposits, i.e. we require that for each P_i there is an additional transaction (that can be redeemed only by him) whose value is $(N-1)d \text{ ₿}$. Of course, in the case of the non-zero fees these values will be slightly smaller, but to keep things simple we assume here that these fees are zero.

To define security, look at the execution of the protocol from the point of view of one party, say P_1 (the case of the other parties is symmetric). Assume P_1 is honest and hence, in particular, the Ledger contains the transactions T^1, U_2^1, \dots, U_N^1 , whose recipient is P_1 and whose value is: 1 ₿ in case of T^1 and $d \text{ ₿}$ in case of the U_j^1 's. Obviously, P_1 has no guarantee that the protocol will terminate successfully, as the other party can, e.g., leave the protocol before it is completed. What is important is that P_1 should be sure that she will not lose money because of this termination (in particular: the other parties should not be allowed to terminate the protocol after he learned that P_1 won). This is formalized as follows: we define the *payoff* of P_1 in the execution of the protocol to be equal to the difference between the money that P_1 invested and the money that he won. More formally, the payoff of P_1 is equal to $X_1 - ((N-1) \cdot d + 1) \text{ ₿}$, where X_1 is defined as the total sum of the values of transactions from the execution of the protocol (including T^1, U_2^1, \dots, U_N^1) that P_1 (and only him) can redeem when the protocol terminates. (The payoff of any other participant P_i is defined symmetrically.)

Ideally we would like to require that the expected payoff of each honest player cannot be negative¹³. However, since the security of our protocol relies on non-perfect cryptographic primitives, such as commitment schemes, we have to take into account a negligible probability of the adversary breaking them. Hence, we require only that these values are “at least negligible”¹⁴ in some security parameter k (that is used in the crypto primitives). Formally, we say that the protocol is *secure* if for any strategy of the adversary, that controls the network and corrupts the other parties, (1) the execution of the protocol terminates in time $O(\max_{\text{Ledger}})$, and (2) the expected payoff of each honest party is at least negligible. The expected values are taken over all the randomness in the experiment (i.e. both the internal randomness of the parties and the adversary). We also note that, of course, a dishonest participant can always interrupt in a very early stage. This is not a problem if the transaction fees are zero. In case of the non-zero transaction fees this may cause the other parties to lose a small amount of money. This problem is addressed in Section VI.

¹³In principle it can be actually positive if the adversary plays against his own financial interest.

¹⁴Formally: a function $\alpha : \mathbf{N} \rightarrow \mathbf{R}$ is *at least negligible* if there exists a function $\beta : \mathbf{N} \rightarrow \mathbf{R}$ such that for every i we have $\alpha(i) \geq \beta(i)$ and β is negligible, i.e. its absolute value is asymptotically smaller than the inverse of any polynomial.

A. The protocol

Our protocol is built on top of the classical coin-tossing protocol of Blum [6] that is based on cryptographic commitments. The Blum’s scheme adapted to N parties is very simple — each party P_i commits herself to an element $b_i \in Z_N$. Then, the parties open their commitments and the winner is P_w where $w = (b_1 + \dots + b_N \bmod N) + 1$. As described in the introduction, this protocol does not directly work for our applications, and we need to adapt it to BitCoin. In particular, in our solution creating and opening the commitments are done by the transactions’ scripts using SHA-256 hashing. Due to the technical limitations of BitCoin scripting language in its current form¹⁵, instead of random numbers b_i , the parties commit themselves to strings s_i sampled with uniformly random from $\mathcal{S}_k^N := \{0, 1\}^{8k} \cup \dots \cup \{0, 1\}^{8(k+N-1)}$, i.e. it is the set of strings of length $k, \dots, (k + N - 1)$ bytes¹⁶, where k is the security parameter. The winner is determined by the *winner choosing function* $f(s_1, \dots, s_N)$ for N players to be equal P_ℓ if $\sum_{i=1}^N |s_i| \equiv (\ell - 1) \bmod N$, where s_1, \dots, s_N are the secret strings chosen from \mathcal{S}_k^N and $|s_i|$ is a length of string s_i in bytes. Honest users first choose length (in bytes) of their strings from the set $\{k, \dots, k + N - 1\}$ and then generate a random string of the appropriate length. It is easy to see that as long as one of the parties draws her string’s length uniformly, then the output of $f(s_1, \dots, s_N)$ is also uniformly random.

1) *First attempt*: For simplicity let us start with the case of $N = 2$ parties, called Alice A and Bob B. Their key pairs are A and B (resp.) and their unredeemed transactions placed on the Ledger before the protocol starts are denoted T^A, U^A and T^B, U^B . We start with presenting a naive and insecure construction of the protocol, and then show how it can be modified to obtain a secure scheme. The protocol starts with Alice and Bob creating their new pairs of keys $\tilde{A} = (\tilde{A}.sk, \tilde{A}.pk)$ and $\tilde{B} = (\tilde{B}.sk, \tilde{B}.pk)$, respectively. These keys will be used during the protocol. It is actually quite natural to create new keys for this purpose, especially since many BitCoin manuals recommend creating a fresh key pair for every transaction. Anyway, there is a good reason to do it in our protocol, e.g. to avoid interference with different sessions of the same protocol. Both parties announce their public keys to each other. Alice and Bob also draw at random their secret strings s_A and s_B (respectively) from the set \mathcal{S}_k^2 and they exchange the hashes $h_A = H(s_A)$ and $h_B = H(s_B)$. Moreover Alice sends to the Ledger the following transaction:

$PutMoney_1^A(\mathbf{in}: T^A)$
in-script: $\text{sig}_A([PutMoney_1^A])$
out-script ($body, \sigma$): $\text{ver}_{\tilde{A}}(body, \sigma)$
val: 1 B

Bob also sends to the Ledger a transaction $PutMoney_1^B$ defined symmetrically (recall that T^A and T^B are standard transactions that can be redeemed by Alice and Bob respectively). If at any point later a party $P \in \{A, B\}$ realizes that the other party is cheating, then the first thing P will do is to “take the money and run”, i.e. to post a transaction that redeems $PutMoney_1^P$. We will call it “halting the execution”. This can clearly be done as long as $PutMoney_1^P$ has not been redeemed by some other transaction. In the next step one of the parties constructs a transaction $Compute_1$ defined as follows:

¹⁵Most of the more advanced instructions (e.g. concatenation, accessing particular bits in a string or arithmetic on big integers) have been disabled out of concern that the clients may have bugs in their implementation. Therefore, computing length (in bytes), hashing and testing equality are the only operations available for strings.

¹⁶The transactions in the protocol will always check if their inputs are from \mathcal{S}_k^N (whenever they are supposed to be from this set). If not, they are considered invalid, and the transaction is not evaluated.

$Compute_1$ (in: $PutMoney_1^A, PutMoney_1^B$)	
in-script₁: $\text{sig}_{\tilde{A}}([Compute_1])$	in-script₂: $\text{sig}_{\tilde{B}}([Compute_1])$
out-script ($body, \sigma_1, \sigma_2, \hat{s}_A, \hat{s}_B$): $(\hat{s}_A, \hat{s}_B \in S_k^2 \wedge H(\hat{s}_A) = h_A \wedge H(\hat{s}_B) = h_B \wedge f(\hat{s}_A, \hat{s}_B) = A \wedge \text{ver}_{\tilde{A}}(body, \sigma_1)) \vee$ $(\hat{s}_A, \hat{s}_B \in S_k^2 \wedge H(\hat{s}_A) = h_A \wedge H(\hat{s}_B) = h_B \wedge f(\hat{s}_A, \hat{s}_B) = B \wedge \text{ver}_{\tilde{B}}(body, \sigma_2))$	
val: 2B	

Note that the body of $Compute_1$ can be computed from the publicly-available information. Hence this construction can be implemented as follows: first one of the players, say, Bob computes $[Compute_1]$, and sends his signature $\text{sig}_{\tilde{B}}([Compute_1])$ on it to Alice. Alice adds her signature $\text{sig}_{\tilde{A}}([Compute_1])$ and posts the entire transaction $Compute_1$ to the Ledger.

The output script of $Compute_1$ is an alternative of two conditions. Since they are symmetric (with respect to A and B) let us only look at the first condition (call it γ). To make it evaluate to true on $body$ one needs to provide as “witnesses” $(\sigma_1, \hat{s}_A, \hat{s}_B)$ where \hat{s}_A and \hat{s}_B are the pre-images of h_A and h_B (with respect to H) from S_k^2 . Clearly the collision-resistance of H implies that \hat{s}_A and \hat{s}_B have to be equal to s_A and s_B (resp.). Hence γ can be satisfied only if the winner choosing function f evaluates to A on input (s_A, s_B) . Since only Alice knows the private key of \tilde{A} , hence only she can later provide a signature on σ_1 that would make the last part of γ (i.e.: “ $\text{ver}_{\tilde{A}}(body, \sigma_1)$ ”) evaluate to true.

Clearly before $Compute_1$ appears on the Ledger each party P can “change her mind” and redeem her initial transaction $PutMoney_1^P$, which would make the transaction $Compute_1$ invalid. As we said before, it is ok for us if one party interrupts the coin-tossing procedure as long as she had to decide about doing it *before* she learned that she lost. Hence, Alice and Bob wait until $Compute_1$ appears on Ledger before they proceed to the step in which the winner is determined. This final step is simple: Alice and Bob just broadcast s_A and s_B , respectively. Now: if $f(s_A, s_B) = A$ then Alice can redeem the transaction $Compute_1$ in a transaction $ClaimMoney_1^A$ constructed as:

$ClaimMoney_1^A$ (in: $Compute_1$)	
in-script: $\text{sig}_{\tilde{A}}([ClaimMoney_1^A]), \perp, s_A, s_B$	
out-script ($body, \sigma$): $\text{ver}_A(body, \sigma)$	
val: 2B	

On the other hand Bob cannot redeem $Compute_1$, as the condition $f(s_A, s_B) = B$ evaluates to false. Symmetrically: if $f(s_A, s_B) = B$ then only Bob can redeem $Compute_1$ by an analogous transaction $ClaimMoney_1^B$.

This protocol is obviously correct. It may also look secure, as it is essentially identical to Blum’s protocol described at the beginning of this section (with the hash functions used as the commitment schemes). Unfortunately, it suffers from the following problem: there is no way to guarantee that the parties always send s_A and s_B . In particular: one party, say, Bob, can refuse to send s_B *after* he learned that he lost (i.e. that $f(s_A, s_B) = A$). As his money is already “gone” (his transaction T^B has already been redeemed in transaction $Compute_1$) he would do it just because of sheer nastiness. Unfortunately in a purely peer-to-peer environment, with no concept of a “reputation”, such behavior can happen, and there is no way to punish it.

This is exactly why we need to use the BitCoin-based commitment scheme from Section III. Later, in Section V we also present another technique for dealing with this problem, which avoids using the deposits. Unfortunately, it suffers from certain shortcomings. First of all, it works only for two parties. Secondly, and more importantly, to achieve full security it needs an assumption that an adversary can not see transactions until they appear on the Ledger.

2) *The secure version of the scheme:* The general idea behind the secure MultiPlayersLottery protocol is that each party first commits to her inputs using the $CS(C, d, t, s)$ commitment scheme, instead of the standard commitment scheme (the parameters d and t will be determined later). Recall that the CS commitment scheme can be opened by sending a value s , and this opening is verified by checking that s hashes to a value h sent by the committer in the commitment phase. So, Alice executes the CS protocol acting as the committer and Bob as a recipient (note that there is only one recipient and hence $n = 1$). Let s_A and h_A be the variables s and h created this way. Symmetrically: Bob executes the CS protocol acting as the committer, and Alice being the recipient, which the corresponding variables s_B and h_B . Once both commitment phases are executed successfully (recall that this includes receiving by each party the signed *PayDeposit* transaction), the parties proceed to the next steps, which are exactly as before: first, each of them posts his transaction *PutMoney* on the Ledger. Once all these transactions appear on the Ledger they create the *Compute*² transaction (in the same way as before), and once it appears on the Ledger they open the commitments. The only difference is obviously that, since they used the CS commitment scheme, they can now “punish” the other party if she did not open her commitment by executing *PayDeposit* after the time t passes, and claim her deposit. On the other hand: each honest party is always guaranteed to get her deposit back, hence she does not risk anything investing this money at the beginning of the protocol.

We also need to comment about the choice of the parameters t and d . First, it is easy to see that the maximum time in which the honest parties will complete the protocol is at most $4 \cdot \max_{\text{Ledger}}$ after time t' is the time when the protocol started. Hence we can safely set $t := t' + 4 \cdot \max_{\text{Ledger}}$.

The parameter d should be chosen in such a way that it will fully compensate to each party the fact that a player aborted. Let us now calculate the payoff of some fixed player P_1 , say, assuming the worst-case scenario, which is as that (a) the protocol is always aborted when P_1 is about to win, and (b) there is only one “aborting party” (so P_i is paid only one deposit). Hence his expected payoff is $-\frac{N-1}{N} \mathfrak{B}$ (this corresponds to the case when he lost) plus $\frac{d-1}{N} \mathfrak{B}$ (the case when the protocol was aborted). Therefore to make the expected value equal to 0 we need to set $d = N \mathfrak{B}$. This implies that the total amount of money invested in deposit by each player has to be equal to $N(N-1) \mathfrak{B}$. In real-life this would be ok probably for small groups $N = 2, 3$, but not for the larger ones.

We now have the following lemma, whose proof will appear in the full version of this paper.

Lemma 2: The MultiPlayersLottery protocol from Figure 2 is a fair lottery protocol for $d = N \mathfrak{B}$ and $t = t' + 4 \cdot \max_{\text{Ledger}}$, where t' is the starting time of the protocol.

V. TWO-PARTY LOTTERY SECURE IN A STRONGER MODEL

In this section we show a construction of a two-party lottery which avoids using the deposits, and hence may be useful for applications where the parties are not willing to invest extra money in the execution of the protocol. The drawback of the protocol presented in this section is that it works only for two parties. Moreover, to achieve full security it needs an assumption that the channel between the parties and the Ledger is private, what means that the adversary can not see the transactions sent by the honest user, before they appear on the Ledger. In reality BitCoin transactions are broadcast via a peer-to-peer network, so it is relatively easy to eavesdrop the transactions waiting to be posted on the Ledger. Another problem related to eavesdropping is malleability of transaction (already described in Section II-B). Recall, that the problem is that an adversary can modify (“maul”) the transaction T eavesdropped in the network in such a way that the modified transaction is semantically equivalent to the original one, but it has a different hash. Then, the adversary can send the modified version of T to the Ledger and if he is lucky it will be

posted on the Ledger and invalidate the original transaction T (its input will be already redeemed). Hence, e.g., the transactions that were created to redeem T will not be able to do it (as the hash of the transaction is different). In order to be secure for the protocol presented in this section, we need to assume that such attacks are not impossible. Technically, we do it by assuming that the channel between from each party to the Ledger is private. The protocols secure in this model will be called *fair under private channel assumption*.

To explain our protocol let us go to the point in Section IV where it turned out that we need the BitCoin commitment schemes. Recall that we observed that the protocol from Section IV-A2 is not secure against a “nasty behavior” of the party that, after realizing that she lost, simply quits the protocol.

3) *An alternative (and slightly flawed) idea for a fix:* Suppose for a moment we are only interested in security against the “nasty Bob”. Our method is to force him to reveal s_B simultaneously with $Compute_1$ being posted on the Ledger, by requiring that s_B is a part of $Compute_1$. More concretely this is done as follows. Recall that in our initial protocol we said that $Compute_1$ is created and posted on the Ledger by “one of the parties”. This was ok since the protocol was completely symmetric for A and B. In our new solution we break this symmetry by modifying the $Compute_1$ transaction (this new version will be denoted $Compute_2$) and designing the protocol in such a way that $Compute_2$ will be always posted on the Ledger by B. First of all, however, we redefine the $PutMoney_1^B$ transaction that B posts on the Ledger at the beginning of the protocol. The modified transaction is denoted $PutMoney_2^B$.

$PutMoney_2^B(\text{in: } T^B)$	
in-script: $\text{sig}_B([PutMoney_2^B])$	
out-script ($body, \sigma, \hat{s}$): $\text{ver}_{\hat{B}}(body, \sigma) \wedge \hat{s} \in \mathcal{S}_k^2 \wedge H(\hat{s}) = h_B$	
val: $1 \text{ } \text{B}$	

The only difference compared to $PutMoney_1^B$ is the addition of the “ $\wedge(H(\hat{s}) = h_B)$ ” part. This trick forces Bob to reveal the pre-image of h_B (which has to be equal to s_B) whenever he redeems $PutMoney_2^B$.

The transaction $PutMoney_1^A$ remains unchanged, i.e.: $PutMoney_2^A := PutMoney_1^A$. Clearly players can still redeem their transactions later in case they discover that the other player is cheating. Transaction $Compute_2$ is the same as $Compute_1$, except that s_B is added to the input script for the second input transaction:

$Compute_2(\text{in: } PutMoney_2^A, PutMoney_2^B)$	
in-script ₁ : $\text{sig}_{\hat{A}}([Compute_2])$	in-script ₂ : $\text{sig}_{\hat{B}}([Compute_2]), s_B$
out-script ($body, \sigma_1, \sigma_2, \hat{s}_A, \hat{s}_B$): $(\hat{s}_A, \hat{s}_B \in \mathcal{S}_k^2 \wedge H(\hat{s}_A) = h_A \wedge H(\hat{s}_B) = h_B \wedge f(\hat{s}_A, \hat{s}_B) = A \wedge \text{ver}_{\hat{A}}(body, \sigma_1)) \vee$ $(\hat{s}_A, \hat{s}_B \in \mathcal{S}_k^2 \wedge H(\hat{s}_A) = h_A \wedge H(\hat{s}_B) = h_B \wedge f(\hat{s}_A, \hat{s}_B) = B \wedge \text{ver}_{\hat{B}}(body, \sigma_2))$	
val: $2 \text{ } \text{B}$	

The parties post their $PutMoney_2$ transactions on the Ledger and construct transaction $Compute_2$ in the following way. First, observe that both parties can easily construct the body of $Compute_2$ themselves, as all the information needed for this is: the transactions $PutMoney_2^A$ and $PutMoney_2^B$ and the hashes of s_A and s_B , which are all publicly available. Hence the only thing that needs to be computed are the input scripts. This computation is done as follows: first Alice computes her input script $\text{sig}_{\hat{A}}([Compute_2])$ and sends it to Bob. Then Bob adds his input script ($\text{sig}_{\hat{B}}([Compute_2]), s_B$), and posts $Compute_2$ on the Ledger.

The $ClaimMoney_2^P$ procedures (for $P \in \{A, B\}$) remain unchanged (except, of course that their input is $Compute_2$ instead of $Compute_1$). Let us now analyze the security of this protocol from

the point of view of both parties. First, observe that Alice does not risk anything by sending $\text{sig}_{\tilde{A}}([Compute_2])$ to Bob. This is because it consists of a signature on the entire body of the transaction, and hence it is useless as long as Bob did not add his input script¹⁷. But, if Bob added a correct input script and posted $Compute_2$ on the Ledger then he automatically had to reveal s_B . Hence, from the point of view of Alice the problem of “nasty Bob” is solved.

Unfortunately, from the point of view of Bob the situation looks much worse, as he still has no guarantee that Alice will post s_A once she learned that she lost. This is why one more modification of the protocol is needed.

4) *The secure version of the scheme:* To fix the problem described above we extend our protocol by adding a special transaction that we denote $Fuse$ and that will be used by Bob to redeem $Compute$ if Alice did not send s_A within some specific time, say, $2 \cdot \max_{\text{Ledger}}$. To achieve this we will use the time-lock mechanism described in the introduction. This requires modifying once again the $Compute_2$ transaction so it can be redeemed by $Fuse$. All in all, the transactions are now defined as follows:

$Compute(\text{in: } PutMoney^A, PutMoney^B)$	
in-script₁: $\text{sig}_{\tilde{A}}([Compute])$	in-script₂: $\text{sig}_{\tilde{B}}([Compute]), s_B$
out-script ($body, \sigma_1, \sigma_2, \hat{s}_A, \hat{s}_B$):	
$(\hat{s}_A, \hat{s}_B \in \mathcal{S}_k^2 \wedge H(\hat{s}_A) = h_A \wedge H(\hat{s}_B) = h_B \wedge f(\hat{s}_A, \hat{s}_B) = A \wedge \text{ver}_{\tilde{A}}(body, \sigma_1)) \vee$	
$(\hat{s}_A, \hat{s}_B \in \mathcal{S}_k^2 \wedge H(\hat{s}_A) = h_A \wedge H(\hat{s}_B) = h_B \wedge f(\hat{s}_A, \hat{s}_B) = B \wedge \text{ver}_{\tilde{B}}(body, \sigma_2)) \vee$	
$(\text{ver}_{\tilde{A}}(body, \sigma_1) \wedge \text{ver}_{\tilde{B}}(body, \sigma_2))$	
val: 2B	

$Fuse(\text{in: } Compute)$
in-script:
$\text{sig}_{\tilde{A}}([Fuse]), \text{sig}_{\tilde{B}}([Fuse]), \perp, \perp$
out-script ($body, \sigma$): $\text{ver}_B(body, \sigma)$
val: 2B
tlock: $t + 2 \cdot \max_{\text{Ledger}}$

(t above refers roughly to the time when $Fuse$ is created, we will define it more concretely in a moment).

The transactions $ClaimMoney^A$ and $ClaimMoney^B$ are almost the same as the transactions $ClaimMoney_2^A$ and $ClaimMoney_2^B$ (except that they redeem $Compute$ transaction instead of $Compute_2$). It is clear that $Compute$ can be generated jointly by Alice and Bob in the same way as before (the only new part of $Compute$ is the last line “ $\text{ver}_{\tilde{A}}(body, \sigma_1) \wedge \text{ver}_{\tilde{B}}(body, \sigma_2)$ ” that can be easily computed by both parties from the public information).

What remains is to describe the construction of the $Fuse$ transaction. Clearly, Bob can create the entire $Fuse$ by himself, except of the signature $\text{sig}_{\tilde{A}}([Fuse])$ in the input script, which has to be computed by Alice, as only she knows her private key. To do this Alice needs to know the body of $Fuse$. It is easy to see that she knows all of it, except of the input transaction $Compute$. Moreover, Bob cannot simply send $Compute$ to Alice, since $Compute$ includes the information on his secret s_B which Alice should not learn at this point.

We solve this problem by exploiting the details of the Bitcoin implementation, namely the fact that the transactions are referenced by their hashes. Hence, to create the body of $Fuse$ Alice only needs to know the *hash* $h_{Compute}$ of $Compute$. Therefore our protocol will contain the following sub-procedure (executed directly after Bob constructs $Compute$, but before he posts it on the Ledger): (1) Bob sends $h_{Compute} = H(Compute)$ to Alice, (2) Alice computes $[Fuse]$, signs it,

¹⁷Recall that the body of a transaction includes also the information about its input transactions, and moreover, a transaction becomes valid only if *all* the input transactions can be redeemed

and sends the signature $\text{sig}_{\tilde{A}}([Fuse])$ to Bob, (3) Bob verifies Alice’s signature and halts if it is incorrect. Time t that is used in the time-lock in $Fuse$ will refer to time when Alice executed Step (2) above. This system guarantees that Bob can always claim his $2\mathfrak{B}$ in time $t + 2 \cdot \max_{\text{Ledger}}$ even if Alice did not execute the last step. Observe that of course Alice should halt her execution if she does not see $Compute$ on the Ledger within time $t + \max_{\text{Ledger}}$, as otherwise Bob could simply post $Compute$ at much later (after time $t + 2 \cdot \max_{\text{Ledger}}$, say) and immediately use $Fuse$ to claim the reward.

There are some issues in this procedure that need to be addressed. Firstly, the reader may be worried that $H(Compute)$ reveals some information on $Compute$. In practice (and in theory if H is a random oracle) this happens only if the set of possible inputs to H is small and known to the adversary. In our case the adversary is the dishonest Alice, and it can be easily seen that from her point of view the set of possible $Compute$ transactions is huge, one reason for this being that $Compute$ includes s_B , which is secret and uniform.

Unfortunately the fact that Alice does not know the complete transaction $Compute$, but only its hash, poses a risk to her. This is because a dishonest Bob can, instead of sending $H(Compute)$, send a hash of some other transaction T in order to obtain the information that can be used to redeem some other transaction used within the protocol, or even outside this session of the protocol. This is actually one of the reasons why we assumed that the keys used by the users in our procedure are fresh and will not be used later: in this way we can precisely know, which transactions can be redeemed if one obtains Alice’s signature on $[Fuse]$ constructed with false $h_{Compute}$.¹⁸ It is easy to see that the only transaction other than $Compute$, that could be potentially redeemed using Alice’s signature is $PutMoney^A$. This transaction cannot be redeemed by “ $Fuse$ with false $h_{Compute}$ ”, for several reasons, one of them being that the value of $PutMoney^A$ is $1\mathfrak{B}$, which is less than the value of $Fuse$ (equal to $2\mathfrak{B}$).

In this way we constructed the TwoPlayersLottery protocol. Its complete description is presented on Figure 3 (page 22). We now have the following lemma (the proof appears in the full version of this paper).

Lemma 3: The TwoPlayersLottery protocol from Figure 3 is a secure lottery protocol under the private channel assumption.

Notice that without the private channel assumption there are two possible attacks, which could harm Bob. One is that Alice could see $Compute$, which contains s_B before it is posted on the Ledger. In case she lost she could react with sending to the Ledger another transaction T , which redeems $PutMoney^A$. If she was lucky and the transaction T was posted on the Ledger before $Compute$, that $Compute$ would become invalidated (one of its inputs would be already redeemed) and Bob would not earn any money. The other possible attack concerns the malleability problem: for the security to hold Bob needs to be sure that the $Fuse$ transaction will redeem the transaction $Compute$. Unfortunately, $Fuse$ has to be created strictly before $Compute$ appears on the Ledger. If the adversary intercepts $Compute$ before it happened (or: if the miner is malicious) then he can post a “mauled” $Compute$ transaction on the Ledger that behaves exactly as the original one, except that it hashes to some other value. Hence $Fuse$ becomes useless.

¹⁸ As a more concrete example what could go wrong without this assumption consider the following scenario. Assume there is a not-redeemed transaction $Compute^*$ on the Ledger whose recipient is Alice and that also can be redeemed by a transaction with an input script $(\perp, \perp, \text{sig}_{\tilde{A}}([T]), \text{sig}_{\tilde{B}}([T]))$ (this can happen, e.g., if two coin-tossing protocol are executed in parallel between Alice and Bob). Then a dishonest Bob can send to Alice $H(Compute^*)$ instead of $H(Compute)$, and redeem $Compute^*$ in time $t + 2 \cdot \max_{\text{Ledger}}$

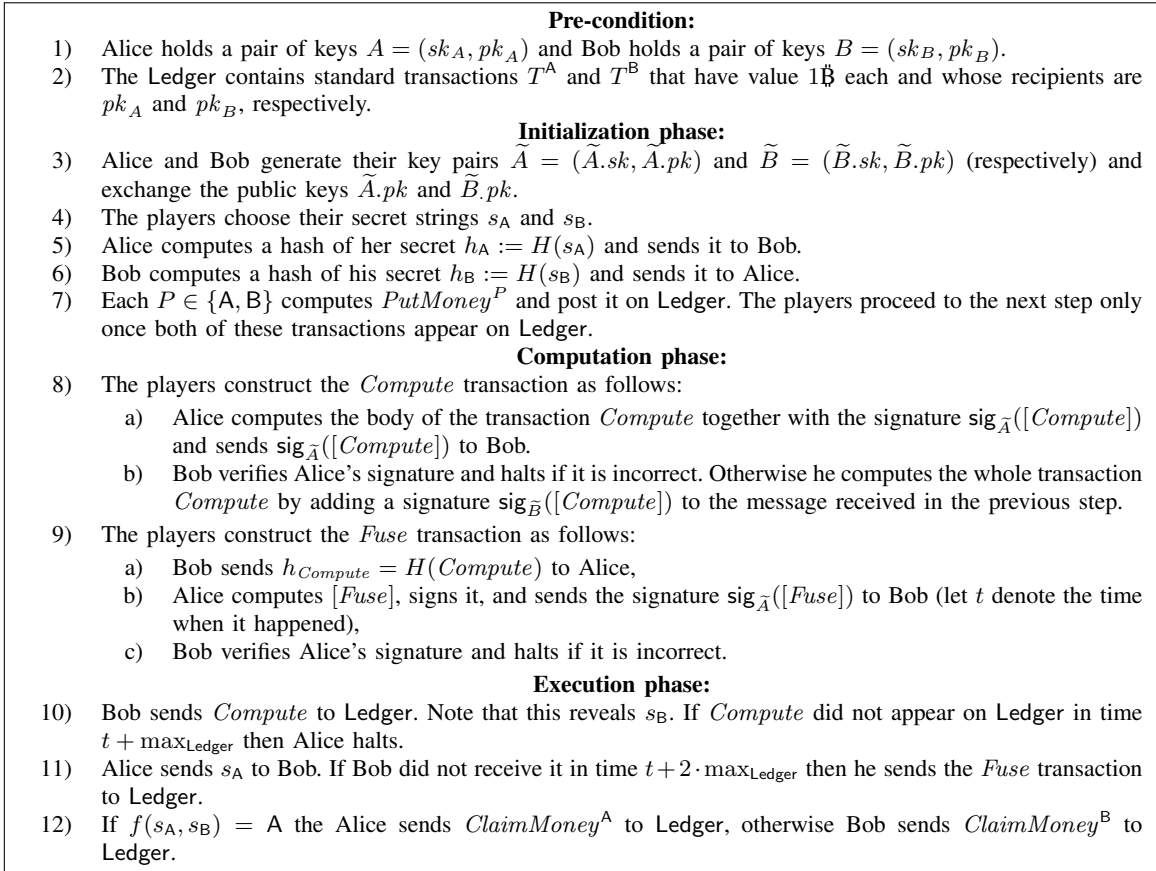
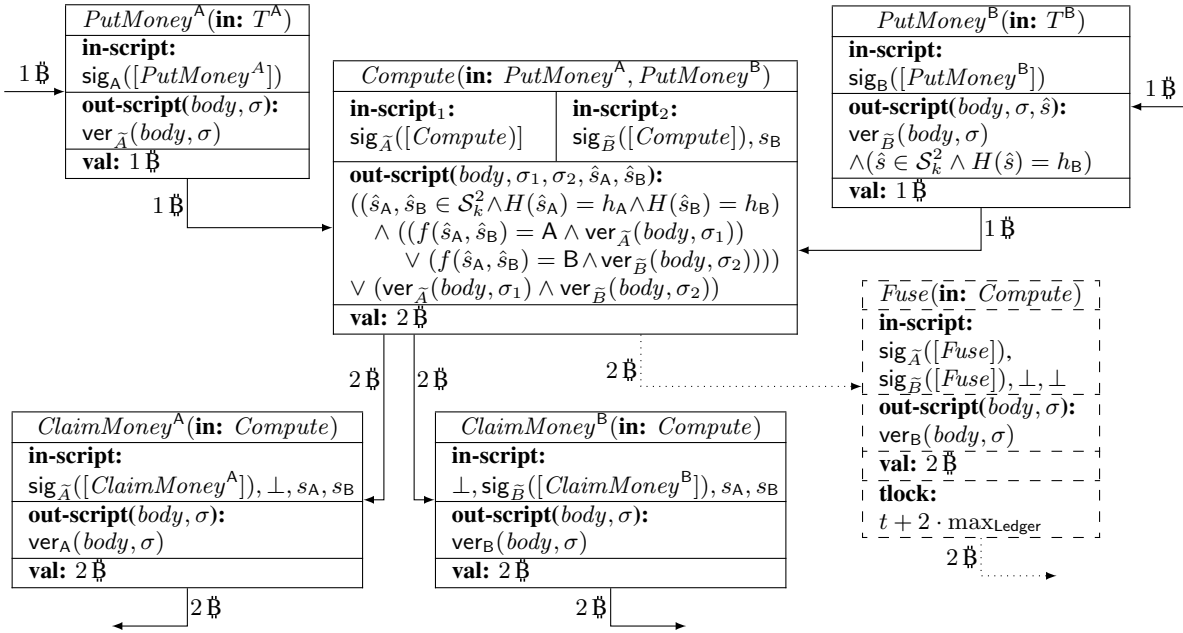


Fig. 3. The TwoPlayersLottery protocol

VI. NON-ZERO TRANSACTION FEES

We now address the problem of the transaction fees, which was ignored in the description above. On a technical level there is no problem with incorporating the fees into our protocol: the transactions can simply include a small fee that has to be agreed upon between the parties before the protocol starts. The expected payoff of the parties will be in this case slightly negative (since the fees need to be subtracted from the outcome). It is straightforward how to modify the security definition to take this into account. One problem that the reader may notice is the issue of the “nasty” behavior of the parties. For example, a malicious Alice can initiate the protocol with Bob just to trigger him to post $PutMoney^B$ on the Ledger. If Alice later aborts then Bob obviously gets his money back, except of the transaction fee. Of course, this does not change his expected payoff, but it still may be against his interests, as he loses some money on a game that from the beginning was planned (by Alice) never to start.

We now describe a partial pragmatic remedy for this problem. The basic idea is to modify the protocol by changing the instructions what to do when the other parties misbehave. Recall, that in our protocols the parties are instructed to simply redeem their all their transactions if they notice a suspicious behavior of the other party. Now, instead of doing this, they could keep these transactions on the Ledger and reuse them in some other sessions of the protocol. Of course, this has to be done with care. For example the timed commitment schemes have to be redeemed within a certain time frame). One also has to be careful to avoid problems with reusing the keys, described above, cf. Footnote 18. To argue formally about the security of this solution one would need to introduce a multi-player mathematical model capturing the fact that several sessions can be executed with shared secrets. This is beyond the scope of this paper, so we just stay on this informal level.

VII. IMPLEMENTATION

As a “proof of concept” we have implemented and executed the presented protocols. The transactions were created using *bitcoinj* Java library as normal BitCoin clients do not allow user to create (nor broadcast) non-standard transactions and sent directly to *Eligius* mining pool. Below we present links to some of these transactions on the blockchain.info website. To save space, all links are relative to the url `blockchain.info/tx-index/` (hence, they are only indices of transactions used by the blockchain.info site).

Commitment scheme CS: Links to all transactions in a correct execution of the commitment scheme with one recipient are as follows: *Commit*: 97079150; *Open*: 97094781.

Here is an example of an execution for two recipients, which finished with *PayDeposit* transactions broadcast: *Commit*: 96947667; *PayDeposit*¹: 96982401; and *PayDeposit*²: 96982398.

Three-party lottery (MultiPlayersLottery): We have performed a correct execution of the three-party lottery protocol, where each player bets 0.0012 ₤ . First, the players perform standard transactions with output value 0.0012 ₤ . The transactions are as follows: $PutMoney^A$: 96946847; $PutMoney^B$: 96946887; and $PutMoney^C$: 96947563. Then the players exchange the hashes h_A , h_B and h_C , and sign and broadcast the *Compute* transaction (96964833). After the revealing of their secrets s_A , s_B and s_C (by opening the commitments), the winner (in this case, player C) broadcasts the $ClaimMoney_C$ transaction (96966124) to get the pot.

Two-party lottery without deposits (TwoPlayersLottery): Below we present links to all transactions in a correct execution of the protocol won by Alice: $PutMoney^A$: 96424665; $PutMoney^B$: 96436412; *Compute*: 96436416; $ClaimMoney^A$: 96436417; In this execution, the players bet 0.04 ₤ each and the transaction fees were set to 0.0001 ₤ for each transaction.

We also performed an execution which finished with the *Fuse* transaction: *PutMoney^A*: 97094615; *PutMoney^B*: 97094780; *Compute*: 97099280; *Fuse*: 97105484.

As an example of a raw transaction we present the *PutMoney^B* transaction from the first of the TwoPlayersLottery protocol executions described above (96436412) in more details. Here is its dump (with some fields omitted):

```
{ "lock_time":0,
  "in":[{"prev_out": {"hash":"a14...096", "n":0},
        "scriptSig":"304...a01 039...443" }],
  "out":[{"value":"0.03990000",
          "scriptPubKey":
            OP_SIZE 32 34 OP_WITHIN OP_VERIFY
            OP_SHA256 f53...226 OP_EQUALVERIFY
            020...e33 OP_CHECKSIG" }]}
```

The meaning of the above is as follows. "lock_time":0 means, that the transaction does not have a time lock. "hash":"a14...096" denotes a hash of the transaction, which is being redeemed in *PutMoney^B* and "n":0 denotes, which output of that transaction is being redeemed. The input script "scriptSig":"304...a01 039...443" consists of the signature (039...443) on *PutMoney^B* under the key $B.pk$ and the public key $B.pk$ itself (304...a01) (in standard transaction's output there is only pk 's hash and pk has to be included in the corresponding input).

The output script expects to get as input an appropriate signature and Bob's secret string. The script consists of three parts: the first part OP_SIZE 32 34 OP_WITHIN OP_VERIFY checks whether the second argument has an appropriate length; the second part OP_SHA256 f53...226 OP_EQUALVERIFY checks if its hash is equal to h_B i.e. f53...226); and the last part 020...e33 OP_CHECKSIG checks if the first argument is an appropriate signature under key $\tilde{B}.pk$ (i.e. 020...e33).

REFERENCES

- [1] I. Abraham, D. Dolev, R. Gonen, and J. Y. Halpern. Distributed computing meets game theory: robust mechanisms for rational secret sharing and multiparty computation. In *PODC*, 2006.
- [2] S. Barber, X. Boyen, E. Shi, and E. Uzun. Bitter to better - how to make Bitcoin a better currency. In *FC*, 2012.
- [3] D. Basin, C. Caleiro, J. Ramos, and L. Vigano. Distributed temporal logic for the analysis of security protocol models. *TCS*, 2011.
- [4] A. Beimel, Y. Lindell, E. Omri, and I. Orlov. $1/p$ -Secure multiparty computation without honest majority and the best of both worlds. In *CRYPTO*, 2011.
- [5] A. Ben-David, N. Nisan, and B. Pinkas. Fairplaymp: a system for secure multi-party computation. In *CCS*, 2008.
- [6] M. Blum. Coin flipping by telephone. In *CRYPTO*, 1981.
- [7] D. Boneh and M. Naor. Timed commitments. In *CRYPTO*, 2000.
- [8] G. Brassard, D. Chaum, and C. Crépeau. Minimum disclosure proofs of knowledge. *J. Comput. Syst. Sci.*, 1988.
- [9] V. Buterin. The bitcoin gambling diaspora, 2013. Bitcoin Magazine.
- [10] V. Buterin. SatoshiDice sold for \$12.4 million, 2013. Bitcoin Magazine.
- [11] C. Cachin and J. Camenisch. Optimistic fair secure computation. In *CRYPTO*, 2000.
- [12] J. Camenisch, S. Hohenberger, and A. Lysyanskaya. Compact e-cash. In *EUROCRYPT*, 2005.
- [13] D. Chaum. Blind signature system. In *CRYPTO*, 1983.
- [14] D. Chaum, A. Fiat, and M. Naor. Untraceable electronic cash. In *CRYPTO*, 1988.
- [15] J. Clark and A. Essex. CommitCoin: Carbon dating commitments with Bitcoin - (short paper). In *FC*, 2012.
- [16] R. Cleve. Limits on the security of coin flips when half the processors are faulty. *STOC*, 1986.
- [17] I. Damgård et al. Practical covertly secure mpc for dishonest majority - or: Breaking the SPDZ limits. In *ESORICS*. 2013.

- [18] D. Dolev and A. C. Yao. On the security of public key protocols. *Information Theory, IEEE Transactions on*, 1983.
- [19] C. Dwork and M. Naor. Pricing via processing or combatting junk mail. In *CRYPTO*, 1992.
- [20] The Economist. Online gambling: Know when to fold, 2013.
- [21] J. A. Garay et al. Rational protocol design: Cryptography against incentive-driven adversaries. In *FOCS*. 2013.
- [22] P. Bogetoft et al. Secure multiparty computation goes live. In *FC*, 2009.
- [23] E. J. Friedman and P. Resnick. The social cost of cheap pseudonyms. *Journal of Economics and Management Strategy*, 10:173–199, 2000.
- [24] J. Garay and M. Jakobsson. Timed release of standard digital signatures. In *FC*, 2003.
- [25] O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game. *STOC*, 1987.
- [26] S. Dov Gordon and J. Katz. Partial fairness in secure two-party computation. In *EUROCRYPT*, 2010.
- [27] J. Y. Halpern and V. Teague. Rational secret sharing and multiparty computation: Extended abstract. In *STOC*, 2004.
- [28] D. Malkhi, N. Nisan, B. Pinkas, and Y. Sella. Fairplay — a secure two-party computation system. In *USENIX Security Symposium*, 2004.
- [29] I. Miers, C. Garman, M. Green, and A. D. Rubin. Zerocoin: Anonymous distributed e-cash from bitcoin. *IEEE S&P*, 2012.
- [30] S. Nakamoto. Bitcoin: A peer-to-peer electronic cash system, 2008.
- [31] B. Pinkas. Fair secure two-party computation. In *Advances in Cryptology – EUROCRYPT 2003*, 2003.
- [32] The Washington Post. Cheating scandals raise new questions about honesty, security of internet gambling, November 30, 2008.
- [33] P. Resnick, K. Kuwabara, R. Zeckhauser, and E. Friedman. Reputation systems. *Commun. ACM*, 43(12):45–48, December 2000.
- [34] D. Ron and A. Shamir. Quantitative analysis of the full bitcoin transaction graph. In *FC*. 2013.
- [35] A. Shamir, R. Rivest, and L. Adleman. Mental poker, April 1979. Technical Report LCS/TR-125, Massachusetts Institute of Technology.
- [36] L. von Ahn, N. J. Hopper, and J. Langford. Covert two-party computation. In *STOC*, 2005.
- [37] A. C.-C. Yao. How to generate and exchange secrets. In *FOCS*, 1986.