

Smart contracts meet quantum cryptography

Andrea Coladangelo

Computing and Mathematical Sciences, Caltech
acoladan@caltech.edu

February 15, 2019

Abstract

We put forward the idea that classical blockchains and smart contracts are potentially useful primitives not only for classical cryptography, but for quantum cryptography as well. Abstractly, a smart contract is a functionality that allows parties to deposit funds, and release them upon fulfillment of algorithmically checkable conditions, and can thus be employed as a formal tool to enforce monetary incentives. In this work, we give the first example of the use of smart contracts in a quantum setting.

We describe a hybrid classical-quantum payment system whose main ingredients are a classical blockchain capable of handling stateful smart contracts, and quantum lightning, a strengthening of public-key quantum money introduced by Zhandry [Zha17]. The resulting hybrid payment system uses quantum states as banknotes and a classical blockchain to settle disputes and to keep track of the valid serial numbers. It has several desirable properties: it is decentralized, requiring no trust in any single entity; payments are as quick as quantum communication, regardless of the total number of users; when a quantum banknote is damaged or lost, the rightful owner can recover the lost value.

Contents

1	Introduction	2
2	Preliminaries	5
2.1	Basic notation	5
2.2	Quantum money and quantum lightning	5
2.3	Universal Composability	7
3	Blockchains and smart contracts	9
4	A payment system based on quantum lightning and a classical blockchain	13
4.1	The payment system and its components	14
5	Security	18
6	Conclusion	21
A	Appendix	23
A.1	Proof of Proposition 2	23

1 Introduction

Cryptocurrencies, along with blockchains and smart contracts, have recently risen to popular attention, with the most well-known examples being Bitcoin and Ethereum [Nak08, But08]. This rise in popularity has been accompanied by increasing discussion about the potential for applications of blockchains and smart contracts beyond simple transactions.

Informally, a blockchain is a public ledger consisting of a sequence of blocks. Each block typically contains information about a set of transactions, and a new block is appended regularly via a consensus mechanism that involves the parties of a network and no a priori trusted authority. A blockchain is endowed with a native currency which is employed in transactions, and whose basic unit is a “coin”. The simplest type of transaction is a *payment*, which transfers coins from one party to another. However, more general transactions are allowed, which are known as *smart contracts*. These can be thought of as contracts stored on a blockchain, and whose consequences are executed upon fulfilment of algorithmically checkable conditions.

The idea of introducing monetary incentives has led to interesting applications in classical cryptography via the study of rational proofs [AM12, GHRV14, GHRV16, CMS16, AGP16]. Smart contracts on a blockchain make it possible to enforce these monetary incentives in a trust-less way, i.e. without having to trust a central entity, or that the verifier pays the correct rewards to the provers. They have found several applications in the construction of multi-party cryptographic primitives like secure multi-party computation [BK14], secure lotteries [BK14], and protocols for playing card-games [BKM17, DDL17, DDL18].

A central issue that needs to be resolved for blockchains to achieve mass-adoption is *scalability*. This refers to the problem of increasing the throughput of transactions (i.e. transactions per second) while maintaining the resources needed for a party to participate in the consensus mechanism approximately constant, and while maintaining security against adversaries that can corrupt constant fractions of the parties in the network. For example, Bitcoin and Ethereum, can currently handle only on the order of 10 transactions per second (Visa for a comparison handles about 3500 per second).

In this work, we show that smart contracts can be combined with tools from quantum cryptography to provide a simple quantum solution to the scalability problem. We clarify that this solution only solves the scalability problem for *payment* transactions, and not for the more general *smart contracts* transactions.

The main quantum ingredient that we employ is a primitive called quantum lightning, introduced by Zhandry [Zha17] and inspired by Lutomirski et al.’s notion of collision resistant quantum money [LAF⁺09]. In a public-key quantum money scheme, a bank is entrusted with generating quantum states (we refer to these as quantum banknotes) with an associated serial number, and a public verification procedure allows anyone in possession of the banknote to check its validity. Importantly, trust is placed in the fact that the central bank will not create multiple quantum banknotes with the same serial number. A quantum lightning scheme has the additional feature that no generation procedure, not even the honest one (and hence not even a bank!), can produce two valid money states with the same serial number, except with negligible probability. This opens to the possibility of having a completely decentralized quantum money scheme. However, if the system ought to be trust-less, some issues need to be addressed: most importantly, who is allowed to mint money? Who decides which serial numbers are valid? A second kind of issue that one might consider when thinking about a future practical deployment of such a money scheme is that quantum states almost inevitably deteriorate with time. Unless there exists a technology reliable enough to store large states for extended periods of time, quantum money would most likely not be feasible unless there is a mechanism that allows a user to replace deteriorated banknotes with fresh ones, or to recover the value that is lost when a banknote is completely damaged.

Our solution leverages smart contracts to provide such a mechanism, and to address all of the above issues. It constitutes the first example of the use of smart contracts in conjunction with quantum tools. We suspect that this is not an isolated example, but that smart contracts could find many other useful applications in quantum cryptography.

Blockchains and smart contracts as ideal functionalities The essential properties of blockchains and smart contracts can be abstracted by modeling them as ideal functionalities

in the Universal Composability (UC) framework of Canetti [Can01]. Such a framework provides both a formal model for multiparty computation, and formal notions of security with strong composability properties. The approach of studying blockchains and smart contracts within such a framework was first proposed by Bentov et al. in [BK14], and explored further in [BKM17]. The main reason why such an approach is desirable is that it abstracts the features of blockchains and smart contracts into building blocks that can be utilized to design more complex protocols in a modular way. The limitation of the works [BK14, BKM17] is that they modify the original model of computation of UC in order to incorporate coins, but they do not prove that a composition theorem holds in this variant. A more natural approach was proposed by Kiayias et al. in [KZZ16]. There, the authors define a global ideal functionality which abstracts the properties of a transaction-ledger and fits into the well-studied framework of Generalized Universal Composability (GUC) of Canetti et al. [CDPW07], and thus inherits its composability properties. Other works [GKL15, PSS17] have operated within a similar framework, and have focused on analyzing security of the Bitcoin protocol. In particular, to the best of our knowledge, Badertscher et al. [BMTZ17] are the first to propose an ideal functionality for a Bitcoin-like transaction ledger within the GUC framework, together with a secure realization on (an abstraction of) the real-world Bitcoin network.

Our contributions Our contribution is two-fold: first, we introduce a new global ideal functionality in the GUC framework. Like the previously mentioned works, this functionality abstracts the features of a transaction ledger, but, additionally, is capable of handling a particular type of smart contracts called *stateful* smart contracts. Our approach is inspired by the previous work of Bentov et al. [BKM17] (which also considers stateful smart contracts) but differs primarily in that our ideal functionality fits into the framework of GUC, for which composition theorems are known. Our ideal functionality is arguably simpler to analyze and more coarse-grained than the functionalities in [KZZ16] and [BMTZ17], but also stronger in that it handles stateful smart contracts. Since Bitcoin’s scripting language does not allow for stateful smart contracts, our functionality is not realizable on the Bitcoin blockchain. We believe, however, that our functionality, or close variations of it, should have a GUC-secure realization on the Ethereum blockchain, but we do not explicitly provide such a realization (as we expect such a realization to require a fine-grained analysis of the workings of Ethereum). Rather, we assume access to such a functionality, and we focus in this work on its application in a quantum setting. We emphasize, however, that the composition theorems of GUC guarantee that one can replace our ideal functionality with any GUC-secure realization of it on any blockchain, while maintaining security of the protocols that employ it. Our analysis is thus agnostic to the particular real-world implementation.

Second, we provide, to the best of our knowledge, the first example of the use of smart contracts in a quantum setting. We design a hybrid classical-quantum payment system that uses quantum states as banknotes and a classical blockchain to settle disputes and to keep track of the valid serial numbers. This has several desirable features:

- It is decentralized, requiring no trust in any single entity.
- Payments are as quick as quantum communication, regardless of the total number of users.
- When a quantum banknote is damaged or lost, the rightful owner can recover the lost value.

As mentioned earlier, the main quantum ingredient that we employ is quantum lightning [Zha17], a strengthening of public-key quantum money with the property that no generation procedure (not even the honest one) can create two banknotes with the same serial number except with negligible probability. As suggested by Zhandry, this property is desirable if one wants to design a decentralized payment system, as it prevents anyone from copying banknotes (even a hypothetical bank). However, since the generation procedure is publicly known, there needs to be a mechanism that regulates the generation of new valid banknotes (to prevent parties from continuously generating banknotes).

We show how stateful smart contracts on a classical blockchain can be used to provide such a mechanism. For instance, they allow to easily keep track of a publicly trusted list of valid serial numbers. We elaborate on this. In our payment system, the native coin of a classical blockchain is used as a baseline classical currency. Any party can spend coins on the classical blockchain to add a serial number of their choice to the list of valid serial numbers. More precisely, they can

deposit any amount (of their choice) d of coins into an appropriately specified smart contract and set the initial value of a *serial number* state variable to whatever they wish (presumably the serial number of a quantum banknote that they have just generated locally). They have thus effectively added to the blockchain a serial number associated to a quantum banknote that, in virtue of this, we can think of as having “acquired” value d . Payments are made by transferring quantum banknotes: party A , the payer, transfers his quantum banknote with serial number s to party B , the payee, and references a smart contract whose serial number state variable is s . Party B then locally verifies that the received banknote is valid. The appeal of such a payment protocol is that transactions only involve the payer and the payee: no consensus mechanism is required to validate the payment, and the payment need not be recorded on the blockchain. The latter is only invoked when a banknote is first generated, and in case of a dispute. Thus, throughput of transactions is no longer a concern. Likewise, long waiting times between when a payment is initiated and when it is confirmed are also no longer a concern since payments are verified immediately by the payee. The full payment system, which is described formally in section 4, includes a mechanism that allows any party in possession of a quantum banknote to recover the coins deposited in the corresponding smart contract. It also includes a mechanism that allows any honest party who has lost or damaged a valid quantum banknote to change the *serial number* state variable of the corresponding smart contract to a fresh value of their choice.

We should note that there is currently only one known construction of quantum lightning, by Zhandry [Zha17], who proved its security under a computational assumption related to the multi-collision resistance of certain degree-2 hash functions (which is arguably not very well studied). Zhandry also shows that secure quantum lightning can be constructed using any non-collapsing hash function, but no example of a provably non-collapsing hash function is currently known (for more details we refer to [Zha17]). In our work, we assume existence of a secure quantum lightning scheme. Moreover, we note that in order to be realized and to be practical, our payment system requires, among other quantum technologies, the ability to store quantum states for extended periods of time, often referred to as “quantum memory”. It also requires the ability of each party to send quantum states to any party that it wishes to make payments to. Hence, it is unlikely that our scheme will be realizable in the near future. The second requirement could be achieved, for example, in a network in which any two parties have the ability to request joint EPR pairs (i.e. maximally entangled pairs of qubits), which is one of the primary components of a “quantum internet” [WEH18]. One can view our scheme as a possible use case of a quantum internet. One drawback (or advantage?) of our payment system is that, much like the classical banknotes that we currently use, it does not provide the payer with a receipt or a proof that the payment has happened. Payments leave no trace on the blockchain. In virtue of this, a malicious party who receives a quantum banknote can always claim that they did not receive it, when in fact they did. This shortcoming can be mitigated by having payers split their payment into quantum banknotes of smaller value, and requiring the payee to provide receipts for previous parts of the payment, before proceeding with the next. (Or alternatively, payers can just be careful with who they send their quantum banknotes to.)

Finally, we point out that, while it is the first one to use quantum resources, our proposal is not the first that tries to increase the transaction throughput by designing a scheme in which most of the transactions happen off-chain. Most notably, Poon and Dryja, proposed a (completely classical) scheme called Lightning Network [PD16] (where the terminology is unrelated to quantum lightning!). While a promising proposal, its feasibility is much more complex to analyze and still under scrutiny. As of February 2019, this proposal is currently still being tested at a small scale, and it is unclear if, when, and in what form the Lightning Network will reach mass-adoption. One example of a drawback of the Lightning Network in its peer-to-peer form is that it requires users to deposit into several smart contracts, and for an extended period of time, more coins than they actually expect to spend.

Outline Section 2 covers preliminaries: 2.1 covers basic notation; 2.2 introduces quantum lightning; 2.3 gives a concise overview of the Universal Composability framework of Canetti [Can01]. Section 3 gives first an informal description of blockchains and smart contracts, followed by a formal definition of our global ideal functionality for a transaction ledger that handles stateful smart contracts. Section 4 describes our payment system. In section 5, we describe an adversarial

model, and then prove security guarantees with respect to it.

2 Preliminaries

2.1 Basic notation

For a function $f : \mathbb{N} \rightarrow \mathbb{R}$, we say that f is *negligible*, and we write $f(n) = \text{negl}(n)$, if for any positive polynomial $p(n)$ and all sufficiently large n 's, $f(n) < \frac{1}{p(n)}$. A *binary* random variable is a random variable over $\{0, 1\}$. We say that two ensembles of binary random variables $\{X_n\}$ and $\{Y_n\}$ are *indistinguishable* if,

$$|\Pr[X_n = 1] - \Pr[Y_n = 1]| = \text{negl}(n).$$

We use the term PPT as an abbreviation for probabilistic polynomial time.

2.2 Quantum money and quantum lightning

Quantum money is a theoretical form of payment first proposed by Wiesner [Wie83], which replaces physical banknotes with quantum states. In essence, a quantum money scheme consists of a generation procedure, which mints banknotes, and a verification procedure, which verifies the validity of minted banknotes. A banknote consists of a quantum state together with an associated serial number. The appeal of quantum money comes primarily from a fundamental theorem in quantum theory, the No-Cloning theorem, which informally states that there does not exist a quantum operation that can clone arbitrary states. A second appealing property of quantum money, which is not celebrated nearly as much as the first, is that quantum money can be transferred almost instantaneously (by quantum teleportation for example). The first proposals for quantum money schemes required a central bank to carry out both the generation and verification procedures. The idea of public key quantum money was later formalized by Aaronson [Aar09]. In public-key quantum money, the verification procedure is public, meaning that anyone with access to a quantum banknote can verify its validity.

In this section, we focus on quantum lightning, a primitive recently proposed by Zhandry [Zha17], and we enhance this to a decentralized quantum payment system. Informally, a quantum lightning scheme is a strengthening of public-key quantum money. It consists of a public generation procedure and a public verification procedure which satisfy the following two properties:

- Any quantum banknote generated by the honest generation procedure is accepted with probability negligibly close to 1 by the verification procedure.
- No adversarial generation procedure (not even the honest one) can generate two banknotes with the same serial number which both pass the verification procedure with non-negligible probability.

As mentioned earlier, there is only one known construction of quantum lightning, by Zhandry [Zha17], who gives a construction which is secure under a computational assumption related to the multi-collision resistance of some degree-2 hash function. Zhandry also proves that any non-collapsing hash function can be used to construct quantum lightning. However, to the best of our knowledge, there are no known hash functions that are proven to be non-collapsing. In this section, we define quantum lightning formally, but we do not discuss any possible construction. Rather, in section 4, we will use quantum lightning as an off-the-shelf primitive.

Definition 1 (Quantum lightning [Zha17]). *A quantum lightning scheme consists of a PPT setup procedure $QL.Setup(1^\lambda)$ (where λ is a security parameter) which samples a pair of polynomial-time quantum algorithms (Gen, Ver). Gen samples states $|\psi\rangle \in \mathcal{H}_s$, which we refer to as “bolts”. Ver takes as input a state in \mathcal{H}_s and outputs either $s \in \{0, 1\}^\lambda$ (the “serial number”) or \perp . They satisfy the following:*

- Let $|\psi\rangle$ be any bolt generated by Gen, and let

$$H_{min}(|\psi\rangle, Ver) := -\log_2 \max_s \Pr[Ver(|\psi\rangle) = s].$$

Then,

$$\mathbb{E}[H_{\min}(|\psi\rangle, \text{Ver})] = \text{negl}(\lambda),$$

where the expectation is taken over $(\text{Gen}, \text{Ver}) \leftarrow \text{QL.Setup}(1^\lambda)$ and $|\psi\rangle \leftarrow \text{Gen}$.

The latter requirement simply asks that for any honestly generated bolt, there is a serial number s which, with overwhelming probability, is the output of the verification procedure applied to the bolt. For security, we require that no adversarial generation procedure can produce two bolts with the same serial number. Formally, we define security via the following game between a challenger and an adversary \mathcal{A} .

- The challenger runs $(\text{Gen}, \text{Ver}) \leftarrow \text{QL.Setup}(\lambda)$ and sends (Gen, Ver) to \mathcal{A} .
- \mathcal{A} produces a state $|\Psi_{12}\rangle \in \mathcal{H}_s^{\otimes 2}$.
- The challenger runs $\text{Ver} \otimes \text{Ver}$ on $|\Psi_{12}\rangle$ obtaining serial numbers s_0, s_1 . The output of the game is 1 if $s_0 = s_1 \neq \perp$, and 0 otherwise.

We let $\text{Copy}(\lambda, \mathcal{A})$ be the random variable which denotes the output of the game.

Definition 2 (Security [Zha17]). *A quantum lightning scheme is secure if, for all polynomial-time quantum adversaries \mathcal{A} ,*

$$\Pr[\text{Copy}(\lambda, \mathcal{A}) = 1] = \text{negl}(\lambda).$$

We define an additional property of a quantum lightning scheme, which in essence establishes that one can trade a quantum banknote for some useful classical information. Intuitively, this is meant to capture the fact that in the construction of quantum lightning based on non-collapsing hash functions proposed by Zhandry, one can measure a bolt with serial number y in the computational basis to obtain a pre-image of y under the hash function. However, doing so damages the bolt so that it will no longer pass verification. In order to define this additional property, we change the procedure $\text{QL.Setup}(1^\lambda)$ slightly, so that it outputs a triple $(\text{Gen}, \text{Ver}, H)$ where $H : \{0, 1\}^{l(\lambda)} \rightarrow \{0, 1\}^\lambda \cup \{\perp\}$ (for some polynomial $l(\lambda)$). The additional property is defined based on the following game between a challenger and an adversary \mathcal{A} .

- The challenger runs $(\text{Gen}, \text{Ver}, H) \leftarrow \text{QL.Setup}(1^\lambda)$ and samples $|\psi\rangle \leftarrow \text{Gen}$. Let $s := \text{Ver}(|\psi\rangle)$. The challenger sends $(\text{Gen}, \text{Ver}, H, |\psi\rangle)$ to \mathcal{A} .
- \mathcal{A} returns $x \in \{0, 1\}^\lambda$ and $|\psi'\rangle$.
- \mathcal{A} wins if the following holds:
 - (i) $H(x) = s$.
 - (ii) $\text{Ver}(|\psi'\rangle) = s$.

Let $\text{Game}_1(\mathcal{A}, \lambda)$ be the random variable that is 1 if condition (i) is satisfied and 0 otherwise. Let $\text{Game}_{12}(\mathcal{A}, \lambda)$ be the random variable that is 1 if both conditions (i) and (ii) are satisfied and 0 otherwise.

Definition 3 (Trading the bolt for a classical certificate). *Let $\lambda \in \mathbb{N}$. We say that a quantum lightning scheme has “bolt-to-certificate capability” if:*

(I) *There exists a polynomial-time quantum algorithm \mathcal{A}_* such that*

$$\Pr[\text{Game}_1(\mathcal{A}_*, \lambda) = 1] = 1 - \text{negl}(\lambda)$$

(II) *For all polynomial-time quantum algorithms \mathcal{A} ,*

$$\Pr[\text{Game}_{12}(\mathcal{A}, \lambda) = 1] = \text{negl}(\lambda).$$

Notice that property (II) also implies that for most H and s it is hard for any adversary to find x such that $H(x) = s$ without access to a valid state whose serial number is s . In fact, if there was an adversary \mathcal{A} which succeeded at that, this could clearly be used to construct an adversary \mathcal{A}' that succeeds in Game_{12} : Upon receiving $(\text{Gen}, \text{Ver}, H, |\psi\rangle)$ from the challenger, \mathcal{A}' computes $s := \text{Ver}(|\psi\rangle)$ (which leaves $|\psi\rangle$ unchanged); then runs \mathcal{A} on input s to obtain some x . \mathcal{A}' returns x and $|\psi\rangle$ to the challenger.

Proposition 1. *Any scheme that uses Zhandry’s construction [Zha17] instantiated with a non-collapsing hash function satisfies the property of Definition 3.*

Proof. In Zhandry’s construction based on a non-collapsing hash function, $\text{QL.Setup}(1^\lambda)$ outputs $(\text{Gen}, \text{Ver}, H)$, where $n \in \mathbb{N}$ is polynomial in the security parameter λ , and $H = \tilde{H}^{\times n}$ is the n -fold product (or concatenation) of a non-collapsing hash-function \tilde{H} , i.e. $H(x_1, \dots, x_n) := (\tilde{H}(x_1), \dots, \tilde{H}(x_n))$. A bolt generated from Gen has the form $|\Psi\rangle = \bigotimes_{i=1}^n |\psi_{y_i}\rangle$, where $y_i \in \{0, 1\}^\lambda$ for all i , and $|\psi_{y_i}\rangle = \sum_{x: \tilde{H}(x)=y_i} |x\rangle$, and Ver has the form of n -fold product of a verification procedure Mini-Ver which acts on a single register.

To show that (I) holds, take \mathcal{A}_* to be the following: \mathcal{A}_* measures each register of the received state in the computational basis to obtain pre-images x_i , for $i = 1, \dots, n$. Then, returns $x = (x_1, \dots, x_n)$ to the challenger (notice that such \mathcal{A}_* clearly destroys the received state, and does not help in winning Game_{12}). For property (II), suppose there exists \mathcal{A} such that $\Pr[\text{Game}_{12}(\mathcal{A}, \lambda) = 1]$ is non-negligible. We use \mathcal{A} to construct an adversary \mathcal{A}' that breaks collision-resistance of a non-collapsing hash function as follows: \mathcal{A}' runs $(\text{Gen}, \text{Ver}, H) \leftarrow \text{QL.Setup}(1^\lambda)$ (where $H = \tilde{H}^{\times n}$ for some non-collapsing hash function \tilde{H}), and samples $|\psi\rangle \leftarrow \text{Gen}$. \mathcal{A}' gives $(\text{Gen}, \text{Ver}, H, |\psi\rangle)$ as input to \mathcal{A} . \mathcal{A} returns x which is parsed as (x_1, \dots, x_n) and $|\psi'\rangle$. \mathcal{A}' then measures each of the n registers of $|\psi'\rangle$ to get $x' = (x'_1, \dots, x'_n)$. If $x_i \neq x'_i$, then \mathcal{A}' outputs (x_i, x'_i) . With non-negligible probability this pair is a collision for H : in fact, since \mathcal{A} wins Game_{12} with non-negligible probability, then $|\psi'\rangle$ must pass Ver with non-negligible probability; and from the analysis of [Zha17], any such state must be such that at least one of the registers is in a superposition which has non-negligible weight on at least two pre-images. \square

Proposition 2. *Zhandry’s construction based on the multi-collision resistance of certain degree-2 hash functions (from section 6 of [Zha17]) satisfies the property of Definition 3.*

Proof. The proof is similar to the proof of Proposition 1. We include it for completeness in the appendix (A.1). \square

2.3 Universal Composability

This subsection is intended as a concise primer about the Universal Composability (UC) model of Canetti [Can01]. We refer the reader to [Can01] for a rigorous definition and treatment of the UC model and of composable security, and to the tutorial [Can06] for a more gentle introduction. At the end, we provide a brief overview of the Generalized UC model (GUC) [CDPW07]. While introducing UC and GUC, we also setup some of the notation that we will employ in the rest of the paper. The reader familiar with these concepts may wish to skip ahead to the next section.

In the universal composability framework (UC), parties are modelled as Interactive Turing Machines (ITM) who can communicate by writing on each other’s externally writable tapes, subject to some global constraints. Informally, a protocol is specified by a code π for an ITM, and consists of various rounds of communication and local computation between instances of ITMs (the parties), each running the code π on their machine, on some private input.

Security in the UC model is defined via the notion of emulation. Informally, we say that a protocol π emulates a protocol ϕ if whatever can be achieved by an adversary attacking π can also be achieved by some other adversary attacking ϕ . This is formalized by introducing simulators and environments.

Given protocols π and ϕ , we say that π emulates (or "is as secure as") ϕ in the UC model, if for any polynomial-time adversary \mathcal{A} attacking protocol π , there exists a polynomial-time simulator \mathcal{S} attacking ϕ such that no polynomial-time distinguisher \mathcal{E} , referred to as the *environment*, can distinguish between π running with \mathcal{A} and ϕ running with \mathcal{S} . Here, the environment \mathcal{E} is allowed to choose the protocol inputs, read the protocol outputs, including outputs from the adversary or the simulator, and to communicate with the adversary or simulator during the execution of the protocol (without of course being told whether the interaction is with the adversary or with the simulator). In this framework, one can formulate security of a multiparty cryptographic task by first defining an *ideal functionality* \mathcal{F} that behaves exactly as intended, and then providing a "real-world" protocol that emulates, or "securely realizes", the ideal functionality \mathcal{F} .

We give more formal definitions for the above intuition. To formulate precisely what it means for an environment \mathcal{E} to tell two executions apart, one has to formalize the interaction between \mathcal{E} and the protocols in these executions. Concisely, an execution of a protocol π with adversary \mathcal{A} and environment \mathcal{E} consists of a sequence of activations of ITMs. At each activation, the active ITM runs according to its code, its state and the content of its tapes, until it reaches a special `wait` state. The sequence of activations proceeds as follows: The environment \mathcal{E} gets activated first and chooses inputs for \mathcal{A} and for all parties. Once \mathcal{A} or a party is activated by an incoming message or an input, it runs its code until it produces an outgoing message for another party, an output for \mathcal{E} , or it reaches the `wait` state, in which case \mathcal{E} is activated again. The execution terminates when \mathcal{E} produces its output, which can be taken to be a single bit. Note that each time it is activated, \mathcal{E} is also allowed to invoke a new party, and assign a unique PID (party identifier) to it. Allowing the environment to invoke new parties will be particularly important in section 5, where we discuss security. There, the fact that the environment has this ability implies that our security notion captures realistic scenarios in which the set of parties is not fixed at the start, but is allowed to change. Moreover, each invocation of a protocol π is assigned a unique session identifier SID, to distinguish it from other invocations of π . We denote by $\text{EXEC}_{\pi, \mathcal{A}, \mathcal{E}}(\lambda, z)$ the output of environment \mathcal{E} initialized with input z , and security parameter λ in an execution of π with adversary \mathcal{A} .

We are ready to state the following (slightly informal) definition.

Definition 4. *A protocol π UC-emulates a protocol ϕ if, for any PPT adversary \mathcal{A} , there exists a PPT simulator \mathcal{S} such that, for any PPT environment \mathcal{E} , the families of random variables $\{\text{EXEC}_{\pi, \mathcal{A}, \mathcal{E}}(\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^{\text{poly}(\lambda)}}$ and $\{\text{EXEC}_{\phi, \mathcal{S}, \mathcal{E}}(\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^{\text{poly}(\lambda)}}$ are indistinguishable.*

Then, given an ideal functionality \mathcal{F} which captures the intended ideal behaviour of a certain cryptographic task, one can define the ITM code $I_{\mathcal{F}}$, which behaves as follows: the ITM running $I_{\mathcal{F}}$ simply forwards any inputs received to the ideal functionality \mathcal{F} . We then say that a “real-world” protocol π securely realizes \mathcal{F} if π emulates $I_{\mathcal{F}}$ according to Definition 4.

A composition theorem The notion of security we just defined is strong. One of the main advantages of such a security definition is that it supports composition, i.e. security remains when secure protocols are executed concurrently, and arbitrary messages can be sent between executions. We use the notation σ^{π} for a protocol σ that makes up to polynomially many calls to another protocol π . In a typical scenario, $\sigma^{\mathcal{F}}$ is a protocol that makes use of an ideal functionality \mathcal{F} , and σ^{π} is the protocol that results by implementing \mathcal{F} through the protocol π (i.e. replacing calls to \mathcal{F} by calls to π). It is natural to expect that if π securely realizes \mathcal{F} , then σ^{π} securely realizes $\sigma^{\mathcal{F}}$. This is the content of the following theorem.

Theorem 1 (Universal Composition Theorem). *Let π, ϕ, σ be polynomial-time protocols. Suppose protocol π UC-emulates ϕ . Then σ^{π} UC-emulates σ^{ϕ} .*

Replacing ϕ by $I_{\mathcal{F}}$ for some ideal functionality \mathcal{F} in the above theorem yields the composable security notion discussed above.

Generalized UC model The formalism of the original UC model is not able to handle security requirements in the presence of a “global trusted setup”. By this, we mean some global information accessible to all parties, which is guaranteed to have certain properties. Examples of this are a public-key infrastructure or a common reference string. Emulation in the original UC sense is not enough to guarantee composability properties in the presence of a global setup. Indeed, one can construct examples in which a UC-secure protocol for some functionality interacts badly with another UC-secure protocol and affects its security, if both protocols make reference to the same global setup. For more details and concrete examples see [CDPW07].

The generalized UC framework (GUC) of Canetti et al. [CDPW07] allows for a “global setup”. The latter is modelled as an ideal functionality which is allowed to interact not only with the parties running the protocol, but also with the environment. GUC formulates a stronger security notion, which is sufficient to guarantee a composition theorem, i.e. ideal functionalities with access to a shared global functionality \mathcal{G} can be replaced by protocols that securely realize them in the presence of \mathcal{G} . Further, one can also replace global ideal functionalities with appropriate protocols realizing them. This kind of replacement does not immediately follow from the previous

composition theorem and requires a more careful analysis, as is done in [CSV16], where sufficient conditions for this replacement are established.

Universal Composability in the quantum setting In our setting, we are interested in honest parties, adversaries and environments that are quantum polynomial-time ITMs. The notion of Universal Composability has been studied in the quantum setting in [BOM04], [Unr04] and [Unr10]. In particular, in [Unr10], Unruh extends the model of computation of UC and its composition theorems to the setting in which polynomial-time classical ITMs are replaced by polynomial-time quantum ITMs (and ideal functionalities are still classical). The proofs are essentially the same as in the classical setting. Although the quantum version of the Generalized UC framework has not been explicitly studied in [Unr10], one can check that the proofs of the composition theorems for GUC from [CDPW07] and [CSV16] also go through virtually unchanged in the quantum setting.

3 Blockchains and smart contracts

In this section, we start by describing blockchains and smart contracts informally. We follow this by a more formal description. As mentioned in the introduction, the essential features of blockchains and smart contracts can be abstracted by modeling them as ideal functionalities in the Universal Composability framework of Canetti [Can01]. In this section, we introduce a global ideal functionality that abstracts the properties of a transaction ledger capable of handling stateful smart contracts. We call this \mathcal{F}_{Ledg} , which we describe in Fig. 1.

Informally, a blockchain is a public ledger consisting of a sequence of blocks. Each block typically contains information about a set of transactions, and a new block is appended regularly via a consensus mechanism that involves the nodes of a network. A blockchain is equipped with a native currency which is employed in transactions, and whose basic unit is referred to as a *coin*.

Each user in the network is associated with a public key (this can be thought of as the user’s address). A typical transaction is a message which transfers coins from a public key to another. It is considered valid if it is digitally signed using the secret key corresponding to the sending address.

More precisely, in Bitcoin, parties do not keep track of users’s accounts, but rather they just maintain a local copy of a set known as “unspent transaction outputs set” (UTXO set). An unspent output is a transaction that has not yet been “claimed”, i.e. the coins of these transactions have not yet been spent by the receiver. Each unspent output in the UTXO set includes a circuit (also known as a “script”) such that any user that can provide an input which is accepted by the circuit (i.e. a witness) can make a transaction that spends these coins, thus creating a new unspent output. Hence, if only one user knows the witness to the circuit, he is effectively the owner of these coins. For a standard payment transaction, the witness is a signature and the circuit verifies the signature. However, more complex circuits are also allowed, and these give rise to more complex transactions than simple payments: smart contracts. A smart contract can be thought of as a transaction which deposits coins to an address. The coins are released upon fulfillment of certain pre-established conditions.

In [BK14] and [BKM17], smart contracts are defined as ideal functionalities in a variant of the Universal Composability (UC) model [Can01]. The ideal functionality that abstracts the simplest smart contracts was formalized in [BK14], and called “Claim or Refund”. Informally, this functionality specifies that a sender P locks his coins and chooses a circuit ϕ , such that a receiver Q can gain possession of these coins by providing a witness w such that $\phi(w) = 1$ before an established time, and otherwise the sender can reclaim his coins. The “Claim or Refund” ideal functionality can be realized in Bitcoin as long as the circuit ϕ can be described in Bitcoin’s scripting language. On the other hand, Ethereum’s scripting language is Turing-complete, and so any circuit ϕ can be described.

“Claim or refund” ideal functionalities can be further generalized to “stateful contracts”. In Ethereum, each unspent output also maintains a *state*. In other words, each unspent output comprises not only a circuit ϕ , but also state variables. Parties can claim partial amounts of coins by providing witnesses that satisfy the circuit ϕ in accordance with the current state variables.

In addition, ϕ also specifies an update rule for the state variables, which are updated accordingly. We refer to these type of transactions as stateful contracts, as opposed to the “stateless” contract of “Claim or Refund”. Stateful contracts can be realized in Ethereum, but not in Bitcoin. From now onwards, we will only work with stateful contracts. We will use the terms “smart contracts” and “stateful contracts” interchangeably.

We emphasize that our modeling is inspired by [BK14] and [BKM17], but differs in the way that coins are formalized. One difference from the model of Bentov et al. is that there, in order to handle coins, the authors augment the original UC model by endowing each party with a *wallet* and a *safe*, and by considering *coins* as atomic entities which can be exchanged between parties. To the best of our knowledge, this variant of the UC framework is not subsumed by any of the previously studied variants, and thus it is not known whether a composition theorem holds for it.

On the other hand, we feel that a more natural approach is to work in the Generalized UC model [CDPW07], and to define a global ideal functionality \mathcal{F}_{Ledg} which abstracts the essential features of a *transaction ledger* capable of handling stateful smart contracts. This approach was first proposed by Kiayias et al. [KZZ16]. The appeal of modeling a transaction ledger as a global ideal functionality is that composition theorems are known in the Generalized UC framework. In virtue of this, any secure protocol for some task that makes calls to \mathcal{F}_{Ledg} can be arbitrarily composed while still maintaining security. This means that one need not worry about composing different concurrent protocols which reference the same transaction ledger. One would hope that it is also the case that a secure protocol that makes calls to \mathcal{F}_{Ledg} remains secure when the latter are replaced by calls to secure real-world realizations of it (on Ethereum for example). This requires a more careful analysis, and Canetti et al. provide in [CSV16] sufficient conditions for this replacement to be possible. We do not prove that a secure real-world realization of \mathcal{F}_{Ledg} on an existing blockchain exists, but we believe that \mathcal{F}_{Ledg} , or a close variant of it, should be securely realizable on the Ethereum blockchain. In any case, we work abstractly by designing our payment system, and proving it secure, assuming access to such an ideal functionality. The appeal of such an approach is that the security of the higher-level protocols is independent of the details of the particular real-world implementation of \mathcal{F}_{Ledg} .

Next, we describe our global ideal functionality \mathcal{F}_{Ledg} . In doing so, we establish the notation that we will utilize in the rest of the paper.

Global ledger ideal functionality We present in Fig. 1 our global ledger ideal functionality \mathcal{F}_{Ledg} . In a nutshell, this keeps track of every registered party’s coins, and allows any party to transfer coins in their name to any other party. It also allows any party to retrieve information about the number of coins of any other party, as well as about any previous transaction. The initial amount of coins of a newly registered party is determined by its PID (recall that in a UC-execution the PID of each invoked party is specified by the environment; see subsection 2.3 for more details). Moreover, \mathcal{F}_{Ledg} handles (stateful) smart contracts: it accepts deposits from the parties involved in a contract and then pays rewards appropriately. Recall that in stateful smart contracts a party or a set of parties deposit an amount of coins to the contract. The contract is specified by a circuit ϕ , together with an initial value for a state variable \mathbf{st} . A state transition is triggered by any party P with PID pid sending a witness w which is accepted by ϕ in accordance with the current state and the current time t . More precisely, the contract runs $\phi(pid, w, t, \mathbf{st})$, which outputs either “ \perp ” or a new state (stored in the variable \mathbf{st}) and a number of coins $d \in \mathbb{N}$ that is released to P . Each contract then repeatedly accepts state transitions until it has distributed all the coins that were deposited into it at the start. Notice that ϕ can accept different witnesses at different times (the acceptance of a witness can depend on the current time t and the current value of the state variable \mathbf{st}). Information about existing smart contracts can also be retrieved by any party.

The stateful-contract portion of \mathcal{F}_{Ledg} resembles closely the functionality \mathcal{F}_{StCon} from [BKM17]. Our approach differs from that of [BKM17] in that we make the smart contract functionality part of the global ideal functionality \mathcal{F}_{Ledg} which also keeps track of party’s coins and transactions. In [BKM17] instead, coins are incorporated in the computation model by augmenting the ITMs with wallets and safes (this changes the model of computation in a way that is not captured by any of the the previously studied variants of the UC framework).

We implicitly assume access to an ideal functionality for message authentication \mathcal{F}_{auth} which

all parties employ when sending their messages, and also to a global ideal functionality for a clock that keeps track of time. We assume implicitly that \mathcal{F}_{Ledg} makes calls to the clock and keeps track of time. Alternatively, we could just have F_{Ledg} maintain a local variable that counts the number of transactions performed, and a local time variable t , which is increased by 1 every time the number of transactions reaches a certain number, after which the transaction counter is reset (this mimics the process of addition of blocks in a blockchain, and time simply counts the number of blocks). From now onwards, we do not formally reference calls to \mathcal{F}_{auth} or to the clock to avoid overloading notation. We are now ready to define \mathcal{F}_{Ledg} .

Global ledger ideal functionality

Initialize the sets $\text{parties} = \{\}$, $\text{contracts} = \{\}$, $\text{AllTransactions} = \{\}$. (Throughout, the variable t denotes the current time.)

Register: Upon receiving a message `Register` from a party with PID $\text{pid} = (id, d)$, send `(Registered, pid)` to the adversary; upon receiving a message `ok` from the adversary, and if this is the first request from pid , add pid to the set parties . Set $\text{pid.id} \leftarrow id$ and $\text{pid.coins} \leftarrow d$.

Retrieve party: Upon receiving a message `(RetrieveParty, pid)` from some party P (or the adversary), output `(RetrieveParty, pid, d)` to P (or to the adversary), where $d = \perp$ if $\text{pid} \notin \text{parties}$, and $d = \text{pid.coins}$ otherwise. (We slightly abuse notation here in that, when taken as part of a message, pid is treated as a string, but, when called by the functionality, pid is a variable with attributes pid.id and pid.coins).

Pay: Upon receiving a message `(pay, pid', d)` from some party P with PID pid , and $\text{pid} \in \text{parties}$, do the following:

- If $\text{pid}' \in \text{parties}$ and $\text{pid.coins} > d$, update $\text{pid}'.\text{coins} \leftarrow \text{pid}'.\text{coins} + d$ and $\text{pid.coins} \leftarrow \text{pid.coins} - d$. Set $\text{trId} = |\text{AllTransactions}| + 1$. Add a variable named trId to AllTransactions , with attribute $\text{trId.transaction} = (\text{pid}, \text{pid}', d, t)$. Send a message `(Executed, trId)` to P .
- Else, return \perp to P .

Retrieve transaction: Upon receiving a message `(RetrieveTransaction, trId)` from some party P (or the adversary), output `(RetrieveTransaction, trId, s)`, where $s = \perp$ if $\text{trId} \notin \text{allTransactions}$, and $s = \text{trId.transaction}$ otherwise.

Smart contract Upon receiving a message `(InitiateSmartContract, Params=(I, D, ϕ, st_0))`, where I is a set of PID's, D is a set $\{(pid, d_{pid}) : pid \in I\}$ of “initial deposits”, with d_{pid} being the amount required initially from the party with PID pid , ϕ is a circuit, and st_0 is the initial value of a state variable st , check that $I \subseteq \text{parties}$. If not, ignore the message; if yes, set $\text{ssid} = |\text{contracts}| + 1$. Add a variable named ssid to contracts with attributes $\text{ssid.Params} = (I, D, \phi, \text{st}_0)$, $\text{ssid.state} = \text{st}$ and $\text{ssid.coins} \leftarrow 0$. Send a message `(RecordedContract, ssid)` to P . Then, do the following:

- **Initialization phase:** Wait to get message `(InitializeWithCoins, ssid, Params = (I, D, ϕ, st_0))` from party with PID pid for all $\text{pid} \in I$. When all messages are received, and if, for all $\text{pid} \in I$, $\text{pid.coins} \geq d_{pid}$, then, for all $\text{pid} \in I$, update: $\text{pid.coins} \leftarrow \text{pid.coins} - d_{pid}$ and $\text{ssid.coins} \leftarrow \text{ssid.coins} + d_{pid}$. Set $\text{st} \leftarrow \text{st}_0$ (We assume that ssid.state changes dynamically with st).
- **Execution phase:** Repeat until termination: Upon receiving a message of the form `(Trigger, ssid, w, d)` at time t from some party with PID $\text{pid} \in \text{parties}$ (where it can also be $d = 0$) such that $\phi(\text{pid}, w, t, \text{st}, d) \neq \perp$, do the following:
 - If $d > 0$, update $\text{pid.coins} \leftarrow \text{pid.coins} - d$ and $\text{ssid.coins} \leftarrow \text{ssid.coins} + d$
 - Update $(\text{st}, e) \leftarrow \phi(\text{pid}, w, t, \text{st}, d)$.
 - If $e = \text{“all coins”}$, let $q := \text{ssid.coins}$. Send the message `(Reward, ssid, q)` to the party with PID pid and update $\text{pid.coins} \leftarrow \text{pid.coins} + q$ and $\text{ssid.coins} \leftarrow 0$. If $e > 0$ and $\text{ssid.coins} \geq e$, send the message `(Reward, ssid, e)` to the party with PID pid , and update $\text{pid.coins} \leftarrow \text{pid.coins} + e$ and $\text{ssid.coins} \leftarrow \text{ssid.coins} - e$. Else, if $\text{ssid.coins} = e' < e$, send the message `(Reward, ssid, e')` to the party with PID pid , and update $\text{pid.coins} \leftarrow \text{pid.coins} + e'$ and $\text{ssid.coins} \leftarrow 0$. Then, terminate.

Retrieve smart contract: Upon receiving a message `(RetrieveContract, ssid)` from some party P (or the adversary), output `(RetrieveContract, ssid, z)`, where $z = \perp$ if $\text{ssid} \notin \text{contracts}$, and $z = (\text{ssid.Params}, \text{ssid.state}, \text{ssid.coins})$ otherwise.

From now on, we will often refer to the number of coins $ssid.coins$ of a contract with session identifier $ssid$ as the coins *deposited* in the contract. When we say that a contract *releases* some coins to a party P with PID pid , we mean more precisely that \mathcal{F}_{Ledg} updates its local variables and moves coins from $ssid.coins$ to $pid.coins$.

4 A payment system based on quantum lightning and a classical blockchain

In this section, we describe our payment system. We give first an informal description, and in section 4.1 we give a formal description.

The building block of our payment system is a quantum lightning scheme, reviewed in detail in section 2.2. Recall that a quantum lightning scheme consists of a generation procedure which creates quantum banknotes, and a verification procedure that verifies them and assigns serial numbers. The security guarantee is that no generation procedure (not even the honest one) can create two banknotes with the same serial number except with negligible probability. As mentioned earlier, this property is desirable if one wants to design a decentralized payment system, as it prevents anyone from cloning banknotes (even the person who generates them). However, this calls for a mechanism to regulate generation of new valid quantum banknotes.

In this section, we describe formally a proposal that employs smart contracts to provide such a mechanism. As we have described informally in the introduction, the high-level idea is to keep track of the valid serial numbers using smart contracts. Any party is allowed to deposit any amount of coins d (of their choice) into a smart contract with specific parameters (see definition 5 below), and with an initial value of his choice for a *serial number* state variable. We can think of the quantum banknote with the chosen serial number as having “acquired” value d . A payment involves only two parties: a payer, who sends a quantum banknote, and a payee who receives it and verifies it locally. As anticipated in the introduction, the full payment system includes the following additional features, which we describe here informally in a little more detail (all of these are described formally in subsection 4.1):

- Removing a serial number from the list of valid serial numbers in order to recover the amount of coins deposited in the corresponding smart contract. This makes the two commodities (quantum banknotes and coins on the blockchain) interchangeable. This is achieved by exploiting the additional property of some quantum lightning scheme from Definition 3. Recall that, informally, this property states that there is some classical certificate that can be recovered by measuring a valid quantum banknote, which no efficient algorithm can recover otherwise. The key is that once the state is measured to recover this certificate, it is damaged in a way that it only passes verification with negligible probability (meaning that it can no longer be spent). We allow users to submit this classical certificate to a smart contract, and if the certificate is consistent with the serial number stored in the contract, then the latter releases all of the coins deposited in the contract to the user.
- Allowing a party to replace an existing serial number with a new one of their choice in case they lose a valid quantum state (they are fragile after all!). We allow a user P to file a “lost banknote claim” by sending a message and some fixed amount of coins d_0 to a smart contract whose serial number is the serial number of the lost banknote. The idea is that if no one challenges this claim, then after a specified time t_{tr} user P can submit a message which changes the value of the serial number state variable to a value of his choice and recovers the previously deposited d_0 coins. On the other hand, if a user P maliciously files a claim to some contract with serial number s , then any user Q who possesses the valid banknote with serial number s can recover the classical certificate from Definition 3, and submit it to the contract. This releases all the coins deposited in the contract to Q (including the d_0 deposited by P to make the claim). As you might notice, this requires honest users to monitor existing contracts for “lost banknote claims”. This, however, is not much of a burden if t_{tr} is made large enough (say a week or a month). The requirement of being online once a week or once a month is easy to meet in practice.

4.1 The payment system and its components

In this section, we describe in detail all of the components of the payment system. It consists of the following: a protocol to generate valid quantum banknotes; a protocol to make a payment; a protocol to file a claim for a lost banknote; a protocol to prevent malicious attempts at filing claims for lost banknotes; and a protocol to trade a valid quantum banknote in exchange for coins.

Recall from the description of \mathcal{F}_{Ledg} that each smart contract is specified by several parameters: I is a set of PIDs of parties who are expected to make the initial deposit, with $\{d_{pid} : pid \in I\}$ being the required initial deposit amounts; a circuit ϕ specifies how the state variables are updated and when coins are released; an initial value st_0 for the state variable st ; a session identifier $ssid$.

Let $\lambda \in \mathbb{N}$. From now onwards, assume that $(\text{Gen}, \text{Ver}, H) \leftarrow \text{QL.Setup}(\lambda)$, where the latter is the setup procedure of a quantum lightning scheme with bolt-to-certificate capability (i.e. a quantum lightning scheme that satisfies the additional property of Definition 3). Let \mathcal{A}_* be as in Definition 3, namely $\text{Game}_1(\mathcal{A}_*, \lambda) = 1$ with all but negligible probability. In Definition 5, we define an instantiation of smart contracts with a particular choice of parameters, which we refer to as banknote-contracts. Banknote-contracts are the building blocks of the protocols that make up our payment system. We describe a banknote-contract informally before giving a formal definition.

A banknote-contract is a smart contract initialized by a single party, and it has a state variable of the form $st = (\text{serial}, \text{ActiveLostClaim})$. The party initializes the banknote-contract by depositing a number of coins d and by setting the initial value of serial to any desired value. The banknote-contract handles the following type of requests:

- As long as $\text{ActiveLostClaim} = \text{“No active claim”}$ (which signifies that there are no currently active lost-banknote claims), any party P can send the message BanknoteLost , together with a pre-established amount of coins d_0 to the contract. This will trigger an update of the state variable ActiveLostClaim to reflect the active lost-banknote claim by party P .
- As long as there is an active lost-banknote claim, i.e. $\text{ActiveLostClaim} = \text{“Claim by } pid \text{ at time } t\text{”}$, any party Q can challenge that claim by submitting a message $(\text{ChallengeClaim}, x, s')$ to the contract, where s' is a proposed new serial number. We say that x is a valid classical certificate for the current value s of serial if $H(x) = s$. Such an x can be thought of as a proof that whoever is challenging the claim actually possessed a quantum banknote with serial number s , and destroyed it in order to obtain the certificate x , and thus that the current active lost-banknote claim is malicious. If x is a valid classical certificate for s , then serial is updated to the new value s' submitted by Q , who also receives all of the coins deposited in the contract (including the d_0 coins deposited by the malicious claim).
- If party P has previously submitted a lost-banknote claim, and his claim stays unchallenged for time t_{tr} , then party P can send a message $(\text{ClaimUnchallenged}, s')$ to the contract, where s' is a proposed new serial number. Then the contract returns to P the d_0 coins he initially deposited when making the claim, and updates serial to s' .
- Any party P can submit to the contract a message $(\text{RecoverCoins}, x)$. If x is a valid classical certificate for the current value of s of serial , then the contract releases to P all the coins currently deposited in the contract. This allows party P to “convert” back his quantum banknote into coins.

Next, we will formally define banknote-contracts, and then formally describe all of the protocols that make the payment system.

ϕ_{\S} ($pid, w, t, (\text{serial}, \text{ActiveLostClaim}), d$) takes as input strings pid and w , where pid is meant to be the PID of some party P , and we refer to w as the “witness”, $t \in \mathbb{N}$ denotes the “current time” maintained by \mathcal{F}_{Ledg} , $(\text{serial}, \text{ActiveLostClaim})$ is the current value of the state variable, and $d \in \mathbb{N}$ is the number of coins that are being deposited to the smart contract with the current message. ϕ_{\S} has hardcoded parameters: $d_0 \in \mathbb{N}$ the amount of coins needed to file a claim for a lost money state, $t_{tr} \in \mathbb{N}$ the time after which an unchallenged claim can be settled (d_0 and t_{tr} are fixed constants agreed upon by all parties, and they are the same for all banknote-contracts). The circuit ϕ_{\S} outputs new values for the state variables and an amount of coins as follows:

On input $(pid, w, t, (\text{serial} = s, \text{ActiveLostClaim}), d)$, ϕ_{\S} does the following:

- If $\text{ActiveLostClaim} = \text{“No active claim”}$:
 - If $w = \text{BanknoteLost}$ and $d = d_0$, then ϕ outputs $((\text{serial} = s, \text{ActiveLostClaim} = \text{“Claim by } pid \text{ at time } t”), 0)$ (to symbolize that at time t party with PID pid has claimed to have lost the money state with serial number s , and that zero coins are being released).
 - If $w = (\text{RecoverCoins}, x)$, where $x \in \{0, 1\}^\lambda$ and $H(x) = s$, then ϕ outputs $((\text{serial} = \perp, \text{ActiveLostClaim} = \perp), \text{“all coins”})$
 - If $\text{ActiveLostClaim} = \text{“Claim by } pid' \text{ at time } t_0”$ for some pid', t_0 :
 - If $w = (\text{ChallengeClaim}, x, s')$, where $x \in \{0, 1\}^\lambda$ and $H(x) = s$, and $s' \in \{0, 1\}^\lambda$ then ϕ outputs $((\text{serial} = s', \text{ActiveLostClaim} = \text{“No active claim”}), d_0)$.
 - If $w = (\text{ClaimUnchallenged}, s')$, $pid = pid'$ and $t - t_0 > t_{tr}$, then ϕ outputs $((\text{serial} = s', \text{ActiveLostClaim} = \text{“No active claim”}), d_0)$.
-

Figure 2: Circuit ϕ_{\S} for banknote-contracts

Definition 5. (*Banknote-contract*) A banknote-contract, is a smart contract on \mathcal{F}_{Ledg} specified by parameters of the following form: $I = \{pid\}$ for some $pid \in [n]$, $D = \{(pid, d_{pid})\}$ for some $d_{pid} \in \mathbb{N}$, $\text{st}_0 = (s, \text{“No active claim”})$ for some $s \in \{0, 1\}^\lambda$, and circuit $\phi = \phi_{\S}$, where ϕ_{\S} is defined as in Fig. 2.

For convenience, we denote by **serial** and **ActiveLostClaim** respectively the first and second entry of the state variable of a banknote-contract.

Generating valid quantum banknotes We describe the formal procedure for generating a valid quantum banknote.

Protocol carried out by some party P with PID pid .

Input of P : An integer d such that $pid.\text{coins} > d$ in \mathcal{F}_{Ledg} (d is the “value” of the prospective banknote).

- Run $|\psi\rangle \leftarrow \text{Gen}$. Let $s := \text{Ver}(|\psi\rangle)$.
 - Send $(\text{InitiateSmartContract}, \text{Params})$ to \mathcal{F}_{Ledg} , where $\text{Params} = (\{pid\}, \{(pid, d)\}, \phi, (s, \text{“No active claim”}))$. Upon receipt of a message of the form $(\text{RecordedContract}, ssid)$, send the message $(\text{InitializeWithCoins}, ssid, \text{Params})$ to \mathcal{F}_{Ledg} .
-

Figure 3: Generating a valid banknote

Making a payment We describe formally the protocol for making a payment in Fig. 4. Informally, the protocol is between a party P , the payer, and a party Q , the payee. In order to pay party Q with a bolt whose serial number is s , party P sends the valid bolt to party Q , the payee, together with the $ssid$ of a smart contract with $\text{serial} = s$. Party Q verifies that

$ssid$ corresponds to a banknote-contract with $serial = s$, and verifies that the banknote passes verification and has serial number s .

The protocol is between some party P with PID pid (the payer) and a party Q with PID pid' (the payee):

Input of P : $|\Psi\rangle$, a valid bolt with serial number s . $ssid$ the session identifier of a smart contract on \mathcal{F}_{Ledg} such that $ssid.state = (s, \text{“No active claim”})$, and $ssid.coins = d$.

- P sends state $|\Psi\rangle$ to Q . P also sends a message $(ssid, s, d)$ to Q .
 - Q sends a message $(\text{RetrieveContract}, ssid)$ to \mathcal{F}_{Ledg} . Upon receiving a message $(\text{RetrieveContract}, ssid, z)$ from \mathcal{F}_{Ledg} (where $z = (ssid.Params, ssid.state, ssid.coins)$ if P is honest), Q does the following:
 - If $z = (Params, (s, \text{“No active claim”}), d)$, then Q checks that the parameters $Params$ are of the form of a banknote-contract (from Definition 5). If so, runs $\text{Ver}(|\Psi\rangle)$ and checks that the outcome is s . If so, sends a message accept to P .
 - Else, Q aborts.
-

Figure 4: Protocol for making and verifying a payment

Recovering lost banknotes As much as we can hope for experimental progress in the development of quantum memories, for the foreseeable future we can expect quantum memories to only be able to store states for a time on the order of days. It is thus important that any payment system involving quantum money is equipped with a procedure for users to recover the value associated to quantum states that get damaged and become unusable. Either users should be able to “convert” quantum money states back to coins on the blockchain, or they should be able, upon losing a quantum banknote, to change the serial number state variable of the associated smart contract to a new serial number (presumably of freshly generated quantum banknote). Here, we describe a protocol for the latter. Later, we describe a protocol for the former.

Informally, a party P who has lost a quantum banknote with serial number s associated to a smart contract with session identifier $ssid$, makes a “lost banknote claim” at time t by depositing a number of coins d_0 to that banknote-contract. Recall the definition of banknote-contracts from Definition 5, and in particular of the circuit $\phi_{\mathcal{S}}$:

- If party P is honest, then after a time t_{tr} has elapsed, he will be able to update the state variable $serial$ of the banknote-contract from s to s' (where s' is presumably the serial number of a new valid bolt that party P has just generated).
- If party P is dishonest, and he is claiming to have lost a banknote that someone else possesses, then the legitimate owner can apply \mathcal{A}_* as in Definition 3 to the legitimate bolt $|\Psi\rangle$ and recover x such that $H(x) = s$. He can then send x to the contract and a new serial number s' (presumably of a freshly generate bolt) and obtain d_0 coins from the contract (the d_0 coins deposited by P in his malicious claim).

We describe the protocol formally in Fig. 5.

One might wonder whether, in practice, an adversary can instruct a corrupt party to make a “lost banknote claim”, and then intercept an honest party’s classical certificate x before this is posted to the blockchain, and have a second corrupt party post it instead. This attack would allow the adversary to “steal” the honest party’s value. In our analysis, we do not worry about this, as we assume access to the ideal functionality \mathcal{F}_{Ledg} , which, by definition, deals with incoming messages in the order that they are received. We also assume in our adversarial model, specified more precisely in Section 5, that the adversary does not have any control over the delivery of messages (and their timing). If one assumes a more powerful adversary (with some control over the timing of delivery of messages), then this kind of issue can still be mitigated, for example, in the following manner: the honest party does not directly post the classical certificate x , but she

instead first posts a commitment to x , and she reveals it only at a latter stage. We thank Or Sattath for pointing out this issue.

Protocol carried out by party P with PID pid for changing the serial number of a smart contract.

P 's input: s the serial number of a (lost) quantum banknote. $ssid$ the session identifier of a banknote-contract such that $ssid.state = (s, \text{"No active claim"})$.

- P sends $(\text{Trigger}, ssid, \text{BanknoteLost}, d_0)$ to \mathcal{F}_{Ledg} . This updates $ssid.state$ to $(s, \text{"Active claim by } pid \text{ at time } t")$ (where t is the current time maintained by \mathcal{F}_{Ledg}), and deposits d_0 coins into the contract.
- After time t_{tr} , P sends $(\text{Trigger}, ssid, (\text{ClaimUnchallenged}, s'), 0)$ to \mathcal{F}_{Ledg} . If P was honest then $ssid.state$ is updated to $(s', \text{"No active claim"})$, and d_0 coins are released to P .

Figure 5: Protocol for changing the serial number of a smart contract

Next, we give a protocol carried out by all parties to prevent malicious attempts at changing the state variable `serial` of a smart contract. Informally, this involves checking the blockchain regularly for malicious attempts at filing lost-banknote claims.

Recall that t_{tr} was defined in Definition 5. Recall also the definitions of \mathcal{A}_* and H from Definition 3.

Protocol carried out by a party P to prevent malicious attempts at changing the state variable `serial` of a smart contract.

Input of P : A triple $(|\Psi\rangle, s, ssid)$, where $|\Psi\rangle$ is a quantum banknote with serial number s , and $ssid$ is the session identifier of a banknote-contract such that $ssid.state = (s, \text{"No active claim"})$

At regular intervals of time $t_r - 1$, do the following:

- Send a message $(\text{RetrieveContract}, ssid)$ to \mathcal{F}_{Ledg} . Upon receiving a message $(\text{RetrieveContract}, ssid, z)$ from \mathcal{F}_{Ledg} , if $z = (\text{Params}, (s, \text{"Claim by } pid' \text{ at time } t"), d)$ for some pid', t, d and for some banknote-contract parameters Params :
 - Apply \mathcal{A}_* to $|\Psi\rangle$ to obtain $x \in \{0, 1\}^\lambda$ (such that $H(x) = s$).
 - Sample $|\Psi\rangle \leftarrow \text{Gen}$. Let $s' := \text{Ver}(|\Psi\rangle)$.
 - Send $(\text{Trigger}, ssid, (\text{ChallengeClaim}, x, s'), 0)$ to \mathcal{F}_{Ledg} . (If P was honest, this updates $ssid.state \leftarrow (s', \text{"No active claim"})$ and releases d_0 coins to P).

Figure 6: Protocol for preventing malicious attempts at changing the state variable `serial` of a smart contract.

Trading a quantum banknote for coins: Finally, we describe a protocol for trading a quantum banknote to recover all the coins deposited in its associated banknote-contract.

Protocol carried out by a party P .

Input of P : A tuple $(|\Psi\rangle, s, ssid, d)$, where $|\Psi\rangle$ is a quantum banknote with serial number s , and $ssid$ is the session identifier of a banknote-contract such that $ssid.state = (z, \text{“No active claim”})$ and $ssid.coins = d$.

- Run \mathcal{A}_* on $|\Psi\rangle$ to get outcome $x \in \{0, 1\}^\lambda$ such that $H(x) = s$.
 - Send message $(\text{Trigger}, ssid, (\text{RecoverCoins}, x), 0)$ to \mathcal{F}_{Ledg} . This releases d coins to P .
-

Figure 7: Protocol for trading a quantum banknote for coins.

5 Security

We first specify an adversarial model, and then use this to prove a security theorem.

Adversarial model We assume that all the messages of honest parties are sent using the ideal functionality for authenticated communication \mathcal{F}_{auth} , and that the adversary sees all messages that are sent (in UC language, we assume that the adversary is activated every time a party sends a message) but has no control over the delivery of messages (whether they are delivered or not) and their timing. Our payment system can be made to work also if we assume that the adversary can delay delivery of honest parties’ messages by a fixed amount of time (see the remark preceding Fig. 5 for more details), but, for simplicity, we do not grant the adversary this power.

The adversary can corrupt any number of parties, and it may do so adaptively, meaning that the corrupted parties are not fixed at the start, but rather an honest party can become corrupted, or a corrupted party can return honest, at any point. The process of corruption is modeled analogously as in the original UC framework, where the adversary simply writes a *corrupt* message on the incoming tape of an honest party, upon which the honest party hands all of its information to the adversary, who can send messages on the corrupted party’s behalf. Our setting is slightly more involved in that corrupted parties also possess some quantum information, in particular the quantum banknotes. We assume that when an adversary corrupts a party he takes all of its quantum banknotes. Importantly, we assume that these are not returned to the party once the party is no longer corrupted. It might seem surprising that we do not upper bound the fraction of corrupted parties. Indeed, such a bound would only be needed in order to realize securely the ideal functionality \mathcal{F}_{Ledg} (any consensus-based realization of \mathcal{F}_{Ledg} would require such a bound). Here, we assume access to such an ideal functionality, and we do not worry about its secure realization. Naturally, when replacing the ideal functionalities with real-world realizations one would set the appropriate bound on the corruption power of the adversary, but we emphasize that our schemes are independent of the particular real-world realization. Note that we do not fix a set of parties at the start, but rather new parties can be created (see below for more details).

We assume that (ITMs of) honest parties run the code π . This represents the “honest” code which executes the protocols from section 4 as specified. The input to π then specifies when and which protocols from section 4 are to be executed. As part of π , we specify that, upon invocation, a party sends a message `Register` to \mathcal{F}_{Ledg} to register itself. We also specify as part of π that an honest party runs the protocol of Fig. 6 (to prevent malicious claims for lost banknotes). Moreover, for notational convenience, we specify as part of π that each party maintains a local variable `banknoteValue`, which keeps track of the total value of the quantum banknotes possessed by the party. `banknoteValue` is initialized to 0, and updated as follows. Whenever a party P successfully receives a quantum banknote (i.e. P is the payee in the protocol from Fig. 4 and does not abort) of value d (i.e. the associated smart contract has d coins deposited), then P updates `banknoteValue` \leftarrow `banknoteValue` + d . Similarly, when P sends a quantum banknote of value d , it updates `banknoteValue` \leftarrow `banknoteValue` - d . Finally, we specify also as part of π , that whenever a party that was corrupted is no longer corrupted, it resets `banknoteValue` = 0 (this is

because we assumed that quantum banknotes are not returned by the adversary). The following paragraph leads up to a notion of security and a security theorem.

Let \mathcal{A} be a quantum polynomial-time adversary and \mathcal{E} a quantum polynomial-time environment. Consider an execution of π with adversary \mathcal{A} and environment \mathcal{E} (see section 2.3 for more details on what an “execution” is precisely). We keep track of two quantities during the execution, which we denote as `AdversaryValueReceived` and `AdversaryValueCurrentOrSpent` (These quantities are not computed by any of the parties, adversary or environment. Rather, they are just introduced for the purpose of defining security). The former represents the amount of value, coins or banknotes, that the adversary has received either by virtue of having corrupted a party, or by having received a payment from an honest party. The latter counts the total number of coins currently possessed by corrupted parties, as recorded on \mathcal{F}_{Ledg} , and the total amount spent by the adversary to honest parties either via coins or via quantum banknotes (it does not count the value of quantum banknotes currently possessed; these only count once they are successfully spent). Both quantities are initialized to 0, and updated as follows throughout the execution:

- (i) When \mathcal{A} corrupts a party P : let d be the number of coins of P according to the global functionality \mathcal{F}_{Ledg} and d' be P 's `banknoteValue` just before being corrupted. Then, `AdversaryValueReceived` \leftarrow `AdversaryValueReceived` + $d + d'$, and `AdversaryValueCurrentOrSpent` \leftarrow `AdversaryValueCurrentOrSpent` + d .
- (ii) When a corrupted party P with d coins and `banknoteValue` = d' ceases to be corrupted and returns honest, `AdversaryValueReceived` \leftarrow `AdversaryValueReceived` - d .
- (iii) When an honest party pays d coins to a corrupted party, `AdversaryValueReceived` \leftarrow `AdversaryValueReceived` + d . Likewise, when an honest party sends a quantum banknote of value d to a corrupted party, through the protocol of Fig. 4, then (even if the corrupted party does not return accept) `AdversaryValueReceived` \leftarrow `AdversaryValueReceived` + d .
- (iv) When \mathcal{A} successfully spends a quantum banknote of value d to an honest party P , i.e. a corrupted party is the payer in the protocol from Fig. 4 and P is the payee and returns accept, or when \mathcal{A} pays d coins to an honest party, then `AdversaryValueCurrentOrSpent` \leftarrow `AdversaryValueCurrentOrSpent` + d .
- (v) When a corrupted party receives d coins from a banknote-contract, then `AdversaryValueCurrentOrSpent` \leftarrow `AdversaryValueCurrentOrSpent` + d . Notice that this can happen only in two ways: \mathcal{A} successfully converts a quantum banknote of value d to coins on \mathcal{F}_{Ledg} (via the protocol of Fig. 7), or a corrupted party successfully challenges a `BanknoteLost` claim (in this case $d = d_0$).

Intuitively, if our payment scheme is secure, then at no point in time should the adversary be able to make `AdversaryValueCurrentOrSpent` - `AdversaryValueReceived` > 0 . This would mean that he has successfully spent/stolen value other than the one he received by virtue of corrupting a party or receiving honest payments. The following theorem formally captures this notion of security. First, we denote by $\mathcal{F}_{Ledg}\text{-EXEC}_{\pi, \mathcal{A}, \mathcal{E}}^{(MaxNetValue)}(\lambda, z)$ the maximum value of `AdversaryValueCurrentOrSpent` - `AdversaryValueReceived` during an execution of π with adversary \mathcal{A} and environment \mathcal{E} , with global shared functionality \mathcal{F}_{Ledg} .

Theorem 2 (Security). *For any quantum polynomial-time adversary \mathcal{A} and quantum polynomial-time environment \mathcal{E} ,*

$$\Pr[\mathcal{F}_{Ledg}\text{-EXEC}_{\pi, \mathcal{A}, \mathcal{E}}^{(MaxNetValue)}(\lambda, z) > 0] = \text{negl}(\lambda).$$

The rationale behind considering executions of π and quantifying over all possible adversaries and environments is that doing so captures all possible ways in which a (dynamically changing) system of honest parties running our payment system alongside an adversary can behave (where the adversary respects our adversarial model).

Recall that, in an execution of π , the environment has the ability to invoke new parties and assign to them new unique PIDs. Since in \mathcal{F}_{Ledg} the PIDs are used to register parties and initialize their number of coins, this means that the environment has the ability to pick the initial number of coins of any new party that it invokes. Moreover, by writing inputs to the parties input tapes, the environment can instruct honest parties to perform the honest protocols from section

4 in any order it likes. Quantifying over all adversaries and environments, in the statement of Theorem 2 means that the adversary and the environment can intuitively be thought of as one single adversary. The statement of the theorem thus captures security against realistic scenarios in which new parties can be adversarially created with an adversarially chosen number of coins, and they can be instructed to perform the honest protocols of the payment system from section 4, in whatever sequence is convenient to the adversary.

Proof of Theorem 2. The proof is straightforward, and we avoid being overly formal.

Suppose for a contradiction that there exists \mathcal{A} and \mathcal{E} such that

$$\Pr[\mathcal{F}_{Ledg}\text{-EXEC}_{\pi,\mathcal{A},\mathcal{E}}^{(MaxNetValue)}(\lambda, z) > 0] \neq \text{negl}(\lambda). \quad (1)$$

Then, we go through all of the possible ways that an adversary can increase its net value, i.e. increase the quantity `AdversaryValueCurrentOrSpent` – `AdversaryValueReceived`: the adversary can do so through actions from items (ii), (iv) and (v) above. Amongst these, it is easy to see that action (ii) never results in `AdversaryValueCurrentOrSpent` – `AdversaryValueReceived` > 0 . Thus, in order for (1) to hold, it must be the case that one of the following happens with non-negligible probability within an execution of π with adversary \mathcal{A} and environment \mathcal{E} .

- An action from item (iv) resulted in a positive net value for \mathcal{A} , i.e. `AdversaryValueCurrentOrSpent` – `AdversaryValueReceived` > 0 . Notice that for this to happen it must be the case that \mathcal{A} has double-spent a banknote, i.e. \mathcal{A} has produced two banknotes with the same serial number that have both been accepted by honest parties in a payment protocol of Fig. 4, and so they have both passed verification. But then, it is straightforward to see that we can use this adversary, together with \mathcal{E} to construct an adversary \mathcal{A}' that breaks the security of the quantum lightning scheme: \mathcal{A}' simply simulates an execution of protocol π with adversary \mathcal{A} and environment \mathcal{E} , and with non-negligible probability the adversary \mathcal{A} in this execution produces two banknotes with the same serial number. \mathcal{A}' uses these banknotes to win the security game of quantum lightning.
- An action from item (v) resulted in a positive net value for \mathcal{A} . Then, notice that for this to happen it must be that either:
 - \mathcal{A} has sent a message `(Trigger, ssid, (RecoverCoins, x), 0)` to \mathcal{F}_{Ledg} for some `ssid` and `x` such that $H(x) = s$, where `ssid.state` = `(s, "No active claim")`, and the last “make a payment” protocol (from Fig. 4) referencing `ssid` had an honest party as payee which remained honest at least up until after \mathcal{A} sent his message (or the banknote-contract was initialized by an honest user and the banknote was never spent). But then, one of the following must have happened:
 - * \mathcal{A} possessed a bolt $|\Psi\rangle$ with serial number `s` at some point, before $|\Psi\rangle$ was spent to the honest user. Then, this adversary would have recovered a good `x` and also spent a bolt with serial number `s` successfully to an honest user. But such an \mathcal{A} , together with \mathcal{E} , can be used to win `Game12` from Definition 3 with non-negligible probability, with a similar reduction to the one above, thus violating the property of Definition 3.
 - * \mathcal{A} recovered `x` such that $H(x) = s$ without ever possessing a valid bolt with serial number `s`. Again, such an adversary could be used, together with \mathcal{E} to win `Game12` from Definition 3).
 - * \mathcal{A} has successfully changed the serial number of contract `ssid` to `s` from some previous `s'` without possessing a bolt $|\Psi\rangle$ with serial number `s'`. This cannot happen since any honest user who possesses the valid bolt with serial number `s'` performs the protocol of Fig. 6.
 - \mathcal{A} has sent a message `(Trigger, ssid, (ChallengeClaim, x), 0)` to \mathcal{F}_{Ledg} for some `ssid` such `ssid.state` = `(s, "Claim by pid at time t")` for some `s, pid, t` with `pid` honest and `x` such that $H(x) = s$. Since `pid` is honest, he must be the last to have possessed a valid bolt with serial number `s`. Then, there are two possibilities:
 - * \mathcal{A} never possessed a valid bolt with serial number `s`, and succeeded in recovering `x` such that $H(x) = s$. Analogously to earlier, this adversary, together with \mathcal{E} , can be used to win `Game12`.

- * \mathcal{A} possessed a bolt $|\Psi\rangle$ with serial number s at some point, before $|\Psi\rangle$ was spent to an honest user. Analogously to earlier, this means such an \mathcal{A} both recovered an x with $H(x) = s$ and spent a bolt with serial number s successfully. Such an \mathcal{A} can be used, together with \mathcal{E} , to win Game_{12} .

□

6 Conclusion

In this work, we gave the first example of the use of classical smart contracts in conjunction with quantum cryptographic tools. Our contribution was two-fold: first, we introduced a global ideal functionality that abstracts the essential features of stateful smart contracts on a transaction ledger, and fits into the well-studied framework of Generalized Universal Composability of Canetti et al. [CDPW07]. Thus, we hope that it will find application within classical protocols that employ stateful smart contracts. We did not provide a real-world realization of this functionality, but we believe that such a realization should be possible in Ethereum. We leave proving such a secure realization as future work. Second, we showed that smart contracts can be combined with quantum tools, in particular quantum lightning, to design a decentralized payment system which solves the problem of scalability of (payment) transactions. There is currently only one known secure construction of quantum lightning, which relies on a computational assumption about multi-collision resistance of certain degree-2 hash functions [Zha17]. Finding alternative constructions of quantum lightning, secure under more well-studied computational assumptions, is a very interesting open problem.

Smart contracts have found several applications in classical cryptographic tasks, but their application to quantum cryptographic tasks is virtually unexplored. We hope that this work will ignite future investigations. Some candidate tasks which might potentially benefit from smart contracts are: generation of public trusted randomness, distributed delegation of quantum computation, secure multi-party quantum computation.

Acknowledgements

The author thanks Andru Gheorghiu, Abel Molina, Mario Larangeira, Or Sattath and Thomas Vidick for valuable comments and discussions on earlier versions of this work. The author also thanks Ran Canetti, Ranjit Kumaresan and Dominique Unruh for helpful email exchanges. The author is supported by the Kortschak Scholars program and AFOSR YIP award number FA9550-16-1-0495.

References

- [Aar09] Scott Aaronson. Quantum copy-protection and quantum money. In *Computational Complexity, 2009. CCC'09. 24th Annual IEEE Conference on*, pages 229–242. IEEE, 2009.
- [AGP16] Pablo Daniel Azar, Shafi Goldwasser, and Sunoo Park. How to incentivize data-driven collaboration among competing parties. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*, pages 213–225. ACM, 2016.
- [AM12] Pablo Daniel Azar and Silvio Micali. Rational proofs. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 1017–1028. ACM, 2012.
- [BK14] Iddo Bentov and Ranjit Kumaresan. How to use bitcoin to design fair protocols. In *International Cryptology Conference*, pages 421–439. Springer, 2014.
- [BKM17] Iddo Bentov, Ranjit Kumaresan, and Andrew Miller. Instantaneous decentralized poker. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 410–440. Springer, 2017.

- [BMTZ17] Christian Badertscher, Ueli Maurer, Daniel Tschudi, and Vassilis Zikas. Bitcoin as a transaction ledger: A composable treatment. In *Annual International Cryptology Conference*, pages 324–356. Springer, 2017.
- [BOM04] Michael Ben-Or and Dominic Mayers. General security definition and composability for quantum & classical protocols. *arXiv preprint quant-ph/0409062*, 2004.
- [But08] Vitalik Buterin. <https://github.com/ethereum/wiki/wiki/White-Paper>, 2013, 2008.
- [Can01] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *Foundations of Computer Science, 2001. Proceedings. 42nd IEEE Symposium on*, pages 136–145. IEEE, 2001.
- [Can06] Ran Canetti. Security and composition of cryptographic protocols: A tutorial. Technical report, Cryptology ePrint Archive, Report 2006/465, 2006. <http://eprint.iacr.org/2006/465>, 2006.
- [CDPW07] Ran Canetti, Yevgeniy Dodis, Rafael Pass, and Shabsi Walfish. Universally composable security with global setup. In *Theory of Cryptography Conference*, pages 61–85. Springer, 2007.
- [CMS16] Jing Chen, Samuel McCauley, and Shikha Singh. Rational proofs with multiple provers. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*, pages 237–248. ACM, 2016.
- [CSV16] Ran Canetti, Daniel Shahaf, and Margarita Vald. Universally composable authentication and key-exchange with global pki. In *IACR International Workshop on Public Key Cryptography*, pages 265–296. Springer, 2016.
- [DDL17] Bernardo David, Rafael Dowsley, and Mario Larangeira. Kaleidoscope: An efficient poker protocol with payment distribution and penalty enforcement. Technical report, Cryptology ePrint Archive, Report 2017/899, 2017. <http://eprint.iacr.org/2017/899>, 2017.
- [DDL18] Bernardo David, Rafael Dowsley, and Mario Larangeira. 21-bringing down the complexity: fast composable protocols for card games without secret state. In *Australasian Conference on Information Security and Privacy*, pages 45–63. Springer, 2018.
- [GHRV14] Siyao Guo, Pavel Hubáček, Alon Rosen, and Margarita Vald. Rational arguments: single round delegation with sublinear verification. In *Proceedings of the 5th conference on Innovations in theoretical computer science*, pages 523–540. ACM, 2014.
- [GHRV16] Siyao Guo, Pavel Hubáček, Alon Rosen, and Margarita Vald. Rational sumchecks. In *Theory of Cryptography Conference*, pages 319–351. Springer, 2016.
- [GKL15] Juan Garay, Aggelos Kiayias, and Nikos Leonardos. The bitcoin backbone protocol: Analysis and applications. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 281–310. Springer, 2015.
- [KZZ16] Aggelos Kiayias, Hong-Sheng Zhou, and Vassilis Zikas. Fair and robust multi-party computation using a global transaction ledger. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 705–734. Springer, 2016.
- [LAF⁺09] Andrew Lutomirski, Scott Aaronson, Edward Farhi, David Gosset, Avinatan Hassidim, Jonathan Kelner, and Peter Shor. Breaking and making quantum money: toward a new quantum cryptographic protocol. *arXiv preprint arXiv:0912.3825*, 2009.
- [Nak08] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. <https://bitcoin.org/bitcoin.pdf>, 2008.
- [PD16] Joseph Poon and Thaddeus Dryja. The bitcoin lightning network: Scalable off-chain instant payments. <https://lightning.network/lightning-network-paper.pdf>, 2016.
- [PSS17] Rafael Pass, Lior Seeman, and Abhi Shelat. Analysis of the blockchain protocol in asynchronous networks. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 643–673. Springer, 2017.

- [Unr04] Dominique Unruh. Simulatable security for quantum protocols. *arXiv preprint quant-ph/0409125*, 2004.
- [Unr10] Dominique Unruh. Universally composable quantum multi-party computation. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 486–505. Springer, 2010.
- [WEH18] Stephanie Wehner, David Elkouss, and Ronald Hanson. Quantum internet: A vision for the road ahead. *Science*, 362(6412):eaam9288, 2018.
- [Wie83] Stephen Wiesner. Conjugate coding. *ACM Sigact News*, 15(1):78–88, 1983.
- [Zha17] Mark Zhandry. Quantum lightning never strikes the same state twice. *Cryptology ePrint Archive, Report 2017/1080*, 2017. <http://eprint.iacr.org/2017/1080>, 2017.

A Appendix

A.1 Proof of Proposition 2

Proof of Proposition 2. We assume some familiarity with Zhandry’s construction (see section 6 of [Zha17] for more details). In his construction, a full bolt is a tensor product of n mini-bolts. A valid mini-bolt with serial number y takes the form $|\Psi\rangle^{\otimes(k+1)}$ where $|\Psi\rangle$ is a superposition of pre-images of y under a certain function f (this is called $f_{\mathcal{A}}$ in Zhandry’s paper, and we do not go into the details of what this function is). The serial number of the full bolt is the concatenation of the serial numbers of the mini-bolts. The computational assumption under which Zhandry’s construction is proved secure is that f is $(2k+2)$ -multi-collision resistant, i.e. it is hard to find $2k+2$ colliding inputs (for this particular function it is easy to find $k+1$ on the other hand). Let \tilde{H} be the function that takes as input $k+1$ pre-images z_1, \dots, z_{k+1} and outputs $f(z_i)$ if this value is the same for all i , and outputs \perp otherwise. Take $H := \tilde{H}^{\times n}$.

For property (I), \mathcal{A}_* again just measures all registers in the computational basis. For property (II), similarly to the proof of Proposition 1, we can construct an adversary \mathcal{A}' that breaks the $(2k+2)$ -multi-collision resistance of f from an adversary \mathcal{A} that wins Game_{12} with non-negligible probability: \mathcal{A}' runs $(\text{Gen}, \text{Ver}, H) \leftarrow \text{QL.Setup}(1^\lambda)$ and samples $|\psi\rangle \leftarrow \text{Gen}$. \mathcal{A}' gives $(\text{Gen}, \text{Ver}, H, |\psi\rangle)$ as input to \mathcal{A} . \mathcal{A} returns x which is parsed as (x_1, \dots, x_n) , where each x_i is a $(k+1)$ -tuple, and $|\psi'\rangle$. \mathcal{A}' then measures each of the n registers of $|\psi'\rangle$ to get n $(k+1)$ -tuples (x'_1, \dots, x'_n) . If there is some i such that (x_i, x'_i) is a $(2k+2)$ -collision, then \mathcal{A}' outputs this. With non-negligible probability \mathcal{A}' outputs a $(2k+2)$ -collision: in fact, since \mathcal{A} wins Game_{12} with non-negligible probability, then $|\psi'\rangle$ must pass Ver with non-negligible probability; from the analysis of Zhandry’s proof, we know that any full bolt that passes verification with non-negligible probability must be such that all mini-bolts have non-negligible weight on at least two pre-images. Since k can be taken to be constant in the size of the security parameter (it is even possible to take $k=1$), then, for any i , there is a non-negligible probability that the $k+1$ entries of x_i are all distinct from the entries of x'_i , which gives a $(2k+2)$ -collision. \square