

A Quantum Money Solution to the Blockchain Scalability Problem

Andrea Coladangelo¹ and Or Sattath²

¹Computing and Mathematical Sciences, Caltech

²Computer Science Department, Ben-Gurion University

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 simple 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 [Zha19]. Our hybrid payment system employs 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.

Note: This work supersedes two previous independent works, [Col19] and [Sat19].

Contents

1	Introduction	3
2	Preliminaries	6
2.1	Notation	6
2.2	Quantum money and quantum lightning	6
2.3	Universal Composability	10
3	Blockchains and smart contracts	12
4	A payment system based on quantum lightning and a classical blockchain	16
4.1	The payment system and its components	17
5	Security	21
6	Practical issues in a less idealized setting	25
6.1	Attacks outside of the idealized setting	25
6.2	Trading a Bolt for a Signature	26
6.3	Security against Sabotage	29
6.4	A resolution of the practical issues	30
7	A practical implementation on Bitcoin and space optimization	32
7.1	Bitcoin to Quantum Money: The Simple Approach	33
7.2	Bitcoin to Quantum Money: Space Optimization	34
8	Comparing our scheme to classical alternatives	34

9 Conclusion	39
A Appendix	43
A.1 Proof of Proposition 2	43
A.2 Standard Security Definitions	43

1 Introduction

Cryptocurrencies, along with blockchains and smart contracts, have recently risen to popular attention, the most well-known examples being Bitcoin and Ethereum [Nak08, But14]. 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.

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 quantum information is inherently well-suited to tackle this problem. We show that a classical blockchain can be leveraged using tools from quantum cryptography (in particular, quantum money) to provide a simple solution to the scalability problem.¹

The main quantum ingredient that we employ is a primitive called *quantum lightning*, formally introduced by Zhandry [Zha19] 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? Our solution leverages a (classical) blockchain to address these questions.

Our contributions 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 is, to the best of our knowledge, the first example of the use of a classical blockchain in combination with quantum cryptographic tools. Our payment system has the following desirable features:

- (i) It is decentralized, requiring no trust in any single entity.
- (ii) Payments involve the exchange of quantum banknotes, and enjoy many of the properties of cash, which cryptocurrencies do not have. For example, transactions are not recorded on the blockchain, and they involve only the payer and the payee. Thus the throughput is unbounded. Payments are as quick as quantum communication, and they do not incur transaction fees.

¹We clarify that this solution only solves the scalability problem for *payment* transactions, and not for the more general *smart contracts* transactions.

- (iii) The rightful owner of a quantum banknote can recover the original value, even if the quantum banknote is damaged or lost; an adversary who tried to abuse this feature would be penalized financially.

Our contribution is primarily conceptual, but our treatment is formal: we work within the Generalized Universal Composability framework (GUC) of Canetti et al. [CDPW07], and we formulate an ideal functionality for a blockchain that supports smart contracts. Note that we do not prove composable security of our payment system. Instead, we prove a “one-shot” version of security, assuming parties have access to such an ideal functionality. Nonetheless, we find it desirable to work within the GUC framework, and we discuss the reason for this choice in more detail below. We also provide an informal construction of our payment system on the Bitcoin blockchain. Its treatment is less rigorous, but it addresses some ways in which the payment system could be optimized.

As a further technical contribution, in order to achieve part (iii), we formalize a novel property of quantum lightning schemes, which we call “bolt-to-signature” capability and is reminiscent of one-time digital signatures. We provide a provably secure construction of a lightning scheme with such a property (assuming a secure lightning scheme which might not satisfy this). The construction is based on the hash-and-sign paradigm as well as Lamport signatures scheme. It allows a user holding a valid quantum lightning state with serial number s to sign a message α *with respect to* s . The security guarantees are that no one who does not possess a lightning state with serial number s can forge signatures *with respect to* s , and once the lightning state is utilized to produce even a single signature, it will no longer pass the lightning verification procedure. We envision that such a primitive could find applications elsewhere and is of independent interest.

How to model a blockchain? 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], which uses the Generalized Universal Composability (GUC) framework by Canetti et al. [CDPW07].

In this work, we define a ledger functionality in the GUC framework which supports universal smart contracts. We use this as an abstraction layer which significantly simplifies exposition and the security analysis. The downside of this approach is that, to the best of our knowledge, there is no known proof of a GUC-secure realization of an ideal functionality for a ledger supporting smart contracts by any cryptocurrency. Proving that a complicated system such as Bitcoin or Ethereum GUC-securely realizes a simple ledger functionality, even without the support for smart contracts, requires already substantial work [BMTZ17]. The upside is that, as soon as one provides a GUC secure realization on some cryptocurrency of our ideal ledger functionality, one can replace the latter in our payment system with its real world implementation, and the security of our payment system would hold verbatim, by the composition properties of the GUC framework. For this reason, we find the approach very desirable. Other approaches that do not fall into the GUC framework include [KMS⁺16] and [DEF18].

We emphasize that we do not prove that our payment system can be composed securely, i.e. we do not define an ideal functionality for our payment system and prove a secure realization of it. Rather, the security we prove is only “one-shot”.

A sketch of our payment system The main ingredient that we employ is quantum lightning. As suggested by Zhandry, this primitive seems well-suited for designing a decentralized payment system, as it is infeasible for anyone to copy 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 quantum banknote is valid. This completes the transaction. Notice that this does not involve interaction with any other third party, and only requires read access to the blockchain. Of course, the same quantum banknote can be successively spent an unlimited number of times in the same manner, without ever posting any new transaction 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 (which are typically in the order of minutes – for example, it is recommended to accept a Bitcoin transaction after 6 confirmations, which takes an hour in expectation) are also no longer a concern since payments are verified immediately by the payee. One can think of the quantum banknotes as providing an off-chain layer that allows for virtually unlimited throughput. We give an extended comparison of our payment system with other off-chain solutions, like Bitcoin’s Lightning Network, in Section 8.

The full payment system includes two additional optional features:

- (i) A mechanism that allows any party in possession of a quantum banknote to recover the coins deposited in the corresponding smart contract. Together with the mechanism outlined earlier, this makes quantum banknotes and coins on the blockchain in some sense interchangeable: one can always convert one into the other by publishing a single transaction on the blockchain.
- (ii) 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.

These two additional desirable features are realizable as long as the quantum lightning scheme employed satisfies an additional property. We formalize this property, which we call “bolt-to-certificate capability”, and show that Zhandry’s proposed constructions satisfy this property. The latter asserts informally that it is possible to measure a valid quantum banknote and obtain a classical certificate, but it is impossible to simultaneously hold both a valid banknote and a valid certificate. In other words, the classical certificate acts as a “proof of destruction”: a certificate which guarantees that no quantum lightning state with the corresponding serial number exists. We also introduce a novel implementation of a one-time signature scheme based on the “bolt-to-certificate capability”, and inspired by Lamport’s signature scheme. This protects honest parties who broadcast a classical certificate to the network from having their classical certificate stolen before it is registered on the blockchain.

We emphasize that the bolt-to-certificate capability of the lightning scheme is only required in order to achieve the two additional features (i) and (ii) above, but is not necessary in order to achieve the basic functionality of our payment system described above.

Limitations of our solution We see two main limitations to our payment system:

- The technologies needed to implement our payment system will not be available in the foreseeable future. For instance, our payment system would require the ability of each party to store large entangled quantum states for extended periods of time. It would also require that each party be able to send quantum states to any party it wishes to make payments to. The latter 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.
- The only known concrete constructions of a quantum lightning scheme rely on non-standard assumptions for security. The construction of Farhi et al. [FGH⁺12a] contains no security proof. The construction of Zhandry [Zha19] is secure based on an assumption related to the multi-collision resistance of certain degree-2 hash functions. Arguably, neither of these assumptions is well-studied.

In Section 8, we extensively discuss the disadvantages (as well as the advantages) of our payment system vis-à-vis with other classical alternatives – namely, a standard cryptocurrency such as Bitcoin and second layer solutions such as the Lightning Network.

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 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 terms PPT and QPT as abbreviations of probabilistic polynomial time and quantum polynomial time respectively.

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 [Zha19], 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 [Zha19], 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 [Zha19]). *A quantum lightning scheme consists of a PPT algorithm $QL.Setup(1^\lambda)$ (where λ is a security parameter) which samples a pair of polynomial-time quantum algorithms ($gen\text{-}bolt$, $verify\text{-}bolt$). $gen\text{-}bolt$ outputs pairs of the form $(|\psi\rangle \in \mathcal{H}_\lambda, s \in \{0, 1\}^\lambda)$. We refer to $|\psi\rangle$ as a “bolt” and to s as a “serial number”. $verify\text{-}bolt$ takes as input a pair of the same form, and outputs either “accept” (1) or “reject” (0) (together with a post-measurement state). They satisfy the following:*

-

$$\Pr[verify\text{-}bolt(|\psi\rangle, s) = 1 : (|\psi\rangle, s) \leftarrow gen\text{-}bolt, (gen\text{-}bolt, verify\text{-}bolt) \leftarrow QL.Setup] = 1 - negl(\lambda)$$

- For a state $|\psi\rangle$ let $\mathcal{M}_{|\psi\rangle}$ be the two outcome measurement which accepts $|\psi\rangle$ and rejects all states orthogonal to it.

$$\Pr[\mathcal{M}_{|\psi\rangle}(|\psi'\rangle) = 1 : (b, |\psi'\rangle) \leftarrow verify\text{-}bolt(|\psi\rangle, s), (|\psi\rangle, s) \leftarrow gen\text{-}bolt, (gen\text{-}bolt, verify\text{-}bolt) \leftarrow QL.Setup(\lambda)] = 1 - negl(\lambda)$$

Here $|\psi'\rangle$ is the post-measurement state upon running $verify\text{-}bolt(|\psi\rangle, s)$.

- For all $s' \in \{0, 1\}^\lambda$,

$$\Pr[verify\text{-}bolt(|\psi\rangle, s') = 1 \wedge s' \neq s : (|\psi\rangle, s) \leftarrow gen\text{-}bolt, (gen\text{-}bolt, verify\text{-}bolt) \leftarrow QL.Setup] = negl(\lambda)$$

The three requirements simply ask that, with overwhelming probability, for any validly generated bolt $|\psi\rangle$, there is a single serial number s such that $(|\psi\rangle, s)$ is accepted by the verification procedure, and that the verification does not perturb the bolt, except negligibly.

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 $(\text{gen-bolt}, \text{verify-bolt}) \leftarrow \text{QL.Setup}(\lambda)$ and sends $(\text{gen-bolt}, \text{verify-bolt})$ to \mathcal{A} .
- \mathcal{A} produces a pair $|\Psi_{12}\rangle \in \mathcal{H}_\lambda^{\otimes 2}, s \in \{0, 1\}^\lambda$.
- The challenger runs $\text{verify-bolt}(\cdot, s)$ on each half of $|\Psi_{12}\rangle$. The output of the game is 1 if both outcomes are “accept”.

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

Definition 2 (Security [Zha19]). *A quantum lightning scheme is secure if, for all polynomial-time quantum adversaries \mathcal{A} ,*

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

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 certificate. Intuitively, this is meant to capture the fact that in the proposed construction of quantum lightning by Zhandry, one can measure a bolt with serial number y in the computational basis to obtain a pre-image of y under some 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 tuple $(\text{gen-bolt}, \text{verify-bolt}, \text{gen-certificate}, \text{verify-certificate})$, where gen-certificate is a QPT algorithm that takes as input a quantum money state and a serial number and outputs a classical string of some fixed length $l(\lambda)$ for some polynomially bounded function l , which we refer to as a *certificate*, and $\text{verify-certificate}$ is a PPT algorithm which takes as input a serial number and a certificate, and outputs “accept” (1) or “reject” (0). The additional property is defined based on the following game Forge-certificate between a challenger and an adversary \mathcal{A} :

- The challenger runs $(\text{gen-bolt}, \text{verify-bolt}, \text{gen-certificate}, \text{verify-certificate}) \leftarrow \text{QL.Setup}(1^\lambda)$ and sends the tuple to \mathcal{A} .
- \mathcal{A} returns $c \in \{0, 1\}^{l(\lambda)}$ and $(|\psi\rangle, s)$.
- The challenger runs $\text{verify-certificate}(s, c)$ and $\text{verify-bolt}(|\psi\rangle, s)$. Outputs 1 if they both accept.

Let $\text{Forge-certificate}(\mathcal{A}, \lambda)$ be the random variable for the output of the challenger in the game above.

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)

$$\Pr[\text{verify-certificate}(s, c) = 1 : (\text{gen-bolt}, \text{verify-bolt}, \text{gen-certificate}, \text{verify-certificate}) \leftarrow \text{QL.Setup}(1^\lambda), \\ (|\psi\rangle, s) \leftarrow \text{gen-bolt}, \\ c \leftarrow \text{gen-certificate}(|\psi\rangle, s)] = 1 - \text{negl}(\lambda)$$

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

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

Notice that property (II) also implies that for most setups and serial numbers s it is hard for any adversary to find a c such that $\text{verify-certificate}(s, c) = 1$ 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 **Forge-certificate**: Upon receiving $(\text{gen-bolt}, \text{verify-bolt}, \text{gen-certificate}, \text{verify-certificate})$ from the challenger, \mathcal{A}' computes $(|\psi\rangle, s) \leftarrow \text{gen-bolt}$; then runs \mathcal{A} on input s to obtain some c . \mathcal{A}' returns $c, |\psi\rangle$ and s to the challenger. We emphasize that in game **Forge-certificate** it is the adversary himself who generates the state $|\psi\rangle$. This is important because when we employ the quantum lightning scheme later on parties are allowed to generate their own quantum banknotes.

Proposition 1. *Any scheme that uses Zhandry’s construction instantiated with a non-collapsing hash function [Zha19] satisfies the property of Definition 3.*

Proof. Zhandry already proved that his scheme satisfies Definitions 1 and 2. As we will see, since our construction does not change **QL.Setup**, **gen-bolt**, and **verify-bolt**, we only need to prove that it satisfies Definition 3. We refer the reader to [Zha19] for a definition of non-collapsing hash function. In Zhandry’s construction based on a non-collapsing hash function, **QL.Setup**(1^λ) outputs $(\text{gen-bolt}, \text{verify-bolt})$. A bolt generated from **gen-bolt** has the form $|\Psi\rangle = \bigotimes_{i=1}^n |\psi_{y_i}\rangle$, where $y_i \in \{0, 1\}^\lambda$ for all i and $|\psi_{y_i}\rangle = \sum_{x: H(x)=y_i} |x\rangle$, where H is a non-collapsing hash-function, and $n \in \mathbb{N}$ is polynomial in the security parameter λ . **verify-bolt** has the form of n -fold product of a verification procedure “Mini-Ver” which acts on a single one of the n registers, though it is not crucial to understand how Mini-Ver for this work. In Zhandry’s construction, the serial number associated to the bolt $|\Psi\rangle$ above is $s = (y_1, \dots, y_n)$.

We define **gen-certificate** as the QPT algorithm that measures $|\psi\rangle$ in the standard basis, and outputs the outcome. When applied to an honestly generated bolt, the outcomes are pre-images x_i of y_i , for $i = 1, \dots, n$. We define **verify-certificate** as the deterministic algorithm which receives a serial number $s = (y_1, \dots, y_n)$ and a certificate $c = (x_1, \dots, x_n)$ and checks that for all $i \in [n]$, $H(x_i) = y_i$.

It is clear that (I) holds.

For property (II), suppose there exists \mathcal{A} such that $\Pr[\text{Forge-certificate}(\mathcal{A}, \lambda) = 1]$ 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 $(\text{gen-bolt}, \text{verify-bolt}) \leftarrow \text{QL.Setup}(1^\lambda)$. Let H be the non-collapsing hash function hard-coded in the description of **verify-bolt**. \mathcal{A}' sends the tuple $(\text{gen-bolt}, \text{verify-bolt}, \text{gen-certificate}, \text{verify-certificate})$ to \mathcal{A} , where the latter two are defined as above in terms of H . \mathcal{A} returns $(c, |\psi\rangle)$, where c is parsed as $c = (x_1, \dots, x_n)$. \mathcal{A}' then measures each of the n registers of $|\psi\rangle$ to get $c' = (x'_1, \dots, x'_n)$. If $x_i \neq x'_i$, then \mathcal{A}' outputs (x_i, x'_i) . We claim that with non-negligible probability \mathcal{A}' outputs a collision for H . To see this, notice that since \mathcal{A} wins **Forge-certificate** with non-negligible probability, then $|\psi\rangle$ must pass **verify-bolt** with non-negligible probability; and from the analysis of [Zha19], 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. For more details, see the proof of Theorem 5 in [Zha19], and specifically the following claim:

If the bolts are measured, two different pre-images of the same y , and hence a collision for $H^{\otimes r}$, will be obtained with probability at least $1/200$.

□

Proposition 2. *Zhandry’s construction based on the multi-collision resistance of certain degree-2 hash functions (from section 6 of [Zha19]) satisfies the property of Definition 3.*

The proof is similar to the proof of Proposition 1. We include it for completeness in Appendix A.1.

Fact 4. *Zhandry’s construction from Proposition 2 only requires a common random string (CRS) for **QL.Setup**.*

This fact will be relevant when we discuss a concrete implementation in a system such as Bitcoin. For more details, see Section 7.

Other Quantum Lightning Schemes? Farhi et al. [FGH⁺12b] constructed a public quantum money scheme which they conjectured to have the following property:

A rogue mint running the same algorithm as the mint can produce a new set of money pairs, but (with high probability) none of the serial numbers will match those that the mint originally produced.

This is less formal, but we believe captures Zhandry’s security definition of a quantum lightning (though it was introduced roughly 7 years earlier).

Very recently, Peter Shor gave a talk on ongoing (unpublished) work for a lattice-based quantum lightning scheme (see <https://youtu.be/8fzLByTn8Xk>).

Here, we briefly compare the two existing published constructions, and we omit Shor’s scheme from this discussion.

- Farhi et al. only provide several plausibility arguments supporting the security of their scheme, but their work does not contain any security proof. In a follow-up work, Lutomirski [Lut11] showed a security reduction from a problem *related* to counterfeiting Farhi et al.’s money scheme. Even though it does not have a (proper) security proof, Farhi et al.’s scheme was first published 8 years ago, and was not broken since then. Zhandry’s construction is proved to be secure under a non-standard hardness assumption, which was first introduced in that work. In his Eurocrypt talk, Zhandry reports that the construction is “broken in some settings” – see <https://youtu.be/fjumbNTZSik?t=1302>.²
- The scheme of Farhi et al. does not require a trusted setup.
- The scheme of Farhi et al. does not have a bolt-to-certificate capability. This property is required in our payment system to transform quantum banknotes back to coins on the blockchain – see Section 4.

2.3 Universal Composability

This section 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. As mentioned in the Introduction, we elect to work in the GUC, because this provides a convenient formal abstraction layer for modeling a ledger functionality that supports smart-contracts, and allows for a clean security proof in the idealized setting in which parties have access to this functionality.

The reader familiar with the UC and GUC framework 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

²Additionally, the hardness assumption was changed between the second and third version on the IACR e-print. This might suggest that the original hardness assumption was also broken. The most updated e-print version does not discuss that discrepancy.

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 5. *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 5.

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 6 (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. We remark that \mathcal{F}_{Ledg} is somewhat of an idealized functionality which allows us to obtain a clean description and security proof for our payment system. However, \mathcal{F}_{Ledg} does not capture, for example, attacks that stem from miners possibly delaying honest parties’ messages from being recorded on the blockchain. We discuss this and other issues extensively in Section 6. We also discuss ways to resolve such issues in detail, but we do not formally incorporate these in the description of our payment system in Section 4: since we view our contribution as primarily conceptual, we elect to keep the exposition and security proof of the basic payment system as clean and accessible as possible.

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 Section 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 \mathcal{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.)

Add party: Upon receiving a message `AddParty` from a party with PID $pid = (id, d)$, send `(AddedParty, 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 $pid.id \leftarrow id$ and $pid.coins \leftarrow d$.

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 $pid \notin \text{parties}$, else $d = pid.coins$. (We abuse notation here in that, when taken as part of a message, pid is treated as a string, but, when called by the functionality, it is a variable with attributes $pid.id$ and $pid.coins$).

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

- If $pid' \in \text{parties}$ and $pid.coins > d$, update $pid'.coins \leftarrow pid'.coins + d$ and $pid.coins \leftarrow pid.coins - d$. Set $trId = |\text{AllTransactions}| + 1$. Add a variable named $trId$ to AllTransactions , with attribute $trId.transaction = (pid, pid', d, t)$. 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 $trId \notin \text{allTransactions}$, and $s = trId.transaction$ otherwise.

Add/trigger smart contract: Upon receiving a message `(AddSmartContract, Params=(I, D, ϕ , st0))`, 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 $ssid = |\text{contracts}| + 1$. Add a variable named $ssid$ to contracts with attributes $ssid.Params = (I, D, \phi, st_0)$, $ssid.state = st$ and $ssid.coins \leftarrow 0$. Send a message `(RecordedContract, ssid)` to P . Then, do the following:

- **Initialization phase:** Wait to get message `(InitializeWithCoins, ssid, Params = (I, D, ϕ , st0))` from party with PID pid for all $pid \in I$. When all messages are received, and if, for all $pid \in I$, $pid.coins \geq d_{pid}$, then, for all $pid \in I$, update: $pid.coins \leftarrow pid.coins - d_{pid}$ and $ssid.coins \leftarrow ssid.coins + d_{pid}$. Set $st \leftarrow st_0$ (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 $pid \in \text{parties}$ (where it can also be $d = 0$) such that $\phi(pid, w, t, st, d) \neq \perp$, do the following:
 - If $d > 0$, update $pid.coins \leftarrow pid.coins - d$ and $ssid.coins \leftarrow ssid.coins + d$
 - Update $(st, e) \leftarrow \phi(pid, w, t, st, d)$.
 - If $e = \text{“all coins”}$, let $q := ssid.coins$. Send the message `(Reward, ssid, q)` to the party with PID pid and update $pid.coins \leftarrow pid.coins + q$ and $ssid.coins \leftarrow 0$. If $e > 0$ and $ssid.coins > e$, send the message `(Reward, ssid, e)` to the party with PID pid , and update $pid.coins \leftarrow pid.coins + e$ and $ssid.coins \leftarrow ssid.coins - e$. Else, if $ssid.coins = e' \leq e$, send the message `(Reward, ssid, e')` to the party with PID pid , and update $pid.coins \leftarrow pid.coins + e'$ and $ssid.coins \leftarrow 0$. 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 $ssid \notin \text{contracts}$, and $z = (ssid.Params, ssid.state, ssid.coins)$ otherwise.

Figure 1: Global ledger ideal functionality \mathcal{F}_{Ledg}

We think of “Retrieve” operations as being fast, or free, as they do not alter the state of the ledger. We think of “Add” and “Trigger” operations as being slow, as they alter the state of the ledger.

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 7 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 Section 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.

Let $\lambda \in \mathbb{N}$. From now onwards, we assume that

$$(\text{gen-bolt}, \text{verify-bolt}, \text{gen-certificate}, \text{verify-certificate}) \leftarrow \text{QL.Setup}(\lambda),$$

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). We are thus assuming a trusted setup for our payment system (note there are practical ways to ensure that such a setup, which is a one-time procedure, is performed legitimately even in a network where many parties might be dishonest). We also assume that all parties have access to an authenticated and ideal quantum channel.

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 .

In Definition 7 below, 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 $\text{st} = (\text{serial}, \text{ActiveLostClaim})$. 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 $(\text{ChallengeClaim}, c, s')$ to the contract, where s' is a proposed new serial number. We say that c is a valid classical certificate for the current value s of serial if $\text{verify-certificate}(s, c) = 1$. Such a c 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 c , and thus that the current active lost-banknote claim is malicious. If c 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 (`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 (`RecoverCoins`, c). If c 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.

$\phi_{\mathfrak{s}}$ ($pid, w, t, (\text{serial}, \text{ActiveLostClaim}), d$) 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. $\phi_{\mathfrak{s}}$ 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), $l(\lambda)$ the length of a certificate in the lightning scheme, a description of `verify-certificate`. The circuit $\phi_{\mathfrak{s}}$ outputs new values for the state variables and an amount of coins as follows:

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

- If `ActiveLostClaim` = “No active claim”:
 - If $w = \text{BanknoteLost}$ and $d = d_0$, then ϕ outputs $((\text{serial} = s, \text{ActiveLostClaim} = \text{“Claim by } pid \text{ at time } t\text{”}), 0)$ (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 $w = (\text{RecoverCoins}, c)$, where $c \in \{0, 1\}^{l(\lambda)}$ and `verify-certificate`(s, c) = 1, then ϕ outputs $((\text{serial} = \perp, \text{ActiveLostClaim} = \perp), \text{“all coins”})$
- If `ActiveLostClaim` = “Claim by pid' at time t_0 ” for some pid', t_0 :
 - If $w = (\text{ChallengeClaim}, c, s')$, where $c \in \{0, 1\}^{l(\lambda)}$ and `verify-certificate`(s, c) = 1, 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 $\phi_{\mathfrak{s}}$ for banknote-contracts

Definition 7. (*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 $d_{pid} \in \mathbb{N}$, $st_0 = (s, \text{“No active claim”})$ for some $s \in \{0, 1\}^\lambda$, and circuit $\phi = \phi_{\mathfrak{s}}$, where $\phi_{\mathfrak{s}}$ 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 $pid.coins > d$ in \mathcal{F}_{Ledg} (d is the “value” of the prospective banknote).

- Run $(|\psi\rangle, s) \leftarrow \text{gen-bolt}$.
 - Send $(\text{AddSmartContract}, \text{Params})$ 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 $\text{serial} = s$, 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 $(|\Psi\rangle, s, ssid, d)$ 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 $z = (\text{Params}, (s, \text{“No active claim”}), d)$, then Q checks that the parameters Params are of the form of a banknote-contract (from Definition 7). If so, runs $\text{verify-bolt}(|\Psi\rangle, s)$ and checks that the outcome is 1. If so, sends the message accept to P .
 - Else, Q sends the message reject and the state $|\Psi\rangle$ back to P .
-

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. After this, we will 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 7, and in particular of the circuit ϕ_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 with serial number s that someone else possesses, then the legitimate owner can run $\text{gen-certificate}(|\psi\rangle, s)$ where $|\Psi\rangle$ is the legitimate bolt, and obtain a valid certificate c . He can then send c 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 c 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. Or alternatively, an adversary could monitor the network for lost-banknote claims by honest parties, and whenever he sees one he will delay this claim, and instruct a corrupt party to make the same claim so that it is registered first on the ledger. In our analysis, we do not worry about such attacks, 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 the first issue can still be resolved elegantly. The second issue has a satisfactory resolution but it is trickier to analyze formally. We discuss this in more detail in Section 6.

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 $(s, \text{“Active claim by } pid \text{ at time } t\text{”})$ (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 7.

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 `(RetrieveContract, ssid)` to \mathcal{F}_{Ledg} . Upon receiving a message `(RetrieveContract, ssid, z)` from \mathcal{F}_{Ledg} , if $z = (\text{Params}, (s, \text{“Claim by } pid' \text{ at time } t \text{”}), d)$ for some pid', t, d and for some banknote-contract parameters `Params`:
 - Run $c \leftarrow \text{gen-certificate}(|\Psi\rangle, s)$.
 - Sample $(|\Psi'\rangle, s') \leftarrow \text{gen-bolt}$.
 - Send `(Trigger, ssid, (ChallengeClaim, c, 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 $c \leftarrow \text{gen-certificate}(|\Psi\rangle, s)$.
 - Send message `(Trigger, ssid, (RecoverCoins, c), 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. Security with respect to this adversarial model is formally captured by Theorem 8. At a high-level, Theorem 8 establishes that, within this adversarial model, no adversary can increase his “value” beyond what he has legitimately spent or received to and from honest parties. This captures, for example, the fact that the adversary will not be able to double-spend his banknotes, or successfully file a “lost banknote claim” for banknotes he does not legitimately possess.

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 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 `AddParty` 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

$\text{AdversaryValueReceived} + d$. 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`) $\text{AdversaryValueReceived} \leftarrow \text{AdversaryValueReceived} + d$.

- (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 $\text{AdversaryValueCurrentOrSpent} \leftarrow \text{AdversaryValueCurrentOrSpent} + d$.
- (v) When a corrupted party receives d coins from a banknote-contract, then $\text{AdversaryValueCurrentOrSpent} \leftarrow \text{AdversaryValueCurrentOrSpent} + d$. 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 $\text{AdversaryValueCurrentOrSpent} - \text{AdversaryValueReceived} > 0$. 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 $\mathcal{F}_{Ledg}\text{-EXEC}_{\pi, \mathcal{A}, \mathcal{E}}^{(\text{MaxNetValue})}(\lambda, z)$ the maximum value of $\text{AdversaryValueCurrentOrSpent} - \text{AdversaryValueReceived}$ during an execution of π with adversary \mathcal{A} and environment \mathcal{E} , with global shared functionality \mathcal{F}_{Ledg} .

Theorem 8 (Security). *For any quantum polynomial-time adversary \mathcal{A} and quantum polynomial-time environment \mathcal{E} ,*

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

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

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

Then, we go through all of the possible ways that an adversary can increase its net value, i.e. increase the quantity $\text{AdversaryValueCurrentOrSpent} - \text{AdversaryValueReceived}$: 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 $\text{AdversaryValueCurrentOrSpent} - \text{AdversaryValueReceived} > 0$. 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. $\text{AdversaryValueCurrentOrSpent} - \text{AdversaryValueReceived} > 0$. 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 (i.e. game Counterfeit): \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 $(\text{Trigger}, ssid, (\text{RecoverCoins}, c), 0)$ to \mathcal{F}_{Ledg} for some $ssid$ and c such that $\text{verify-certificate}(s, c) = 1$, where $ssid.\text{state} = (s, \text{“No active claim”})$, 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 valid c 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 game Forge-certificate 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 c such that $\text{verify-certificate}(s, c) = 1$ without ever possessing a valid bolt with serial number s . Again, such an adversary could be used, together with \mathcal{E} to win Forge-certificate 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 $(\text{Trigger}, ssid, (\text{ChallengeClaim}, c), 0)$ to \mathcal{F}_{Ledg} for some $ssid$ such $ssid.\text{state} = (s, \text{“Claim by } pid \text{ at time } t\text{”})$ for some s, pid, t with pid honest and c such that $\text{verify-certificate}(s, c) = 1$. 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 c such that $\text{verify-certificate}(s, c) = 1$. Analogously to earlier, this adversary, together with \mathcal{E} , can be used to win Forge-certificate.
 - * \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 a c with $\text{verify-certificate}(s, c) = 1$ and spent a bolt with serial number s successfully. Such an \mathcal{A} can be used, together with \mathcal{E} , to win Forge-certificate.

□

Remark: The security guarantee of Theorem 8 establishes that an adversary cannot end up with more value than he started with (after taking into account the amount he received from honest parties and the amount he successfully spent to honest parties). However, we do not analyze formally attacks which do not make the adversary gain value directly, but

which “sabotage” honest parties, making them lose value. We believe that it should be possible to capture such attacks within our model by modifying the way we keep track of the adversary’s value, but we leave this analysis for future work. We discuss and formalize the notion of “sabotage” in detail in Section 6.3.

6 Practical issues in a less idealized setting

The ideal functionality \mathcal{F}_{Ledg} defined in Section 3 does not capture adversaries that are allowed to see messages sent by honest parties to \mathcal{F}_{Ledg} before they are registered on the ledger, and who could try to use this information to their advantage: by definition of \mathcal{F}_{Ledg} , messages are processed and registered on the ledger in exactly the order that they are sent, and are not seen by the adversary until they make it onto the ledger. While such a definition of \mathcal{F}_{Ledg} makes for a clear exposition and clean proof of security, it is in practice unrealistic. In typical real-world implementations of blockchains, miners (and hence potentially adversaries) can see a pool of pending messages which have not yet been processed and registered on the blockchain, and can potentially delay the processing of certain messages, while speeding up the processing of others. This makes the system susceptible to the attacks described in the following subsection.

6.1 Attacks outside of the idealized setting

- (i) An adversary could file a malicious “lost banknote claim” (running the protocol of Fig. 5) corresponding to a serial number s of a quantum banknote that he does not possess. This would prompt an honest party who possesses a valid quantum banknote with serial number s to publish the corresponding classical certificate x , in order to stop the malicious claim. If the adversary could read this message before it is published on the ledger, it could instruct a corrupt party to also publish x , and attempt to have this message appear first on the ledger. This would effectively result in the adversary having stolen the honest party’s value associated to the serial number s .
- (ii) Suppose an honest party wants to trade their quantum banknote with serial number s (registered on the ledger) for the coins deposited in the corresponding contract. The honest party executes the protocol of Fig. 7. This includes publishing the classical certificate x associated to s . An adversary who sees x before it is registered on the ledger, could instruct a corrupt party to make the same claim for the coins in the contract associated to s . If the corrupt party’s message is processed faster than the honest party’s message, the adversary has succeeded in stealing the coins deposited in the contract.
- (iii) Suppose an honest party has lost a valid quantum banknote with serial number s associated to some contract on the ledger. The honest party files a “lost banknote claim” by executing the protocol of Fig. 5. An adversary who hears this before the claim is registered on the ledger could instruct a corrupt party to make a similar claim, and have it appear on the ledger before the honest claim. This would result in the corrupt party obtaining a valid quantum banknote associated to the above contract.
- (iv) The unforgeability property alone is not enough for *public* quantum money, as it allows sabotage: an attacker might be able to burn other people’s money, without a direct gain from the attack. Consider an adversary that wants to harm its competitor. The adversary does not follow the honest protocol that generates the quantum money state. Instead, it could create a tweaked quantum money which has a noticeable probability to pass verification once, and fail the second time. This way, the adversary could buy some merchandise using this tweaked quantum money. When the merchant will try to pay to others using this quantum money, the verification will fail – the receiver will run the verification (and this is exactly the second verification), which will cause it to fail. This is not an issue with schemes in which the verification is a

projective (or very close to being projective), but, for example, the scheme by Farhi et al. [FGH⁺12b] is *not* projective.

Indeed, even though our security proof (see Theorem 8) guarantees that an adversary cannot end up with more money than he was legitimately given, it does not rule out sabotage.

- (v) Similarly to the sabotage attack on the verification of quantum money, an attacker might sabotage the ability to produce a valid certificate from the verified money.

In the next Section 6.2, we will introduce a novel property of quantum lightning schemes, which we call *bolt-to-signature capability*, we give a provably-secure construction of it. We will employ this in Section 6.4 to give a somewhat elegant resolution to issues (i) and (ii) above. In Section 6.3, we show that our scheme is secure against sabotage attacks, hence resolving issues (iv) and (v). In Section 6.4 we discuss a practical approach to resolving issue (iii) above, though it does not contain a formal analysis for reasons which are made clear there.

6.2 Trading a Bolt for a Signature

We define a new property of a quantum lightning scheme which we call “trading a bolt for a signature” (and we say that a lightning scheme has *bolt-to-signature* capability). This is an extension of the property of “trading a bolt for a certificate” defined in Section 2.2. The known concrete constructions of quantum lightning do not possess this property, but we will show (Fig. 8) that a lightning scheme with *bolt-to-certificate* capability can be bootstrapped to obtain a lightning scheme with *bolt-to-signature* capability. We envision that this primitive could find application elsewhere, and is of independent interest.

We start with an informal description of this property, highlighting the difference between the *certificate* and *signature* properties. Suppose Alice shows Bob a certificate with a serial number s . Bob can conclude that Alice cannot hold a bolt with the same serial number. In particular, if she held such bolt, it means she must have measured and destroyed it to produce the certificate. For the *bolt-to-signature* property, we ask the following:

- There is a way for Alice, who holds a bolt with serial number s , to produce a signature of any message α of her choice, *with respect to* serial number s .
- The signature should be verifiable by anyone who knows s .
- Just like a certificate, anyone who accepts the signature of α with respect to s can conclude that Alice can no longer hold a quantum money with serial number s ;
- (one-time security) As long as Alice signs a *single* message α , no one other than Alice should be able to forge a signature for a message $\alpha' \neq \alpha$ (even though her state is no longer a valid bolt, Alice can still sign more than one message, but the unforgeability guarantee no longer holds).

A quantum lightning scheme with *bolt-to-signature* capability is a quantum lightning scheme with two additional algorithms: **gen-sig** is a QPT algorithm which receives as input a quantum state, a serial number, and a message of any length, and outputs a classical signature. **verify-sig** is a PPT algorithm which receives a serial number, a message of any length and a signature, and either accepts or rejects. Thus, we modify the setup procedure of the lightning scheme so that **QL.Setup** outputs a tuple (**gen-bolt**, **verify-bolt**, **gen-sig**, **verify-sig**).

The definition of the property has a completeness and a soundness part. The latter is formally defined through the following game **Forge-sig**. The game **Forge-sig** is similar in spirit to the game for *onetime* security of a standard digital signature scheme – see, e.g. [KL14, Definition 12.14], [Gol04, Definition 6.4.2].

- The challenger runs (**gen-bolt**, **verify-bolt**, **gen-sig**, **verify-sig**) \leftarrow **QL.Setup**(1^λ) and sends this tuple to \mathcal{A} .
- \mathcal{A} sends $(|\psi\rangle, s)$ and α to the challenger.

- The challenger runs `verify-bolt`($|\psi\rangle, s$). If this rejects, the challenger outputs “0”. Else it proceeds to the next step.
- Let $|\psi'\rangle$ be the leftover state. The challenger runs $\sigma \leftarrow \text{gen-sig}(|\psi'\rangle, s, \alpha)$ and sends σ to \mathcal{A} .
- \mathcal{A} returns a pair (α', σ') .
- The challenger checks that $\alpha \neq \alpha'$ and runs `verify-sig`(s, α', σ'). If the latter accepts, the challenger outputs “1”.

Let $\text{Forge-sig}(\mathcal{A}, \lambda)$ be the random variable for the outcome of the game.

Definition 9 (Trading a bolt for a signature). *We say that a quantum lightning scheme has bolt-to-signature capability if the following holds:*

(I) For every α :

$$\begin{aligned} \Pr[\text{verify-sig}(s, \sigma, \alpha) = 1 : (\text{gen-bolt}, \text{verify-bolt}, \text{gen-sig}, \text{verify-sig}) \leftarrow \text{QL.Setup}(1^\lambda), \\ (|\psi\rangle, s) \leftarrow \text{gen-bolt}, \\ \sigma \leftarrow \text{gen-sig}(|\psi\rangle, s, \alpha)] = 1 - \text{negl}(\lambda) \end{aligned}$$

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

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

Informally, the security definition based on game `Forge-sig` guarantees that:

- As long as Alice signs only one message with respect to s , no one except her can forge a signature of another message with respect to s . This property is very similar to one-time unforgeability for digital signatures, if one views the bolt as a secret key, and the serial number s as a public key. The difference is that the “secret key” in this case has meaning beyond enabling the signing of messages: it can be spent as quantum money. This is what make the next property important.
- Signing a message destroys the bolt, i.e. it is infeasible to simultaneously produce both a valid signature with respect to a serial number s and a bolt with serial number s which passes verification (an adversary who can succeed at this is easily seen to imply an adversary who wins `Forge-sig`). This property is unique to the quantum lightning setting. It says that signing a message with respect to serial number s inevitably destroys the bolt with serial number s . We remark that it is possible to sign more messages with the leftover state, but such state will no longer pass the quantum lightning verification procedure, i.e. it can no longer be spent. One can think of the bolt with serial number s as being “burnt” once the owner decides to use it to sign a message.

We are now ready to present our construction of a quantum lightning scheme with bolt-to-signature capability. The construction is based on the hash-and-sign paradigm (see, e.g., [KL14]), as well as Lamport signatures [Lam79] (familiarity with those is helpful, although our presentation is self-contained). For convenience, we use \mathcal{H} to denote a family of fixed-length hash functions, and we use the notation $H \leftarrow \mathcal{H}(1^n)$ to denote an efficient random sampling of a hash function H with output length n from the family \mathcal{H} .

Given: A quantum lightning scheme with bolt-to-certificate capability with setup procedure QLC.Setup . A family \mathcal{H} of fixed-length collision-resistant hash functions.

QLDS.Setup : takes as input a security parameter λ , and outputs a tuple of (descriptions of) algorithms (QLDS.Gen , QLDS.Ver , QLDS.gen-sig , QLDS.verify-sig).

- Let $n = \text{poly}(\lambda)$. Sample $(\text{gen-bolt}, \text{verify-bolt}, \text{gen-bolt}, \text{verify-certificate}) \leftarrow \text{QLC.Setup}()$. Sample $H : \{0, 1\}^* \rightarrow \{0, 1\}^n \leftarrow \mathcal{H}(1^n)$.
- QLDS.Gen : Run gen-bolt $2n$ times to obtain bolts $(|\psi_i\rangle \in \mathcal{H}_\lambda, s_i)$, for $i = 1, \dots, 2n$. Let $|\Psi\rangle = \bigotimes_{i=1, \dots, 2n} |\psi_i\rangle \in \mathcal{H}_\lambda^{\otimes 2n}$. Let $s = s_1 || \dots || s_{2n}$. Output $(|\Psi\rangle, s)$.
- QLDS.Ver : Takes as input a state $|\Psi\rangle \in \mathcal{H}_\lambda^{\otimes 2n}$. Applies QLDS.Ver to each of the $2n$ factors, and outputs “accept” if all $2n$ verifications output “accept”.
- QLDS.gen-sig : Takes as input a state $|\Psi\rangle \in \mathcal{H}_\lambda^{\otimes 2n}$, a serial number s and a message α of any length. Lets $\beta = H(\alpha) \in \{0, 1\}^n$. For $i = 1, \dots, n$:
Run gen-certificate on the $(\beta_i \cdot n + i)$ -th factor to obtain a certificate x_i . Outputs $\sigma = x_1 || \dots || x_n$.
- QLDS.verify-sig : Takes as input a serial number s , a message α , a signature σ . Parses s as $s = s_1 || \dots || s_{2n}$, σ as $\sigma = x_1 || \dots || x_n$. Computes $\beta = H(\alpha)$. For $i = 1, \dots, n$:
Let $r_i = \text{verify-certificate}(s_{\beta_i \cdot n + i}, x_i)$. Output “accept” if $r_i = 1$ for all i .

Figure 8: Our construction of a quantum lightning scheme with bolt-to-signature capability QLDS

We state and prove the main theorem of this section.

Theorem 10. *If there exists a secure quantum lightning scheme which has a bolt-to-certificate capability, and a family of fixed length quantum-secure collision-resistant hash functions, then there exists a quantum lightning scheme with bolt-to-signature capability.*

Specifically, under the assumptions that QLC is a secure quantum lightning scheme with a bolt-to-certificate capability and that \mathcal{H} is a fixed-length family of collision-resistant hash functions, the construction QLDS of Fig. 8 is a secure quantum lightning scheme with bolt-to-signature capability.

Proof. First of all, it is clear that QLDS is still a secure quantum lightning scheme according to Definition 2.

The fact that QLDS satisfies the correctness requirement of a quantum lightning scheme with bolt-to-signature capability ((I) in Definition 9) follows from the correctness property of QLC (namely (I) in Definition 3) and that n is at most polynomial in λ .

Assume towards a contradiction that there exists a QPT adversary \mathcal{A} that wins the Sig-Forge game with non-negligible probability $\epsilon(\lambda)$. We construct an adversary \mathcal{B} that wins the game Forge-certificate (from Definition 3) with probability at least $\epsilon(\lambda)$. \mathcal{B} runs as follows:

- \mathcal{B} receives a tuple $(\text{gen-bolt}, \text{verify-bolt}, \text{gen-certificate}, \text{verify-certificate})$ from the challenger.
- \mathcal{B} constructs algorithms gen-sig and verify-sig from gen-certificate and $\text{verify-certificate}$. Sends the tuple $(\text{gen-bolt}, \text{verify-bolt}, \text{gen-sig}, \text{verify-sig})$ to \mathcal{A} .
- \mathcal{A} returns a pair $(|\psi\rangle, s)$ and a message α . Let $\beta = H(\alpha)$. \mathcal{B} simulates the next steps of the challenger in Forge-sig with the following modification: it runs $\text{verify-bolt}(|\psi\rangle, s)$ but only measures the registers corresponding to indexes associated with β (i.e. $\beta_i \cdot n + i$ for $i \in [n]$). It then runs $\sigma \leftarrow \text{gen-sig}(|\psi'\rangle, s, \alpha)$, where $|\psi'\rangle$ is the leftover state after verification. \mathcal{B} sends σ to \mathcal{A} , and \mathcal{A} returns a pair (α', σ') , where $\sigma' = x'_1 || \dots || x'_n$.
- \mathcal{B} computes $\beta = H(\alpha)$ and $\beta' = H(\alpha')$. If $\beta \neq \beta'$, let i be the first index such that $\beta_i \neq \beta'_i$. \mathcal{B} outputs $(|\psi_{\beta'_i \cdot n + i}\rangle, s_{\beta'_i \cdot n + i})$ and x'_i .

We analyze the winning probability of \mathcal{B} in game Forge-certificate . With probability at least $\epsilon(\lambda)$ it is $\alpha \neq \alpha'$ (otherwise \mathcal{A} simply loses). Moreover, with overwhelming probability,

it must be $\beta \neq \beta'$, otherwise it is immediate such an \mathcal{A} could be used to break collision-resistance of H . Hence, with non-negligible probability, there is an index i such that $\beta_i \neq \beta'_i$. Using the definition of QLDS.verify-sig , we deduce that with probability at least ϵ it must be that $\text{verify-certificate}(s_{\beta'_i, n+i}, x'_i) = 1$. Moreover, the state $(|\psi_{\beta'_i, n+i}\rangle)$ was not measured, and it must pass verification with probability at least $\epsilon(\lambda)$. \square

6.3 Security against Sabotage

To address Item (iv) in Section 6.1, we define two security games that capture the notion of sabotage; the first is denoted **Sabotage-Money**:

- The challenger runs $(\text{gen-bolt}, \text{verify-bolt}) \leftarrow \text{QL.Setup}(\lambda)$ and sends $(\text{gen-bolt}, \text{verify-bolt})$ to \mathcal{A} .
- \mathcal{A} outputs a quantum state $|\psi\rangle$ and sends it to the challenger.
- The challenger runs **verify-bolt** two consecutive times on the quantum state $|\psi\rangle$.
- The adversary wins if the first verification accepts with a serial number s and the second rejects, or accepts with a serial number $s' \neq s$. Let $\text{Sabotage-Money}(\mathcal{A}, \lambda)$ be the random variable that is 1 if the adversary \mathcal{A} wins, and is 0 otherwise.

Definition 11 (Security against sabotage). *A quantum lighting scheme is secure against sabotage if for every QPT \mathcal{A} there exists a negligible function negl such that:*

$$\Pr(\text{Sabotage-Money}(\mathcal{A}, \lambda) = 1) = \text{negl}(\lambda)$$

The security against sabotage was first defined in the context of quantum money in [BS16] (though, the term sabotage was not used).

We extend the notion of sabotage in the natural way for schemes with bolt-to-certificate or bolt-to-signature capability. Our goal is to avoid a scenario in which an adversary gives a user a quantum lightning state which passes verification, but later fails to produce a valid certificate or signature. We define the following experiment, **Sabotage-Certificate**:

1. The challenger runs $(\text{gen-bolt}, \text{verify-bolt}, \text{gen-certificate}, \text{verify-certificate}) \leftarrow \text{QL.Setup}(\lambda)$ and sends that tuple to \mathcal{A} .
2. \mathcal{A} sends $|\psi\rangle, s$ to the challenger.
3. The challenger runs $\text{verify-bolt}(|\psi\rangle, s)$. If verification fails, set $r = 1$. Otherwise, the challenger uses the post-measured state $|\psi'\rangle$ to generate a certificate $c \leftarrow \text{gen-certificate}(|\psi'\rangle, s)$, and checks whether it is a valid certificate: $r \leftarrow \text{gen-certificate}(|\psi'\rangle, s)$.
4. \mathcal{A} wins if $r = 0$. Let $\text{Sabotage-Certificate}(\mathcal{A}, \lambda)$ be the random variable that is 1 if the adversary \mathcal{A} wins, and is 0 otherwise.

Definition 12. *A quantum lighting scheme with a bolt-to-certificate capability is secure against sabotage if, in addition to the requirement in Definition 11, for every QPT \mathcal{A} there exists a negligible function negl such that:*

$$\Pr(\text{Sabotage-Certificate}(\mathcal{A}, \lambda) = 1) = \text{negl}(\lambda)$$

We suspect that Zhandry's construction based on non-collapsing hash functions, as well as the construction by Farhi et al. (see p. 10) does not satisfy the security against sabotage. Fortunately, the construction based on the multi-collision resistance is secure against sabotage:

Proposition 3. *The quantum lightning construction with the bolt to certificate capability discussed in Proposition 2 is secure against sabotage.*

Proof. Zhandry’s scheme discussed in Proposition 2 has the property that $\text{verify-bolt}(s, \cdot)$, is (exponentially close to) a rank-1 projector, and therefore upon one successful verification, it will continue to pass verifications and therefore satisfy Definition 11. In fact, it holds even against unbounded adversaries. This rank-1 projector is such that the state that it accepts only has support on x ’s that are valid certificates. Since the **gen-certificate** algorithm is simply a measurement in the standard basis, we conclude that Definition 12 holds. \square

The Sabotage-Signature experiment is defined in an analogous fashion:

1. The challenger runs $(\text{gen-bolt}, \text{verify-bolt}, \text{sign}, \text{verify-sig}) \leftarrow \text{QL.Setup}(\lambda)$ and sends $(\text{gen-bolt}, \text{verify-bolt}, \text{sign}, \text{verify-sig})$ to \mathcal{A} .
2. \mathcal{A} sends $|\psi\rangle$, and (a document to be signed) α to the challenger.
3. The challenger runs $\text{verify-bolt}(|\psi\rangle)$. If verification fails, set $r = 1$. If verification accepts the challenger uses the post-measured state $|\psi'\rangle$ to generate a signature of α : $\sigma \leftarrow \text{sign}(|\psi'\rangle, s, \alpha)$. The challenger runs $r \leftarrow \text{verify-sig}(s, \sigma, \alpha)$.
4. \mathcal{A} wins if $r = 0$. Let $\text{Sabotage-Signature}(\mathcal{A}, \lambda)$ be the random variable that is 1 if the adversary \mathcal{A} wins, and is 0 otherwise.

Definition 13. *A quantum lightning scheme with a bolt-to-signature capability is secure against sabotage if, in addition to the requirement in Definition 11, for every QPT \mathcal{A} there exists a negligible function negl such that:*

$$\Pr(\text{Sabotage-Signature}(\mathcal{A}, \lambda) = 1) = \text{negl}(\lambda)$$

Proposition 4. *The construction QLDS in Fig. 8 is secure against sabotage, assuming the underlying quantum lightning scheme with a bolt-to-certificate capability QLC which it uses is secure against sabotage.*

Proof. Given a QPT adversary \mathcal{A} that wins the Sabotage-Signature experiment with probability ϵ with respect to the QLDS scheme, we can construct an adversary \mathcal{B} that wins the Sabotage-Certificate experiment with probability $\frac{\epsilon(\lambda)}{2n}$. By our assumption that QLC is secure against sabotage, we conclude that $\epsilon(\lambda)$ is necessarily a negligible function.

The adversary \mathcal{B} receives from his challenger $(\text{Gen}, \text{verify-bolt}, \text{gen-certificate}, \text{verify-certificate})$. \mathcal{B} will sample $H : \{0, 1\}^* \rightarrow \{0, 1\}^n \leftarrow \mathcal{H}^{1^n}$. \mathcal{B} will play the role of the challenger in the Sabotage-Signature experiment, and will also simulate \mathcal{A} . \mathcal{A} will receive the tuple above and H as part of the setup, and will produce a state $|\psi\rangle$ over $2n$ registers, serial number s_1, \dots, s_n a document α . \mathcal{B} will sample $i \in [2n]$ uniformly at random, and send the i ’th register of $|\psi\rangle$ to his challenger.

Next we show that indeed \mathcal{B} wins with probability at least $\frac{\epsilon}{2n}$.

We know that with probability $\epsilon(\lambda)$, \mathcal{A} ’s challenger will verify all the $2n$ states, generate a certificate from n of these states, and at least one of these certificates will verification. Recall that i was sampled uniformly at random. \mathcal{B} ’s challenger performs exactly the same procedure – the challenger generates a certificate from the state it receives and checks whether it passes the verification as a certificate. The probability that i is one of the failed certificates is therefore at least $\frac{\epsilon}{2n}$. \square

6.4 A resolution of the practical issues

We first employ a quantum lightning scheme with bolt-to-signature capability to resolve issues (i) and (ii) of Section 6.1.

We make a simple modification to the protocols of Fig. 6 (preventing and challenging malicious “lost banknote claims”) and Fig. 7 (trading quantum banknotes for coins).

We upgrade the quantum lightning scheme with bolt-to-certificate capability to one with bolt-to-signature capability. To deal with issue (i), we make two modifications to our payment system of Section 4.

- We modify the protocol of Fig. 6 as follows: when party P notices a malicious “lost banknote claim” for a serial number s associated to a quantum banknote $|\psi\rangle$ that he possesses, he does not simply compute the classical certificate associated to s and send it in clear to \mathcal{F}_{Ledg} . Rather, P computes a signature $\sigma \leftarrow \text{gen-sig}(|\psi\rangle, s, \alpha)$ where α is a message saying “Party P challenges the claim”.
- We modify the definition of banknote-contract (Definition 7) so that whenever there is an active lost claim, the coins deposited in a contract with state variable `serial` = s are released to a party P only upon receipt of a signature σ with respect to s of a message “Party P challenges the claim”.

We do not give a formal proof of security of the modified scheme, as this would require first defining formally the modified ledger functionality. Instead we argue informally: it is straightforward to see that the attack described in point (i) of Section 6.1 is impossible. Any adversary that is able to carry out that attack could be used to create an adversary that is able to forge a signature with respect to a serial number s of a bolt that they do not possess. This violates the security of the bolt-to-signature property of the lightning scheme.

To deal with issue (ii), we make a similar simple modification:

- We modify the protocol of Fig. 7 as follows: in order to trade a valid quantum banknote with serial number s for the coins deposited in a smart-contract with state variable `serial` = s , party P does not simply compute the classical certificate associated to s and send it in clear to \mathcal{F}_{Ledg} . Instead, P computes a signature $\sigma \leftarrow \text{gen-sig}(|\psi\rangle, s, \alpha)$ where α is a message saying “Party P wishes to recover the coins in the contract”.
- We modify the definition of banknote-contract (Definition 7) so that whenever there is no active claim, the coins deposited in a contract with state variable `serial` = s are released to a party P only upon receipt of a signature σ with respect to s of a message “Party P wishes to recover the coins in the contract”.

For a similar reasoning as for point (i), the attack in point (ii) of Section 6.1 is no longer possible.

Dealing with issue (iii) of Section 6.1 is trickier. In this case, an honest party P has lost a valid quantum banknote with serial number s , and there is no way for P to recover any certificate or signature proving possession of the banknote. The only difference between P and everyone else, as far as owning the coins deposited in the associated contract, is that P knows that the banknote has been damaged and lost, and no one else does. The “lost banknote claim” protocol of Fig. 5 requires P to send a message to \mathcal{F}_{Ledg} declaring the loss, and the only reason why P is able to recover the coins deposited in the contract in the idealized setting is that he is the first to make this claim, and no one has the ability to challenge it. The situation changes dramatically if we allow the adversary to delay the processing of honest parties’ messages to \mathcal{F}_{Ledg} in favour of its own. The adversary could simply notice a “lost banknote claim” and take advantage of this by making its own claim and ensuring that it is registered first on the ledger. We propose the following modification to handle this issue:

- Instead of directly making a “lost banknote claim”, a party P posts a commitment to a message of the form “ P is filing a lost banknote claim associated to the smart contract with identifier `ssid`”. The commitment contains a deposit of d_0 coins.
- P has to reveal that commitment after no more than t_0 blocks.

- The coins are released to user P only t_1 blocks after the reveal phase, and provided no reclaim was made during that time.
- In case there are two or more reveals to two or more commitments to the same *ssid* – the one which was committed to the earliest (i.e., the commitment appears in an earlier block) receives the coins.

Intuitively, this modification resolves issue (iii). This is because the commitment is hiding, and hence the adversary does not learn anything about the claim it contains prior to the reveal. After the reveal phase, it is too late for the adversary to start the commitment phase. On the other hand, if an adversary simply tries to guess what the content of a claim is (i.e. which quantum banknote was lost), and tries to make a claim of his own, the adversary will most likely lose coins, assuming the frequency at which “lost banknote claims” are made is low (recall that making a claim requires staking some coins, which are lost if the claim is successfully challenged).

As much as this resolution to issue (iii) seems to work in theory, the adversary could in practice possess some side information regarding the claim that is hidden in the commitment, and it becomes difficult to model this side information in a way that is both rigorous and faithful to what is possible in practice. We illustrate this with some practical examples.

First, it might be hard to know, for an honest user who lost their quantum banknote, whether the latter was indeed lost, or whether it was stolen. Therefore, an honest user who wrongfully believes that his quantum banknote was lost might end up losing an extra d_0 coins, to the thief who stole it, when trying to recover it.

Additionally, if an adversary could guess a serial number of a quantum money state that was or will be lost, and post it before the honest user, the adversary will be the one who will eventually receive these coins. Such a setting might be plausible, for example, during (or even before) power outages.

Another attack vector is the following. Indeed, the commitment does not reveal which quantum banknote was lost. However, each commitment reveals that 1 banknote was lost. Suppose there is a claim for 1743 coins at some time t . There is a good chance that these 1743 coins belong to one person who lost all his 1743 coins. An adversary that can figure out who owns exactly 1743 coins that were recently lost, and what are the serial numbers of these coins, can effectively steal these coins. Since our construction does not claim any guarantees about the privacy of users, this information might be readily available to the adversary. Various attacks on the users’ privacy are known in the “classical” Bitcoin literature – see Item 6 in Section 8.

Therefore, we elect to stir away from formal security claims, and we leave this as a proposed resolution which requires further investigation.

7 A practical implementation on Bitcoin and space optimization

In this section, we informally discuss a practical implementation of our payment system using Bitcoin. In Section 4, we have elected to describe our payment system in a model with access to a ledger functionality that supports stateful universal smart contracts. This was for two main reasons: clarity of exposition and flexibility of the framework. Nevertheless, it is possible to implement our payment system on the Bitcoin blockchain.

Bitcoin uses a special purpose scripting language for transactions. The basic building block of a script is called an *opcode*. The opcodes that Bitcoin supports are extremely limited. For example, the opcode `OP_ADD`, which adds two inputs, is supported, but even a simple operation such as multiplication or raising a number to a power are *not* supported – as they are not strictly needed, yet they increase the attack surface (for example, they may be used to preform a memory-based attack). More information about the scripting language and opcodes can be found in Ref. [NBF⁺16]. The only required adjustment to the Bitcoin protocol needed to implement our payment system is to add two opcodes

to the Bitcoin scripting language. These opcodes are utilized once when transforming Bitcoin to quantum banknotes, and once when the quantum banknotes are transformed back to Bitcoin. We provide more detail about this implementation by describing a possible improvement to the space efficiency of our payment system. In Section 7.1, we informally describe an implementation on Bitcoin of the original payment system. In Section 7.2, we informally describe a modification that drastically improves space efficiency.

7.1 Bitcoin to Quantum Money: The Simple Approach

In order to use a quantum lightning scheme, of course, we need to first run $\text{QL.Setup}(1^\lambda)$. In some scenarios, where we can assume some trusted setup, this is not an issue. But in an existing distributed system, such as Bitcoin, this could be considered contentious. One of Zhandry’s construction only requires a Common Random String (CRS) – see Fact 4.

In the context of Bitcoin, such a CRS could be generated by a multiple-source randomness extractor, applied on the headers of the first n blocks, where n is determined by the amount of randomness required for the quantum lightning scheme, and the min-entropy each header contains. More details regarding using Bitcoin as a source of randomness can be found in Ref. [BCG15]³. Bonneau et al. have argued that the min-entropy in each block header is at least 32 bits [BCG15, Table 1].

To transform x bitcoins associated to a signing key sk , to quantum money, the user acts as follows. First, the user mints a quantum lightning state, together with some serial number: $(|\$\rangle, s) \leftarrow \text{mint}_{pk}$. Second, the user signs the concatenation⁴ of a special “quantum money” opcode, the serial number, and the value y using the secret key:

$$\sigma \leftarrow \text{sign}_{sk}(\text{OP_BITCOIN_TO_QUANTUM_MONEY}||s||y).$$

Third, the user propagates the signed message $(\text{OP_BITCOIN_TO_QUANTUM_MONEY}||s||y, \sigma)$ to the Bitcoin miners. Here y is the value which the quantum money holds, and $x - y$ represents the fee which incentivizes miners to include that message in the blockchain, in line with the current fee mechanism in Bitcoin (which we have not covered in this work – see [NBF⁺16] for more details).

A miner would include the message $(\text{OP_BITCOIN_TO_QUANTUM_MONEY}||s||y, \sigma)$ in their block, as long as (i) the signature is legitimate, i.e., $\text{verify}_{vk}(\text{OP_BITCOIN_TO_QUANTUM_MONEY}||s||y, \sigma) = 1$, (ii) $y < x$, where x is the value of the bitcoins associated with the verification key (bitcoin address) vk , and (iii) as long as the miner fee is attractive enough.

Under the above conditions, eventually, a block which contains that transaction would be mined, verified by all other nodes, and would become part of the longest chain.

At this point, the original owner of the x bitcoins does not possess any “standard” bitcoins, since they are considered spent. Instead, she holds the quantum state $(|\$\rangle, s)$ with the “value” y . She could then spend the quantum money by sending it to other others, and it would be treated as having the same value as y bitcoins. To verify the quantum money $(|\psi\rangle, s, y)$, the receiver would first check that $\text{OP_BITCOIN_TO_QUANTUM_MONEY}||s||y$ indeed appears in the blockchain. This step requires storing the blockchain. Any old version that contains the signed message will do. Techniques such as Simplified Payment Verification (SPV) could be used to eliminate the need to store that information, with some downsides in terms of privacy and security [Nak08, NBF⁺16]. Then, the receiver would accept the quantum money if $\text{verify}_{pk}(|\psi\rangle, s) = 1$. The receiver could in turn spend the received quantum money to other users via the same procedure.

³Though, for our purposes, we do not need a recurring form of randomness usually called a randomness beacon, which is the focus of their work. In this context the source of randomness is only used once.

⁴We use $||$ to denote concatenation.

7.2 Bitcoin to Quantum Money: Space Optimization

The approach mentioned above has a limitation. Unlike bitcoin, which is divisible almost infinitely, quantum money is not (see Item 13 in p. 38 for more details). The most naive workaround is to divide the value of the x bitcoins between several quantum money states, with values y_1, \dots, y_n so that $\sum_{i \in [n]} y_i < x$. The disadvantage is in terms of space: each serial number is recorded on the blockchain, which causes high fees and somewhat of a scalability problem.

We now present a much more efficient approach, in which the space required on the blockchain is independent of the number of quantum money states the user creates. Suppose the user wants to split the x bitcoins into 2^n quantum money states, each having a value of $2^{-n}y$. The user first creates 2^n quantum money states with serial numbers s_1, \dots, s_n , and calculates the Merkle Hash Tree [Mer80]⁵ of these serial numbers. The user then signs and publishes only the root of the Merkle tree r , the total value y (which must satisfy $y < x$ as before), and n . Each quantum money state now consists of three parts: the quantum part $|\$ \rangle$, the serial number s_i , and a Merkle path from s_i to the root r . Note that the information recorded on the block-chain is independent of n , and one could create very small denominations using this approach.

Another advantage of this approach is that it allows a slightly weak variant of a Proof of Payment – see Item 17 in p. 39.

One could even take this one step further. Miners could include a Merkle root of the Merkle tree of all existing serial numbers, at the end of every calendar year. This would require users that wish to verify quantum money to come online once every year to store the Merkle root. Every owner of quantum money would calculate the Merkle path from the serial number of the quantum money state, to that root, and append it to their quantum money state. That would allow these users to verify the quantum money state generated in previous years using only one hash, typically, e.g., 256 bits. Of course, this technique does not reduce the space requirement for quantum money which was generated in a year which has not ended yet.

8 Comparing our scheme to classical alternatives

In this section, we discuss several trade-offs between Bitcoin and our payment system based on quantum money. The current roadmap for solving the scalability problem of Bitcoin goes through a second layer solution called the Lightning Network (LN)⁶. This section discusses the trade-offs between the following three different alternatives of employing Bitcoin: (a) a standard Bitcoin transaction, (b) a LN transaction and (c) a quantum money transaction as in our payment system (we will refer to this as QM for short, from now on). Before doing so, we give a brief overview of the LN.

The LN improves upon Bitcoin, and provides a way to transact off-chain using *bidirectional payment channels* between two parties [PD16] (see also [BDW18]). The opening and closing of the channel is done on-chain (i.e., both of these operations require 1 transaction that would appear on the Bitcoin blockchain) while updating the balance of the channel is done completely off-chain, and thus allows for an unlimited number of transactions which do not affect the limited throughput of the Bitcoin blockchain, and do not incur the long confirmation times of on-chain transactions. At any point, each of the two parties involved can “close” the channel, by submitting the most up to date state of the channel to the blockchain. Crucially, the lightning network supports routing. What this means is that the graph induced by all bi-directional payment channels allows Alice to pay Bob, without any on-chain transaction, as long there is a path in the graph from Alice to Bob. The

⁵We note that the original motivation of Merkle was very different than the one we use here. See [KL14, Section 5.6.2] and [NBF⁺16, Section 1.2] for the specific application we need.

⁶Note that there is no connection between the terms quantum lightning and the lightning network

objectives of the LN are to increase the effective transaction throughput of the Bitcoin blockchain (by having most of the transactions happen off-chain), make transactions be effective immediately, and reduce the transaction costs. Since early 2018, the LN has been active on the Bitcoin main-net. The capacity of the LN is shown in Fig. 9, and as of August 2019, it stores only 0.004% of the total bitcoins in circulation, and is still considered experimental⁷.

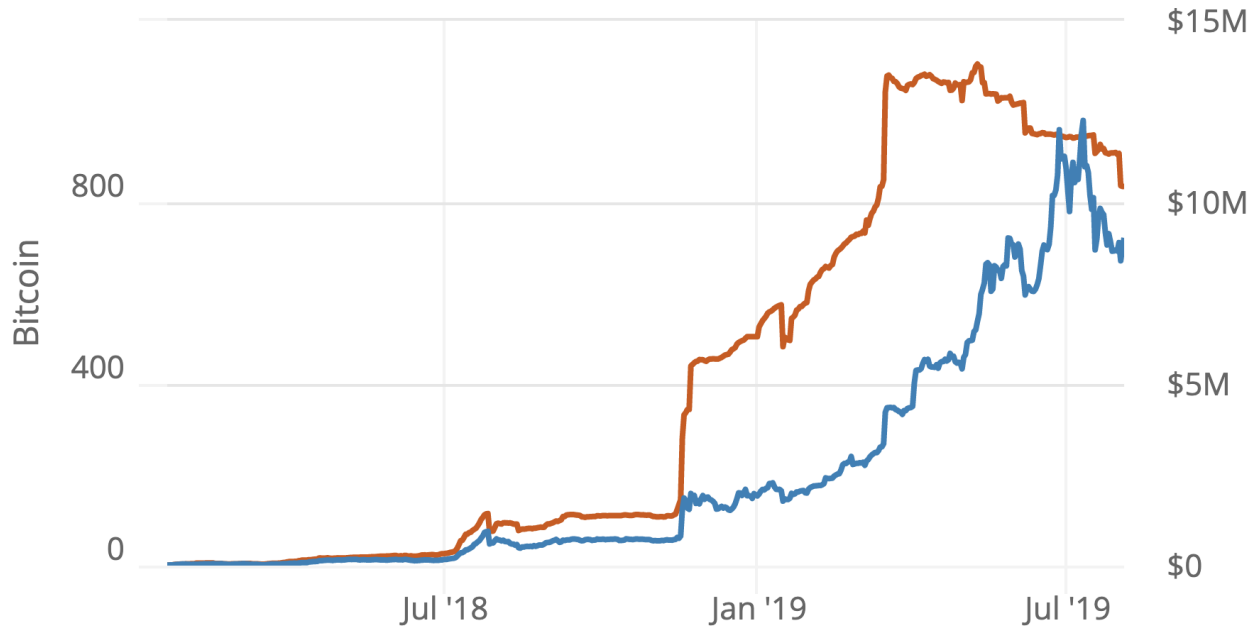


Figure 9: The total capacity of the lightning network, starting on January 2018, brown in BTC, blue in USD. Source: <https://bitcoinvisuals.com/ln-capacity>

The following list provides an exhaustive comparison between the three alternative modes of operation mentioned above. It is worth noting that quantum money, used as in our payment system, in many ways resembles (digital) cash banknotes, both in terms of the advantages and the disadvantages, and is thus the closest of the three modes to ideal digital cash. Items 1-9 present the aspects in which quantum money outperform Bitcoin and the LN, while the disadvantages of quantum money are presented in Items 10-17.

1. **Throughput.** The throughput of Bitcoin is strongly capped. On average, in the first 6 months of 2019, the Bitcoin network had 3.9 transaction per second⁸. This throughput cannot increase dramatically without changing the Bitcoin protocol.

To receive bitcoins through the LN, Alice must have an open channel, which she will eventually need to close. This requires at least two on-chain transaction, though, the number of uses is not bounded. In this regard, transforming quantum money to Bitcoin is very similar to opening a quantum channel, and transforming the quantum money back to Bitcoin is similar to closing a channel.

The balance of a LN channel can be updated, but is always bounded by the transaction that opened that channel. For example, if Alice locks 10 bitcoins in a channel with Bob, then initially she has 10 bitcoins and Bob has zero; Alice and Bob could then update it, but Bob could never receive more than 10 bitcoins using the channel.

A user could receive and send quantum money without ever making an on-chain transaction. Quantum money has no limit on the value transferred.

⁷For example, as of August 2019, the maximal capacity of a LN channel is 0.16 bitcoin, see https://github.com/lightningnetwork/lightning-rfc/blob/master/02-peer-protocol.md#the-open_channel-message, and the requirements regarding the parameter `funding_satoshis`.

⁸Source: <https://www.blockchain.com/charts/n-transactions>.

2. **Liquidity.** Suppose Alice wants to pay Bob. In Bitcoin, she could do that, assuming she has enough funds, and connection to the Bitcoin network. In the LN, this is not always the case: there needs to be a way to route money among the open channels. Sometimes no such route exists, or is inaccessible (using these channels for routing requires the cooperation of the parties along the channel). This may have an impact on the effective throughput of the LN. A quantum money state can always be sent from the receiver to the sender.
3. **Latency.** The recommended best practice for a Bitcoin transaction is waiting for 6 confirmations, which takes 1 hour on average.
Both the LN and QM need one transaction – in order to open a LN channel, or to transform bitcoin to QM. That part can be done in advance, and is only needed once, but it suffers from the same latency as a Bitcoin transaction. A LN has no inherent latency other than the delays caused by the routing. For example, a single transaction might involve several hops between distant nodes, which takes on the order of a second. QM has a slightly better performance, especially when the sender and receiver are physically close to each other – the latency in this case is only due to the quantum verification algorithm. Overall, the LN and QM have comparable latencies, which are much better than Bitcoin’s.
4. **Fees.** Each Bitcoin transaction needs to pay a *fee*, in order to get approved. This fee is collected by the miner, which included the transaction. The average fee per transaction in the first 6 months of 2019 was 1.41 USD⁹. To encourage LN nodes to provide liquidity, the protocol uses routing fees. These fees are expected to be smaller than on-chain transaction fees, but still non-zero. No fees are needed to transact with QM.
5. **Dependence on mining.** The Bitcoin protocol and implicitly, the LN, are based on mining (also known as Proof-of-Work [DN92]). This approach suffers from two main drawbacks: (a) As of August 2018, Bitcoin mining (also known as Proof-of-Work [DN92]) consumed slightly under 1% of the world’s electricity [Nar18]. Mining is required to secure the network. Proof-of-Stake is a competing approach with somewhat different trade-offs [KRDO17, DGKR18], with the main advantage that it does not spend so much energy. (b) The security model related to mining is convoluted, and is based on incentives, without clear assumptions nor a full security analysis. In particular, it is known that the Bitcoin protocol is *not* incentive compatible [ES14]. In PoW and PoS, a corrupted majority can double-spend. In addition, the Bitcoin protocol does not provide finality – in principle, at any point, a roll-back (known as re-organization) could occur.
Quantum money does not require PoW or PoS, and provides finality. For example, if all the bitcoins were transformed to quantum money, PoW or PoS would not be needed at all. Of course, coordinating such a change is highly non-trivial.
6. **Privacy.** The users’ privacy is not guaranteed in Bitcoin, and certainly not by default. All the transactions that have ever occurred are publicly accessible. The common approach to provide some level of privacy is to use a new bitcoin address per transaction. Yet, there are various techniques to relate these addresses to real-world identities [RS13, RS14, MPJ⁺16]. There are several commercial services that deanonymize Bitcoin transactions¹⁰. More recent works tackle this privacy issue – see [vS13, MGGR13, BCG⁺14, Poe16] and references therein.
There is a trade-off between privacy and concurrency in the LN, see [MMK⁺17] and

⁹Sources: <https://www.blockchain.com/charts/n-transactions> and <https://www.blockchain.com/charts/transaction-fees-usd>.

¹⁰E.g., <https://www.chainalysis.com>. They claim to have “the biggest database in the world of connections between real world entities and transactions on the blockchain”, see <https://youtu.be/yNpNz-FvSYQ?t=154>.

references therein. In addition, this aforementioned work explicitly leaves the question of privacy preserving routing for the LN open.

Quantum money enjoys superior privacy, as only the sender and receiver take part in the transaction, and the transaction leaves no trace. The privacy of QM is analogous to that of physical banknotes. Bear in mind that banknotes have serial numbers, and these could potentially be traced, although arguably the effect on privacy is negligible. In this sense, coins provide better privacy for the users, since they are indistinguishable. *Quantum Coins* have been formally studied: these are indistinguishable quantum states which provide an analogous level of privacy as physical coins - see [MS10, JLS18] (improving upon [TOI03]). Unfortunately, these constructions are for private quantum money (rather than public), and do not constitute a quantum lightning scheme, which is crucial for our construction, and thus cannot be used in our payment system. We leave it as an open question whether the same level of anonymity achieved using quantum coins, or the construction mentioned above, could be achieved in the context of quantum money for Bitcoin.

7. **Risk of Locked funds.** In the LN, the parties lock funds in a channel. When both parties are cooperative, these funds can be easily unlocked, and used in the next block. In case one party is uncooperative, and refuses to close the channel, these funds are effectively locked for some period of time, typically, for a few days.

Quantum money does not require users to lock any funds.

8. **Connectivity requirements.** A Bitcoin transaction requires communication with the Bitcoin network. At the very least, the receiver needs a communication channel with at least one other Bitcoin node. A LN transaction needs even more resources, since the LN uses source routing – therefore, the sender has to be well connected to the entire network in order to find the route. A QM transaction only requires quantum communication between the sender and the receiver.
9. **Liveliness.** For technical reasons which are outside the scope of this work, both parties participating in each channel have to be on-line occasionally, and monitor the Bitcoin network, in order to revoke a transaction, in case the other party cheats. This task can be outsourced to a *Watchtower*, which has to be trusted to perform its job (and of course the watch-tower has to be on-line occasionally). Online availability is not required for QM, if one opts for a version of our payment system in which there is no procedure to recover the value corresponding to lost or damaged quantum banknotes. If one opts to include such a recovery mechanism, then a similar level of online availability as in the LN is required.
10. **Technological requirements.** Transactions with quantum money require 3 main resources: long term quantum memory to store the quantum money, a universal quantum computer to verify the quantum money, and a quantum network to transmit the quantum money between users.
11. **Smart contracts.** Bitcoin provides a scripting language which can be used to design *smart contracts*. Notable examples include *multi-sig transactions*, in which m -out-of- n signatures are needed to spend the money, and *atomic cross-chain swap* which allows Alice and Bob to trade two crypto-currencies without trusting each other [NBF⁺16]. Both capabilities naturally extend to the LN¹¹. Other crypto-currencies such as Ethereum have a richer programming language, which allows constructing Decentralized Applications (DApps) [Woo14, But14].
QM does not solve the consensus problem or any variant of it, and does not provide the functionality of smart contracts.

12. **Backup.** It is pretty easy to backup a Bitcoin wallet. Typically, all that a user needs is a fairly short string called a *seed*. This seed is used to generate a Hierarchical

¹¹The video in <https://youtu.be/cBVcgzEuJ7Q> demonstrates a LN atomic swap between Bitcoin and Litecoin.

Deterministic Wallet [Wui13]. A backup can be done once, and never needs to be updated in the lifetime of a Bitcoin wallet. LN channels are slightly harder to backup, since the protocol is stateful, and therefore, currently, backing up requires having the most up-to-date state of the channel.

By definition, it is impossible to backup a quantum money state. In this work, we proposed a mechanism to *recover* lost banknotes. In essence, a user can claim that her quantum money was lost, by depositing d bitcoins. The user would receive these bitcoins after some period of time (e.g., one year). If the party that claimed that the coins were lost is dishonest, then the person holding the legitimate quantum money state can produce a certificate of this fact, and claim the d bitcoins, in addition to the value of the quantum money that she originally held. To avoid theft, users that want this option available have to be on-line occasionally (in this example, at least once a year).

13. **Divisibility.** One of the advantages of Bitcoin and the LN is that any amount, down to 10^{-8} bitcoin can be sent, and in principle, even smaller amounts could be used. Quantum money, on the other hand, is *not* divisible. The user must decide, at the time of the quantum minting, what is the denomination of the quantum money, and it remains the same for the lifetime of the quantum money.

14. **Hidden inflation.** Consider a computationally powerful adversary that attacks the Bitcoin network. Such an adversary could, for example, break the digital signature scheme and steal other people’s Bitcoins. Yet, such an adversary couldn’t “print” bitcoin from thin air, without others noticing it¹².

When we use quantum money, the situation is different. A powerful adversary could create new quantum money from thin air, without others being able to notice it. This might be a threat either because of invalid computational assumptions, or flaws in the implementation. This threat is not unique to quantum money. In fact, such a flaw in the implementation occurred in ZCash, a crypto-currency which is based on the ZeroCash protocol [BCG⁺14], see <https://electriccoin.co/blog/zcash-counterfeiting-vulnerability-successfully-remediated>. Inevitably, there is no definitive way to know whether that bug was exploited.

15. **Finality of the security parameters.** One interesting feature of Bitcoin is that the level of security that is achieved can, in principle, be increased. If the level of security seems insufficient due to technological advancement, the protocol may allow users to transition to more secure schemes. This is exactly the case for the proposed post-quantum secure digital signature schemes [But13, LRS18]. It is the responsibility (and incentive) of each individual user to transition – otherwise, her funds might be lost.

The quantum money can be used too increase security in the same manner: users with QM in circulation can create new QM with the improved parameters, sign the new serial number using their the bolt-to-signature capability (and by that destroying it), and adding that signed message to the blockchain. Yet, the incentives here a slightly different: an adversary could steal the bitcoins of some user, if the security parameters are poorly chosen. In the quantum money setting, the adversary could print money from thin air – see Item 14. This makes the system, as a whole, insecure. Therefore, it is advisable to make such a transition mandatory.

16. **Optional Transparency.** Bitcoin transactions are publicly available, and organizations or individuals that wish to, can make their accounting book completely transparent. See Bitcoin Improvement Proposal (BIP) 32 [Wui13] and Ref. [NBF⁺16, Chapter 4.2] for more details.

Quantum money transactions leave no trace, and therefore it seems impossible to achieve this sort of transparency.

¹²Unless there is some bug that is unrelated to cryptography

17. **Proof of payment.** Consider the following scenario. You go to a store, and pay for an item. You hand over a valid banknote to the seller. The seller takes it to the bill checking machine, and secretly replaces your valid bill with a fake one. The seller then gives you back the fake money, and blames you for trying to fraud him.

This kind of attack cannot happen in Bitcoin, if used appropriately. The seller can ask for a signed payment request, and the buyer can then verify the authenticity of that message, using the seller’s public key. After payment, the seller cannot argue that the payment was not received – the buyer can prove that the bitcoins were sent to the seller’s address, by showing the payment on the blockchain.

A Bitcoin payment can be done in such a way that an honest user would have a proof of payment [AH13]. A similar functionality might be possible to achieve in the LN, though currently, as far as the authors are aware, the LN does not provide such functionality.

On the other hand, QM transactions leaves no trace, and proof of payment seems harder to achieve. A possible workaround (which works for the LN as well) could be the following. Suppose Alice wants to send 10 bitcoins worth of quantum money to Bob. Instead of sending it all at a time, she could divide the payment into 100 iterations. In each iteration she would send 0.1 bitcoin, and expect a digital signature in return approving the payment in return. If Bob fails to provide such a signature, she would abort. The worst case scenario in this case is that she would not have a proof of payment for 0.1 bitcoins.

9 Conclusion

In this work, we gave the first example of the use of classical smart contracts in conjunction with quantum cryptographic tools. 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 [Zha19]. 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.

Acknowledgments A.C is supported by the Simons Institute for the Theory of Computing. O.S. is supported by the Israel Science Foundation (ISF) grant No. 682/18 and 2137/19, and by the the Cyber Security Research Center at Ben-Gurion University.

References

- [Aar09] S. Aaronson. Quantum copy-protection and quantum money. In *Computational Complexity, 2009. CCC’09. 24th Annual IEEE Conference on*, pages 229–242. IEEE, 2009.
- [AH13] G. Andresen and M. Hearn. Payment Protocol. Bitcoin Improvement Proposal (BIP) 70 <https://github.com/bitcoin/bips/blob/master/bip-0070>, 2013.
- [BCG⁺14] E. Ben-Sasson, A. Chiesa, C. Garman, M. Green, I. Miers, E. Tromer, and M. Virza. *Zerocash: Decentralized Anonymous Payments from Bitcoin*. In

2014 *IEEE Symposium on Security and Privacy, SP 2014, Berkeley, CA, USA, May 18-21, 2014*, pages 459–474. IEEE Computer Society, 2014.

- [BCG15] J. Bonneau, J. Clark, and S. Goldfeder. [On Bitcoin as a public randomness source](#). *IACR Cryptology ePrint Archive*, 2015:1015, 2015.
- [BDW18] C. Burchert, C. Decker, and R. Wattenhofer. [Scalable funding of Bitcoin micropayment channel networks](#). *Royal Society Open Science*, 5(8):180089, August 2018.
- [BK14] I. Bentov and R. Kumaresan. How to use bitcoin to design fair protocols. In *International Cryptology Conference*, pages 421–439. Springer, 2014.
- [BKM17] I. Bentov, R. Kumaresan, and A. Miller. Instantaneous decentralized poker. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 410–440. Springer, 2017.
- [BMTZ17] C. Badertscher, U. Maurer, D. Tschudi, and V. Zikas. Bitcoin as a transaction ledger: A composable treatment. In *Annual International Cryptology Conference*, pages 324–356. Springer, 2017.
- [BOM04] M. Ben-Or and D. Mayers. General security definition and composability for quantum & classical protocols. *arXiv preprint quant-ph/0409062*, 2004.
- [BS16] S. Ben-David and O. Sattath. Quantum Tokens for Digital Signatures, 2016, arXiv: [1609.09047](#).
- [But13] V. Buterin. Bitcoin Is Not Quantum-Safe, And How We Can Fix It When Needed. <https://bitcoinmagazine.com/articles/bitcoin-is-not-quantum-safe-and-how-we-can-fix-1375242150>, <http://www.webcitation.org/6wDiIPU31>, 2013.
- [But14] V. Buterin. A next-generation smart contract and decentralized application platform. <https://github.com/ethereum/wiki/wiki/White-Paper>, 2014.
- [Can01] R. 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] R. 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] R. Canetti, Y. Dodis, R. Pass, and S. Walfish. Universally composable security with global setup. In *Theory of Cryptography Conference*, pages 61–85. Springer, 2007.
- [Col19] A. Coladangelo. Smart contracts meet quantum cryptography, 2019, arXiv: [1902.05214](#).
- [CSV16] R. Canetti, D. Shahaf, and M. Vald. Universally composable authentication and key-exchange with global PKI. In *IACR International Workshop on Public Key Cryptography*, pages 265–296. Springer, 2016.
- [DEF18] S. Dziembowski, L. Eeckey, and S. Faust. Fairswap: How to fairly exchange digital goods. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, pages 967–984. ACM, 2018.
- [DGKR18] B. David, P. Gazi, A. Kiayias, and A. Russell. [Ouroboros Praos: An Adaptively-Secure, Semi-synchronous Proof-of-Stake Blockchain](#). In J. B. Nielsen and V. Rijmen, editors, *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part II*, volume 10821 of *Lecture Notes in Computer Science*, pages 66–98. Springer, 2018.
- [DN92] C. Dwork and M. Naor. [Pricing via Processing or Combatting Junk Mail](#). In *Advances in Cryptology - CRYPTO '92, 12th Annual International Cryptology*

Conference, Santa Barbara, California, USA, August 16-20, 1992, Proceedings, pages 139–147, 1992.

- [ES14] I. Eyal and E. G. Sirer. [Majority Is Not Enough: Bitcoin Mining Is Vulnerable](#). In *Financial Cryptography and Data Security - 18th International Conference, FC 2014, Christ Church, Barbados, March 3-7, 2014, Revised Selected Papers*, pages 436–454, 2014, arXiv: [1311.0243](#).
- [FGH⁺12a] E. Farhi, D. Gosset, A. Hassidim, A. Lutomirski, and P. Shor. Quantum money from knots. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, pages 276–289, 2012.
- [FGH⁺12b] E. Farhi, D. Gosset, A. Hassidim, A. Lutomirski, and P. Shor. [Quantum money from knots](#). In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, pages 276–289. ACM, ACM, 2012, arXiv: [1004.5127](#).
- [Gol04] O. Goldreich. *The Foundations of Cryptography - Vol. 2, Basic Applications*. Cambridge University Press, 2004.
- [JLS18] Z. Ji, Y. Liu, and F. Song. [Pseudorandom Quantum States](#). In H. Shacham and A. Boldyreva, editors, *Advances in Cryptology - CRYPTO 2018 - 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2018, Proceedings, Part III*, volume 10993 of *Lecture Notes in Computer Science*, pages 126–152. Springer, 2018, arXiv: [1711.00385](#).
- [KL14] J. Katz and Y. Lindell. *Introduction to Modern Cryptography, Second Edition*. CRC Press, 2014.
- [KMS⁺16] A. Kosba, A. Miller, E. Shi, Z. Wen, and C. Papamanthou. Hawk: The blockchain model of cryptography and privacy-preserving smart contracts. In *2016 IEEE Symposium on Security and Privacy (SP)*, pages 839–858. IEEE, 2016.
- [KRDO17] A. Kiayias, A. Russell, B. David, and R. Oliynykov. [Ouroboros: A Provably Secure Proof-of-Stake Blockchain Protocol](#). In J. Katz and H. Shacham, editors, *Advances in Cryptology - CRYPTO 2017 - 37th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 20-24, 2017, Proceedings, Part I*, volume 10401 of *Lecture Notes in Computer Science*, pages 357–388. Springer, 2017.
- [KZZ16] A. Kiayias, H.-S. Zhou, and V. 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] A. Lutomirski, S. Aaronson, E. Farhi, D. Gosset, A. Hassidim, J. Kelner, and P. Shor. Breaking and making quantum money: toward a new quantum cryptographic protocol. *arXiv preprint arXiv:0912.3825*, 2009.
- [Lam79] L. Lamport. Constructing digital signatures from a one-way function, 1979.
- [LRS18] T. Lee, M. Ray, and M. Santha. On a quantum Bitcoin mining contest. private communication, 2018.
- [Lut11] A. Lutomirski. Component mixers and a hardness result for counterfeiting quantum money, 2011, arXiv: [1107.0321](#).
- [Mer80] R. C. Merkle. [Protocols for Public Key Cryptosystems](#). In *Proceedings of the 1980 IEEE Symposium on Security and Privacy, Oakland, California, USA, April 14-16, 1980*, pages 122–134. IEEE Computer Society, 1980.
- [MGGR13] I. Miers, C. Garman, M. Green, and A. D. Rubin. [Zerocoin: Anonymous Distributed E-Cash from Bitcoin](#). In *2013 IEEE Symposium on Security and Privacy, SP 2013, Berkeley, CA, USA, May 19-22, 2013*, pages 397–411. IEEE Computer Society, 2013.

- [MMK⁺17] G. Malavolta, P. Moreno-Sanchez, A. Kate, M. Maffei, and S. Ravi. *Concurrency and Privacy with Payment-Channel Networks*. In B. M. Thuraisingham, D. Evans, T. Malkin, and D. Xu, editors, *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS 2017, Dallas, TX, USA, October 30 - November 03, 2017*, pages 455–471. ACM, 2017.
- [MPJ⁺16] S. Meiklejohn, M. Pomarole, G. Jordan, K. Levchenko, D. McCoy, G. M. Voelker, and S. Savage. *A fistful of Bitcoins: characterizing payments among men with no names*. *Commun. ACM*, 59(4):86–93, 2016.
- [MS10] M. Mosca and D. Stebila. *Quantum coins*, volume 523 of *Contemp. Math.*, pages 35–47. Amer. Math. Soc., 2010, arXiv: 0911.1295.
- [Nak08] S. Nakamoto. *Bitcoin: A peer-to-peer electronic cash system*, 2008.
- [Nar18] A. Narayanan. *Hearing on Energy Efficiency of Blockchain and Similar Technologies*, Committee on Energy and Natural Resources, United States Senate, August 2018.
- [NBF⁺16] A. Narayanan, J. Bonneau, E. W. Felten, A. Miller, and S. Goldfeder. *Bitcoin and Cryptocurrency Technologies - A Comprehensive Introduction*. Princeton University Press, 2016.
- [PD16] J. Poon and T. Dryja. *The bitcoin lightning network: Scalable off-chain instant payments*, 2016.
- [Poe16] A. Poelstra. *Mimblewimble*, 2016.
- [RS13] D. Ron and A. Shamir. *Quantitative Analysis of the Full Bitcoin Transaction Graph*. In *Financial Cryptography and Data Security - 17th International Conference, Japan*, pages 6–24, 2013.
- [RS14] D. Ron and A. Shamir. *How Did Dread Pirate Roberts Acquire and Protect his Bitcoin Wealth?* In *Financial Cryptography and Data Security - FC 2014 Workshops, BITCOIN and WAHC 2014, Barbados*, pages 3–15, 2014.
- [Sat19] O. Sattath. *Transforming Bitcoin to Quantum Money, and Back*. Unpublished Manuscript, 2019.
- [TOI03] Y. Tokunaga, T. Okamoto, and N. Imoto. *Anonymous quantum cash*, 2003.
- [Unr04] D. Unruh. *Simulatable security for quantum protocols*. *arXiv preprint quant-ph/0409125*, 2004.
- [Unr10] D. Unruh. *Universally composable quantum multi-party computation*. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 486–505. Springer, 2010.
- [vS13] N. van Saberhagen. *CryptoNote v 2.0*, 2013.
- [WEH18] S. Wehner, D. Elkouss, and R. Hanson. *Quantum internet: A vision for the road ahead*. *Science*, 362(6412):eaam9288, 2018.
- [Wie83] S. Wiesner. *Conjugate coding*. *ACM Sigact News*, 15(1):78–88, 1983.
- [Woo14] G. Wood. *Ethereum: a secure decentralised generalised transaction ledger*. <http://gavwood.com/paper.pdf>, 2014.
- [Wui13] P. Wuille. *Hierarchical deterministic wallets*. Bitcoin Improvement Proposal (BIP) 32 <https://github.com/bitcoin/bips/blob/master/bip-0032>, 2013.
- [Zha19] M. Zhandry. *Quantum Lightning Never Strikes the Same State Twice*. In Y. Ishai and V. Rijmen, editors, *Advances in Cryptology - EUROCRYPT 2019 - 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Darmstadt, Germany, May 19-23, 2019, Proceedings, Part III*, volume 11478 of *Lecture Notes in Computer Science*, pages 408–438. Springer, 2019, arXiv: 1711.02276.

A Appendix

A.1 Proof of Proposition 2

Proof of Proposition 2. We assume some familiarity with Zhandry’s construction (see section 6 of [Zha19] 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 H (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. `verify-bolt` has H hardcoded in its description. The computational assumption under which Zhandry’s construction is proved secure is that H 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).

Similarly to the proof of Proposition 1, we define `gen-certificate` to be the QPT algorithm that measures each mini-bolt in the computational basis and outputs the concatenation of the outcomes. We define `verify-certificate` to be the deterministic algorithm which receives as input a serial number (y_1, \dots, y_n) , and the concatenation of $(z_1^{(i)}, \dots, z_{k+1}^{(i)})$ for $i = 1, \dots, n$, and checks that $H(z_1^{(i)}) = \dots = H(z_{k+1}^{(i)}) = y_i$ for all i .

It is clear that (I) holds.

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 H from an adversary \mathcal{A} that wins game `Forge-certificate` with non-negligible probability: \mathcal{A}' runs $(\text{gen-bolt}, \text{verify-bolt}) \leftarrow \text{QL.Setup}(1^\lambda)$ and sends $(\text{gen-bolt}, \text{verify-bolt}, \text{gen-certificate}, \text{verify-certificate})$ to \mathcal{A} , where the latter two are defined as above in terms of the function H hardcoded in `verify-bolt`. \mathcal{A} returns c 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. We claim that with non-negligible probability \mathcal{A}' outputs a $(2k + 2)$ -collision: in fact, since \mathcal{A} wins `Forge-certificate` with non-negligible probability, then $|\psi'\rangle$ must pass `verify-bolt` 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 most mini-bolts have non-negligible weight on most pre-images (in fact they should be close to uniform superpositions over all pre-images). Indeed, in Section 6.3, p. 435, Zhandry argues:

Conditioned on acceptance, by the above arguments the resulting mini bolts must all be far from singletons when we trace out the other bolts. This means that if we measure the mini-bolts, the resulting superpositions will have high min-entropy.

□

A.2 Standard Security Definitions

Definition 14 (Digital Signature Scheme). *A digital signature scheme consists of 3 PPT algorithms `key-gen`, `sign` and `verify`. The scheme is complete if the following holds. When a document is signed using the private key, the signature is accepted by the verification algorithm using the public key. Formally, for every $\alpha \in \{0, 1\}^*$:*

$$\Pr \left[(vk, sk) \leftarrow \text{key-gen}(1^\lambda); \sigma \leftarrow \text{sign}_{sk}(\alpha); r \leftarrow \text{verify}_{vk}(\alpha, \sigma) : r = 1 \right] = 1$$

The scheme satisfies Post-Quantum Existential Unforgeability under a Chosen Message Attack (PQ-EU-CMA) if the following holds. Consider a QPT adversary with the capability of adaptively requesting documents to be signed by a signing oracle. The scheme is secure if such an adversary cannot generate a signature for any fresh document – a

document which the adversary did not ask the oracle to sign. Formally, for every QPT adversary \mathcal{A} there exists a negligible function negl such that

$$\Pr \left[(vk, sk) \leftarrow \text{key-gen}(1^\lambda); (\alpha, \sigma) \leftarrow \mathcal{A}^{\text{sign}_{sk}}(1^\lambda, vk) : \text{verify}_{vk}(\alpha, \sigma) = 1 \wedge \alpha \notin Q_{\mathcal{A}}^{\text{sign}_{sk}} \right] \leq \text{negl}(\lambda),$$

where $\mathcal{A}^{\text{sign}_{sk}}$ is a QPT algorithm with access to the signing oracle, and $Q_{\mathcal{A}}^{\text{sign}_{sk}}$ is the set of queries it made to the oracle.