

Title: Quantum Money

Date: Jan 27, 2010 02:00 PM

URL: <http://pirsa.org/10010078>

Abstract: Ever since there's been money, there have been people trying to counterfeit it, and governments trying to stop them. In 1969, the physicist Stephen Wiesner raised the remarkable possibility of money whose authenticity would be guaranteed by the laws of quantum mechanics. However, the question of whether one can have secure quantum money that anyone (not only the bank) can verify has remained open for forty years. In this talk, I'll tell you about progress on the question over the last two years.

(1) I'll show that no publicly-verifiable quantum money scheme can have security based on quantum physics alone: like in most cryptography, one also needs a computational hardness assumption.

(2) I'll show that one can have quantum money that remains hard to counterfeit, even if a counterfeiter gains access to a "black box" for verifying the money.

(3) I'll describe a candidate quantum money scheme I proposed last spring, and how that scheme was recently broken by Lutomirski et al. I'll also discuss a new class of schemes that might evade the existing attacks -- schemes with the bizarre property that not even the bank can prepare the same bill twice.

The talk is designed to be accessible to those without a quantum information background.

Reference for (1)-(2): S. Aaronson, "Quantum copy-protection and quantum money," in Proceedings of CCC'2009, <http://www.scottaaronson.com/papers/noclone-ccc.pdf>.

Reference for (3): 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, Proceedings of Innovations in Computer Science (ICS), 2010. <http://arxiv.org/abs/0912.3825>.

Quantum Money



Scott Aaronson (MIT)

Based partly on joint work with Ed Farhi, David Gosset, Avinatan Hassidim, Jon Kelner, Andy Lutomirski, and Peter Shor

Ever since there's been money, there've been people trying to counterfeit it

One of the oldest “security problems” facing human civilization; has to be solved reasonably well before a market economy becomes possible

Ever since there's been money, there've been people trying to counterfeit it

One of the oldest "security problems" facing human civilization; has to be solved reasonably well before a market economy becomes possible


In his capacity as Master of the Mint, Isaac Newton added milled edges to English coins to make them harder to counterfeit



Ever since there's been money, there've been people trying to counterfeit it

One of the oldest "security problems" facing human civilization; has to be solved reasonably well before a market economy becomes possible

In his capacity as Master of the Mint, Isaac Newton added milled edges to English coins to make them harder to counterfeit



Today: Holograms, embedded strips,
“microprinting,” special inks...

Leads to an arms race with no
obvious winner



U.S. Treasury Department via Reuters

Today: Holograms, embedded strips,
“microprinting,” special inks...

Leads to an arms race with no
obvious winner



U.S. Treasury Department via Reuters

Problem: From a computer science perspective,
uncopyable cash seems impossible for trivial reasons!

Today: Holograms, embedded strips, “microprinting,” special inks...

Leads to an arms race with no obvious winner



U.S. Treasury Department via Reuters

Problem: From a computer science perspective, uncopyable cash seems impossible for trivial reasons!

Any printing technology the good guys can build, bad guys can in principle build also

$x \rightarrow (x,x)$ is a polynomial-time operation

What's done in practice: Have a trusted third party (the bank) authorize every transaction



What's done in practice: Have a trusted third party (the bank) authorize every transaction



OK, but there are some cases where you want the convenience, privacy, and anonymity of cash, and it seems you can never make **cash** cryptographically secure






Uncertainty Principle: You can measure a particle's position, or its momentum, but not both to unlimited precision



Uncertainty Principle: You can measure a particle's position, or its momentum, but not both to unlimited precision

Logical consequence: **No-Cloning Theorem**




First Idea in the History of Quantum Info

Wiesner 1969: Money that's impossible to counterfeit, assuming only the validity of quantum mechanics

First Idea in the History of Quantum Info


Wiesner 1969: Money that's impossible to counterfeit, assuming only the validity of quantum mechanics



Each bill includes a few hundred qubits (say electrons), secretly polarized in one of four random directions

First Idea in the History of Quantum Info


Wiesner 1969: Money that's impossible to counterfeit, assuming only the validity of quantum mechanics



Each bill includes a few hundred qubits (say electrons), secretly polarized in one of four random directions

In a giant database, the bank remembers how it polarized every electron on every bill

Want to verify a bill? Take it to the bank. Bank uses its knowledge of the polarizations to measure each electron in the appropriate basis:



Theorem: A counterfeiter who doesn't know a bill's state can copy it with probability at most $(5/6)^n$
(where n is the number of electrons per bill)

Theorem: A counterfeiter who doesn't know a bill's state can copy it with probability at most $(5/6)^n$
(where n is the number of electrons per bill)

Drawbacks of Wiesner's scheme?

1. Need to keep bills from decohering in your wallet!

Theorem: A counterfeiter who doesn't know a bill's state can copy it with probability at most $(5/6)^n$
(where n is the number of electrons per bill)

Drawbacks of Wiesner's scheme?

1. Need to keep bills from decohering in your wallet!
2. Bank needs to maintain a giant polarization database

Solution (Bennett et al. '82): Pseudorandom functions

Theorem: A counterfeiter who doesn't know a bill's state can copy it with probability at most $(5/6)^n$
(where n is the number of electrons per bill)

Drawbacks of Wiesner's scheme?

1. Need to keep bills from decohering in your wallet!
2. Bank needs to maintain a giant polarization database

Solution (Bennett et al. '82): Pseudorandom functions

3. Only the bank knows how to authenticate the bills
No analogue of a convenience-store clerk holding up a bill to the light

Theorem: A counterfeiter who doesn't know a bill's state can copy it with probability at most $(5/6)^n$
(where n is the number of electrons per bill)

Drawbacks of Wiesner's scheme?

1. Need to keep bills from decohering in your wallet!
2. Bank needs to maintain a giant polarization database

Solution (Bennett et al. '82): Pseudorandom functions

3. Only the bank knows how to authenticate the bills
No analogue of a convenience-store clerk holding up a bill to the light

Which brings us to...

Public-Key Quantum Money

(Secure Quantum Money That Anyone Can Authenticate)

Which brings us to...

Public-Key Quantum Money

(Secure Quantum Money That Anyone Can Authenticate)

Overview of Results

[A., CCC 2009]

Public-key quantum money requires computational assumptions

Which brings us to...

Public-Key Quantum Money

(Secure Quantum Money That Anyone Can Authenticate)

Overview of Results

[A., CCC 2009]

Public-key quantum money requires computational assumptions

Which brings us to...

Public-Key Quantum Money

(Secure Quantum Money That Anyone Can Authenticate)

Overview of Results

[A., CCC 2009]

Public-key quantum money requires computational assumptions

Secure public-key quantum money is possible, **if** counterfeiters only have black-box access to checking device

(Already nontrivial: “Complexity-Theoretic No-Cloning Theorem”)

[AFGHKLS, submitted, 2009]

[AFGHKLS, submitted, 2009]

Break of Aaronson's scheme

Possible new candidate schemes, where not even the bank can duplicate a bill

[AFGHKLS, submitted, 2009]

Break of Aaronson's scheme

Possible new candidate schemes, where not even the bank can duplicate a bill

(Security assumption: These schemes can't be broken)

Related task [A., CCC'09]:

Quantum software copy-protection

[AFGHKLS, submitted, 2009]

Break of Aaronson's scheme

Possible new candidate schemes, where not even the bank can duplicate a bill

(Security assumption: These schemes can't be broken)

Related task [A., CCC'09]:

Quantum software copy-protection

"Generic" copy-protection secure against black-box adversaries

Explicit candidate schemes for copy-protecting the family of **point functions**

[AFGHKLS, submitted, 2009]

Break of Aaronson's scheme

Possible new candidate schemes, where not even the bank can duplicate a bill

(Security assumption: These schemes can't be broken)

Related task [A., CCC'09]:

Quantum software copy-protection

"Generic" copy-protection secure against black-box adversaries

Explicit candidate schemes for copy-protecting the family of **point functions**

Definition of Quantum Money Schemes

- n:** Security parameter (all computations should be polynomial in n)
- B:** Poly-size quantum circuit (the “bank”), which maps a secret key $s \in \{0,1\}^n$ to a public key e_s and quantum banknote ρ_s
- A:** Poly-size quantum circuit (the “authenticator”), which takes (e, ρ) as input and either accepts or rejects

Definition of Quantum Money Schemes

n: Security parameter (all computations should be polynomial in n)

B: Poly-size quantum circuit (the “bank”), which maps a secret key $s \in \{0,1\}^n$ to a public key e_s and quantum banknote ρ_s

A: Poly-size quantum circuit (the “authenticator”), which takes (e, ρ) as input and either accepts or rejects

(B, A) has *completeness error* ε if for every s ,

$$\Pr[A(e_s, \rho_s) \text{ accepts}] \geq 1 - \varepsilon.$$

Definition of Quantum Money Schemes

n: Security parameter (all computations should be polynomial in n)

B: Poly-size quantum circuit (the “bank”), which maps a secret key $s \in \{0,1\}^n$ to a public key e_s and quantum banknote ρ_s

A: Poly-size quantum circuit (the “authenticator”), which takes (e, ρ) as input and either accepts or rejects

(B, A) has *completeness error* ε if for every s ,

$$\Pr[A(e_s, \rho_s) \text{ accepts}] \geq 1 - \varepsilon.$$

(B, A) has *soundness error* δ if for every poly(n)-size quantum circuit C (the “counterfeiter”) mapping $\rho_s^{\otimes k}$ to $r > k$ output

registers $\sigma_s^1, \dots, \sigma_s^r$,

$$\sum_{i=1}^r \Pr[A(e_s, \sigma_s^i) \text{ accepts}] \leq k + \delta.$$

Counterfeiter only gets ρ_s : scheme is **private-key**

Counterfeiter gets both ρ_s and e_s : scheme is **public-key**

Counterfeiter only gets ρ_s : scheme is **private-key**

Counterfeiter gets both ρ_s and e_s : scheme is **public-key**

Goal: A public-key scheme where completeness error ε and soundness error δ are both exponentially small

Theorem: No public-key quantum money scheme can be information-theoretically secure.

Theorem: No public-key quantum money scheme can be information-theoretically secure.

Proof Sketch: A counterfeiter with unlimited computation time can do this...

Let U be an ensemble of possible quantum money states
Initially, U_0 contains ρ_s for **every** $s \in \{0,1\}^n$

For $t:=0$ to $n-1$ {

 If the legitimate authenticator A_{s^*} accepts a random state from U_t with high probability, we're done!

 Otherwise, get a legitimate quantum money state ρ_{s^*}

 Find an authenticator A_s that rejects **most** states in U_t , but accepts ρ_{s^*}

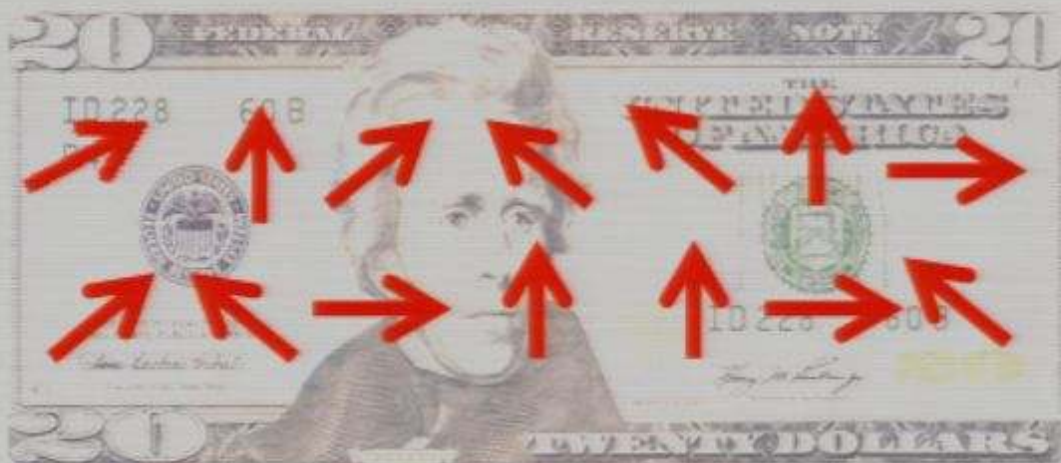
 Let U_{t+1} be the set of states in U_t that A_s accepts w.h.p.

Public-Key Quantum Money Secure Against **Black-Box** Adversaries

Public-Key Quantum Money Secure Against **Black-Box** Adversaries

Doesn't Wiesner's scheme already provide this?

No! A counterfeiter could copy a bill, by using the checking device to figure out the polarization of one qubit at a time...




Public-Key Quantum Money Secure Against **Black-Box** Adversaries

Doesn't Wiesner's scheme already provide this?

Public-Key Quantum Money Secure Against **Black-Box** Adversaries

Doesn't Wiesner's scheme already provide this?


No! A counterfeiter could copy a bill, by using the checking device to figure out the polarization of one qubit at a time...



Public-Key Quantum Money Secure Against **Black-Box** Adversaries

Doesn't Wiesner's scheme already provide this?

No! A counterfeiter could copy a bill, by using the checking device to figure out the polarization of one qubit at a time...



Solution: The bank chooses an n -qubit quantum money state $|\psi\rangle$ uniformly at random under the Haar measure

Solution: The bank chooses an n -qubit quantum money state $|\psi\rangle$ uniformly at random under the Haar measure

The checking device, U , accepts $|\psi\rangle$ and rejects every state orthogonal to $|\psi\rangle$

Key Question: Can a counterfeiter create additional copies of $|\psi\rangle$, using $k=\text{poly}(n)$ copies of $|\psi\rangle$ **together with** $\text{poly}(n)$ queries to U ?

Solution: The bank chooses an n -qubit quantum money state $|\psi\rangle$ uniformly at random under the Haar measure

The checking device, U , accepts $|\psi\rangle$ and rejects every state orthogonal to $|\psi\rangle$

Key Question: Can a counterfeiter create additional copies of $|\psi\rangle$, using $k = \text{poly}(n)$ copies of $|\psi\rangle$ **together with** $\text{poly}(n)$ queries to U ?

If the counterfeiter only had $|\psi\rangle^{\otimes k}$, and not U :

No, by the No-Cloning Theorem

If the counterfeiter only had U , and not $|\psi\rangle^{\otimes k}$:

No, by the optimality of Grover's search algorithm

U must be queried $\Omega(2^{n/2})$ times to find $|\psi\rangle$

Complexity-Theoretic No-Cloning Theorem

Let $|\psi\rangle$ be an n -qubit state. Suppose we're given $|\psi\rangle^{\otimes k}$, **as well as** a black box U that accepts $|\psi\rangle$ and rejects all states orthogonal to $|\psi\rangle$. Then to prepare $r > k$ states ρ_1, \dots, ρ_r such that

$$\sum_{i=1}^r \langle \psi | \rho_i | \psi \rangle \geq k + \delta,$$

we need this many queries to U : $\Omega\left(\frac{\delta^2 \sqrt{2^n}}{r^2 k \log k} - r\right)$

Complexity-Theoretic No-Cloning Theorem

Let $|\psi\rangle$ be an n -qubit state. Suppose we're given $|\psi\rangle^{\otimes k}$, **as well as** a black box U that accepts $|\psi\rangle$ and rejects all states orthogonal to $|\psi\rangle$. Then to prepare $r > k$ states ρ_1, \dots, ρ_r such that

$$\sum_{i=1}^r \langle \psi | \rho_i | \psi \rangle \geq k + \delta,$$

we need this many queries to U : $\Omega\left(\frac{\delta^2 \sqrt{2^n}}{r^2 k \log k} - r\right)$

Proof requires generalizing Ambainis's adversary method, to the case where the quantum algorithm's initial state already encodes some information about the target state

Complexity-Theoretic No-Cloning Theorem

Let $|\psi\rangle$ be an n -qubit state. Suppose we're given $|\psi\rangle^{\otimes k}$, **as well as** a black box U that accepts $|\psi\rangle$ and rejects all states orthogonal to $|\psi\rangle$. Then to prepare $r > k$ states ρ_1, \dots, ρ_r such that

$$\sum_{i=1}^r \langle \psi | \rho_i | \psi \rangle \geq k + \delta,$$

we need this many queries to U : $\Omega\left(\frac{\delta^2 \sqrt{2^n}}{r^2 k \log k} - r\right)$

Proof requires generalizing Ambainis's adversary method, to the case where the quantum algorithm's initial state already encodes some information about the target state

Explicit Candidate Scheme

A **stabilizer state** is a state obtainable from $|0\dots 0\rangle$ by applying Hadamard, Controlled-NOT, and Phase gates only:

$$\frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}, \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix}$$

Explicit Candidate Scheme

A **stabilizer state** is a state obtainable from $|0\dots 0\rangle$ by applying Hadamard, Controlled-NOT, and Phase gates only:

$$\frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}, \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix}$$

These states can always be efficiently prepared!

In my scheme, a dollar bill consists of:

- L random stabilizer states $|C_1\rangle, \dots, |C_L\rangle$ on n qubits each
- A table of measurements to apply to the $|C_i\rangle$'s
- A (conventional) digital signature of the table

The table:

$ C_1\rangle$	$ C_2\rangle$	$ C_3\rangle$...
M_{11}	M_{21}	M_{31}	...
M_{12}	M_{22}	M_{32}	...
M_{13}	M_{23}	M_{33}	...
M_{14}	M_{24}	M_{34}	...
\vdots	\vdots	\vdots	\ddots

For each $|C_i\rangle$, we have lots of random garbage measurements, but also a secret ε fraction that commute with $|C_i\rangle$

The table:

$ C_1\rangle$	$ C_2\rangle$	$ C_3\rangle$...
M_{11}	M_{21}	M_{31}	...
M_{12}	M_{22}	M_{32}	...
M_{13}	M_{23}	M_{33}	...
M_{14}	M_{24}	M_{34}	...
\vdots	\vdots	\vdots	\ddots

For each $|C_i\rangle$, we have lots of random garbage measurements, but also a secret ε fraction that commute with $|C_i\rangle$

To verify a bill:

1. Verify the table's digital signature
2. For each i , apply a random measurement M_{ij} to $|C_i\rangle$
3. Accept if more than $\frac{1}{2} + \frac{\varepsilon}{2}$ of the measurements do

The table:

$ C_1\rangle$	$ C_2\rangle$	$ C_3\rangle$...
M_{11}	M_{21}	M_{31}	...
M_{12}	M_{22}	M_{32}	...
M_{13}	M_{23}	M_{33}	...
M_{14}	M_{24}		
\vdots	\vdots		

For each $|C_i\rangle$, we have lots of random garbage measurements, but also a secret ε fraction that commute with $|C_i\rangle$

Hope: Learning classical descriptions of the $|C_i\rangle$'s, or copying them in any other way, is computationally intractable (a "noisy parity problem")

1. Verify the
2. For each i , apply a random measurement M_{ij} to $|C_i\rangle$

3. Accept if more than $\frac{1}{2} + \frac{\varepsilon}{2}$ of the measurements do

Breaking Aaronson's Scheme

Breaking Aaronson's Scheme

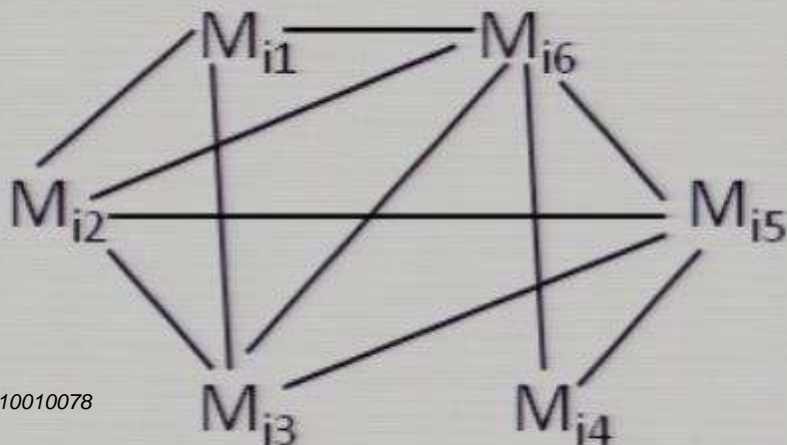
Two cases:

1. ε is extremely small. Then the test is "too weak," and we can guess our own states $|C_i\rangle$ that pass the test

Breaking Aaronson's Scheme

Two cases:


1. ε is extremely small. Then the test is "too weak," and we can guess our own states $|C_i\rangle$ that pass the test
2. ε is reasonably large. Then for each $|C_i\rangle$, consider a graph of the possible measurements, with an edge between M_{ij} and M_{ik} iff they commute with each other:



Breaking Aaronson's Scheme

Two cases:

1. ε is extremely small. Then the test is "too weak," and we can guess our own states $|C_i\rangle$ that pass the test
2. ε is reasonably large. Then for each $|C_i\rangle$, consider a graph of the possible measurements, with an edge between M_{ij} and M_{ik} iff they commute with each other:



The "secret" measurements that commute with $|C_i\rangle$ also commute with each other. Thus, the problem reduces to finding a "planted clique" in a random-looking graph.

Breaking Aaronson's Scheme

Two cases:

1. ϵ is extremely small. Then the test is "too weak," and we can guess our own states $|C_i\rangle$ that pass the test

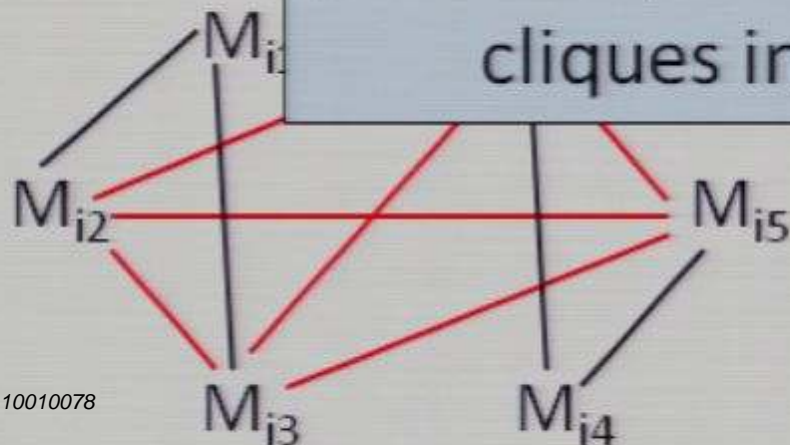
2. ϵ is reasonably large. Then we can use a graph coloring algorithm between

Here we're able to adapt an eigenvector-based algorithm of Alon, Krivelevich, and Sudakov (SODA'98) for finding large planted cliques in random graphs

for a
ge
other:

ements
 $|C_i\rangle$ also

commute with each other. Thus, the problem reduces to finding a "planted clique" in a random-looking graph.



Our New Scheme

$$\frac{1}{2^{n/2}} \sum_{x \in \{0,1\}^n} |x\rangle$$

1. Start with an equal superposition over all n-bit strings

Our New Scheme

$$\frac{1}{2^{n/2}} \sum_{x \in \{0,1\}^n} |x\rangle$$

1. Start with an equal superposition over all n-bit strings

Our New Scheme

$$\frac{1}{2^{n/2}} \sum_{x \in \{0,1\}^n} |x\rangle |h_1(x), \dots, h_m(x)\rangle$$

1. Start with an equal superposition over all n-bit strings
2. Compute randomly-chosen hash functions $h_1, \dots, h_m: \{0,1\}^n \rightarrow \{0,1\}$ (with $m \sim \sqrt{n}$)

Our New Scheme

$$|\psi\rangle |r_1, \dots, r_m\rangle$$

1. Start with an equal superposition over all n -bit strings
2. Compute randomly-chosen hash functions $h_1, \dots, h_m: \{0,1\}^n \rightarrow \{0,1\}$ (with $m \sim \sqrt{n}$)
3. Measure $h_1(x), \dots, h_m(x)$, leaving a superposition $|\psi\rangle$ over all x 's for which h_1, \dots, h_m take on prescribed values r_1, \dots, r_m

To verify a bill $|\psi\rangle|r\rangle|\text{sig}(r)\rangle$:

To verify a bill $|\psi\rangle|r\rangle|\text{sig}(r)\rangle$:

1. Verify r 's digital signature.
2. Construct a Markov chain M , whose stationary distribution is uniform over the set $S = \{x : h_1(x)=r_1, \dots, h_m(x)=r_m\}$. Using M , verify that $|\psi\rangle$ is an equal superposition over S .

Conjecture: Any quantum algorithm needs exponential time to copy $|\psi\rangle$

To verify a bill $|\psi\rangle|r\rangle|\text{sig}(r)\rangle$:

1. Verify r 's digital signature.
2. Construct a Markov chain M , whose stationary distribution is uniform over the set $S = \{x : h_1(x)=r_1, \dots, h_m(x)=r_m\}$. Using M , verify that $|\psi\rangle$ is an equal superposition over S .

Conjecture: Any quantum algorithm needs exponential time to copy $|\psi\rangle$

Striking feature of this scheme: The **bank** can't copy $|\psi\rangle$, any more than a counterfeiter can!!

Nor (we believe) can the bank efficiently create two bills with the same "serial number" r

Unlike with the stabilizer scheme, here there's no obvious "classical secret" that lets you copy a bill if you learn it

Quantum Software Copy-Protection



Finally, a serious use for quantum computing

Copy-Protecting Point Functions

Point function:

$$f_s(x) = \begin{cases} 1 & \text{if } x = s \\ 0 & \text{otherwise} \end{cases}$$

Think: The UNIX password program

Except, given the quantum program $|\psi_s\rangle$, we want it to be hard not merely to learn the password s , but even to create more programs able to **recognize** s !

Possible Solution: Use s to generate a *pseudorandom quantum circuit* U_s , then set $|\psi_s\rangle := U_s|0 \cdots 0\rangle$

Copy-Protecting Point Functions

Point function:

$$f_s(x) = \begin{cases} 1 & \text{if } x = s \\ 0 & \text{otherwise} \end{cases}$$

Think: The UNIX password program

Except, given the quantum program $|\psi_s\rangle$, we want it to be hard not merely to learn the password s , but even to create more programs able to **recognize** s !

Possible Solution: Use s to generate a *pseudorandom quantum circuit* U_s , then set $|\psi_s\rangle := U_s|0\dots 0\rangle$

To compute $f_s(x)$, measure $U_x^{-1}|\psi_s\rangle$ in the standard basis, and see if you get back the all-0 string

Summary

Unforgeable money (and copy-protected software, etc.) remains one of the most striking potential applications of quantum mechanics to computer science

Summary

Unforgeable money (and copy-protected software, etc.) remains one of the most striking potential applications of quantum mechanics to computer science

So we've been revisiting this 40-year-old idea using the arsenal of modern CS theory

Biggest challenge: Secure quantum money that anyone can verify (not just the bank)

Summary

Unforgeable money (and copy-protected software, etc.) remains one of the most striking potential applications of quantum mechanics to computer science

So we've been revisiting this 40-year-old idea using the arsenal of modern CS theory

Biggest challenge: Secure quantum money that anyone can verify (not just the bank)

I showed how to achieve this in the 'black-box world'

But in the 'real' world, finding a scheme that withstands attack is harder than it looks!



Open Problems



Can we base the security of public-key quantum money on a “standard” cryptographic assumption? How about copy-protection?

Can we copy-protect anything besides point functions?

Can we get provably-secure public-key quantum money, with the help of only a **classical** black box?

Other “non-cloneable functionalities”: keys? ID cards?

Summary

Unforgeable money (and copy-protected software, etc.) remains one of the most striking potential applications of quantum mechanics to computer science

So we've been revisiting this 40-year-old idea using the arsenal of modern CS theory

Biggest challenge: Secure quantum money that anyone can verify (not just the bank)

Quantum Software Copy-Protection



Finally, a serious use for quantum computing

We know copy-protection is **fundamentally impossible** in the classical world (not that that's stopped people from trying...)

Question: Can you have a quantum state $|\psi_f\rangle$ that lets you efficiently compute an unknown Boolean function $f:\{0,1\}^n \rightarrow \{0,1\}$, but **can't** be efficiently used to prepare more states that also let you efficiently compute f ?

To verify a bill $|\psi\rangle|r\rangle|\text{sig}(r)\rangle$:

1. Verify r 's digital signature.
2. Construct a Markov chain M , whose stationary distribution is uniform over the set $S = \{x : h_1(x)=r_1, \dots, h_m(x)=r_m\}$. Using M , verify that $|\psi\rangle$ is an equal superposition over S .

Conjecture: Any quantum algorithm needs exponential time to copy $|\psi\rangle$

Striking feature of this scheme: The **bank** can't copy $|\psi\rangle$, any more than a counterfeiter can!!

Nor (we believe) can the bank efficiently create two bills with the same "serial number" r

Unlike with the stabilizer scheme, here there's no obvious "classical secret" that lets you copy a bill if you learn it

To verify a bill $|\psi\rangle|r\rangle|\text{sig}(r)\rangle$:

1. Verify r 's digital signature.
2. Construct a Markov chain M , whose stationary distribution is uniform over the set $S = \{x : h_1(x)=r_1, \dots, h_m(x)=r_m\}$. Using M , verify that $|\psi\rangle$ is an equal superposition over S .

Our New Scheme

$$|\psi\rangle|r\rangle|\text{sig}(r)\rangle$$

1. Start with an equal superposition over all n -bit strings
2. Compute randomly-chosen hash functions $h_1, \dots, h_m: \{0,1\}^n \rightarrow \{0,1\}$ (with $m \sim \sqrt{n}$)
3. Measure $h_1(x), \dots, h_m(x)$, leaving a superposition $|\psi\rangle$ over all x 's for which h_1, \dots, h_m take on prescribed values r_1, \dots, r_m
4. As the dollar bill, distribute $|\psi\rangle$, $r=(r_1, \dots, r_m)$, and a conventional digital signature of r