

Κατανεμημένα Συστήματα

Bitcoin and Blockchain

2016-2017

<http://www.cslab.ece.ntua.gr/courses/distrib>

Blockchain Defined

Simply defined a Blockchain is little more than a:

- Distributed
- Secure
- Ledger (logfile)

A digital currency was in a lot of ways the first demonstrable use

What is Bitcoin

- A **protocol** that supports a decentralized, pseudo-anonymous, peer-to-peer digital currency
- A **publicly** disclosed linked **ledger** of transactions stored in a blockchain
- A **reward** driven system for achieving **consensus** (mining) based on “Proofs of Work” for helping to secure the network
- An economy with an eventual cap of about 21M bitcoins

Bitcoin Whitepaper – 2008.10.31

Bitcoin: A Peer-to-Peer Electronic Cash System

Satoshi Nakamoto
satoshin@gmx.com
www.bitcoin.org

Abstract. A purely peer-to-peer version of electronic cash would allow online payments to be sent directly from one party to another without going through a financial institution. Digital signatures provide part of the solution, but the main benefits are lost if a trusted third party is still required to prevent double-spending. We propose a solution to the double-spending problem using a peer-to-peer network. The network timestamps transactions by hashing them into an ongoing chain of hash-based proof-of-work, forming a record that cannot be changed without redoing the proof-of-work. The longest chain not only serves as proof of the sequence of events witnessed, but proof that it came from the largest pool of CPU power. As long as a majority of CPU power is controlled by nodes that are not cooperating to attack the network, they'll generate the longest chain and outpace attackers. The network itself requires minimal structure. Messages are broadcast on a best effort basis, and nodes can leave and rejoin the network at will, accepting the longest

Features of Bitcoin

- Essentially it's "deflationary" – the reward is cut in half every four years
- Nearly infinitely divisible currency units supporting eight decimal places 0.00000001 (known as a Satoshi)
- Nominal transaction fee's paid to the network
 - Same cost to send \$.01 as \$1,000,000
- Consensus driven – no central authority
- Counterfeit resilient
 - Cannot add coins arbitrarily
 - Cannot be double-spent
- Non-repudiation – aka "gone baby gone" – no recourse and no one to appeal to return sent tokens

5

When did it start?

- “Satoshi Nakamoto” created the reference implementation that began with a Genesis Block of 50 coins
- **2008**
 - **August 18** Domain name "bitcoin.org" registered^[1].
 - **October 31** Bitcoin design paper published
 - **November 09** Bitcoin project registered at SourceForge.net
- **2009**
 - **January 3** Genesis block established at 18:15:05 GMT
 - **January 9** Bitcoin v0.1 released and announced on the cryptography mailing list
 - **January 12** First Bitcoin transaction, in block 170 from Satoshi to Hal Finney

<https://en.bitcoin.it/wiki/History>




6

Why does it have value?

*The worth of a thing
is the price it will bring.*

Why does it matter?

16 Billion Dollar Market Cap!

▲#	Name	Market Cap	Price	Available Supply	Volume (24h)
1	 Bitcoin	\$14,427,302,148	\$895.58	16,109,487 BTC	\$132,286,000
2	 Ethereum	\$903,906,323	\$10.27	88,010,820 ETH	\$17,578,600
3	 Ripple	\$255,764,406	\$0.006956	36,771,322,652 XRP *	\$1,800,960

<http://coinmarketcap.com>

BitCoin: Challenges

- All virtual currency must address the following challenges:
 - Creation of a virtual coin/note
 - How is it created in the first place?
 - How do you prevent inflation? (What prevents anyone from creating lots of coins?)
 - Validation
 - Is the coin legit? (proof-of-work)
 - How do you prevent a coin from double-spending?
- BitCoin takes a infrastructure-less approach
 - Rely on proof instead of trust
 - No central bank or clearing house

BitCoin: Motivation

- Rely on proof instead of trust
 - Current online transactions rely on a trusted party (e.g, VISA)
 - They take some risk, manage fraud, and get paid a fee.
- Buyer and Seller protection in online transactions
 - Buyer pays, but the seller doesn't deliver → Solved by using an escrow (Buyer protection)
 - Seller delivers, buyer pays, but the buyer makes a claim. VISA refunds; the payment is reversed. Either the seller is penalized and/or VISA charges more fee to handle these cases. Some behaviors are fraudulent.
 - BitCoin gets rid of this trusted middleman, by being able to directly show the cryptographic proof that the money is transferred.

Four components in secure communication

- Authentication
- Confidentiality
- Integrity
- Availability

What do we want to secure?

- Authentication (Who am I talking to?)
 - Identification and assurance of the origin of information
- Confidentiality (Is my data hidden?)
 - Concealment of information
- Integrity (Has my data been modified?)
 - Prevent improper and unauthorized changes
- Availability (Can I use the resources?)
 - The ability to use the information or resource desired

From the perspective of BitCoin


- **Authentication**
 - Am I paying the right person? Not some other impersonator?
- **Integrity**
 - Is the coin double-spent?
 - Can an attacker reverse or change transactions?
- **Availability**
 - Can I make a transaction anytime I want?
- **Confidentiality**
 - Not very relevant. But privacy is important.

From the perspective of BitCoin

- **Authentication** → Public Key Crypto: Digital Signatures
 - Am I paying the right person? Not some other impersonator?
- **Integrity** → Digital Signatures and Cryptographic Hash
 - Is the coin double-spent?
 - Can an attacker reverse or change transactions?
- **Availability**
 - Can I make a transaction anytime I want?
- **Confidentiality**
 - Not very relevant. But privacy is important.

Cryptographic Hash Functions

- **Consistent:** $H(X)$ always yields same result
- **One-way:** given Y , hard to find X s.t. $H(X) = Y$
- **Collision resistant:** given $H(W) = Z$, hard to find X such that $H(X) = Z$



Collision resistant

- Find a collision:
 - Try 2^{130} randomly chosen inputs
 - 99,8% chance that two of them collide

- Takes too long to matter

SHA256

In practice, we hope that SHA256 behaves “like a random oracle”.

SHA256: TextFiles $\rightarrow \{0, \dots, 2^{256} - 1\}$

Calculation: If we made *all* computers on the world compute SHA256...

It takes \sim “ $40 \times 14 \cdot 10^9$ years” to find $x_1 \neq x_2$ s.t.
 $\text{SHA256}(x_1) = \text{SHA256}(x_2)$.


Application: Hash as a message digest

- $H(x) = H(y)$ then safe to assume $x=y$
- To recognize a file that we saw before, just remember its hash
- Useful because it's small


One way

- Given $H(x)$, infeasible to find x
- Distributions of values should be very spread out (e.g., uniform)

Digital Signature

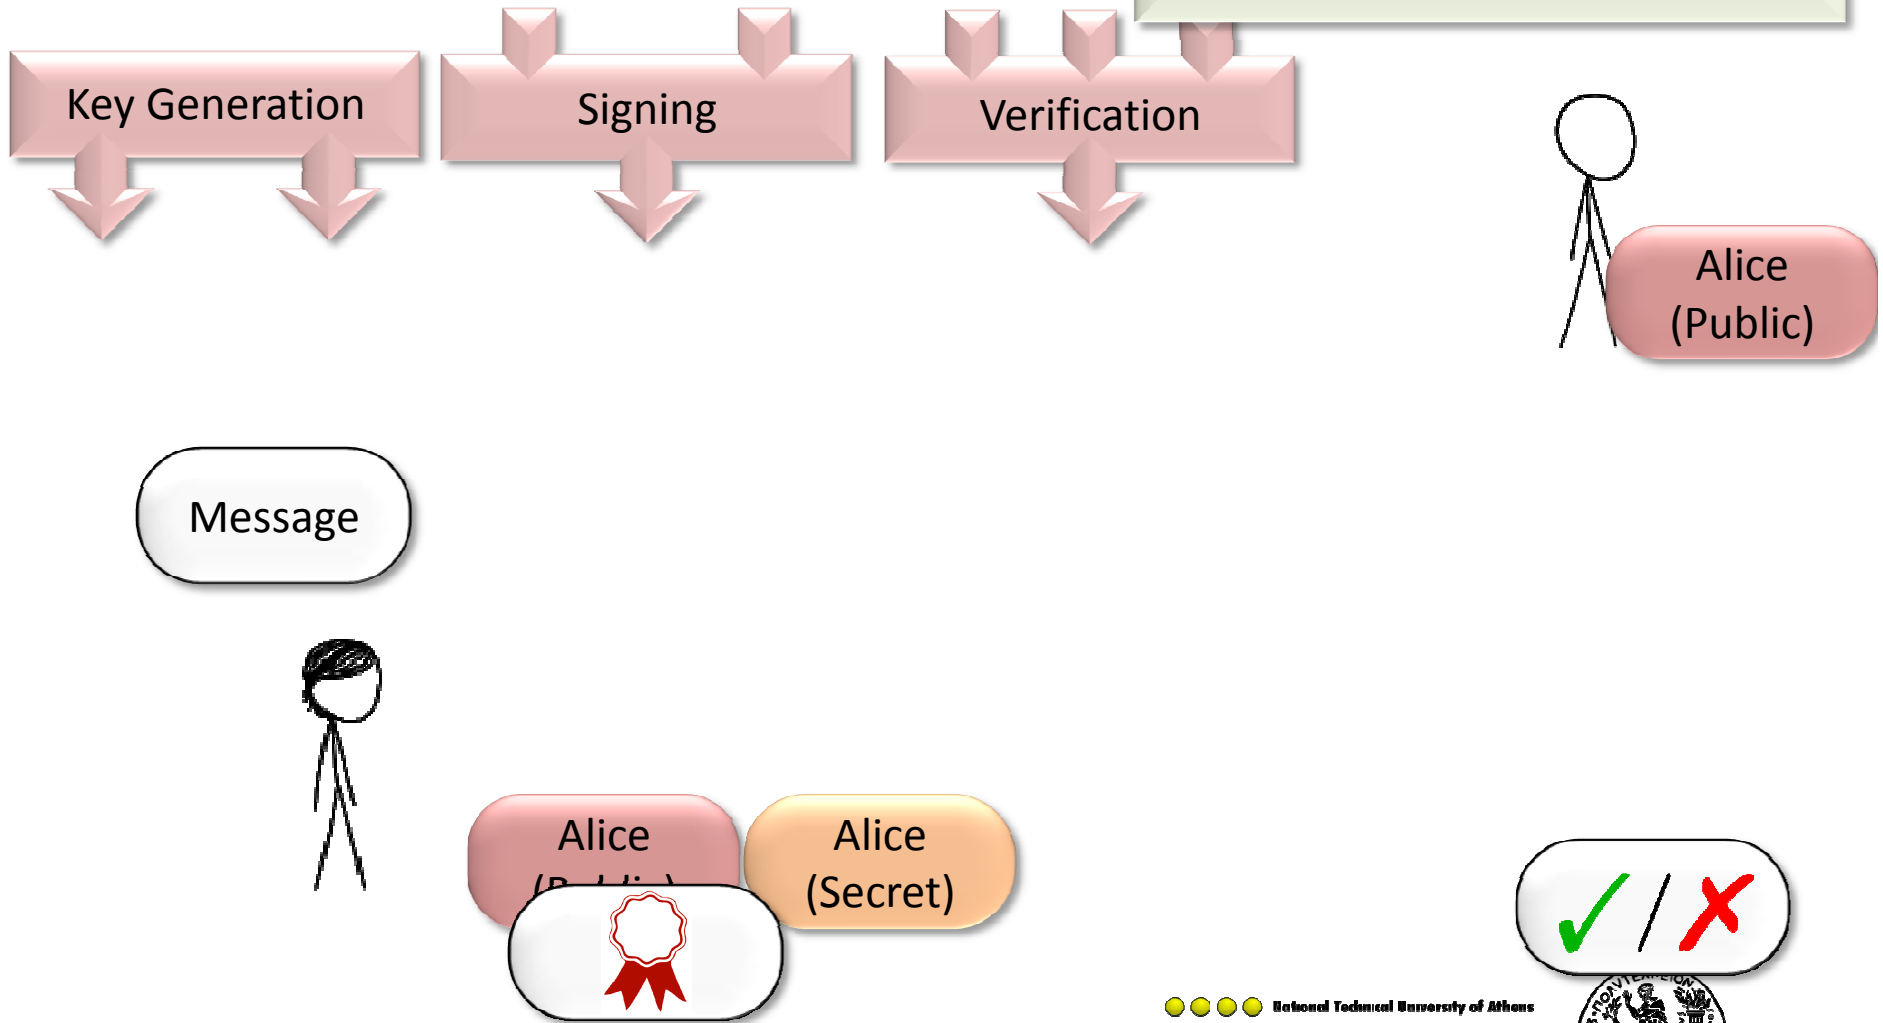


Digital Signature



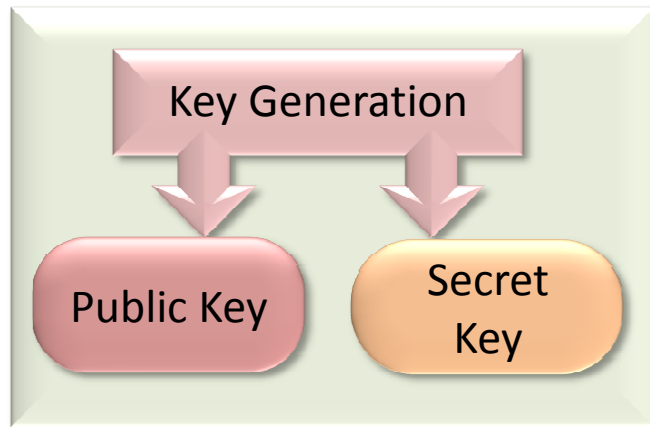


Digital Signatures

Goal: Bob should be sure that the message originates from Alice.



Digital Signature



Security (informal): You cannot produce valid signatures without the secret key.

Back to BitCoins

- Validation
 - Is the coin legit? (proof-of-work) → Use of **Cryptographic Hashes**
 - How do you prevent a coin from double-spending? → **Broadcast to all nodes**
- Creation of a virtual coin/note
 - How is it created in the first place? → Provide **incentives for miners**
 - How do you prevent inflation? (What prevents anyone from creating lots of coins?) → **Limit the creation rate of the BitCoins**

ATTEMPT #1

We now try to build bitcoin...


... but we will fail.

Goals

- We want some kind of “digital money”.
 - *Everyone* can participate.
 - No central instance – no bank.

Setting

- A network of computers.



- Every computer can send messages to *some* other computers.

Basic idea


- Every computer maintains a table: “who owns what?”
- We will need: *all* computers have the *same* table.

Alice (Public)	10 BTC
Bob (Public)	0.2 BTC
Charlie (Public)	17 BTC
Dora (Public)	0.001 BTC
Eliza (Public)	2 BTC


Remark: The public keys are just bit strings.

Sending Bitcoins




To send money, we use **transactions**. These are messages like this:



In “short”, transactions look like this:



Main Transaction pic





Sending Bitcoins

I'll send 0.1 Bitcoin to Bob.



Alice



Protocol: **sending BTC**

1. Craft a transaction.
2. Give it to your computer.

Protocol: **participating**


On valid transactions:

1. Update ledger
2. Relay transaction


Double Spending

Thanks!

I can exploit this!



These transactions spend previously spent bitcoins!








Black Hat prepares *two* transactions:

: Give BTC from Black Hat to Alice

: Give BTC from Black Hat to Bob

Double Spending




- The bad guy spends the *same* Bitcoins with two different transactions  and .
- Computers receiving transaction  will have a *different* ledger than computers receiving transaction .

Consensus Protocols

- We need a protocol to *agree* on a transaction.
- “Consensus protocols”. Studied since 1980, starting with Pease, Shostak, Lamport.
- Huge literature!

What transaction are you using?
Main idea for protocols:

Protocols work if (say) $> 70\%$ of the computers follow the protocol.



This solution does not help us!

Design goal: *Everyone* can participate.

I will gladly participate...

With 1 000 virtual machines!



By running a special program, a bad guy controls many virtual computers.

Like this, he can make different participants believe different things.

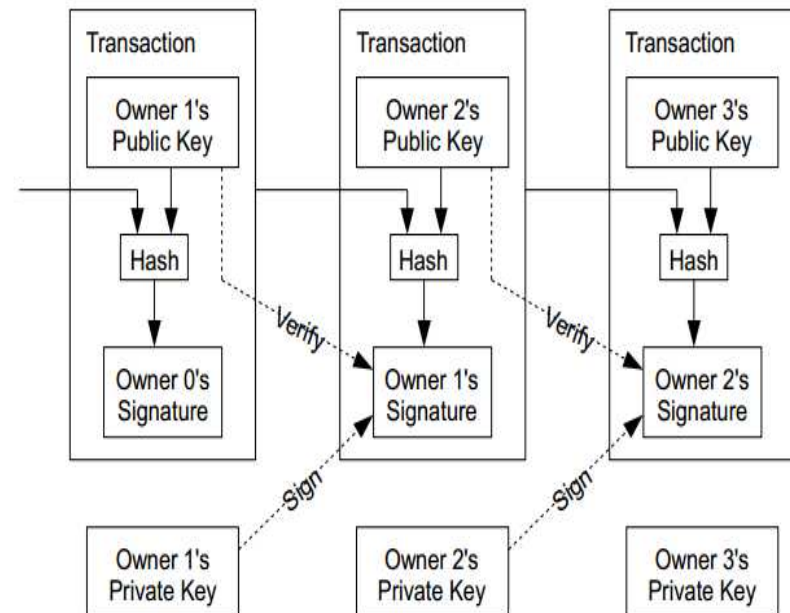
BITCOIN'S CONSENSUS PROTOCOL

Step 1: How does the protocol look like?

Step 2: What happens if people cheat?

BitCoin

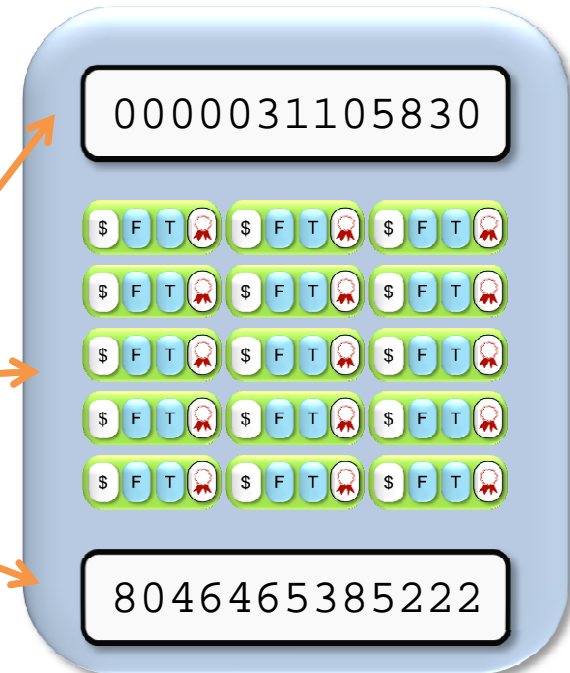
- Electronic coin == chain of digital signatures
- BitCoin transfer: $\text{Sign}(\text{Previous transaction} + \text{New owner's public key})$
- Anyone can verify (n-1)th owner transferred this to the nth owner.
- Anyone can follow the history given a BitCoin



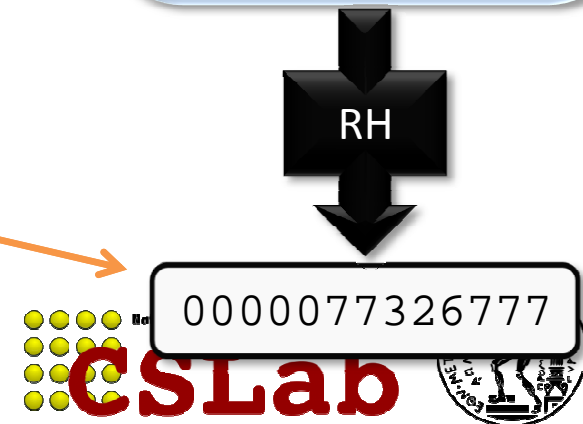
Blocks

A block B contains

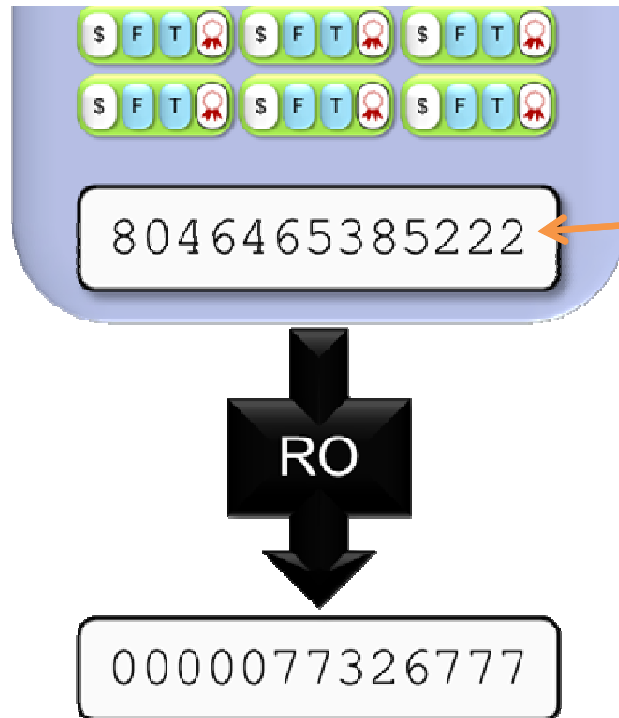
- $\text{RH}(B')$ for another block B' ,
- a list of transactions,
- and an arbitrary number



Block B is **valid** if the first $d = 5$ digits of the hash of B are all



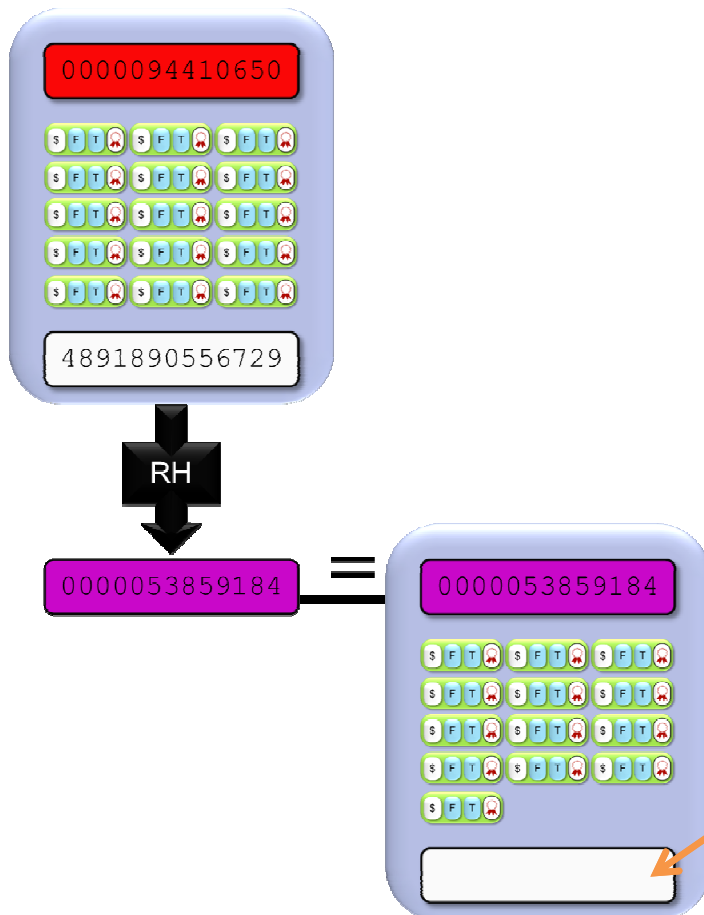
Blocks



Block B is **valid** if the first $d = 5$ digits of $RO(B)$ are all zero.

- To **find** a valid block, we try different values for this string (“nonce”).
- On average, after $10^d = 100000$ tries, we find a valid block.
- Bitcoin chooses d on the fly such that this takes about 10 minutes.

Blocks

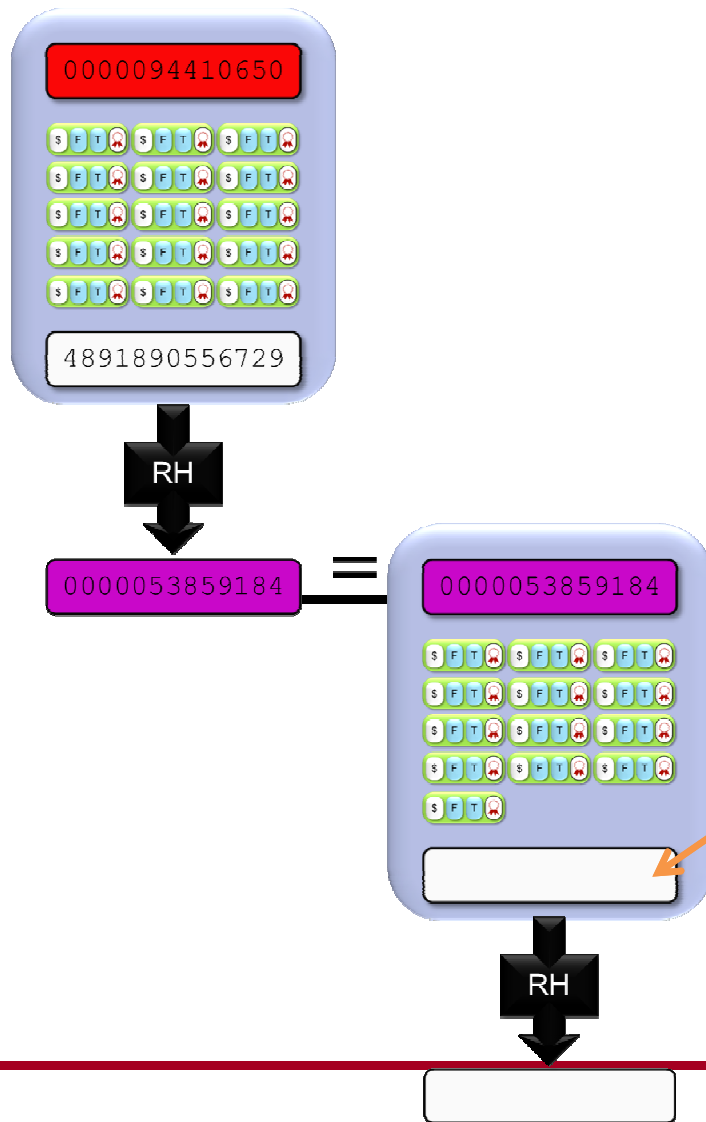


If we have a block, we can find a “next block”:

Take $RH(B')$ from the previous block B' . Add transactions.

Try different values for this string until the hash starts with d zeros.

Blocks



If we have a block, we can find a “next block”:

Take $RH(B')$ from the previous block B' . Add transactions.


Try different values for this string until the hash starts with d zeros.

Bitcoin chooses d such that this takes ~10 minutes

BitCoin Network

- Each P2P node runs the following algorithm [bitcoin]:
 - New transactions are broadcast to all nodes.
 - Each node collects new transactions into a block.
 - Each node works on finding a proof-of-work for its block.
(Hard to do. Probabilistic. The one to finish early will probably win.)
 - When a node finds a proof-of-work, it broadcasts the block to all nodes.
 - Nodes accept the block only if all transactions in it are valid (digital signature checking) and not already spent (check all the transactions).
 - Nodes express their acceptance by working on creating the next block in the chain, using the hash of the accepted block as the previous hash.

A Tree of Blocks



If we have a block, with a bit of work, we can find a “next block” ...

...and yet another “next block” ...


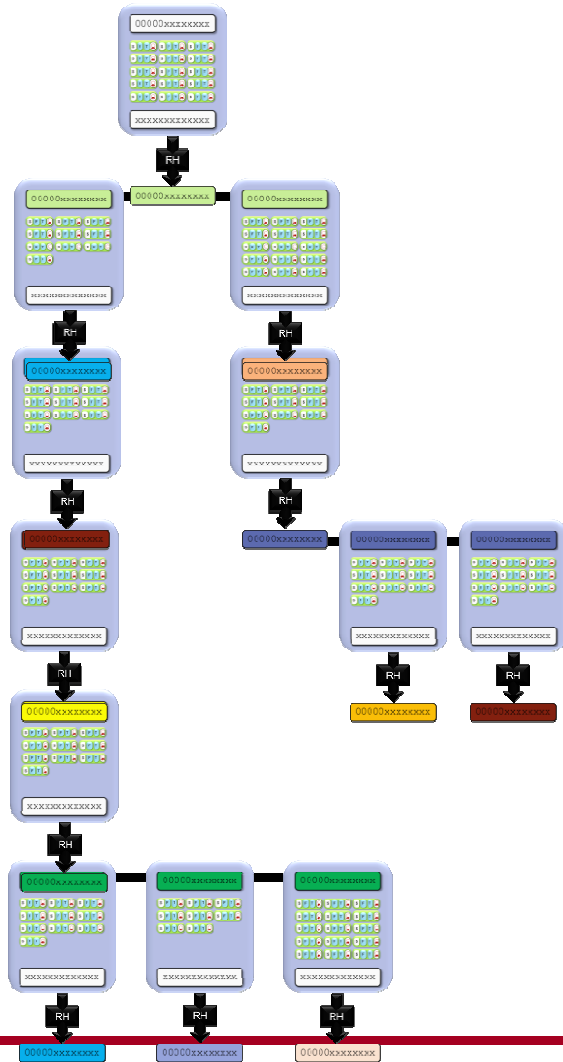
...or a block which continues here...

... and so on

A Tree of Blocks

In general, we can build a tree of blocks like this.

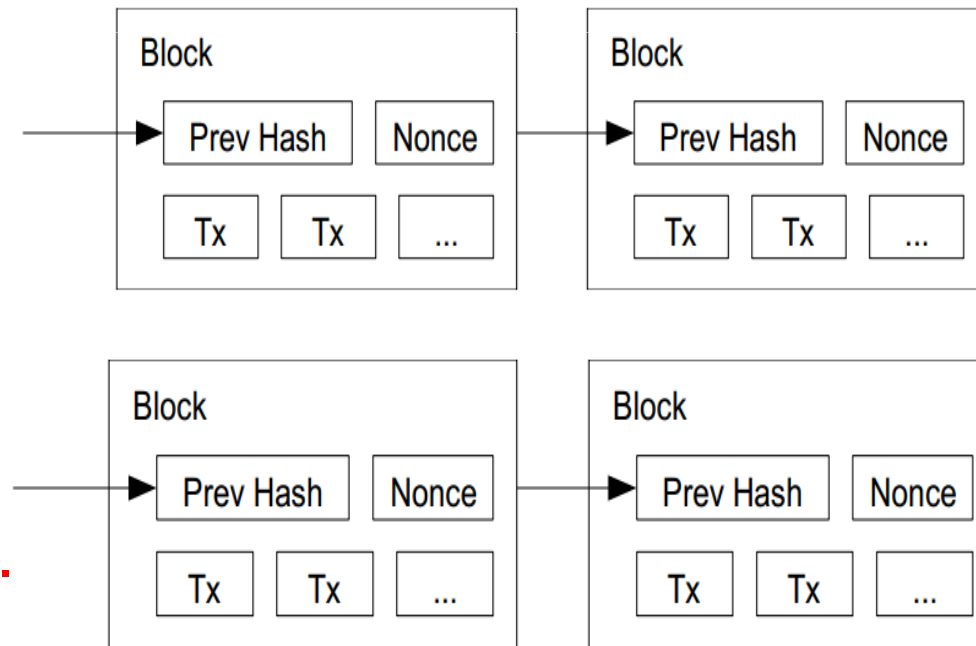
But only ever downwards!



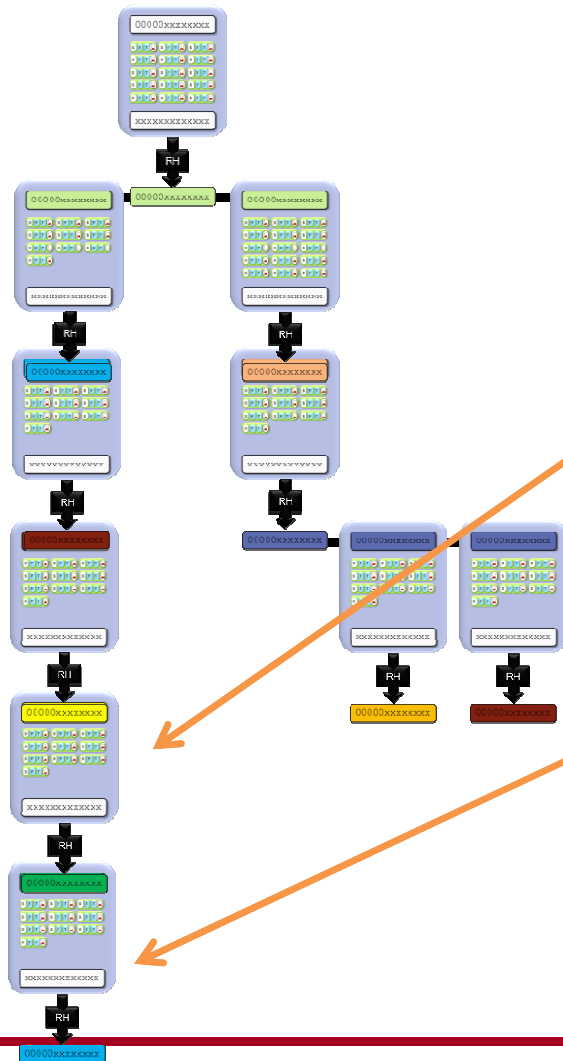
Tie breaking

- Two nodes may find a correct block simultaneously.
 - Keep both and work on the first one
 - If one grows longer than the other, take the longer one

Two different block chains (or blocks) may satisfy the required proof-of-work.



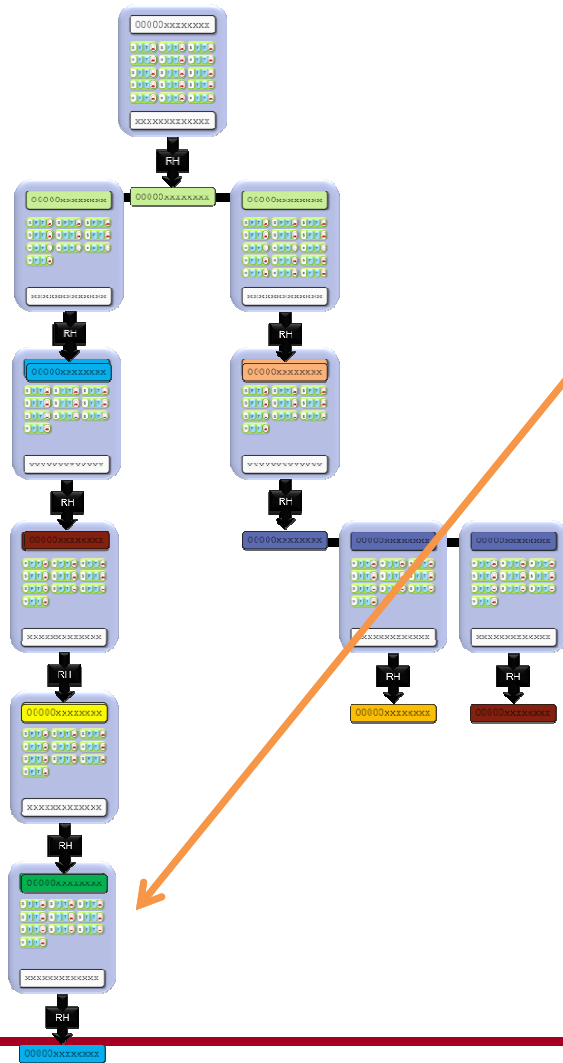
The Protocol for Finding Blocks



Protocol: finding blocks

1. Take the longest chain you can find.
2. Collect transactions.
3. Find a new valid block here.
4. Publish it.

The Protocol for Participants

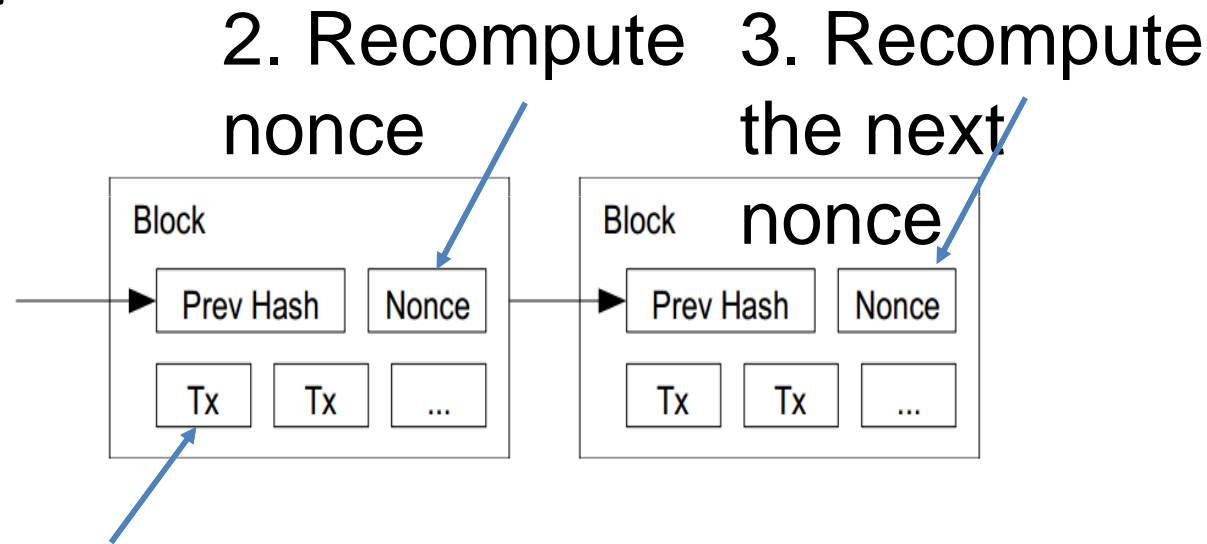


Protocol: To know who owns BTC

1. Take the longest chain you can find.
2. Process the transactions in this chain in order.

Reverting is hard...

- Reverting gets exponentially hard as the chain grows.



1. Modify the transaction
(revert or change the
payer)

Practical Limitation

- At least 10 mins to verify a transaction.
 - Agree to pay
 - Wait for one block (10 mins) for the transaction to go through.
 - But, for a large transaction (\$\$\$) wait longer. Because if you wait longer it becomes more secure. For large \$\$\$, you wait for six blocks (1 hour).

Why work to find blocks?

Many
uses a

A *real* lot!

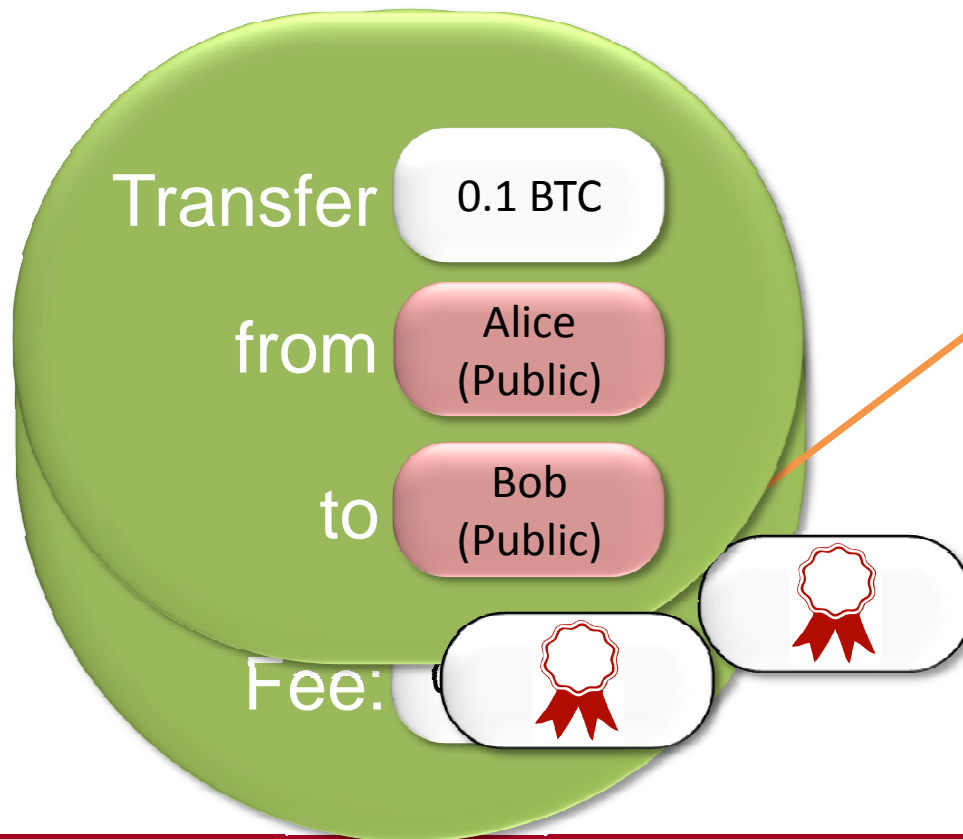


This is called "mining"
CSLab



Block reward

If you find a block, you get bitcoins as a reward.



Every transaction specifies a fee. It goes to the person who puts the transaction into a valid block.

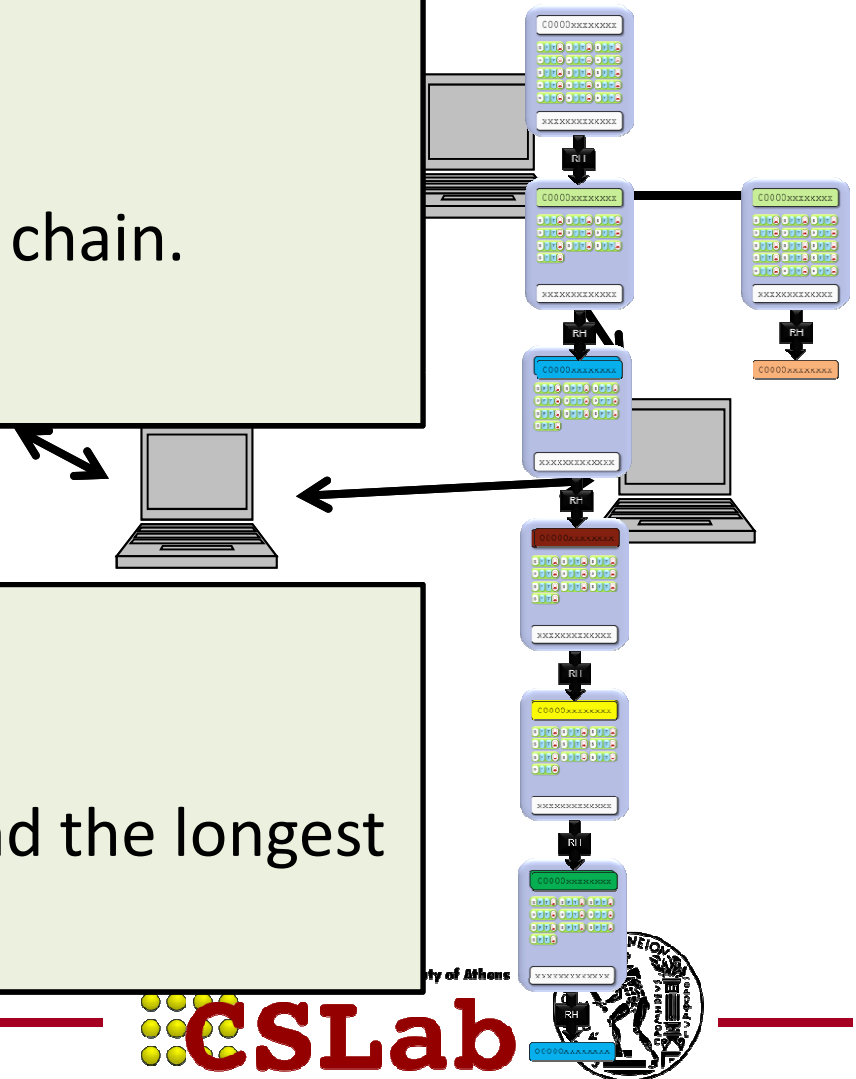
Recap: The Bitcoin Protocol

Protocol: **participate**

- ❑ Relay valid transactions.
- ❑ Relay valid blocks in the longest chain.
- ❑ Work with the longest chain.

Protocol: **miners**

- ❑ Collect valid transactions.
- ❑ Publish valid blocks which extend the longest chain.



Step 1: How does the protocol look like?

BITCOIN'S CONSENSUS PROTOCOL


Step 2: What happens if people cheat?

Double Spends found a valid block!

I can exploit this!



Black Hat

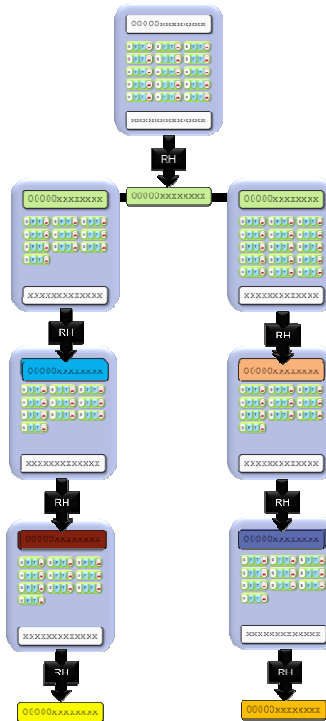


Once a block is found, the double spends vanish.

Occasionally, two people find blocks at around the same time... but typically the problem disappears.

Build an Alternate Chain?

Maybe I should build another chain?




¶ The more RH-calls are devoted to a chain, the faster it grows.

Thus, intuitively: to build a chain as fast as the rest, you need as many RH-calls as the rest.

Hardware War

Bitcoin Mining Hardware Comparison

	Miner	Hash Power	Price	Buy
	Antminer S5	1.16 TH/s	\$139.99	
	Antminer S7	4.73 TH/s	\$489.99	
	Antminer S9	14.0 TH/s	\$3,000	
	Avalon 6	3.50 TH/s	\$559.95	
	SP20 Jackson	1.3-1.7 TH/s	\$90.00	

Summary

- BitCoin combined techniques from crypto and the right incentives.
 - Nice design
 - A trait for popular systems
- BitCoin is becoming industrialized.
 - Miners form a pool.
 - Mining hardware becomes sophisticated.
 - BitCoin exchange
 - Derivative market, etc.
 - Government agencies are keeping an eye on them.
- Who will control BitCoin in the end?

More uses of blockchain?

- If the blockchain technology works, it gives a new consensus algorithm. What else can we use it for?
- Ideas:
 - Multiparty computation protocols based on the blockchain.
 - Timestamping
 - Crowdfunding
 - Have your shares in the blockchain
 - Smart payments
 - etc...

References

- Bitcoin and Cryptocurrency Technologies, Arvind Narayanan, Joseph Bonneau, Edward Felten, Andrew Miller, Steven Goldfeder
- Bitcoin: A primer by François R. Velde, senior economist FRB
- Bitcoin: A Peer-to-Peer Electronic Cash System, Satoshi Nakamoto
- <http://bitcoinbook.cs.princeton.edu/>