CPSC 367: Cryptography and Security

Michael J. Fischer

Lecture 24
April 23, 2019

Blockchains

Blockchains

The theoretical concept of a *Blockchain* is a distributed database built using cryptography. It has several nice properties:

- It is a decentralized, anonymous and public digital *ledger*, managed by collectively trusted servers.
- Entries are indelible, that is permanent and unchangeable.
- Additions to the blockchain are immediately visible to everyone.
- A blockchain gives a unique, public, global, and consistent view of its data.

The problem is that real implementations cannot fully implement the theory.

How a blockchain works

A blockchain (*ledger*) is a record of transactions.

- It consists of a list of *blocks*, linked together using cryptographic hash functions.
- Each block contains a list of *transactions*, the author ID, and a cryptographic hash of the prior block in the chain.
- The chain is managed by a peer-to-peer network of maintainers (often called *miners*) who follow a protocol for communication and validation.
- Once recorded, the data in a block cannot be changed without changing the prior-block hashes of all subsequent blocks.
- As time passes and more blocks are added, it becomes very unlikely, but not impossible, that a given block will change.

How blockchains grow

From time to time, a miner will process a batch of properly signed transaction requests by placing them in a new block and attaching that block to the current blockchain, thereby *extending* it.

The miner checks the validity of the transactions before placing them in the block, and anyone can subsequently do the same in order to verify the validity of the new chain.

The miner then sends the updated chain to all other miners.

Any miner receiving a valid chain that is longer than the one it currently knows about will discard the current chain and accept the new one as current.

Forking

- A *fork* happens when two different blocks (call them A and B) are created at the same time in different parts of the network. Both will be used to extend the current chain, and both will be propagated to the other miners.
- Propagation across the net takes time. During that time, both new blocks are circulating among the miners. Some miners will receive A first; other will get B first.
- At this point in time, some miners have the new chain ending in block A, whereas others have the new chain ending in block B.
- The chain has thus forked into two parts that differ in their last block.

Reaching consensus

To achieve a single consistent ledger, the miners must somehow agree on which of the new chains to accept as genuine.

There are two different mechanisms in use for doing so.

1. They can run a *consensus protocol* in order to agree among themselves which of the two chains to accept as valid.
2. They can do nothing and hope that the next time one of the forks is extended (either A or B), the new longer chain will reach all miners before another block is added. In this case, all miners will adopt the new chain and discard all shorter chains. This only works if the rate of adding new blocks is sufficiently low.

Transaction commitment

A transaction to a database is said to be *committed* if it is permanently in the system and cannot be rolled back.

If a transaction appears in block A of a fork but not in block B, then the transaction might not be committed to the ledger and could later disappear from the blockchain altogether.

If the miners reach consensus, then the transactions in the agreed-upon block are committed.

But if the miners simply hope agreement happens because one chain overtakes all the others, no one can ever be assured that a transaction has committed.

Bitcoins

Bitcoin history

- Bitcoin was invented by an unknown person writing under the pseudonym Satoshi Nakamoto. A link to his original whitepaper can be found here.
- Bitcoin was the first popular cryptocurrency; Ethereum branched from it and uses slightly different protocols. Many current cryptocurrencies use the same or similar software protocols.
- Blockchains were introduced along with *Bitcoin* and are used to implement it. They address the forking problem by using a mechanism called *proof of work*.

Some properties of Bitcoin

Bitcoins are a kind of virtual digital currency based on cryptography.

- They exist only in the cloud.
- Their supply is limited.
- Like commodities, their value goes up and down depending on market forces.
- They can be used for online transactions without involving a central party, but most users transact with a Bitcoin exchange.
- Bitcoin transactions are (in some circumstances) anonymous and are a favored medium of exchange for illegal transactions.

Bitcoin exchange

Bitcoins are bought and sold on exchanges such as CEX·IO or coinbase, where today's bitcoin price is around $5380 USD.

Bitcoins are volatile. The next slide shows the price for Bitcoin (BTC) since Aug 16, 2010.

Bitcoin Price History Chart

Image from https://www.buybitcoinworldwide.com/price/

Bitcoins in the news

Bitcoins are frequently in the news. Depending on who is talking, they are:

- the transaction medium of the future.
- a threat to national security.
- of use primarily to drug dealers and criminals.
- a good investment.
- the analog of cash on the internet.
- secure because they don't rely on trust in a government.
- a kind of Ponzi scheme that will sooner or later collapse.

Bitcoin Implementation

How to understand Bitcoins

To evaluate these claims, one must first understand how they work.

- Physically, a Bitcoin is an entry in a blockchain as described above.
- A Bitcoin is identified by an *address*, which is a public cryptographic key. The owner holds the corresponding private key in her *Bitcoin wallet*.
- From a user's perspective, Bitcoin is like a mobile app that provides a personal Bitcoin wallet.

Storing the records: Bitcoin

Bitcoin is a cryptocurrency built on a blockchain.

- It is maintained by many *miners*, all of whom own Bitcoins.
- Copies of the blockchain are distributed among the maintainers.
- The copies are stored and controlled by miners.
- In January 2017, the blockchain file for Bitcoin was 100 GB and growing.
- Simply storing the file permanently costs money. Also, transmitting it, or even part of it, on the internet costs both time and money.

Bitcoin miners

- The Bitcoin blockchain is said to be *permissionless* since anyone can become a miner without requiring permission from a central authority.
- A new miner need only obtain a copy of the current blockchain and Bitcoin software in order to participate in the protocol.
- The consensus network is claimed to be the first decentralized peer-to-peer payment system that is powered by its users with no central authority or middlemen.
- Bitcoin proponents make a big deal about eliminating the need for trust. We will see that this is not true in practice, although who must be trusted is different than in traditional financial systems.

Bitcoin creation

- Miners incur costs in storing and maintaining the blockchain.
- To offset these costs, the Bitcoin protocol rewards miners who successfully add blocks to the blockchain with newly-minted Bitcoin. This happens approximately once every 10 minutes. It is the only way that new Bitcoin can be created.
- The total number of Bitcoin that will ever be mined is 21 million. Approximately 80% have already been mined.
- The reward for solving a block was initially 50 Bitcoin. It was halved in 2012 to 25 and again in 2016 to 12.5. It will be halved again every 4 years until all 21 million Bitcoins have been created.
- The last will be mined in approximately 2140.

Controlling rate of blockchain growth

- To slow the rate at which the blockchain is extended, a miner must solve a compute-intensive *puzzle* before adding a block.
- The solution to the puzzle is added to the new block so that others can verify that the miner was indeed authorized to add the block.
- The successful miner receives her reward and a new puzzle is initiated.
- The difficulty of the new puzzle is adjusted periodically to maintain the rate of approximately one solution per 10 minutes, regardless of the number of miners or the amount of compute power they possess.

SHA-256 puzzles

The puzzle that Bitcoin uses is based on the SHA-256 hash function.

It consists of finding a number y (called the *nonce*) such that the SHA-256 hash of the proposed new block together with y yields a hash value (when interpreted as a binary number) that is smaller than the current specified target value.

The smaller the target, the harder the puzzle.

Thus, if the target is 2^{30}, then the hash value must begin with 30 zeros. Because SHA-256 is assumed to be hard to invert, the only known way to solve the puzzle is to try approximately 2^{30} different nonces using that same number of SHA-256 computations.

Bitcoin transactions

- A transaction takes one or more Bitcoins as inputs and produces one or more Bitcoins as outputs. Fractional inputs and outputs are permitted.
- The total value of the input Bitcoins equals the total value of the output Bitcoins.
- The transaction specifies the address(es) at which the output Bitcoin(s) are stored.
- The parties involved in a Bitcoin transaction work together to create the transaction record. The owners of the input Bitcoins must all sign the transaction. The receiving parties must create addresses for the new coins.

Avoiding double spending

A transaction can only be committed to a new block if the source Bitcoins have not already been spent.

The blockchain is examined before a new transaction is accepted to make sure the input addresses control Bitcoins of sufficient value for the transaction. This is how Bitcoins cope with the problem of double spending.

Because of decentralization, it is possible for the same Bitcoin to be used in multiple conflicting transactions, where those transactions are sent to multiple miners.

Committing a transaction

Once a transaction record has been created, it must be entered into the database.

- The transaction creators broadcast the transaction to all miners.
- Each miner enters the transaction into a list of pending transactions.
- Every now and then, a miner solves the current puzzle, incorporates all valid pending transactions into a new block, and extends the blockchain by appending the new block.
- The new blockchain is broadcast to all other miners.

Mine's longer

In the world of Bitcoin, "consensus" means that all miners agree on one version of the blockchain.

- When a Bitcoin miner records a transaction by including it in a new block, the new block may or may not be permanent.
- If a fork happens, one version or the other of the blockchain will eventually grow longer than the other.
- Unless both versions grew at more or less the same time, the one that grew will cause receiving miners to discard the shorter one.

When is my transaction fully committed?

The answer is, "never".

A dishonest miner can go back in time to an earlier block and try to extend the block chain in a new direction, perhaps one that alters or excludes an earlier transaction.

To do so, he must solve a sequence of puzzles before any other miner solves the "current" puzzle.

The further back in time, the more unlikely it is that the dishonest miner can succeed.

So a given transaction becomes more and more secure as more and more new blocks are successfully added to the block chain.

Summary

Summary of the transfer protocol

Here's how Alice transfers a Bitcoin to Bob:

1. Alice creates and signs a transaction request giving the coin to Bob.
2. The transaction is then broadcast to all of the miners.
3. Each miner first verifies the validity of the transaction by using its current most-recent copy of the database.
4. If valid, the miner attempts to create a new certified database incorporating the new transaction (along possibly with others) into the current database.
5. To certify a database requires solving a computationally-intensive puzzle.

Outline of the transfer protxocol (cont.)

6. The puzzle consists of finding a nonce y such that the SHA-256 hash of the database together with y yields a hash value beginning with a long string of 0's.
7. A successful miner broadcasts the new database to all other miners.
8. Each miner upon receiving a new certified database discards all older ones and begins working with the newer one.
9. The system never reaches consensus, but the probability of a certified database being discarded in favor of another decreases exponentially over time.

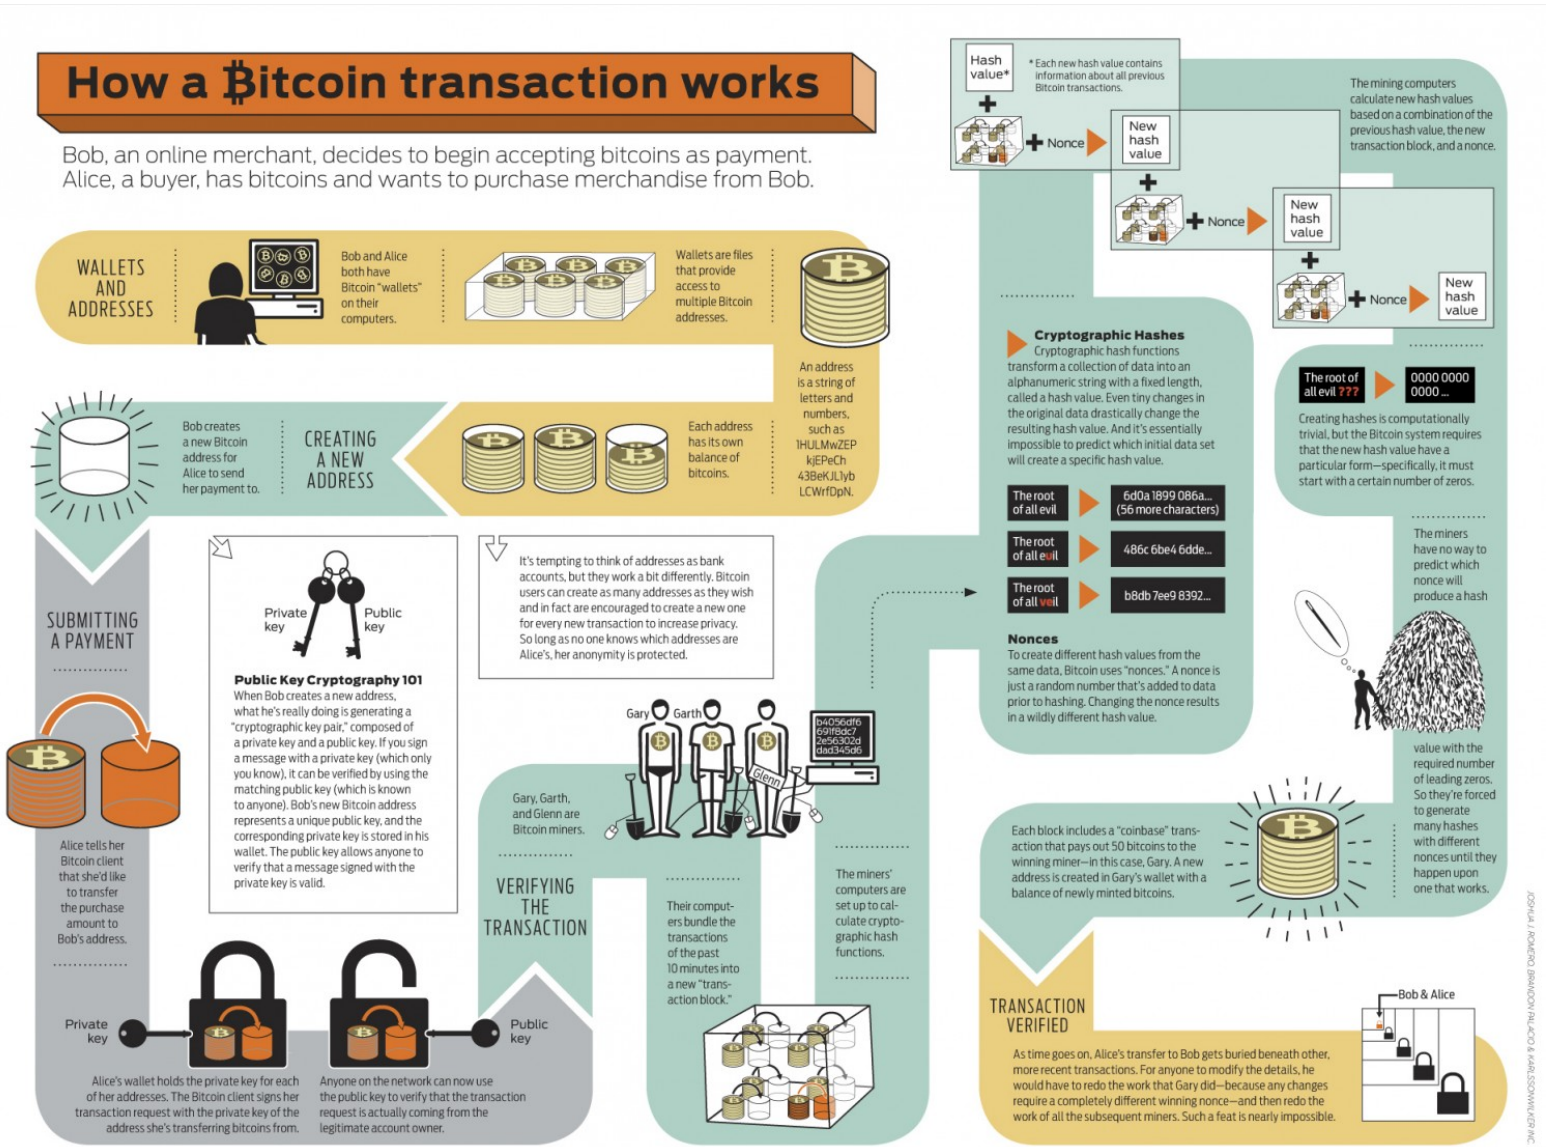

From https://visual.ly/community/infographic/technology/bitcoin-infographic