CPSC 467: Cryptography and Security

Michael J. Fischer

Lecture 25
December 3, 2020

Bitcoin Implementation

Security

Bitcoin Transfer Protocol

Trust

Waste

Bitcoin Implementation

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.

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.

Security

Security

Any miner or mining poll who controls 51% or more of the mining capacity can change past transactions or even rewrite the blockchain entirely. Here's how that can happen.

Imagine a race between two runners, Alice and Bob. Alice starts first but Bob runs slightly faster than Alice. Eventually Bob will pass Alice.

The "51%" attacker privately grows a bogus blockchain as fast as possible. Eventually it will overtake the legitimate one. At that point, the attacker can broadcast the bogus chain to the other miners and will all accept it.

Has this already happened?

People believe some mining pools have already passed the 51% mark, giving them the capability of corrupting the current blockchain. If they haven't done so, its because they deem it in their best interests not to, not because the lack the capability.

Anonymity

Bitcoin transactions are often said to be anonymous.

It's true that they contain only a public key, not the user's personal information.

However, if the key can be linked to an individual, the anonymity is lost.

There are many ways for such linking to take place, e.g., by confiscating the user's wallet, or by monitoring a user's transaction in progress.

Even if the underlying transaction protocol is anonymous, users of a Bitcoin exchange will likely be identified by the exchange in order to enable the conversion of Bitcoins to or from conventional currencies.

Analysis

Why is consensus almost-certainly reached?

- Suppose two miners solve a puzzle simultaneously.
- Both broadcast their versions of the new database D and D'.
- Perhaps half of the miners work on D and half on D'.
- Most miners are likely attempting to incorporate Alice's transaction into a new database.
- Suppose some miner working on D solves the puzzle and sends out the new database D''.
- All miners receiving D'' discard the old D or D' and begin working on D''.
- Now an overwhelming majority of them believe D'' is the current database. They will only change their minds if a yet-longer certified database shows up.

Where's my money?

A good question to ask is, "Where is my money?".

It's obviously in the cloud, but where it is exactly is in the miners' computers.

Security relies on there being many honest miners.

Successful miners are currently rewarded with new Bitcoins, but as time goes on, the rewards are programmed to diminish.

What happens when miners no longer have the incentive to solve the computationally-intensive puzzles?

Other potential problems

There are other potential problems as well.

- What happens if Alice's private signing key gets compromised?
- What happens to Bitcoins that are lost?
- What happens if the puzzle turns out to be not as hard as expected?
- What happens if people turn their attention to a competing digital cash scheme?
- Is this another Ponzi scheme? Why or why not?
- Bitcoins have been compared to gold. Is that comparison valid?

Bitcoin Transfer Protocol

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 protocol (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.

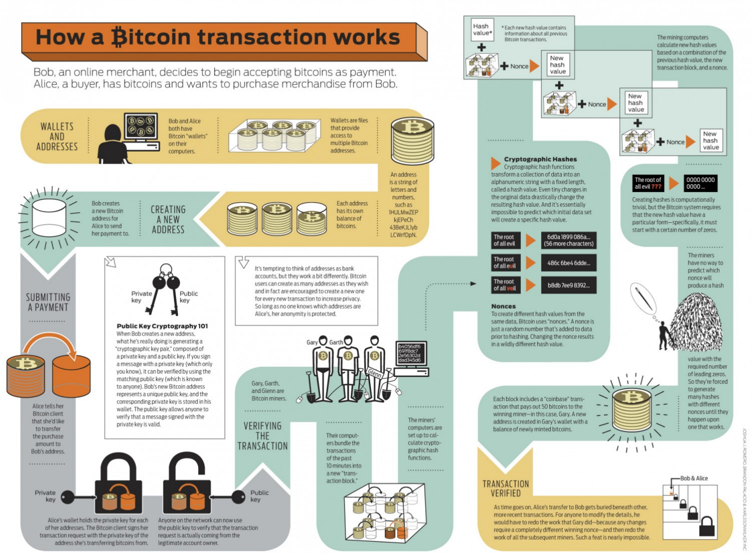

How a ₿itcoin transaction works

Bob, an online merchant, decides to begin accepting bitcoins as payment. Alice, a buyer, has bitcoins and wants to purchase merchandise from Bob.

WALLETS AND ADDRESSES

Bob and Alice both have Bitcoin "wallets" on their computers.

Wallets are files that provide access to multiple Bitcoin addresses.

An address is a string of letters and numbers, such as 1HULMwZEPkjEPeCh43BeKJL1ybLCWrfDpN.

Each address has its own balance of bitcoins.

CREATING A NEW ADDRESS

Bob creates a new Bitcoin address for Alice to send her payment to.

Public Key Cryptography 101
When Bob creates a new address, what he's really doing is generating a "cryptographic key pair," composed of a private key and a public key. If you sign a message with a private key (which only you know), it can be verified by using the matching public key (which is known to anyone). Bob's new Bitcoin address represents a unique public key, and the corresponding private key is stored in his wallet. The public key allows anyone to verify that a message signed with the private key is valid.

It's tempting to think of addresses as bank accounts, but they work a bit differently. Bitcoin users can create as many addresses as they wish and in fact are encouraged to create a new one for every new transaction to increase privacy. So long as no one knows which addresses are Alice's, her anonymity is protected.

SUBMITTING A PAYMENT

Alice tells her Bitcoin client that she'd like to transfer the purchase amount to Bob's address.

Alice's wallet holds the private key for each of her addresses. The Bitcoin client signs her transaction request with the private key of the address she's transferring bitcoins from.

Anyone on the network can now use the public key to verify that the transaction request is actually coming from the legitimate account owner.

VERIFYING THE TRANSACTION

Gary, Garth, and Glenn are Bitcoin miners.

Their computers bundle the transactions of the past 10 minutes into a new "transaction block."

The miners' computers are set up to calculate cryptographic hash functions.

Cryptographic Hashes
Cryptographic hash functions transform a collection of data into an alphanumeric string with a fixed length, called a hash value. Even tiny changes in the original data drastically change the resulting hash value. And it's essentially impossible to predict which initial data set will create a specific hash value.

Nonces
To create different hash values from the same data, Bitcoin uses "nonces." A nonce is just a random number that's added to data prior to hashing. Changing the nonce results in a wildly different hash value.

* Each new hash value contains information about all previous Bitcoin transactions.

The mining computers calculate new hash values based on a combination of the previous hash value, the new transaction block, and a nonce.

Creating hashes is computationally trivial, but the Bitcoin system requires that the new hash value have a particular form—specifically, it must start with a certain number of zeros.

The miners have no way to predict which nonce will produce a hash value with the required number of leading zeros. So they're forced to generate many hashes with different nonces until they happen upon one that works.

Each block includes a "coinbase" transaction that pays out 50 bitcoins to the winning miner—in this case, Gary. A new address is created in Gary's wallet with a balance of newly minted bitcoins.

TRANSACTION VERIFIED

As time goes on, Alice's transfer to Bob gets buried beneath other, more recent transactions. For anyone to modify the details, he would have to redo the work that Gary did—because any changes require a completely different winning nonce—and then redo the work of all the subsequent miners. Such a feat is nearly impossible.

JOSHUA J. ROMERO, BRANDON PALACIO & KARLSSONWILKER INC.

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

Trust

We Have to Trust the Institutions Around Us

Multiple organizations are involved in making our monetary system work:

- Banks
- Charge-card services
- PayPal
- The Federal Reserve
- Congress and the treasury department of the U.S.A.

We expect them to respect our privacy and not steal from us.

Cryptocurrency advocates believe we should not trust them because they can regulate things, enforce laws, and manipulate the currency.

We Have Permanent Records

We trust our governments to keep accurate permanent public records of:

- Births and deaths
- Land Sales
- Citizenship and voters registration
- Taxes owed and paid
- Military service and status

These records have been made, preserved, and made accessible for hundreds of years—all without using cryptography.

Do you trust these systems? Why or why not?
Would you trust a blockchain system more? Why or why not?

We Have Business Records

We trust our large companies to keep accurate records of:

- Investments
- Cash in the bank
- Telephone numbers (paper and online)
- Accounts and payments
- Academic records and degrees

These ordinary business services are the foundation of our way of life.

The point is, we cannot live without trusting third parties. Trust is the basis of all civilization.

We Rely on the Internet

In doing so, we rely on governments and big companies.

- ISP's and Internet Exchanges, some operated by the government, others by universities or businesses.
- Infrastructure maintained by big companies and open to spying and government control.
- Domain name servers.
- Internet access points that connect subnets. Many are owned or controlled by governments.

The proper functioning of blockchains assume that the Internet is trustworthy, open, and free from government intervention. Not so!

Some people do not trust anybody.

Libertarians, like many blockchain proponents, distrust authority, state power, and “big business”.

One of the claims about blockchain is that it eliminates the need for trusting third parties.

Blockchain does make it more difficult to pursue criminals or regulate shady practices!

The protocols have been designed to replace trusted third parties by collective trust.

Collective trust

In place of trusting a third party (bank, government, business), blockchains are based on collective trust in the large set of people who maintain the blockchain. The requirements for enabling that trust are:

- The parties maintaining the cryptocurrency must be independent, and no party should hold a majority of the power.
- All maintainers must have a stake in the integrity of the blockchain.
- The majority of maintainers will act in their own self interest by following the rules that support the integrity of the blockchain.

Why trust the miners?

In both Bitcoin and Ethereum, miners are collectively trusted.

- Miners invest money (electricity, computer hardware, big buildings) in the mining process. They are (eventually, probably) rewarded for solving a puzzle by getting some bitcoin.
- The self-interest of these miners is to maintain the value of that bitcoin.
- These same people commit transactions (and earn fees for it), manage the blockchain, maintain its validity, and store all or part of it.
- The wasteful mining mechanism makes it expensive to be a miner. This helps to prevent a single miner or group from having enough control to corrupt the blockchain.

Why these assumptions may not hold in the real world

- Mining is not really open to an individual any more. It is only large groups (companies, governments) that can succeed often enough to make it worth the investment.
- A large enough party (e.g., China) that is also strictly authoritarian does not have the same motivations as the individual miners.
- In much of the world, individuals have no choice about obeying their government.
- The internet and access to it is controlled by governments and large companies.

Cryptocurrencies do not eliminate third parties

The blockchain only keeps track of the record of all transactions. By itself, that is not enough to support a currency system.

- There must be a reliable way to locate other miners in order to interact with them and stay synchronized.
- There must be a reliable way to locate the full blockchain if you have been offline for a while.
- Converting to and from real currencies requires trusting an exchange.

Trusted third parties within the Bitcoin community are used for all of these purposes.

Waste

Reality: Protocols

The protocols that implement block chain are far from perfect.

- They can have, and have had, bugs. They are written, maintained, and managed by true-believers acting in good faith. They have bugs anyway, some disastrous. (See the Bitcoin and Ethereum articles.)
- There are a variety of known attacks on the systems.
- Bitcoin enterprises depend on the internet being reliable, open, and not under anyone's control.
- From an environmental point of view, Bitcoin and Ethereum are irresponsible wastes of energy and resources.

Reality: Environmental Impact

Environmentalists believe in avoiding wasteful use of resources.

- By design, earning a Bitcoin is compute-intensive. The computations are totally useless except for competing for Bitcoins. They do not produce useful goods or knowledge.
- If another miner posts a solution a split second before you do, you lose. Any energy you have used trying to solving that puzzle is totally wasted.
- This is repeated millions of times for every bitcoin mined.
- More miners equals more waste and a higher cost for mining each bitcoin.
- As more mining computers are added to the pool, the rate of waste per bitcoin mined grows.

It's a really BIG waste

These statistics involve only Bitcoin. Ethereum and other currencies add to the numbers.

- Whole warehouses of specially designed computers work 24/7 trying to solve these puzzles.
- Currently, worldwide mining uses about 0.5% of the world's electricity. This is more than is used in the entire Czech Republic and almost as much as in Ireland.
- Miners are looking for cheap electricity. Some are investing in renewable energy sources.
- The cheap coal-powered plants in Sichuan and government subsidy of fossil-fuel in Canada make those sites attractive.