How Bitcoin Works

Sam Patterson
October 2014

Abstract. Bitcoin is a technological and monetary innovation that allows for the peer to peer transfer of value online without trusting third parties. The system uses a distributed public ledger to reach consensus among network participants about the legitimacy of transactions. This ledger isn't denominated with fiat currencies, but instead uses its own units called bitcoins. The system relies on various elements of cryptography to ensure that this public ledger is accurate instead of trusting a central institution. The code is open-source and the network is decentralized, resulting in low barriers to entry and censorship-resistance. New bitcoins are issued by a predictable algorithm that is intended to replicate mining for a scare resource.

Double-spending problem

Bitcoin is both a currency and a payment network. While most attention is paid to the currency, the technological innovation that makes Bitcoin unique is in its payment network. To understand why Bitcoin is valuable, we must understand the context of digital payments previous to this innovation.

Bitcoin is not the first digital currency. The majority of the traditional fiat money supply isn't physically printed, but instead is represented digitally. Services such as PayPal, WebMoney, or Liberty Reserve allowed individuals to send and receive money digitally.

All of these digital currencies were centralized, meaning they have an issuing body and organization who is able to control transactions on the network. This was necessary because of the double-spending problem. Because data is trivially easy to replicate, using digital currency has one inherent drawback. The recipient cannot be certain their digital representation of value hasn't been duplicated, and spent elsewhere. This conundrum is called the double-spending problem, and it was dealt with by having a central organization maintain the ledger balances of the currencies. This meant that this organization—whether a governmental body or a private company—would be able to directly control the balances and transactions occurring on the network.

Bitcoin solves the double-spending problem without relying on a central organization to maintain the ledger. Instead, it makes the transaction ledger public, distributes it among network participants, and relies on consensus to ensure that balances are accurate. This distributed ledger—called the blockchain—contains the history of every transaction which has ever occurred on the network, allowing any participant to know the exact ledger balances at any moment and verify if proposed transactions are legitimate. Once the first transaction has been recorded in the blockchain, any attempt at double-spending would be automatically rejected by network participants as invalid.

Public key cryptography

Public key cryptography is used for the digital signatures. Public key cryptography is a technology that allows for secure digital signatures by creating a pair of keys, a private key and a corresponding public key. The keys are seemingly random strings of alphanumeric characters that correspond to each other. Using public key cryptography allows a user to sign messages with their private key, and for other users to verify that the message's signature is accurate by checking it against the corresponding public key. As an example, consider sending a digitally signed email message. I would first distribute my

public key to the recipients. I would then digitally sign the message with my private key, and send the email message. The recipients would then check the digital signature against the public key to verify accuracy. Only the owner of the private key would be able to create a signature that can be verified by the public key.

Bitcoin uses an elliptical curve—specifically the secp256k1 curve—in order to create the key pair. The public key in the pair is used to create the Bitcoin address. These addresses are identifiers for where the bitcoins reside in the blockchain. The private key in the pair that corresponds to the public address is then used as proof of ownership over that entry in the ledger. In this sense, owning bitcoins in a public address actually means having control of the corresponding private key.

In order for network participants (commonly referred to individually as nodes) to determine that a user is legitimately in control over a certain entry in the public ledger, the node which initiates a transaction will sign a message with their private key stating that they want to send bitcoins from their public address to another public address. If other nodes verify the signature is valid and the amount is valid, that transaction is then put into a type of queue—called a block—until the block is entered into the blockchain. If the signature doesn't correspond to the public address, then the transaction is rejected by local nodes and will not enter into a block. New blocks are added to the blockchain approximately every ten minutes, which is how the blockchain derives its name—it is a chronological chain of blocks containing all the transactions on the network back to the first transaction, dubbed the genesis block. As of mid-October 2014, there are approximately 326,000 blocks in the blockchain.

Proof-of-work

There are cryptographic rules determining how blocks are created and entered into the blockchain. These are necessary to make sure that blocks are created in chronological order, and at a somewhat regular time interval. If blocks didn't follow these rules, it would be nearly impossible for network participants to agree to a single version of the blockchain; transactions with different blocks in the blockchain wouldn't be compatible, and the entire network would be unusable.

This is solved by incorporating proof-of-work for those maintaining the blockchain. Proof-of-work means that creating a block is intentionally made more computationally difficult, thus taking more time. This is achieved by only allowing a node to issue a block after it has done extensive cryptographic hashing. Hashing is a one-way function that takes arbitrary information and changes it to a fixed size and format. Creating a hash of data is simple, but recreating the data from only the hash is difficult.

Bitcoin uses hashing as a proof-of-work, specifically the SHA-256 hash function. In order for a node to broadcast a block to the network for confirmation and acceptance into the blockchain, it must also broadcast proof that they did the extensive hashing necessary to find the solution to an arbitrary puzzle. All nodes who are attempting to add a block into the blockchain are simultaneously hashing SHA-256 in search for an output that begins with a certain number of zeros. Since the output of the hash function is essentially random, nodes rapidly iterate through values in search of this elusive output. The first node to find the correct value broadcasts the answer along with the block, and if other nodes verify the answer is correct, the block is entered into the blockchain. This is called “finding a block.”

Since more hashing power can easily be added to the network, Bitcoin uses an algorithm to ensure that the blocks are still created approximately ten minutes apart. Hashing difficult is increased automatically by adding more zeros onto the hash output required, thus making finding the value computationally more difficult.

It's possible for two blocks to be found nearly simultaneously. In this case, the network can temporarily be said to have two competing blockchains. This is solved as soon as one of the blockchains finds the

next block. The code automatically considers the longest blockchain to be valid, so miners which were working on the shorter blockchain will switch over immediately. The block which didn't enter the longest chain is called an orphaned block and is abandoned.

Mining

Nodes which participate in hashing in order to create blocks and maintain the blockchain are called miners. Miners are essential to keeping the network functioning, and their efforts are rewarded by receiving a block reward if they successfully find a block. This reward is in bitcoins. Originally, a block reward was 50 bitcoins, but the reward is halved approximately every 4 years. Currently at 25 bitcoins, the block reward is projected to halve again in late 2016. This continual halving is meant to replicate the mining of a scare resource. Apart from the block rewards, miners also receive small transaction fees for verifying transactions on the network.

Because of these incentives, the hashing power of the Bitcoin network is substantial. As of mid-October 2014, the network is performing 261 million billion hashes every second.

When the Bitcoin network was young, miners used normal CPUs in desktops or laptops to hash for coins. As more users joined the network, some realized they could hash more quickly by using GPUs instead. Driven partially by the increase in value, hardware eventually was built with the sole purpose of hashing SHA-256 in order to mine bitcoin. Application specific integrated circuit (ASIC) hardware is now necessary to do any effective mining on the network.

Because there are so many miners, the chances of individually finding a block are small. To improve their chances, miners have grouped together and created mining pools. These pools allow everyone to coordinate their hashing power and increase their chances for finding a block. The block reward is then split based on how much hashing power a miner contributed to the pool, minus a fee if the operator of the pool charges for the service.

Mining pools have become the standard way for miners to search for block rewards. This has raised some concerns around centralization of the network. If one pool of miners reached 51% of the network they would be able to negatively impact the creation of new blocks, though it would appear to be against their own interests to do so. This “51% attack” has never occurred on the Bitcoin network.

This method of issuing currency is substantially different than fiat currency models. Instead of being directly controlled by a central issuing organization, the Bitcoin network creates new currency based on a predictable algorithm, and in a decentralized fashion. Any network participant can become a miner and attempt to find a block reward. Arbitrary increases in the money supply aren't possible, making inflation unlikely. While Bitcoin is too young to make any definitive statements on how it functions as a monetary system, it appears that because of the limited supply that increasing demand will effectively make it a deflationary currency.

The network

Only a small portion of network participants are miners. Most are nodes in order to use the payment network. Becoming a node means downloading and running software locally on their computer that connects a user to the network, and downloading the blockchain in order to be up to date with the network.

The network is open and decentralized. The code is open source and not controlled by any organization. As a result, the only barriers to entry into the network are technical ones; a user must have some hard drive space and an internet connection. Unlike a credit card network, or the SWIFT system for banking transfers, there are no fees to apply and no applications at all to use the network.

The code is not fixed, but changes over time based on the consensus of those running nodes.

Developers suggest changes to the code, and if those changes are adopted by the majority of nodes then it becomes a part of the core code. If nodes don't agree with the changes, they won't update their code. In this sense, the network is highly resistant to malicious code changes; nodes can see the code and wouldn't accept new code that is harmful. It also allows for new innovations to be widely adopted if they are clearly beneficial.

Estimating the size of the network is difficult. One reason this is difficult is that a single user typically has numerous bitcoin addresses. Many bitcoin wallets now automatically generate new addresses for each transaction, which gives the user more privacy. If an address is reused consistently, an outside observer could gain knowledge of a users entire financial history via the blockchain. Estimates tend to vary from 500,000 users to several million worldwide. There are approximately 70,000 transactions per day on the blockchain.

Wallets

As mentioned earlier, owning bitcoins means having control of the private key that gives access to an entry (or multiple entries) on the blockchain. The private key is a small piece of data comprising of a string of letters and numbers. Keeping that private key secure is essential. A lost private key means those funds are forever unspendable. A stolen private key will almost certainly lead to the loss of the bitcoins, since the attacker will use the key to transfer funds to an address under their control.

There are numerous implementations of software to manage users private keys and connect to the network. These are called bitcoin clients or wallets. These wallets have varying features and levels of security. The most secure wallets are actually not software at all; private keys can be physically printed onto paper (or other materials) and cannot be stolen electronically. Cold wallets also offer strong security. A cold wallet is a method for storing private keys offline where they cannot easily be attacked.

Less secure wallets offer more convenience. These are hot wallets, meaning the private keys are held on an internet accessible device. Currently the blockchain stands at approximately 23 gigabytes, which deters many users from keeping a full copy on their computers. This has led to bitcoin clients that are able to operate by only downloading a portion of the blockchain, or even simply connecting to other nodes that hold the blockchain for them. Mobile applications are becoming an increasingly common way to manage bitcoin wallets as well.

All bitcoin clients have some core functionality in common. They hold the users private keys, sign transactions and broadcast them to the network, and monitor the blockchain to keep funds updated. To send bitcoins, the user inputs the public address of the recipient. This public address is a string of 26-34 alphanumeric characters. Since it is inconvenient to type this, QR codes are commonly used on mobile devices to rapidly import the bitcoin address and payment amount for transactions. Once the public address is entered, the user selects how many bitcoins they wish to send and confirms their transaction. The client then creates a message with these details, signs it, and broadcasts it out to the network.

Bitcoin Improvement Proposals

Bitcoin is open source, and the code is not stagnant. New features to improve the code are called Bitcoin Improvement Proposals (BIPs). Since the first BIP in 2011, there have been over 70 new BIPs. Many never became widely accepted by the network, but there are some notable improvements.

M-of-N transactions, also called multisignature or multisig, requires multiple parties to agree to a transaction before it can be completed. Instead of a sole user having control over the bitcoins, ownership is now joint between any number of parties. One of the common uses for multisig is a 2-of-3 transaction, meaning that two of three parties must sign a transaction for it to be valid.

Using 2-of-3 allows for a secure escrow system. The three parties are the buyer, seller, and a trusted

third party. Instead of sending bitcoins directly to the seller, the buyer will send funds to the multisig address, which is created with the public keys of the three parties. The seller can then ship the product or otherwise provide a service, and then sign one of the two signatures needed to release funds from the multisig to himself. If the buyer receives the good or service, and is content, they can then sign the second signature, and the funds will then be released from the multisig to the seller. If the seller never shipped the product, or otherwise tries to scam the buyer, the buyer can simply work with the trusted third party to sign two signatures to return the funds from the multisig back to the buyer. If the buyer receives the product but doesn't sign to release funds, the seller can also work with the third party to obtain the necessary two signatures to release funds from the multisig to the seller.

Hierarchical deterministic (HD) wallets are another notable improvement. Instead of creating a new bitcoin key pair for each address, a HD wallet creates a single master key which can more easily create new key pairs that are also controlled by the master key. With an HD wallet, you can automatically create a new key pair for each transaction, but not have to worry about maintaining the private keys for all keys pairs, only the master private key is needed. Also, as part of BIP 32, HD wallets can be backed up by using a random string of 12 words, called a seed.

BIP 38 created private keys that could be password protected at the time of generation. BIP 70 created a payment protocol that allows merchants to offer more options for customers using bitcoin, such as messages noting payment received and the creation of refund addresses.

Altcoins

Bitcoin was the first digital currency to use a blockchain system. These systems are now called cryptocurrencies. After Bitcoin was started in early 2009, many cryptocurrencies have been created. Cryptocurrencies other than Bitcoin are commonly referred to as altcoins.

Almost all of these are Bitcoin clones, typically with some minor changes to the code. It is debatable if these changes are significant enough for the new currencies to overcome the network effects of the Bitcoin network, which is the largest cryptocurrency network by a considerable margin. Many of these altcoins are created solely for speculative purposes, and disappear quickly. Some altcoins are used as a testing ground to try new features that are considered too risky to test with the Bitcoin network itself.

It's important to note that blockchain technology isn't limited to managing a currency. While the first application was Bitcoin, other blockchains have since been created to manage other ledgers. Namecoin is a notable example, which maintains a blockchain that records arbitrary names, as an alternative to Domain Name System (DNS) and other centralized identity systems.

Conclusion

Bitcoin's distributed public ledger and foundation on cryptography allow for the peer to peer transaction of value without trusting in traditional institutions.

Banks and payment processors aren't necessary in the Bitcoin ecosystem; the trade is peer to peer and no third party has to be trusted with the funds. Central banks aren't necessary; the issuance of new currency is determined by a predictable algorithm. Governmental bodies aren't necessary; the transaction isn't secured by the rule of law but instead by mathematical laws in the cryptography.

This new system has low barriers to entry and cannot be manipulated by any central organization. The cost for sending bitcoins is significantly lower than existing systems, and it is also substantially faster. If the technological infrastructure in developing nations grows, Bitcoin could give previously unbanked individuals the opportunity to connect to a global trade network without needing to pay fees or submit applications.

Bitcoin continues to innovate and implement new features, such as multisig, that open up applications

for money that weren't previously possible. Because it is open source, the options for this programmable money are only limited by what the community can come to consensus on.