Bitcoin

Digital peer-to-peer cryptocurrency

Αλέξανδρος Αφεντούλης
el06150 [at] mail.ntua.gr

Κοινότητα Ελεύθερου Λογισμικού ΕΜΠ

free and open source software community / national technical university of athens

Cryptography

Bitcoin is heavily using crypto :

- Digital signatures – ECDSA is used, an elliptic curve variant of DSA
- Hashing – SHA-256 and RIPEMD-160 are used

If we wanna make a distributed peer-to-peer currency with no central authority, we still have to trust someone...

… let it be crypto!

Digital Signatures

- Digital signatures is an application of asymmetric cryptography
- Prove the authenticity (and integrity) of a piece of data.
- We use them every day in HTTPS, in software packages and more.
- Alice owns some data. Creates a pair of public and private key. Digitally signs data with her private key. Gives the data along with its digital signature to Bob.
- Bob knows Alice's public key. Bob is able to verify that data is digitally signed by Alice.

Digital Signatures

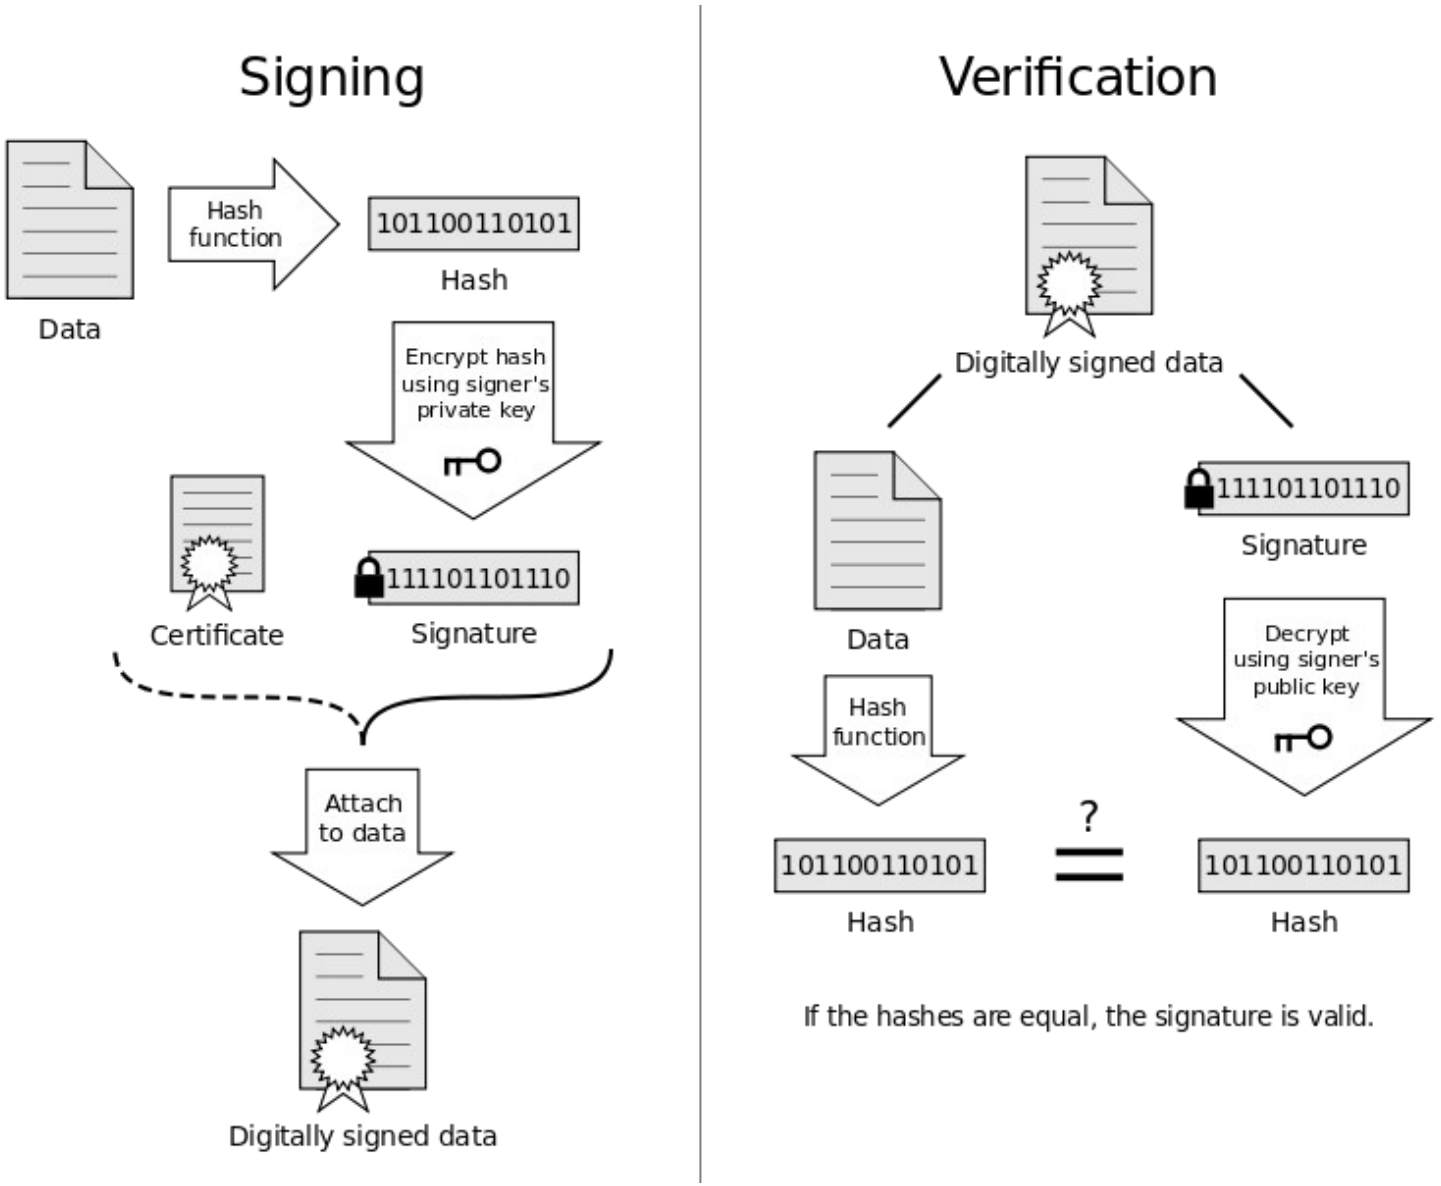

Digital signatures & elliptic curve crypto

- In Bitcoin, transactions are signed using ECDSA, i.e. elliptic curve DSA.
- Elliptic curve crypto is based on the difficulty of solving the discrete logarithm problem over a certain elliptic curve.
- ECC's main advantage is considered to be smaller key size while maintaining the same level of security with crypto based in prime factorization.
- Bitcoin is using Secp256k1 elliptic curve.

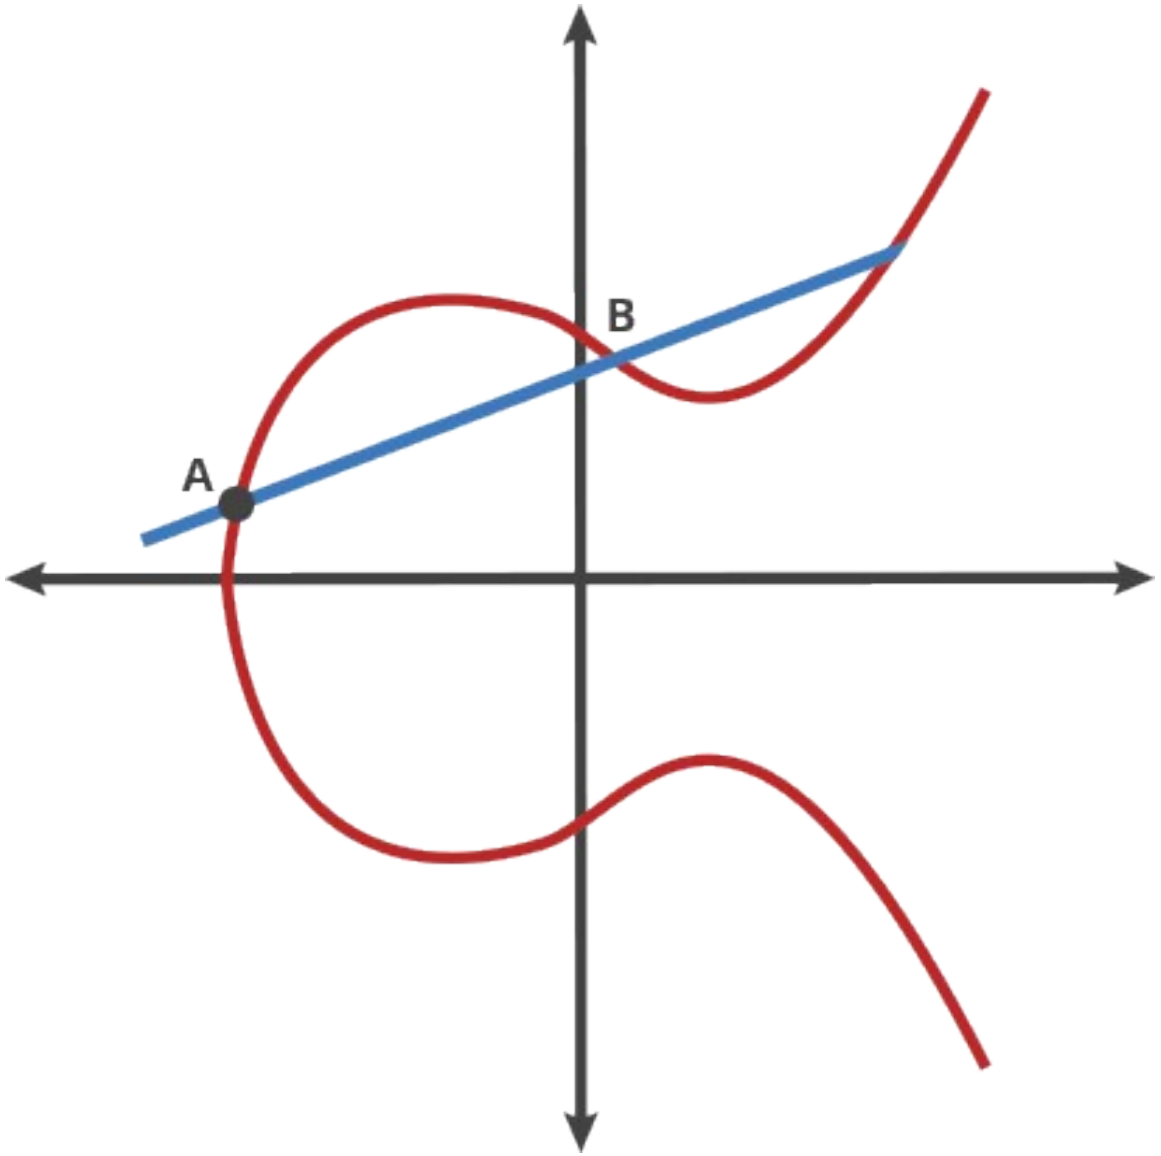
A
B

Hashing

- Hashing functions take an arbitrary block of data as input and return an output of fixed length.
- Have the following properties:
  - output is easily computed for any input
  - if input is slightly modified output will be completely different
  - we can't (easily) find a collision, i.e. same output for different inputs.
- Hashing functions are widely used in crypto, but also in other applications that require message integrity.
- Bitcoin uses SHA-256 (outputs 256 bits) and RIPEMD-160 (outputs 160 bits)

Monetary problems

Bitcoin -as any currency- has to solve some problems :

- Double-spending, Alice cannot send the same bitcoin to Bob and Charlie at the same time
- Forgery, Alice cannot produce arbitrary amount of bitcoins, a hard problem especially when currency is digital
- Validity of transactions, someone somehow has to guarantee the validity of the transactions, especially in a decentralized network
- Inflation?

Bitcoin

- Introduced by Satoshi Nakamoto back in 2009 with the relevant paper. Nakamoto also wrote an initial open source implementation of the idea in C++.
- Bitcoin is a protocol which enables the existence of distributed peer-to-peer digital currency. Transactions are almost free, almost instant and truly irreversible.
- Bitcoin is secured by cryptography
- Bitcoin is open source and nowadays developed by a community
- Has gained a lot attention by a variety of people globally
- Also, is now considered valuable...

Bitcoin - addresses

- Bitcoin addresses can be seen as the “bank account” of Alice. Alice's bitcoin address is where someone will send bitcoins to Alice.
- Alice can create arbitrary number of bitcoin addresses. Multiple bitcoin addresses consist a bitcoin wallet.
- Essentially bitcoin addresses correspond to the public part, of a public-private key pair. Alice creates such key pairs at will.
- A bitcoin address is in fact the hash of a ECDSA public key.

Elliptic-Curve Public Key to BTC Address conversion

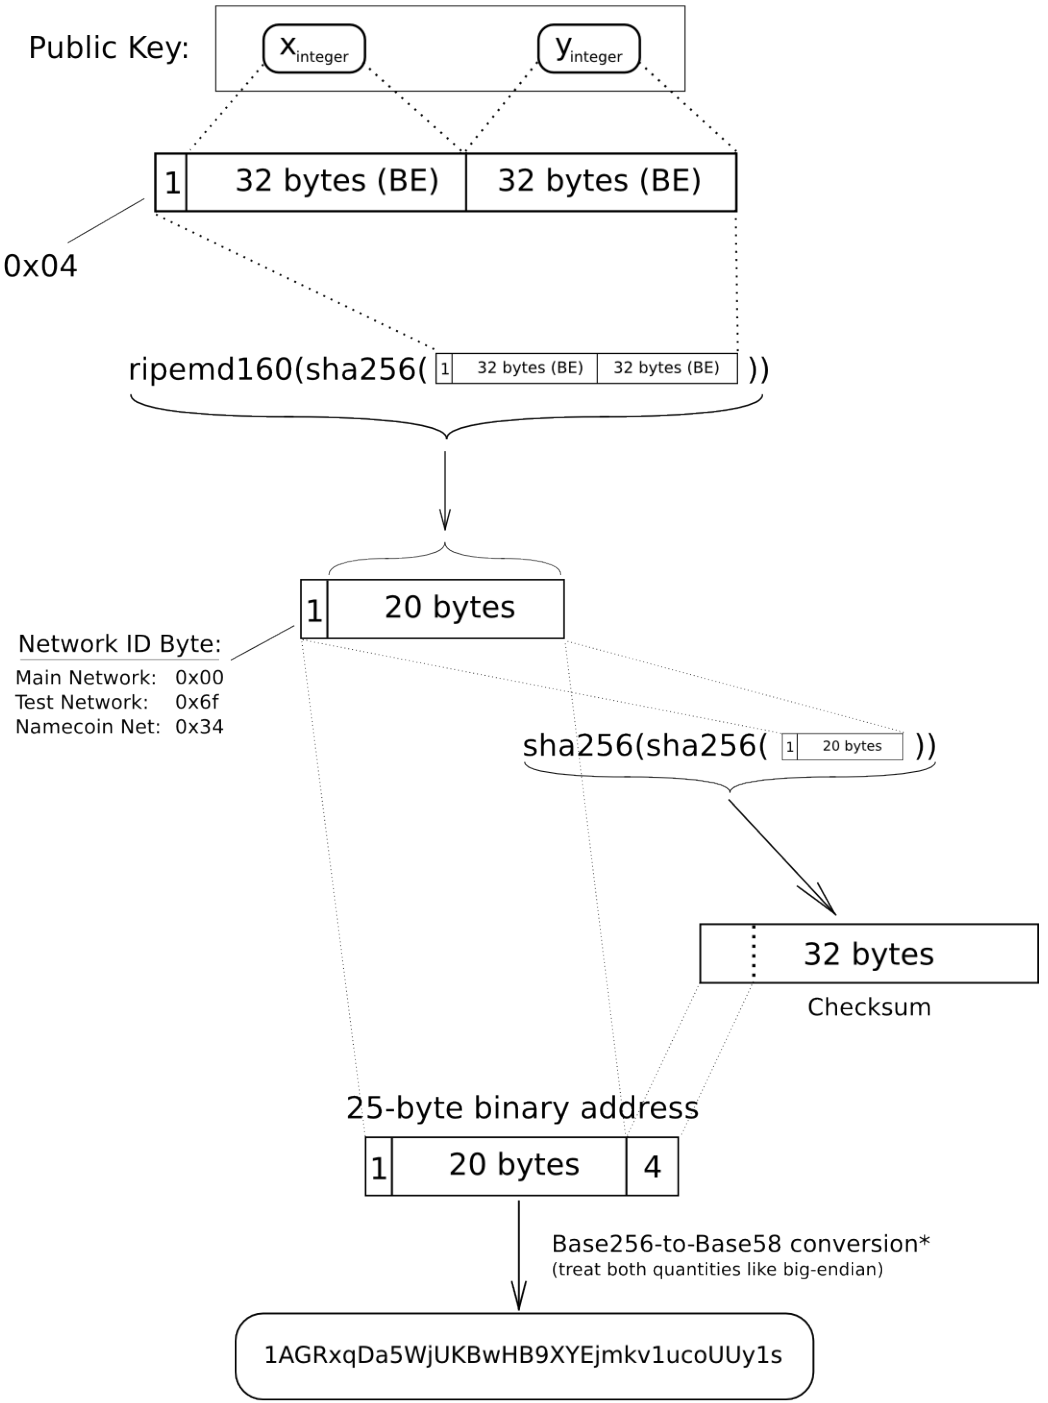

Bitcoin - transactions

- A bitcoin transaction is a public stament that some bitcoins belonging to Alice's bitcoin address are now belonging to Bob's bitcoin address.
- Every transaction is digitally signed with sender's private key, the pair of the public key which holds the coins.
- Every transaction is broadcasted through a network of peers.
- Each transaction will have to be validated by the so called 'miners' (we'll talk later on about mining). Transactions are validated in a variable time window of some minutes.
- Transactions are grouped in blocks.
- As long as a transaction is validated, it is part of the blockchain.

Bitcoin - transactions

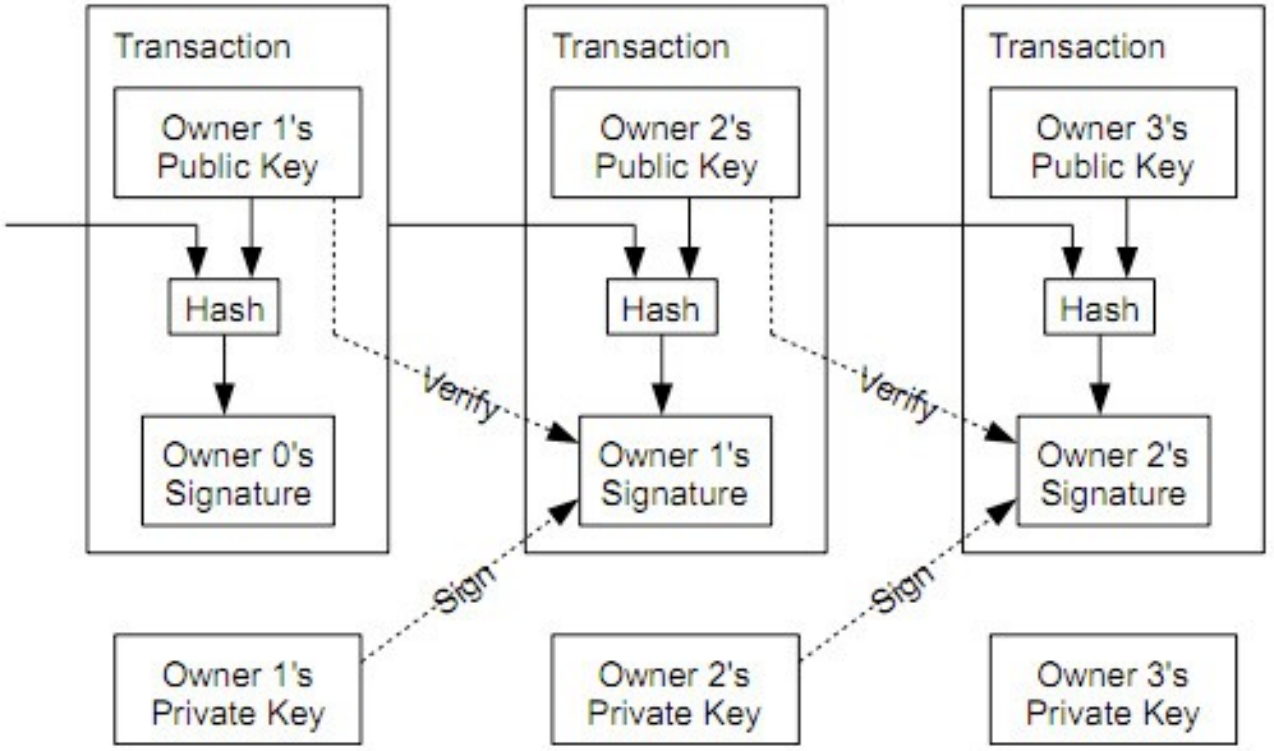

Bitcoin – blocks & blockchain

- Miners are trying to find bitcoin blocks which will include transactions waiting to be validated.
- A block contains an amount of bitcoin transactions which were at some point of time validated.
- Every block, apart from transactions, contains a unique identification of the previously found block.
- Approximately every 10 minutes a block is found.
- Thus a chain of blocks is created, the blockchain. Blockchain is essentially the complete history of every single transaction ever happened in the bitcoin network .
- Blockchain is the serial binding of many blocks.

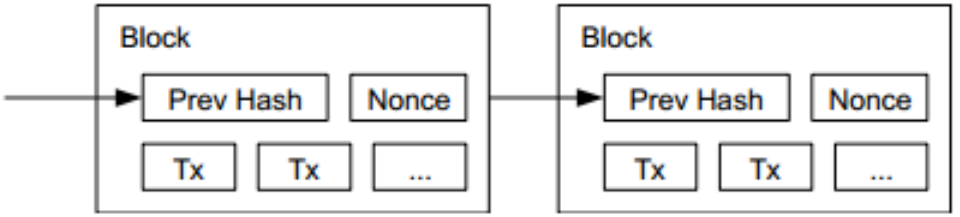

Bitcoin - Mining

- Mining is the procedure of finding a bitcoin block.
- Miners try to solve a hard problem. Finding a solution is hard but any solution can be instantly verified by the whole network.
- Miner takes as input an amount of broadcasted transactions waiting to be validated and a unique identification of the last discovered block.
- The problem is to find a hash of the above input with a specific number of leading '0'. The more zeros, the more difficult is the problem. Difficulty is adjusted by the network, in order to have a solution every 10 minutes independently of the network's total hashrate.
- As long as a miner finds a solution, they broadcast it, network verifies it and the block is added to blockchain. Transactions in that block are then considered valid.
- Finding a solution serves as a proof of work. One has to spend resources and time to find a block. Remember a block contains a part of bitcoin transaction's history. An adversary must have more than 50% of network's hashing rate to change the history (create another version of blockchain).

Bitcoin - Mining & Rewards

- Mining validates transactions and secures bitcoin, thus it's very important.
- Mining is a resource expensive procedure.
- As an incentive to keep mining and keep the network stable, miners finding a block are rewarded with an amount of bitcoins.
- Esentially, miner finding a block is allowed to perform a special transaction, sending to himself a predefined amount of bitcoins out of nowhere.
- That's how bitcoins are created.
- Mining reward is currently at 25BTC. Reward is halving every 210.000 blocks, so eventually only 21 million btc can ever be found
- This was taken as measure against inflation. Remember that in Bitcoin there is no central authority printing money. This is subject of an ongoing debate.
- Miners also get the fee for every transaction they include in a block.

Bitcoin & the generals' problem

- Bitcoin is an elegant solution to the Byzantine Generals' Problem
- The problem in bitcoin, is how can a network of entities that don't trust each other mutually agree on a certain thing
- This thing that bitcoin peers agree on is the blockchain, the entire history of bitcoin transactions.
- Every peer has a local copy of the blockchain and needs to contantly update it with new transactions. That's essential in order to know how many coins each address has.
- Bitcoin manages to solve the problem by introducing the proof of work, i.e. hashing (and by making every transaction public).
- Each miner repeatedly hashes queued transactions until they find a valid solution which will be broadcasted. Then the whole network agrees to rebase their blockchain history with this block included.

Bitcoin - Clients

- A bitcoin client is needed to perform bitcoin transactions
- Start here : https://bitcoin.org
- 'Bitcoin core' is the original client, stable and secure, but harvest more resources since it downloads the whole blockchain
- Electrum is a lightweight client. Uses a set of remote servers which hold the entire blockchain. Still keys are kept and transactions are created client-side. http://electrum.org/
- Web wallets must be avoided. Bitcoin gives the advantage of having your coins locally, there is no reason to give them to the “cloud”.

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.

CREATING A NEW ADDRESS

Each address has its own balance of bitcoins.

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

SUBMITTING A PAYMENT

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

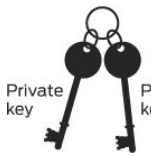

Public Key Cryptography 101
When Bob creates a new address, what he's really doing is generating a "cryptographic key pair," composed of a private key and a public key. If you sign a message with a private key (which only you know), it can be verified by using the matching public key (which is known 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.

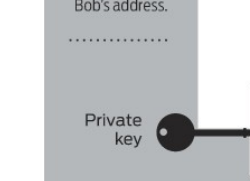

Alice's wallet holds the private key for each of her addresses. The Bitcoin client signs her transaction request with the private key of the address she's transferring bitcoins from.

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."

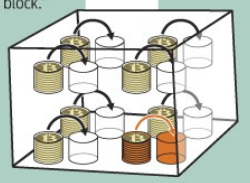

The miners' computers are set up to calculate cryptographic hash functions.

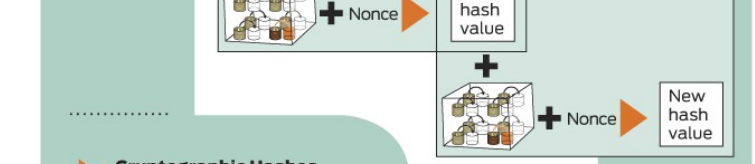

The mining computers calculate new hash values based on a combination of the previous hash value, the new transaction block, and a nonce.

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.

The root of all evil	6d0a 1899 086a... (56 more characters)
The root of all evil	486c 6be4 6dde...
The root of all veil	b8db 7ee9 8392...

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.

The root of all evil ???	0000 0000 0000 ...

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.

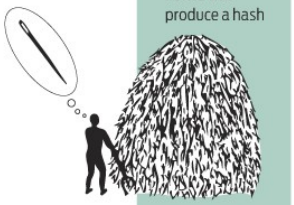

The miners have no way to predict which nonce will produce a hash value with the required number of leading zeros. So they're forced to generate many hashes with different nonces until they happen upon one that works.

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

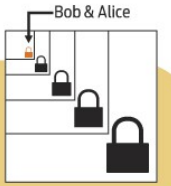

As time goes on, Alice's transfer to Bob gets buried beneath other, more recent transactions. For anyone to modify the details, he would have to redo the work that Gary did—because any changes require a completely different winning nonce—and then redo the work of all the subsequent miners. Such a feat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

JOSHUA J. ROMERO, BRANDON PALACIO & KARL SOWYLKER INC.

Bitcoin - What's so great about it?

- Peers do not trust each other, yet are able to reach a consensus
- Transactions cannot be blocked (e.g. banks, paypal, visa etc)
- Transactions are almost instant, have very low fees, are independent of sender's/receiver's location, contrary to banking system
- Every peer has full control of its coins, no entity can "freeze" them
- Multi-signature addresses and transactions. Funds can be transferred when n of m keys are used.
- We can create bitcoin keys/addresses offline, then print them on paper, keeping the coins in "cold storage".
- Bitcoin is not anonymous, actually every transaction is public. Bitcoin can under certain conditions be more anonymous than traditional banking, but certainly less anonymous than cash.

Please gimme moar

- https://en.wikipedia.org/wiki/Elliptic_curve_cryptography
- http://blog.cloudflare.com/a-relatively-easy-to-understand-primer-on-elliptic-curve-cryptography
  - djb's curve25519 : http://cr.yp.to/ecdh.html
  - agl's curve25519-donna : https://github.com/agl/curve25519-donna
- https://en.wikipedia.org/wiki/ECDSA
- https://en.wikipedia.org/wiki/SHA-2
- https://en.bitcoin.it/wiki/Protocol_specification
- https://github.com/bitcoin/bips
- https://en.wikipedia.org/wiki/Two_Generals%27_Problem
- https://en.wikipedia.org/wiki/Byzantine_fault_tolerance
- https://github.com/spesmilo/electrum (client in python)