A Technical Introduction to Bitcoin

Niklas Fors, 2018-02-20

Market Price (USD)

$10,503.30

Bitcoin

- ***Decentralized*** digital currency
  - Anyone can be part of the network
- **Global distributed ledger** called ***blockchain***

First Appearance

- **Bitcoin: A Peer-to-Peer Electronic Cash System**
  by Satoshi Nakamoto, November 2008
- First implementation: January 2009

Centralized vs decentralized

Centralized database

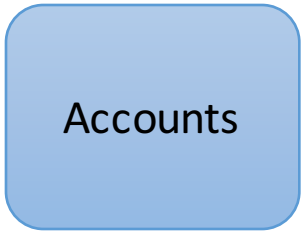

Centralized control
A central authority decides which nodes are part of the network

Decentralized database

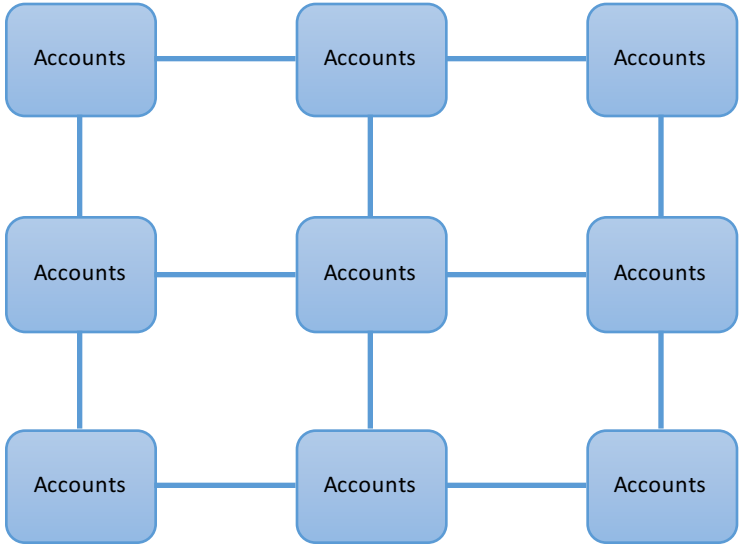

Decentralized control
Anyone can join the network

Blockchain Size

157.2 GB

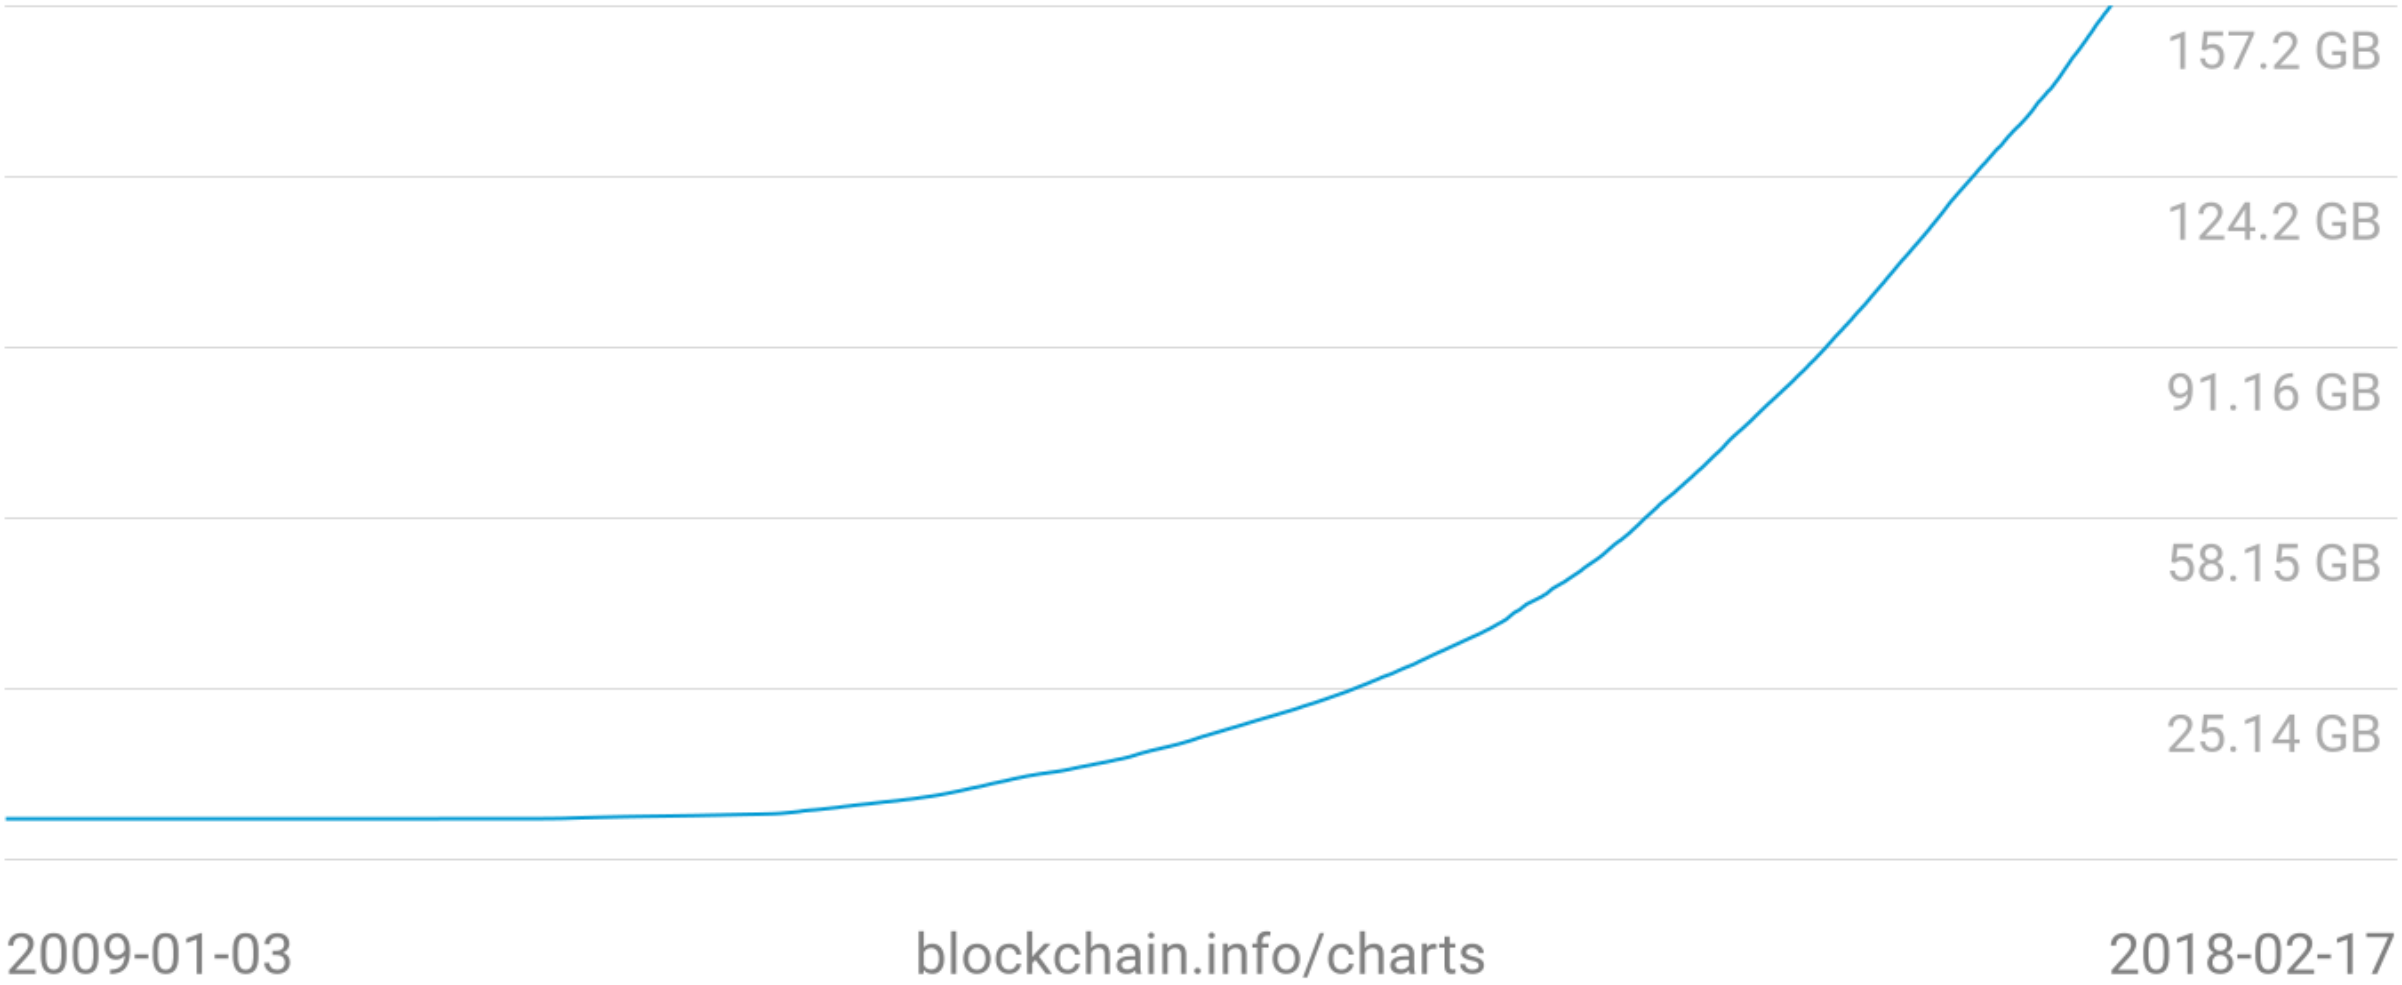

Cryptographic Background

Important concepts from cryptography:

- **Cryptographic hash functions**
  - Applications: message/file integrity, hash pointers, storing passwords…
- **Digital signatures**
  - Applications: email signatures (PGP), …

Cryptographic Hash Functions

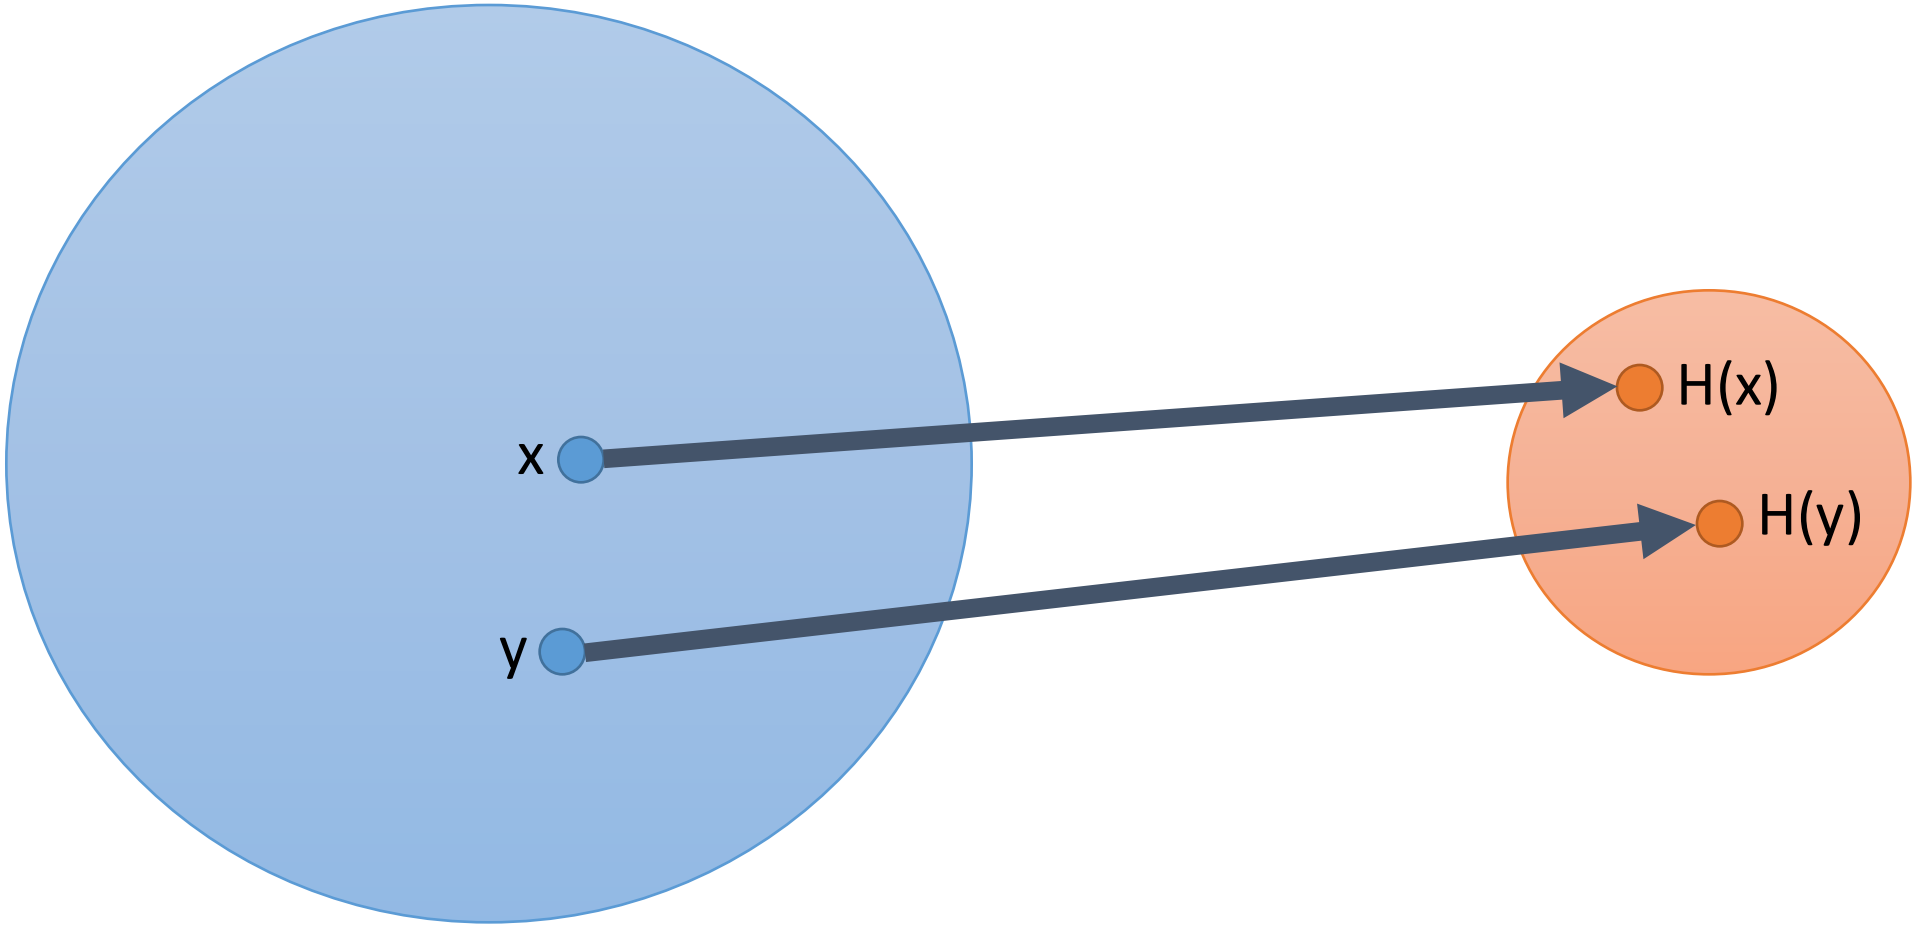

Infinite set of values
(all possible strings)

Finite set of values
(e.g., using 256 bits)

Hash Collision

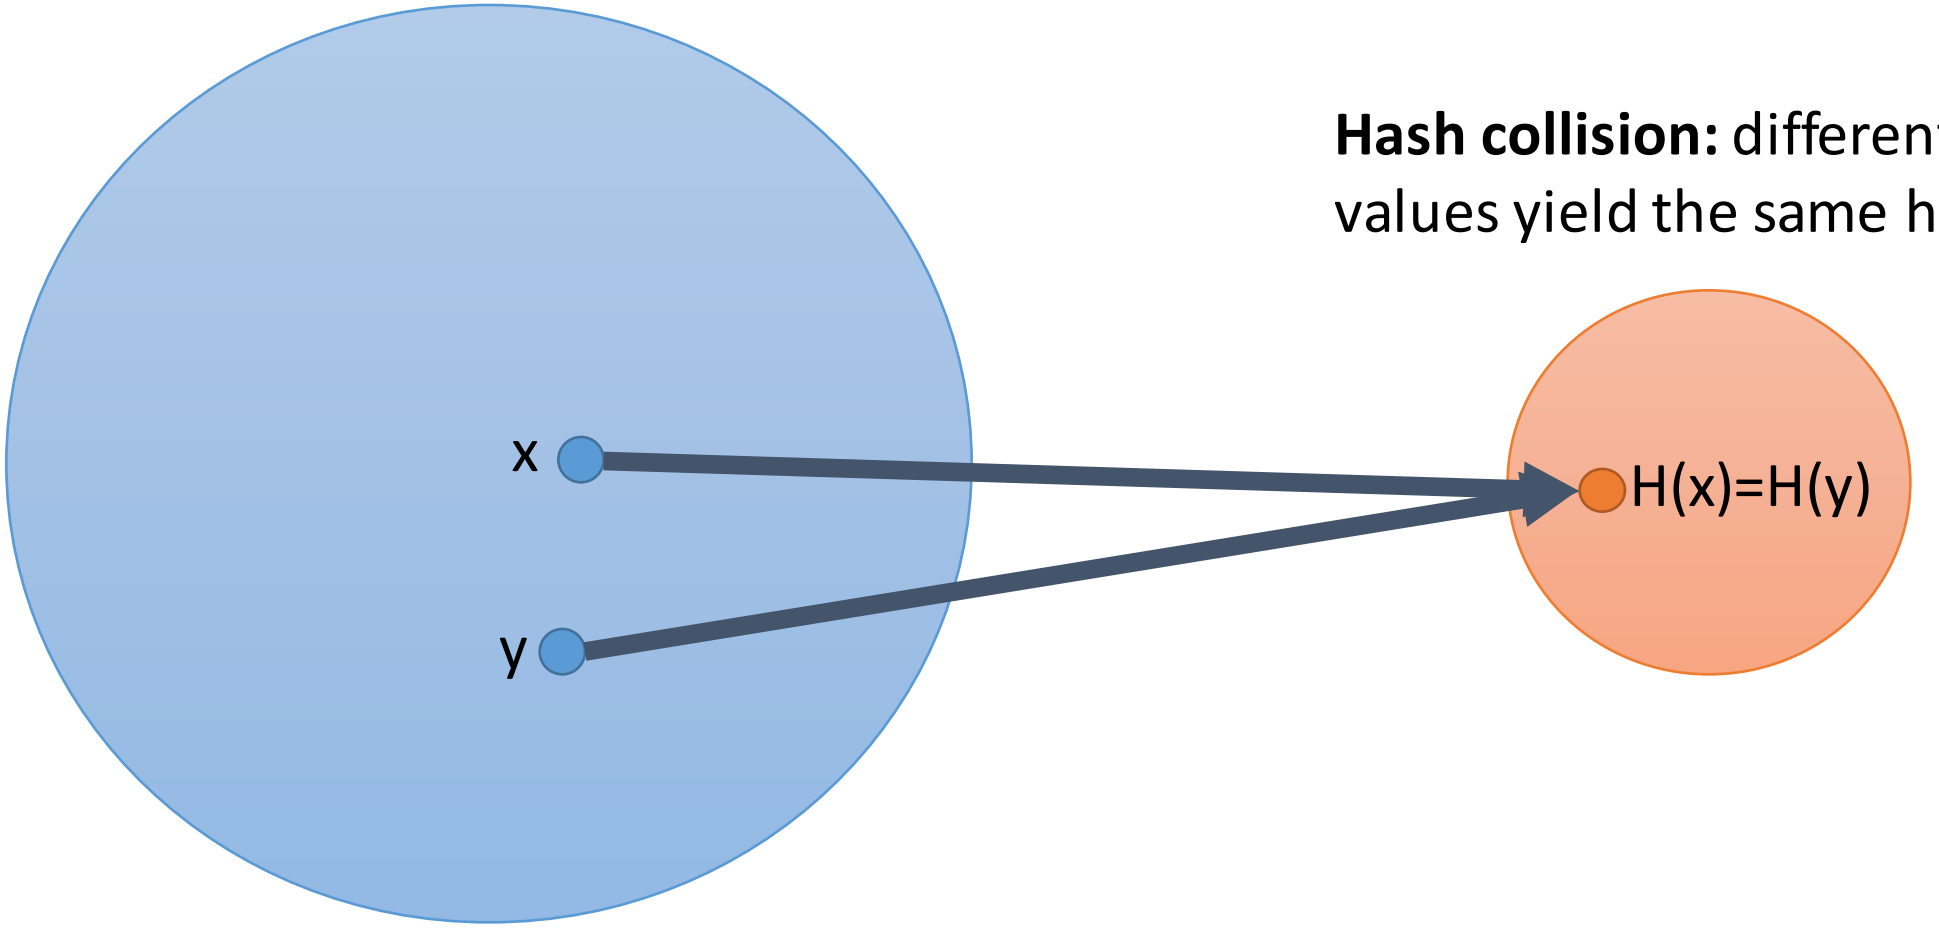

Hash collision: different input values yield the same hash value

Infinite set of values
(all possible strings)

Finite set of values
(e.g., using 256 bits)

Important Properties for Bitcoin

1) Collision-resistance
A hash function H is said to be **collision resistant** if it is infeasible to find two values, x and y, such that $x \neq y$, yet $H(x)=H(y)$.

2) Hiding
Given $y = H(x)$, it should be infeasible to figure out x.

3) Puzzle friendliness
Can be used for puzzles where the only solving strategy is bruteforcing

SHA256

Bitcoin uses the hash function SHA256 (from SHA-2 family).
The output uses 256 bits => 2^256 different values
You will get a hash collision when computing 2^128 hashes (on average)

Examples

sha256(niklas) =
760dcecfbe1ce8c36f9ac03686d3ad74e4c4f08978648677aa62b87014c27365

sha256(niklaz) =
1f5fd1befbf9da49d1fc5f8c241fc932800aa907358742155d091d880c2b18d8

|| is concatenation

Hash Pointers

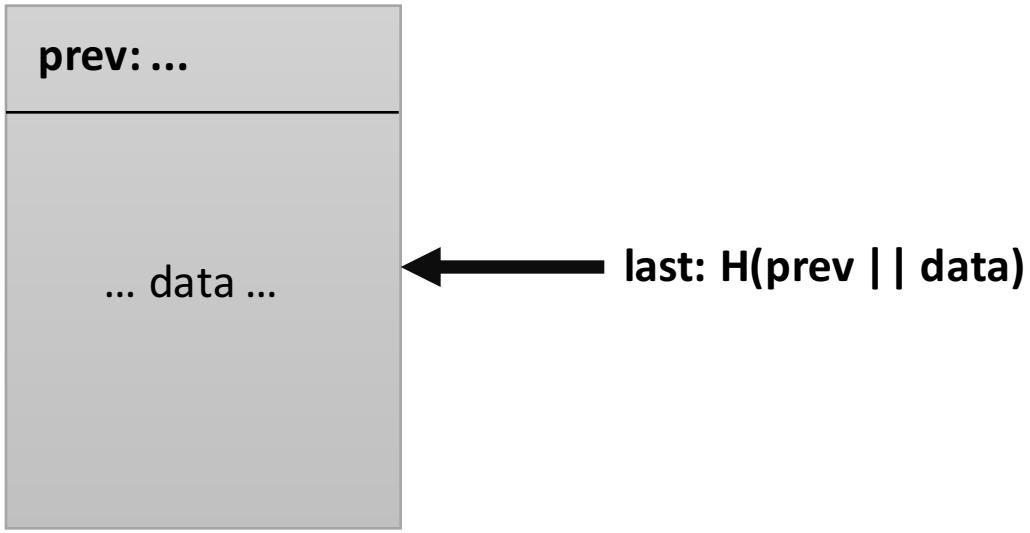

Last is a hash pointer, which is the hash of the content of **B1**.
If we change the data in **B1**, the value of **last** will change.
Thus, given the hash pointer, we can verify that B1 has not changed (probabilistic).

A Linked Chain of Blocks

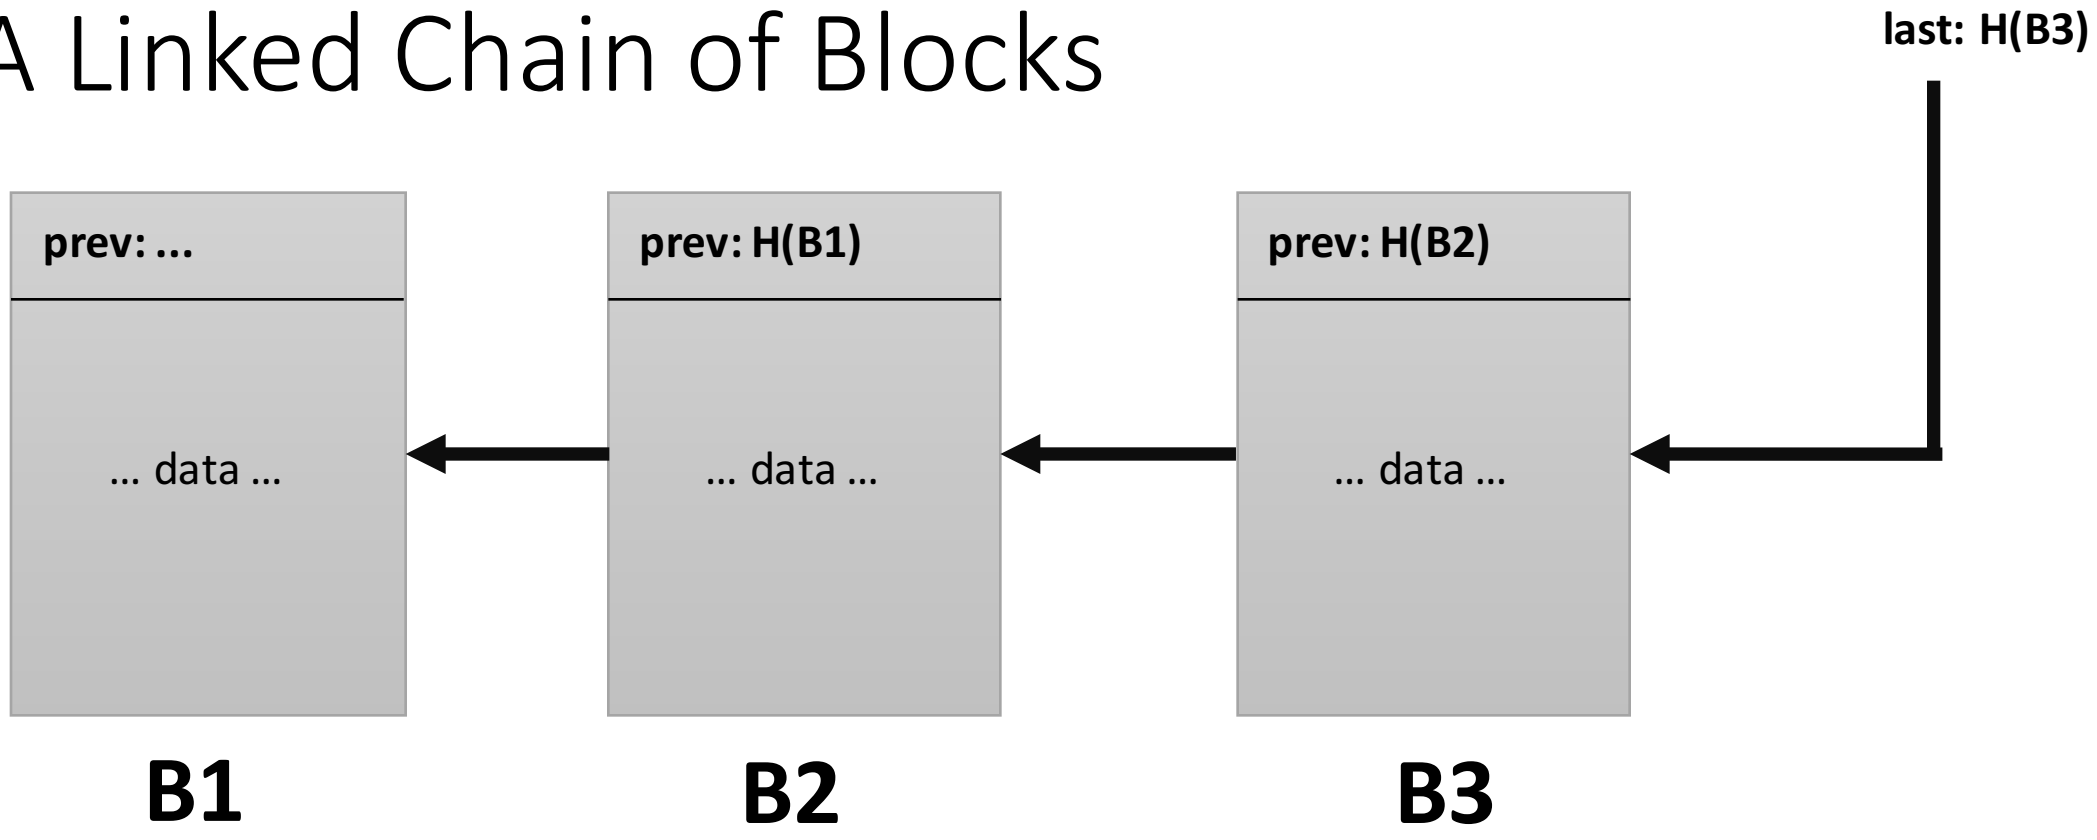

Given the value of **last**, it's very difficult to change the data of **B1,** without changing the value of **last**.

Digital Signatures

Signing messages that can be verified.

API

(privateKey, publicKey) <- generateKeys()
signature <- sign(privateKey, message)
verify(publicKey, message, signature)

Property:

verify(publicKey, message, sign(privateKey, message)) == true

Bitcoin

- Addresses
- Transaction-based ledger
- Blocks – a collection of transactions
- Mining – verifying blocks
- Double-spend problem

Public Keys as Identities

In Bitcoin, public keys are used as identities.

Coins are sent to *addresses*, which is the hash of the public key.

To use a coin:
Create a new transaction and sign it with the corresponding private key.

Transactions-based ledger

The ledger is transaction-based (no accounts)

- A transaction has input coins and output coins (index from 0)
- Inputs are consumed in the transaction (cannot be used again)
- Outputs are produced from the inputs, thus, sum(inputs) >= sum(outputs)
- The inputs reference outputs from previous transactions

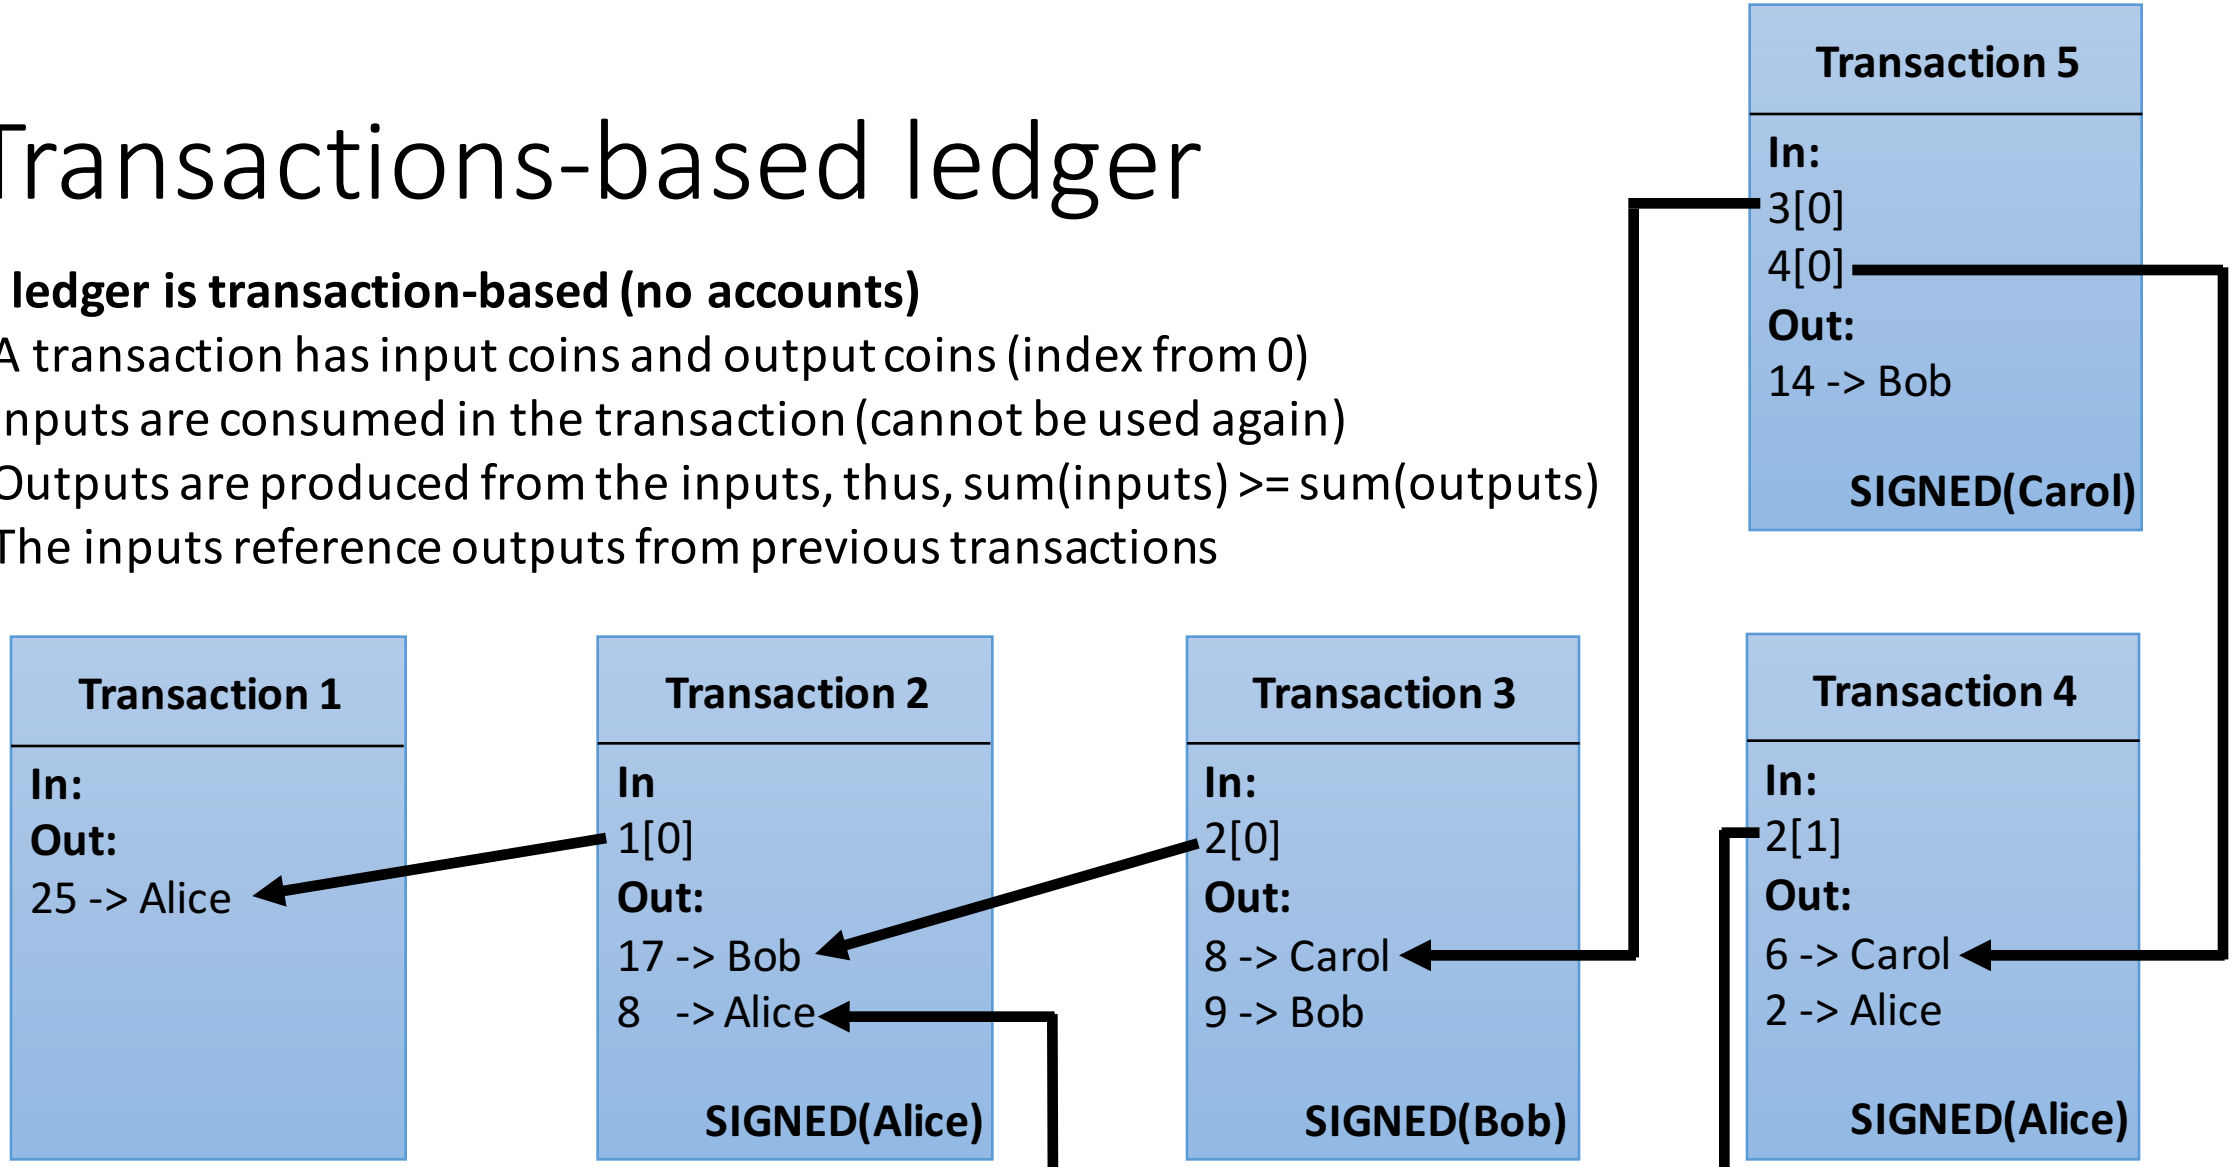

End result:
Alice: 2
Bob: 23

UTXO: unspent transaction output

Example Transactions

Change address
A(2) → B(1), A(1)

Merging
B(1), B(1) → B(2)

Joint payment
A(1), B(1) → C(2)

Splitting
B(2) → B(1), B(1)

Don’t Lose Your Private Key!

Quest for lost hard drive with £4m stored bitcoins

A Newport man has visited a landfill site in south Wales hoping to find a computer hard drive he threw away which is now worth over £4m.

James Howells' hard drive contains 7,500 bitcoins - which is a virtual form of currency for use online. This week, a single bitcoin's value hit $1,000 (£613) for the first time, making his collection worth $7.5m (£4.6m).

BBC correspondent Hywel Griffiths spoke with Mr Howells about what happened.

28 Nov 2013 | Technology

Today worth (approximately)**:**
7500*10000 = **75 000 000 USD**

Example of Transaction Data

```
{ "hash":"1b4890246...",
  "vin_sz":1,
  "vout_sz":1
  "size":223,
  "inputs":[
      {"prev_out":{
           "hash":"76a91496b..."
           "n":0},
       "scriptSig":"47304402201420..."}
  ],
  "out":[
      {"value":2298949,
       "scriptPubKey": "OP_DUP ... <pubKeyHash>..."}
  ]
}
```

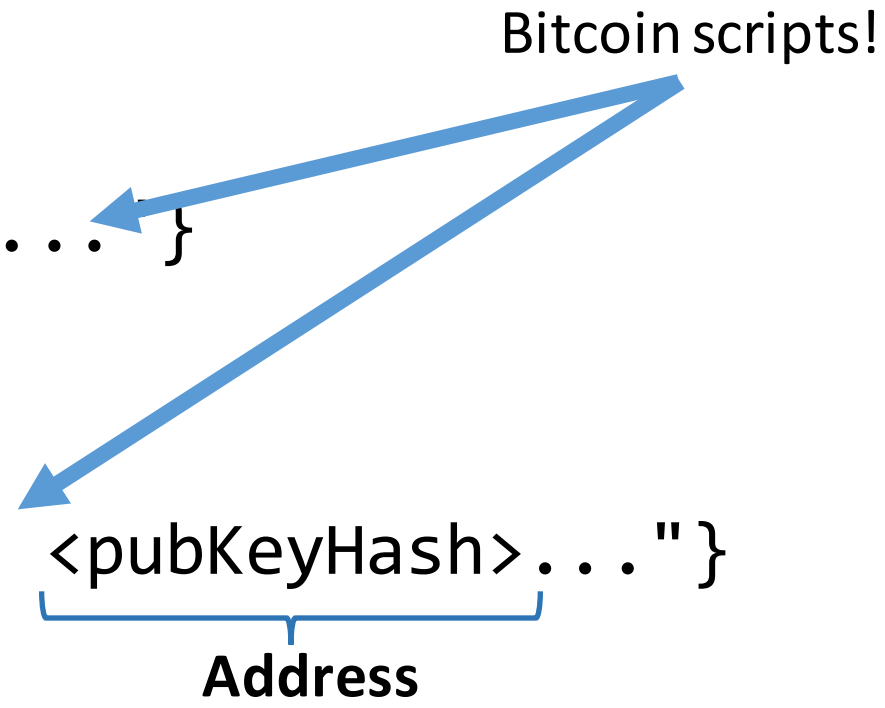

Example Transaction Verification

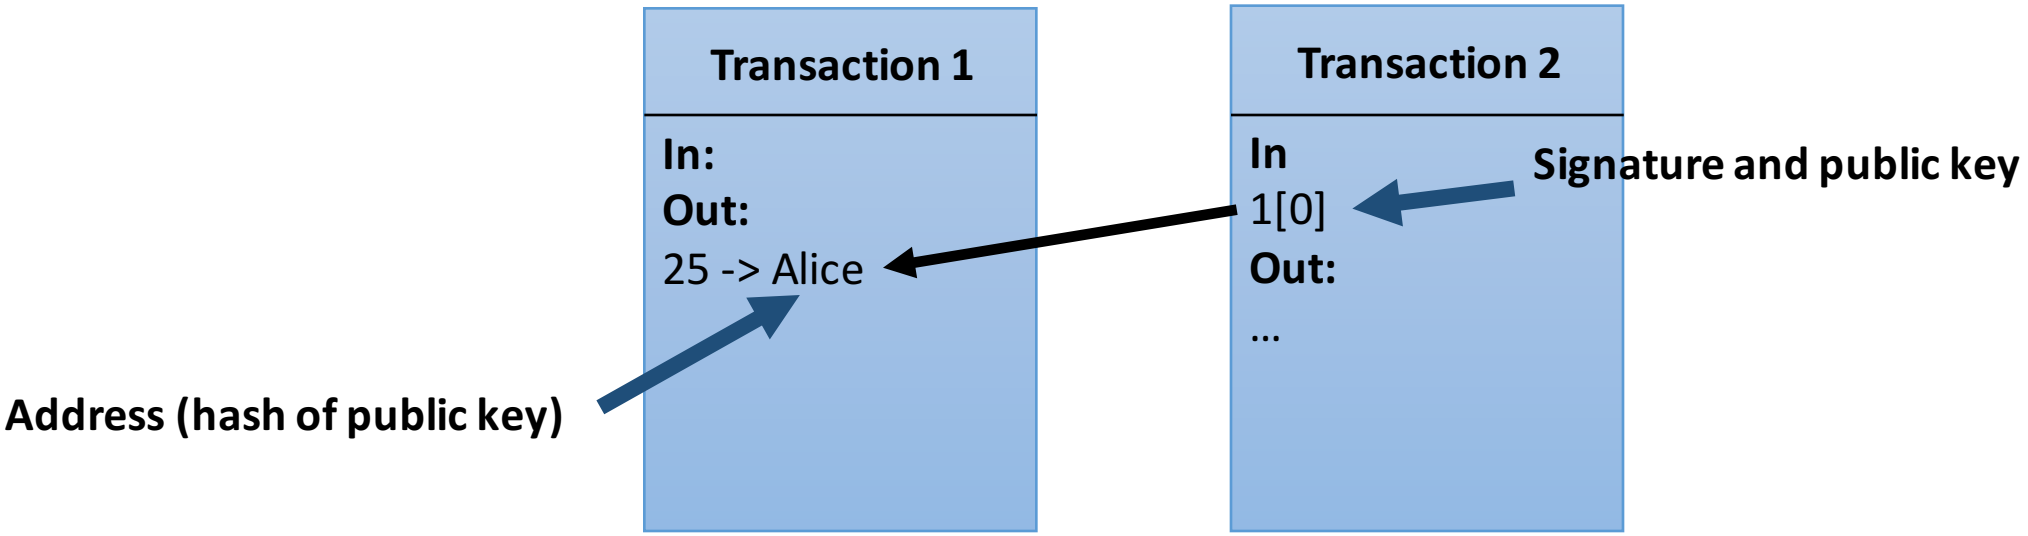

To verify an input

1. Find the referenced output
2. Hash the public key (**h**) given in the input
3. Compare **h** with address specified in referenced output
4. Verify signature with public key

Bitcoin Scripts (Pay-to-PubkeyHash script)

Script in referenced output (earlier transaction):

```
scriptPubKey:
OP_DUP
OP_HASH160
<pubKeyHash>
OP_EQUALVERIFY
OP_CHECKSIG
```

Script in input (new transaction)

```
scriptSig:
<sig>
<pubKey>
```

The scripts are concatenated:

```
<sig>
<pubKey>
OP_DUP
OP_HASH160
<pubKeyHash>
OP_EQUALVERIFY
OP_CHECKSIG
```

Script Execution

	Command	Stack	Description
From input	<sig>	<sig>	Push
	<pubKey>	<sig> <pubKey>	Push
From referenced output	<OP_DUP>	<sig> <pubKey> <pubKey>	Duplicate top of stack
	<OP_HASH160>	<sig> <pubKey> <hashOfPubKey>	Hash top of stack
	<pubKeyHash>	<sig> <pubKey> <hashOfPubKey> <pubKeyHash>	Push
	OP_EQUALVERIFY	<sig> <pubKey>	Top of stack should be equal
	OP_CHECKSIG	true	Verify signature of public key

Scripting Languages

- The scripting language in Bitcoin is limited
- However, other cryptocurrencies (**Ethereum**,…) have scripting languages that are Turing-complete
  => making it possible to write arbitrary programs
- A way to implement **smart contracts** (contracts specified in code)

Blockchain

- A block is a collection of transactions (some thousands transactions)
- A new block is created every 10 minutes (on average)
- The blocks are put in a blockchain

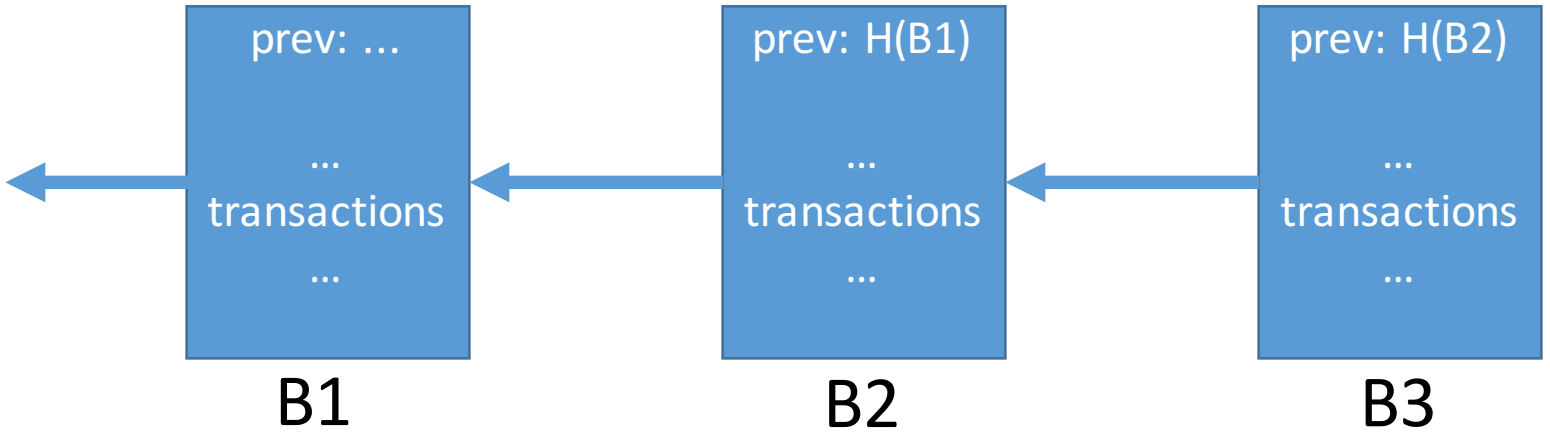

Double Spend Attempt

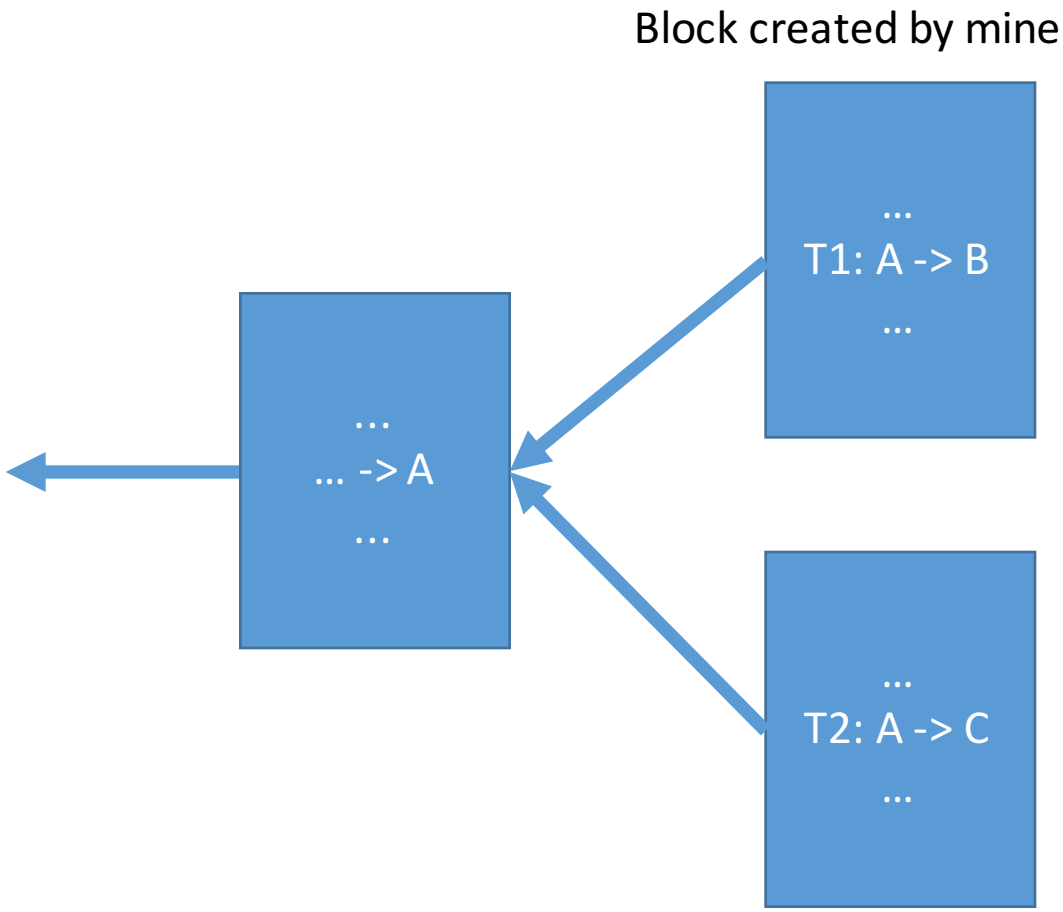

Alice creates two transaction that uses the same output, thus, a **double spend attempt**!

Two block are created simultaneously by two different miners.

Which transaction is valid? T1 or T2? Both?

Answer: we don't know yet

Which Block to Extend? (1)

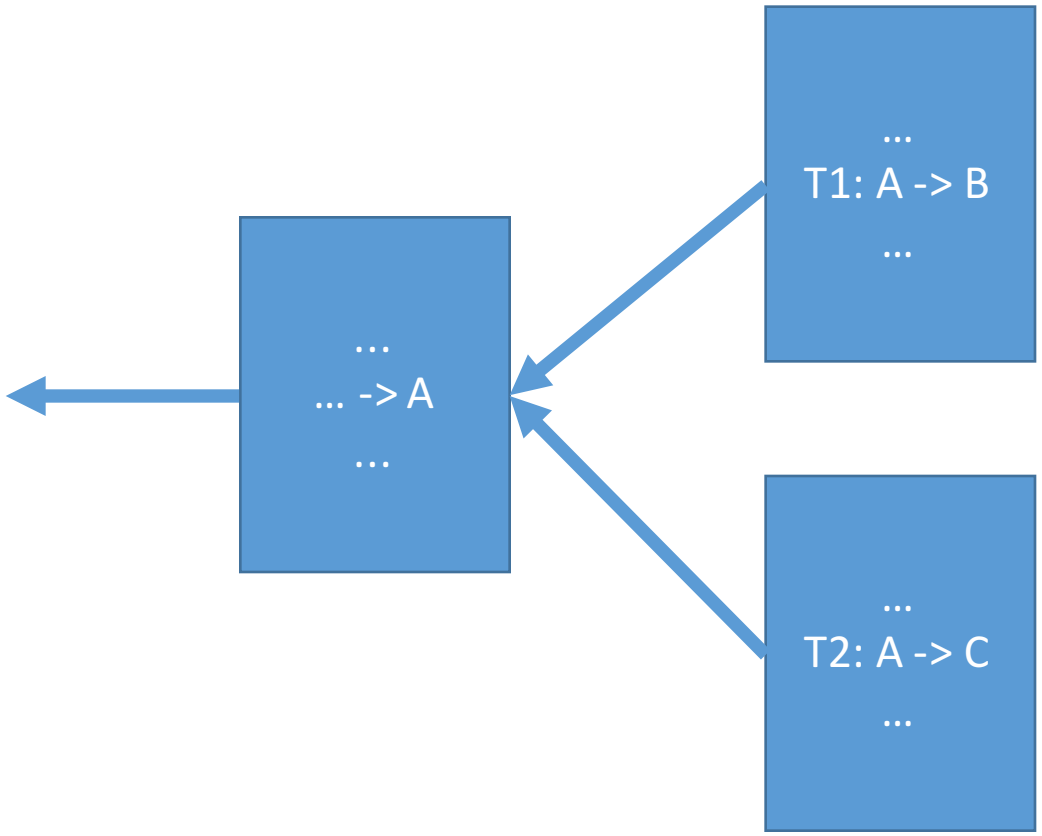

A new block is created by a miner.
Which previous block to extend?

The miner decides that!
(probably the block that the miner observed first)

Which Block to Extend? (1)

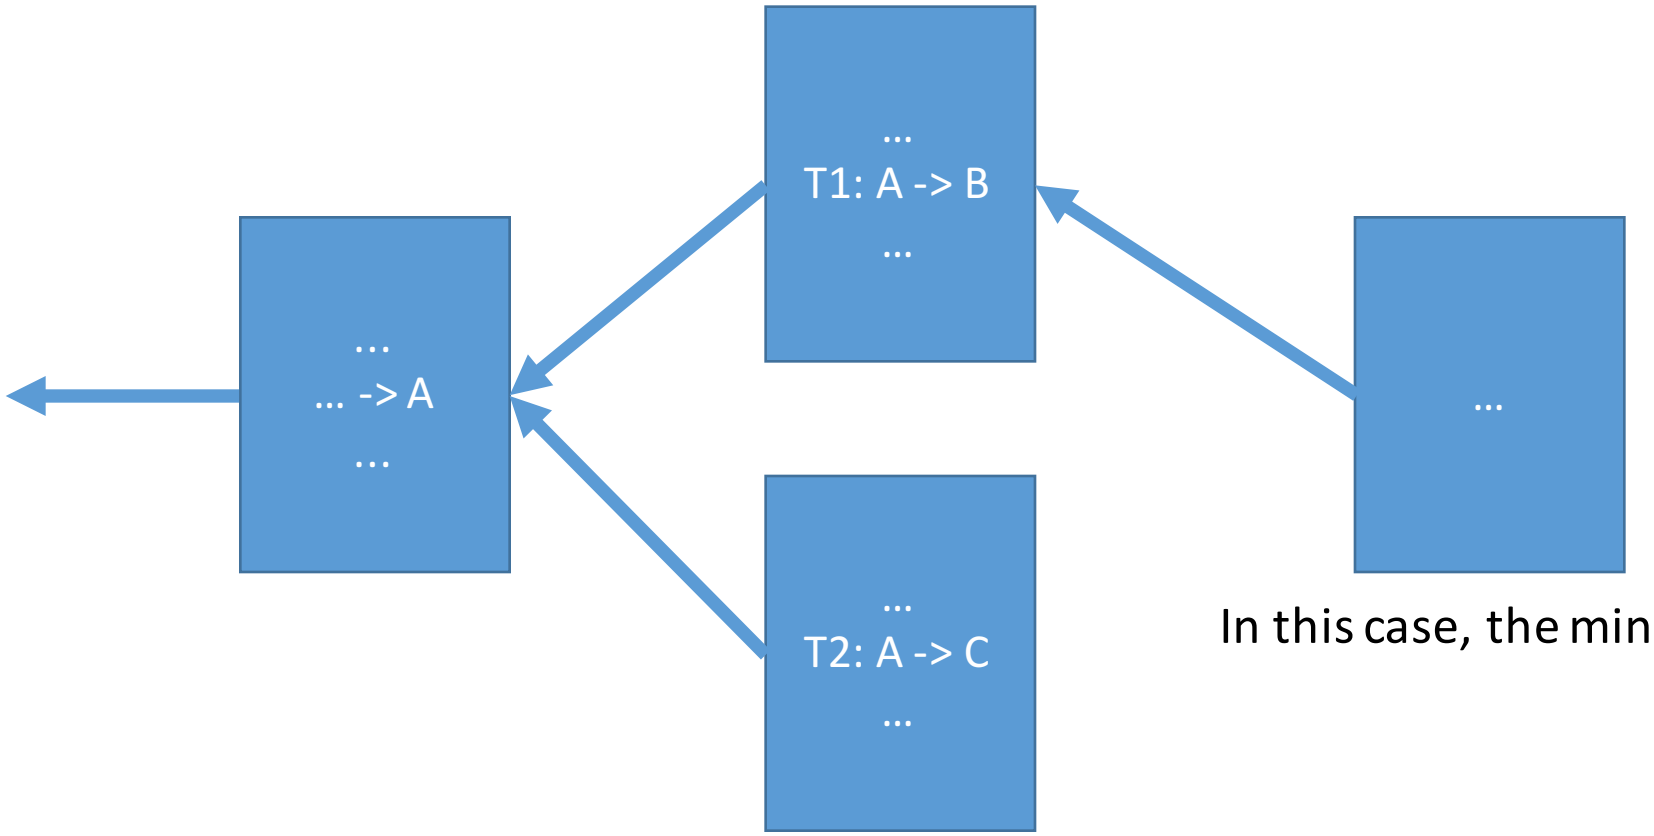

In this case, the miner selected the top block.

Which Block to Extend? (2)

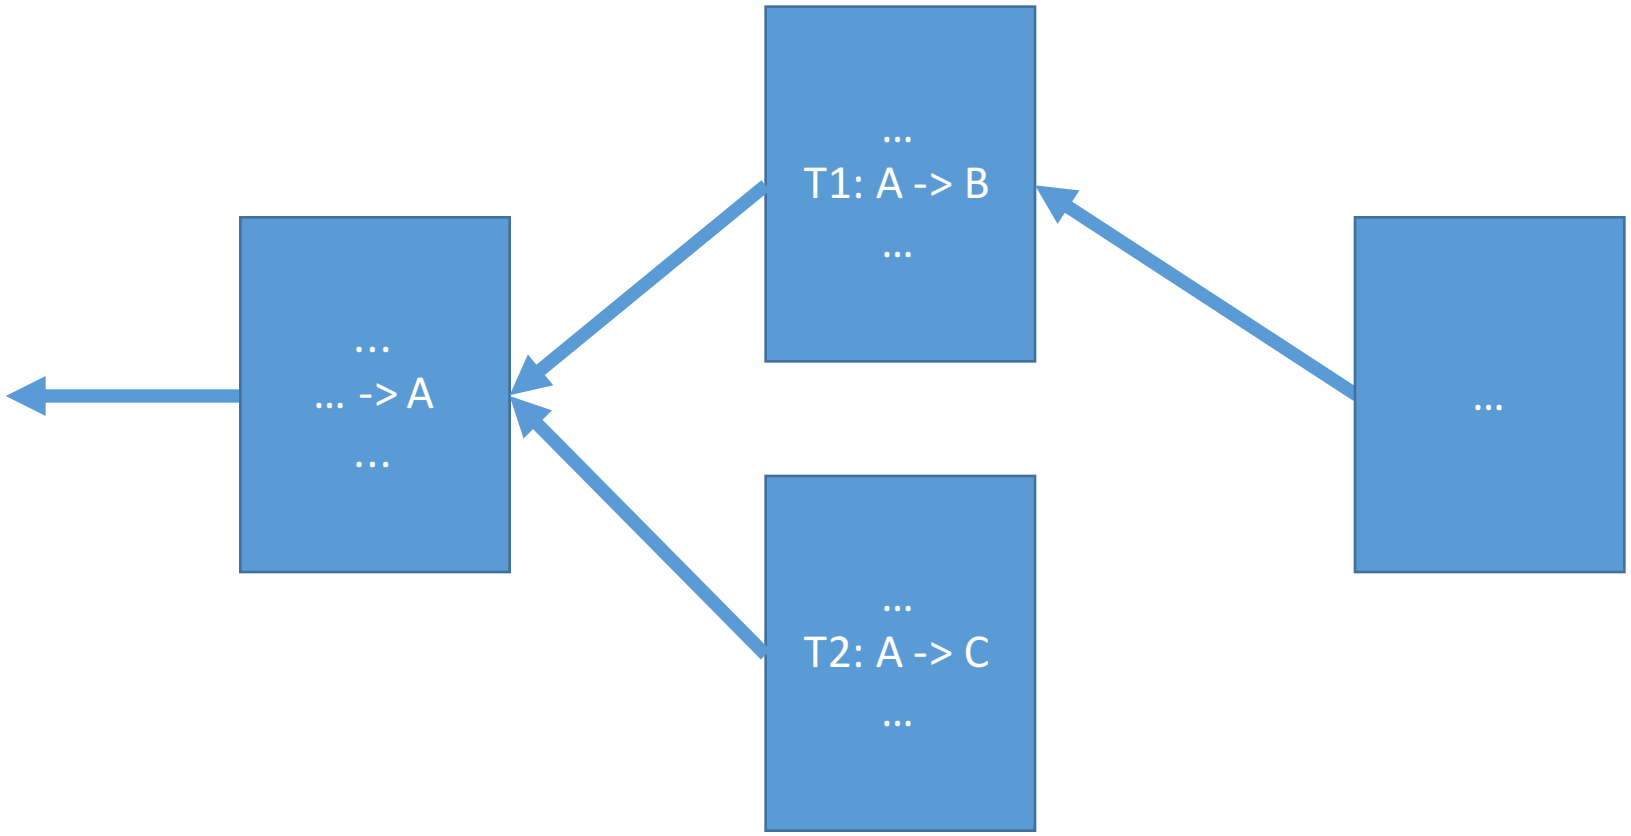

A new block is created.
Which block to extend?

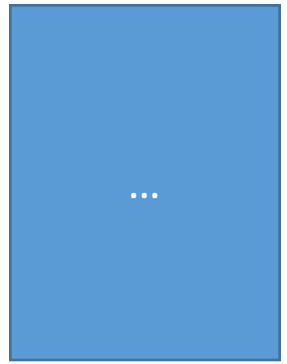

Longest Chain is Extended!

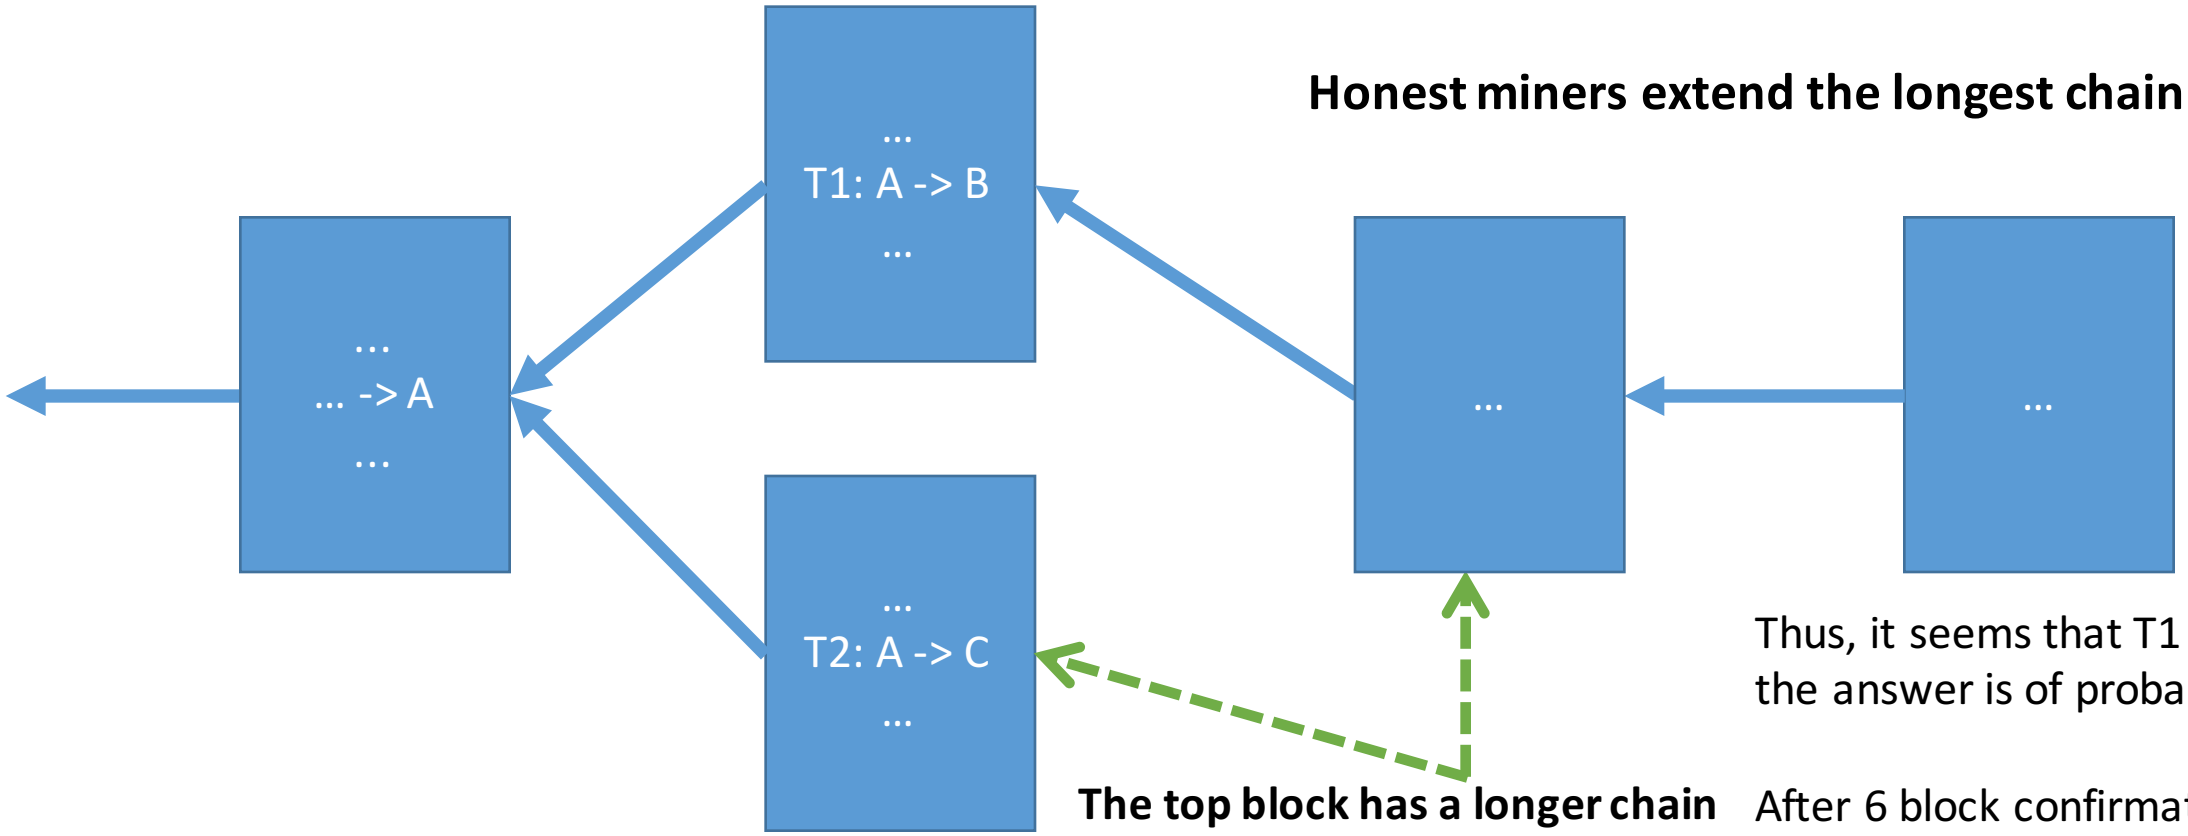

Thus, it seems that T1 succeeded, but the answer is of probabilistic nature.

After 6 block confirmations, it's very likely that the transaction succeeded.

Block Creation (1)

How is a block created?
Miners need to solve a cryptographic puzzle!

For the whole network, it takes an average of 10 minutes to solve the puzzle.

Block Creation

The puzzle requires a solution to:

H(nonce || prev_hash || …) < difficultyTarget

The hash should have a leading number of zero bits (difficulty decides how many)

The miner tries different values of the **nonce** to meet the target (by bruteforcing).

The puzzle is hard to solve, but very easy to verify.

Proof of Work

This technique is called ***Proof of Work (PoW)***, an approach for ***distributed consensus***

It can be thought of **as one-CPU-one-vote.**

PoW prevents attacks on the network, or rather, it makes them very costly.

If you own 10% of all hash power of the network,
then you will on average create 10% of the blocks.

(There are other consensus mechanisms: Proof of Stake, …)

Hash Rate

23.15 EH/s

Exa=10^18

21 290 000 000 000 000 000 hashes/s

Requires a lot of energy!

How long time before we get a hash collision with this hash rate?

$\frac{2^{128}}{23*10^{18}}$/(86400*365) = 469 142 742 209 years

13 799 000 000 years
(the age of the universe)

Answer: 34 times the age of the universe

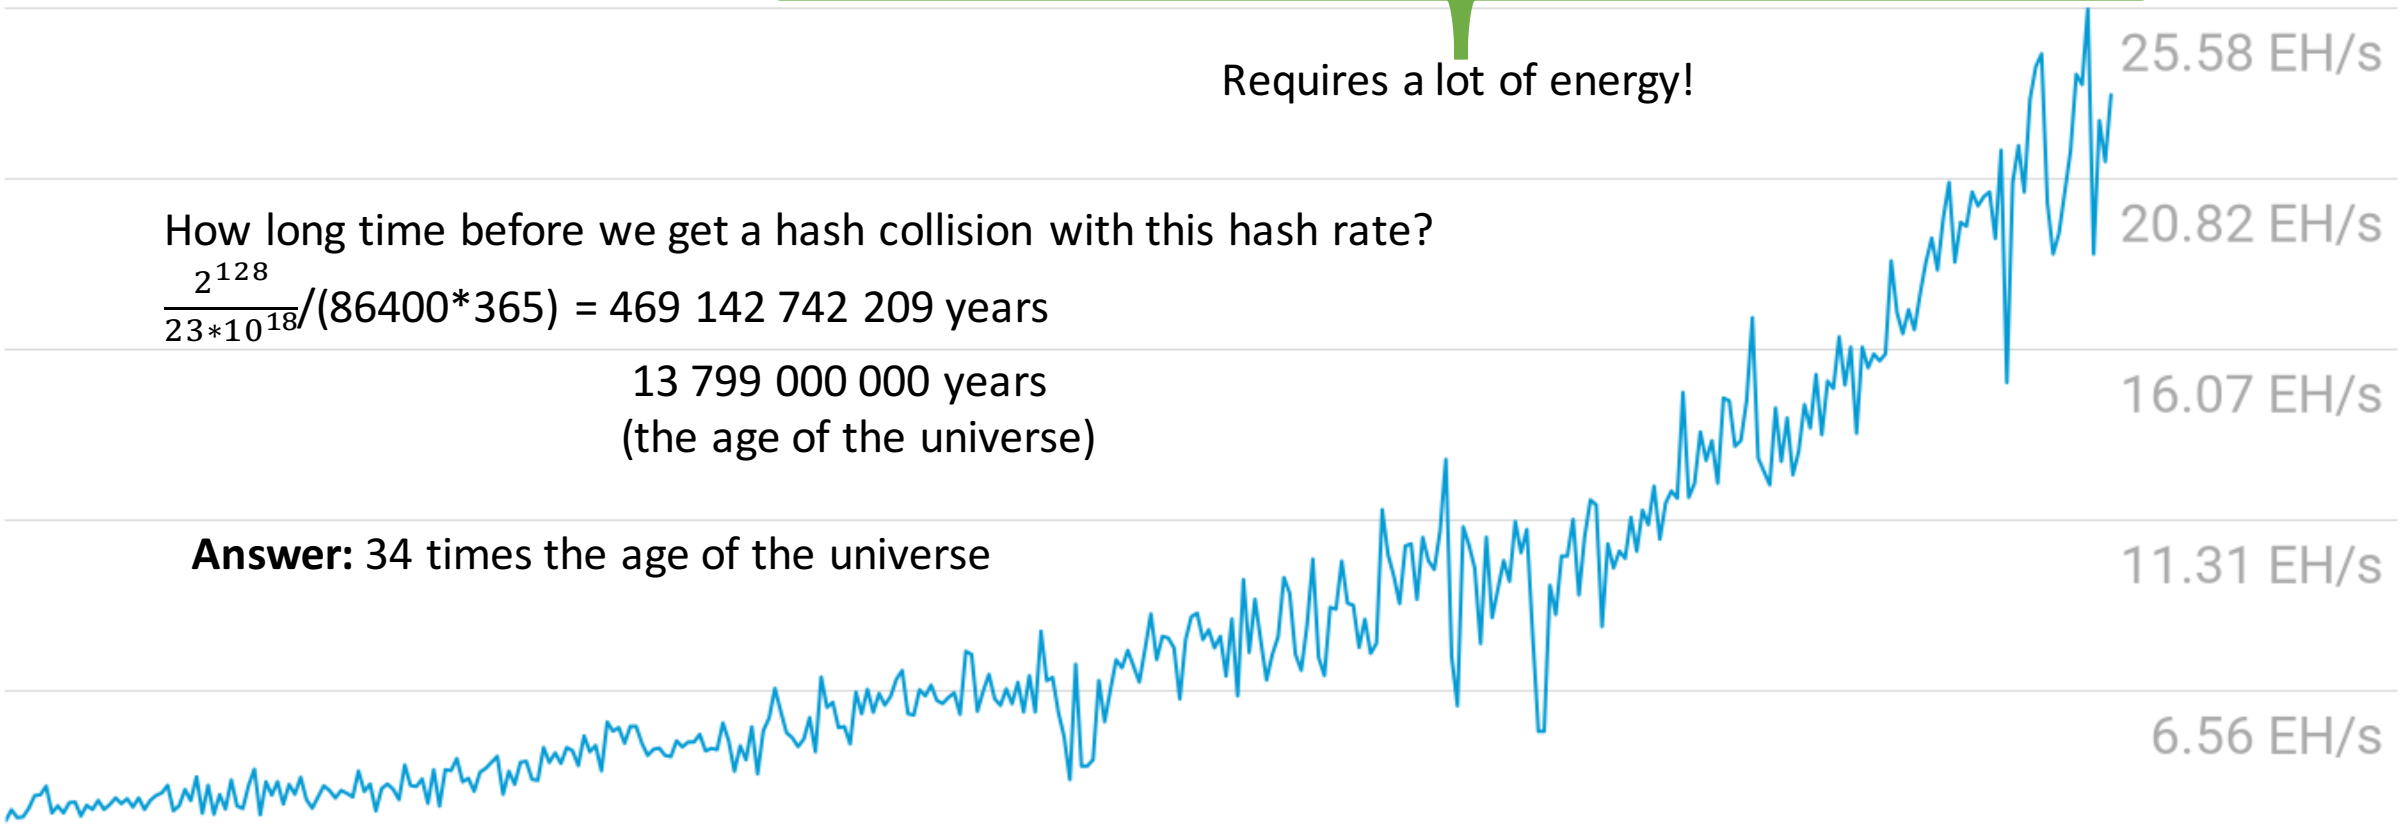

2017-02-19 blockchain.info/charts 2018-02-18

Difficulty

2,874,674,234,415

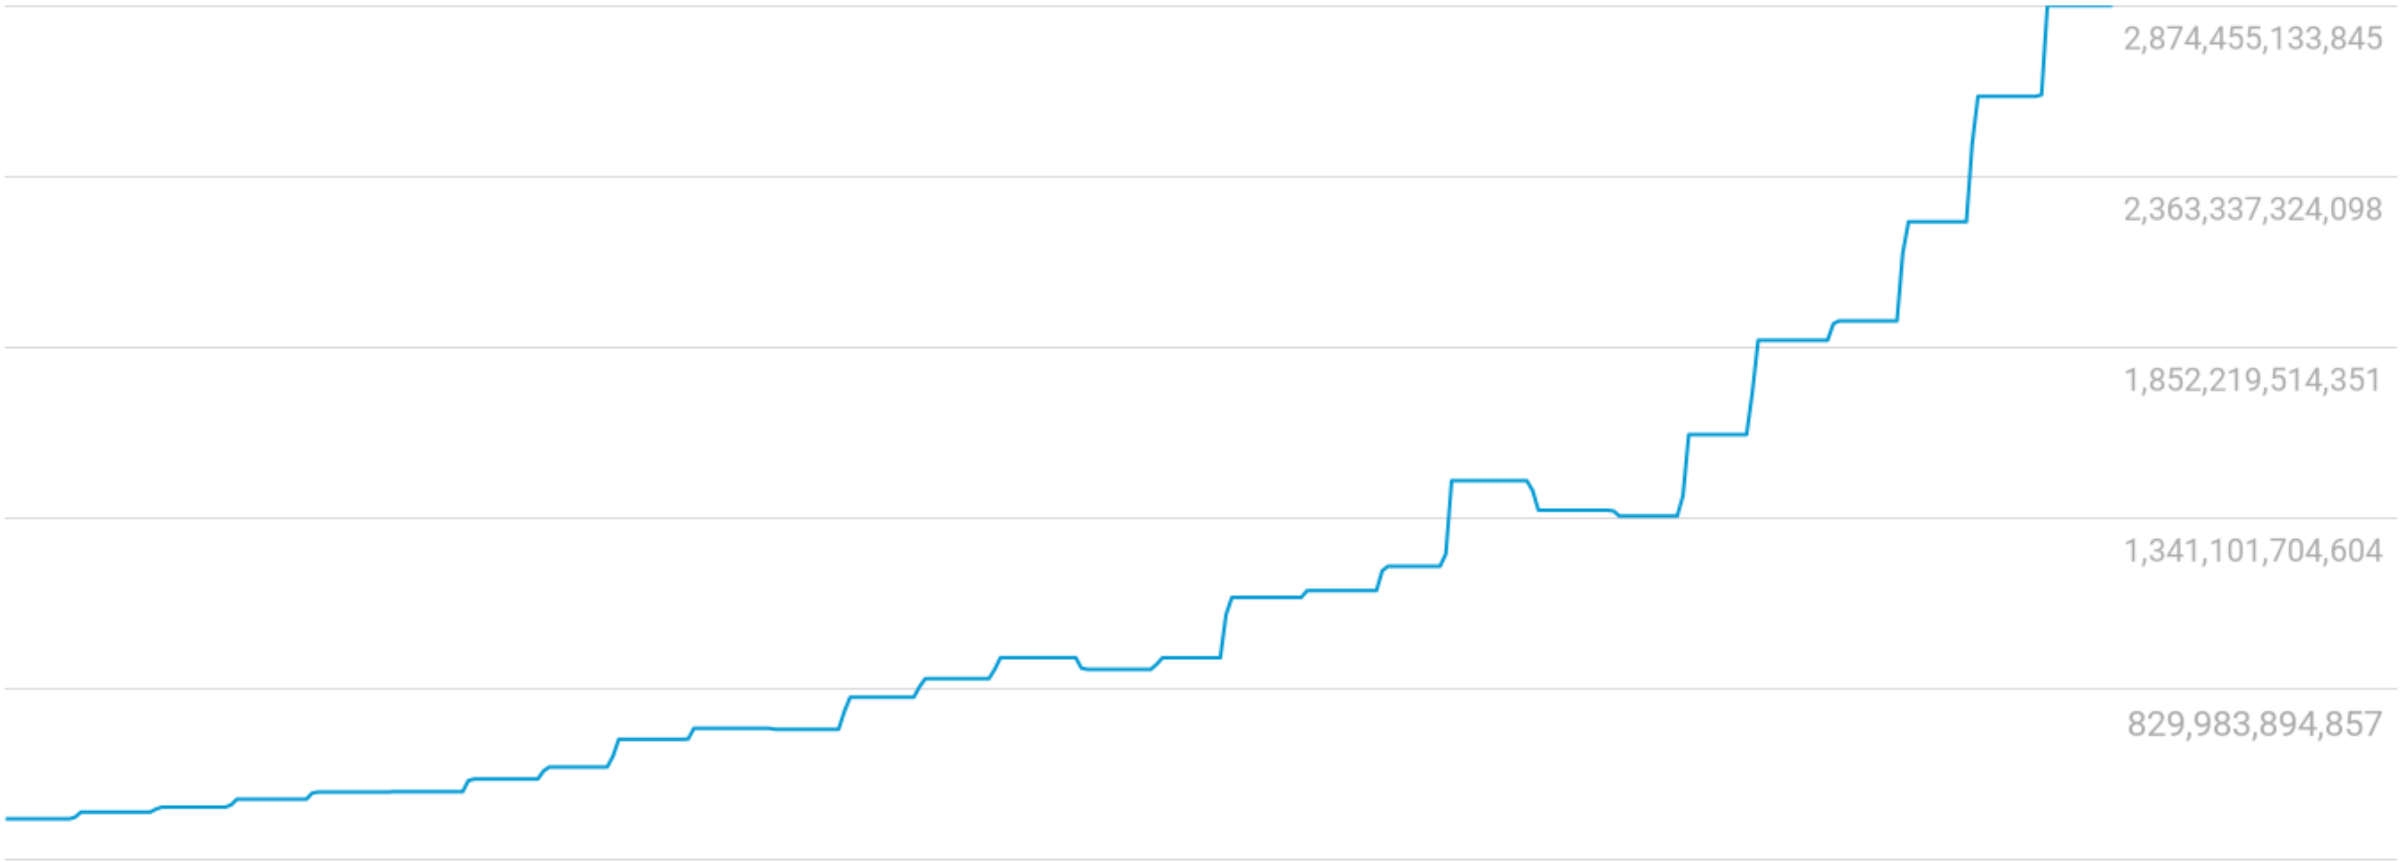

Network (from Bitcoin paper)

The steps to run the network are as follows:

1. New transactions are broadcast to all nodes.
2. Each node collects new transactions into a block.
3. Each node works on finding a difficult proof-of-work for its block.
4. When a node finds a proof-of-work, it broadcasts the block to all nodes.
5. Nodes accept the block only if all transactions in it are valid and not already spent.
6. Nodes express their acceptance of the block by working on creating the next block in the chain, using the hash of the accepted block as the previous hash.

Merkle Tree

The transactions in a block are stored in a Merkle tree

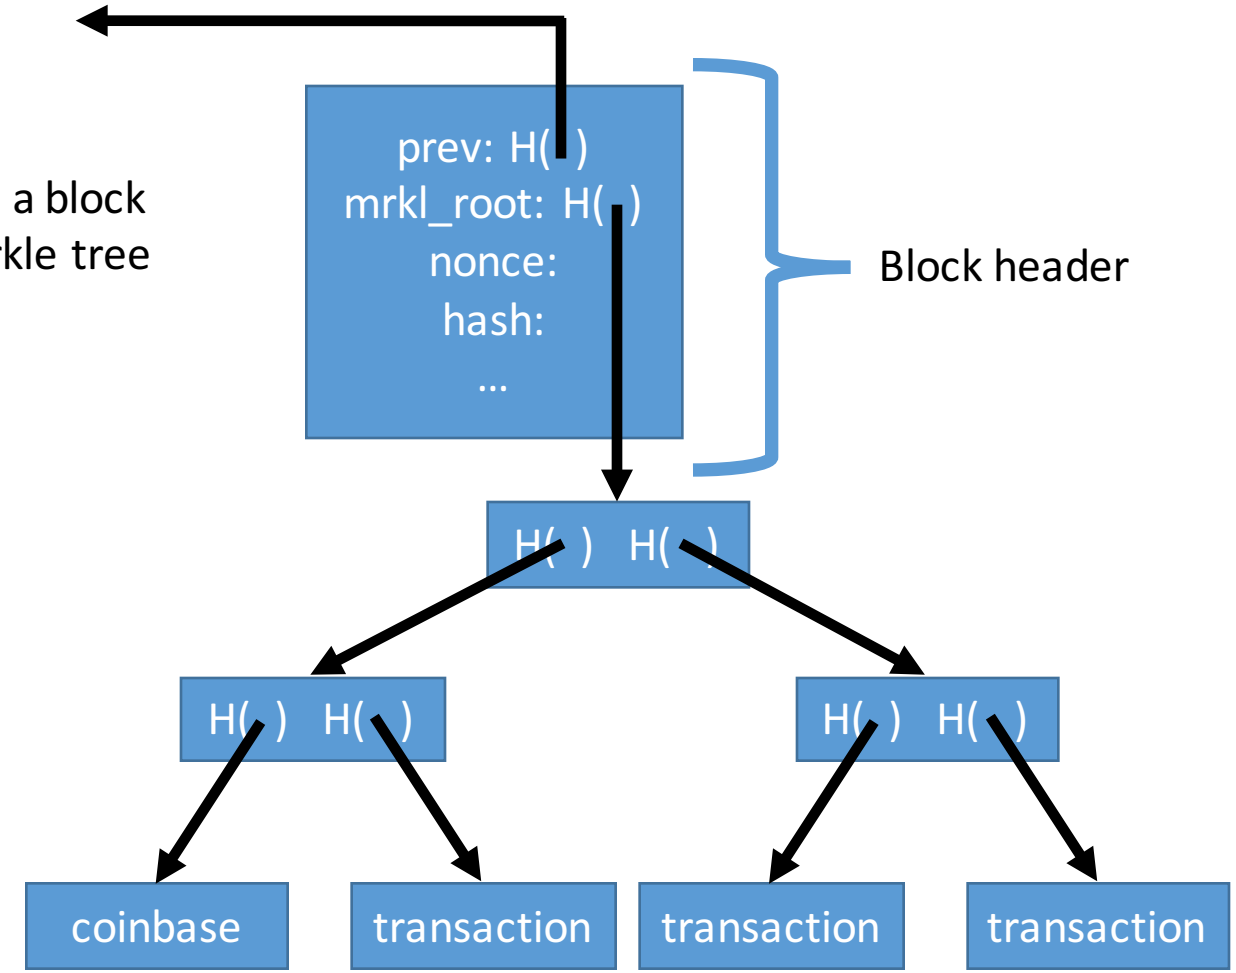

CPU mining pseudocode

```
TARGET=(65535<<208)/DIFFICULTY;
coinbase_nonce=0;
while(1){
  header=makeBlockHeader(transactions,coinbase_nonce);
  for(header_nonce=0;header_nonce<(1<<32); header_nonce++){
    if(SHA256(SHA256(makeBlock(header,header_nonce))) < TARGET)
      break;//block found!
  }
  coinbase_nonce++;
}
```

Mining Incentive

Why do miners mine?

Because they are rewarded!
The rewards encourage them stay honest.

Block rewards

- New coins are created in each block (called the **coinbase transaction**)
  - The number decreases over time
- Transaction fees (when sum(inputs) > sum(outputs))

The Genesis Block

The Gensis block contains the following text in its coinbase transaction:

The Times 03/Jan/2009 Chancellor on brink of second bailout for banks

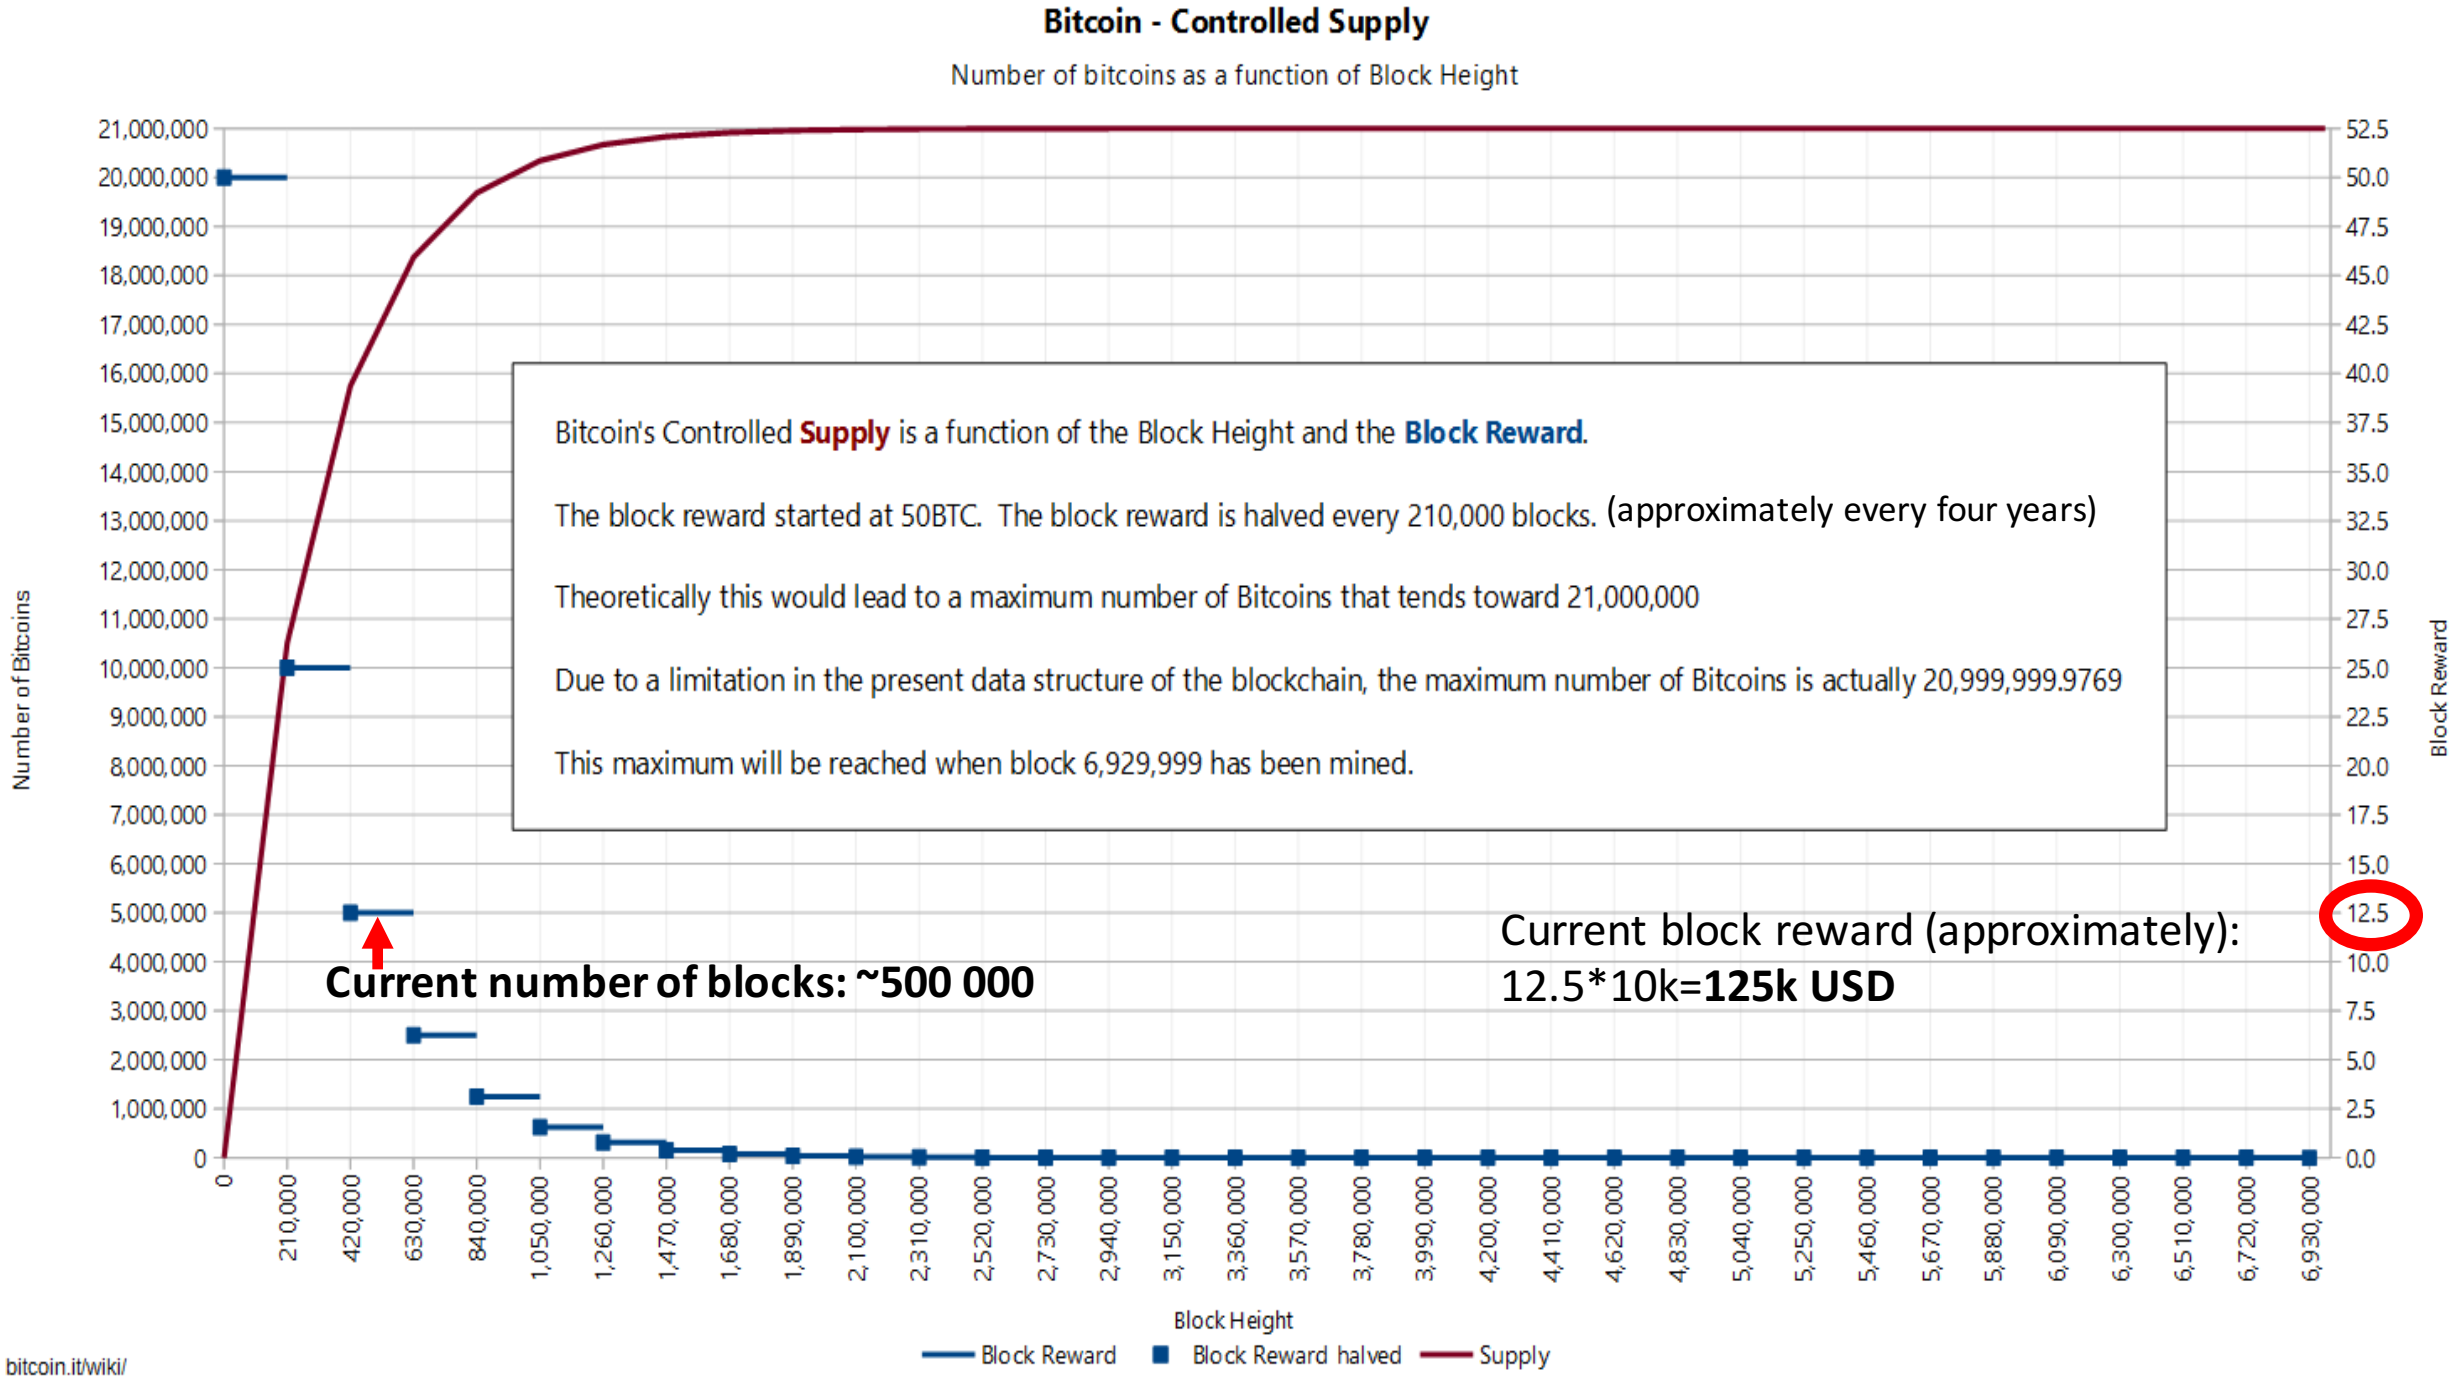
Bitcoin - Controlled Supply
Number of bitcoins as a function of Block Height
Bitcoin's Controlled Supply is a function of the Block Height and the Block Reward.
The block reward started at 50BTC. The block reward is halved every 210,000 blocks. (approximately every four years)
Theoretically this would lead to a maximum number of Bitcoins that tends toward 21,000,000
Due to a limitation in the present data structure of the blockchain, the maximum number of Bitcoins is actually 20,999,999.9769
This maximum will be reached when block 6,929,999 has been mined.
Current number of blocks: ~500 000
Current block reward (approximately):
12.5*10k=125k USD
Number of Bitcoins
Block Reward
Block Height
Block Reward
Block Reward halved
Supply
bitcoin.it/wiki/

Bitcoins in circulation

16,871,012.50 BTC

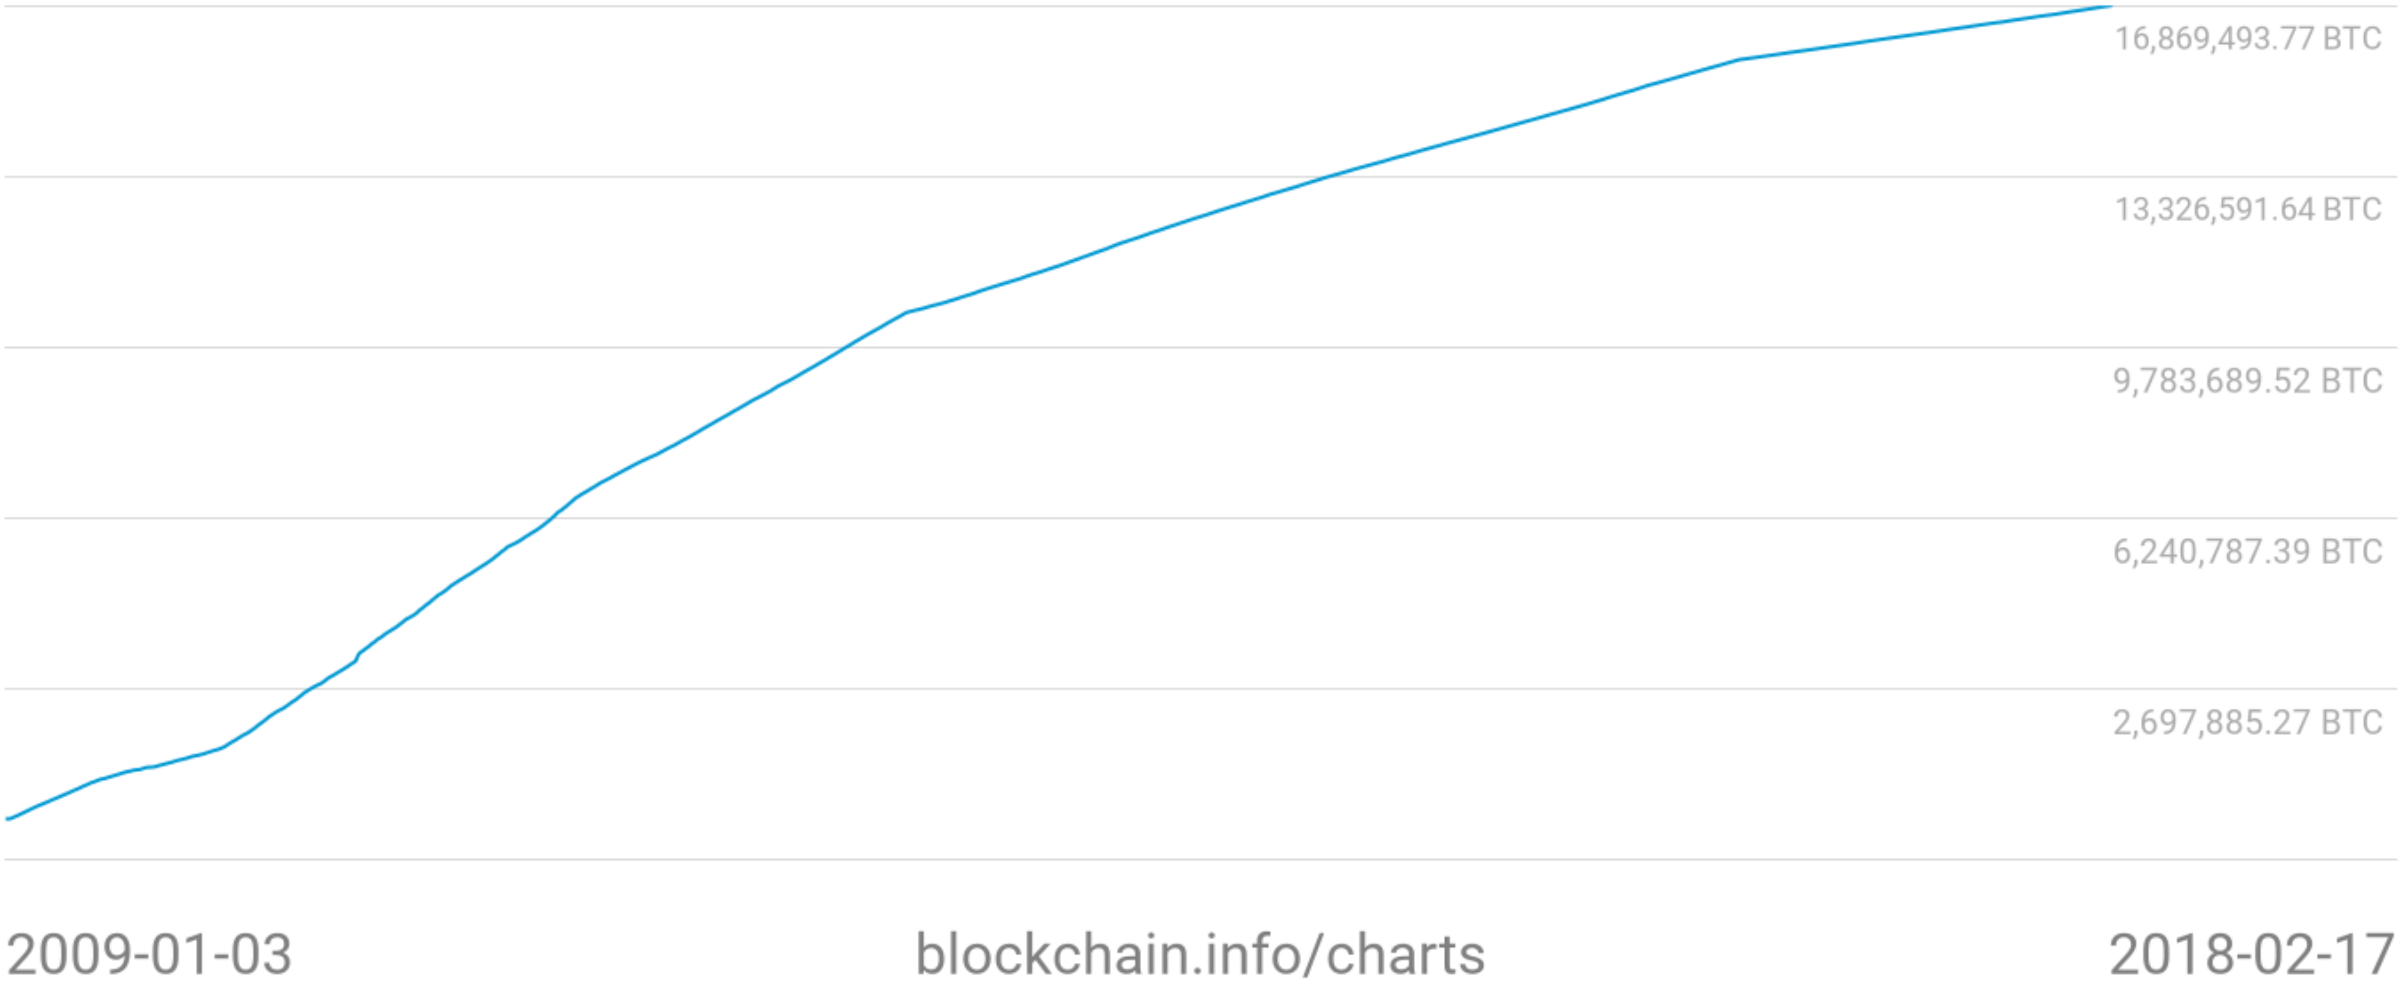

The Cost of Mining

If mining reward > mining cost
miner profits

where
mining reward = block reward + transaction fees
mining cost = hardware cost + operating costs (electricity, cooling, etc.)

Mining Hardware

The miners are increasingly using more efficient hardware:

1. CPU
2. GPU
3. FPGA
4. ASIC

Mining Pools

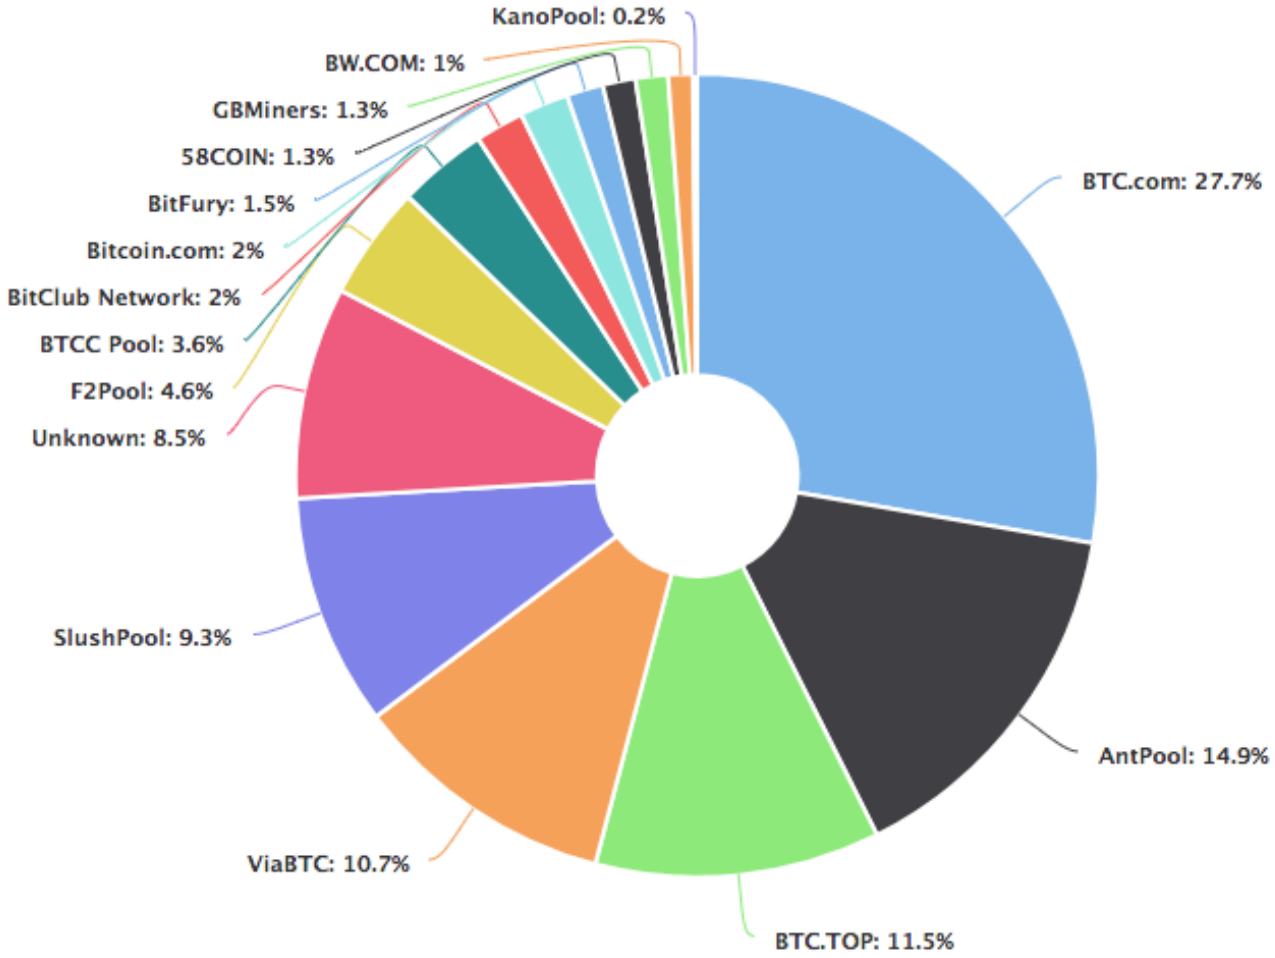

Source: blockchain.info

To get a more stable stream of income, be a member of a mining pool.

Scalability?

- A new block is created every 10 minutes
- The **max block size is 1 MB**
- Number of transactions per second:
  ~average transaction size/1 MB/60*10
- The current limit is about **7 transactions/second** => 604 800/day

Ongoing work

- **SegWit:** roughly doubling the block size
- **Lightning network:** second layer on top of Bitcoin blockchain for micropayments

Confirmed Transactions Per Day

159,495

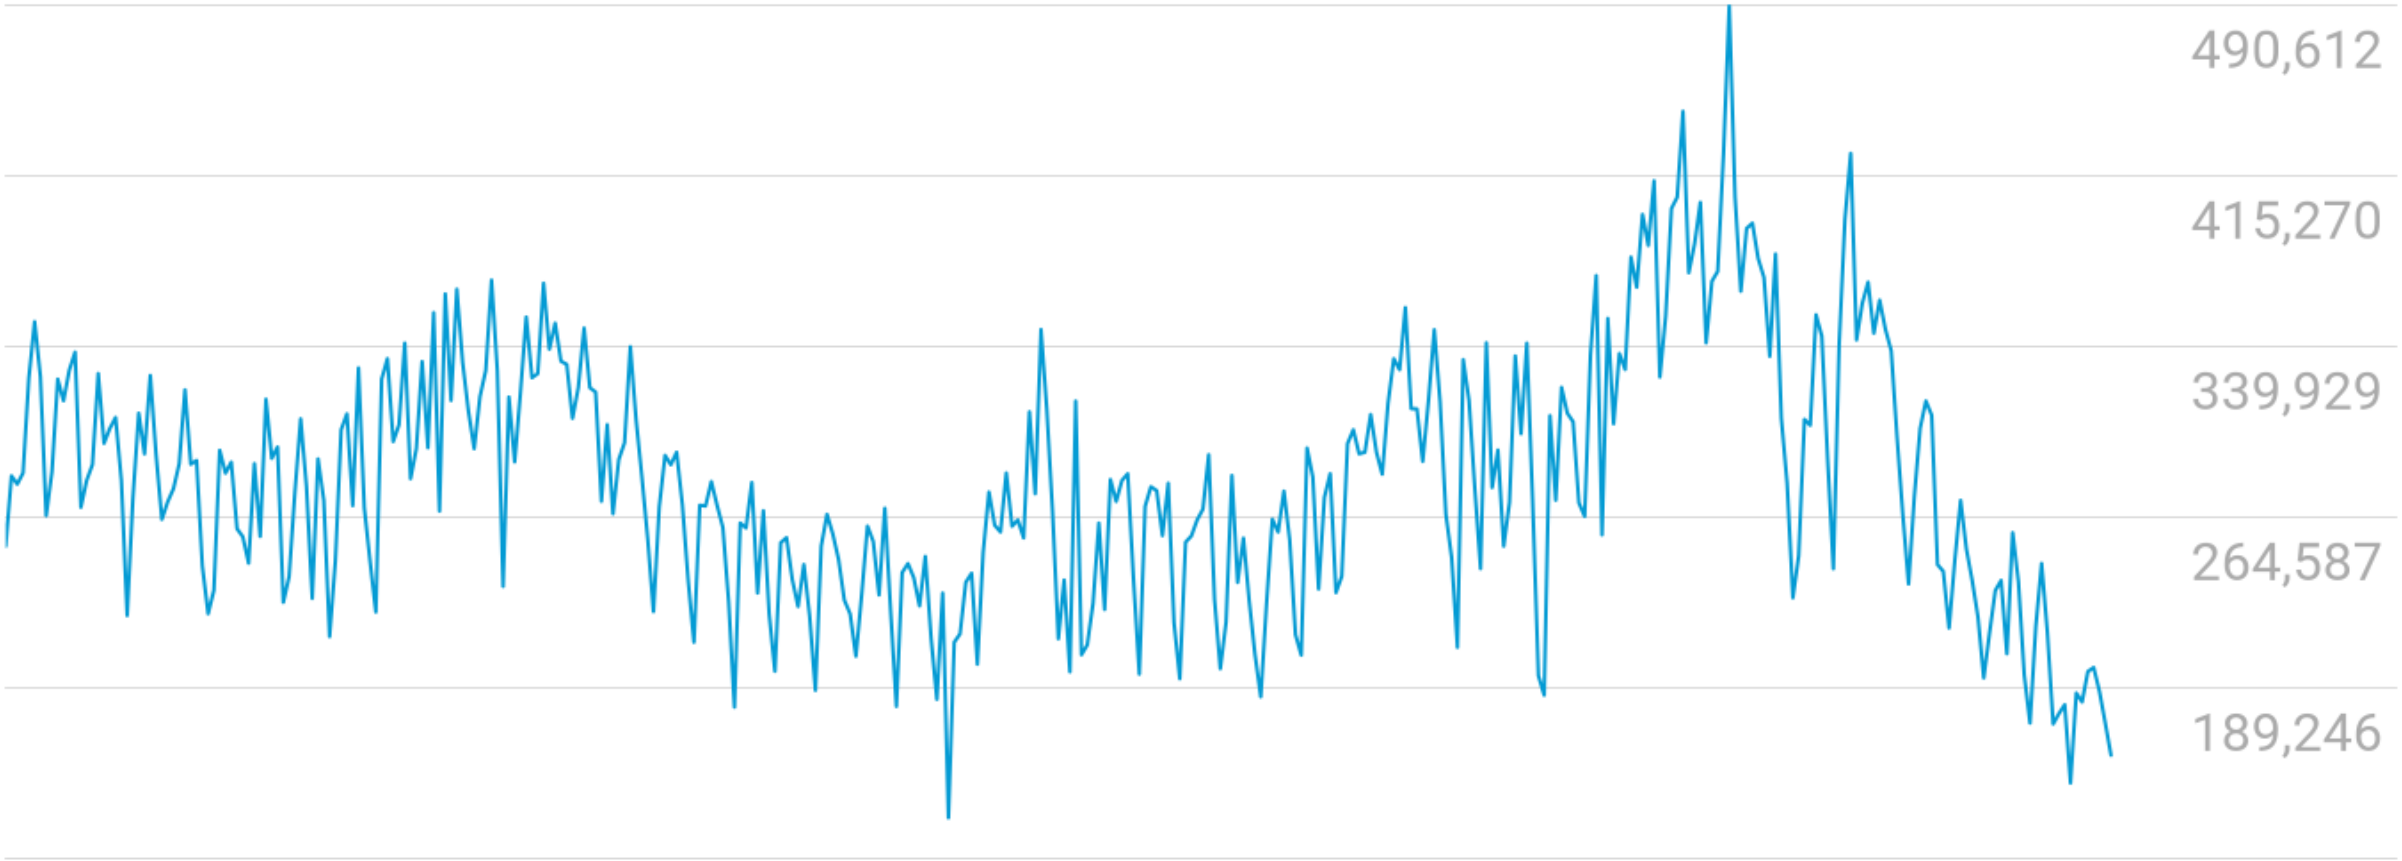

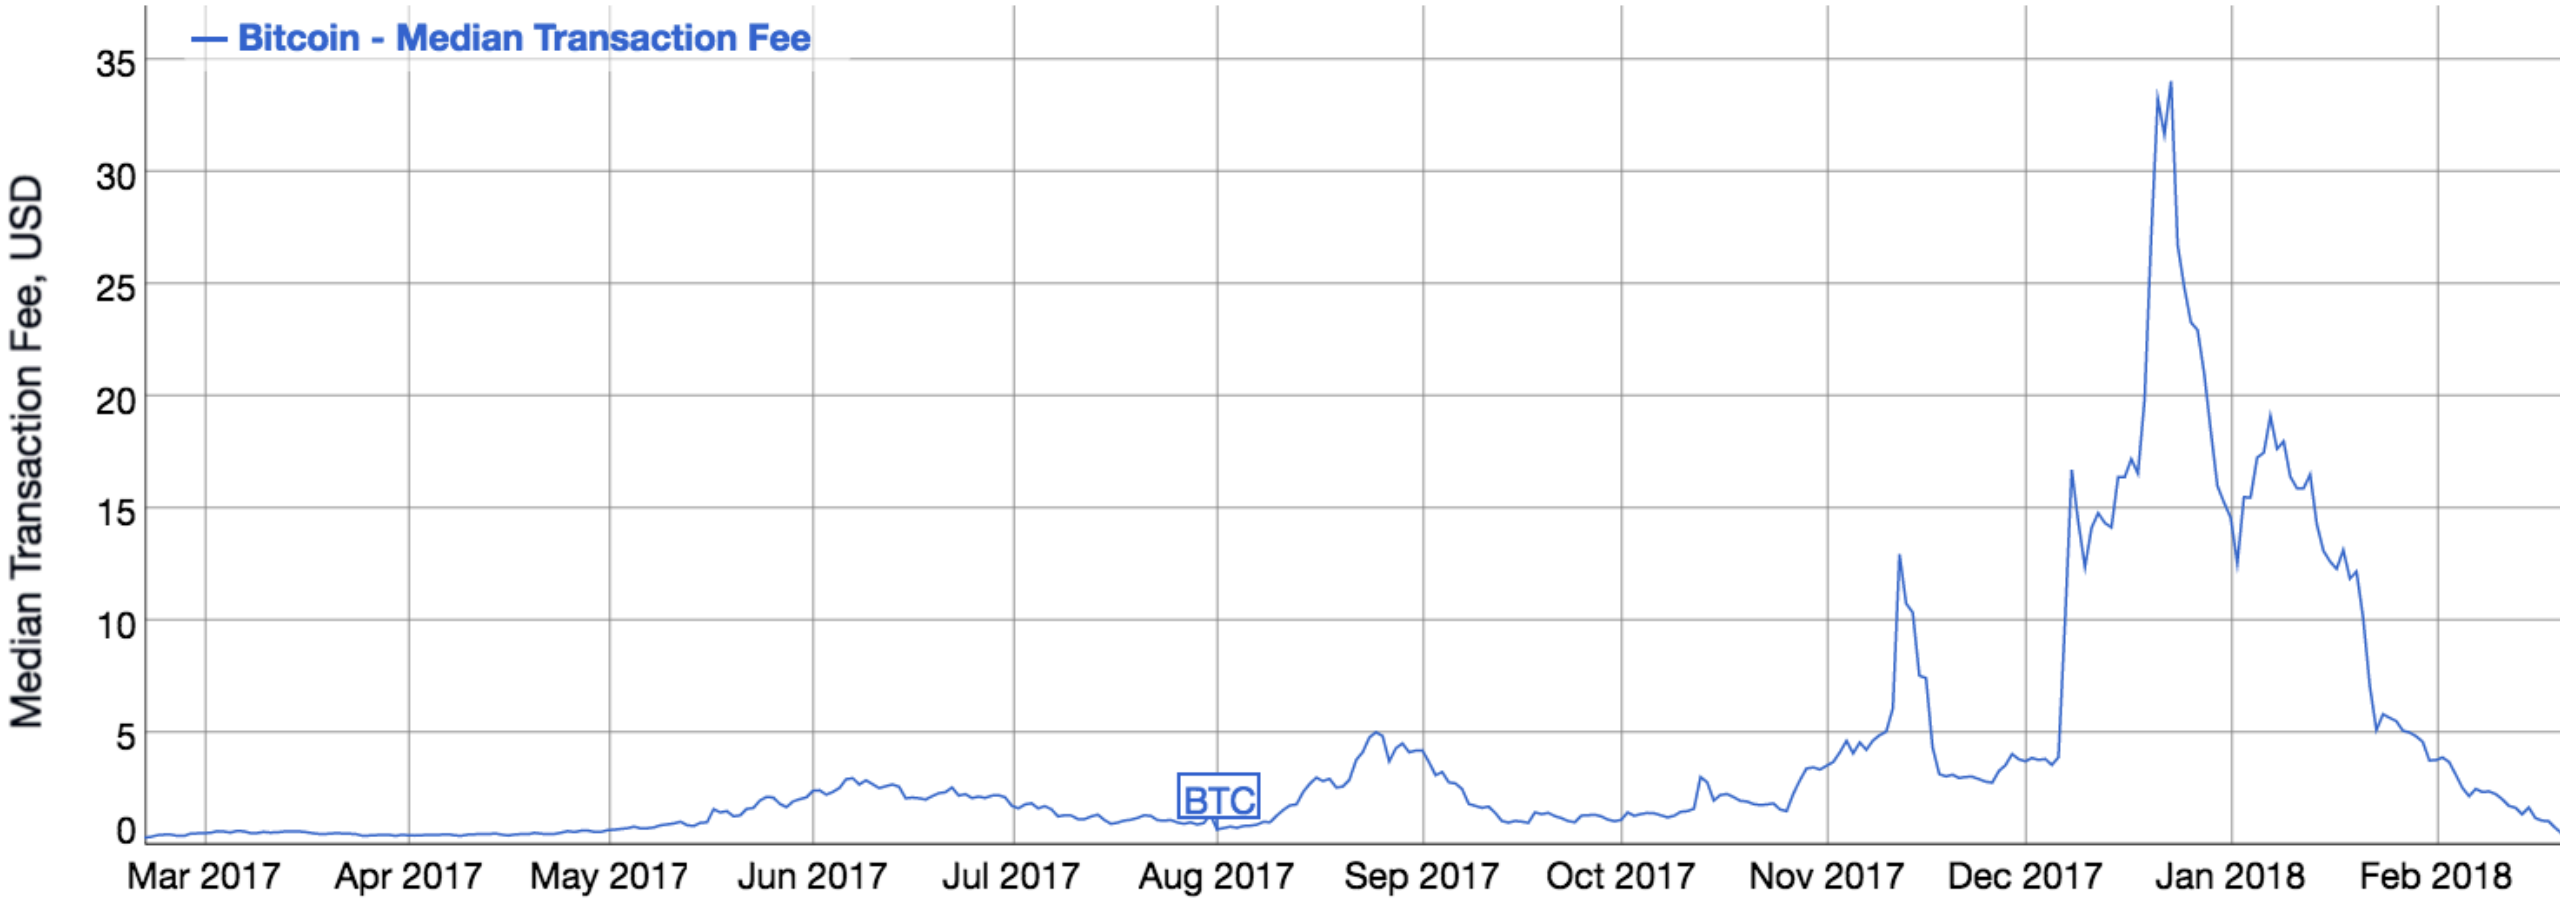

Current median transaction fee: 0.5-1 USD

Source: bitinfocharts.com

Read More

- The content of this lecture is based on the book: **Bitcoin and Cryptocurrency Technologies**
- The authors also have a course on Coursera