UCI University of California, Irvine

Blockchain Technical Overview

ML+BC Seminar

Oct 21 2019

Adrian Dabrowski

A data structure and protocol,...

- Creates trust between untrusted parties
  - Decentralized
  - Honest majority required
- Consensus Protocol
  - Eventual Consistency
  - Selection Protocol, e.g., Proof-of-Work
  - Peer-to-peer Network
- An append-only (Immutability) distributed database (ledger)
  - With terrible write performance
- Anonymous(?)

- Application: Cryptocurrency, Domain-name registration, ...

Cryptocurrency (e.g., Bitcoin)

- Proof of Ownership
- Proof of transfer (loss of ownership)
- Double-spending protection
- Chargeback-fraud
- In variable denomination
- Between untrusted (or dishonest) parties
- Permissionless

- Satoshi Nakamoto's Bitcoin 2008, Genesis-block Jan 3rd 2009
  - Global Limit 21M
  - 1 Bitcoin = 10^8 Satoshis
  - Bad Metaphor "Coin": no (virtual) token of value exchanged.

Blockchain

- New Blocks are “mined”, collecting transactions
- Based on probabilistic Proof-of-Work (PoW)
- Longest chain of valid blocks is “agreed” state
  - Starting at the genesis block
- Accidental forks
  - Two or more blocks are mined at the same time and broadcast
  - Network nodes randomly choose one chain to work on, until there is a “winner”
  - Some parties require a transaction to be “buried” at least *n* blocks deep to be “confirmed”
- Intentional forks (“hard” and “soft” forks)
  - Software or rule change

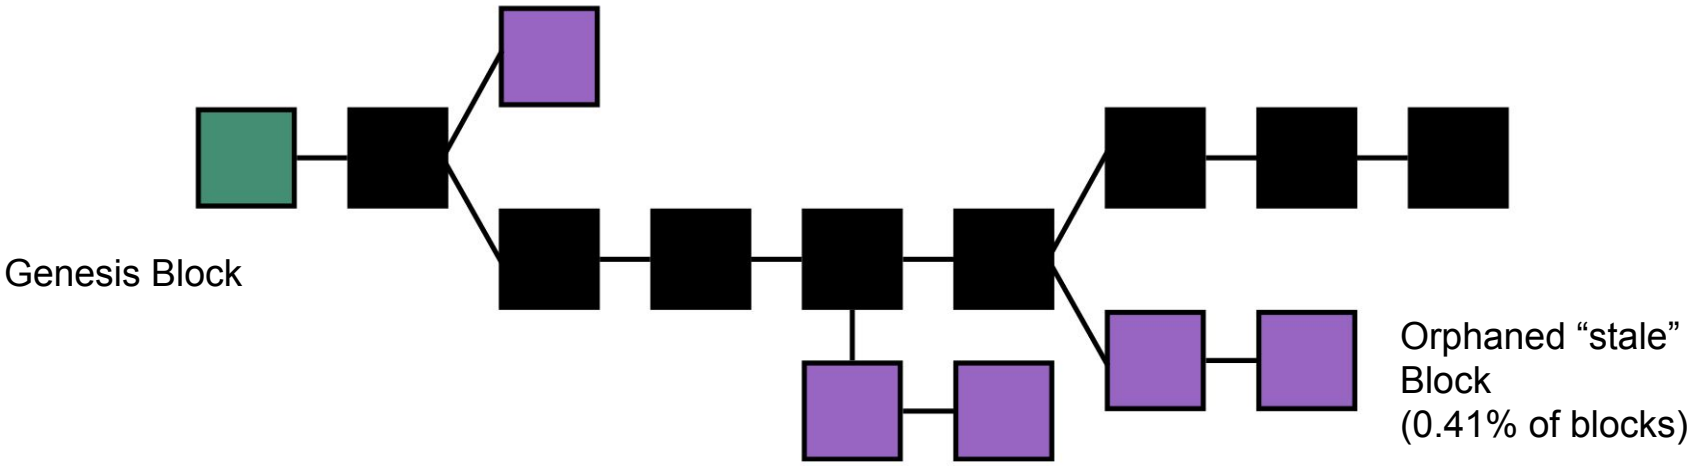

Transaction, Block, Address

- Transaction
  - N Inputs, M Outputs
  - Sum of Inputs >= Sum of Outputs
    - Exception: “Coinbase Transaction” coll. “Mining reward”
    - Leftovers usually parked at a “change address”
    - Difference between ∑Inputs and ∑Outputs is a transaction fee for the miner
  - Script (Bitcoin)
    - Stack-based non-turing-complete script, no loops
    - Typ. verifies signatures, before clearing of transfers

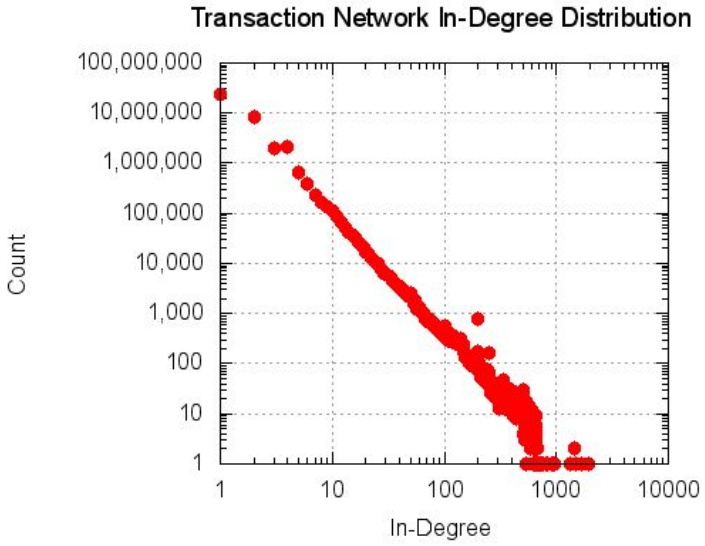

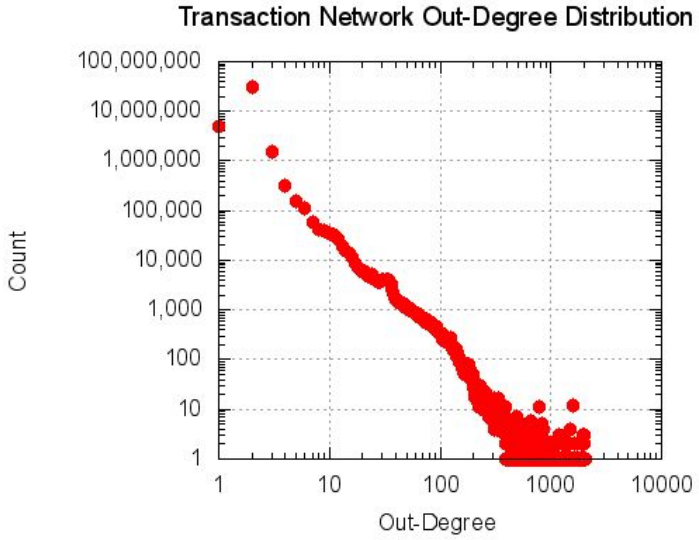

Transaction, Block, Address

- Block
  - Collection of transactions (organized as Merkle Tree)
  - One Coinbase Transaction
  - Hash of the parent block
  - Proof-of-Work Hash, adjustable complexity
  - ~1 MB per Block (Bitcoin)
- Address
  - Public/Private key pair
  - (Double-Hash of) public key is an “address”
  - Private key used to “unlock” the output of another transaction.
  - Often organized in **Wallets**
    - The wallet stores the keys or deterministically generates them on-the-fly based on one master-seed/secret.
  - Textual representation written in Base-58 representation with checksum
  - “Vanity Address”

Main Hash: $H_M(x) = SHA256(SHA256(x)) = SHA256^2(x)$

Address Hash: $H_A(x) = RIPEMD160(SHA256(x))$

Transaction Structure (Coinbase, Block)

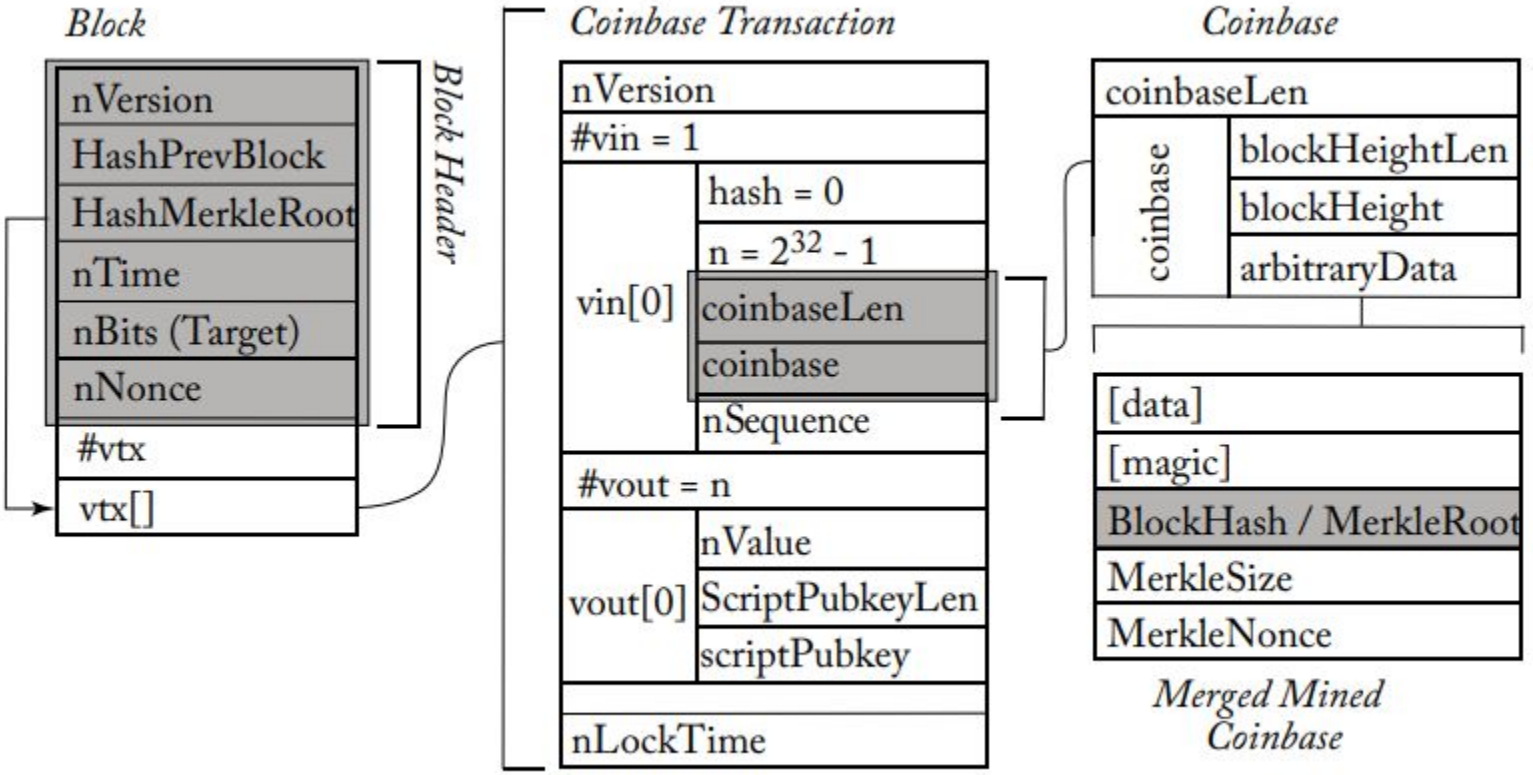

Source: Judmayer et al., Blocks and Chains: Introduction to Bitcoin, Cryptocurrencies, and Their Consensus Mechanisms, 2017

Transaction Structure (Simplified)

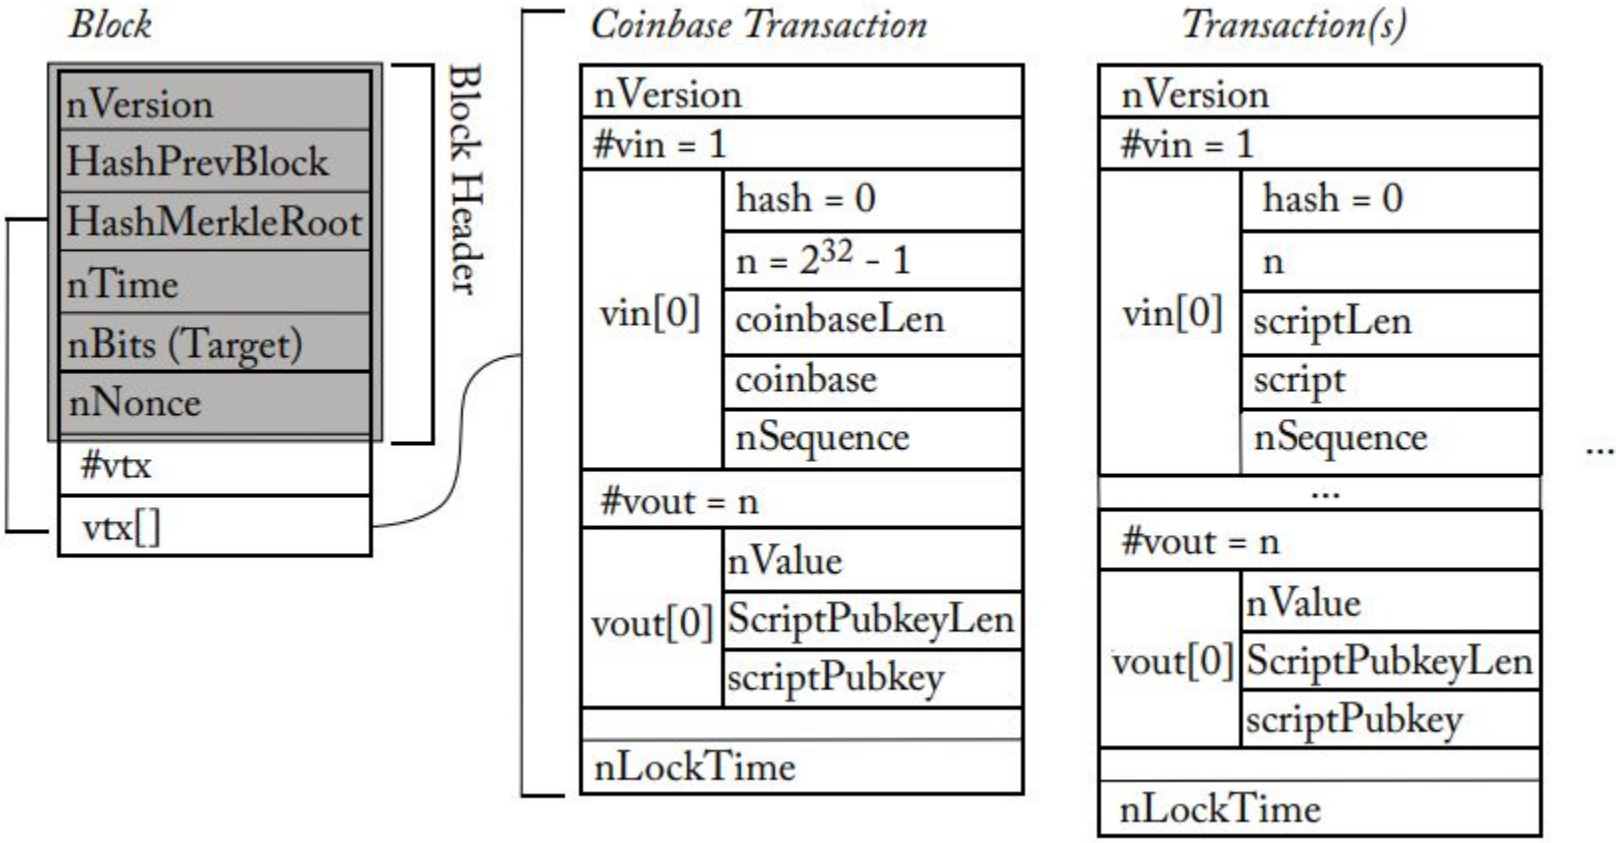

Source: Judmayer et al., Blocks and Chains: Introduction to Bitcoin, Cryptocurrencies, and Their Consensus Mechanisms, 2017

Transaction (Script)

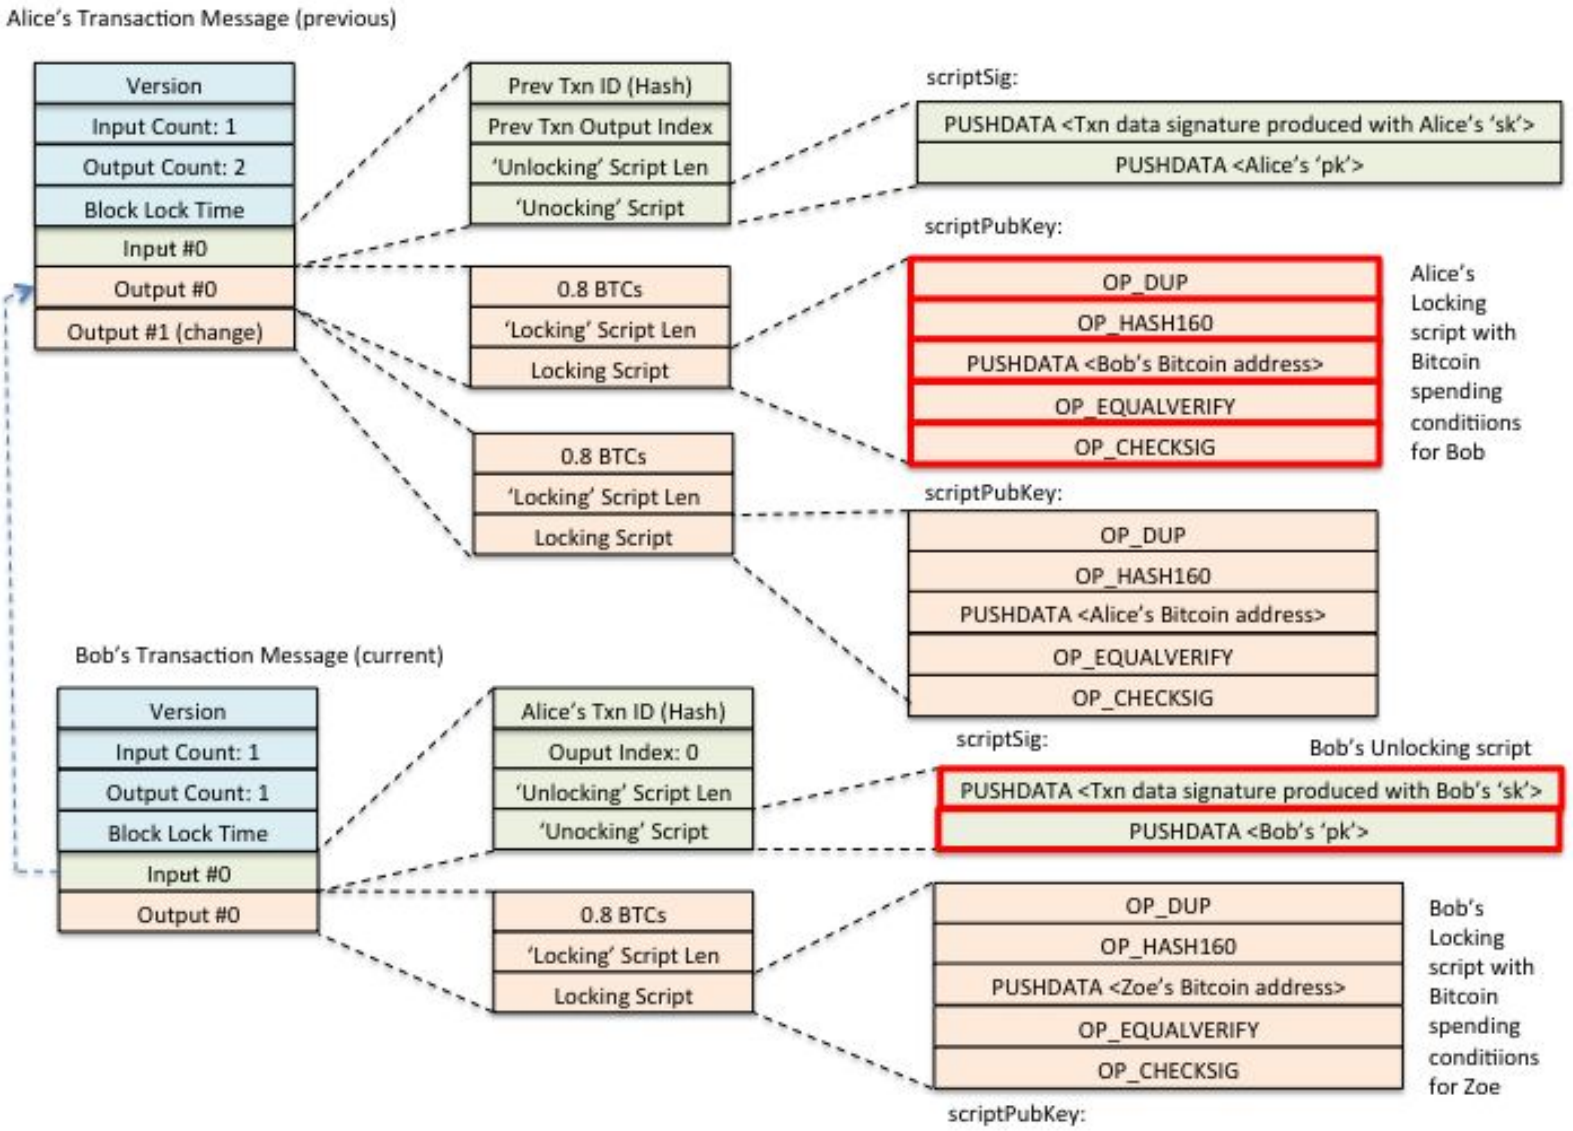

Pay-to-Public-Key-Hash (P2PKH) Script

Bitcoin scripting language example execution of P2PKH:

scriptPubKey (locks output)

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

scriptSig (unlocks output within input)

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

Language Reference: https://en.bitcoinwiki.org/wiki/Script

Nakamoto Consensus

- Trust without trusted third parties
- Distributed trust
- Decentralized trust
- Dynamic membership
- Fault tolerance

Proof of Work (PoW)

Properties

1. Easy to verify
2. Hard to generate
3. Difficulty is parameterizable
4. No reuse of previously generated PoW
5. No generation of PoW ahead of time (and use them later)

Bitcoin implementation:

- Hash puzzle
- Find a hash over the block with z leading zero bits (z: difficulty)
- Z adjusted periodically based on hash rate to have one block every 10 minutes
- Effectively a round master election/lottery based on hash rate

Difficulty

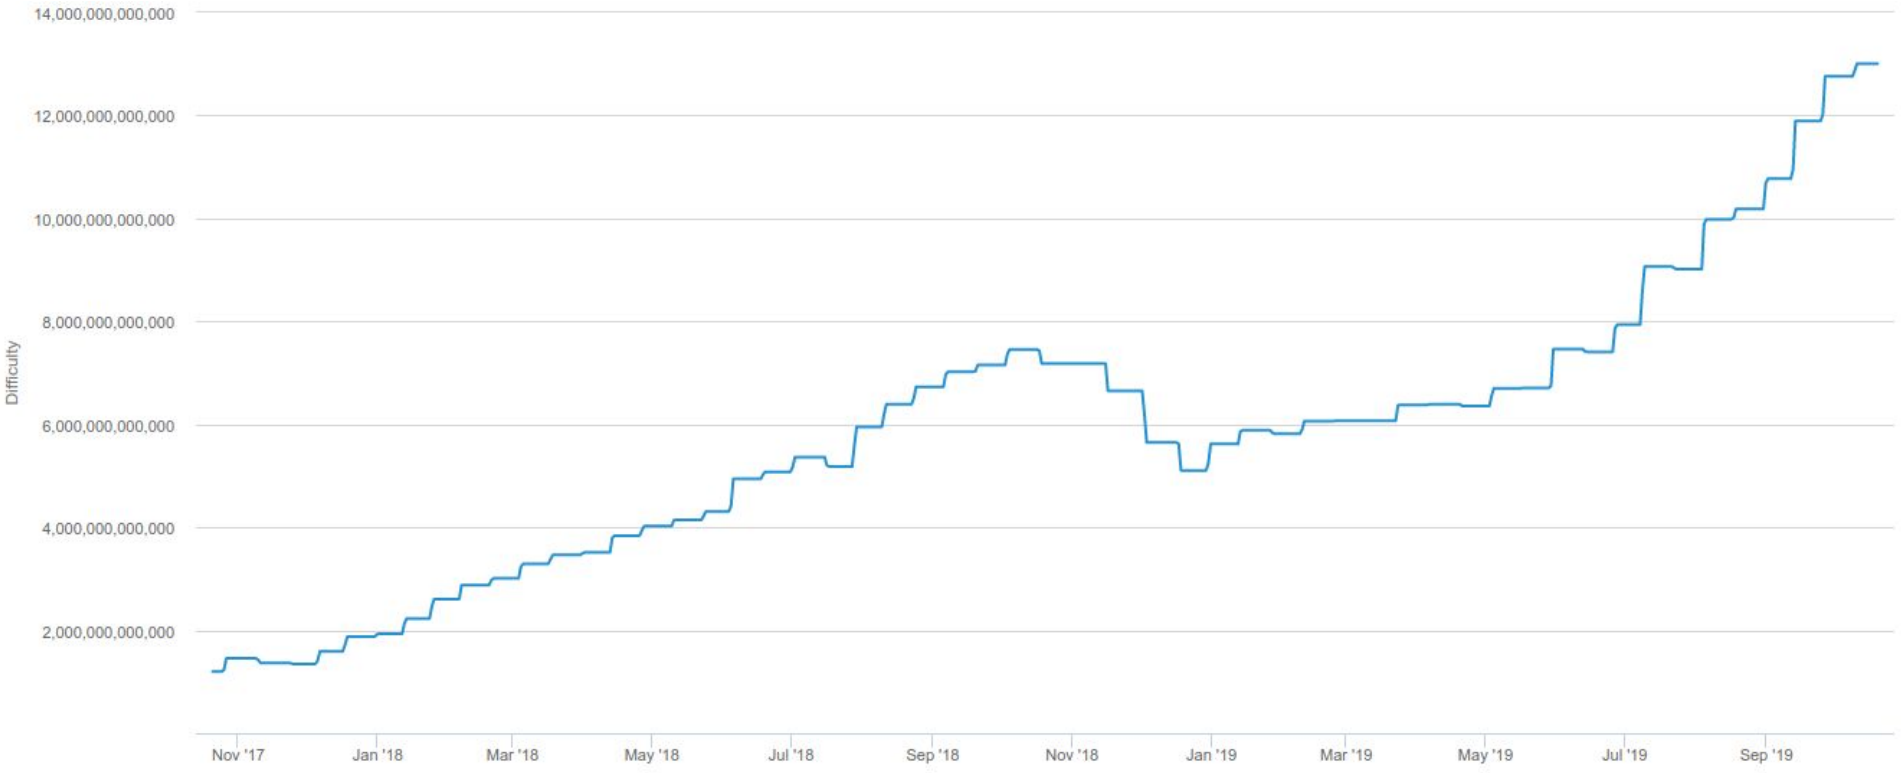
14,000,000,000,000
12,000,000,000,000
10,000,000,000,000
8,000,000,000,000
6,000,000,000,000
4,000,000,000,000
2,000,000,000,000
Difficulty
Nov '17
Jan '18
Mar '18
May '18
Jul '18
Sep '18
Nov '18
Jan '19
Mar '19
May '19
Jul '19
Sep '19

Mining

- New transactions and mined blocks are flooded into the P2P Network
- Transactions are stored in "**Mempool**" until used by miner
  - Miners sort by transaction fee
- Income stream
  - Block reward
  - Transaction fees
- Difficulty is adjustable
- Block reward is halved every 210,000 blocks
  - Finite supply of 21M Bitcoins
- Since finding the right solution is effectively a lottery, why not create lottery clubs? -> "Mining pool"
- Used to be on CPUs
  - Then GPUs
  - Today: Mining pools only accept ASIC miners

Era	Reward	Date
1	50 BTC	2009-01-03
2	25 BTC	2012-11-28
3	12.5 BTC	2016-07-09
4	6.25 BTC	-
...	...	-
33	0.00000001 BTC	-

Hash Rate

(Estimated)

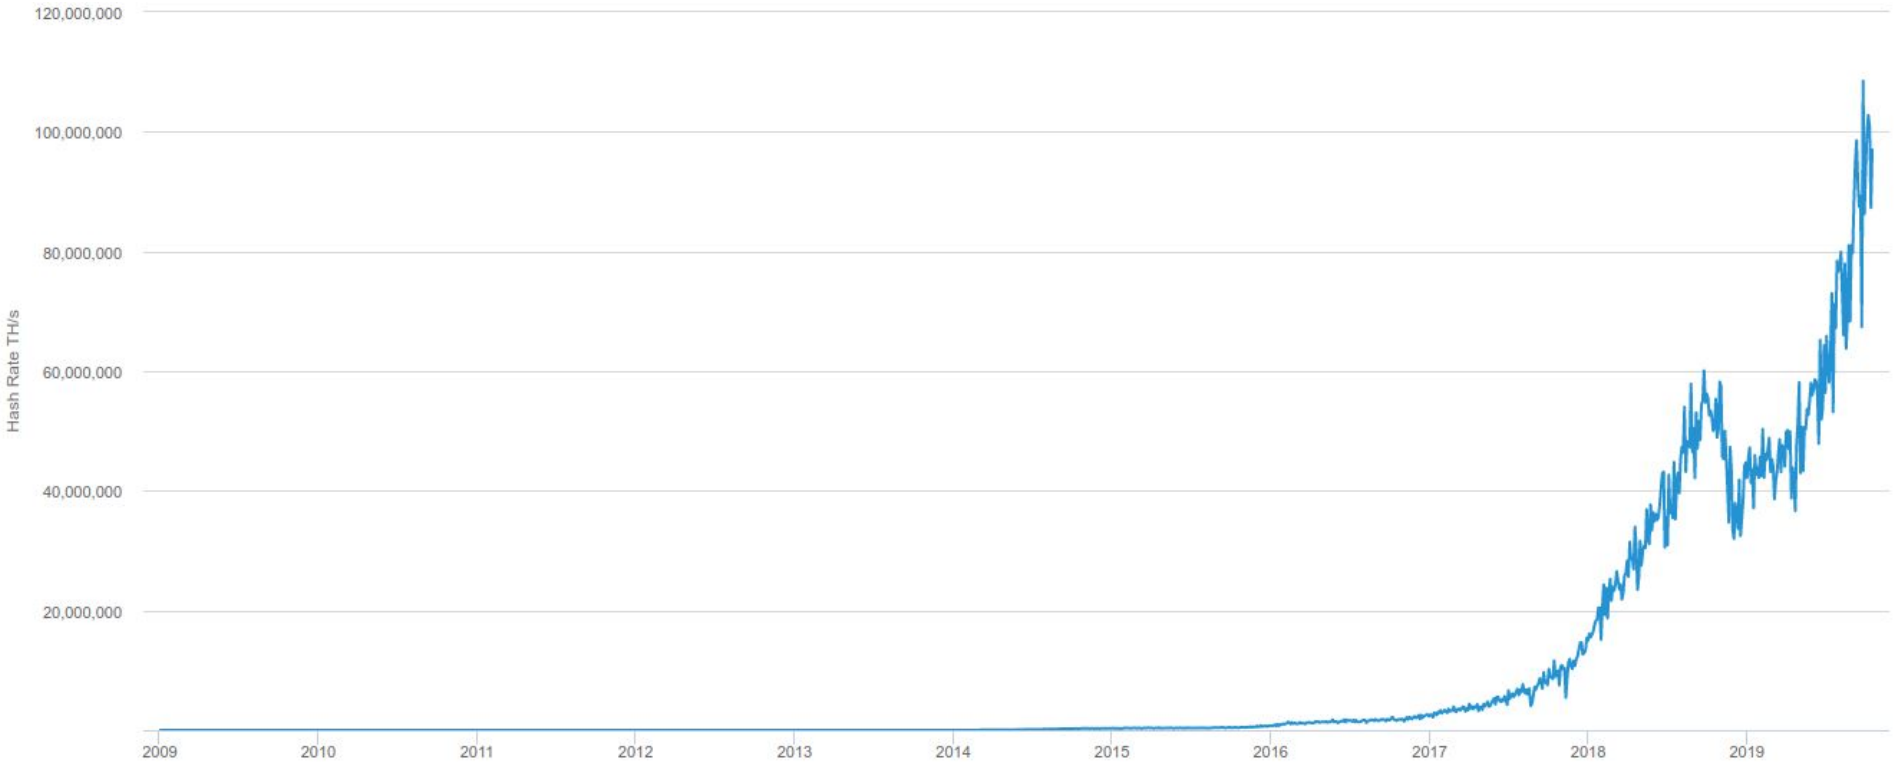

Mining Reward

(Block reward + transaction fees) * exchange value

Cost per Transaction

Mining Revenue / #Transactions

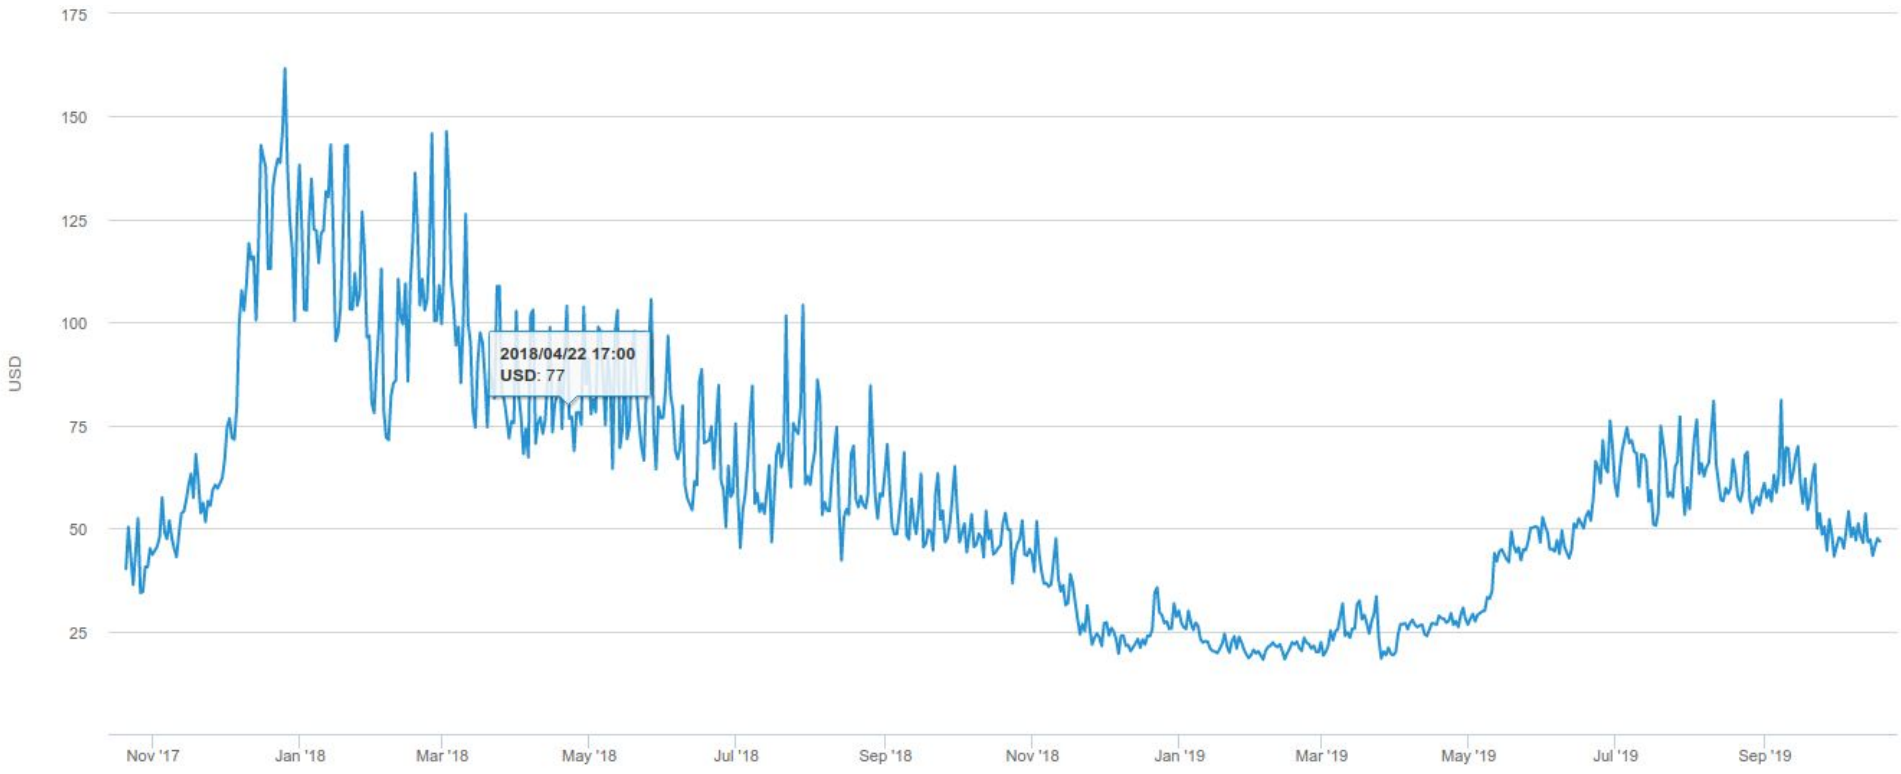

Mempool size

Number of transactions waiting for inclusion

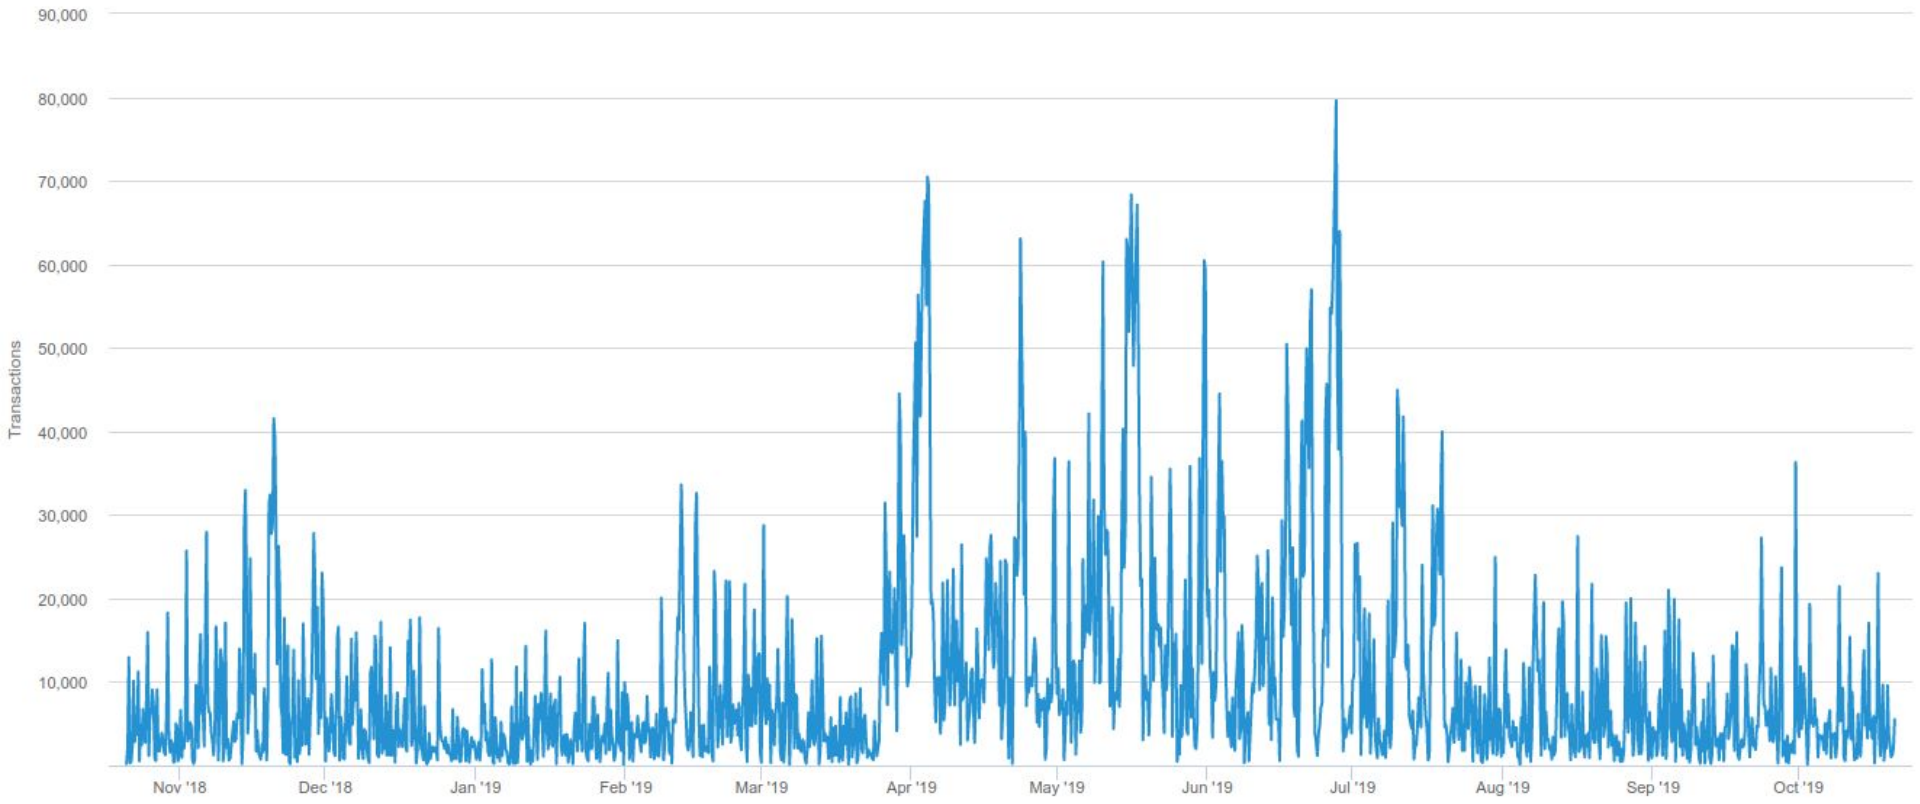

Bitcoin Client Types

P2P Network

- Initial seed via DNS
  - than P2P discovery
- Initial block download
- Broadcasting of transactions
- Broadcasting of Blocks

- 14 degrees of separation
- Default client accepts only one connection per /16

Transaction visibility

- On-chain transactions
  - persistent & visible
  - Two formats: Classic and SegWit
- Mempool
  - Non-persistent, but visible
- Off-chain transaction systems
  - E.g, lightning network
- Exchange-based transactions (or other trusted third parties)
  - Transaction traded entirely within a market system of an exchange
  - Transfers into and out of an exchange can happen from different pool accounts collecting money from thousands of customers
  - Majority of fiat money exchange

Scalability & Resource Consumption Problems

- Transaction-rate ~5 per second
- More Hashpower does not translate into higher transaction rates
  - Difficulty adjusted
  - But uses more power
  - Tragedy of the commons
- Multiple solutions proposed
  - Segwit accepted

Replace of PoW

- Essentially a selection algorithm to choose a round master
- Alternatives
  - Useful PoW
  - Proof of Stake
  - Distributed verifiable random number generation

Bitcoin Attacks (Selection)

- Insufficient randomness on Wallets
- DNS Seed
- 51% Attack
  - Unhonest miner(s) with enough hashing power can “rewrite” history
  - Mining pools limit their hash power
- Eclipse Attack
  - Have a number of (powerful) miners connect only to attacker-controlled nodes.
  - Attacker can control on which state and transactions of the blockchain the miners dedicate their power. Allows double-spending
- Sybil Attack
  - Variant of above, attacker spreads malicious nodes over the network
- Race Attack
  - Try to outrun a posted transaction by (faster) broadcasting a competing transaction, also: increase transaction fee
- Selfish / Stubborn Mining
  - A found block is not broadcasted but used to find the next block, later revealed.
  - Maximizing block reward
  - -> **n confirmation attack**