HS18

DISTRIBUTED SYSTEMS ADVANCED & BLOCKCHAIN

IOTA / PoS / Bazo

T. Bocek

thomas.bocek@hsr.ch

HSR
HOCHSCHULE FÜR TECHNIK
RAPPERSWIL

FHO Fachhochschule Ostschweiz

Distributed Systems - In the *News*

- **29.10.2019 Facebooks Digitalwährung Libra und die Zentralbanken: Regulieren oder selbst mitmischen?**
  - «Die Zentralbanken müssten vielleicht selbst ins Stablecoin-Geschäft einsteigen.»
  - «Ganz vorn dabei ist die People's Bank of China»
  - «A former high-ranking Congressional official said China's central bank is likely to be the first to issue its own national digital currency, as he contended that Libra was doomed to fail.» (28.10.2019)
- **28.10.2019 Blockchain-Technik jenseits von Kryptogeld**
  - Grundlagen, Industriepraxis, Eigenbau
- **China crypto currency news increase**
  - 28.10.2019 China Wants Communist Party Members to Pledge Loyalty on Blockchain
  - 28.10.2019 Chinese Official Warns Libra Could Abet Illegal Cross-Border Transfers
  - 28.10.2019 Chinese Central Bank Official Calls for Commercial Bank Blockchain Adoption
  - 28.10.2019 From Banking Giants to Tech Darlings, China Reveals Over 500 Enterprise Blockchain Projects
  - 26.10.2019 China's Congress Passes Cryptography Law, Effective Jan. 1, 2020
  - 25.10.2019 President Xi Says China Should 'Seize Opportunity' to Adopt Blockchain
- **25.10.2019 Ethereum Targets Dec. 4 for Istanbul Mainnet Activation**
  - ~~Scheduled for 04.12.2019 (changes)~~ (ProgPoW)

Introduction

- **What is IOTA?**
  - “A permissionless distributed ledger for a new economy” - for IoT
  - Zero transaction fees / feeless microtransactions and data integrity for machines
    - A node can only issue a transaction after solving a cryptographic puzzle
  - Quantum-proof algorithms
  - “It is safe to receive any number of transactions to a given address until an outgoing transfer (a "send" transfer) is made. After that, this address should no longer be re-used!” (source)
    - Sending IOTAs - part of the private key of that address is revealed (brute force)
- **An Open-Source Distributed Ledger**
  - Uses a Decentralized Acyclic Graph (DAG) instead of a blockchain (Tangle)
  - Tangle: transaction must approve two previous transactions in order to be considered valid.
    - Perform a small amount of computational work

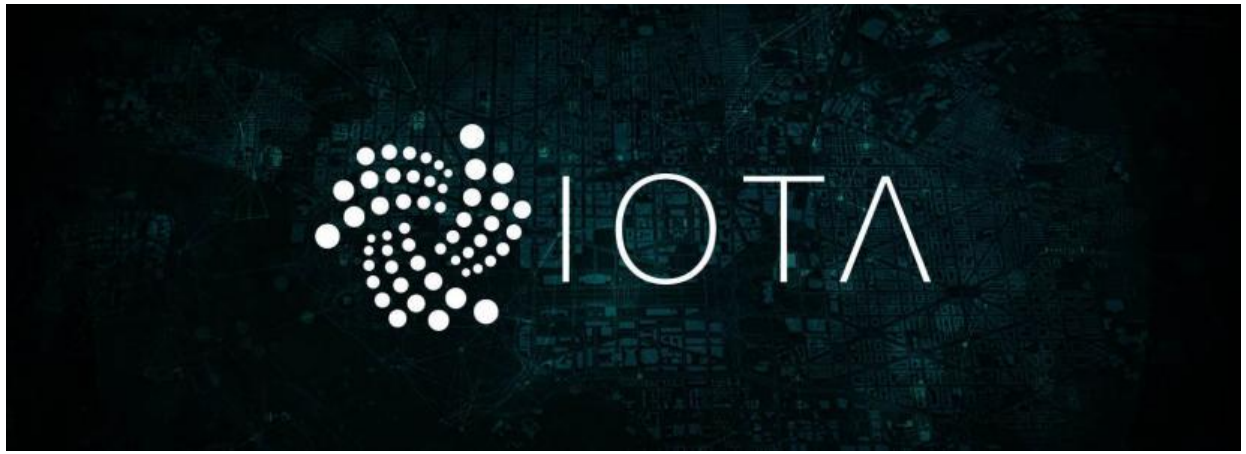

Tangle

- **Blockchain**
  - Blocks with transactions linked to each other
  - Decentralized, used in: Bitcoin, Ethereum (uncle blocks), NEO
    - Uncle blocks: reward for duplicate but non-conflicting blocks
  - Transactions / sec limitations
  - The older a tx in a block, the more secure

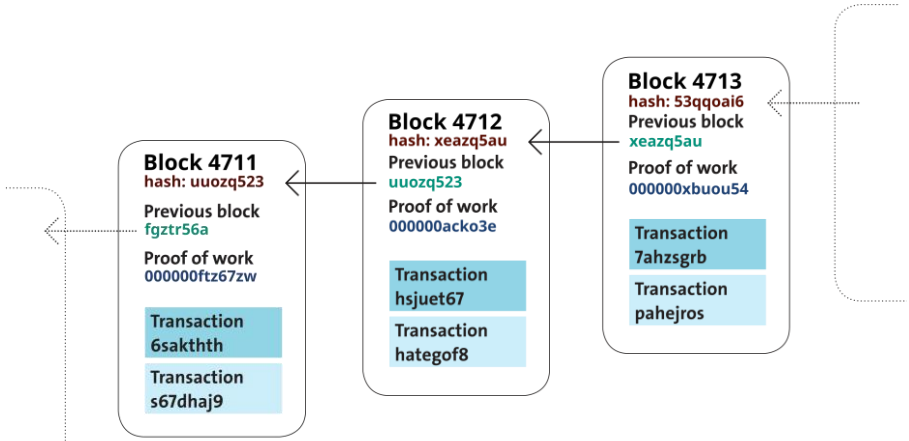

- **DAG**
  - Transactions linked to two previous tx
  - Regular milestones (centralized? coordicide?)
  - More activity = more validation (but spam attacked happened in the past)
  - The more confidence in a tx, the more secure
    - Tip selection: weighted random walk from the genesis towards the tips
    - Run the tip selection algorithm 100 times.
    - Count how many of those 100 tips approve our transaction, and call it A.
    - The confirmation confidence of our transaction is “A percent”.

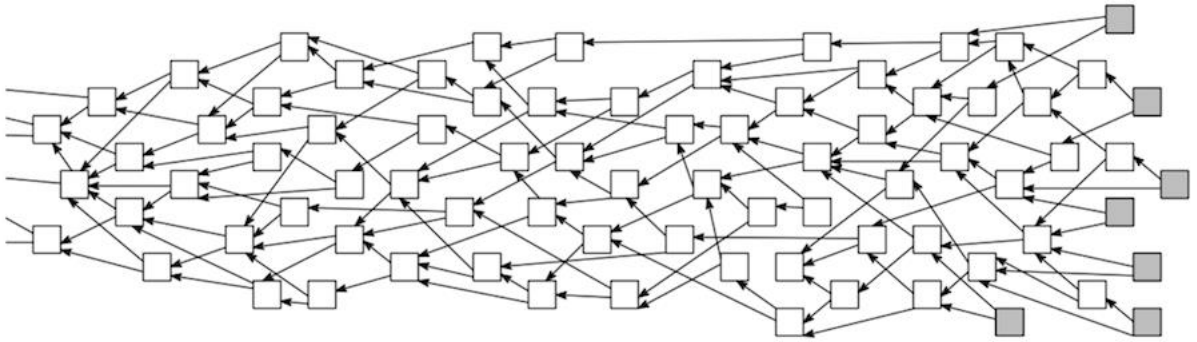

IOTA Coordinator / Snapshots

- **Milestone: every ~64s the Coordinator makes a normal transaction, confirming 2 tx**
- **Want to know if your tx is verified, get latest milestone and check**
- **Eventually will be switched off, now active due to small network**
  - "The only role the Coordinator serves is to protect against attacks in this temporary infancy stage of the Tangle ledger,.."
- **"IOTA is currently in what should be considered a 'transition period' towards large scale deployment and standardization."**

- **IOTA Snapshot**
  - Pruning: "This is essentially a "pruning" of the ledger - it removes all events and addresses on the ledger which do not have a positive balance."
  - "As you probably already know, we are performing global snapshots of the IOTA network on a more or less regular basis. The main motivation for global snapshots is to prune the database and allow smaller nodes to keep participating in the Mainnet."
  - Local snapshots in future

IOTA Wallet

- **Never write your own crypto algorithms**

- **Wallet: https://trinity.iota.org/**
  - Reference implementation in Java: https://github.com/iotaledger/iri
  - API in many languages: C#, Javascript, …

- **More documentation**
  - https://docs.iota.org/docs/getting-started/0.1/introduction/what-is-iota
  - https://blog.iota.org/the-tangle-an-illustrated-introduction-79f537b0a455

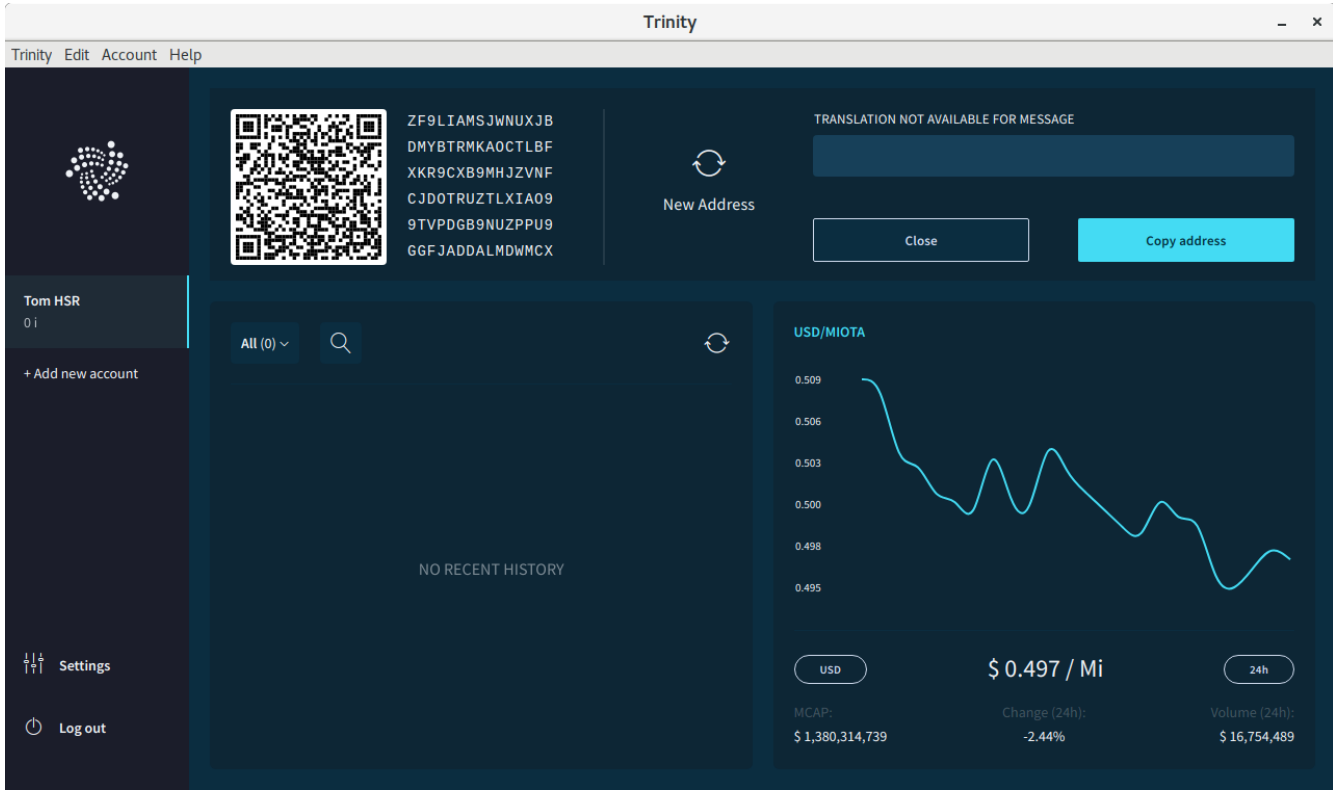

URL: b.socrative.com/student/

Room: HSR

Agenda

- **Consensus**
  - Definition
  - Consensus Types
  - Blockchain Consensus (PBFT, PoW, PoS)
  - Consensus in the Bazo Blockchain
- **Scalability**
  - Definition
  - Motivation and Problems
  - Solutions (Plasma, State Channels, Sharding)
  - Introduction to Sharding
  - Scalability in the Bazo Blockchain

What is Consensus?

- **Definition: Consensus decision-making is a group decision-making process in which group members develop, and agree to support a decision in the best interest of the whole.**
- **Consensus may be defined professionally as an acceptable resolution, one that can be supported, even if not the "favourite" of each individual.**
- **Different ways to reach consensus:**
  - Person in-charge decides
  - Executive committee decides
  - Simple Majority (>50%)
  - Super Majority thresholds (90%, 80%, 75%, two-thirds, …)
  - ...

Blockchain Consensus

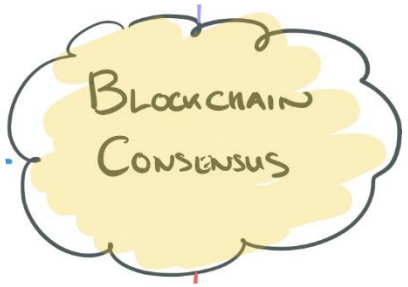

- **Classical Consensus Models**
  - Crash failure models → honest nodes failing
  - Byzantine failure model → able to tolerate a portion of malicious nodes
- **Elected Leader**
  - Probabilistic elected leader (e.g., who can find the hash first?)
  - Most known Proof-of-Work (PoW)
  - Also, Proof-of-Stake (value held on the chain and its variants), Proof-of-Capacity (PoC), Proof-of-Burn (PoB), Proof-of-Authority (PoA), etc.
- **Hybrid Consensus Models**
  - Single consensus has many limitations
  - Combine different consensus mechanisms
- **Scalability**
  - System is divided into shards (communities)
  - Cross-chain communications

Source

Blockchain Consensus

- **Nodes must evaluate and agree on all data before they are permanently incorporated into the blockchain - it is vital that all new information must be reviewed and confirmed before being accepted.**
- **Algorithms to discuss:**
  - Practical Byzantine Fault Tolerance (PBFT)
    e.g. Hyperledger, Stellar
  - Proof of Work (PoW)
    e.g. Bitcoin, Ethereum
  - Proof of Stake (PoS)
    e.g. Peercoin, Ethereum's Casper, Bazo

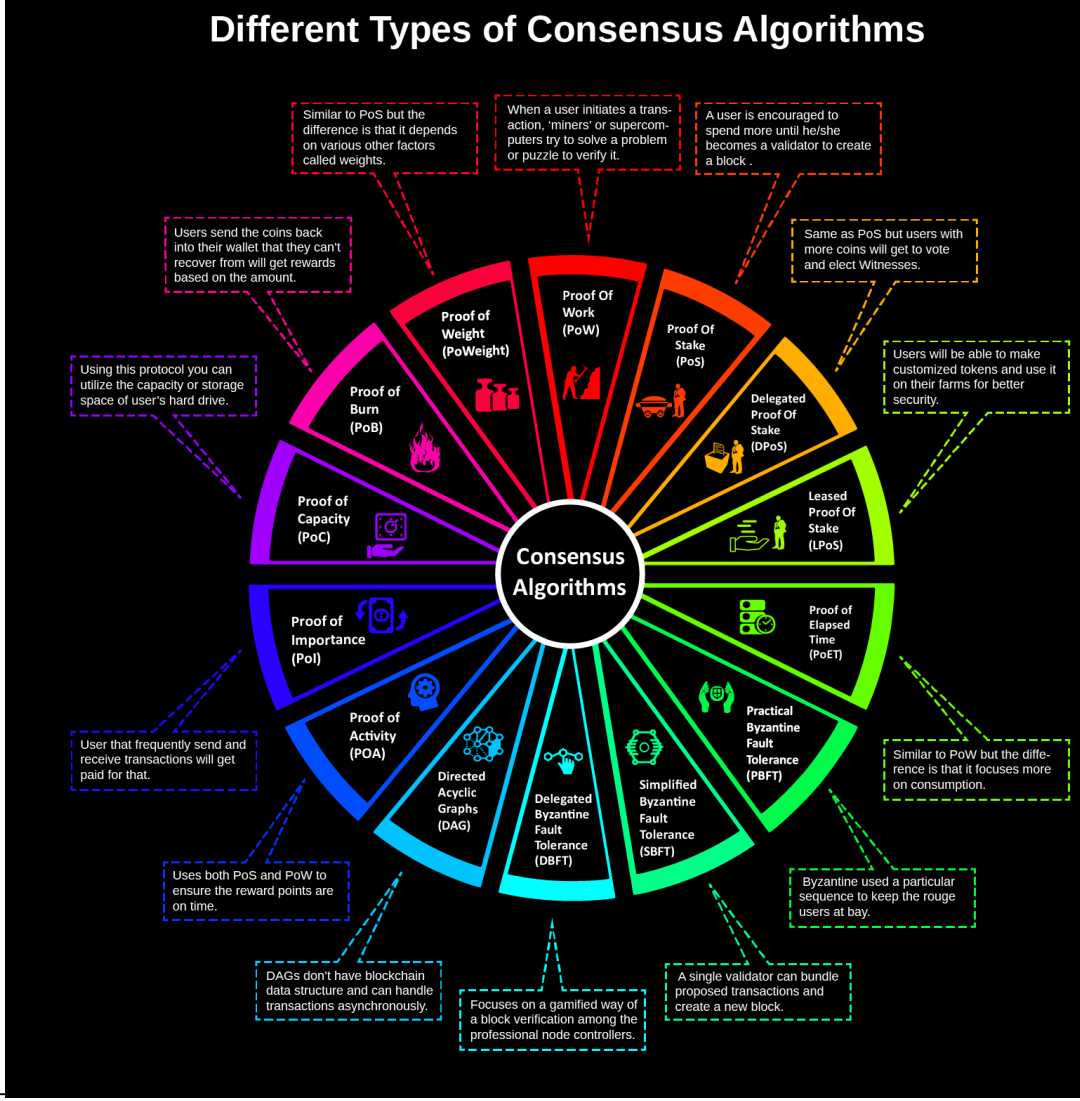

Source: 101blockchains.com

URL: b.socrative.com/student/

Room: HSR

The Byzantine Generals Problem

- **First referenced in the paper "The Byzantine Generals Problem", published in 1982**
- **How do you make sure that multiple entities, which are separated by distance, are in absolute full agreement before an action is taken?**

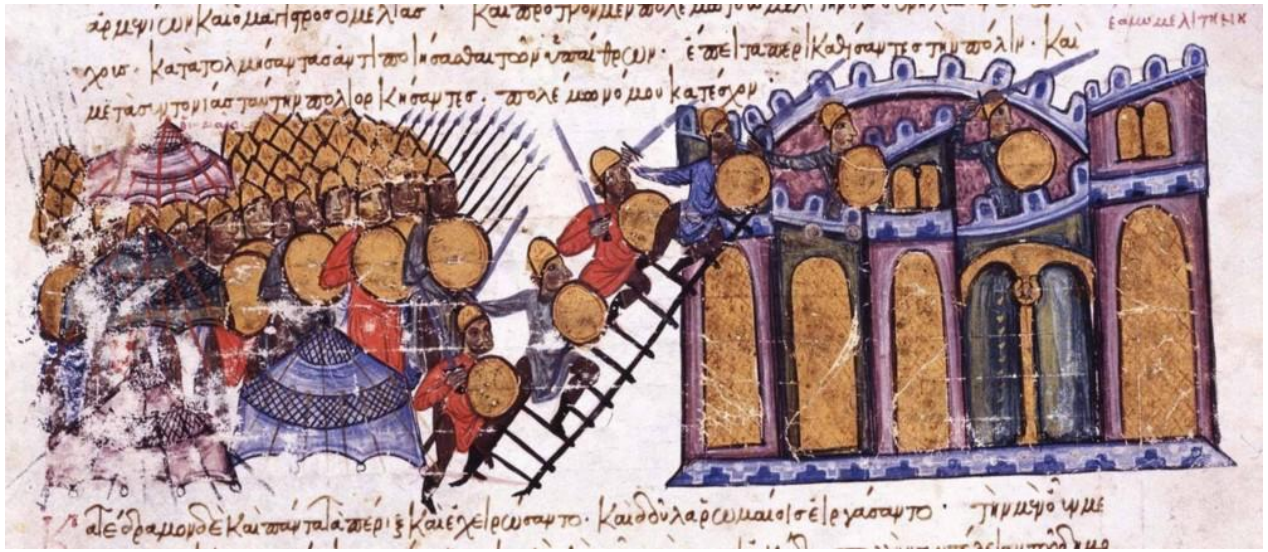

The Byzantine Generals Problem

LESLIE LAMPORT, ROBERT SHOSTAK, and MARSHALL PEASE
SRI International

Reliable computer systems must handle malfunctioning components that give conflicting information to different parts of the system. This situation can be expressed abstractly in terms of a group of generals of the Byzantine army camped with their troops around an enemy city. Communicating only by messenger, the generals must agree upon a common battle plan. However, one or more of them may be traitors who will try to confuse the others. The problem is to find an algorithm to ensure that the loyal generals will reach agreement. It is shown that, using only oral messages, this problem is solvable if and only if more than two-thirds of the generals are loyal; so a single traitor can confound two loyal generals. With unforgeable written messages, the problem is solvable for any number of generals and possible traitors. Applications of the solutions to reliable computer systems are then discussed.

Categories and Subject Descriptors: C.2.4. [**Computer-Communication Networks**]: Distributed Systems—*network operating systems*; D.4.4 [**Operating Systems**]: Communications Management—*network communication*; D.4.5 [**Operating Systems**]: Reliability—*fault tolerance*

General Terms: Algorithms, Reliability

Additional Key Words and Phrases: Interactive consistency

1. INTRODUCTION

A reliable computer system must be able to cope with the failure of one or more of its components. A failed component may exhibit a type of behavior that is often overlooked—namely, sending conflicting information to different parts of the system. The problem of coping with this type of failure is expressed abstractly as the Byzantine Generals Problem. We devote the major part of the paper to a discussion of this abstract problem and conclude by indicating how our solutions can be used in implementing a reliable computer system.

We imagine that several divisions of the Byzantine army are camped outside an enemy city, each division commanded by its own general. The generals can communicate with one another only by messenger. After observing the enemy, they must decide upon a common plan of action. However, some of the generals

The Byzantine Generals Problem (Repetition)

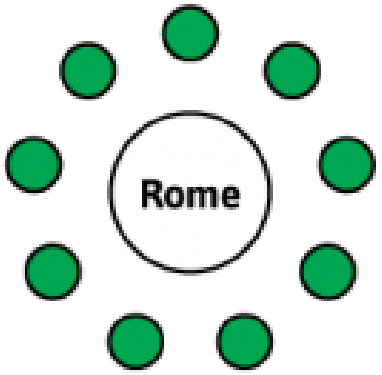

Imagine Rome being besieged by **nine armies**, each commanded by a (Byzantine) general.

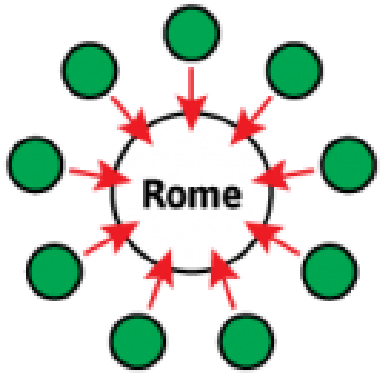

In order to launch a succesful **attack** or retreat, all armies have to **do the same**, otherwise they will be decimated by Rome's armies.

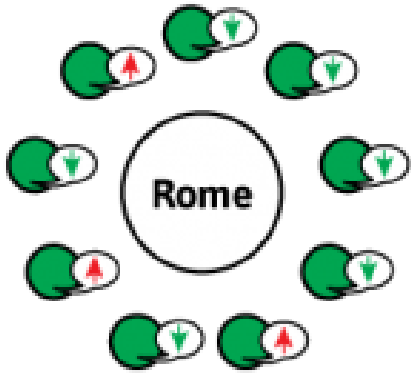

The decision to either **attack** or retreat is put up to a **vote**. Whichever option receives **more than 50%** of the votes, that's what the Generals will do (**retreat** in the example above).

Source: CryptoGraphics.info

The Byzantine Generals Problem (Repetition)

Imagine Rome being besieged by **nine armies**, each commanded by a (Byzantine) general.

In order to launch a succesful **attack** or retreat, all armies have to **do the same**, otherwise they will be decimated by Rome's armies.

The decision to either **attack** or retreat is put up to a **vote**. Whichever option receives **more than 50%** of the votes, that's what the Generals will do (**retreat** in the example above).

Problem 1
The generals communicate by using **couriers**, who have to cross unknown areas controlled by the Romans, risking capture or their message becoming corrupt.

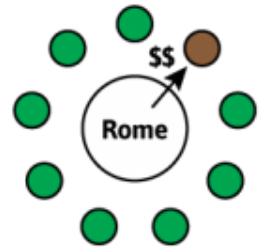

Problem 2
Each of the generals could be bribed by the Romans: **Traitorous Generals**.

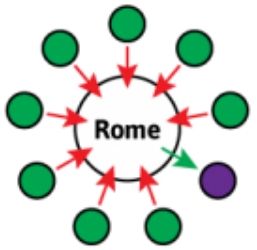

Problem 3
Any of the Generals can make the wrong decision, regardless of the vote: **Improperly Functioning Generals**.

- **"The proof-of-work chain is a solution to the Byzantine Generals' Problem." – Satoshi Nakamoto**
- **It's a probabilistic solution to the Byzantine Generals Problem, which means the confidence that a consensus is reached is growing with every block added to the chain, but it never reaches 100%.**
- **Practical Byzantine Fault Tolerance (PBFT): If more than ⅓ of the network's nodes are malicious, it does not work (sybil attack)**

Source: CryptoGraphics.info

Blockchain and the Byzantine Generals Problem

1. **A General solves the PoW problem, creating a block that is broadcasted to the network**
2. **Following receipt of this block, each General verifies and works on solving the next PoW problem, incorporating the prior solution into it, so that their plan adds on to the previous resolution**
3. **Each time a General solves a PoW problem, a block is generated and the chain begins to grow. In time, any General working on a different solution will switch over to the longest chain. This is the one most Generals are contributing to and therefore has the greatest chance of success**
4. **As the Generals know roughly how long a PoW solution takes to solve, after a set amount of time they will know if enough of the other Generals are also working on the same chain**

Through this process, the Generals can arrive at a consensus of when to attack, can estimate their chances of successfully doing so, and can prevent multiple different signals to attack being sent simultaneously.

Proof of Work (Repetition)

- **Each node of the network has to solve a difficult mathematical problem to propose a block**
- **Since it takes significant computing power to solve this problem, it proves that the node has done sufficient work to produce the block**
- **Bitcoin: The digest of the hash function must start with a number of zeroes (simplified!)**

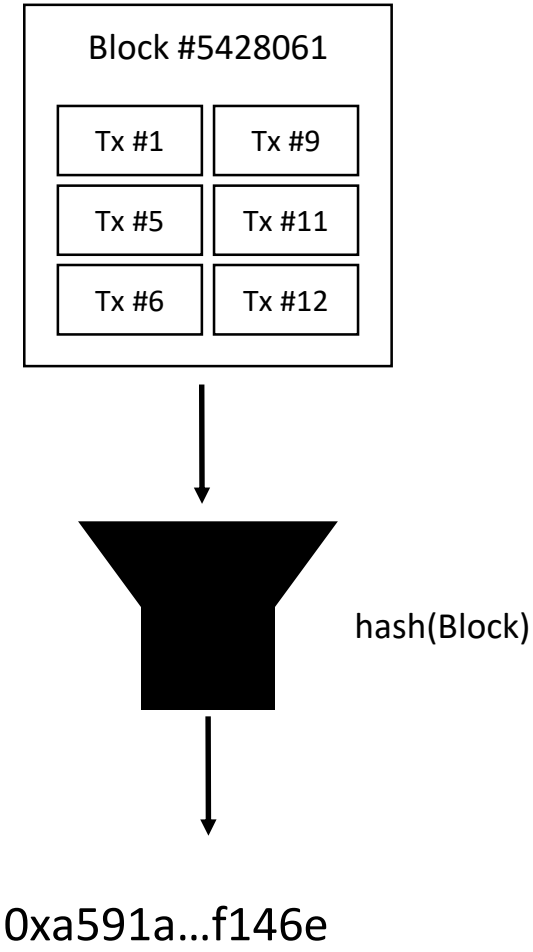

Recall: Hash Functions (Repetition)

- **A hash function (like SHA-256) takes a block of data in and produces an effectively random fixed size integer**
- **For the same input, the hash function always yields the same output.**
- **Any change to the input randomizes it**
  - “Simple” input changes, e.g., one char, lead to “dramatic” result changes

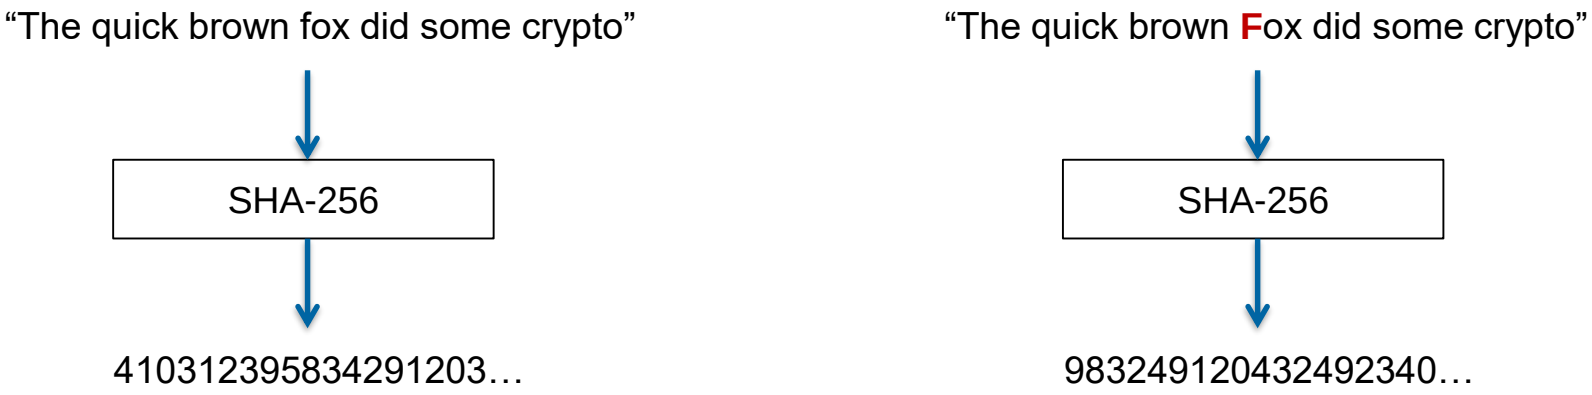

Recall: Bitcoin Header

- **Version number: Used to keep track of upgrades and changes in the protocol**
- **Previous block header hash: Linkage to the previous block, secures the chain**
- **Timestamp: Number of seconds since the first of January 1970**
- **Merkle Root: Root hash of the generated Merkle tree**
- **Nonce: Used to vary the output of a cryptographic function. It is one of the values that is altered by the miners to try different permutations to achieve the difficulty level required**

```
hash(block w/ nonce = 1)  => a80a8…8d732
hash(block w/ nonce = 2)  => f7bc9…5345f
hash(block w/ nonce = 3)  => ea758…ae629
hash(block w/ nonce = 13) => 0ebc5…37a66
```

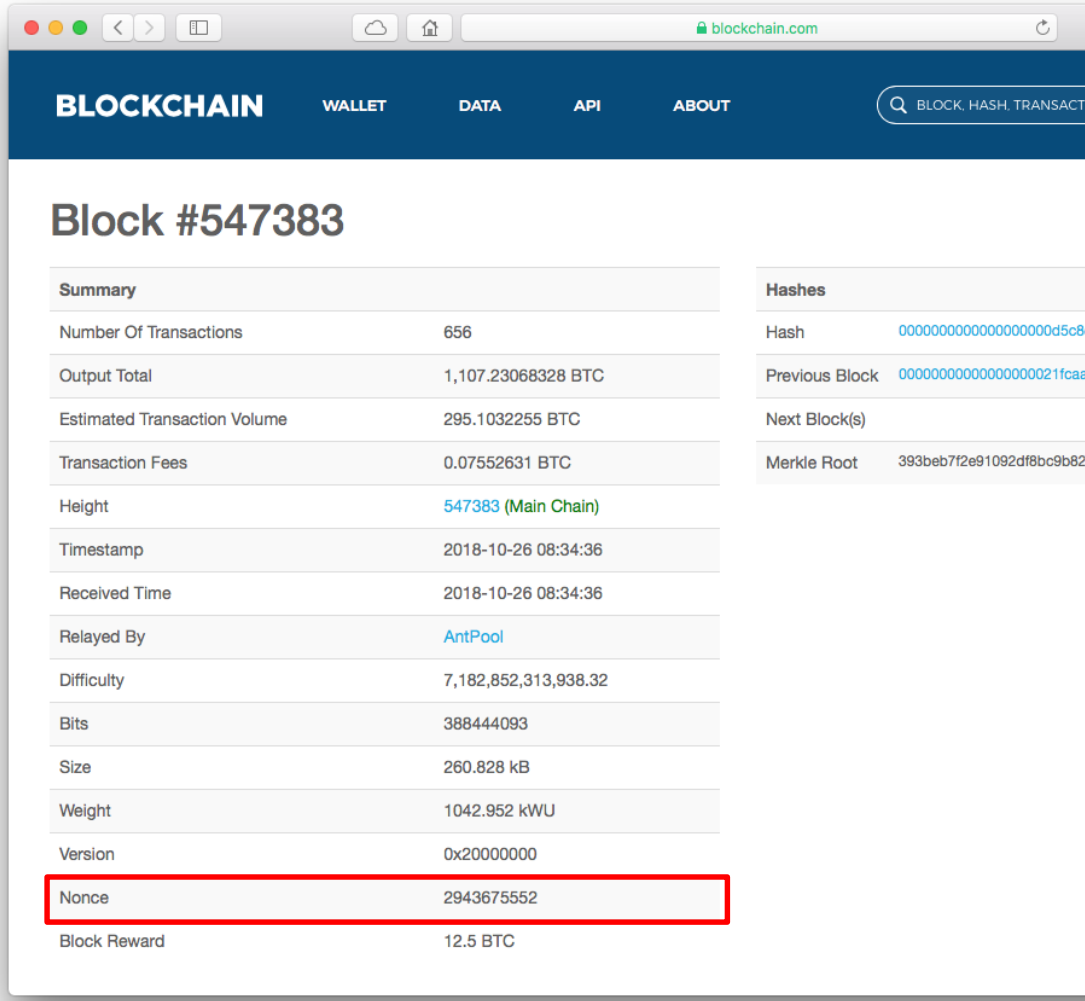

URL: b.socrative.com/student/

Room: HSR

Hash-based Proof of Work

- **To find a hash with N zeros (in hex) at input start, it requires $2^{4^N} = 16^N$ computations, which proves computational work performed**

```
in 3e-05 seconds, nonce = 0 yielded 0 zeros.          value = 4c8f1205f49e70248939df9c7b704…
in 0.000138 seconds, nonce = 12 yielded 1 zeros.      value = 05017256be77ad2985b36e75e486a…
in 0.000482 seconds, nonce = 112 yielded 2 zeros.     value = 00ae7e0956382f55567d0ed9311cf…
in 0.014505 seconds, nonce = 3728 yielded 3 zeros.    value = 000b5a6cfc0f076cd81ed3a606820…
in 0.595024 seconds, nonce = 181747 yielded 4 zeros.  value = 0000af058b74703b55e27437b89b1…
in 3.491151 seconds, nonce = 1037701 yielded 5 zeros. value = 00000e55bd0d2027f3024c378e0cc…
in 32.006105 seconds, nonce = 9913520 yielded 6 zeros. value = 000000 77a77854ee39dc0dc996dea…
in 590.89462 seconds, nonce = 186867248 yielded 7 zeros. value = 0000000225060b16117b23dbea9ce…
in 4686.171007 seconds, nonce = 1424462909 yielded 8 zeros. value = 00000000 2dd743724609a9f57260e…
```

Proof of Stake

- **Motivation: Mining requires a great deal of computing power to run different cryptographic calculations to unlock the computational challenges**
  - Bitcoin's current estimated annual electricity consumption* (TWh): 73.12
  - Country closest to Bitcoin in terms of electricity consumption: Austria
- **Proof-of-Stake algorithms achieve consensus by requiring users to stake an amount of their tokens so as to have a chance of being selected to validate blocks of transactions, and get rewarded for doing so.**
  - First cryptocurrency to implement PoS was Peercoin
  - Ouroboros was the first scientifically peer-reviewed PoS protocol
- **The more a user stakes, the better their chance of being selected since they'd have more skin in the game—acting maliciously would see them set back by a greater amount than someone who stakes less.**

Proof of Stake

- **Nodes wishing to participate in the validation process put money on *stake*, i.e., joining a set of validators ("lottery").**

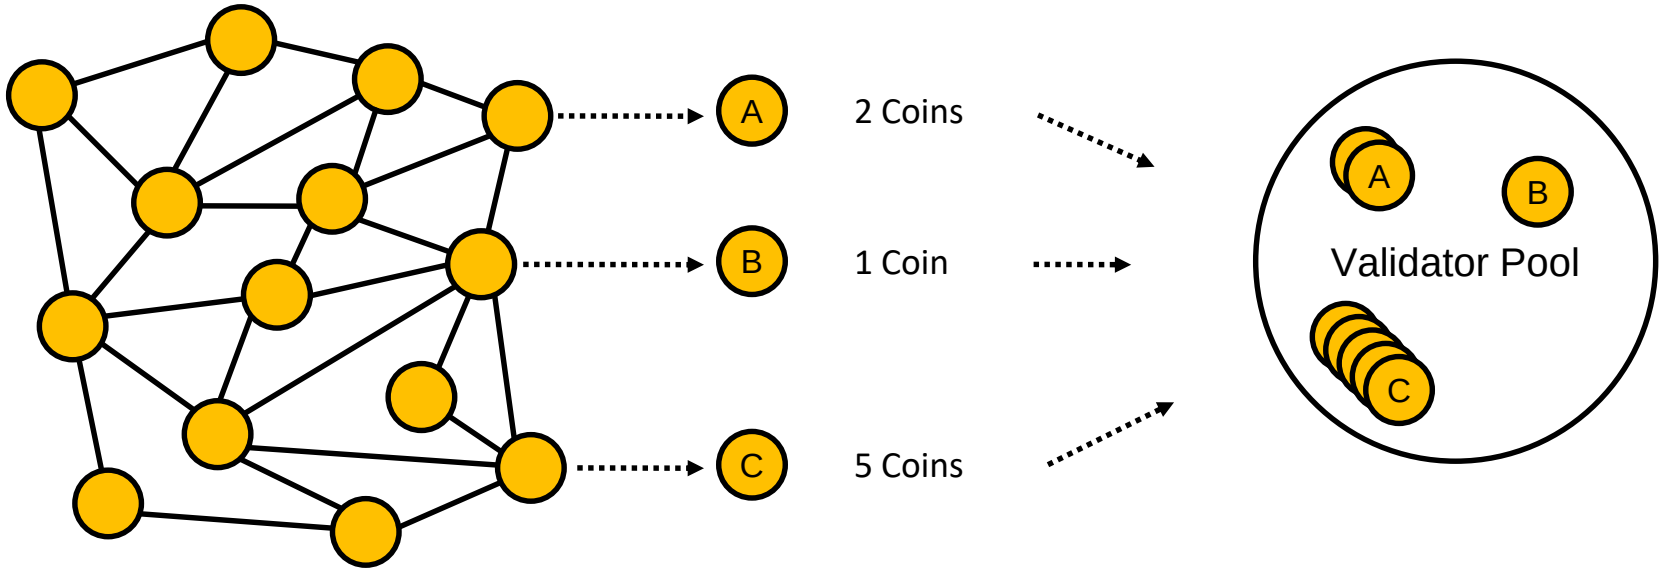

Proof of Stake

- **The more coins a node invests, the higher its stake; and the higher the chance of winning the "lottery".**
- **If the validator turns out to be adverse, and other validators notice it, his or her money on stake will be slashed, i.e., the adversary's money is gone**

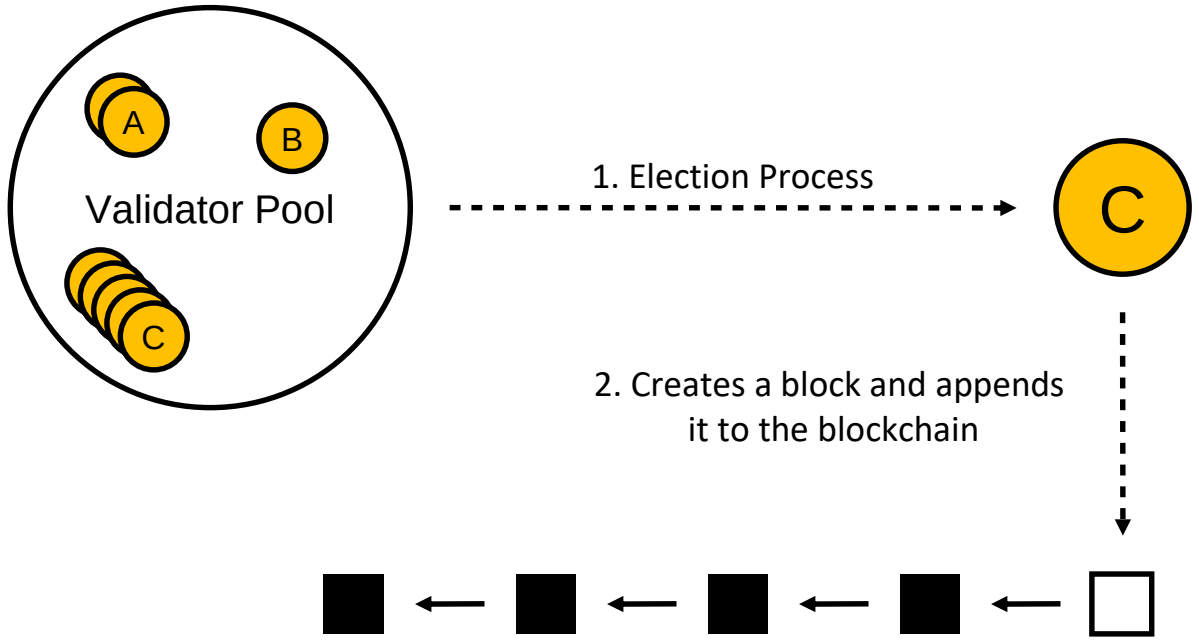

Note: In some blockchains, after node C got elected and appended a block, it must wait for a certain amount of time.

Proof of Stake in Bazo

- **Bazo uses a chain-based PoS approach, that is, the algorithm randomly selects a validator and assigns him or her the right to append a new block to the blockchain**
- **In order to become a validator, a user generates an RSA keypair $(pkco_{mm}, skcomm)$ and publishes a stake transaction (StakeTx) containing the public key pk_{comm} to the network**
- **A transaction of this type is accepted by the network if the issuer fulfills the minimum staking amount that is needed**
- **After the minimum waiting time elapsed, each validator participates in the election process in a non-interactive way**

Proof of Stake in Bazo

Public
PubKey: 4dc4…b904
RSA PubKey: 3ae7…32cd

A

Private
PrivKey: 8063…aeac
RSA PrivKey: bd75…3428

Publish Stake Transaction to the network

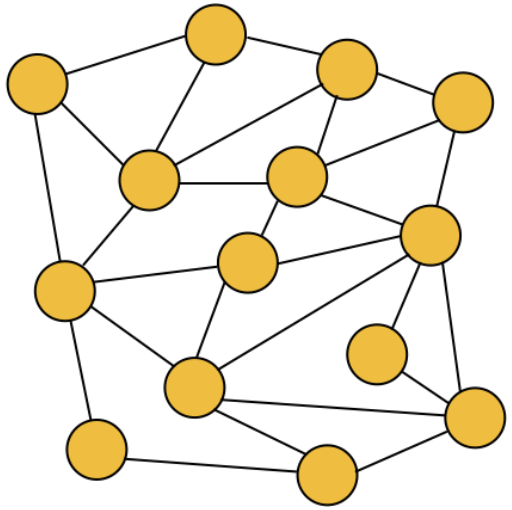

StakeTx #4'301	
Fee	2
Is Staking	True
Account	4dc4…b904
Timestamp	2018-04-27 23:11:54
Key Commitment	3ae7…32cd
Signature	ffa1…43e1

Note: This has to be done in advance. After a StakeTx is published, the user has to wait for X blocks until he or she can participate in the validation process.

Proof of Stake in Bazo

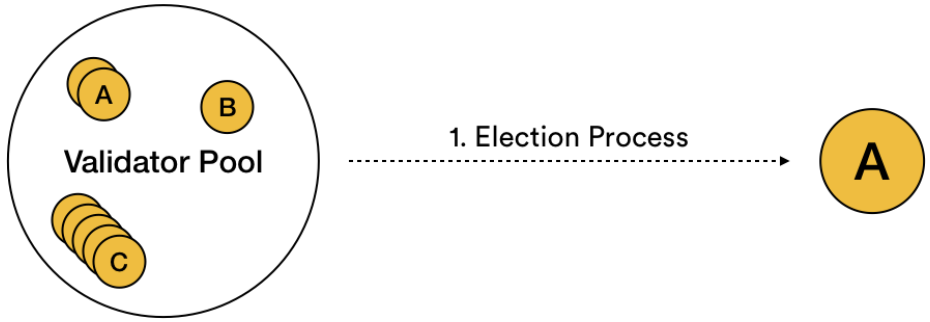

4.3 PoS Condition

The introduced PoS condition in Section 2 is updated to

$$\frac{\text{SHA-256}([Proof_{PrevBlocks}] \cdot Proof_{Local} \cdot BlockHeight \cdot TimeInSeconds)}{Coins} \leq Target. \quad (15)$$

The following parameters changed as opposed to Equation 1:

List of the Previous Proofs ($Proof_{PrevBlocks}$) By including a list of the previous proofs a stake grinding attack becomes infeasible.

Local Proof ($Proof_{Local}$) All parameters in Equation 15 are the same for each validator in the network except the local proof. The local proof individualises the PoS condition to each validator and therefore, a validator with a low amount of coins also has the possibility to append a block to the blockchain.

Note: Contrary to PoW, a validator is limited to exactly 1 H/s (hash per second) by providing a valid TimeInSeconds to fulfill the PoS condition

Proof of Stake in Bazo

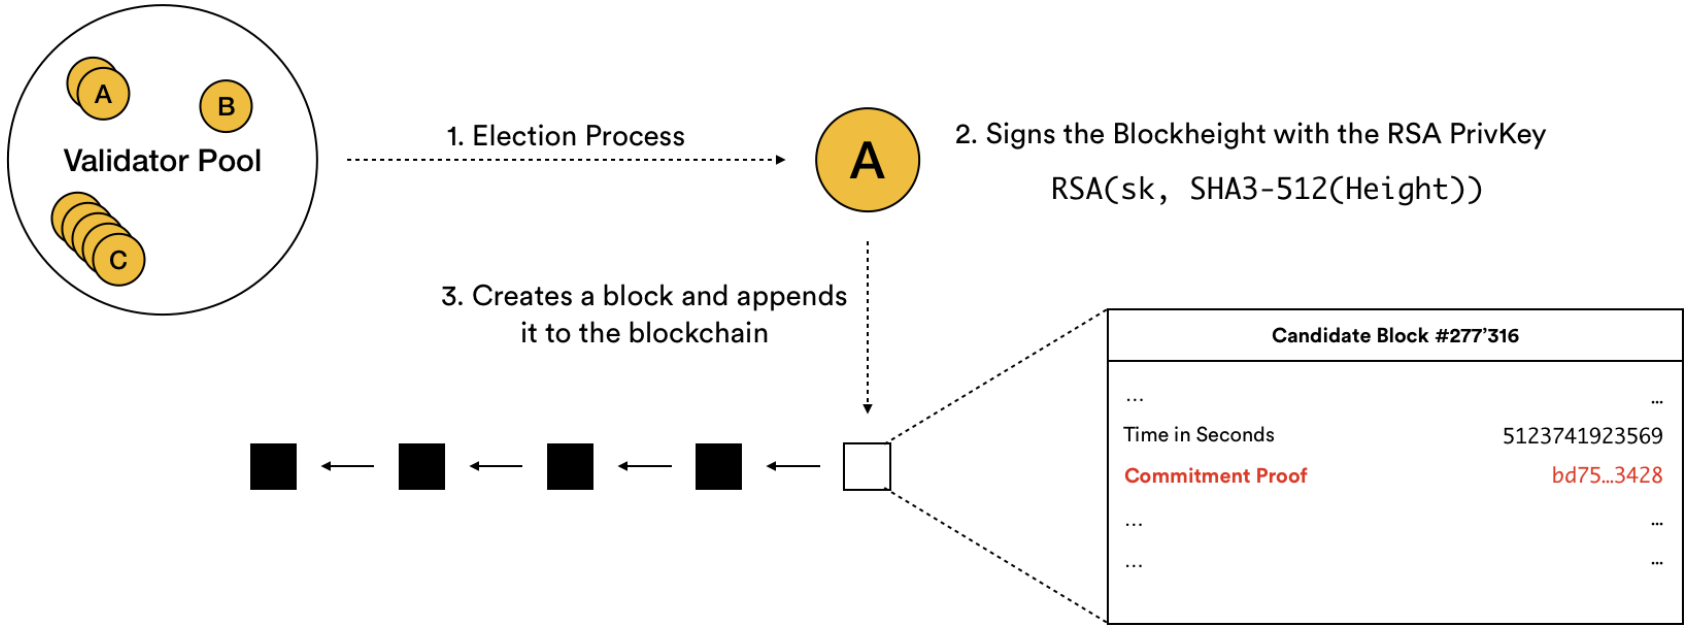

Proof of Stake in Bazo

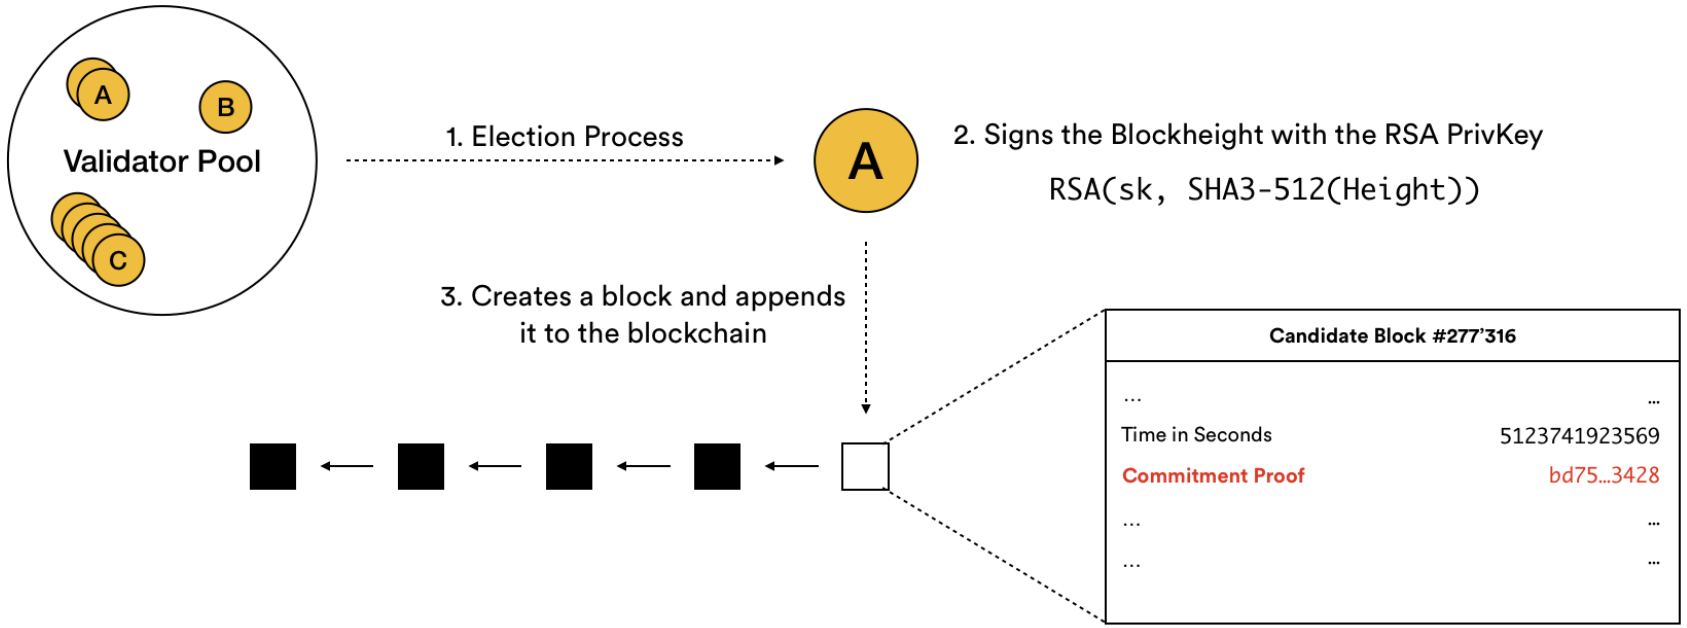

Note: A validator who appended a block to the blockchain has to wait for X blocks until he or she can take part in the election process again.

URL: b.socrative.com/student/

Room: HSR

What is Scalability?

- **“A system is said to be *scalable* if it can handle the addition of users and resources without suffering a noticeable loss of performance or increase in administrative complexity.”**
- **If a system is expected to grow, its ability to scale must be considered when the system is designed (naming, authentication, authorization, accounting, communication, use of remote resources)**

Two ways to scale computer systems:

- **Scale up (vertical): add more resources to a single computation unit, i.e., better CPU, more RAM**
- **Scale out (horizontal): add additional computation units to split workload across units**

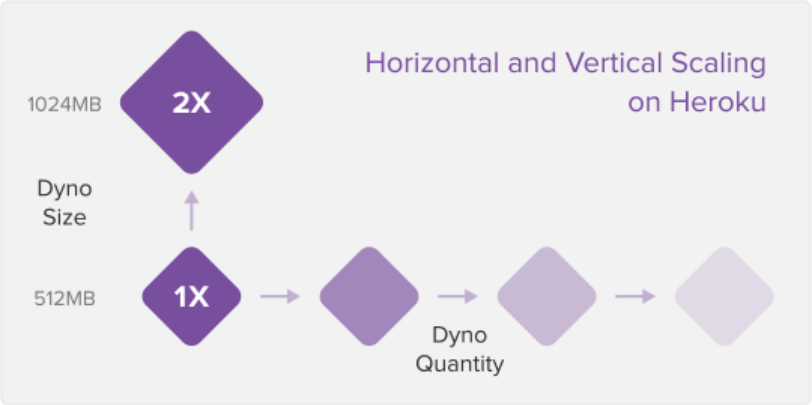

Motivation and Problems

- **Situation: Bitcoin can process roughly 8 transactions per second (tps), Ethereum around 15 tps**
  In comparison, Visa could process up to 65k tps
- **More nodes = more power. So more speed, right? Not exactly.**
- **Every single node must process every single transaction.**
  - Higher security and decentralization, less scalability
- **Divide transactions to different groups of nodes**
  - Smaller group of nodes are more vulnerable
  - Increases scalability, decreases security and decentralization

The (Blockchain) Trilemma

- **A trilemma between decentralization, scalability and security.**
- **Examples of distributed systems:**
  - Bitcoin: decentralized, secure, but not scalable
  - Google: scalable, secure, but not decentralized
  - BitTorrent: decentralized, scalable, but not secure

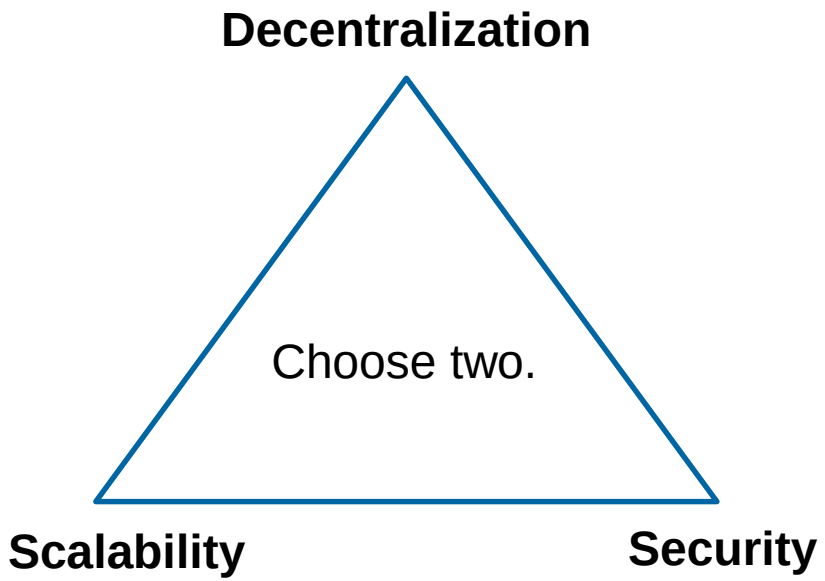

The (Blockchain) Trilemma

- **A trilemma between decentralization, scalability and security.**
- **Examples of distributed systems:**
  - Bitcoin: decentralized, secure, but not scalable
  - Google: scalable, secure, but not decentralized
  - BitTorrent: decentralized, scalable, but not secure
- **Possible solution: Increasing the block size**
  - More transactions fit into the block
    => increased transactions per second
  - But: Nodes need more resources
    => increased centralization
  - Leads to a hard fork, e.g., Bitcoin and Bitcoin Cash

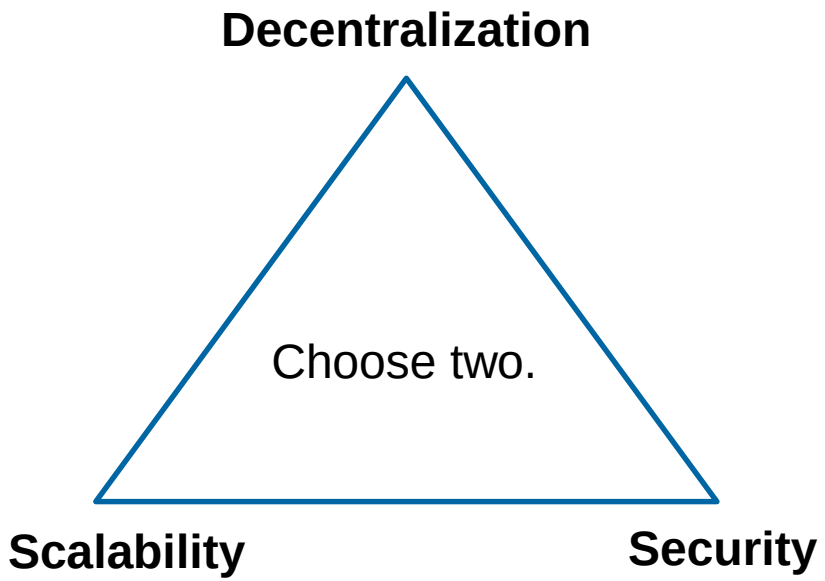

Blockchain Scalability Solutions

- **On-chain**
  - Segwit (efficient use of storage)
  - Sharding
    - Divides the network state into partitions, where each node is responsible for a single partition
- **Side-chain**
  - Observable relationships between chains
- **Off-chain**
  - Payment channels
    - 2-parties sign the latest state of the balances in the channel and can provide proofs and signatures to the smart contract at the end of the process.
  - State channels
    - Same as payment channels but not only for balances
  - Plasma
    - Same as payment channels but for multi-party interaction

https://ethereum.stackexchange.com/questions/71664/offchain-vs-sidechain-vs-statechannel

Database Sharding

- **Compute a hash of the shard key field's value**
- **Each chunk is assigned a range based on the hashed shard key values**
- **Each chunk is held on a separate database server instance**

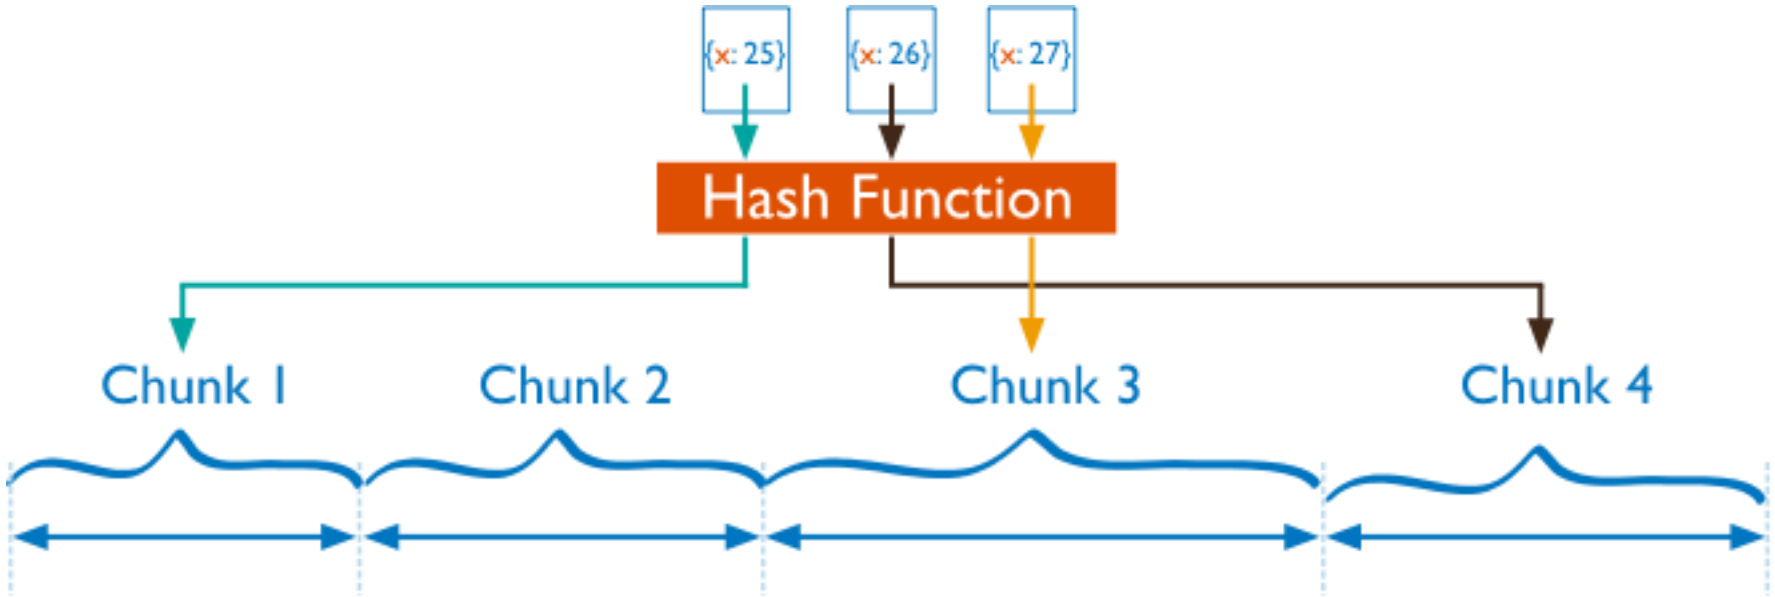

Blockchain Sharding

- **The transaction load on the network is divided into different shards**
- **Only designated nodes validate the transaction, and not the entire network**
- **The blockchain should scale linearly with every new node added to the network.**
- **Higher scalability without sacrificing security or decentralization**

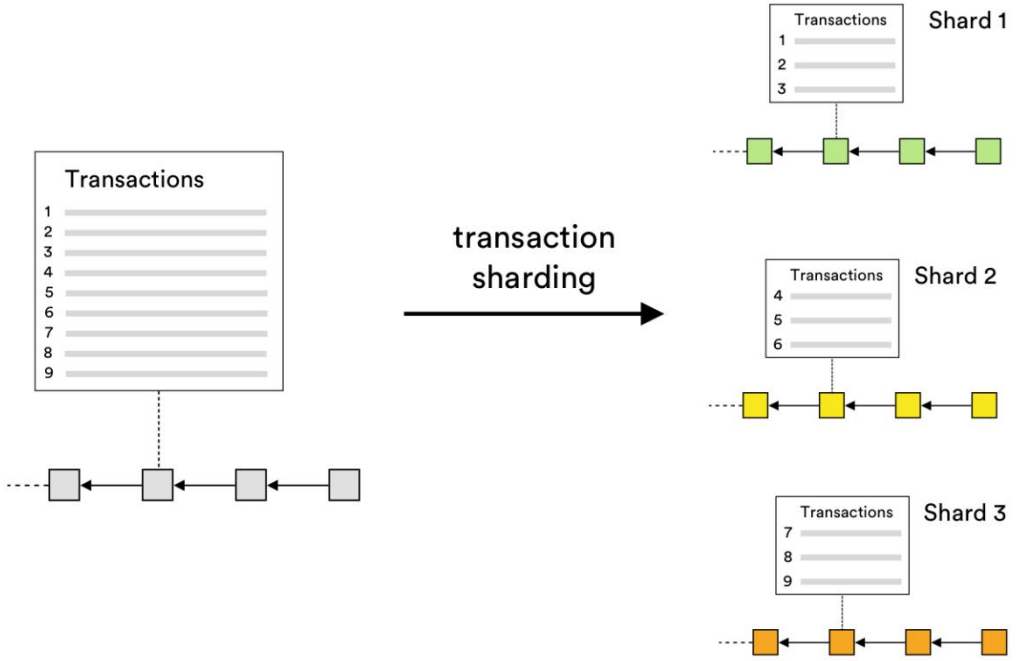

Single-Shard Take-Over Attack

- **Problem: Traditional sharding protocols do not take into account a Byzantine setting**
  Nodes must be resilient against malicious attackers and independent nodes have to reach consensus over the current state of the network.

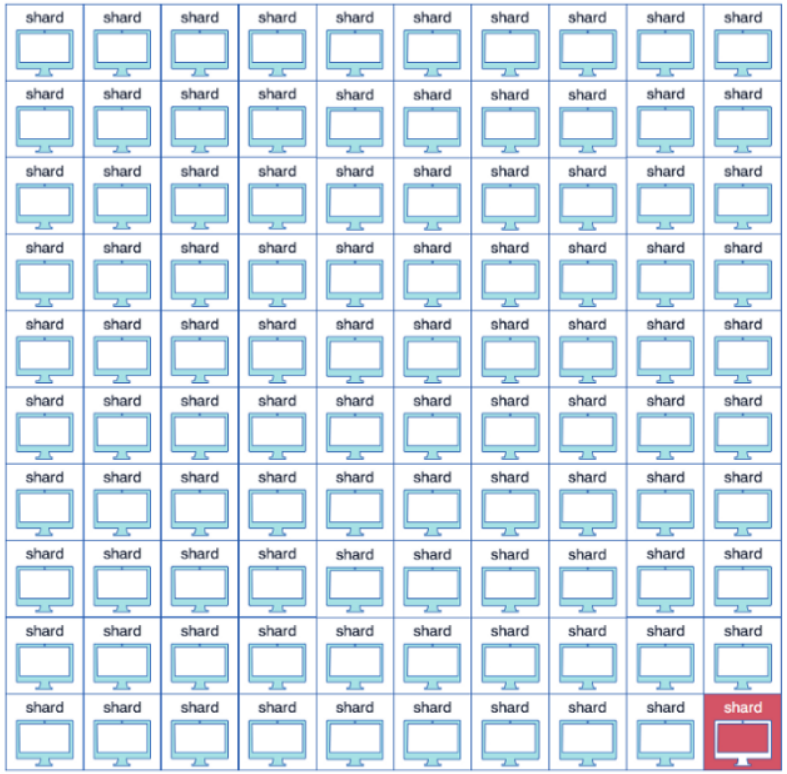

1% Attack

> In 100 shards system, it takes only 1% of network hash rate to dominate the shard.

What is Bazo

- **Bazo is a permissionless blockchain for <u>research</u>**
  - Its **not** an ICO, not looking for funding, it will never be a productive blockchain
  - Looking for great ideas!
- **Test new ideas, new mechanisms, new algorithms**
  - Blockchain from scratch, started with simple and solid, well-tested mechanism
  - Gradually brought in new ideas and concepts
- **Not a single unified blockchain**
  - Fork with self contained proofs, sharding, pruning, smart contracts
- **https://github.com/bazo-blockchain**

- Self-Contained Proofs
  - Miners can verify the validity of transaction without the blockchain history
- Superlight
  - Reduce data on-chain

Recall: Merkle Trees

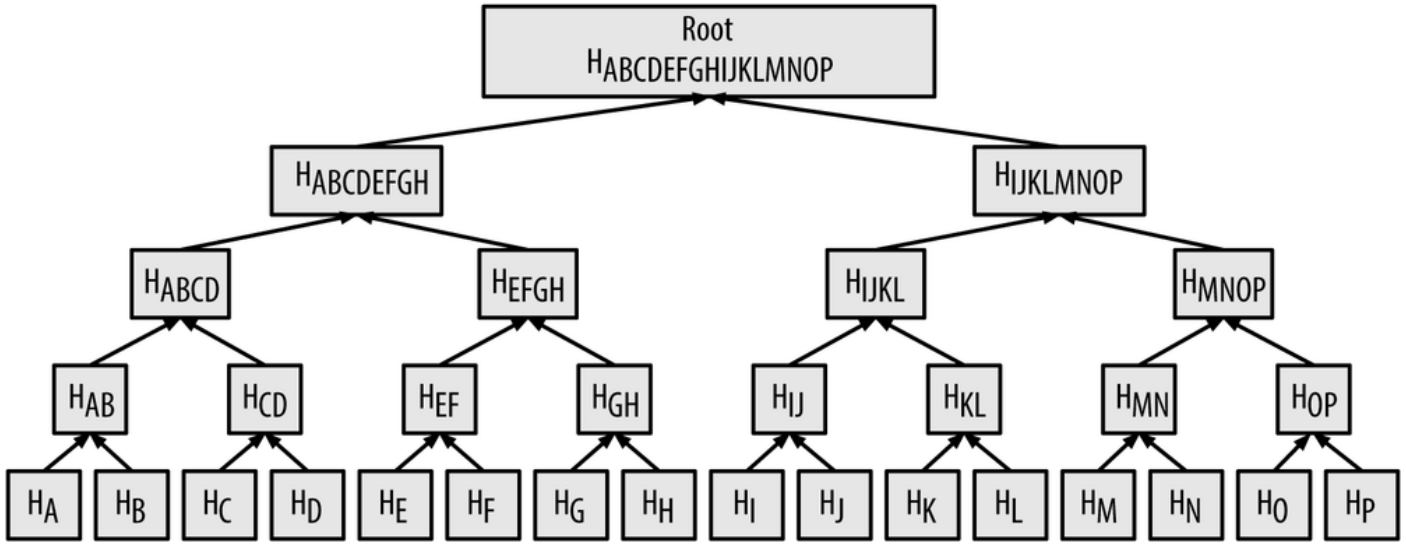

- **A Merkle tree is a *binary hash tree* containing N leaf nodes**
- **Constructed bottom-up, i.e., $H_{AB} = Hash(HA + HB)$**
- **Used to summarize all transactions in a block**
- **To prove that a specific transaction is included in a block, a node only needs to produce $log_2(N)$ hashes, constituting a merkle path connecting the specific transaction to the root of the tree.**

Recall: Merkle Proofs

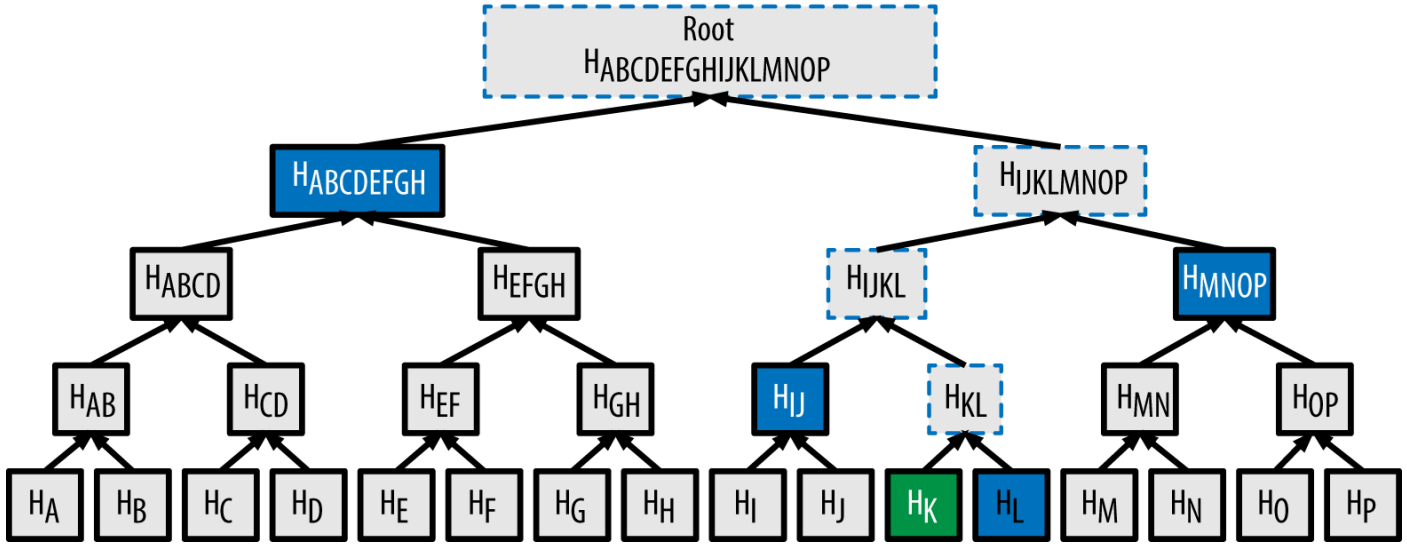

- **In above example, a node can prove that transaction K is included in the block by producing a merkle path that is only $log_2(16) = 4\ hashes$ long, i.e., $Proof = \{HL, HIJ, HMNOP, HABCD_{EFGH}\}$**

Self-Contained Proofs (SCPs)

- **Traditionally, miners require the full blockchain to validate a transaction of a user**
  e.g. does the user have sufficient funds to execute the transaction
- **Instead of validators requiring the full blockchain to validate a transaction, a user has to provide all required proofs in the transaction in order for validators to verify a transaction, independent of the blockchain.**
- **A SCP is a list of Merkle proofs. A user has to provide one Merkle proof for each block that contains one or more transactions where he or she sent or received funds**

Example of a SCP

- **Assume a user *A* with past transactions *F3* and *F6***
- **A new transaction of user *A* has to include two Merkle proofs, i.e.,**

$$SCP = \{Proof_1, Proof_2\}$$

where

$$Proof_1 = \{F5, F6, B2\}$$
$$Proof_2 = \{F3, F4, A1\}$$

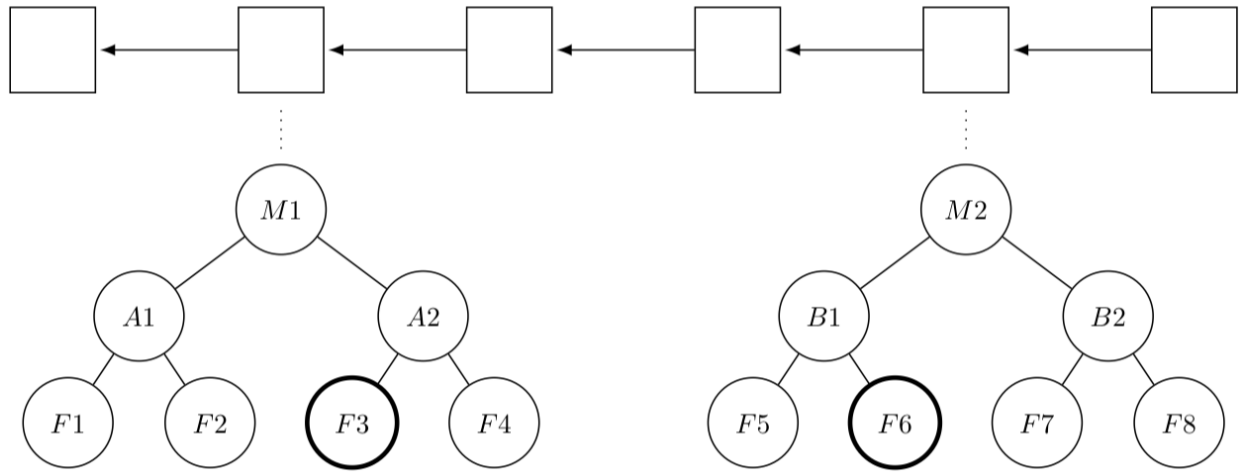

Superlight

- **Header**
- **New addresses found in transactions for the current block header**
  - All block headers = knows all addresses in the blockchain
- **Perfect Bloom filter of sender and receiver address of involved transactions in that block header**
- **BLS signatures from all senders and receivers of involved transactions in that block header to verify that a sender or receiver was involved**
  - Signature aggregation
- **Merkle root or transaction hashes.**

- **Advantages**
- **Only Block headers need to be stored**
  - Reduce storage size
  - BLS signature – constant size
  - Merkle root – constant size / hashes can grow
  - Perfect bloom filter / address involved in TX – can grow
  - New addresses for perfect bloom filter – can grow
- **Hybrid approach**
  - Mark global state, store in block
    - E.g. introduce “global” keyword in Lazo

Superlight

```
version 1.0

contract SimpleContract{
  [Local]
  Map<address, int> payments

  [Global]
  int totalAmount

  [Payable]
  [Pre: msg.coins > 0]
  function void pay() {
    payments[msg.sender] += msg.coins
  }
}
```

Drawbacks

- Interactive protocol
  - Miner proposes block, needs signature from receiver and receiver (only Merkle root)
  - Receiver must acknowledge a TX, otherwise no funds received, both signatures required
    - Receiver must be online
- Losing one proof = losing access to funds
- Computational complexity of BLS

Questions?

Dr. Thomas Bocek thomas.bocek@hsr.ch

References

- Scale in Distributed Systems
  B. Clifford Neuman, 1994
- Notes on Scalable Blockchain Protocols (version 0.3.2)
  Vitalik Buterin, 2015
- Introduction to Parallel Computing 2nd Edition
  A. *Grama, A. Gupta, G. Karypis, V. Kumar, 2003*
- The Byzantine Generals Problem
  L. Lamport, R. Shostak, M. Pease, 1982
- Cryptographic Sortition for Proof of Stake in Bazo
  R. Blum, 2018
- Scalability for the Bazo Blockchain with Sharding
  R. Blum, 2018
- Sharding in MongoDB, Link:
  https://docs.mongodb.com/manual/sharding/
  MongoDB Documentation, visited 14th Aug 2018
- Mastering Bitcoin
  Andreas M. Antonopoulos, 2014
- Slides partially by R.Blum for HS18