Red Hat

CAN'T WE JUST AGREE?

ONDŘEJ CHALOUPKA, VOJTĚCH JURÁNEK

DevConf.CZ 2019

WHAT IS THIS TALK ABOUT

WHY CONSENSUS ALGORITHM IS AN ESSENTIAL PART
IN DISTRIBUTED LEDGER SYSTEMS

A BRIEF OVERVIEW OF CONSENSUS ALGORITHMS
IN DISTRIBUTED LEDGER SYSTEMS

DIFFERENCES BETWEEN CONSENSUS ALGORITHMS
AND HOW THE CHOICE IMPACTS CAPABILITIES OF DISTRIBUTED SYSTEM

TRANSFER OF THE MONEY

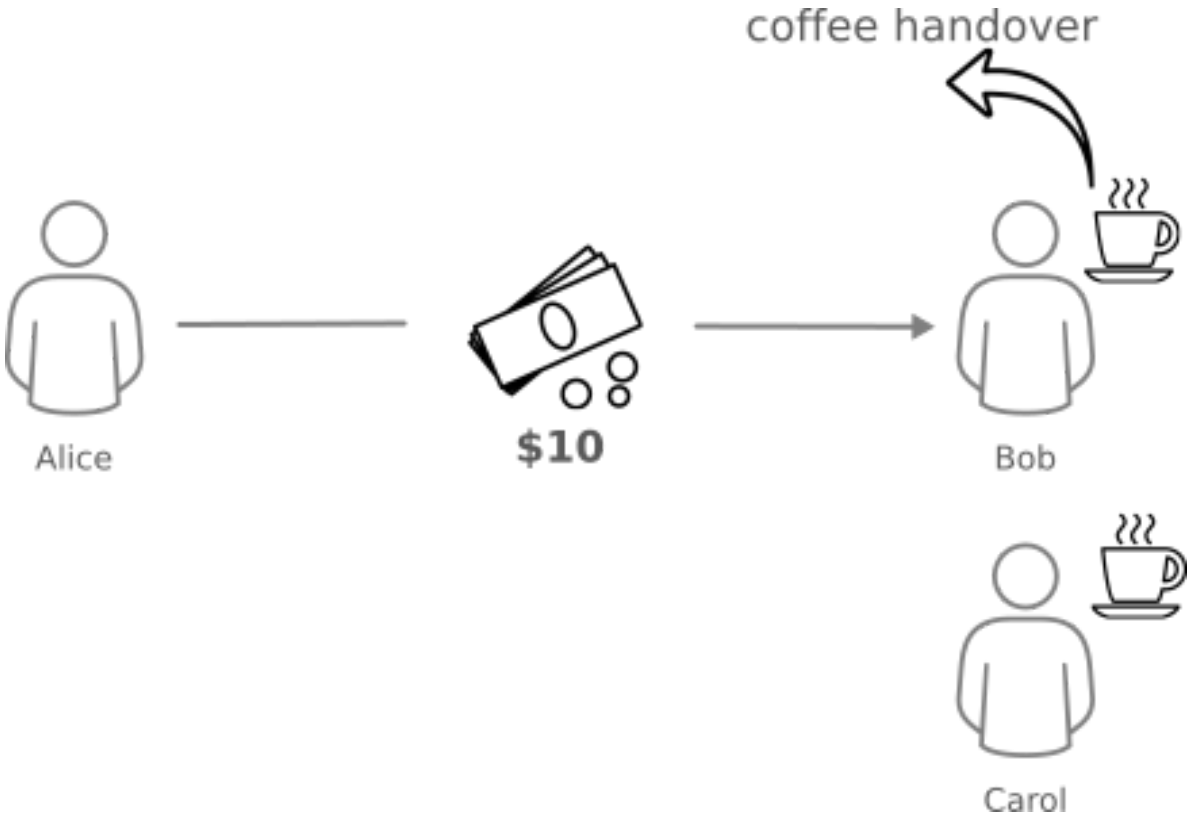

Beware: distributed ledgers are not limited only to cryptocurrencies, there are lots of other applications!

CENTRAL AUTHORITY

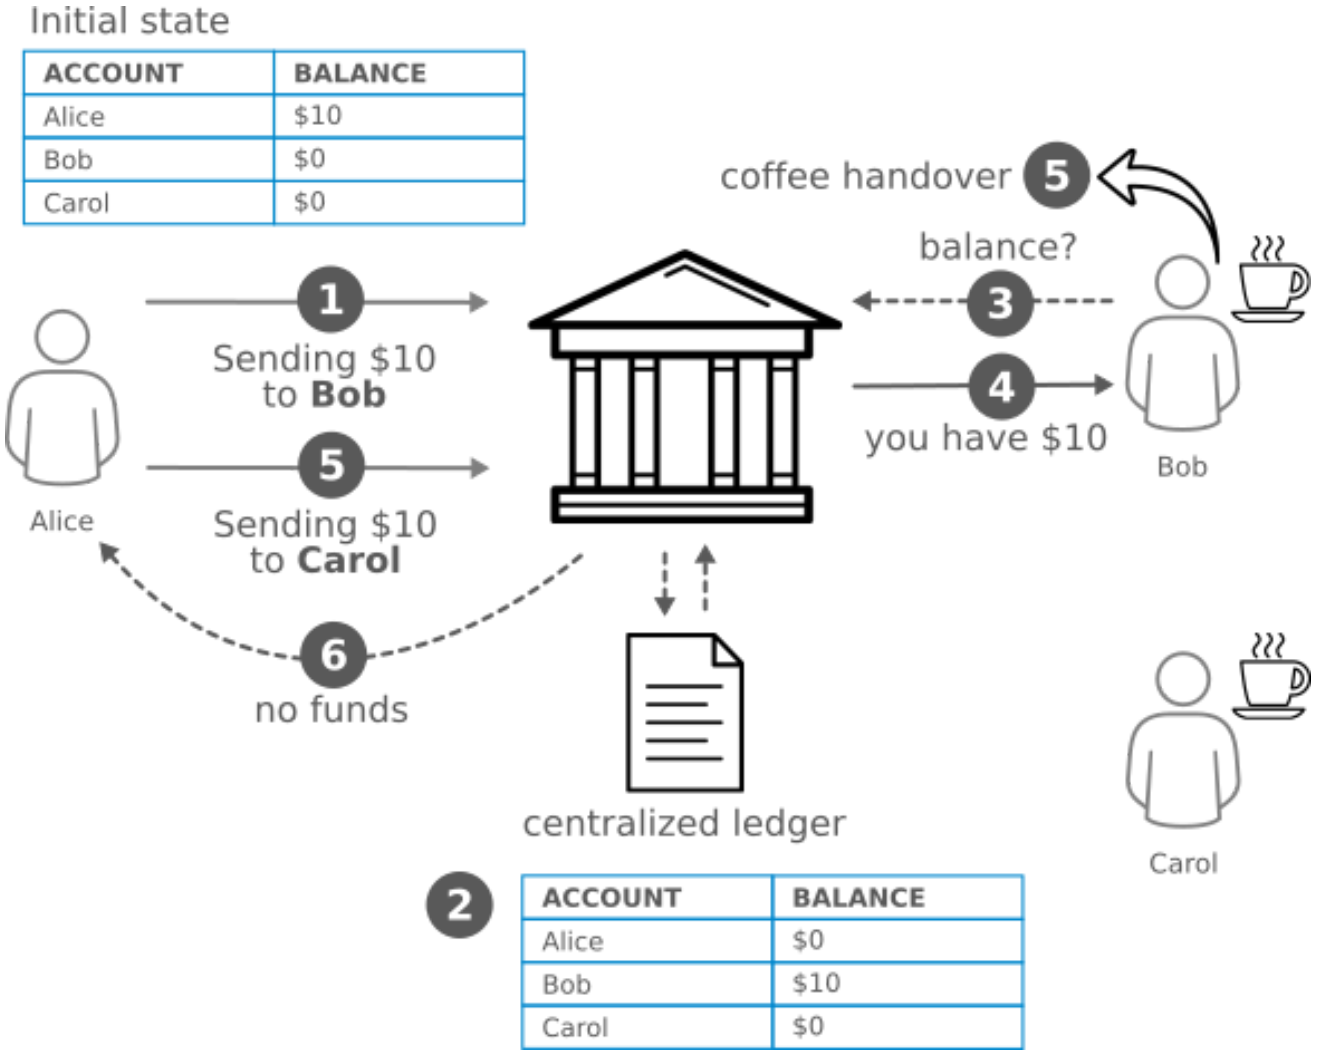

DOUBLE SPEND PROBLEM

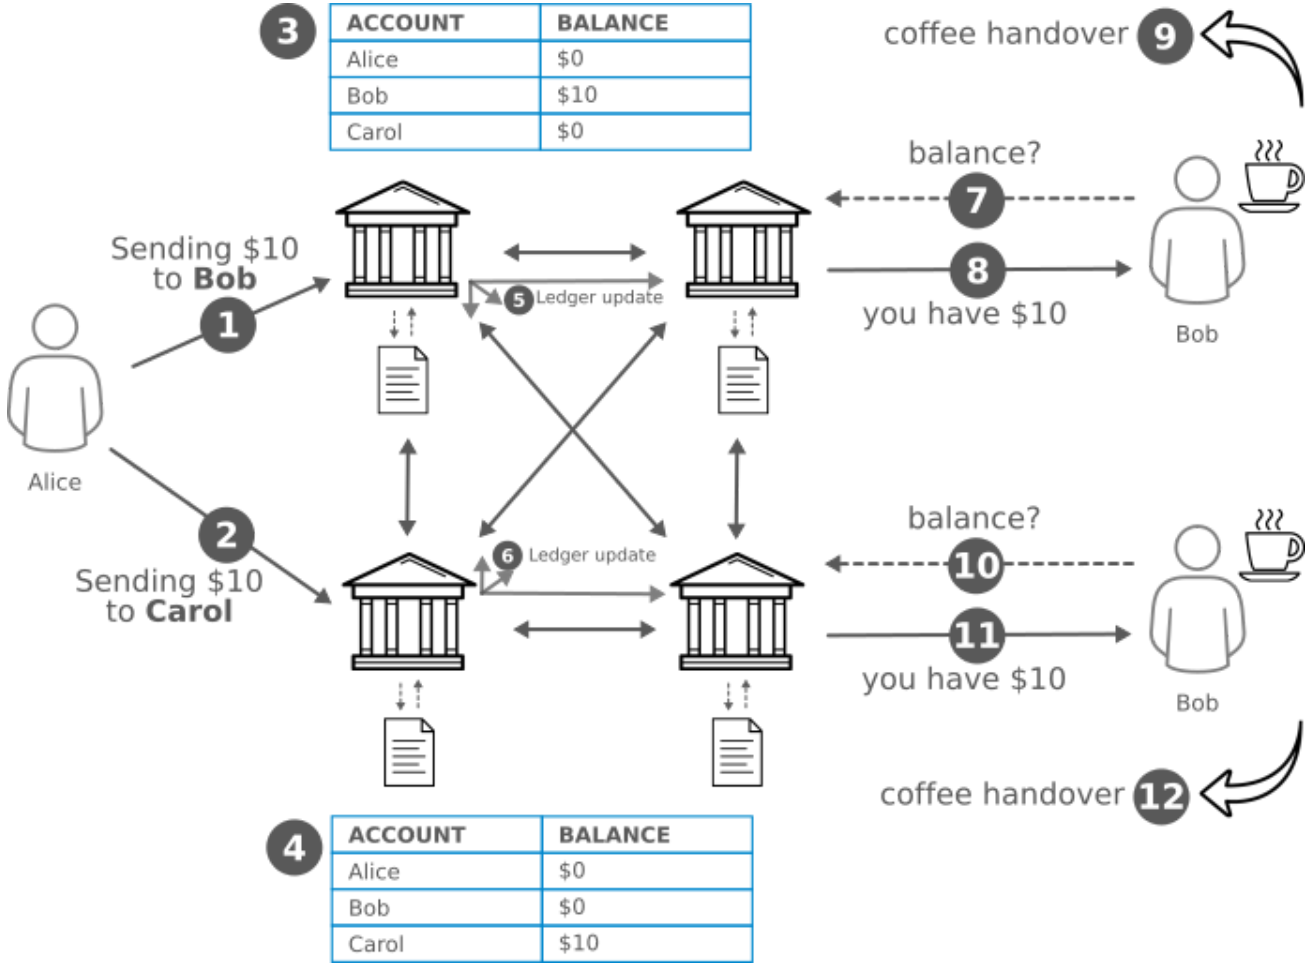

In distributed systems, we need agreement between participants, which transaction is valid.

CONSENSUS

- Agreement between nodes on something, e.g. some value
- Another example: whether to commit a (distributed) transaction to a database
- Hard, network is unreliable. T there can be delays or failures in communication
- Consensus has to have two properties: **safety** and **liveness**

CONSENSUS

Emin Gün Sirer 

@el33th4xor

Follow 

Ok, there is a terribly wrong framework emerging around consensus protocols. People think that PoW and PoS are consensus protocols, and that they are the only two consensus protocols out there.

This is false. Let me explain.

9:09 AM - 13 Jun 2018

https://twitter.com/el33th4xor/status/1006931658338177024

NON-BYZANTINE

FAULT TOLERANT CONSENSUS

- can withstand failures but not a cheating participant
- trust in all involved parties
- protocols like Paxos (Google Spanner), Raft (etcd), Zab (Zookeeper)
  - leader-based vs. leader-less

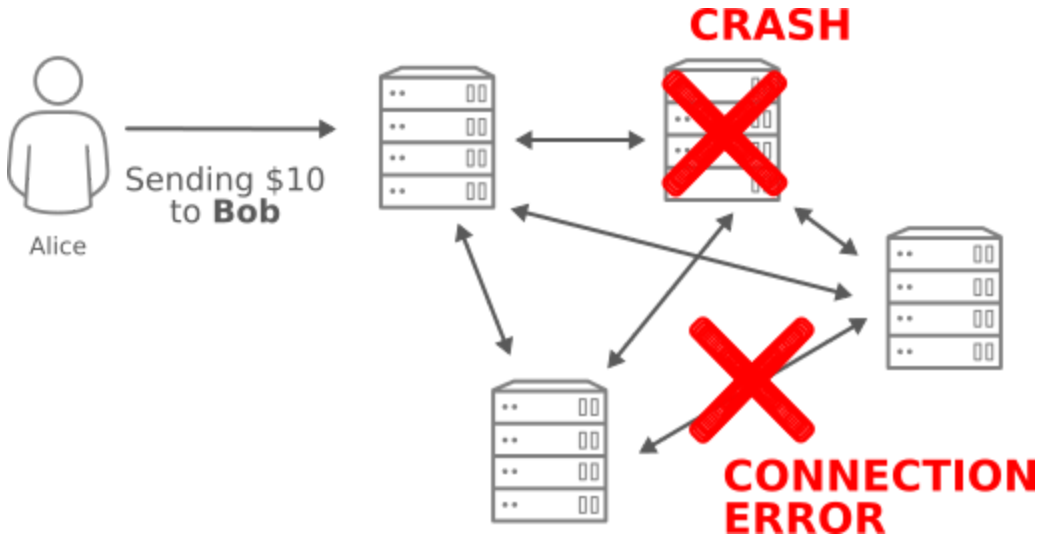

RAFT ALGORITHM

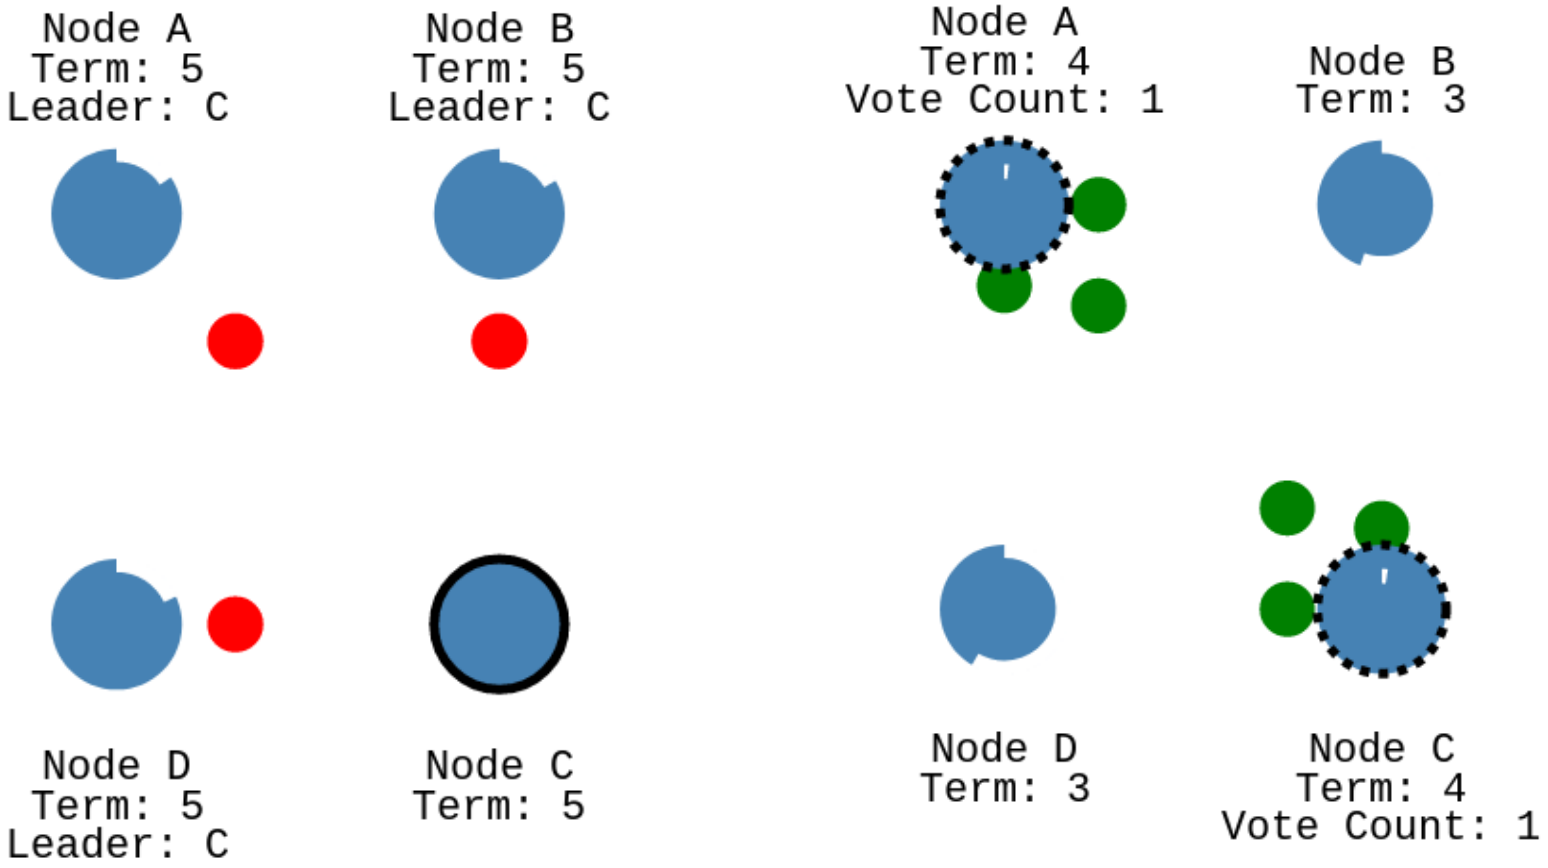

Leader based

Leader election

http://thesecretlivesofdata.com/raft/

RAFT ALGORITHM

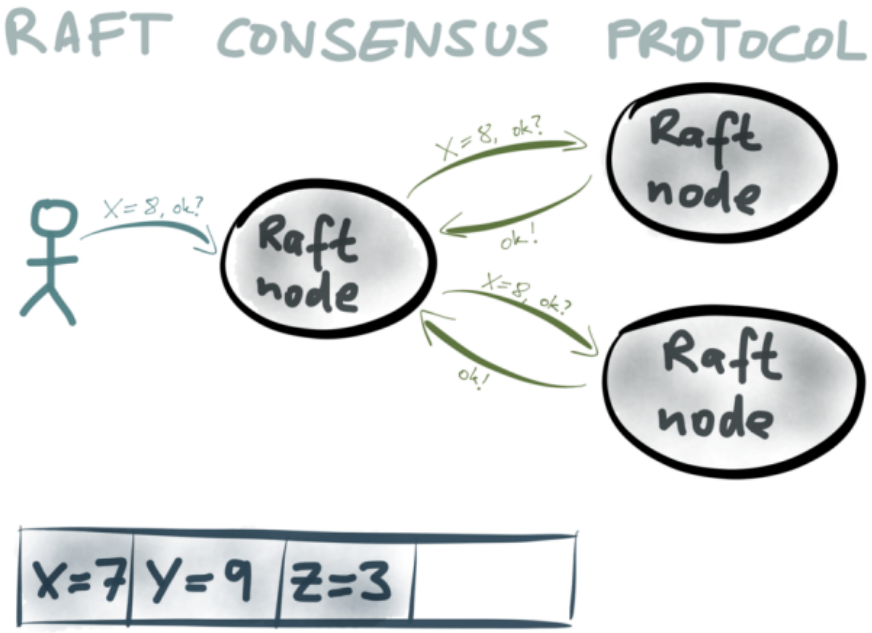

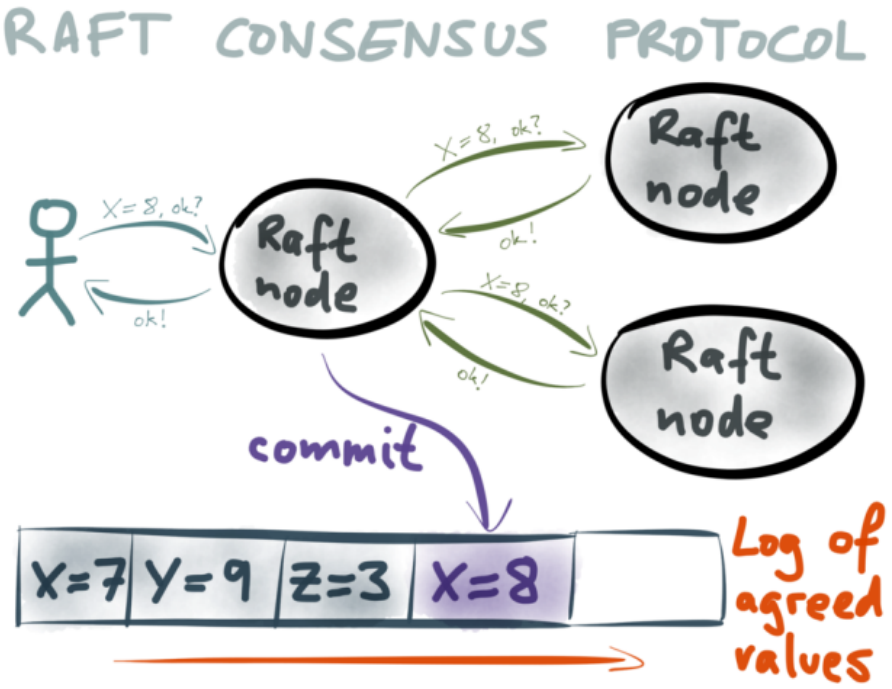

https://www.confluent.io/blog/using-logs-to-build-a-solid-data-infrastructure-...

PRIVATE BLOCKCHAINS

- all nodes are under control of a **single organization**
- number of participant is small, they know about each other and **trust each** other
  - Raft in Hyperledger (Sawtooth, on roadmap for Fabric)
- **Permissioned** : nodes needs some permission to join the network
- **Permissionless** : nodes can join the network without any permission (public blockchains)

BYZANTINE FAILURE

- Besides delays and failures in communication over network, the situation can be even worse - there can be malicious participants!
- A Byzantine failure is any fault presenting different symptoms to different observers.
- E.g. attempt to double spend money.

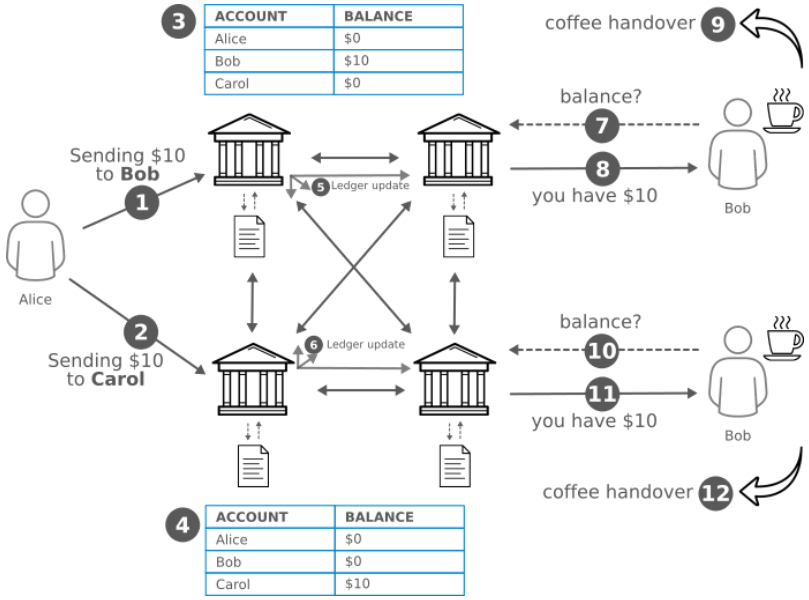

BFT ALGORITHMS

BYZANTINE FAULT TOLERANT

- PBFT, Tendermint, Stellar ...
- Agreement can be imaged like a three-phase commit (propose a value, pre-prepare commit, prepare, commit)..

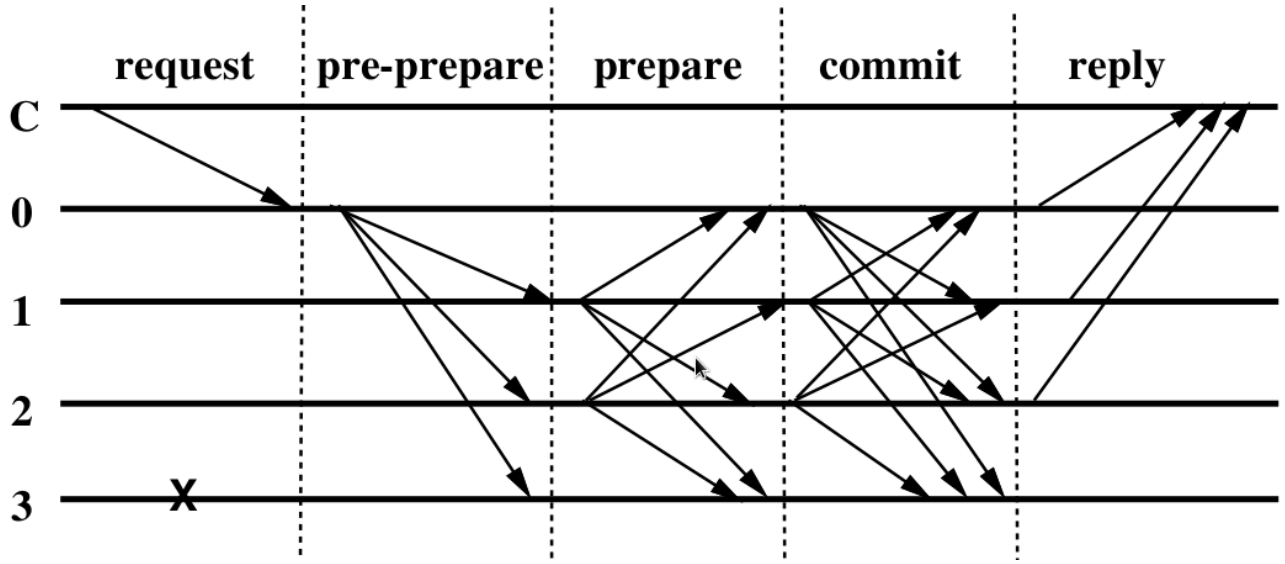

http://pmg.csail.mit.edu/papers/osdi99.pdf

FEDERATED DISTRIBUTED LEDGERS

- Still to some extend centralized.
- Usually to some extend permissioned or have to be combined e.g. with PoS (Tendermint).

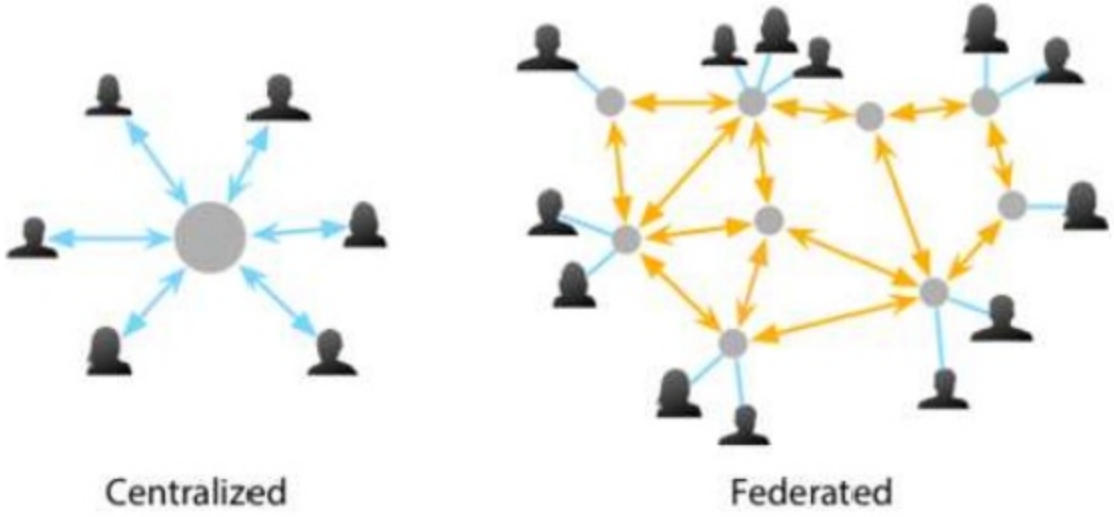

https://mc.ai/whats-new-in-deep-learning-research-understanding-federated-learning/

SIBYLS ATTACK

- Forging the identity to subverted the result

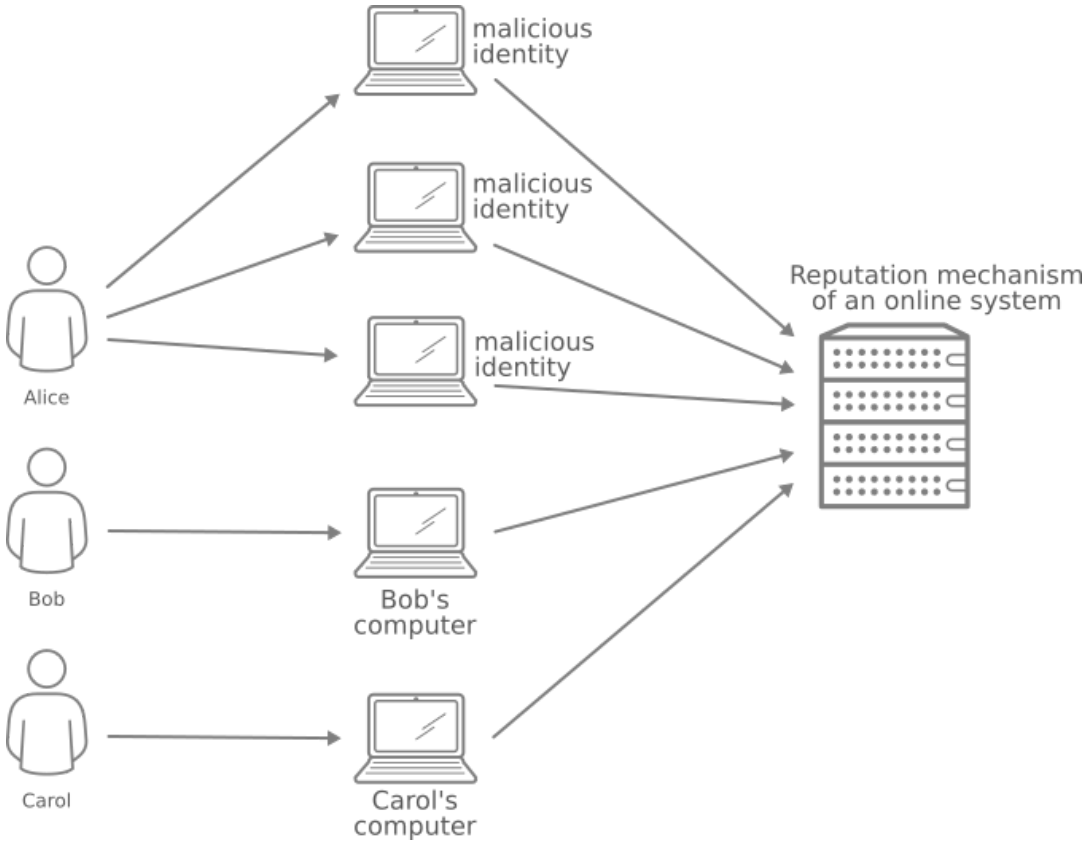

PROOF-OF-WORK

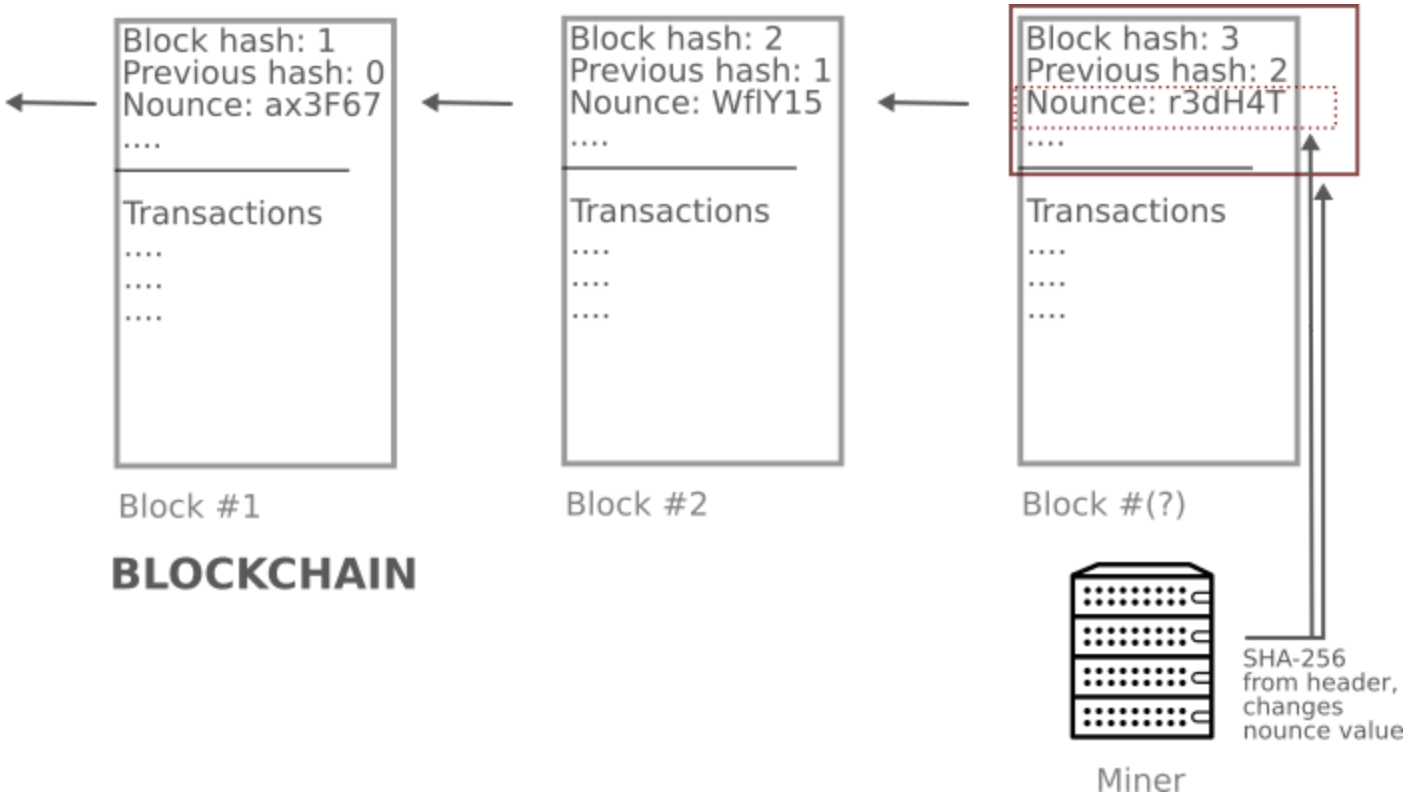

BLOCKCHAIN FORK

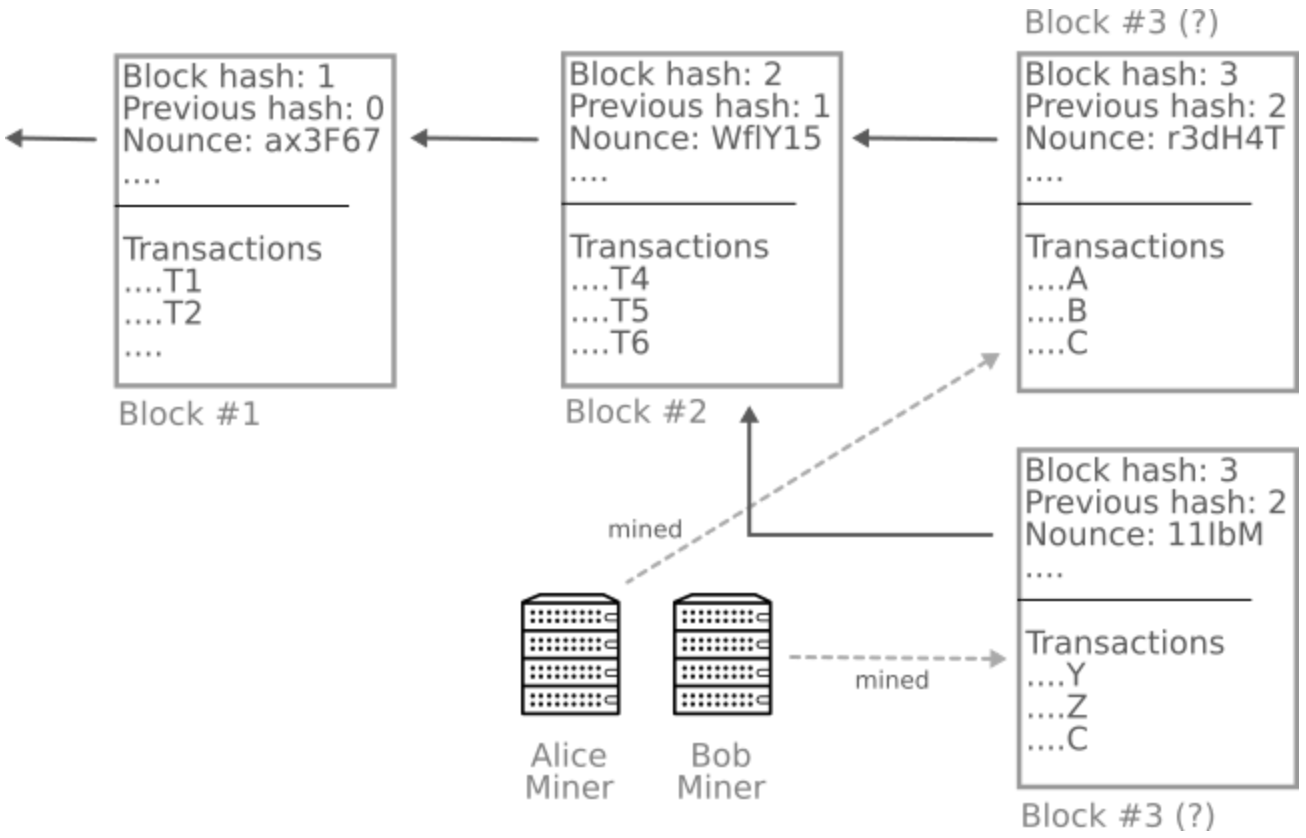

NAKAMOTO CONSENSUS

- Bitcoin is (almost) **Byzantine fault tolerant (BFT)** and also resistant to **Sybil** attack
- Proof-of-Work (PoW) is mechanism how to prevents Sybil attacks
- The truth (agreement) is determined by the longest chain (created by Proof-of-Work) usually called **Nakamoto consensus algorithm**
- **Probabilistic**: probability of consensus is less than 1
- Proof-of-Work finding a new block is stable to 10 minutes
  - Bitcoin network is essentially **synchronous**

WHAT ABOUT POS?

- Proof-of-Stake (PoS) is again **not a consensus** algorithm
- Proof-of-Stake **solves Sybil** attack
- **Chain-based** vs. **PBFT-style**
  - **Hybrid PoS**: Proof-of-Work with Proof-of-Stake on top

Consensus algorithms on top of DAG

Blockchain

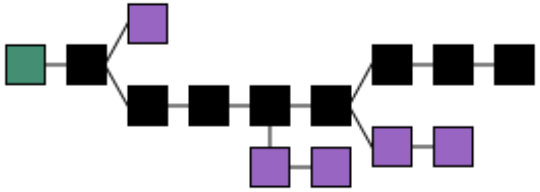

https://en.wikipedia.org/wiki/Blockchain#/media/File:Blockchain.svg

DAG

- Hashgraph, Avalanche, Tangle...
- Usually are not resistant to Sibyl attack (needs PoS or something else)

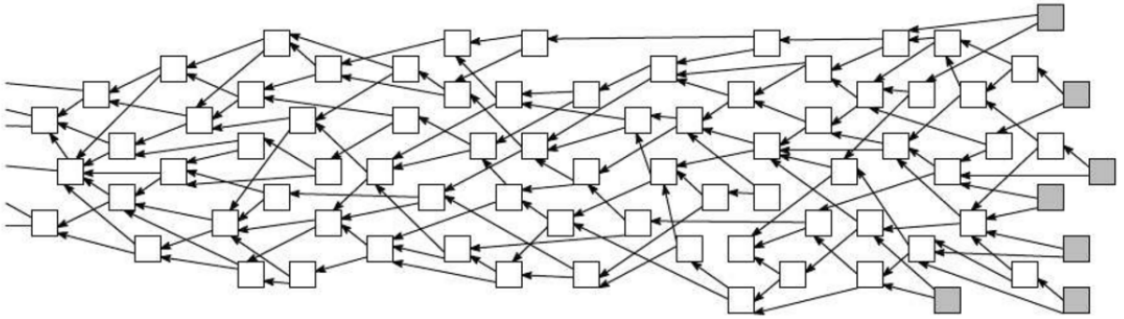

https://ministryofblockchain.io/is-directed-acyclic-graph-dag-blockchains-new-competitor/

Hashgraph

- Gossip about gossip
- Virtual voting

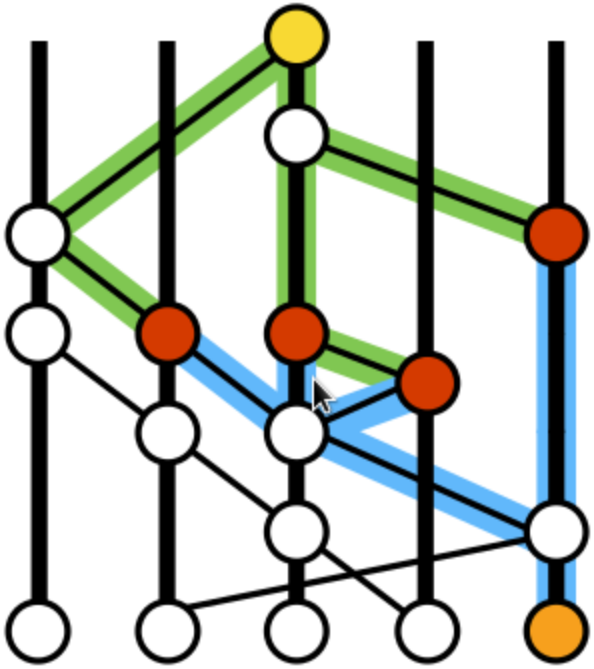

https://swirlds.com/downloads/SWIRLDS-TR-2016-01.pdf

Avalanche

- Gossip protocol
- Metastability

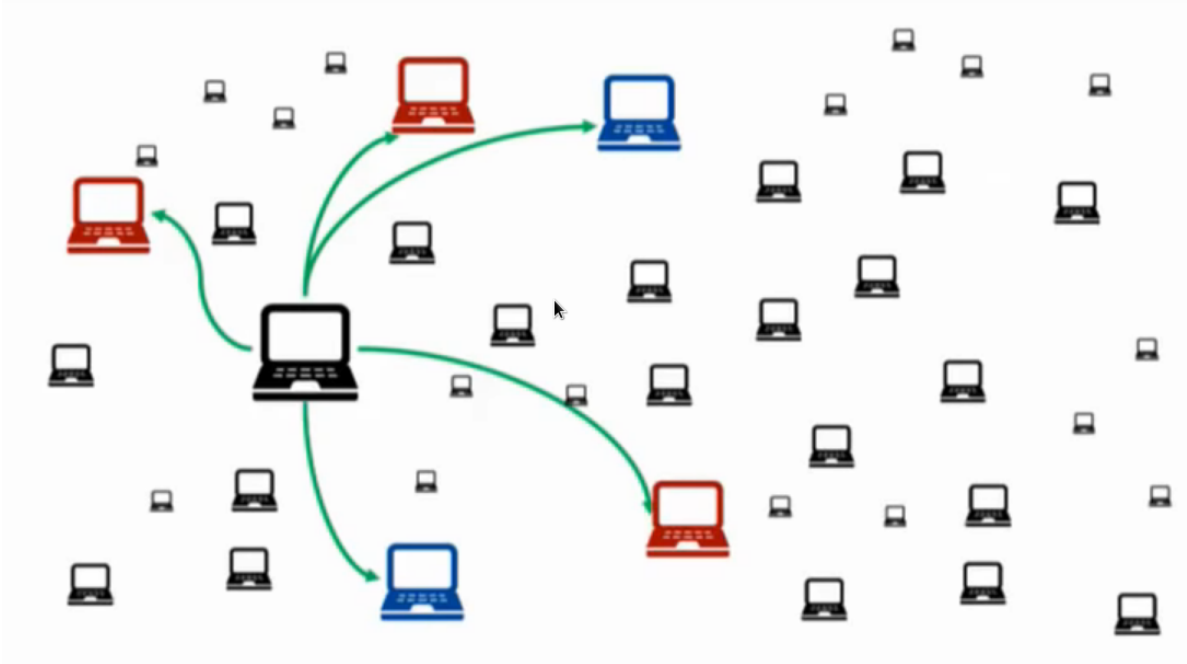

https://www.youtube.com/watch?v=AXrrqtFIGow

PRACTICAL EXCERCISE

Chaloupka-Juranek almost-consensus algorithm with (usually) positive social side-effects:

PRACTICAL EXCERCISE

Chaloupka-Juranek almost-consensus algorithm with (usually) positive social side-effects:

1. Go to https://sched.co/Jcj3
2. Scroll down to Feedback form and click on this icon 
3. Talk to 5 randomly chosen people on the corridor during the conference
4. Introduce yourselves, explain them the algorithm and ask them to execute points 1.-4. (and continue chatting with the person)

Expected result: if most of the DevConf attendees are not Byzantine, we should win best talk competition!

CHALOUPKA-JURÁNEK TAXONOMY

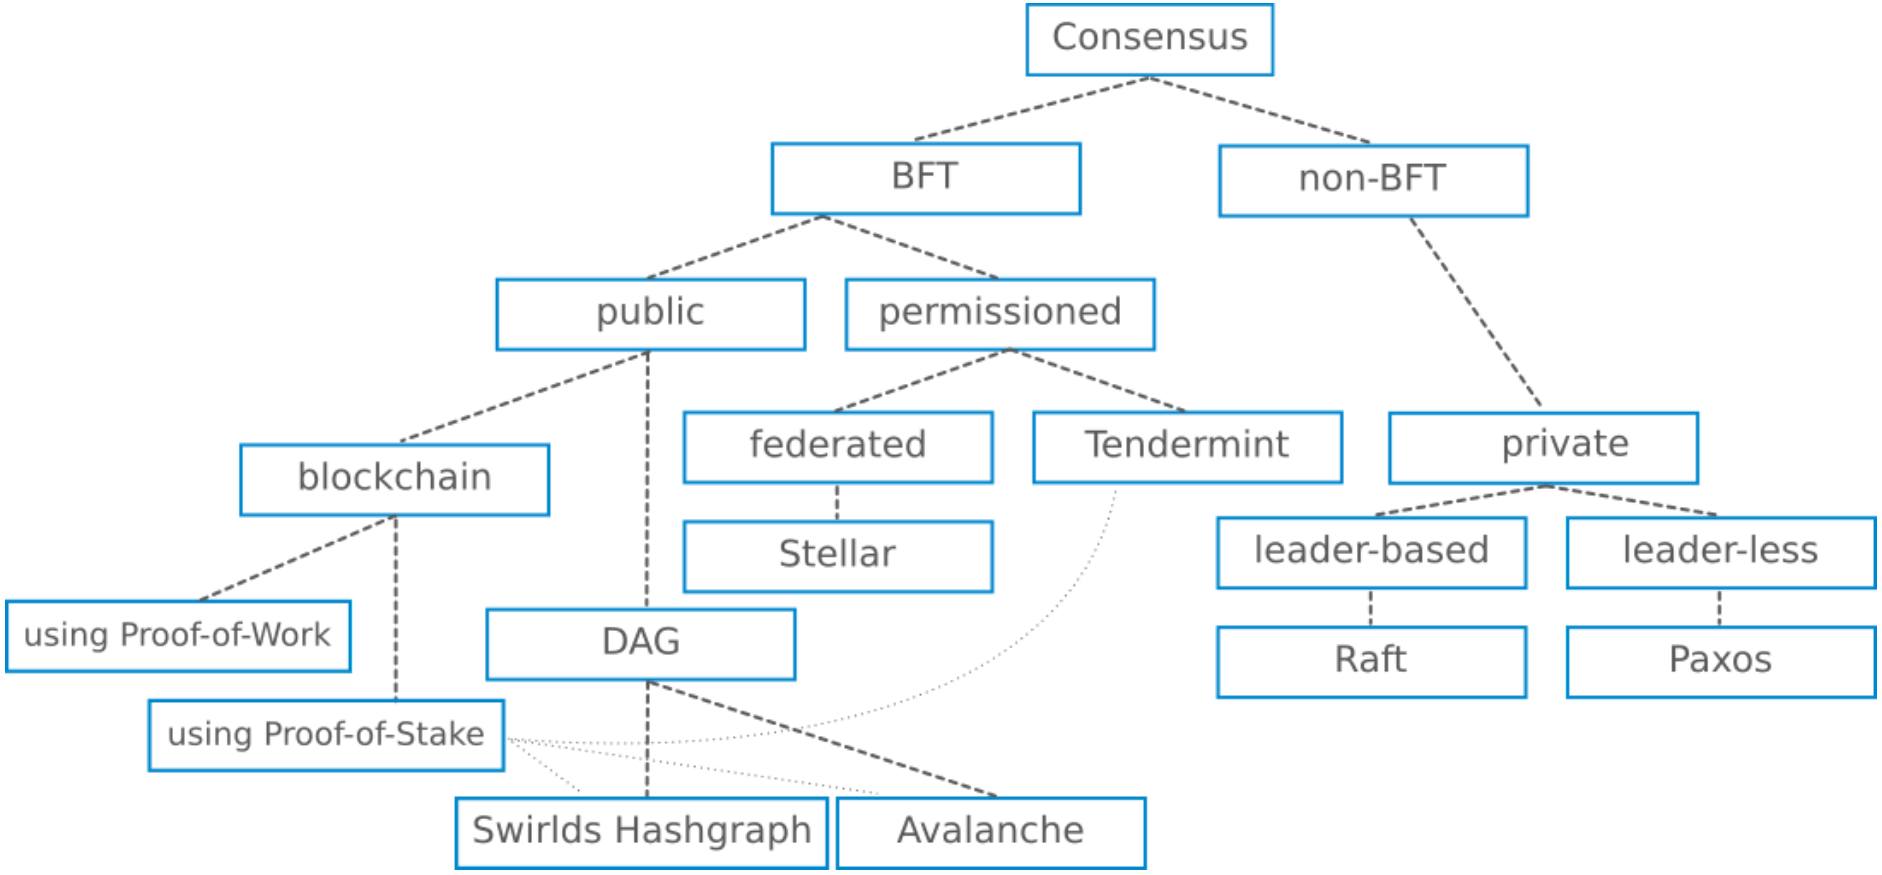

TAKE-OFFS

- Consensus protocol is a crucial part of any distributed ledger.
- Choice of consensus protocol influences heavily many characteristic of distributed ledger (including performance and security).
- There are several types of distributed ledgers, several families of consensus algorithms and not every consensus algorithm is suitable for every distributed ledger.

QUESTIONS

LINKS

- S. Nakamoto, Bitcoin: A Peer-to-Peer Electronic Cash System
- D. Ongaro, J. Ousterhout, In Search of an Understandable Consensus Algorithm
- M. Castro, B Liskov, Practical Byzantine Fault Tolerance
- D. Mazieres, The Stellar Consensus Protocol
- The latest gossip on BFT consensus - Tendermint consensus algorithm
- L. Baird, The Swirlds Hashgraph Consensus Algorithm
- Team Rocket, Snowflake to Avalanche

Thank you for your attention!

BACKUP SLIDES

BITCOIN 51% ATTACK

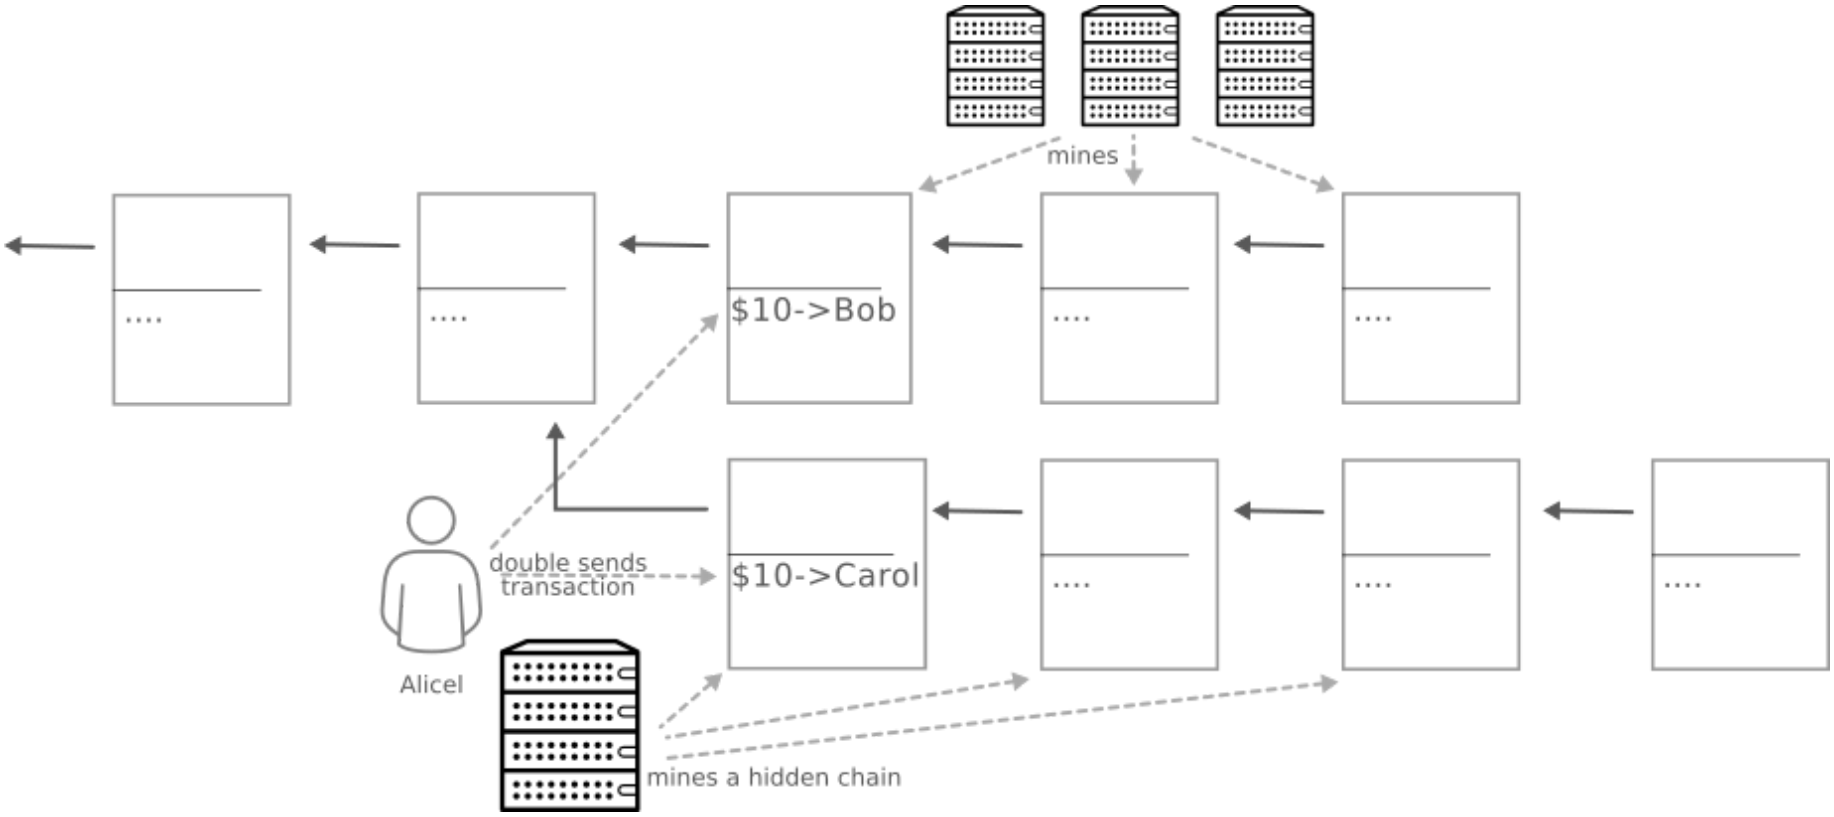

POS: NOTHING AT STAKE

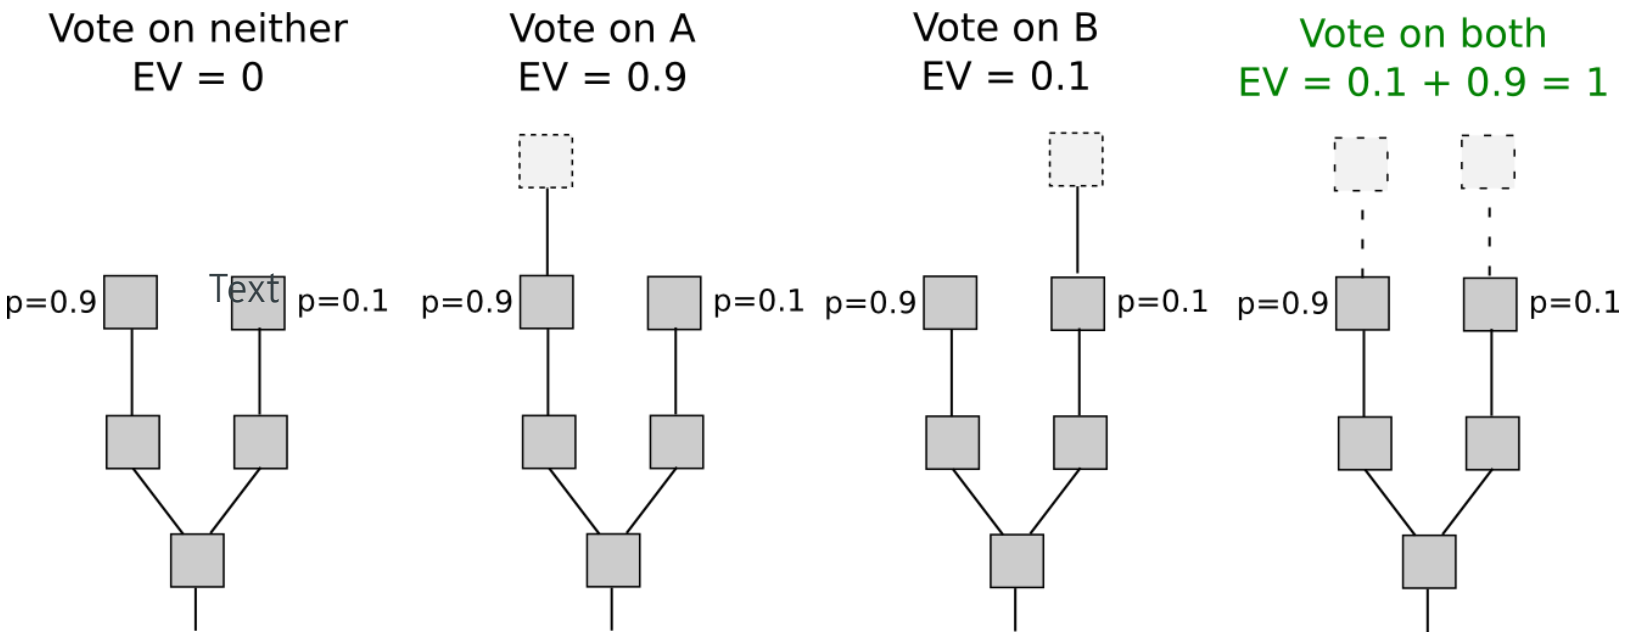

https://github.com/ethereum/wiki/wiki/Proof-of-Stake-FAQs

Proof-of-X

There are many:

- Proof of Capacity
- Proof of Elapsed Time (PoET)
- Proof of Authority
- Proof of Activity
- Proof of Burn
- Proof of Weight
- ...

Beware: Proof-of-X doesn't mean it's similar to PoW, actually in many cases it's quite different (e.g. centralized).

TAXONOMY TABLE

Proof-of-Work	Bitcoin	Public Permissionless	Byzantine tolerant	Probabilistic
Proof-of-Stake	Ethereum 2.0	Public Permissionless	Byzantine tolerant	Finite
Delegated PoS	Stellar	Permissioned public	Byzantine tolerant	Finite
Raft	Hyperledger	Permissioned private	non-Byzantine tolerant	Finite
Tendermint	Cosmos	Permissionless public	Byzantine tolerant	Finite
Swirlds	Hedera Hashgraph	Permissionless public	Byzantine tolerant	Finite
Avalanche	Ava	Permissionless public	Byzantine tolerant	Probabilistic