A Quick Introduction to How Blockchains Work

Robbert van Renesse
Cornell University

(some slides due to former postdoc
Ittay Eyal, now at Technion)

Blockchain's Promise and Limits

Promises

- Global currency
- *Smart contracts*
- Notarization
- Accountability
- …

State Machine Replication (Lamport'78)

- A generic way to tolerate failures
- Simply start multiple *replicas* (copies) of a state machine, and keep them in sync by agreeing on the transitions (operations) and the order in which to apply them

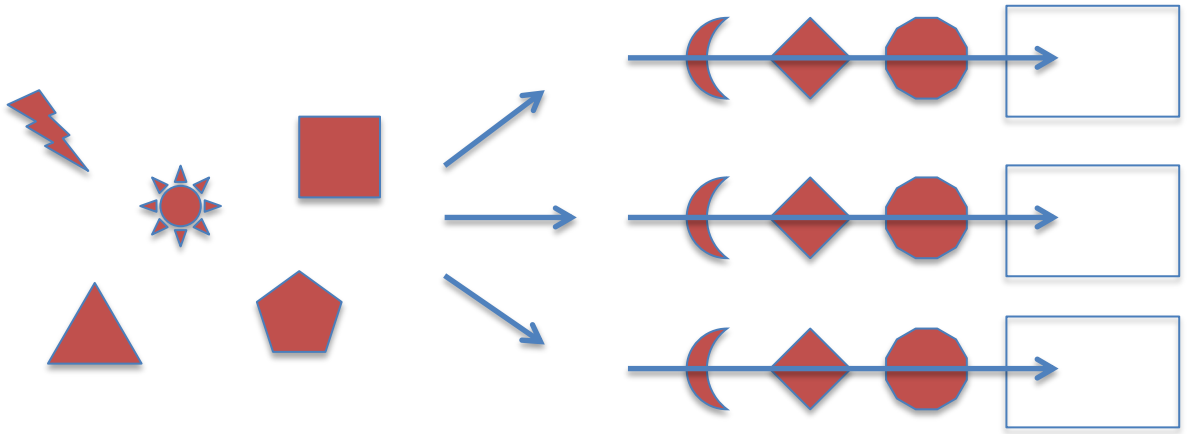

A Replicated Ledger of Transactions

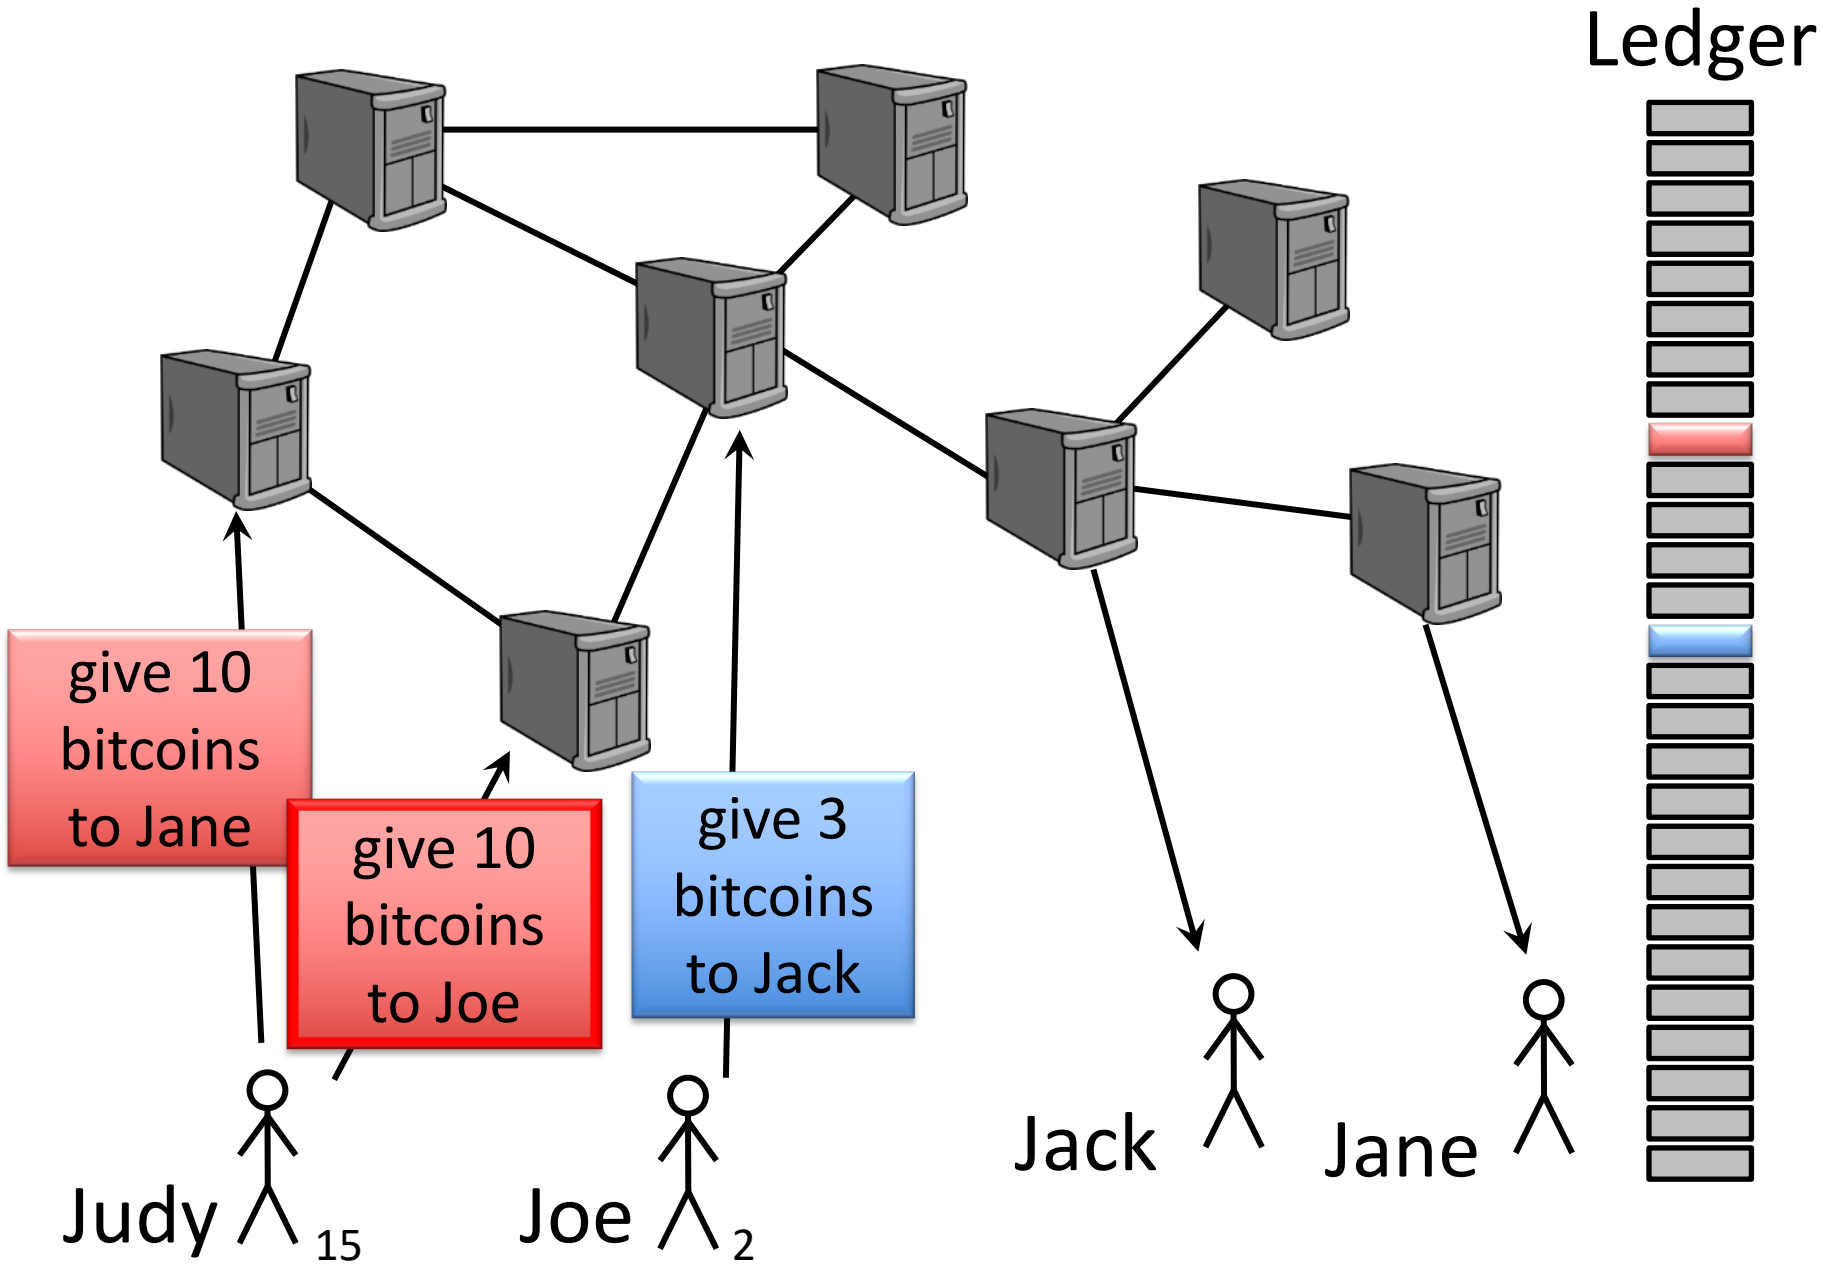

The Blockchain

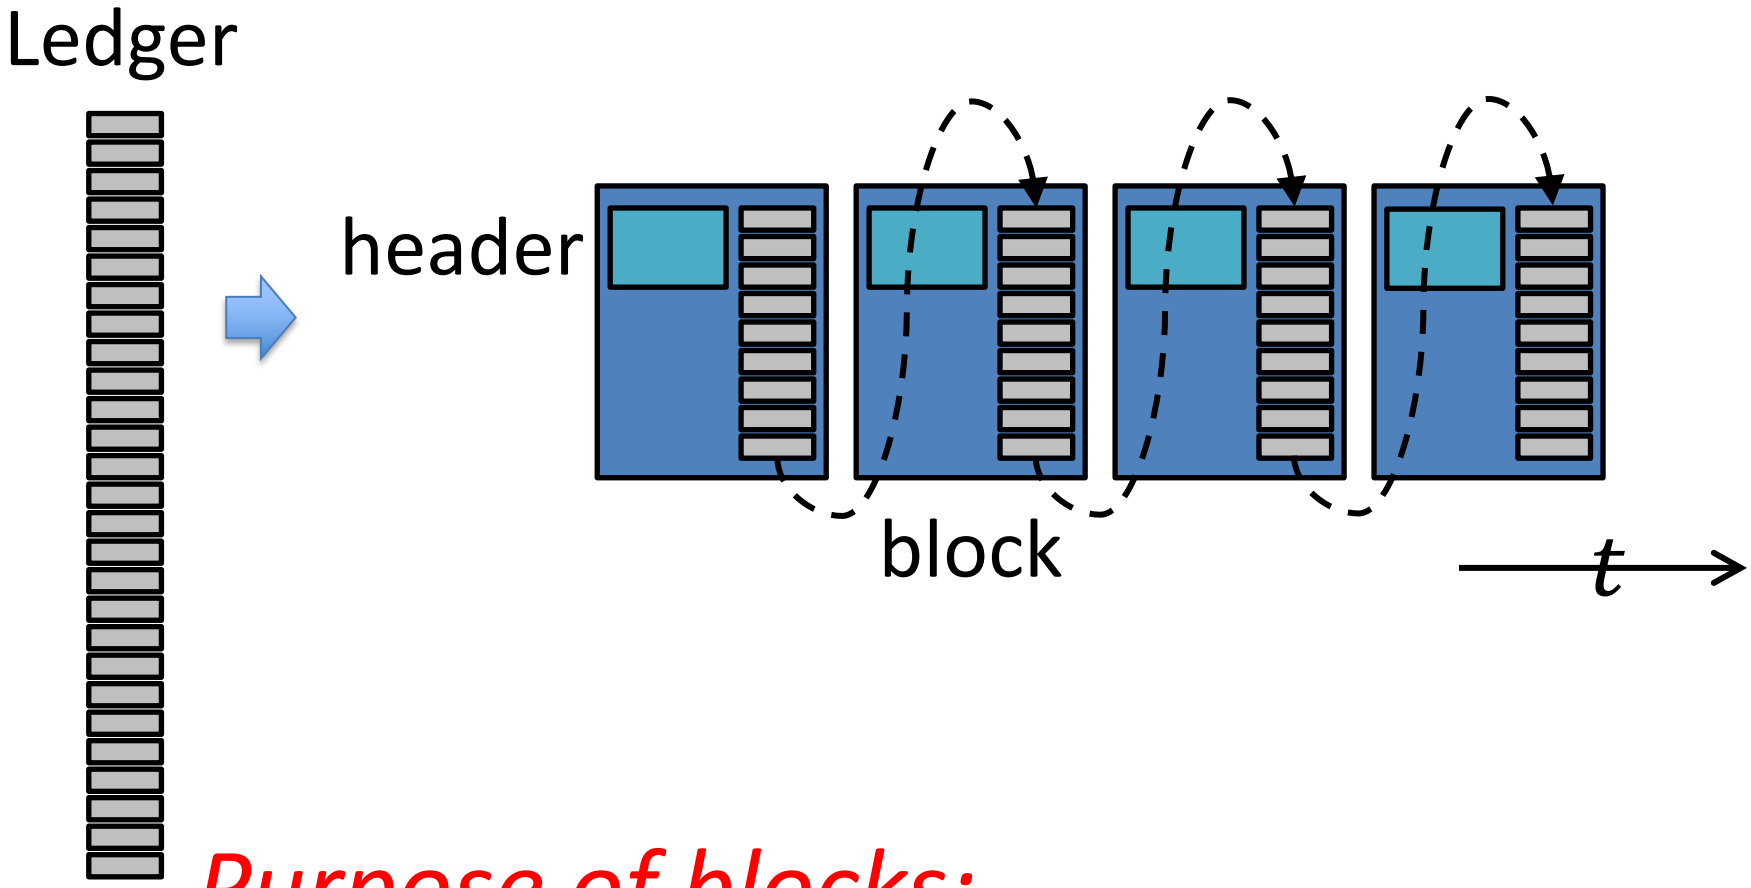

Purpose of blocks:

***Batching** transactions for efficiency!*

Cryptographic One-Way Hash Function

hash(X) = Y

- Given X it is easy to compute Y (the *digest*)
- Given Y it is computationally infeasible to find X
  - *unless you already know X, of course*
- It is computationally infeasible to find X_1 and X_2, $X_1 \neq X_2$, such that hash(X_1) = hash(X_2)
- In some sense, *Y identifies X*

Examples: SHA-256, SHA-3

Note: unlike an ordinary hash function where you typically have fewer buckets than objects and thus multiple objects per bucket, with cryptographic hash functions you typically have many more "virtual buckets" than objects, and at most one object in a bucket

The Blockchain

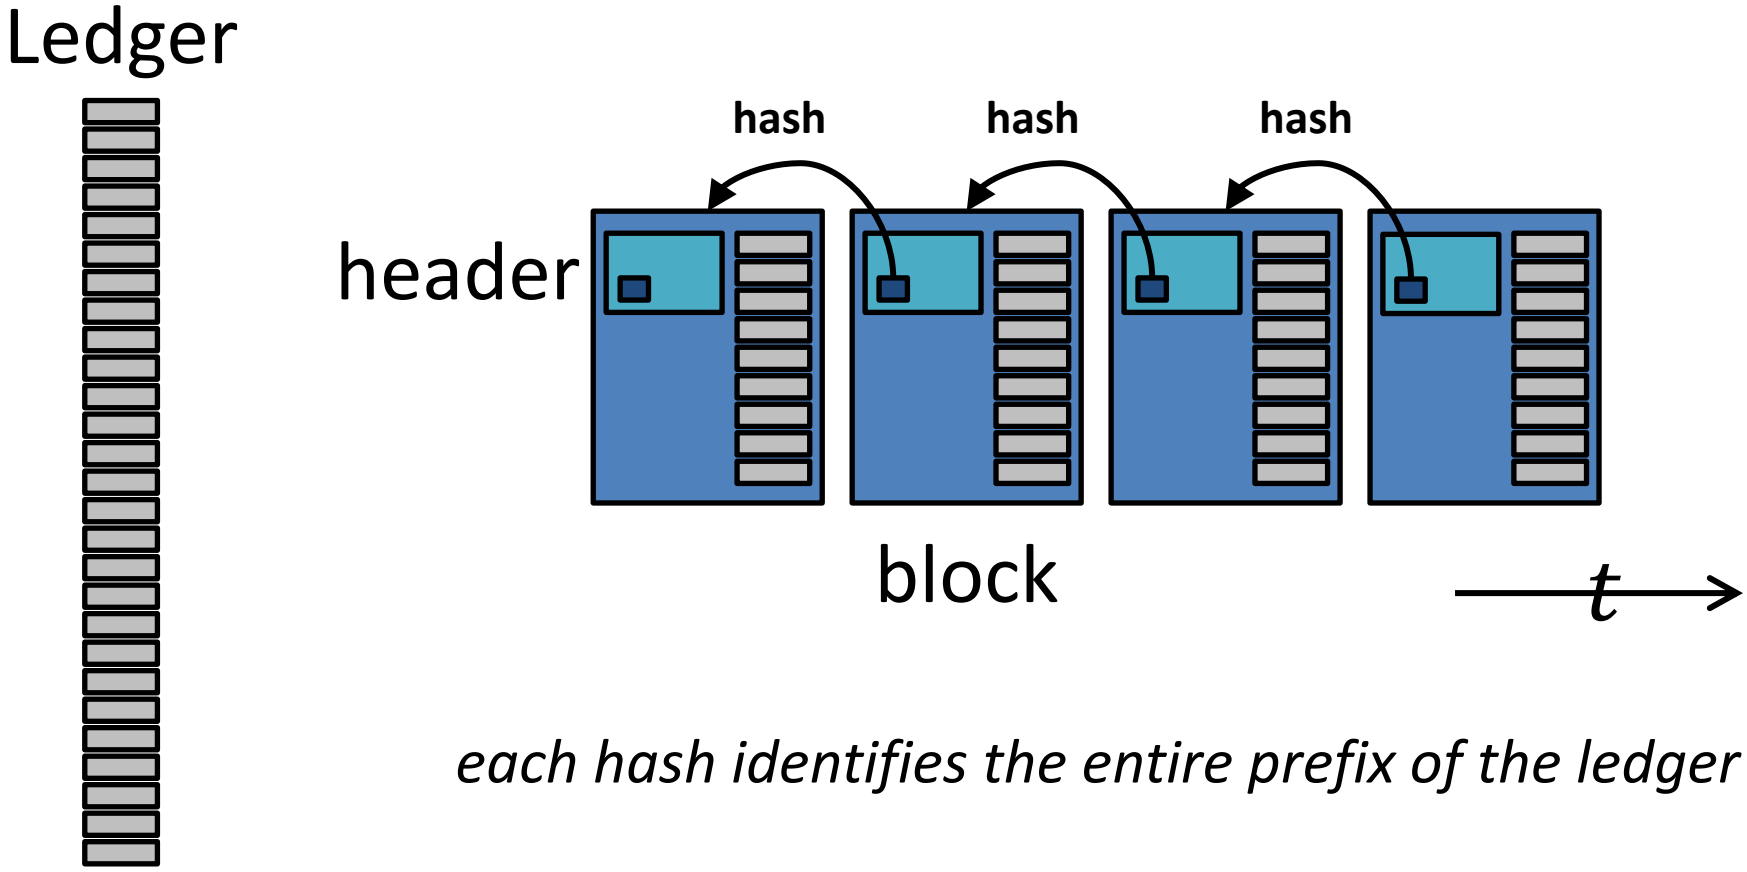

each hash identifies the entire prefix of the ledger

Blockchain Desirables

- Performance:
  - High Throughput, Low Latency
  - Energy-Efficient
- Security:
  - Always available for reading and appending
  - Fair
  - Tamperproof (Integrity)
  - Possibly confidentiality as well
- No Single Administrative Domain
- Open membership (or not)

Open Membership is Hard

- Traditional replication is based on voting
- Problem: "Sybil" or impersonation attacks
  - a participant may try to vote multiple times
  - with open membership, anybody can create identities and vote many times

Permissionless vs Permissioned Blockchains

	Permissionless	Permissioned
Approach	Competitive	Cooperative
Basic technique	Proof-of-Resource	Byzantine Consensus
Trust requirements	Crypto (+ peers…)	Peers (+ crypto…)
Membership	Open	Closed
Energy-efficiency	Often terrible	Excellent
Transaction rate	At best hundreds / sec	Many thousands per second
Txn latency	As high as many minutes	Less than a second

Bitcoin Blockchain

- Permissionless, open membership
- Proof-of-Work
- There are thousands of Bitcoin miners
  - often use cheap, dirty coal-based energy
  - they use ASIC hardware to compute SHA256 hashes
  - use about as much energy as an average European country
- Overall rate is a few transactions per second

The Blockchain

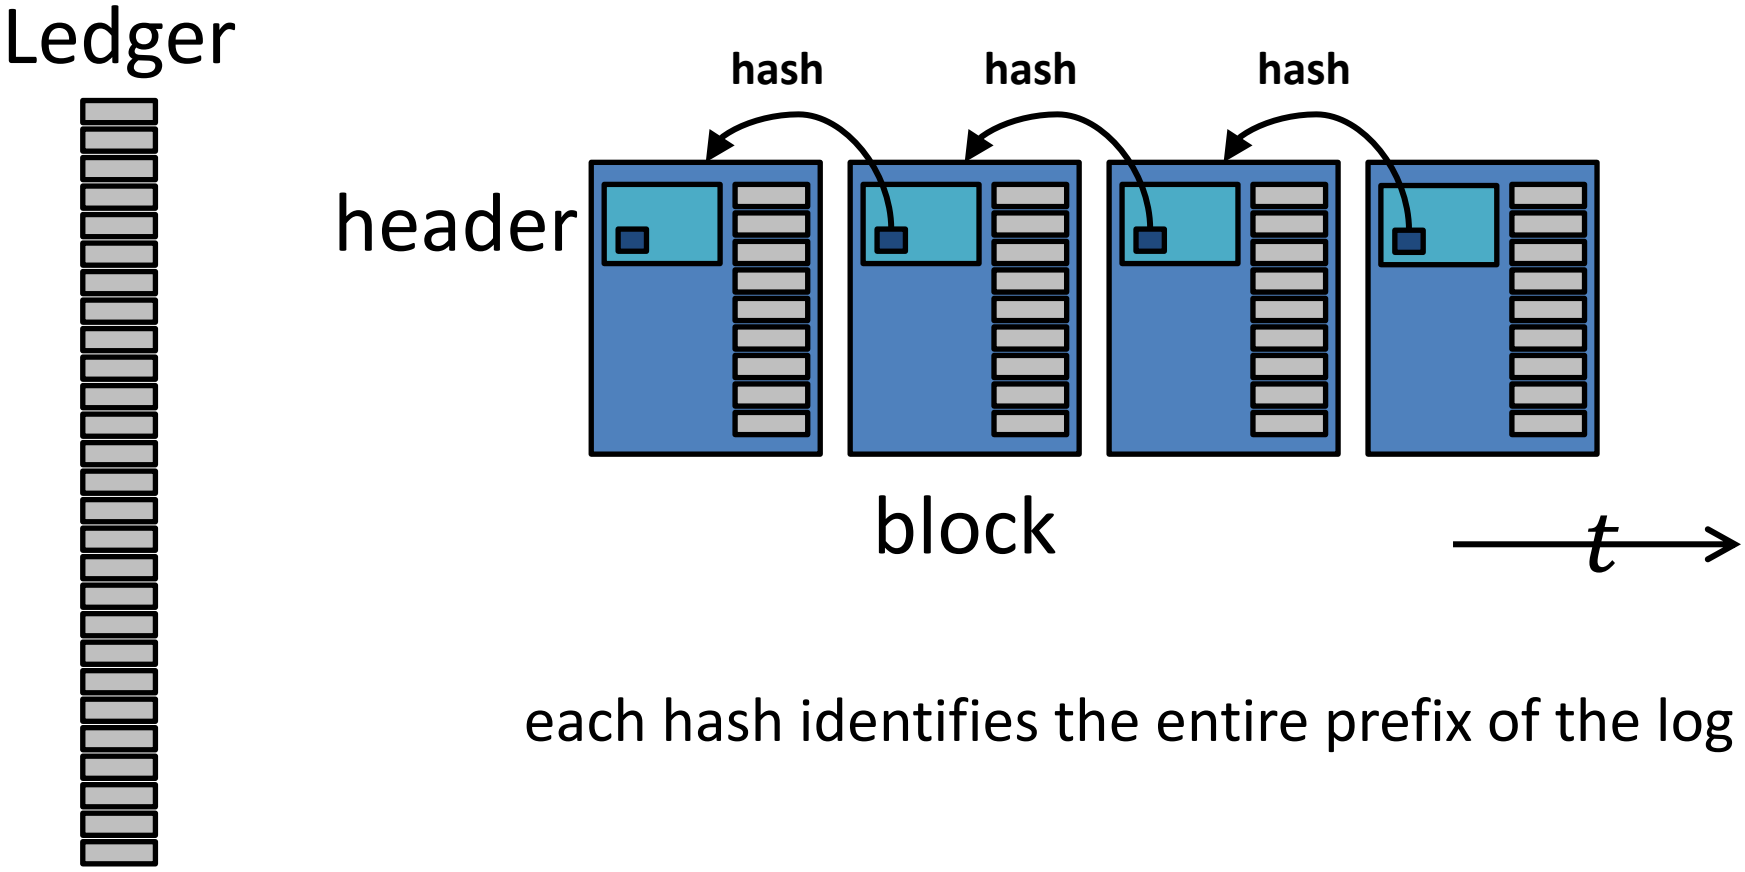

each hash identifies the entire prefix of the log

The Blockchain

Ledger

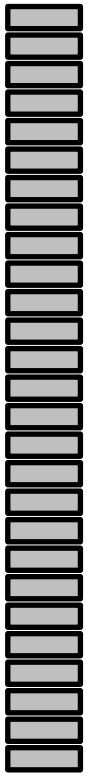

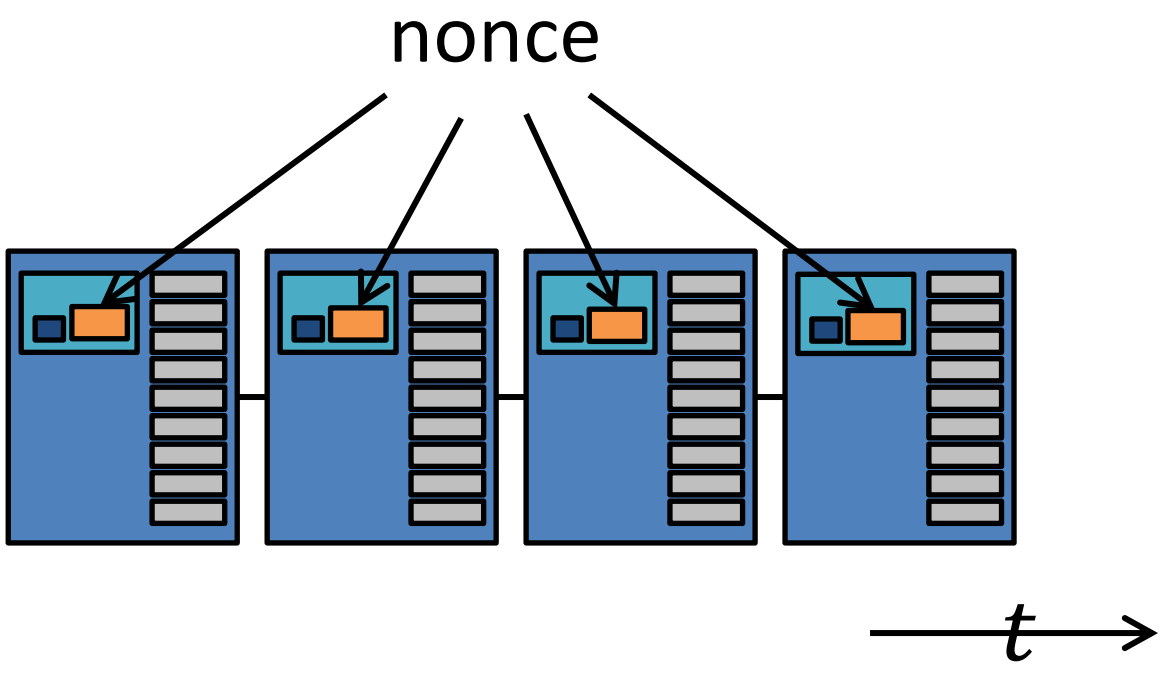

HASH(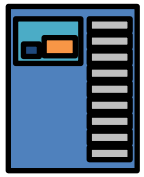) < ***target***

"cryptopuzzle"

The Blockchain

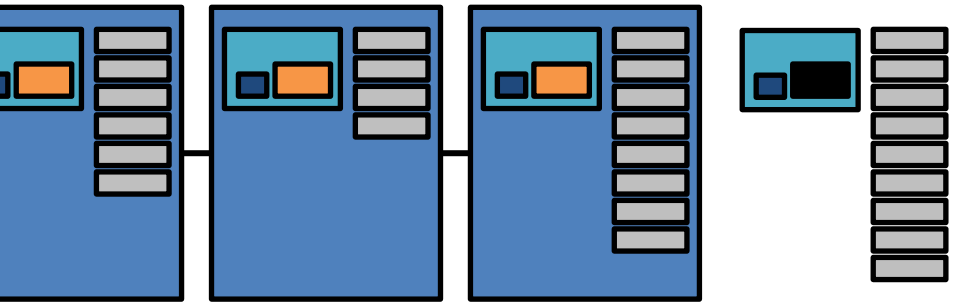

The Blockchain

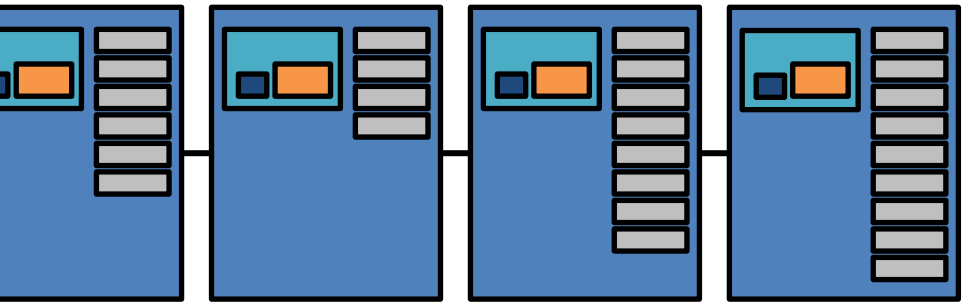

The Blockchain

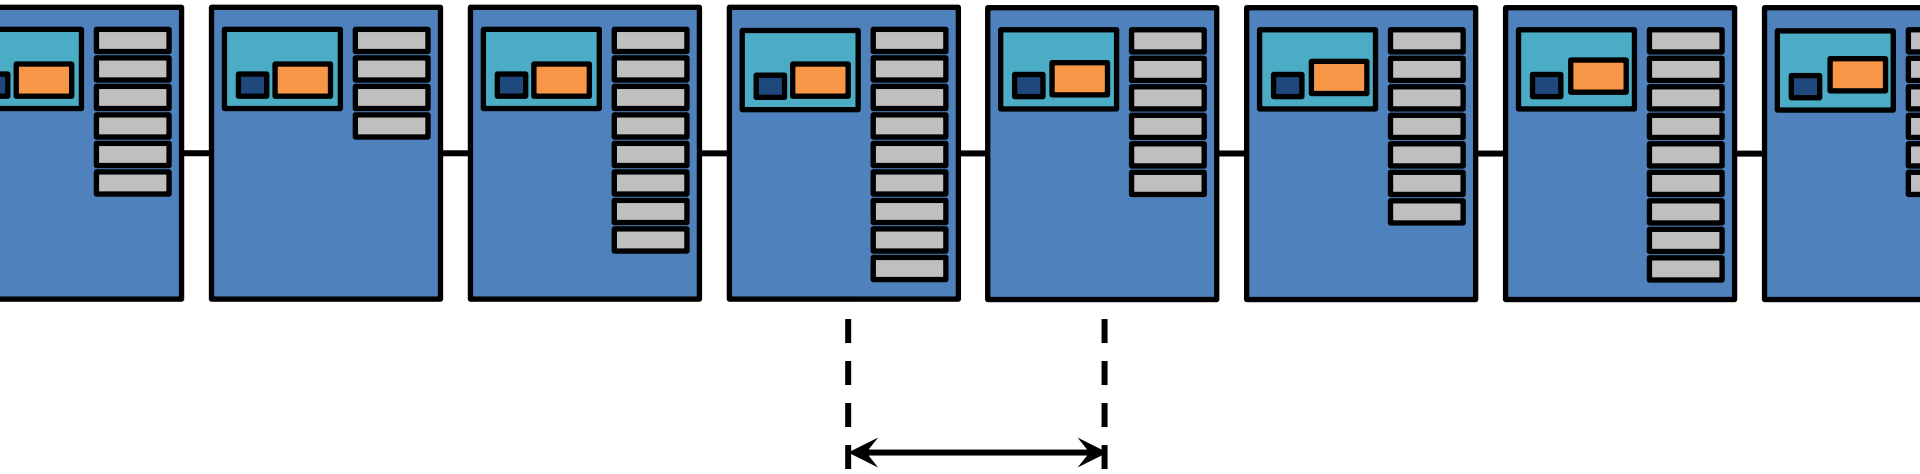

Exponentially distributed, with
constant mean interval

target automatically adjusted every 2016 blocks so that mean interval is **10 minutes**

Incentives for Mining

Prize:

- **“Minting”**
- **Transaction Fees**

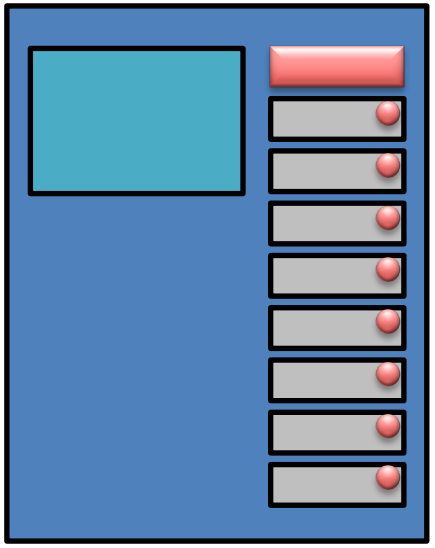

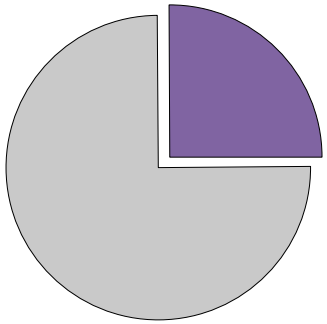

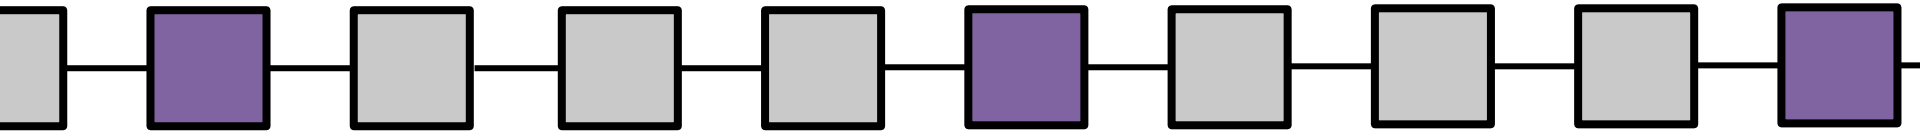

Wins *proportional to computation power*

Forks

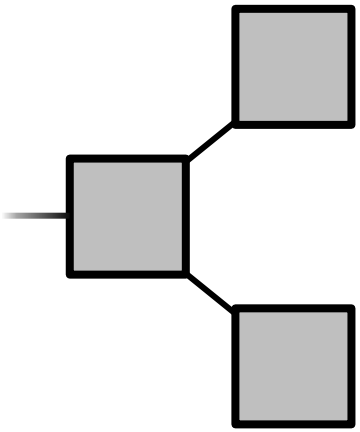

Two blocks “mined” at approximately the same time by two different miners

Fork Resolution

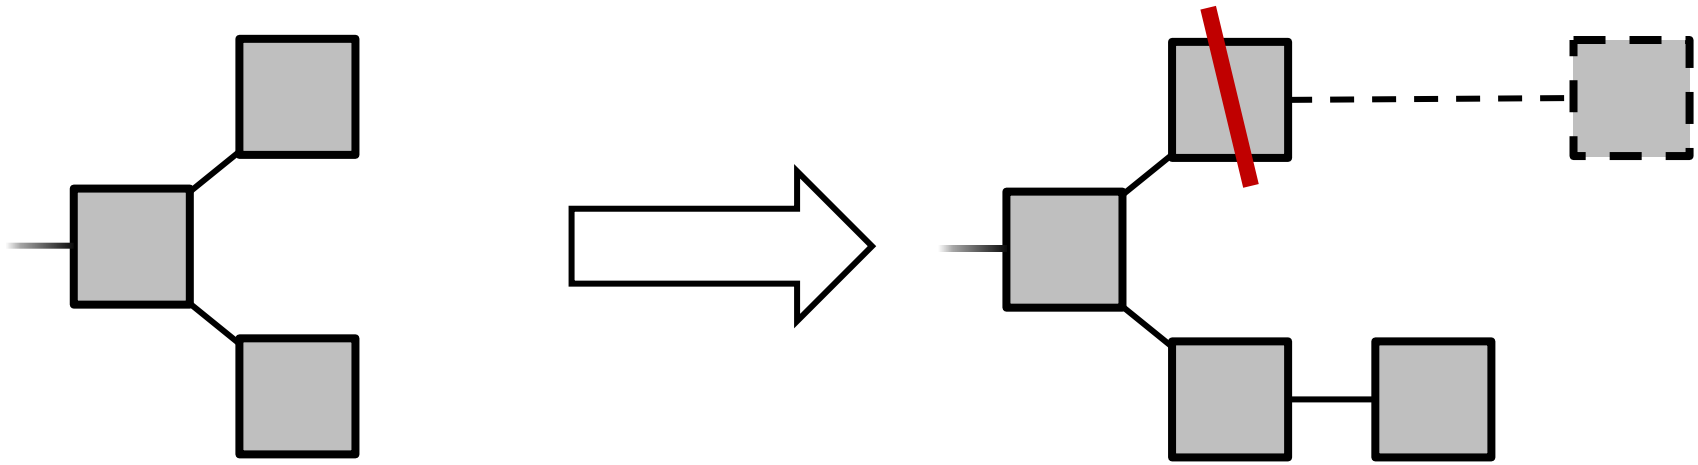

- **Longest** chain wins
- Transactions on short chain are reverted

Fork Resolution

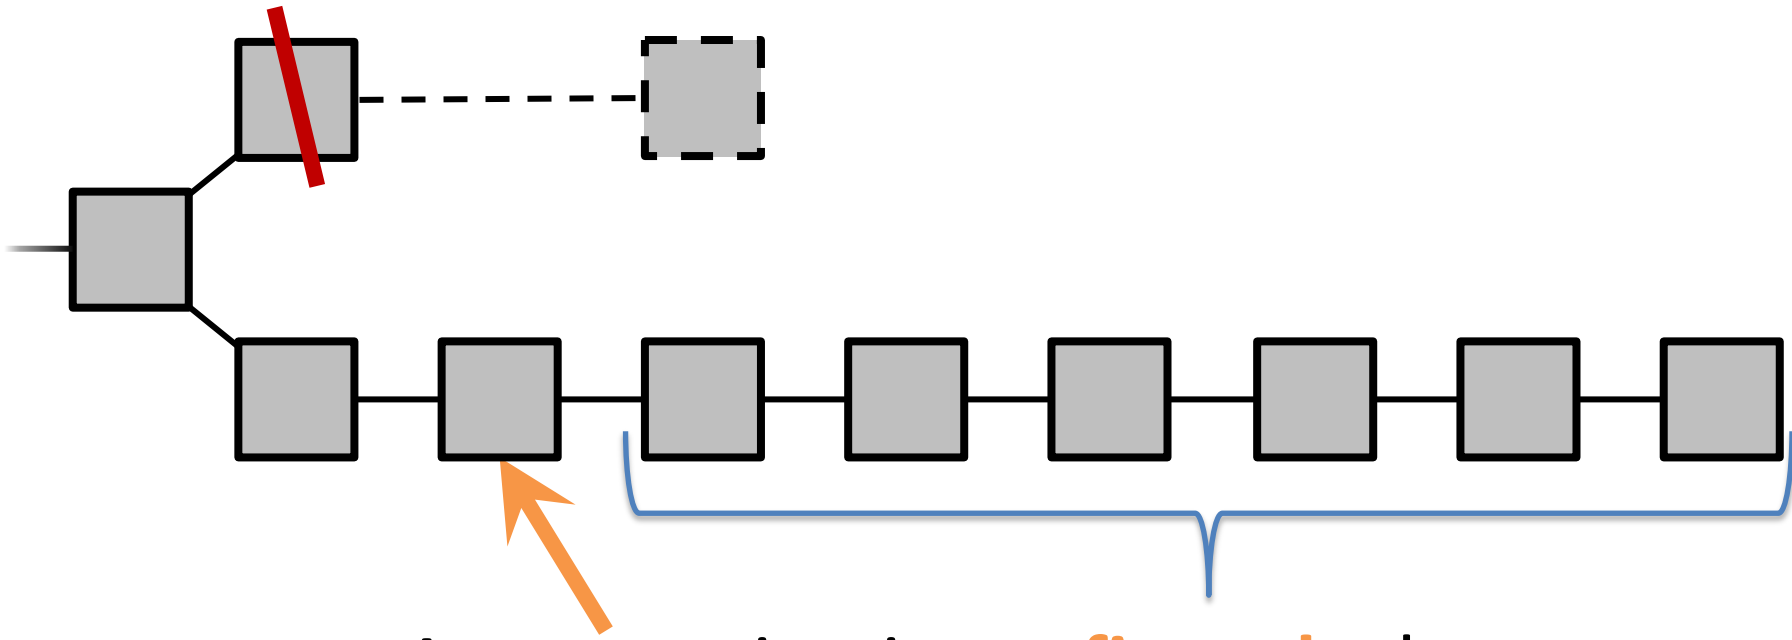

A transaction is **confirmed** when
it is **buried** "deep enough"
(typically 6 blocks – i.e., one hour)

Security Threat!

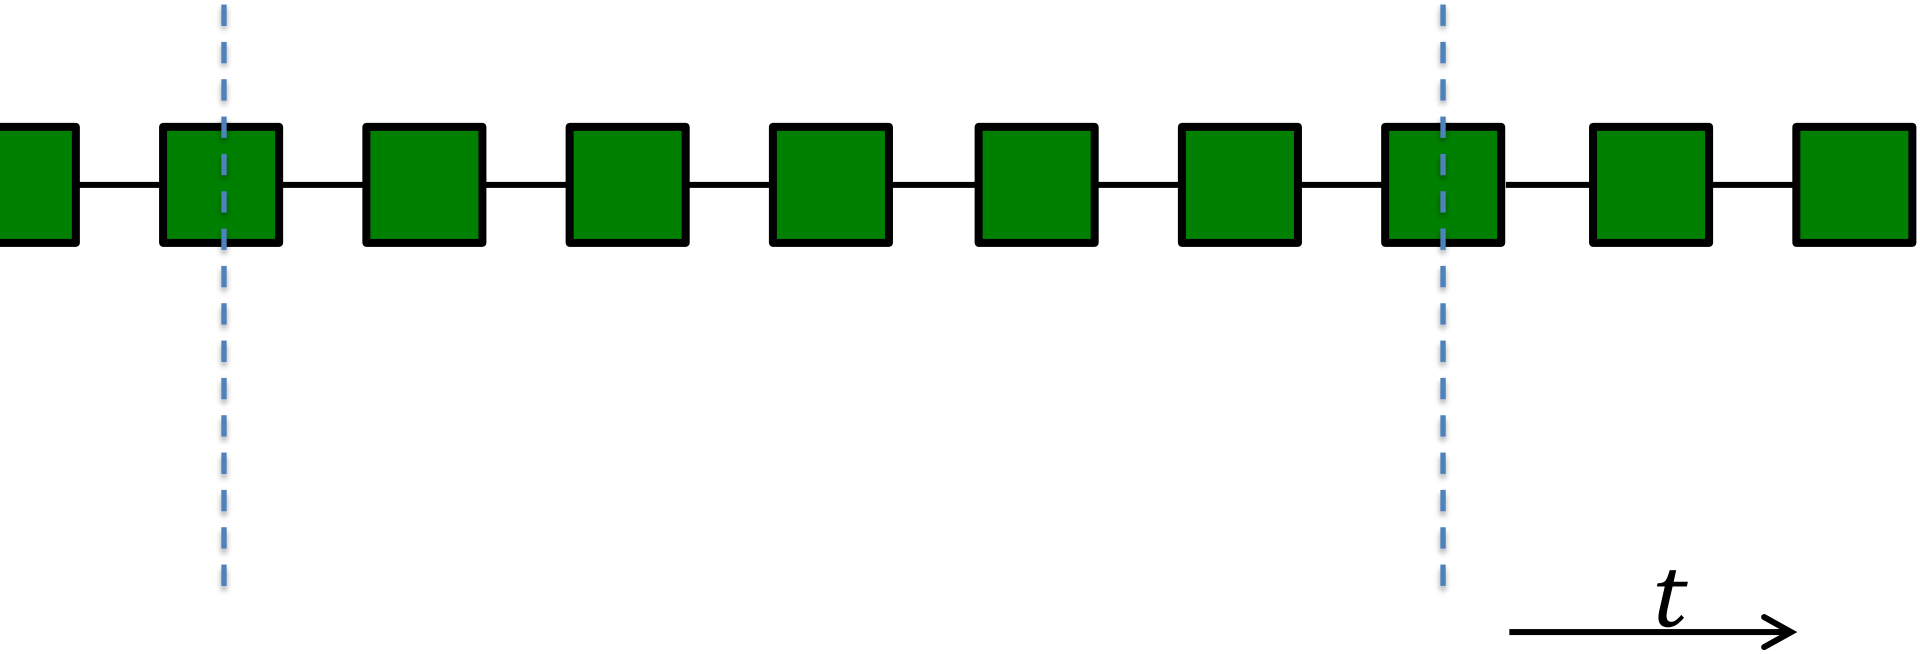

Security Threat!

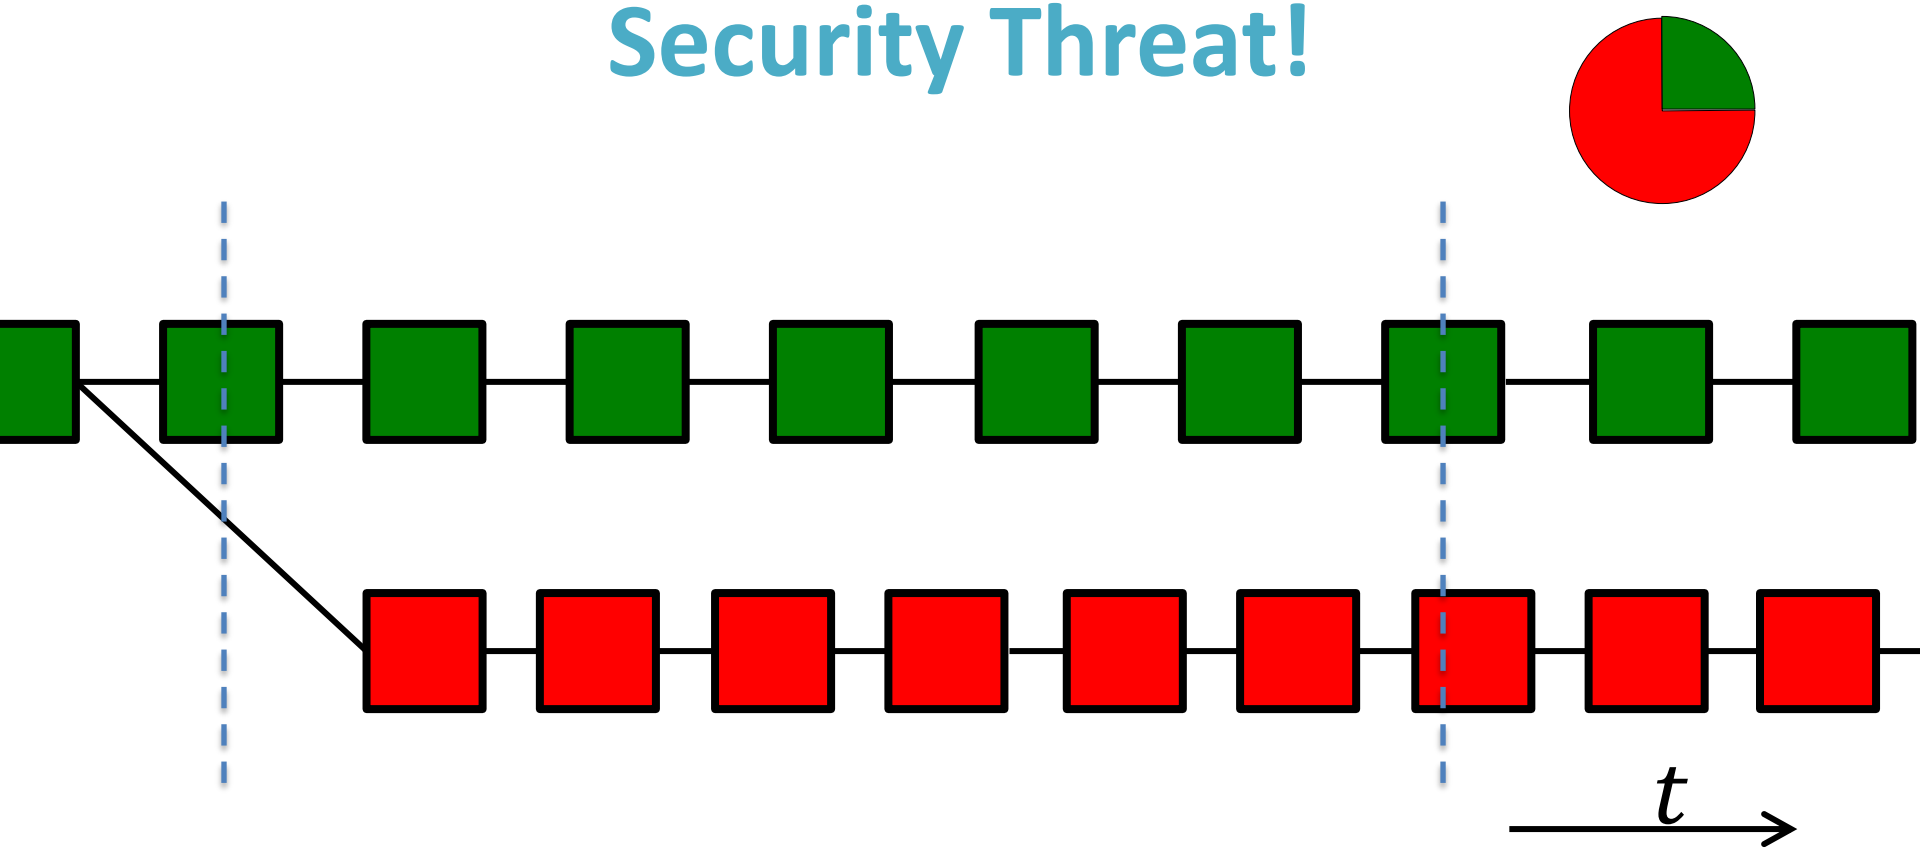

Threat: attacker outruns good miners

Security Threat!

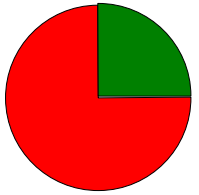

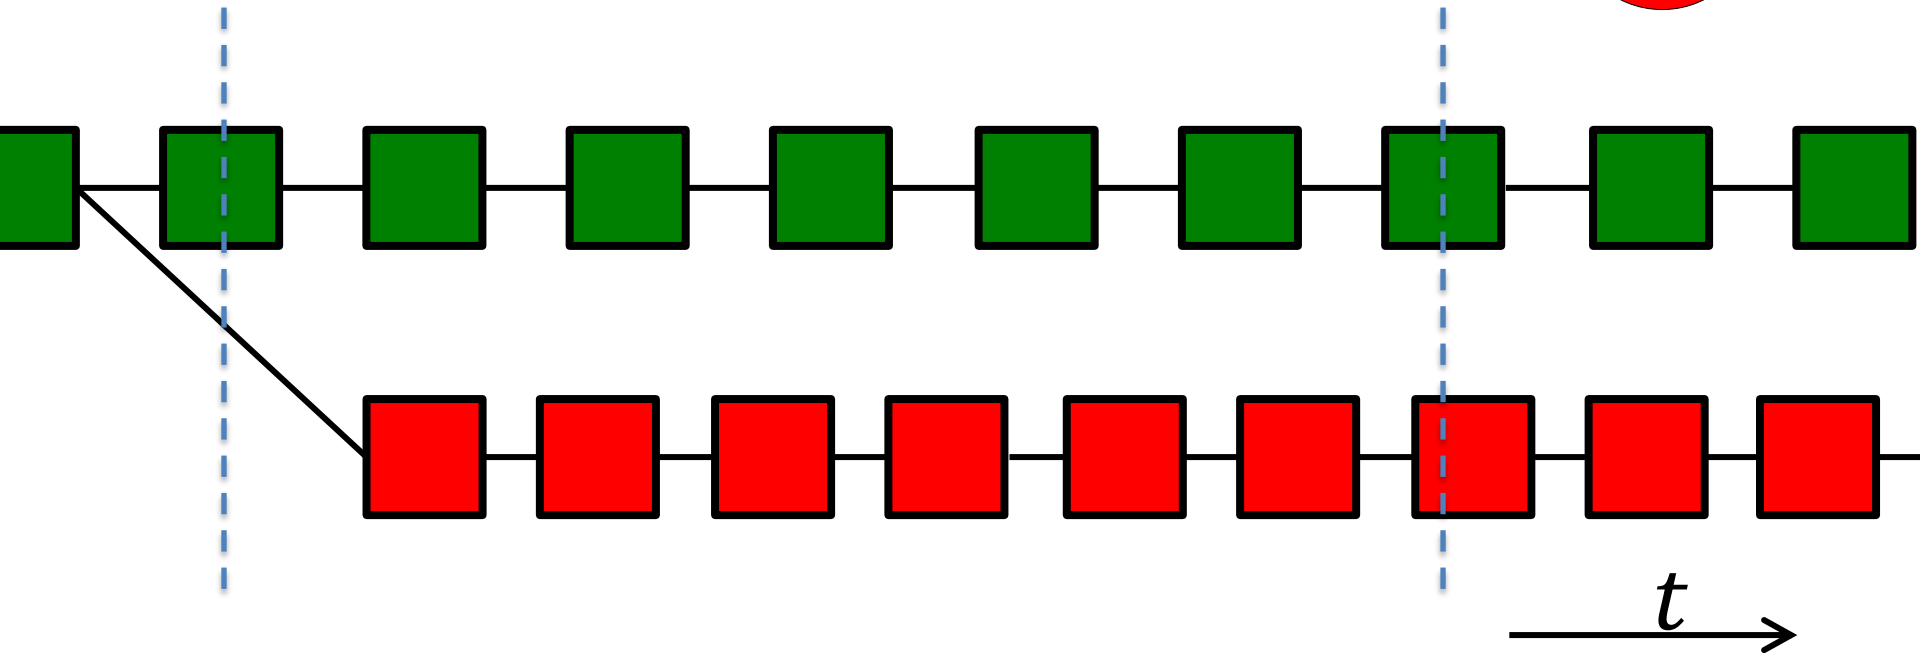

Threat: attacker outruns good miners

➔ **Security Assumption**: *good miners own >.5 of the total compute power*

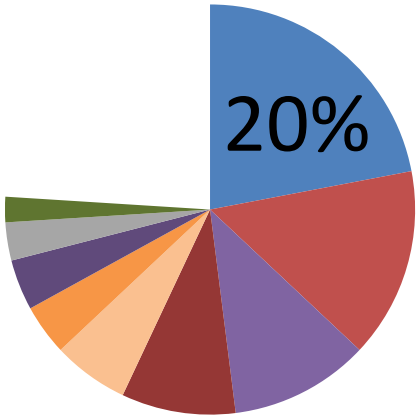

[blockchain.info, April 2015]

Bitcoin Parameters

- block size = 1 MByte
- *target* set such that E(block interval) = 10 minutes
  - the block size is small enough and the block interval is large enough such that all miners can learn about a new block (through “gossiping”) before the next block is mined
- Results in fewer than 10 transactions per second
- *Why not a larger block or a smaller interval?*

Security-Performance Tradeoff

Nakamoto's Blockchain exhibits a tradeoff:
[Sompolinsky+'15, Lewenberg+'15]

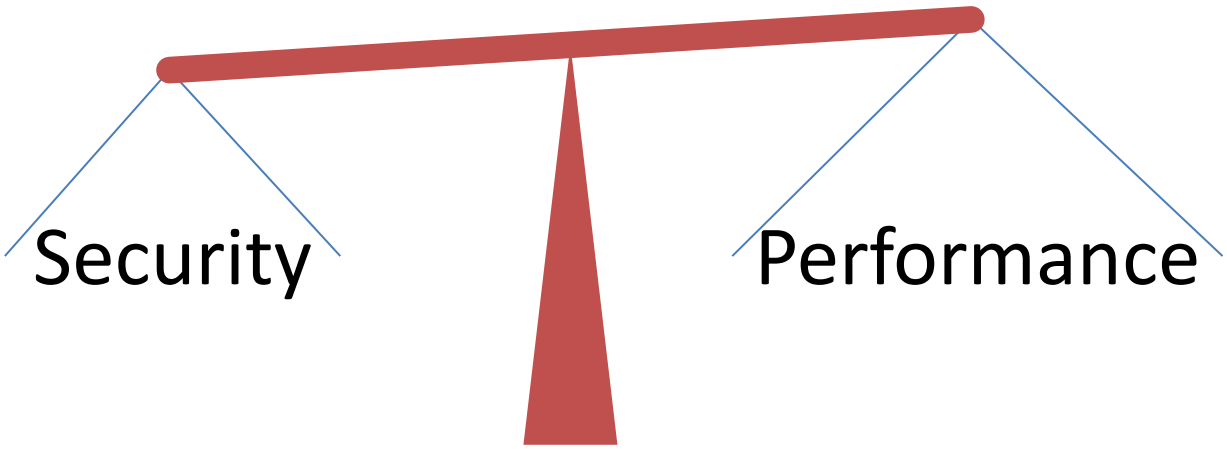

Metrics

Performance:

- Throughput
- Latency

Security:

- Mining power utilization
- Fairness

Mining Power Utilization

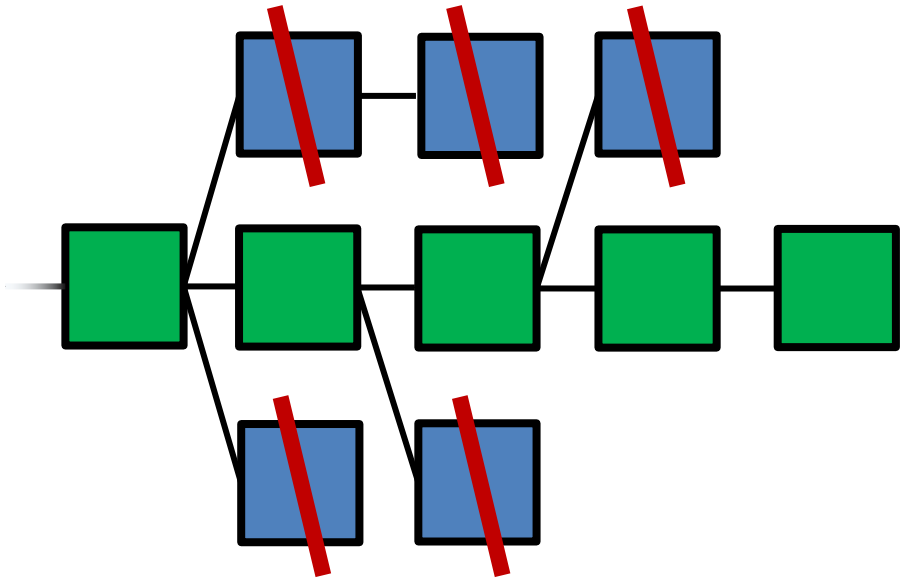

Attacker only has to out-run the main chain

Block Frequency

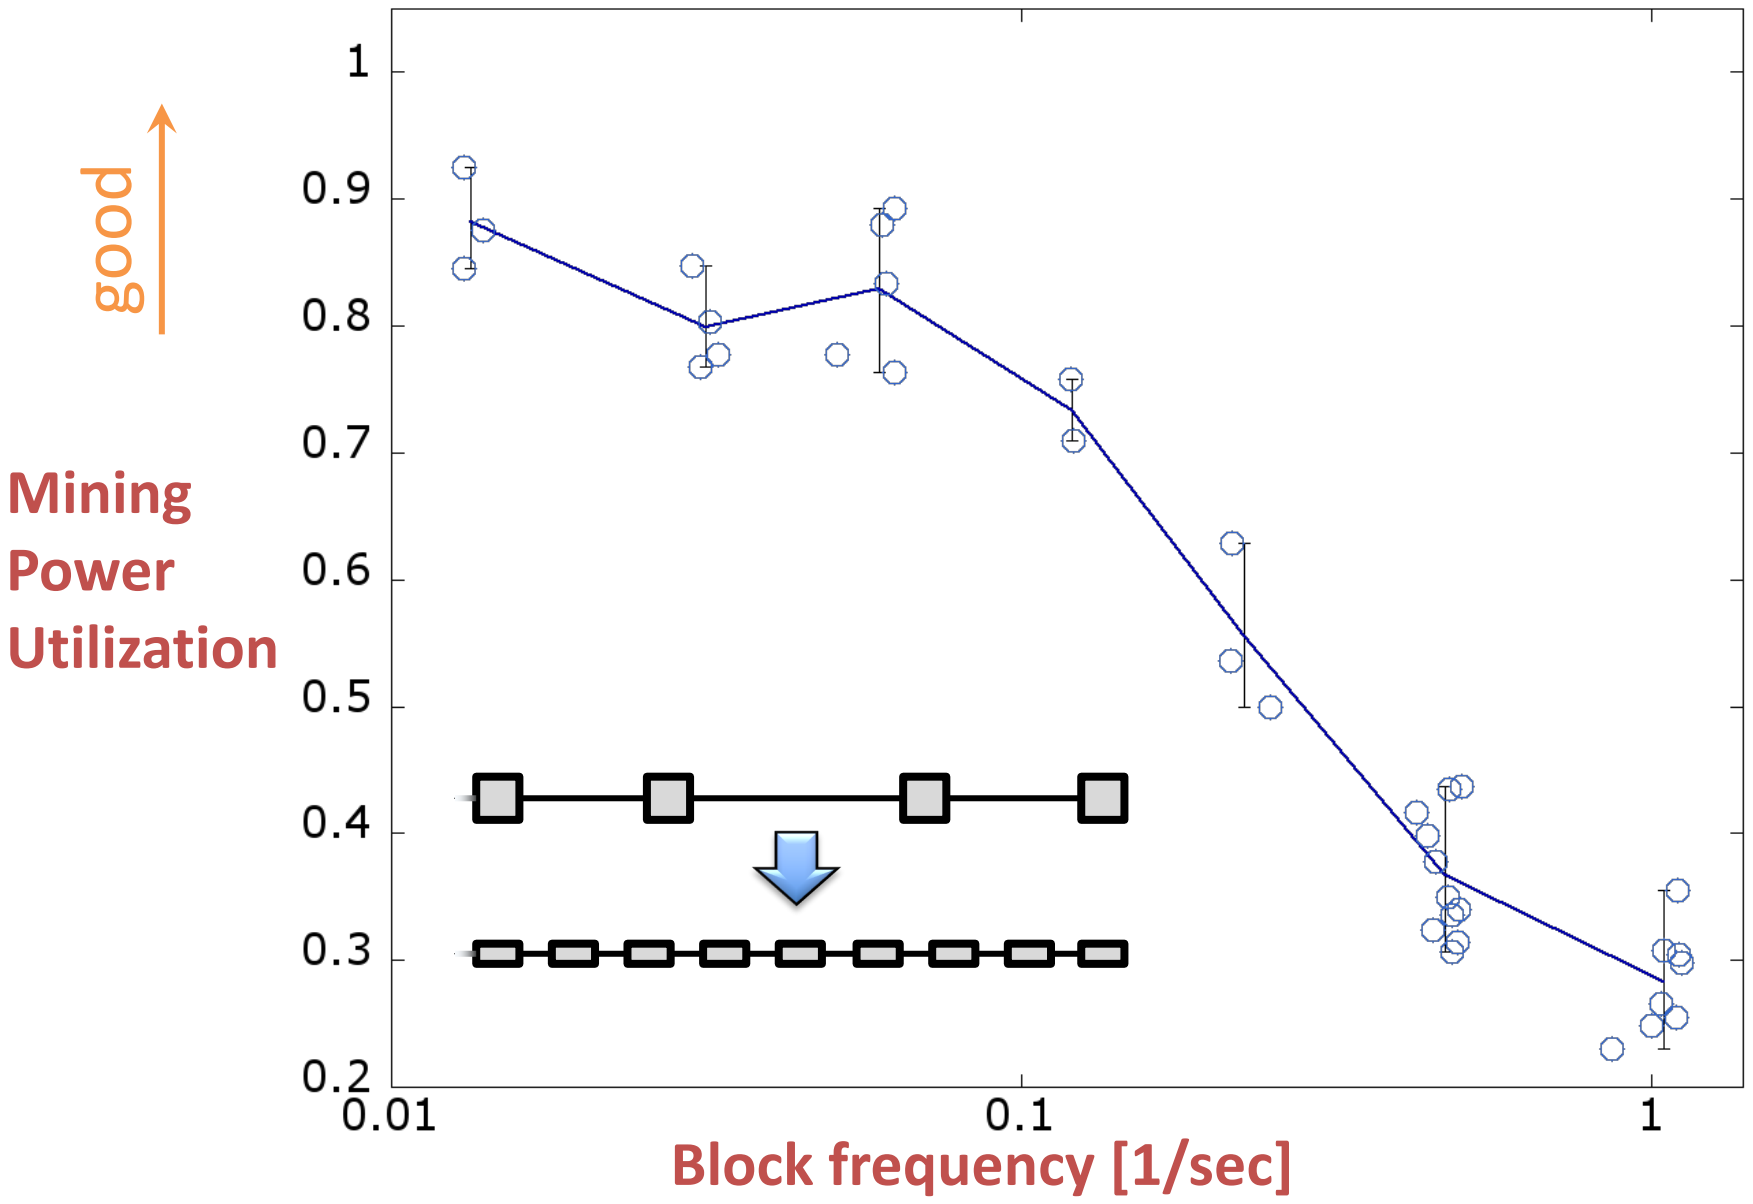

Block Size

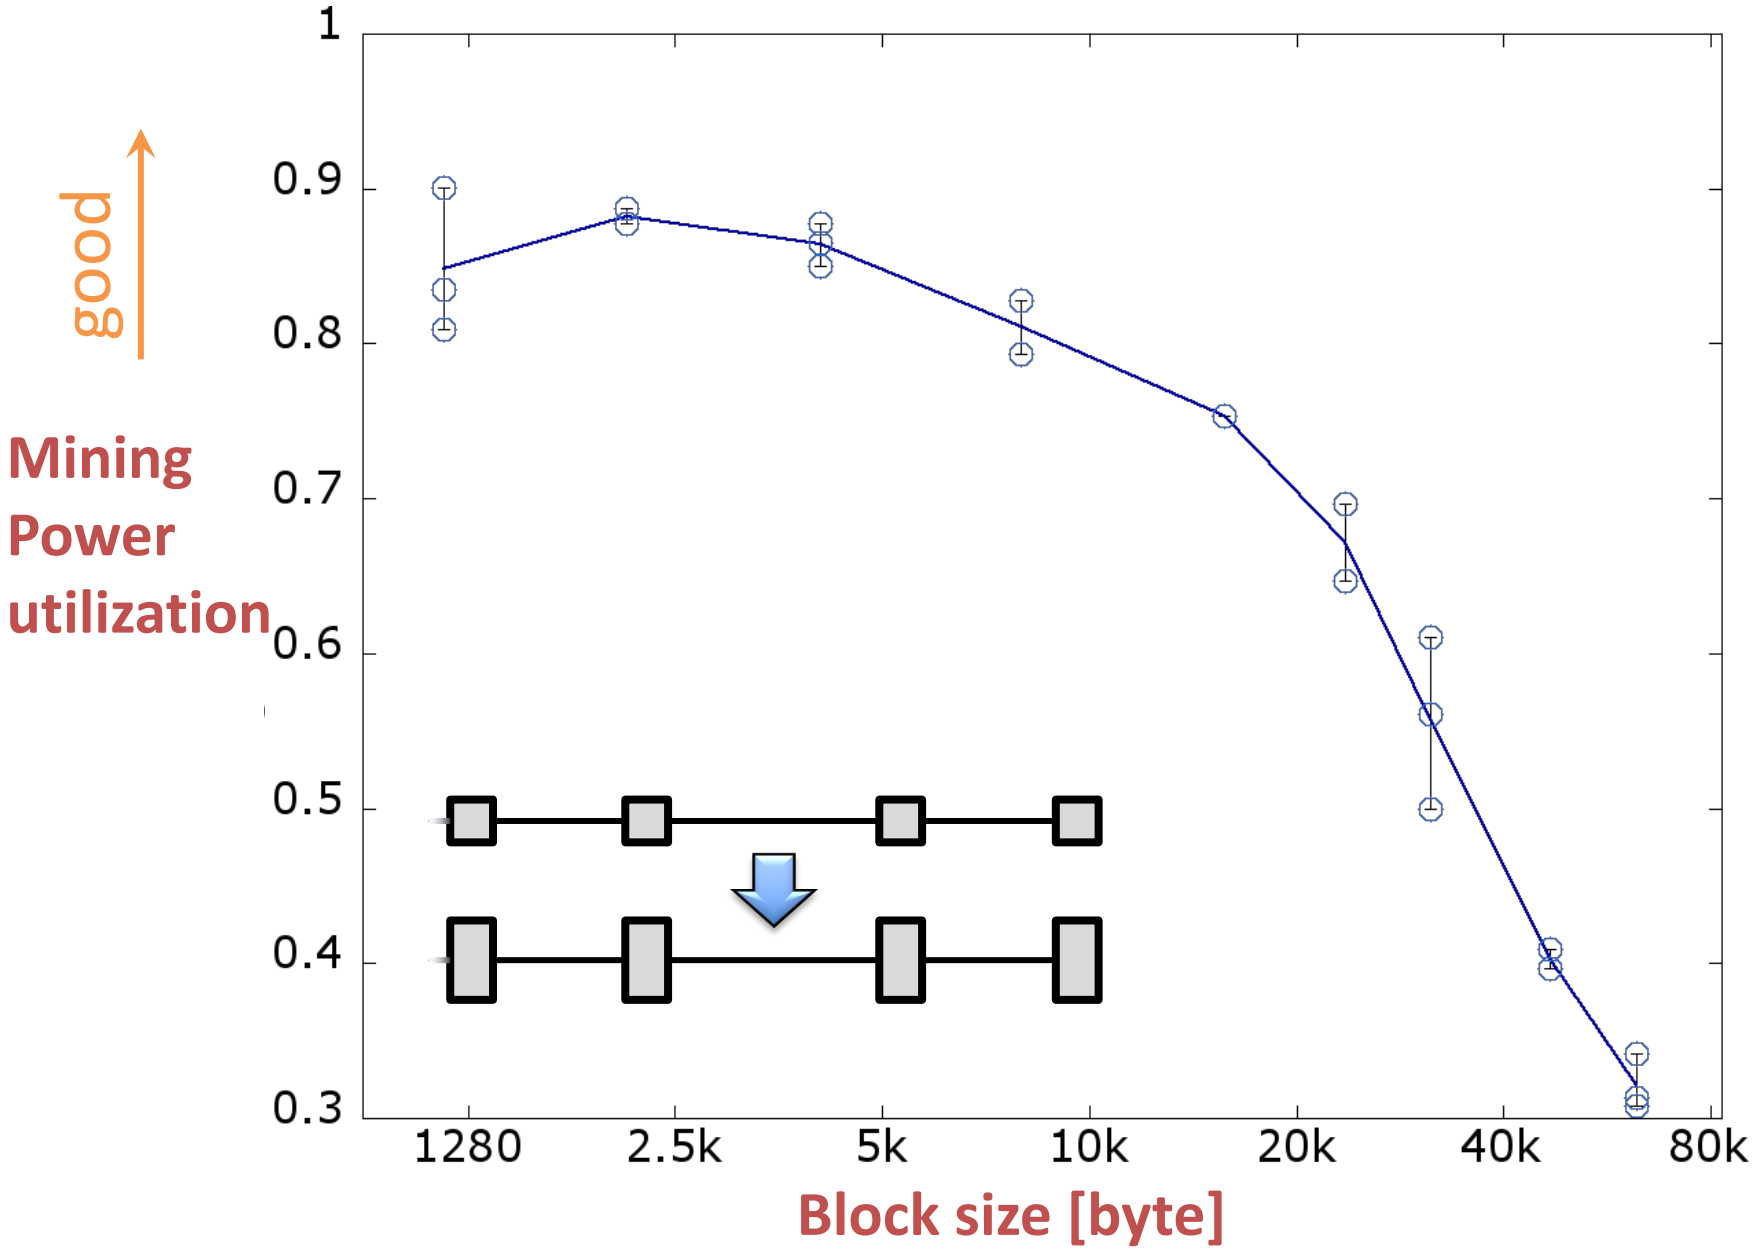
good
Mining Power utilization
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
1280
2.5k
5k
10k
20k
40k
80k
Block size [byte]

Fairness

Known Miner Sizes

[blockchain.info, April 2015]

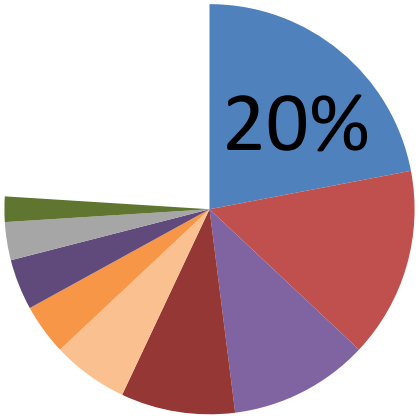

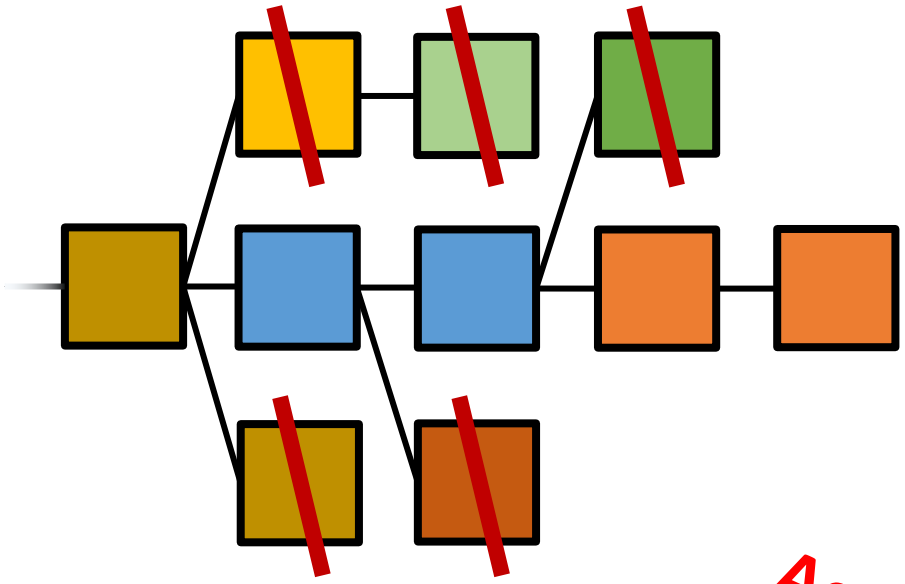

Presence:

Fair

$$\frac{\sum_{all} \neg \blacksquare}{\sum_{all}} = 80\% \qquad \square$$

Actual

$$\frac{\sum_{main} \neg \blacksquare}{\sum_{main}} = 60\% \qquad \square$$

→ tendency towards centralization

Permissionless vs Permissioned Blockchains

	Permissionless	Permissioned
Approach	Competitive	Cooperative
Basic technique	Proof-of-Resource	Byzantine Consensus
Trust requirements	Crypto (+ peers…)	Peers (+ crypto…)
Membership	Open	Closed
Energy-efficiency	Often terrible	Excellent
Transaction rate	At best hundreds / sec	Many thousands per second
Txn latency	As high as many minutes	Less than a second