The Power of Random Symmetry-Breaking in Nakamoto Consensus

Lili Su (Northeastern U.), Quanquan C. Liu (MIT),
and Neha Narula (MIT)

DISC 2021

Nakamoto Consensus

Motivations:

- A cryptocurrency needs a ledger to record transactions and to trace the ownership of a coin
- Decentralization: Each maintains a local copy of an append-only ledger

consensus problem

Protocol

In each round r, node i:

1) updates its local chain to be one of the longest chain it accessed;
2) successfully mines a block with probability p;
3) extends its local chain with this mined block;
4) “broadcasts” updated local chain to others;

Nakamoto Consensus

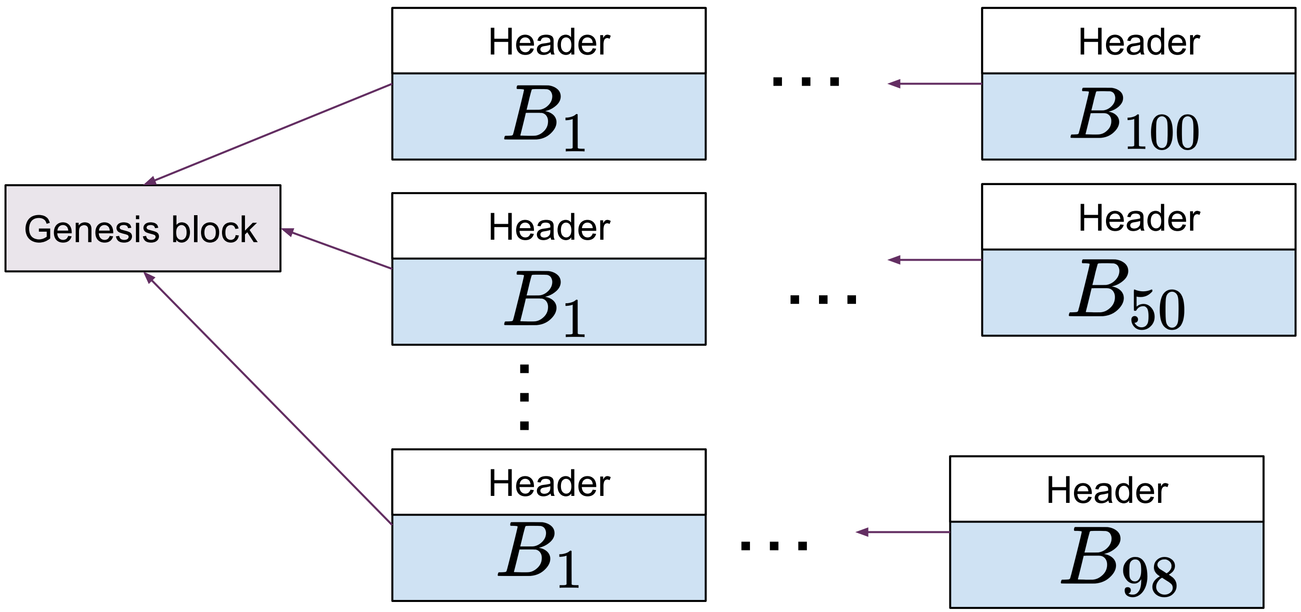

Nakamoto Consensus

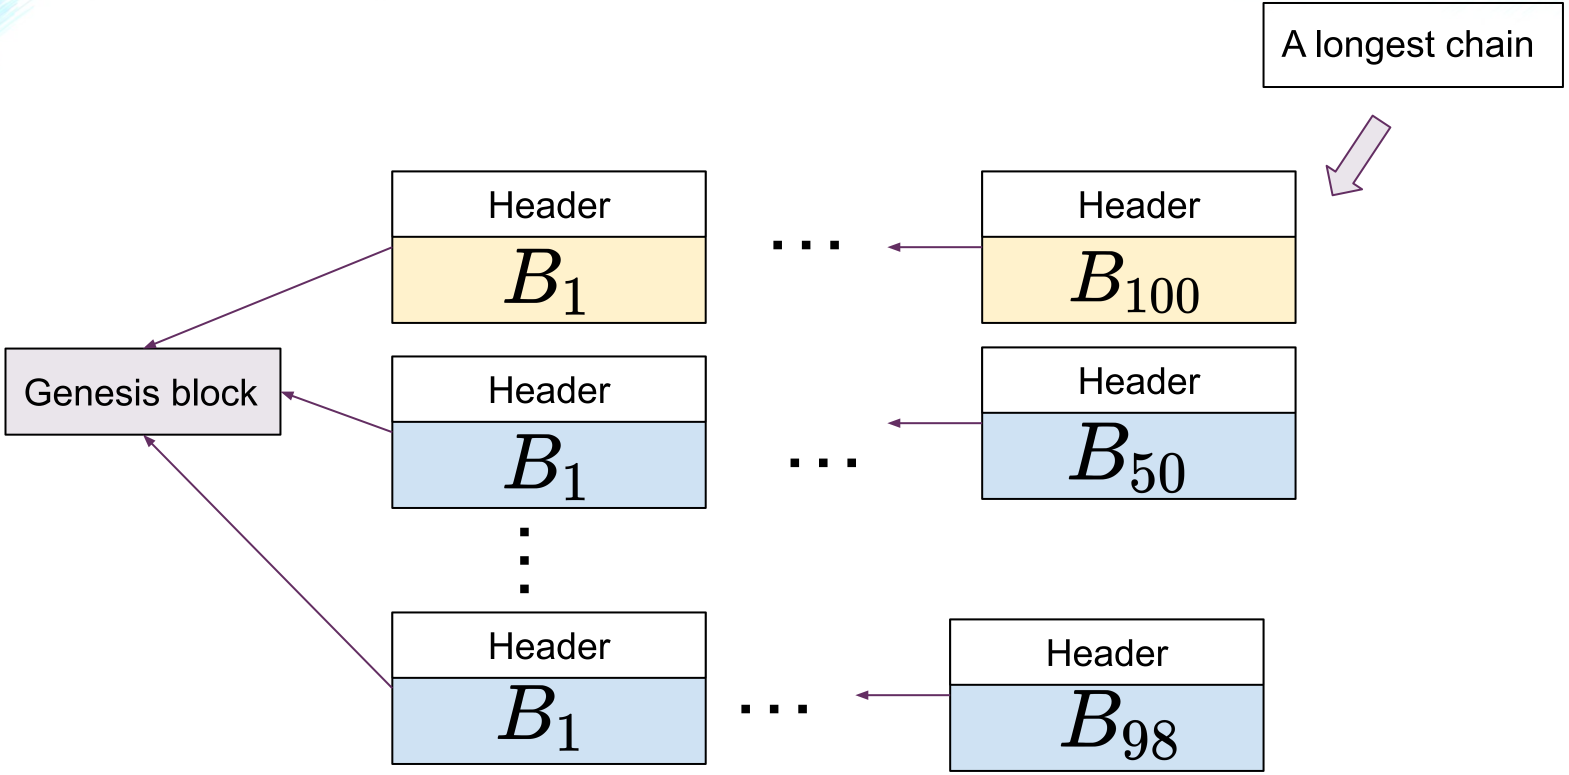

Nakamoto Consensus

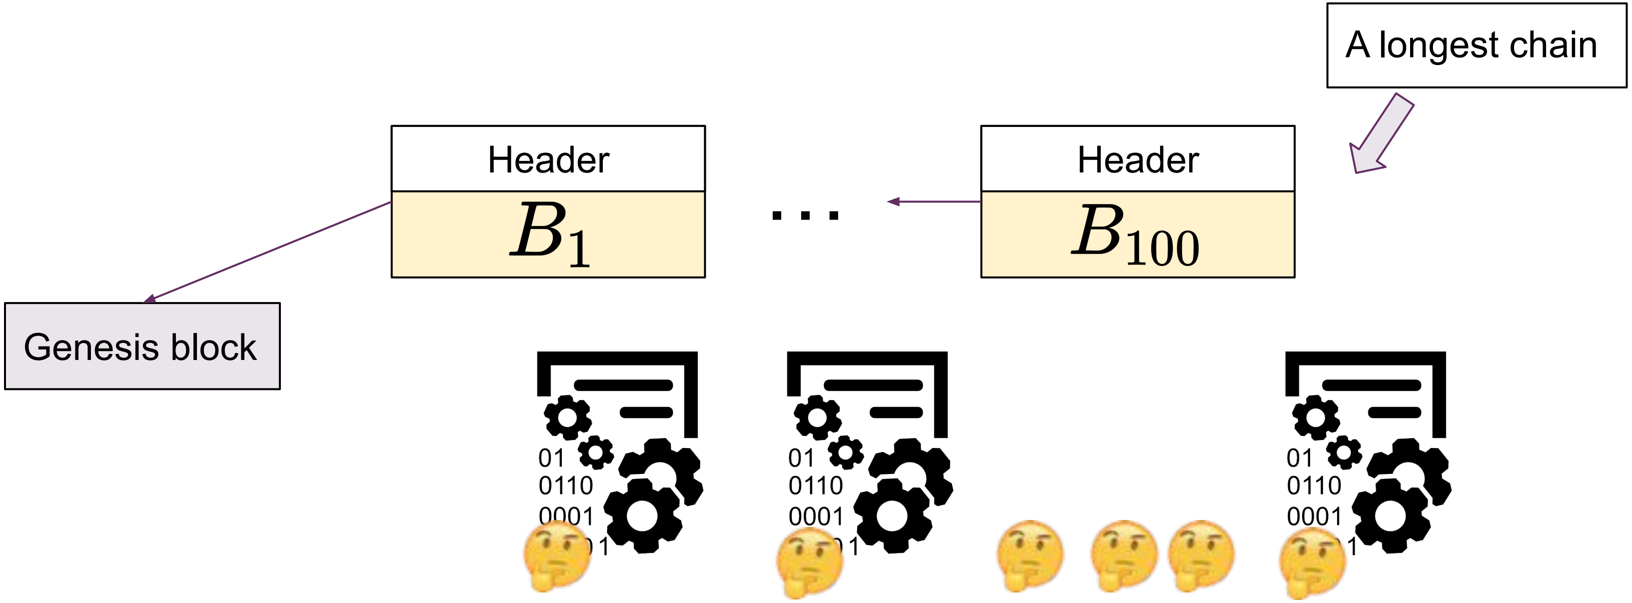

Miners extend longest chains by solving puzzles

Nakamoto Consensus

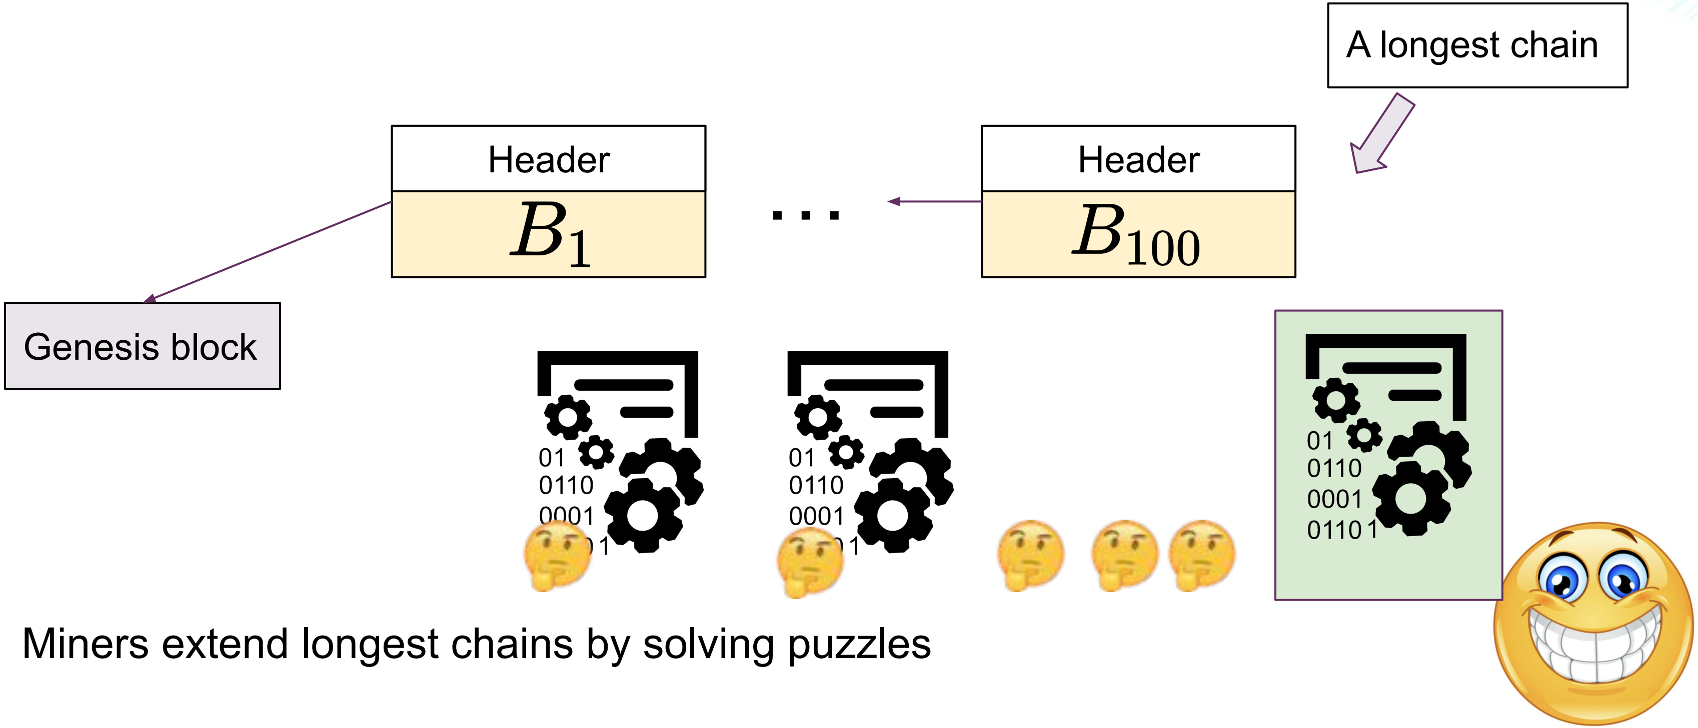

Miners extend longest chains by solving puzzles

Nakamoto Consensus

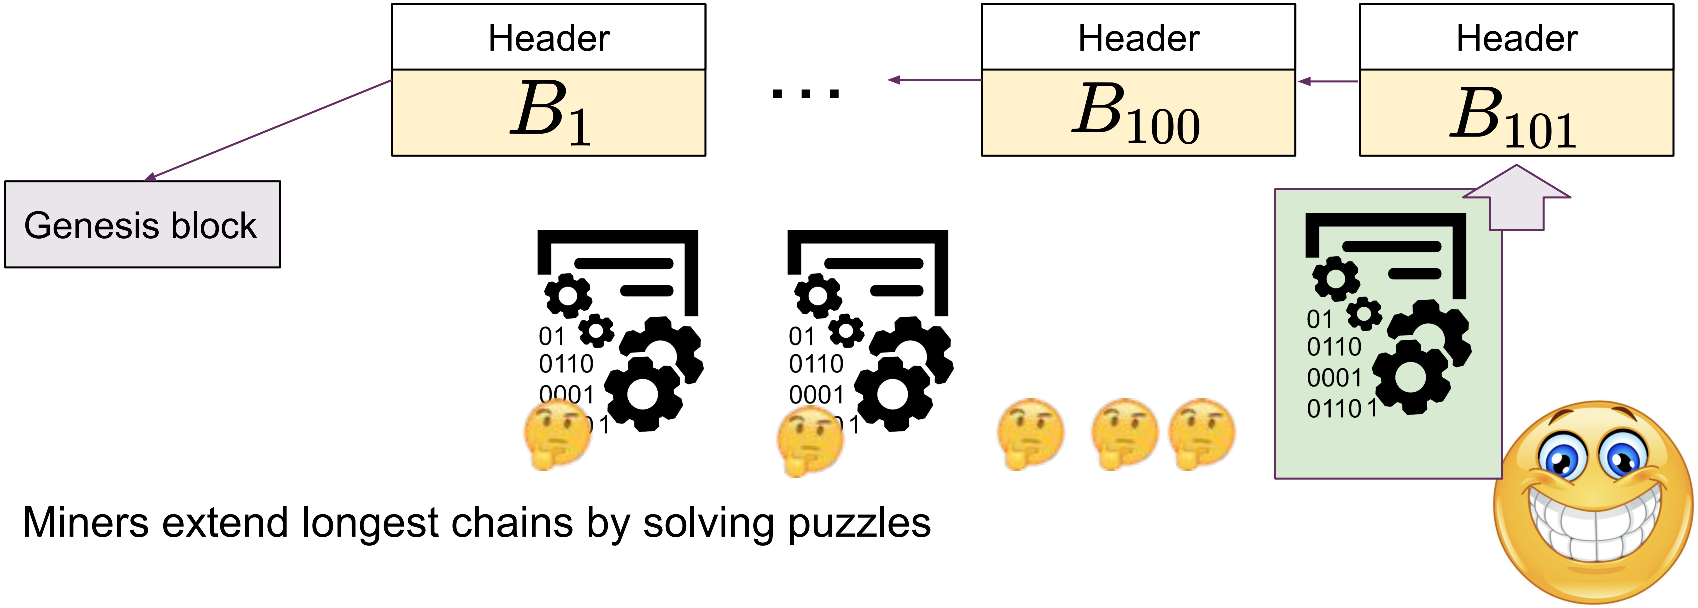

Miners extend longest chains by solving puzzles

Nakamoto Consensus

- Prevents Sybil attacks using **proofs-of-work**

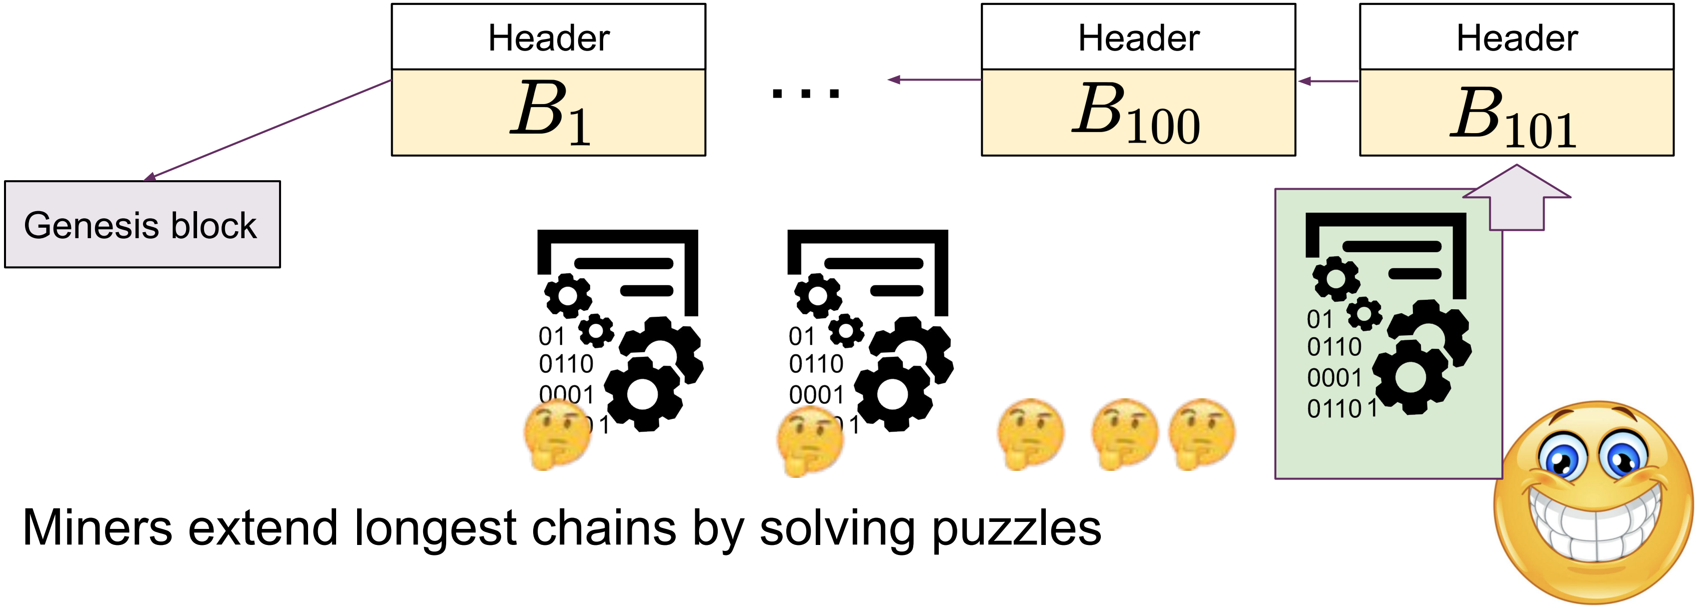

Miners extend longest chains by solving puzzles

Main Challenges

- An adversarial miner might not necessarily extend on a longest chain
-

Main Challenges

- An adversarial miner not necessarily extends on a longest chain
- Multiple longest chains

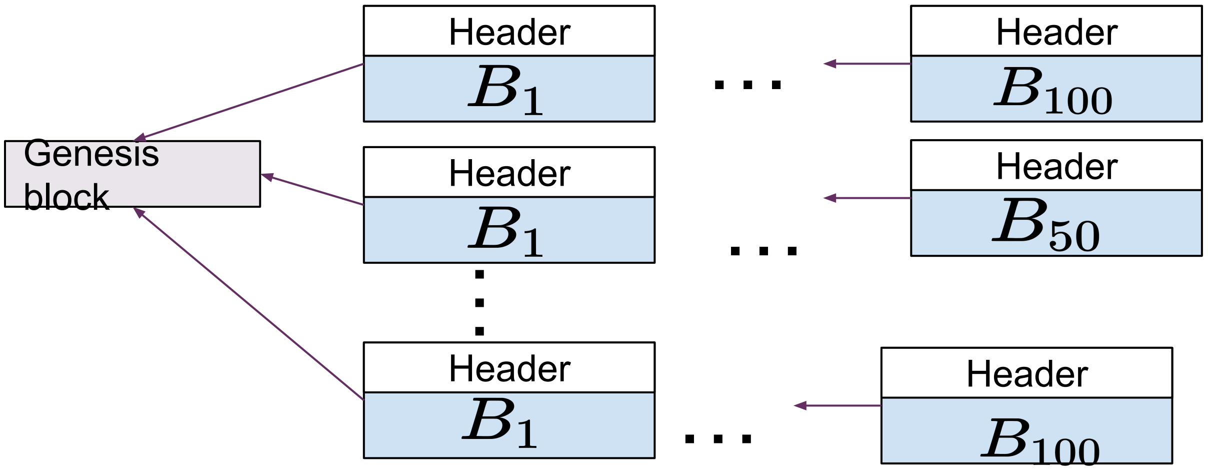

Main Challenges

- An adversarial min
  longest chain
- Multiple longest chains

(1) Puzzles are too easy?

(2) Adversarial attacks

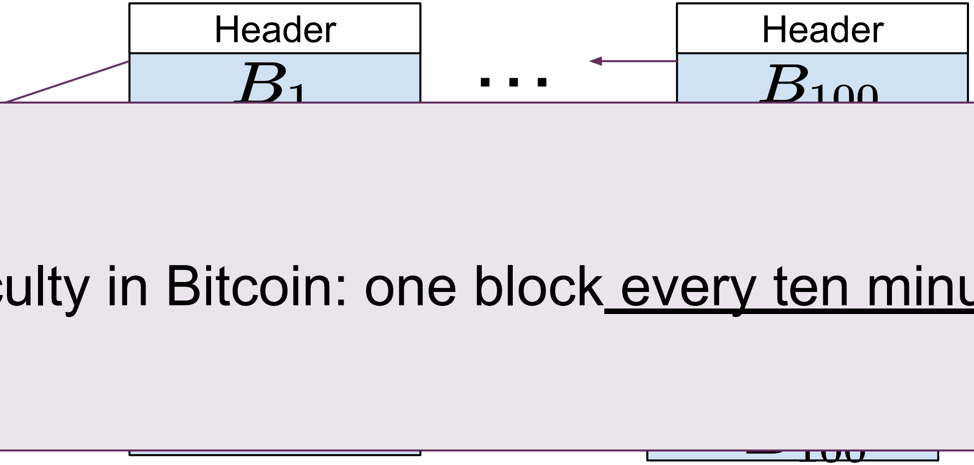

Mining difficulty in Bitcoin: one block <u>every ten minutes</u>.

Limitations of Existing Work

- Showed common-prefix and chain growth **<u>when</u>** the puzzle difficulty very high **[GKL, 2015] [PSS, 2017]**

The honest majority assumption in [GKL, 2015] implies that

$$p \leq \frac{n-2b}{2(n-b)^2}$$

- ➢When n-2b=O(1), p =O(1/n^2);
- ➢When b=0, p = O(1/n)

- p: the probability that any miner will solve the puzzle in a given round
- n: the number of active miners;
- b: the upper bound of the adversarial miners;

Limitations of Existing Work (cont.)

- Common belief is that easy puzzles fundamentally constrain chain growth, even in the absence of an adversary, due to the potential of increased forking.

Thus, should be avoid in practice

- Another common conjecture [GKL, 2015] is that the choice of symmetry-breaking strategies is not relevant to correctness.

In this paper, we revisit these two beliefs and exam their correctness

Our Contributions

- **Insights:** In the absence of adversary, the forking caused by large p *itself* does not prevent chain growth if we break symmetry uniform-at-random*

  (--* choosing among chains of equal length randomly)

- **Analysis:**
  - •Analyze Nakamoto consensus under a wide range of p including the existing well-studied region
  - • Introduce a new analysis method:(existing) quantifying # of convergence opportunities **[GKL, 2015, 2017a,b, 2020] [PSS, 2017]**

    ⇨ (ours) coupling + coalescing random walks
  - •New notion: *adversarial advantages* and *coalescing opportunities*

Protocol

Synchronous network

In each round r, node i:

1) updates its local chain to be one of the longest chain it accessed;

   1.1) If multiple exist, chooses one uniformly at random

2) successfully mines a block with probability p;
3) extends its local chain with this mined block;
4) “broadcasts” updated local chain to others;

Maximal Common Prefix and Inconsistency

- $\mathcal{C} = \{\tilde{C}_1, \cdots, \tilde{C}_m\}$: a collection of chains;
- Maximal common-prefix $P_{\mathcal{C}}$: the longest common-prefix of chains in $\mathcal{C}$
- Maximal inconsistency $I_{\mathcal{C}}$: $I_{\mathcal{C}} = \max_{1 \leq i \leq m} |\tilde{C}_i - P_{\mathcal{C}}|$

Simple generalization of [GKL, 2015] [PSS, 2017]

➢ $\tilde{C}_i - P_{\mathcal{C}}$ is the sub-chain after removing $P_{\mathcal{C}}$

➢ |. |: the length of a chain

Maximal inconsistency: the length of the longest fork

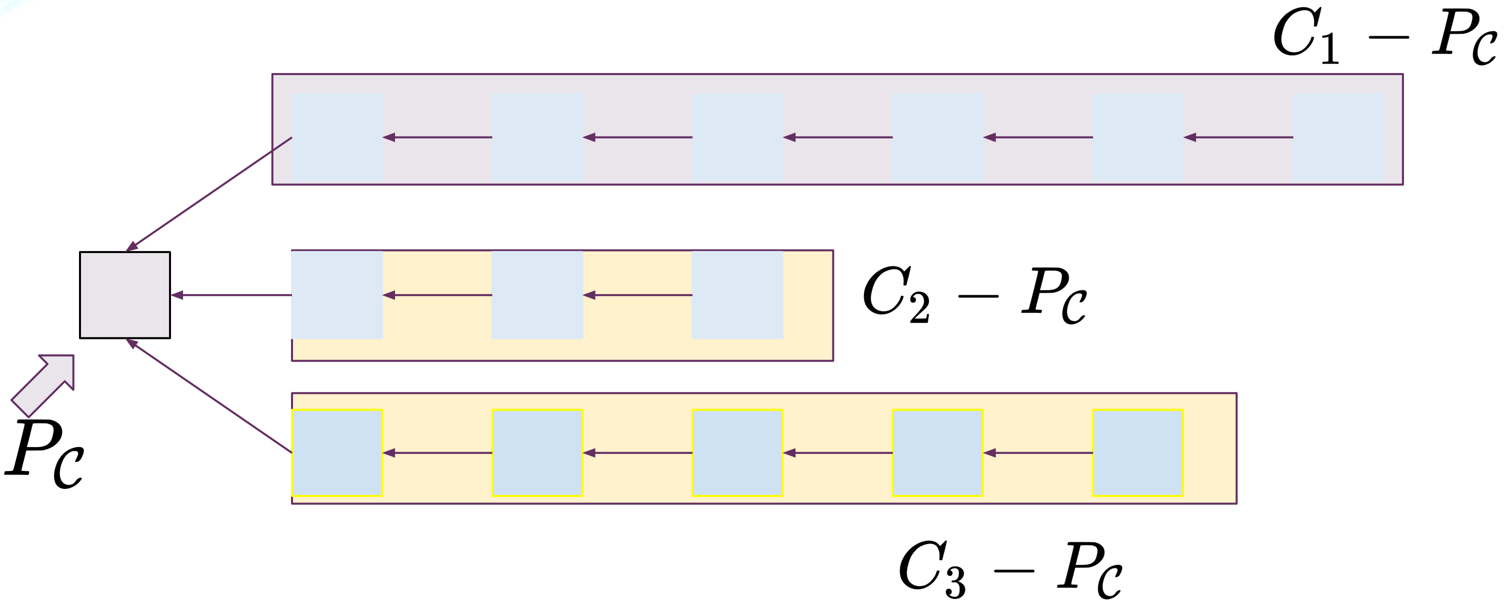

p=1, b=0: Theorem

Theorem 1: Suppose that p = 1 and b = 0. Then for any given round index t ≥ 1, in expectation, the local chains share a common prefix of length t + 1 − O(n).

``Expected'' chain length

Expected maximal inconsistency

Remarks:

- Expectation: is w. r. t. the randomness in the symmetry breaking strategy.
- Large p indeed boosts the growth of the common prefix;
- Though temporal forking exists, such forking can be quickly resolved by repetitive symmetry-breaking across rounds.

Build up connection of coalescing random walk and maximal inconsistency

General p < 1: Adversary-Free Theorem

Theorem 2: Suppose that $np = \Omega(1)$. If $p < (4 \ln 2)/n$, in expectation, at the end of round t, the length of a common prefix is

$$(1 + (1 - (1-p)^n)t) - O(1/npe^{-np}).$$

If $p \geq 4 \ln 2 / n$, in expectation, at the end of round t, the local chains at the nodes share a common prefix of length

$$(1 + (1 - (1-p)^n)t) - O\left(\frac{2np}{(1-2exp(-13np))}\right).$$

Expected chain length

Expected maximal inconsistency

Remarks:

- Maximum prefix growth rate in terms of t. Second term is maximal inconsistency
- Maximal inconsistency is independent of t

General p: Adversary-Prone

Assumption: In each round, a chain can be extended by at most 1 block.

Can be ensured via new *VDF-based scheme*.

General p: Adversary-Prone

Theorem 3: For any given *t* ≥ *1* and $M \geq \frac{4}{\beta(p_{+1}-p_{-1})}$ where $\beta = \frac{(n-b)p}{2(3np)^2}$,

at the end of round *t*, with probability at least

$$1 - \exp(-\frac{(p^*)^2 M}{2}) - \exp(-\frac{(p_{+1}-p_{-1})^2 M}{16p^*})$$
$$- \frac{2}{\beta}\exp(-\frac{1}{2}(n-b))$$

the expected maximal inconsistency among a given pair of honest nodes is < M

$$p^* = p_{-1} + p_{+1}$$

p_{+1} : the probability at in a round only honest miners found block;

p_{-1} : the probability at in a round only adversarial miners found block;

Conclusion & Open Questions

- Showed convergence opportunities *not necessary* to make chain progress

- **Open:** Providing a scheme that is not based on VDFs for removing assumption in general *p*, adversary-prone case
- **Open:** Explicit trade-off of system parameters n, b, p, etc
- **Open:** investigating Nakamoto consensus with more complex symmetry-breaking strategies

Model and Definitions [GKL, 2015] [PSS, 2017]

- Synchronous network
- All Byzantine nodes are controlled by a probabilistic polynomial time (PPT) adversary $\mathcal{A}$;

Bounded computation power

- At any time, $\mathcal{A}$ can corrupt up to b nodes;
- A corrupted node remains corrupted;

Warmup: p=1 and b=0

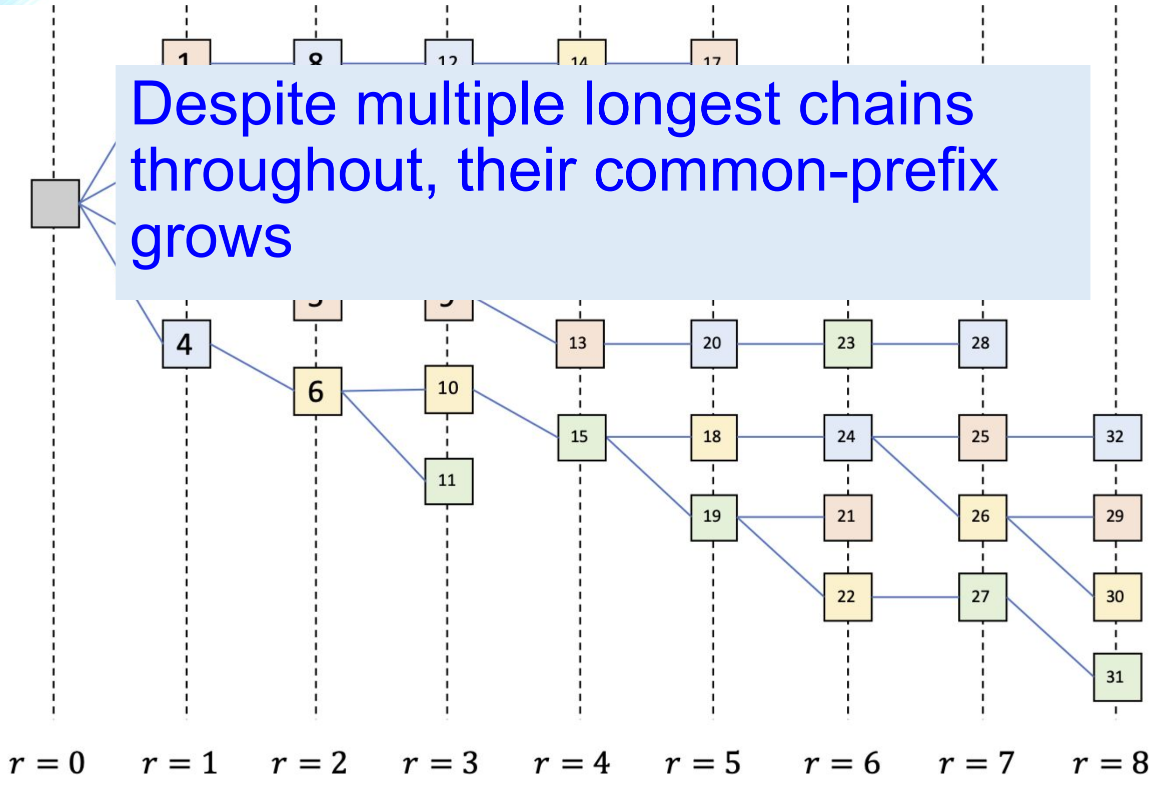

Despite multiple longest chains throughout, their common-prefix grows

Illustrating example: n=4, p=1, b=0

- Each color represents a different miner;
- As p=1, every miner mines a block in each round;
- At the beginning of each round, there are four longest chains;
- Each miner chooses one chain to extend uniformly at random.

Coalescing Random Walks

- Given a undirected graph;
- Given a set of particles;
- Each particle independent random walks until they meet;
- Whenever two or more particles meet, they unite to form a single particle, then continues the random walk.

Particles on vertices of an undirected graph G = (V, E);

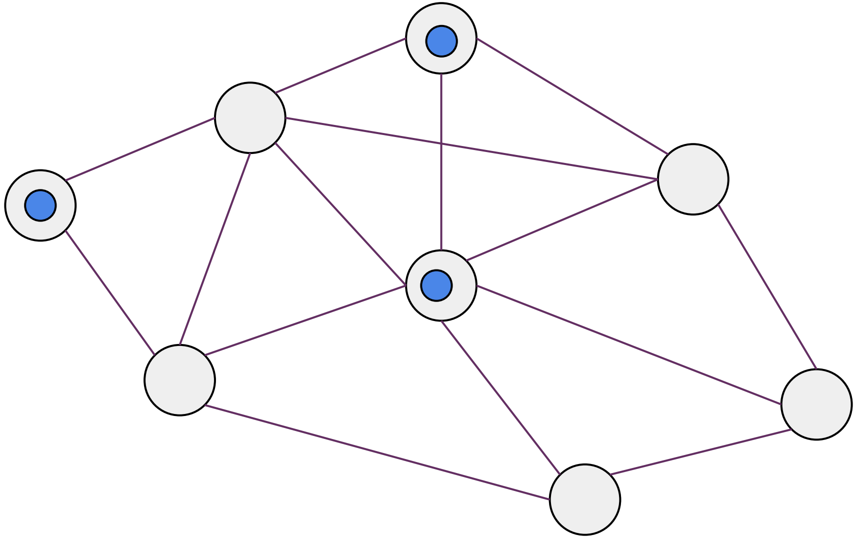

Coalescing Random Walks

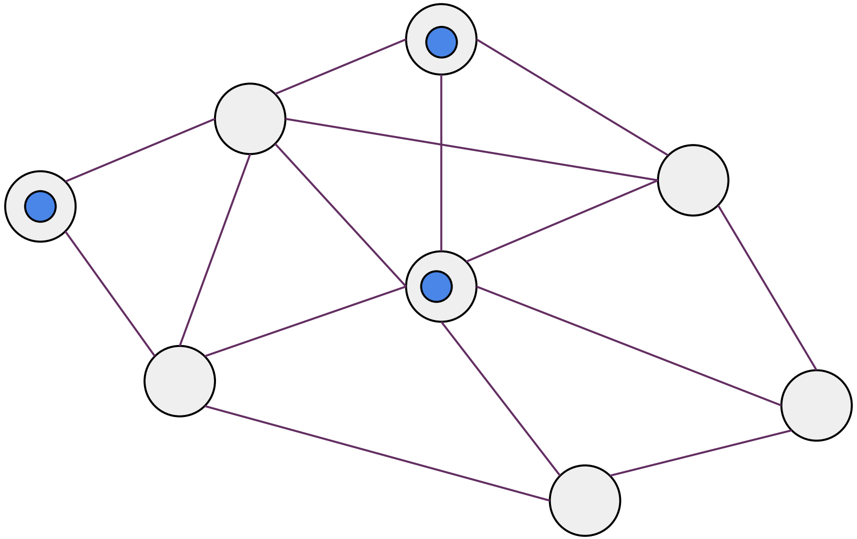

Particles on vertices of an undirected graph

Coalescing Random Walks

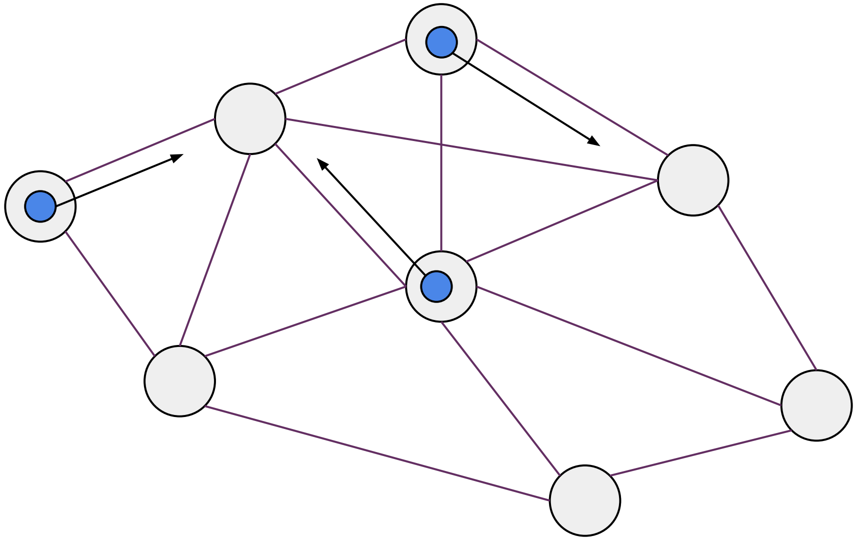

Particles perform random walk on graph

Coalescing Random Walks

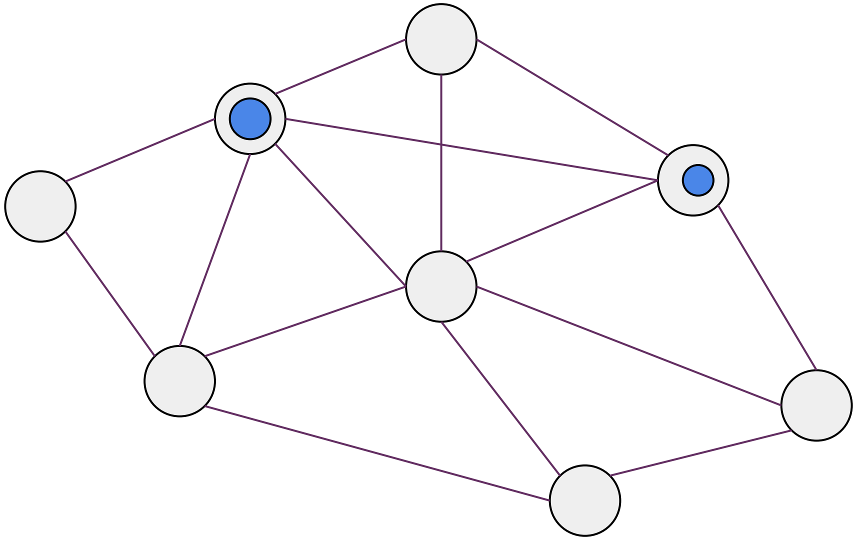

When two or more particles land on same vertex, they merge

Coalescing Random Walks

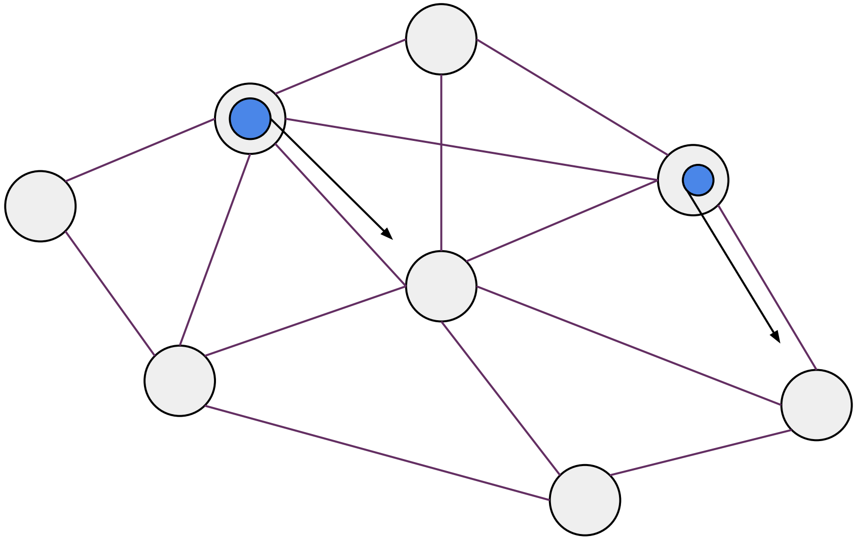

Continue performing random walks

Coalescing Random Walks

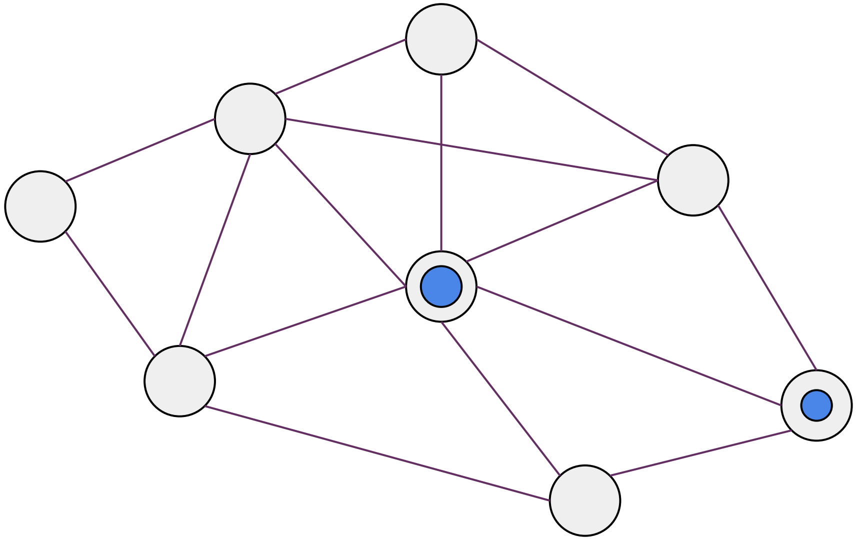

Continue performing random walks

Coalescing Random Walks

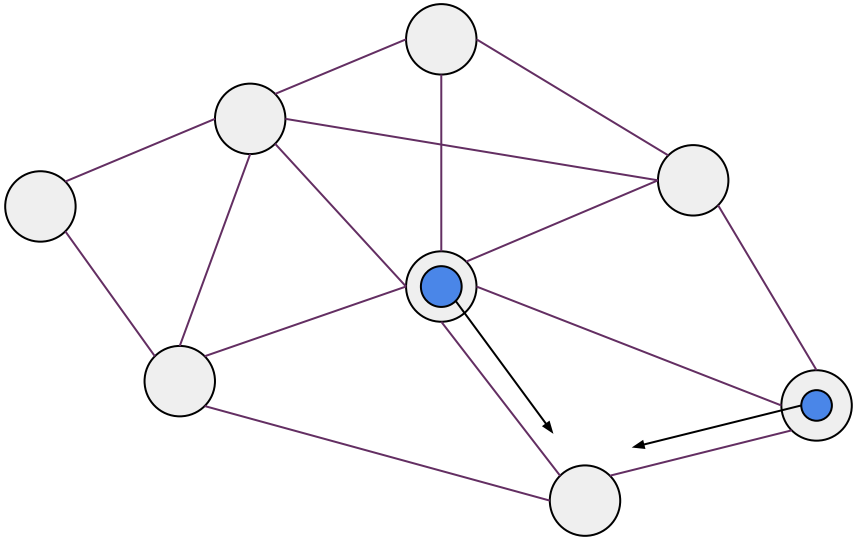

Continue performing random walks

Coalescing Random Walks

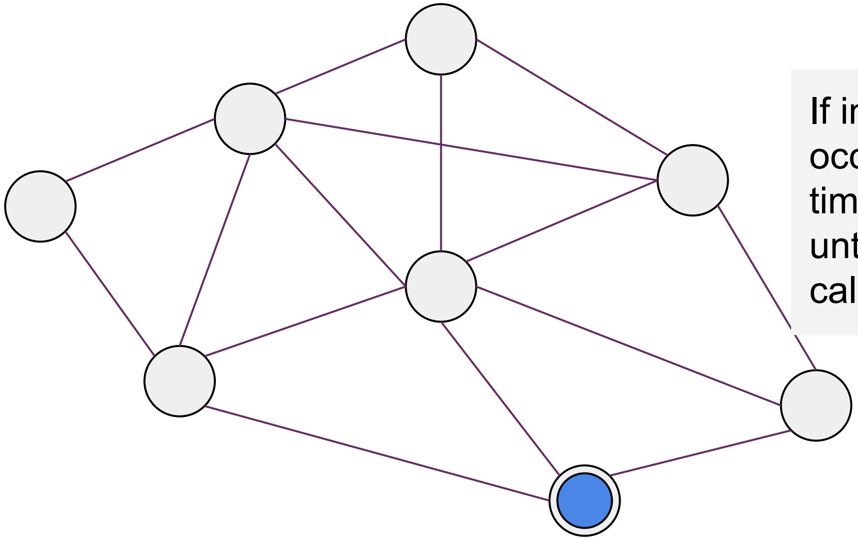

If initially every vertex is occupied with a particle, the time takes
until all particles merge is called ***coalescing time***

Illustrating example (n=4, b=0, p=1)

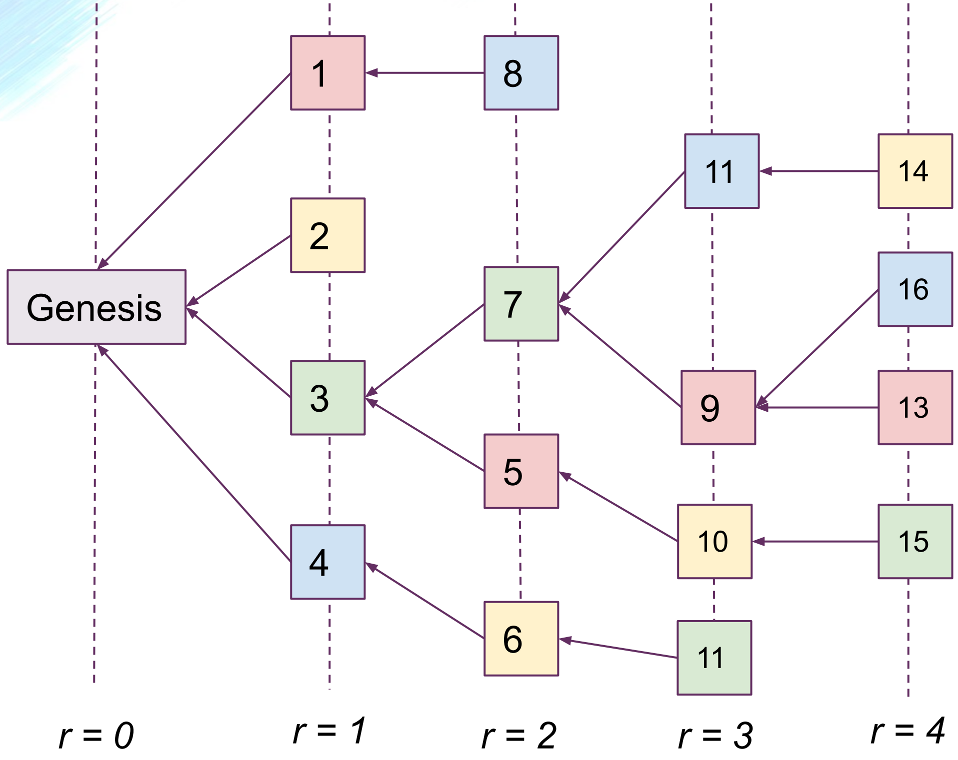

- Each color represents a different miner;

Illustrating example (n=4, b=0, p=1)

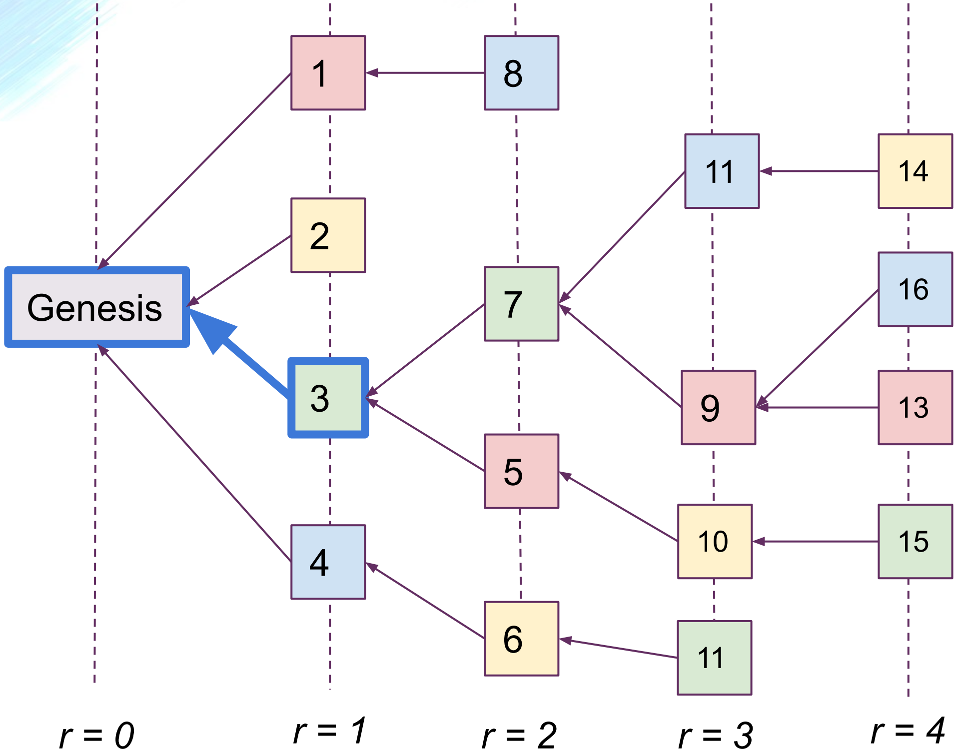

All longest chains have Genesis and Block 3 as common prefix

Illustrating example (n=4, b=0, p=1)

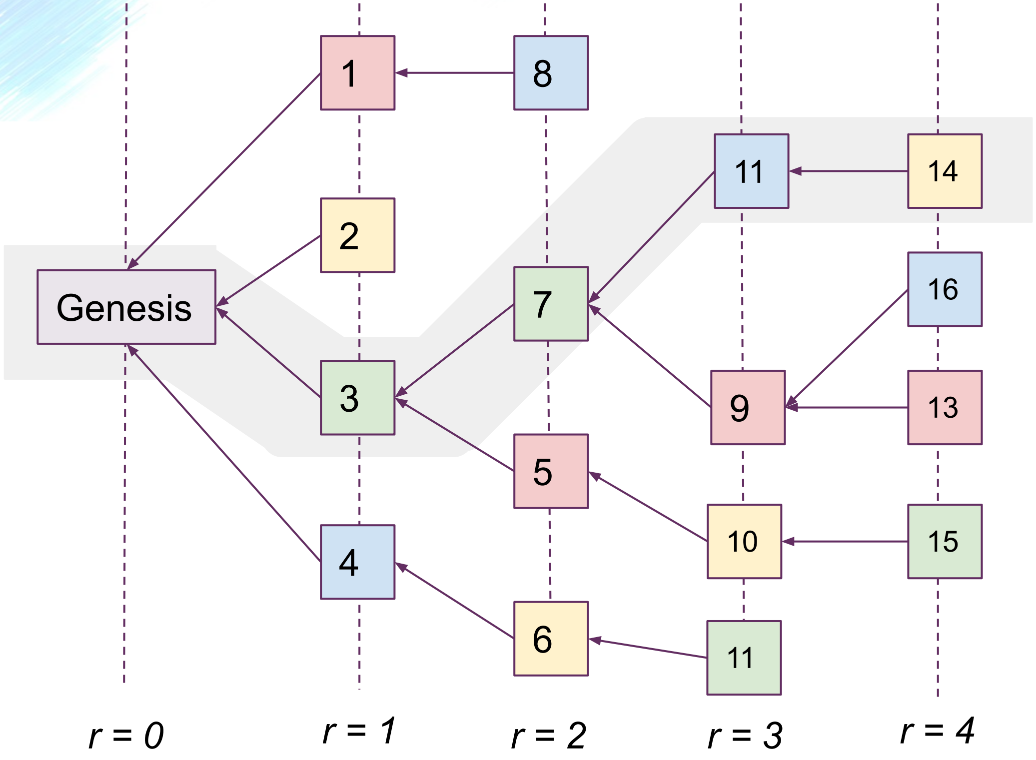

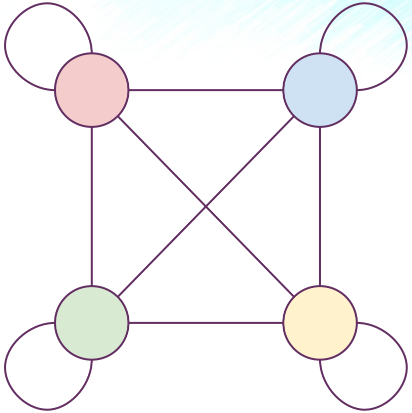

- Each backward chain modeled as random walk on complete graph (with self-loops) with number of vertices equal to number of miners

Illustrating example (n=4, b=0, p=1)

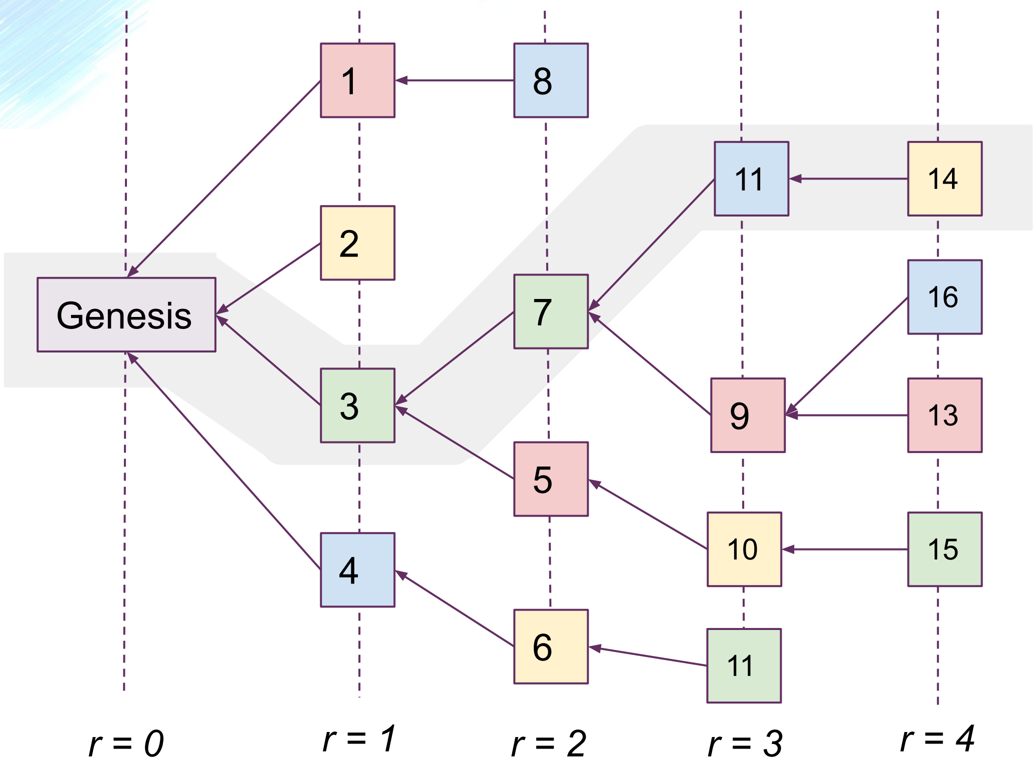

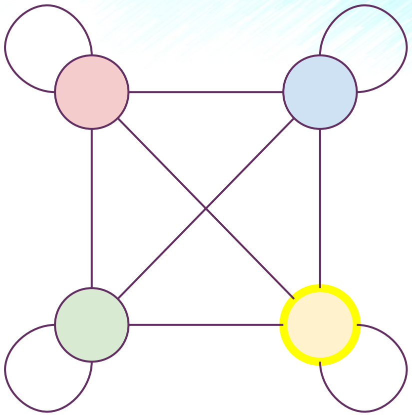

- Visits (yellow, 14), (blue, 11), (green, 7), and then (green, 3)

Illustrating example (n=4, b=0, p=1)

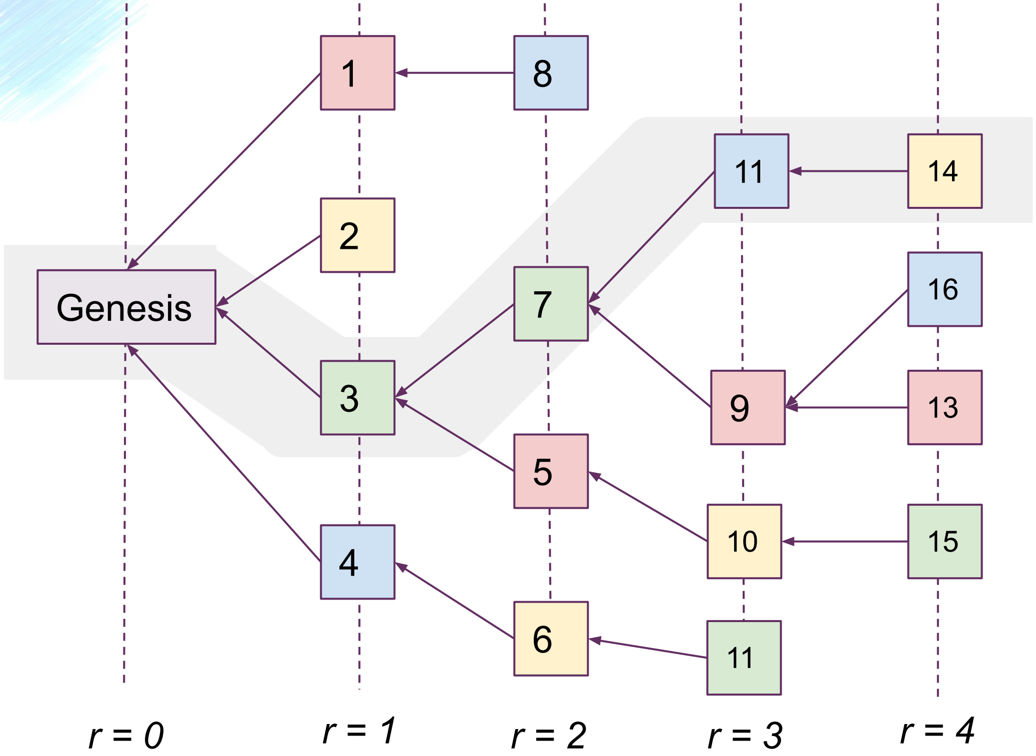

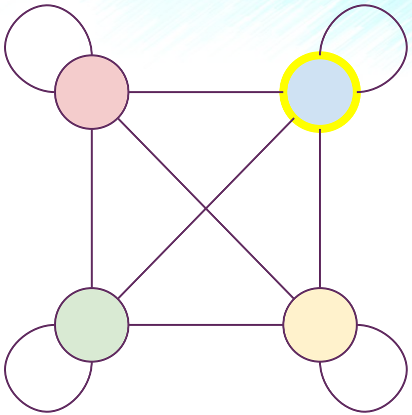

- Visits (yellow, 14), (blue, 11), (green, 7), and then (green, 3)

Illustrating example (n=4, b=0, p=1)

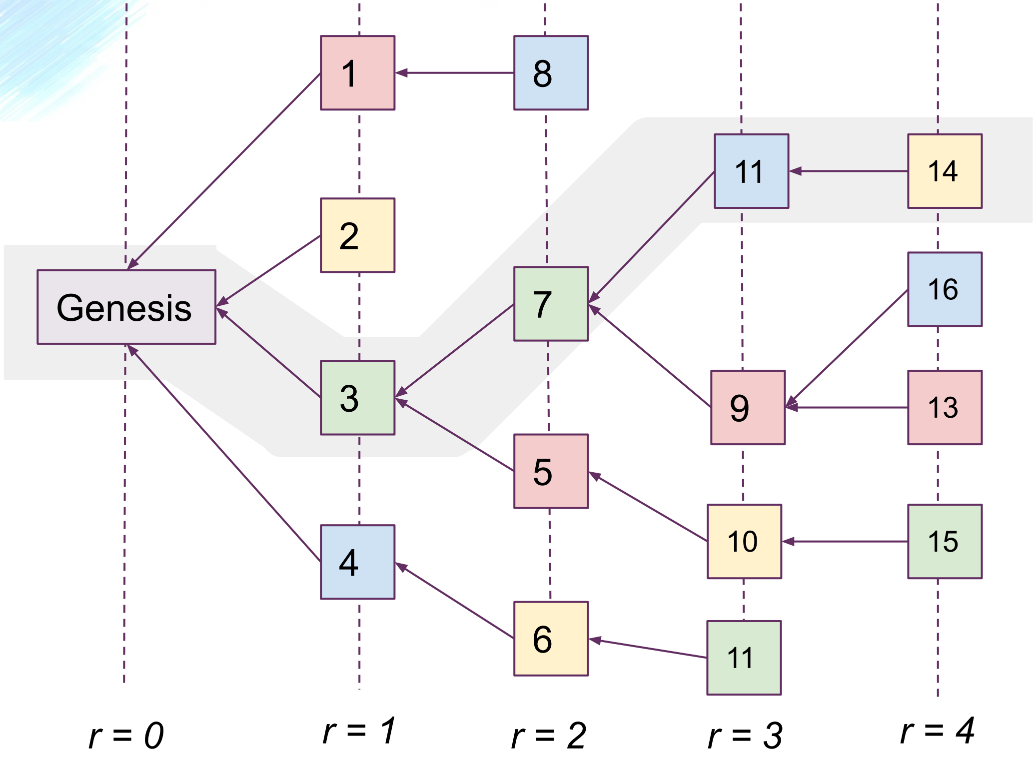

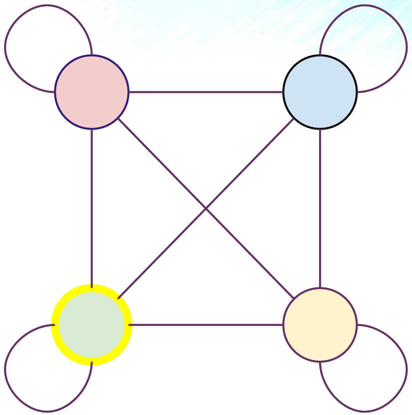

- Visits (yellow, 14), (blue, 11), (green, 7), and then (green, 3)

Illustrating example (n=4, b=0, p=1)

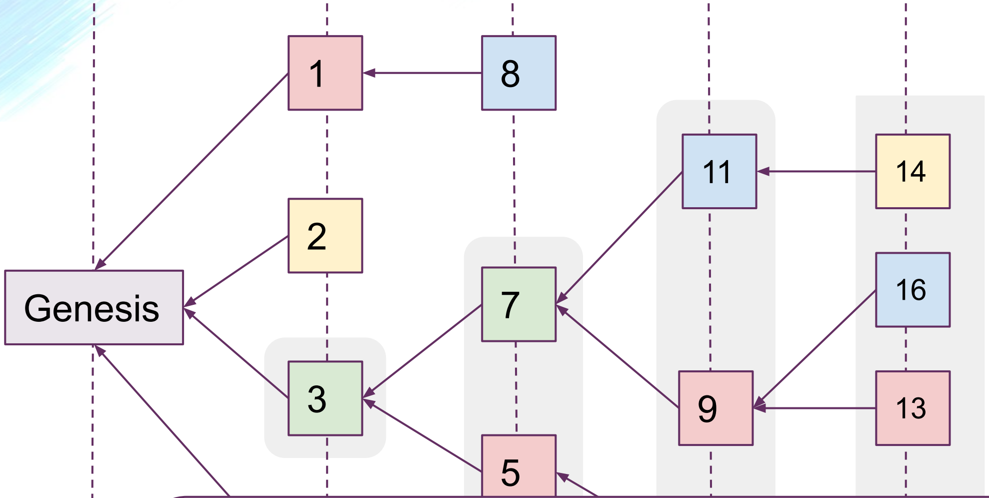

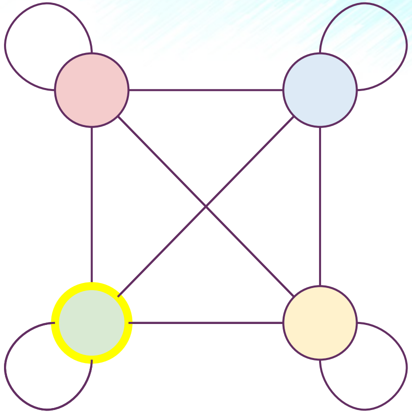

If G = (V, E) is complete, then the expected coalescing time is O(n).
[Aldous and Fill, 2002] [Cooper, Frieze, and Radzik, 2010]

- Visits (yellow, 14), (blue, 11), (green, 7), and then (green, 3)

General p < 1: Adversary-Free

Key challenges: the number of longest chains are time-varying

- **Proof Sketch:**
  - Use lazy coalescing random walk
  - No fixed correspondence between color and vertex
  - Use stochastic dominance to bound maximal inconsistency

***u*-Lazy coalescing random walk:** each step with probability *(1-u)* stay at the current vertex; probability *u* moves to an adjacent vertex, picked uniformly at random

General p: Adversary-Prone

Theorem 3: For any given *T ≥ 1* and $M \geq \frac{4}{\beta(p+1-p-1)}$ where $\beta = \frac{(n-b)p}{2(3np)^2}$,

at the end of round *T*, with probability at least

$$1 - \exp\left(-\frac{(p*)^2 M}{2}\right) - \exp\left(-\frac{(p+1-p-1)2M}{16p^*}\right) - \frac{2}{\beta}\exp\left(-\frac{1}{2}(n-b)\right)$$

over the randomness in the block mining, the expected maximal inconsistency among a given pair of honest nodes is less than *M*, where the expectation is taken over the randomness in the symmetry breaking.

Nakamoto Consensus (cont.)

Observations:

Depending on the identity of participants is vulnerable to Sybil attacks

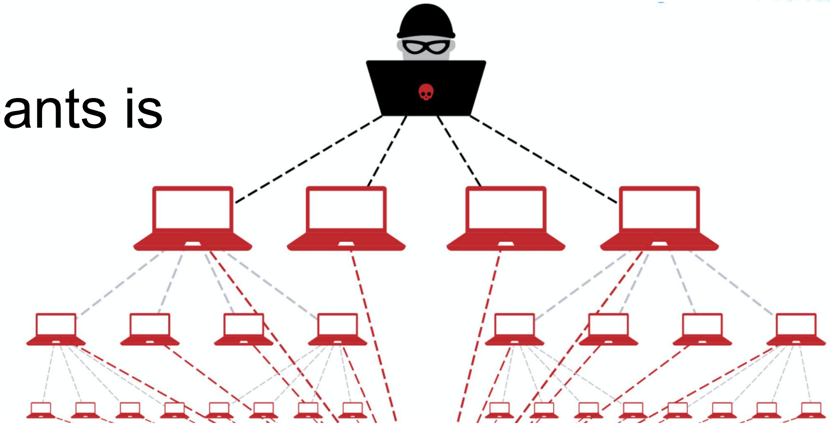

Key ideas:

incorporating <u>computational puzzles</u> + a simple <u>longest-chain</u>

(proof-of-work/mining)

(most work)

Correctness and Liveness

Characterized via three properties:

- *Common prefix*: any two honest miners share a common prefix of consecutive blocks
- *Chain-growth*: the rate at which the common-prefix grows over time
- *Chain quality*: the fraction of blocks created by the honest miners

[Garay, Kiayia, and Leonardas, 2015] [Pass, Seeman, and Shelat, 2017]

Random Symmetry-Breaking

- Among all chains of equal, longest length, randomly pick one

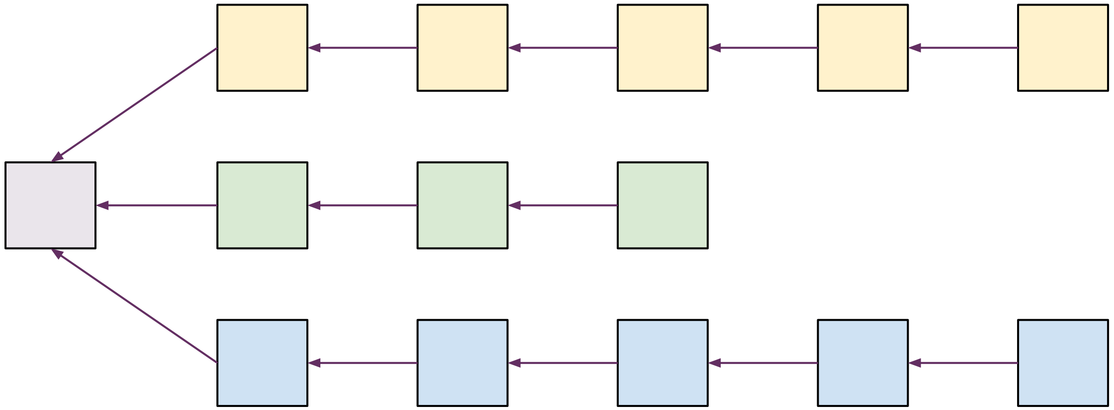

Random Symmetry-Breaking

- Among all chains of equal, longest length, randomly pick one

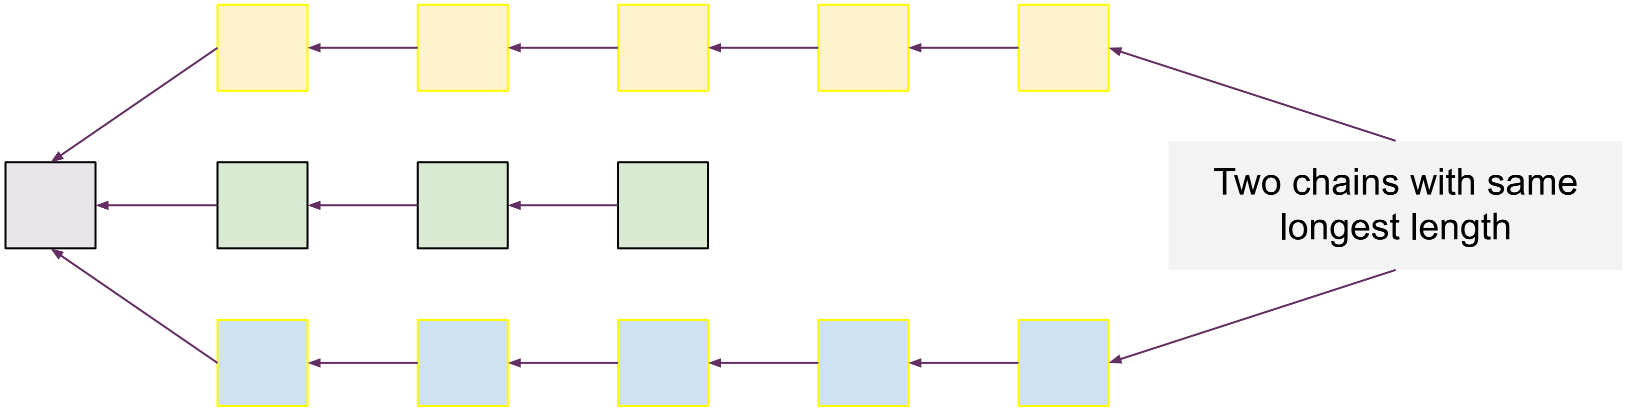

Random Symmetry-Breaking

- Among all chains of equal, longest length, randomly pick one

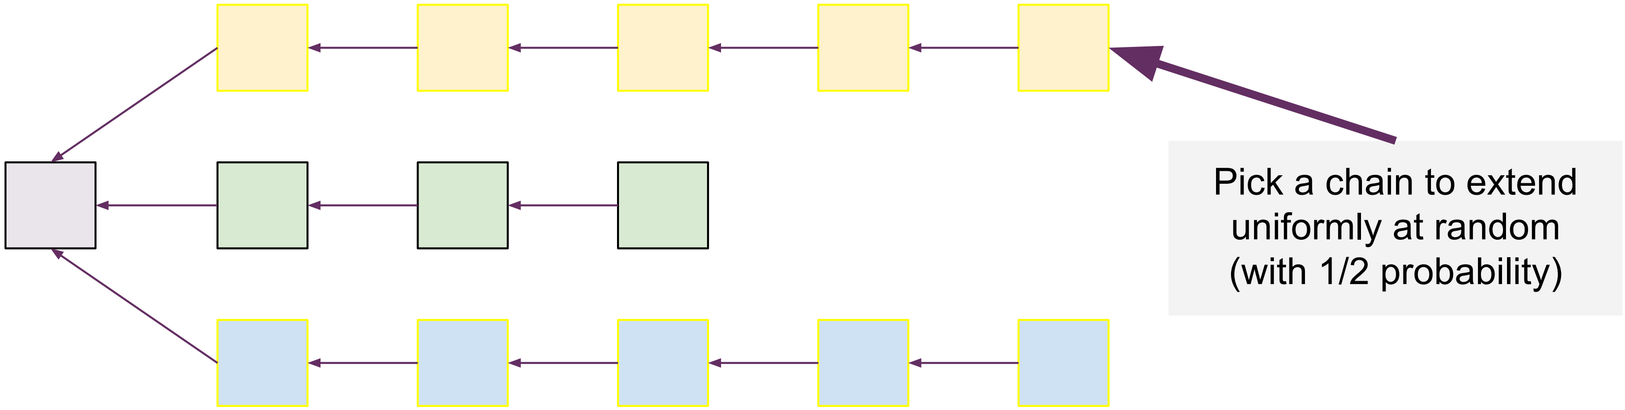

Random Symmetry-Breaking

- Among all chains of equal, longest length, randomly pick one

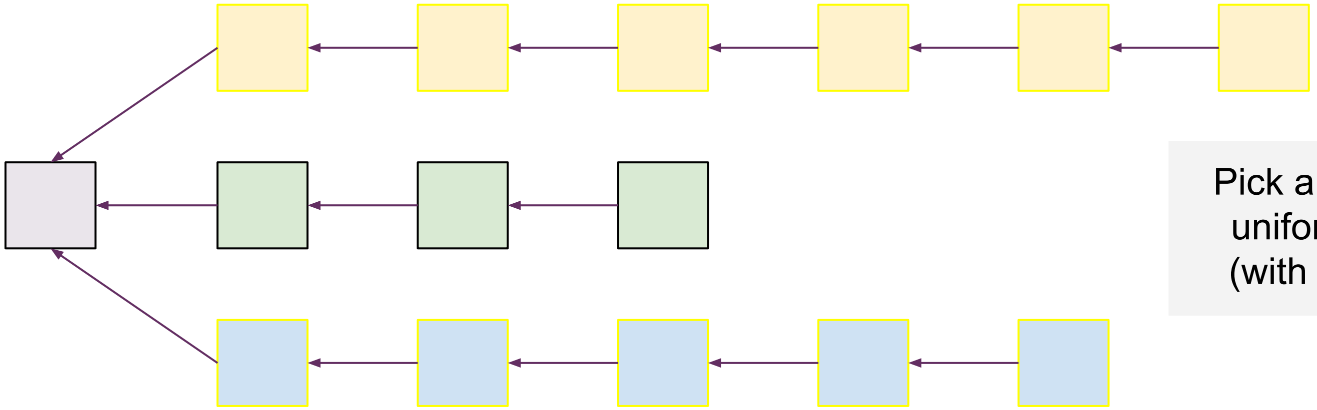

Pick a chain to extend uniformly at random (with 1/2 probability)

Model and Definitions [GKL, 2015] [PSS, 2017]

- Synchronous network: Messages are exchanged in synchronous rounds, messages sent in round r-1 will be delivered at the beginning of round r (i.e., $\Delta = 1$)

$$\Longleftrightarrow$$

$\exists$ a global clock and the time is evenly slotted into rounds

- Permissionless system:
  - ➢miners/nodes have identical computation power
  - ➢miners can join and leave at any time but the number of active miners <u>remains to be n</u>