Stellar

Fast, Deterministically-Safe Proof-of-Work Consensus

Ali Farahbakhsh*, Giuliano Losa§, Youer Pu*, Lorenzo Alvisi*, Ittay Eyal+
S&P 2026

* Cornell University
§ Stellar Development Foundation
+ Technion

Nakamoto consensus: fully permissionless, but tricky probabilistic guarantees

- Fully permissionless:
  - Parties can join or leave without any synchronization
  - No identifying information is required for participation
- Probability of block reverts depends on depth, adversarial work ratio β, and the message-delay bound
- Guarantees are (very) hard to analyze (the “6-deep” is a rule-of-thumb)

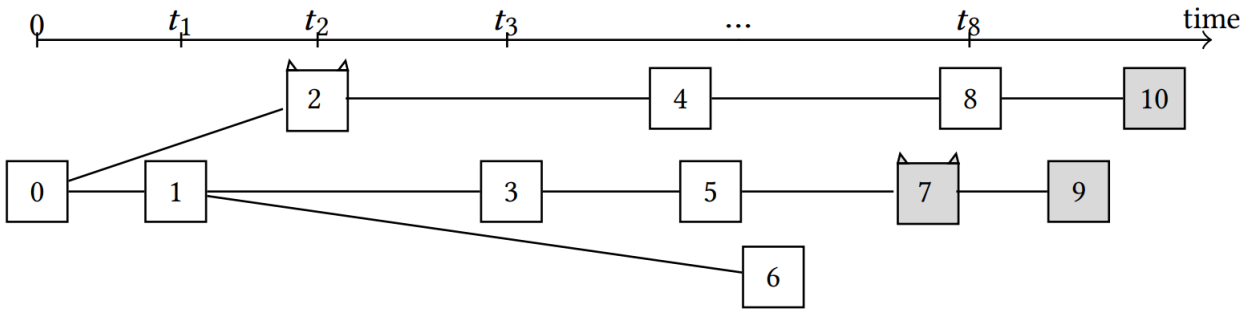

16 years after Bitcoin's birth, in *How to Beat Nakamoto in the Race,* CCS 2025, Shujie Cao and Dongning Guo write:

8 Conclusion

This paper resolves the longstanding question of block safety in PoW Nakamoto consensus protocols by identifying an optimal attack strategy, termed "bait-and-switch," and computing the precise probability of a safety violation after a given number of confirmations. Our analysis considers a baseline model where an omniscient adversary controls message propagation delays up to a maximum of Δ. Central to our approach is the formulation of a tractable discrete-time Markov decision process that concisely captures the system state and the adversary's permissible actions.

The bait-and-switch attack naturally extends to scenarios with more constrained adversaries. Specifically, when the adversary has partial control over the views of honest nodes, it can delay the propagation of honest blocks within its capabilities and release the bait once it observes a second-best branch matching the public height. Determining the precise conditions under which this generalized attack remains optimal is an open question.

Reinforcement learning offers a promising framework for identifying attack strategies in settings where the adversary has incomplete state information or reduced capabilities. However, the

Sleepy protocols: simple guarantees, dynamic participation, but registration required

- Parties must register (e.g. PoS protocols maintain a list of stakers)
- Parties can “sleep” (i.e. be *inactive*), but *active* honest participants must sufficiently outnumber *registered* adversaries
- Recent sleepy protocols like MMR have *deterministic safety* and *constant latency*, which makes them particularly simple to reason about

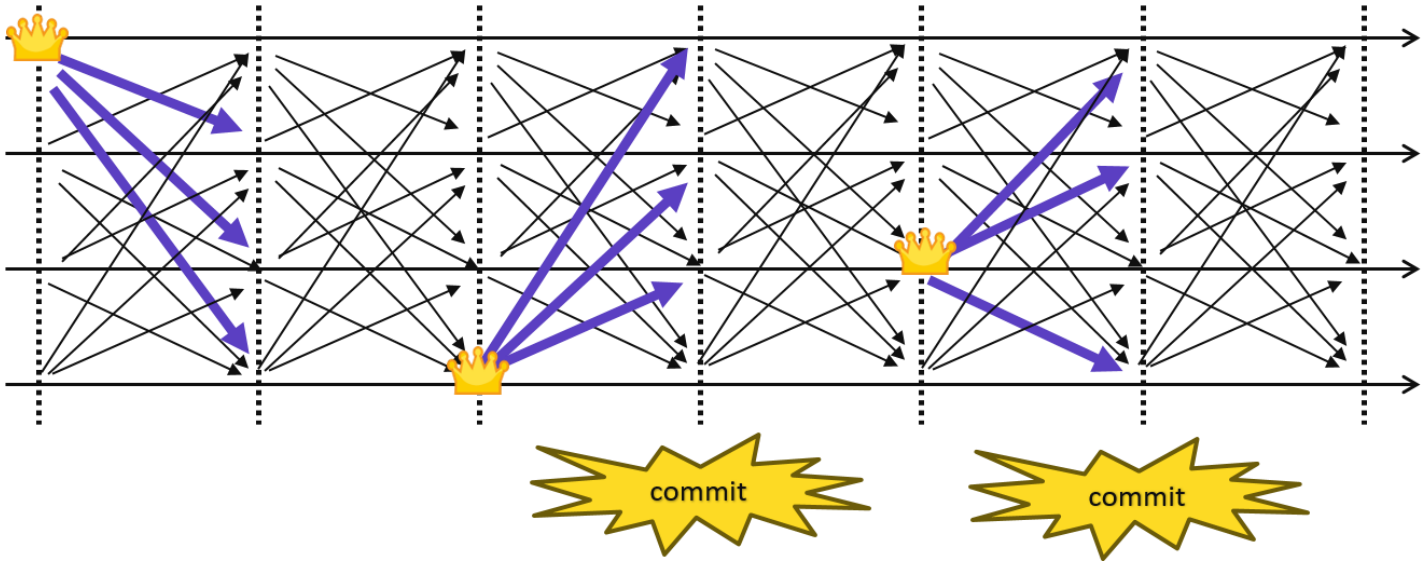

Can we have deterministic safety, constant latency, and fully-permissionless participation?

Nakamoto consensus

- Fully-permissionless
- Probabilistic safety that is hard to reason about

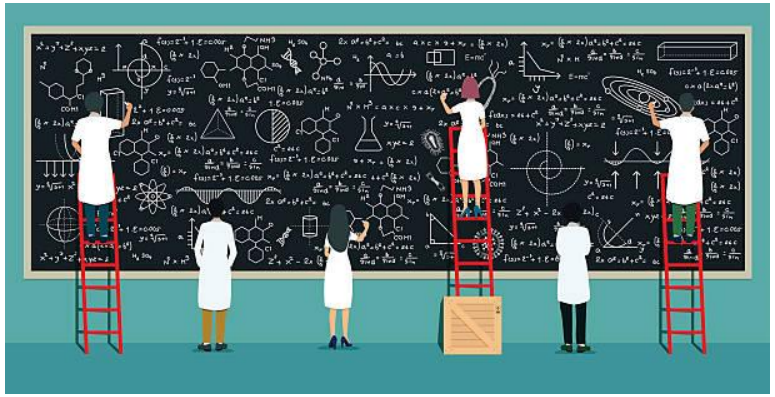

Sleepy protocols

- Allow dynamic participation, but require registration
- Can have deterministic safety and constant latency

  ⇒ simple to reason about

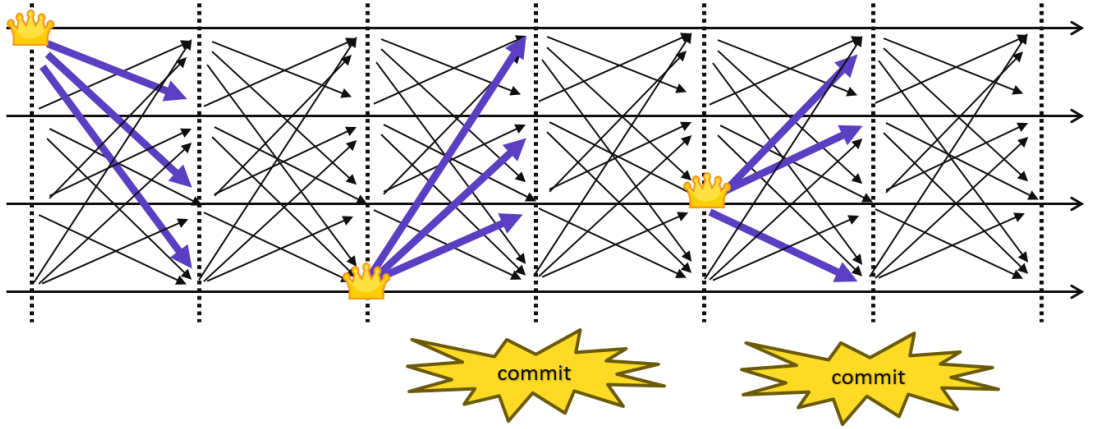

MMR solves sleepy consensus simply under a 1/3 adversary, but crucially depends on enforcing one message per party per round

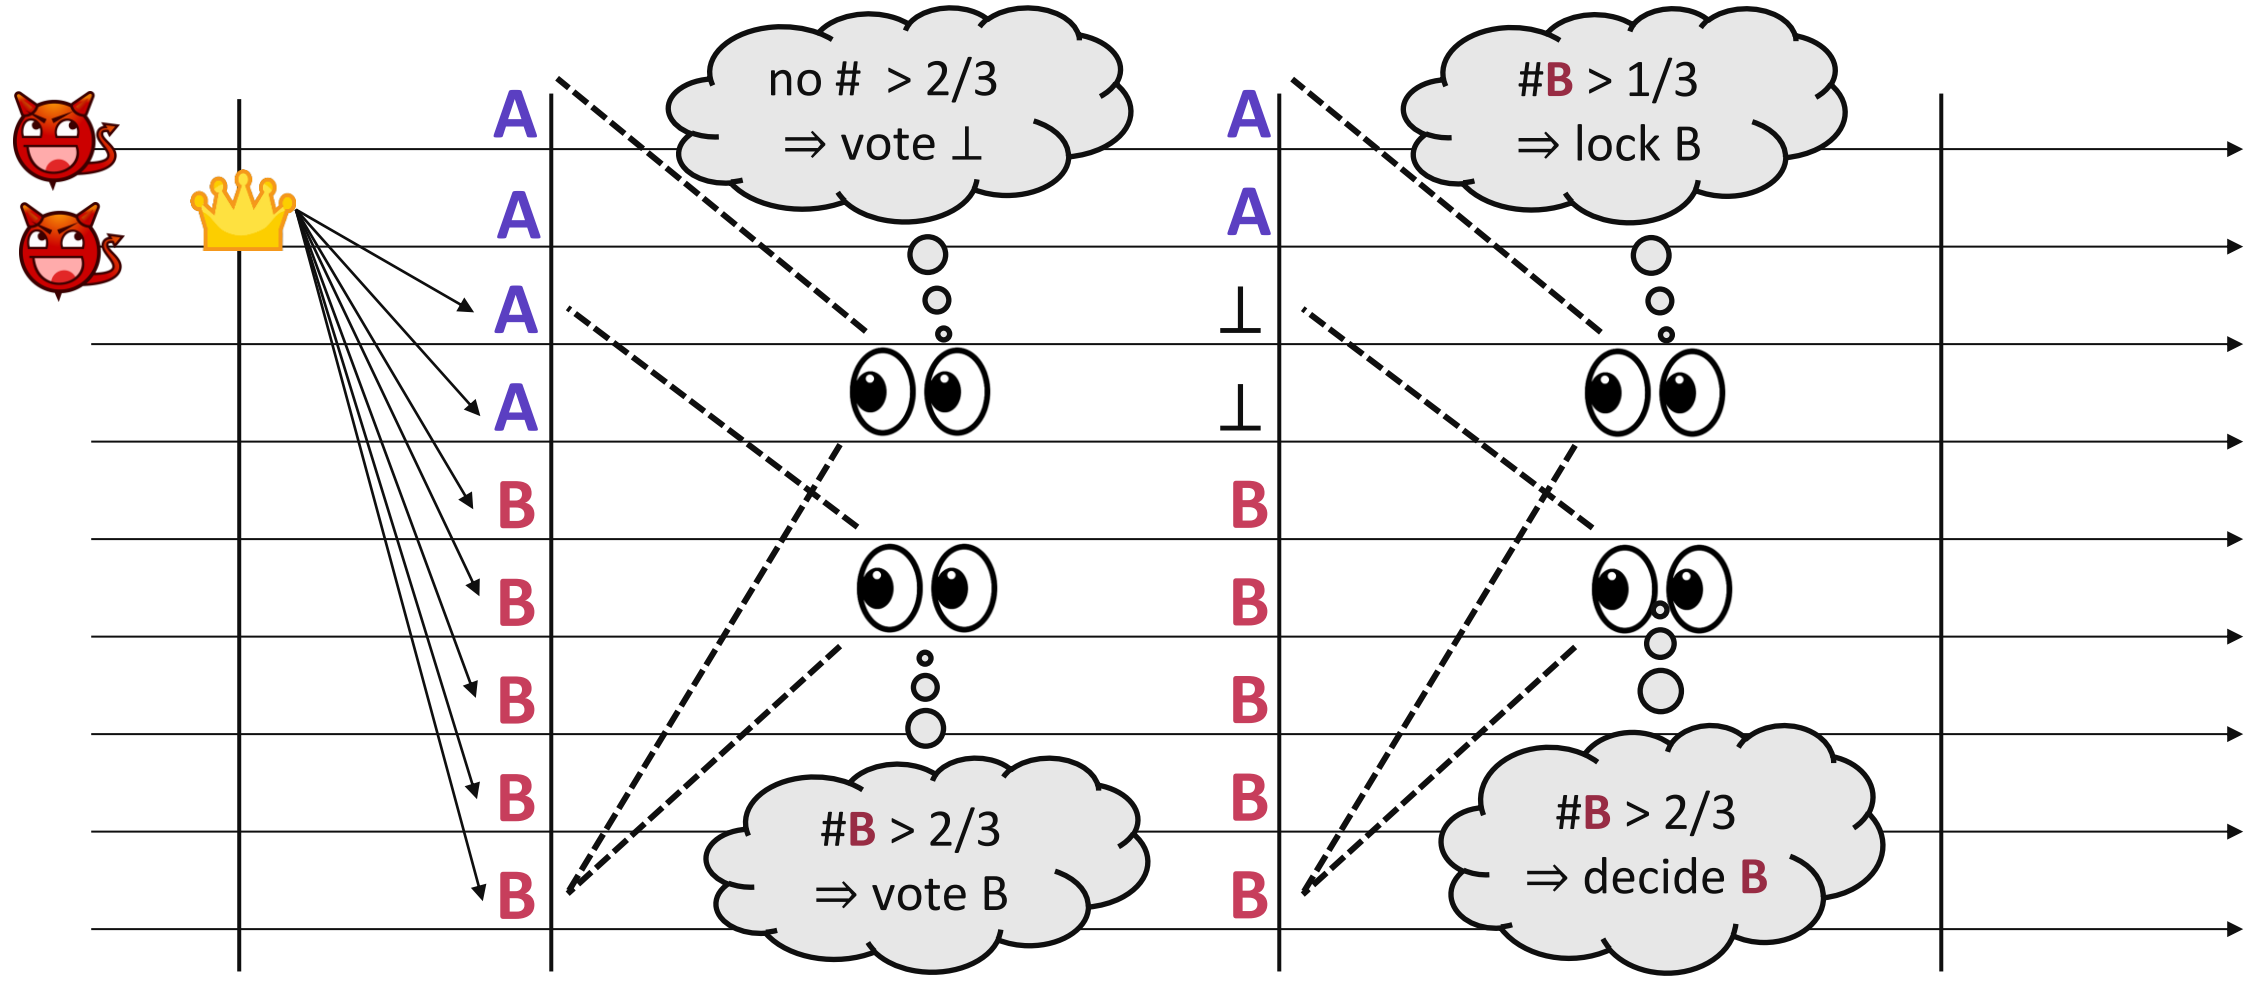

Correct parties vote for A or B

Correct parties vote B or ⊥

Correct parties decide B or lock B

To run MMR in the fully-permissionless model, we must deal with Sybil attacks

- There is no list of participants
  - Anybody on the Internet can send a vote…
  - Attackers can send as many votes as they want using pseudonymous identities (a Sybil attack)
  - How do we know which votes to count?

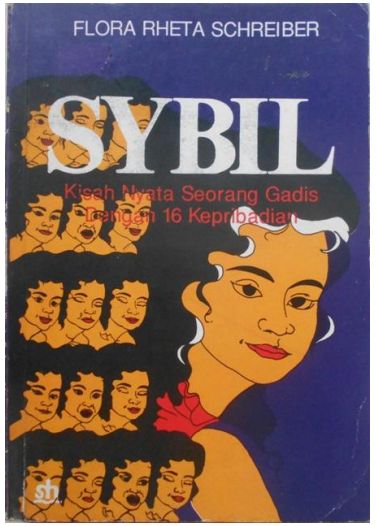

Idea: one PoW = one message (no PoW ⇒ message rejected)

If adversaries control less than 1/3 mining power, will MMR work?

“One PoW = one message” does not work: adversaries can accumulate PoWs and reveal them at a later step

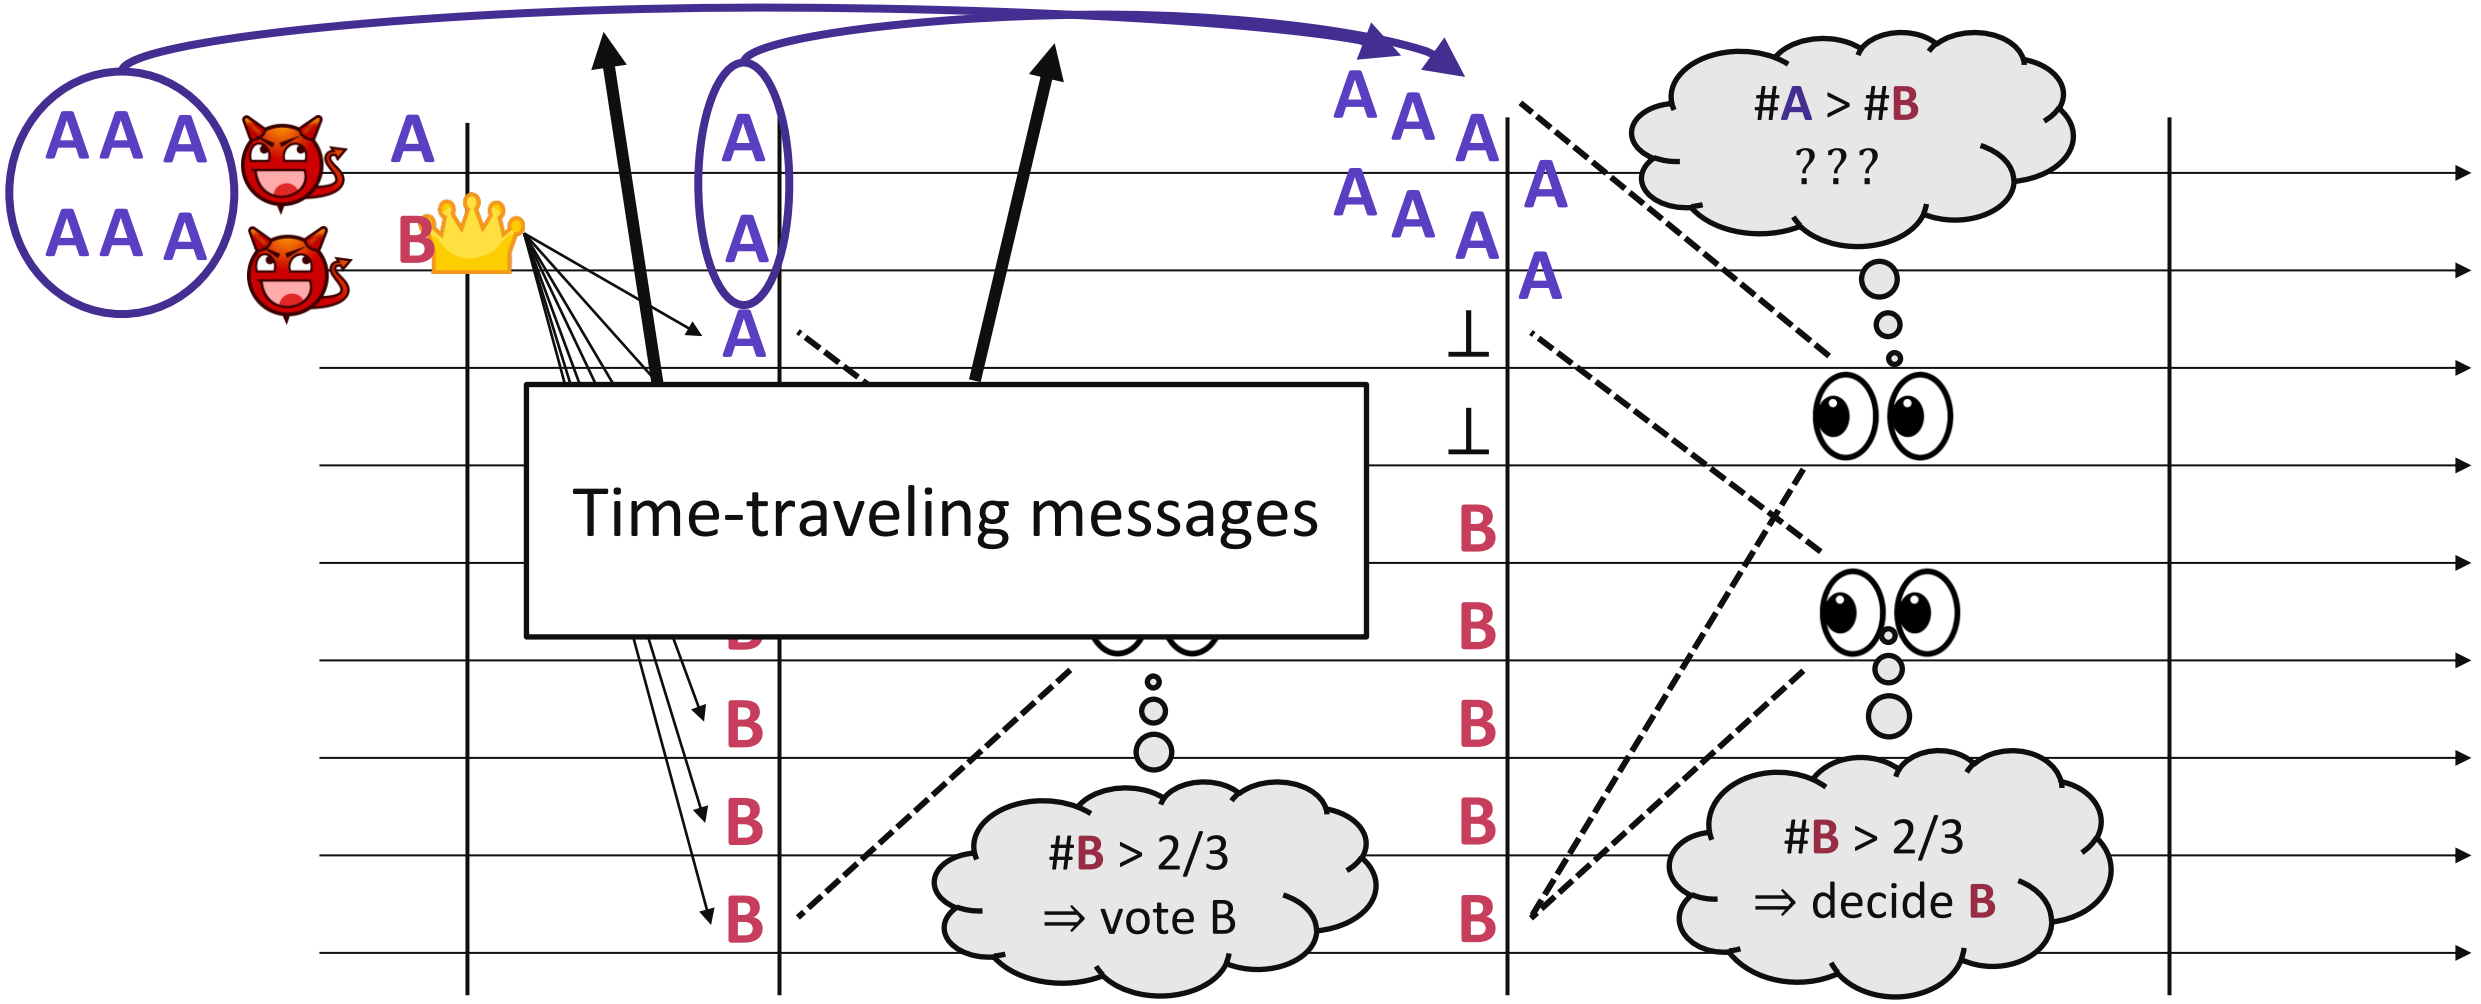

Contribution: Sieve, a (deterministic) messaging layer implementing Time-Travel Resilient Broadcast (TTRB)

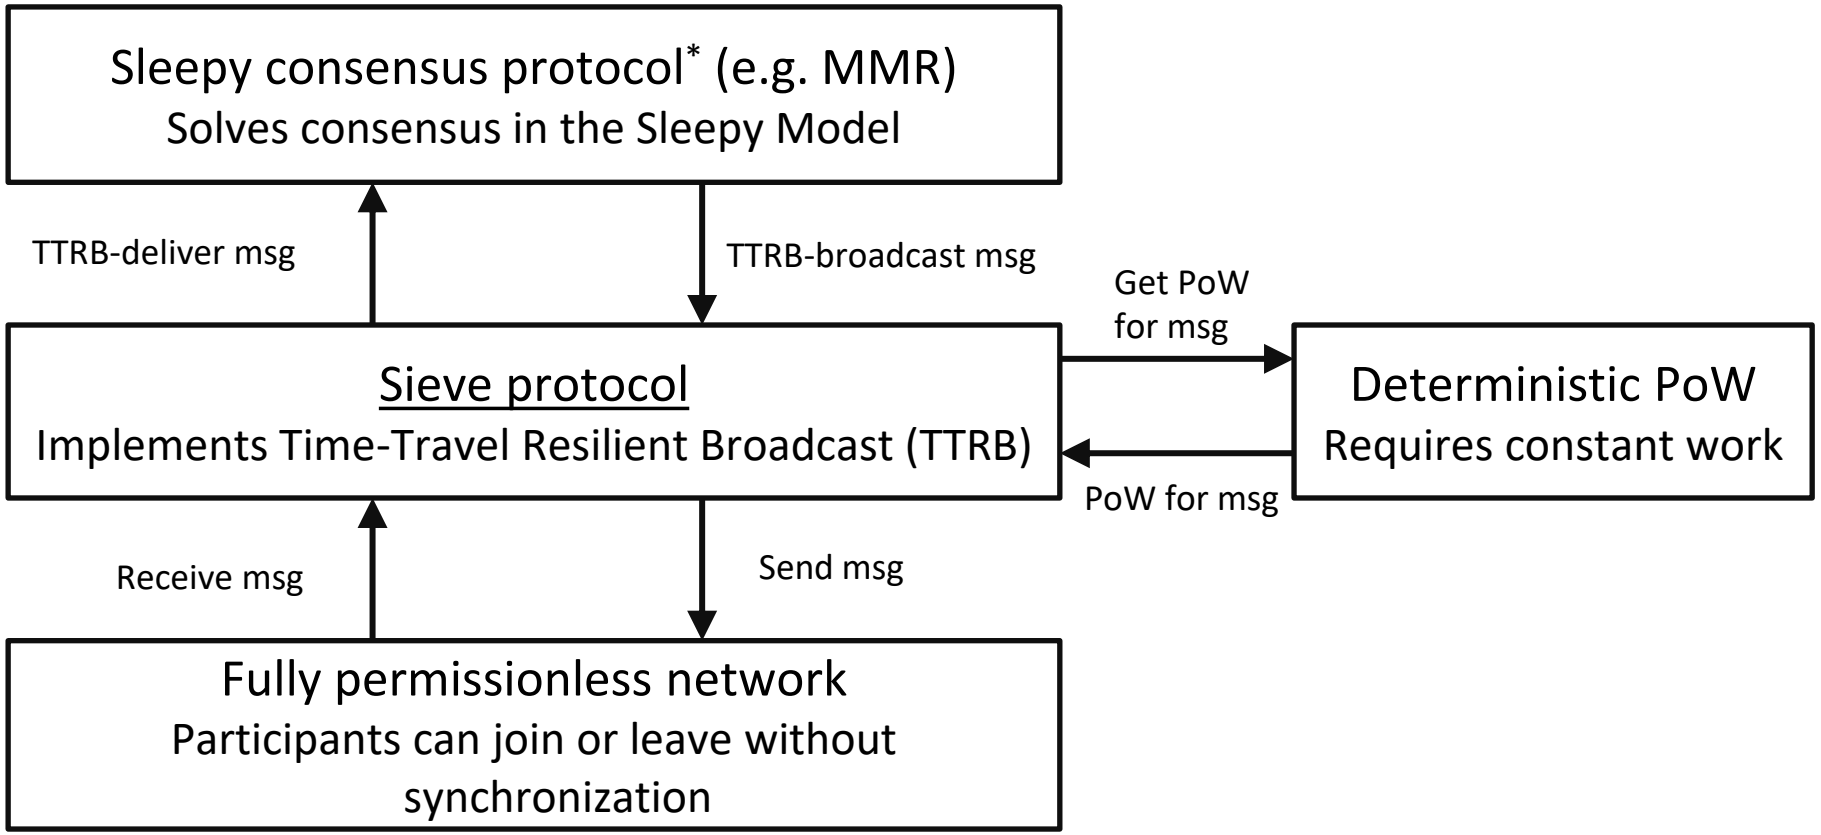

* with 1-round look-back

To identify time-traveling messages, Sieve analyzes the causal history of messages

- Each message must contain a “coffer” consisting of:
  - the set of all the messages received from the previous step
  - a PoW tied to the set

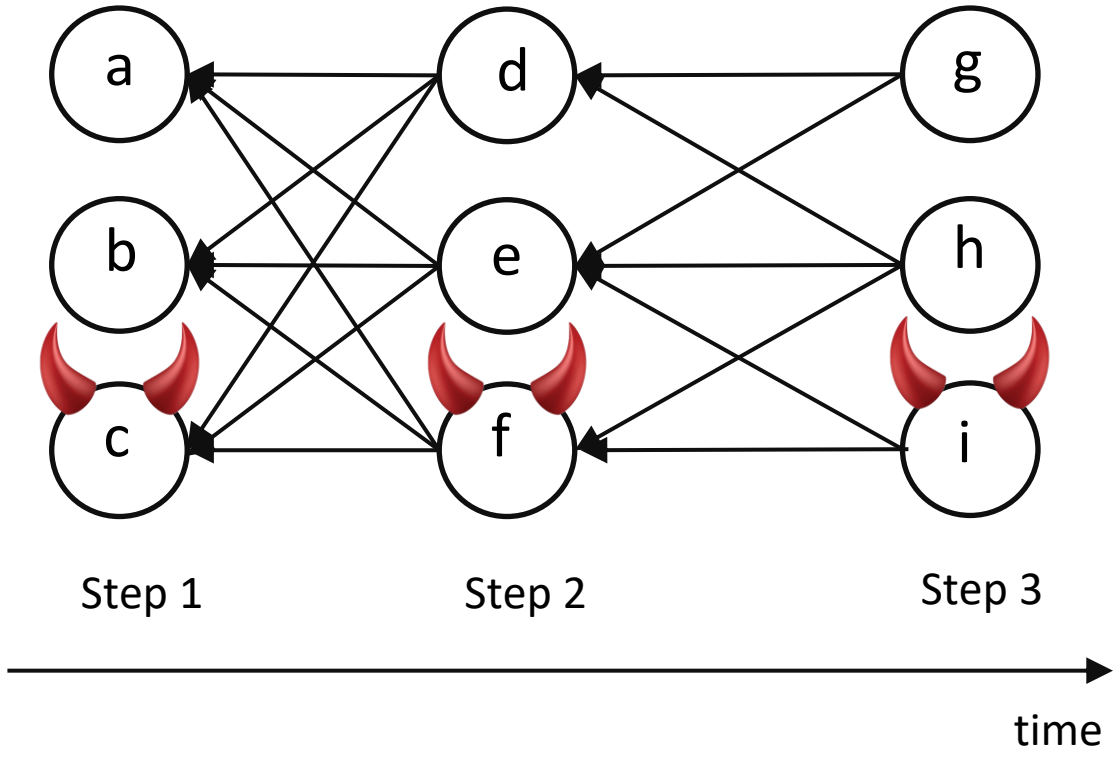

To identify time-traveling messages, Sieve analyzes the causal history of messages

- Each message must contain a “coffer” consisting of:
  - the set of all the messages received from the previous step
  - a PoW tied to the set
- Time-traveling messages cannot point to enough non-time-traveling previous-step messages!

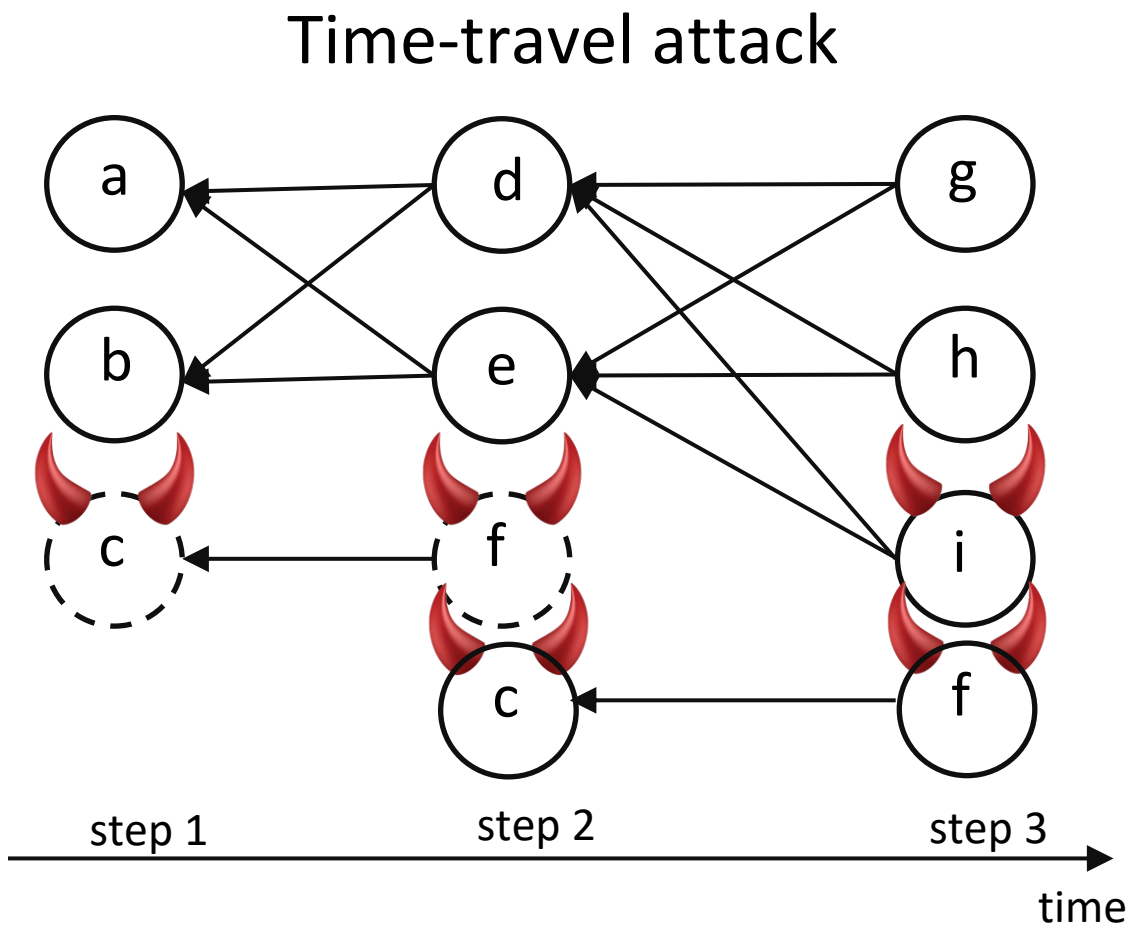

We use a PoW puzzle that takes a fixed, determined amount of work

Honest provers expend fixed work

Adversaries do less work with probability decreasing exponentially with k

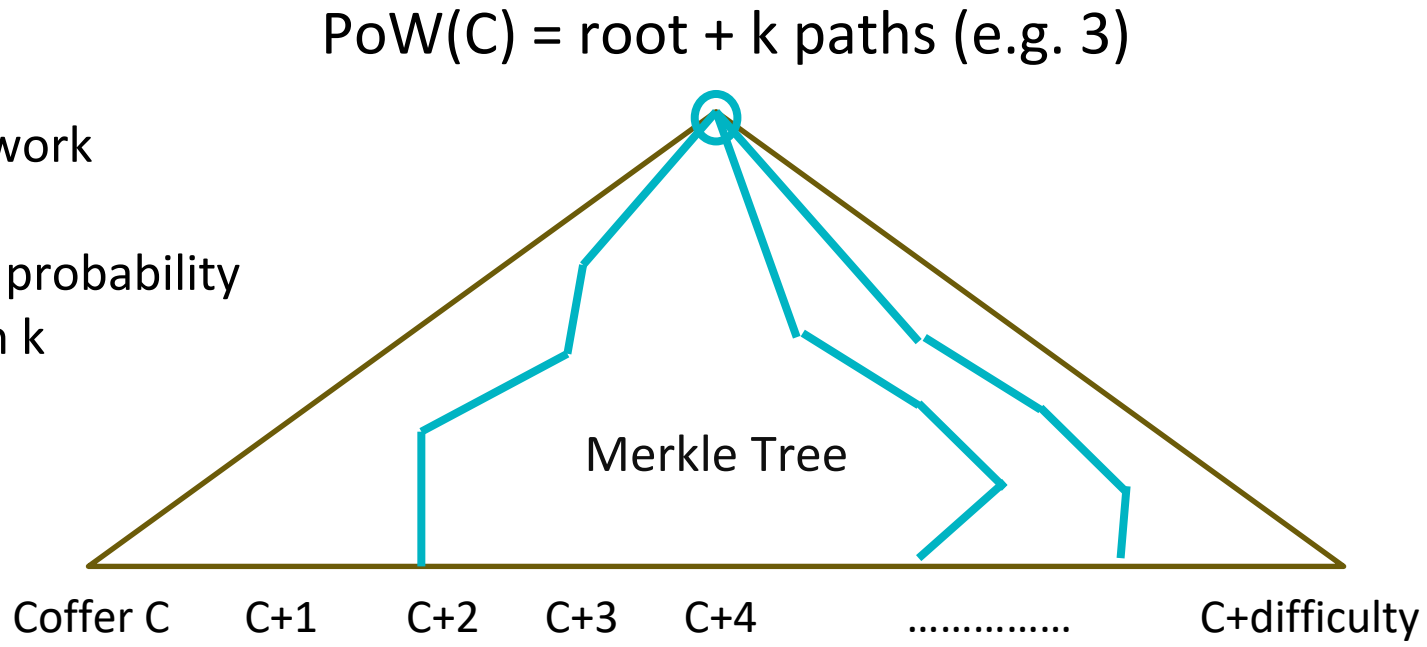

Due to Fabien Coelho in *An (Almost) Constant-Effort Solution-Verification Proof-of-Work Protocol Based on Merkle Trees*. Progress in Cryptology – AFRICACRYPT 2008

Sieve-MMR: consensus with deterministic safety, constant latency, and fully-permissionless participation

- Sieve-MMR solves consensus in the fully-permissionless PoW model
  - With deterministic safety
  - Termination with probability 1 (uses randomized leader election for liveness)
- Resilience: 1/3 adversarial computing power
  - Sieve is ½ resilient, but MMR only 1/3
- Best-case latency of 3 steps and 7 steps in expectation
- Future work: make it more efficient!
  - Can we have less than n^2 communication per round?
  - Sieve has worst-case local computation exponential in the size of the causal graph when a party first joins. Can we do better?