Better Consensus In The Bitcoin Model

Prateek Saxena

Asst. Professor, Computer Science, NUS

Blockchain Winter School, Switzerland

Blockchains: Origin & Today

Bitcoin: A Peer-to-Peer Electronic Cash System

Satoshi Nakamoto
satoshin@gmx.com
www.bitcoin.org

Abstract. A purely peer-to-peer version of electronic cash would allow online payments to be sent directly from one party to another without going through a financial institution. Digital signatures provide part of the solution, but the main benefits are lost if a trusted third party is still required to prevent double-spending. We propose a solution to the double-spending problem using a peer-to-peer network. The network timestamps transactions by hashing them into an ongoing chain of hash-based proof-of-work, forming a record that cannot be changed without redoing the proof-of-work. The longest chain not only serves as proof of the sequence of events witnessed, but proof that it came from the largest pool of CPU power. As long as a majority of CPU power is controlled by nodes that are not cooperating to attack the network, they'll generate the longest chain and outpace attackers. The network itself requires minimal structure. Messages are broadcast on a best effort basis, and nodes can leave and rejoin the network at will, accepting the longest proof-of-work chain as proof of what happened while they were gone.

1. Introduction

Commerce on the Internet has come to rely almost exclusively on financial institutions serving as trusted third parties to process electronic payments. While the system works well enough for most transactions, it still suffers from the inherent weaknesses of the trust based model. Completely non-reversible transactions are not really possible, since financial institutions cannot avoid mediating disputes. The cost of mediation increases transaction costs, limiting the minimum practical transaction size and cutting off the possibility for small casual transactions, and there is a broader cost in the loss of ability to make non-reversible payments for non-reversible services. With the possibility of reversal, the need for trust spreads. Merchants must be wary of their customers, hassling them for more information than they would otherwise need. A certain percentage of fraud is accepted as unavoidable. These costs and payment uncertainties can be avoided in person by using physical currency, but no mechanism exists to make payments over a communications channel without a trusted party.

What is needed is an electronic payment system based on cryptographic proof instead of trust, allowing any two willing parties to transact directly with each other without the need for a trusted third party. Transactions that are computationally impractical to reverse would protect sellers from fraud, and routine escrow mechanisms could easily be implemented to protect buyers. In this paper, we propose a solution to the double-spending problem using a peer-to-peer distributed timestamp server to generate computational proof of the chronological order of transactions. The system is secure as long as honest nodes collectively control more CPU power than any cooperating group of attacker nodes.

2008

Top 100 Cryptocurrencies by Market Capitalization

Cryptocurrencies ▾ | Exchanges ▾ | Watchlist | USD ▾ | Next 100 → | View All

#	Name	Market Cap	Price	Volume (24h)	Circulating Supply	Change (24h)	Price Graph (7d)
1	Bitcoin	$59,580,761,374	$3,399.59	$4,939,435,528	17,525,862 BTC	-0.49%	
2	XRP	$12,001,134,334	$0.291508	$356,976,506	41,169,202,069 XRP *	-0.25%	
3	Ethereum	$10,974,571,873	$104.75	$2,280,623,059	104,766,118 ETH	-0.48%	
4	EOS	$2,126,001,619	$2.35	$472,575,374	906,245,118 EOS *	-0.64%	
5	Bitcoin Cash	$2,041,982,753	$115.96	$198,734,194	17,609,650 BCH	0.05%	
6	Tether	$2,026,509,895	$1.00	$3,511,890,558	2,021,103,317 USDT *	0.10%	
7	Litecoin	$2,000,776,268	$33.14	$636,413,250	60,369,927 LTC	0.12%	
8	TRON	$1,712,362,099	$0.025684	$136,340,725	66,671,422,606 TRX	-0.61%	
9	Stellar	$1,422,240,776	$0.074196	$114,737,113	19,168,570,823 XLM *	0.28%	
10	Binance Coin	$1,101,770,088	$7.80	$84,783,682	141,175,490 BNB *	-4.79%	
11	Bitcoin SV	$1,091,169,120	$61.97	$83,503,727	17,608,711 BSV	-1.61%	
12	Cardano	$940,904,576	$0.036290	$12,298,410	25,927,070,538 ADA *	-0.51%	
13	Monero	$727,519,153	$43.36	$46,156,605	16,777,423 XMR	0.27%	

2019

The Blockchain Consensus Problem

The Problem

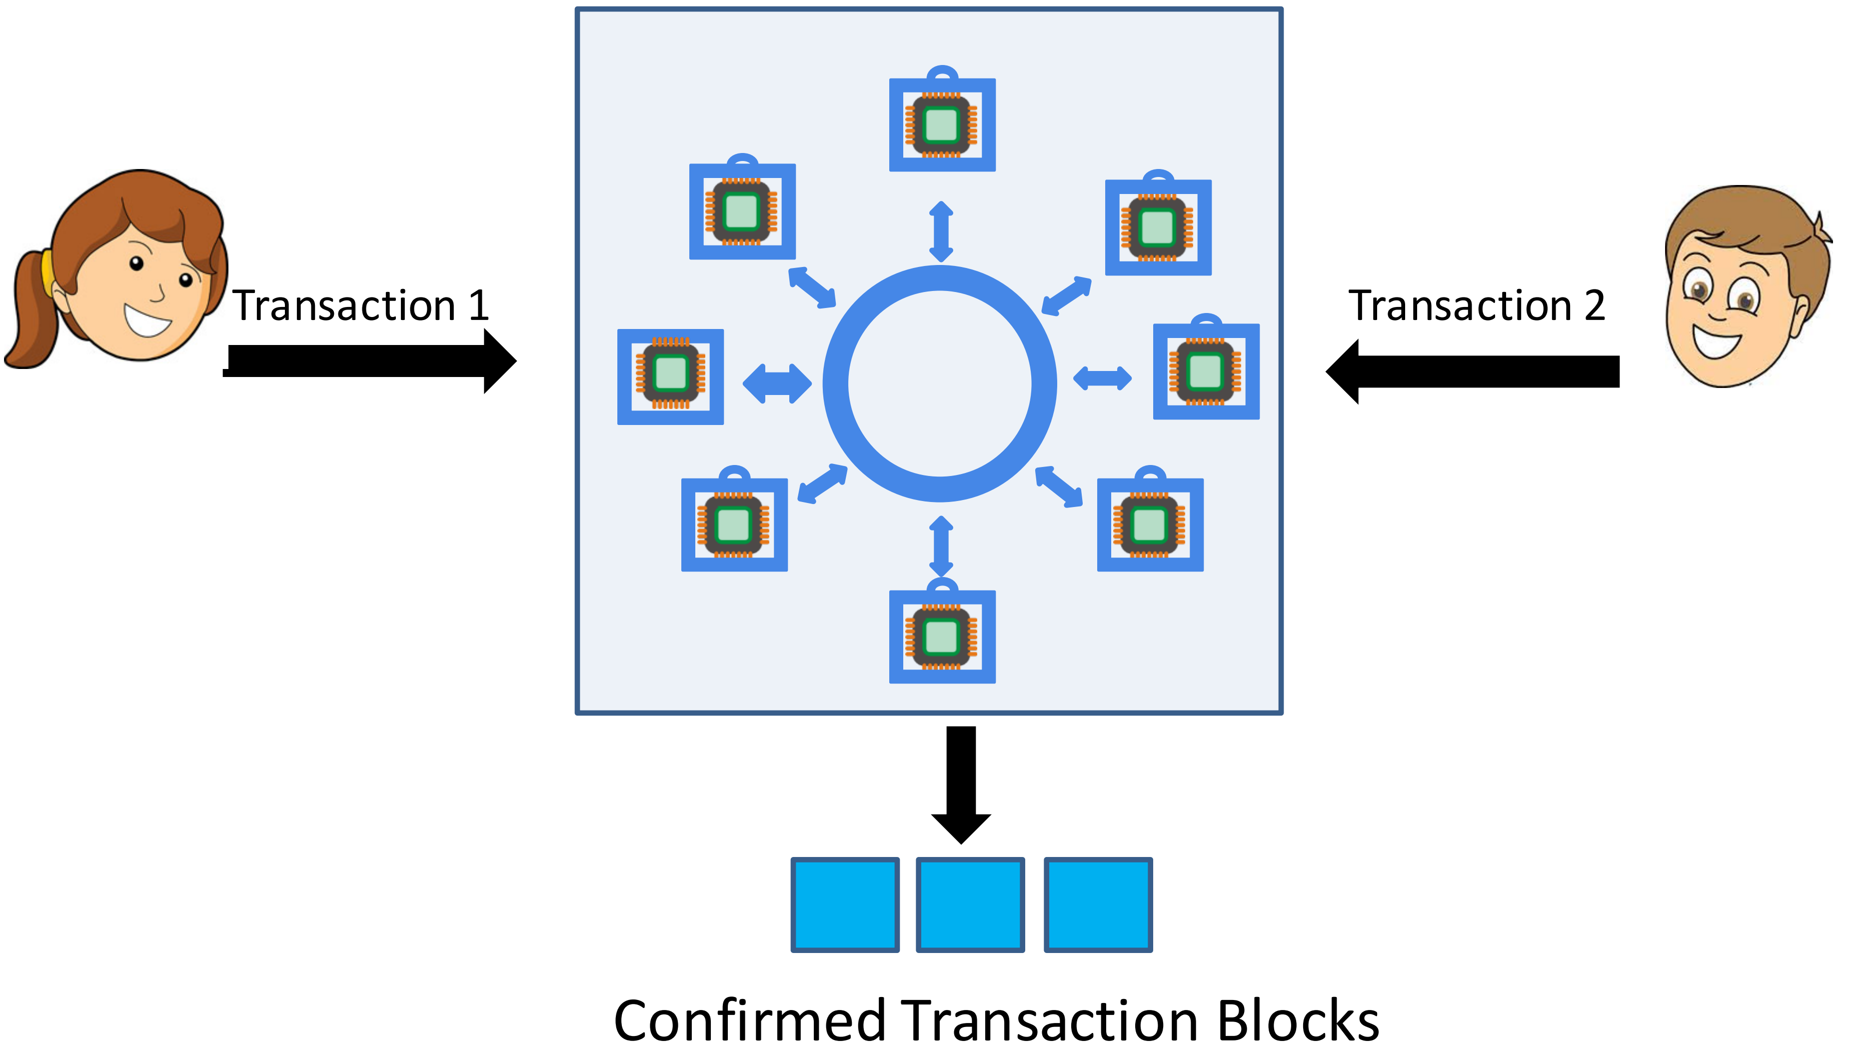
Transaction 1
Transaction 2
Confirmed Transaction Blocks

Key Challenge:
Agreement over Transaction Ordering

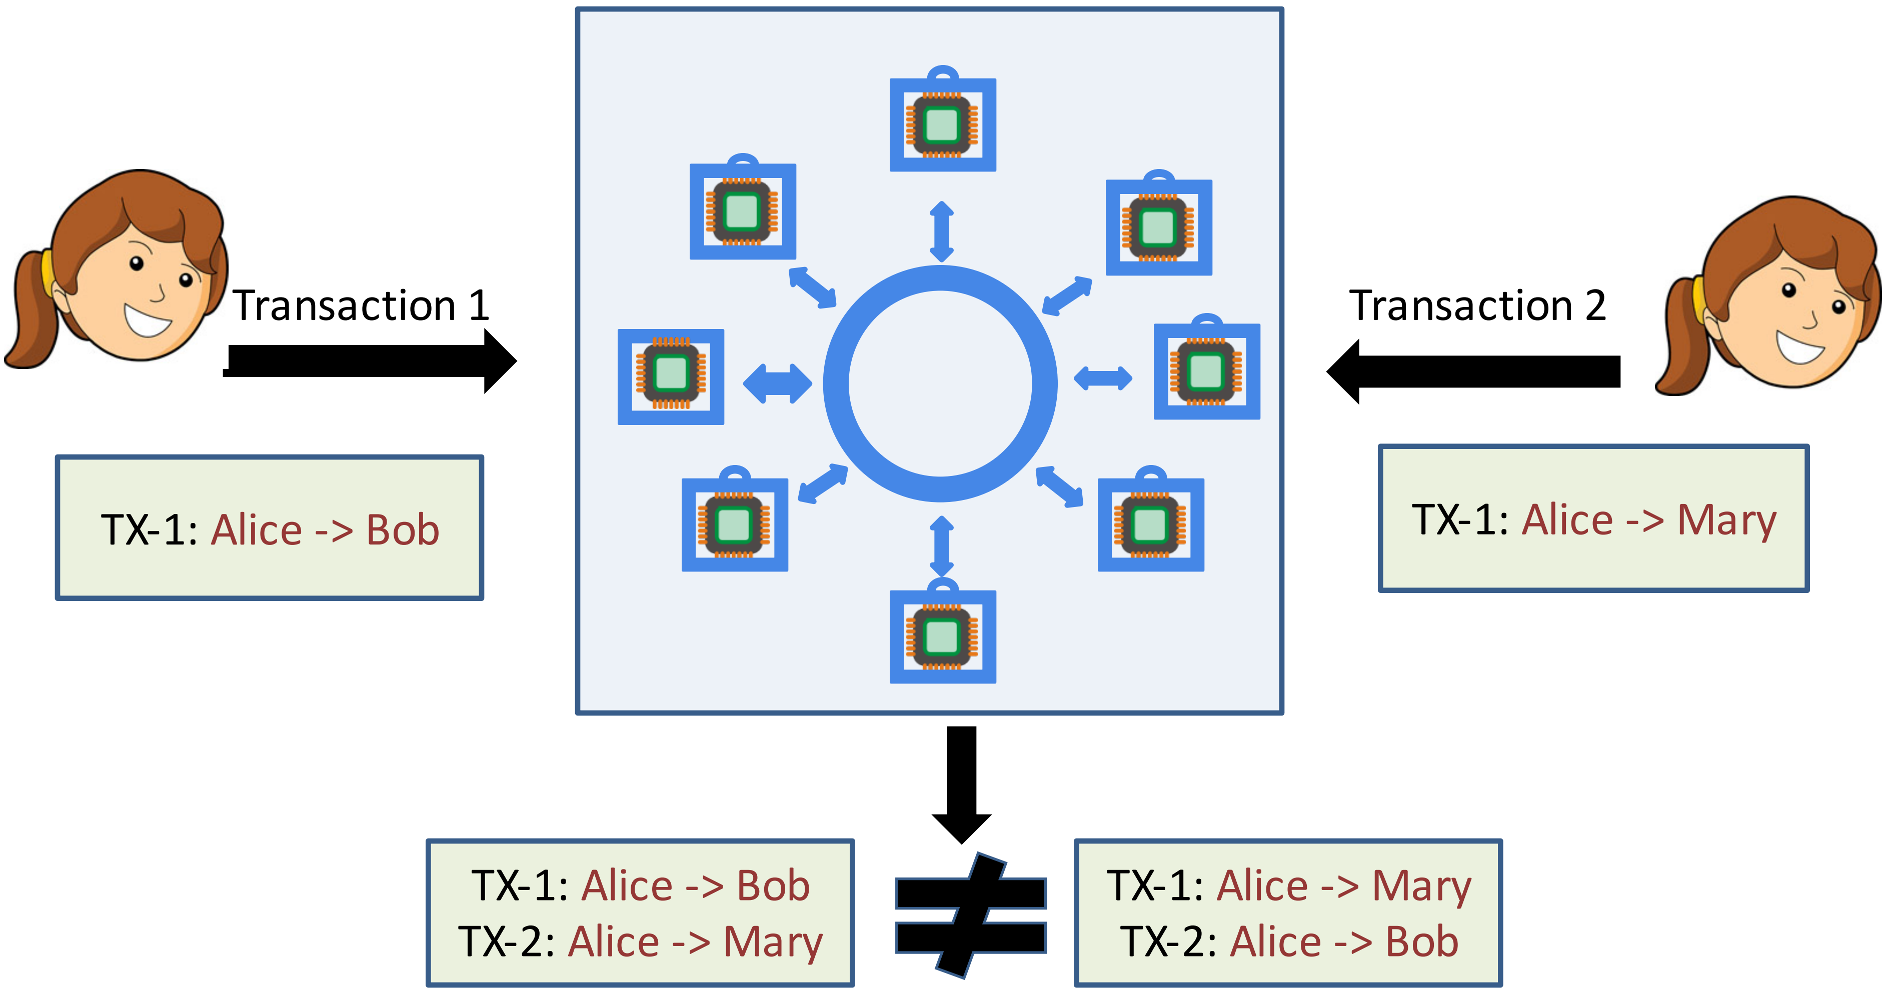

Ordering Transactions is sufficient to prevent double-spends!

Why Total Order?

- Replicated State Machines *[Lamport84, Schnieder90]*
  - Useful for backups, snapshots, distributed locks, …
  - A sequence of commands transition from state to state

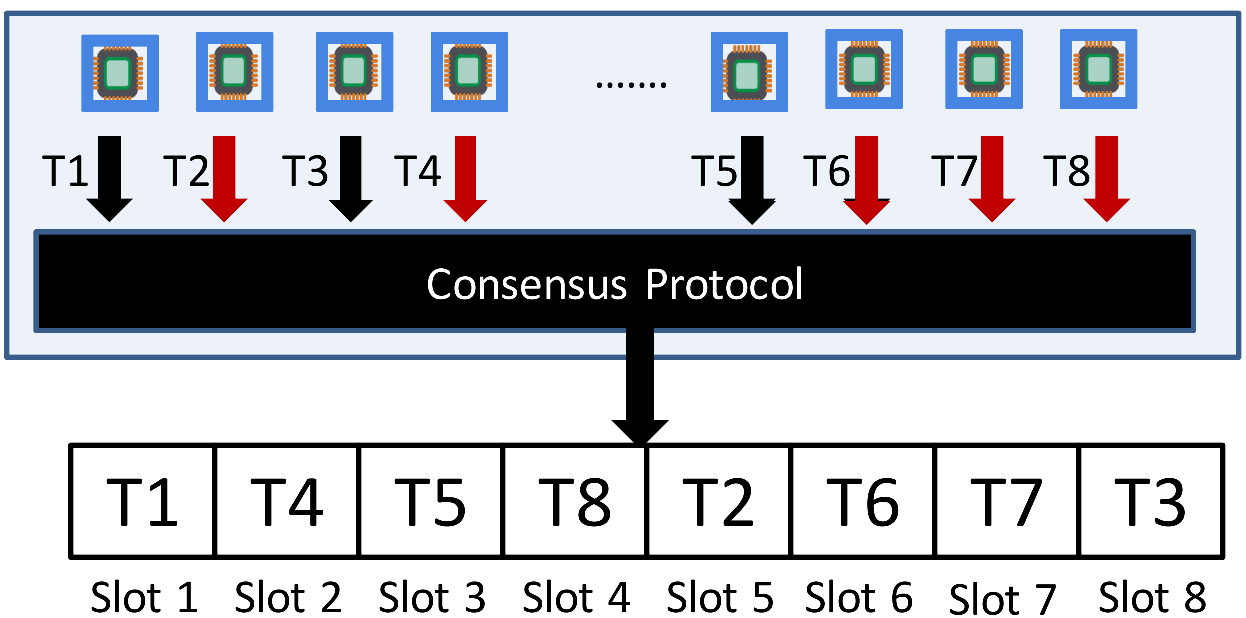

Replicated Log

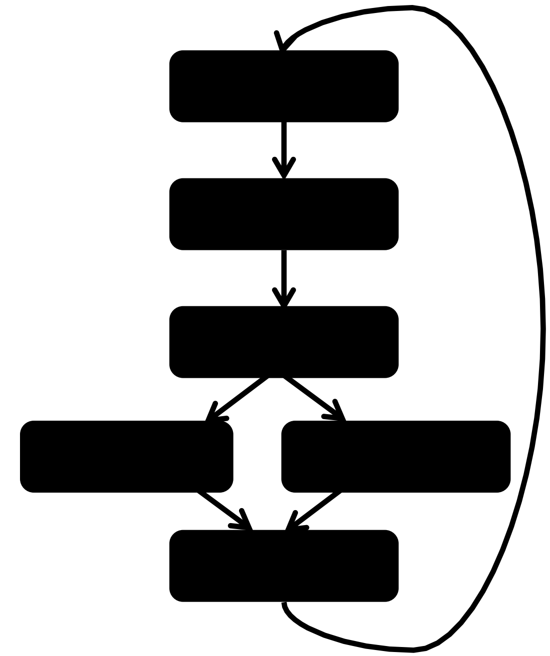

Deterministic State Machine

Enables General-Purpose Computing

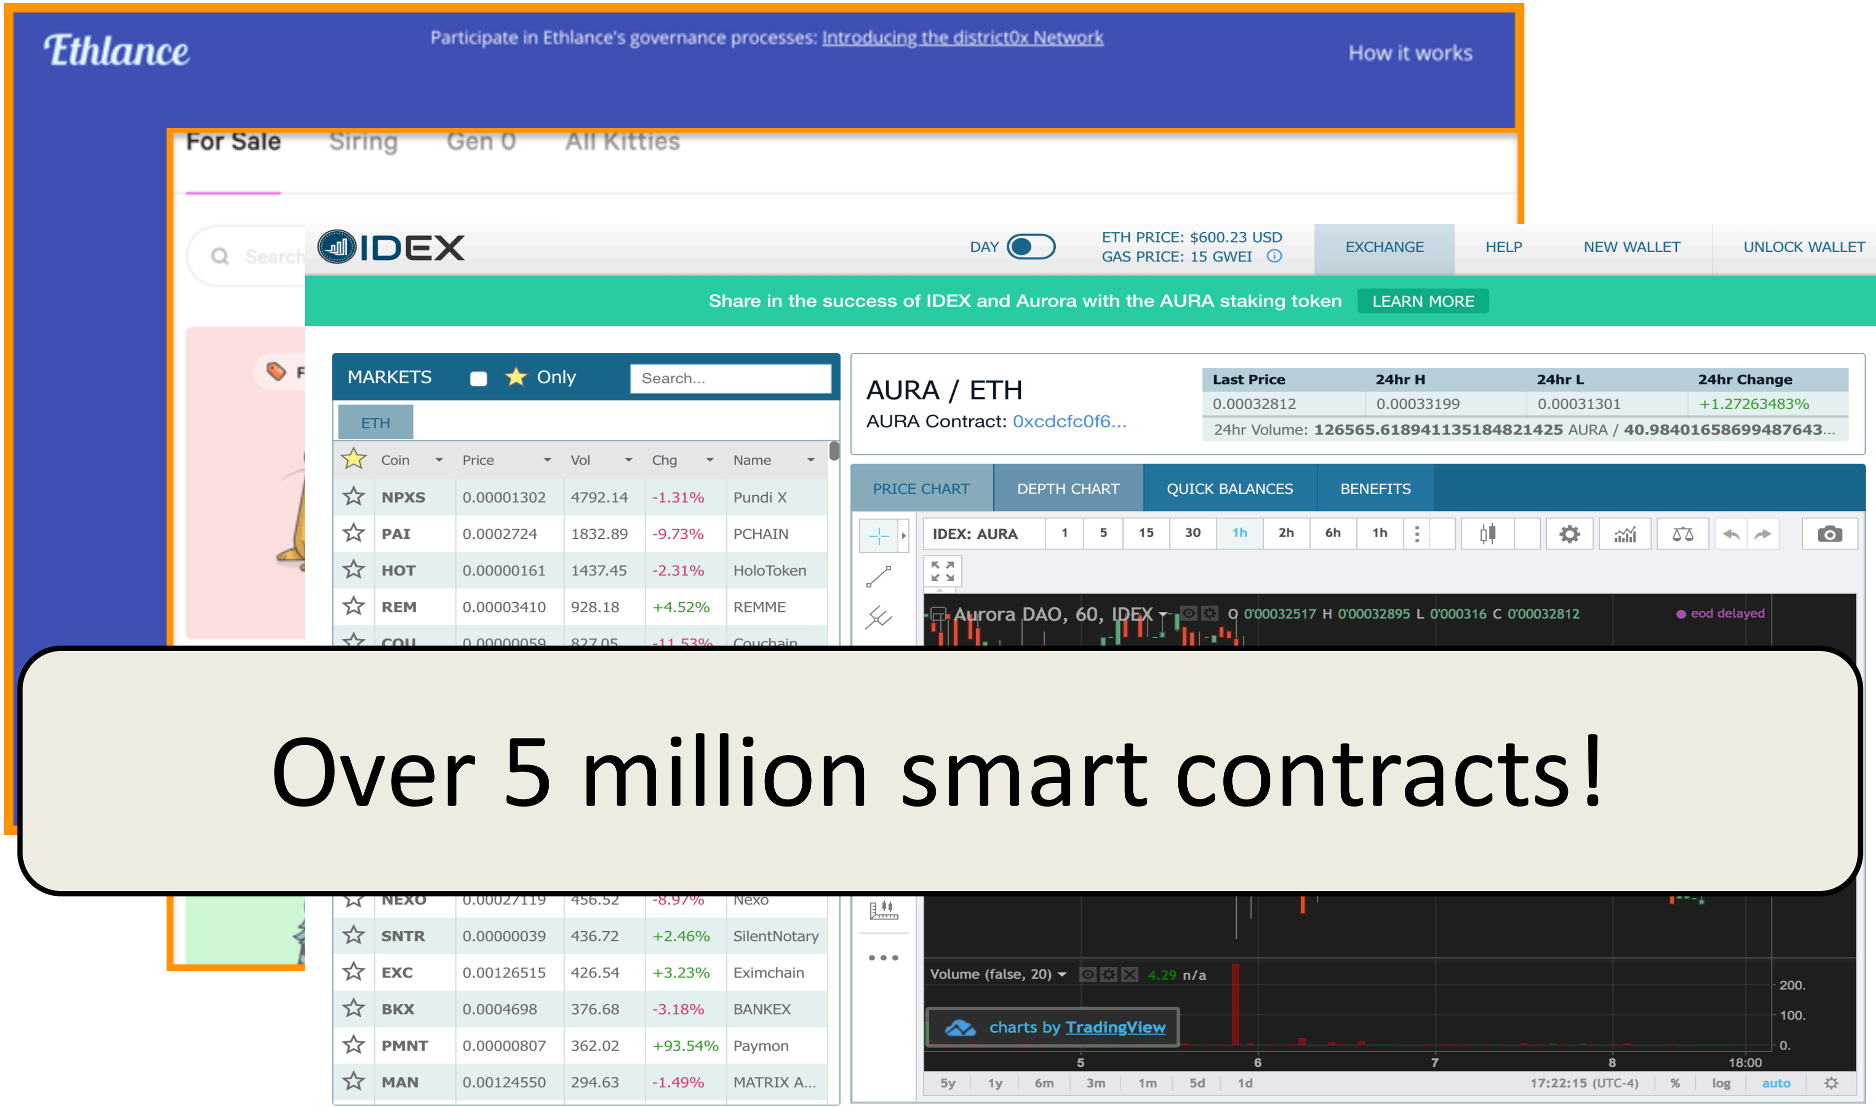

Over 5 million smart contracts!

The Bitcoin Model

- Assumptions:
  - A trusted “genesis” block
  - No pre-established identities, joining is **permissionless**
  - Network is synchronous (Blocks transmitted within some delay)
- Security Properties:
  - **Safety:** Nothing bad happens
    - **Stability:** A block once confirmed can’t be changed
    - **Agreement:** All miners order blocks same way
  - **Liveness:** Honest blocks are accepted eventually
  - **Fairness:** Your confirmed blocks are proportional to your computational power

Nakamoto Consensus Protocol

- Miners keep a local copy of the blockchain
- Miners solve a computational Proof-of-Work puzzle:

- Successful miners (usually <u>one</u>) broadcast solution
- Miners check the received solutions, and if valid:
  - Extend their chain with that block
- Confirm block on the <u>longest chain</u> after it is k-deep
    - Bitcoin proposes k = 6

Nakamoto Consensus: Overview

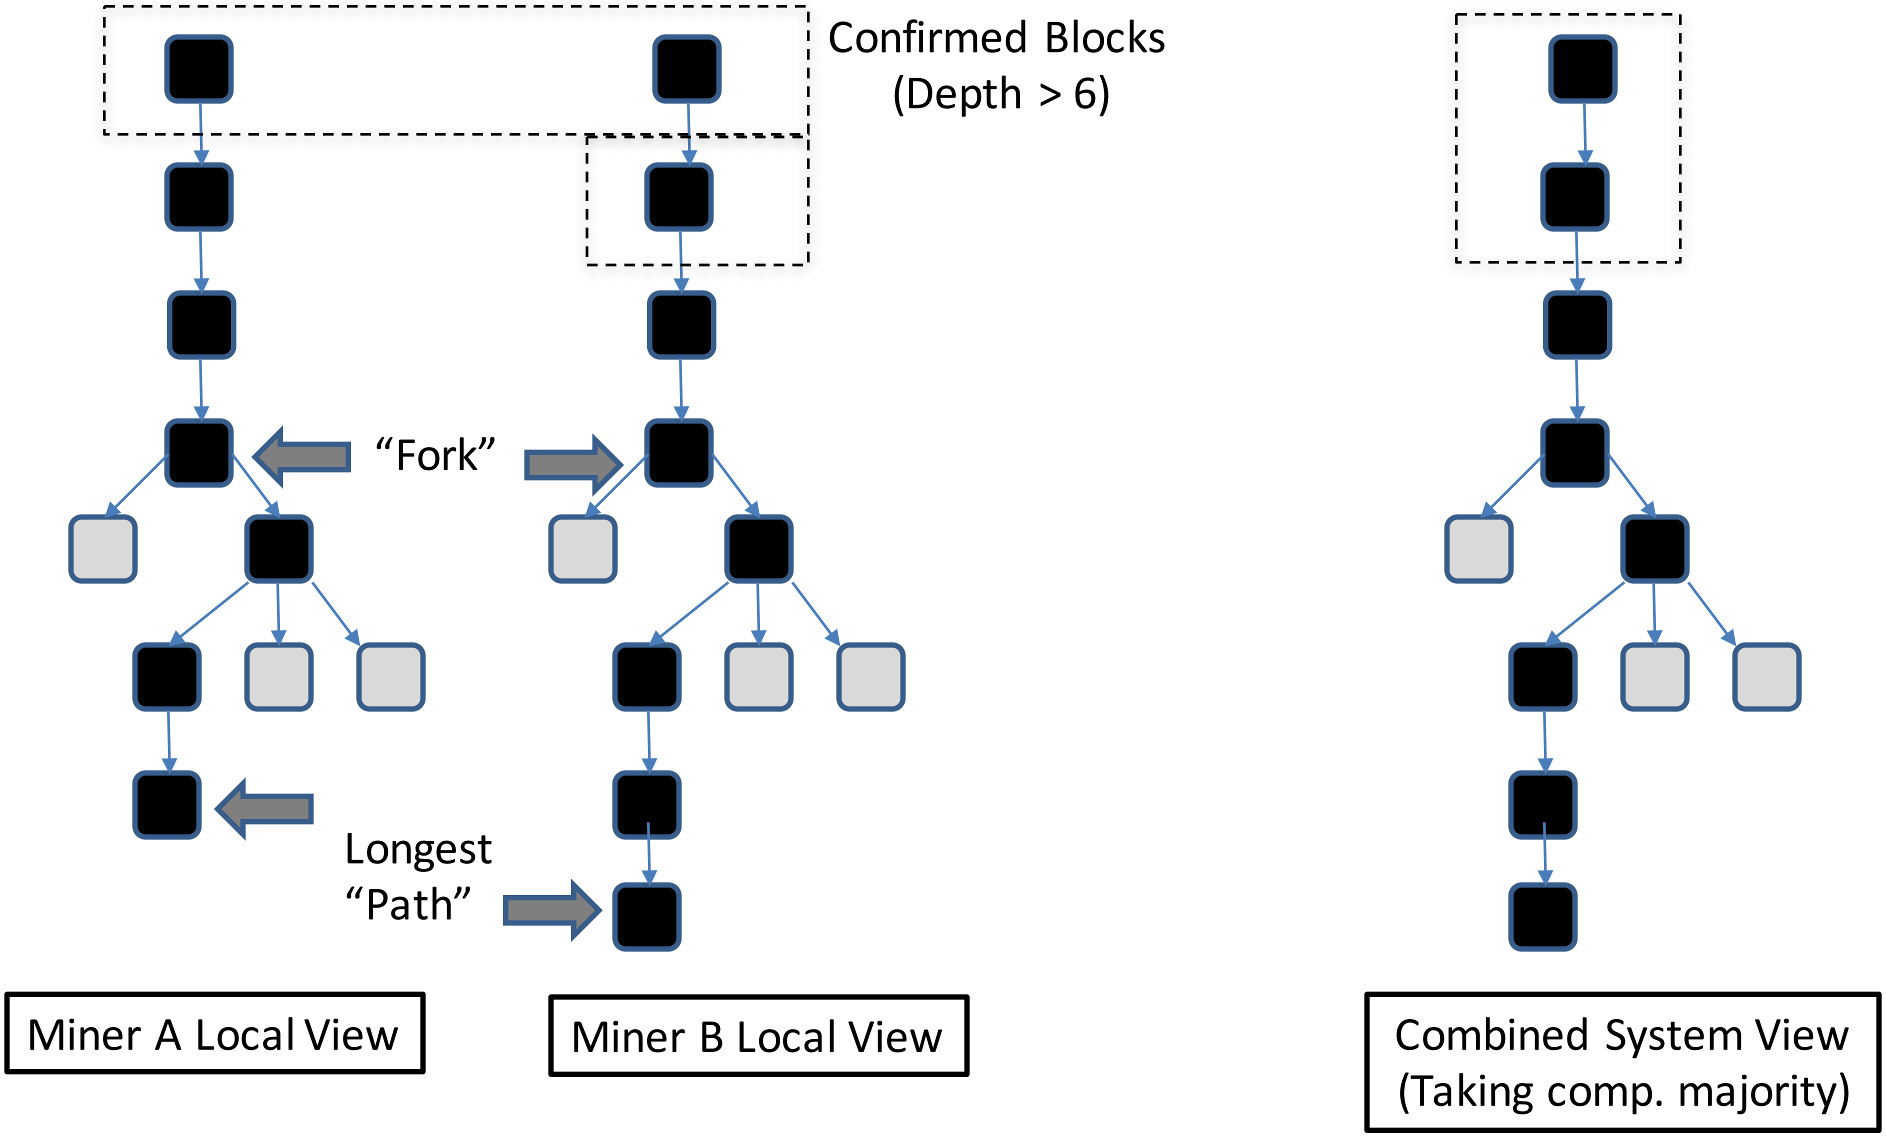

Can We Do Better?

Fundamental Limits & Optimality: **Latency**

- Limit 1: *Block Propagation Delay* (Δ)
- Optimal Transaction confirmation latency is $\Theta(\Delta)$
- A random (ER) graph with N nodes and degree d
  - Avg. hops between nodes = $(\log N)/(\log d)$
    - Bitcoin N = 12,000, d = 16, avg. hops = 3.36
    - Ethereum N = 35000, d = 25, avg. hops = 3.25
    - (Hypothetical) N = 1M, d=40, avg. hops = 4.29
  - Δ ~= $avg.\, hops \times hop\ latency$
    - On Amazon EC2 (geo-distributed) about 1-2 seconds
    - Changes minimally with (N,d)

Nakamto Consensus: Safety & Liveness are Near-Optimal

- For Nakamoto consensus,
  - **Resilience** (f) is "near-optimal" at blk. interval > 3Δ

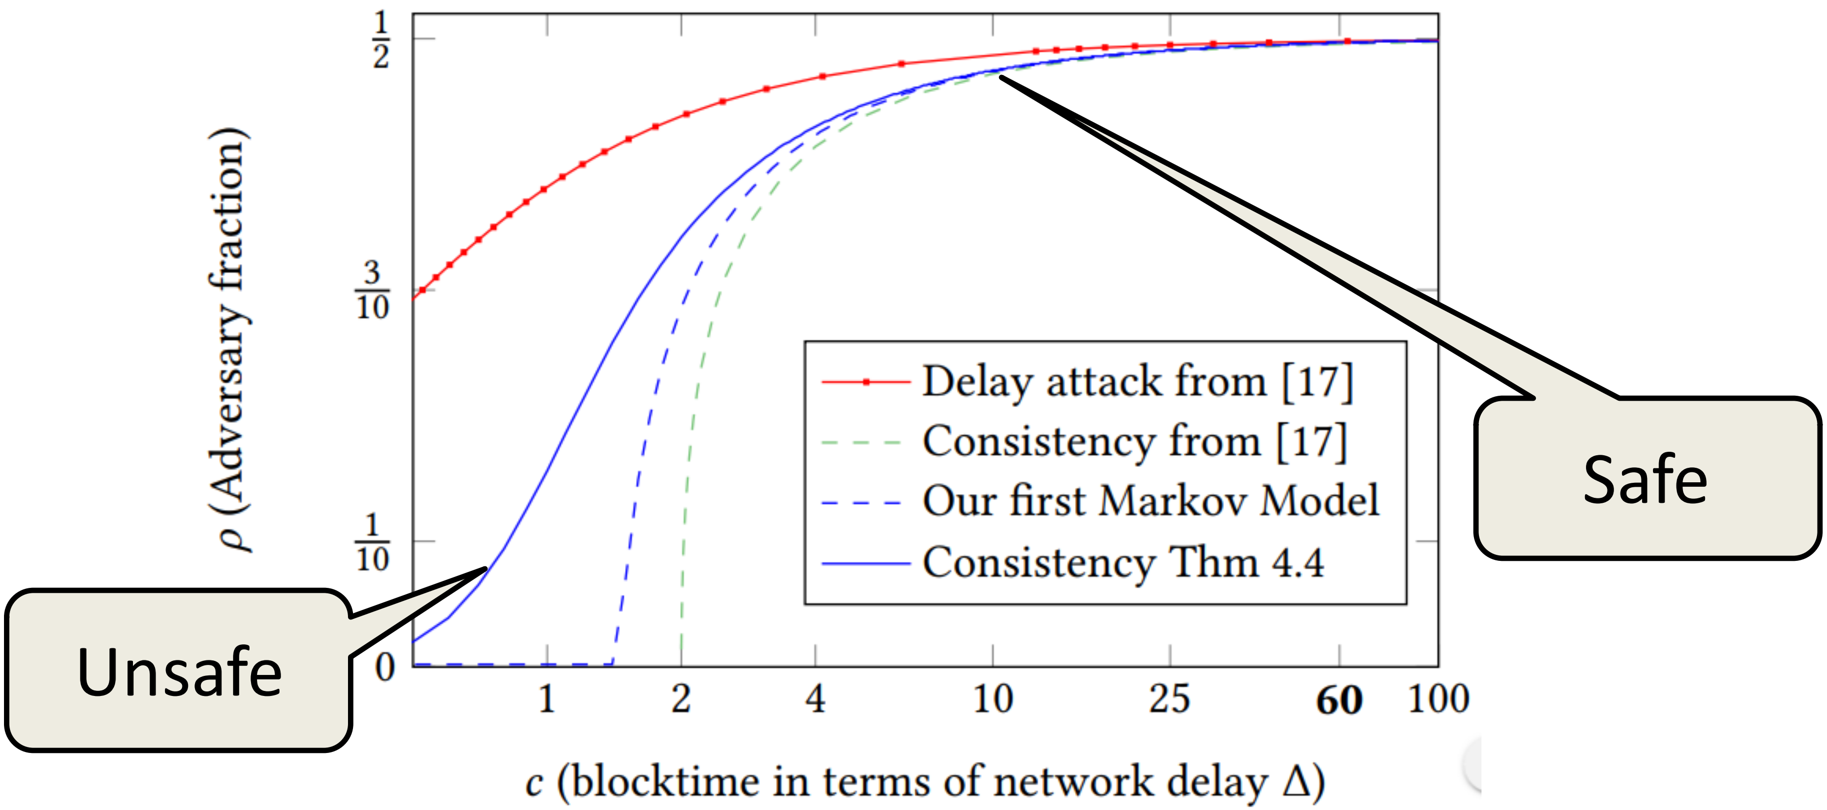

A Better Method to Analyze Blockchain Consistency" – KRS'18. (Also see GK'15, GKL'17, PSS'17)

Fundamental Limits & Optimality: **Throughput**

- Limit 2: Broadcast Throughput (β)
- Transactions per second = β / transaction size
- An experiment showing Δ (un)correlation with β

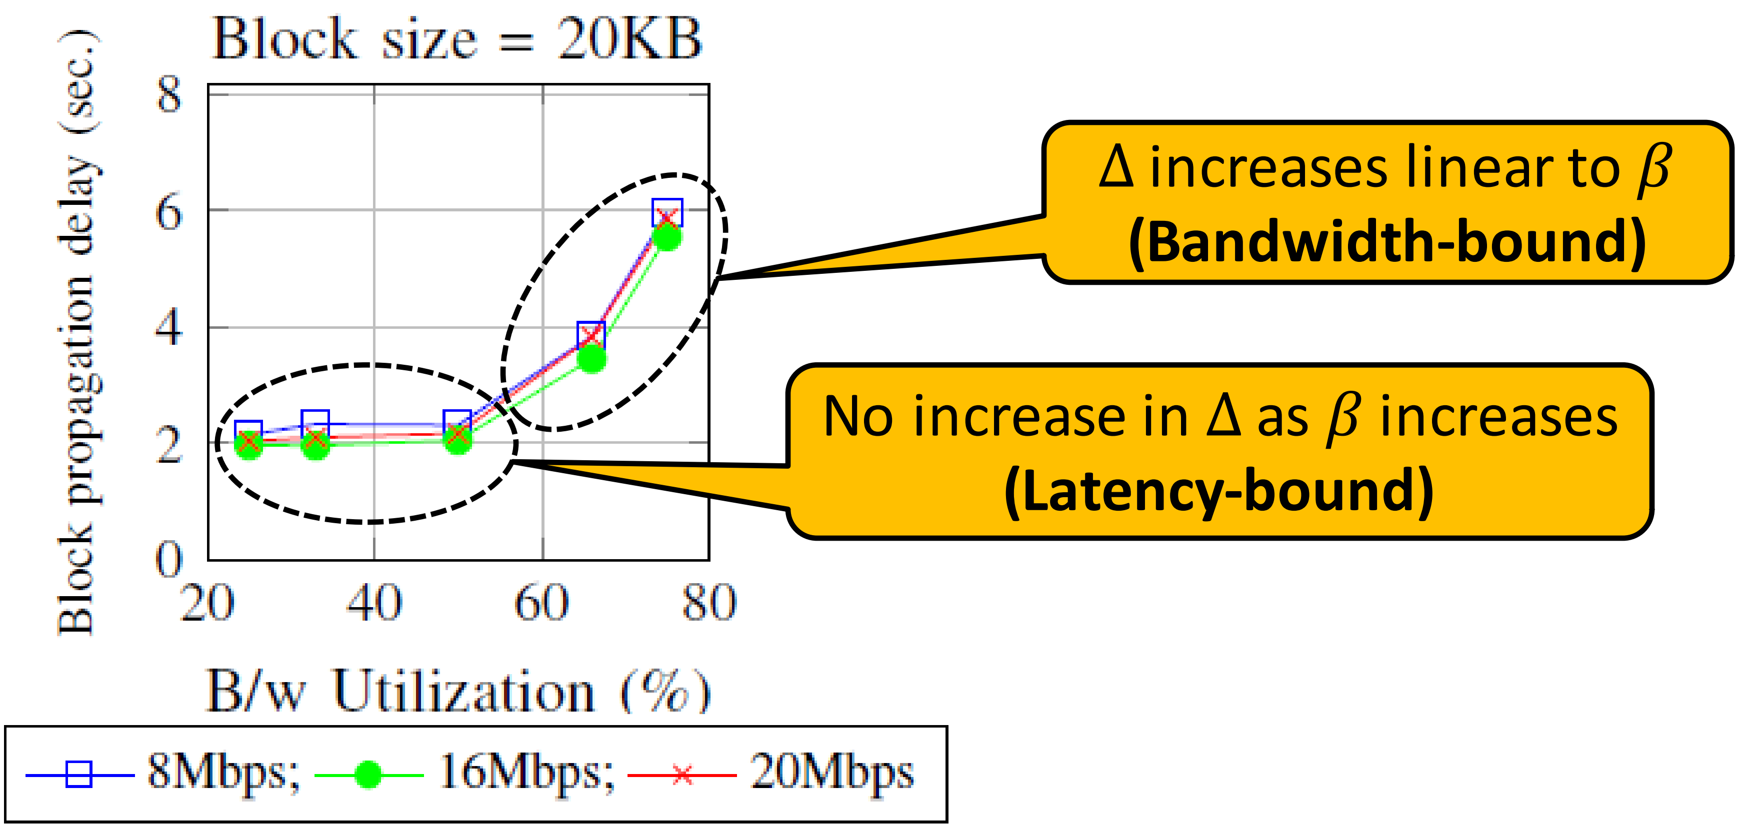

Fundamental Limits & Optimality: **Decentralization**

- In anonymous, permissionless setup
  - Mining concentration reflects “real” wealth distribution

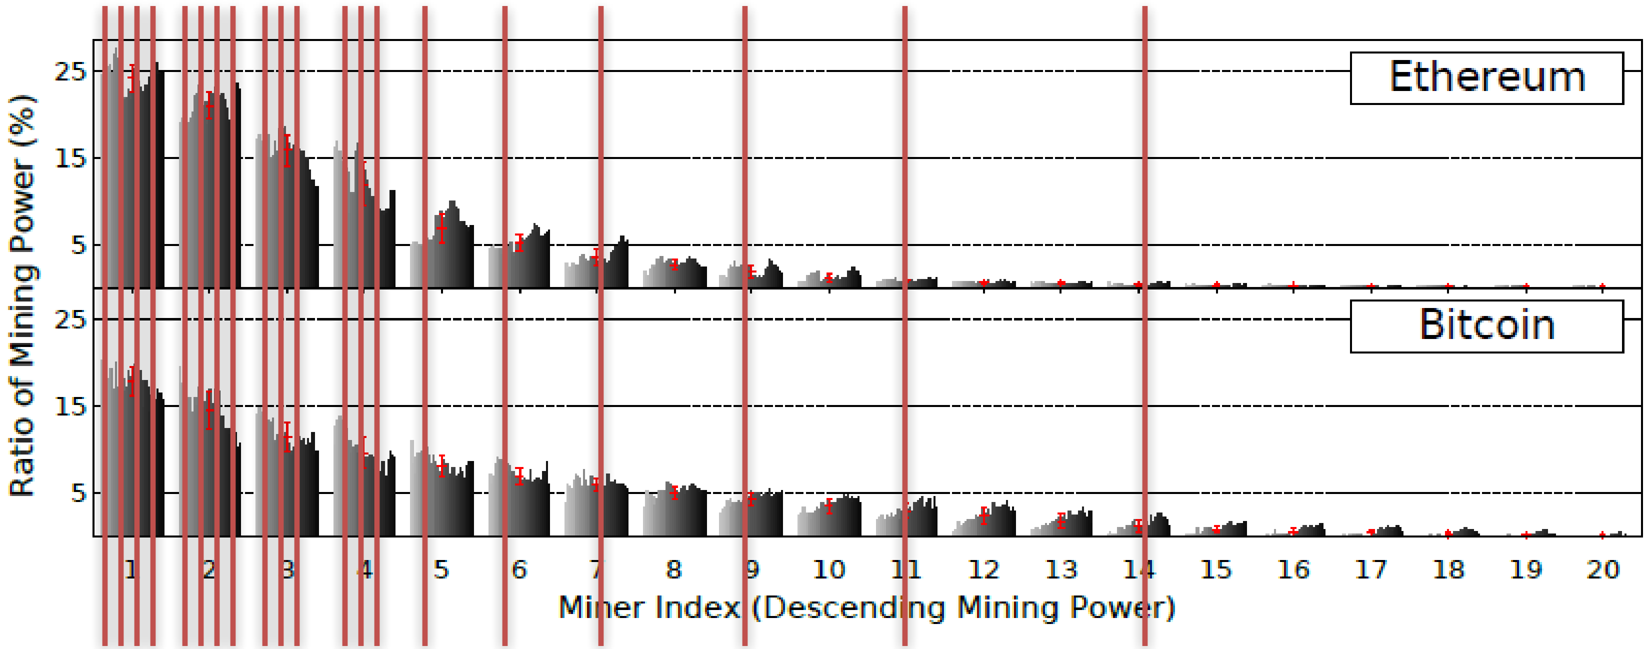

- Goal of decentralization: Maximize miners/sec
- Optimal Decentralization is $\Theta(\beta)$

Decentralization in Bitcoin and Ethereum Networks – Gencer et. al.

Nakamoto Consensus:
Not Optimal In Throughput & Decentralization

bitcoin

- 2-4 Kilobytes / second
- 6-12 TXs per second
- 3-60 minutes latency

ethereum

- Support limited computations
- Outages and Unavailability
- A cryptoKitties app clogged the entire network

Demand from Practice: 1,200 - 50,000 TXs/s

PayPal™ VISA MasterCard

Towards Better Consensus Protocols

Extending Nakamoto: With Large Blocks

- Increase block size (e.g Bitcoin-NG)
  - May achieve near-optimal throughput, latency, resilience
    - Needs a careful implementation
  - Poor decentralization:
    - A single block proposer broadcasts tens of thousands of TXs
    - Number of miners participating is <u>not</u> $\Theta(\beta)$

Extending Nakamoto With Smaller Block Interval

The GHOST protocol

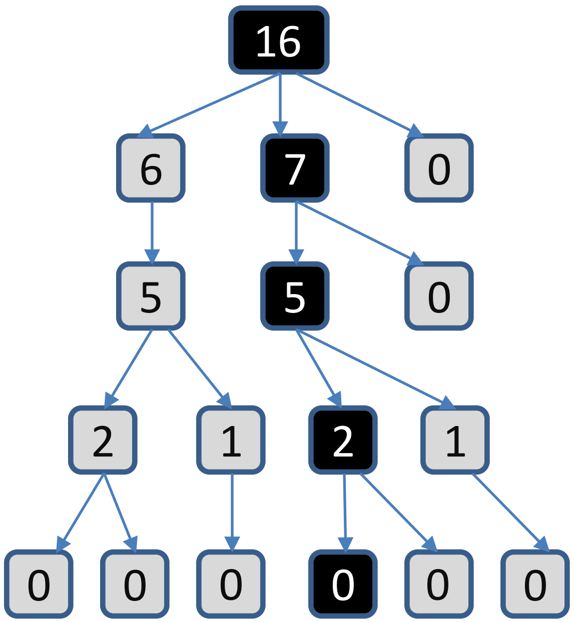

“Heaviest” rather than longest chain

Active Balancing Attack on GHOST

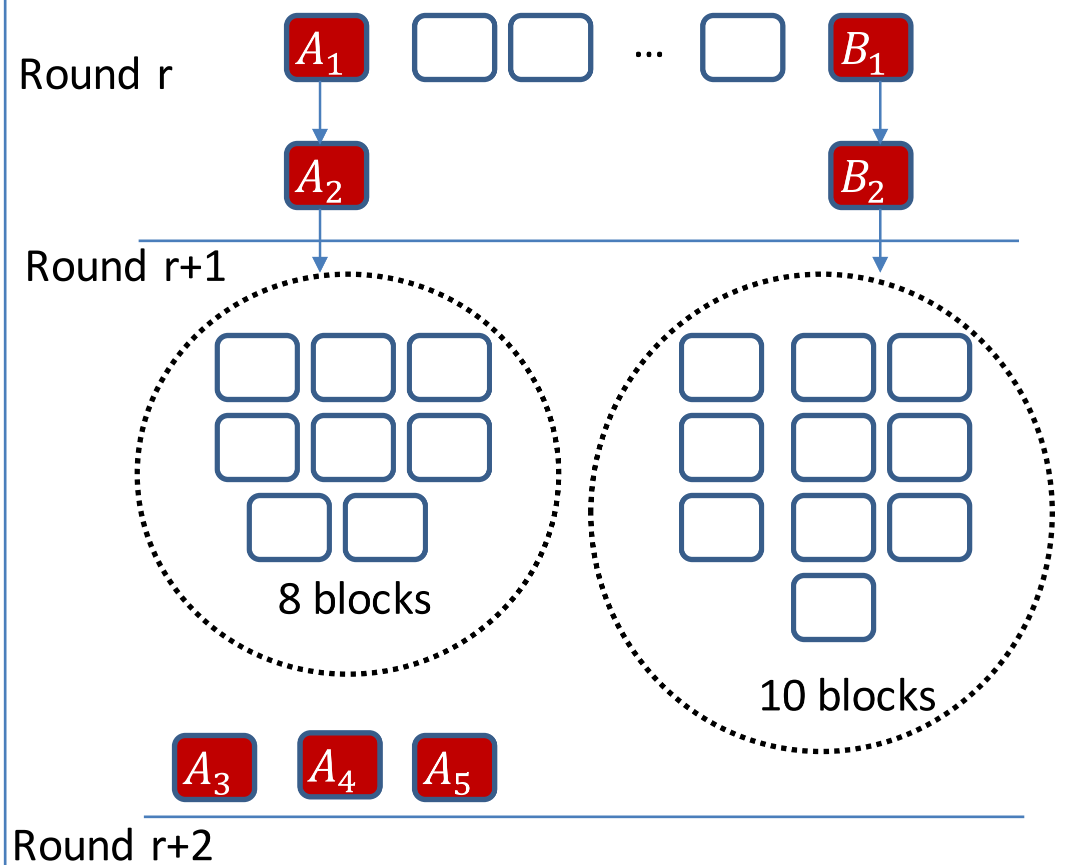

Honest “A miners” Will extend A

Honest “B miners” Will extend B

Secure High-Rate Transaction Processing in Bitcoin – SZ13

Attack Effectiveness on GHOST

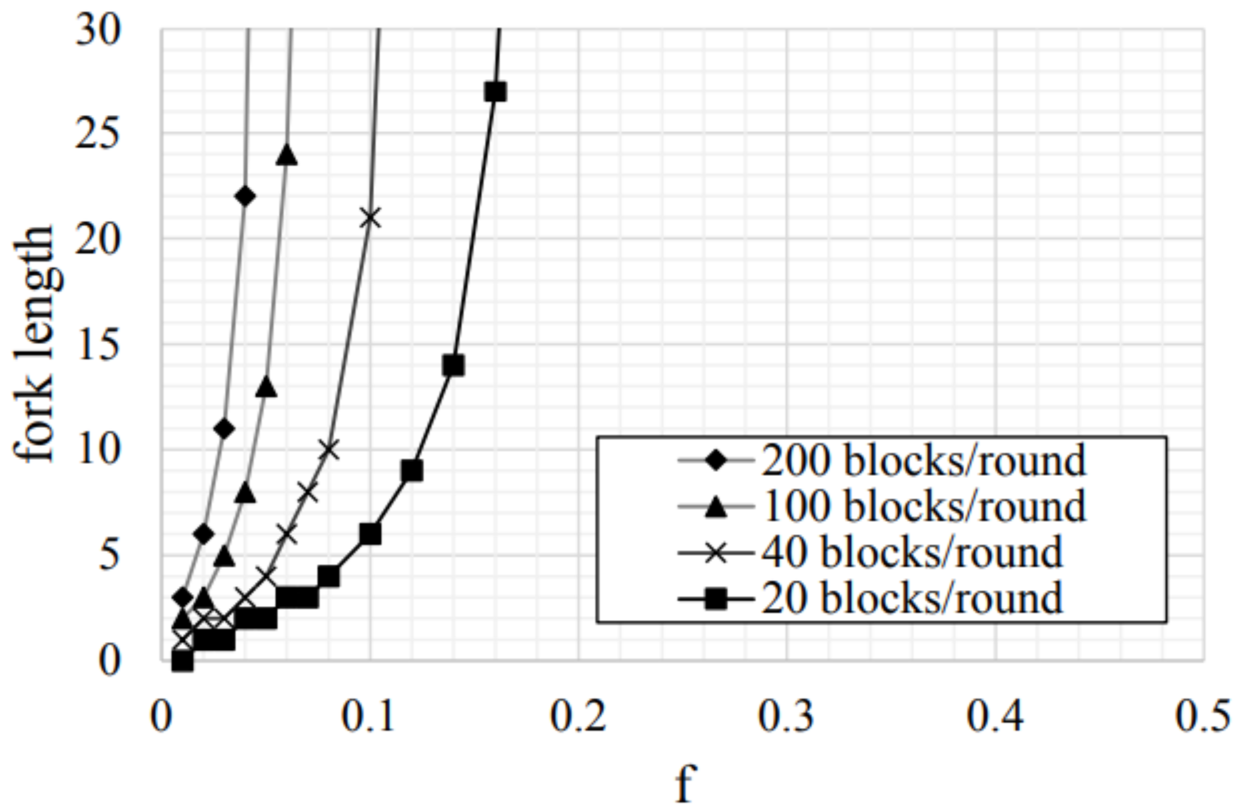

OHIE: Blockchain scaling made simple – Yu, Nikolic, Saxena, Hou

A Principled Approach To Scale Nakamoto

Key Observations

- There is a safe way to run Nakamoto
  - Tolerates f ~ 0.5 if block interval exceeds 3Δ
  - Have established proofs from prior works

- Independence of Design Parameters
  - Block interval depends only on desired f and Δ
  - Confirmation latency depends only on block interval
  - Throughput depends only on available bandwidth (β)
  - Decentralization depends only on number of blocks/sec.

OHIE: Blockchain scaling made simple – Yu, Nikolic, Saxena, Hou

The OHIE Protocol: Run "k" parallel chains!

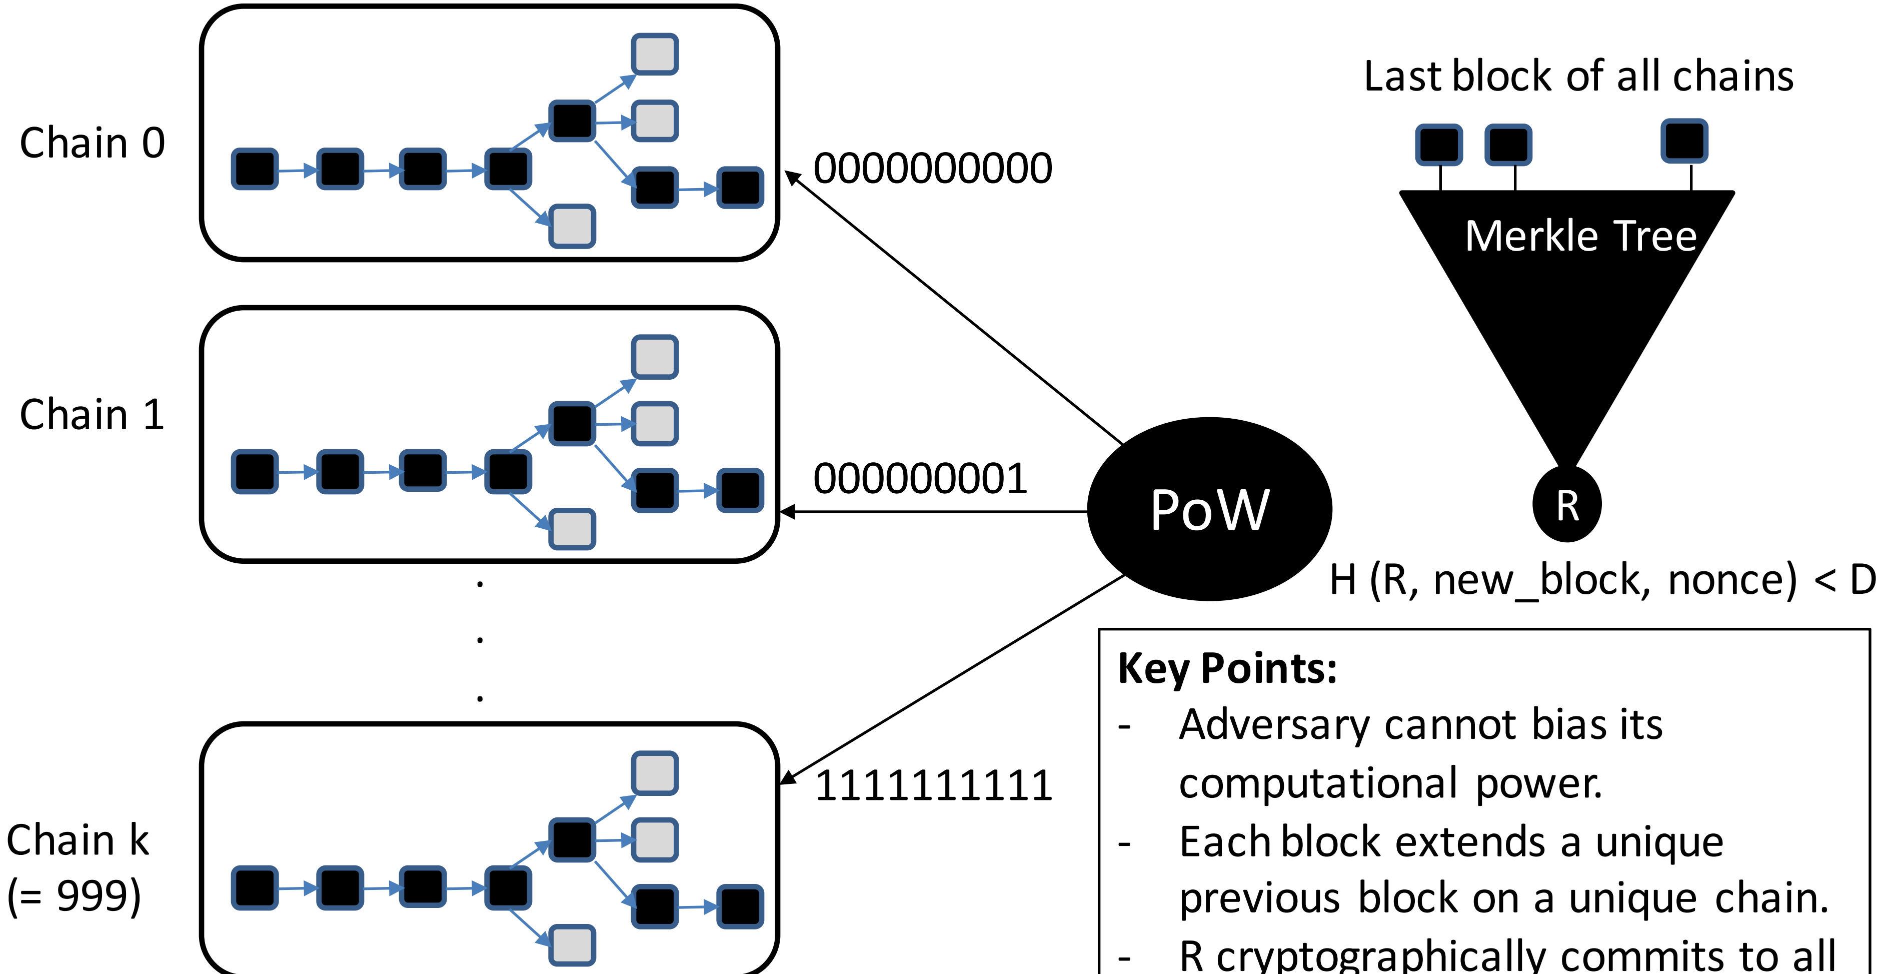

Key Points:

- Adversary cannot bias its computational power.
- Each block extends a unique previous block on a unique chain.
- R cryptographically commits to all "k" blocks that the miner sees.

[OHIE: Blockchain scaling made simple]() – Yu, Nikolic, Saxena, Hou

The OHIE Protocol

- Construction is simple and modular
- Safety and Liveness Proof:
  - Reduces to that of Bitcoin backbone protocol
  - Intuition:
    - Probabilistic process on each chain is identical to Bitcoin
    - Each block extends a single prior block
    - The state that the block extends can't be forged
  - Takes Θ (log k) more confirmation blocks (union bound)

OHIE: Blockchain scaling made simple – Yu, Nikolic, Saxena, Hou

Total Ordering Across Chains?

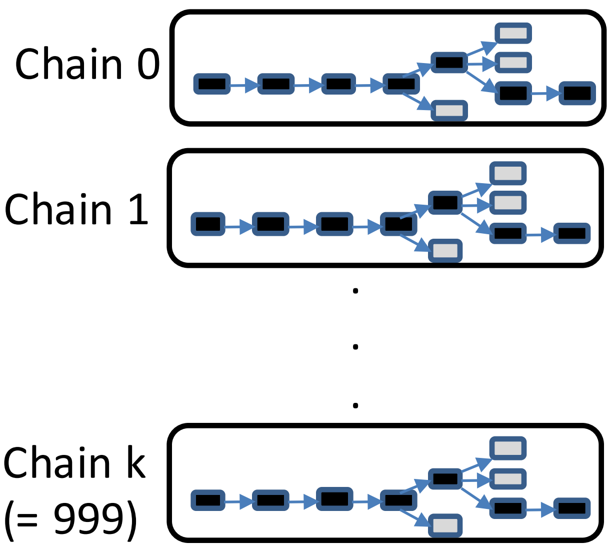

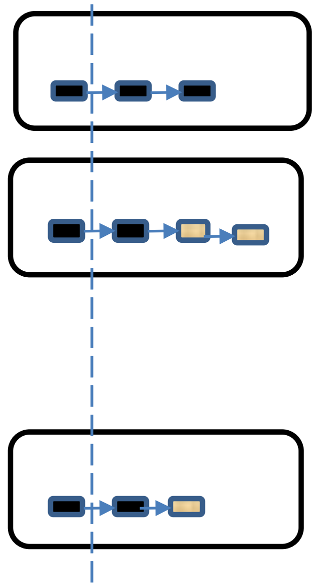

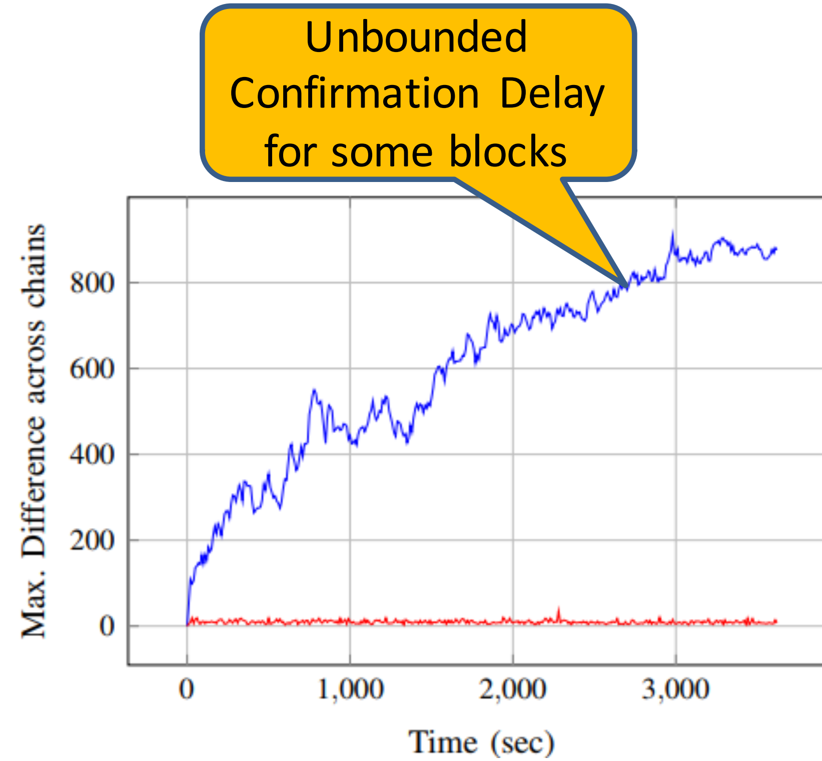

Order by chain ID

Order by Position, tie-breaking on chain ID

[OHIE: Blockchain scaling made simple]() – Yu, Nikolic, Saxena, Hou

Total Ordering Scheme In OHIE

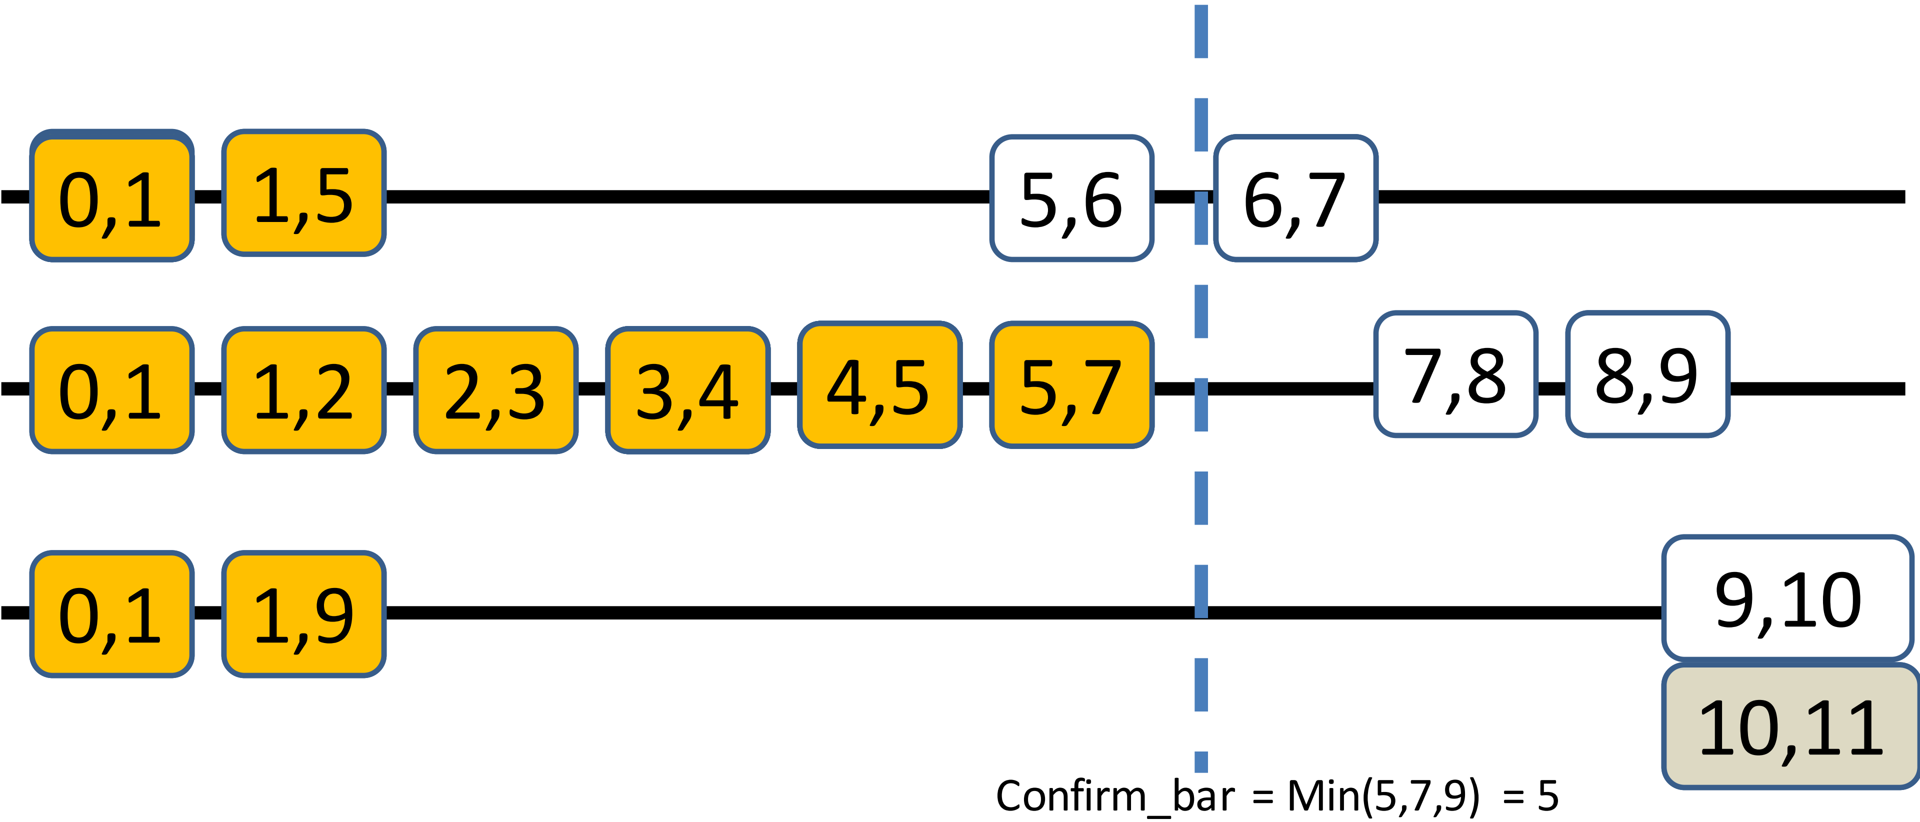

Num. of confirmations per chain (T) = 2

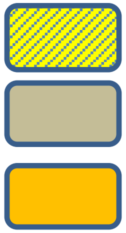

- Partially confirmed (embedded T blocks deep)
- Trailing (max next_rank)
- Fully confirmed

Safety & Liveness
Proofs in paper

OHIE: Blockchain scaling made simple – Yu,Nikolic, Saxena,Hu

Prototype & Micro Experiments

- Less than 5 KLOC of code
- Micro Experiments
  - 1000 miners, 20 Mbps
- Critical Observations:
  - Block propagation delay (Δ) proportional to graph diameter (1-2 seconds)
  - **Parallel broadcasts don't impact latency (Δ)**

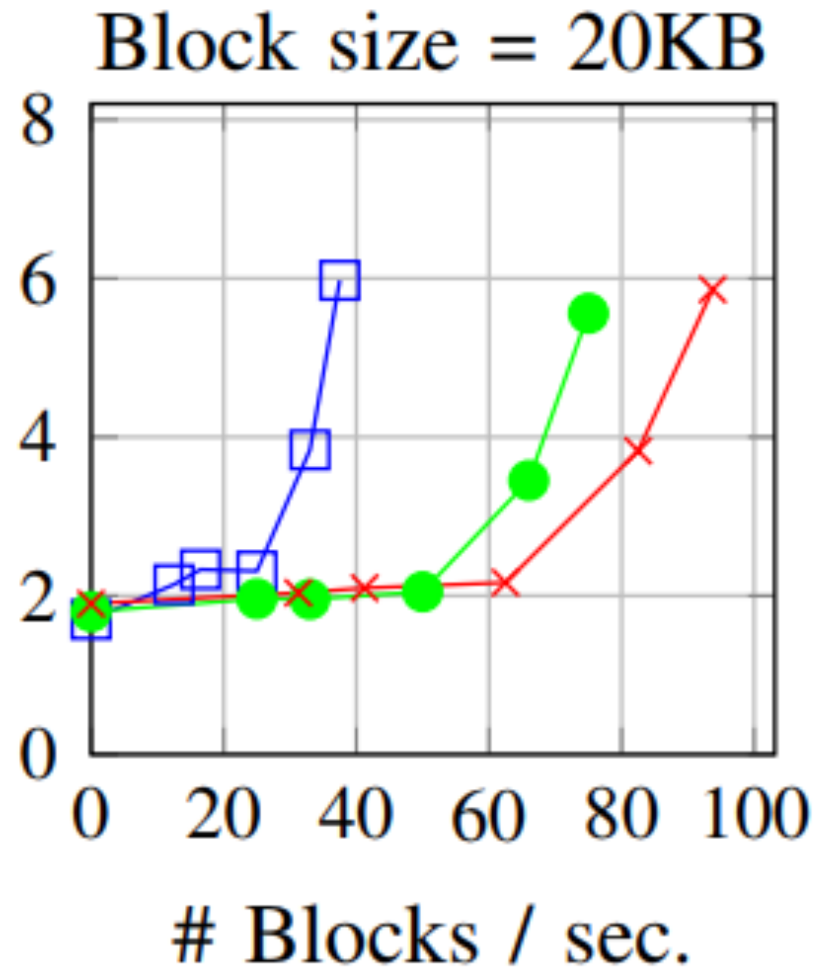

[OHIE: Blockchain scaling made simple]() – Yu, Nikolic, Saxena, Hou

Macro Experiments: Linear Scaling with Available Bandwidth

- 50,000 miners, 20 Mbps, resilience (f) ~ 0.46

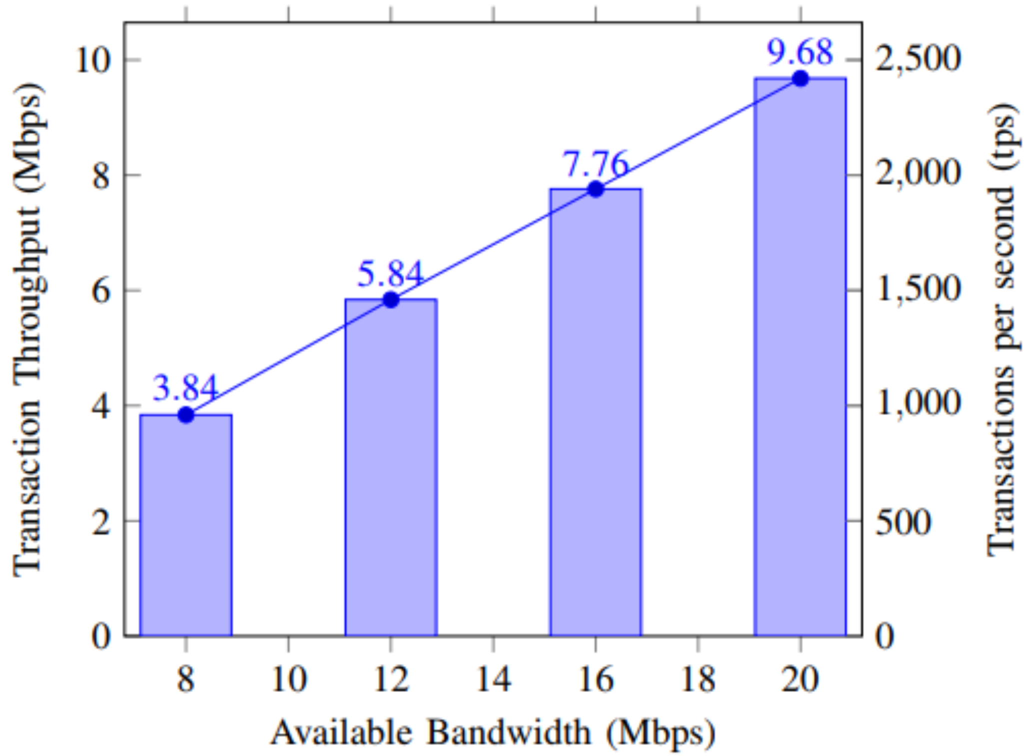

Num. of chains (k) = $0.5 \cdot 3\Delta \cdot \frac{\beta}{block\,size}$

3Δ *implies* $f \sim 0.46$

OHIE: Blockchain scaling made simple – Yu, Nikolic, Saxena, Hou

Macro Experiments: Decentralization

- 50,000 miners, 20 Mbps, f ~ 0.46
- Decentralization: Scales linearly with bandwidth
  - k>60 blocks per second

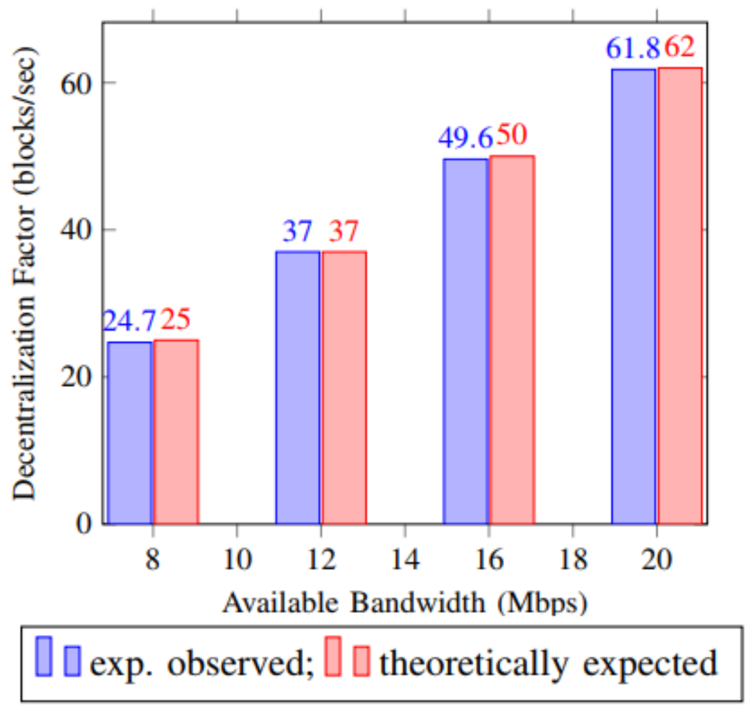

OHIE: Blockchain scaling made simple – Yu, Nikolic, Saxena, Hou

Macro Experiments: Confirmation Delay

- 50,000 miners, f ~ 0.46
- Confirmation Delay
  - **Under 10 minutes** $(3\Delta T)$
  - Independent of throughput! (once we fix "k")

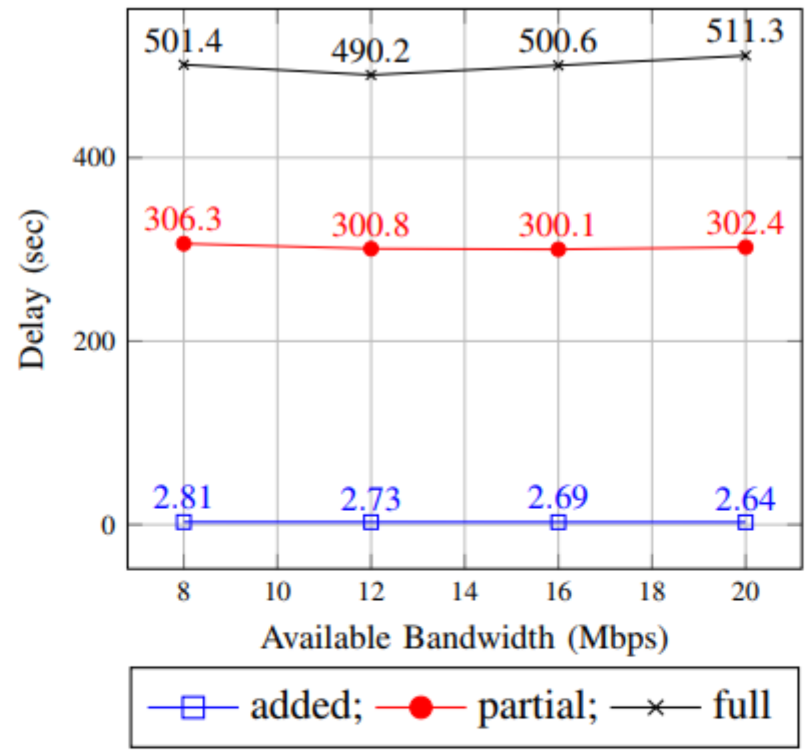

- Conf. Blks (T) = 15 - 30
  - $T_{BTC} + \Theta$(log k)

OHIE: Blockchain scaling made simple – Yu, Nikolic, Saxena, Hou

Security vs. Performance: State-of-the-art

Approach	Resilience	Throughput	Decentralization	Latency
Nakamoto with reduced block intervals	$f < \frac{1}{3}$	Low	Medium	Good
Nakamoto with large blocks	$f < \frac{1}{2}$	High	Low	Medium
AlgoRand (with BA) [SOSP'17]	$f < \frac{1}{5}$	High	Low	Good

60 proposers per sec

30 secs.

10 mins

State-of-the-practice

Repurposing BA Protocols?

- Agree on 1 block per round with standard BFT / BA
- Honest miners sign that block with round id.

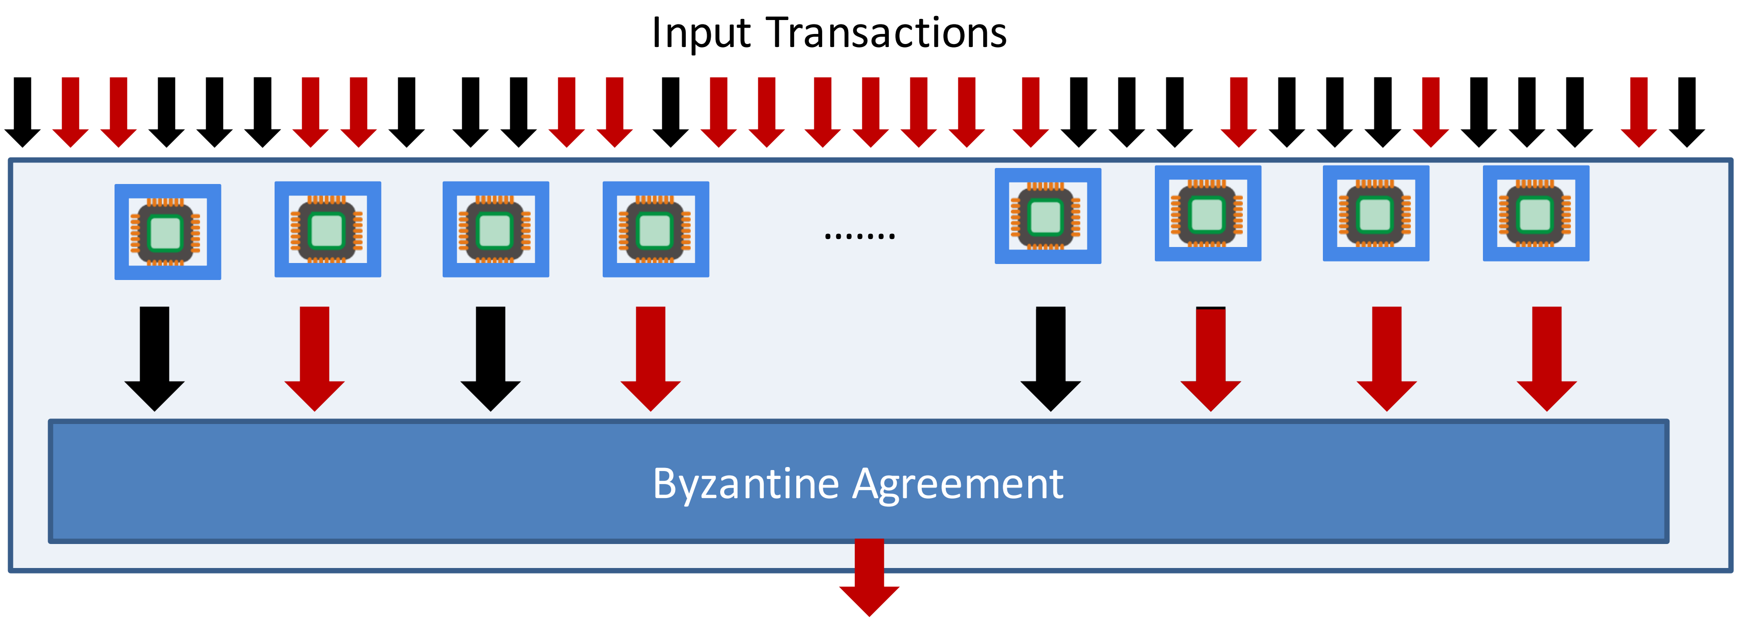

- Challenge: Participants must be <u>known a-priori</u>
  – Chicken-n-egg: Agreeing on participants is itself…

Proof-of-work for Sharding

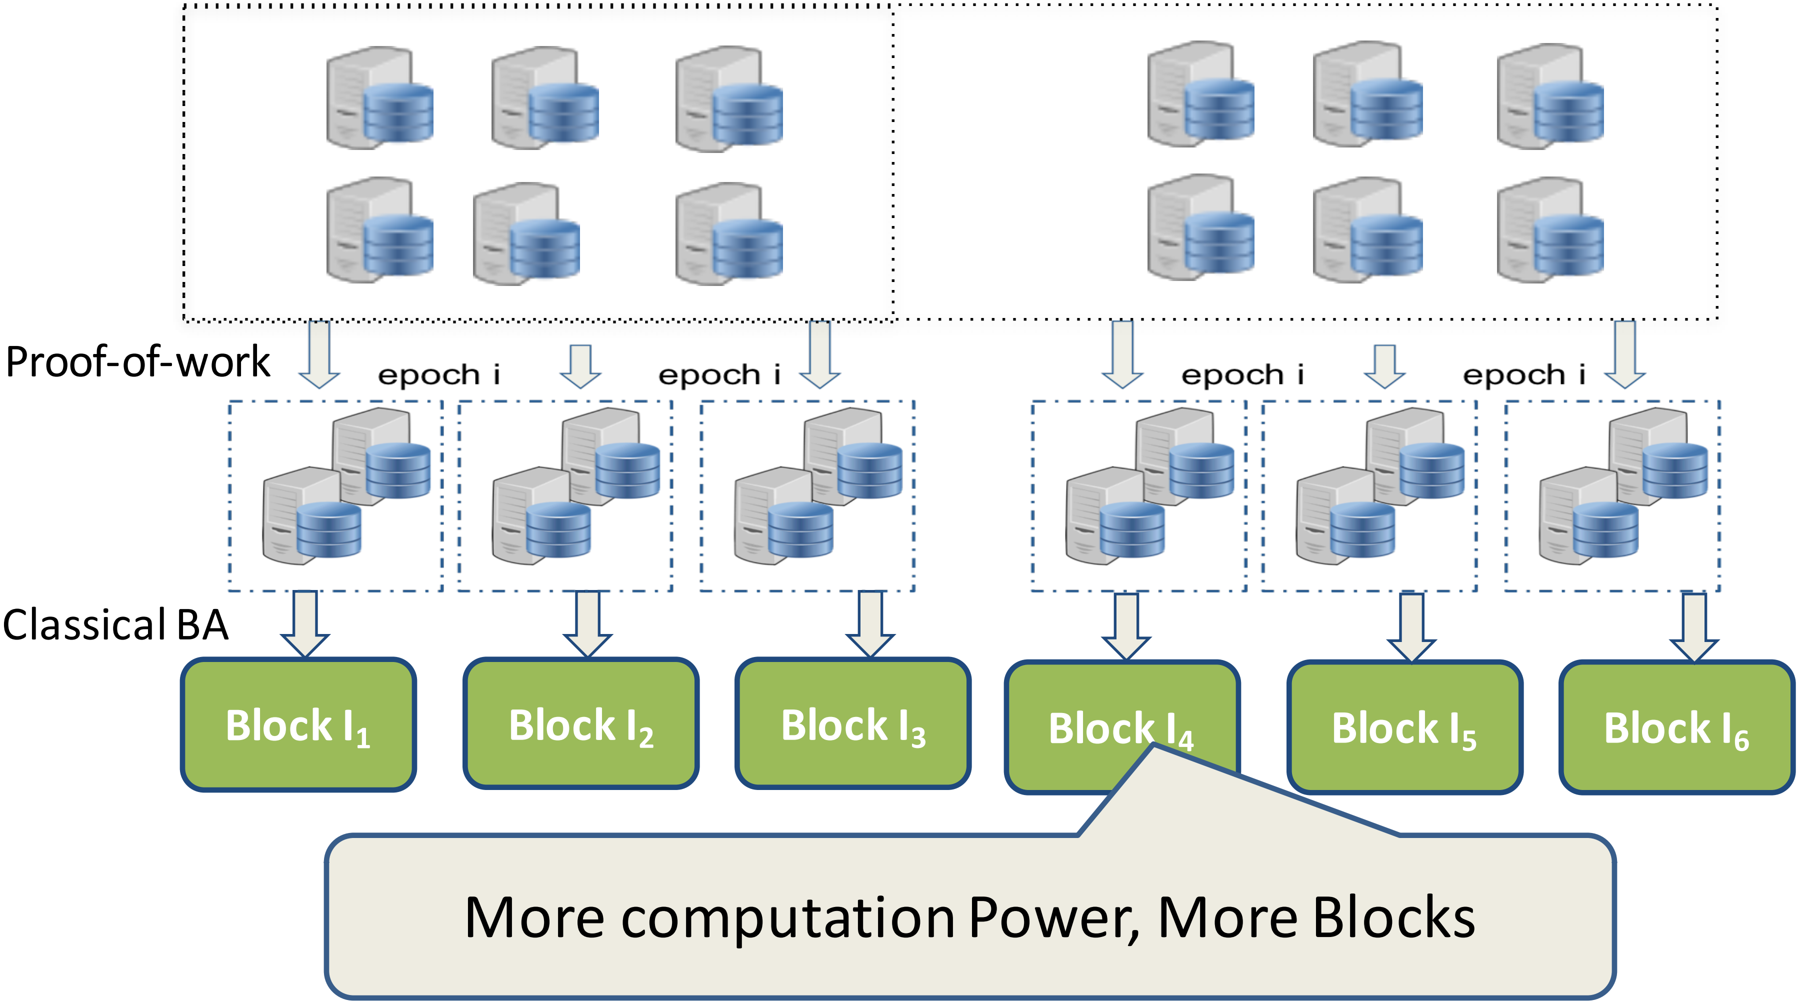

Elastico – CCS'16 (Also see Omniledger – Oakland'18, RapidChain-CCS'18)

Commercialized as the Zilliqa public blockchain platform

Open to public mining (Feb 2019)

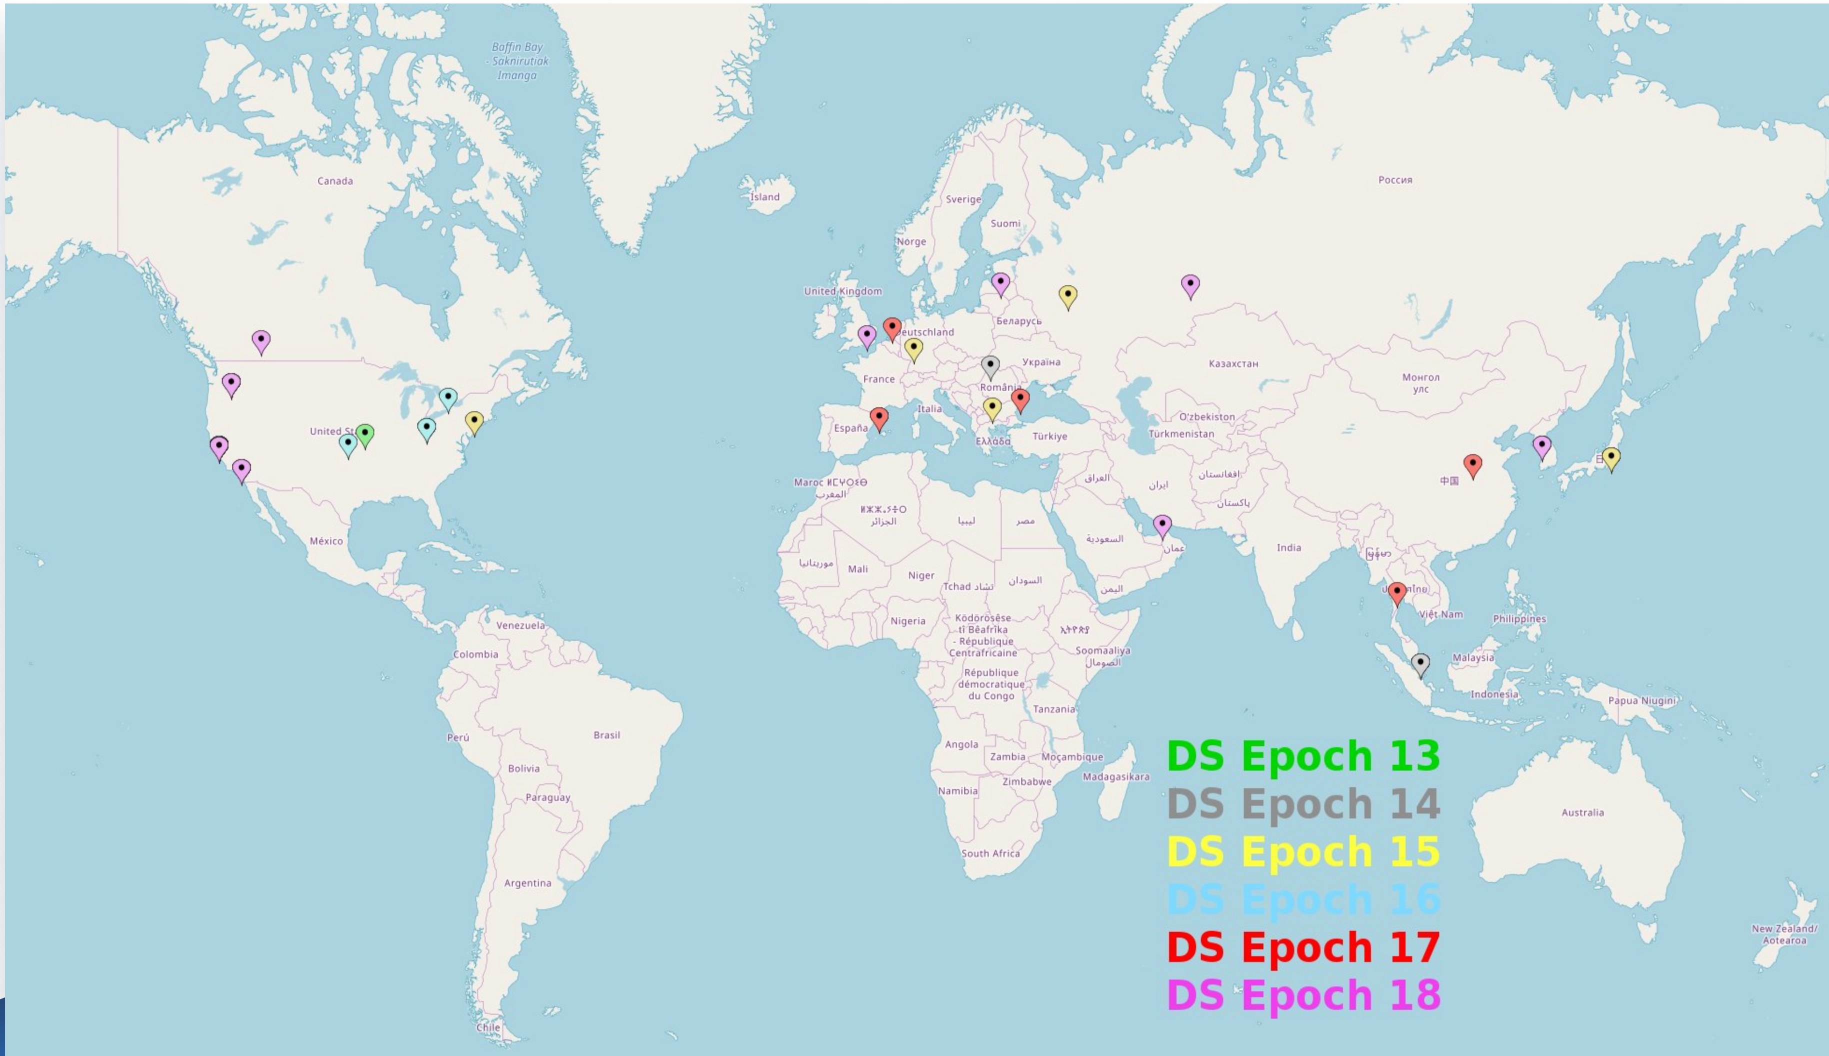

Takeaways

- Decentralized Systems propose exciting algorithmic problems
  - Build better crypto, distributed algorithms, verification tools, …
- Is there an Optimal Consensus Protocol?
  - Latency $\Theta(\Delta)$, Throughput $\Theta(\beta)$, Decentralization $\Theta(\beta)$, Res. f ~ 0.5
  - Simplicity
  - Improve the constants
- Need for new models and drawing new connections:
  - Consistency & Isolation properties offered by blockchains
  - Sybil resistance mechanisms: Proof-of-Stake vs. Proof-of-Work
  - Incentive mechanism design: Fairness, Variance,…
  - Trusting Off-chain computations

Thank you!

Collaborators:

- Loi Luu (PhD, NUS & CEO – Kyber Network)
- Haifeng Yu (Prof, NUS)
- Ivica Nikolic (Postdoc, NUS)
- Seth Gilbert (Prof., NUS)
- Hrishi Olickel (UG, Yale-NUS)
- Roumu Hou (UG, NUS)