Mysticeti

Reaching the Limits of Latency with Uncertified DAGs

NDSS 2025

Andrey Chursin
MystenLabs

Starting concepts

Do you know?

- How blockchain(roughly) works?
- What is Byzantine Fault Tolerance?
- What is consensus protocol?
- Proof-of-Work and Proof-of-Stake
- What is mempool in a blockchain?

Proof-of-Stake consensus

- List of validators is known and each validator has assigned “stake”
- Network and machine failures are tolerated
- Up to ~33%(f) validator stake can run arbitrary code
- At least ~66% (2f+1) behave according to the protocol
- All correct validators eventually agree on the same “state”
- “State” can mean different things
  - List of transactions and result of the execution
  - Just list of transactions
  - Something else..

“Traditional consensus”

- PBFT (~TowerBFT)
- Tendermint (CometBFT, …)
- HotStuff (DiemBFT, AptosBFT, MonadBFT, …)

- Low leader commit latency
- Commit transactions from the leader only
- Does not specify transaction dissemination
- Transaction latency depends on mempool implementation

Transaction finality

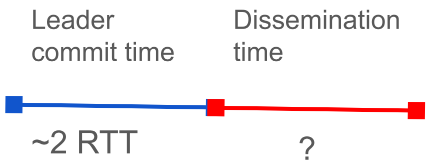

DAG-based mempool: Narwhal

Traditional consensus(HotStuff) + Narwhal mempool: Diem, Aptos

- Two different kind of blocks - DAG mempool blocks and consensus blocks
- Block in DAG contains transactions from validator and links to other blocks
- DAG-mempool provides Byzantine fault-tolerant way to fetch DAG block content and dependencies by hash
- Consensus agrees on opaque “hash” provided by DAG-mempool
- High throughput(100k+ TPS), latency consensus + mempool

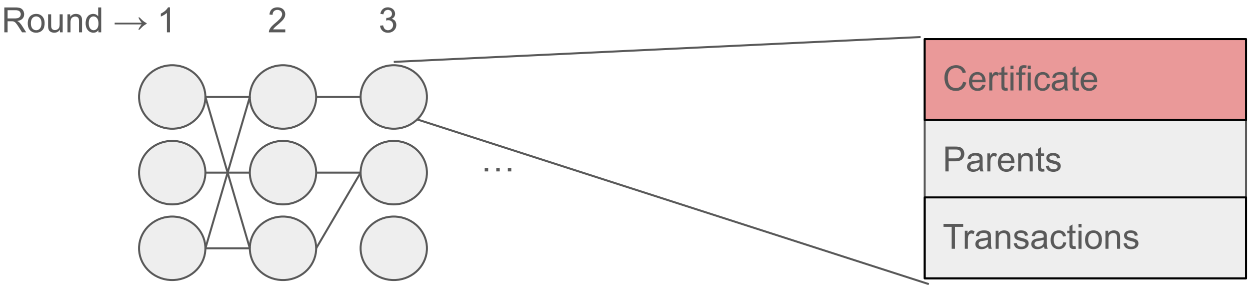

Dag mempool + consensus: Narwal/Bullshark

- Apparently, once you have DAG-mempool you don’t need separate consensus!
- Bullshark takes narwhal DAG and derives total order
- No additional network messages, all you need is DAG
- High throughput - 200k+ TPS, latency in seconds though

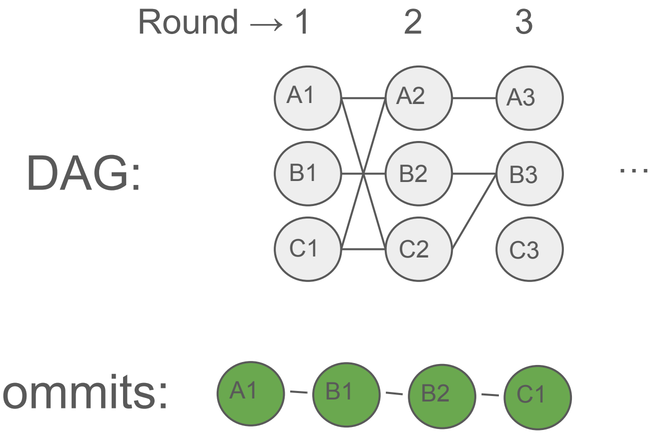

Uncertified DAG: Mysticeti

Key differences:

- Individual DAG elements are not certified - 0.5 RTT to make DAG element (vs 1.5+ RTT in Narwhal)
- Need minimum 3 rounds to make commit (vs 2 for Narwhal)
- Overall anchor commit latency is lower - 0.5*3=1.5, 1.5*2=3

Results:

- High TPS of DAG-based mempool
- Low transaction commit latency of traditional consensus
- More complex consensus

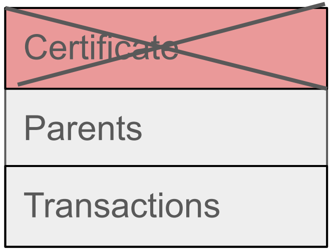

Transaction finality

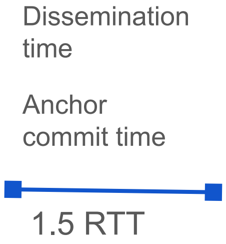

TLDR - “Traditional” and DAG-based consensus

“Traditional” BFT consensus(Tendermint, HotStuff, etc):

- Mempool is a separate component(out of scope for consensus)
- Good consensus commit latency, without accounting for mempool latency
- Throughput very limited

DAG-based consensus(Narwhal[2021]):

- Very high throughput comparing to traditional consensus
- Integrated mempool
- Latency much higher than HotStuff/Tendermint

TLDR - Mysticeti combines benefits of both!

- As a DAG-based consensus Mysticeti has integrated mempool and provides even higher throughput then narwhal (500K transactions and higher)
- With a few techniques(more on this later) Mysticeti achieves latency similar to that of well implemented traditional consensus protocol such as Tendermint/HotStuff.
- Mysticeti has integrated FastPath - replacement for the FastPay[2020] consensus-less protocol.

Implementation details

- Implementation written in Rust
- (Almost) entire protocol is just one(!) long poll RPC
- Simple networking - TCP sockets, async IO
- Synchronous consensus core
- Ed25519 signatures
- Wal based storage to fully utilize disk IO and reduce write amplification
- Tested in a real cluster of 100+ nodes distributed across the world

Consensus-only performance

- Sub-second latency with up to 300k TPS
- Up to 500K TPS with 2s latency

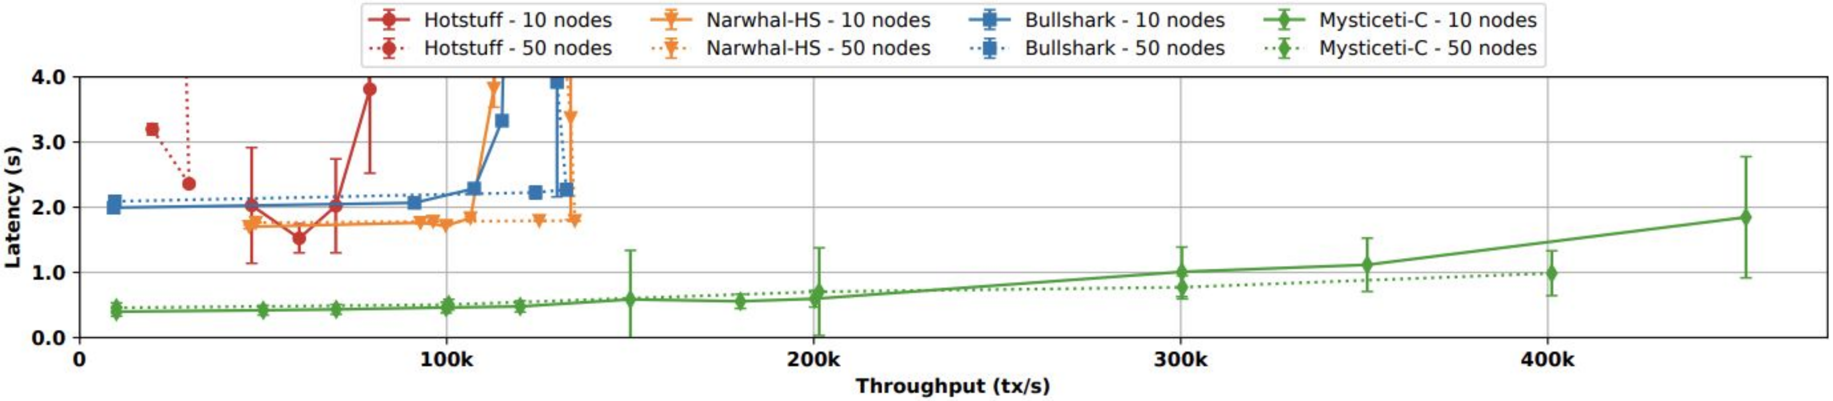

Fig. 5: Throughput-Latency graph comparing MYSTICETI-C performance with state-of-the-art consensus protocols.

Practical application

- SUI is layer 1 blockchain, #12 market cap as of today.
- SUI launched with Narwhal in May 2023
- Mysticeti research started June 2023, promising results by Sep 2023
- SUI fully migrated to Mysticeti in fall 2024
- 0 forks and 1 availability incident since then

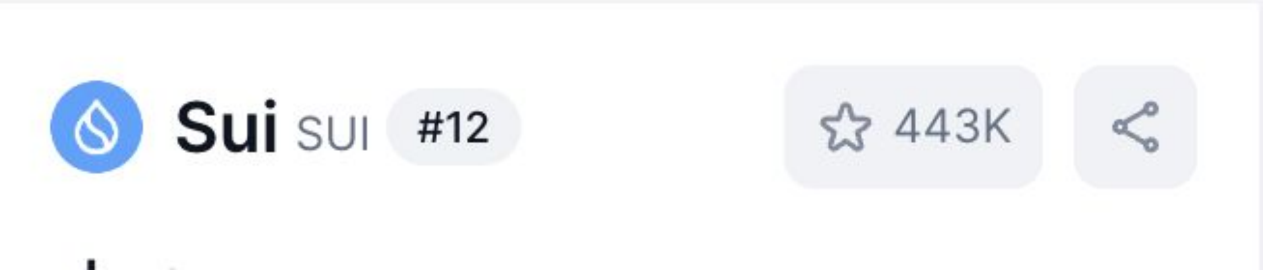

Sui SUI #12

☆ 443K

Production performance in SUI blockchain

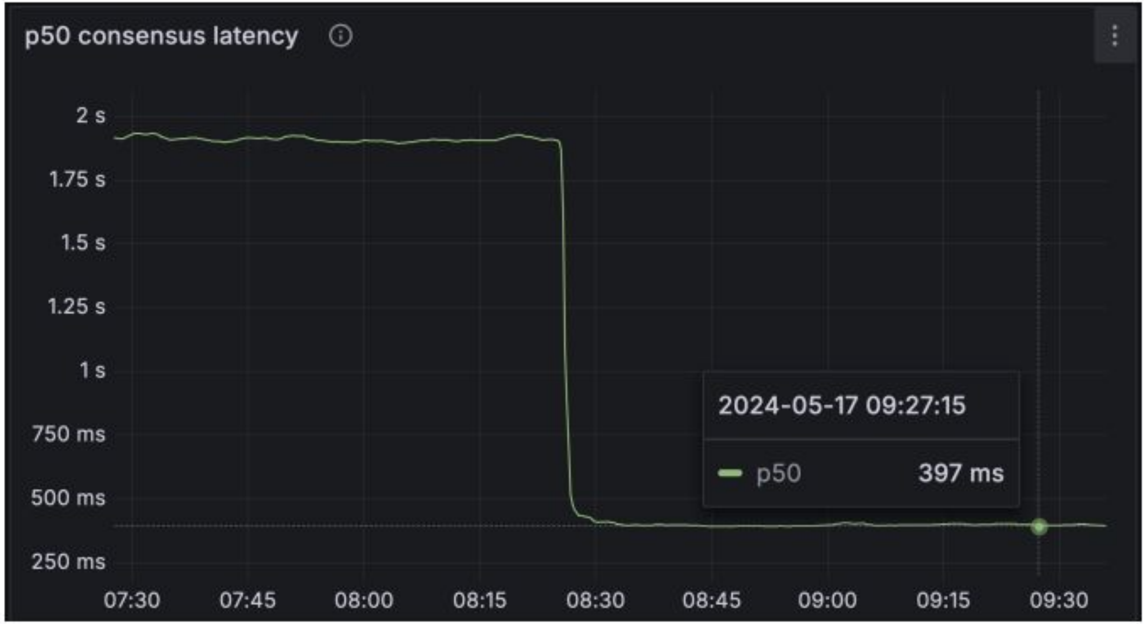

Fig. 1: P50 latency of a major blockchain switching from Bullshark (1900ms) to MYSTICETI-C (390ms) consensus on 106 independently run validators

Extra: not just consensus

- Mysticeti has a variant Mysticeti-FPC (FastPath + Consensus)
- Can finalize some transactions even before consensus finality is reached
- Based on generalisation of idea from FastPay[2020] paper
- Introduces new “messages” inside DAG blocks and leverages existing DAG structure
- No additional RPC/network communication aside from adding more data into the DAG

Thank you!