

# Algorand: Scaling Byzantine Agreements for Cryptocurrencies

Ikechi Iwuagwu




# Introduction

- Current cryptocurrencies suffer a trade-off between latency and confidence in a transaction.
- Applications that require low latency cannot be sure that their transaction will be confirmed, so they must trust the payer to not double-spend.



# Introduction

- Algorand is a cryptocurrency designed to confirm transactions in under a minute.
- It uses a Byzantine agreement protocol called *BA★* that is scalable and allows Algorand to reach consensus on a new block with low latency and without the possibility of forks.



# Introduction

- A key technique that makes BA★ suitable for Algorand is its use of verifiable random functions (VRFs) to randomly select users in a private and non-interactive way.



# Challenges

- Algorand must avoid Sybil attacks
- BA★ must scale to millions of users
- Algorand must be resistant to DoS attacks and continue to function even when a malicious user disconnects some of the other users.



# Techniques

- Weighted users
  - Assigns weights to users based on the money in their account.
  - $BA^\star$  guarantees consensus as long as a weighted fraction (greater than  $2/3$ ) of users are honest.



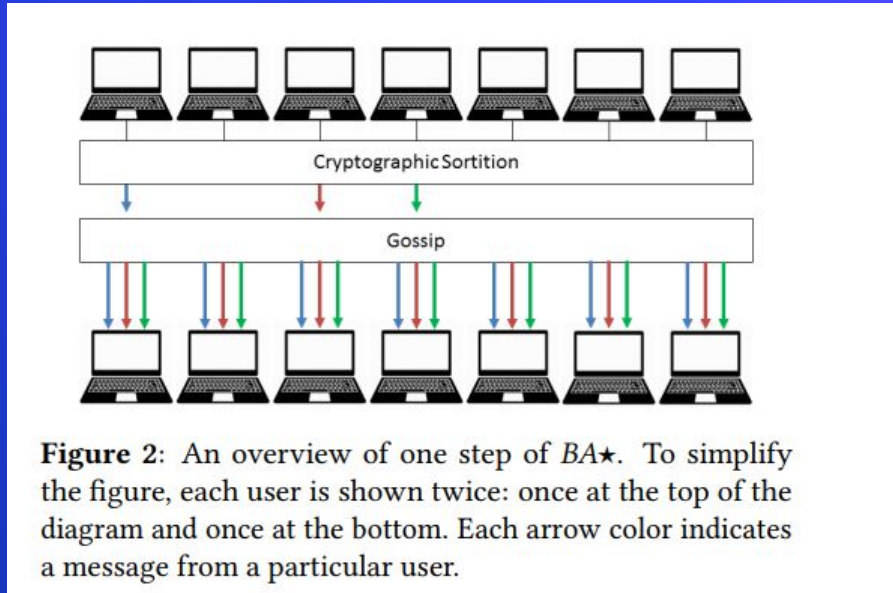
# Techniques

- Consensus by committee
  - BA★ achieves scalability by choosing a committee to run each step of its protocol.
  - BA★ chooses committee members randomly among all users based on the users' weights.



# Techniques

- Cryptographic sortition





# Techniques

- Participant replacement
  - $BA\star$  requires committee members to only speak once.



# Goals

- Safety goal
  - All users agree on the same transactions.
  - Same holds for disconnected users.



# Goals

- Liveness goal
  - Algorand allows new transactions to be added to the blockchain.
  - The aim is to reach consensus on a new set of transaction in about 1 minute.




# Assumptions

- Strong synchrony assumption
  - Algorand assumes that most honest users can send messages that will be received by other other honest users.



# Assumptions

- Weak synchrony assumption
  - The network can be controlled by an adversary for a long period of time.
  - But after, the network must be strongly controlled by honest users again for a long period of time to ensure safety.



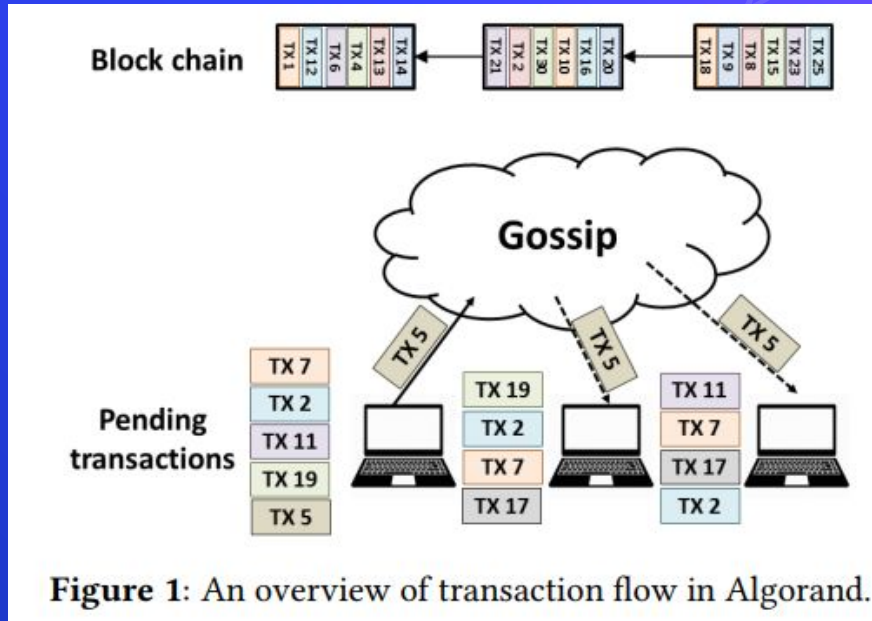
# Assumptions

- Loosely synchronized clocks assumption
  - Clocks must be close enough in order to start recovery protocol.



# Gossip Protocol

- Algorand users communicate through a gossip protocol.



# Consensus

- BA★ can produce two kinds of consensus
  - Final Consensus
  - Tentative Consensus






# Latency


- Algorand can confirm transactions in under a minute.
- The latency is constant as the number of users grows.



# Latency




**Figure 5:** Latency for one round of Algorand, with 5,000 to 50,000 users.




**Figure 6:** Latency for one round of Algorand in a configuration with 500 users per VM, using 100 to 1,000 VMs.

# Throughput

- Algorand's agreement time is independent of the block size.
- As Algorand's block size grows, Algorand achieves higher throughput at the cost of some increase to latency.



# Throughput



**Figure 7:** Latency for one round of Algorand as a function of the block size.

# CPU, Bandwidth, Storage

- The CPU costs of running are modest.
- Bandwidth costs are about 10 Mbit/second.
- Algorand stores block certificates in addition to blocks themselves.
  - Each certificate is 300 Kbytes.




# Adversaries/Timeout Parameters

- Algorand is not significantly affected by misbehaving users.
- BA★ finishes in under 20 seconds regardless of malicious users.



# Adversaries/Timeout Parameters



**Figure 8:** Latency for one round of Algorand with a varying fraction of malicious users, out of a total of 50,000 users.

# References

Yossi Gilad, Rotem Hemo, Silvio Micali, Georgios Vlachos, and Nickolai Zeldovich. 2017.

Algorand: Scaling Byzantine Agreements for Cryptocurrencies. In Proceedings of the 26th Symposium on Operating Systems Principles (SOSP '17). Association for Computing Machinery, New York, NY, USA, 51–68.

DOI:<https://doi.org/10.1145/3132747.3132757>