Constellation: Peer-to-Peer Overlays for Federated Byzantine Agreement Systems

Giuliano Losa[*], Yifan Mao[+], Shaileshh Bojja Venkatakrishnan[+], and Yunqi Zhang[+]

[*]Stellar Development Foundation, [+]Ohio State University

Overlay networks allow scalable communication in large-scale blockchains

Blockchain protocols often need to broadcast

- Sending to all one-by-one is not scalable
- A solution: gossip protocol in a logical overlay network (each nodes sends to its neighbors, which send to their neighbors, etc)

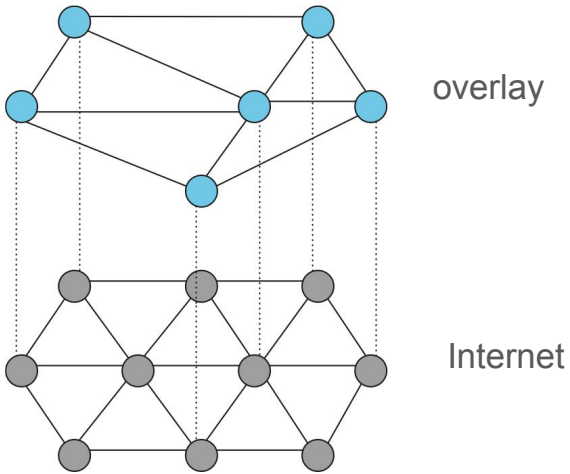

Figure credits: https://book.systemsapproach.org/applications/overlays.html

Overlay networks allow scalable communication in large-scale blockchains

Blockchain protocols often need to broadcast

- Sending to all one-by-one is not scalable
- One solution: gossip protocol in a logical overlay network (each nodes sends to its neighbors, which send to their neighbors, etc)

Attacks that disconnect the overlay (e.g. eclipse attacks) can be disastrous:

- In fully-permissionless/dynamically-available systems (e.g. Bitcoin, Ouroboros), they can cause forks
- In partially-synchronous systems (e.g. Tendermint, Stellar), they can stop the whole system

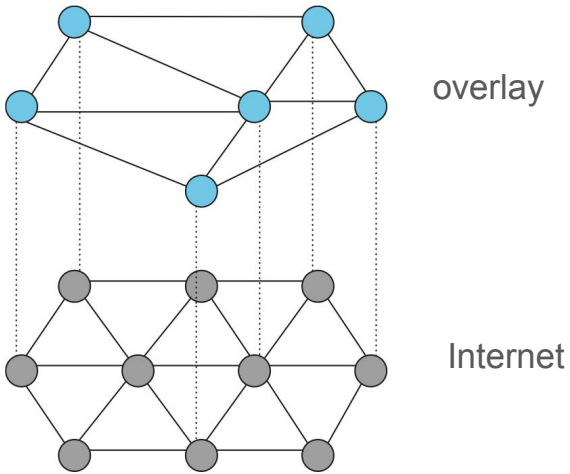

Figure credits: https://book.systemsapproach.org/applications/overlays.html

Goal: construct efficient overlays matching the fault-tolerance of consensus algorithms

Efficiency goals

- Minimize degree, to minimize redundant traffic in gossip dissemination algorithms
- Keep diameter low to minimize latency

Fault tolerance/security goals

- Remain connected under all possible failures allowed by the consensus model
  - E.g., if consensus assumes *f* out of *n* nodes may be faulty, then removing any *f* nodes should not disconnect the overlay

Goal: construct efficient overlays matching the fault-tolerance of consensus algorithms

Efficiency goals

- Minimize degree, to minimize redundant traffic in gossip dissemination algorithms
- Keep diameter low to minimize latency

Fault tolerance/security goals

- Remain connected under all possible failures allowed by the consensus model
  - E.g., if consensus assumes *f* out of *n* nodes may be faulty, then removing any *f* nodes should not disconnect the overlay
- In this work, we consider blockchains using the Federated Byzantine Agreement (FBA) model

In the Federated Byzantine Agreement (FBA) model, participants collectively determine tolerated failures

In the Federated Byzantine Agreement (FBA) model, participants collectively determine tolerated failures

- Each node makes unilateral agreement requirements

In the Federated Byzantine Agreement (FBA) model, participants collectively determine tolerated failures

- Each node makes unilateral agreement requirements

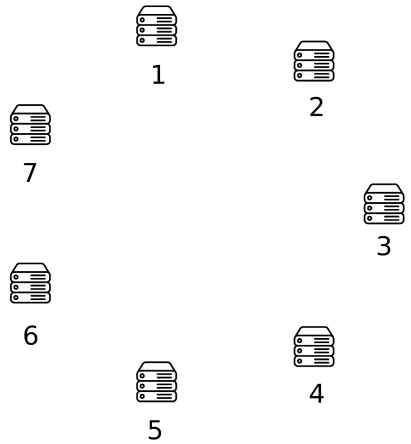

In the Federated Byzantine Agreement (FBA) model, participants collectively determine tolerated failures

- Each node makes unilateral agreement requirements

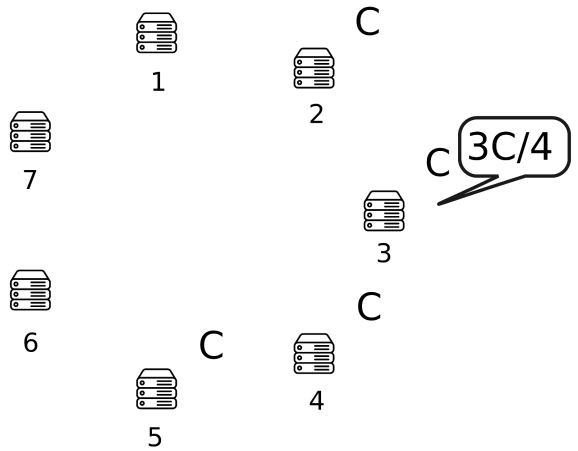

In the Federated Byzantine Agreement (FBA) model, participants collectively determine tolerated failures

- Each node makes unilateral agreement requirements

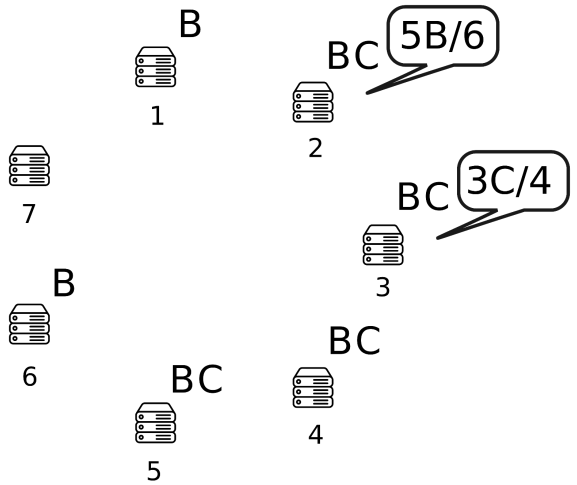

In the Federated Byzantine Agreement (FBA) model, participants collectively determine tolerated failures

- Each node makes unilateral agreement requirements

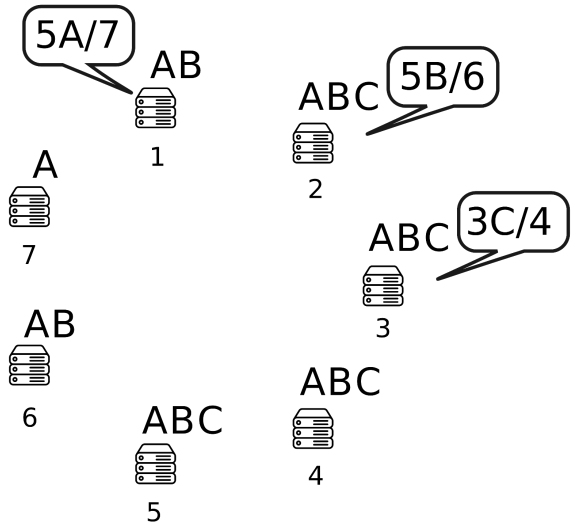

In the Federated Byzantine Agreement (FBA) model, participants collectively determine tolerated failures

- Each node makes unilateral agreement requirements

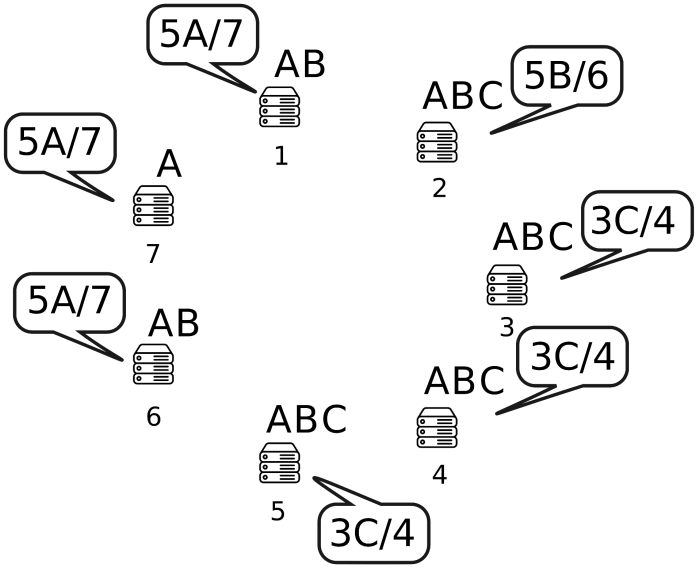

In the Federated Byzantine Agreement (FBA) model, participants collectively determine tolerated failures

- Each node makes unilateral agreement requirements
- Define a quorum as a set that satisfies the requirements of all its members

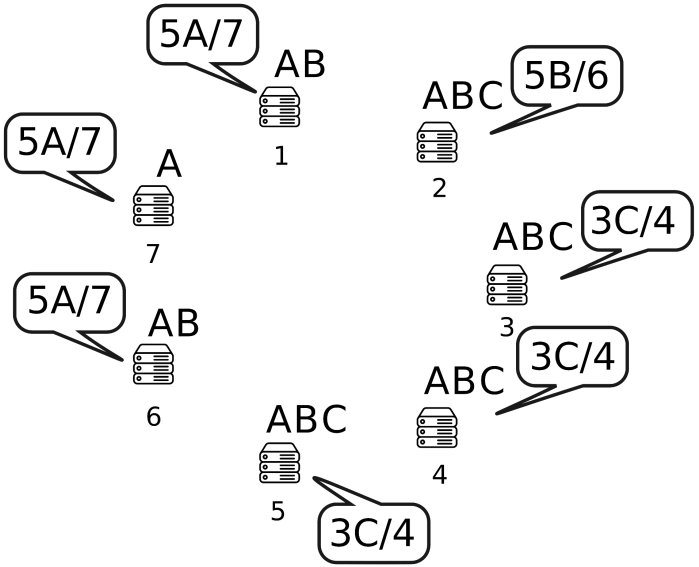

In the Federated Byzantine Agreement (FBA) model, participants collectively determine tolerated failures

- Each node makes unilateral agreement requirements
- Define a quorum as a set that satisfies the requirements of all its members

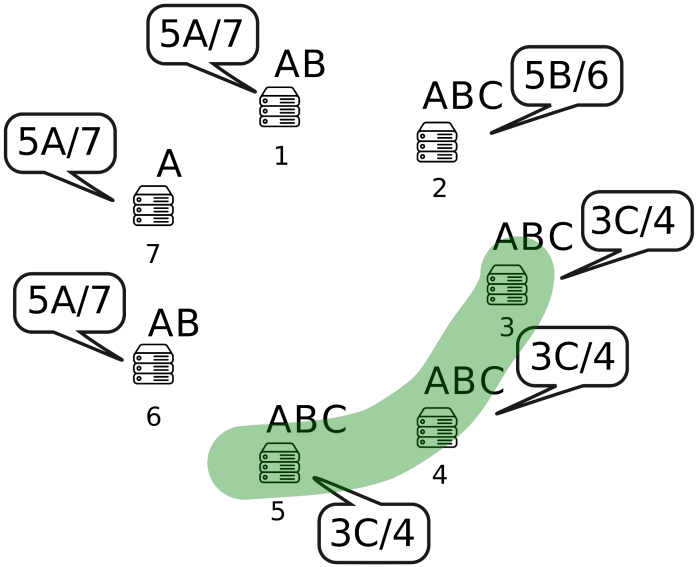

In the Federated Byzantine Agreement (FBA) model, participants collectively determine tolerated failures

- Each node makes unilateral agreement requirements
- Define a quorum as a set that satisfies the requirements of all its members

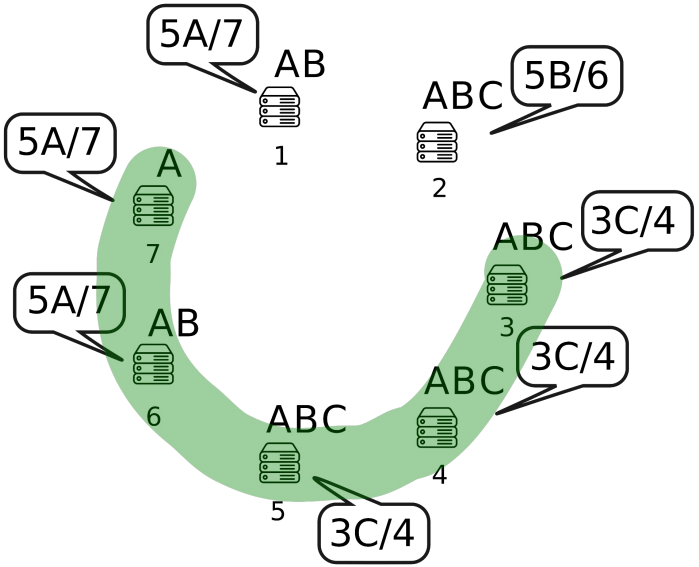

In the Federated Byzantine Agreement (FBA) model, participants collectively determine tolerated failures

- Each node makes unilateral agreement requirements
- Define a quorum as a set that satisfies the requirements of all its members

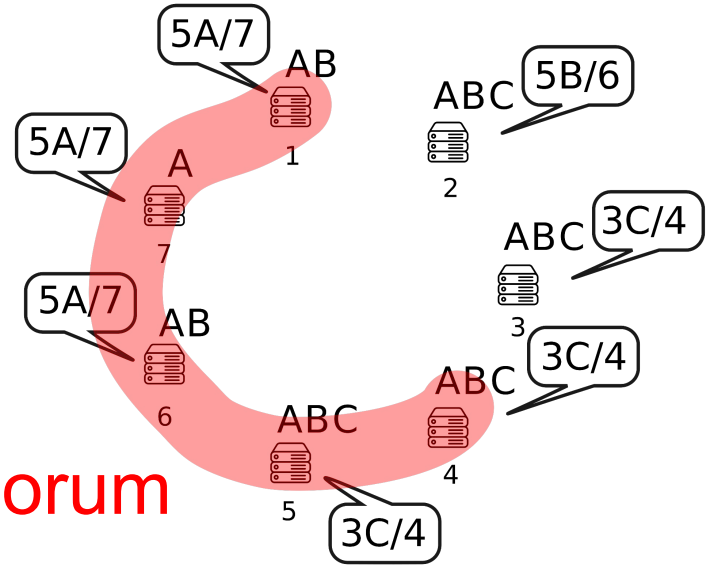

In the Federated Byzantine Agreement (FBA) model, participants collectively determine tolerated failures

- Each node makes unilateral agreement requirements
- Define a quorum as a set that satisfies the requirements of all its members

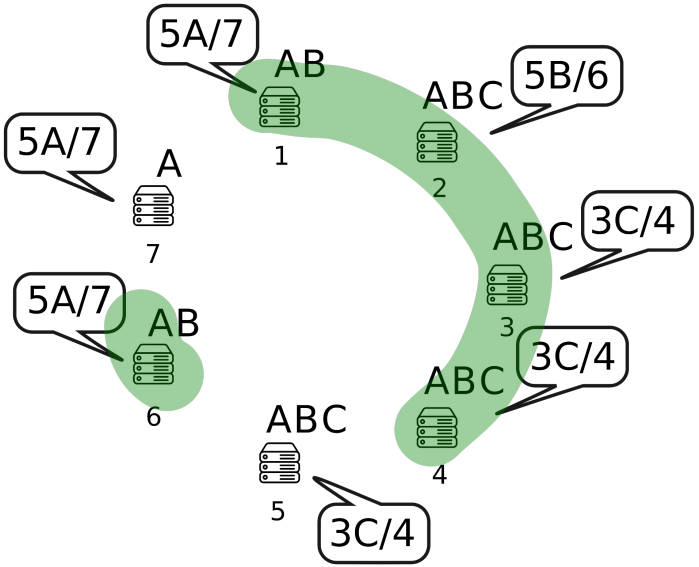

In the Federated Byzantine Agreement (FBA) model, participants collectively determine tolerated failures

- Each node makes unilateral agreement requirements
- Define a quorum as a set that satisfies the requirements of all its members
- Failure assumption: at least one quorum is honest and well-behaved

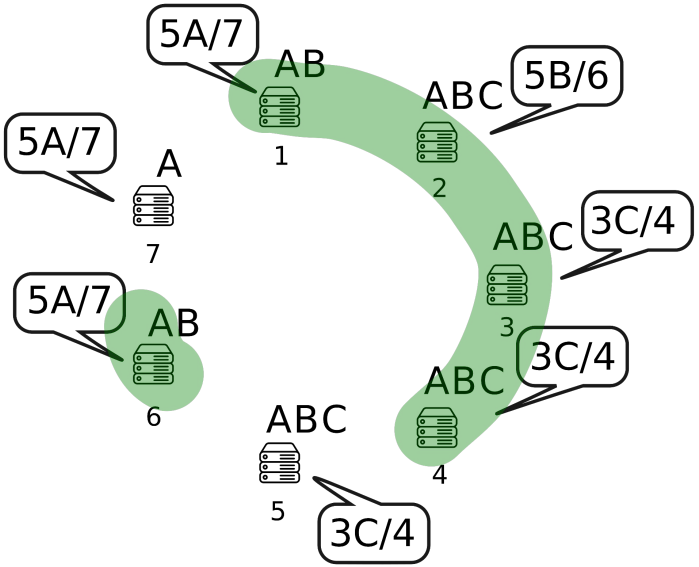

In the Federated Byzantine Agreement (FBA) model, participants collectively determine tolerated failures

- Each node makes unilateral agreement requirements
- Define a quorum as a set that satisfies the requirements of all its members
- Failure assumption: at least one quorum is honest and well-behaved

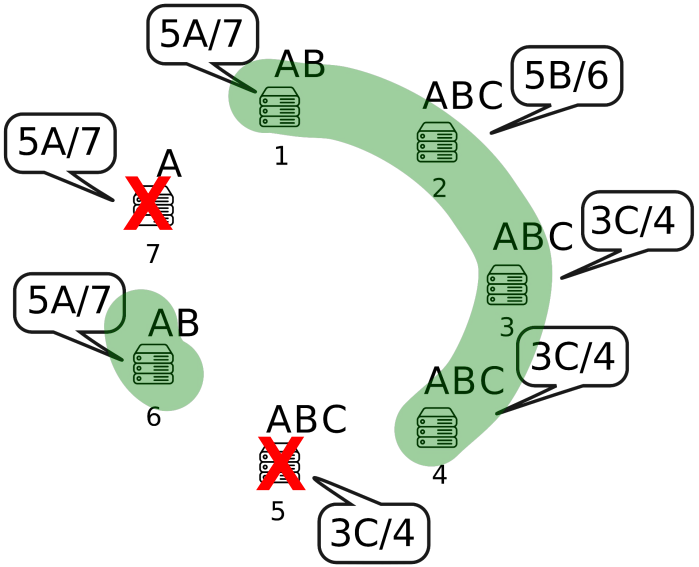

An FBA-resilient overlay must remain connected despite any failures allowed by the FBA model

The overlay must remain connected even if we take an arbitrary quorum and remove its complement

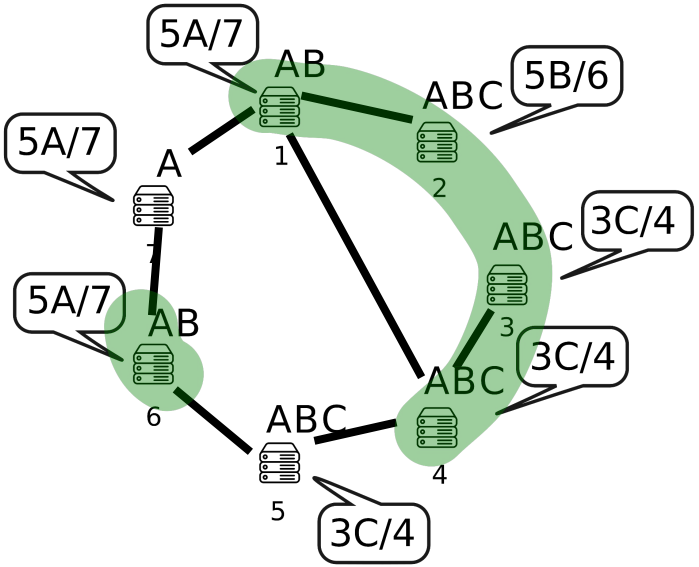

An FBA-resilient overlay must remain connected despite any failures allowed by the FBA model

The overlay must remain connected even if we take an arbitrary quorum and remove its complement

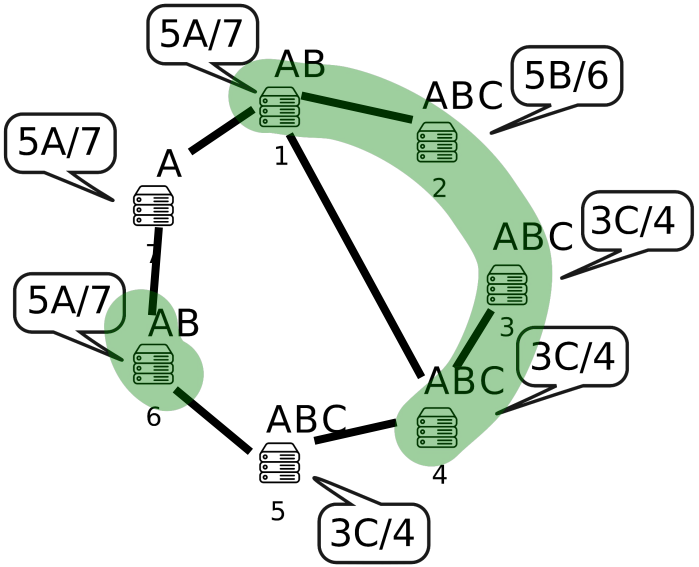

An FBA-resilient overlay must remain connected despite any failures allowed by the FBA model

The overlay must remain connected even if we take an arbitrary quorum and remove its complement

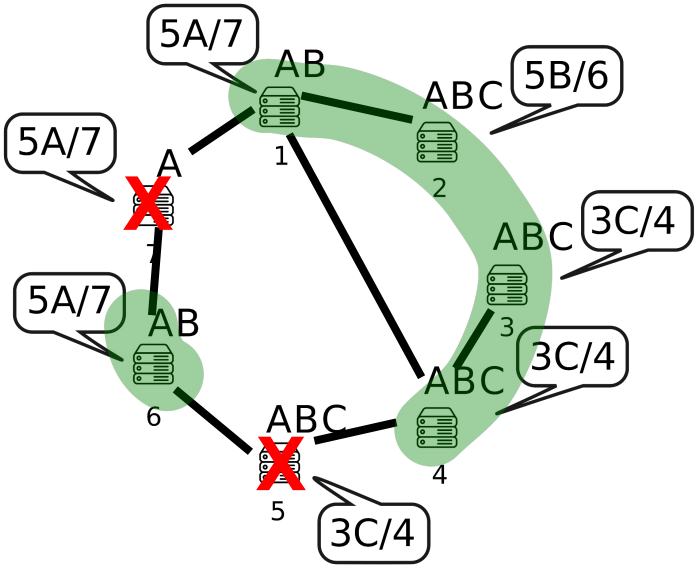

An FBA-resilient overlay must remain connected despite any failures allowed by the FBA model

The overlay must remain connected even if we take an arbitrary quorum and remove its complement

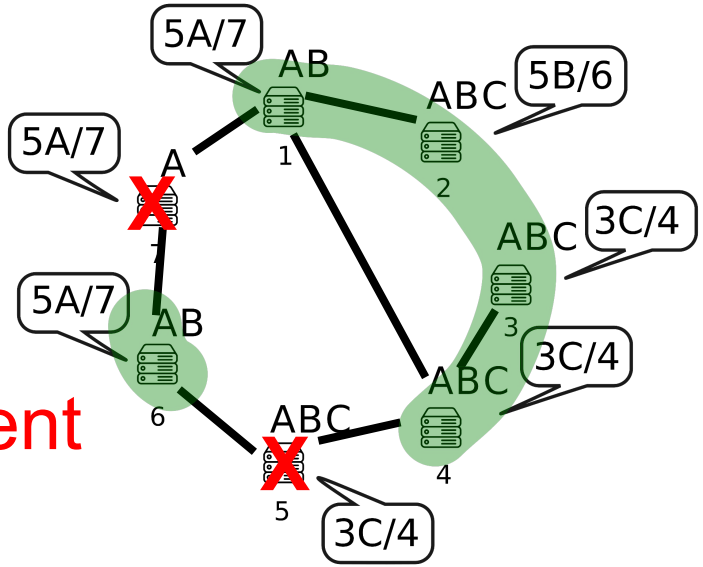

Not FBA-Resilient

An FBA-resilient overlay must remain connected despite any failures allowed by the FBA model

May be FBA-resilient

The overlay must remain connected even if we take an arbitrary quorum and remove its complement

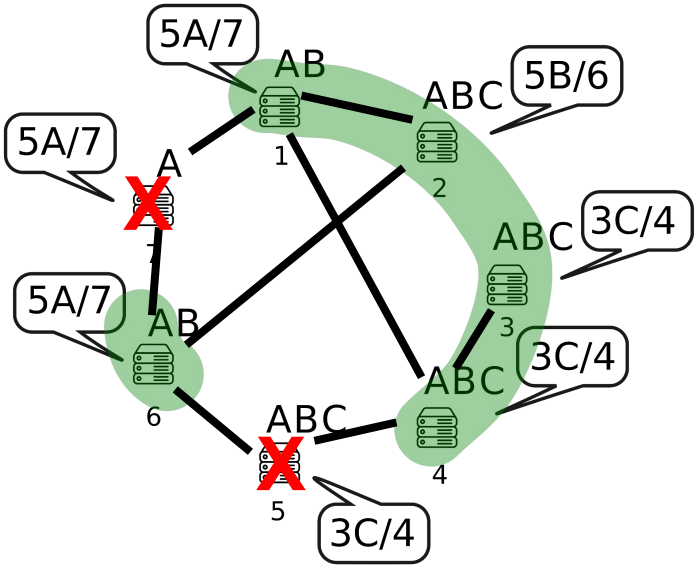

Given agreement requirements, can we compute an optimal FBA-resilient overlay graph?

The problem

Find an FBA-resilient overlay of *minimal* degree and diameter 2

Combinatorial explosion!

May be FBA-resilient

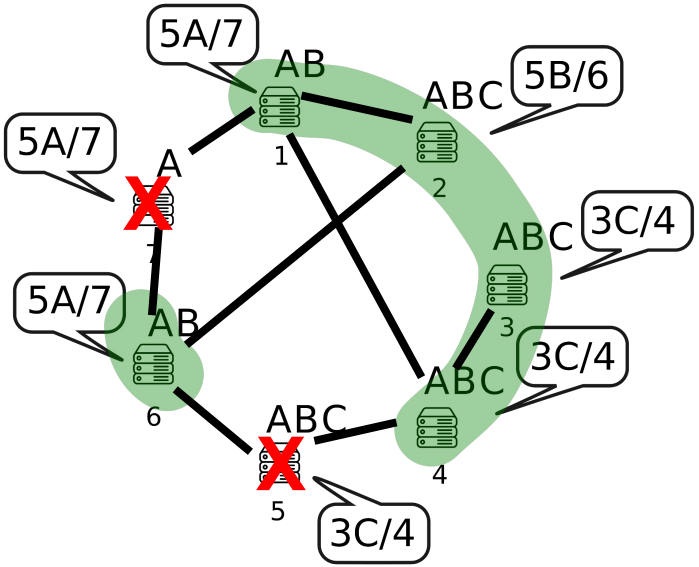

We want to apply this to the Stellar network

Stellar is the largest FBA deployment to date

Roughly 500 overlay nodes in total, with about 200 “validator nodes” taking part in consensus

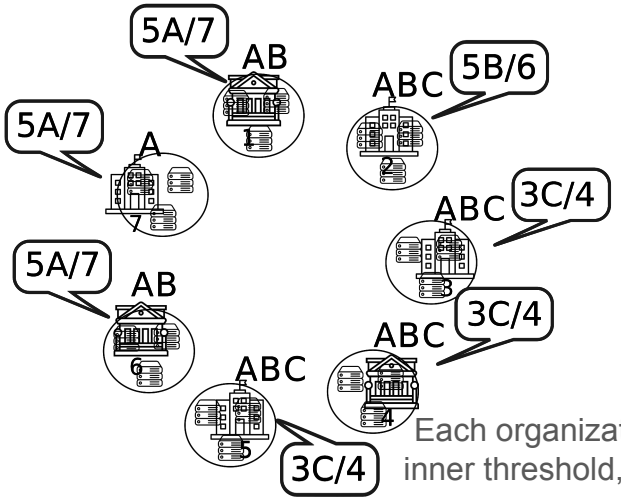

Each organization has an inner threshold, usually 2/3

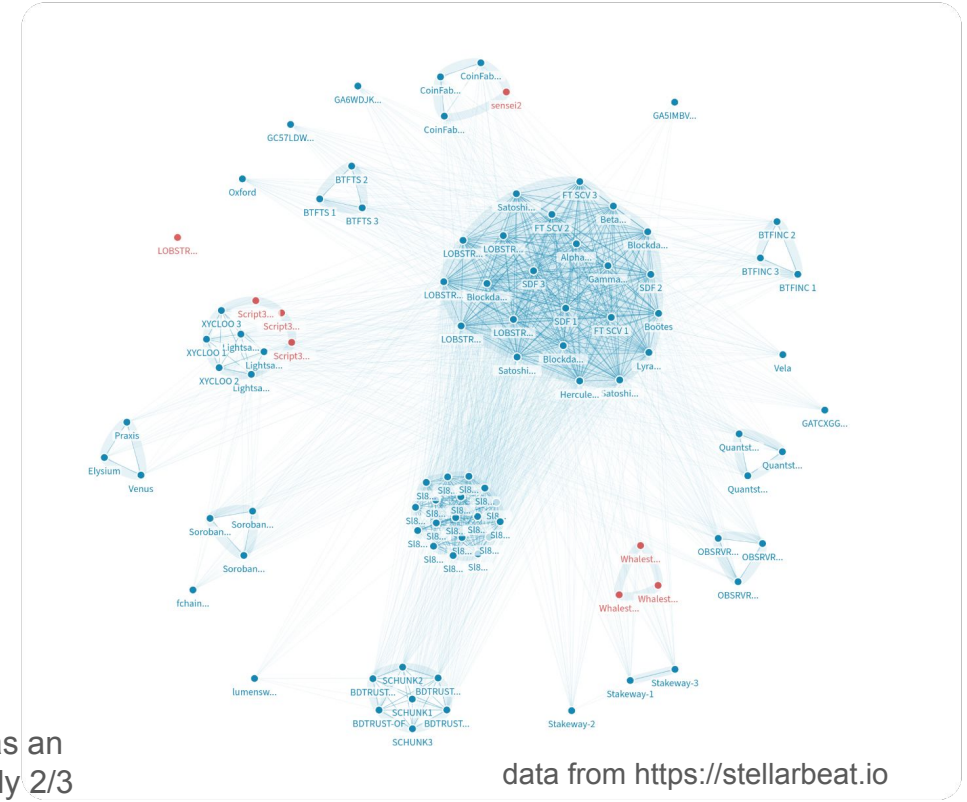

data from https://stellarbeat.io

Constellation curbs combinatorial explosion by searching for optimal instances of the Constellation graph template

- The template has two parameters:
  - k, number of clusters
  - f, an assignment of organizations to clusters

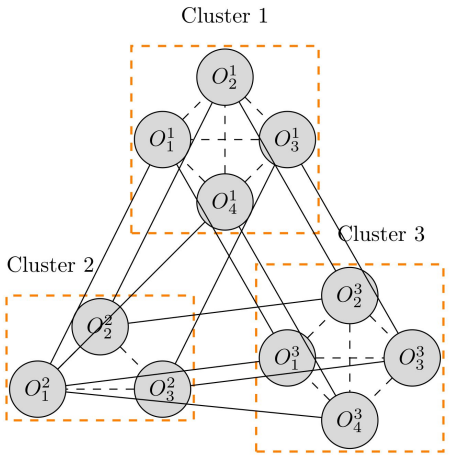

Example template instantiation for 11 organizations (k=3 clusters)

Constellation curbs combinatorial explosion by searching for optimal instances of the Constellation graph template

- The template has two parameters:
  - k, number of clusters
  - f, an assignment of organizations to clusters
- Given assignment of organizations to clusters:
  - Create complete intra-cluster graph
  - Match each organization with a peer organization in each other cluster

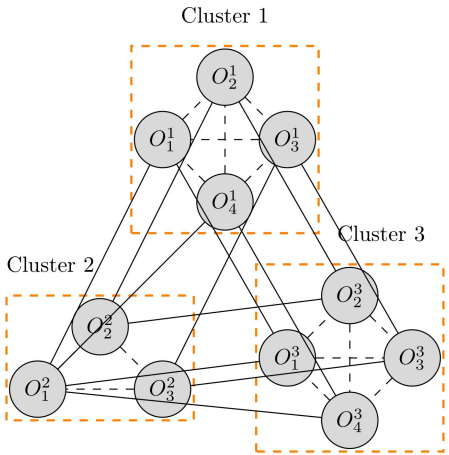

Example template instantiation for 11 organizations (k=3 clusters)

Constellation curbs combinatorial explosion by searching for optimal instances of the Constellation graph template

- The template has two parameters:
  - k, number of clusters
  - f, an assignment of organizations to clusters
- Given assignment of organizations to clusters:
  - Create complete intra-cluster graph
  - Match each organization with a peer organization in each other cluster
- Dotted and full edges correspond to two node-level connection patterns:

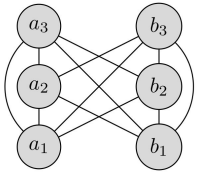

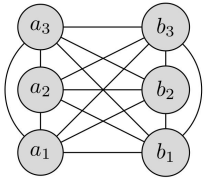

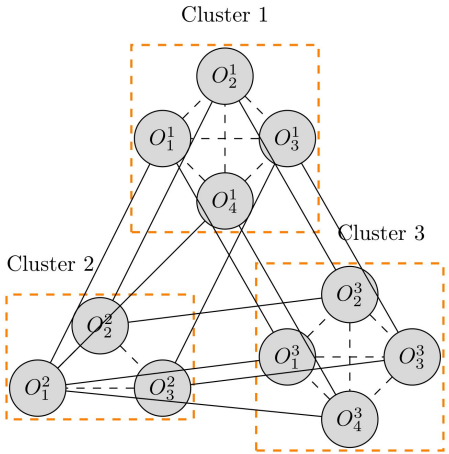

Example template instantiation for 11 organizations (k=3 clusters)

Constellation curbs combinatorial explosion by searching for optimal instances of the Constellation graph template

- The construction guarantees diameter 2
- We use brute-force search over k and f to find an FBA-resilient graph with minimal degree
  - Template instances = partitions of the multiset of thresholds used
  - We use Knuth's Algorithm M to enumerate multi-set partitions

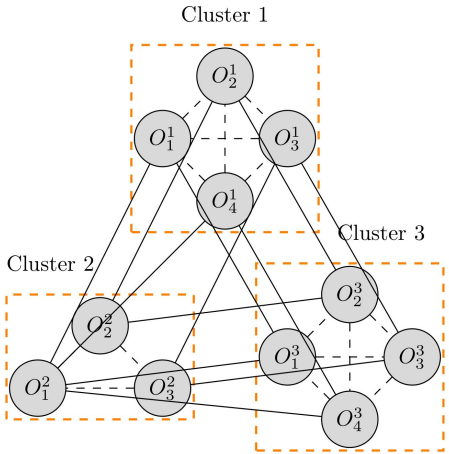

Example template instantiation for 11 organizations (k=3 clusters)

In symmetric cases, we get products of complete graphs, which have well-studied connectivity properties

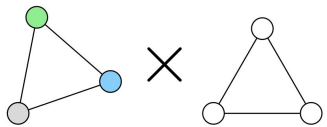

Example product of complete graphs K3 x K3

In symmetric cases, we get products of complete graphs, which have well-studied connectivity properties

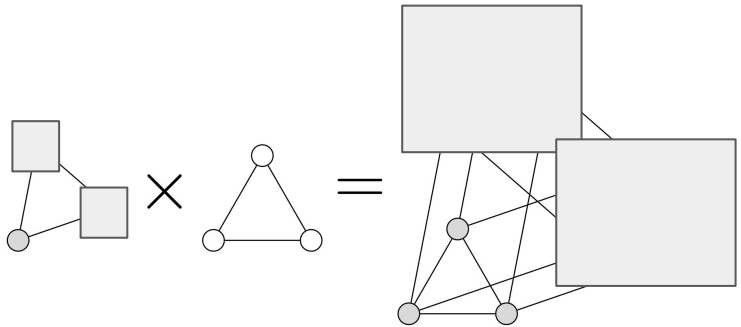

Example product of complete graphs K3 x K3

In symmetric cases, we get products of complete graphs, which have well-studied connectivity properties

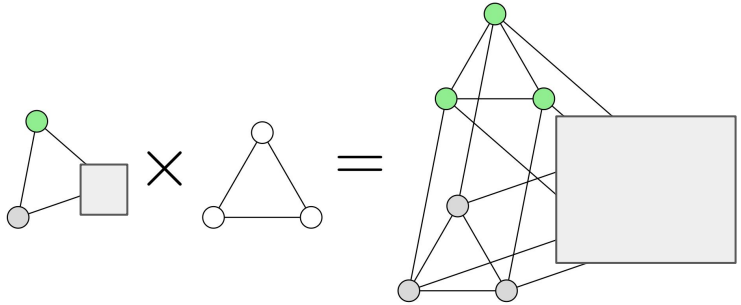

Example product of complete graphs K3 x K3

In symmetric cases, we get products of complete graphs, which have well-studied connectivity properties

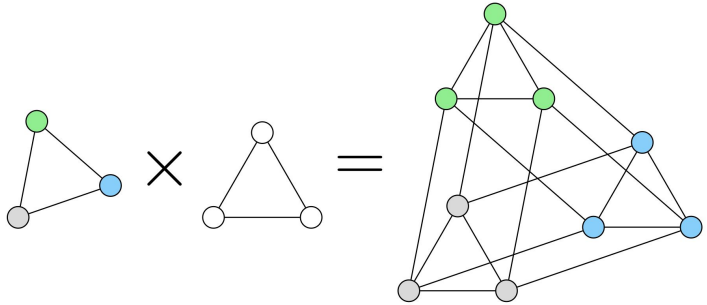

Example product of complete graphs K3 x K3

In symmetric cases, we get products of complete graphs, which have well-studied connectivity properties

Liouville's formula* gives a lower bound on the connectivity of products of complete graphs

Theorem: In a symmetric FBAS with n nodes with threshold k, K_k x K_{n-k} is close to optimal

Theorem: upper bound on the number of failures to keep diameter under 3

If not symmetric, simulations show good results

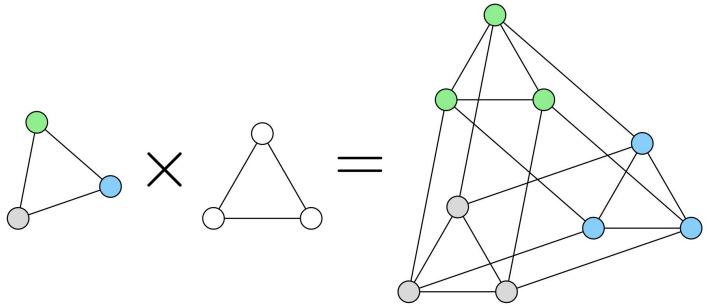

Example product of complete graphs K3 x K3

*Liouville, B. Sur la connectivité des produits de graphes, 1978

Constellation achieves better degree than random and greedy fault-tolerant overlays

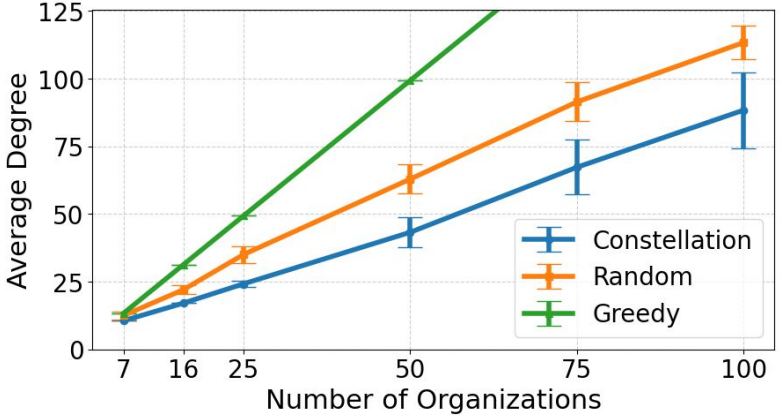

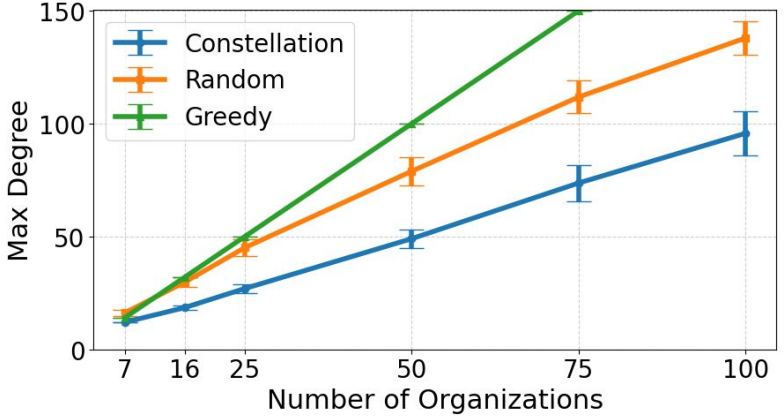

In practice in the Stellar network, we expect a top-tier/second-tier structure

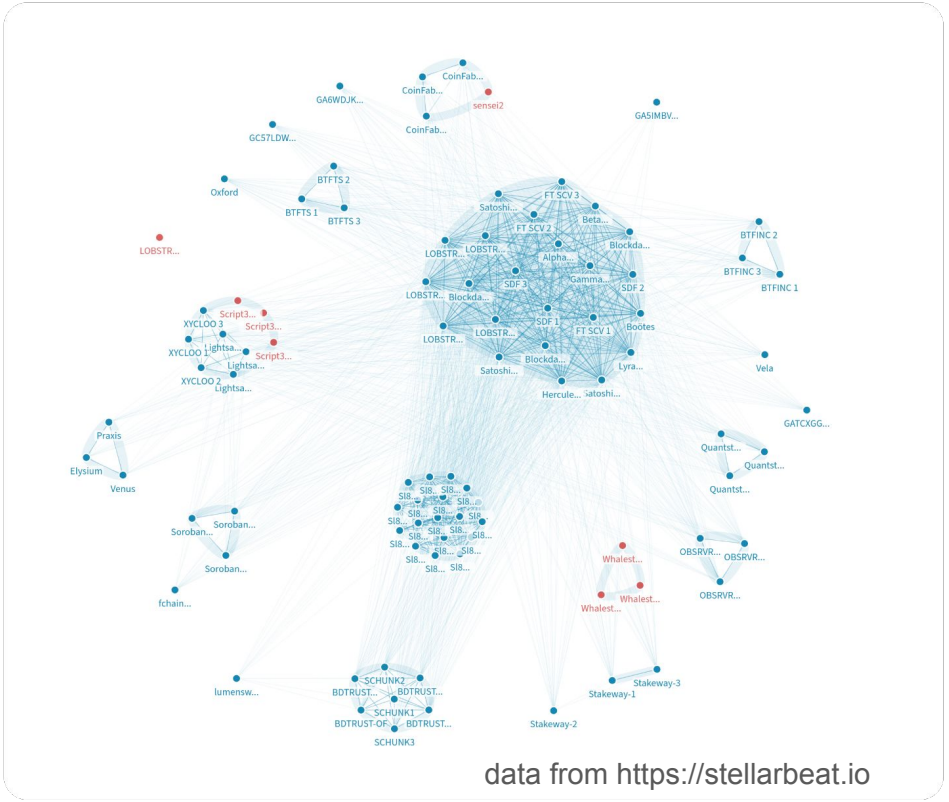

data from https://stellarbeat.io

In practice in the Stellar network, we expect a top-tier/second-tier structure

Top-tier organizations require agreement from each other

Second-tier organizations require agreement from the top-tier, and rarely from other second-tier

From the top-tier point of view, second-tier nodes may as well be Sybils

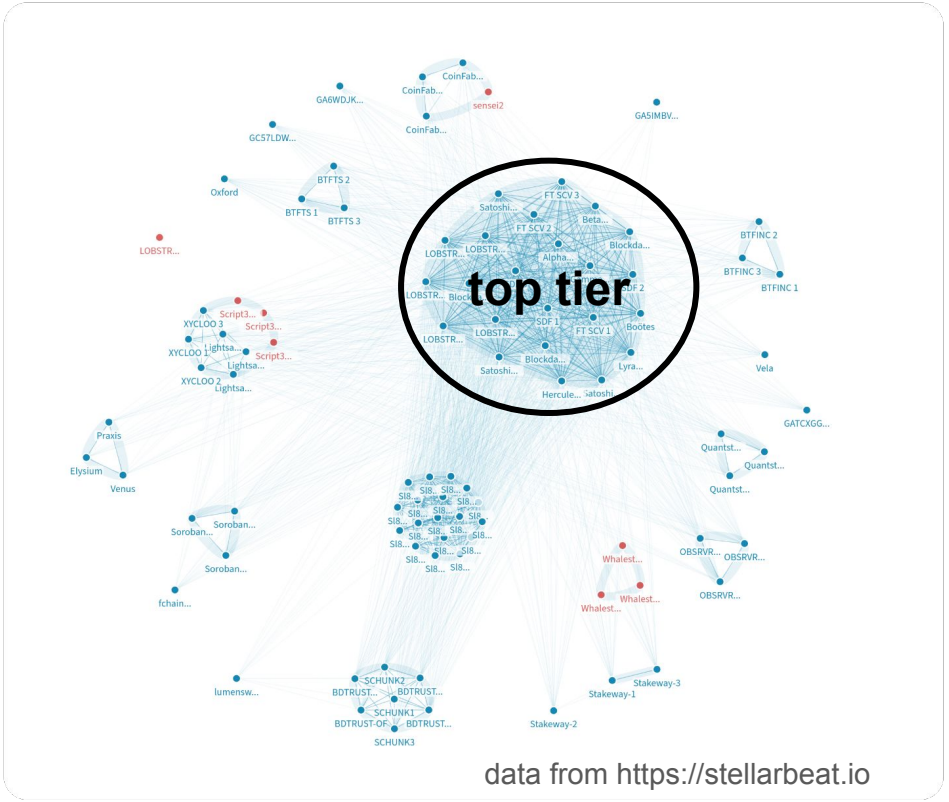

Constellation uses a structured overlay for the top-tier and a best-effort random overlay for the second-tier

Agreement requirements are posted on chain

Compute top-tier

Compute a structured, FBA-resilient overlay for the top-tier

Add best-effort random overlay connecting the 2nd tier

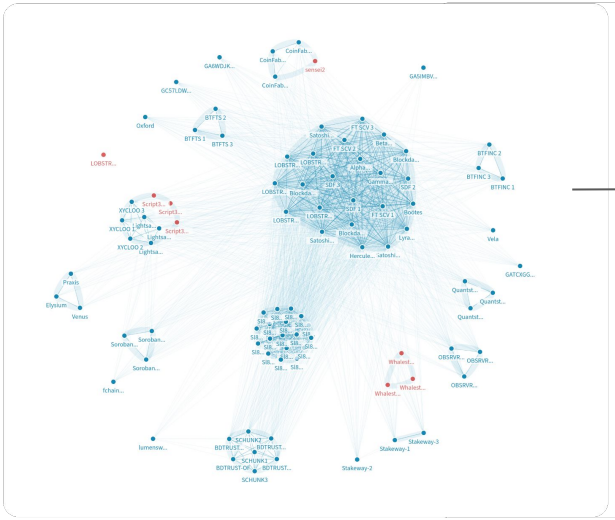

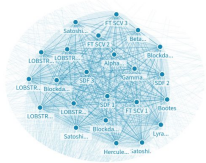

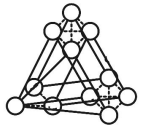

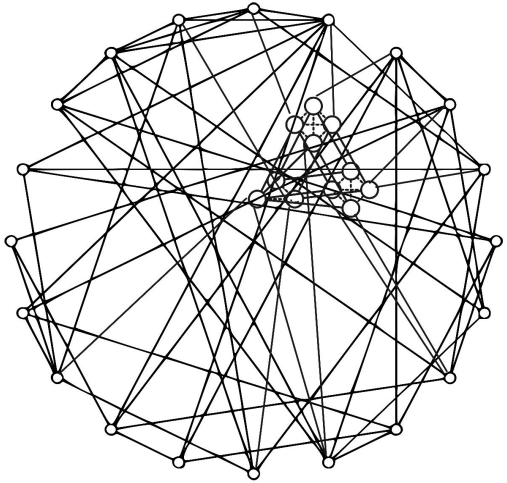

Testbed evaluation: throughput of stellar-core on a replica of the Stellar network

Constellation guarantees fault-tolerance and matches the performance of the best non-fault-tolerant random overlay

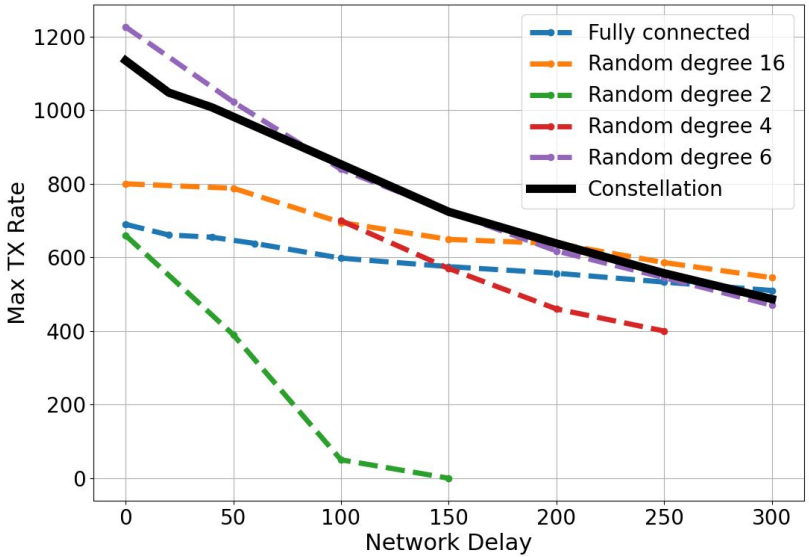

Q & A