Resilience of Blockchain Overlay Networks

Aristodemos Paphitis, Nicolas Kourtellis, Michael Sirivianos

am.paphitis@edu.cut.ac.cy

Cyprus University of Technology

Telefonica Research

Motivation

- Blockchains for critical infrastructure
- Impacts a lot of people
- Their security and resilience depends on the P2P network
- The P2P network has not been studied in depth

Why this knowledge gap?

- Assumption of reliable Internet communications
- Decentralization $\neq$ Safe & Robust
- The network topology is unknown

This work

- Seven distinct blockchain overlays
- Structural resilience
  - Against random failures
  - Targeted attacks
  - Spatial centralization in Ases
  - Inter-dependencies (common nodes in different networks)

Selected networks

Well known, established cryptocurrencies.
Frequently listed in top50 by CoinMarketCap

Bitcoin

BitcoinCash

DASH

Dogecoin

Ethereum

Litecoin

ZCash

Challenges

Topology is unknown

Inferring the topology is very hard

Topology hiding techniques are used

Topology knowledge threats:

Eclipse Attacks,

Facilitate network partitioning,

User anonymity

Main Idea

- Peer Address propagation helps discovery process
- Construct **connectivity graphs** that contain ALL POSSIBLE links

Key Idea: Exact topology is hard

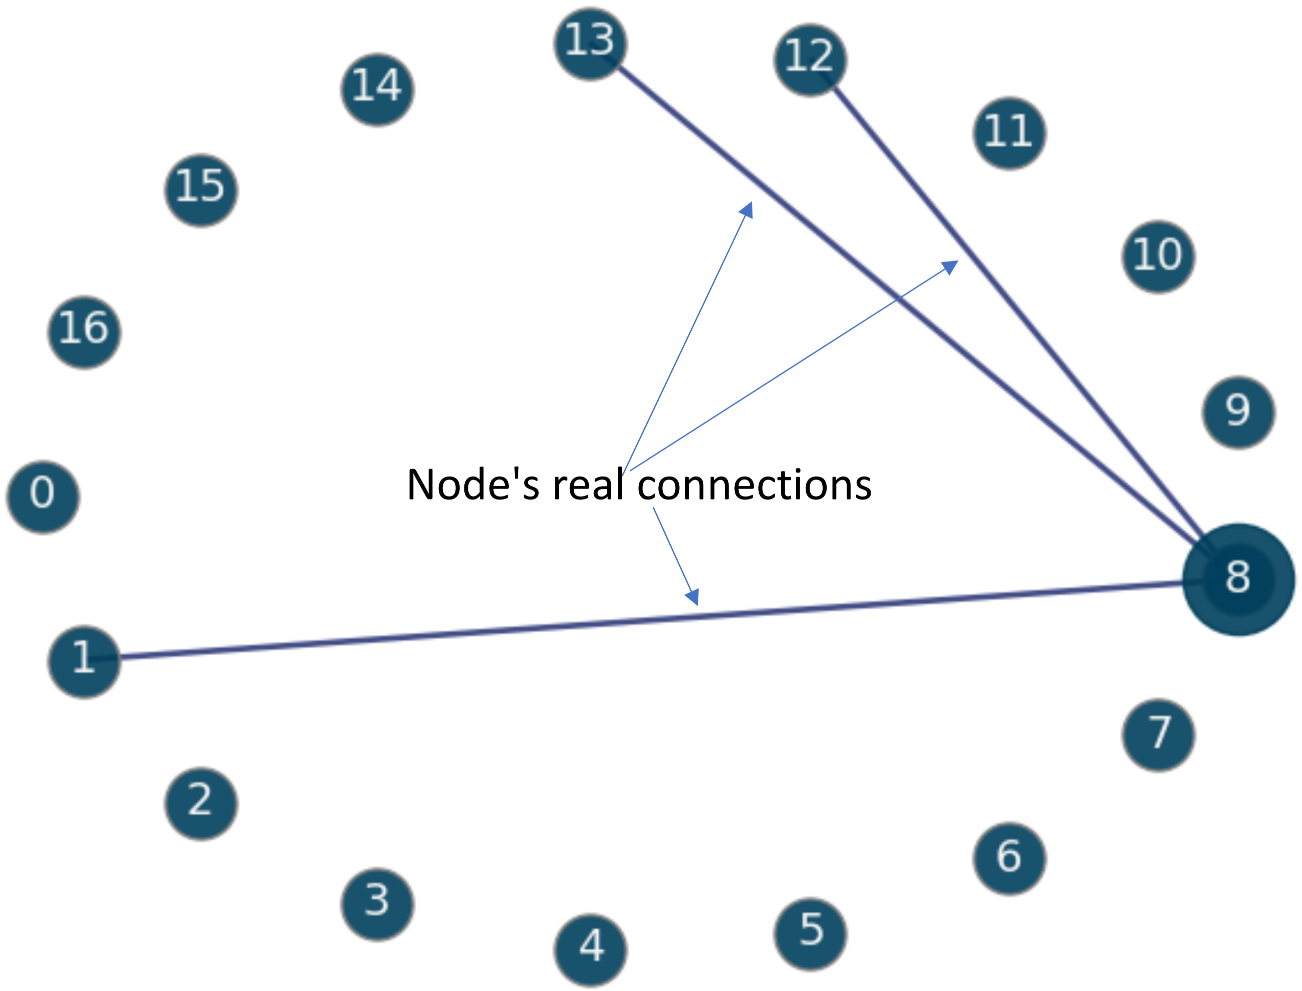

8th Node's network view:

Nodes 0 – 16 are know to Node 8

Node 8 has active connections to nodes: #1, #12, #13

Key Idea: Exact topology is hard

Study all possible paths!

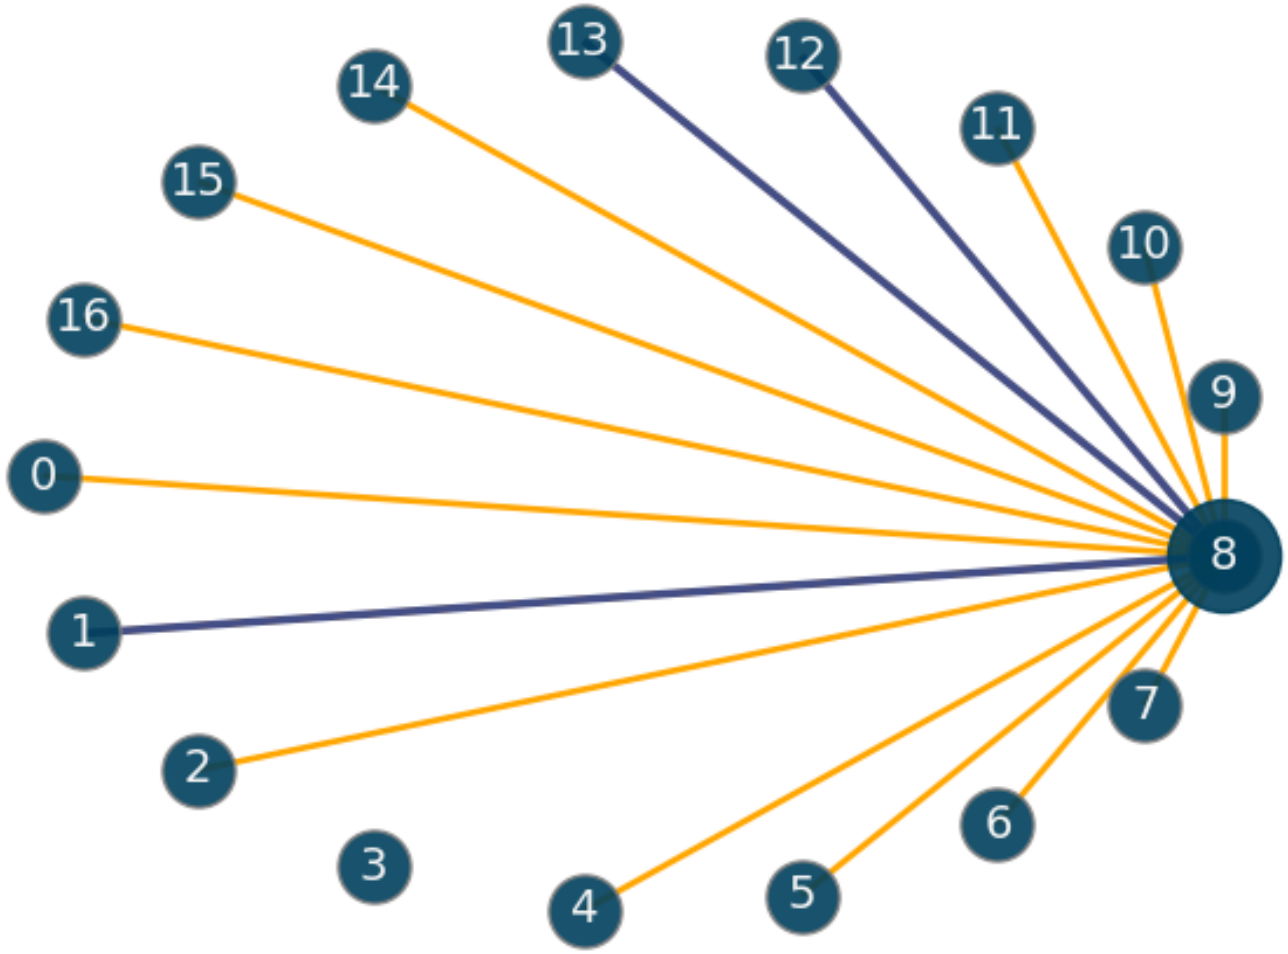

Node's view of the network

Combine views from all nodes

Connectivity Graph

2 – hour snapshots x 28 days x 7 Blockchains

2-hour snapshots aggregated to 24-hour snapshots

Key Idea: Exact topology is hard
Study all possible paths!

GOAL: Ask all nodes for the addresses they know

- Avoids the need of accurate topology
- Results are more robust against measurement inaccuracies
- Strengthens Resilience study

We trade accuracy for completeness:

the actually realized topology of an overlay is highly unlikely to be resilient if our inferred topology of possible connections is not

Real vs Potential connections

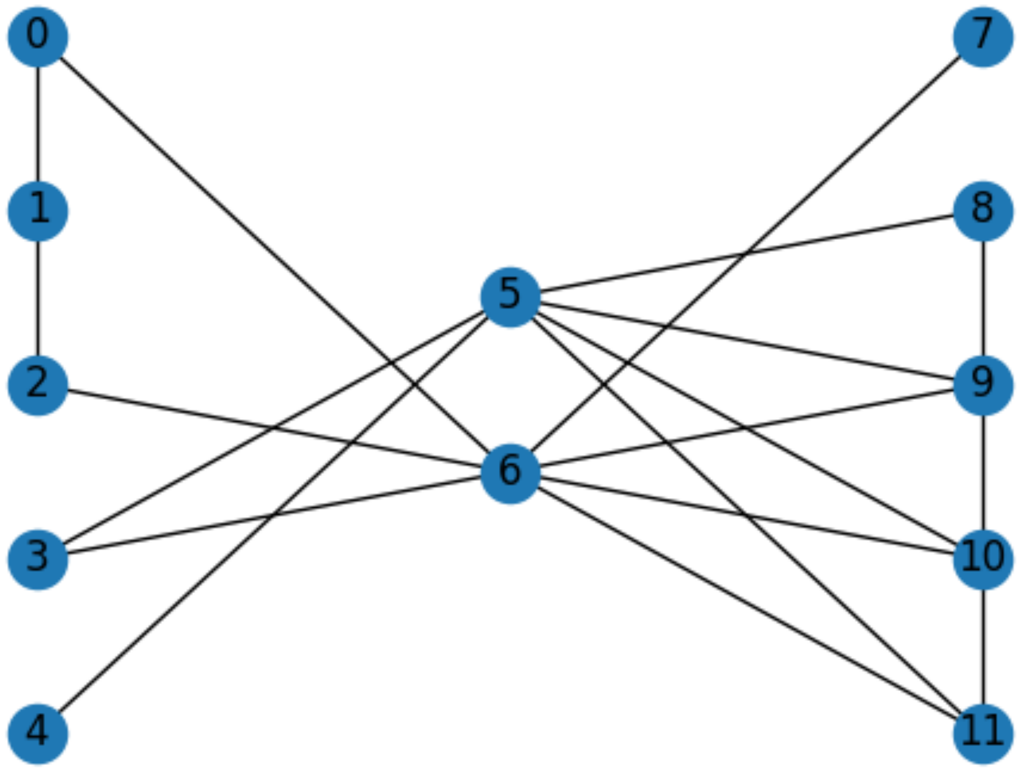

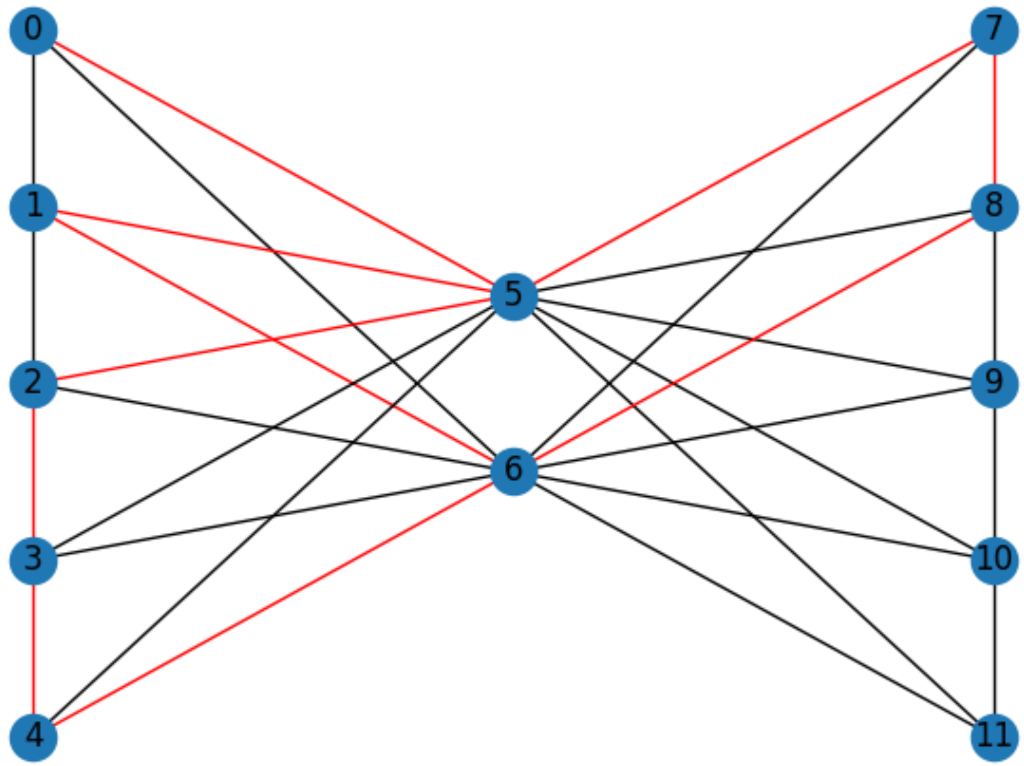

Real vs Potential connections

If R is spanning subgraph of G then:

$$G : k(R) \leq k(G)$$

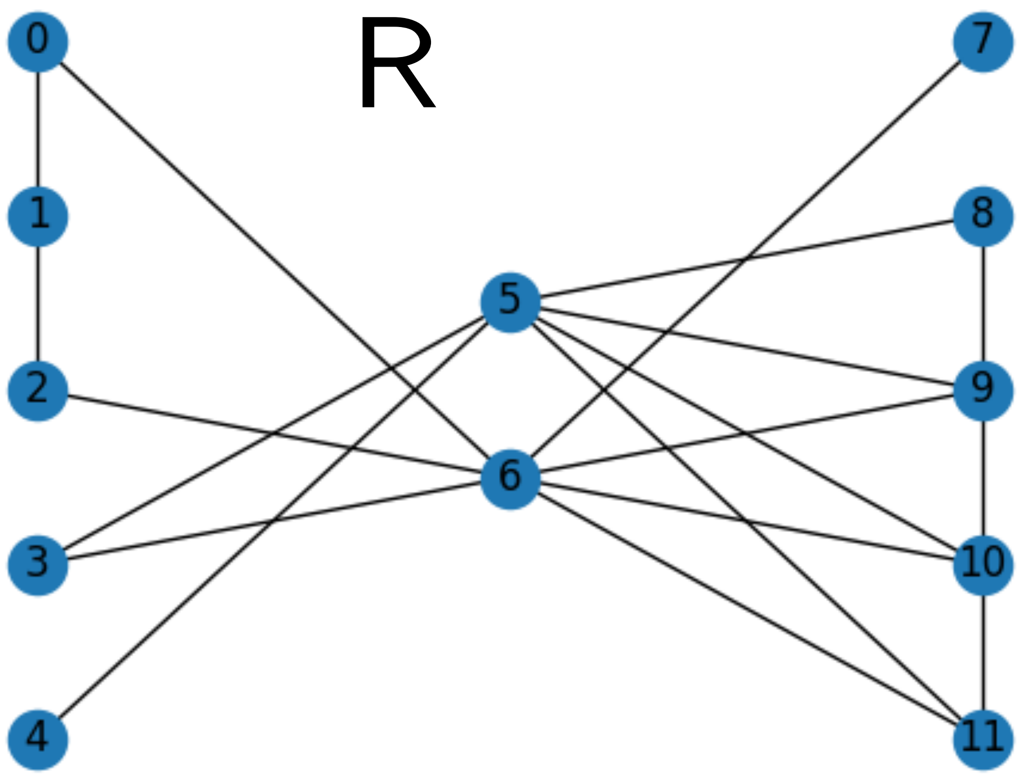

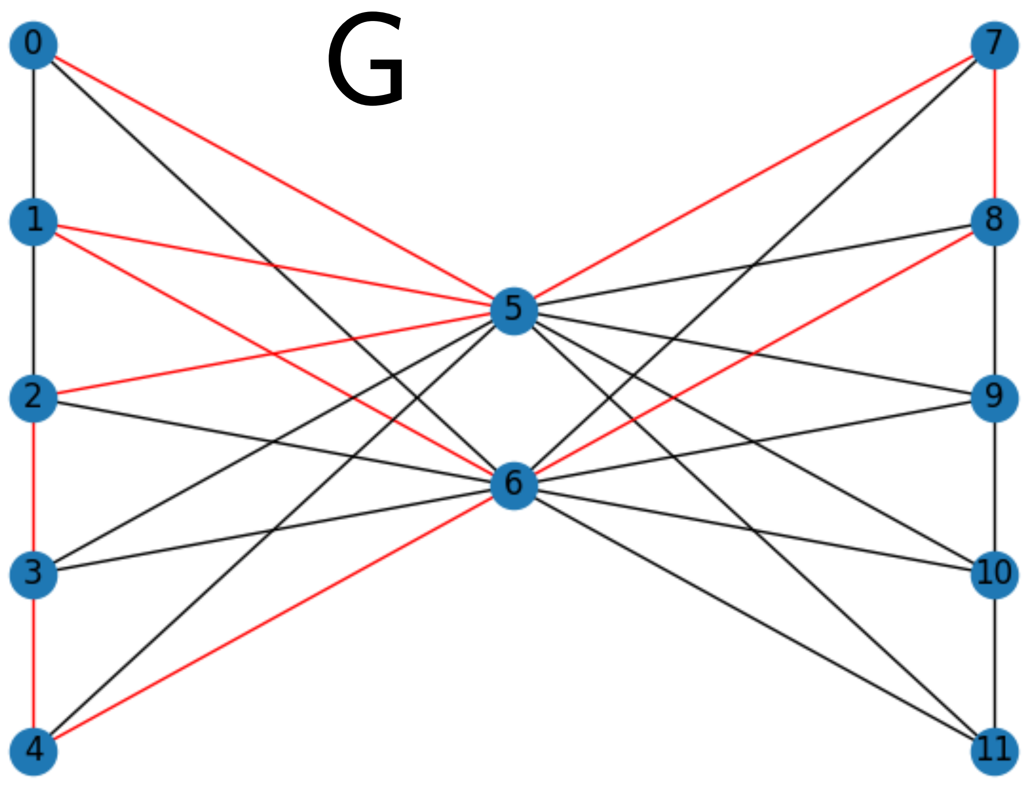

Lemma 1 by Harary [36]: Harary, F.: The maximum connectivity of a graph. Proceedings of the National Academy of Sciences of the United States of America 48(7) (1962)

- ***R*** => the real graph
- ***G*** => connectivity graph, reconstructed from our data collection

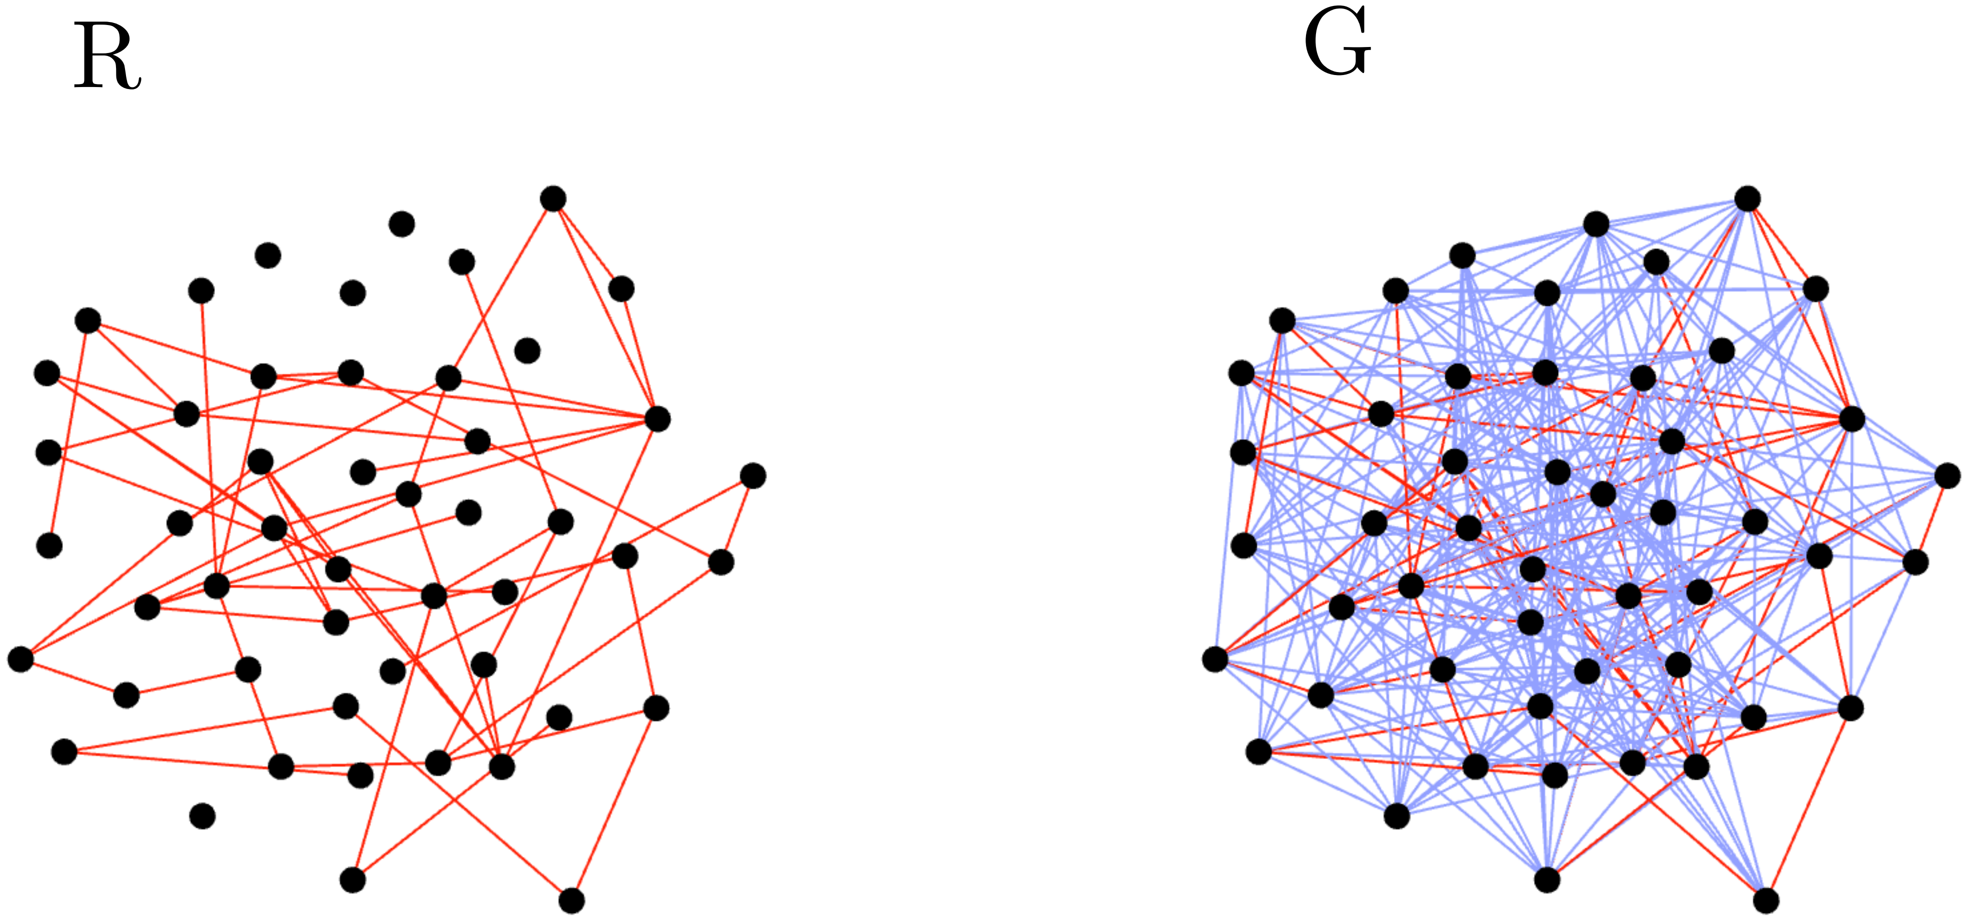

Sample random graphs for example purposes.

Limitations

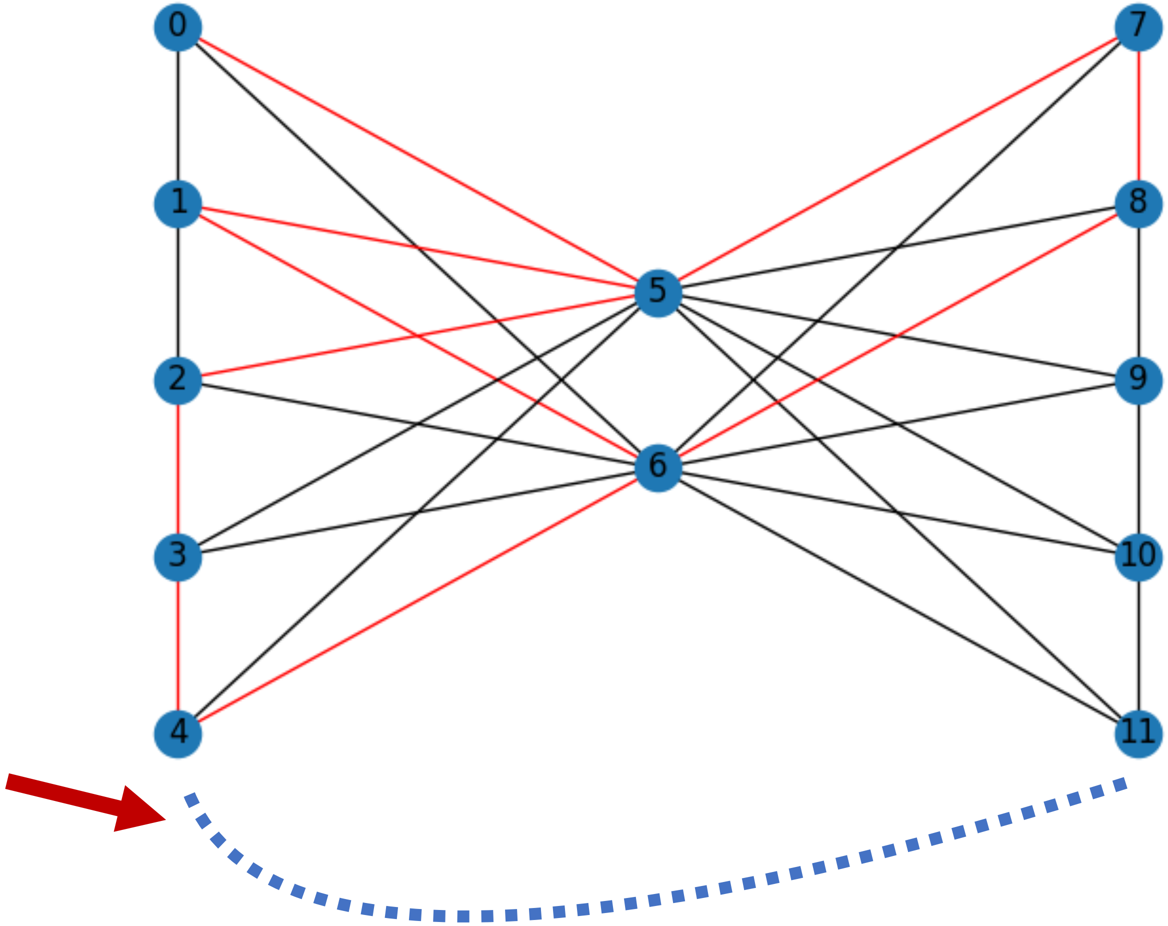

Our validations indicate that such misses are unlikely

Attack Strategies

Remove nodes randomly – simulating failures (baseline)

Remove nodes in order according to a metric (targeted attack)

- Degree
- Betweenness Centrality
- Page Rank

Static setting: metrics are not recalculated after node removal

Measuring failure/attack effects

Size of Largest Weakly Connected Component (in the undirected graph)

Number of Connected Components

Network Diameter

Results – Targeted Attacks

Bitcoin shown; similar results for BitcoinCash & Ethereum

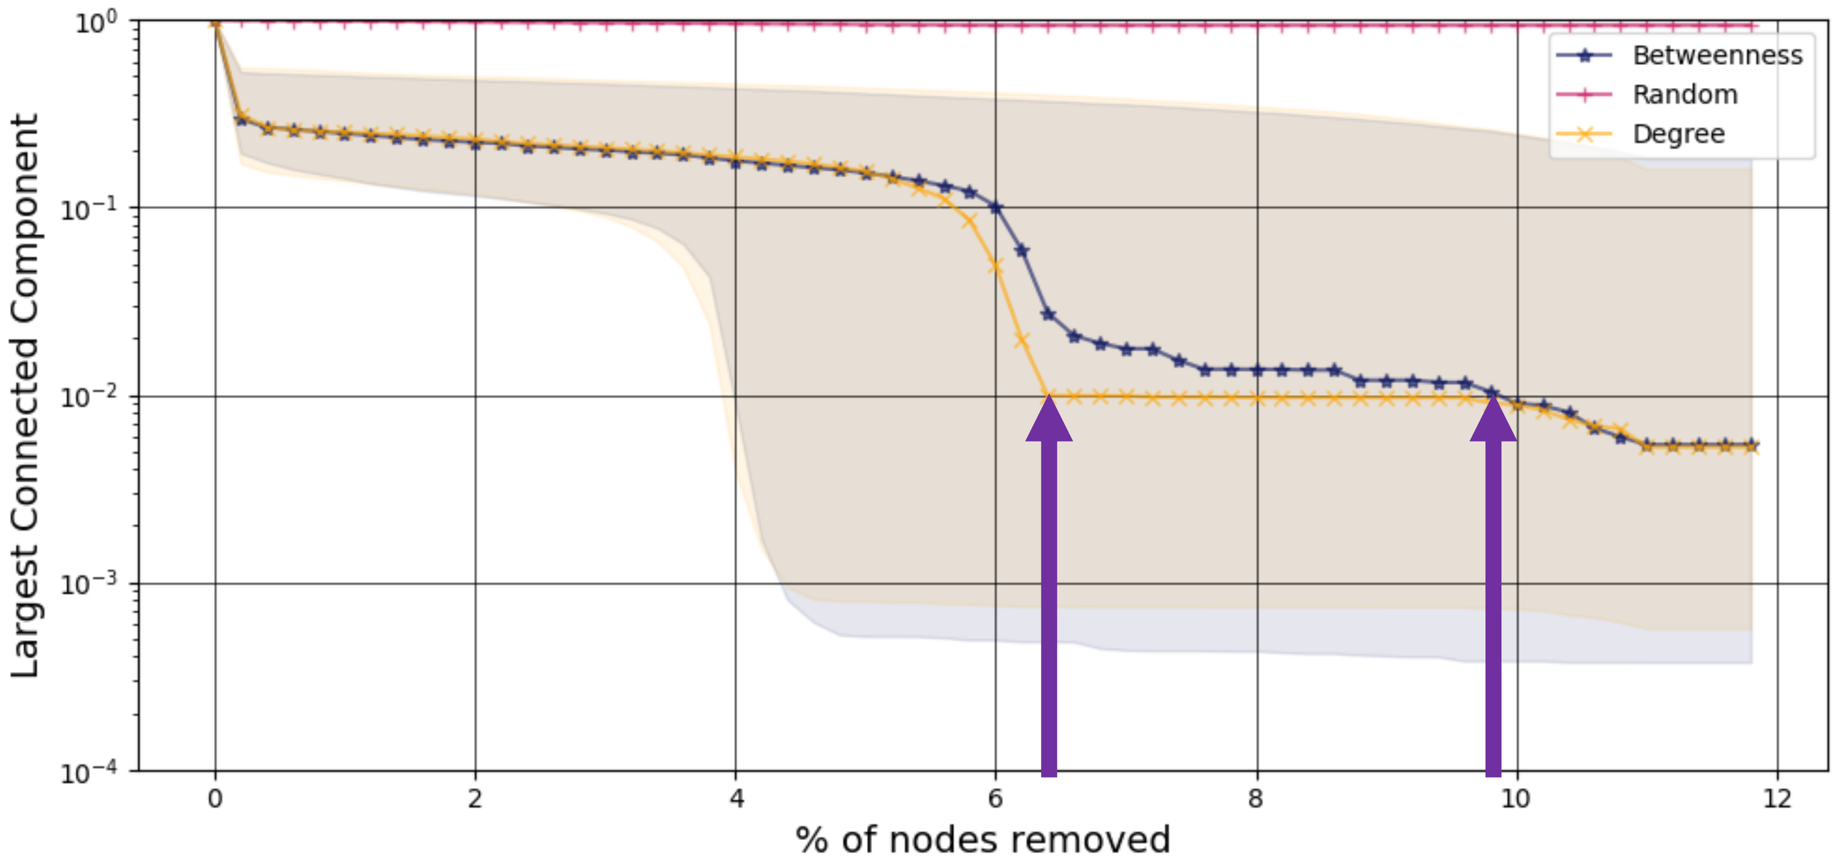

Results – Targeted Attacks

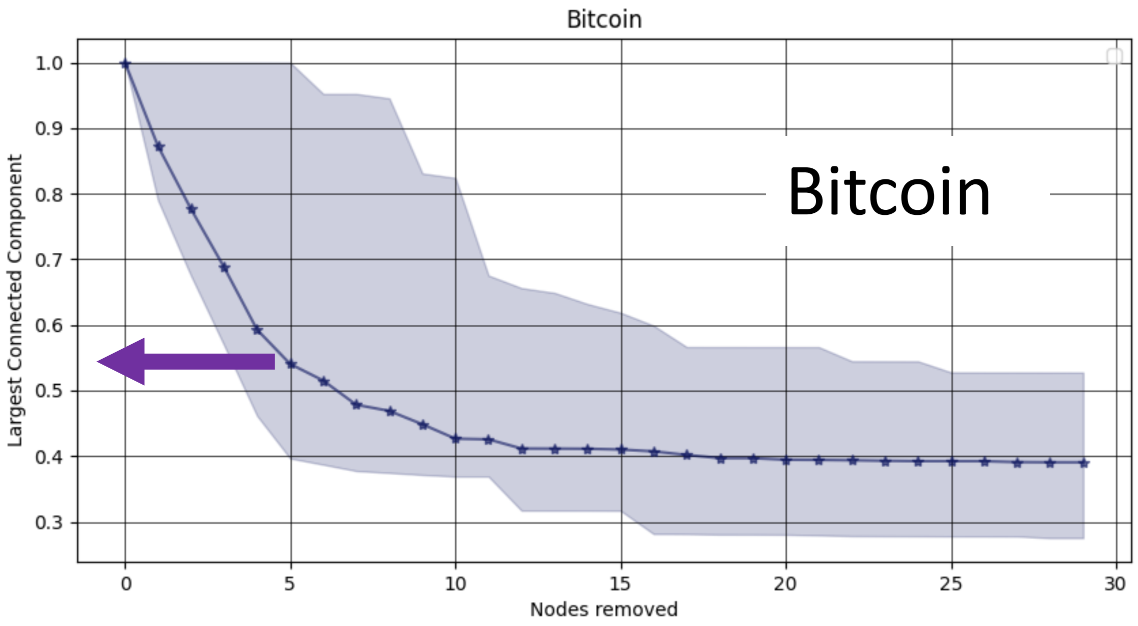

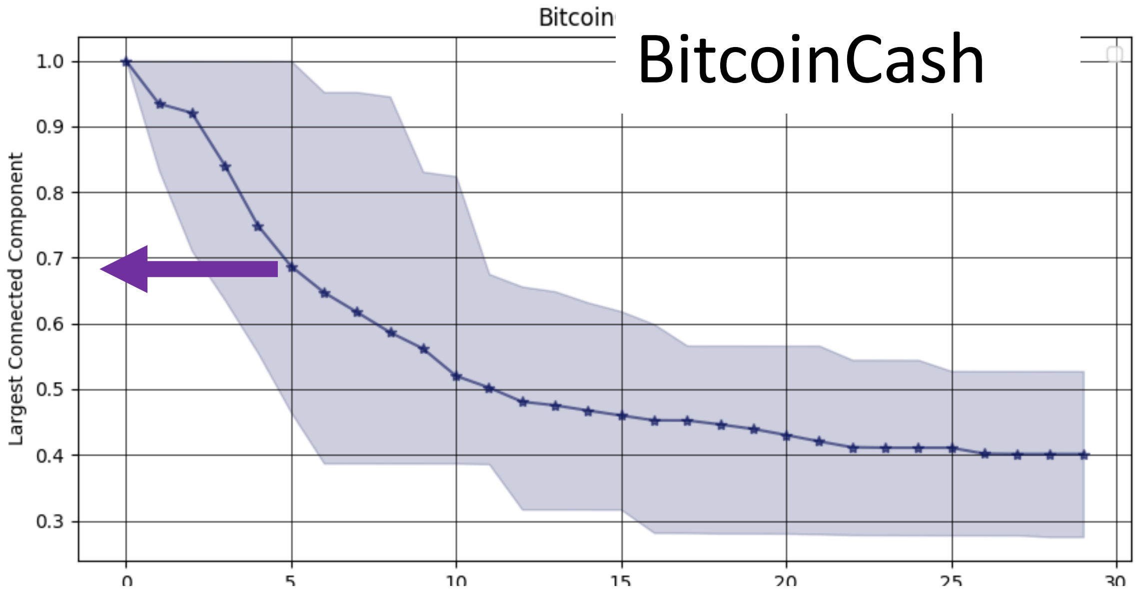

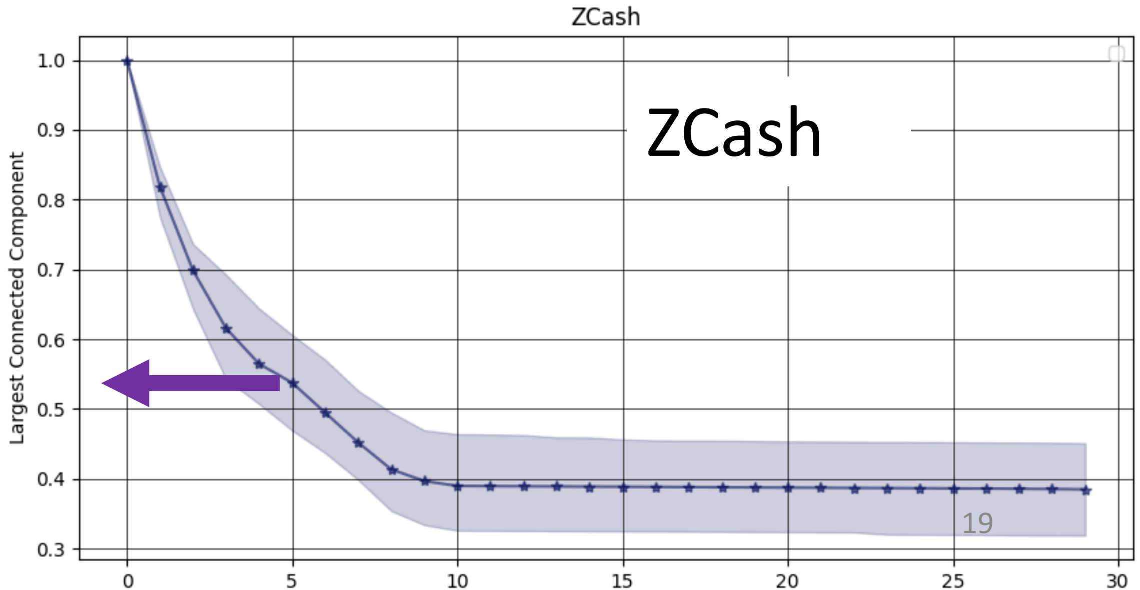

Network-layer inter-dependencies

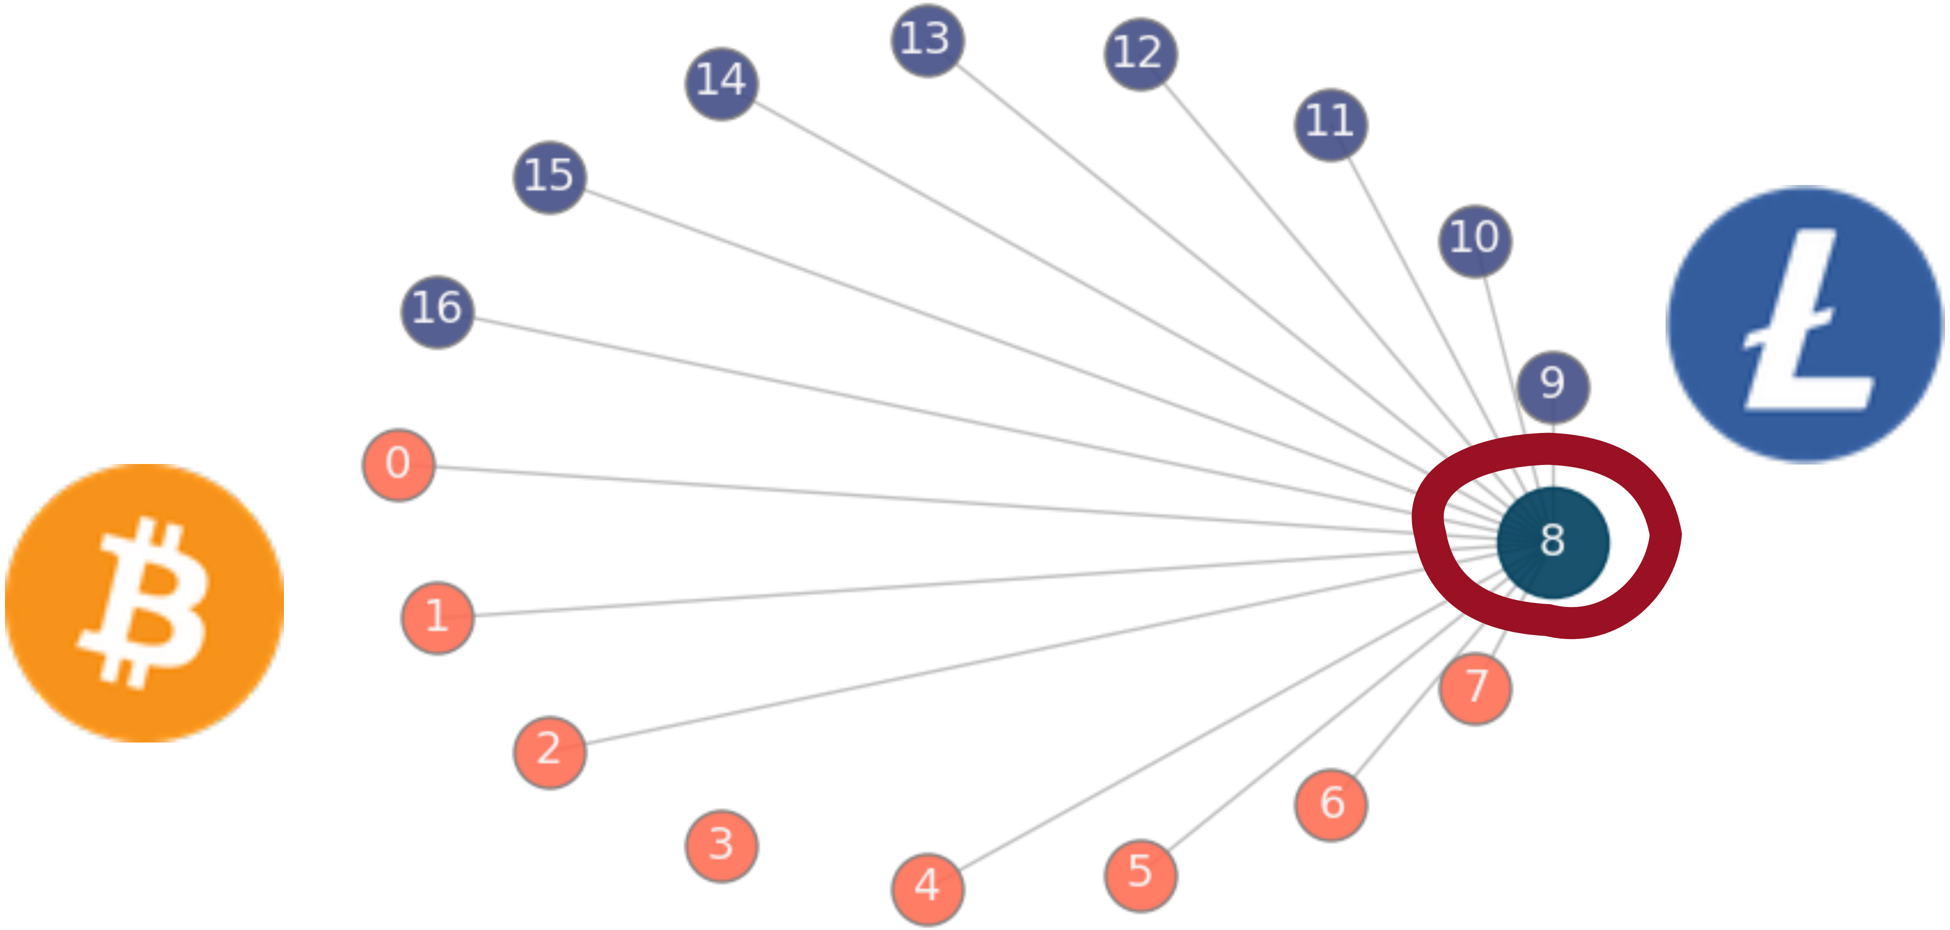

Results
Network-layer inter-dependencies

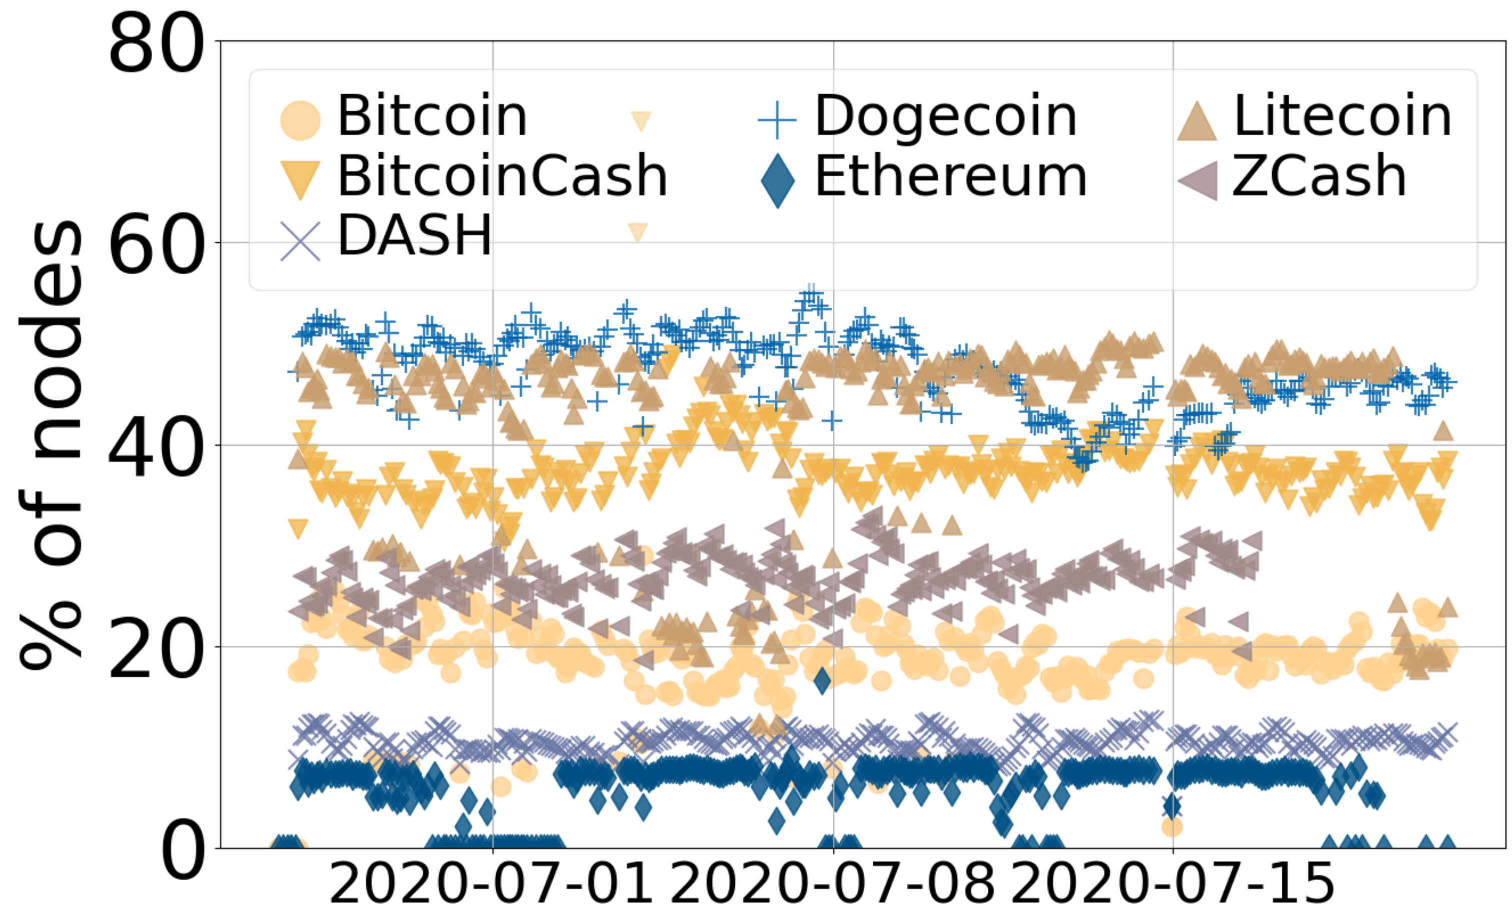

Results
Network-layer inter-dependencies

Bitcoin, Bitcoin Cash, Litecoin, and Zcash share a significant number of nodes.

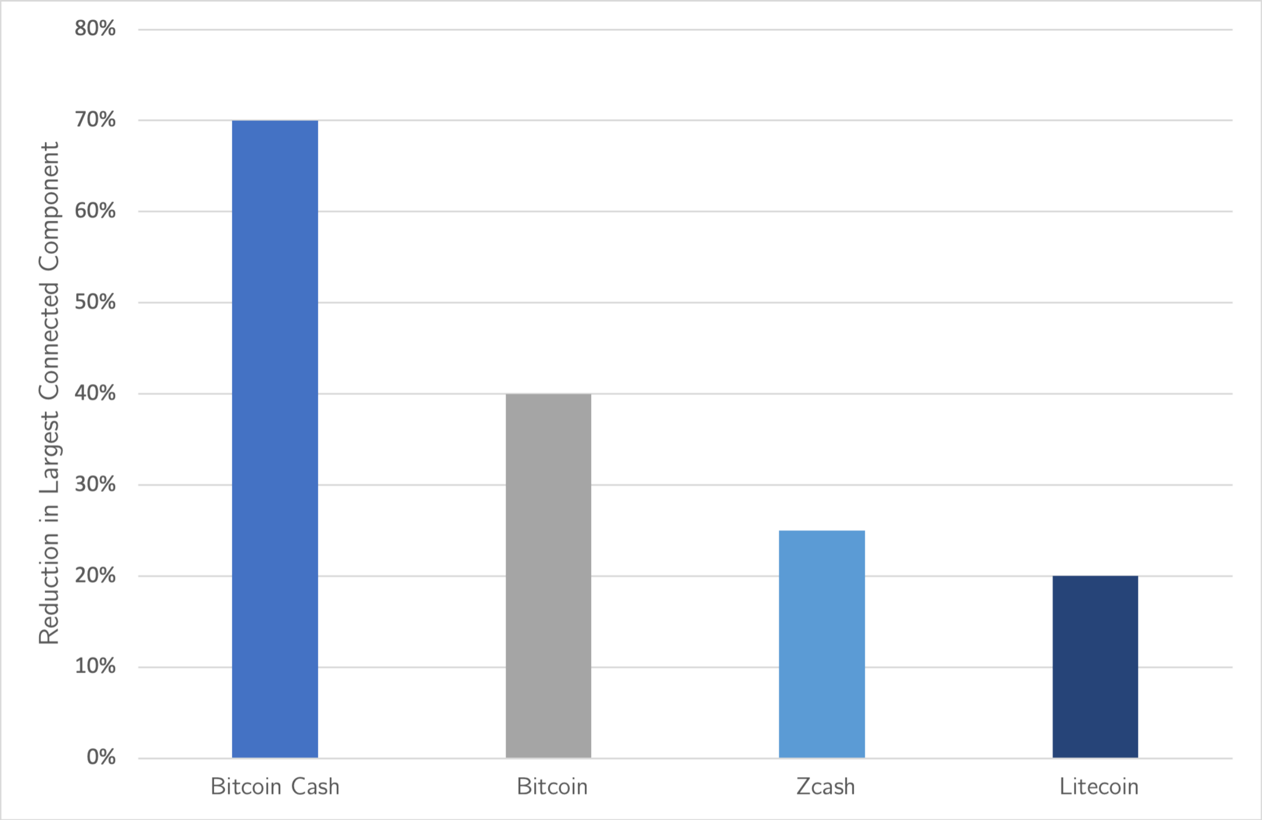

Results
Spatial Centralization

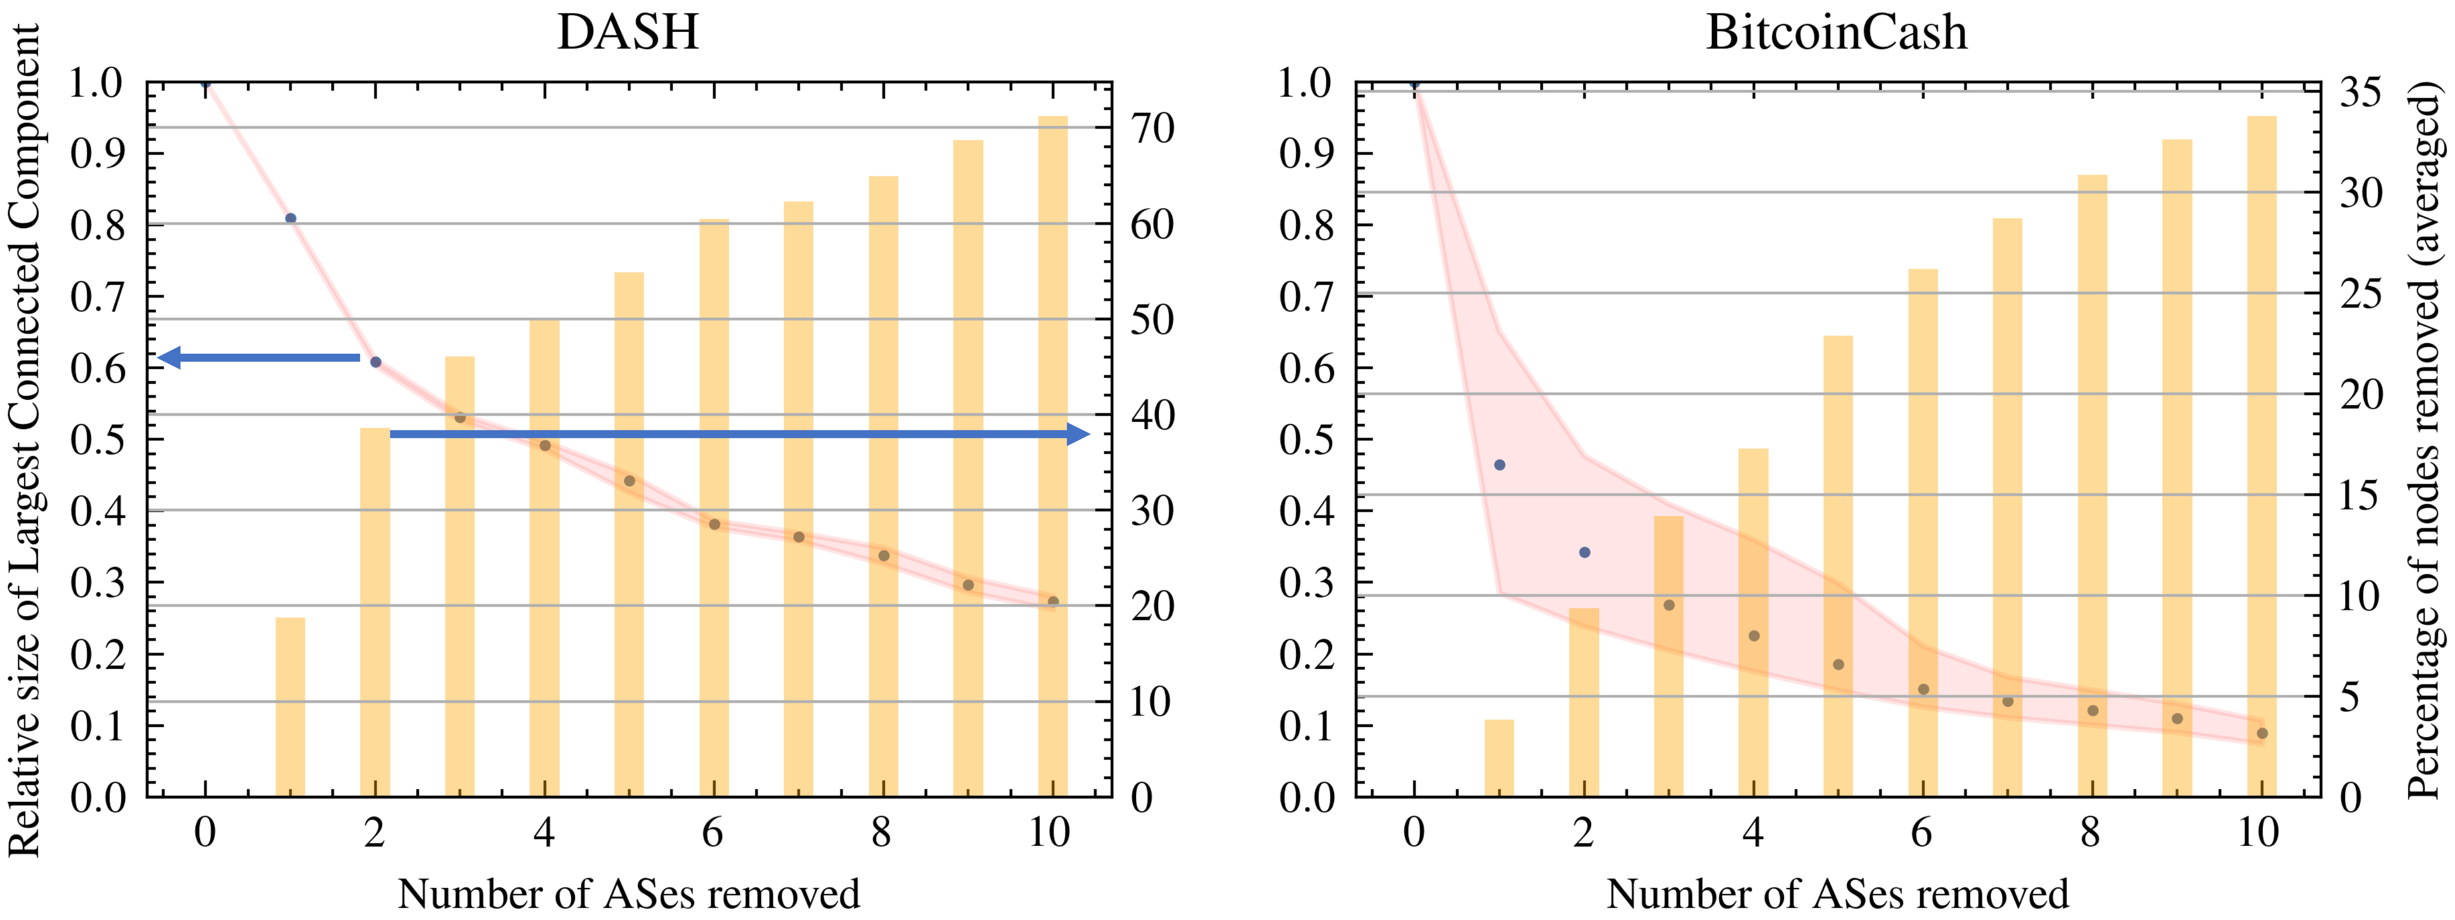

Results
Spatial Centralization + Interdependencies

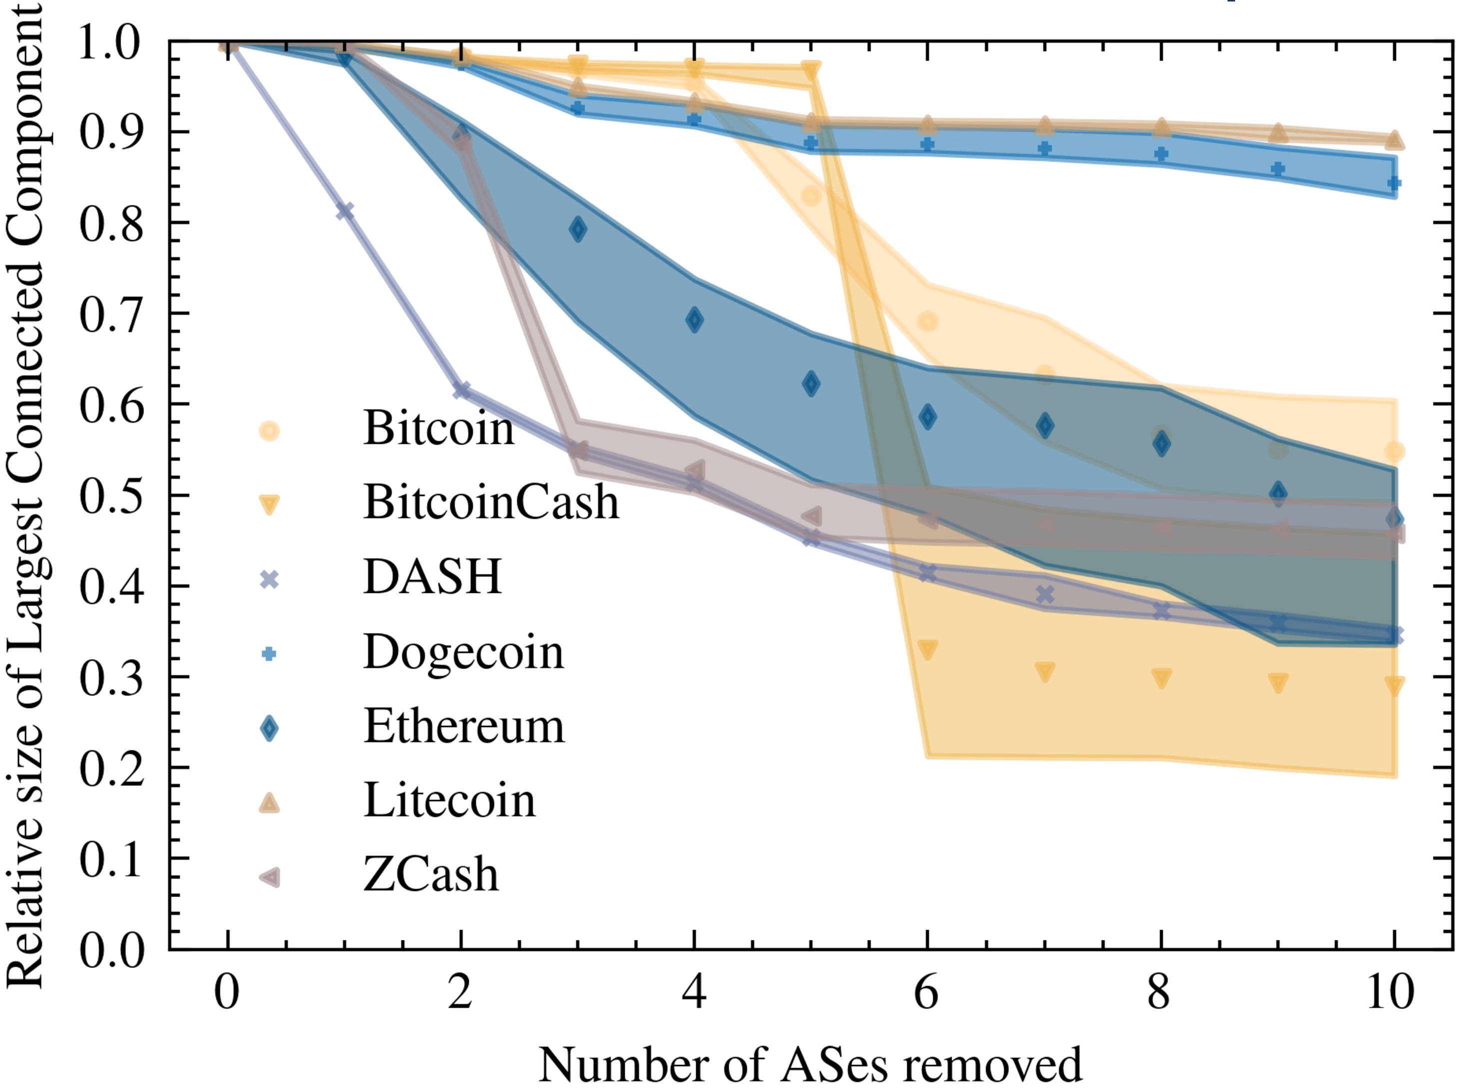

Key insights

Blockchain P2P overlays are:

- Robust against failures
- Weak against targeted attacks
- Not random, contrary to their intended design

Different networks are interconnected

Significant co-location in ASes

→ Simultaneous Disruption of many blockchains

Resilience of Blockchain Overlay Networks

Aristodemos Paphitis, Nicolas Kourtellis, Michael Sirivianos

am.paphitis@edu.cut.ac.cy

Cyprus University of Technology

Telefonica Research

Results - RQ1
Network Characteristics

- Well connected networks with low diameter
- Larger Networks have a smaller Strongly Connected Component
- Highly Dynamic
- NOT Random Graphs
- Small – world property not satisfied

Are they structured in a similar fashion?

- No

Results – RQ2
Network-layer inter-dependencies

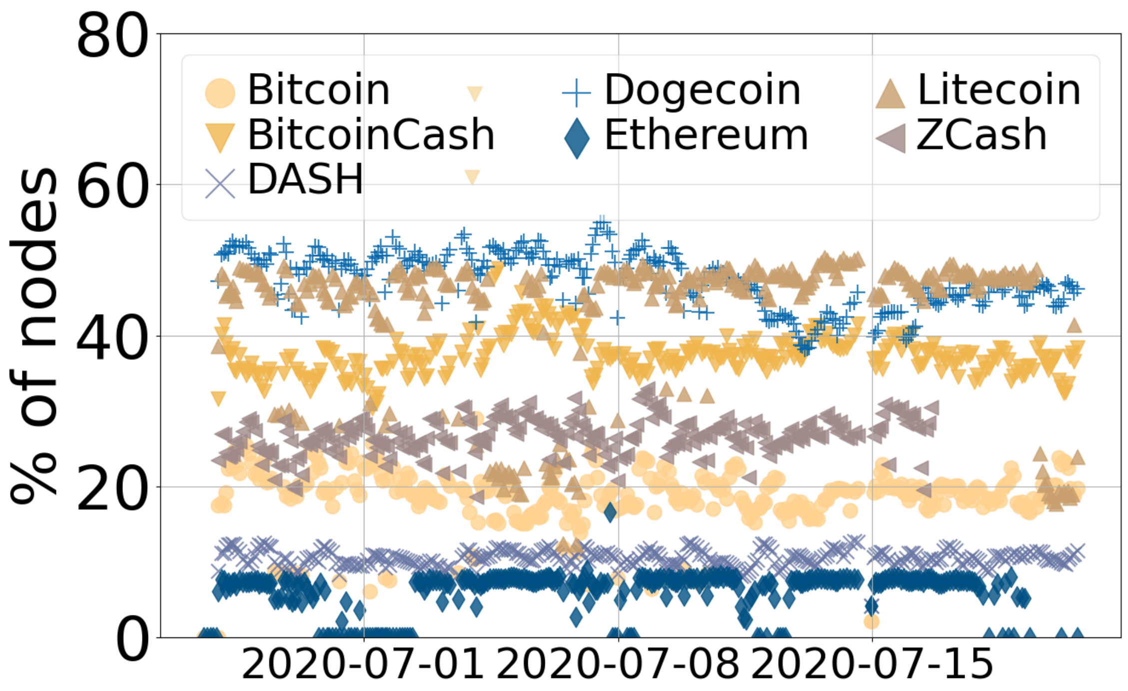

Overlapping Nodes identified by:

- In-Degree
- Page-Rank
- Betweenness centrality

Results – RQ3
Targeted Attacks - Strategy

Sort nodes according to Betweenness Centrality metric (descending)

Remove nodes one by one

Calculate size of Largest Connected Component

Targeting overlapping nodes

Removal of less than 10% of overlapping nodes

Network	Bitcoin	Bitcoin Cash	Litecoin	ZCash
Largest Connected Component Reduction	40%	70%	20%	25%

Summary

- Blockchain Overlay Networks follow are **structurally different**
- Significant number of **overlapping nodes**
- Resilience to random failures is **high**
- Resilience to targeted attacks is questionable
- Network connectivity is paramount for security => **New protocols are needed.**

Pre-print: `arxiv.org/abs/`**`2104.03044`**

Dataset: **`https://drive.google.com/drive/folders/1115O8SY8U9NLZARzhcO1Q-8Vzdn3WaSy`**

Related Work

- Maya Dotan et al. 2020. SOK: cryptocurrency networking context, state-of-the-art, challenges.
- Matthias Grundmann, Till Neudecker, Hannes Hartenstein.
  2018. Exploiting Transaction Accumulation and Double Spends for Topology Inference in Bitcoin.
- Sergi Delgado Segura, et al. 2019. TxProbe: Disco vering Bitcoin's Network Topology Using Orphan Transactions.

Related Work

- Andrew Miller et al. 2015. Coinscope: Discovering Bitcoin's Network Topology and Influential Nodes
- T. Neudecker, P. Andelfinger and H. Hartenstein. 2016. "Timing Analysis for Inferring the Topology of the Bitcoin Peer-to-Peer Network"
- Wang, Liang, and Ivan Pustogarov. "Towards better understanding of bitcoin unreachable peers."

Related Work

- Qawi K Telesford et al. 2011. The Ubiquity of Small-world Networks.
- TopoShot: uncovering Ethereum's network topology leveraging replacement transactions
  - 100% accuracy -> Very high cost ($15000 to map 1000 Ethereum nodes)

Methodology

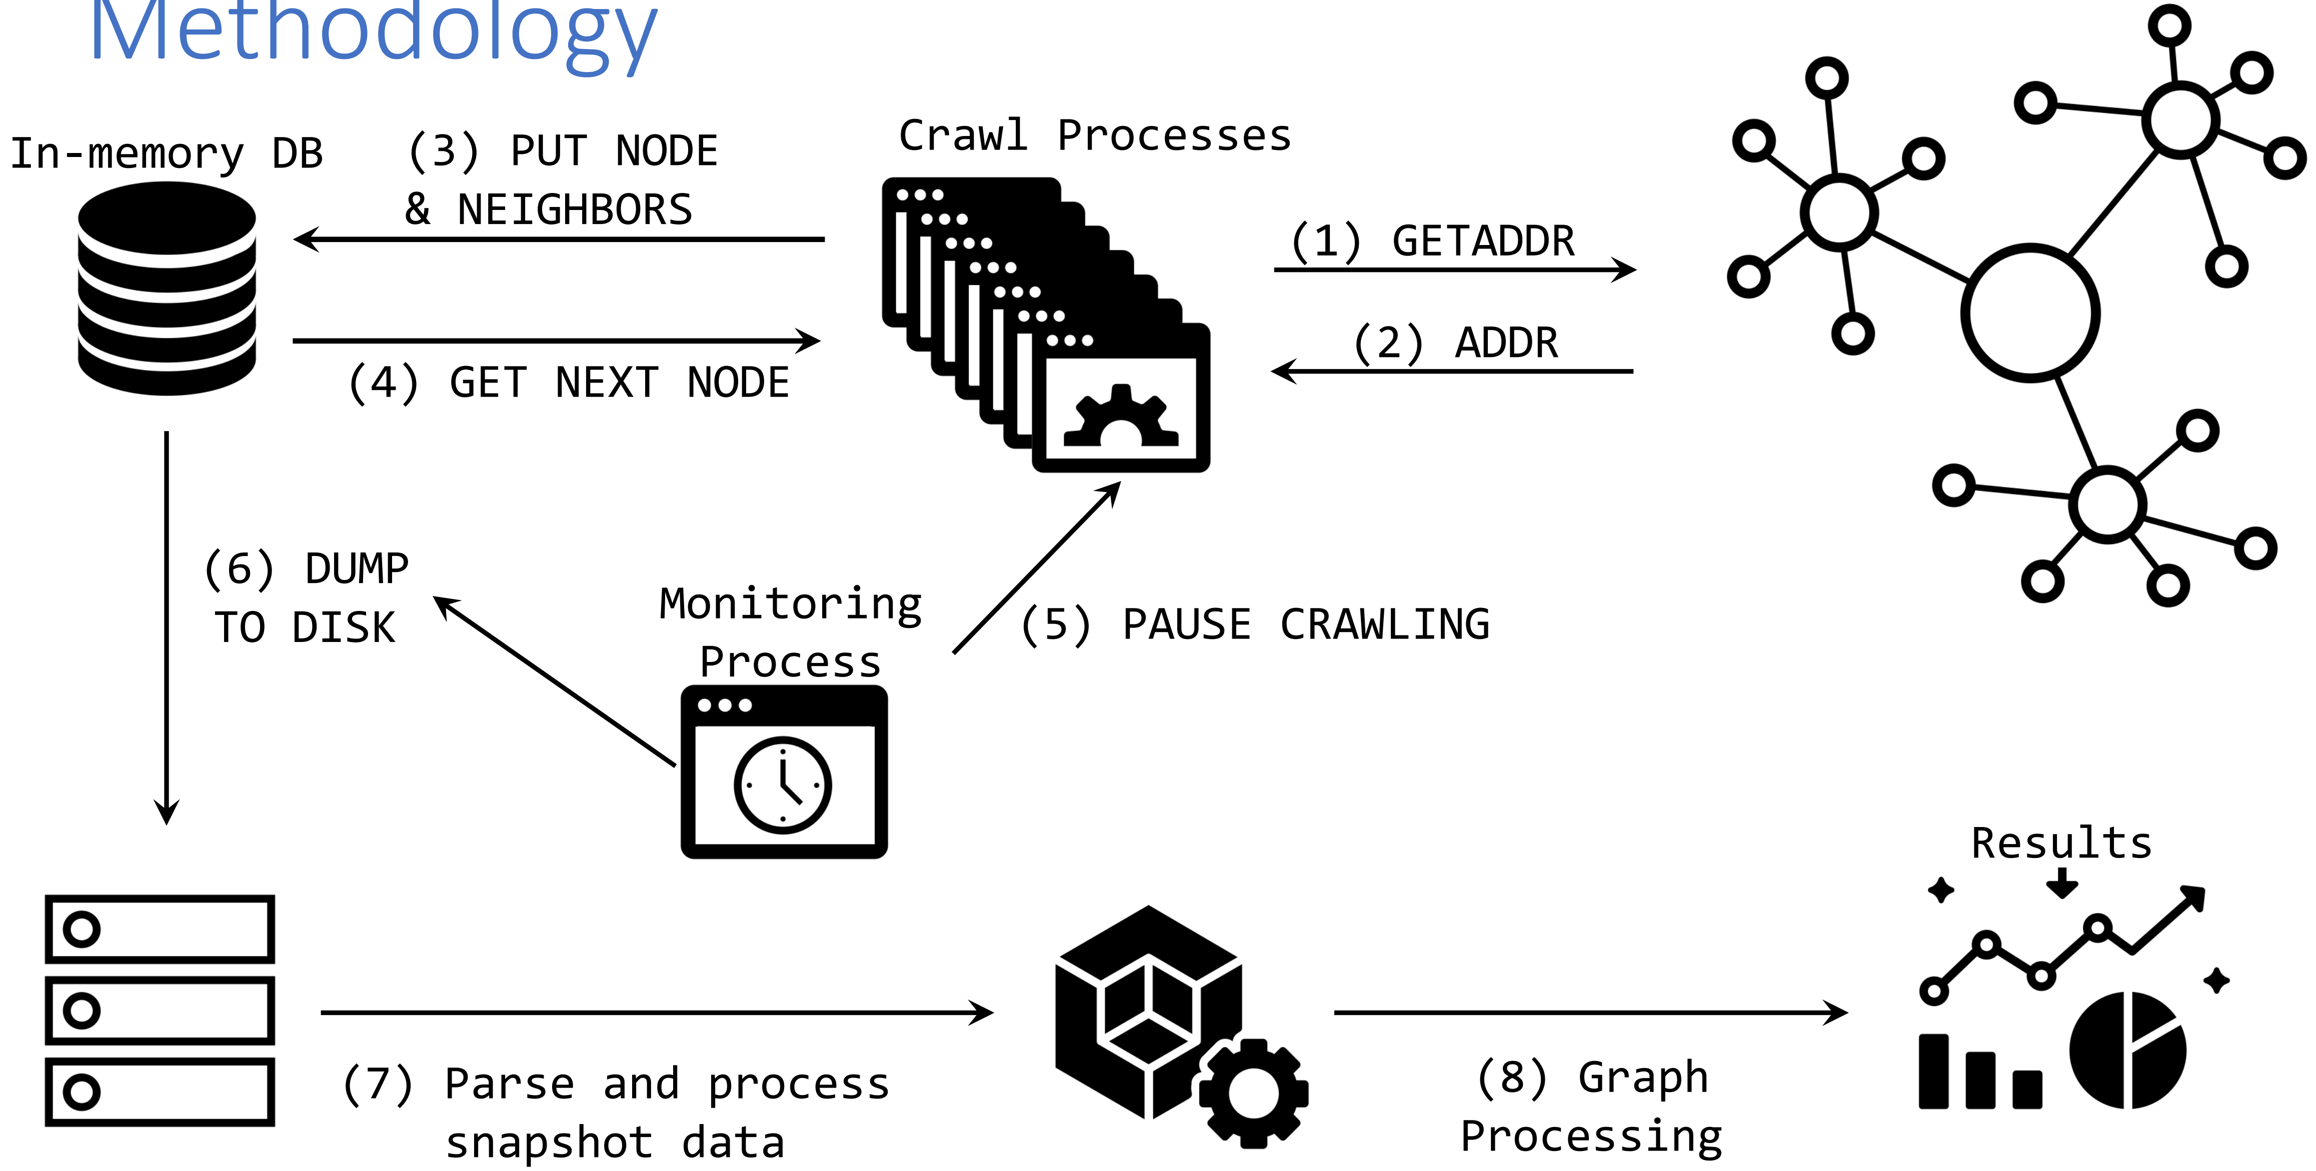

Network	Bitcoin	Bitcoin Cash	DASH	Doge	Ethereum	Litecoin	ZCash
Nodes	120k	33k	9k	2.1k	17.5k	11.7k	4.1k
Edges	37M	748k	29M	330k	556k	3.7M	231k
Conn. Comp.	1	1	1	1	0.99	1	1
SCC	0.06	0.03	0.75	0.2	0.13	0.14	0.06
Diameter	4	4	3	3	5	3	4
Density	0.004	0.001	0.5	0.11	0.004	0.047	0.024

Network	Bitcoin	Bitcoin Cash	DASH	Doge	Ethereum	Litecoin	ZCash
Avg. Degree	254	20	2370	126	31	278	48
Assortativity	-0.2	-0.64	-0.06	-0.13	-0.02	-0.01	-0.22
Reciprocity	0.32	0.21	0.49	0.34	0.02	0.27	0.25
Global CC	0. 049	0.011	0.166	0.286	0.002	0.07	0.3
Avg. Shortest Path	2.5	2.8	1.9	1.7	3.7	1.9	1.7

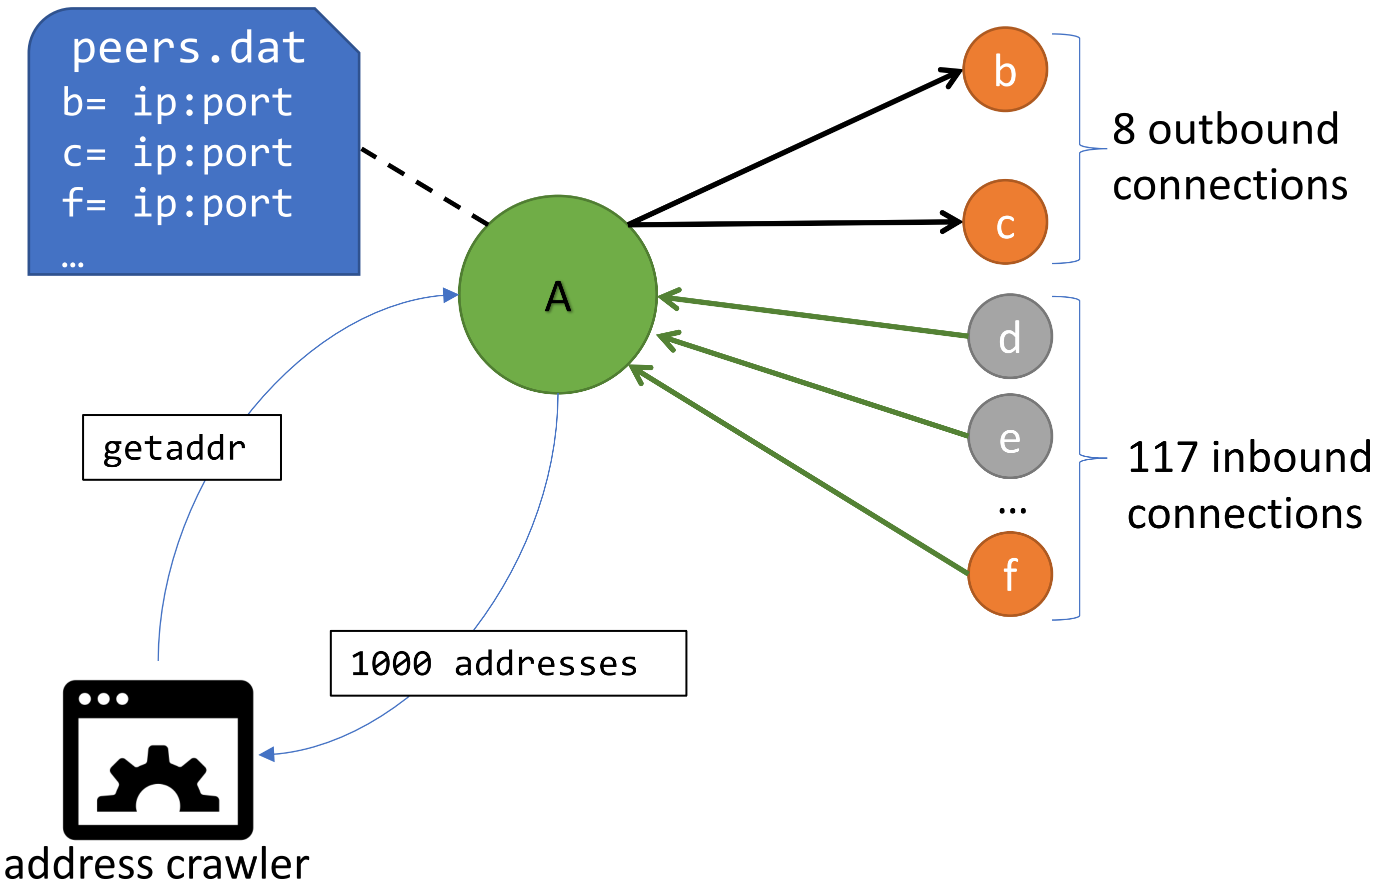

Bitcoin address management

Sample view of synthesized Connectivity Graph

BITCOIN

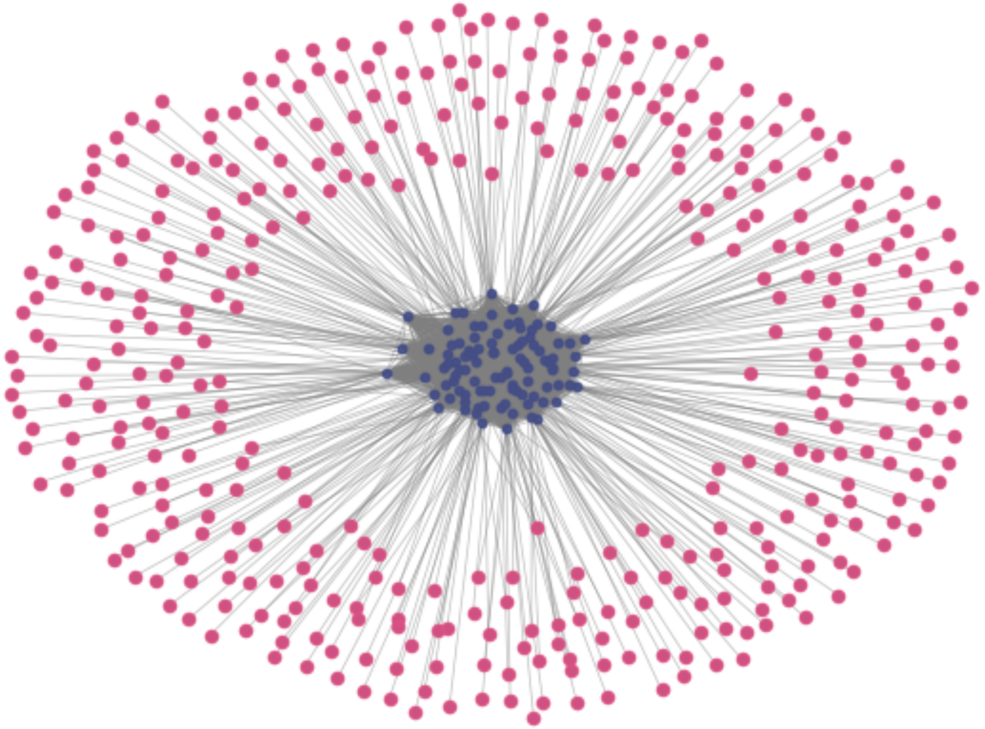

Results
Network-layer inter-dependencies

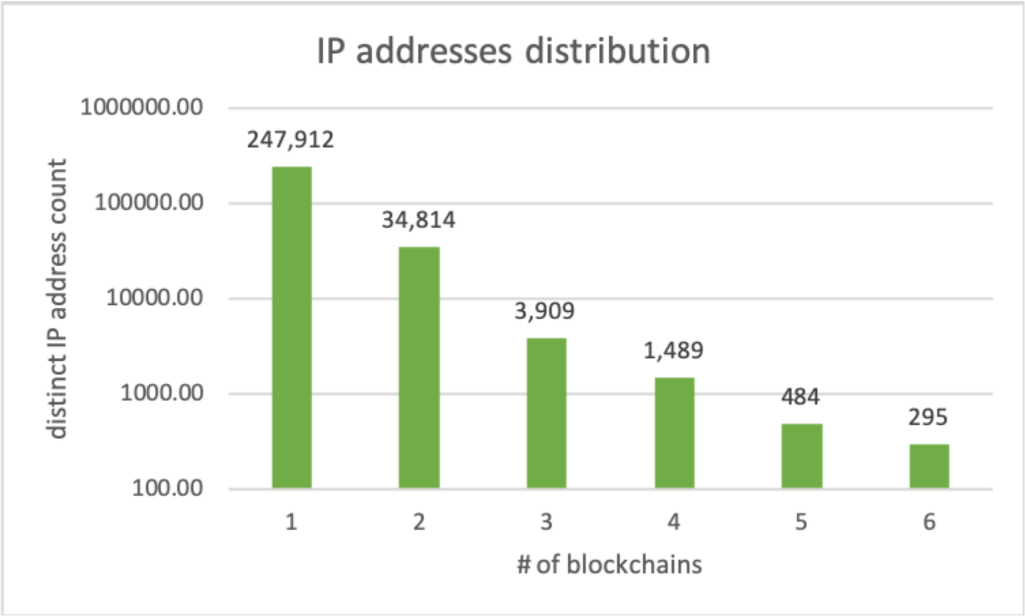

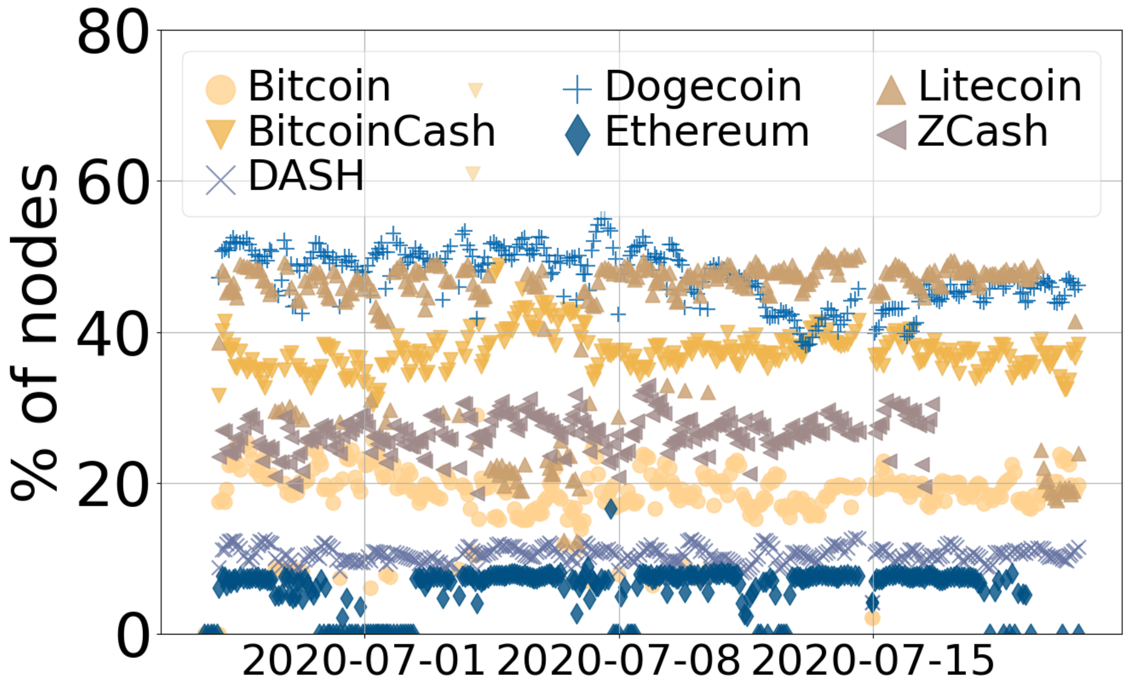

Spatial placement of nodes

- 20% of highly connected Bitcoin nodes collocated in the same AS
- Highly connected overlapping nodes collocated in a single AS
- Distribution per blockchain
  - Ethereum highly connected nodes are spread in 500 ASes
  - Bitcoin in 200
  - BitcoinCash / DASH / Dogecion in 160
  - ZCash / Litecoin in 65

Sample view of synthesized connectivity graph

BITCOIN

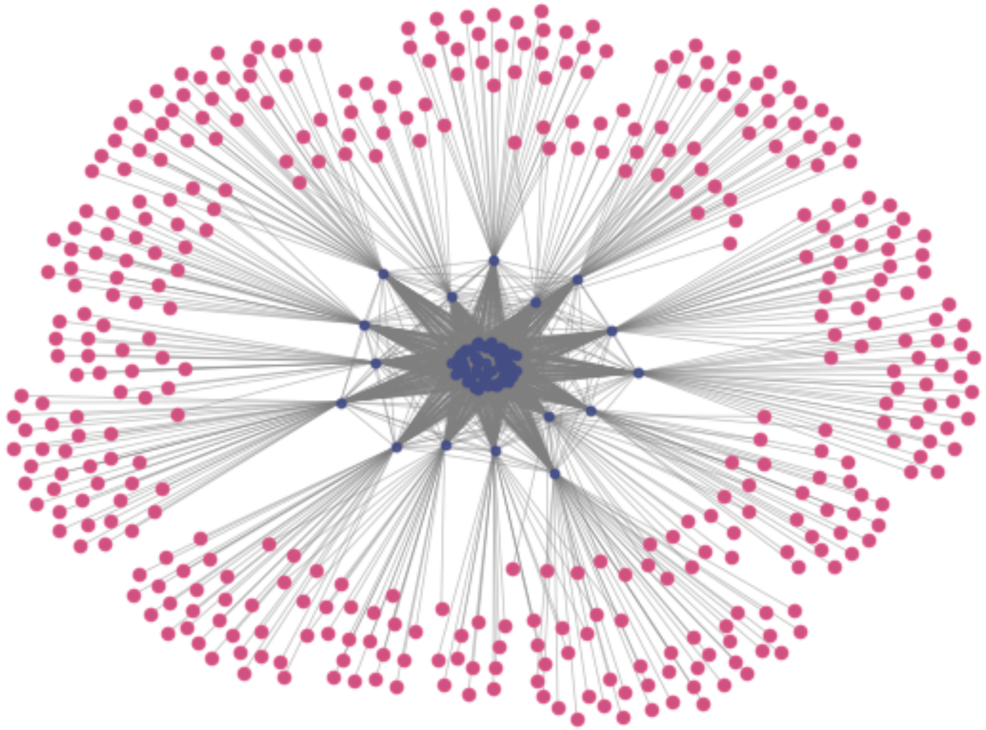

Strongly Connected Core

10x unreachable nodes to the perimeter

Removal of a single node in the core

BITCOIN

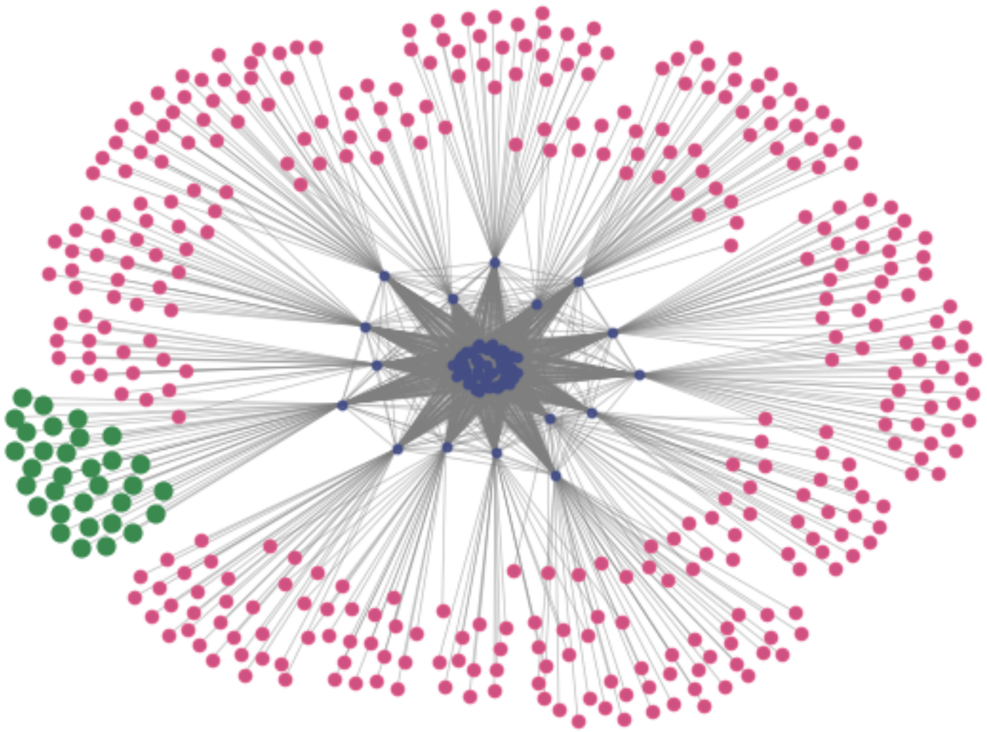

Strongly Connected Core

10x unreachable nodes to the perimeter

Unreachable peers:

establish 3.5 connections (avg)

are involved in propagation of

43% of transactions

[Wang and Pustogarov '17]

Risks in network partitioning

- Facilitate 51% attacks
- Selfish – mining
- Double spending
- Increased fork rate
- Node / Transaction censoring
- Attack based on external incentives

Summary

- Network connectivity is paramount
- Significant number of overlapping nodes
- Resilience of an artificial network with increased connectivity is easily disrupted
- Even if nodes increase connections not a great benefit is expected
- New protocols are needed!

Selected networks

Well known, established cryptocurrencies.
Frequently listed in top50 by CoinMarketCap

Bitcoin

BitcoinCash

DASH

Dogecoin

Ethereum

Litecoin

ZCash

[1] S. Delgado-Segura, C. Pérez-Solà, J. Herrera-Joancomartí, G. Navarro-Arribas, and J. Borrell, 'Cryptocurrency Networks: A New P2P Paradigm', *Mobile Information Systems*, vol. 2018, p. 2159082, Mar. 2018.

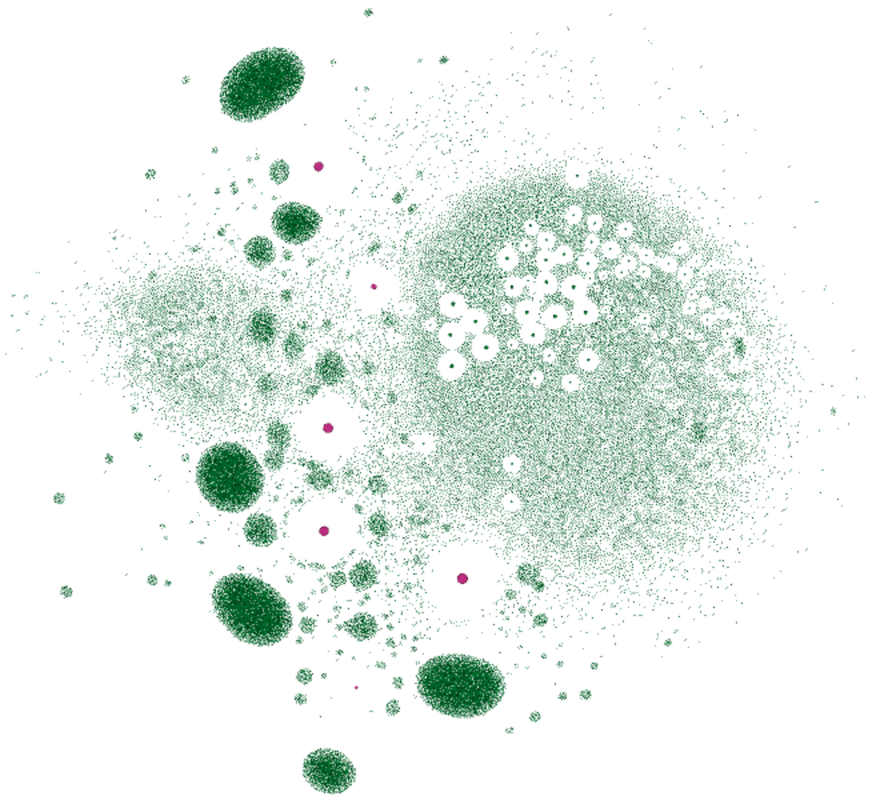

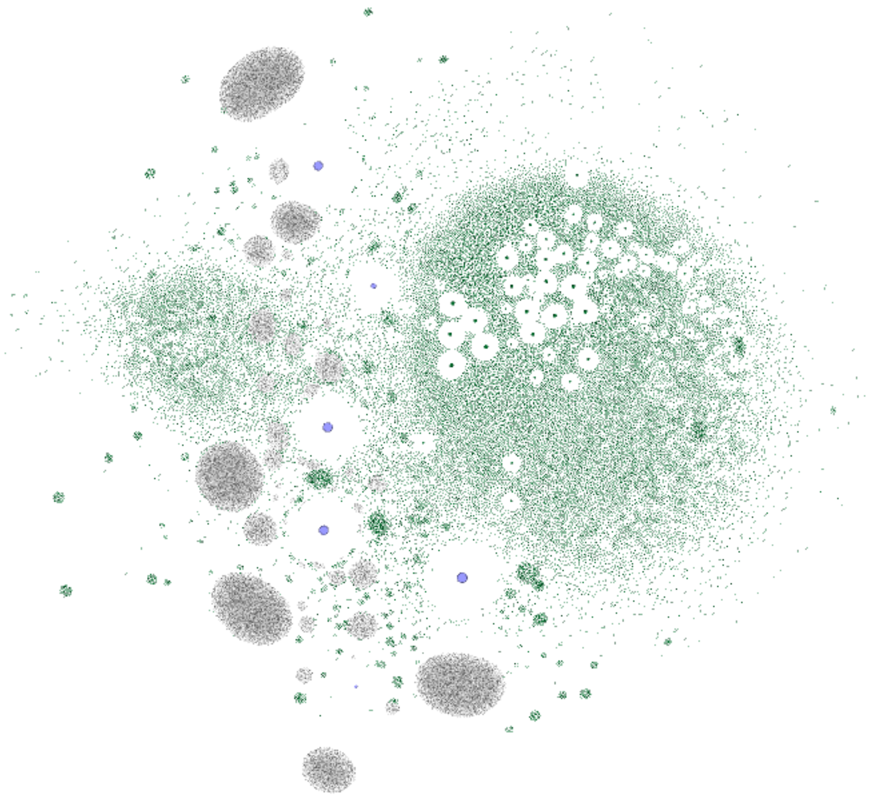

drawn using the Yifan Hu Multilevel layout algorithm, Gephi