Foundational Layer(s) for Blockchain

Blockchain Engineering 2023

Bulat Nasrulin

TU Delft

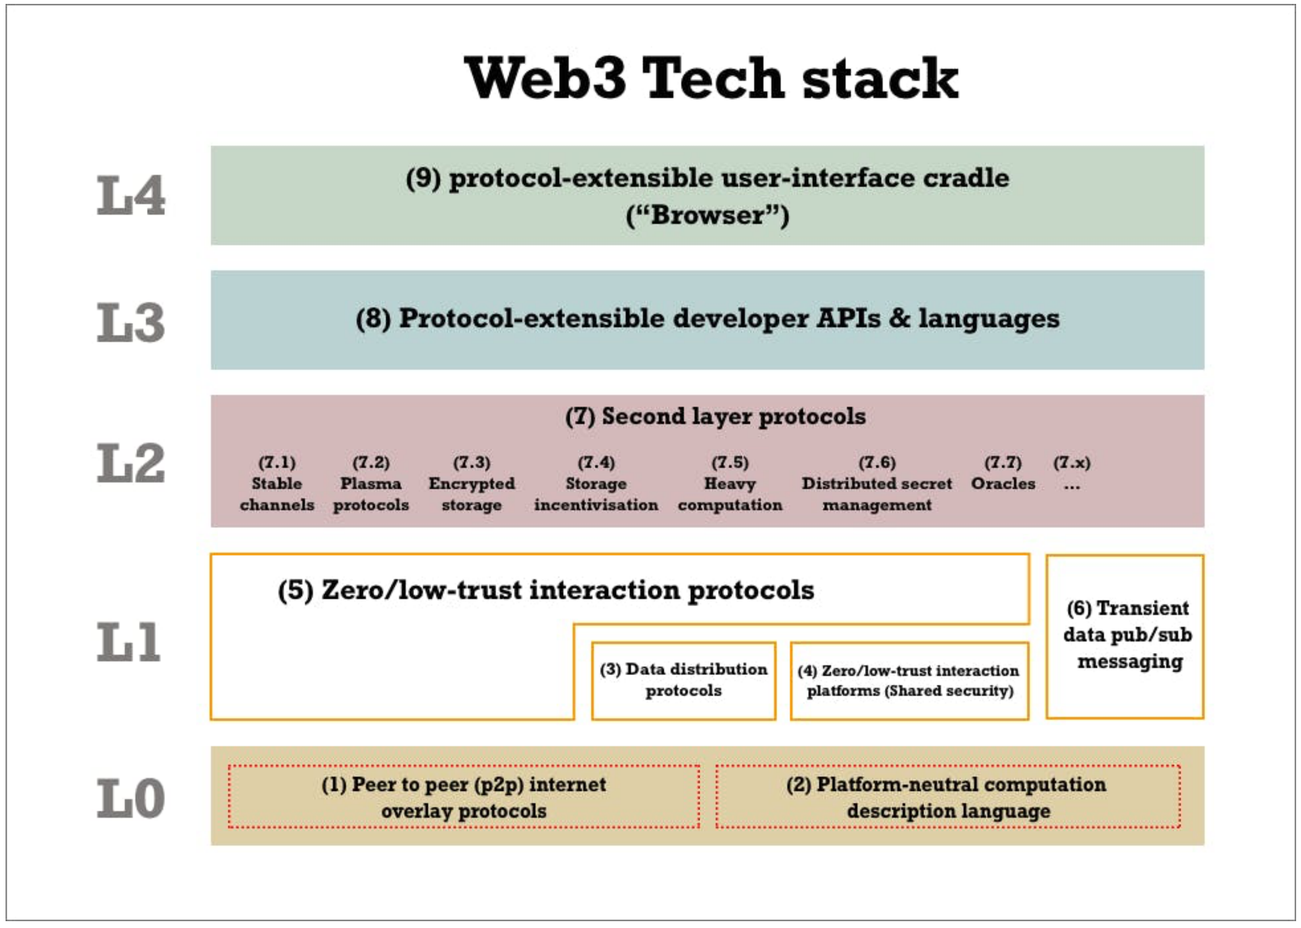

Layer 2

Scaling: channels, sidechains, rollups

Layer 1

Execution protocols

Data Consistency (Consensus) Protocols

Layer 0

Data Delivery (Gossip) Protocols

P2P Overlay Protocols

TUDelft

Blockchain simplified

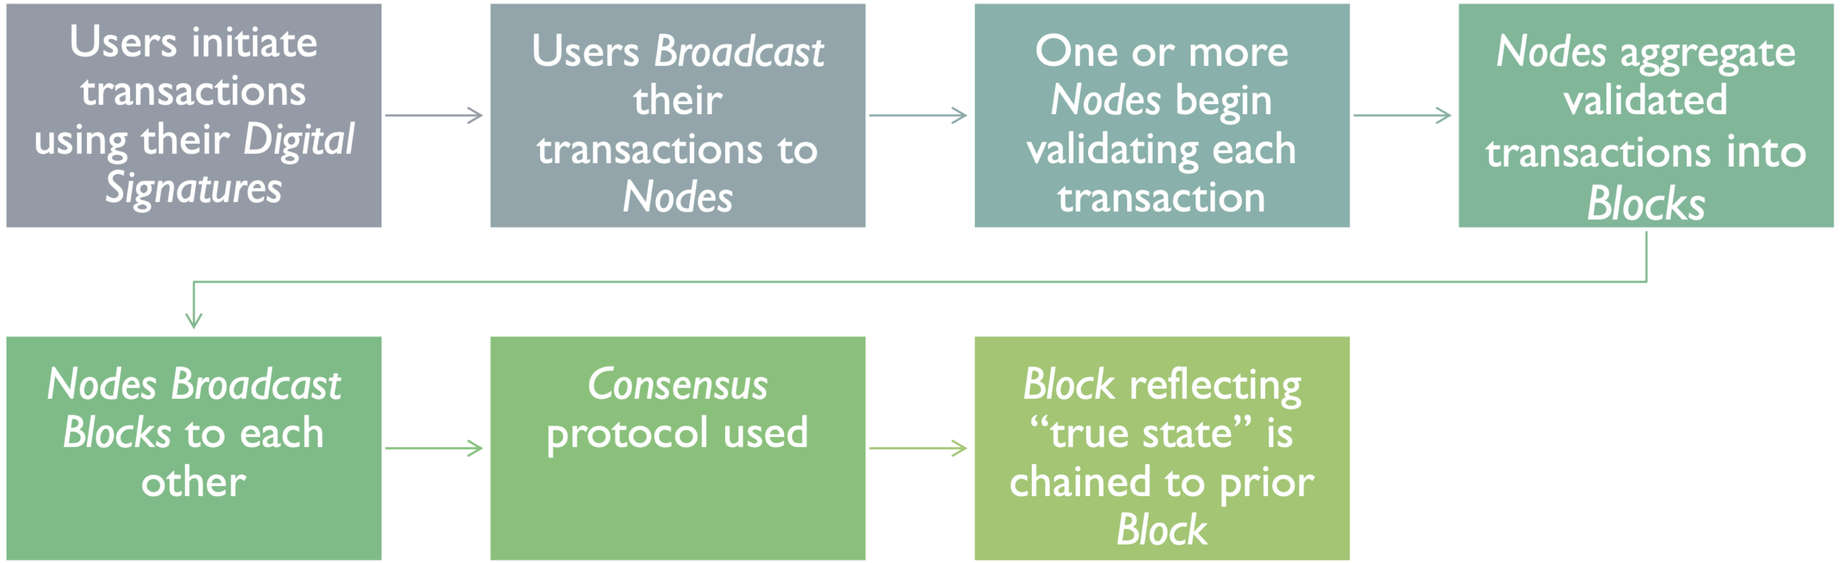

Layer 0 questions

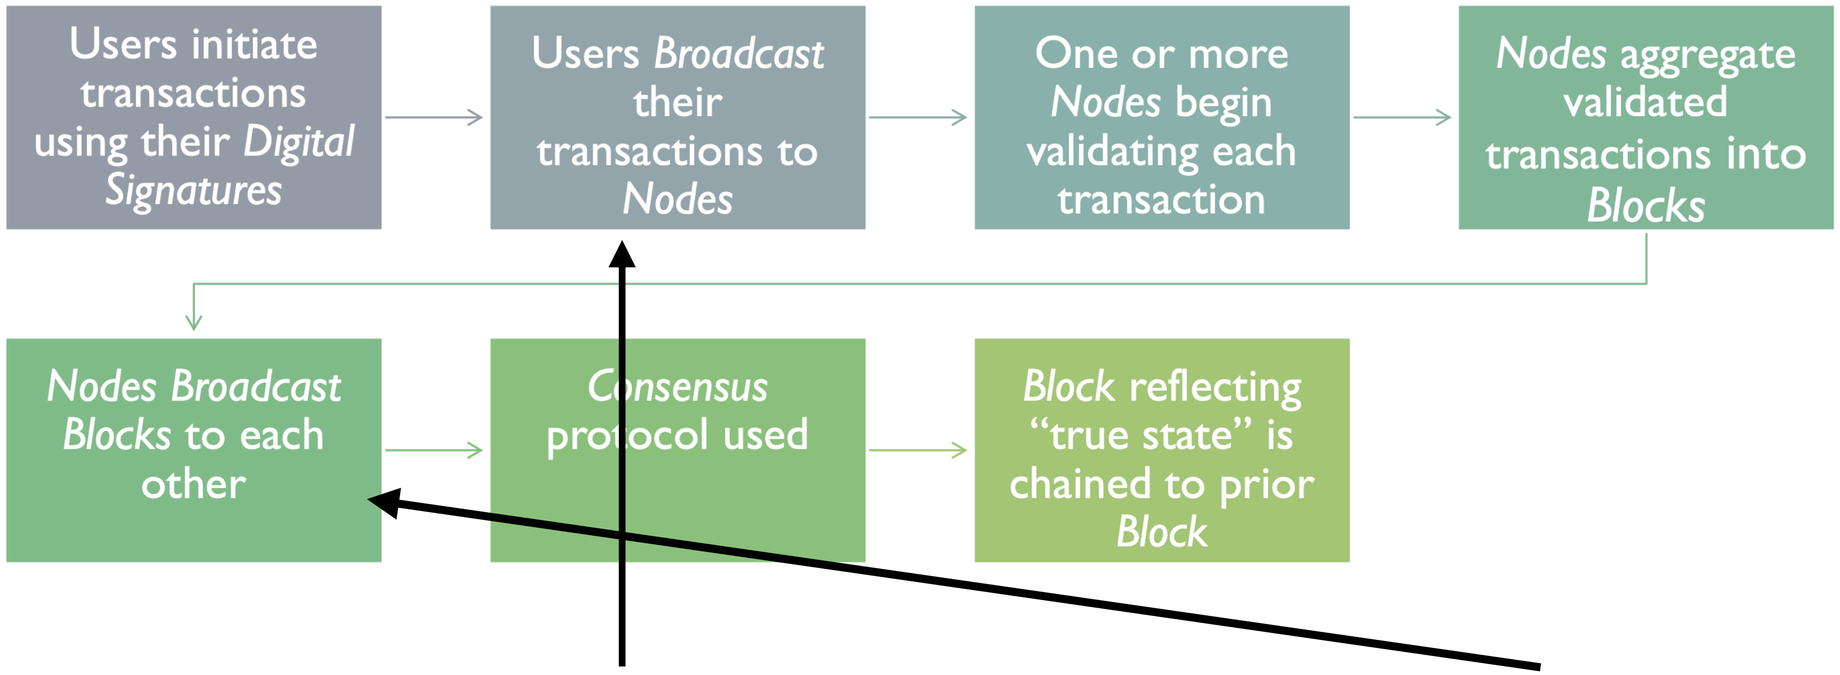

Which nodes?
How many?
All nodes need to see transactions?

How nodes are connected?
Sync all blocks? Or push one?

TUDelft

Network Architectures

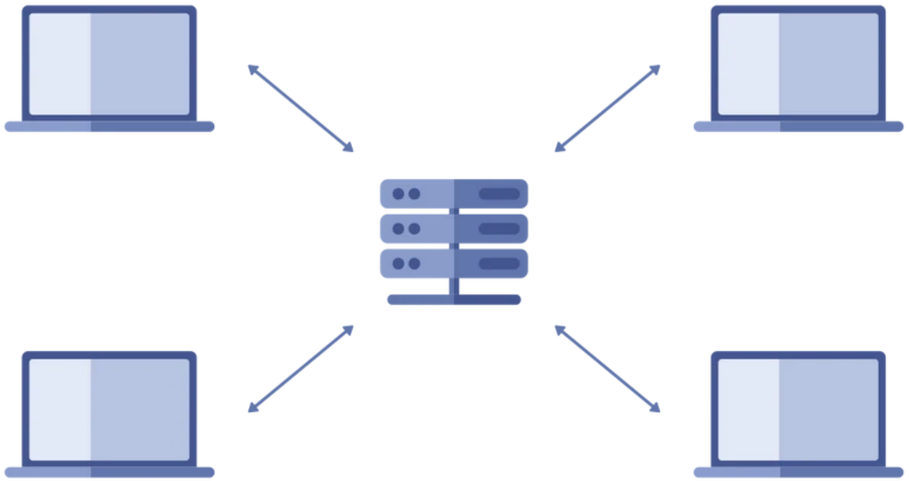

Centralized network

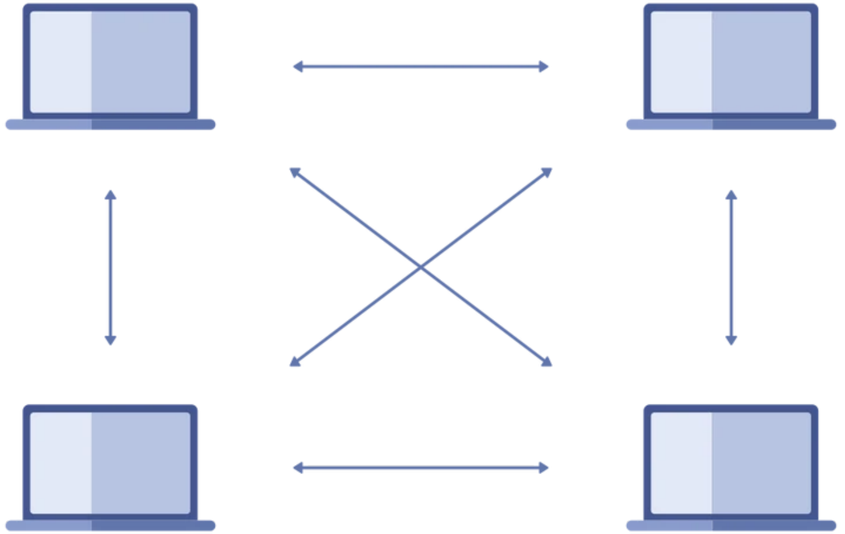

Peer-to-Peer network

TUDelft

Work together with strangers.

Get evil corporations out of your technology stack.

TU Delft

Desirable Properties of P2P Collaboration

Open join and leave

Scalable

Minimal overhead

Robust to Churn, Failures

P2P Evolution

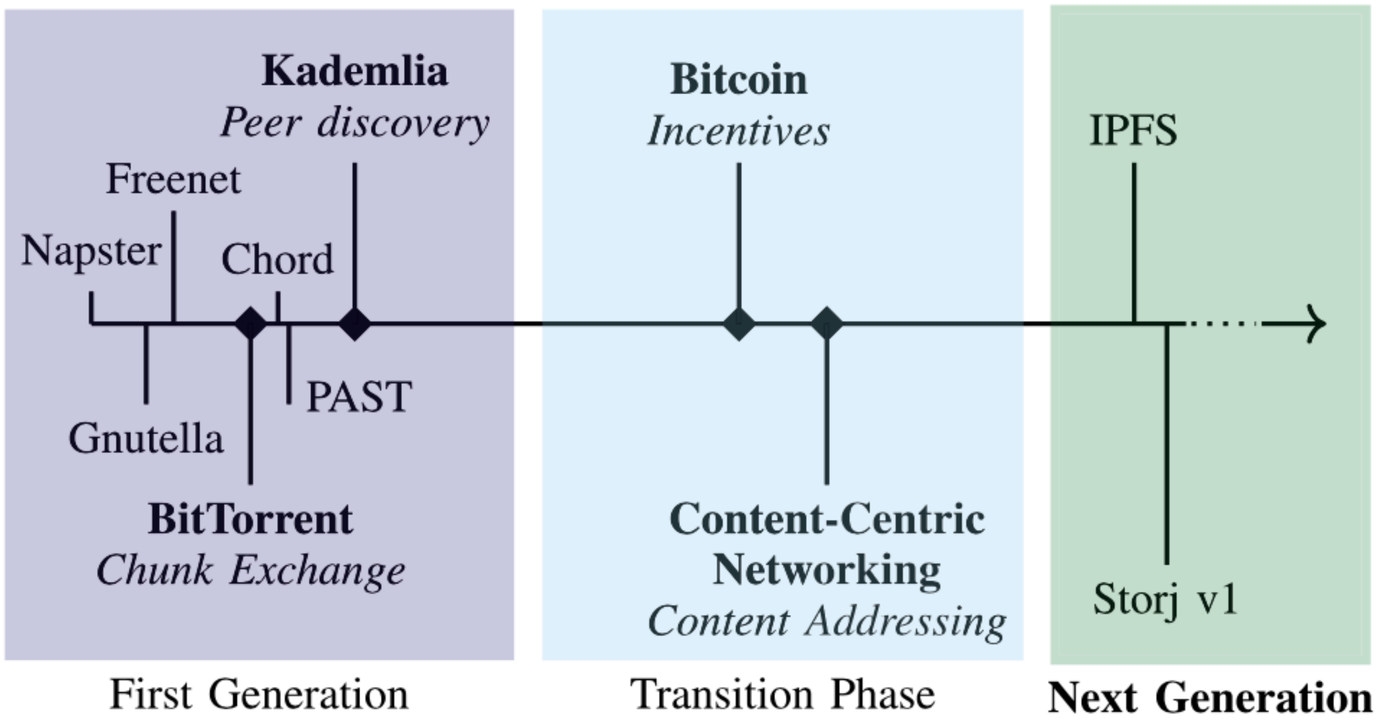

Rich History and many proposals

Two main ways to build overlays

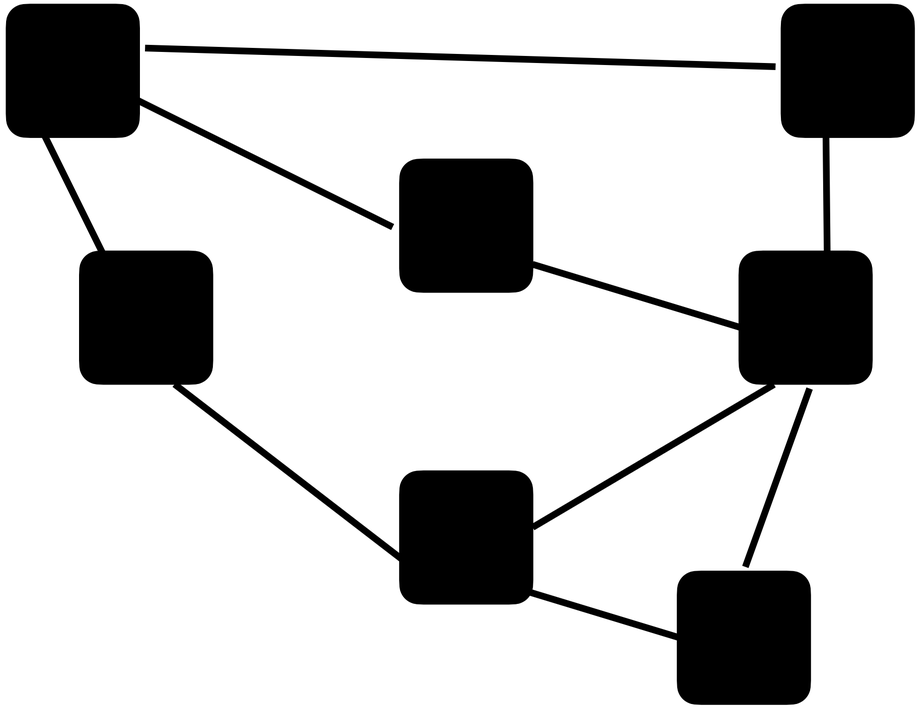

Structured overlay

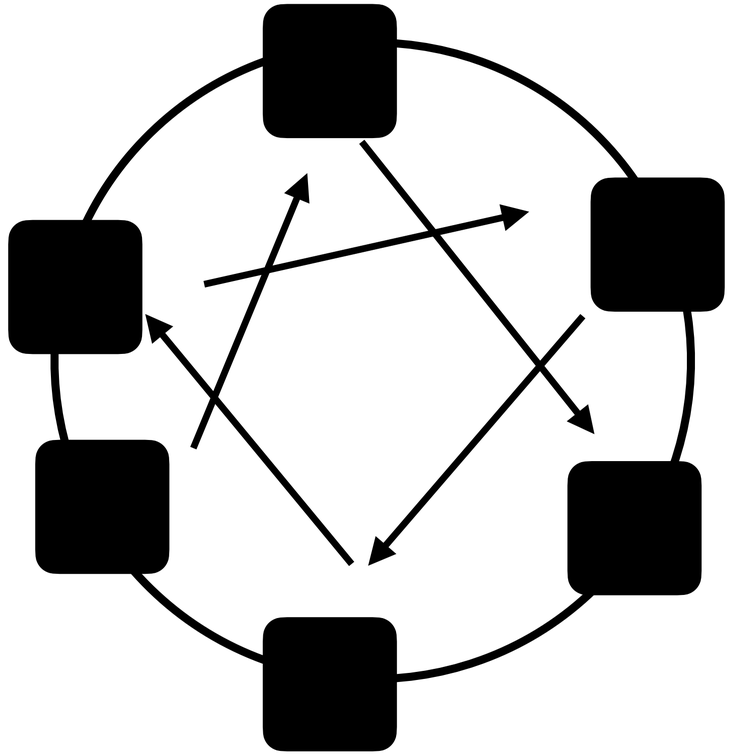

Unstructured overlay

Two main ways to build overlays

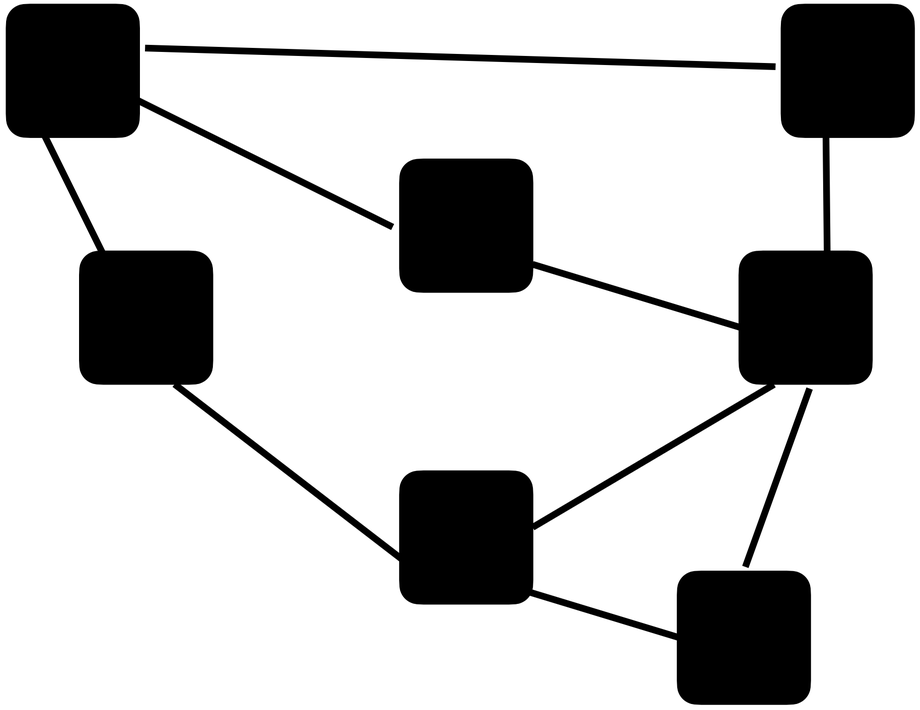

Unstructured overlay

- Each peer independently connects to random k peers
- Random network with ad-hoc protocols for search and storage

Two main ways to build overlays

- Use id of the peers to build structured overlay (like closest ids)
- Structured index
- Rigid organizational principles for search, storage etc.

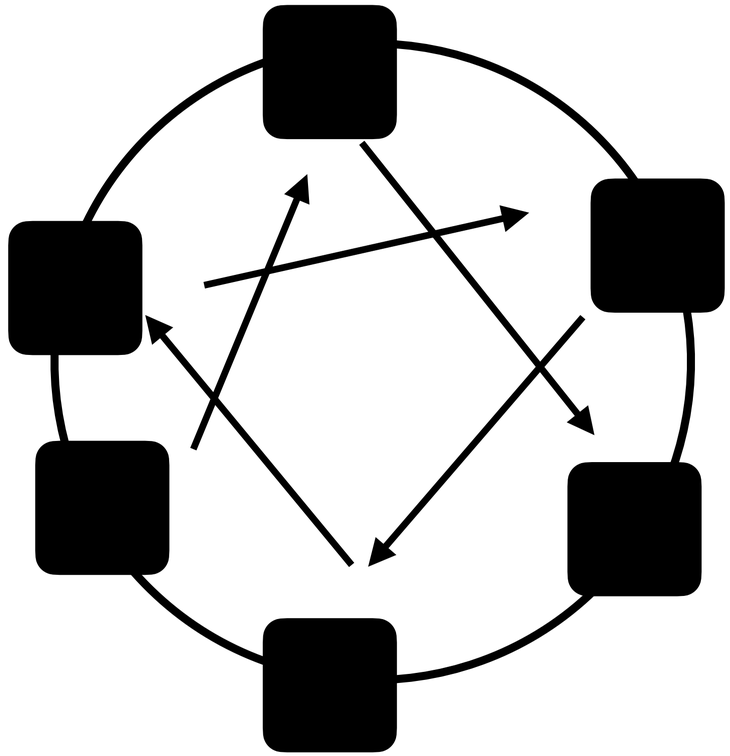

Unstructured overlay

What is better for blockchain?

Structured overlay **Unstructured overlay**

Open join and leave?

Minimal overhead?

Robust to Churn, Failures?

Scalable?

Bitcoin Network

Bitcoin Network

REACHABLE BITCOIN NODES

Updated: Mon Feb 27 15:56:30 2023 CET

16308 NODES CHARTS

IPv4: +3.0% / IPv6: +3.4% / .onion: +21.8%

Top 10 countries with their respective number of reachable nodes are as follows.

RANK	COUNTRY	NODES
1	n/a	9618 (58.98%)
2	United States	1838 (11.27%)
3	Germany	1388 (8.51%)
4	Netherlands	405 (2.48%)
5	France	371 (2.27%)
6	United Kingdom	304 (1.86%)
7	Canada	277 (1.70%)
8	Finland	242 (1.48%)
9	Russian Federation	185 (1.13%)
10	Switzerland	128 (0.78%)

All (91) »

NOTE / The data above includes reachable Bitcoin nodes only. View combined estimation of reachable and unreachable Bitcoin nodes »

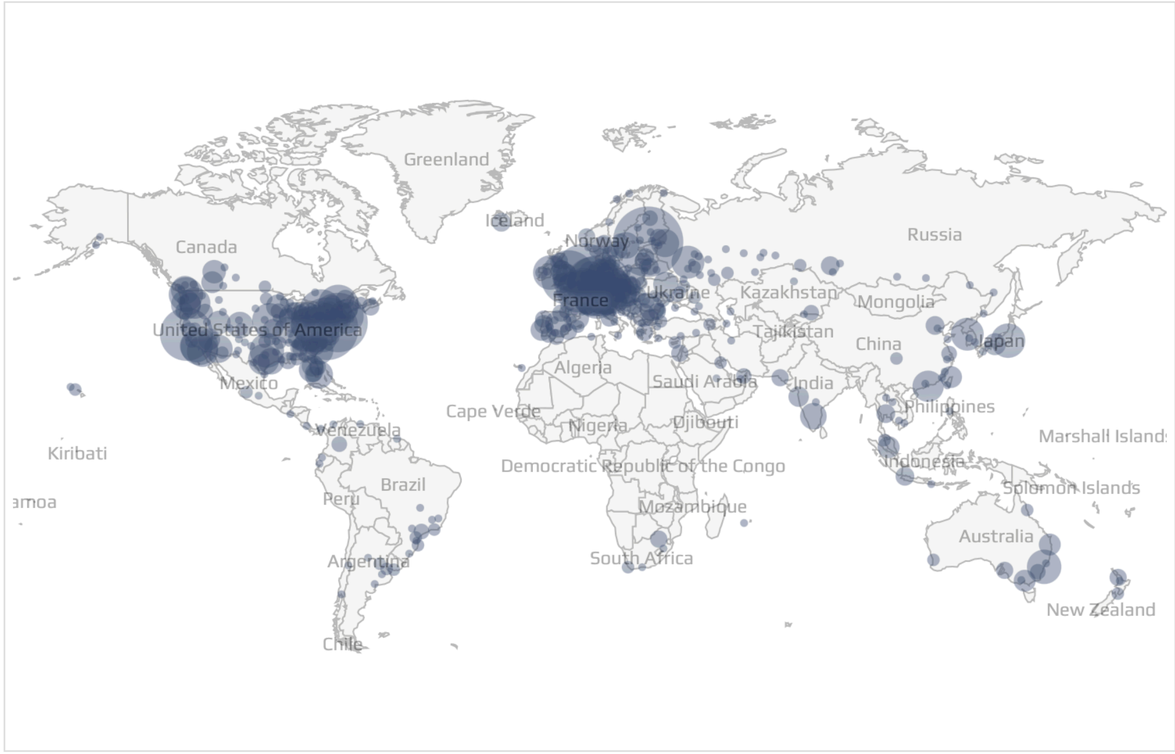

Map shows concentration of reachable Bitcoin nodes found in countries around the world. LIVE MAP

TUDelft

Bitcoin Connection and Bootstrap

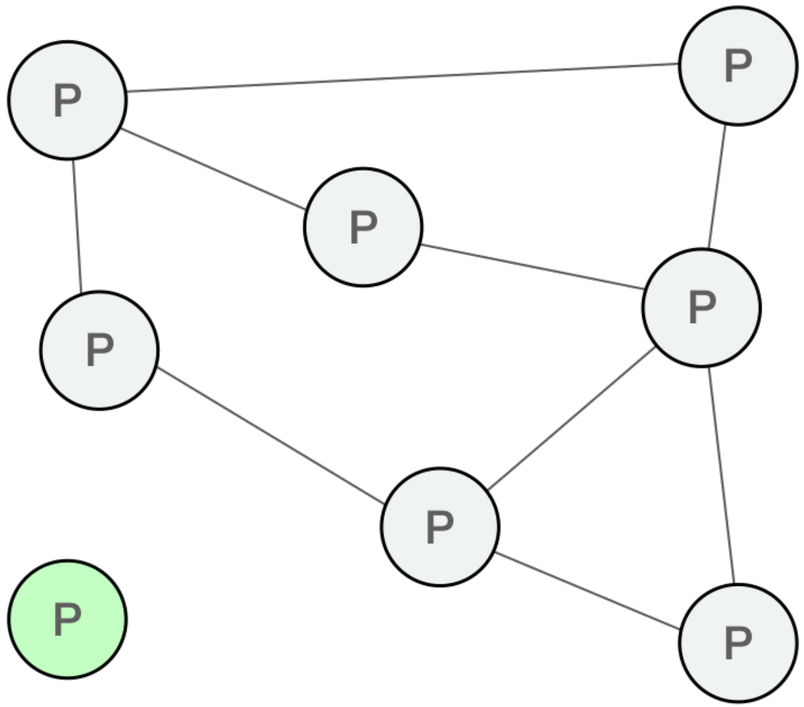

Initialization of any p2p network

Handshake Introduction

```
// From: https://github.com/bitcoin/bitcoin/blob/master/src/chainparams.cpp

vSeeds.emplace_back("seed.bitcoin.sipa.be"); // Pieter Wuille, only supports
vSeeds.emplace_back("dnsseed.bluematt.me"); // Matt Corallo, only supports x
vSeeds.emplace_back("dnsseed.bitcoin.dashjr.org"); // Luke Dashjr
vSeeds.emplace_back("seed.bitcoinstats.com"); // Christian Decker, supports
vSeeds.emplace_back("seed.bitcoin.jonasschnelli.ch"); // Jonas Schnelli, onl
vSeeds.emplace_back("seed.btc.petertodd.org"); // Peter Todd, only supports
vSeeds.emplace_back("seed.bitcoin.sprovoost.nl"); // Sjors Provoost
vSeeds.emplace_back("dnsseed.emzy.de"); // Stephan Oeste
```

Hardcoded DNS servers for Bootstrap

Bitcoin Connection and Bootstrap

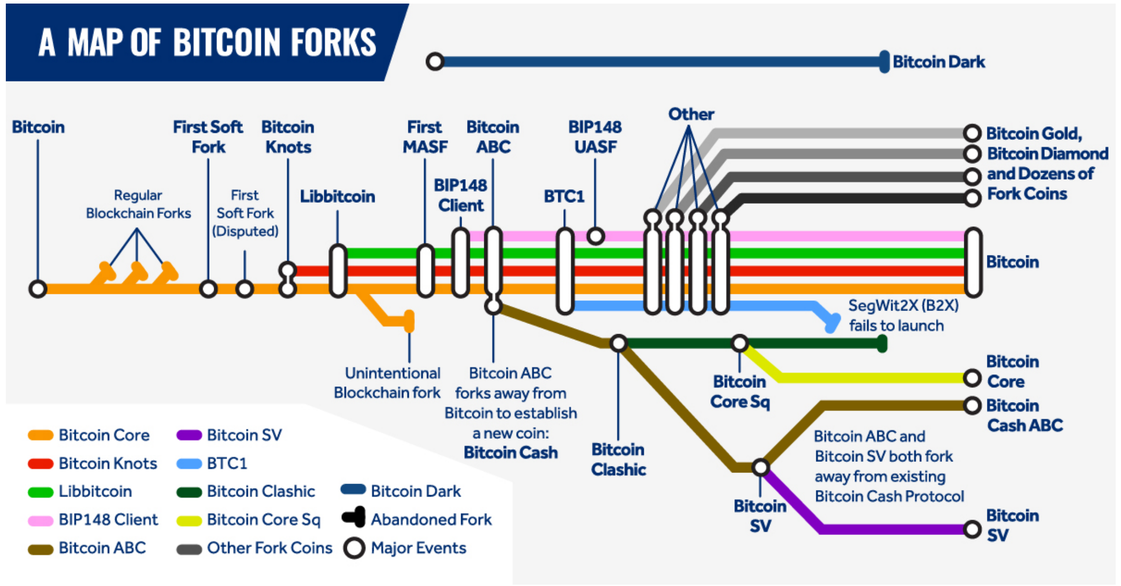

1. Exchange versions. Connect if compatible

2. Peers keep a list of active peers

Connected peers

> In order to maintain a connection with a peer, nodes by default will send a message to peers before 30 minutes of inactivity. If 90 minutes pass without a message being received by a peer, the client will assume that connection has closed

Known active peers

> The typical presumption is that a node is likely to be active if it has been sending a message within the last three hours.

3. Request ‘getaddr’ list of active peers from other

TUDelft

Bitcoin network messages

Messages

- *version* - Information about program version and block count. Exchanged when first connecting.
- *verack* - Sent in response to a version message to acknowledge that we are willing to connect.
- *addr* - List of one or more IP addresses and ports.
- *inv* - "I have these blocks/transactions: ..." Normally sent only when a *new* block or transaction is being relayed. This is only a list, not the actual data.
- *getdata* - Request a single block or transaction by hash.
- *getblocks* - Request an *inv* of all blocks in a range.
- *getheaders* - Request a *headers* message containing all block headers in a range.
- *tx* - Send a transaction. This is sent only in response to a *getdata* request.
- *block* - Send a block. This is sent only in response to a *getdata* request.
- *headers* - Send up to 2,000 block headers. Non-generators can download the headers of blocks instead of entire blocks.
- *getaddr* - Request an *addr* message containing a bunch of known-active peers (for bootstrapping).
- *submitorder*, *checkorder*, and *reply* - Used when performing an IP transaction.
- *alert* - Send a network alert.
- *ping* - Does nothing. Used to check that the connection is still online. A TCP error will occur if the connection has died.

Bitcoin network messages

Only 4 classes

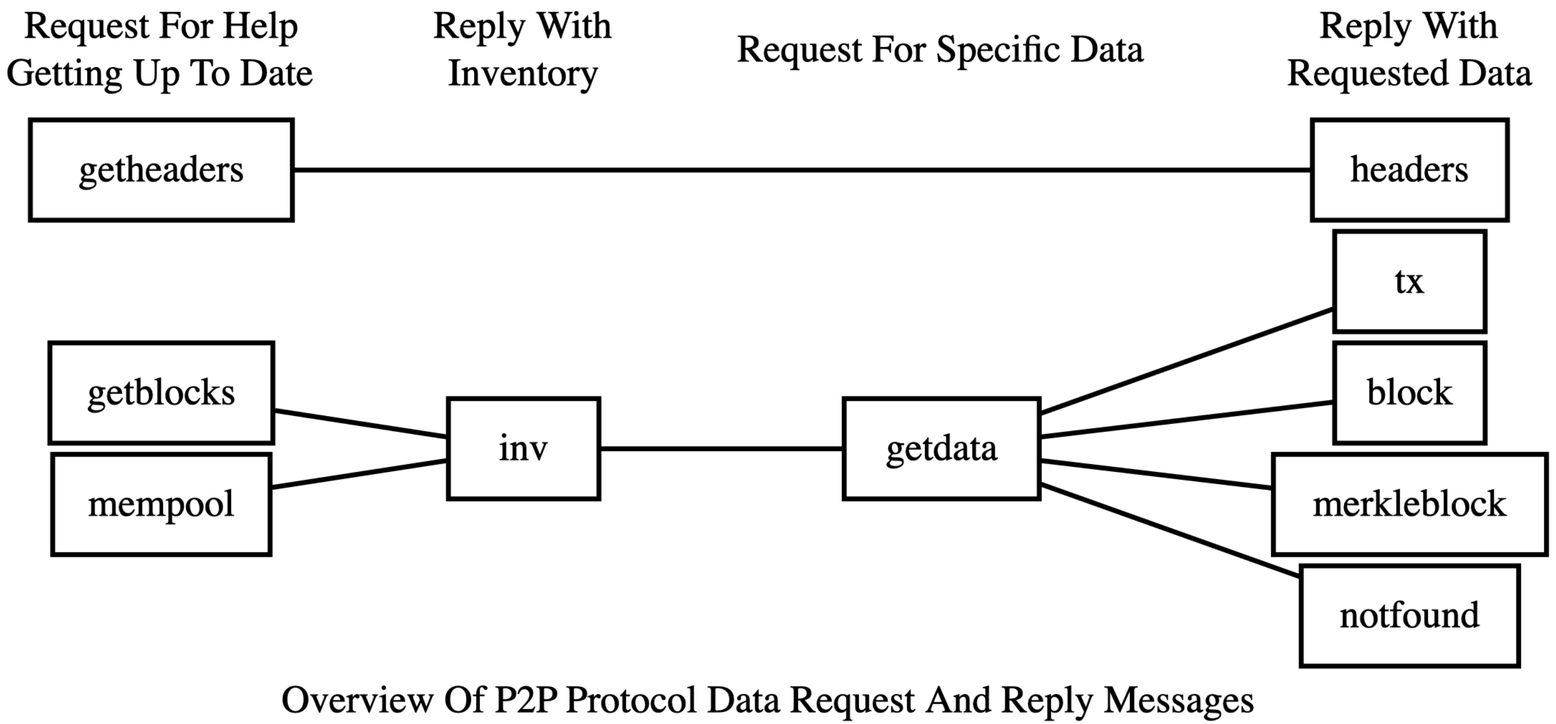

Overview Of P2P Protocol Data Request And Reply Messages

Same network used for transaction and block propagation

Block Propagation

Two methods of blockchain sync

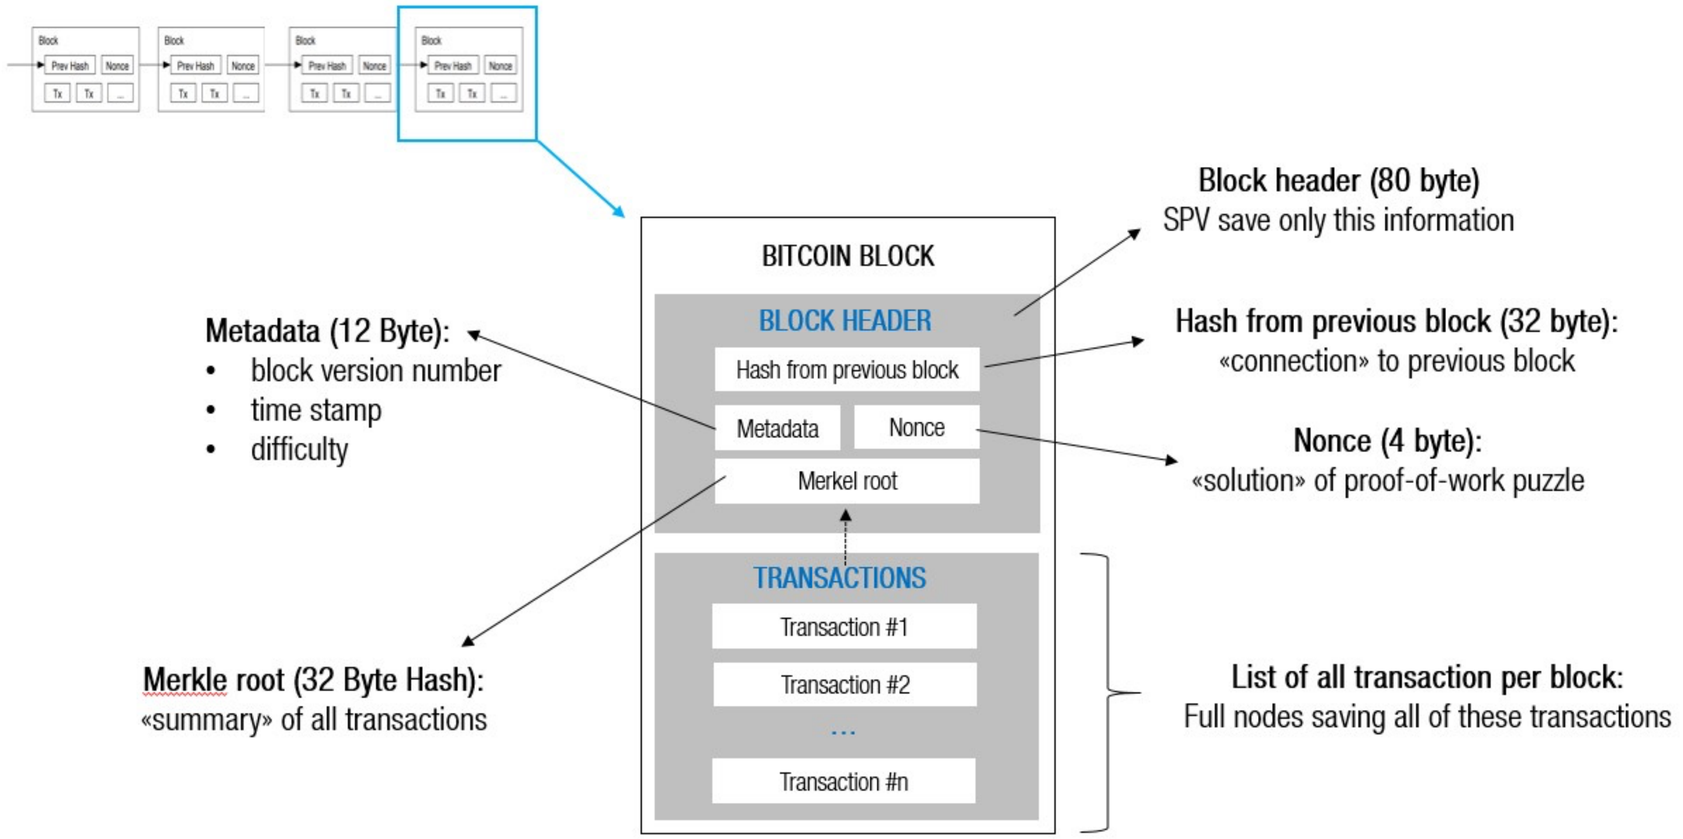

Block First

Sequentially download blocks

Headers First

First download all headers

Request block if needed

TUDelft

First Block Sync: Introduction Sync

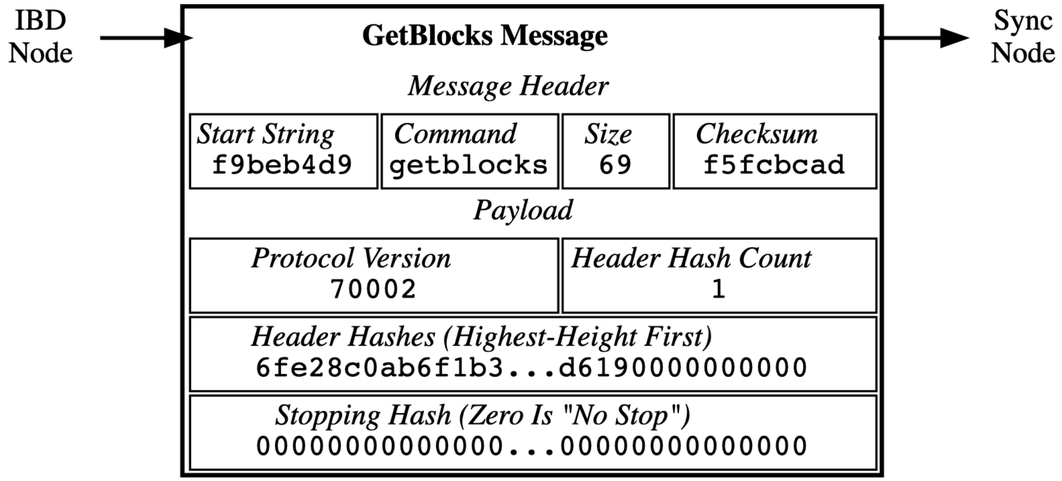

First getblocks message sent from Initial Blocks Download (IBD) node

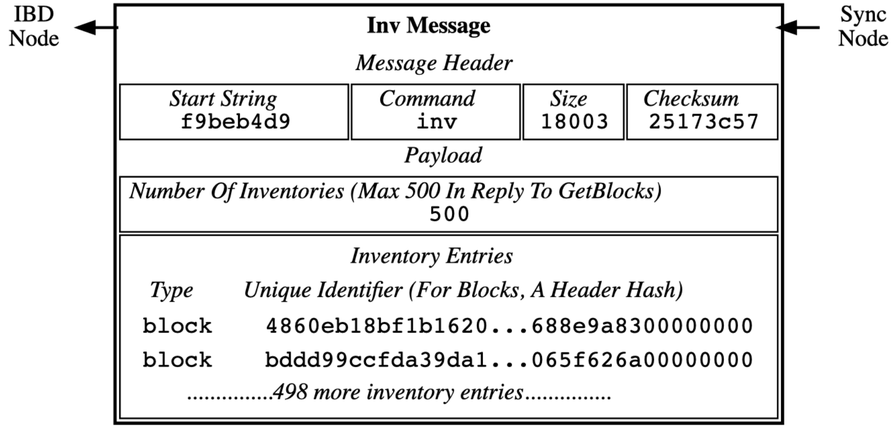

First inv message reply sent to Initial Blocks Download (IBD) node

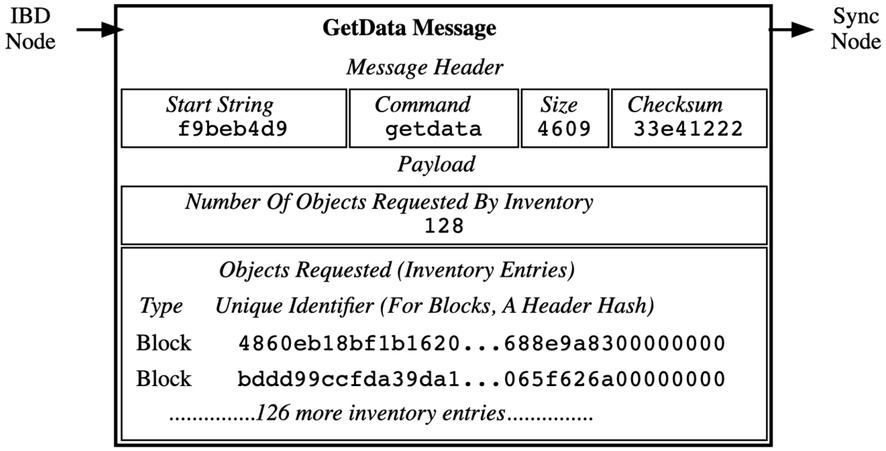

First getdata message sent from Initial Blocks Download (IBD) node

Initial Block Download

Active in Network: Block Broadcasting

Unsolicited Block Push

Low Latency
Just block with content
Push Gossip

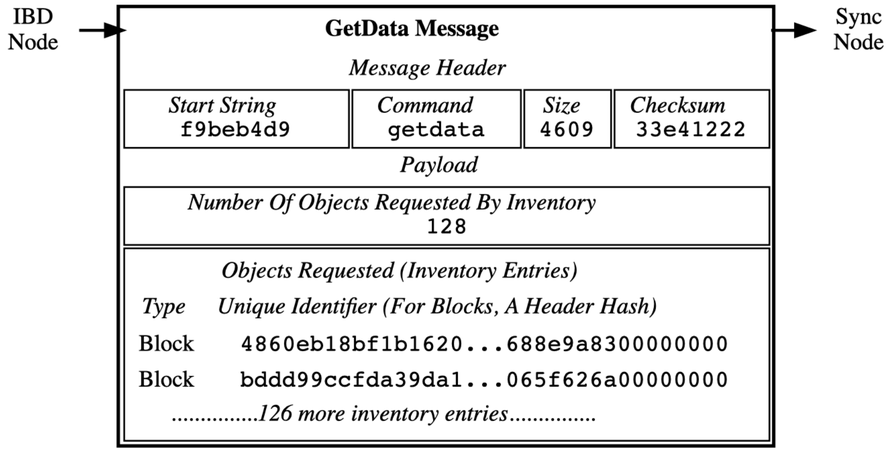

First getdata message sent from Initial Blocks Download (IBD) node

Standard Relay

High Latency
Additional announcement round
Reconciliation Gossip

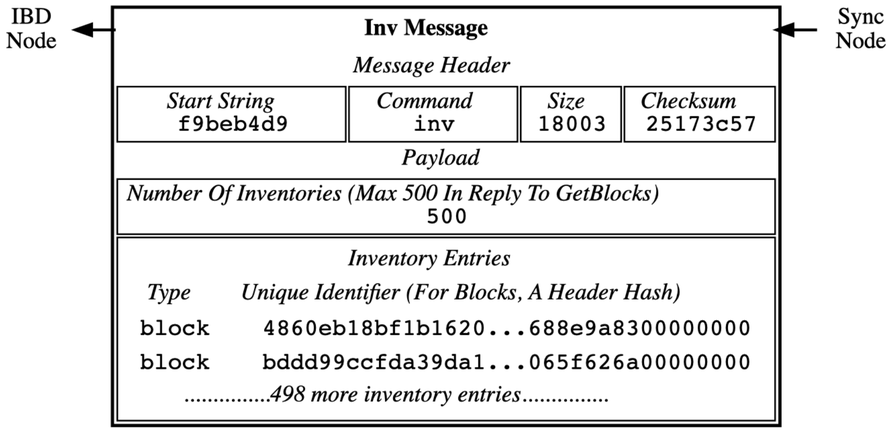

First inv message reply sent to Initial Blocks Download (IBD) node

TUDelft

Transaction Propagation

Bitcoin Transaction LifeCycle

Simple story

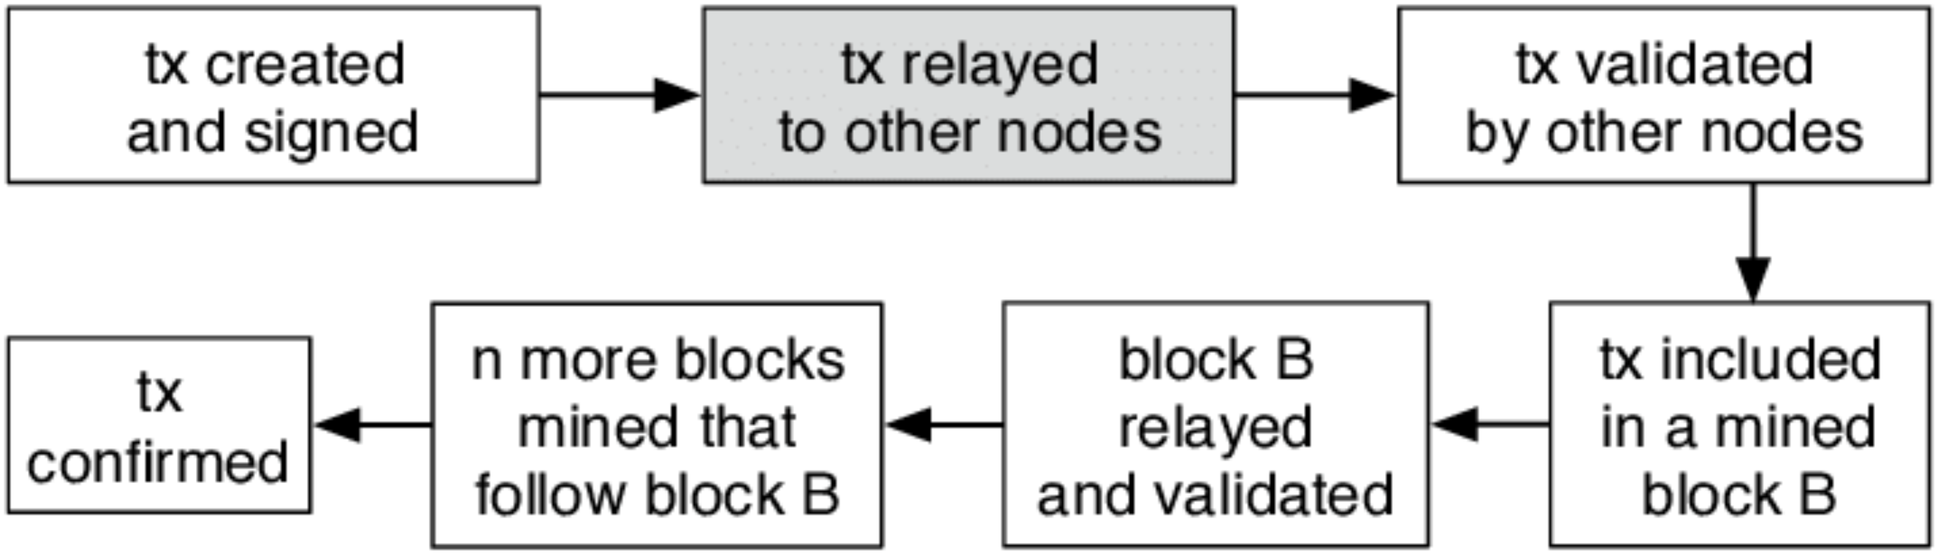

Transaction Broadcasting: Mempool

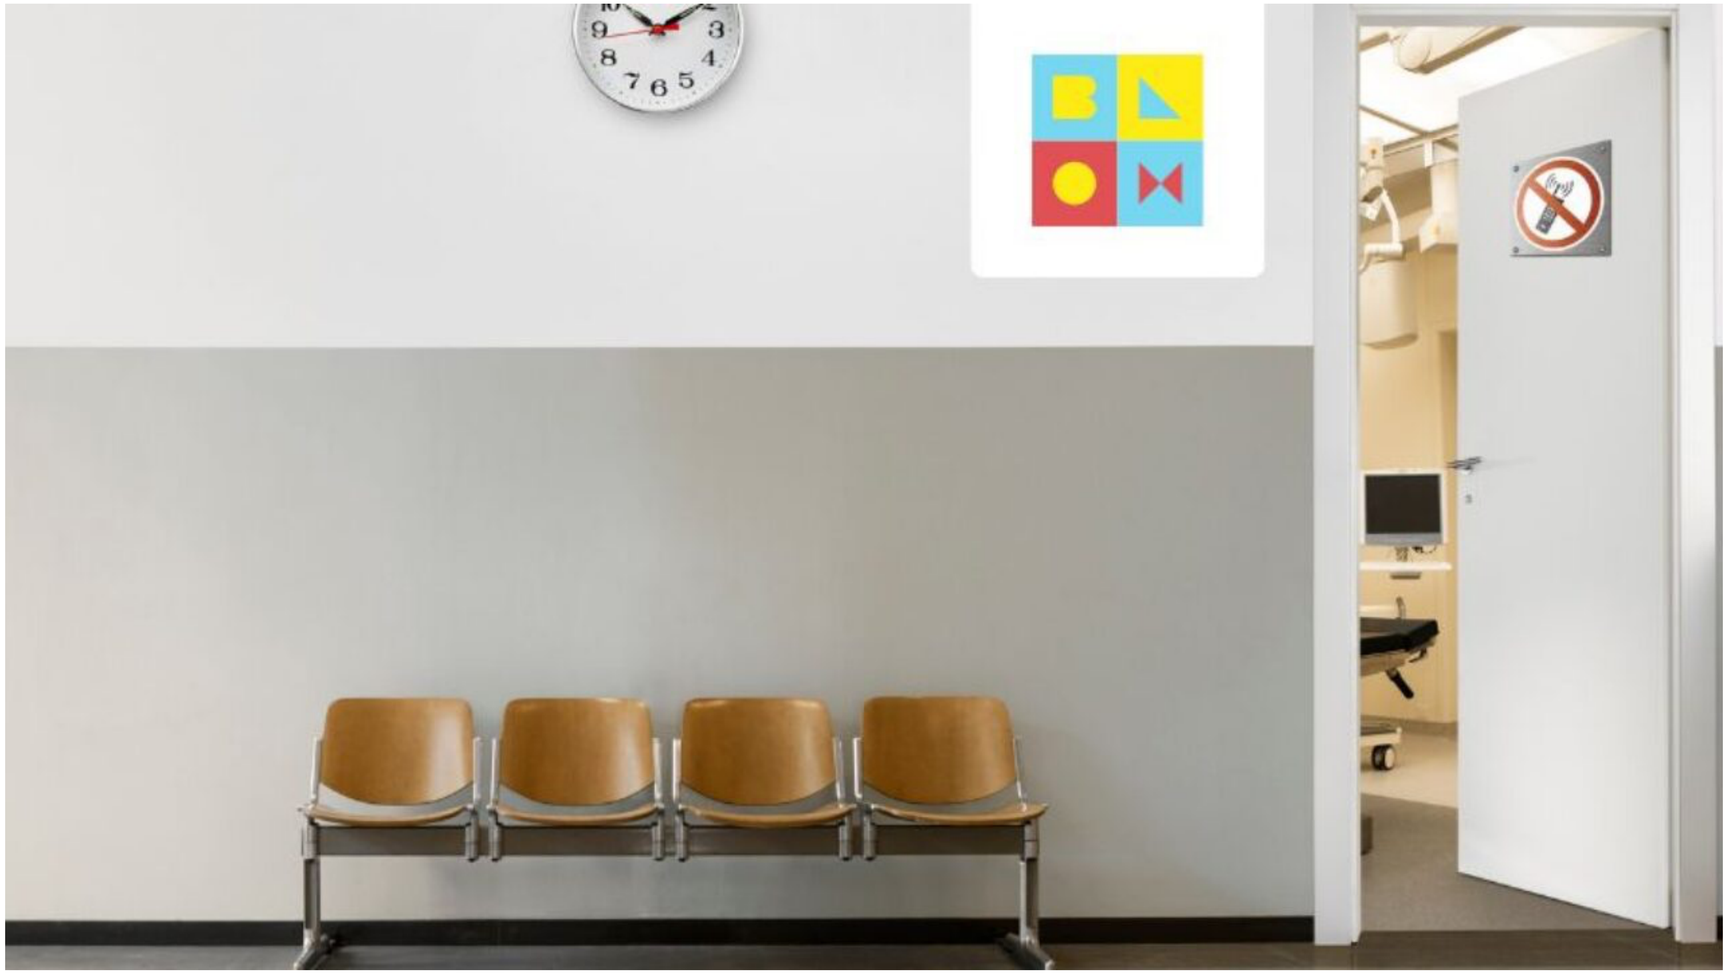

TUDelft

Transaction Broadcasting: Mempool

TUDelft

Other Networks

Complexity increase

Ethereum transaction requires more actions

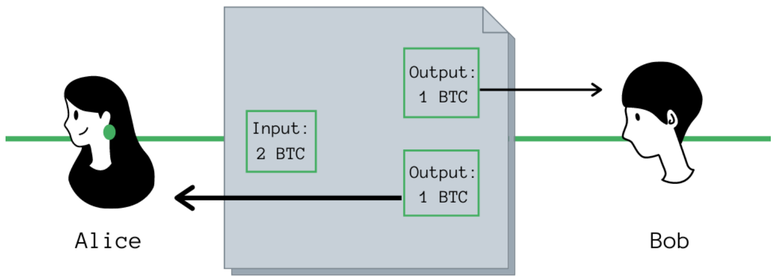

Bitcoin transaction

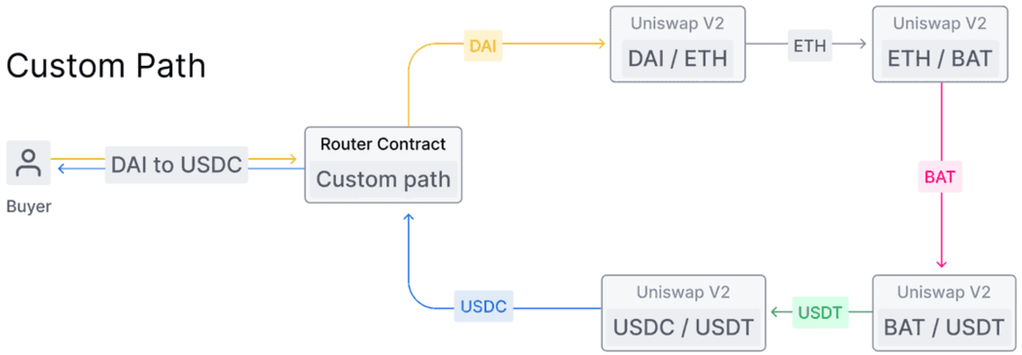

Ethereum transaction (Uniswap)

TUDelft

The Web3 Transaction Lifecycle

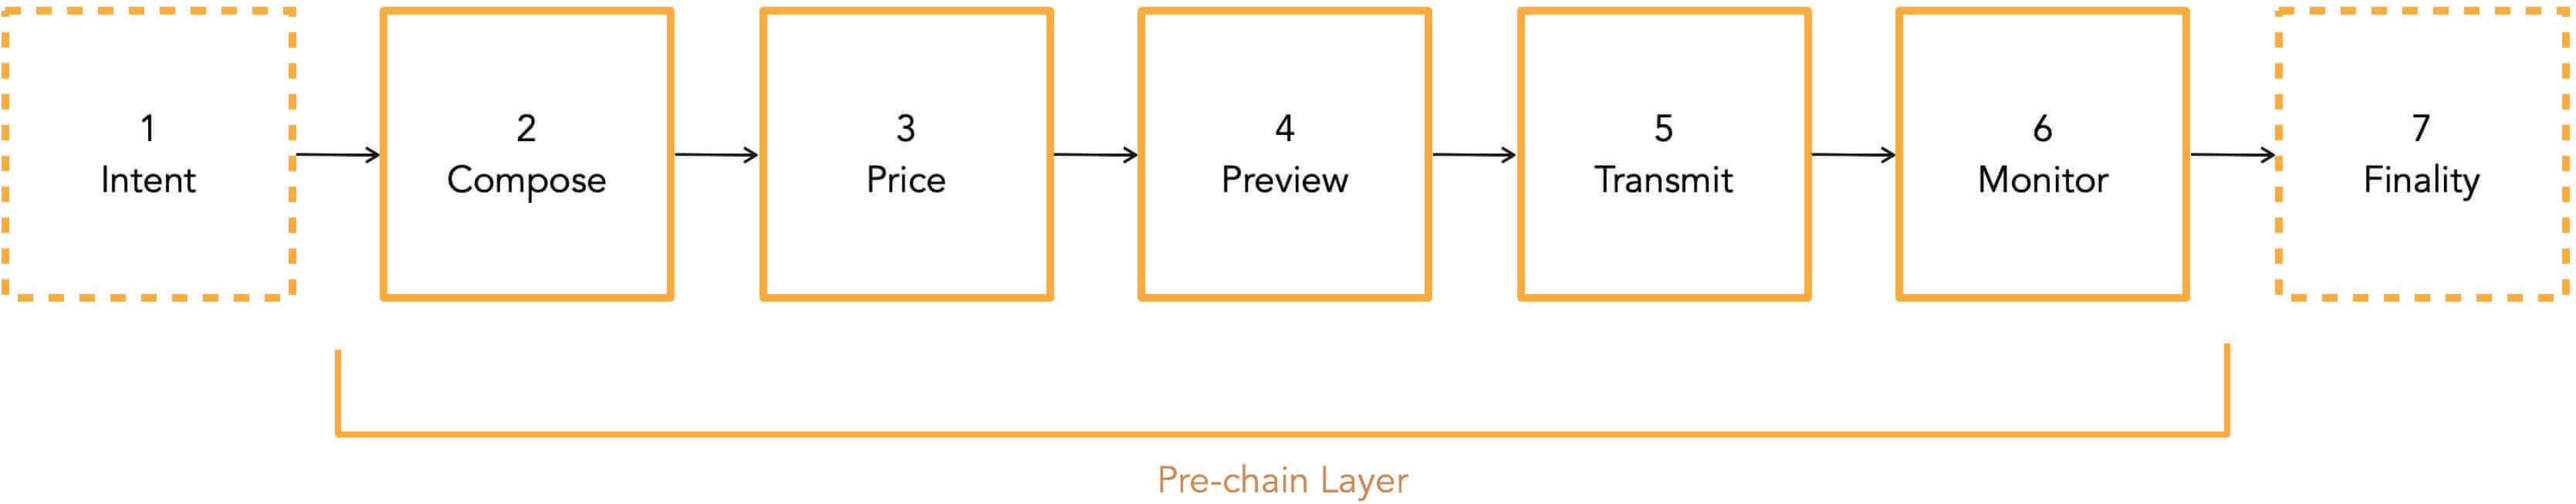

blocknative

TUDelft

The Web3 Transaction Lifecycle: Not Always Linear

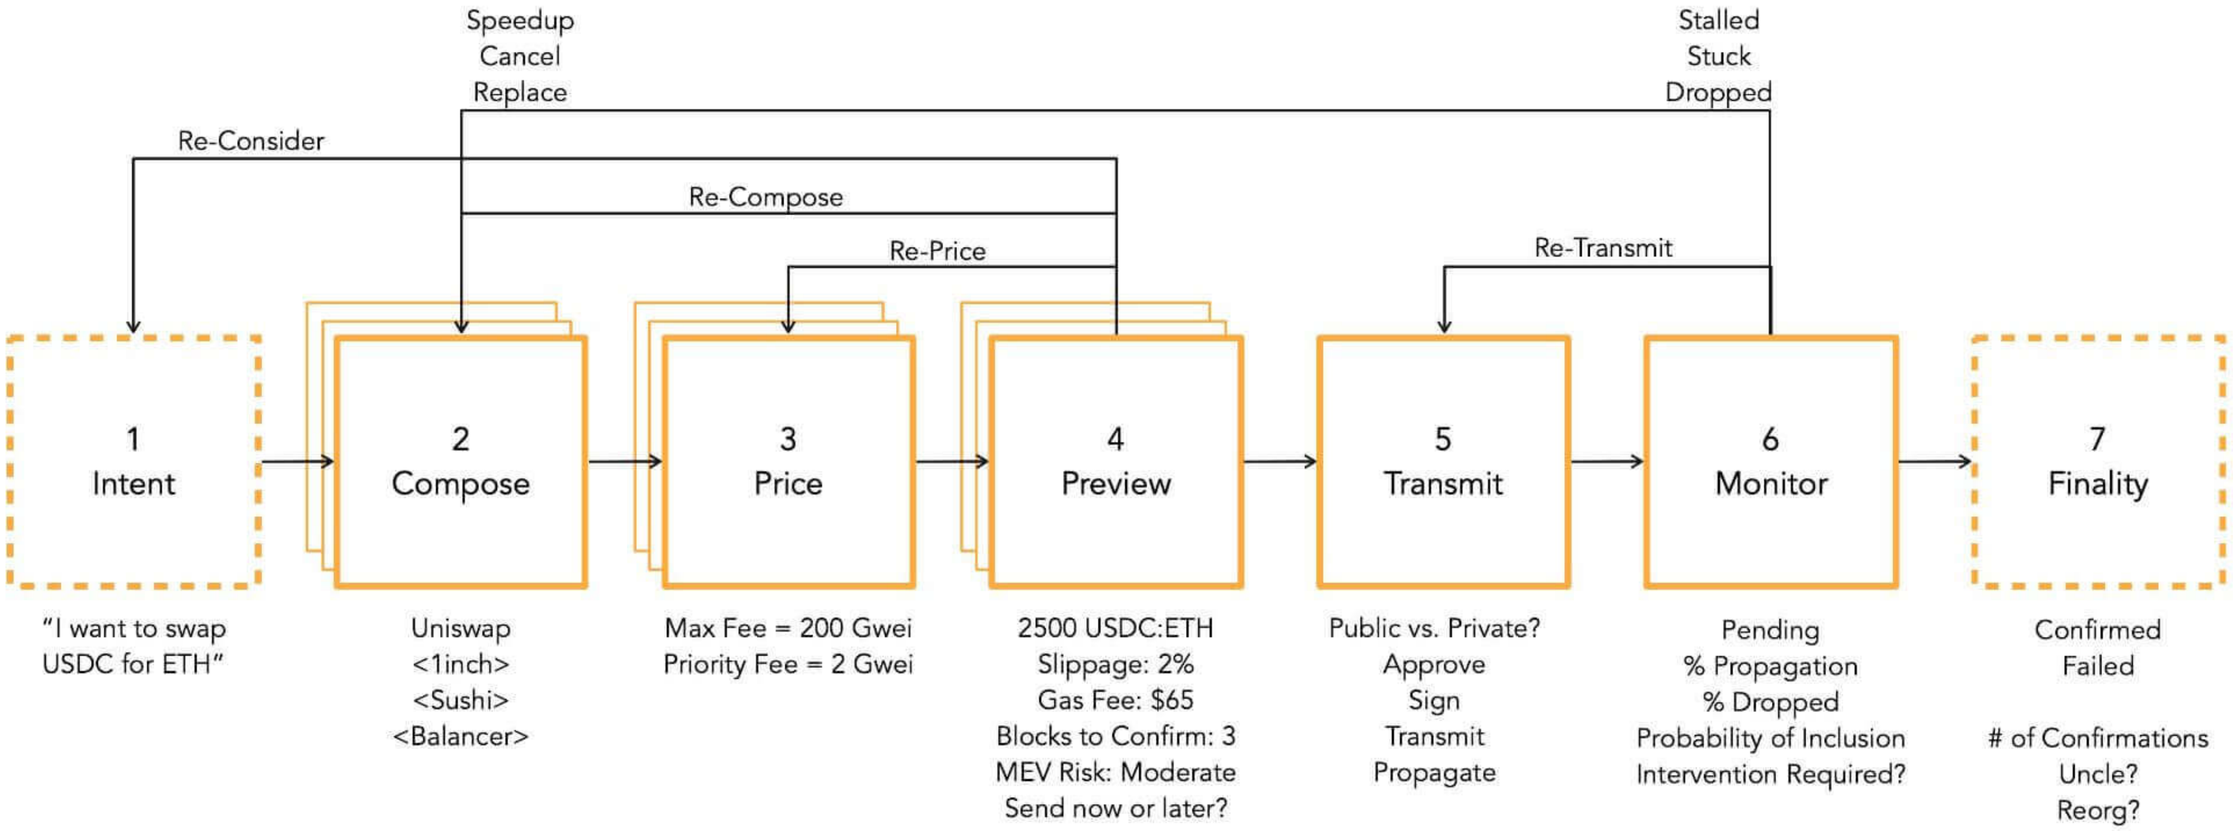

TUDelft

blocknative

Ethereum multiple networks

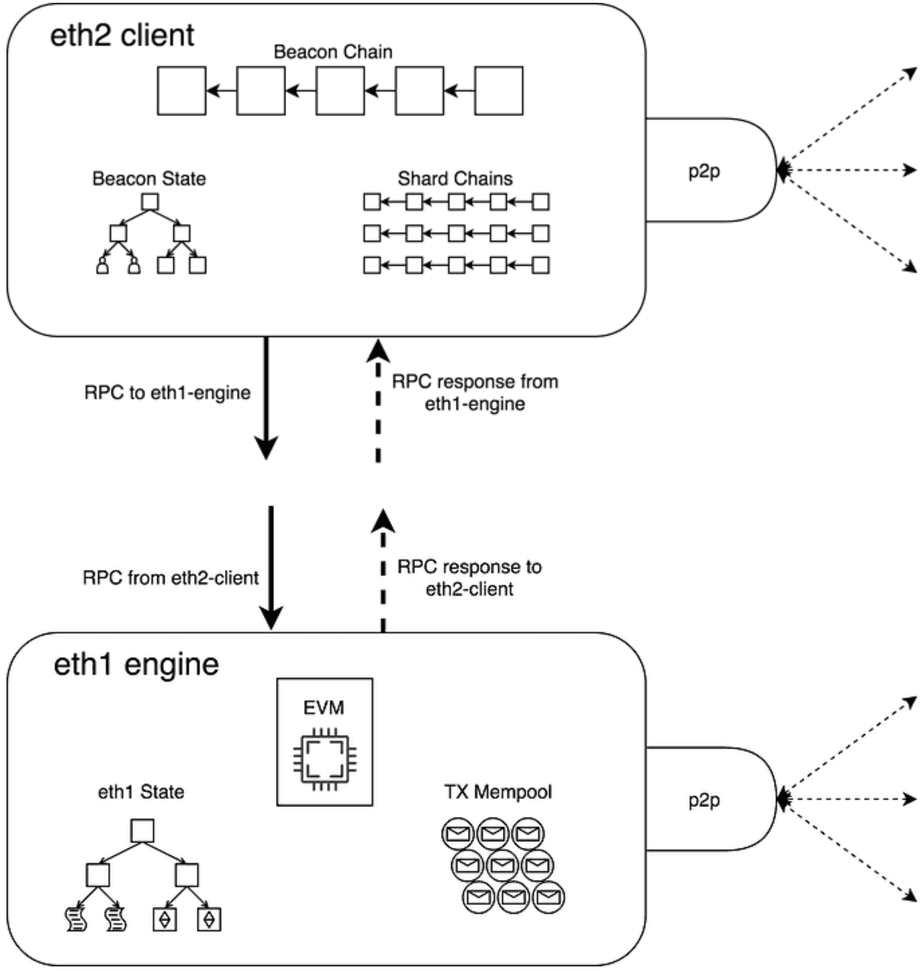

TUDelft

Multiple subnetworks

More complexity

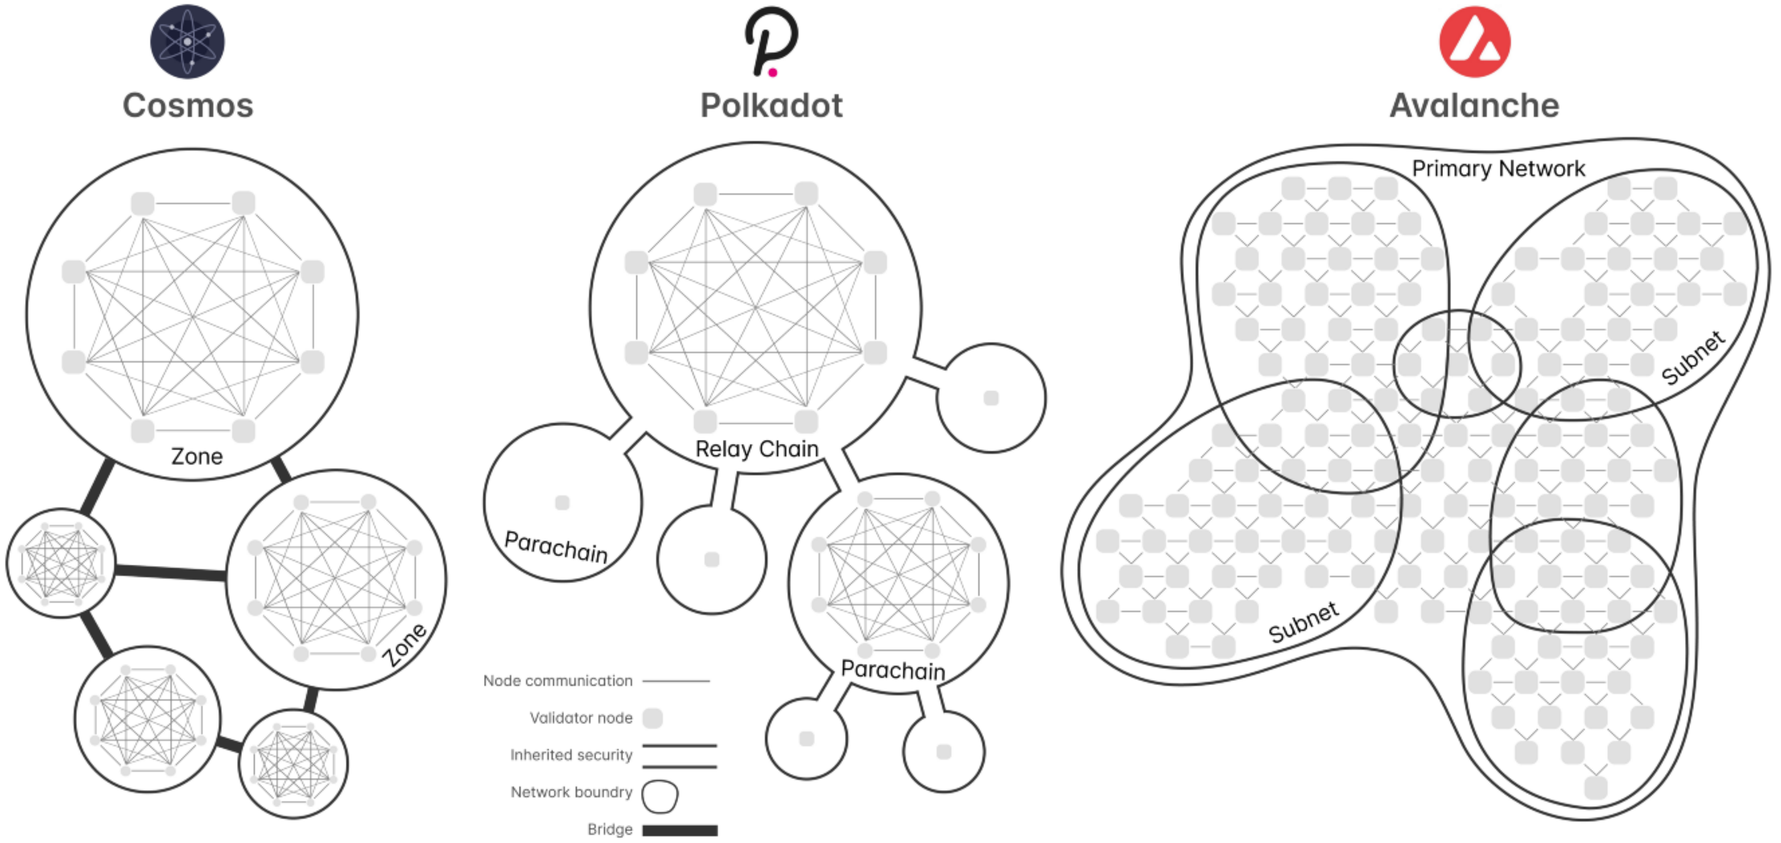

TUDelft

Robustness of Networks rely on Incentives

- Incentives play big role
- Why do I stay online?
- Why download and store you data?
- Covered in next lecture

IPv8: Networking Library

TU Delft

IPv8

Simple Networking Stack

Identity out of the box

Hooks for you Ledger

NAT traversal

Supports Semantic Clustering of Multiple Networks

TUDelft

Practice Time