Homework 2: Information Propagation and the Network Layer

Patrick McCorry

Kings College London, UK
patrick.mccorry@kcl.ac.uk

Abstract. We'll focus on the network layer for cryptocurrencies like Bitcoin which is responsible for spreading transactions and blocks to all peers on the network. It is recommended to make notes on this homework sheet for future use.

1 Who maintains the infrastructure?

Last week, we discussed how cryptography protects the blockchain's integrity (*cryptographic hash functions*) and it ensures only the rightful owner can spend their coins (*digital signatures*). This led us down the rabbit-hole to discover that fundamentally, the blockchain, is just just a cryptographic audit trail to let us recompute an entire database from scratch (and new blocks are simply responsible for performing batch updates). The next problem to investigate is how a user can discover the *one true blockchain* for a cryptocurrency and how to ensure new transactions are publicly available so miners can include them in their new blocks. This brings us to the network layer for a cryptocurrency.

There is no appointable central institution or organisation who can take on the role of propagating transactions (or blocks). Instead, cryptocurrencies rely on *peer-to-peer networks* which lets altrusic community members connect with each other on a global scale. This should form a *mesh network* where all peers and miners are connected with each other (which we typically call *Vanilla Bitcoin*). If everyone is connected to everyone else, then we can be assured that all new transactions are propagated to miners (so they are aware of transactions to include in their blocks) and all new blocks are propagted to all other peers (including other miners). Before propagating a transaction, each peer *should* validate new transactions/blocks according to a list of consensus rules (and the state of their local database). This global replication of computation (and verification) is the blockchain's cornerstone of security.

For this homework, we'll focus on the peer discovery (and bootstrapping) process, the announcement protocol (i.e. IP addresses, transactions, and blocks), the type of peers on the network, eclipse attacks and finally why Bitcoin is no longer like a vanilla ice cream.

2 Peer-to-peer network

Let's focus on questions that will help us better understand the peer-to-peer network protocols.

2.1 Bootstrapping new peers

To begin, every user must download the cryptocurrency software (i.e. bitcoind, geth, etc) and find an entry point into the peer-to-peer network. As we mentioned in class, bootstrapping new peers is a non-trivial problem and cryptocurrencies like Bitcoin still rely on a trusted bootstrapping process. In practice, it is maintained/distributed amongst seven influential software developers. While it isn't a great prospect that seven people have the power to eclipse new users on the network, it appears to be the best we can do at the moment. Let's dive a bit deeper:

- Why does the client-server infrastructure not work for cryptocurrencies like Bitcoin? [2 marks]
- At a high level, what is the process for a new peer to find other peers on the network? In your answer, please consider the bootstrapping process and the last-resort option if finding new peers fails [6 marks]
- Once a peer has connected to other peers on the network, what is the process of finding new peers? [4 marks]
- What is the distinction between outgoing and incoming peers? [2 marks]
- How does a node announce their presence to the rest of the network? [4 marks]
- What is the difference between a validating peer and a miner node? [4 marks]
- What is an eclipse attack? [4 marks]

2.2 Information Propagation

Users must *publish their transactions* via the peer-to-peer network to other peers (and miners). As we learnt in class, the peer-to-peer network is a gossip protocol that alone is sufficient for a cryptocurrency to *scale-out* and it also unintentionally helps network adversaries to learn more information than they should know about individual transactions. Again, let's dive a bit deeper to understand information propagation.

- Please explain the three-step protocol for exchanging new transactions/blocks. [3 marks]
- What is the purpose of a memory pool? [2 marks]
- Why does hosting a new peer harm the network's scalability in vanilla Bitcoin and what impact does it have upon the network's fork rate? Please explain why and you may provide a drawing to help illustrate the idea. [6 marks]
- What is a fast relay network (falcon, fibre, etc) and why does it improve the network's scalability? [4 marks]

– What is SPV mining? Why does it speed up the propagation of new blocks? And why is it dangerous? [6 marks]
– A peer may receive thousands of blocks on the network. How does a peer know they are following the *one true blockchain*? In your answer, please consider the mining puzzle and how blocks are chained together. [6 marks]

2.3 0-confirmation transactions

Mallory may sign two transactions:

– TransactionA sends 1 coin from Mallory to Alice,
– TransactionB sends 1 coin from Mallory to Bob.

Typically this is OK. But what if both transaction A and transaction B are spending the same unspent transaction output?[1]

– While both transactions are unconfirmed and propagating across the peer to peer network, which transaction should be considered invalid? [2 marks]
– What does it mean to say all users should wait 6 confirmations before considering a transaction final? [4 marks]
– The heuristic in Bitcoin is to wait 6 blocks for confirming a transaction, why do we need to wait for more blocks in Ethereum? [4 marks]
– An SPV peer does not validate every transaction in the blockchain and instead only relies on the chain of block headers with the most proof of work. How can we prove to an SPV client that a 0-confirmation transaction has recently been included in a block? [4 marks]

While the answers will be released in a week's time, if you found any of the questions difficult then please visit Patrick McCorry during his office hours.

[1] Or if mallory's balance is only 1 coin in the account-based model.