Homework 4: Problems scaling on-chain transactions

Patrick McCorry

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

Abstract. As we mentioned in class, cryptocurrencies like Bitcoin and Ethereum can only support a small magnitude of transactions (10 tx per second). We'll focus on the problems surrounding scaling cryptocurrencies and some proposals on how to alleviate the issue. It is recommended to make notes on this homework sheet for future use.

1 Who maintains the infrastructure?

We have covered the blockchain as a data structure, how transactions/blocks propagate across the network and how we can deploy and run censorship-resistant smart contracts.

Over the past few weeks, we have tried to highlight the beauty of *public verifiability*, thanks to cryptography, any peer can re-validate every transaction (and thus the network's current state). The transparency and accountability of the network is the cornerstone of why we can *trust* smart contracts. While publicly verifiable smart contracts have promising use-cases for replacing trusted third parties, so far they have failed to garner real-world use as blockchain-based cryptocurrencies do not scale. Bitcoin and Ethereum can only support a small magnitude of 10 transactions per second, and the network's transaction fee skyrockets from 20 cents to $20 (US Dollar) whenever the transaction throughput ceiling is reached. Sometimes it appears counter-intuitive why cryptocurrencies do not scale when there are 10k peers providing the network's infrastructure. As we covered in class, this is because the peer-to-peer network does not distribute work of validating transactions amongst themselves, but instead every peer will replicate the work of others. This global replication (and the desire for public verifiability) is at the heart of the scalability issues facing cryptocurrencies.

For this homework, we will deep-dive into some of the scalability issues that arise for cryptocurrencies. This will focus on information propagation for the peer-to-peer network, some game-theoretic attacks if miners do not honestly follow the protocol, and finally some of the leading research directions pursued by the community.

2 Myth of Decentralisation

Cryptocurrencies are often touted as being *decentralised*. Some of the beliefs include the lack of a single point of control (of the blockchain or smart contracts),

anyone can use the network, and it is global in nature and censorship-resistant. While these are admirable goals to pursue, it is not entirely true that the network is *fully decentralised*:

- How can big blocks (i.e. 1 GB every 10 minutes) harm the network's public verifiability? [2 marks]
- How can small blocks (i.e. 1 MB every 10 minutes) harm the diversity of users who can afford to use the network? [2 marks]
- What is a mining pool? What is the role of a mining pool operator? And why is it desirable for small miners? [6 marks]
- Name three attacks miners can perform as their hash rate increases from 1% to 51% of the network's total computational power? [6 marks]

2.1 Information Propagation

All transactions and blocks are propagated via the peer-to-peer network and this propagation serves two purposes. It ensures all information is available to let peers validate the blockchain in real-time and it ensures that miners receive blocks from other miners. As we have leart in class, the peer-to-peer network does not help it scale-out as the task of validation is replicated (and not distributed) amongst the peers. Let's dive a bit deeper:

- What is the blockchain fork rate? And what are the two factors that impact it? [4 marks]
- What is the honest mining strategy? [2 marks]
- Assuming a 1MB block takes 2.4 minutes to propagate across the network (and the block interval is set to 10 minutes), what is the maximum block size the network can handle and why? [2 marks]
- There is approximately 10k peers in Bitcoin's peer-to-peer network. Why does adding more peers to the peer-to-peer network harm its scalability? [2 marks]
- What is a compact block? Why does set reconcilation help improve information propagation? And what needs to be widely propagated for it to work? [6 marks]
- What is the relay network and why does it help improve information propagation? [4 marks]
- What is the meaning of *Vanilla Bitcoin*? [2 marks]
- At a high level, how does Bitcoin-NG better utilize the peer-to-peer network's bandwidth? And why is it better than relying on 0-confirmation transactions? [4 marks]

3 Dishonest miners

Satoshi Nakamoto proposed the 51% attack as a security threshold for Bitcoin and it was long believed that Bitcoin was secure as long as no miner achieved 51%

of the network's computational power. In 2013, Eyal and Sirer proposed *selfish mining* which demonstrated that *majority was not enough* to protect Bitcoin from dishonest miners. While selfish mining has not truly been witnessed in large cryptocurrencies like Bitcoin and Ethereum, it does provide a new threat that must be considered when trying to scale on-chain transactions. Let's dive a bit deeper:

- Why does selfish mining waste the time of other miners? And how does it work? [8 marks]
- Why does the selfish miners communication speed with the rest of the network impact the success of selfish mining? [2 marks]
- Explain why nakamoto consensus is potentially unstable if the block subsidy is removed? [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.