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]
- The issue with big blocks reduces the diversity of peers with the computational, bandwidth and storage resources who can validate every transaction. Only large data centres will have the ability to do it.
- 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]
- Every block has limited capacity (i.e. 7 tps) and users will have to compete via transaction fees for inclusion. This is why the fees in Bitcoin skyrocketed from 20 cents to $20 (US Dollars) in December 2017.
- What is a mining pool? What is the role of a mining pool operator? And why is it desirable for small miners? [6 marks]
- A mining pool lets a group of miners work together towards solving a block. The operator is responsible for co-ordinating tasks to all workers and ensuring all works are fairly paid. Mining pools are desirable as they let smaller miners get rewarded more fequently (although each reward is smaller).
- Name three attacks miners can perform as their hash rate increases from 1% to 51% of the network's total computational power? [6 marks]
- **Delay or Censor Transactions**. Miners can simply not accept transactions (and try to encourage other miners to do the same). If they have 51% of the networks hashrate, they can always censor transactions.
  Withhold blocks. Miners can withhold blocks they have mined from the network. This can be useful for selfish mining, or if the block hash is used as a random beacon (i.e. bias random number for gambling game.)
  Impact execution order of transactions. Miners decide the order of transactions and can always ensure their transaction is executed before others. This enables front-running attacks.
  Reverse transactions. When a miner has 51% of the networks hashrate, they can also remove transactions from the blockchain by mining a new fork that excludes it. Real-world example includes the attack on Ethereum Classic.

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]

- The fork rate is the ratio of stale blocks (i.e. did not make it into the longest chain) and the total number of blocks. It is typically caused when two independent miners create new blocks at approximately the same time, and thus the blocks have to compete for acceptance. Two factors that impact the fork rate include the size of blocks and the frequency of new blocks.
- What is the honest mining strategy? [2 marks]
- Miners will immediately publish all newly minted blocks to the other miners. This prevents other miners wasting their proof of work cycles on creating a competing block.
- 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]
- Around 4MB blocks. This will ensure peers always receive new blocks before the next one is minted. Thus they can validate the network in real-time.
- 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]
- Every new peer adds an additional hop for transactions/blocks. This slows down the network as it increases the network's work as a whole as opposed to distributing the work. (i.e. this peer will perform an extra validation).
- 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]
- A compact block is simply a block with a list of truncated transaction hashes (i.e. we only send the transaction hash and not the entire network). This improves information propagation as we can send "hints" to other peers on how to re-construct the entire block as opposed to re-sending the same transaction across the network. In terms of performance, it is claimed a compact block is around 9kb-20kb which is significantly smaller than a 1MB block. Of course, unconfirmed transactions must be widely propagated. Otherwise peers will need to ask their neighbours for the available unconfirmed transaction which would increase the communication/roundtrip overhead.
- What is the relay network and why does it help improve information propagation? [4 marks]
- The relay network is an exclusive and private network for miners (alongside economically important players like exchanges). It effectively lets miners connect directly with other miners which speeds up information propagation as the blocks no longer need to be sent across the peer-to-peer network.
- What is the meaning of *Vanilla Bitcoin*? [2 marks]
- Vanilla Bitcoin assumes that peers and miners rely upon the same peer-to-peer network. There is no private relay network for miners.
- 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]
- Bitcoin-NG seperates the idea of electing a leader (i.e. key block) and confirming transactions (i.e. micro block). All miners solve the proof of work to

produce a key block and elect themselves as leader. Once elected, the miners can continuously publish micro blocks (i..e during the 10 minute interval) until a new key block is created. Unlike 0-confirmation transactions, the micro block provides a good signal that the transaction will likely be accepted into the blockchain. We can consider it a 'pending' transaction waiting final approval by the next leader (which is better than a normal 'unconfirmed' tansaction).

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]
- A selfish miner will withhold their block to get a heard start on the next block. Thus it wastes the time for other miners as they are tryign to extend an old block. Any answer describing selfish mining should include the '1 block ahead' rule (i.e. keep mining private chat), the 'block tie if network catches up' (i.e. publish withheld block immediately) and the '2 block ahead rule' (i.e. selfish miner will always win the fork).
- Why does the selfish miners communication speed with the rest of the network impact the success of selfish mining? [2 marks]
- If we assume that honest miners always extend the first block they receive, then the selfish miner is guaranteed to win the block tie if their block is communicated/sent to all other miners before the competing (and honest) block.
- Explain why nakamoto consensus is potentially unstable if the block subsidy is removed? [4 marks]
- Without a block subsidy, miners will only get rewarded from transaction fees. The issue arises when there is a transaction with a large fee (lets say 25 BTC) that is ready to be accepted into the blockchain. If miner A includes the transaction with the large fee in their Block, then Miner B may decide to create a competing fee in order to steal the large tansaction fee.

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.