



Computer Systems

— Solution to Assignment 11 —

1 Eventual Consistency & Bitcoin

1.1 Delayed Bitcoin

- a) It is true that naturally occurring forks of length l decrease exponentially with l , however this covers naturally occurring blockchain forks only. As there is no information how much calculation power exists in total, it is always possible a large blockchain fork exists. This may be the result of a network partition or an attacker secretly running a large mining operation.

This is a general problem with all “open-membership” consensus systems, where the number of existing consensus nodes is unknown and new nodes may join at any time. As it is always possible a much larger unknown part of the network exists, it is impossible to have strong consistency.

In the Bitcoin world an attack where an attacker is secretly mining a second blockchain to later revert many blocks is called a 51% attack, because it was thought necessary to have a majority of the mining power to do so. However later research showed that by using other weaknesses in Bitcoin it is possible to do such attacks already with about a third of the mining power.

- b) The delay in this case prevents coins from completely vanishing in the case of a fork. Newly mined coins only exist in the fork containing the block that created them. In case of a blockchain fork the coins would disappear and transactions spending them would become invalid as well. It would therefore be possible to taint any number of transactions that are valid in one fork and not valid in another. Waiting for maturation ensures that it is very improbable that the coins will later disappear accidentally.

Note that this is however only a protection against someone accidentally sending you money that disappears with a discontinued fork. The same thing can still happen, if someone with evil intent double spends the same coins on the other side of the fork. You will not be able to replay a transaction of a discontinued fork on the new active chain if the old owner spent them in a different transaction in the meantime. To prevent theft by such an attacker you need to wait enough time to regard the chance of forks continuing to exist to be small enough. A common value used is about one hour after a transaction entered a block (~ 6 blocks).

1.2 Double Spending

- a) Figure ?? depicts the final situation. 7 nodes have seen T first and 5 nodes have seen T' first. The 5 nodes at the edge cut between the green and the red cut have seen both transactions.




Figure 1: Random Bitcoin network

- b) Each node has $1/12$ of all computational resources, hence the probability of T being confirmed is $7/12 \approx 58\%$, while T' has a $5/12 \approx 42\%$ chance of being confirmed. The higher connectivity from the first node seeing T resulted in the transaction spreading faster, increasing the probability of winning the double-spend.
- c) The first node that sees T' now has 20% of the computational resources. T' therefore has a probability to win of $2/10 + 1/11 \cdot 8/10 \cdot 4 \approx 49\%$. The distribution of computational resources in the network therefore matters. The goal of an attacker is to spread the transaction that she wants to have confirmed to a majority of the computational resources, which may not be the same as spreading it to a majority of nodes.

1.3 The Transaction Graph




Figure 2: Transaction Graph. The red outputs are UTXOs.

- a) See Figure ??.
- b) See red outputs in Figure ??.
- c) Fully spending an output simplifies the bookkeeping considerably as an output can only be in two possible states: spent or unspent. This means that it is easy to detect conflicts, because two transactions spending the same output are conflicts. If we were to partially spend outputs, allowing multiple transactions to spend the same output until the coins on that output were completely spent, then the conflicts become more complicated. Assume an output with value 1 bitcoin. When partially spending outputs we could create 3 transactions claiming 0.5 bitcoins from that output, two of them are valid and the third will be invalid,

but there are 3 possible combinations that are valid. So the simple answer is: it makes conflicts evident and reduces combinations for conflicts. The number of possible combinations increases rapidly with the number of transactions.

There are many points in the Bitcoin software where this would complicate things. A miner needs to construct a valid block from the known transactions, however this becomes more difficult if arbitrary combinations of transactions suddenly conflict with each other. Furthermore with complex conflicts attackers can create a situation where most nodes do not agree on the transactions which will be in the next block. However nodes are optimized to quickly be able to forward new blocks which look “expected”. If the nodes do not agree at least loosely on the transactions to be committed in the next block, the propagation delays become much larger as many transactions need to be retransmitted, which finally results in smaller total transaction processing capability.

Note: As you might have realized the flow of money can be nicely followed with the transaction data from the Bitcoin blockchain. If some hacker steals your valuable coins, you can watch him buy things with it and see where the vendors are spending this money too! For this reason it is often possible to “buy” larger amounts of Bitcoin with less Bitcoin. The larger amount you can “buy” has very likely been involved in some crime and is being tracked by the police, thus the owner is eager to exchange them for other coins. Look for “Bitcoin doubling” in your favorite list of darknet services!

1.4 Bitcoin Script

- a) Transactions are instantly finalized, so the large confirmation delay of the blockchain is irrelevant. Only the signatures of both parties are needed, then the money has effectively changed the owner. Furthermore no transaction fees have to be paid to miners for replacing a transaction.
- b) Without the opening transaction A could just spend the money with a transaction without a timelock to a different address owned by himself. Requiring both signatures prevents this and gives security to B. In this construction B can trust that the funds will be available after the first timelock runs out.

Note that if B wants to access the funds earlier, it is still possible for A and B together to sign a transaction which directly executes the latest state. As long as both agree it is thus not necessary to wait for the timelocks. The timelocks are only necessary to ensure the last state in case there is disagreement.

- c) A “kickoff” transaction can be introduced after the opening. Only the opening is executed (i.e., sent to the blockchain) at the beginning to secure the funds. Now transactions can be replaced and if someone wants to close the channel he can execute the kickoff. This starts the timers on the subsequent transactions. See Figure ?? for the new transactions.

In detail the protocol is the following.

Setup:

- (a) A creates all transactions of the setup (opening, kickoff and first state).
- (b) A sends these together with signatures for the kickoff and first state to B.
- (c) B signs the kickoff and first state and sends the signatures to A.
- (d) A signs the opening transaction and executes it on the blockchain.

Updating:

- (a) A creates a new transaction spending the kickoff output with a lower timelock.
- (b) A signs it and sends it to B.

(c) B sends his signature to A.

After the update both have a signed version of the new state and can terminate the channel in this state.

Closing:

(a) A or B proposes a settlement transaction that directly spends the locked-in funds in the opening output.

(b) The other party signs and sends it to the blockchain.

This is the cooperative closing case. If some dispute happens, either of the two parties can always send the kickoff and latest state to the blockchain.




Figure 3: A “payment channel”. A and B both have to sign to spend the output in the middle. The upper transaction can only be committed starting from blockheight 100, the lower one starting from blockheight 90.

2 Advanced Blockchain

2.1 Randomness from Previous Block

The validator who has been assigned the hash range starting from 000000... to, say, 000000FFFFFF... has an incentive to construct a block such that the hash of that block will land in their range of hash. In other words, they control the randomness of the process. In this way, they will iterate over many combinations of block content until they find the right block that will result in a specific hash that is favorable to them in the next round. QED.

2.2 DAG-Blockchain

- a) *A B C D E F G L I H J*
- b) No. Since none of the nodes *A–J* have any reference to *K* (otherwise it would not be a DAG), the total order for all of *K*'s parents remains unchanged by the creation of *K*.

2.3 Selfish mining

Here's our (somewhat subjective and non-conclusive) answer to this question:

Until the next difficulty adjustment, the block generation rate with respect to the longest chain is reduced in the presence of a selfish miner as the total hashing power would be distributed on two forks of the chain. Here, the selfish miner might even get unlucky and make less profit than by following the protocol.

However, due to the difficulty adjustment, the rewards distributed among the miners finding blocks will be reestablished to match a constant rate. From Theorem we know that it is rational (profitable) for a miner to mine selfishly if it has enough hashing power (e.g. $\alpha \geq \frac{1}{3}$ or even some smaller α with $\gamma > 0$).

Further thoughts:

- Assume that miners reinvest some of their income in buying new mining equipment. As the global rewards per hour are kept at a constant level due to the difficulty adjustment and a selfish miner would get disproportionately high rewards, the honest miners must be awarded disproportionately small rewards. This would mean that the selfish miner will have a greater increase in his hashing power in comparison to the honest miners. Consequently, his share of block rewards will increase even more.
- Assume that other miners are rational (miners trying to maximize their own profit). They would be better off joining the selfish miner. This poses the potential for a majority attack.

Remark: There are papers that compare selfish mining rewards vs. honest mining rewards over a longer term taking Bitcoin's difficulty adjustment into account. For example: <https://arxiv.org/abs/1912.01798> and <https://eprint.iacr.org/2018/1084.pdf>.

2.4 Smart Contracts

In our opinion, it's easiest to use the following tools for this example smart contract. YMMV.

- Use Metamask to get a Ropsten test network address.
- Use <https://faucet.ropsten.be> to get Ether into your Ropsten address.
- Use Ganache to instantiate a local blockchain.
- Use Truffle to build/test the smart contract on the local blockchain.
- Use Truffle/Infura to deploy the smart contract on the Ropsten network.

The following is the diff between the example smart contract and the one with additional minters:

```
7a8
>     mapping (address => bool) public additional_minters;
22c23
<         require(msg.sender == minter);
---
>         require(msg.sender == minter || additional_minters[msg.sender] == true);
24a26,30
>     function add_additional_minter(address additional_minter) public {
>         require(msg.sender == minter);
>         additional_minters[additional_minter] = true;
>     }
```

The above tools show that the actual building/deploying of smart contracts on Ethereum is quite straightforward. Smart contracts themselves can be complicated and need to be extensively reviewed/tested/formally verified. Many third party libraries provide standard functionality that is tested extensively beforehand. Most of the times, it's recommended to use these libraries. But sometimes, it can be quite catastrophic - see <https://www.parity.io/a-postmortem-on-the-parity-multi-sig-library-self-destruct/>.