Quantum Coin

Blockchain Data Availability

Authors : The Quantum Coin Community

Community Publication Date: January 2022

Revision: October 2024

Introduction

Data Availability is of paramount importance to the secure and resilient functioning of blockchains. Data Availability problems fall into two categories; One set of problems is the availability of current blocks, to enable clients to validate transactions. The other is the historical ledger data availability; Some may call this a problem of data retrievability instead of data availability, but we choose to use the same terminology for both.

Quantum Coin blockchain will be a combined multi-fork of Bitcoin, Ethereum, Dogecoin and DogeP. These blockchains which Quantum Coin will multi-fork are three of the largest. Because of this, the ledger size of Quantum Coin will be huge, to begin with, after the multi-fork. Hence data availability is critical to the Quantum Coin blockchain, to make it more secure and resilient. Some of the problems detailed in this whitepaper will also fall under the topic of Disaster Recovery.

In this whitepaper, we will first describe what these two problems are in detail, why they need to be solved and how we solve them.

Current Data Availability

About the problem

One of the important factors in the quick and inexpensive processing of transactions (such as point-of-sale transactions) is the ability to run light nodes (light clients). Light nodes do not have to download the entire blockchain, but only download block headers, using which they can verify the state of the blockchain. Without light nodes, the blockchain will become more centralized, because running full nodes would require a lot of processing power, storage, and bandwidth which very few can run.

The security assurances however are weaker for light nodes since they do not validate the entire state of the blockchain, but rather rely on the blockchain's consensus protocol. In addition, in a sharded architecture, nodes in other chains or validators of the main chain (such as the beacon chain) may not be able to keep track of all state data of the other shards, because of hardware limitations.

Attack Vectors

Malicious full nodes may relay invalid blocks or selectively withhold certain blocks. An adversarial attacker may selectively target light clients for a transaction of interest (such as a high-value transaction) and prevent access to honest full nodes but allow access to dishonest full nodes (for example, using side-channel attacks on the client, by selective denial of service or other means).

Malicious block producers may also withhold data selectively to full nodes, especially on a targeted attack. When blockchain usage in the payment industry grows, these types of attacks may become common.

These types of attacks can undermine the security of the blockchain. Therefore it is critical to improve the ability of light nodes to perform deeper checks on blockchain transactions. While light clients might not be able to perform as deep validation as full nodes, increasing this ability is a step up for the security of the blockchain in general.

How do other blockchains attempt to solve it?

Ethereum

Ethereum's proposal to solve this problem is using Erasure Codes [10] to generate data availability fraud proofs[1]. This works under the assumption that in a network with a large number of honest light nodes, the assurances of data availability increase.

Polkadot

Polkadot uses Erasure Codes [2] for validating data availability in its parachain architecture. The assumption here is that relay chain validators would sign blocks only if they received their assigned part of the erasure-coded block.

Near Protocol

Near Protocol [3] also follows a similar approach to Polkadot for data availability. In addition, Near Protocol also deals with lazy block producers who attempt to sign blocks without waiting for their assigned part of the erasure code, by bit-masking chunks of the block. Block producers will be slashed if they produce blocks with invalid bitmasks.

Reed Solomon Codes

Erasure Codes allows recovering a message, even if parts of the message are lost. Erasure codes work by transforming a message of M symbols into a longer one

with N symbols. These N symbols can then be transmitted or shared. Even if some of the N symbols are lost or not available, the original message M can be retrieved from just a subset of N.

Reed Solomon code is a popular error correction code that has been traditionally used in storage media such as DVDs and to improve performance using FEC (Forward Error Correction) in lossy networks. Quantum Coin will use Reed Solomon codes like the other blockchains as detailed in the previous section, to create erasure-coded versions of the block. Light nodes will use a similar scheme as detailed in Ethereum's fraud-proof model [1], to improve assurance of block validity. The exact parameters of the Reed Solomon code to be used will be determined closer to implementation.

Historical Ledger Data Availability

About the problem

Over time, the ledger of the blockchain gets larger, as each block is produced, and transactions are added to it. The ledger size becomes larger soon, especially for blockchains that support a large number of transactions per second (TPS). For example, the Solana blockchain can produce up to 4 petabytes of data every year if transactions are committed at full capacity [6]. Even Full Node providers might not be able to store all the historical ledger data at this scale.

One solution may be for Full Nodes to spread the data across volumes, but this can be afforded by fewer nodes, thus causing centralization. Blockchain sharding might appear to be a solution, but at a higher number of sustained TPS, the problem would remain even with sharding. In addition, sharding also comes with its trade

offs, such as increased complexity and reduced security model (because fewer validators per-shard, as opposed to all validators in a single chain; this is debatable though).

Over years, (or decades), depending on the ledger size, gradually node operators and validators might either prefer not to operate full nodes (because of economic feasibility) which is detrimental to the blockchain. Historical data is very important to validate the blockchain, hence the data just cannot be discarded.

Multi-Fork Challenge for Quantum Coin

For Quantum Coin, the historical ledger data problem is even more important to

solve, since it will be a multi-fork of three major blockchains; Bitcoin, Ethereum and Dogecoin, DogeP. The historical blocks of these three blockchains have to be additionally signed with the Quantum Coin blockchain validator's keys (like Falcon), as part of this multi-fork process. In addition, Quantum Coin will support a model in which validators can vote on node hardware requirements, bandwidth and block

gas limit. The higher these values are, the higher the TPS will be and the quicker the ledger size grows (provided there are enough transactions carried out on the network).

Solana Archivers

Solana proposes a solution involving achievers that uses a modified version of Proof of Replication (PoR)[7] to replicate data across archivers nodes. PoR is a replication system used in FileCoin [8]. Archiver nodes have an economic incentive by getting a percentage of the block rewards. Archiver nodes do not need to have heavy hardware requirements (since data is sharded across various archivers). In addition, Solana has also built an interoperability platform with ArWeave[6], to use Arweave's storage solution.

How will Quantum Coin solve it?

Quantum Coin will follow a multi-pronged approach to solve this historical ledger data availability problem.

Ledger Replication

Similar to Solana, Quantum Coin will also use a modified form of Proof-of Replication to store historical data. Archivers will have economic incentives to store this data. The actual implementation details, economic incentives will be detailed in a separate whitepaper. At a high level, the replication system will guard against common attacks on such replication schemes(as detailed in FileCoin's whitepaper), including:

- Sybil attacks: where an archiver node (bad actor) creates multiple identities that claim to store individual copies of the data to be replicated whereas it is one actor that stores only one copy.

- Outsourcing attack: where an archiver node claims to store replicated data but outsources the data storage to a 3rd party. If multiple such nodes outsource to the same provider, the replication system is broken, because

only a few copies exist than there should be.

Offchain

Despite the on-chain data storage, there could be unforeseen circumstances including software bugs, large scale natural disasters that could wipe out ledger data for many if not all nodes.

For example, a software bug either in the operating system or a side-channel security attack may delete ledger data in all nodes, causing a catastrophic impact to the Quantum Coin blockchain. For example, if a zero-day vulnerability is identified in Golang which is used to build and run Geth (Ethereum's client node software), then Ethereum will become an easy attack target. If such a vulnerability gives RCE privilege, an extremely bad actor (such as a highly funded state actor) may choose to delete all ledger data in all the public Ethereum nodes. Such high risk situations are not entirely rare, as can be seen from the recent Log4J vulnerability.

Under these circumstances, it may be impossible to rebuild the blockchain. Hence it is critical to store the ledger offline at periodic intervals. Some of the caveats with offline backup are that it might not be automated and proving ledger validity also becomes tricky. They also tend to deviate towards centralized approaches.

However, having a break-glass mechanism to support ledger backup offline is better than having none. If blockchains do get wide adoption for payment processing (like credit cards are now), not having this break-glass mechanism is a high risk, because such problems can cause largescale economic impact worldwide.

Quantum Coin will support snapshots of the ledger that will be signed with validators as of the snapshot block. This signed piece of data will be shared widely by the community in various media, and other decentralized storage platforms. The hash of the data and other metadata will also be provided as a reference. The major caveat here is that it requires community support, recognition and benevolent actors of the blockchain to make untampered data available in other media, but it is still an option to consider.

Like Solana, Quantum Coin will also look into integrating directly with other storage solution providers like ArWeave, in the future.

Time Capsule

Quantum Coin takes a long-term approach to solving the blockchain availability problem and preserving historical ledger data. Let's say 30 to 50 years down the line, the community handling the project might be different. A lot of contexts might have been lost. Add a few more decades and historical ledger data might get skewed or lost. The original off-chain backup storage providers might no longer exist.

In addition, even a quantum-resistant crypto-scheme might become vulnerable in the future due to newly found algorithms that break them. Long-range attacks is another possibility. In this case, the historical ledger can potentially be rebuilt and spoofed to represent the new one, because validator account keys should be considered compromised if the underlying crypto-scheme is vulnerable.

A break-glass approach would be required for client nodes to identify the right valid chain. This would involve a managed community-driven approach to switch to a different digital signature algorithm and to bring back the original blockchain ledger.

Quantum Coin proposes a scheme wherein Time Capsules [9] will be created with wide community observation and publicity; the ledger data will be signed by validators (at that time) and preserved in a physical time capsule, possibly more than one across the world.

This will also be recorded in AV (and possibly VR) and shared across various media. In addition, a unique identifier that contains a timestamp, a UUID, latitude, longitude of the place will be stored as an identifier of the Time Capsule. This identifier will then be signed and included back in the blockchain for cross reference. This process will be repeated periodically, though the frequency is dependent on the economic viability of this approach.

During a break-glass event such as vulnerability in the cryptography scheme, a community-driven approach can be taken to validate the ledger from the Time Capsule and cross-reference the hash in the on-chain blockchain, to verify the original blockchain ledger. While this method is not fool-proof, it is one step towards formulating a disaster recovery plan. This process is important when we speak of historical ledger data created over multiple decades.

In addition, Time Capsules also will help verify the integrity of blockchains many decades or centuries in the future and also serve as a historical reference. One of the goals post mainnet after the multi-fork is to create such a Time Capsule, to serve as an off-chain integrity record of the multi-fork of these three blockchains: Bitcoin, Ethereum and Dogecoin.

A follow-up whitepaper on disaster recovery will cover more scenarios and how these Time Capsules will help.

Summary

We covered two types of data availability problems and how the Quantum Coin blockchain will solve them. While there will continue to be edge cases wherein data availability can still become a problem, Quantum Coin will continue to evolve to protect from such problems. The Time Capsule is a new off-chain solution for blockchains in general. While the off-chain solutions are not fool-proof, they are a step in the direction of making blockchains robust and will set some standards to be used for payment processing and other use-cases.

Addendum

"Quantum Coin" and "Quantum Coin Community" were previously known under the monikers "Doge Protocol" and "Doge Protocol Community" respectively.

Appendix

1. Fraud and Data Availability Proofs: Maximizing Light Client Security and Scaling Blockchains with Dishonest Majorities https://arxiv.org/pdf/1809.09044.pdf

2. Polkadot Availability and Validity: https://research.web3.foundation/en/latest/polkadot/Availability_and_Validity.html

3. Near Protocol Nightshade: https://near.org/downloads/Nightshade.pdf

4. Erasure Codes: https://en.wikipedia.org/wiki/Erasure_code

5. Reed Solomon Codes:

https://www.cs.cmu.edu/~guyb/realworld/reedsolomon/reed_solomon_codes.html

6. Solana Archivers: https://solana.com/news/announcing-the-solana-arweave-interoperability-hack

7. Solana Archivers Replication: https://medium.com/solana-labs/replicators-solanas-solution-to-petabytes-of-blockchain-data-storage-ef79db053fa1

8. FileCoin: https://filecoin.io/proof-of-replication.pdf

9. Time Capsule: https://en.wikipedia.org/wiki/Time_capsule

10. Erasure Codes: https://en.wikipedia.org/wiki/Erasure_code