Bitcoin in research

Topic presentation for IN5420 Distributed Blockchain Technologies
Department of Informatics, University of Oslo
Presented by Michael Eikeland

References

A. Zohar, “**Bitcoin: Under the Hood**,” Commun. ACM, vol. 58, no. 9, pp. 104–113, Aug. 2015.

J. Bonneau et al., “**SoK: Research Perspectives and Challenges for Bitcoin and Cryptocurrencies,**” in 2015 IEEE Symposium on Security and Privacy, 2015, pp. 104–121.

K. Croman et al., “**On Scaling Decentralized Blockchains,**” in Financial Cryptography and Data Security, 2016, pp. 106–125.

Researching Bitcoin

- No formal structure or specification
    - White paper with initial design philosophy
    - *bitcoind*, de facto specification with ad hoc binary formats/protocols
- Scattered design discussions and choices
    - Forum postings, wiki articles, mailing lists, IRC logs
- Complex system
    - hard to model
    - relies on an unknown combination of factors
    - socio economics

Stability

“When is Bitcoin stable?” - a question of definition

Bonneau et al. (SoK) finds that a common definition of Bitcoin’s stability is when:
“the system will continue to behave in a way that facilitates a function currency as it grows and participants attempt novel attacks”

To model this Bitcoin has to be split in to components.

Consensus

- The Bitcoin network needs to be synchronized
- Specifically, to agree upon what transactions are valid and confirmed
- Known as *consensus* in distributed computing

Byzantine fault tolerance

- Trust needs to be implicit in the protocol due to the decentralized nature of Bitcoin
- The ability to detect and resist failure or false information in distributed computing is known as *Byzantine fault tolerance*
- An impossible task under certain conditions

Nakomoto consensus

- Bitcoin “circumvents” the BFT problem
- Group pending transactions into a new block
- The new block represents the consensus and is the base for the next block
- There could however exist conflicting blocks at the same time
- All participants chose the longest/most *difficult* chain of blocks
- Still, a longer chain may be discovered at any time
- Thus consensus is not deterministic but probabalistic

Stability of the consensus protocol

These properties of the consensus layer are required for Bitcoin to be considered *stable*:

- **Eventual consensus:** "At any time, all compliant nodes will agree upon a prefix of what will become the eventual valid blockchain. We cannot require that the longest chain at any moment is entirely a prefix of the eventual blockchain, as blocks may be discarded (become “stale”) due to temporary forks."
- **Exponential convergence:** "The probability of a fork of depth n is O(2^−n). This gives users high confidence that a simple “k confirmations” rule will ensure their transactions are permanently included with high confidence."
- **Liveness:** "New blocks will continue to be added and valid transactions with appropriate fees will be included in the blockchain within a reasonable amount of time."
- **Correctness:** "All blocks in the longest chain will only include valid transactions."
- **Fairness:** "On expectation, a miner with a proportion α of the total computational power will mine a proportion ~ α of blocks (assuming they choose valid blocks)."

Incentives in the consensus protocol

Incentive compatibility can be modeled with game theory. Introducing a notion of a **compliant miner**, one who follows the default mining rules - a **non-compliant miner** doesn't

Simply put: There is a Nash equilibrium if no participant has anything to gain from deviating from the *compliant* mining strategy

Nakomoto's argument: Bitcoin is stable as long as all miners follow their own economic incentives

Research seems to disagree. *Non-compliant mining strategies* may give better economic incentives.

Bitcoin-denominated incentives

- Simple majority compliance may not ensure fairness
- Majority compliance is an equilibrium with perfect information
- Majority compliance implies convergence, consensus and liveness
- With a majority miner, stability is not guaranteed
- If miners collude, stability is not known
- Stability is not known as mining rewards decline

Externally-denominated incentives

- Liquidity limits
- Exchange rates in the face of an attack
- Long term stake in bitcoin-denominated mining rewards

Alternative consensus protocols

SoK categorize and suggest alternative consensus protocols in three different groups

- **Computation puzzles:** Variants of the *proof-of-work* model
- **Virtual mining:** Variants of the *proof-of-stake* model
- **Designated authority:** Efficient consensus protocols at the cost of weakened decentralization

Computational puzzles

- **ASIC-resistant puzzles:** Puzzles that are hard to optimize for ASICs in order to restore the democratic principle of Bitcoin (“one CPU - one vote”)
- **Useful puzzles:** Reduce the wasteful nature of *proof-of-work* puzzles
- **Nonoutsourcable puzzles:** Prevent potentially hidden miner collusion in order to strengthen decentralization

Virtual mining

Also known as *proof-of-stake*. Principals may spend or demonstrate ownership of funds to generate new blocks.

Some implementation variations:

- proof-of-coin-age
- proof-of-deposit
- proof-of-burn
- proof-of-activity

Designated authority

- Decide on a set of authorities that receive, validate and sign transactions
- Not very specific on suggested implementations
- Would make consensus significantly easier to achieve
- Threatens decentralization

Stability of the peer-to-peer layer

Most models assume very optimistic terms of information propagation. Research shows that this might not be accurate.

The peer-to-peer layer may not be *incentive compatible*: Selfish miners may benefit from not propagating transactions.

Compared to traditional distributed systems, Bitcoin massively replicated data and workload. Might be unnecessary wasteful of resources.

Storage

All users are encouraged to keep a copy of the ledger and verify incoming transactions to strengthen the network. As of february 2018 the entire ledger has a size of *155 GB*.

Should Bitcoin reach volumes comparable to mainstream payment processors such as Visa at *2000 transactions/sec* the ledger would grow by *2.5 TB* a month.

Lacks proper incentivization. Might favor large mining entities who can maintain a single copy.

Scalability

Bitcoin today:

- *10 minutes* or longer to confirm transactions
- *7 transactions/sec* as maximum throughput

Croman et al. correctly estimated 2017 as the year the transaction throughput reaches capacity due to maximum block size. Ideally cryptocurrencies, such as Bitcoin, should be able to handle the same challenges and volumes as a mainstream payment processor. They state that in a broad sense their finding is that: *“fundamental protocol redesign is needed for blockchains to scale significantly while retaining their decentralization”*.

Key metrics in scalability

Scalability is not a single metric, but something that depends on several complex factors. Three key metrics are however identified:

- **Maximum throughput:** Constrained by the maximum block size and the block interval. Currently *3.3-7 transaction/sec.*
- **Latency:** The expected average time for a transaction to get confirmation. Currently *10 minutes.*
- **Bootstrap time:** The time needed for a node to download and process the ledger in order to be be ready to validate the system in the current state. Currently *4 days.*

Scaling by reparametrization

Bitcoin will to some degree be able to scale by reparametrization only. The limiting factors and they lower bounds are for *90% effective throughput* are:

- **Throughput limit:** Block size should not exceed *4 MB* at *10 min.* block interval.
- **Latency limit:** Block interval should not be smaller than *12s.*

These lower bounds do not take all properties into account and may affect the consensus protocol’s stability, such as the fairness property.

Cost per Confirmed Transaction

An interesting way to look at the scalability of Bitcoin is to look at the cost per confirmed transaction.

All costs tied to confirming a transaction in terms of *mining*, *transaction validation*, *bandwidth* and *storage*. were calculated based on transaction data from october 2015. At a throughput of *1.57 transactions/sec*, the cost per confirmed transaction totalled at *$6.2*. At maximum throughput it would be in the range of *$1.4 - $2.9*.

Network propagation

Some measurements were performed on the network's performance in 2012 and 2016.

- **2012:** Block propagation median and 90% time at *6.5s* and *26s* respectively. Average block size of *87 KB*
- **2016:** Block propagation 10%, median and 90% time at *0.8s*, *8.7* and *79s* respectively. Average block size of *540 KB*.
- **Estimates for 1 MB block size:** Block propagation 10%, median and 90% time at *1.5s*, *15.7s* and *2.4 min*.

Bottleneck

Network propagation rate for *90% effective throughput* was at *55 Kbps*. Significantly lower than what bandwidth analysis of individual nodes showed.

No single factor stand out as the reason for this.

Scaling beyond reparametrization

In order to scale beyond the scope of what parameter changes can achieve, radical changes were suggested. The paper, for this purpose, models Bitcoin in a hierarchy of *planes*: Network, Consensus, Storage, View and Side planes.

Network plane

As was implied, the underlying network's bandwidth isn't fully utilized. Two inefficiencies are pointed out:

- The local validation of all received transactions "contributes significantly" to the overall propagation times
- All transactions are effectively propagated twice: once on announcement and once more after inclusion in a block

Proposals for the network plane

- Avoid propagating each transaction twice by implementing a reconciliation protocol
- Use a dedicated, centralized, high-speed relay network for inter-miner communications
- Improve the network layer's function as a broadcast channel which have known low-latency protocols

Also this paper points out that the network layer lacks proper incentives. It relies on voluntary participation and requires ad hoc defenses to attacks.

Proposals for the consensus plane

- **Improving Proof-of-Work:**
  - GHOST has different chain selection rules that improves fairness and improves mining power utilization.
  - Bitcoin-NG eliminates the tradeoffs of Bitcoin with an alternative blockchain protocol
- **Proof-of-Stake:** Principals deposit funds in order to be able to create blocks. Lacks formal guarantees for convergence and potentially threatens decentralization.
- **Consortium consensus:** A model that relies on stronger assumptions for trust (such as external factors). In other words, increased centralization. However, it carries significant performance increases. Might not be viable for cryptocurrency.
- **Sharding:** Split the consensus task among participants. May however incur substantial overhead in a Byzantine setting.
- **Sidechains:** Introduce lower-tier conesensus instances that have a lower degree of decentralization. Has challenges such as the need for independent security, miner coordination, inter-chain transactions, potential high latency.

Proposals for the storage and view planes

A few, nonspecific suggestions are made, but the essence seems to be that it might not be necessary to hold on to the entire history of transactions.

A suggestion is to discard transactions after they have been spent, only keeping track of *unspent transaction outputs* (UTXO). Keeping a track of the state instead of building it from a history transactions. Could open for sharding in the distributed network.

Side plane

Introduce off-chain functionalities such as overlay payment networks (i.e. Lightning Network).
Introduces a whole range of new factors that haven’t been properly analyzed.

Untouched topics

- Stability of transaction validity rules
- Stability with incentives other than mining income
  - Goldfinger attack
  - Feather-forking
- Stability of mining pools
- Client-side security
  - Simplified Payment Verification Security
  - Key management
- Modifying Bitcoin
  - Altcoins
- Anonymity & Privacy
- Extending Bitcoin's functionality

Discussion

- Discuss the the topics and their relevance
- Discuss the stability model
- Discuss the scalability model
- Discuss incentive compatibility
- Were there factors that were unjustifiable left out?
- What would you have liked to see more of?
- Is the probabilistic nature of consensus in Bitcoin problematic?
- Discuss the alternative consensus protocols - any favorites?

Discussion

- To what extent is Bitcoin still aligned with its design philosophy?
  - decentralized
  - democratic
  - privacy
- Discuss whether or not these are a threat to Bitcoin:
  - lack of strong identities
  - mining pools
- What do you identify as the biggest problem Bitcoin faces?