Abstract

This work analyzes the consensus problem in the context of the Bitcoin peer-to-peer network.

1 The Consensus Problem

The consensus problem is a fundamental problem in distributed computing with many real-world applications in areas such as database systems or autonomous multi-agent systems (e.g., robots). The basic question is how a set of independent agents or entities can reach consensus over a data value. For example, large distributed database systems with multiple replicas must be able to reach consensus over values stored in the database, assuming that any node can fail at any time. To solve this problem, the Paxos algorithm [5] is utilized in many large-scale database systems such as Google Spanner [1]. As another example, planetary exploration robots or sensor networks, where every entity can only communicate with its immediate neighbors, may want to reach consensus about physical properties such as temperature. In the simplest case, consensus can be reached by iteratively receiving temperature values from neighbors, computing weighted average values, and propagating those average values to all neighbors.

A variety of definitions of consensus can be found in literature. Typically, we assume that every entity can propose (multiple) potential values, but must make a decision at some point. Entities can fail, in which case they may "misbehave", send bogus values to other entities, or even masquerade as "correct" entities (*Byzantine failure*). According to a popular definition, a consensus protocol should exhibit the following properties [2].

- **Termination:** Every entity is making a decision eventually.
- **Validity:** If all entities propose the same value, then all correct entities must decide on that value.
- **Integrity:** Every entity may only make one final decision. The decided value must have been proposed by some entity.
- **Agreement:** Every correct entity is deciding on the same value.

In this work, we describe the consensus protocol of the Bitcoin cryptocurrency. Bitcoin is a decentralized currency, i.e., there is no central authority or bank, and all participants must agree on what amount of money belongs to whom. Since Bitcoin is a currency, it is a lucrative target for attackers, who may want to conduct fraudulent transactions for their own benefit. The consensus protocol in Bitcoin was designed in such a way that it can handle Byzantine failures as long as the majority of the computing power belongs to correct (honest) entities (participants).

2 Bitcoin

Bitcoin is a decentralized, anonymous cryptocurrency created in 2009 by a person or organization under the pseudonym Satoshi Nakamoto [6]. It is accepted as a payment method by a small number of shops and websites, but currently used mostly for investment and speculation [4]. Bitcoin is considered a fiat currency; its value is not backed by a physical commodity but based on supply and demand. In December 2017, Bitcoin reached a value of 17,900 USD per Bitcoin.

From a technical perspective, Bitcoin is a distributed (peer-to-peer) ledger, which contains the transactions of all Bitcoin users. Users can make transactions with their private key and other users can validate those transaction with the corresponding public keys. One main problem that Bitcoin solves is the *double spending problem*: How can the system ensure that the same money is not spent twice? In paper-based currencies, this problem is solved by adding security features to banknotes or coins; those features prevent criminals from duplicating banknotes. However, data (or messages) can easily be copied. This a typical consensus problem: All entities (Bitcoin network participants) must agree on what amount of Bitcoins is owned by each entity. Some entities may misbehave in order to gain a benefit for themselves, e.g., by submitting fraudulent transactions. Such transactions should not be accepted by correct entities. In the following paragraphs, we will first briefly describe the basics of Bitcoin and then focus on how consensus is achieved in Bitcoin.

2.1 Transactions

Transactions are used to send money to other Bitcoin users. Every transaction has a set of inputs and a set outputs (Figure 1). Every input must reference an output from a previous transaction and every output has a destination (receiver) and amount of money in Satoshi (1 Bitcoin = 100,000,000 Satoshi) and must only be referenced once by another input. The amount of money in the output must be equal or smaller than the amount of money in the input. The difference between input and output is a transaction fee that is paid to miners (explained later). In the simplest case, users can only use those transactions as inputs which specify themselves as destination; this is enforced using public-key cryptography. Bitcoin transactions have two interesting properties: First, the user creating a transaction can choose the transaction fee arbitrarily. Second, since every output can be referenced only once, if users received a certain amount of Bitcoins during a previous transaction at some point, but want to transfer only a fraction of that amount to someone else, they must specify two outputs; one for the actual receiver and one for themselves.

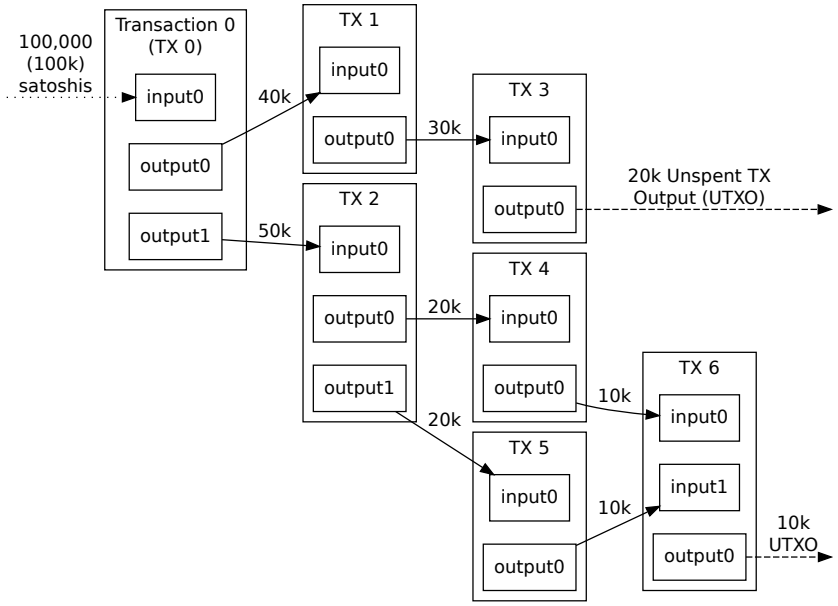

Figure 1: Example: Network of Transactions. Every transaction can have an arbitrary number of inputs and outputs. The amount of money entering a transaction must be greater or equal to the amount of money leaving a transaction. Source: `https://bitcoin.org`

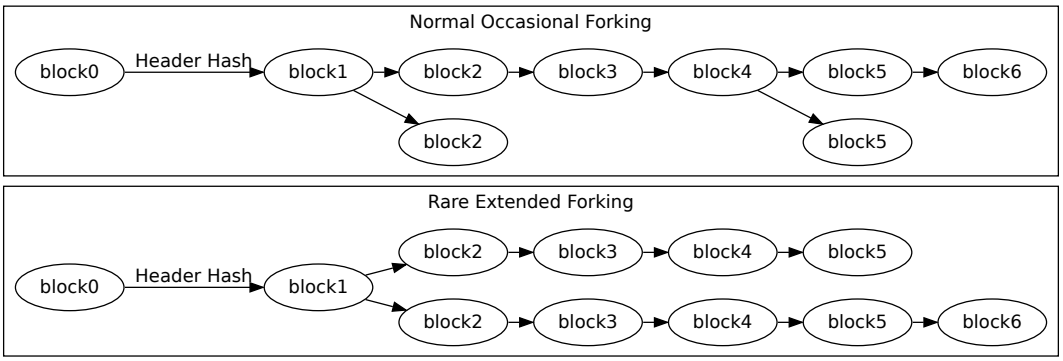

Figure 2: Example: Block Chain Forks. Even though any block may have multiple children, only the longest chain produced with the most computing power according to proof-of-work is the correct one. Bitcoin is a peer-to-peer network and nodes may not always be able to see the entire network, but in the absence of a considerable number of fraudulent peers, blockchain forks are unlikely to grow long. Source: `https://bitcoin.org`

2.2 Blockchain

Transactions created by Bitcoin users are propagated in the Bitcoin peer-to-peer network. They are picked up by *miners*, who combine multiple transactions into a *block*. Only then, a transaction takes effect. Every block references the previous block, establishing a temporal order among all blocks (and thus transactions). All blocks together form the ledger. Creating a block is computationally expensive: It does not only involve combining a few transactions but also requires solving a cryptographic "puzzle" (*proof of work* [3]): Miners must add a nonce (4-byte value) to a block, such that the hash value of the block starts with n many zeros. This number n is automatically adjusted every two weeks based on how many blocks were created. It is adjusted such that the average block creating time is 10 minutes. Reducing n makes block creation easier and faster. Increasing n makes block creation more difficult and slower. Miners may add a special transaction to every block that assigns them the transaction fee plus a certain amount of money that is created with each transaction[1].

Since transactions are sent out to the entire Bitcoin network, multiple miners may include the same transaction. More importantly, multiple miners may attempt to continue the blockchain with a different block, which would result in a fork of the ledger. This is forbidden, because there may only be one *correct* version of the ledger. Only one miner can succeed and the Bitcoin network must decide which one will. The Bitcoin protocol defines a mechanism for reaching such a consensus.

2.3 Consensus

Bitcoin defines a simple rule for determining which block out of all existing ones is the one representing the current state of the ledger. First, invalid blocks that cannot be validated (e.g., because they were forged by an attacker) are discarded. Note that the entire history of transactions and blocks can be validated using the participants' public keys. Second, among all

[1] This amount of money decreases gradually, limiting the amount of Bitcoins that can be created.

valid blocks, the one that took the most computing time to produce (taking into account the chain of all previous blocks) wins. Although a bit misleading, this is called the *longest chain* (Figure 2).

Consensus in Bitcoin is different from other consensus mechanisms in one important aspect: Instead of giving every participant a single vote, it is based on computing power (longest chain). Therefore, having a large number of fraudulent/zombie Bitcoin peers (or IP addresses) alone is not enough to compromise the ledger. For example, attackers may try to steal back Bitcoins from previous purchases by trying to establish a new block as current state of the ledger which does not contain their earlier transactions used for payment. Such attackers must generate a chain that is longer than the honest/correct one. During that time, the other honest miners will increase the length of the honest chain by adding new blocks. Consequently, attackers can take over the Bitcoin network only if they possess the majority of the computing power. However, they could also use their computing power to generate new Bitcoins in an honest way. In his paper, Nakamoto describes the process and probability of an attacker being able to generate an alternate chain faster than the honest chain in more detail.

References

[1] James C. Corbett, Jeffrey Dean, Michael Epstein, Andrew Fikes, Christopher Frost, J. J. Furman, Sanjay Ghemawat, Andrey Gubarev, Christopher Heiser, Peter Hochschild, Wilson Hsieh, Sebastian Kanthak, Eugene Kogan, Hongyi Li, Alexander Lloyd, Sergey Melnik, David Mwaura, David Nagle, Sean Quinlan, Rajesh Rao, Lindsay Rolig, Yasushi Saito, Michal Szymaniak, Christopher Taylor, Ruth Wang, and Dale Woodford. Spanner: Google's globally distributed database. *ACM Trans. Comput. Syst.*, 31(3):8:1–8:22, August 2013.

[2] George Coulouris, Jean Dollimore, Tim Kindberg, and Gordon Blair. *Distributed Systems: Concepts and Design.* Addison-Wesley Publishing Company, USA, 5th edition, 2011.

[3] Cynthia Dwork and Moni Naor. Pricing via processing or combatting junk mail. In Ernest F. Brickell, editor, *Advances in Cryptology — CRYPTO' 92*, pages 139–147, Berlin, Heidelberg, 1993. Springer Berlin Heidelberg.

[4] Alexander DAlfonso, Peter Langer, and Zintis Vandelis. The future of cryptocurrency: An investor's comparison of bitcoin and ethereum, 2016.

[5] Leslie Lamport. The part-time parliament. *ACM Trans. Comput. Syst.*, 16(2):133–169, May 1998.

[6] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system, http://bitcoin.org/bitcoin.pdf.