Consensus Mechanism Overview

Jianing Wu, WisdomTree

A decentralized system implies that no single participant has control over the system's rules, inputs, and outputs. Security therefore becomes the biggest challenge to any decentralized system. This is especially true when participants don't trust each other, and the system provides a record of transactions that ascribe value (like on a public blockchain).

Without third-party verification, how can participants validate transactions and prevent malicious actors from imposing fake and fraudulent information?

Satoshi Nakamoto provided a solution to this question by combining[1] various ideas to create a distributed, immutable, and cryptographic ledger of transactions. At its core is the proof-of-work consensus mechanism – a way to verify transactions by proving to others that considerable computing efforts were spent for the information to be appended to the ledger.

Since Bitcoin's birth there have been many other consensus mechanisms created. Each of them has its own characteristics that determine the associated network's attributes. In this blog post, we review several existing consensus mechanisms and provide an overview on how they differ.

What's a Consensus Mechanism?

A consensus mechanism is an algorithm to approve transactions or records onto a decentralized ledger such that fake or fraudulent records are rejected.

The algorithm is run when new blocks are being appended to the existing chain of blocks, which is how the blockchain gets updated as an append-only ledger.

The idea is that by imposing a requirement of certain effort spent (or risk taken), malicious actors would refrain from tempering with the ledger as they deem the effort (or loss) to be unprofitable. The very first purpose of proof of work's invention was to filter email spam.

Hashcash, a proof-of-work system proposed by Adam Back in 1997, requires email senders to create and attach stamps on email headers to prove to receivers that they spent CPU power to generate emails. These stamps are one-way encryption algorithms that are easy to verify by the receiver but hard (in computing terms) to generate by the sender. In this model, spammers would be reluctant to send out large quantities of email as it becomes unprofitable to use a large amount of CPU power to create stamps. However, the price of sending a single email is still affordable by regular users.

Since consensus mechanisms in the blockchain world are generally referred to as activities of "mining" and "staking," they are frequently regarded as methods to issue new coins. However, their primary purpose is to secure the decentralized network, whereas rewards in the form of coins are an added economic incentive for workers to maintain the network.

[1] Narayanan, Arvind and Clark, Jeremy. Bitcoin's Academic Pedigree. ACM Queue Vol 15, No. 4. August 29, 2017.

Comparison of Major Consensus Mechanisms					
	Invented Year	**Nouns**	**Pros**	**Cons**	**Blockchain**
Proof-of-Work (PoW)	1993	mining I miners	• secure • simple • relatively long history	• energy intensive • susceptible to be centralized	Bitcoin, Litecoin, Ethereum (up to Serenity)
Proof-of-Stake (PoS)	2012	minting I validators	• energy efficient • less centralized • better designed for attack recovery	• shorter track record • nothing-at-stake problem • long-range attacks	Ethereum (planned to be implemented in Serenity), Cardano
Delegated Proof-of-Stake (DPoS)	2014	minting I witnesses, delegates	• same as PoS • but more democratic • faster	• susceptible to be centralized	Bitshares, Steemit, Ark, Lisk

Proof-of-Work (PoW)

Proof-of-work is the oldest and the most popular consensus mechanism. It accounts for more than 75% of the market cap of blockchain protocols.[2] It is used by Bitcoin, Ethereum (up to Serenity), and Litecoin etc.

In PoW, **miners** race to generate a block of data containing three things: new transactions waiting for verification, a record of the previous block, and a new transaction (called the 'coinbase') paying themselves a reward (and in the process increasing the supply of currency).

This block must satisfy a certain mathematical requirement when the block of data is '**hashed'**. Miners are free to choose any transactions they want to verify from a pool of unverified transactions (i.e., not in a block already on the blockchain) maintained by the network. Before adding a transaction to their block, the miner checks the sending party's digital signature on each transaction, and that the party sending coins has previously received enough coins in a prior transaction recorded in a block on the blockchain.

A **hash function** takes an arbitrary amount of input data and computes a fixed length numerical output that has several important properties: for the same input data, anyone using the same hash function will get the same hash output; the output is unpredictable and cannot be guessed; and there is no way to get from the hash output back to the inputs. Blockchains may use different hash functions (e.g., bitcoin uses the SHA256 function), but whichever they choose typically has these properties.

For a block to be valid the hash output must be a 256-bit long number that starts with a certain number of leading zeros (known as the **difficulty**). In binary notation this number has 256 ones or zeros. A number this big can uniquely describe all of the atoms in the universe. Therefore, a valid hash output is usually a very small number relative to all possible numbers with 256-bits.

[2] As of 7/23/2021. Calculated using CoinMarketCap data.

Since the hash function generates this output in an unpredictable manner, the miner cannot know what the result will be, and chances are that the hash output won't start with enough zeros. The miner repeatedly changes and hashes the block by incrementing a number known as the **nonce** and adding the nonce to the block's data – thereby guaranteeing a different hash output each time.

When the miner finds a nonce and set of transactions that, when hashed, satisfy the difficulty requirement it broadcasts that block to other miners so they can verify and accept that block. Once other miners verify a block's validity, they add that block to their copy of the blockchain and start mining a new block including a hash of the previous block in their new block. There is no incentive to continue verifying transactions contained in a new block as other miners will only add to the longest chain of blocks.

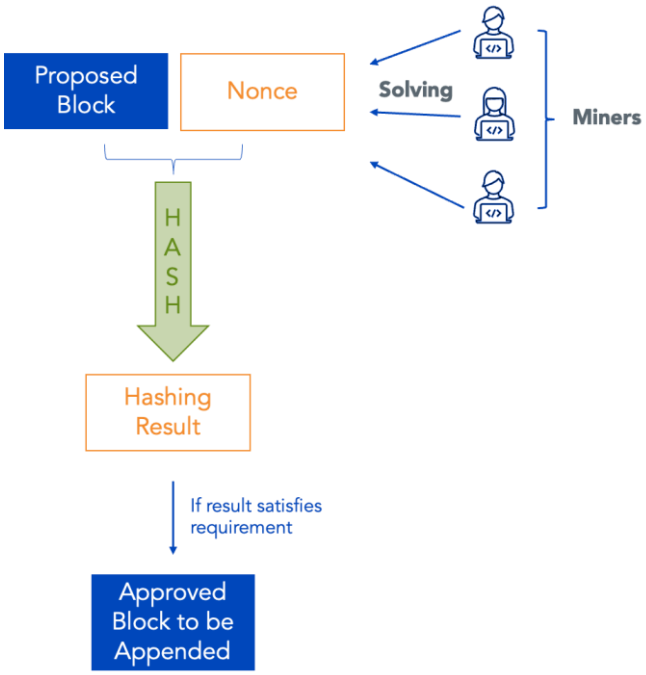

The advantage of this mechanism is its security and simplicity as a system that only relies on computing power. It also has a relatively long history – it was proposed in 1993 and put into use in Bitcoin's protocol in 2009.

But the downside is that it is energy intensive as each miner uses computing power to hash blocks trillions of times a second. It also tends to be more centralized as the mechanism requires advanced machines to run the operation. Miners with aggregated advanced machines have an advantage in the number of hashes they can run and have more chance of producing a valid block.

This leads to the possibility of a **51% attack**, in which hackers can control the protocol by amassing more than 50% of the hash rate. However, for Bitcoin, it is expensive and nearly impossible to do so given its size and high hash rate. Satoshi thus dismisses this concern in the Bitcoin whitepaper. 51% attack poses a general threat to all blockchains despite their underlying consensus mechanisms, since hackers just need to collect 51% of the verifying assets of the protocol. The difference lies in the difficulty of doing so.

Proof-of-Stake (PoS)

Proof-of-stake is the second most popular consensus mechanism. Ethereum is currently transitioning from PoW to PoS to be more efficient, scalable, and sustainable.

Unlike PoW, PoS doesn't require participants to use computing power to hash blocks and solve a mathematical requirement, but it requires them to stake ether. The participants are called **validators** as opposed to miners, and the action of appending is referred as **minting** instead of mining.

PoS can also be described in two processes: **selection** and **attestation**.

To become a validator, participants need to set a certain amount of the cryptocurrency aside, and depending on the blockchain, validators are selected differently. One of the more popular selection methods is **randomized block selection**, which will be used by Ethereum. The mechanism randomly selects a group of validators based on the amount of cryptocurrency they staked. The more wealth one stakes, the more likely they are selected.

After being selected, the attestation process begins.

Validators need to stake an amount of cryptocurrency that covers the transaction fee and their potential reward until the block is successfully appended. Validators individually attest to the block and broadcasts their decision to the network. If a certain number of validators approve it, the block gets appended, and a proportionated reward is given to the validators. If the block ends up being rejected by the group of validators, the block will not be appended.

Fraud is detected through a pre-coded set of rules that will be triggered by inconsistent, absent, and abnormal behaviors. Dishonest participants could lose their stakes and be banned from the network.

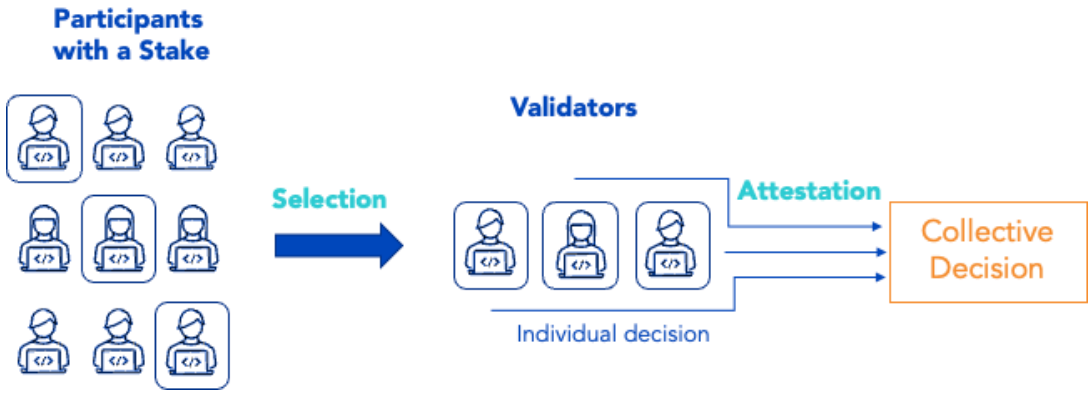

PoS reduces mining energy as it only requires staking. It also doesn't require validators to have advanced machines, which lowers barrier for entry and makes it less concentrated than PoW[3] However, PoS faces several potential exploits from the nothing-at-stake problem[4] and long-range attacks that demand mitigations from other areas.

[3] Buterin, Vitalik. Why Proof of Stake. https://vitalik.ca/. November 6, 2020.

[4] Sharma, Abhishek. Understanding Proof of Stake through it's Flaws (Pt. 1). Medium.com. January 15, 2018.

Delegated Proof-of-Stake (DPoS)

Delegated proof-of-stake is a variation of proof-of-stake. It changes the selection process in PoS from randomized algorithms to a more democratic approach.

Coin holders with a stake in a DPoS network constantly select block validators by voting. Voting power increases with one's wealth. All coin holders have the potential to be elected by convincing other participants to vote for them. Some of the considerations that play into the voting decision includes hardware robustness, dedicated teams, etc.

Elected validators are called **delegates** or **witnesses**, and they are responsible for signing and verifying new blocks. To prevent manipulation, a group of witnesses (normally 21-100)[5] are selected for a certain period. During that period (called an **epoch**), witnesses take turns verifying and signing new blocks with their private keys. These signed blocks are then left unconfirmed until a majority of the witness group approves.

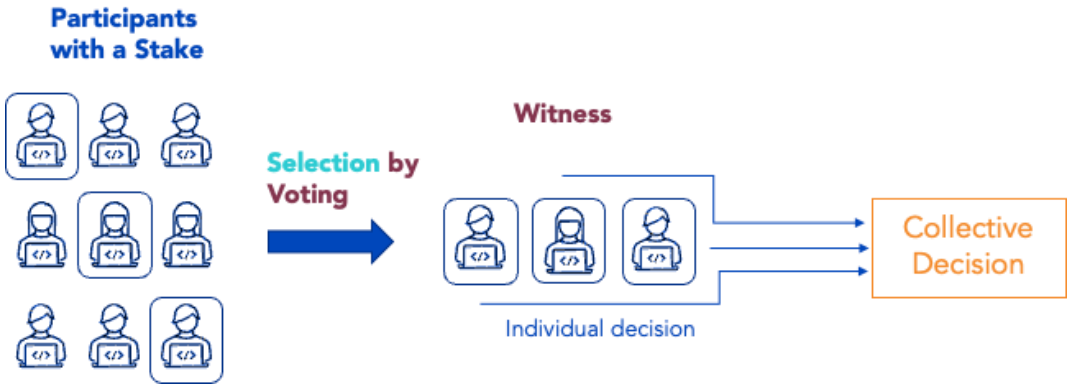

Approved blocks reward witnesses, and these rewards are usually shared with voters. Failed blocks leave no reward. Malicious actors would be quickly voted down once discovered.

DPoS introduces collective human judgement to replace a competition in pure computational power. It can process many more transactions even when compared to PoS. It also reduces the locking time for collateral being staked. By allowing participants to vote, it makes the network more democratic than PoW and PoS.

However, the voting process may also centralize the network as coin holders with a small stake would forfeit their votes given their insignificance.

Other Consensus Mechanisms

Besides PoW, PoS, and DPoS, there are many proof-of-X mechanisms that try to establish a decentralized and secure network. They include proof-of-capacity, proof-of-elapsed time, proof-of-importance, etc.

[5] Walter, Cedric. Delegated Proof of Stake (DPoS). https://tokens-economy.gitbook.io/consensus/. 2018.

Another major family of consensus mechanisms is **Byzantine Fault Tolerance**. It has several variations such as **practical Byzantine Fault Tolerance (pBFT)**, which is currently used by Hyperledger Fabric, and its improved version is used by the People's Bank of China (PBoC) to develop its Central Bank Digital Currency (CBDC). Another variation is called **delegated Byzantine Fault Tolerance (dBFT)**, which is used by Neo. The Stellar network's model of consensus leverages a **federated Byzantine agreement (FBA) model**, and it seeks to build upon these models to establish an open network for storing and moving money.

These mechanisms are more commonly used in permissioned protocols, which requires prior permissions to join. They are early solutions of the **Byzantine General Problem**, which is the same problem Bitcoin tries to solve.

The advantages of these mechanisms include energy efficiency, transaction finality which requires less validations from nodes to confirm a transaction, and increased ability to coordinate in a closed system. However, they have a lower tolerance for malicious nodes in the network and may be hard to scale, so they're not an ideal candidate for public blockchains.

Technical Comparison of Major Consensus Mechanisms				
	PoW	**PoS**	**DPoS**	**pBFT**
Type	Probabilistic-finality	Probabilistic-finality	Probabilistic-finality	Absolute-finality
Fault Tolerance	50%	50%	50%	33%
Power Consumption	Large	Less	Less	Negligible
Application	Public	Public	Public	Permissioned

Source: Zhanga, Shijie and Lee, Jong-Hyouk. Analysis of the main consensus protocols of blockchain. Science Direct Vol. 6, Issue 2. June 2020. This chart summarizes the differences of discussed consensus mechanisms. Finality type refers to the model of how committed blocks are confirmed and irreversible. Probabilistic finality means that blocks are increasingly difficult to be reverted as the blockchain gets longer. Absolute finality means that blocks are finalized as soon as they are appended to the blockchain. Fault tolerance refers to a system's tolerance of malfunctioned or malicious components that would prevent it from operating. Power consumption refers to if the system consumes large amount of energy. Scalability refers to how easy the system can grow and expand. Application refers to the ideal type of blockchain the consensus mechanism should be utilized in. Public refers to blockchains that can be accessed to everyone. Private refers to blockchains that require permissions to join.

Conclusion

The consensus mechanism is a key component to a decentralized network. It not only secures the system but also affects its efficiency and scalability.

When examining a consensus mechanism, it is important to consider its functionality based on three aspects: degree of decentralization, security, and scalability. Most of the consensus mechanisms can only optimize on two of the three factors – security comes at a cost of scalability and the risk of centralization, and scalable networks might be more susceptible to attacks.

Therefore, different consensus mechanisms must be analyzed based on the needs of the particular network.

Additional Risks

There are risks associated with investing, including the possible loss of principal. Crypto assets, such as bitcoin and ether, are complex, generally exhibit extreme price volatility and unpredictability, and should be viewed as highly speculative assets. Crypto assets are frequently referred to as crypto "currencies," but they typically operate without central authority or banks, are not backed by any government or issuing entity (*i.e.,* no right of recourse), have no government or insurance protections, are not legal tender and have limited or no usability as compared to fiat currencies. Federal, state or foreign governments may restrict the use, transfer, exchange and value of crypto assets, and regulation in the U.S. and worldwide is still developing.

Crypto asset exchanges and/or settlement facilities may stop operating, permanently shut down or experience issues due to security breaches, fraud, insolvency, market manipulation, market surveillance, KYC/AML (know your customer / Anti-Money Laundering) procedures, non-compliance with applicable rules and regulations, technical glitches, hackers, malware or other reasons, which could negatively impact the price of any cryptocurrency traded on such exchanges or reliant on a settlement facility or otherwise may prevent access or use of the crypto asset. Crypto assets can experience unique events, such as forks or airdrops, which can impact the value and functionality of the crypto asset. Crypto asset transactions are generally irreversible, which means that a crypto asset may be unrecoverable in instances where: (i) it is sent to an incorrect address, (ii) the incorrect amount is sent, or (iii) transactions are made fraudulently from an account. A crypto asset may decline in popularity, acceptance or use, thereby impairing its price, and the price of a crypto asset may also be impacted by the transactions of a small number of holders of such crypto asset. Crypto assets may be difficult to value and valuations, even for the same crypto asset, may differ significantly by pricing source or otherwise be suspect due to market fragmentation, illiquidity, volatility and the potential for manipulation.

Crypto assets generally rely on blockchain technology and blockchain technology is a relatively new and untested technology which operates as a distributed ledger. Blockchain systems could be subject to internet connectivity disruptions, consensus failures or cybersecurity attacks, and the date or time that you initiate a transaction may be different then when it is recorded on the blockchain. Access to a given blockchain requires an individualized key, which, if compromised, could result in loss due to theft, destruction or inaccessibility. In addition, different crypto assets exhibit different characteristics, use cases and risk profiles. Information provided by WisdomTree regarding digital assets, crypto assets or blockchain networks should not be considered or relied upon as investment or other advice, as a recommendation from WisdomTree, including regarding the use or suitability of any particular digital asset, crypto asset, blockchain network or any particular strategy. WisdomTree is not acting and has not agreed to act in an investment advisory, fiduciary or quasi-fiduciary capacity to any advisor, end client or investor, and has no responsibility in connection therewith, with respect to any digital assets, crypto assets or blockchain networks.

Value: Characterized by lower price levels relative to fundamentals, such as earnings or dividends. Prices are lower because investors are less certain of the performance of these fundamentals in the future. This term is also related to the Value Factor, which associates these stock characteristics with excess returns vs the market over time. Blockchain: a distributed ledger system in which a record of transactions made in cryptocurrencies are maintained across computers linked in a peer-to-peer network Bitcoin: A digital currency (also called a cryptocurrency) created in 2009, which is operated by a decentralized authority as opposed to a traditional central bank or monetary authority. Proof-of-capacity: a consensus mechanism algorithm used in blockchains that allows for mining devices in the network to use their available hard drive space to decide mining rights and validate transactions. Proof-of-time-elapsed: a consensus mechanism algorithm that is often used on the permissioned blockchain networks to decide the mining rights or the block winners on the network. Proof-of-importance: a blockchain consensus technique – essentially, proof of importance works to prove the utility of nodes in a cryptocurrency system, so that they can create blocks.