MULTICOIN CAPITAL

Delegated Proof of Stake: Features & Tradeoffs

By Myles Snider,
Tushar Jain, & Kyle Samani

Multicoin Capital is a thesis-driven cryptofund that invests in tokens reshaping entire sectors of the global economy. We rigorously research blockchain protocols, teams, and market opportunities to deliver venture capital economics with public market liquidity.

Report Disclosures

***Disclosures**: As of the publication date of this report, Multicoin Capital Management LLC and its affiliates (collectively "Multicoin"), others that contributed research to this report and others that we have shared our research with (collectively, the "Investors") may have long or short positions in and may own options on the token of the project covered herein and stand to realize gains in the event that the price of the token increases or decreases. Following publication of the report, the Investors may transact in the tokens of the project covered herein. All content in this report represent the opinions of Multicoin. Multicoin has obtained all information herein from sources they believe to be accurate and reliable. However, such information is presented "as is," without warranty of any kind – whether express or implied.*

This document is for informational purposes only and is not intended as an official confirmation of any transaction. All market prices, data and other information are not warranted as to completeness or accuracy, are based upon selected public market data, and reflect prevailing conditions and Multicoin's views as of this date, all of which are accordingly subject to change without notice. Multicoin has no obligation to continue offering reports regarding the project. Reports are prepared as of the date(s) indicated and may become unreliable because of subsequent market or economic circumstances.

Any investment involves substantial risks, including, but not limited to, pricing volatility, inadequate liquidity, and the potential complete loss of principal. This report's estimated fundamental value only represents a best efforts estimate of the potential fundamental valuation of a specific token, and is not expressed as, or implied as, assessments of the quality of a token, a summary of past performance, or an actionable investment strategy for an investor.

This document does not in any way constitute an offer or solicitation of an offer to buy or sell any investment or token discussed herein.

The information contained in this document may include, or incorporate by reference, forward-looking statements, which would include any statements that are not statements of historical fact. These forward-looking statements may turn out to be wrong and can be affected by inaccurate assumptions or by known or unknown risks, uncertainties and other factors, most of which are beyond Multicoin's control. Investors should conduct independent due diligence, with assistance from professional financial, legal and tax experts, on all tokens discussed in this document and develop a stand-alone judgment of the relevant markets prior to making any investment decision.

Note: Most of this analysis refers to DPoS as implemented in BitShares, Steem, and EOS. Other platforms use a similar DPoS framework but alter certain features. Much of this analysis will focus specifically on EOS's consensus algorithm. We will publish a full analysis and valuation of EOS in the future.

Introduction

Distributed ledgers don't easily scale. That fact has become readily apparent in the last few years as Bitcoin, Ethereum, and others have faced serious challenges as they attempt to increase the speed and throughput of their platforms.

This problem can be best understood as a scalability trilemma (this idea was first formalized by Vitalik Buterin and Trent McConaghy). The scalability trilemma posits that any blockchain system in which every node validates every transaction can have only two of three potential properties: decentralization of block production (DBP), safety, and scalability. These properties can be defined as follows:

- DBP can be quantified as the number of block producers.
- Safety can be quantified as the cost of mounting a Byzantine attack that affects liveness or transaction ordering. Note that safety does not refer to the integrity of cryptographic signatures, or the ability of a 3rd party to derive a set of private keys from public keys.
- Scalability can be quantified as the number of transactions per unit of time that the system can process.

While projects like Ethereum, Dfinity, Polkadot, and Kadena are attempting to *solve* the scalability trilemma via sharding, alternative consensus schemes, and other techniques, we don't yet have a live platform that has solved this trilemma. Even if one of these projects does manage to solve the scalability trilemma, the market may not care. It is entirely possible that users are willing to accept tradeoffs in decentralization of block production or safety in the name of better performance and easier user experience for certain use cases.

Decentralization is valuable to ensure that any given party cannot alter the database. More decentralization means it is harder to collude to alter the database. There are different levels of protection which are necessary for different use cases. Bitcoin, being censorship-resistant money, is designed for sovereign-grade protection; it is designed to withstand an attack by a large nation-state. However this isn't necessary for most decentralized applications (dApps). These dApps need platform-grade protection; global, neutral databases uncontrolled by any one party.

Delegated Proof of Stake (DPoS) concentrates block production in the hands of just a few, known, semi-trusted entities in order to achieve orders of magnitude more scalability than proof-of-work (PoW) or other proof-of-stake (PoS) blockchains. In this analysis, we'll examine the features and tradeoffs of DPoS.

Delegated Proof of Stake

Delegated proof of stake (DPoS) is a consensus algorithm invented by Dan Larimer in 2013. DPoS was originally invented to power BitShares, Larimer's first blockchain project. He refined it in his second project, Steem, and is refining it further in EOS, which he's been working on for about one year. While Larimer invented DPoS and continues to evolve the algorithm, several other projects have adopted DPoS and made changes. Current blockchains utilizing DPoS include:

- EOS, BitShares, Steem, Golos, Ark, Lisk, PeerPlays, Nano (formerly Raiblocks), and Tezos
- Cosmos/Tendermint, Cardano, and a few others use consensus algorithms loosely based on DPoS

In DPoS, those who hold the network token are able to cast votes to elect block producers; votes are weighted by the voter's stake, and the block producer candidates that receive the most votes are those who produce blocks. Users can also delegate ("proxy") their voting power to another user who can vote on their behalf. **DPoS is a liquid, representative democracy with token holder suffrage.** DPoS can also be thought of as a formalized, digital version of a traditional organizational hierarchy that operates in a completely transparent way. While there are problems with both democracy and corporate governance that are beyond the scope of this paper, one compelling features of DPoS is that the open-source nature of these protocols means that users can fork if they disagree with the majority. The same cannot be said of democracies, corporations, and other organizational structures. DPoS adopts ideas from many traditional governance models, but is ultimately far more flexible and transparent.

Block producers can be voted in or out at any time, so the threat of loss of income and reputation is one of the major incentives against bad behavior. Additionally, slashing conditions can be implemented in DPoS rather trivially. Most traditional PoS implementations allow users to produce blocks proportional to their stake in the network. DPoS allows users to cast votes proportional to their stake to decide who produces blocks. Block producers themselves do not necessarily need to have a large stake, but they must compete to receive votes from users.

DPoS can power entire blockchains, or it can be used as a consensus algorithm for child chains, sidechains, private blockchains, and more. DPoS could be used to power consensus within Ethereum Plasma chains, and DPoS bears many similarities to the "Proof of Authority" consensus mechanism formalized by Parity. It could also be a solution for application-specific chains like those in Cosmos zones.

For other in-depth documentation on DPoS, see this link and this link from BitShares, and this white paper from Larimer.

DPoS Features and Tradeoffs

The core elements of DPoS are the following:

BLOCK PRODUCERS

Like other PoS chains, DPoS doesn't include miners who run hashes to produce blocks . Instead, an elected subset of users is chosen to perform the work of validating the chain. We'll simply refer to these users are block producers, though they are sometimes called delegates, notaries, validators, forgers, or witnesses.

Because the token holders decide on a elected subset of users to produce blocks, mechanisms to allow a wider group of computers to participate in block production (which ultimately slow down block production) can be removed. DPoS can be seen as a form of "controlled semi-centralization" that gets the benefits of semi-centralization (efficiency and speed) while still maintaining some calculated measure of decentralization (X # of independent block producers that can be voted in and out by token holders).

The multi-billion-dollar question then becomes "How many block producers are necessary in order to be sufficiently decentralized?" This is a loaded question and one that divides the crypto community. It's also the single feature that is often presented as a game-ending criticism of DPoS: DPoS is not decentralized enough.

Importantly, decentralization is a spectrum, and increased decentralization often also incurs higher costs. There are many different measures of decentralization. Some like Balaji Srinivasan have attempted to quantify decentralization, but even he admits that his proposals need more work. He breaks down networks into subsystems, each of which can be measured differently and each of which contributes to an overall notion of system-wide decentralization. Some people may have opposing views about which subsystems to include and how much weight to assign to each. Ultimately, measures of decentralization can produce different results when different variables are examined, and they don't produce a binary result. Systems are not *decentralized or not*; rather, some are more decentralized than others, though this measure can be somewhat subjective.

The goal is not decentralization for its own sake. **Decentralization is a feature of systems that allows them to achieve other goals: censorship resistance, open participation, immunity from certain attacks, and elimination of single points of failure.** While some features that contribute to decentralization can be quantified, the phenomenon as a whole cannot be. The number of block producers is just one metric, but it fails to capture all of the relevant subtleties. It also doesn't describe how much decentralization is necessary to accomplish the underlying goals, or how to think about theoretical vs practical decentralization.

DPoS doesn't attempt to "find" a balance between the number of block producers needed to ensure that control is sufficiently decentralized and the number of block producers that can easily be monitored for bad behavior. Rather, it explicitly sets the balance, though it can be modified later.

There does not seem to be a *single* number that encapsulates this. In our view, 20 block producers all located in China is less decentralized than 10 block producers spread across various jurisdictions around the globe. Cornell's IC3 team recently wrote a paper that attempted to quantify decentralization in Bitcoin and Ethereum. They ultimately found that block production in Bitcoin and Ethereum was *far more* concentrated than commonly thought. They noted:

> "These results show that a Byzantine quorum system of size 20 could achieve better decentralization than proof-of-work mining at a much lower resource cost. This shows that further research is necessary to create a permissionless consensus protocol without such a high degree of centralization."

DPoS is one potential solution to this problem; we look forward to more research on this front. The benefits of decentralization can't be accurately measured by a single number—they are emergent properties that must be observed in an iterative, dynamic, unpredictable, real-world environment. **Token voters in DPoS systems will have to take into account not just how many block producers there are, but also in which jurisdiction they are located, with whom they are affiliated, and more. If voters do enough due diligence to ensure that 21 individual entities located in different jurisdictions are producing blocks, then validation in DPoS has the potential to be far more decentralized than almost any other blockchain.**

For an argument of why DPoS is more decentralized *in practice* than PoW, see this excellent post from Ian Grigg. In it, he describes how the Chinese government could launch an attack on Bitcoin miners within the country, either by directly taking them over, forcing them to shut down, or exploiting the great fire wall that controls the in/out pipes of the internet in China. This would be hugely disruptive for the Bitcoin network, and would render it severely compromised for a period of time, even if it were eventually able to recover. DPoS, on the other hand, could easily avoid this problem by voting out Chinese block producers and replacing them with block producers in other jurisdictions as soon as (or even before) the government took action.

One interesting feature of DPoS is that while block producer *candidates* compete for token holder votes, elected block producers actually *cooperate* to secure the network during rounds. Block producers share block rewards (from inflation) equally, and they are only allowed to produce one block per round. There are no incentives for block producers to compete to try to produce more blocks than the other producers,

since this is impossible. In EOS and Steem, block producers are funded entirely by inflation; there are no transaction fees. Block producers aren't incentivized to order transactions based on who pays the highest fee; rather, priority is given relative to overall stake—ownership in the network gives users a claim to bandwidth within the network when the network is at capacity. Block producers compete off-chain to get votes, but once they are voted in they cooperate to secure the chain. And since the number of block producers is fixed, block-producing power doesn't concentrate, even with economies of scale.

CENTRALIZATION OF BLOCK PRODUCTION

The most notable feature of any DPoS-powered protocol is that the number of block producers is explicitly limited. The number of block producers varies in different implementations—101 in Lisk, 21 in EOS, etc. Often the number itself is a parameter that can be changed by a token holder vote.

DPoS validation happens in rounds; each round consists of a period in which each block producer is given one slot to produce a block. In Lisk, for example, a round would consist of 101 blocks. At the beginning of each round, each block producer is assigned a slot. Each slot corresponds to one block and is the only time in that round in which the assigned block producer can produce a block. If a producer fails to create a block during their slot, then that block is skipped and the transactions from that slot are included in the next block. Token holder votes are tallied each round, and block producers can be voted in or out each round. In every DPoS implementation, the number of potential block producers is always greater than the number of producers allowed in each round. Thus, there are always block producers ready to step in if a malicious producer is voted out. It is also possible to configure DPoS parameters to compensate backup block producers in order to incentivize them to be ready to fill the spots of producers that are voted out (see total witnesses in Steem).

Collusion and censorship by validators is always a concern with any blockchain protocol. If three of the largest Ethereum mining pools colluded, it would be possible for them to pierce the safety of the network. DPoS is aware of these risks and attempts to mitigate them through transparency. In DPoS, token holders are directly responsible for deciding who controls validation. While this puts more responsibility on individual token holders, it also means that the owners of the network have recourse if validators behave badly. If Ethereum mining pools colluded, individual miners participating in the pool would have to point their miners elsewhere, or the community would have to hard fork the network. Both scenarios require some form of off-chain coordination. If DPoS block producers colluded, the community could vote them out within a single round and replace them with honest block producers. This still requires a form of ex-protocol coordination (token holders deciding how to re-allocate their votes), but it is more formalized and arguably more trivial to enact. This architecture allows DPoS to introduce a form of centralization while still maintaining security. The implications of this (as well as potential attack vectors) will be explored later in this document.

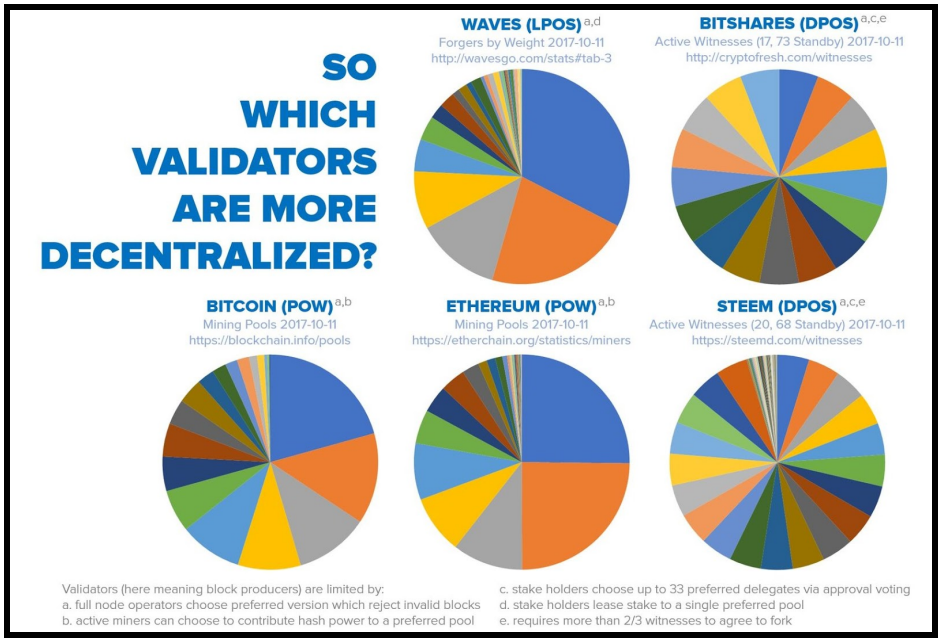

Source

SCALABILITY

Known and limited block producers means that blocks can be propagated through the network much more efficiently, enabling significant scalability increases. Blocks can also be consistently and reliably produced in a much smaller time frame (BitShares currently produces a block every 3 seconds). Furthermore, finality can be reached as soon as ⅔ of the block producers have confirmed a transaction, with strong guarantees that a transaction is on a valid chain even before that.

While DPoS is certainly more scalable than PoW, exactly how fast it can go depends on a variety of factors. EOS is attempting to optimize for extremely high throughput by using the WASM virtual machine and employing a message-based architecture instead of a state-based one. Early results showed 50,000 tps on an EOS smart contract, but those results are not indicative of full-scale main net performance. Recently, the community-run EOS test net hit 600 tps with beta software and non-specialized machines, and Larimer recently claimed in an update that the software could be set to debut with 5,000 tps in its single-threaded architecture before moving to parallel execution.

It is worth noting that three of the top 5 blockchains with the most operations performed daily are DPoS blockchains, according to Blocktivity.info. And while Ethereum is consistently operating at near capacity, both BitShares and Steem have no pending transactions and plenty of bandwidth to spare. Smart contract platforms that host large-scale dApps need to be able to handle many thousands of operations per second, whether they are simple likes or million-dollar value transfers. This chart provides a great breakdown of

operation type and count for BitShares, while this link provides statistics for Steem. For more information on how BitShares achieves its performance, see detailed explanations here and here.

NETWORK INFRASTRUCTURE AS A SERVICE

Every blockchain network can be divided into two major subsets of actors—those who perform operations on the network and those who validate the operations. We can refer to the first group as users and the second group as validators.

In Bitcoin, Ethereum, Monero and other proof of work-based blockchains, the validators are miners. Miners race to solve computationally intensive puzzles, and the first to solve the puzzle is allowed to produce a block (and collect transaction fees and the block reward). In traditional PoS, users must bond their tokens, and they are awarded the ability to produce blocks proportional to their bonded stake.

In each of these scenarios, validators are providing key network infrastructure: They are collecting transactions, ordering them, and preventing double-spends. In PoW, validators are those with access to the best hardware and cheapest electricity. In PoS, validators are those with large ownership stakes in the network. In DPoS, however, the validators can be thought of as contractors that are hired by the owners of the network (token holders). They are hired (voted in), given a job (to produce blocks), paid (through inflation or transaction fees), and can be fired (voted out) for not performing their duties.

This structure means that DPoS network token holders, as the owners of the network, ultimately have control over who provides infrastructure, while Bitcoin and Ethereum give token holders no choice. If miners misbehave, as they have incentive to do, users have no recourse. DPoS is the *only* algorithm that allows infrastructure providers to be easily *fired* (no slashing required) and replaced for not providing a good service. Reasons for firing could include any of the following:

DISHONESTY

- The appearance of dishonesty, or even a mere lack of sufficient transparency compared to other candidates
- Greed (e.g. trying to demand higher block rewards than other producers)
- Censorship
- Malicious collusion
- Support of controversial or malicious changes to the network
- Failure to support community-backed changes
- Regulatory fear based on jurisdiction (for example, Chinese block producers could be voted out if China announced a crackdown on crypto)

ON-CHAIN GOVERNANCE

DPoS is inherently a form of on-chain governance; it uses stake-weighted voting to allow the owners of the network (token holders) to make decisions about the network. DPoS is a form of liquid representative democracy where voting power can be allocated to other participants and votes can be changed at any time. While voting for block producers is the primary governance use case, token holder voting can also be used to decide on things like development funding, monetary policy, network parameters, hard forks and more.

Blockchain governance is still very much a nascent field, and there exists a lot of disagreement over which approach to blockchain governance is the most promising. Some, like Ethereum researcher Vlad Zamfir, have argued that on-chain governance is a bad idea because, among other reasons, it negates the role of non-block-producing full nodes in the governance process. Analyst Nic Carter similarly concluded that Bitcoin's informal off-chain governance, which consists of several different social and technological layers, is the ideal form of governance for a decentralized network. Fred Ehrsam, co-founder of Coinbase, argued for on-chain governance as a way to bring formal structure to these messy interactions. Many point to the Bitcoin scaling debate, the Ethereum DAO fork, and the recent Parity wallet bug debate as three examples of situations where on-chain governance could have more easily rectified very messy and uncertain situations.

DPoS chains embrace the fact that all blockchain networks are inherently political and seek to formalize the political process. While there are certainly issues with on-chain governance and token voting (which we'll explore later), both are key features of DPoS. DPoS is a community-owned operational hierarchy that operates in a fully transparent, decentralized way. While it is not clear that on-chain governance is better than other forms of blockchain governance, it certainly isn't clear that it is worse. We believe strongly that this approach should be tried and tested.

SELF-FUNDING THROUGH INFLATION

Almost every major blockchain pays for infrastructure with inflation. In the case of Bitcoin and Ethereum, miners receive block rewards as compensation for validating the blockchain. When block rewards run out in the future, infrastructure will have to be supported through fees alone. This raises questions about how high fees will get in the future, how that will affect incentives to mine, and whether that will affect the security of the chain.

EOS and Steem, because they have no transaction fees, use an entirely different model. Not only do they use inflation to pay block producers to provide infrastructure, but they also use inflation to fund the

development of the platform itself. Token holders can vote on a maximum annual inflation rate, initially set at 5%. This number can be changed, as it has been several times in Steem. Token holders also vote on how much of the annual inflation is paid directly to block producers. If the token price increases, users can decide whether to keep block producer pay steady (by lowering block rewards) or to allow block producers to capture additional profits that can be used to scale up their infrastructure. That which is not paid to block producers can be paid to a set of community smart contracts that can be used in a wide variety of ways. A contract could be a development fund that pays out developers based on community votes; it could go directly to a company that is actively working on development; it could be used to fund hackathons; it could be burned; and more. DPoS actually makes it possible for developers, marketers, and others building community tools to be funded *by the blockchain itself*.

Many people struggle with the concept of inflation and are adverse to relying on inflation funding. This shouldn't be the case. Inflation is perhaps the only method by which blockchains can be funded in a fair way because it solves the tragedy of the commons problem. Some blockchains, like Monero, rely entirely on community crowdfunding of development initiatives. While the generosity of the Monero community is remarkable, it remains unclear whether that is a sustainable way to fund development. Everyone benefits from the development advances made, but only a minority of users are willing to contribute their own money to funding. With inflation funding, all users *collectively* fund development and security ratably, and they all collectively reap the benefits. As Fred Ehrsam points out, inflation funding can actually be a net positive for token holders:

> "If Ether holders believed an upgrade (ex: sharding) would make the price go up by >10%, they'd be happy to pay close to 10% of their tokens for it. That means Ethereum could crowdfund a $3bn feature bounty by inflating the number of ETH by 10% and pay the newly created tokens to the creator(s) of the upgrade. This is somewhat analogous to taxes: everyone in the community chips in to fund common infrastructure (ex: roads) which no one would build alone."

Every blockchain must pay to secure its network through either transaction fees, inflation, or both. Transaction fees force active users to pay, while passive users (hodlers) don't. A user could secure her entire life savings in Bitcoin or Monero without almost ever contributing to the security of the platform through transaction fees. This creates a free-rider problem. Transaction fees are also variable and unpredictable, and may need to be astronomically high in order to pay for network security. Inflation is a more equitable and user-friendly way of securing the network.

Similar approaches have been tried by other projects in a number of ways. Zcash collects a "Founders' Reward," that sends 10% of the total money supply to the Zcash Company and its shareholders. Dash collects a portion of block rewards for a masternode-vote development fund. DPoS formalizes these arrangements, bakes them directly into the protocol, and allows for maximum flexibility and accountability.

DPoS Attack Vectors

Below, we'll outline the major DPoS attack vectors and evaluate the threat they present.

NOTHING-AT-STAKE

The "Nothing at Stake" problem is a flaw in some PoS schemes. Specifically, it refers to the fact that in the case of a fork, there is very little additional cost to the validator to validate on both chains. This is the optimal strategy for validators, since it is likely the most profitable.

DPoS does not suffer the nothing-at-stake problem in practice. Token holders in DPoS are using their stake to vote on validators, not on blocks. DPoS is a longest-chain-wins algorithm. Because the number of validators is fixed and the order is decided each round, it would be impossible for a minority subset of validators to produce a fork that overtook the main chain.

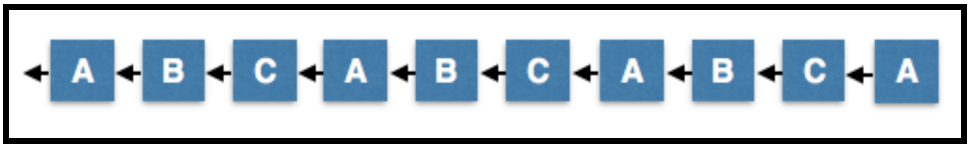

DPoS block production under normal network conditions (source)

If a single block producer were to produce blocks on multiple forks (as shown below), the main chain would still advance with the rest of the honest block producers. The longest chain is considered the canonical chain, so the producer is unable to do any harm.

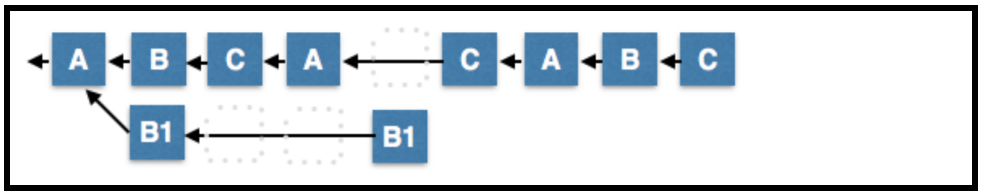

Source

Even if a majority of block producers colluded to produce blocks on several forks (as shown below), the honest minority would still determine the longest chain. In each of these instances, there would be clear cryptographic evidence that block producers had created blocks on multiple conflicting chains, and they could be voted out for doing so. It would also be possible to introduce slashing conditions and other *protocol-level* punishments based on cryptographic evidence of Byzantine behavior, but these additional features aren't always necessary.

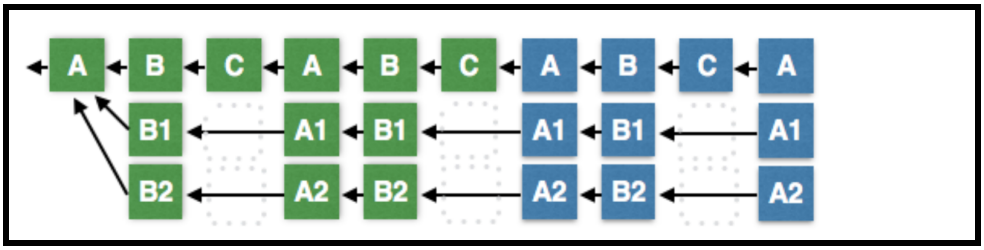

Source

It is possible for block producers to create blocks on multiple forks at little additional cost, but they do have some things at stake—their job, reputation, and future income stream. Because Byzantine behavior can be detected, it is a risky move. To actually corrupt the integrity of the chain would require collusion among a strong majority of delegates—at that point the attack becomes less about "nothing-at-stake" and more about traditional Byzantine fault tolerance.

For a detailed explanation, see Dan Larimer's DPoS white paper.

EXPLOIT LOW VOTER TURNOUT

This is the most obvious attack against a DPoS blockchain. The core of this attack is the fact that in any voting system, very few participants actually show up and vote. In blockchain token voting systems, anyone with a small stake is unlikely to actually influence the direction of the platform with their vote. The time spent researching on how to vote may not be worth the effort for what they view as a minimal impact. Voters with a small stake may practice rational ignorance— the time spent researching how to vote may be more costly to them than the value that voting brings. DPoS attempts to at least partially rectify this by allowing proxy voting, in which a user can lend their voting power to another user who they consider more informed. In this case, the effort of deciding to whom to delegate voting power is likely less than the value gained. Still, the natural result of these systems is often that overall voter turnout is low, and voting is mostly done by whales, exchanges, and wallet providers. This problem has been explored by Vitalik in his blog post on blockchain governance.

Say, for example, that only 10% of the total supply of tokens was being used to vote. A whale (or group of whales) with more than 5% of the total supply could step in and take over governance. For context, the top block producer on BitShares has about 33% approval. Most DPoS systems use approval voting in which users split their votes among all candidates, and the top candidates by total approval become the block producers. This makes it much more difficult for a whale to take over the voting process.

One important caveat is that attacks on the voting system are bad for the network as a whole, and successful ones will likely result in a decrease in the price of each token. Anyone with a significant stake in the network should be incentivized to vote to protect the value of their tokens. While it is entirely possible

that the percentage of token holders who vote may be small, the ones who do will be those with the largest stake. That means that an attacker would still have to purchase a very significant stake in order to take over governance. Like all other PoS systems, DPoS will likely find follow the Pareto Principle, where a small subset of large stakeholders does most of the work related to decision making. This is not necessarily a bad thing, as large stakeholders have incentives to improve the network.

In fact, participation is a problem that plagues all proof of stake systems, including Ethereum's proposed Casper PoS. Because ETH has utility outside of just staking, it is likely that only a fraction of all ETH in circulation will actually be staked to secure the network. If this is a very small percentage of total ETH, then an ETH whale could step in at any time, stake their tokens, and take over the validation process. This threat is described well in this Cosmos blog post. One advantage to the DPoS model is that, in most designs, DPoS tokens can delegate their voting power (either directly to a block producer or to a proxy voter) and still retain all of their utility. So participation in the voting system has no cost other than the time spent deciding how to allocate one's voting power. While small stakeholders in Ethereum may use staking pools to earn a passive return on their assets, the process of joining a staking pool may not be super easy, and there are capital lockup costs associated with doing so. DPoS token holders can quite easily delegate their voting power to someone they trust to do proper research into block producers, and they are still free to utilize their tokens in any way they choose to do so.

Some models do require DPoS voters to lock up their tokens for a period of time when voting in order to incentivize votes that have some "skin in the game." This pushes users to cast more informed votes that take into account the long-term success of the platform, but it also limits the number of participants willing to vote. This is currently a subject of debate within the EOS community.

Voter participation will likely depend on a few things—voting participation as part of the social contract, how well the community encourages it, how easy it is for users to vote or delegate voting power, and more. Education and easy-to-use voting interfaces will be especially important in this regard.

BRIBING ATTACKS

Another attack, and one that has been observed in practice, is that of block producers paying for votes. This issue has affected Lisk and Steem recently. Whether this should be considered an attack depends on perspective and may hinge upon what the elected block producers do with their purchased validation power. Still, this situation is not desirable, so it will be examined in the context of maliciousness.

Lisk, for example, has two pools (LiskElite and LiskGDT) that promise a portion of block rewards will be paid back to the users who vote for their delegates. There is even a website dedicated to helping users find the delegates with the best payouts and encouraging users to vote out those that don't share rewards.

The sustainability and efficacy of these attacks depends on other elements of the protocol. The attack is possible in any implementation of DPoS (or any implementation of on-chain voting, in general), but other features can make it much less likely. In system where the requirements for a block producer are simply to run a cloud instance that does the validation work, profit sharing is an easy option. Blockchains like Lisk, which are still in their infancy and don't have much usage, don't require a lot from block producers.

In EOS, however, the requirements for block producers go far beyond just running software. EOS block producers are also expected to provide storage, participate in governance, and gradually use their profits to scale up their hardware in order to increase the total capacity of the network. Larimer has even said that eventually EOS will scale to 21+ data centers with gigabit connections speaking directly to one another. The operational costs for block producers will be much higher than in other systems, and paying for votes will cut into profit margins. Block producers who pay for votes will also have less resources to scale up their systems, so the network itself will suffer. Voters who care about the long-term health of the network (and price of the token) will prefer block producers who contribute to increased network capacity rather than providing short-term kickbacks. It will be in the interest of businesses built on top of EOS (who will likely be among the largest token holders) to do vote for block producers who don't pay voters.

If a block producer offers nothing but a share in the block reward, it is unlikely that he or she could complete for very long without becoming a tax on the network. In Steem, for example, witnesses campaign with plans about how they will work to improve the network. Similar campaigns are already happening in EOS (see EOS New York, EOS SoCal, EOSYS, and more). Token holders should realize that the network would stagnate with vote-bribing block producers essentially taxing the system, and eventually they should be voted out. Further, EOS also implements a vote decay system, whereby more recent votes carry more weight. Voters who recast their votes every month will have the full weight of their votes counted, while older votes will slowly decay until they have minimal impact after 2 years. This will encourage participation and also slowly discount those who simply cast votes once. In the case of block producers offering reward sharing, this means that those who simply cast votes one time and expect to earn passive income will gradually have their voting power removed. Other methods, like requiring voters to lock up their tokens for a period of time, could discourage vote bribing, as well.

Because DPoS is a community-driven consensus algorithm, the response to block producer bribes will ultimately be in the hands of the community. One option for the community would be to incorporate a ban on bribes and profit sharing directly into the EOS constitution. The constitution has been described as a "peer-to-peer terms of service agreement or a binding contract among those users who sign it." Every transaction contains a hash of the constitution and expresses the user's endorsement of the contract (the full implications of this digital constitution will be examined in our upcoming full analysis and valuation report on EOS). In Lisk, profit-sharing block producers had to make their intentions very public, so it would be very easy to identify offenders.

Ultimately, we believe that block producers paying for votes is a bad thing. It encourages voting based not on what is best for the network but rather who pays the highest returns. This does not align with the long-term incentives of the network, or even the long-term incentives of token holders. We hope that the communities that emerge around DPoS chains create norms where paying for votes is very much frowned upon, and we plan to encourage such a norm ourselves. We also encourage further research into in-protocol mechanisms, like voting lockups and voting decay, that might discourage such behavior.

ATTACKS AT SCALE

One interesting attack vector that has not yet been observed in practice involves assumptions about what an industrial-scale DPoS blockchain looks like. Larimer has said that EOS is likely to scale in a way such that large data centers act as block producers in order to provide the level of bandwidth and speed the network requires. This outcome may be several years away, but the implications are worth considering.

If block producers are expected to be in dedicated data centers, this limits the number of potential block producers and especially limits the number of entities that could step in to replace block producers that are voted out. Validator churn may be quite low as a result. If there aren't any block producers with sufficient resources to replace block producers that have been voted out, then the network may suffer as a result. Voters would have to decide between punishing a misbehaving block producer and lowering the overall resources of the network.

Importantly, DPoS networks can continue to run with a smaller number of validators until a new peer is ready to join the block producer quorum. This is not ideal, but it at least allows network operations to continue as normal during the transition period. New block producers may not have the same resources as existing producers, but they can campaign for election on the promise to use their block rewards to quickly scale up their efforts.

BLOCK PRODUCERS COLLUDE

In any blockchain system, the threat of block producers colluding is looming. In DPoS, this threat is often presented as especially dangerous; because the number of validators is small, it should theoretically be easy to organize collusion among them. While block producer collusion is obviously not desired, it is important to think about what kind of damage colluding block producers could do, and what the recourse is in the case that they attempt to do so.

In DPoS systems, there are three major attacks that colluding block producers (meaning more than ⅔ of all block producers) could launch:

1. Censorship
2. Changing System Parameters
3. Double Spends

Censorship in the context of DPoS means that a block producer refuses to process valid transactions. If a single block producer censors an individual or entity, it will be futile. Not only will the next block producer validate the transaction in the following block, but the single block producer's censorship attempt will be visible on-chain, and repeated infractions would get the producer voted out. For individual block producers (or even a minority group of block producers), the most damage they could cause is delaying transactions by not processing them in their blocks. Those transactions will still be processed by the honest majority, so it is unlikely that block producers will even attempt censorship since it won't amount to much (they'd simply delay a transaction for a few seconds until the next block producer's turn). And while it is certainly possible to do damage by delaying transactions, block producers are risking their reputation, future income stream, and possibly even arbitration every time they attempt to do so (see page 6 of this paper). An effective implementation of DPoS will likely see block producers quickly voted out for any attempt at censorship.

Another attack that block producers could launch would be an attempt to change the protocol parameters. This could mean many things—changing the constitution, increasing their block rewards, forking out certain stakeholders, and many other options. Luckily, DPoS is designed in such a way that these attacks are not possible without implicit voter approval. In EOS, for example, changes to system parameters have time delays before they are actually incorporated. Changes to the constitution require approval by 17/21 block producers, and they must maintain that approval for 30 consecutive days before the changes take place. If users disagree with the changes, they can vote out those block producers during that time and replace them with producers that don't support the changes. More information on EOS protocol updates can be found here. Ultimately, changes to the system must be endorsed by token holders through, at the very least, passive approval. Parameter changes cannot be enacted without a time delay during which they could be negated. This severely limits the damage that malicious block producers could cause.

Finally, block producers may be able to coordinate double spend attacks by majority collusion, though these situations are highly unlikely in practice. DPoS uses a concept of "last irreversible block" that provides finality once more than two thirds of block producers have built on the same chain. Users who need strong guarantees of finality for their transactions can wait for this confirmation.

DISTRIBUTED DENIAL-OF-SERVICE ATTACK (DDOS)

In most DPoS implementations, block producers are known entities. In some implementations they may even be highly public individuals or organizations whose whereabouts and IP addresses are known. Because the order of block production is fixed during each round, attackers could identify who will be producing blocks at a given time and launch DDoS attacks on the producers.

Such an attack would be difficult to pull off in practice. While an attacker may be able to target a single block producer, targeting several different block producers simultaneously would be unlikely. The network may suffer temporary delays, depending on how many producers the attacker was able to target, but it seems near impossible that an attacker could simultaneously DDoS the majority of the nodes. Further, block producers can campaign on their ability to resist DDoS attacks by using backup servers in other locations and a variety of other means. Finally, if a single block producer (or even a few) were consistently failing to produce blocks because of such an attack, they could be voted out and replaced by backup producers within a single round.

Conclusion

DPoS is an elegant, robust, and most importantly, *practical and proven* solution to the blockchain scalability problem. It also offers solutions for blockchain governance, funding, the nothing at stake problem, and more.

Blockchains based on DPoS architecture achieve high scalability by compromising on the "decentralization of block production" in the context of the scalability trilemma.

Of the three properties comprising the trilemma, two are ends in and of themselves: scalability and safety. These are essential features of any blockchain. Decentralization, on the other hand, is a means to an end. Therefore, it makes sense to compromise on decentralization of block production if the desired ends can still be achieved. The goals of decentralization are censorship resistance, openness, and no single point of failure. We believe that DPoS still achieves all of these goals.

Some but not all aspects of decentralization can be quantified; the number of validators is simply one aspect. There are other factors outside of the number of individual entities that must be considered, and DPoS token voters will have to vote to ensure that the system retains all of the desired effects of decentralization.

There are some drawbacks to DPoS—mostly that it places more requirements on token holders to monitor the health of the network, watch for bad behavior, and decide upon what makes for sufficient

decentralization. In exchange, however, DPoS gives huge performance enhancements and has other very advantageous attributes. We believe that DPoS has a very compelling set of features and tradeoffs that make it a great solution for many types of decentralized applications.

DPoS recognizes that decentralization has a cost—both economically and in terms of performance—and it opts for semi-centralization in exchange for scalability. If DPoS systems can still offer the requisite levels of censorship resistance, permissionless-ness, and trustlessness, then DPoS is better for a huge range of decentralized applications. For certain use cases—absolutely censorship-resistant digital gold, peer-to-peer digital money, etc., a tradeoff in favor of decentralization at the expense of performance may make sense. For the vast majority of applications, scalability is far more pragmatic.

DPoS is not the only consensus algorithm that could succeed at scale. It may not be the right fit for every type of decentralized application, but it is highly likely to have a place in the world. Even if we assume the worst—that 21 known entities control the database, or that a concerted international government effort could result in censorship—DPoS still provides a set of features that may be highly desirable for certain use cases. Businesses want a neutral database that offers scalability, low latency, and maybe even desire some notion of government endorsement. The size of this market could be measured in the trillions.

Still, we estimate that DPoS, *in practice*, will be far more resilient than described above. We look forward to learning from the incredible social experiment that is DPoS.

Thanks to Jesse Walden, Denis Nazarov, Trent McConaghy, Sam Kazemian, Malcolm Mason Rodriguez, Thomas Cox, Ian Grigg, and others for their input and feedback.