

Transactions as Proof-of-Stake

by Daniel Larimer

dlarimer@invictus-innovations.com

November, 28th 2013

Abstract

The concept behind Proof-of-Stake is that a block chain should be secured by those with a financial interest in the chain. This paper will introduce a new approach to Proof-of-Stake that utilizes coin-days-destroyed by every transaction as a substitute for the vast majority of the security currently provided by Proof-of-Work. Unlike prior Proof-of-Stake systems in which only some nodes contribute to the proof-of-stake calculation, we present a new approach to Proof-of-Stake whereby all nodes generating transactions contribute to the security of the network. We speculate that such a network would be immune to known attacks against Bitcoin or Peercoin.

Background

The concept of Proof-of-Stake was first introduced as a means to counter known attacks on the Bitcoin network, primarily the 51% attack.

Existing Proof-of-Stake systems such as Peercoin are based upon 'proof blocks' where the target the miner must meet is inversely related to the coin-days-destroyed. Someone who owns Peercoins must choose to become a Proof-of-Stake miner and commit some of their coins for a period of time to secure the network.

The creators of Peercoin recognized that Proof-of-Stake in this form was insufficient, so they rely upon a hybrid system whereby both Proof-of-Stake and Proof-of-Work are used to secure the network. As the difficulty in the Proof-of-Work increases the block reward decreases which should work to automatically throttle the amount of Proof-of-Work mining.

Despite using Proof-of-Stake, Peercoin still relies on 'mining' with respect to Proof-of-Stake which inherently limits the number of people and the percentage of the money supply available to secure the network via Proof-of-Stake. The incentive given for Proof-of-Stake mining is an average 1% return on your stake. This incentive is currently dwarfed by the 8% inflation paid to the Proof-of-Work miners.

Peercoin isn't the only proof-of-stake system proposed. Others involve various forms of signature blocks, lottery selection of signers proportional to transaction size, etc. None of these ideas have been as thoroughly adopted as Peercoin in the market.

Despite Peercoins success, its application of Proof-of-Stake does not fully solve the problem of double spending or denial of service. Ultimately the network is still secured by Proof-of-Work and it is still possible to mine secret alternative chains that could be used to perform a double spend. An actor such as a government could still acquire enough hashing power to win over 51% of the proof-of-stake blocks and all of the proof-of-work blocks. All Peercoin has achieved is to increase the cost of attacking the network without changing means by which the network can be attacked.

The Purpose of Mining

In most literature regarding crypto-currencies and mining, the focus has been primarily on coin distribution and securing the network from a monopoly of hashing power. Mining rewards are seen as necessary fees used to purchase 'security' and the general theory is that the more mining the better. Existing Proof-of-Stake systems carry over this mentality and pay people who 'mine' on a Proof-of-Stake basis.

For the purposes of this paper, we would like to focus on the more critical role that mining plays that is entirely independent of security or currency creation. This role is to determine in a decentralized manner who will publish the next block. A purely proof-of-stake system must still determine which of 1000's of computers will build and publish the next block. These blocks must come at regular intervals large enough that the network may reach consensus between every block.

There is no cost for producing or broadcasting a block and in a centralized system and a single computer could be assigned the task of publishing blocks every couple of minutes. In a decentralized system this task may be distributed far and wide. Because the cost of producing and broadcasting a block is essentially 0, block rewards worth as little as a couple of cents would be enough to motivate this behavior.

The challenge for a decentralized system is to find a way to limit broadcasting to at most one or two nodes every couple of minutes. We will present such a solution that does not depend upon mining.

The Reason Proof-of-Work Provides Security

If we are to remove Proof-of-Work from a block chain's security model then we must fully understand the nature of the security provided by Proof-of-Work and find a suitable replacement. On network's like Bitcoin the computational power used to find hashes is extremely high. This high level of hash power can be used as a proxy for total combined investment into a particular block chain. A rational individual can generally assume that the largest investment represents majority consensus as to the truth.

Making this assumption opens the door for attackers to abuse the trust placed in this metric and perform several known attacks including: double spend, denial of service, and selfish mining. All of these attacks operate based upon the creation of secret alternative chains. The security provided by the decentralized hash power is via making

counterfeit chains expensive to produce. With this security model there is a direct relationship between the cost of the security and the cost of attacking the network.

The denial of service attack is perhaps the most destructive on the network, while the double spend attack is potentially profitable. Both attacks disrupt the smooth operation of the network and therefore detract from the value of the currency. A government that wished to legalize, but control a crypto-currency could mine at levels that are not profitable for any private organization and thereby achieve a monopoly that allowed them to filter transactions at will.

If we are to replace proof-of-work then the replacement must be capable of preventing counterfeit chains from being easily produced. Lets look at the key to using hash power to produce counterfeit chains. Both double spending and selfish mining depend upon the attacker maintaing a secret alternative block chain that is longer than the public chain. All that is required to stop these attacks is to make the production of secret chains impossible with less than 50% of the money supply.

The denial of service attack does not require a secret block chain. To prevent this attack the selection of the longest block chain must be based upon a metric other than Proof-of-Work.

Transactions as Proof-of-Stake

Every transaction on the network carries with it an implicit Proof-of-Stake in the network. The creator of the transaction wants the network to accept it and the receiver of the transaction is making decisions on whether or not to ship goods based upon whether or not the network has accepted the transaction. It is clear that those behind the transaction have a stake in the health of the network. After all, the network is worthless if transactions cannot be executed as expected. A well functioning network will have thousands of transactions every single block. This represents thousands of stake holders who could be contributing to the security of the network.

A coin-day represents the number of days since a particular coin was last transacted on the network. At any given point in time there exist a limited number coin-days and they accrue in the hands of those who hold large balances for a long period of time. As a result coin-days can be seen as a proxy for stake in the network. Coin-days are destroyed every time there is an transaction involving those coins and therefore cannot be reused.

In order for a 51% attack to be successful in a Proof-of-Work system, the attacker must keep their alternative chain secret. Once they have locked in the profits from their first spend, they can broadcast the longer secret block chain which will invalidate the original transaction. Keeping solved blocks secret is also used in the selfish-mining attack which can be effective with much less than 51% of the hashing power.

In order to prevent this kind of behavior we must make it impractical for miners to maintain secret block chains. If every transaction that is broadcast contains the hash of a recent block and the block chain enforces the rule that the transaction's proof-of-stake can only be credited in block chains that build off of that block then no one will be able to build secret block chains that leverage the coin-days-destroyed of transactions in the public chain.

Now that the transactions are committed to a certain public chain, the best block chain can be measured by total coin-days destroyed rather than total work.

Relative Security of Bitcoin Proof of Work vs Proof-of-Stake

As of November 2013, Bitcoin is valued at about \$1000 per coin with a market capitalization of \$12 billion. Every block pays the miner 25 BTC worth \$25,000 and an efficient market will tend to push the costs of mining a block toward \$25,000. For an attacker to perform a brute-force double spend attack they must produce 7 blocks in secret while the rest of the network produces 6 blocks. This would represent about \$175,000 in electricity consumed in 60 minutes plus the one-time cost of capital associated with consuming that much electricity. Once the capital is owned, the network is effectively compromised and the attacker can interfere with the network continuously.

In a proof-of-stake system on a network worth \$12 billion dollars and producing 50,000 blocks per year with the requirement that all money be moved at a least once per year, an average of \$240,000 per block will be spent destroying a proportional number of coin-days. On average, an attacker would have to accumulate enough coin-days to build a secret chain that destroyed \$1.68 million worth of coin-days over 7 blocks. Once the attack was complete, the attacker will have to wait another year before they could reuse their \$1.68 million in capital to attempt a second attack.

To compare the cost of attacking the two networks depends upon the cost of acquiring and deploying enough ASIC chips to consume \$175,000 in electricity and then compensating for the fact that the ASIC chips can be used continuously and likely generate significant profits that more than offset the costs.

Another factor to Proof-of-Work based security is that for most crypto-currencies the mining rewards fall over time. When this happens the amount spent on securing the network relative to the value of the network falls. For example, assuming no changes in the value of Bitcoin, when the mining reward falls by 50% so does the security of the network. The value of the network must double to \$24 billion to maintain the same level of security it has today after the difficulty adjustment.

Under a proof-of-stake system, security of the network would double along with the networks valuation. It can be seen from this that Proof-of-Stake can offer an order-of-magnitude greater security in the long-run than Proof-of-Work can. Bitcoin is paying over \$1.4 billion dollars per year for Proof-of-Work to secure the network. This cost is

being funneled directly into electric costs and provides no additional value to society. Meanwhile Proof-of-Stake is able to achieve greater security with no cost at all.

Occasional Double Spend Attack

There exists the potential for someone to save a large number of coin-days and then use these coin-days to execute a single double spend attack unless certain protective actions are taken. Before getting into some proposed protective measures, I would like to address the economic consequences of a successful double-spend attack.

First the attacker can only profit by the amount he is able to double-spend of his own funds in an anonymous manner. If the attacker is not anonymous then he would face significant criminal charges for theft if he stole any significant amount of money with his double spend. The attack would have to be well timed because at any time another player could randomly produce above average Proof of Stake and the attempted double-spend would fail. Lastly, if a large double-spend did occur everyone on the network would know and in theory could cooperate to add more proof-of-stake to the weaker chain with your original double-spend.

Unlike Proof-of-Work systems, it is very easy for a few honest nodes belonging to large stake holders to act as guardians of the chain. When they see an attempted double spend attack they can use their own savings to squash the attack in minutes. With a Proof-of-Work system honest nodes are already giving the network everything they have and they have no capacity to correct obvious double spend attempts.

Large chain forks are unusual, especially without any major disruptions to internet infrastructure. As a result, most active and fully connected nodes should not automatically switch to a longer Proof of Stake chain unless it was significantly longer. The attempted double spend would be detected, published, and likely ignored by the overall network.

Lastly, given the cost of a double spend attack, the reality that such an attack would significantly reduce the value of the coins, and the difficulty in performing large double spends anonymously it becomes clear that the attacker would lose more value in depreciation of their assets than they would gain if they were successful.

Anyone with that kind of money would not bother attempting a double-spend over anything trivial. Therefore, I submit that for most ordinary transactions a double spend is very unlikely and the losses from such a double spend attempt would be minimal. Furthermore, the attacker could only perform it once per year.

Unlike Bitcoin where your confirmation time is entirely dependent upon miners finding blocks, someone wishing to accelerate the confirmation time of one transaction can do so by confirming it with some of their own coin-days. Large transactions can also be broken up into multiple parts where the later parts confirm the earlier parts.

Overall the combination of knowing your customers and waiting longer for larger amounts can easily make the potential for a double spend attack essentially 0.

Offline Transactions

Offline Transactions would not necessarily have access to the current head of the block chain at the time they are signed. Therefore, they would be unable to verify the current head block at the time they are signed. The only coin-days that count for the purposes of the transaction are those between the output and head block included in the transaction.

Migrating Transactions from Minority Forks

With existing block chain designs, transactions are relatively independent from the blocks which contain them. In the event that there is a chain fork, transactions from the minority fork can be migrated from the minority chain to the majority chain when the networks reconnect. The only transactions that are invalidated are those transactions which depend upon one of the coinbase transactions from the minority fork. It is for this reason that Bitcoin requires coinbase transactions to mature 120 blocks before they can be spent.

In practice it is rare for there to be a fork longer than a couple of blocks. For all practical purposes the 6 blocks required for a normal confirmation is enough to be secure against a chain fork. The current recommendation for anyone who notices a chain fork is to stop all transactions until minority network can rejoin the majority or people could be the victim of double spend attacks.

Under our approach, transactions that migrate from a minority fork would not contribute to the coin-days-destroyed. This will insure that chain forks do not require individuals to re-issue transactions.

Decentralized Block Generation and Broadcast

At this point it is clear that Proof of Stake is sufficient for identifying the best chain and providing better and cheaper network security. By eliminating mining a decentralized network needs a new approach to reaching a consensus about what block to add next. In a mature network transactions could occur as quickly as 12 per second while network propagation delay could be over 60 seconds. This means that new blocks that destroy more coin-days can be generated at a rate faster than the blocks can propagate. Without mining there is no lottery system able to select just one node every 10 minutes that is granted the authority to include (or not) transactions in the block.

The network must have a way to regulate the rate that transactions and blocks are produced. This can be achieved by adjusting the minimum fees per block that must be paid like Bitcoin would adjust difficulty. These fees will slow transaction volume to the desired data rate and result in delaying the time it takes to build the first candidate

block; however, once the first candidate is built additional candidates will still be produced faster than the first one can propagate.

The next challenge is to decide who gets to broadcast the block when all nodes could generate the new block at the same time. We propose that the owner of the single input that destroys the most coin-days of all transactions in the block is the only one who may broadcast the block. This owner will sign the block header and broadcast the block. If someone else would like to compete to decide the block they must destroy more coin-days which effectively bids up the cost of earning the right to produce the block and in the process increasing the security per block.

To motivate nodes to perform this function in a manner that includes as many transactions as possible with as much proof-of-stake as possible as quickly as possible the broadcaster earns a cut of the transaction fees proportional to their coin-days destroyed.

In the event that there are many transactions that have accumulated for a block that have paid many fees but lack significant coin-days, a large stake-holder can use the opportunity to collect some sizable fees by spending to herself while claiming the largest Proof of Stake input. Before she did so she would want to carefully consider whether any other players will 'out-bid' her and build a better block in order to collect the fees.

We suspect that there is ample opportunity for speciality algorithms to be developed that could earn block creators some revenue for securing the network with their stake. These same algorithms would likely also defend the network against double spend attacks when ever they are observed.

While the rate of transactions might be sustained at 12 per second, the rate of new individual inputs capable of destroying more coin-days than all other inputs in a block will become less and less frequent as the number of inputs increases. This provides a natural unique selection of the node that should sign the block. Considering the likely size of the transaction, they have a vested interest in both the fees earned and in the verification of their transaction. In fact, if they do not sign the block and no one comes along with a larger input, then the transaction will have to be removed from the block so the second-largest input can become the largest and sign the block. In effect, this metric creates an upper limit on how long it takes for someone to produce a block while the transaction fees place a lower limit.

Conclusion

In this paper we have provided a simplified Proof-of-Stake algorithm that allows all users of the network to contribute to the security of the network against attacks. It should be economically infeasible for any actor to maintain a secret block chain that contains more coin-days destroyed than the public transaction ledger. The techniques

presented solve the 51% attack, the selfish-miner attack, and provide protection against double-spending all while requiring no mining at all.

Because Proof-of-Stake can eliminate the need for mining rewards, they also eliminate the need for inflation while also eliminating wasteful consumption of energy for proof-of-work. In a world with an increasing number of block chains, security through proof-of-work becomes fragmented without merged-mining and merged-mining has its own overhead. With this new approach to Proof-of-Stake merged mining is no longer required and an unlimited number of block chains can be supported without compromising the potential security of any individual chain.