

Zano

Andrey N Sabelnikov

October, 2019

v.1.2

Table of content

Hybrid PoS — PoW Consensus Mechanism	3
Classic Proof-of-Stake	3
Zano’s PoS Implementation	3
Design Considerations	4
Cumulative chain weighing	4
PoS power multiplication through PoW power problem	7
Aliases	8
Encrypted Attachments	8
Additional Data Fields	8
Encryption	9
Multi-Sig	10
Indexing	11
Consolidated Transaction Terms	12
Escrow	13
Escrow Proposal	15
Escrow Response	16
Difficulty Adjustment	17
Code Design	18
References	20

Zano leverages the proven and time-tested cryptographic primitives that were first introduced with CryptoNote. Transactions are made both untraceable, and unlinkable by using stealth addresses and ring-signatures. As first implemented in Boolberry, downstream sender privacy is guaranteed by using output flags.

Hybrid PoS — PoW Consensus Mechanism

Zano uses a hybrid PoS — PoW consensus mechanism. This makes double-spend attacks both unfeasible and improbable. PoS was implemented to complement and enhance the security provided by traditional PoW blockchains. With the ability to "rent" hashing power, many PoW coins have been vulnerable to 51% attacks, as an attacker must only have a sufficient amount of hashing power to rewrite the longest chain. This requires a certain investment and the objective is to double spend the coins without regard to the impact of reputational damage, and consequent price decline. However, in a 51% attack on a PoW coin, the attacker has no interest or concern regarding the coin's value beyond the initial attack. Similarly, pure PoS coins have their weaknesses. While there is an inherent disincentive to attack a network that you have a financial interest in, pure PoS protocols face certain pitfalls such as "nothing at stake", "stake grinding", among others. Proposed solutions include slashing, improving validator randomization, staking specific tokens, and forgers, yet there remain a variety of challenges with each of them. Any attack on Zano's hybrid parity protocol, would require an attacker to obtain not only a majority of hashing power, but also a majority of the coins involved in staking. As a result, the risk/reward profile shifts, such that a 51 % attack on a hybrid network has a negligible probability of economic benefit.

Classic Proof-of-Stake

PoS mining is typically implemented such that a random coin owner obtains the right to sign a new block with their secret key. The process of verification compromises full anonymity, because anyone can verify the signature using this owner's public key. Nevertheless, it is possible to preserve untraceability and unlinkability.

Zano's PoS Implementation

Ring signatures allow the transaction creator to provide a set of possible public keys for signature verification, thus keeping their identity indistinguishable from other users. The concept of an anonymous, secure PoS mechanism seemed to be unachievable. The basis of PoS is a so-called *kernel*, which is a data structure that depends on the transaction output and includes:

- Current timestamp with 15-second granularity.
- Key Image, which corresponds to each transaction output. A *keyimage* is comprised of 32 pseudo-random bytes derived from the key in such a way that it is impossible to reconstruct the key, given only its image.
- Stake Modifier. An additional 64 pseudo-random bytes derived from the last PoS and PoW blocks, which disallows any predictability of the *stakemodifier* in the blocks ahead.

During mining, a user is allowed to sign a block, if $Hash(kernel) < CoinAmount * PosTarget$, where *CoinAmount* is the amount of coins in a particular output, and *PosTarget* is an adaptive parameter that works to keep the block creation rate constant.

The PoS miner iterates the timestamp (within the allowed boundaries) for each of their unspent outputs (UTXO) and checks to see if they possess a UTXO that's *keyimage* satisfies the *PoSTarget* formula above. In the event of a "winning" result, they spend this particular output, anonymously, with a ring-signature.

Note: The signature includes the *keyimage* (used in the kernel), but not the key itself, which is why it does not compromise anonymity. The miner signs both the transaction and the block and broadcasts the new block to the network.

Design Considerations

Here are some of the factors that were incorporated in our design decisions:

- a. **Kernel Structure:** Unlike the classic PoS where the kernel includes a direct link to the winning output, an anonymous transaction system (Zano) cannot use this approach, since it compromises privacy. This data was replaced with a *keyimage*, which literally is a hash of an output public key. Thus, the key image contains 32 bytes of pseudo-random data related to the *winning* output.
- b. **Stake Modifier:** The purpose of *stakemodifier* is to make the kernel unpredictable, thereby protecting it from machinations. The *stakemodifier* is recalculated for each new block. The value is the concatenation of the last known PoS block's kernel hash and the last known PoW *block ID*. Both structures are essentially hashes, i.e. unpredictable bytes. It should be noted that a pure PoS (without PoW) network would not be secure under such rules, (typical attacks are described here: <https://download.wpsoftware.net/bitcoin/old-pos.pdf>), but our hybrid model is safe.
- c. **Age of Coins:** The age of coins has no impact on the probability of mining a PoS block. Restricting coins eligibility to participate in staking has two direct impacts. While reducing the coins in circulation can increase scarcity and possibly price, projects that employ these requirements tend to be illiquid, which tends to increase volatility, and makes them less attractive than a free-flowing market with price discovery and liquidity. Additionally, this creates a more equitable environment for users; we consider egalitarianism a condition precedent to acceptance and wide scale adoption, both on PoS mining and PoW mining through our ASIC-resistant hash function, ProgPoWZ.

Cumulative chain sequencing and weighing

The nature of PoW difficulty and PoS difficulty is completely different. PoW difficulty reflects the estimated number of hashing operations needed to reach the target. PoS difficulty reflects the number of coins involved at a particular moment in PoS-mining ("staking"). Each type of difficulty is controlled independently, and to ensure that the network will choose a chain with greater efforts from both sides (PoS and PoW) we need to establish an algorithm which evaluates summary efforts in a proper way.

Sequencing

To mitigate a range of potential attacks, the first step is to ensure that the core consistently prioritizes the most diverse chain. This requires interweaving different types of blocks as closely as possible. Ideally, each Proof of Work (PoW) block should be followed by a Proof of Stake (PoS) block, and next PoW again, so on. To achieve this, a *sequence factor* is introduced in the core. This *sequence factor* is a specific coefficient applied to the current block's difficulty before it is added to the cumulative difficulty of same type of blocks (PoW or PoS) for that chain:

$$CD_i = CD_{i-1} + Sf * d_i^{adj}$$

where:

CD_i — cumulative difficulty that being calculated on height i

CD_{i-1} — cumulative difficulty calculated on height $i - 1$

Sf — “sequence factor”, coefficient which set penalty for a block as:

$$Sf = 0.75^{n-1}$$

where n is "sequence number", number of blocks of the same type goes in a row, and 0.75 is sequence decrease rate parameter.

If the blocks form a perfect sequence where the block types alternate with each new block, then the coefficient will always be 1. If two blocks go in a row, then the last one's difficulty will be reduced by multiplying on 0.75, if 3 goes in a row - second would be cut on 0.75 and third by 0.5625, and so on.

Note, that *sequence factor* does not affect difficulty target in any way, it's only applied to the value that is added to cumulative difficulty for the chain, while actual difficulty target is controlled only by difficulty adjustment algo.

Weighing

The problem is that the balance between PoW and PoS difficulty is constantly changing and primitive algorithms don't provide correct protection against manipulation.

The idea is to put into analysis only a work which is made on the distance of split and for evaluation of PoS work to take the ratio of PoW difficulty and PoS difficulty at the split point between two competing chains. With this approach adjusted cumulative difficulty exists only for a given sub-chain started from the point of the split. Let's define the adjustment coefficient at point of split K' as:

$$K' = \frac{d_i^{pow}}{d_i^{pos}}$$

where:

d_i^{pow} — current PoW difficulty at the point of the split

d_i^{pos} — current PoS difficulty at the point of the split

Then cumulative difficulty for given sub-chain S split from the main chain after block B' is:

$$CD_s = CD_s^{pow} + K' \cdot CD_s^{pos}$$

where $CD_s^{pow} = \sum_i d_i^{pow}$ and $CD_s^{pos} = \sum_j d_j^{pos}$

Illustration provided down below:




Figure 1

For the given example cumulative difficulties for both chains would be calculated as the following:

$$\begin{aligned}
 CD_A^{pow} &= 1000 + 1000 + 1000 + 1000 = 4000 \\
 CD_A^{pos} &= 100000 + 100000 + 100000 + 100000 = 400000 \\
 CD_A^{adj} &= 4000 + 400000 \cdot K' = 4000 + 400000 \cdot \left(\frac{1000}{100000}\right) = 8000
 \end{aligned}$$

For chain “B” against “A” we are getting:

$$\begin{aligned}
 CD_B^{pow} &= 1000 + 1010 + 1010 + 1010 = 4030 \\
 CD_B^{pos} &= 100000 + 100 + 10 + 10 = 100120 \\
 CD_B^{adj} &= 4030 + 100120 \cdot K' = 4000 + 100120 \cdot \left(\frac{1000}{100000}\right) = 5001
 \end{aligned}$$

CD_A^{adj} now properly reflects PoS contribution and looks way greater than CD_B^{adj} . Let’s try to see how much resources would be needed to commit double spend attack with CD_s formula. If we do analysis based on numbers from the example above, then we can describe it with this equation:

$$8000 = CD_{PoS} + \frac{1}{100} \cdot CD_{PoW}$$

With the graph presented above, you can see the red zone which represents an area of potential double spends and the green area is safe. In this graph we assumed that amount of coins used in PoS mining in chain A was only 40% of total coins in circulation, just to show that PoS “hash power” is a limited resource.




Figure 2

Figure 2 shows that if a potential attacker has nearly 100% of one of the types of resources (PoW or PoS) and a minority of the other - then he can commit double spend attack by creating an alternative chain with superior CD_s . In order to make the formula more resistant to the manipulation of difficulty through only one of the resources type (PoW or PoS) to the detriment of the other, let's introduce two balancing coefficients for the PoW and PoS and respectively, which will take into account the ratio of cumulative difficulties on two alternative chains. In this case, the total cumulative difficulty can only be calculated relative to another alternative chain, but cannot exist by itself.

$$K' = \frac{d_i^{pow}}{d_i^{pos}} ;$$

$$K_{PoW(X \rightarrow Y)} = \frac{CD_{PoW(X)}}{CD_{PoW(Y)}}$$

$$K_{PoS(X \rightarrow Y)} = \frac{CD_{PoS(X)}}{CD_{PoS(Y)}}$$

$$\begin{cases} CD_{A \rightarrow B} = K_{PoW(A \rightarrow B)} \cdot K_{PoS(A \rightarrow B)} \cdot (CD_{PoW(A)} + CD_{PoS(A)} \cdot K') \\ CD_{B \rightarrow A} = K_{PoW(B \rightarrow A)} \cdot K_{PoS(B \rightarrow A)} \cdot (CD_{PoW(B)} + CD_{PoS(B)} \cdot K') \end{cases}$$

Now, to see resulting curve which describe boundary of potential "double spend" attack, based on the same numbers as in the example above, we can use both sides of the system of equations in the following inequality:

$$K_{PoW(A \rightarrow B)} \cdot K_{PoS(A \rightarrow B)} \cdot (CD_{PoW(A)} + CD_{PoS(A)} \cdot K') < K_{PoW(B \rightarrow A)} \cdot K_{PoS(B \rightarrow A)} \cdot (CD_{PoW(B)} + CD_{PoS(B)} \cdot K')$$

$$K_{PoW(A \rightarrow B)} \cdot K_{PoS(A \rightarrow B)} \cdot \left(4000 + 400000 \cdot \frac{1}{100} \right) < K_{PoW(B \rightarrow A)} \cdot K_{PoS(B \rightarrow A)} \cdot \left(CD_{PoW(B)} + CD_{PoS(B)} \cdot \frac{1}{100} \right)$$




Figure 3

PoS power multiplication through PoW power problem

Let's assume PoW miner has the majority of hashrate and can use his/her hashpower to let him/her multiply his chances to find PoS block by not announcing all mined PoW blocks but announcing only those which let him/her win PoS blocks. In this case having 100% hashrate (twice more than needed for performing double spend in pure PoW system) let him do double spend by having only 25% of PoS power, with 200% needed only 12.5% of PoS power and so on. Nevertheless, such manipulations do not create the possibility to go beyond the curve defined by the formula of relative complexity:




Figure 4

In our opinion, this approach creates a stable equilibrium between PoW and PoS, ensuring decentralization and high level of resistance to 51% attacks.

Aliases

Each Zano user can register with an alias, for example: @easytouse, a human-readable name associated with a payment address and text comment, which is stored in the blockchain. This alias provides a short, easy-to-remember name rather than a long and confusing string of random characters. Blockchain storage ensures that aliases are protected from being altered or commandeered.

An alias registration is a special kind of transaction that includes a record of an alias assignment. The system core validates the record and checks the availability of the name. Once validation is successful, the transaction reaches the blockchain and a new record is created in the alias database.

Zano users can easily send transactions to an alias: their wallet automatically checks whether the name is registered in the blockchain, and then obtains the associated public keys to which the transaction will be sent. In the event of a missing alias, the wallet will generate an error report and no funds will be sent or lost.

Aliases can be used for more than just Zano transactions. Think of them as a decentralized address book with universal IDs that can be used for various services based on the Zano platform.

To reduce possibility of phishing we set limitations on alias registrations. Users can use any combination of the lower-case Latin letters a-z and the Arabic numerals 0-9. Additionally, there is a length minimum of 6 characters, and a maximum length of 12.

Encrypted Attachments

Additional Data Fields

Zano allows for including arbitrary data in transactions; this data can be stored in the blockchain permanently or removed after passing checkpoints, similar to the pruning mechanism that was first implemented in Boolberry. The purpose of pruning was initially to manage blockchain bloat, but we have extended this to include arbitrary data that similar to ring signatures, have a diminished relevance over time. This gives developers the power to build a variety of applications based on the Zano blockchain platform.




Figure 5

For this purpose, we implemented two types of transaction data fields:

- **Transaction Extra:** This field contains data that is included in the transaction prefix. The prefix is the first part of any transaction; it stores essential information about the financial transaction. Data in the prefix is completely hashed during the generation/validation of transaction signatures, which is why it cannot be later removed from the blockchain.
- **Transaction Attachments:** This field contains additional data that does not need to be kept in the blockchain permanently. A couple of examples are transactional data (including escrow proposals), comments on transactions, or random data used by third-party services. Old transaction attachments are deleted (voluntarily removed from the local hard drive) after every checkpoint has been passed, meaning the core will not check them again. To protect data from being altered, we have put the transaction attachment hash into the transaction extra field, so that the deleted data can always be verified, even after passing checkpoints.

Encryption

For many tasks, such as text comments in transactions, special fields in escrow proposals, etc., we need to encrypt data included in transactions (attachments/extra). However, all keys should be retrievable, for both the sender and recipient, in the event of hard drive failure or recovery of a wallet from a seed phrase (brain wallet). Essentially, they should be able to be derived from the main wallet key and data in the blockchain.




Figure 6

In each transaction we already have a common secret element for both participants, known as a derivation. This is part of an ephemeral key, which is produced using the Diffie-Hellman protocol for unlinkable payments [2][4]. All transferred data is encrypted via a block cipher ChaCha using the key obtained from the derivation. (Figure 6)

The Challenge is that while the recipient can always calculate a derivation, given only an incoming transaction, the sender cannot. Usually, the sender's first step when preparing a transaction is to generate a random number (without saving it) to produce the derivation. The next time the sender sees this transaction in the blockchain, he cannot recalculate the same derivation without that randomly generated number.

To overcome this, we allow senders to store the derivation in the transaction and encrypt it using their own secret key. As a result, both sender and recipient can reconstruct the derivation and decrypt their shared data from the transaction, and no private information is ever published.

Multi-Sig

Multi-signature, or simply *multisig*, is a commonly used term for special types of transactions in which money can only be redeemed jointly by several recipients. M of N multisig means that a transaction (namely, the identifier of its specific output) contains N keys, which belong to the multisig users, and at least M of these users (key owners) must cooperate by providing their signatures in order for the cryptocurrency to be released.

Zanos multisig feature is extremely useful for user security. For example, this feature allows a user to hold their coins in a 2-of-2 multisig, where both keys belong to this user, but are stored on different devices. In this case, the user obtains a two-factor authentication mechanism, which protects their wallet from a single key compromise.

Another popular feature of Zano is the escrow service, which will be described later in "Escrow" section.




Figure 7

Technically speaking, a multisig output differs only slightly from a standard output: it contains N keys (instead of one key) and an M number. Developers' work with multisig usually relates to 1) the consolidated creation of transactions, or how to make the process of joint creation and signing more user friendly; and 2) the development of a key and output indexing structure.

Indexing

In the interest of simplicity, we have introduced a new method of multi-signature output identification. Currently, existing cryptocurrencies employ one of the two most commonly used identification methods:

- Identification by the transaction hash plus the order number of an output inside the transaction: Here we would need to maintain a global index of all transactions. The downside to this approach is that if a transaction is changed (new inputs or outputs are added), its hash also changes. Therefore, we cannot refer to an output (using its identifier) until the transaction has been finalized, which disallows any subsequent changes.
- The global index of outputs method (CryptoNote, [2]): The amount of the output plus its order number among other outputs with the same amount. This method optimizes the support of the ring signature feature, which ensures the untraceability of payments. The challenge here is that we do not know the output order number until

the transaction reaches the blockchain (another transaction can be confirmed prior to it, thus altering the global index of outputs).




Figure 8

To ensure the system supports the escrow service, (see next section) we must be able to make a reference to each multisig output before its transaction is fused into the blockchain and even before its transaction is in its final form. Specifically, we want to include a template for the future transaction B as an attachment to transaction A, so that the embedded transaction B refers to transaction A "from within".

As a unique multi-sig output *ID* we use:

$$ID = H(txPubKey||Out)$$

where:

- *ID* is a multisig output identifier;
- *H* is a cryptographic hash function;
- *txPubKey* is a one-time transaction public key (see section 1);
- *Out* is the output body;
- "||" is the concatenation operation

Due to the random nature of txPubKey and hash mapping properties, we get non-repeating identifiers for multisig outputs (Figure 8). In order to prevent fraudulent operations, where attackers generate identical public keys and outputs to create identical IDs, the program's core tracks the uniqueness of each multisig ID.

Consolidated Transaction Terms

One of the most important conditions for enabling the escrow service on the platform is to provide the ability for two or more users to jointly create a transaction. This way, payments are atomic (all or nothing) and each participant's funds are protected from theft, being frozen, or intentional money burn.




Figure 9 Creation of Consolidated Transaction

Figure 9 illustrates the basic principles in making a consolidated transaction. For example, Alice and Bob agree to make a consolidated transaction, which sends money from each of them to a 2-of-2 multisig (meaning that only together, can Alice and Bob release the funds).

- a. Alice prepares the main multisig output, adds all relevant data (her inputs, outputs for balance of payments, attachments, etc.) and signs the transaction. She also uses a special flag TX FLAG SIGNATURE MODE SEPARATE, which will activate the core's special advanced mode of transaction signature protocol verification.

- b. At this point the transaction is incomplete: the total output sum is greater than the total input sum. Bob needs to fund his portion of the transaction, so Alice sends a "transaction template" (transaction information) to Bob as a special attachment.
- c. If Bob accepts this proposal (transaction template), then Bob funds his portion of the transaction and signs it. Once the transaction template is fully funded, and signed, the transaction is considered valid and can be broadcast to the network and subsequently included in the blockchain.

Security notes:

- To protect Alice's attachment's from being altered, the data hash is stored in the corresponding transaction input, which is sealed by and with Alice's signature.
- Until the transaction has been confirmed by the network, none of the participant's funds are considered spent. Once the transaction proposal is created, the funds used in this transaction are temporarily locked in the wallet, this mitigates potential conflicts regarding other coins being spent from the same wallet. If the proposal (transaction template) is rejected by the counterparty, or expires, the coins will be unlocked in the wallet from which they were sent, without any further action.
- If Bob (or anyone else) were to attempt to interfere with or obstruct Alice (for example, locking Alice's money by publishing the transaction with incorrect data), they would need to make the transaction semantically correct, which would require adding their own coins to the inputs. This requirement creates a disincentive that detours malicious behavior.

Escrow

Escrow, like its name, is a mechanism that was designed to facilitate secure, anonymous, payments between counter-parties. Traditionally, the term "escrow" refers to an independent "trusted" third party that is tasked with acting as an intermediary to oversee and provide mutual assurances to each party that their agreement is executed in a manner that is satisfactory to all parties. Their fiduciary duty includes carrying out the agreement as it was intended, but also allows for modifications with mutual consent. Finally, the "escrow" is responsible for disposition, transfer, or distribution as a final step in the process.

Trust and reputation create value for participants, and justify fees commensurate with such escrows, along with warranties, both expressed and implied. Certain aspects of an escrow are difficult to augment or mechanize. However, we have developed a system that is intended to let participants chose, from a variety of variables, the structure that best suits them and the nature of their respective transaction.

Zano provides the framework for a secure and private transaction without the need for a trusted third party. We anticipate market participants are best suited to choose how the framework is applied to facilitate their objectives relative to consummating a transaction. Our Escrow system, as proposed, will require participants to make additional deposits, which they will forfeit if there is any attempt to act maliciously, or in a way that is contemptuous toward their counterparty. In the absence of a rating or feedback system, similar to eBay or other escrow types, we must rely on financial incentives to augment credibility and integrity in the absence of a trusted third party.

The main idea is as follows:

Alice wants to buy a laptop from Bob for 100 ZANO:




Figure 10. Creation of Consolidating-Transaction

- Alice wants to buy an item from Bob for 100 ZANO.
- Alice prepares an escrow transaction template that will transfer a total of 300 ZANO from both participants on a 2-of-2 multisig. The first key belongs to Alice and the second key to Bob.
- Alice can include details, including delivery instructions, messages etc., all of which will be encrypted and only retrievable by Bob. This data type was described above as transaction attachments and while the hash of the attachment will be verifiable indefinitely on the blockchain, the actual attachment data will fall off after passing checkpoints to manage blockchain bloat as this information at some point in the future loses its relevance. For example, the shipping details including name, address etc. on a package you received a year ago is no longer relevant.
- Alice adds her inputs with 200 ZANO and sends the proposal transaction template (which is not yet semantically correct) to Bob for his consideration.
- Upon acceptance, Bob adds his input for 100 ZANO and broadcasts the transaction (which now becomes semantically correct) to the network. At this point the funds involved in the transaction are locked.
- Bob, pursuant to their agreement, sends the item(s) to Alice and waits for a response.
- If Alice is satisfied with the order, both users sign a new transaction, which unlocks the money and sends 100 ZANO back to Alice and 200 ZANO to Bob.
- Other scenarios may include: Alice sends the goods back to Bob, both get their money back, or Alice keeps defective goods, asks for compensation and Alice and Bob share the money 50/50, and so on.

- i. If anyone breaks the escrow (Bob does not send any goods or Alice receives the shipment, but refuses to sign the money unlocking transaction), both will lose more than they would receive (Bob loses 100 ZANO or Alice pays twice for the goods). Therefore, it is in everyone's best interest to play fair.

Note: The size of locked deposits is determined by mutual agreement between the merchant and the customer.

Escrow Proposal

The main steps here are an escrow transaction. The buyer sends an incomplete "transaction template" to the merchant / counterparty (we refer to this as the "proposal"). They cannot do so directly, since establishing direct connections tends to be inconvenient and, more importantly, destroys anonymity. Additionally, these transactions cannot be sent to the network using broadcast packets, as this would flood the network with excessive traffic. This is where "transaction attachments" come in (see section "Encrypted Attachments").




Figure 11

The buyer prepares a "proposal" also referred to as a "transaction template" (Figure 12), packs it, and stores it in the attachment field of a special "transport/carrier" transaction (the transferred amount of money in such a transaction can be as low as the buyer wants, or even without any outputs). Additionally, the buyer encrypts the attachment using the mechanism described in section "" in this paper, to preserve privacy. Once this is complete, the transaction enters the

blockchain and the merchant can then decrypt the attachment and proceed with the escrow transaction which remains in the first transaction's input.

IMPORTANT: Attachment data is not stored in the blockchain forever but is removed after passing each checkpoint. The only "excess data" is the 32-byte hash of the attachments. At the same time, the escrow attachments can be arbitrarily large (the only limit is the systems current maximum transaction size) and can contain any necessary information: order ID, delivery address, deposit sums, payment IDs, etc.




Figure 12

Escrow Response

The last steps are simple. Once the merchant receives the escrow template, they can decline it (no money will be locked, no additional actions are needed) or accept it (by completing the escrow transaction and shipping the order to the

buyer). Prior to sending purchased goods to the buyer, the seller must prepare a "release transaction template" which will send 100 ZANO back to Alice and 200 ZANO to Bob, then encrypt it with Alice's key and store it in the attachment field of the escrow transaction (see figure 13).

The buyer receives the template of the second transaction and then gives their signature upon delivery of the goods purchased from the merchant to release the money. Altogether, we employ a three-round protocol with no direct connection between participants.




Figure 13

Difficulty Adjustment

Some miners were doing one of the typical "greedy mining" strategies - abruptly raised hashrate, keep mining until difficulty get adjusted, and then dropped all hashrate, leaving a network stuck with one blocks/hour or even worse. For miners/pools this strategy is profitable because it let them get "cheap" blocks(mined with low difficulty), but this is definitely a problem for cryptocurrency - in such situation transactions confirmations takes hours, and other fair miners are quitting because mining on high difficulty is getting unprofitable.

To analyze this type of attacks we reviewed Boolberry blocks history and picked up a representative period, where hashrate was raised and dropped in this typical manner. Then, using difficulty associated with the blocks and their timestamps estimated hashrate was derived. Here what I've got:




Figure 14

That was a first week of June 2014, and as it seen on marked areas - there was a hard time for the network, blocks were coming once per hour or even worse.

Then, using this hashrate numbers, associated with time, became possible to calculate estimate blocks flow with timestamps and difficulty values progression, and this simulation had been run for bunch different variations of adjustment functions.

At this graph showed original cryptonote adjustment function(the blue one), couple variations with “adjustment window”, and currently preferred version(green one):




Figure 15

It calculated as combination of original adjustment functions for three different periods:

$$D = (D720(\text{historical_data}) + D240(\text{historical_data})) / 2$$

Where:

D720 - function which calculates difficulty based on last 720 blocks window with 60 / 60 cuts

D240 - function which calculates difficulty based on last 240 blocks window with 7 / 5 cuts

This approach leads to prioritisation for a newest blocks influence to difficulty adjustment algorithm over older blocks. We do this prioritization to achieve faster adaptation for a hashrate drops and spikes, but keep the older blocks in the range to get smooth behavior and protection from timestamps manipulation. For PoS adjustment also we going to use only *D240* and *D720*.

Code Design

Zano's methodologies are aimed at maximizing reliability, while maintaining a flexible and secure architecture.

- **Component-based Modular Structure:** There are three main software components: the daemon, miner, and the wallet. The daemon connects to the network, validates and exchanges new blocks and transactions, and maintains the local blockchain database. The daemon is the main component, sometimes referred to as the node. The miner is a subprogram, that performs mining tasks and produces new blocks. The wallet stores users' private keys and provides an interface for creating transactions. The wallet connects to the daemon independently, which provides the following advantages:
 - The daemon is the only component connected to the network, but it does not store private keys. This means it is easier to ensure security since keys are separated from the network.
 - A single daemon can serve two or more wallets. This is particularly important for enterprise solutions, such as web-wallets or payment processors that work with multiple keys at the same time.
 - Any modifications affecting the mining algorithm or user wallet features (GUI, keys encryptions) do not affect the daemon's operation.
- **Forward and Backward Compatibility:** When updating the software and enabling new features, it is of great importance to maintain compatibility between old client software versions and new versions. Usually when a "new client" sends a modified (new) command or message to an "old client," the old client software does not recognize it. Zano's protocol allows older versions of client software to identify portions of data in the new versions of messages that correspond to previous versions of messages (older protocol). Obviously, the client software of older versions ignores the new data in new versions of messages, because it cannot handle such data. Nevertheless, this feature helps to avoid the typical problems of evolving the protocol and adding new product features.
- **Asynchronous Core Architecture:** Most Cryptonote projects use an interlocked multithreading model: each incoming network or RPC call, is handled by the core, and passes through a mutex lock, which only allows one thread to work with the core simultaneously. In the event of multiple requests from the network or RPC calls (high transaction flow or high-load service requests) all requests will be handled consecutively, one-by-one, which significantly limits network throughput. Our new core is a fully asynchronous architecture. This allows multiple network / RPC requests to be handled in parallel. For example: Even if the atomic operation of "adding a block" to the core is processing, the core still can handle parallel RPC/network requests simultaneously, and without limitations. This was accomplished by a deep refactoring and improvements to the multithreading model of the core. Another strength of the new asynchronous core is that the transaction pool is no longer blocked during the operation of adding a new block to the core. In classic CryptoNote implementation, transaction pool is locked for the entire "Add New Block" operation (which can take several seconds, in case of big blocks). During times of high transaction flow, while the core is processing larger size blocks, validation of these blocks can cause the transaction pool to be blocked for up to several seconds. During this seemingly short period, connections from other nodes while relaying new transactions would be blocked, which can cause connection time-outs, which are damaging to network connectivity, throughput, and impact network propagation. While in this sub-optimal state, transfers can be rejected, and chain splits / orphan blocks become more likely. We were faced with this exact situation while conducting stress tests and we were able to solve this issue by making transaction pool lock independently, and for a particularly fast, individual, operation, thereby eliminating the bottle neck associated with the "Add New Block" operation. To illustrate this how significant of an improvement this represents, we will show screenshots from our diagnostics tool, which monitors connectivity of the network: Colored dots are network nodes and lines between them represent connections. The color and width of the lines reflect connection age:
 - Newer connections are illustrated with a thin red line.

- Older connections are shown with thin blue line.
- Stable long-term connections are highlighted by bold blue lines.




Figure 14

- On the left side you can see the result of stress-testing one of our old test networks with, literally, hundreds of thousands of transactions. Most connections are timing out and reconnecting, and the network is suffering from these time outs and the bad relaying of transactions / blocks. On the right side is an example that illustrates a network in perfect condition all connections are stable, ongoing transactions and blocks are being relayed without any glitches or gaps.
- Note regarding Core tests: The risk of a "bug" in the currency core has the potential to be devastating, and with ongoing development, even the most cautious approach introduces the possibility of human error. With this fact in mind, we expended an extraordinary amount of effort on rigorous "core tests" –large set of project specific unit-tests, which cover most of the currency core and wallet codebase. Almost every currency rule that is validated by the core has a corresponding core test

References

- [1] Satoshi Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," URL: <https://bitcoin.org/bitcoin.pdf>
- [2] Nicolas van Saberhagen, "CryptoNote v 20," URL: [https:// cryptonote.org/whitepaper.pdf](https://cryptonote.org/whitepaper.pdf), October 17, 2013
- [3] Nicolas van Saberhagen, Johannes Meier, Antonio M Juarez, "CryptoNote Signatures," URL: <https://cryptonote.org/cns/ cns001.txt>, December 2011
- [4] Nicolas van Saberhagen, Seigen, Johannes Meier, Richard Lem, "CryptoNote One-Time Keys," URL: <https://cryptonote.org/cns/cns006.txt>, November 2012
- [5] Sunny King, Scott Nadal, "PPCoin: Peer-to-Peer Crypto-Currency with Proof-of-Stake," URL: <https://peercoin.net/assets/paper/peercoin-paper.pdf>, August 19th, 2012
- [6] URL: <https://mintpeercoinexploremet/chart>
- [7] Whitfield Diffie and Martin E. Hellman, "New directions in cryptography," URL: <https://eestanfordedu/~hellman/publications/24.pdf>, November 1976
- [8] Daniel J Bernstein, "ChaCha, a variant of Salsa20," URL: [https://crypto/chacha/chacha-20080128.pdf](https://crypto.chacha/chacha-20080128.pdf)
- [9] Andrew Poelstra, "Distributed Consensus from Proof of Stake is Impossible," URL: <https://download.wpsoftware.net/bitcoin/old-pos.pdf>, May 28, 2014
- [10] https://boolberry.com/files/Boolberry_Solves_CryptoNote_Flaws.pdf
- [11] https://boolberry.com/files/Block_Chain_Based_Proof_of_Work.pdf

