

# *Autolykos*: The Ergo Platform PoW Puzzle

Alexander Chepurnoy, Vasily Kharin, Dmitry Meshkov

April 6, 2019  
v1.0

## 1 Introduction

Security of Proof-of-Work blockchains relies on multiple miners trying to produce new blocks by participating in PoW puzzle lottery, and the network is secure if the majority of them are honest. However, the reality becomes much more complicated than the original one-CPU-one-vote idea from the Bitcoin whitepaper[1].

The first threat to decentralization came from mining pools – miners tend to unite in mining pools. Regardless of the PoW algorithm number of pools controlling more than 50% of computational power is usually quite small: 4 pools in Bitcoin, 2 in Ethereum, 3 in ZCash, etc. This problem led to the notion of non-outsourcable puzzles [2, 3]. These are the puzzles constructed in such a way that if a mining pool outsources the puzzle to a miner, miner can recover pool’s private key and steal the reward with a non-negligible probability. However the existing solutions either have too large solution size (kilobyte is already on the edge of acceptability for distributed ledgers) or very specific and can not be modified or extended in any way without breaking non-outsourcability.

The second threat to cryptocurrencies decentralization is that ASIC-equipped miners are able to find PoW solutions orders of magnitude faster and more efficiently than miners equipped with the commodity hardware. In order to reduce the disparity between the ASICs and regular hardware, memory-bound computations were proposed in [4]. The most interesting practical examples are two asymmetric memory-hard PoW schemes which require significantly less memory to verify a solution than to find it [5, 6]. Despite the fact that ASICs already exist for both of them [7, 8], they remain the only asymmetric memory-hard PoW algorithms in use.

In this paper we propose *Autolykos* — new asymmetric memory-hard non-outsourcable PoW puzzle. In Section 2 we provide a full specification of *Autolykos*, while in Section 3 we discuss its properties. Few auxiliary algorithms are placed in Appendix.

## 2 Ergo PoW puzzle

The proposed scheme requires following components:

1. Cyclic group  $\mathbb{G}$  of prime order  $q$  with fixed generator  $g$  and identity element  $e$ . Secp256k1 elliptic curve is used for this purposes.
2. Number of elements  $k$  required in the solution. Value  $k = 32$  is used in implementation.
3. Number  $N$  of elements in the list  $R \subset \mathbb{Z}/q\mathbb{Z}$  to be stored in miner's memory. Value  $N = 2^{26}$  is used in implementation.
4. Hash function  $H$  which returns the values in  $\mathbb{Z}/q\mathbb{Z}$ . Particular implementation is based on Blake2b256 and is described in Alg.3.
5. Hash function  $genIndexes$  which returns a list of numbers from  $0 \dots (N - 1)$  of size  $k$ . It is based on Blake2b256 and is described in Alg.4.
6. Target interval parameter  $b$ , that is recalculated via difficulty adjustment rules.
7. Constant message  $M = [0, \dots, 1023].flatMap(i => Longs.toByteArray(i))$  that is used to enlarge message size and increase elements calculation time.

*Autolykos* is based on one list  $k$ -sum problem: miner should find  $k$  elements from the pre-defined list  $R$  of size  $N$ , such that  $\sum_{j \in J} r_j - sk = d$  is in the interval  $\{-b, \dots, 0, \dots, b \bmod q\}$ . In addition, we require set of element indexes  $J$  to be obtained by one-way pseudo-random function  $genIndexes$ . This prevents optimizations as soon as it is hard to find such a seed, that  $genIndexes(seed)$  returns the desired indexes.

Thus we assume that the only option for miner is to use the simple brute-force algorithm 1 to create a valid block.

---

**Algorithm 1** Block mining

---

```
1: Input: upcoming block header hash  $m$ , key pair  $pk = g^{sk}$ 
2: Generate randomly a new key pair  $w = g^x$ 
3: Calculate  $r_{i \in [0, N)} = H(j || M || pk || m || w)$ 
4: while true do
5:    $nonce \leftarrow \text{rand}$ 
6:    $J := genIndexes(m || nonce)$ 
7:    $d := \sum_{j \in J} r_j \cdot x - sk \bmod q$ 
8:   if  $d < b$  then
9:     return  $(m, pk, w, nonce, d)$ 
10:  end if
11: end while
```

---

Note that although the mining process utilizes private keys, solution itself only contains public keys. Solution verification can be performed by Alg. 2.

---

**Algorithm 2** Solution verification

---

- 1: **Input:**  $m, pk, w, nonce, d$
  - 2: require  $d < b$
  - 3: require  $pk, w \in \mathbb{G}$  and  $pk, w \neq e$
  - 4:  $J := genIndexes(m||nonce)$
  - 5:  $f := \sum_{j \in J} H(j||M||pk||m||w)$
  - 6: require  $w^f = g^d pk$
- 

### 3 Discussion

First, notice that in Algorithm 1 we refer to construction  $f(m, nonce, w, pk) = \sum_{j \in genIndexes(m||nonce)} H(j||M||pk||w)$  as a hash function. Public key plays a role of commitment. Therefore, the pair  $(pk, d)$  is a Schnorr signature with a public key  $w$  over the message  $(m, nonce)$  with a hash function  $f$ . If one denotes  $e$  the corresponding value of  $f$ , and pass to more common notations:  $e = f(m, nonce, w, w^e g^{-d})$ . The puzzle consists in trying different nonces and keys in order for signature to satisfy  $d \in \{-b, \dots, 0, \dots, b\}$ . Security follows from the security of Schnorr signatures, and outsourcing the puzzle is equivalent to outsourcing the signature (or parts of signature creation routine). The only difference from conventional setup is the design of function  $f$ . It must be constructed in such a way that efficient massive evaluations with different nonces require allocating large amount of memory (benefitting from data reuse), whereas single evaluation on verifier's side can be done "on fly".

To achieve this, algorithm 1 requires to keep the whole list  $R$  during the main loop. Every pre-calculated hash occupies 32 bytes, so the whole list of  $N$  elements occupies  $N \cdot 32 = 2Gb$  of memory. For sure, a miner can recalculate these elements "on fly" during the main loop and thus reduce memory requirements. However in such a case the number of calls of  $H$  will significantly grow up (e.g. assuming GPU hashrate  $G = 2^{30} H/s$  [9] and block interval  $t = 120 s$ , every element will be used  $(G/N) \cdot k \cdot t = 3 \cdot 10^4$  times on average.) reducing miner's efficiency and profit.

While list  $R$  is quite big, it's filling consumes quite a lot of time: our initial implementation [10] consumed 25 seconds on Nvidia GTX 1070 to fill list  $R$ . This part, however, may be sufficiently optimized if miner in addition stores a list of unfinalized hashes  $u_{i \in [0, N]} = H(i||M||pk)$  in memory, consuming 5 more Gigabytes of it. In such a case this work to calculate unfinalized hashes should be done only once during mining initialization while finalizing them and filling the list  $R$  for the new header will only consume few milliseconds (for about 50 ms on Nvidia GTX 1070).

The protocol is quite efficient in terms of solution size and verification time: it consists of 2 public keys of size 32 bytes, number  $d$  that is at most 32 bytes (but contains a lot of leading zeros in case of the small target  $b$ ) and an 8-bytes long nonce. Header verification requires verifier to calculate 1 *genIndexes* hash,  $k$  hashes  $H$  and perform two exponentiations in the group. Reference Scala

implementation [11] allows verifying block header in 2 milliseconds on Intel Core i5-7200U, 2.5GHz.

## References

- [1] S. Nakamoto, “Bitcoin: A peer-to-peer electronic cash system,” 2008.
- [2] A. Miller, A. Kosba, J. Katz, and E. Shi, “Nonoutsourcable scratch-off puzzles to discourage bitcoin mining coalitions,” in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2015, pp. 680–691.
- [3] I. E. P. Daian, E. G. Sirer, and A. Juels, “Piecework: Generalized outsourcing control for proofs of work,” in *BITCOIN Workshop*, 2017.
- [4] C. Dwork, A. Goldberg, and M. Naor, “On memory-bound functions for fighting spam,” in *Annual International Cryptology Conference*. Springer, 2003, pp. 426–444.
- [5] A. Biryukov and D. Khovratovich, “Equihash: Asymmetric proof-of-work based on the generalized birthday problem,” *Ledger*, vol. 2, pp. 1–30, 2017.
- [6] Ethash. [Online]. Available: <https://github.com/ethereum/wiki/wiki/Ethash/6e97c9cea49605264c6f4d1dc9e1939b1f89a5a3>
- [7] Bitmain confirms release of first ethereum asic miners. [Online]. Available: <https://www.coindesk.com/bitmain-confirms-release-first-ever-ethereum-asic-miners>
- [8] Bitmain's latest crypto asic can mine zcash. [Online]. Available: <https://www.coindesk.com/bitmain-latest-crypto-asic-can-mine-zcash>
- [9] Non-specialized hardware comparison. [Online]. Available: [https://en.bitcoin.it/wiki/Non-specialized\\_hardware\\_comparison](https://en.bitcoin.it/wiki/Non-specialized_hardware_comparison)
- [10] Autolykos gpu miner. [Online]. Available: <https://github.com/ergoplatform/Autolykos-GPU-miner>
- [11] Autoleakus scala implementation. [Online]. Available: <https://github.com/ergoplatform/ergo/tree/master/src/main/scala/org/ergoplatform/mining>

## Appendix

Implementation of hash function  $H$  which returns the values in  $\mathbb{Z}/q\mathbb{Z}$ :

---

**Algorithm 3** Numeric hash

---

```
1: function H(input)
2:   validRange :=  $(2^{256}/q) \cdot q$ 
3:   hashed := Blake2b256(input)
4:   if hashed < validRange then
5:     return hashed.mod(q)
6:   else
7:     return H(hashed)
8:   end if
9: end function
```

---

Implementation of hash function *genIndexes* which returns a list of size *k* with numbers in  $0 \dots (N - 1)$ :

---

**Algorithm 4** Index generator

---

```
1: function GENINDEXES(seed)
2:   hash := Blake2b256(seed)
3:   extendedHash := hash||hash
4:   return  $(0 \dots k - 1).map(i => extendedHash.slice(i, i + 4).mod(N))$ 
5: end function
```

---