

Asymmetric proof-of-work based on the Generalized Birthday problem

Alex Biryukov
University of Luxembourg
alex.biryukov@uni.lu

Dmitry Khovratovich
University of Luxembourg
khovratovich@gmail.com

Abstract—The proof-of-work is a central concept in modern cryptocurrencies, but the requirement for fast verification so far made it an easy prey for GPU-, ASIC-, and botnet-equipped users. The attempts to rely on memory-intensive computations in order to remedy the disparity between architectures have resulted in slow or broken schemes.

In this paper we solve this open problem and show how to construct an *asymmetric* proof-of-work (PoW) based on a computationally hard problem, which requires a lot of memory to generate a proof (called “memory-hardness” feature) but is instant to verify. Our primary proposal is a PoW based on the generalized birthday problem and enhanced Wagner’s algorithm for it. We introduce the new technique of *algorithm binding* to prevent cost amortization and demonstrate that possible parallel implementations are constrained by memory bandwidth. Our scheme has tunable and steep time-space tradeoffs, which impose large computational penalties if less memory is used.

Our solution is practical and ready to deploy: a reference implementation of a proof-of-work requiring 700 MB of RAM runs in 30 seconds on a 1.8 GHz CPU, increases the computations by the factor of 1000 if memory is halved, and presents a proof of just 148 bytes long.

I. INTRODUCTION

Request of intensive computations as a countermeasure against spam and denial of service (DoS) was first proposed by Dwork and Naor in [17]. Amount of work is certified by a proof, thus called *proof-of-work*, which is feasible to get by an ordinary user, but at the same time slows down multiple requests from the single machine or a botnet. Perhaps the simplest scheme is Hashcash [7], which requires a hash function output to have certain number of leading zeros and is adapted within the Bitcoin cryptocurrency. Nowadays, to earn 25 Bitcoins a miner must make an average of 2^{67} calls to a cryptographic hash function.

Long before the rise of Bitcoin it was realized [16] that the dedicated hardware can produce a proof-of-work much faster and cheaper than a regular desktop or laptop. Thus the users equipped with such hardware have an advantage over others, which eventually led the Bitcoin mining to concentrate in a few hardware farms of enormous size and high electricity consumption. An advantage of the same order of magnitude is given to “owners” of large botnets, which nowadays often accommodate hundreds of thousands of machines.

a) Memory-hard computing: In order to remedy the disparity between the ASICs and regular CPUs, Dwork et al. [16] proposed memory-intensive computations, and later memory-hard computations [18], i.e. those using a lot of

memory so that this amount can be reduced only at the large computational cost to the user. As memory is a very expensive resource in terms of area and the amortized chip cost, ASICs would be only slightly more efficient than regular x86-based machines. In addition, botnets are far less comfortable with computations that require GBytes of RAM, as they are very noticeable.

However, no scheme in [16], [18] has been adapted for the practical use as a PoW. Firstly, they are too slow for reasonably large amount of memory, and must use too little memory if required to run in reasonable time (say, seconds). The first memory-hard candidates [18] were based on superconcentrators [33] and similar constructions explored in the theory of pebbling games on graphs [23]. To fill N blocks in memory a superconcentrator-based functions make $N \log N$ hash function calls, essentially hashing the entire memory dozens of times. Better performance is achieved by memory-hard constructions among the finalists of the Password Hashing Competition [4], but their time-space tradeoffs have been explored only recently [14].

Secondly, the proposed schemes (including the PHC constructions) are *symmetric* with respect to the memory use. To initialize the protocol in [16], [18], a verifier must use the same amount of memory as the prover. This is in contrast with the Bitcoin proof-of-work: whereas it can be checked instantly (thus computational asymmetry) without precomputation, virtually no memory-hard scheme offers *memory asymmetry*. Thus the verifier has to be almost as powerful as the prover, and may cause DoS attacks by itself. In the cryptocurrency setting, fast verification is crucial for network connectivity. Even one of the fastest memory-hard constructions, scrypt [32], had to be taken with memory parameter of 128 KB to ensure fast verification in Litecoin. As a result, Litecoin is now mined on ASICs with 100x efficiency gain over CPU [2].

Finally, these schemes have not been thoroughly analyzed for possible optimizations and amortizations. To prove the work, the schemes should not allow any optimization (which adversaries would be motivated to conceal) nor should the computational cost be amortizable over multiple calls [17].

A reasonably fast and memory-asymmetric schemes would become a universal tool and used as an efficient DoS countermeasure, spam protection, or a core for a new egalitarian cryptocurrency.

b) Recent asymmetric candidates: There have been two notable attempts to solve the symmetry and performance problems. The first one by Dziembowski et al. [19] suggests

an interactive protocol, called a *proof-of-space*, where the prover first computes a memory-hard function and then a verifier requests a subset of memory locations to check whether they have been filled by a proper function. The verification can thus be rather quick. However, the memory-hard core of the scheme is based on a stack of superconcentrators and is expected to be prohibitively slow even in the range of a few MBytes (the actual scheme was not implemented). Moreover, the computational penalties are guaranteed only after the memory reduction by a logarithmic factor (say, 30 for 1 GB), which defeats the memory-hardness of the proof. Finally, this scheme is not amortization-free [31]: producing N proofs costs as much as producing one.

A more promising scheme was proposed by Tromp [37] as the *Cuckoo-cycle* PoW. The prover must find a cycle of certain length in a directed bipartite graph with N vertices and $O(N)$ edges. It is reasonably efficient (only 10 seconds to fill 1 GB of RAM with 4 threads) and allows very fast verification. The author claimed prohibitive time-memory tradeoffs. However, the original scheme was broken by Andersen [6]: a prover can reduce the memory by the factor of 50 with time increase by the factor of 2 only. Moreover, Andersen demonstrated a simple time-memory tradeoff, which allows for the constant time-memory product (reduced by the factor of 25 compared to the original). Thus the actual performance is closer to 3-4 minutes per GB¹). Apart from Andersen’s analysis, no other tradeoffs were explored for the problem in [37], there is no evidence that the proposed cycle-finding algorithm is optimal, and its amortization properties are unknown. Finally, Andersen’s tradeoff allows to parallelize the computations independently thus reducing the time-memory product and the costs on dedicated hardware.

Finally, the scheme called Momentum [27] simply looks for a collision in 50-bit outputs of the hash function with 26-bit input. The designer did not explore any time-space tradeoffs, but apparently they are quite favourable to the attacker: reducing the memory by the factor of q imposes only \sqrt{q} penalty on the running time [38] (more details in Appendix A).

c) Our contributions: We propose a family of fast, memory-asymmetric, optimization/amortization-free, limited parallelism proofs of work based on hard and well-studied computational problems. First we show that a wide range of hard problems (including a variety of NP-complete problems) can be adapted as an asymmetric proof-of-work with tunable parameters, where the ASIC and botnet protection are determined by the time-space tradeoffs of the best algorithms.

Our primary proposal is the PoW based on the generalized birthday problem, which has been explored in a number of papers from both theoretical and implementation points of view [12], [13], [26], [29], [39]. To make it amortization-free, we develop the technique called **algorithm binding** by exploiting the fact that Wagner’s algorithm carries its footprint on a solution.

In our scheme a user can independently tune time, memory, and time-memory tradeoff parameters. In a concrete setting,

¹The project webpage [37] claims Andersen’s optimizations be integrated into the miner, but the performance numbers are mainly unchanged since before the cryptanalysis appeared.

our 500 MB-proof is 142 bytes long and can be found in 31 seconds on a desktop PC using a reference implementation. Adversary trying to use 250 MB of memory would pay 1000-fold in computations using the best tradeoff strategy, whereas a memoryless algorithm would require prohibitive 2^{75} hash function calls. These properties and performance are unachievable by existing proposals. We have implemented and tested our scheme in several settings, with the code available at request.

To increase confidence in our proposal, we review and improve the best existing tradeoff strategies for the generalized birthday problem.

We also show how to build a PoW from two different hard problems generically and get the best of the their tradeoffs when the algorithm binding technique is not applicable. In this context we explore the hard knapsack problem, for which the best algorithms have been scrutinized in the recent papers [9], [15], [24].

d) Outline: This paper is structured as follows. First, we review the properties required from an asymmetric proof of work and show how to adapt a computationally hard problem for a PoW (Section II). We review the generalized birthday problem and Wagner’s algorithm in Section III and outline our primary proposal in Section IV. The new results on the time-space tradeoffs and parallelism are proven in Sections V and VI. Generic problem composition is left for Appendix.

II. PROOFS OF WORK AND HARD COMPUTATIONAL PROBLEMS

In this section we list the properties that we require from a proof-of-work and explain in a generic way how to turn a hard computational problem into a proof-of-work and what are the necessary conditions for such a problem. A reader interested in a concrete proposal with technical details may immediately proceed to Section IV.

A. Properties

We define a *problem*

$$\mathcal{P} : \mathcal{R} \times \mathcal{I} \times \mathcal{S} \rightarrow \{\text{true}, \text{false}\}.$$

as a hard predicate, where \mathcal{R} is the set of parameters that determine the hardness, \mathcal{I} is the set of inputs conforming to \mathcal{R} and \mathcal{S} is the set of possible solutions. We assume that there is an algorithm (or a family of algorithms) \mathcal{A}_R that solves \mathcal{P}_R on any I , i.e. finds S such that $\mathcal{P}(R, I, S) = \text{true}$.

A *proof-of-work scheme* based on \mathcal{P} (and implicitly on the algorithm \mathcal{A} for it) is an interactive protocol, which operates as follows:

- 1) The Verifier sends a challenge input $I \in \mathcal{I}$ and parameters $R \in \mathcal{R}$ to the Prover.
- 2) The Prover finds solution S such that $\mathcal{P}(R, I, S) = \text{true}$ and sends it to the Verifier.
- 3) The Verifier checks that S is a solution to $\mathcal{P}(R, I)$.

A non-interactive version (e.g., for cryptocurrencies) can be derived easily. In this setting I contains some public value (last block hash in Bitcoin) and prover’s ID. The prover publishes S so that every party can verify the proof.

Informally, \mathcal{A} should be moderately hard to impose significant costs on the prover. We also want that all the provers, equipped with sufficient memory, be in equal position so that no secret optimization or amortization can be applied. We summarize these requirements to \mathcal{P} and \mathcal{A} and the other properties it must satisfy in order to become ASIC- and botnet-resistant in the following list (cf. [17], [25]).

Progress-free process. In Bitcoin-like cryptocurrencies the mining process is usually a race among the miners who finds the proof first. To avoid centralization and mitigate the network delays, the mining must be a stochastic process, where the probability of the proof generation at any given time is non-zero and independent of the previous events. Therefore, the mining must be close to Poisson process with the number of proofs found in given timeframe following the Poisson distribution and the running time of the algorithm \mathcal{A}_R following the exponential distribution:

$$T(\mathcal{A}_R) \sim \text{Exp}(\lambda(R)).$$

The Poisson process is often emulated by the difficulty filter: a fixed-time algorithm \mathcal{B} , which additionally takes some nonce N as input, is concatenated with a hash function G , whose output should have a certain number of trailing zeros. In this case, the algorithm \mathcal{B} must also be *amortization-free*, i.e. producing q outputs for \mathcal{B} should be q times as expensive.

Large AT cost. We expect that the ASIC or FPGA implementation of algorithms with large area requirements and high area-time product (AT) would not be much better than desktop implementations by the Price-Performance parameter. The area is maximized by the memory requirements. Therefore, for a regular user, the optimal implementation of \mathcal{A}_R should require sufficiently large memory M . Most desktops and laptops can handle 1 GB of RAM easily, whereas 1 GB of memory on chip is expensive².

Small proof size and instant verification. The solution must be short enough and verified quickly using little memory in order to prevent DoS attacks on the verifier. We assume that some verifiers may use lightweight hardware such as smartphones with limited RAM and network bandwidth.

Step time-space tradeoffs. Memory requirements are worthwhile as long as the memory reduction disproportionately penalizes the user. Many memory-intensive algorithms can run with reduced memory. Suppose that \mathcal{A}_R is a family of algorithms using different amounts of memory. Then we can think of it as a single algorithm taking the available memory amount as a parameter. Let $T_R(M)$ be the average running time of \mathcal{A}_R with memory M , and let M_0 be the default memory requirements. Let us consider the running time growth when only M/q memory is used, $q > 1$:

$$C_R(q) = \frac{T_R(M/q)}{T_R(M_0)}.$$

²An increase in the AT cost for ASICs can be illustrated as follows. A compact 50-nm DRAM implementation [21] takes 500 mm² per GB, which is equivalent to about 15000 10 MHz SHA-256 cores in the best Bitcoin 40-nm ASICs [1] and is comparable to a CPU size. Therefore, an algorithm requiring 1 GB for 1 minute would have the same AT cost as an algorithm requiring 2⁴² hash function calls, whereas the latter can not finish on a PC even in 1 day. In other words, the use of memory can increase the AT cost by a factor of 1000 and more at the same time cost for the desktop user.

We say that the time-space tradeoff for \mathcal{A}_R is *polynomial with steepness s* if $C_R(q)$ can be approximated by a polynomial of q of degree s . We say that the tradeoff is exponential if $C_R(q)$ is exponential of q . We note that polynomial steepness with $s = 1$ (which we call *linear*) implies that the memory can be equally traded for time. This keeps the AT cost constant, but reduces the design and production cost of a potential chip. Thus higher steepness is desirable. Finding time-space tradeoffs for most hard problems is a non-trivial task [8], [20], as the best algorithms are usually optimized for computational complexity rather than for space requirements.

Flexibility. To account for further algorithm improvements and architecture changes, the time, memory, and steepness of the PoW must be tunable independently. For cryptocurrencies this helps to sustain constant mining rate. We recommend the following procedure (Figure 1). To adjust M , we change R and get a new complexity (T', M') . To increase T by the factor 2^d , we harden \mathcal{P} in the style of the Bitcoin difficulty filter: $H(S)$ must have d leading zero bits, where H is a cryptographic hash function.

Parallelism-constrained. When a large portion of ASIC is occupied by memory, adding a few extra cores does not increase the area. If \mathcal{A} can be parallelized, then the total time may be reduced and thus the AT cost can be lowered. However, if these cores have to communicate with each other, then the parallelism is limited by the network bandwidth. Thus if a specific PoW allows parallelism, the parallel version of algorithm \mathcal{A}_R should be *bandwidth-hard*, i.e. it quickly encounters the bottleneck in the network or RAM bandwidth.

Optimization-free. To avoid a clever prover getting advantage over the others, \mathcal{A} must be the most efficient algorithm to date, already employing all possible optimizations and heuristics, and it should be hard to find better algorithms.

We can now identify the problems with the Cuckoo-cycle PoW [37]. It can be parallelized, which lowers the AT cost. Dramatic optimizations were identified [6]. Its time-space tradeoff has steepness 1. Finally, it has not been explored for amortization.

B. Memory-hard proof-of-work based on a hard problem

Several hard computational problems suitable for a proof-of-work (PoW) were studied by Dwork and Naor [17] and later by Back [7]. These were PoWs with computational shortcuts: a verifier spends much less time than the prover. One could hope for memory shortcuts as well, i.e. verification requiring much less memory than the generation. However, a memory-hard PoW with a memory shortcut has been an open problem for quite a long time, as the existing proposals implicitly require both the verifier and the prover to allocate the same amount of memory.

Nevertheless, almost any hard problem can be turned into a proof-of-work in the framework outlined in Section II-A. Any reader with an expertise in a particular problem is thus encouraged to create his own PoW. The well-known NP(-complete) problems (including their search and optimization analogs) are the most natural candidates, since the best algorithms for them run in exponential time, i.e. $\log T = O(|I|)$. On the other hand, the verification is polynomial in $|I|$, so

it is polylogarithmic in T . Thus the verification for NP-complete-based PoW should be fast compared to the running time. Moreover, the best algorithms for NP-complete problems usually require non-negligible amount of memory and exhibit non-trivial time-space tradeoffs [20]. SAT solving, cliques, and Hamiltonian paths are all candidates for a PoW. The problems not demonstrated to be NP-complete but with best algorithms running in exponential time like factoring and discrete logs are natural proposals as well.




Fig. 1. Tuning time and memory requirements for proof-of-work. The problem with parameters R can be solved in time T and memory M . In order to change M to M' , we replace R with R' . To increase time by the factor of q , we add the difficulty filter q in addition to R .

It turns out that the properties that we listed in Section II-A are hard to satisfy simultaneously. A great difficulty lies in the optimization-free requirement, as the complexity of the most algorithms for hard problems is not evaluated with sufficient precision. Many algorithms are inherently amortizable. The existing implementations contain a number of heuristics. We concluded that the problem and the best algorithm must be very simple. So far we identified three problems, for which the best algorithms are explored, scrutinized, and implemented:

- The generalized-birthday, or k -XOR problem, which looks for a set of n -bit strings that XOR to zero. The best existing algorithm is due to Wagner [39] with minor modifications in [29]. The algorithm was implemented in [13], and its time-space tradeoffs were explored in [12]. This problem is the most interesting, as we can manipulate the tradeoff by changing k .
- The hard-knapsack problem, which looks for a subset of signed integers summing to 0. Whereas earlier instances of the knapsack problem can be solved in polynomial time [36], certain parameters are considered hard. For the latter the best algorithms are given in [9], [24]. This problem is appealing as its solution is likely to be unique, and the time and memory complexity are roughly the same.
- The information set decoding problem, which looks for a codeword in random linear code. Many algorithms were proposed for this problem [10], and many were implemented so we expect them to be well scrutinized. However, in the typical setting the memory complexity is significantly lower than the time complexity.

Among these, the generalized birthday problem appeared

the most suitable as its tradeoff steepness is tunable. In the next sections we introduce the problem and build our primary PoW proposal on it.

III. GENERALIZED-BIRTHDAY PROOF-OF-WORK

In this section we expose the generalized birthday problem and the algorithm for it by Wagner [39].

a) *Problem:* The generalized birthday problem for one list is formulated as follows: given list L of n -bit strings $\{X_i\}$, find distinct $\{X_{i_j}\}$ such that

$$X_{i_1} \oplus X_{i_2} \oplus \dots \oplus X_{i_{2k}} = 0.$$

We consider the setting where X_i are outputs of some PRNG, e.g. a hash function H in the counter mode. Thus we have to find $\{i_j\}$ such that

$$H(i_1) \oplus H(i_2) \oplus \dots \oplus H(i_{2k}) = 0. \quad (1)$$

For $k = 1$ this problem is the collision search, and can be solved trivially by sorting in $2^{n/2}$ time and space complexity if $|L| > 2^{n/2}$. However, for $k > 1$ and smaller lists the problem is harder. For instance, from the information-theoretic point of view we expect a solution for $k = 2$ in a list of size $2^{n/4}$, but no algorithm faster than $2^{n/3}$ operations is known.

Wagner demonstrated an algorithm for $k > 1$ and the lists are large enough to contain numerous solutions. It has time and space complexity of $O(2^{\frac{n}{k+1}})$ for lists of the same size. Wagner's algorithm generalizes easily to some operations other than XOR (e.g., to the modular addition). We also note that for $k \geq \log_2 n$ a XOR solution can be found by the much faster Gaussian elimination [11] with complexity of $O(2^k)$ string operations.

b) *Wagner's algorithm:* The basic algorithm to find a solution to Equation (1) is described in Algorithm 1.

Algorithm 1 Basic Wagner's algorithm for the generalized birthday problem.

Input: list L of N n -bit strings ($N \ll 2^n$).

- 1) Enumerate the list as $\{X_1, X_2, \dots, X_N\}$ and store pairs (X_j, j) in a table.
- 2) Sort the table by X_j . Then find all unordered pairs (i, j) such that X_i collides with X_j on the first $\frac{n}{k+1}$ bits. Store all tuples $(X_{i,j} = X_i \oplus X_j, i, j)$ in the table.
- 3) Repeat the previous step to find collisions in $X_{i,j}$ on the next $\frac{n}{k+1}$ bits and store the resulting tuples $(X_{i,j,k,l}, i, j, k, l)$ in the table.
- ... Repeat the previous step for the next $\frac{n}{k+1}$ bits, and so on until only $\frac{2n}{k+1}$ bits are non-zero.
- $k + 1$ At the last step, find a collision on the last $\frac{2n}{k+1}$ bits. This gives a solution to the original problem.

Output: list $\{i_j\}$ conforming to Equation (1).

c) *Analysis*: For the further text, assume that sorting $l = O(N)$ elements is computationally equivalent³ to l calls to the hash function H . Let a single call to H be our time unit.

Proposition 1: For $N = 2^{\frac{n}{k+1}+1}$ Algorithm 1 produces two solutions (on average) using $(2^{k-1} + n)N/8$ bytes of memory in time $(k+1)N$.

Proof: Suppose we store $N = 2^{\frac{n}{k+1}+1}$ tuples at the first step. Then after collision search we expect

$$2^{\frac{2n}{k+1}+1-\frac{n}{k+1}} = 2^{\frac{n}{k+1}+1}$$

entries for the second table, then the same number of entries for the third table and so on. Before the last (k -th) collision search we have $2^{\frac{n}{k+1}+1}$ entries, thus on average we obtain two solutions.

The computational complexity is dominated by the complexity of list generation (N hash calls) and subsequent k sortings of N elements. Therefore, the total computational complexity is equivalent to

$$(k+1)N = (k+1)2^{\frac{n}{k+1}+1}$$

hash function calls. This ends the proof \blacksquare

If larger lists are used, the table will grow in size over the steps. We have taken the list size exactly so that the expected number of solutions is small and the table size does not change much.

d) *Algorithm binding*: The generalized birthday problem in its basic form lacks some necessary properties as a proof-of-work. The reason is that Wagner's algorithm can be iterated to produce multiple solutions by selecting other sets of colliding bits or using more sophisticated techniques [13]. If more memory is available, these solutions can be produced at much lower amortized cost (Proposition 3). Since this property violates the non-amortization requirement for the PoW (Section II-A), we suggest modifying the problem so that only two solutions can be produced on average.

Our modification is inspired by the fact that a solution found by Wagner's algorithm carries its footprint. Namely, the intermediate 2^l -XORs have leading $\frac{nl}{k+1}$ bits, for example $X_{i_4} \oplus X_{i_5} \oplus X_{i_6} \oplus X_{i_7}$ collide on certain $\frac{2n}{k+1}$ bits. Therefore, if we pre-fix the positions where 2^l -XORs have zero bits, we bind the user to a particular algorithm flow. Moreover, we can prove that the total number of possible solutions that conform to these restrictions is only 2 on average, so that the problem becomes amortization-free for given input list L . We only have to take care of duplicate solutions which appear if we swap 2^{l-1} -XORs within the 2^l -XOR, for any l . We simply require that every 2^l -XOR is ordered as a pair, e.g. with lexicographic order. We stress that a certain order is a prerequisite as otherwise duplicate solutions (produced by swaps in pairs, swaps of pairs, etc.) would be accepted.

With this modification the Gaussian elimination algorithm [11] does not apply anymore, so we can use larger k with no apparent drop in complexity. Moreover,

e) *Time-space tradeoffs*: The time-space tradeoffs for Wagner's algorithm are explored in details in Section V-B. Here we report the main results. First, we consider optimizations, which are based on methods from [13].

Proposition 2: Optimized Wagner's algorithm (Algorithm 2) for $N = 2^{\frac{n}{k+1}+1}$ runs in $M(n, k) = 2^{\frac{n}{k+1}}(2^k + \frac{n}{2(k+1)})$ bytes of memory and $T(n, k) = k2^{\frac{n}{k+1}+2}$ time.

The next proposition is a corollary from results in [12].

Proposition 3: Using $qM(n, k)$ memory, a user can find $2q^{k+1}$ solutions with cost $qT(n, k)$, so that the amortized cost drops by q^{k-1} .

Our novel result is the following tradeoffs for standard and algorithm-bound problems.

Proposition 4: Using $M(n, k)/q$ memory, a user can find 2 solution in time $C_1(q)T(n, k)$, where

$$C_1(q) \approx \frac{3q^{\frac{k-1}{2}} + k}{k+1}.$$

Therefore, Wagner's algorithm for finding 2^k -XOR has a tradeoff of steepness $(k-1)/2$. At the cost of increasing the solution length, we can increase the penalty for memory-reducing users.

Proposition 5: Using constant memory, a user can find one algorithm-bound solution in time

$$2^{\frac{n}{2}+2k+\frac{n}{k+1}}.$$

Proposition 6: Using $M(n, k)/q$ memory, a user can find 2 algorithm-bound solutions in time $C_2(q)T(n, k)$, where

$$C_2(q) \approx 2^k q^{k/2} k^{k/2-1}.$$

Therefore, the algorithm-bound proof-of-work has higher steepness $(k/2)$, and the constant is larger.

f) *Parallelism*: So far we equalized the time and computational complexity, whereas an ASIC-equipped user or the one with a multi-core cluster would be motivated to parallelize the computation if this reduces the AT cost. The following result, also stated in [12], is explained in details in Section VI.

Proposition 7: With p processors and $M(n, k)$ shared memory a user can find 2 algorithm-bound solutions in time $\frac{T(n, k)}{p}(1 - \log_N p)$. Additionally, the memory bandwidth grows by the factor of p .

For fixed memory size, memory chips with bandwidth significantly higher than that of typical desktop memory (such as DDR3) are rare. Assuming that a prover does not have access to memory with bandwidth higher than certain Bw_{max} , we can efficiently bound the time-memory (and thus the time-area) product for such implementations.

Corollary 1: Let the reference memory of size M have bandwidth Bw , and let the prover be equipped with memory chips of bandwidth Bw_{max} . Then the time-area product for the prover can be reduced by the factor $\frac{Bw_{max}}{Bw}$ using $\frac{Bw_{max}}{Bw}$ parallel sorting processors.

To the best of our knowledge, the highest reported bandwidth in commercial products does not exceed 200 GB/s

³The actual ratio depends on the hash function and the sorting algorithm.

(Playstation 4 has 172 GB/s bandwidth [5]), whereas the desktop DDR3 can have as high as 17 GB/s bandwidth [3]. Thus we conclude that the highest advantage a prover can get from parallel hardware does not exceed the factor of 15. This may mean that our proof-of-work is GPU- and desktop-friendly, whereas it is also ASIC-resistant in the sense that an ASIC implementation does not yield smaller time-area product.

IV. OUR PRIMARY PROPOSAL

A. Specification

To generate an instance for the proof protocol, a verifier selects a cryptographic hash function H and integers n, k, d , which determine time and memory requirements as follows:

- Memory M is $2^{\frac{n}{k+1}+k}$ bytes;
- Time T is $(k+1)2^{\frac{n}{k+1}+d}$ calls to the hash function H .

Then he selects a seed I (which may be a hash of transactions, block chaining variable, etc.) and asks the prover to find 160-bit nonce V and $(\frac{n}{k+1}+1)$ -bit x_1, x_2, \dots, x_{2^k} such that

$$\begin{cases} H(I||V||x_1) \oplus H(I||V||x_2) \oplus \dots \oplus H(I||V||x_{2^k}) = 0, \\ H(I||V||x_1||x_2||\dots||x_{2^k}) \text{ has } d \text{ leading zeros,} \\ H(I||V||x_{w2^l+1} \oplus \dots \oplus H(I||V||x_{w2^l+2^l}) \\ \text{has } \frac{nl}{k+1} \text{ leading zeros for all } w, l \\ (x_{w2^l+1}||x_{w2^l+2}||\dots||x_{w2^l+2^{l-1}}) < \\ < (x_{w2^l+2^{l-1}+1}||x_{w2^l+2^{l-1}+2}||\dots||x_{w2^l+2^l}). \end{cases} \quad (2)$$

Here the order is lexicographical. A prover is supposed to run Wagner's algorithm and then H (Figure 2).

The analysis in Section III shows that Wagner's algorithm produces 2 solutions per call on average for each V . Thus to produce a hash with d leading zeros it must be called 2^{d-1} times with distinct V , which yields the time complexity $(k+1)2^{\frac{n}{k+1}+d}$. The memoryless verifier checks all the conditions. Note that computations for $V = V_0$ can not be reused for another $V \neq V_0$. Note that the order of two (or more solutions) produced by Wagner's algorithm is not important; we do not have to enumerate them; only the one that passes the d -zero test is needed. Also note that shorter ($2^l, l < k$) solutions are not allowed neither full solutions based on them (the order prohibits this).

Our proposal fulfills all the properties from Section II-A. The large AT cost is ensured by $M \geq 2^{30}$. The implementation can be made fast enough to use $M = 2^{29}, n = 144, k = 5$ in 30 seconds. The verification is instant, as it requires only 2^k hashes. The tradeoff has steepness $(k-1)/2$ and a large constant. Parallelism is restricted due to the use of sequential and memory-hard function \mathcal{G} . Optimizations are explored, and amortization does not reduce costs as the number of solutions is small on average. Finally, time, memory, and steepness can be adjusted independently.

B. Implementation and concrete parameters

Varying n and k we can reach a wide range of the memory and time complexity of our proof-of-work proposal. From the




Fig. 2. Proof-of-work based on the generalized birthday problem.

implementation point of view, it is convenient to have $\frac{n}{k+1}$ as multiples of 8 so that we can work with integer number of bytes.

We suggest a wide range of parameters, which cover different memory requirements and tradeoff resistance (Table I). The memory and time requirements are taken from Proposition 2 with $\epsilon = 0$ and indices trimmed to 8 bits. The solution size also includes input and output of the memory-hard function \mathcal{G} , which take 40 bytes altogether.

n	k	Complexity				Solution size
		Memory-full		Memoryless		
		Peak memory	Time	Time		
96	5	2.5 MB	$2^{19.2}$	2^{74}	116 B	
128	7	8.5 MB	2^{20}	2^{94}	404 B	
160	9	32.5 MB	$2^{20.3}$	2^{114}	1.5 KB	
176	10	64.5 MB	$2^{20.4}$	2^{124}	3 KB	
192	11	128.5 MB	$2^{20.5}$	2^{134}	6 KB	
96	3	320 MB	2^{27}	2^{78}	52 B	
144	5	704 MB	$2^{27.5}$	2^{106}	148 B	
192	7	2.2 GB	2^{28}	2^{134}	532 B	
240	9	8.3 GB	$2^{28.2}$	2^{162}	1.6 KB	
96	2	82 GB	$2^{34.5}$	2^{84}	40 B	
288	8	11 TB	2^{36}	2^{192}	1.3 KB	

TABLE I. CONCRETE PARAMETERS AND THEIR SECURITY LEVEL FOR THE PROOF-OF-WORK BOUND TO WAGNER'S GENERALIZED BIRTHDAY ALGORITHM. MEMORY-FULL COMPLEXITIES ARE TAKEN FROM THE ANALYSIS OF ALGORITHM 2 (PROPOSITION 2). MEMORYLESS COMPLEXITY IS TAKEN FROM PROPOSITION 5.

As a proof of concept, we have implemented and tested our proof-of-work with $n = 144, k = 5$ and several smaller versions. Our implementation is written in C without assembly optimization but with multi-threading enabled. We used standard library's quicksort and a small amount of extra memory to buffer the produced collisions. We trimmed only the upper $\frac{n}{k+1} + 1$ -th bit of each index to fit into integer number of bytes rather than aggressively optimize with very short indices. As we did not exploit the uniformity of the hash values, we expect that the sorting step can be optimized. However, such an implementation would have to be tuned for concrete (n, k) .

V. TIME-SPACE TRADEOFFS FOR THE GENERALIZED BIRTHDAY PROBLEM

There can be two types of time-space tradeoffs for the generalized birthday algorithm. First, there could be small optimizations in storing indices, the hash outputs, and sorting

Threads	Time
1	67 sec
2	49 sec
4	33 sec
8	31 sec

TABLE II. PERFORMANCE OF THE GENERALIZED BIRTHDAY ALGORITHM WITH $n = 144, k = 5$ (PROOF OF 500 MB) USING 1.5 GB OF RAM ON 1.8 GHZ CPU WITH 4 CORES.

algorithms. We will show that the straightforward implementation of the generalized birthday algorithm allows the memory reduction by a small factor (2 or 3 depending on the parameters) with about the same increase in the time complexity. However, these optimizations are limited.

If the prover wants to save the memory further, he would have to reduce the total number of tuples. We will show that this approach would cause him harsh computational penalties.

A. Optimizations

In Algorithm 1 the index size doubles in size at each step, whereas we can trim $\frac{n}{k+1}$ of intermediate sum per step. There can be two types of optimizations, partly explored in [13]:

- Not storing the intermediate XOR value but recomputing it at each step from the indices. This approach was taken in [13]. However, for large k this approach becomes too expensive.
- Storing only a fraction of index bits, e.g. t bits only per index. Then after the last step we have to figure out the missing bits for all 2^k -XOR candidates. For large t this can be done by simply checking the 2-XOR of all subpairs of the 2^k -XOR. For smaller t we essentially have to repeat our algorithm with 2^k different lists, which gives an overhead time factor about k , and $2^{\frac{n}{k+1}+1-t}$ values in each list.

These two optimizations are illustrated in Figure 3 for $n = 144, k = 5$. It appears that the combination of both optimizations yields the best results, where we recompute the hashes from at first steps, and trim the indices at later steps (Algorithm 2).

The optimal pattern depends on n and k , so we checked it manually for all suggested parameters assuming that we trim the indices to 8 bits. For all parameters the peak memory use is in the last step, and for all parameters except $n = 96, k = 2, 3$ it is optimal to switch to index trimming before the last step. Therefore, the peak memory is upper bounded by 2^{k-1} 8-bit indices per tuple at the last step. Therefore, at the last step the tuple length is $\frac{2n}{k+1} + 8 \cdot 2^{k-1}$ bits, or $2^{\frac{n}{k+1}}(2^k + \frac{n}{2(k+1)})$ bytes in total. To recover the trimmed bits, we generate $2^{\frac{n}{k+1}-7}$ hash values for each i_j and run the algorithm again, now with 2^k lists, each 2^8 times as small. In the multi-list version of the algorithm, it suffices to keep only k lists in memory simultaneously [26]. The time complexity of the multi-list step is dominated by generating the 2^k values, as later lists are much smaller. Thus the multi-list phase runs in $2^{\frac{n}{k+1}-7+k}$ time, which is smaller than $(k-1)2^{\frac{n}{k+1}+1}$ for all our parameters. This implies Proposition 2.

Algorithm 2 Optimized Wagner’s algorithm for the generalized birthday problem.

Input: list L of N n -bit strings.

- 1) Enumerate the list as $\{X_1, X_2, \dots, X_N\}$ and store pairs (j) in a table.
- 2) Sort the table by X_j , computing it on-the-fly. Then find all unordered pairs (i, j) such that X_i collides with X_j on the first $\frac{n}{k+1}$ bits. Store all such pairs (i, j) in the table.
- 3) Repeat the previous step to find collisions in $X_{i,j}$ (again recomputing it on the fly) on the next $\frac{n}{k+1}$ bits and store the resulting tuples (i, j, k, l) in the table.
- ... Repeat the previous step for the next $\frac{n}{k+1}$ bits, and so on. When indices trimmed to 8 bits plus the length $X_{i,j,\dots}$ becomes smaller than the full index tuple, switch to trimming indices.
- $k+1$ At the last step, find a collision on the last $\frac{2n}{k+1}$ bits. This gives a solution to the original problem.

Output: list $\{i_j\}$ conforming to Equation (1).




Fig. 3. Tuple length in bits over the steps of Wagner’s algorithm in the basic and the optimized implementations.

B. Generic tradeoffs

If a prover wants to save even more memory than the indices allow, he would have to store fewer tuples at each step. The first time-space tradeoffs of this kind were explored by Bernstein in [12]. Suppose that the adversary stores only $2^{\frac{n}{k+1}+1}/q$ elements at each step, $q > 1$. Then Bernstein suggests truncating H to $n - (k+1)\log q$ bits and apply Wagner’s algorithm to $n' = n - (k+1)\log q$. After the last step we check if the remaining $(k+1)\log q$ bits are zero, which succeeds with probability q^{-k-1} . Therefore, the algorithm must be repeated q^{k+1} times, but each step is cheaper by the factor of q . Thus the computational penalty is computed as $C(q) = q^k$. We note that the same reasoning applies to the case where more memory is available: if we have $qM(n, k)$ memory for $q > 1$ then we obtain $2q^{k+1}$ solutions in the end, but we spend q times as many computations. This implies Proposition 3.

In the later paper Bernstein et al. [13] suggested applying the memoryless collision search method at the last step of the algorithm. They viewed the first $(k-1)$ steps as a single function \mathcal{H} that outputs $\frac{2n}{k+1} + (k+1)\log q$ bits. The complexity

of the memoryless collision search is about the square root of the output space size [30], [38] with more precise estimate of $4 \cdot 2^{l/2}$ for l -bit function. Therefore, the total computational complexity is

$$4 \cdot 2^{\frac{n}{k+1} + \frac{(k+1) \log q}{2}} \cdot k \cdot 2^{\frac{n}{k+1} + 1 - \log q} = 4k2^{\frac{2n}{k+1} + 1} \cdot q^{(k-1)/2}$$

the total computational penalty is computed as

$$C(q) \approx 4 \cdot 2^{\frac{n}{k+1}} q^{\frac{k+1}{2}}.$$

Later, Kirchner suggested [26] using available memory for the last step, getting a slightly better tradeoff⁴. The drawback of both approaches is that it requires multiple iterations of the algorithm and thus is very demanding in terms of memory and network bandwidth in practical implementations.

a) Parallel collision search: The following result is of great importance for our tradeoff exploration. The parallel collision search algorithm for the function that maps m bits to n bits, $m \geq n$, due to van Oorschot and Wiener [38], finds T collisions using $2wm$ bits of memory in time

$$\frac{2.5T \cdot 2^{n/2}}{\sqrt{w}} \quad (3)$$

. Here the memory is occupied by w distinguished points, and about $2^{n/2} \sqrt{w}$ calls must be spent to generate those points. Every point is a pair of inputs plus some additional information.

b) Proof of Proposition 4: Our idea to improve Bernstein's tradeoffs is to increase the number of colliding bits at the first step to $\frac{n}{k+1} + k \log q$ and generate $2^{\frac{n}{k+1} + 1} / q$ collisions. The next $(k-2)$ steps require collisions on $\log q$ fewer bits each, and the final step — on $2 \log q$ fewer bits. Then on average we obtain the same 2 collisions at the last step, thus being as lucky as in the memory-full approach. Therefore, we ran Wagner's algorithm only once, as soon as we get the necessary amount of inputs. However, the first step is more expensive.

A straightforward approach to generate that many collisions would be to carry Bernstein's idea to the first step of Algorithm 2 and keep only those inputs that yield $k \log q$ leading zero bits. However, this gives the computational penalty factor of q^{k-1} . A better approach is to use the parallel collision search method, where the average cost of one collision grows sublinearly as a function of memory (Equation (3)). Each distinguished point is smaller than the output of H . Using $2^{\frac{n}{k+1} + 1} / q$ distinguished points to generate $2^{\frac{n}{k+1} + 1} / q$ collisions on $\frac{n}{k+1} + k \log q$ bits, we make

$$5 \cdot 2^{\frac{n}{2(k+1)} + \frac{k \log q}{2} + \frac{n}{2(k+1)} + \frac{1 - \log q}{2}} \approx 3 \cdot q^{\frac{k-1}{2}} \cdot 2^{\frac{n}{k+1} + 1}.$$

calls to H . The next steps calls the equivalent of $\frac{k}{q} 2^{\frac{n}{k+1} + 1}$ hash calls. The success rate is the same. Thus the total computational penalty is estimated as

$$C(q) \approx \frac{3q^{\frac{k-1}{2}} + k}{k+1}.$$

⁴A rigorous formula is difficult to extract from Kirchner's manuscript.




Fig. 4. Computation-memory tradeoffs for the generalized birthday algorithm: ours (Proposition 4), Bernstein's [12], and "Bernstein-Kirchner" [13], [26]. The complexities are given in \log_2 .

C. Algorithm-bound tradeoffs

In our proof-of-work proposal we explicitly specify that the solution must carry the footprint of Wagner's algorithm, in particular all intermediate 2^l -XORs must have $\frac{nl}{k+1}$ zeros at certain positions. First we show that the expected total number of solutions that conform to this property is 2. Indeed, there are $2^{\frac{2^k n}{k+1} + 2^k}$ ordered 2^k -tuples of $\frac{n}{k+1}$ -bit values. There are $2^k - 1$ intermediate 2^l -XORs that must have lexicographic order, which reduces the total number of tuples to $2^{\frac{2^k n}{k+1} + 1}$. The restriction of having zeros at certain position contains one $\frac{2n}{k+1}$ -bit condition at the last step, two $\frac{n}{k+1}$ -bit conditions at the step before last, ..., 2^{k-1} $\frac{n}{k+1}$ -bit conditions at the first step, or $2^k \frac{n}{k+1}$ filtering bits in total. Thus there are 2 possible solutions left on average.

Let us now explore the time-space tradeoffs. It is not evident that the tradeoffs that we obtained in Section V-B can be carried out to the algorithm-bound setting. Since the inputs to H are limited to $2^{\frac{n}{k+1} + 1}$ values, it is not trivial even to find a memoryless attack with complexity smaller than 2^n . Surprisingly, the tradeoffs for algorithm-bound adversaries are only slightly worse than the original ones.

a) Proof of Proposition 5 : First we show how to find an algorithm-bound solution with very low memory. Recall that the memoryless collision search for f works as the ρ -method: we iterate f till we detect a cycle so that the two different entry points to the cycle constitute a collision. A cycle is detected by either iterating $f()$ and $f(f())$ simultaneously or by storing a few distinguished points along the iteration. The time complexity is about $4 \cdot 2^{l/2}$ for l -bit f for the success rate close to 1 [30]. However, we might have to truncate f first to ensure that its domain is as large as the range.

It is important to know that the basic memoryless algorithm is seeded, i.e. it starts at some point and the eventual complexity is determined by this point. We can imagine an oracle \mathcal{O}_f that takes S as seed and outputs a collision for f .

Consider Equation (2). The algorithm binding requires us

to find a solution such that the intermediate 2^l -XORs collide on certain $\frac{nl}{k+1}$ bits. Let us denote such XORs by separate functions:

$$f^{2^l} = H(I||x_1) \oplus H(I||x_2) \oplus \dots \oplus H(I||x_{2^l})$$

Therefore for the original problem we have to find a collision in $f^{2^{k-1}}$, and in the algorithm-bound setting each of the colliding inputs must itself be a collision for $f^{2^{k-2}}$ on fewer bits and so on. At each subsequent step we require a collision on $\frac{n}{k+1}$ bits only, as the colliding bits accumulate from the nodes of the recursion to the root. Note that $f^{2^{k-1}}$ has only $2^{\frac{n}{k+1}+1}$ inputs but a $\frac{2n}{k+1}$ -bit output, i.e. it is an expanding function.

This suggests the memoryless solution search as a recursive memoryless collision search. The last step makes $9 \cdot 2^{\frac{3n}{2(k+1)}}$ calls to $\mathcal{O}_{f^{2^{k-1}}}$ (Appendix A) where the seeds are intermediate inputs to f^{2^k} . Oracle $\mathcal{O}_{f^{2^{k-1}}}$ makes $2^{\frac{n}{2(k+1)}+2}$ calls to $\mathcal{O}_{f^{2^{k-2}}}$, and so on. In total we make $2^{\frac{n}{2}+2k+\frac{n}{k+1}}$ calls to f to find one collision. This ends the proof.

b) Proof of Proposition 6: The memoryless solution search can be adapted for the reduced memory settings. The idea is to replace the memoryless collision search at each level but the last one with the collision search with distinguished points. In contrast to Section V-B, we apply the method recursively, and do not store the solutions. However, we have to store distinguished points for each recursion level simultaneously, which affects the computation complexity as follows.

Suppose we have memory sufficient to store only $2^{\frac{n}{k+1}+1}/q$ tuples for some $q > 1$. We split the available memory evenly between k sets of distinguished points (for the sake of simplicity assume that each point is about the same size as the tuple), so that each set contains $\frac{2^{\frac{n}{k+1}+1}}{qk}$ points. The sets are available at all levels. The amortized collision cost using w distinguished points is $2.5 \cdot \frac{2^{n/2}}{\sqrt{w}}$ (Equation (3)). Thus we obtain that oracle $\mathcal{O}_{f^{2^l}}$ makes on average

$$2.5 \cdot 2^{\frac{n}{2(k+1)} - \frac{n}{2(k+1)} - 0.5} \sqrt{qk} \approx 2\sqrt{qk}$$

calls to $\mathcal{O}_{f^{2^{l-1}}}$ to produce a solution. The final oracle makes $2.5 \cdot 2^{\frac{3n}{2(k+1)}+1.5}/\sqrt{w} \approx 2^{\frac{n}{k+1}+2} \sqrt{qk}$ calls to $\mathcal{O}_{f^{2^{k-1}}}$. Therefore, the total computational complexity is

$$2^{\frac{n}{k+1}+1} \cdot 2^k q^{k/2} k^{k/2},$$

and the penalty should be computed as

$$C(q) \approx 2^k q^{k/2} k^{k/2-1}.$$

This ends the proof.

Let us take one of our concrete parameters as an example. Let $n = 144, k = 5$, i.e. we suggest finding 32-XOR on 144 bits. A straightforward implementation of the generalized birthday algorithm would require 1.6 GBytes of RAM and about 1 minute of a single CPU core. Recomputing the hash values for the first two steps and truncating the indices to 8 bits at the last two steps, we can decrease the peak tuple length to 176 bits, thus in total requiring 704 MBytes, or aggressively trim to 4 bits, reaching 500 MBytes. However,

further reductions are more expensive. Using 2^{24} instead of 2^{25} tuples would cause the computational penalty factor of 2^{10} , and factor of 2^{20} for using 2^{20} tuples ($q = 1/32$). We summarize that for large memory reductions the computational penalty would be prohibitive even for adversaries equipped with a number of parallel computational cores.

VI. PARALLELISM

It is rather easy to analyze Wagner's algorithm from the parallelism point of view [12], since it consists of well-known procedures: batch hashing and collision search via sorting. Suppose we have p processors with $M(n, k)$ shared memory. The hashing step is straightforward to parallelize: the processors merely fill their own block of memory.

Parallel sorting algorithms have been explored for decades, and full exposition of these results is beyond the scope of this work. Whereas the quicksort is traditional choice for single-thread applications, a number of its variations as well as that of bucket sort, radix sort, sample sort, and many others have been proposed, as they all differ in scalability, computational time, communication complexity, and memory bandwidth. The implementations on a CPU, a multi-core cluster [22], a GPU [40], and even FPGA have been reported.

a) Proof of Proposition 7: For our purpose a modification of bucket sort, also called a sample sort, suffices. The idea is the following (for the sake of simplicity assume that p is a power of 2). Let us denote the total number of tuples by N . We partition the entire memory into p^2 equal cells and represent it as a $(p \times p)$ matrix $M[i, j]$. At the first step the processors operate row-wise. They generate the hashes and place them in one of p cells according to the leading $\log p$ bits. This takes time N/p . Thus column j has only entries starting with j (in the bit representation), so there is no collision between the entries from different columns. Then the processors operate column-wise and sort the columns simultaneously. Then each processor goes over the sorted column, identifies collisions, and overwrites the column with placing the collisions into different buckets, so now row j has only entries starting with j . At the next step the processors operate rowwise, and so on. This method requires a small buffer to store the collisions, but due to uniformity of entries it can be rather small (1% in our experiments).

Sorting each column with quicksort requires $O(\frac{N}{p} \log \frac{N}{p})$ time, and this can be done independently for each step of the algorithm. Therefore each collision search step of Wagner's algorithm is faster by the factor

$$\frac{p \log N}{\log \frac{N}{p}} = \frac{p}{1 - \log_N p}.$$

We note that each processor must have access to one entire column and one entire row, so that it is not possible to restrict a processor to its own memory. Therefore, memory conflicts are unavoidable, and the memory chip bandwidth becomes the bottleneck as p increases.

The total bandwidth is calculated as follows. Assuming that the amount of memory operations in sorting is (almost) a linear function of array length, we get that if one core needs R memory reads/writes to sort N entries, then p cores need

R/p operations each. In addition, the cores must distribute the collisions into different buckets, spending $O(N)$ memory operations for that. If the running time decreases by the factor of p , then the bandwidth grows at least by the same factor.

This ends the proof.

b) Parallel sorting in practice: The observable speedup on multi-core CPU and GPU is not that big. The fastest GPU sortings we are aware of have been reported in [28], where radix sort was implemented and tested on a number of recent GPUs. The best performance was achieved on GTX480, where 2^{30} 32-bit keys were sorted in 1 second⁵. The same keys on the 3.2 GHz Core-i7 were sorted with rate 2^{28} keys per second [34], i.e. only 4 times as slow. Thus the total advantage of GPU over CPU is about the factor of 4, which is even smaller than bandwidth ratio (134 GB/s in GTX480 vs 17 GB/s for DDR3). This supports our assumption of very limited parallelism advantage due to restrictions of memory bandwidth ([34] also mentions high GPU memory latency as a slowing factor).

VII. ACKNOWLEDGEMENT

Various people from the Bitcoin community, in particular Gregory Maxwell for some of the properties that a useful PoW should satisfy. We thank John Tromp for reporting some mistakes in the discussion on Momentum and clarifying some points in the Cuckoo cycle scheme.

VIII. CONCLUSION

We have described a general approach to construction of asymmetric proofs of work from hard problems. Given a list of requirements for an asymmetric and ASIC-resistant PoW we identified the generalized birthday problem as the one with a scrutinized algorithms decently studied for tradeoffs.

We showed that the running time of the generalized birthday algorithm can be amortized over multiple solutions. We have introduced a technique called *algorithm binding* that prevents solution amortization by making solutions almost unique. Moreover, we investigated the time-memory tradeoffs and demonstrated that the new technique gives time-space tradeoffs that are better for the defender at negligible verification overhead. Thanks to the solution length parameter k in the generalized birthday problem, we may vary the tradeoffs so we suggest a wide range of time, memory, and tradeoff parameters for a variety of applications.

We also demonstrated that even though Wagner's algorithm is inherently parallel, any parallel implementation of its components quickly exhaust available memory bandwidth and thus has only limited advantage while running on GPU or ASIC. Altogether, we get a memory-hard, ASIC- and botnet-resistant PoW with extremely fast verification and very small proof size.

To demonstrate the practicality of our solution, we implemented our PoW on a PC for a 500-MB proof. A reference, non-optimized implementation runs in 30 seconds with 4 threads, verification takes microseconds, and the proof length is tiny, just 148 bytes.

⁵Closer to our $n = 144, k = 5$ proposal, where 48-bit keys are sorted together with 128-bit values, GTX480 sorts $32+128$ -bit entries with rate $2^{27.5}$ per second, but there is no such benchmark in [34]

REFERENCES

- [1] Avalon asic's 40nm chip to bring hashing boost for less power, 2014. <http://www.coindesk.com/avalon-asics-40nm-/chip-bring-hashing-boost-less-power/>.
- [2] Litecoin: Mining hardware comparison, 2015. available at https://litecoin.info/Mining_hardware_comparison and <http://cryptomining-blog.com/3106-quick-//comparison-of-the-//available-30-mhs-scrypt-asic-miners/>.
- [3] Memory deep dive: Memory subsystem bandwidth, 2015. <http://frankdenneman.nl/2015/02/19/memory-deep-dive-memory-subsystem-bandwidth/>.
- [4] Password Hashing Competition, 2015. <https://password-hashing.net/>.
- [5] The xbox one: Hardware analysis & comparison to playstation 4, 2015. <http://www.anandtech.com/show/6972/xbox-one-hardware-compared-to-playstation-4/2>.
- [6] David Andersen. A public review of cuckoo cycle. <http://www.cs.cmu.edu/~dga/crypto/cuckoo/analysis.pdf>, 2014.
- [7] Adam Back. Hashcash – a denial of service counter-measure, 2002. available at <http://www.hashcash.org/papers/hashcash.pdf>.
- [8] Paul Beame, Allan Borodin, Prabhakar Raghavan, Walter L. Ruzzo, and Martin Tompa. Time-space tradeoffs for undirected graph traversal. In *FOCS'90*, pages 429–438. IEEE Computer Society, 1990.
- [9] Anja Becker, Jean-Sébastien Coron, and Antoine Joux. Improved generic algorithms for hard knapsacks. In *EUROCRYPT'11*, volume 6632 of *Lecture Notes in Computer Science*, pages 364–385. Springer, 2011.
- [10] Anja Becker, Antoine Joux, Alexander May, and Alexander Meurer. Decoding random binary linear codes in $2^n/20$: How $1 + 1 = 0$ improves information set decoding. In *EUROCRYPT'12*, volume 7237 of *Lecture Notes in Computer Science*, pages 520–536. Springer, 2012.
- [11] Mihir Bellare and Daniele Micciancio. A new paradigm for collision-free hashing: Incrementality at reduced cost. In Walter Fumy, editor, *Advances in Cryptology - EUROCRYPT '97, International Conference on the Theory and Application of Cryptographic Techniques, Konstanz, Germany, May 11-15, 1997, Proceeding*, volume 1233 of *Lecture Notes in Computer Science*, pages 163–192. Springer, 1997.
- [12] Daniel J Bernstein. Better price-performance ratios for generalized birthday attacks. In *Workshop Record of SHARCS*, volume 7, page 160, 2007.
- [13] Daniel J. Bernstein, Tanja Lange, Ruben Niederhagen, Christiane Peters, and Peter Schwabe. Fsbday. In *INDOCRYPT'09*, volume 5922 of *Lecture Notes in Computer Science*, pages 18–38. Springer, 2009.
- [14] Alex Biryukov and Dmitry Khovratovich. Tradeoff cryptanalysis of memory-hard functions. Technical report, 2015. available at <http://eprint.iacr.org/2015/227>.
- [15] Itai Dinur, Orr Dunkelman, Nathan Keller, and Adi Shamir. Efficient dissection of composite problems, with applications to cryptanalysis, knapsacks, and combinatorial search problems. In *CRYPTO'12*, volume 7417 of *Lecture Notes in Computer Science*, pages 719–740. Springer, 2012.
- [16] Cynthia Dwork, Andrew Goldberg, and Moni Naor. On memory-bound functions for fighting spam. In *CRYPTO'03*, volume 2729 of *Lecture Notes in Computer Science*, pages 426–444. Springer, 2003.
- [17] Cynthia Dwork and Moni Naor. Pricing via processing or combatting junk mail. In *CRYPTO'92*, volume 740 of *Lecture Notes in Computer Science*, pages 139–147. Springer, 1992.
- [18] Cynthia Dwork, Moni Naor, and Hoeteck Wee. Pebbling and proofs of work. In *CRYPTO'05*, volume 3621 of *Lecture Notes in Computer Science*, pages 37–54. Springer, 2005.
- [19] Stefan Dziembowski, Sebastian Faust, Vladimir Kolmogorov, and Krzysztof Pietrzak. Proofs of space. *IACR Cryptology ePrint Archive 2013/796*. to appear at Crypto'15.
- [20] Lance Fortnow. Time-space tradeoffs for satisfiability. *J. Comput. Syst. Sci.*, 60(2):337–353, 2000.
- [21] B. Giridhar, M. Cieslak, D. Duggal, R. G. Dreslinski, H. Chen, R. Patti, B. Hold, C. Chakrabarti, T. N. Mudge, and D. Blaauw. Exploring DRAM organizations for energy-efficient and resilient exascale memories. In *International Conference for High Performance Computing, Networking, Storage and Analysis 2013*, pages 23–35. ACM, 2013.

- [22] David R. Helman, David A. Bader, and Joseph Jájá. A randomized parallel sorting algorithm with an experimental study. *J. Parallel Distrib. Comput.*, 52(1):1–23, 1998.
- [23] John E. Hopcroft, Wolfgang J. Paul, and Leslie G. Valiant. On time versus space. *J. ACM*, 24(2):332–337, 1977.
- [24] Nick Howgrave-Graham and Antoine Joux. New generic algorithms for hard knapsacks. In *EUROCRYPT’10*, volume 6110 of *Lecture Notes in Computer Science*, pages 235–256. Springer, 2010.
- [25] Markus Jakobsson and Ari Juels. Proofs of work and bread pudding protocols. In Bart Preneel, editor, *Secure Information Networks’99*, volume 152 of *IFIP Conference Proceedings*, pages 258–272. Kluwer, 1999.
- [26] Paul Kirchner. Improved generalized birthday attack. *IACR Cryptology ePrint Archive*, 2011:377, 2011.
- [27] Daniel Lorimer. Momentum – a memory-hard proof-of-work via finding birthday collisions.
- [28] Duane Merrill and Andrew Grimshaw. High performance and scalable radix sorting: A case study of implementing dynamic parallelism for GPU computing. *Parallel Processing Letters*, 21(02):245–272, 2011.
- [29] Lorenz Minder and Alistair Sinclair. The extended k -tree algorithm. In *SODA’09*, pages 586–595. SIAM, 2009.
- [30] Gabriel Nivasch. Cycle detection using a stack. *Inf. Process. Lett.*, 90(3):135–140, 2004.
- [31] Sunoo Park, Krzysztof Pietrzak, Joël Alwen, Georg Fuchsbauer, and Peter Gazi. Spacecoin: A cryptocurrency based on proofs of space. *IACR Cryptology ePrint Archive*, 2015:528, 2015.
- [32] Colin Percival. Stronger key derivation via sequential memory-hard functions. 2009. <http://www.tarsnap.com/scrypt/scrypt.pdf>.
- [33] Nicholas Pippenger. Superconcentrators. *SIAM J. Comput.*, 6(2):298–304, 1977.
- [34] Nadathur Satish, Changkyu Kim, Jatin Chhugani, Anthony D. Nguyen, Victor W. Lee, Daehyun Kim, and Pradeep Dubey. Fast sort on CPUs and GPUs: A case for bandwidth oblivious SIMD sort. In *Proceedings of the 2010 ACM SIGMOD International Conference on Management of Data*, SIGMOD ’10, pages 351–362, 2010.
- [35] Richard Schroeppel and Adi Shamir. A $T = O(2^{n/2}), S = O(2^{n/4})$ algorithm for certain np-complete problems. *SIAM J. Comput.*, 10(3):456–464, 1981.
- [36] Adi Shamir. On the cryptocomplexity of knapsack systems. In *STOC’79*, pages 118–129. ACM, 1979.
- [37] John Tromp. Cuckoo cycle: a memory bound graph-theoretic proof-of-work. Cryptology ePrint Archive, Report 2014/059, 2014. available at <http://eprint.iacr.org/2014/059>, project webpage <https://github.com/tromp/cuckoo>.
- [38] Paul C. van Oorschot and Michael J. Wiener. Parallel collision search with cryptanalytic applications. *J. Cryptology*, 12(1):1–28, 1999.
- [39] David Wagner. A generalized birthday problem. In *CRYPTO’02*, volume 2442 of *Lecture Notes in Computer Science*, pages 288–303. Springer, 2002.
- [40] Xiaochun Ye, Dongrui Fan, Wei Lin, Nan Yuan, and Paolo Ienne. High performance comparison-based sorting algorithm on many-core gpus. In *IPDPS’10*, pages 1–10. IEEE, 2010.

APPENDIX

A. Memoryless collision search for expanding functions and analysis of Momentum

The memoryless collision search algorithm for functions that maps n bits to n bits is well known [30]. It runs in $O(2^{n/2})$ calls to f . However, it is much less efficient if the domain of f is smaller. For instance, suppose that f maps n bits to $2n$ bits (so that only one collision is expected), and we truncate $f(x)$ to certain n bits to iterate it in the Pollard-rho fashion. Then each found collision has only 2^{-n} chance to be the right collision, and we have to rerun the algorithm 2^n times to find the right collision. Therefore, the memoryless algorithm runs in $O(2^{3n/2})$ time in this case.

To explore the full time-memory tradeoff, we turn to an alternative view on this collision search. Finding a collision in an expanding function f mapping n to $n + m$ bits is the same as finding the *golden collision* (i.e. one specific collision) in f truncated to n bits. The golden collision search in a n -bit function has complexity $5 \cdot 2^{3n/2} / \sqrt{w}$ if we have enough memory to store w distinguished points [38], where a distinguished point has size about $2wn$ bits. For very small w , this converges to an algorithm with time complexity $9 \cdot 2^{3n/2}$ for an n -bit function.

Consider now the Momentum PoW [27], where a single collision in F mapping n bits to $2n$ bits is the proof of work. We immediately obtain the following results.

Proposition 8: There is an algorithm that finds the Momentum PoW in $T_0 = M_0 = O(2^n)$ time and memory.

Proposition 9: The time increase in the Momentum PoW is a sublinear function of the memory reduction factor:

$$T(M_0/q) = \sqrt{q}T(M_0); \quad C(q) = \sqrt{q}.$$

Therefore, the Momentum PoW allows a large reduction in the time-area product as the time grows slower than the area decreases.

Note that both propositions can be viewed as special cases ($k = 1$) of Propositions 5 and 6.

B. Generic problem composition

Our primary proposal consists of two independent steps: Wagner’s algorithm \mathcal{A} and the difficulty filter H . We achieved amortization-free and tradeoff steepness just by manipulating \mathcal{A} . Now we consider generic **problem composition** as a tool to get steeper tradeoffs and restrict the number of solutions. It can be used when the algorithm binding method is not applicable.

1) Averaging tradeoffs: Our idea is to cascade two (or more) problems so that the solution to the first is the input to the second. Interestingly, the resulting time-space tradeoff is better for the verifier than either of original tradeoffs.

Formally let \mathcal{P}_1 and \mathcal{P}_2 be two problems with the following properties:

- \mathcal{P}_1 can be solved in time T with memory M and has strong tradeoffs: any memory reduction causes a large computational penalty.
- \mathcal{P}_2 can be solved in time αT and memory M and has a small, fixed number of solutions.

Let us investigate the time-space tradeoffs for $\mathcal{P}_2 \circ \mathcal{P}_1$. Suppose that $C_1(q)T$ is the amortized cost of finding a solution for \mathcal{P}_1 given qM memory, and $\alpha C_2(q)T$ is the amortized cost of finding a solution for \mathcal{P}_2 given qM memory. Then the amortized cost for the composition $\mathcal{P}_2 \circ \mathcal{P}_1$ of problems is

$$T(q) = C_1(q)T + C_2(q)\alpha T.$$

Since for $q = 1$ the time complexity is $(1 + \alpha)T$, the computational penalty is computed as

$$C(q) = \frac{C_1(q)}{1 + \alpha} + \frac{\alpha C_2(q)}{1 + \alpha}$$

For $\alpha \approx 1$ we get $C(q) > \max(C_1(q), C_2(q))/2$, i.e. is at least the half of the largest penalty.

We may want to change the default memory parameters from M to M' , where M' is the maximal memory value such that all memory reductions come at cost above some threshold. This is illustrated in Figure 5 with $\alpha = 1$: both \mathcal{P}_1 and \mathcal{P}_2 need time T and memory M to be solved but have different tradeoffs. It is worth to increase M to M' so that the decrease from M' would be penalized by \mathcal{P}_1 whereas the increase from M' would be penalized by \mathcal{P}_2 .




Fig. 5. Time-memory tradeoff for the composition of hard problems with different tradeoffs.

If α is too large or too small, the tradeoff will be largely determined by one of the two problems. In order to balance them, we suggest iterating the faster problem $1/\alpha$ times.

2) Composition with the generalized birthday problem:

Let us investigate which problems can be composed with the generalized birthday problem. The latter with parameters (n, k) can be solved with $2^{k + \frac{n}{k+1}}$ bytes of memory and time equivalent to $2^{1 + \log k + \frac{n}{k+1}}$ calls to the hash function H . Thus the gap between the time and memory exponents is very close; in other words we need as much time as if we'd hash the entire memory a few times.

For the second problem \mathcal{P}_2 we have to choose the parameters so that the memory requirements 2^l bytes would be very close to the memory needed by Wagner's algorithm, i.e.

$$l \approx k + \frac{n}{k+1}.$$

Secondly, the time complexity must be of the same order of magnitude, i.e. if solving \mathcal{P}_2 with 2^l memory requires $2^{\beta l}$ time, then

$$\beta \approx \frac{1 + \log k + \frac{n}{k+1}}{k + \frac{n}{k+1}}.$$

Therefore, β must be slightly smaller than 1.

We have searched over several hard problems for such ratio, but the β value is often much larger. For instance, the best information set decoding algorithm [10] on random linear codes of length n with the full decoding setting have time complexity $2^{0.1n}$ and space complexity $2^{0.076n}$. If we set $n = 400$, then the memory requirements would be 2^{30} and time would be 2^{40} , which is much higher than 2^{28} time we

get for the generalized birthday problem with 2^{30} memory. A better candidate might be the McEliece decoding parameters, for which the algorithm in [10] obtains the time complexity $2^{0.067n}$ and space complexity $2^{0.59n}$.

Although the information set decoding algorithms are well studied, we found the hard knapsack problem [24], [35] a bit more scrutinized. In the further text we explain the problem, show how to instantiate a proof-of-work with it, and describe existing time-memory tradeoffs.

C. Hard knapsack problem and the proof-of-work based on it

The computational knapsack problem is described as follows. We are given positive numbers a_1, a_2, \dots, a_k, S of length n and have to find $\epsilon_i \in \{0, 1\}$ such that

$$\epsilon_1 a_1 + \epsilon_2 a_2 + \dots + \epsilon_k a_k = S.$$

The problem is known to be NP-hard [24], though for a subset of parameters a fast algorithm exists. If $k < 0.94n$ a solution can be found fast using lattice reduction algorithms. Similar algorithms apply when k is much larger than, i.e. there are multiple solutions. The hardest setting is $k = n$ [24], where the best algorithms are exponential.

In order to construct a proof-of-work protocol, we reformulate the problem similarly to the PoW based on the generalized birthday. We consider the hash function output $H(i)$ as an integer of n bits. We have to find $\epsilon_i \in \{0, 1\}$ such that

$$\epsilon_1 H(1) + \epsilon_2 H(2) + \dots + \epsilon_n H(n) = H(n+1) \pmod{2^n}$$

and $\sum_i \epsilon_i = n/2$.

The best existing algorithms so far have been proposed in [24] and [9]. Though the second algorithm is asymptotically better, for practical n it is outperformed by the algorithm from [24].

The algorithm in [24] works as follows:

- 1) Choose integer $M \approx 2^{0.51n}$ and random $R < M$. Let us denote $\bar{H}(n+1)$ by S .
- 2) Solve the original knapsack modulo M with $S = R$ with a set of solutions L_1 and $\sum_i \epsilon_i = n/4$.
- 3) Solve the original knapsack modulo M with $S = S - R$ with a set of solutions L_2 and $\sum_i \epsilon_i = n/4$.
- 4) Merge two solution sets and filter out pairs of solutions that activate the same ϵ_i .

The smaller knapsacks are solved with the algorithm from [35] so that the $2^{0.31n}$ solutions are produced in time $2^{0.31n}$. Then the solutions are merged with the total complexity $2^{0.337n}$ (corrected value from [9]) and $2^{0.3}$ memory.

The algorithm [35] works similarly: it chooses $M = 2^{n/2}$, then splits the knapsack in two and solves the left part for M and the right part for $S - M$, then merges the two solutions.

Reducing memory by q results into smaller lists L_1, L_2 and thus the quadratic decrease in the success rate. Since the time complexity per iteration also reduces by q , we obtain the simple time-memory tradeoff $TM = 2^{0.63n}$ [9] up to small M . The best memoryless algorithm found so far has the complexity $2^{0.72n}$ [9].

It is reported [9] that the algorithm above runs for $n = 96$ in 15 minutes and requires 1.6 GB of RAM. This memory requirements correspond to $n = 192, k = 7$ in the generalized birthday algorithm, where the time complexity is around 2^{28} hash function calls or about 2 minutes on a single thread. Thus we conclude that these problems couple well, but to equalize the time we would have to run the problem \mathcal{P}_1 several times.