


Proofs of Space

Stefan Dziembowski¹, Sebastian Faust², Vladimir Kolmogorov³,
and Krzysztof Pietrzak³

¹ University of Warsaw, Warszawa, Poland
`stefan@dziembowski.net`

² Ruhr-University Bochum, Bochum, Germany

³ IST Austria, Klosterneuburg, Austria

Abstract. Proofs of work (PoW) have been suggested by Dwork and Naor (Crypto'92) as protection to a shared resource. The basic idea is to ask the service requestor to dedicate some non-trivial amount of computational work to every request. The original applications included prevention of spam and protection against denial of service attacks. More recently, PoWs have been used to prevent double spending in the Bitcoin digital currency system.

In this work, we put forward an alternative concept for PoWs – so-called *proofs of space* (PoS), where a service requestor must dedicate a significant amount of disk space as opposed to computation. We construct secure PoS schemes in the random oracle model (with one additional mild assumption required for the proof to go through), using graphs with high “pebbling complexity” and Merkle hash-trees. We discuss some applications, including follow-up work where a decentralized digital currency scheme called Spacecoin is constructed that uses PoS (instead of wasteful PoW like in Bitcoin) to prevent double spending.

The main technical contribution of this work is the construction of (directed, loop-free) graphs on N vertices with in-degree $O(\log \log N)$ such that even if one places $\Theta(N)$ pebbles on the nodes of the graph, there's a constant fraction of nodes that needs $\Theta(N)$ steps to be pebbled (where in every step one can put a pebble on a node if all its parents have a pebble).

1 Introduction

Proofs of Work (PoW). Dwork and Naor [16] suggested *proofs of work* (PoW) to address the problem of junk emails (aka. Spam). The basic idea is to require that an email be accompanied with some value related to that email that is

S. Dziembowski—Supported by the Foundation for Polish Science WELCOME/2010-4/2 cohort founded within the framework of the EU Innovative Economy (National Cohesion Strategy) Operational Programme.

V. Kolmogorov—VK is supported by European Research Council under the European Unions Seventh Framework Programme (FP7/2007-2013)/ERC grant agreement no 616160.

K. Pietrzak—Research supported by ERC starting grant (259668-PSPC).

moderately hard to compute but which can be verified very efficiently. Such a proof could for example be a value σ such that the hash value $\mathcal{H}(\text{Email}, \sigma)$ starts with t zeros. If we model the hash function \mathcal{H} as a random oracle [8], then the sender must compute an expected 2^t hashes until she finds such a σ .¹ A useful property of this PoW is that there is no speedup when one has to find many proofs, i.e., finding s proofs requires $s2^t$ evaluations. The value t should be chosen such that it is not much of a burden for a party sending out a few emails per day (say, it takes 10 s to compute), but is expensive for a Spammer trying to send millions of messages. Verification on the other hand is extremely efficient, the receiver will accept σ as a PoW for **Email**, if the hash $\mathcal{H}(\text{Email}, \sigma)$ starts with t zeros, i.e., it requires only one evaluation of the hash function. PoWs have many applications, and are in particular used to prevent double spending in the Bitcoin digital currency system [38] which has become widely popular by now.

Despite many great applications, PoWs suffer from certain drawbacks. Firstly, running PoW costs energy – especially if they are used on a massive scale, like in the Bitcoin system. For this reason Bitcoin has even been labelled an “environmental disaster” [3]. Secondly, by using dedicated hardware instead of a general purpose processor, one can solve a PoW at a tiny fraction of the hardware and energy cost, this asymmetry is problematic for several reasons.

Proofs of Space (PoS). From a more abstract point of view, a proof of work is simply a means of showing that one invested a non-trivial amount of effort related to some statement. This general principle also works with resources other than computation like real money in micropayment systems [37] or human attention in CAPTCHAs [12, 46]. In this paper we put forward the concept of *proofs of space* where the resource in question is disk space.

PoS are partially motivated by the observation that users often have a significant amount of free disk space available, and in this case using a PoS is essentially for free. This is in contrast to a PoW: even if one only contributes computation by processors that would otherwise be idle, this will still waste energy which usually does not come for free.

A PoS is a protocol between a prover P and a verifier V which has two distinct phases. After an initialisation phase, the prover P is supposed to store some data \mathcal{F} of size N , whereas V only stores some small piece of information. At any later time point V can initialise a proof execution phase, at the end of which V outputs either reject or accept. We require that V is highly efficient in both phases, whereas P is highly efficient in the execution phase providing he stored and has random access to the data \mathcal{F} .

As an illustrative application for a PoS, suppose that the verifier V is an organization that offers a free email service. To prevent that someone registers a huge number of fake-addresses for spamming, V might require users to dedicate some nontrivial amount of disk space, say 100 GB, for every address registered. Occasionally, V will run a PoS to verify that the user really dedicates this space.

¹ The hashed **Email** should also contain the receiver of the email, and maybe also a timestamp, so that the sender has to search for a fresh σ for each receiver, and also when resending the email at a later point in time.

The simplest solution to prove that one really dedicates the requested space would be a scheme where the verifier V sends a pseudorandom file \mathcal{F} of size 100 GB to the prover P during the initialization phase. Later, V can ask P to send back some bits of \mathcal{F} at random positions, making sure V stores (at least a large fraction of) \mathcal{F} . Unfortunately, with this solution, V has to send a huge 100 GB file to P , which makes this approach pretty much useless in practice.

We require from a PoS that the computation, storage requirement and communication complexity of the verifier V during initialization and execution of the PoS is very small, in particular, at most polylogarithmic in the storage requirement N of the prover P and polynomial in some security parameter γ . In order to achieve small communication complexity, we must let the prover P generate a large file \mathcal{F} locally during an initialization phase, which takes some time I . Note that I must be at least linear in N , our constructions will basically² achieve this lower bound. Later, P and V can run executions of the PoS which will be very cheap for V , and also for P , assuming it has stored \mathcal{F} .

Unfortunately, unlike in the trivial solution (where P sends \mathcal{F} to V), now there is no way we can force a potentially cheating prover \tilde{P} to store \mathcal{F} in-between the initialization and the execution of the PoS: \tilde{P} can delete \mathcal{F} after initialization, and instead only store the (short) communication with V during the initialization phase. Later, before an execution of the PoS, \tilde{P} reconstructs \mathcal{F} (in time I), runs the PoS, and deletes \mathcal{F} again once the proof is over.

We will thus consider a security definition where one requires that a cheating prover \tilde{P} can only make V accept with non-negligible probability if \tilde{P} *either* uses N_0 bits of storage in-between executions of the PoS *or* if \tilde{P} invests time T for every execution. Here $N_0 \leq N$ and $T \leq I$ are parameters, and ideally we want them to be not much smaller than N and I , respectively. Our actual security definition in Sect. 2 is more fine-grained, and besides the storage N_0 that \tilde{P} uses in-between initialization and execution, we also consider a bound N_1 on the total storage used by \tilde{P} during execution (including N_0 , so $N_1 \geq N_0$).

High Level Description of Our Scheme. We described above why the simple idea of having V send a large pseudorandom file \mathcal{F} to P does not give a PoS as the communication complexity is too large. Another simple idea that comes to mind is to let V send a short description of a “randomly behaving” permutation $\pi : \{0, 1\}^n \rightarrow \{0, 1\}^n$ to P , who then stores a table of $N = n2^n$ bits where the entry at position i is $\pi^{-1}(i)$. During the execution phase, V asks for the preimage of a random value y , which P can efficiently provide as the value $\pi^{-1}(y)$ is stored at position y in the table. Unfortunately, this scheme is not a good PoS because of time-memory trade-offs [27] which imply that one can invert a random permutation over N values using only \sqrt{N} time and space.³ For random functions (as opposed to permutations), it’s still possible to invert in time and space $N^{2/3}$. The actual PoS scheme we propose in this paper is based

² One of our constructions will achieve the optimal $I = \Theta(N)$ bound, our second construction achieves $I = O(N \log \log N)$.

³ And initialising this space requires $O(N \log(N))$ time.

on hard to pebble graphs. During the initialisation phase, V sends the description of a hash function to P , who then labels the nodes of a hard to pebble graph using this function. Here the label of a node is computed as the hash of the labels of its children. V then computes a Merkle hash of all the labels, and sends this value to P . In the proof execution phase, V simply asks P to open labels corresponding to some randomly chosen nodes.

Outline and Our Contribution. In this paper we introduce the concept of a PoS, which we formally define in Sect. 2. In Sect. 3 we discuss and motivate the model in which we prove our constructions secure (It is basically the random oracle model, but with an additional assumption). In Sect. 4 we explain how to reduce the security of a simple PoS (with an inefficient verifier) to a graph pebbling game. In Sect. 5 we show how to use hash-trees to make the verifier in the PoS from Sect. 4 efficient. In Sect. 6 we define our final construction and prove its security in Sects. 6.1 and 6.2.

Our proof uses a standard technique for proving lower bounds on the space complexity of computational problems, called *pebbling*. Typically, the lower bounds shown using this method are obtained via the *pebbling games* played on a directed graph. During the game a player can place pebbles on some vertices. The game starts with some pebbles already on the graph. Informally, placing a pebble on a vertex v corresponds to the fact that an algorithm keeps the label of a vertex v in his memory. Removing a pebble from a vertex corresponds therefore to deleting the vertex label from the memory. A pebble can be placed on a vertex v only if the vertices in-going to v have pebbles, which corresponds to the fact that computing v 's label is possible only if the algorithm keeps in his memory the labels of the in-going vertices (in our case this will be achieved by defining the label of v to be a hash of the labels of its in-going vertices). The goal of the player is to pebble a certain vertex of the graph. This technique was used in cryptography already before [17, 19, 20]. For an introduction to the graph pebbling see, e.g., [44].

In Sect. 6.1 we consider two different (infinite families of) graphs with different (and incomparable) pebbling complexities. These graphs then also give PoS schemes with different parameters (cf. Theorem 3). Informally, the construction given in Theorem 1 proves a $\Omega(N/\log N)$ bound on the storage required by a malicious prover. Moreover, no matter how much time he is willing to spend during the execution of the protocol, he is forced to use at least $\Omega(N/\log N)$ storage when executing the protocol. Our second construction from Theorem 2 gives a stronger bound on the storage. In particular, a successful malicious prover either has to dedicate $\Theta(N)$ storage (i.e., almost as much as the N stored by the honest prover) or otherwise it has to use $\Theta(N)$ time with every execution of the PoS (after the initialization is completed). The second construction, whose proof can be found in the full version of this paper [18], is based on superconcentrators, random bipartite expander graphs and on the graphs of Erdős, Graham and Szemerédi [21] is quite involved and is the main technical contribution of our paper.

More Related Work and Applications. Dwork and Naor [16] pioneered the concept of proofs of work as easy-to-check proofs of computational efforts. More concretely, they proposed to use the CPU running time that is required to carry out the proof as a measure of computational effort. In [1] Abadi, Burrows, Manasse and Wobber observed that CPU speeds may differ significantly between different devices and proposed as an alternative measure the number of times the memory is accessed (i.e., the number of cache misses) in order to compute the proof. This approach was formalized and further improved in [2, 15, 17, 47], which use pebbling based techniques. Such memory-hard functions cannot be used as PoS as the memory required to compute and verify the function is the same for provers and verifiers. This is not a problem for memory-hard functions as the here the memory just has to be larger than the cache of a potential prover, whereas in a PoS the storage is the main resource, and will typically be in the range of terabytes.

Originally put forward to counteract spamming, PoWs have a vast number of different applications such as metering web-site access [22], countering denial-of-service attacks [6, 30] and many more [29]. An important application for PoWs are digital currencies, like the recent Bitcoin system [38], or earlier schemes like the Micromint system of Rivest and Shamir [42]. The concept of using bounded resources such as computing power or storage to counteract the so-called “Sybil Attack”, i.e., misuse of services by fake identities, has already mentioned in the work of Douceur [14].

PoW are used in Bitcoin to prevent double spending; honest *miners* must constantly devote more computational power towards solving PoWs than a potential adversary who tries to double spend. This results in a gigantic waste of energy [3] devoted towards keeping Bitcoin secure, and thus also requires some strong form of incentive for the miners to provide this computational power.⁴ Recently a decentralized cryptocurrency called Spacecoin [39] was proposed which uses PoS instead of PoW to prevent double spending. In particular, a miner in Spacecoin who wants to dedicate N bits of disk space towards mining must just run the PoS initialisation phase once, and after that mining is extremely cheap: the miner just runs the PoS execution phase, which means accessing the stored space at a few positions, every few minutes.

A large body of work investigates the concepts of *proofs of storage* and *proofs of retrievability* (cf. [5, 9, 13, 24, 25, 31] and many more). These are proof systems where a verifier sends a file \mathcal{F} to a prover, and later the prover can convince the verifier that it really stored or received the file. As discussed above, proving that one stores a (random) file certainly shows that one dedicates space, but these proof systems are not good PoS because the verifier has to send at least $|\mathcal{F}|$ bits to the verifier, and hence does not satisfy our polylogarithmic bound on the communication complexity.

⁴ There are two mechanisms to incentivise mining: miners who solve a PoW get some fixed reward, this is currently the main incentive, but Bitcoin specifies that this reward will decrease over time. A second mechanism are transactions fees.

Proof of Secure Erasure (PoSE) are related to PoS. Informally, a PoSE allows a space restricted prover to convince a verifier that he has erased its memory of size N . PoSE were suggested by Perito and Tsudik [41], who also proposed a scheme where the verifier sends a random file of size N to the prover, who then answers with a hash of this file. Dziembowski, Kazana and Wichs used graph pebbling to give a scheme with small communication complexity (which moreover is independent of N), but large $\Omega(N^2)$ computation complexity (for prover and verifier). Concurrently, and independently of our work, Karvelas and Kiayias [32], and also Ateniese et al. [4] construct PoSE using graphs with high pebbling complexity. Interestingly, their schemes are basically the scheme one gets when running the initialisation and execution phase of our PoS (as in Eq. (7) in Theorem 3).⁵ References [32] and [4] give a security proof of their construction in the random oracle model, and do not make any additional assumptions as we do. The reason is that to prove that our “collapsed PoS” (as described above) is a PoSE it is sufficient to prove that a prover uses much space *either* during initialisation or during execution. This follows from a (by now fairly standard) “ex post facto” argument as originally used in [17]. We have to prove something much stronger, namely, that the prover needs much space (or at least time) in the execution phase, even if he makes an unbounded amount of queries in the initialisation phase (we will discuss this in more detail in Sect. 3.1). As described above, a PoS (to be precise, a PoS where the execution phase requires large space, not just time) implies a PoSE, but a PoSE does not imply a PoS, nor can it be used for any of the applications mentioned in this paper. The main use-case for PoSE we know of is the one put forward by Perito and Tsudik [41], namely, to verify that some device has erased its memory. A bit unfortunately, Ateniese et al. [4] chose to call the type of protocols they construct also “proofs of space” which led to some confusion in the past.

Finally, let us mention a recent beautiful paper [10] which introduces the concept of “catalytic space”. They prove a surprising result showing that *using* and *erasing* space is not the same relative to some additional space that is filled with random bits and must be in its original state at the end of the computation (i.e., it’s only used as a “catalyst”). Thus, relative to such catalytic space, proving that one has access to some space as in a PoS, and proving that one has erased it, like in PoSE, really can be different things.

2 Defining Proofs of Space

We denote with $(\text{out}_V, \text{out}_P) \leftarrow \langle V(\text{in}_V), P(\text{in}_P) \rangle(\text{in})$ the execution of an interactive protocol between two parties P and V on shared input in , local inputs⁶ inp

⁵ There are some differences, the bounds in [4] are somewhat worse as they use hard-to-pebble graphs with worse parameters, and [32] do not use a Merkle hash-tree to make the computation of the verifier independent of N .

⁶ We use the expression “local input/output” instead the usual “private input/output”, because in our protocols no values will actually be secret. The reason to distinguish between the parties’ inputs is only due to the fact that P ’s input will be very large, whereas we want V to use only small storage.

and in_V , and with local outputs out_V and out_P , respectively. A proof of space (PoS) is given by a pair of interactive random access machines,⁷ a prover P and a verifier V . These parties run the PoS protocol in two phases: a PoS initialization and a PoS execution as defined below. The protocols are executed with respect to some statement id , given as common input (e.g., an email address in the example from the previous section). The identifier id is only required to make sure that P cannot reuse the same space to execute PoS for different statements.

Initialization is an interactive protocol with shared inputs an identifier id , storage bound $N \in \mathbb{N}$ and potentially some other parameters, which we denote with $\text{prm} = (\text{id}, N, \dots)$. The execution of the initialization is denoted by $(\Phi, S) \leftarrow \langle V, P \rangle(\text{prm})$, where Φ is short and S is of size N . V can output the special symbol $\Phi = \perp$, which means that it aborts (this can only be the case if V interacts with a cheating prover).

Execution is an interactive protocol during which P and V have access to the values stored during the initialization phase. The prover P has no output, the verifier V either accepts or rejects.

$$(\{\text{accept}, \text{reject}\}, \emptyset) \leftarrow \langle V(\Phi), P(S) \rangle(\text{prm})$$

In an honest execution the initialization is done once at the setup of the system, e.g., when the user registers with the email service, while the execution can be repeated very efficiently many times without requiring a large amount of computation.

To formally define a proof of space, we introduce the notion of a (N_0, N_1, T) (dishonest) prover \tilde{P} . \tilde{P} 's storage after the initiation phase is bounded by at most N_0 , while during the execution phase its storage is bounded to N_1 and its running time is at most T (here $N_1 \geq N_0$ as the storage during execution contains at least the storage after initialization). We remark that \tilde{P} 's storage and running time is unbounded during the the initialization phase (but, as just mentioned, only N_0 storage is available in-between the initialization and execution phase).

A protocol (P, V) as defined above is a (N_0, N_1, T) -proof of space, if it satisfies the properties of completeness, soundness and efficiency defined below.

Completeness: We will require that for any honest prover P :

$$\Pr[\text{out} = \text{accept} : (\Phi, S) \leftarrow \langle V, P \rangle(\text{prm}), (\text{out}, \emptyset) \leftarrow \langle V(\Phi), P(S) \rangle(\text{prm})] = 1.$$

Note that the probability above is exactly 1, and hence the completeness is perfect.

Soundness: For any (N_0, N_1, T) -adversarial prover \tilde{P} the probability that V accepts is negligible in some statistical security parameter γ . More precisely, we have

$$\Pr[\text{out} = \text{accept} : (\Phi, S) \leftarrow \langle V, \tilde{P} \rangle(\text{prm}), (\text{out}, \emptyset) \leftarrow \langle V(\Phi), \tilde{P}(S) \rangle(\text{prm})] \leq 2^{-\Theta(\gamma)} \quad (1)$$

⁷ In a PoS, we want the prover P to run in time much less than its storage size. For this reason, we must model our parties as random access machines (and not, say Turing machines), where accessing a storage location is assumed to take constant (or at most polylogarithmic) time.

The probability above is taken over the random choice of the public parameters prm and the coins of \tilde{P} and V .⁸

Efficiency: We require the verifier V to be efficient, by which (here and below) we mean at most polylogarithmic in N and polynomial in some security parameter γ . Prover P must be efficient during execution, but can run in time $\text{poly}(N)$ during initialization.⁹

In the soundness definition above, we only consider the case where the PoS is executed only once. This is without loss of generality for PoS where V is stateless (apart from Φ) and holds no secret values, and moreover the honest prover P uses only read access to the storage of size N holding S . The protocols in this paper are of this form. We will sometimes say that a PoS is (N_0, N_1, T) -secure if it is a (N_0, N_1, T) -proof of space.

It is instructive to observe what level of security trivially cannot be achieved by a PoS. Below we use small letters n, t, c to denote values that are small, i.e., polylogarithmic in N and polynomial in a security parameter γ . If the honest prover P is an $(N, N + n, t)$ prover, where t, n denote the time and storage requirements of P during execution, then there exists *no*

1. $(N, N + n, t)$ -PoS, as the honest prover “breaks” the scheme by definition, and
2. $(c, I + t + c, I + t)$ -PoS, where c is the number of bits sent by V to P during initialization. To see this, consider a malicious prover \tilde{P} that runs the initialization like the honest P , but then only stores the messages sent by V during initialization instead of the entire large S . Later, during execution, \tilde{P} can simply emulate the initialization process (in time I) to get back S , and run the normal execution protocol (in time t).

3 The Model

We analyze the security and efficiency of our PoS in the random oracle (RO) model [8], making an additional assumption on the behavior of adversaries, which we define and motivate below. Recall that in the RO model, we assume that all parties (including adversaries) have access to the same random function $\mathcal{H} : \{0, 1\}^* \rightarrow \{0, 1\}^L$. In practice, one must instantiate \mathcal{H} with a real hash function like SHA3. Although security proofs in the RO model are just a heuristic argument for real-world security, and there exist artificial schemes where this heuristic fails [11, 23, 34], the model has proven surprisingly robust in practice.

⁸ Our construction is based on a hash-function \mathcal{H} , which will be part of prm and we require to be collision resistant. As assuming collision resistance for a fixed function is not meaningful [43], we must either assume that the probability of Eq. (1) is over some distribution of identities id (which can then be used as a hash key), or, if we model \mathcal{H} as a random oracle, over the choice of the random oracle.

⁹ As explained in the introduction, P 's running time I during initialization must be at least linear in the size N of the storage. Our construction basically match this $I = \Omega(N)$ lower bound as mentioned in Footnote 2.

Throughout, we fix the output length of our random oracle $\mathcal{H} : \{0, 1\}^* \rightarrow \{0, 1\}^L$ to some $L \in \mathbb{N}$, which should be chosen large enough, so it is infeasible to find collisions. As finding a collision requires roughly $2^{L/2}$ queries, setting $L = 512$ and assuming that the total number of oracle queries during the entire experiment is upper bounded by, say $2^{L/3}$, would be a conservative choice.

3.1 Modeling the Malicious Prover

In this paper, we want to make statements about adversaries (malicious provers) \tilde{P} with access to a random oracle $\mathcal{H} : \{0, 1\}^* \rightarrow \{0, 1\}^L$ and bounded by three parameters N_0, N_1, T . They run in two phases:

1. In a first (initialization) phase, \tilde{P} makes queries¹⁰ $\mathcal{A} = (a_1, \dots, a_q)$ to \mathcal{H} (adaptively, i.e., a_i can be a function of $\mathcal{H}(a_1), \dots, \mathcal{H}(a_{i-1})$). At the end of this phase, \tilde{P} stores a file S of size N_0L bits, and moreover he must commit to a subset of the queries $\mathcal{B} \subseteq \mathcal{A}$ of size N (technically, we'll do this by a Merkle hash-tree).
2. In a second phase, $\tilde{P}(S)$ is asked to output $\mathcal{H}(b)$ for some random $b \in \mathcal{B}$. The malicious prover $\tilde{P}(S)$ is allowed a total number T of oracle queries in this phase, and can use up to N_1L bits of storage (including the N_0L bits for S).

As \mathcal{H} is a random oracle, one cannot compress its uniformly random outputs. In particular, as S is of size N_0L , it cannot encode more than N_0 outputs of \mathcal{H} . We will make the simplifying assumption that we can explicitly state which outputs these are by letting $S_{\mathcal{H}} \subseteq \{0, 1\}^L, |S_{\mathcal{H}}| \leq N_0$ denote the set of all possible outputs $\mathcal{H}(a), a \in \mathcal{A}$ that $\tilde{P}(S)$ can write down during the second phase without explicitly querying \mathcal{H} on input a in the 2nd phase.¹¹ Similarly, the storage bound N_1L during execution implies that \tilde{P} cannot store more than N_1 outputs of \mathcal{H} at any particular time point, and we assume that this set of $\leq N_1$ inputs is well defined at any time-point. The above assumption will allow us to bound the advantage of a malicious prover in terms of a pebbling game.

The fact that we need the additional assumption outlined above and cannot reduce the security of our scheme to the plain random oracle model is a bit unsatisfactory, but unfortunately the standard tools (in particular, the elegant “ex post facto” argument from [17]), typically used to reduce pebbling complexity

¹⁰ The number q of queries in this phase is unbounded, except for the huge exponential $2^{L/3}$ bound on the total number of oracle queries made during the entire experiment by all parties mentioned above.

¹¹ Let us stress that we do not claim that such an $S_{\mathcal{H}}$ exists for every \tilde{P} , one can easily come up with a prover where this is not the case (as we will show below). All we need is that for every (N_0, N_1, T) prover \tilde{P} , there exists another prover \tilde{P}' with (almost) the same parameters and advantage, that obeys our assumption.

An adversary with $N_0 = N_1 = T = 1$ not obeying our assumption is, e.g., a \tilde{P} that makes queries 0 and 1 and stores $S = \mathcal{H}(0) \oplus \mathcal{H}(1)$ in the first phase. In the second phase, $\tilde{P}(S)$ picks a random $b \leftarrow \{0, 1\}$, makes the query b , and can write down $\mathcal{H}(b), \mathcal{H}(1 - b) = S \oplus \mathcal{H}(b)$. Thus, $\tilde{P}(S)$ can write $2 > N_0 = 1$ values $\mathcal{H}(0)$ or $\mathcal{H}(1)$ without querying them in the 2nd phase.

to the number of random oracle queries, cannot be applied in our setting due to the auxiliary information about the random oracle the adversary can store. We believe that a proof exploiting the fact that random oracles are incompressible using techniques developed in [26, 45] can be used to avoid this additional assumption, and we leave this question as interesting future work.

3.2 Storage and Time Complexity

Time Complexity. Throughout, we let the *running time* of honest and adversarial parties be the *number of oracle queries* they make. We also take into account that hashing long messages is more expensive by “charging” k queries for a single query on an input of bit-length $L(k - 1) + 1$ to Lk . Just counting oracle queries is justified by the fact that almost all computation done by honest parties consists of invocations of the random-oracle, thus we do not ignore any computation here. Moreover, ignoring any computation done by adversaries only makes the security proof stronger.

Storage Complexity. Unless mentioned otherwise, the storage of honest and adversarial parties is measured by the number of outputs $y = \mathcal{H}(x)$ stored. The honest prover P will only store such values by construction; for malicious provers \tilde{P} this number is well defined under the assumption from Sect. 3.1.

4 PoS from Graphs with High Pebbling Complexity

The first ingredient of our proof uses graph pebbling. We consider a directed, acyclic graph $G = (V, E)$. The graph has $|V| = N$ vertices, which we label with numbers from the set $[N] = \{1, \dots, N\}$. With every vertex $v \in V$ we associate a value $w(v) \in \{0, 1\}^L$, and extend the domain of w to include also ordered tuples of elements from V in the following way: for $V' = (v_1, \dots, v_n)$ (where $v_i \in V$) we define $w(V') = (w(v_1), \dots, w(v_n))$. Let $\pi(v) = \{v' : (v', v) \in E\}$ denote v 's predecessors (in some arbitrary, but fixed order). The value $w(v)$ of v is computed by applying the random oracle to the index v and the values of its predecessors

$$w(v) = \mathcal{H}(v, w(\pi(v))) . \tag{2}$$

Note that if v is a source, i.e., $\pi(v) = \emptyset$, then $w(v)$ is simply $\mathcal{H}(v)$. Our PoS will be an extension of the simple basic PoS $(P_0, V_0)[G, A]$ from Fig. 1, where A is an efficiently samplable distribution that outputs a subset of the vertices V of $G = (V, E)$. This PoS does not yet satisfy the efficiency requirement from Sect. 2, as the complexity of the verifier needs to be as high as the one of the prover. This is because, in order to perform the check in Step 3 of the execution phase, the verifier needs to compute $w(C)$ himself. In our model, as discussed in Sect. 3.1, the only way a malicious prover $\tilde{P}_0(S)$ can determine $w(v)$ is if $w(v) \in S_{\mathcal{H}}$ is in the encoded set of size at most N_0 , or otherwise by explicitly making the oracle query $\mathcal{H}(v, w(\pi(v)))$ during execution. Note that if $w(i) \notin S_{\mathcal{H}}$ for some

Parameters $\text{prm} = (\text{id}, N, G = (V, E), A)$, where G is a graph on $|V| = N$ vertices and A is an efficiently samplable distribution over V^β (we postpone specifying β as well as the function of id to Sect. 6).

Initialization $(S, \emptyset) \leftarrow \langle P_0, V_0 \rangle(\text{prm})$ where $S = w(V)$.

Execution $(\text{accept/reject}, \emptyset) \leftarrow \langle V(\emptyset), P(S) \rangle(\text{prm})$

1. $V_0(\emptyset)$ samples $C \leftarrow A$ and sends C to P_0 .
2. $P_0(S)$ answers with $A = w(C) \subset S$.
3. $V_0(\emptyset)$ outputs **accept** if $A = w(C)$ and **reject** otherwise.

Fig. 1. The basic PoS $(P_0, V_0)[G, A]$ (with inefficient verifier V_0).

$i \in \pi(v)$, then $\tilde{P}_0(S)$ will have to make even more queries recursively to learn $w(v)$. Hence, in order to prove (N_0, N_1, T) -security of the PoS $(P_0, V_0)[G, A]$ in our idealized model, it suffices to upper bound the advantage of Player 1 in the following pebbling game on $G = (V, E)$:

1. Player 1 puts up to N_0 initial pebbles on the vertices of V .
2. Player 2 samples a subset $C \leftarrow A$ of size α of challenge vertices.
3. Player 1 applies a sequence of up to T steps according to the following rules:
 - (i) it can place a pebble on a vertex v if (1) all its predecessors $u \in \pi(v)$ are pebbled and (2) there are currently less than N_1 vertices pebbled.
 - (ii) it can remove a pebble from any vertex.
4. Player 1 wins if it places pebbles on all vertices of C .

In the pebbling game above, Step 1 corresponds to a malicious prover \tilde{P}_0 choosing the set $S_{\mathcal{H}}$. Step 3 corresponds to \tilde{P}_0 computing values according to the rules in Eq. (2), while obeying the N_1 total storage bound. Putting a pebble corresponds to invoking $y = \mathcal{H}(x)$ and storing the value y . Removing a pebble corresponds to deleting some previously computed y .

5 Efficient Verifiers Using Hash Trees

The PoS described in the previous section does not yet meet our Definition from Sect. 2 as V_0 is not efficient. In this section we describe how to make the verifier efficient, using hash-trees, a standard cryptographic technique introduced by Ralph Merkle [35].

Using Hash Trees for Committing. A hash-tree allows a party P to compute a commitment $\phi \in \{0, 1\}^L$ to N data items $x_1, \dots, x_N \in \{0, 1\}^L$ using $N - 1$ invocations of a hash function $\mathcal{H} : \{0, 1\}^* \rightarrow \{0, 1\}^L$. Later, P can prove to a party holding ϕ what the value of any x_i is, by sending only $L \log N$ bits. For example, for $N = 8$, P commits to x_1, \dots, x_N by hashing the x_i 's in a tree like structure as

$$\phi = \mathcal{H}(\mathcal{H}(\mathcal{H}(x_1, x_2), \mathcal{H}(x_3, x_4)), \mathcal{H}(\mathcal{H}(x_5, x_6), \mathcal{H}(x_7, x_8)))$$

We will denote with $\mathcal{T}^{\mathcal{H}}(x_1, \dots, x_N)$ the $2N - 1$ values of all the nodes (including the N leaves x_i and the root ϕ) of the hash-tree, e.g., for $N = 8$, where we define $x_{ab} = \mathcal{H}(x_a, x_b)$

$$\mathcal{T}^{\mathcal{H}}(x_1, \dots, x_8) = \{x_1, \dots, x_8, x_{12}, x_{34}, x_{56}, x_{78}, x_{1234}, x_{5678}, \phi = x_{12345678}\}$$

The prover \mathbf{P} , in order to later efficiently open any x_i , will store all $2N - 1$ values $\mathcal{T} = \mathcal{T}^{\mathcal{H}}(x_1, \dots, x_N)$, but only send the single root element ϕ to a verifier \mathbf{V} . Later \mathbf{P} can “open” any value x_i to \mathbf{V} by sending x_i and the $\log N$ values, which correspond to the siblings of the nodes that lie on the path from x_i to ϕ , e.g., to open x_3 \mathbf{P} sends x_3 and $\text{open}(\mathcal{T}, 3) = (x_{12}, x_4, x_{5678})$ and the prover checks if

$$\text{vrfy}(\phi, 3, x_3, (x_{12}, x_4, x_{5678})) = \left(\mathcal{H}(x_{12}, \mathcal{H}(x_3, x_4)), x_{56789} \stackrel{?}{=} \phi \right)$$

As indicated above, we denote with $\text{open}(\mathcal{T}, i) \subset \mathcal{T}$ the $\log N$ values \mathbf{P} must send to \mathbf{V} in order to open x_i , and denote with $\text{vrfy}(\phi, i, x_i, o) \rightarrow \{\text{accept}, \text{reject}\}$ the above verification procedure. This scheme is correct, i.e., for ϕ, \mathcal{T} computed as above and any $i \in [N]$, $\text{vrfy}(\phi, i, x_i, \text{open}(\mathcal{T}, i)) = \text{accept}$.

The security property provided by a hash-tree states that it is hard to open any committed value in more than one possible way. This “binding” property can be reduced to the collision resistance of \mathcal{H} : from any $\phi, i, (x, o), (x', o'), x \neq x'$ where $\text{vrfy}(\phi, i, x, o) = \text{vrfy}(\phi, i, x', o') = \text{accept}$, one can efficiently extract a collision $z \neq z', \mathcal{H}(z) = \mathcal{H}(z')$ for \mathcal{H} .

We add an initialization phase to the graph based PoS from Fig. 1, where the prover $\mathbf{P}(\text{prm})$ commits to $x_1 = w(v_1), \dots, x_N = w(v_N)$ by computing a hash tree $\mathcal{T} = \mathcal{T}^{\mathcal{H}}(x_1, \dots, x_N)$ and sending its root ϕ to \mathbf{V} . In the execution phase, the prover must then answer a challenge c not only with the value $x_c = w(c)$, but also open c by sending $(x_c, \text{open}(\mathcal{T}, c))$ which \mathbf{P} can do without any queries to \mathcal{H} as it stored \mathcal{T} .

If a cheating prover $\tilde{\mathbf{P}}(\text{prm})$ sends a correctly computed ϕ during the initialization phase, then during execution $\tilde{\mathbf{P}}(\text{prm}, S)$ can only make $\mathbf{V}(\text{prm}, \phi)$ accept by either answering each challenge c with the correct value $w(c)$, or by breaking the binding property of the hash-tree (and thus the collision resistance of the underlying hash-function).

We are left with the challenge to deal with a prover who might cheat and send a wrongly computed $\tilde{\phi} \neq \phi$ during initialization. Some simple solutions are

- Have \mathbf{V} compute ϕ herself. This is not possible as we want \mathbf{V} 's complexity to be only polylog in N .
- Let \mathbf{P} prove, using a proof system like computationally sound (CS) proofs [36] or universal arguments [7], that ϕ was computed correctly. Although these proof systems do have polylogarithmic complexity for the verifier, and thus formally would meet our efficiency requirement, they rely on the PCP theorem and thus are not really practical.

Dealing with Wrong Commitments. Unless $\tilde{\mathbf{P}}$ breaks the collision resistance of \mathcal{H} , no matter what commitment $\tilde{\phi}$ the prover \mathbf{P} sends to \mathbf{V} , he can later only

open it to some fixed N values which we will denote $\tilde{x}_1, \dots, \tilde{x}_N$.¹² We say that \tilde{x}_i is consistent if

$$\tilde{x}_i = \mathcal{H}(i, \tilde{x}_{i_1}, \dots, \tilde{x}_{i_d}) \text{ where } \pi(i) = \{i_1, \dots, i_d\} \tag{3}$$

Note that if *all* \tilde{x}_i are consistent, then $\tilde{\Phi} = \phi$. We add a second initialization phase to the PoS, where V will check the consistency of α random \tilde{x}_i 's. This can be done by having \tilde{P} open \tilde{x}_i and \tilde{x}_j for all $j \in \pi(i)$. If \tilde{P} passes this check, we can be sure that with high probability a large fraction of the \tilde{x}_i 's is consistent. More concretely, if the number of challenge vertices is $\alpha = \varepsilon t$ for some $\varepsilon > 0$, then \tilde{P} will fail the check with probability $1 - 2^{-\Theta(t)}$ if more than an ε -fraction of the \tilde{x}_i 's are inconsistent.

A cheating \tilde{P} might still pass this phase with high probability with an $\tilde{\Phi}$ where only $1 - \varepsilon$ fraction of the \tilde{x}_i are consistent for some sufficiently small $\varepsilon > 0$. As the inconsistent \tilde{x}_i are not outputs of \mathcal{H} , \tilde{P} can chose their value arbitrarily, e.g., all being 0^L . Now \tilde{P} does not have to store this εN inconsistent values \tilde{x}_j while still knowing them.

In our idealized model as discussed in Sect. 3.1, one can show that this is already all the advantage \tilde{P} gets. We can model an εN fraction of inconsistent \tilde{x}_i 's by slightly augmenting the pebbling game from Sect. 4. Let the pebbles from the original game be *white* pebbles. We now additionally allow player 1 to put εN *red* pebbles (apart from the N_0 white pebbles) on V during step 1. These red pebbles correspond to inconsistent values. The remaining game remains the same, except that player 1 is never allowed to remove red pebbles.

We observe that being allowed to initially put an additional εN red pebbles is no more useful than getting an additional εN white pebbles (as white pebbles are strictly more useful because, unlike red pebbles, they later can be removed.) Translated back to our PoS, in order prove (N_0, N_1, T) -security of our PoS allowing up to εN inconsistent values, it suffices to prove $(N_0 - \varepsilon N, N_1 - \varepsilon N, T)$ -security of the PoS, assuming that the initial commitment is computed honestly, and there are no inconsistent values (and thus no red pebbles in the corresponding game).

6 Our Main Construction

Below we formally define our PoS (P, V) . The common input to P, V are the parameters $\text{prm} = (\text{id}, 2N, \gamma, G, A)$, which contain the identifier $\text{id} \in \{0, 1\}^*$, a storage bound $2N \in \mathbb{N}$ (i.e., $2NL$ bits),¹³ a statistical security parameter γ , the description of a graph $G(V, E)$ on $|V| = N$ vertices and an efficiently samplable distribution A which outputs some “challenge” set $C \subset V$ of size $\alpha = \alpha(\gamma, N)$.

Below \mathcal{H} denotes a hash function, that depends on id : given a hash function $\mathcal{H}'(\cdot)$ (which we will model as a random oracle in the security proof), throughout

¹² Potentially, \tilde{P} cannot open some values at all, but wlog. we assume that it can open every value in exactly one way.

¹³ We set the bound to $2N$, so if we let N denote the number of vertices in the underlying graph, we must store $2N - 1$ values of the hash-tree.

we let $\mathcal{H}(\cdot)$ denote $\mathcal{H}'(\text{id}, \cdot)$. The reason for this is simply so we can assume that the random oracles $\mathcal{H}'(\text{id}, \cdot)$ and $\mathcal{H}'(\text{id}', \cdot)$ used in PoS with different identifiers $\text{id} \neq \text{id}'$ are independent, and thus anything stored for the PoS with identifier id is useless to answer challenges in a PoS with different identifier id' .

Initialization $(\Phi, S) \leftarrow \langle V, P \rangle(\text{prm})$:

1. **P sends V a commitment ϕ to $w(V)$**
 - P computes the values $x_i = w(i)$ for all $i \in V$ as in Eq. (2).
 - P’s output is a hash-tree $S = \mathcal{T}^{\mathcal{H}}(x_1, \dots, x_N)$, which requires $|S| = (2N - 1)L$ bits) as described in Sect. 5.
 - P sends the root $\phi \in S$ to V.
2. **P proves consistency of ϕ for $\alpha = \alpha(\gamma, N)$ random values**
 - V picks a set of challenges $C \leftarrow \Lambda$ where the size of C is α and sends C to P.
 - For all $c \in C$, P opens the value corresponding to c and all its predecessors to V by sending, for all $c \in C$

$$\{(x_i, \text{open}(S, i)) : i \in \{c, \pi(c)\}\}$$

- V verifies that P sends all the required openings, and they are consistent, i.e., for all $c \in C$ the opened values \tilde{x}_c and $\tilde{x}_i, i \in \pi(c) = (i_1, \dots, i_d)$ must satisfy $\tilde{x}_c = \mathcal{H}(c, \tilde{x}_{i_1}, \dots, \tilde{x}_{i_d})$, and the verification of the opened commitments passes. If either check fails, V outputs $\Phi = \perp$ and aborts. Otherwise, V outputs $\Phi = \phi$, and the initialization phase is over.

Execution $(\text{accept/reject}, \emptyset) \leftarrow \langle V(\Phi), P(S) \rangle(\text{prm})$:

- P proves it stores the committed values by opening a random $\beta = \Theta(\gamma)$ subset of them**
- V picks a challenge set $C \subset V$ of size $|C| = \beta$ at random, and sends C to P.
 - P answers with $\{o_c = (x_c, \text{open}(S, c)) : c \in C\}$.
 - V checks for every $c \in C$ if $\text{vrfy}(\Phi, c, o_c) \stackrel{?}{=} \text{accept}$. V outputs **accept** if this is the case and **reject** otherwise.

6.1 Constructions of the Graphs

We consider the following pebbling game, between a player and a challenger, for a directed acyclic graph $G = (V, E)$ and a distribution λ over V .

1. Player puts initial pebbles on some subset $U \subseteq V$ of vertices.
2. Challenger samples a “challenge vertex” $c \in V$ according to λ .
3. Player applies a sequence of steps according to the following rules:
 - (i) it can place a pebble on a vertex v if all its predecessors $u \in \pi(v)$ are pebbled.

(ii) it can remove a pebble from any vertex.

4. Player wins if it places a pebble on c .

Let $S_0 = |U|$ be the number of initial pebbles, S_1 be the total number of used pebbles (or equivalently, the maximum number of pebbles that are present in the graph at any time instance, including initialization), and let T be the number of pebbling steps given in 3i). The definition implies that $S_1 \geq S_0$ and $T \geq S_1 - S_0$. Note, with $S_0 = |V|$ pebbles the player can always achieve time $T = 0$: it can just place initial pebbles on V .

Definition 1. Consider functions $f = f(N, S_0)$ and $g = g(N, S_0, S_1)$. A family of graphs $\{G_N = (V_N, E_N) \mid |V_N| = N \in \mathbb{N}\}$ is said to have *pebbling complexity* $\Omega(f, g)$ if there exist constants $c_1, c_2, \delta > 0$ and distributions λ_N over V_N such that for any player that wins the pebbling game on (G_N, λ_N) (as described above) with probability 1 it holds that

$$\Pr[S_1 \geq c_1 f(N, S_0) \wedge T \geq c_2 g(N, S_0, S_1)] \geq \delta \tag{4}$$

Let $\mathcal{G}(N, d)$ be the set of directed acyclic graphs $G = (V, E)$ with $|V| = N$ vertices and the maximum in-degree at most d . We now state our two main pebbling theorems:

Theorem 1. *There exists an explicit family of graphs $G_N \in \mathcal{G}(N, 2)$ with pebbling complexity*

$$\Omega(N/\log N, 0) \tag{5}$$

In the next theorem we use the *Iverson bracket* notation: $[\phi] = 1$ if statement ϕ is true, and $[\phi] = 0$ otherwise.

Theorem 2. *There exists a family of graphs $G_N \in \mathcal{G}(N, O(\log \log N))$ with pebbling complexity*

$$\Omega(0, [S_0 < \tau N] \cdot \max\{N, N^2/S_1\}) \tag{6}$$

for some constant $\tau \in (0, 1)$. It can be constructed by a randomized algorithm with a polynomial expected running time that produces the desired graph with probability at least $1 - 2^{-\Theta(N/\log N)}$.

Complete proofs of these theorems are given in the full version of this paper [18]; here we give a brief summary of our techniques. For Theorem 1 we use the construction of Paul, Tarjan and Celoni [40], and derive the theorem as a corollary of their Lemma 2. For Theorem 2 we use a new construction which relies on three building blocks: (i) random bipartite graphs $R_{(m)}^d \in \mathcal{G}(2m, d)$ with m inputs and m outputs; (ii) superconcentrator graphs $C_{(m)}$ with m inputs and m outputs; (iii) graphs $D_t = ([t], E_t)$ of Erdős, Graham and Szemerédi [21] with *dense long paths*. These are directed acyclic graphs with t vertices and $\Theta(t \log t)$ edges (of the form (i, j) with $i < j$) that satisfy the following for some constant $\eta \in (0, 1)$ and a sufficiently large t : for any subset $X \subseteq [t]$ of size at most ηt graph D_t contains a path of length at least ηt that avoids X . We show that family D_t can be chosen so that the maximum in-degree is $\Theta(\log t)$. The main component of our construction is graph $\tilde{G}_{(m,t)}^d$ defined as follows:

- Add mt nodes $\tilde{V} = V_1 \cup \dots \cup V_t$ to $\tilde{G}_{(m,t)}^d$ where $|V_1| = \dots = |V_t| = m$. This will be the set of challenges.
- For each edge (i, j) of graph D_t add a copy of graph $R_{(m)}^d$ from V_i to V_j , i.e. identify the inputs of $R_{(m)}^d$ with nodes in V_i (using an arbitrary permutation) and the outputs of $R_{(m)}^d$ with nodes in V_j (again, using an arbitrary permutation).

We set $d = \Theta(1)$, $t = \Theta(\log N)$ and $m = \Theta(N/t)$ (with specific constants), then $\tilde{G}_{(m,t)}^d \in \mathcal{G}(mt, O(\log \log N))$.

Note that a somewhat similar graph was used by Dwork, Naor and Wee [17]. They connect bipartite graphs $R_{(m)}^d$ consecutively, i.e. instead of graph D_t they use a chain graph with t nodes. Dwork et al. give an intuitive argument that removing at most τm nodes from each layer V_1, \dots, V_t (for some constant $\tau < 1$) always leaves a graph which is “well-connected”: informally speaking, many nodes of V_1 are still connected to many nodes of V_t . (We give a formal treatment of their argument in the full version of this paper [18].) However, this does not hold if more than $m = \Theta(N/\log N)$ nodes are allowed to be removed: by placing initial pebbles on, say, the middle layer $V_{t/2}$ one can completely disconnect V_1 from V_t .

In contrast, in our construction removing any $\tau' N$ nodes still leaves a graph which is “well-connected”. Our argument is as follows. If constant τ' is sufficiently small then there can be at most ηt layers with more than τm initial pebbles (for a given constant $\tau < 1$). By the property of D_t , there exists a sufficiently long path P in D_t that avoids those layers. We can thus use the argument above for the subgraph corresponding to P . We split P into three parts of equal size, and show that many nodes in the first part are connected to many nodes in the third part.

In this way we prove that graphs $\tilde{G}_{(m,t)}^d$ have pebbling complexity $\Omega(0, [S_0 < \tau N] \cdot N)$. To get complexity $\Omega(0, [S_0 < \tau N] \cdot \max\{N, N^2/S_1\})$, we add mt extra nodes V_0 and a copy of superconcentrator $C_{(mt)}$ from V_0 to \tilde{V} . We then use a standard “basic lower bound argument” for superconcentrators [33].

Remark 1. As shown in [28], any graph $G \in \mathcal{G}(N, O(1))$ can be entirely pebbled using $S_1 = O(N/\log N)$ pebbles (without any initial pebbles). This implies that expression $N/\log N$ in Theorem 1 cannot be improved upon. Note, this still leaves the possibility of a graph that can be pebbled using $O(N/\log N)$ pebbles only with a large time T (e.g. superpolynomial in N). Examples of such graph for a non-interactive version of the pebble game can be found in [33]. Results stated in [33], however, do not immediately imply a similar bound for our interactive game.

6.2 Putting Things Together

Combining the results and definitions from the previous sections, we can now state our main theorem.

Theorem 3. *In the model from Sect. 3.1, for constants $c_i > 0$, the PoS from Sect. 6 instantiated with the graphs from Theorem 1 is a*

$$(c_1(N/\log N), c_2(N/\log N), \infty)\text{-secure PoS.} \tag{7}$$

Instantiated with the graphs from Theorem 2 it is a

$$(c_3N, \infty, c_4N)\text{-secure PoS.} \tag{8}$$

Efficiency, measured as outlined in Sect. 3.2, is summarized in the table below where γ is the statistical security parameter

	Communication	Computation P	Computation V
PoS Eq. (7) Initialization	$O(\gamma \log^2 N)$	$4N$	$O(\gamma \log^2 N)$
PoS Eq. (7) Execution	$O(\gamma \log N)$	0	$O(\gamma \log N)$
PoS Eq. (8) Initialization	$O(\gamma \log N \log \log N)$	$O(N \log \log N)$	$O(\gamma \log N \log \log N)$
PoS Eq. (8) Execution	$O(\gamma \log N)$	0	$O(\gamma \log N)$

Equation (8) means that a successful cheating prover must either store a file of size $\Omega(N)$ (in L bit blocks) after initialization, or make $\Omega(N)$ invocations to the RO. Equation (7) gives a weaker $\Omega(N/\log N)$ bound, but forces a potential adversary not storing that much after initialization, to use at least $\Omega(N/\log N)$ storage during the execution phase, no matter how much time he is willing to invest. This PoS could be interesting in contexts where one wants to be sure that one talks with a prover who has access to significant memory during execution.

Below we explain how security and efficiency claims in the theorem were derived. We start by analyzing the basic (inefficient verifier) PoS $(P_0, V_0)[G, A]$ from Fig. 1 if instantiated with the graphs from Theorems 1 and 2.

Proposition 1. *For some constants $c_i > 0$, if G_N has pebbling complexity $\Omega(f(N), 0)$ according to Definition 1, then the basic PoS $(P_0, V_0)[G_N, A_N]$ as illustrated in Fig. 1, where the distribution A_N samples $\Theta(\gamma)$ (for a statistical security parameter γ) vertices according to the distribution λ_N from Definition 1, is*

$$(S_0, c_1f(N), \infty)\text{-secure (for any } S_0 \leq c_1f(N)) \tag{9}$$

If G_N has pebbling complexity $(0, g(N, S_0, S_1))$, then for any S_0, S_1 the PoS $(P_0, V_0)[G_N, A_N]$ is

$$(S_0, S_1, c_2g(N, S_0, S_1))\text{-secure.} \tag{10}$$

Above, secure means secure in the model from Sect. 3.1.

(The proof of appears in the full version [18].) Instantiating the above proposition with the graphs G_N from Theorems 1 and 2, we can conclude that the simple (inefficient verifier) PoS $(P_0, V_0)[G_N, A_N]$ is

$$(c_1N/\log N, c_2N/\log N, \infty) \quad \text{and} \quad (S_0, S_1, c_3 \cdot [S_0 \leq \tau N] \cdot \max\{N, N^2/S_1\}) \tag{11}$$

secure, respectively (for constants $c_i > 0$, $0 < \tau < 1$ and $[S_0 < \tau N] = 1$ if $S_0 \leq \tau N$ and 0 otherwise). If we set $S_0 = \lfloor \tau N \rfloor = c_4 N$, the right side of Eq. (11) becomes $(c_4 N, S_1, c_3 \cdot \max\{N, N^2/S_1\})$ and further setting $S_1 = \infty (c_4 N, \infty, c_3 N)$ As explained in Sect. 5, we can make the verifier V_0 efficient during initialization, by giving up on εN in the storage bound. We can choose ε ourselves, but must check $\Theta(\gamma/\varepsilon)$ values for consistency during initialization (for a statistical security parameter γ). For our first PoS, we set $\varepsilon = \frac{c_1}{2 \log N}$ and get with $c_5 = c_1/2$ using $c_2 \geq c_1$

$$\left(\underbrace{c_1 \cdot N / \log N - \varepsilon \cdot N}_{=c_5 N / \log N}, \underbrace{c_2 \cdot N / \log N - \varepsilon \cdot N}_{\geq c_5 N / \log N}, \infty \right)$$

security as claimed in Eq. (7). For the second PoS, we set $\varepsilon = \frac{c_4}{2}$ which gives with $c_6 = c_4/2$

$$\left(\underbrace{c_4 N - \varepsilon N}_{\geq c_6 N}, \infty - \varepsilon N, c_3 N \right)$$

security, as claimed in Eq. (8). Also, note that the PoS described above are PoS as defined in Sect. 6 if instantiated with the graphs from Theorems 1 and 2, respectively.

Efficiency of the PoS Eq.(7). We analyze the efficiency of our PoS, measuring time and storage complexity as outlined in Sect. 3.2. Consider the $(c_1 N / \log N, c_2 N / \log N, \infty)$ -secure construction from Eq. (7). In the first phase of the initialization, P needs roughly $4N = \Theta(N)$ computation: using that the underlying graph has max in-degree 2, computing $w(V)$ according to Eq. (2) requires N hashes on inputs of length at most $2L + \log N \leq 3L$, and P makes an additional $N - 1$ hashes on inputs of length $2L$ to compute the hash-tree. The communication and V 's computation in the first phase of initialization is $\Theta(1)$ (as V just receives the root $\phi \in \{0, 1\}^L$).

During the 2nd phase of the initialization, V will challenge P on α (to be determined) vertices to make sure that with probability $1 - 2^{-\Theta(\gamma)}$, at most an $\varepsilon = \Theta(1/\log N)$ fraction of the \hat{x}_i are inconsistent. As discussed above, for this we have to set $\alpha = \Theta(\gamma \log N)$. Because this PoS is based on a graph with degree 2 (cf. Theorem 1), to check consistency of a \hat{x}_i one just has to open 3 values. Opening the values requires to send $\log N$ values (and the verifier to compute that many hashes). This adds up to an $O(\gamma \log^2 N)$ communication complexity during initialization, V 's computation is of the same order.

During execution, P opens ϕ on $\Theta(\gamma)$ positions, which requires $\Theta(\gamma \log N)$ communication (in L bit blocks), and $\Theta(\gamma \log N)$ computation by V .

Efficiency of the PoS Eq. (8). Analyzing the efficiency of the second PoS is analogous to the first. The main difference is that now the underlying graph has larger degree $O(\log \log N)$ (cf. Theorem 2), and we only need to set $\varepsilon = \Theta(1)$.

References

1. Abadi, M., Burrows, M., Wobber, T.: Moderately hard and memory-bound functions. In: NDSS 2003. The Internet Society, February 2003
2. Alwen, J., Serbinenko, V.: High parallel complexity graphs and memory-hard functions. In: Symposium on Theory of Computing, STOC 2015 (2015)
3. Anderson, N.: Mining Bitcoins takes power, but is it an “environmental disaster”? April 2013. <http://tinyurl.com/cdh95at>
4. Ateniese, G., Bonacina, I., Faonio, A., Galesi, N.: Proofs of space: when space is of the essence. In: Abdalla, M., De Prisco, R. (eds.) SCN 2014. LNCS, vol. 8642, pp. 538–557. Springer, Heidelberg (2014)
5. Ateniese, G., Burns, R.C., Curtmola, R., Herring, J., Kissner, L., Peterson, Z.N.J., Song, D.: Provable data possession at untrusted stores. In: Ning, P., De Capitani di Vimercati, S., Syverson, P.F. (eds.) ACM CCS 2007, pp. 598–609. ACM Press, October 2007
6. Back, A.: Hashcash - a denial of service counter-measure (2002). <http://www.hashcash.org/papers/hashcash.pdf>
7. Barak, B., Goldreich, O.: Universal arguments and their applications. SIAM J. Comput. **38**(5), 1661–1694 (2008)
8. Bellare, M., Rogaway, P.: Random oracles are practical: a paradigm for designing efficient protocols. In: Ashby, V. (ed.) ACM CCS 1993, pp. 62–73. ACM Press, November 1993
9. Bowers, K.D., Juels, A., Oprea, A.: Proofs of retrievability: theory and implementation. In: CCSW, pp. 43–54 (2009)
10. Buhrman, H., Cleve, R., Koucký, M., Loff, B., Speelman, P.: Computing with a full memory: catalytic space. In: Symposium on Theory of Computing, STOC 2014, May 31 - June 03 2014, pp. 857–866, New York, NY, USA (2014)
11. Canetti, R., Goldreich, O., Halevi, S.: The random oracle methodology, revisited (preliminary version). In: 30th ACM STOC, pp. 209–218. ACM Press, May 1998
12. Canetti, R., Halevi, S., Steiner, M.: Mitigating dictionary attacks on password-protected local storage. In: Dwork, C. (ed.) CRYPTO 2006. LNCS, vol. 4117, pp. 160–179. Springer, Heidelberg (2006)
13. Di Pietro, R., Mancini, L.V., Law, Y.W., Etalle, S., Havinga, P.: Lkhw: a directed diffusion-based secure multicast scheme for wireless sensor networks. In: 2003 Proceedings of the International Conference on Parallel Processing Workshops, pp. 397–406 (2003)
14. Douceur, J.R.: The sybil attack. In: Druschel, P., Kaashoek, M.F., Rowstron, A. (eds.) IPTPS 2002. LNCS, vol. 2429, pp. 251–260. Springer, Heidelberg (2002)
15. Dwork, C., Goldberg, A.V., Naor, M.: On memory-bound functions for fighting spam. In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 426–444. Springer, Heidelberg (2003)
16. Dwork, C., Naor, M.: Pricing via processing or combatting junk mail. In: Brickell, E.F. (ed.) CRYPTO 1992. LNCS, vol. 740, pp. 139–147. Springer, Heidelberg (1993)
17. Dwork, C., Naor, M., Wee, H.M.: Pebbling and proofs of work. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 37–54. Springer, Heidelberg (2005)
18. Dziembowski, S., Faust, S., Kolmogorov, V., Pietrzak, K.: Proofs of space. Cryptology ePrint Archive, Report 2013/796 (2013). <http://eprint.iacr.org/2013/796>
19. Dziembowski, S., Kazana, T., Wichs, D.: Key-evolution schemes resilient to space-bounded leakage. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 335–353. Springer, Heidelberg (2011)

20. Dziembowski, S., Kazana, T., Wichs, D.: One-time computable self-erasing functions. In: Ishai, Y. (ed.) TCC 2011. LNCS, vol. 6597, pp. 125–143. Springer, Heidelberg (2011)
21. Erdős, P., Graham, R.L., Szemerédi, E.: On sparse graphs with dense long paths. Technical report STAN-CS-75-504, Stanford University, Computer Science Department (1975)
22. Franklin, K.M., Malkhi, D.: Auditable metering with lightweight security. In: Luby, M., Rolim, J.D.P., Serna, M. (eds.) FC 1997. LNCS, vol. 1318, pp. 151–160. Springer, Heidelberg (1997)
23. Goldwasser, S., Kalai, Y.T.: On the (in)security of the Fiat-Shamir paradigm. In: 44th FOCS, pp. 102–115. IEEE Computer Society Press, October 2003
24. Golle, P., Jarecki, S., Mironov, I.: Cryptographic primitives enforcing communication and storage complexity. In: Blaze, M. (ed.) FC 2002. LNCS, vol. 2357, pp. 120–135. Springer, Heidelberg (2003)
25. Gratzer, V., Naccache, D.: Alien vs. quine. *IEEE Secur. Priv.* **5**(2), 26–31 (2007)
26. Haitner, I., Hoch, J.J., Reingold, O., Segev, G.: Finding collisions in interactive protocols - a tight lower bound on the round complexity of statistically-hiding commitments. In: 48th FOCS, pp.669–679. IEEE Computer Society Press, October 2007
27. Hellman, M.E.: A cryptanalytic time-memory trade-off. *IEEE Trans. Inf. Theory* **26**(4), 401–406 (1980)
28. Hopcroft, J., Paul, W., Valiant, L.: On time versus space. *J. ACM* **24**(2), 332–337 (1977)
29. Jakobsson, M., Juels, A.: Proofs of work and bread pudding protocols. In: Preneel, B. (ed.) Proceedings of the IFIP Conference on Communications and Multimedia Security, vol. 152, pp. 258–272. Kluwer (1999)
30. Juels, A., Brainard, J.G.: Client puzzles: a cryptographic countermeasure against connection depletion attacks. In: NDSS 1999. The Internet Society, February 1999
31. Juels, A., Kaliski Jr., B.S.: Pors: proofs of retrievability for large files. In: Ning, P., De Capitani di Vimercati, S., Syverson, P.F. (eds.) ACM CCS 07, pp. 584–597. ACM Press, October 2007
32. Karvelas, N.P., Kiayias, A.: Efficient proofs of secure erasure. In: Abdalla, M., De Prisco, R. (eds.) SCN 2014. LNCS, vol. 8642, pp. 520–537. Springer, Heidelberg (2014)
33. Lengauer, T., Tarjan, R.E.: Asymptotically tight bounds on time-space trade-offs in a pebble game. *J. ACM* **29**(4), 1087–1130 (1982)
34. Maurer, U.M., Renner, R.S., Holenstein, C.: Indifferentiability, impossibility results on reductions, and applications to the random oracle methodology. In: Naor, M. (ed.) TCC 2004. LNCS, vol. 2951, pp. 21–39. Springer, Heidelberg (2004)
35. Merkle, R.C.: Method of providing digital signatures. US Patent 4309569, 5 January 1982
36. Micali, S.: Computationally sound proofs. *SIAM J. Comput.* **30**(4), 1253–1298 (2000)
37. Micali, S., Rivest, R.L.: Micropayments revisited. In: Preneel, B. (ed.) CT-RSA 2002. LNCS, vol. 2271, pp. 149–163. Springer, Heidelberg (2002)
38. Nakamoto, S.: Bitcoin: a peer-to-peer electronic cash system (2009). <http://bitcoin.org/bitcoin.pdf>
39. Park, S., Pietrzak, K., Alwen, J., Fuchsbauer, G., Gazi, P.: Spacecoin: a cryptocurrency based on proofs of space. *Cryptology ePrint Archive*, Report 2015/528 (2015). <http://eprint.iacr.org/2015/528>

40. Paul, W.J., Tarjan, R.E., Celoni, J.R.: Space bounds for a game on graphs. *Math. Syst. Theory* 10(1), 239–251 (1976–1977)
41. Perito, D., Tsudik, G.: Secure code update for embedded devices via proofs of secure erasure. In: Gritzalis, D., Preneel, B., Theoharidou, M. (eds.) *ESORICS 2010*. LNCS, vol. 6345, pp. 643–662. Springer, Heidelberg (2010)
42. Rivest, R.L., Shamir, A.: Payword and micromint: two simple micropayment schemes. In: *CryptoBytes*, pp. 69–87 (1996)
43. Rogaway, P.: Formalizing human ignorance. In: Nguyễn, P.Q. (ed.) *VIETCRYPT 2006*. LNCS, vol. 4341, pp. 211–228. Springer, Heidelberg (2006)
44. Savage, J.E.: *Models of Computation: Exploring the Power of Computing*, 1st edn. Addison-Wesley Longman Publishing Co. Inc., Boston (1997)
45. Simon, D.R.: Findings collisions on a one-way street: can secure hash functions be based on general assumptions? In: Nyberg, K. (ed.) *EUROCRYPT 1998*. LNCS, vol. 1403, pp. 334–345. Springer, Heidelberg (1998)
46. Von Ahn, L., Blum, M., Hopper, N.J., Langford, J.: CAPTCHA: Using hard AI problems for security. In: Biham, E. (ed.) *EUROCRYPT 2003*. LNCS, vol. 2656, pp. 246–256. Springer, Heidelberg (2003)
47. Waters, B., Juels, A., Halderman, J.A., Felten, E.W.: New client puzzle outsourcing techniques for dos resistance. In: *Proceedings of the 11th ACM Conference on Computer and Communications Security, CCS 2004*, pp. 246–256. ACM, New York (2004)