

# Proofs of Space: When Space is of the Essence

Giuseppe Ateniese<sup>1,2</sup>, Ilario Bonacina<sup>1</sup>, Antonio Faonio<sup>1</sup>, and Nicola Galesi<sup>1</sup>

<sup>1</sup> Sapienza - University of Rome, Italy

{ateniese, bonacina, faonio, galesi}@di.uniroma1.it

<sup>2</sup> Johns Hopkins University, USA

**Abstract.** Proofs of computational effort were devised to control denial of service attacks. Dwork and Naor (CRYPTO '92), for example, proposed to use such proofs to discourage spam. The idea is to couple each email message with a proof of work that demonstrates the sender performed some computational task. A proof of work can be either CPU-bound or memory-bound. In a CPU-bound proof, the prover must compute a CPU-intensive function that is easy to check by the verifier. A memory-bound proof, instead, forces the prover to access the main memory several times, effectively replacing CPU cycles with memory accesses.

In this paper we put forward a new concept dubbed *proof of space*. To compute such a proof, the prover must use a specified amount of space, i.e., we are not interested in the number of accesses to the main memory (as in memory-bound proof of work) but rather on the amount of actual memory the prover must employ to compute the proof. We give a complete and detailed algorithmic description of our model. We develop a comprehensive theoretical analysis which uses combinatorial tools from Complexity Theory (such as pebbling games) which are essential in studying space lower bounds.

**Keywords:** Space Complexity, Proof of Work, Pebbling Game, Random Oracle Model.

## 1 Introduction

Space has a special meaning in Computer Science. It refers to the number of cells of the working tape used by a Turing Machine (TM). While a TM computes a function, it will make several steps (relevant to time complexity) and use a certain number of tape cells (relevant to space complexity).

In [14], Dwork and Naor proposed to employ proof of work (PoW) to discourage spam and, in general, to hinder denial of service attacks. Before any action (such as sending an email), the prover must perform some work and generate a proof of it that can be efficiently verified. Proofs of work are currently being used to implement a publicly verifiable ledger for Bitcoin, where transactions are registered and verified by a community of users to avoid the double-spending problem [28]. The work performed by the prover can be CPU-bound, in which the work represents the number of steps made by a TM, or memory-bound, in which the work represents the times a TM access the working tape. The motivation behind memory-bound PoW is that, while CPU speed may differ significantly among distinct platforms, memory latencies vary much less across machines and may prove to be more equitable and egalitarian. We stress that memory-bound function complexity measures the number of memory accesses and does not take into account the actual amount of memory employed. That is, a TM may read and mark a single cell several times to reach a certain complexity level but it will still end up using only one cell.

In this work we define the notion of *Proof of Space* (PoSpace). PoSpace forces the prover to use at least a specified amount of memory. This means, for instance, that a TM must now use a predetermined number of distinct tape cells to be able to respond to a challenge. We will show that our PoSpace construction is also a memory-bound PoW under the definition provided in [13, 15], while in general a PoW cannot be a PoSpace under our definition. The state of the art memory-bound PoW was described in [15] by Dwork, Goldberg, and Naor. Their scheme requires both the prover and the verifier to store a large table  $T$  but they devised ways to mitigate this problem via either hash trees or public-key signatures. We view PoSpace as a valid alternative to various flavors of PoWs. In PoSpace the spotlight is turned on the amount of space rather than on CPU cycles or memory accesses as in PoWs. In addition, PoSpace solves certain problems where PoW is

not applicable. For instance, we believe PoSpace can be employed in forensic analysis or device attestation to confirm remotely that an embedded device has been successfully wiped. That is, a remote device could be instructed to respond to a wipe command with PoSpace as evidence that its functional memory is now overwritten (cf. [30]).

*Straw Man Solutions.* Memory-bound functions were first introduced by Abadi et al. [1]. In the main construction of memory-bound PoW given in [13], both the prover and the verifier share a large random table  $T$ . The prover must compute a function by making several memory accesses to uniformly random positions in  $T$ . Through a proper tuning of the parameters, it is also possible to force the prover to reserve a specific amount of memory. In another construction, the authors of [13] show that the verifier does not have to store  $T$ . The idea is to sign all pairs  $(i, T[i])$  and then challenge the prover on  $\ell$  positions of  $T$ . The prover will return the  $\ell$  values  $T[i]$  along with an aggregate signature that can be checked by the verifier to ensure the prover is holding the table  $T$ .

We first remark that it is possible to harness recent advances in proof of storage schemes, such as Provable Data Possession (PDP) [4] and (compact) Proof of Retrievability (POR) [33], to reduce the message complexity from  $O(\ell)$  to essentially constant. This solution improves upon the one in [13] and, as long as the initialization phase is performed only once, would meet our *efficiency* requirements for PoSpace. However, proof of storage does not satisfy our definition of PoSpace since the running time of the verifier depends on the size of  $T$ . The only way to avoid linear dependency is to run a PDP-based scheme with spot checking [4], but then the prover either must use more space than required or will not access all the memory locations. Intuitively, the reason why a solution based on proof of storage will not work rests upon the interpretation of what *proof of space* really means. Proof of storage applied to our context satisfies the notion that “*the prover can access space*”. PoSpace instead captures the stronger notion that “*the prover can handle space*”, i.e., the prover possesses, controls, and manipulates space directly. In particular, we distinguish between a prover that can only read memory and a prover that can read and write memory. This is important because, among other things, write operations cannot be parallelized within classical computer architectures. We will provide a formal definition later and make this intuition rigorous.

We also remark that the adversarial model considered in [1, 13] contemplates the existence of a small but fast cache memory that must be saturated to force the prover to dispense with the cache and use traditional RAM memory. Thus, the constructions in [1, 13, 15] do provide a form of proof of space where the space coincides with the cache memory. But, as for proof-of-storage schemes, these schemes satisfy the weaker notion of PoSpace where the prover can only read memory.

*Other Related Work.* A proof of work is also known as a *cryptographic puzzle* in the computer security literature. Puzzles were devised to improve on the proposal by Back [6] and employed to thwart denial of service attacks. In particular, it is important to make them hard to precompute (see [23] and references therein). Waters et al. [36] suggest to outsource the creation of puzzles to an external secure entity. Abliz and Tznati [2] introduce the concept of network-bound puzzles where clients collect tokens from remote servers before querying the service provider. They argue that network latency provides a good solution to the resource disparity problem.

All solutions above deal with proof of effort and cannot be adapted to prove possession of space in the way it is meant and defined in this paper.

Litecoin ([litecoin.org](http://litecoin.org)) is a variant of Bitcoin that employs *scrypt* [29] to certify a public ledger. *scrypt* is defined as a sequential memory-hard function and originally designed as a key derivation function, but it is used as a proof of effort in litecoin to hinder the use of specialized hardware. Technically, *scrypt* is not a memory-bound function as defined in [1] since no lower bound on the memory used by the function can be guaranteed. Thus, it is not even a PoSpace. We also note it requires both the prover and the verifier to dedicate a possibly large amount of memory, while ideally only the prover should reserve and use actual memory (as in our construction to be presented later).

Dziembowski et al. [19] have independently suggested a notion of proof of space. Their original construction generalizes the hash-based PoW of Cash [6] and does not employ the pebbling framework of [15] (cf. Appendix A of [16]). A major overhaul version of their paper later appeared on the IACR Crypto Eprint

repository [16], along with ours [3]. Their new version does use pebbling and adopts techniques similar to ours.

However, there are two main differences between our work and [16] that make their work quite compelling:

1. The definition of proof of space in [16] is stronger in that it allows a two-stage protocol. The first stage can be executed once while the second stage can potentially be executed many times after the first. There are several applications that would benefit from this two-stage notion (see for example the Gmail scenario described in [16]). At the same time, however, this stronger notion is achieved in a more idealized model than the Random Oracle Model. To emphasize the differences between these two notions, we will often refer to ours as a one-stage **PoSpace**.
2. In [16] the authors provide two main constructions. The one that can be compared with our work achieves better pebbling complexity ( $N/(\log(N))$  vs.  $N/(\log(N) \cdot k)$ ). Namely, they are able to remove the extra  $k$  (a security parameter) through a clever technique.

The pebbling framework introduced in [15] has been used successfully in many other contexts (see for example [16–18, 24, 34]). Dziembowski et al. [17] built a leakage resilient key evolution scheme based on the pebbling framework. Their model allows an internal memory-bounded adversary that can control the update operation and leak bounded amounts of information. Smith and Zhang obtain a more efficient scheme [34], specifically, their update operation runs in time quasilinear in the key length, rather than quadratic. The key-evolution scheme can be adapted to obtain a proof of space with efficient communication complexity but the space complexity of the verifier would not satisfy the efficiency constraint of **PoSpace**.

In a recent paper [24], Karvelas and Kiayias provide two efficient constructions for Proof of Secure Erasure (PoSE) as introduced by Perito and Tsudik [30]. Informally, in PoSE the prover must convince a verifier that a certain amount of memory has been *erased*. Both schemes in [24] are ingenious and one of them uses the pebbling framework. PoSE and **PoSpace** are closely related notions but have different requirements as stated in [24]. In addition, in **PoSpace** the prover must show that he can access (read/write) memory while in PoSE, intuitively, there should be the extra and necessary requirement that the memory contents before and after the protocol execution are uncorrelated.

*General ideas behind our protocol.* We cast **PoSpace** in the context of delegation of computation, where a *delegator* outsources to a *worker* the computation of a function on a certain input. Securely delegating computation is a very active area [10, 11, 20, 22] thanks also to the popularity of cloud computing where weak devices use the cloud to compute heavy functions. In the case of **PoSpace**, a function  $f$  is first selected satisfying the property that there exists a space lower bound for any TM computing it. Then, the verifier chooses a random input  $x$  and delegates to the prover the computation of  $f(x)$ .

Specifically, we will turn to the class of functions derived from the “*graph labeling problem*” already used in cryptography (see for example [15, 17, 18]). Important tools in delegation of computation are interactive proof (and argument) systems. The delegation problem, as described in [27], can be solved as follows: The worker computes  $y := f(x)$  and proves using an interactive proof system that indeed  $y = f(x)$ . While interactive proofs with statistical soundness and non-trivial savings in verification time are unlikely to exist [8, 9, 21], Kilian showed [25] that proof systems for **NP** languages with only computational soundness (i.e., argument systems [7]) and *succinctness*<sup>3</sup> do exist. The construction in [25] relies on Merkle Trees and PCP proofs. However, in the context of **PoSpace**, the PCP machinery is an overkill. In fact, the prover does not have to prove the statement  $f(x) = y$ , but only that a space-consuming computation was carried out. Therefore, we replace the PCP verification with an ad-hoc scheme that results in a very efficient overall construction (while PCP-based constructions are notoriously impractical).

*Our Contributions.* We introduce a formal definition of Proof of Space (**PoSpace** Definition 2) capturing the intuitive idea of proving to be able to handle (read / write) at least a specified amount of space. We provide two **PoSpace** protocols in the ROM. In addition, we provide a weaker form of **PoSpace** (**wPoSpace**)

<sup>3</sup> I.e., the total amount of communication and the verification time are both less than their respective values required to transmit and check the NP-witness.

and prove that it is indeed a separate notion. Most of previous work on proof of storage [4, 5, 33] and on memory-bound PoW, as defined in [1, 13, 15], can somehow be adapted to meet this weaker definition but we will not elaborate on this any further in this paper.

*Structure of the paper.* Section 2 contains some preliminary definitions. Section 3 contains the formal definition of PoSpace and provide two PoSpace protocols. Section 5 contains the definition of a *weak* variant of PoSpace that captures read-only provers and a separation result between PoSpace and this weaker variant.

*Open Problems.* It is not clear whether in general PoSpace implies the standard definition of PoW in the sense of [14] (i.e., whether PoSpace is also a PoW). The main obstacle to proving a positive result consists in showing *non-amortizability*. Roughly speaking, non-amortizability means that the “price” of computing the function for  $l$  different inputs is comparable to  $l$  times the “price” of computing the same function once. In the context of PoWs, the “price” is measured in terms of computational time. A PoSpace prover for space  $S$  requires a computational time proportional to  $S$ , thus we would like to reduce an adversary for  $l$  different protocol executions to an adversary for a single execution which spends less than  $S$  computational time. The point is that in order to carry out  $l$  executions, the prover needs  $S$  space and we need to ensure that he spends  $l \times S$  computational time. But space is reusable and it may well happen that something already computed for one instance is reused to compute the proof for another instance. Nevertheless, our second construction does satisfy the definition of PoW. This is simply because we resort to the Random Oracle to ensure that two instances of the protocol are uncorrelated.

## 2 Notations and Preliminary Definitions

*Graph Notation.* Given a directed acyclic graph (DAG)  $G$ , the set of successors and predecessors of  $v$  in  $G$  are respectively  $\Gamma^+(v)$  and  $\Gamma^-(v)$ . We will implicitly assume a topological ordering on its vertex set and a boolean encoding of the vertices respecting that ordering. If  $\Gamma^+(v) = \emptyset$  then  $v$  is an *output-node* and  $\mathbf{T}(G)$  is the set of all output-nodes of  $G$ . Analogously if  $\Gamma^-(v) = \emptyset$  then  $v$  is an *input-node* and  $\mathbf{S}(G)$  is the set of all input-nodes of  $G$ .

*Sampling, Interactive Execution, and Space.* Let  $\mathcal{A}$  be a probabilistic TM. We write  $y \leftarrow \mathcal{A}(x)$  to denote  $y$  sampled from the output of  $\mathcal{A}$  on input  $x$ . Moreover, given  $\sigma \in \{0, 1\}^*$ , we write  $y := \mathcal{A}(x; \sigma)$  to denote the output of  $\mathcal{A}$  on input  $x$  fixing the random coins to be  $\sigma$ . We write PPT for the class of probabilistic polynomial time algorithms. Let  $\mathcal{A}, \mathcal{B}$  be two probabilistic interactive TMs. An interactive joint execution between  $\mathcal{A}, \mathcal{B}$  on common input  $x$  is specified via the next message function notation: let  $b_1 \leftarrow \mathcal{B}(x)$ ,  $a_i \leftarrow \mathcal{A}(x, b_1, \dots, b_i)$ , and  $b_{i+1} \leftarrow \mathcal{B}(x, a_1, \dots, a_i)$ , then  $\langle \mathcal{A}, \mathcal{B} \rangle(x)$  denotes a joint execution of  $\mathcal{A}$  and  $\mathcal{B}$  on common input  $x$ . If the sequence of messages exchanged is  $(b_1, a_1, \dots, b_k, a_k, b_{k+1})$  we say that  $k$  is the number of *rounds* of that joint execution and we denote with  $\langle \mathcal{A}, \mathcal{B} \rangle(x) = b_{k+1}$  the last output of  $\mathcal{B}$  in that joint execution of  $\mathcal{A}$  and  $\mathcal{B}$  on common input  $x$ .

With  $\text{Space}(\mathcal{A}, x)$ , we denote the maximal amount of space used by the deterministic TM  $\mathcal{A}$  on input  $x$  without taking into account the length of the input and output. More formally, we say that  $\mathcal{A}(x)$  has space  $s$  (or  $\text{Space}(\mathcal{A}, x) = s$ ) if and only if at most  $s$  locations on  $\mathcal{A}$ ’s work tapes (excluding the input and the output tapes) are ever written by  $\mathcal{A}$ ’s head during its computation on  $x$ . In the case of probabilistic TM  $\mathcal{A}$ , the function  $\text{Space}(\mathcal{A}, x)$  is a random variable that depends on  $\mathcal{A}$ ’s randomness.

Similarly, given two interactive TMs  $\mathcal{A}$  and  $\mathcal{B}$ , we can define the space occupied by  $\mathcal{A}$  during a joint execution on common input  $x$  as follows. Let  $(b_1, a_1, \dots, b_k, a_k, b_{k+1})$  the sequence of messages exchanged during  $\langle \mathcal{A}, \mathcal{B} \rangle(x)$ , then:

$$\text{Space}_{\mathcal{A}}(\langle \mathcal{A}, \mathcal{B} \rangle, x) := \max_{i=1, \dots, n} \{\text{Space}(\mathcal{A}, x, b_1, \dots, b_i)\}$$

As before, if  $\mathcal{A}$  and  $\mathcal{B}$  are probabilistic interactive TM then the function  $\text{Space}_{\mathcal{A}}(\langle \mathcal{A}, \mathcal{B} \rangle, x)$  is a random variable that depends on the randomness of both  $\mathcal{A}$  and  $\mathcal{B}$ . For simplicity, when the inputs of  $\mathcal{A}$  (or  $\langle \mathcal{A}, \mathcal{B} \rangle$ ) are clear from the context, we write  $\text{Space}(\mathcal{A})$  (or  $\text{Space}_{\mathcal{A}}(\langle \mathcal{A}, \mathcal{B} \rangle)$ ).

*Merkle Trees.* Merkle Trees (MT) are classical tools in cryptography. A MT enables a party to succinctly commit itself to a string  $\mathbf{l} = (l_1, \dots, l_n)$  with  $l_i \in \{0, 1\}^k$ .

The term “succinctly” here means that the MT-commitment has size  $k$  which is independent with respect to the size of  $\mathbf{l}$ . In a later stage, when a party opens the MT-commitment, it is guaranteed that the “opening” can yield only the string  $\mathbf{l}$  committed before (the binding property). Moreover  $\mathbf{l}$  can be succinctly opened location-by-location: the party can open  $l_i$  for any  $i$  giving the certificate attesting the value of  $l_i$ . Here we use the term succinctly to mean that the “opening”, i.e., the certificate, has size  $\log n \times k$  (sub-linear in  $n$ ) with respect to the size of  $\mathbf{l}$ .

Abstractly, we define a Merkle Tree as a tuple of three algorithms  $(\text{Gen}_{\text{CRH}}, \text{MT}, \text{Open})$  where the first algorithm is a key generation algorithm for a collision resistance hash (CRH) function. Suppose that  $\text{Gen}_{\text{CRH}}$  outputs a key  $s$ . Then MT takes as input  $s$  and a sequence of strings  $\mathbf{l}$  and outputs the *commitment*  $C$  for  $\mathbf{l}$ , i.e.,  $C = \text{MT}_s(\mathbf{l})$ . The algorithm Open takes as input the key  $s$ , a sequence of strings  $\mathbf{l}$ , and an index  $i$  (it is denoted as  $\text{Open}_s(\mathbf{l}, i)$ ), and outputs the string  $l_i$  in  $\mathbf{l}$ .

Usually, the term “commitment” refers to a scheme that is both *hiding* (i.e., the receiver cannot infer any knowledge on  $\mathbf{l}$  from the commitment) and *binding*. The MT scheme that we use does not provide the *hiding* property but we still refer to it as a commitment.

A full description of the Merkle Tree is deferred to Appendix B.

*Pebbling Games with wildcards.* The following definition of *pebbling game with wildcards* is a modification of the standard definition of pebbling game that can be found, for instance, in [26]. Given a DAG  $G = (V, E)$ , we say that a sequence  $\mathcal{P} = (P_0, \dots, P_T)$  of subsets of  $V$  is a *pebbling sequence on  $G$  with  $m$  wildcards* if and only if  $P_0 = \emptyset$  and there exists a set  $W \subseteq V$  of size  $m$  such that, for each  $i \in \{1, \dots, T\}$ , exactly one of the following holds:

- $P_i = P_{i-1} \cup \{v\}$  if  $\Gamma^-(v) \subseteq P_{i-1} \cup W$  (*pebbling*) or
- $P_i \subseteq P_{i-1}$  (*unpebbling*).

If a set of vertexes  $\Gamma$  is such that  $\Gamma \subseteq \bigcup_{i=0}^T P_i$ , we say that  $\mathcal{P}$  *pebbles*  $\Gamma$ . If  $\mathcal{P}$  pebbles  $\mathbf{T}(G)$  then we say that  $\mathcal{P}$  is a *pebbling game on  $G$  with  $m$  wildcards*. Moreover we say that the *pebbling time* of  $\mathcal{P}$  is  $T$  and the *pebbling space* of  $\mathcal{P}$  is  $\max_i |P_i|$ . Intuitively, a pebbled node is a node for which we have made some computations. Instead,  $W$  represents *complementary* nodes, for which we have made no computations.

One of the main ingredients for the correctness of our constructions is the Pebbling Theorem [26] that proves that stacks of *superconcentrators* graphs (Pippenger [31]) have an exponential pebbling space-time trade off.

**Theorem 1 (Pebbling Theorem).** *There exists a family of efficiently sampleable directed acyclic graphs  $\{G_{N,k}\}_{k,N \in \mathbb{N}}$  with constant fan-in  $d$ ,  $N$  input nodes,  $N$  output nodes and  $kN \log N$  nodes in total, such that:*

1. *Any pebbling sequence that pebbles  $\Delta$  output nodes has pebbling space at most  $S$  pebbles and  $m$  wildcards, where  $|\Delta| \geq 4S + 2m + 1$ , and pebbling time  $T$  such that*

$$T \geq |\Delta| \left( \frac{N - 2S - m}{2S + m + 1} \right)^k.$$

2. *There exists a pebbling sequence that pebbles the graph  $G_{N,k}$  which has pebbling space  $N + 2$  and needs time  $O(kN \log N)$ .*
3. *For any node  $v \in V(G_{N,k})$ , the incoming nodes of  $v$  and the position of  $v$  in first lexicographic topologically order of  $G_{N,k}$  are computable in  $O(k \log N)$ .*

More details about the construction and the proof of this theorem are given in Appendix A.

*Graph Labeling Problem with faults.* We adopt the paradigm where the action of pebbling a node in a DAG  $G$  is made equivalent to the action of having calculated some labeling on it. This paradigm was introduced in [15] and also recently used in [17, 18]. We make use of a Random Oracle (RO)  $\mathcal{H}$  to build a labeling on  $G$  according to the pebbling rules.

**Definition 1 ( $\mathcal{H}$ -labeling with faults).** *Given a DAG  $G$  with a fixed ordering of the nodes and a Random Oracle function  $\mathcal{H} : \{0, 1\}^* \rightarrow \{0, 1\}^k$ , we say that  $\ell : V(G) \rightarrow \{0, 1\}^k$  is a (partial)  $\mathcal{H}$ -labeling of  $G$  with  $m$  faults if and only if there exists a set  $M \subseteq V(G)$  of size  $m$  such that for each  $v \in V(G) \setminus M$*

$$\ell(v) := \mathcal{H}(v \parallel \ell(v_1) \parallel \dots \parallel \ell(v_d)) \text{ where } \{v_1, \dots, v_d\} = \Gamma^-(v). \quad (1)$$

*Given a label  $\ell$  and a node  $v$ , we say that  $\ell$  well-label  $v$  if only if the equation (1) holds for  $\ell$  and  $v$ .*

Our framework generalize the paradigm of [15] by introducing the concept of “faults”. As it is shown in Section 3.1, dealing with “faults” is necessary because an adversary challenged on a labeling function could cheat by providing an inconsistent label on some nodes (which, indeed, are then referred to as “faults”).

The use of a Random Oracle  $\mathcal{H}$  provides two important benefits to our construction: First, the incompressibility of any output given the input and the evaluation of the function in many different points. Specifically, for any  $x$  and  $x_1, \dots, x_m$ , the value of  $\mathcal{H}(x)$  is uniformly random and independent of  $\mathcal{H}(x_1), \dots, \mathcal{H}(x_m)$ . Therefore, to store  $\mathcal{H}(x)$ , an adversary needs space equal to the minimum between the shortest description of the input and the length of the output. In particular, notice that we do not require that the *entire* function be incompressible (this holds for a Random Oracle but it is trivially false for any *real-world* instantiation of it). Second, in order to  $\mathcal{H}$ -label the graph, any TM must follow a pebble strategy. In particular, to label a node  $v$ , a TM must necessarily calculate and store the label values of all the predecessors  $\ell(v_1), \dots, \ell(v_d)$  of the node  $v$ . If the graph  $G$  needs at least  $S$  pebbles to be pebbled efficiently in a pebbling game, then a TM needs to store at least  $S$  labels (i.e., RO outputs) to compute an  $\mathcal{H}$ -labeling of  $G$ . This general strategy is proven sound in [15] and referred to as the *Labeling Lemma*. In our context, however, we provide  $m$  degrees of freedom and, given a partial  $\mathcal{H}$ -labeling  $\ell$  of  $G$  with  $m$  faults and a  $\mathcal{H}$ -labeling  $\ell'$  of  $G$ , it will likely be the case that  $\ell(v) \neq \ell'(v)$  for each node  $v$  that is a descendant of a not well-labeled node. For this reason, we must state a more general version of the Labeling Lemma in [15].

**Lemma 1 (Labeling with Faults Lemma).** *Consider a DAG  $G$  with degree  $d$ , a TM  $\mathcal{A}$  with advice  $h$ , and a Random Oracle  $\mathcal{H} : \{0, 1\}^* \rightarrow \{0, 1\}^k$  that computes an  $\mathcal{H}$ -labeling  $\ell$  with  $m$  faults of  $G$ . If  $h$  is independent of  $\mathcal{H}$ , with overwhelming probability, there exists a pebbling sequence  $\mathcal{P} = (P_1, \dots, P_T)$  for the DAG  $G$  with  $m$  wildcards having pebbling space  $S$  such that:*

- $S \leq \frac{1}{k} \text{Space}(\mathcal{A}) + d$ ,
- $T \leq (d + 2)\sigma$ , where  $\sigma$  is the number of queries of  $\mathcal{A}$  to  $\mathcal{H}$ .

*In particular  $T$  is a lower bound for the execution time of  $\mathcal{A}$ .*

Before proceeding with the proof of the Lemma, we recall the definitions of min-entropy and conditional average min-entropy as stated in [12]. For any random variable  $X$  we define the min-entropy of  $X$  as

$$\mathbf{H}_\infty(X) := -\log \left( \max_{\mathcal{B}} \Pr[\mathcal{B}(X) = X] \right)$$

and we define the conditional average min-entropy of  $X$  given another random variable  $Y$  as

$$\tilde{\mathbf{H}}_\infty(Y|X) := -\log \left( \max_{\mathcal{B}} \Pr[\mathcal{B}(X) = Y] \right).$$

The following chaining rule for conditional average min-entropy holds [12]:

$$\tilde{\mathbf{H}}_\infty(Y|X) \geq \mathbf{H}_\infty(Y) - |X|.$$

*Proof.* Let  $W$  be the set of not well-labeled nodes by the partial labeling  $\ell$ , in particular  $|W| = m$ . We look at the ordered set of queries  $\mathcal{Q} = \{Q_1, \dots, Q_\tau\}$  of  $\mathcal{A}$  to  $\mathcal{H}$  of the form  $Q_i := (w_i \| \ell(w_i^1) \| \dots \| \ell(w_i^d))$ , where  $\forall i \ w_i \in V(G)$  and  $\forall j \ w_i^j \in \Gamma^-(w_i)$ . Notice that by construction  $\tau \leq \sigma$ .

Let  $S_i = S_{i-1} \cup \{w_i\} \cup \Gamma^-(w_i)$ . As  $\ell$  is a partial  $\mathcal{H}$ -labeling and  $\mathcal{H}$  a RO, we have that with overwhelming probability  $\Gamma^-(w_i) \subseteq S_{i-1} \cup W \cup \mathbf{S}(G)$ . Hence, the sequence  $(S_1, \dots, S_\sigma)$  can be interpolated to a pebbling sequence adding source vertexes when needed. This would make a pebbling sequence  $\tilde{P}$  of length at most  $(d+1)\tau$ . By interpolating  $\tilde{P}$  with suitable unpebbling moves after each  $S_i$  (i.e., removing all unnecessary pebbles following the informal rule “unpebble as soon as possible”), we can obtain a pebbling sequence  $\mathcal{P}$  using minimal space. This is the desired pebbling sequence  $\mathcal{P} = (P_1, \dots, P_T)$  with  $W$  as the set of wildcards. Notice that by construction  $T = (d+2)\tau \leq (d+2)\sigma$ , hence we have only to focus on bounding the pebbling space of  $\mathcal{P}$ .

Let  $\tilde{S} = \text{Space}(\mathcal{A})$ . Suppose for simplicity that  $k$  divides  $\tilde{S}$  and, by contradiction, suppose that there are  $\tilde{S}/k + d + 1$  pebbles at some time  $t^*$  in  $\mathcal{P}$ . This time  $t^*$  should correspond to the time a query to  $\mathcal{H}$  is made. After the unpebble operation, by the rule “unpebble as soon as possible”, there are left at least  $\tilde{S}/k + 1$  pebbles at time  $t^* + 1$ . We look now at the labels of the nodes of  $G$  having at time  $t^* + 1$  one of the  $\tilde{S}/k + 1$  pebbles: by the use of a RO  $\mathcal{H}$ , their labels correspond to a uniformly random string  $L := \ell_1 \| \dots \| \ell_{\tilde{S}/k+1} \in \{0, 1\}^{\tilde{S}+k}$ .

Let  $M$  be the content of the memory of  $\mathcal{A}$  at time  $t^* + 1$ . Note that  $M$  is a random variable in  $\{0, 1\}^{\tilde{S}-1}$ . By the chaining rule of the conditional average min-entropy, we have that

$$\tilde{\mathbf{H}}_\infty(L|M, h) = \tilde{\mathbf{H}}_\infty(L|M) > \mathbf{H}_\infty(L) - S = k, \quad (2)$$

where the first equality comes from the independence of  $\mathcal{H}$  from  $h$ .

We want to upperbound  $\tilde{\mathbf{H}}_\infty(L|M, h)$ . To do this we define  $\mathcal{B}$  a TM that on input  $M$  and  $h$  executes  $\mathcal{A}$  from the state where she was at time  $t^* + 1$ .  $\mathcal{B}$  writes the string  $M$  in  $\mathcal{A}$ 's working tape and the hint  $h$  in  $\mathcal{A}$ 's input tape. In addition,  $\mathcal{B}$  simulates the random oracle  $\mathcal{H}$ . Notice that, as we unpebble as soon as possible then, for any  $1 \leq i \leq \tilde{S}/k + 1$ , the following cases occur:

1.  $\mathcal{A}$  is going to query (as substring of queries to  $\mathcal{H}$ ) the value  $\ell_i$  with overwhelming probability. In fact, if this is not the case then it would be already deleted in  $\mathcal{P}^4$ . The overwhelming probability comes from the fact that  $\mathcal{A}$  could simply guess  $\ell_i$ . Hence  $\mathcal{B}$  can collect all of those labels  $\ell_1, \dots, \ell_{\tilde{S}/k+1}$  by looking at the queries made by  $\mathcal{A}$  to  $\mathcal{H}$ . Moreover,  $\mathcal{B}$  can reply to all the queries to  $\mathcal{H}$  by uniformly choosing strings in  $\{0, 1\}^k$  and storing those values to reply accordingly in the future. This simulation is possible as  $h$  is independent from  $\mathcal{H}$ .
2.  $\mathcal{A}$  will not query  $\mathcal{H}$  with a string which corresponds to a successor of  $\ell_i$  before having queried  $\ell_i$  to  $\mathcal{H}$  with overwhelming probability. Otherwise it would have already been deleted in  $\mathcal{P}$ . The overwhelming probability comes from the fact that there could be a collision in the RO  $\mathcal{H}$ . Hence the simulation of  $\mathcal{H}$  stores the labels  $\ell_1, \dots, \ell_{\tilde{S}/k+1}$  before using them as replies to  $\mathcal{A}$ . Therefore the simulation is indistinguishable from  $\mathcal{A}$ 's point of view. Eventually  $\mathcal{B}$  outputs all the labels collected.

The existence of  $\mathcal{B}$  implies that  $-\log(1 - \nu(k)) \geq \tilde{\mathbf{H}}_\infty(L|M, h)$ , where  $\nu(k)$  is a negligible function. This gives an immediate contradiction with equation (2). Hence the pebbling space of  $\mathcal{P}$  is at most  $\tilde{S}/k + d$ .  $\square$

### 3 Proof-of-Space Protocols

In this section we define the notion of **PoSpace**, then we provide two constructions that meet the definition. We later define a second notion of a *weak* form of **PoSpace** and show a separation result between the two notions. In our definition below, we allow the adversary to access extra information to model the case in which the adversary may outsource storage and computation to an external provider.

<sup>4</sup> This is the reason why we apply this information-theoretical argument to time  $t^* + 1$  and not to time  $t^*$ . In fact it would be the case that at time  $t^*$  we are pebbling a vertex  $w$  but none of the vertexes in  $\Gamma^-(w)$  would be used again.

We model the write permission on the storage by providing a precomputation phase to the adversary. That is, the adversary can use as much space as needed to produce an hint. The hint, that may depend on the public parameters of PoSpace, can be read during the interactive phase (i.e., when the protocol is started). In contrast, wPoSpace does not provide the adversary with a precomputation phase.

**Definition 2 (PoSpace).** Consider  $\Sigma = (\text{Gen}, \text{P}, \text{V})$ , where  $\text{Gen}$  is a PPT TM, and  $\text{P}, \text{V}$  are interactive PPT TMs. Let  $k \in \mathbb{N}$  be a security parameter. Suppose that the following points hold :

- (Completeness) For all  $\text{pk} \in \text{Gen}(1^k)$  and for all  $S \in \mathbb{N}$ ,  $S > k$  it holds  $\langle \text{P}, \text{V} \rangle(\text{pk}, 1^S) = 1$ , time complexity of  $\text{P}$  is  $O(\text{poly}(k, S))$ ;
- (Succinctness) For all  $S \in \mathbb{N}$ ,  $S > k$  the time and space complexity of  $\text{V}$  and the message complexity of  $\langle \text{P}, \text{V} \rangle$ , as functions of  $k$  and  $S$ , are  $O(\text{poly}(k) \cdot \text{poly} \log S)$ ;
- (Soundness) For any PPT adversary  $\mathcal{A}$  and for any PPT TM with advice  $\mathcal{A}'$  such that  $\text{pk} \leftarrow \text{Gen}(1^k)$  and  $h \leftarrow \mathcal{A}'(\text{pk}, 1^S)$ , the following event:

$$\mathbf{Snd}_{\Sigma, S}^{\mathcal{A}, \mathcal{A}'}(k) := \langle \mathcal{A}(h), \text{V} \rangle(\text{pk}, 1^S) = 1 \wedge \text{Space}_{\mathcal{A}}(\langle \mathcal{A}(h), \text{V} \rangle, \text{pk}, 1^S) < S$$

has negligible probability (as a function of  $k$ ) for all  $S \in \mathbb{N}$ ,  $S > k$ .

Then, we say that  $\Sigma = (\text{Gen}, \text{P}, \text{V})$  is a (one-stage) Proof of Space (PoSpace). To be concise, we could say informally that  $\mathcal{A}$  wins when the event  $\exists S \in \mathbb{N} : \mathbf{Snd}_{\Sigma, S}^{\mathcal{A}, \mathcal{A}'}(k)$  occurs.

Notice that in the completeness part we set just a very mild upper bound on the space complexity of  $\text{P}$ . This is done on purpose to allow comparison among different PoSpace protocols. In particular a useful measure on a PoSpace protocol  $(\text{Gen}, \text{P}, \text{V})$  is the following *space gap*: the ratio  $\text{Space}(\text{P}(\text{pk}, 1^S))/S$ .

Notice that  $\mathcal{A}'$  is a space-unbounded PPT TM that models the fact that there might be information that can be efficiently computed that the space-bounded adversary  $\mathcal{A}$  can exploit somehow to compromise PoSpace.

It is easy to see that the PoSpace definition implicitly provides *sequential composability*. In fact, the adversary  $\mathcal{A}'$  gives to  $\mathcal{A}$  a hint which is a function of the public key, therefore the adversary  $\mathcal{A}'$  can compute all the previous executions of the protocol “in his head”.

We provide next a 4-messages PoSpace protocol in the Random Oracle Model (ROM) without any computational assumption. By applying the Fiat-Shamir paradigm to the scheme, we obtain a 2-message non-interactive PoSpace.

### 3.1 A 4-message PoSpace protocol

The protocol  $\Sigma_4 = (\text{Gen}, \text{V}, \text{P})$  is described in Figure 3.1 and it is a 4-message protocol.

The protocol follows in some way Kilian’s construction of argument systems [25]. For any string  $\alpha \in \{0, 1\}^*$ , let  $\mathcal{H}_\alpha(\cdot)$  be defined as  $\mathcal{H}(\alpha \parallel \cdot)$ . The verifier chooses a random  $\alpha$  and asks the prover to build a  $\mathcal{H}_\alpha$ -labeling of graph  $G_{N, k}$ , where  $N$  depends on  $S$ . The purpose of  $\alpha$  is to “reset” the Random Oracle. That is, any previous information about  $\mathcal{H}$  is now useless with overwhelming probability. The labeling provides evidence that the prover has handled at least  $S$  memory cells. The prover then commits the labeling and sends the commitment to the verifier. At this point, the verifier asks the prover to open several random locations in the commitment and then it checks locally the integrity of the labeling. For a commitment  $C$  and for any node that the verifier has challenged, the prover sends what we call a *C-proof* for the node (defined next).

**Definition 3 (C-proof).** Given a DAG  $G$ , a commitment  $C$ , and a Random Oracle  $\mathcal{H}$ , we say that a string  $\pi = (\pi_0, \dots, \pi_d)$  is a *C-proof* for a vertex  $v \in V(G)$  w.r.t.  $\mathcal{H}$  if only if given  $\Gamma_G^-(v) = \{w_1, \dots, w_d\}$  the following points hold:

1.  $\pi_0$  is a *C-opening* for  $v$ , let  $x$  be the value  $\pi_0$  is opening to;
2. for each  $i = 1, \dots, d$ ,  $\pi_i$  is a *C-opening* for  $w_i$ , let  $x_i$  be the value  $\pi_i$  is opening to;



---

**Generator Gen** takes as **input**  $1^k$  and outputs  $\mathbf{pk} := (\{G_{N,k}\}_{N \in \mathbb{N}}, \mathcal{H}, s)$ , where  $\{G_{N,k}\}_{N \in \mathbb{N}}$  is a family of graphs satisfying the Pebbling Theorem (Theorem 1), equipped with the natural lexicographic topological ordering on its vertex set,

$\mathcal{H}$  is a RO and  $s$  is a key for CRH  $H$ .

**Common input:**  $k, S, \mathbf{pk}$

$N := \lceil 4\gamma(d + S/k) + \gamma \rceil$ , where  $d$  is the degree of  $G_{N,k}$  and  $\gamma \in \mathbb{R}$ ,  $\gamma > 1$ .

**Verifier V**

**Prover P**

1. **Pick**  $\alpha \leftarrow \{0, 1\}^k$

$\xrightarrow{\alpha}$

Let  $\ell$  be a  $\mathcal{H}_\alpha$ -labeling of  $G_{N,k}$ , where  $\mathcal{H}_\alpha(\cdot) := \mathcal{H}(\alpha \parallel \cdot)$ .

$\xleftarrow{C}$

**Commit**  $C := \text{MT}_s(\ell)$ .

2. **Pick**  $(v_1, \dots, v_l) \leftarrow V^l$  uniformly at random, where  $l = \lfloor k \ln^2 k \log N \rfloor$  and  $V = V(G_{N,k})$

$\xrightarrow{(v_1, \dots, v_l)}$

For any  $v_i$ ,  $\pi_i := \text{Open}_s(\ell, v_i)$ :  
For any  $v_i^j \in \Gamma^-(v_i)$ :  
 $\pi_i^j := \text{Open}_s(\ell, v_i^j)$   
 $\Pi_i := (\pi_i, \pi_i^1, \dots, \pi_i^d)$

$\xleftarrow{(\Pi_1, \dots, \Pi_l)}$

**Send**  $(\Pi_1, \dots, \Pi_l)$

3. **Check** for any  $i \leq l$  if  $\Pi_i$  is a  $C$ -proof for  $v_i$  wrt the  $\mathcal{H}_\alpha$ -labeling  $\ell$

---

**Fig. 3.1.** The 4-message PoSpace protocol  $\Sigma_4$ .

3.  $x = \mathcal{H}(v \parallel x_1 \parallel \dots \parallel x_d)$ .

We omit  $C$ ,  $G$ , and  $\mathcal{H}$ , when they are clear from the context, by saying that  $\pi$  is simply a proof.

In the definition of  $C$ -proof, the points 1 and 2 refer to the commitment  $C$  while point 3 ensures the integrity of the labeling. Note that the size of  $\pi$  is  $O(kd \log N)$ .

We remark that when TM  $\mathcal{B}$  takes as input a RO  $\mathcal{H}$ , it is intended that  $\mathcal{B}$  has oracle access to  $\mathcal{H}$  and the length of the input of  $\mathcal{B}$  does not take into account the (exponentially long) length of  $\mathcal{H}$ .

**Theorem 2.** *The protocol  $\Sigma_4$  in Figure 3.1 is a (one-stage) PoSpace.*

We start by giving the intuition behind the proof of the Theorem. Completeness is trivial. Succinctness follows easily from point (3) of the Pebbling Theorem (Theorem 1). For Soundness, we first prove that the protocol  $\Sigma_4$  is a Proof-of-Knowledge (PoK) of a partial labeling with “few” faults. By the PoK property, we can extract from a winning adversary a partial labeling with few faults. Thus, with overwhelming probability, the adversary  $\mathcal{A}$  computes a partial labeling. We then exploit the binding property of the Merkle-Tree and show that the adversary has computed the labeling during the round (1) of the protocol. By appropriately fixing the randomness of the protocol, we ensure that the hint  $h$  is independent of  $\mathcal{H}$ , thus meeting all the hypotheses of the Labeling Lemma (Lemma 1). We obtain then a pebbling strategy that has pebbling space and pebbling time roughly upper-bounded by the respective space and time complexity of the adversary  $\mathcal{A}$ . Since the adversary uses strictly less space than  $S$  and  $\mathcal{A}$  is PPT then, by our choice of the parameters, there is a contradiction with the point (1) of the Pebbling Theorem (Theorem 1).

*Proof.* We focus on the *Soundness*. We divide the proof in two parts. First, in the PoK Lemma, we prove that the Protocol  $\Sigma_4$  is a PoK for a partial labeling with “few” faults, and that the partial labeling has been computed during the round (1), i.e., after the verifier  $V$  sent the message  $\alpha$  and before the adversary sent the commitment message  $C$ . Then we combine the PoK Lemma with the Labeling Lemma and the Pebbling Theorem to reach a contradiction assuming that there exists a winning PPT adversary.

We stress that the knowledge extractor doesn’t need to be efficient since we do not rely on any computational assumption.

**Lemma 2 (PoK Lemma).** *Consider the protocol  $\Sigma_4 = (\text{Gen}, V, P)$ , a PPT adversary  $\mathcal{A}$ , a PPT Turing Machine with advice  $\mathcal{A}'$ , a space parameter  $S$ , and a security parameter  $k$ . Sample a random  $\text{pk} \leftarrow \text{Gen}(1^k)$  and let  $h \leftarrow \mathcal{A}'(\text{pk}, 1^S)$ . If  $\Pr \left[ \text{Snd}_{\Sigma_4, S}^{\mathcal{A}, \mathcal{A}'}(k) \right]$  is noticeable (where the probability is taken over the randomness of  $\text{pk}, h$  and all the randomness used during the protocol execution between  $\mathcal{A}$  and the verifier) then, for a noticeable probability over the choice of  $\text{pk}$ , there exist a first verifier message  $\tilde{\alpha}$  and an adversary’s randomness  $\tilde{\rho}$  such that  $\mathcal{A}(h)$  calculates an  $\mathcal{H}_{\tilde{\alpha}}$ -labeling  $\ell$  of  $G_{N, k}$  with  $m$  faults such that the following holds:*

- the hint  $h$  is independent of  $\mathcal{H}_{\tilde{\alpha}}$ ,
- $m = O\left(\frac{N}{\ln k}\right)$  and
- for any well-labeled node  $u$  in  $\ell$  we have  $(\tilde{\alpha} \| u \| \ell(u_1) \| \dots \| \ell(u_d)) \in \mathcal{Q}$ ,

where  $\mathcal{Q}$  is the set of queries to  $\mathcal{H}$  made by  $\mathcal{A}$  when fed with randomness  $\tilde{\rho}$  during round (1) of the protocol  $\Sigma_4$ , and  $\{u_1, \dots, u_d\} = \Gamma^-(u)$ .

*Proof.* Consider the event

$$\mathcal{Q}' := \{\alpha \in \{0, 1\}^k \mid \mathcal{A}'(\text{pk}, 1^S) \text{ made at least a query to } \mathcal{H} \text{ with prefix } \alpha\}.$$

We say that a first message  $\alpha$  is *good* iff

$$\Pr [\langle \mathcal{A}(h), V \rangle(\text{pk}, 1^S) = 1 \mid V \text{ sends } \alpha \text{ as first message}]$$

is noticeable. By Markov’s inequality, the probability that “a uniformly random first message  $\alpha$  is good” is noticeable. On the other hand, the probability that a uniformly random first message  $\alpha$  is in  $\mathcal{Q}'$  is bounded by  $\text{poly}(k, S)/2^k$ . Thus, there exists  $\tilde{\alpha}$  that is both good and not in  $\mathcal{Q}'$ . Since  $\tilde{\alpha} \notin \mathcal{Q}'$  then  $h \leftarrow \mathcal{A}'(\text{pk}, 1^S)$  is independent of  $\mathcal{H}_{\tilde{\alpha}}$ . Let  $V = V(G_{N, k})$  be the vertex set of  $G_{N, k}$  and let  $W_C(\mathbf{v})$  be the following event:

- “ $\mathcal{A}(h, \text{pk}, \tilde{\alpha})$  sends the commitment  $C$  at the end of round (1)
- and makes the verifier  $V$  accept
- when challenged with the vector  $\mathbf{v} \in V^l$  during round (2)”.

The event  $W_C(\mathbf{v})$  depends on the public key  $\text{pk}$ ,  $\mathcal{A}$ ’s randomness, and the challenge  $\mathbf{v} \leftarrow V^l$ , where  $l$  is the length of the message as defined in the protocol  $\Sigma_4$ .

A commitment string sent by  $\mathcal{A}(h, \text{pk}, \tilde{\alpha})$  to the verifier  $V$  at the end of round (1) is a *good commitment* iff  $\Pr[W_C(\mathbf{v})]$  is noticeable. We have that with noticeable probability  $\mathcal{A}(h, \text{pk}, \tilde{\alpha})$  sends a good commitment at the end of round (1).

It might be the case that  $\mathcal{A}(h, \text{pk}, \tilde{\alpha})$ , instead of computing the labeling through the RO queries, simply guesses the correct one. However,  $\mathcal{A}$ ’s guess is correct with probability at most  $2^{-k}$ . Hence, with noticeable probability during round (1):

- $\mathcal{A}(h)$  is not making any guesses and
- $\mathcal{A}(h)$  sends a good commitment.

Thus, there exists a randomness  $\tilde{\rho}$  such that  $\mathcal{A}(h, \text{pk}, \tilde{\alpha})$  satisfies the two properties above when fed with  $\tilde{\rho}$ . To simplify the notation, we call such an adversary  $\tilde{\mathcal{A}}$ .

Thus, there exists a deterministic adversary  $\tilde{\mathcal{A}}$  that, with noticeable probability over  $\mathbf{pk}$  and over the challenge  $\mathbf{v}$ , makes the verifier accept. Let  $C$  be the good commitment sent by  $\tilde{\mathcal{A}}$ . Now we focus on round (2). Let

$$E := \left\{ u \in V \mid \Pr_{\mathbf{pk}, \mathbf{v}} \left[ \begin{array}{c} \tilde{\mathcal{A}} \text{ challenged on } \mathbf{v} \in V^l \\ \text{sends a } C\text{-proof for } u \end{array} \wedge u \in \mathbf{v} \right] < r(k) \right\},$$

where  $r(k) \in \text{poly}(k)$ .

There exists a procedure to compute the set  $E$  (recall that the extraction process doesn't need to be efficient). Notice that by construction the nodes of  $V \setminus E$  can be  $C$ -opened with noticeable probability, hence they  $C$ -open to fixed labels with noticeable probability (otherwise an efficient sampling procedure can invalidate the binding property of the underlying Merkle-Tree).

For each  $u \in V \setminus E$ , there exists a randomness  $\rho_u$  such that  $\tilde{\mathcal{A}}$  sends to the verifier  $\mathbf{V}$  a  $C$ -proof for  $u$  that  $C$ -opens  $u$  to the most probable label. Here, it is assumed that  $\mathcal{A}(h)$  is using  $\rho_u$  as source of randomness during round (2).

If we collect all those values, we obtain an  $\mathcal{H}_{\tilde{\alpha}}$ -labeling with at most  $|E|$  faults.

We now prove the second item of the lemma, i.e., the upper-bound for  $m = |E|$ . Notice that  $m$  is a random variable that depends on  $\mathbf{pk}$ . For any  $\tilde{\mathbf{pk}}$  let  $\tilde{m} := m(\tilde{\mathbf{pk}})$ , the following chain of inequalities hold:

$$\begin{aligned} & \Pr_{\mathbf{v}} \left[ W_C(\mathbf{v}) \mid \mathbf{pk} = \tilde{\mathbf{pk}} \right] \\ &= \Pr_{\mathbf{v}} \left[ W_C(\mathbf{v}) \wedge \forall v_i \in \mathbf{v} \ v_i \notin E \mid \mathbf{pk} = \tilde{\mathbf{pk}} \right] + \\ & \quad + \Pr_{\mathbf{v}} \left[ W_C(\mathbf{v}) \wedge \exists v_i \in \mathbf{v} \ v_i \in E \mid \mathbf{pk} = \tilde{\mathbf{pk}} \right] \\ &\leq \Pr_{\mathbf{v}} \left[ \forall v_i \in \mathbf{v} \ v_i \notin E \mid \mathbf{pk} = \tilde{\mathbf{pk}} \right] + \sum_{u \in E} \Pr_{\mathbf{v}} \left[ W_C(\mathbf{v}) \wedge u \in \mathbf{v} \mid \mathbf{pk} = \tilde{\mathbf{pk}} \right] \\ &\leq \left( 1 - \frac{\tilde{m}}{|V|} \right)^l + \sum_{u \in E} \Pr_{\mathbf{v}} \left[ \begin{array}{c} \tilde{\mathcal{A}} \text{ challenged on } \mathbf{v} \in V^l \\ \text{sends a } C\text{-proof for } u \end{array} \wedge u \in \mathbf{v} \mid \mathbf{pk} = \tilde{\mathbf{pk}} \right] \\ &\leq e^{-\frac{\tilde{m}}{|V|} \cdot l} + \sum_{u \in E} \Pr_{\mathbf{v}} \left[ \begin{array}{c} \tilde{\mathcal{A}} \text{ challenged on } \mathbf{v} \in V^l \\ \text{sends a } C\text{-proof for } u \end{array} \wedge u \in \mathbf{v} \mid \mathbf{pk} = \tilde{\mathbf{pk}} \right] \end{aligned}$$

Taking the expected value over  $\mathbf{pk}$ , we obtain:

$$\begin{aligned} \Pr_{\mathbf{pk}, \mathbf{v}} [W_C(\mathbf{v})] &\leq \mathbb{E}_{\mathbf{pk}} \left[ e^{-\frac{m}{|V|} \cdot l} \right] + \sum_{u \in E} \Pr_{\mathbf{pk}, \mathbf{v}} \left[ \begin{array}{c} \tilde{\mathcal{A}} \text{ challenged on } \mathbf{v} \in V^l \\ \text{sends a } C\text{-proof for } u \end{array} \wedge u \in \mathbf{v} \right] \\ &< \mathbb{E}_{\mathbf{pk}} \left[ e^{-\frac{m}{|V|} \cdot l} \right] + |V| r(k). \end{aligned}$$

There exists a constant  $c \in \mathbb{N}$  such that, for  $k \gg 1$ ,  $\Pr[W_C(\mathbf{v})] \geq k^{-c}$ . Hence, if we select  $r(k) = (|V|2k^c)^{-1}$ , we obtain that

$$\mathbb{E}_{\mathbf{pk}} \left[ e^{-\frac{m}{|V|} \cdot l} \right] > \frac{1}{2k^c}.$$

Via the Paley-Zygmund inequality [32], we have:

$$\Pr_{\mathbf{pk}} \left[ e^{-\frac{m}{|V|} \cdot l} > \frac{1}{4k^c} \right] \geq \frac{1}{8k^{2c}}.$$

Thus, with noticeable probability over  $\mathbf{pk}$ :

$$\frac{m}{|V|} \cdot l < \ln 4 + c \ln k.$$

Notice that, in the graph family  $\{G_{N,k}\}$ ,  $|V| = kN \log N$ . Thus, since  $l = k \ln^2 k \log N$ , we achieve the required upper-bound for  $m$ .

We now prove the third item of the Lemma, i.e., for any assignment  $\rho$  of the randomness of  $\mathcal{A}$  for the round (1) such that (a) the commitment  $C$  is a good commitment and (b)  $\mathcal{A}$  has not made any guess, we have that for any  $u \in V \setminus E$ ,  $Q_u := (\tilde{\alpha} \| u \| \ell(u_1) \| \dots \| \ell(u_d)) \in \mathcal{Q}$ , where  $\mathcal{Q}$  is the set of queries made by  $\mathcal{A}$  to  $\mathcal{H}$  until the end of round (1) and  $\{u_1, \dots, u_d\} = \Gamma^-(u)$ .

If that would not be the case, we can define a reduction that violate the binding property of the Merkle Tree. The probability that a uniformly random assignment of the randomness of  $\mathcal{A}$  has the properties (a) and (b) and a uniformly random first message of  $V$  is good and not in  $\mathcal{Q}'$  is noticeable. Let  $p(k) \in \text{poly}(k)$  be a lower-bound for that probability.

By contradiction suppose there exists a  $u$  such that the query  $Q_u \notin \mathcal{Q}$ . By the choice of  $\tilde{\alpha}$ ,  $Q_u \notin \mathcal{Q}'$  hence it must have been made in the round (2) of the protocol.

Suppose for now that the reduction knows the values  $\ell(u_1), \dots, \ell(u_d)$ , then:

- (a) it can choose an assignment of the randomness  $\rho$  of  $\mathcal{A}$  for the round (1) and a first message  $\alpha \in \{0, 1\}^k$  both uniformly at random. Then, it executes  $\mathcal{A}(h, \text{pk}, 1^S, \alpha; \rho)$  to obtain a commitment  $C$ ;
- (b)  $V$  challenges  $\mathcal{A}$  in round (2) with a uniformly random vector  $\mathbf{v} \in V^l$  containing  $u$ .  $\mathcal{A}$  will make the query  $Q_u$  to  $\mathcal{H}$  and the reduction responds to it with a uniformly random string in  $\{0, 1\}^k$ ;
- (c) it rewinds  $\mathcal{A}$  to the point corresponding to the end of round (1) of the protocol, then runs again point (b).

$\mathcal{A}$  will  $C$ -open the node  $u$  in points (b) and (c) of the reduction with two different values with a probability at least  $p(k) \cdot (r(k)^2 - 2^{-k})$ .

In general the reduction doesn't know the values  $\ell(u_1), \dots, \ell(u_d)$  and it can't figure out which of the queries made by  $\mathcal{A}$  in the round (2) of the protocol is  $Q_u$ . But  $\mathcal{A}$  is PPT so we can find a  $q(k) \in \text{poly}(k, S)$  that upper-bounds the number of queries made by  $\mathcal{A}$  to  $\mathcal{H}$  during round (2).

$\mathcal{A}$  will  $C$ -open the node  $u$ , by guessing the exact point in time  $Q_u$  is queried, with two different values with a probability at least  $q(k)^{-1} \cdot p(k) \cdot (r(k)^2 - 2^{-k})$ . Hence, this reduction violates the binding property of the Merkle Tree.  $\square$

We are ready now to prove the Theorem.

By contradiction, suppose there exists an adversary  $\mathcal{A}$  for the protocol  $\Sigma_4$ . By applying Lemma 2, we extract with noticeable probability a partial  $\mathcal{H}_{\tilde{\alpha}}$ -labeling, ensure that  $\mathcal{A}(h)$  computed that labeling during round (1) and that the hint  $h$  is independent of  $\mathcal{H}_{\tilde{\alpha}}$ . Hence, satisfying the hypotheses of Lemma 1.

This gives us, with overwhelming probability, a pebbling sequence  $\mathcal{P}$  of  $G_{N,k}$  with  $m$  wildcards having pebbling space  $S' = \frac{S}{k} + d$ ,  $m = O(\frac{N}{\ln k})$ . In addition, the pebbling time of  $\mathcal{P}$  is a lower bound for the execution time of  $\mathcal{A}$  during round (1). To apply Theorem 1, we fix  $N$  such that

$$N - l \geq 4S' + 2m + 1, \quad (3)$$

where  $l$  denotes the number of wildcards that are in  $\mathbf{T}(G_{N,k})$ .

Given  $c \in (0, 1)$  and  $\epsilon > 0$ , we have that  $(2 + \epsilon)m + l \leq cN$ . If we find  $N$  such that

$$N \geq 4S' + cN + 1,$$

(hence  $N \geq 4\gamma S' + \gamma$ , where  $\gamma = 1/(1 - c)$ ) then the inequality (3) will follow. (This is main reason why we have set  $N := \lceil 4\gamma S' + \gamma \rceil$  in the Protocol  $\Sigma_4$ .)

By applying Theorem 1, the execution time  $T$  of  $\mathcal{A}$  is such that

$$T \geq (N - l) \left( \frac{N - 2S' - m}{2S' + m + 1} \right)^k \geq c'(1 + \epsilon)^k, \quad (4)$$

where  $c'$  is a constant that is a lower bound for  $N - l$ . The value  $(1 + \epsilon)^k$  derives from the fact that, eventually,

$$N - 2S' - m \geq (1 + \epsilon)(2S' + m + 1).$$

Equation (4) shows that the execution time of  $\mathcal{A}$  is exponential in  $k$ . This is not possible as, by hypothesis,  $\mathcal{A}$  is PPT.  $\blacksquare$

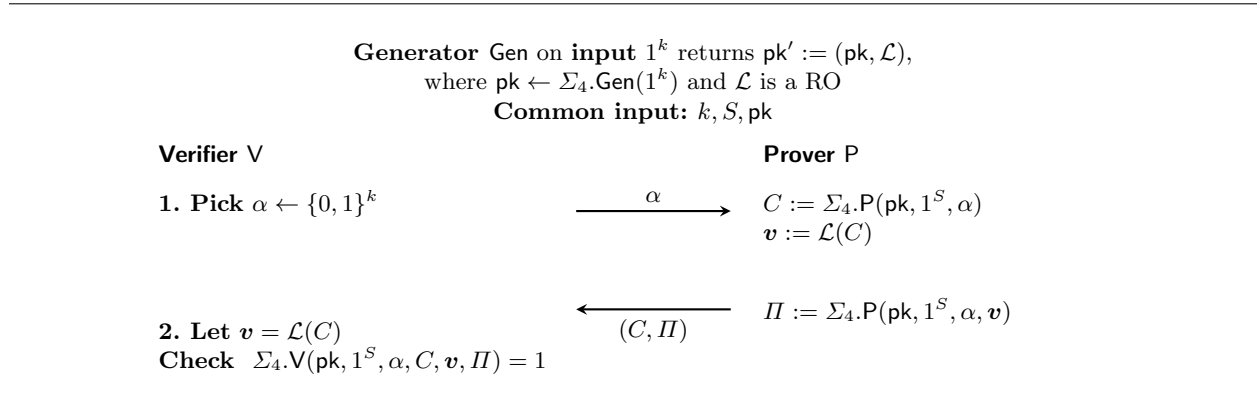
*On the Space Gap of the Protocol.* The prover algorithm can be implemented basically in two ways. The most natural implementation is to first build the labeling for all the nodes in the graph and then apply the Merkle Tree, while keeping in memory the labeling which is reused during the second phase of the prover algorithm. Through this algorithm, the space gap of the prover is  $O(\log S)$  while the time complexity is essentially dominated by the one labeling phase. Another way to implement the prover algorithm is by computing the labeling and the Merkle Tree *simultaneously*. In this way the algorithm can reuse space resulting in a strategy with space gap  $O(1)$ . Note however that, in this implementation, the prover must build the labeling twice (once for the commitment and then during the challenge phase).

## 4 A 2-messages PoSpace protocol

We apply the standard *Fiat-Shamir* paradigm to the PoSpace scheme given in Section 3.1 by using two independent Random Oracles  $\mathcal{H}, \mathcal{L}$ :

- $\mathcal{H} : \{0, 1\}^* \rightarrow \{0, 1\}^k$  is used by the prover for the labeling of the graph;
- $\mathcal{L} : \{0, 1\}^k \rightarrow V^l$  given the commitment  $C$  as input, it yields the second verifier's message (of the protocol  $\Sigma_4$ ).

Let  $(\Sigma_4.\text{Gen}, \Sigma_4.\text{P}, \Sigma_4.\text{V})$  the PoSpace defined in Figure 3.1. We define in Figure 4.1 a 2-messages PoSpace: we call  $\Sigma_2 = (\text{Gen}, \text{P}, \text{V})$  this protocol. Furthermore, let  $\mathcal{H}' : \{0, 1\}^* \rightarrow \{0, 1\}^k$  be a RO then the function  $f(x) := \Sigma_2.\text{P}(\text{pk}, 1^S, \mathcal{H}'(x))$  is a non-interactive PoSpace that satisfies the syntactic definition of PoWs in [14].



**Fig. 4.1.** The 2-messages (one-stage) PoSpace protocol  $\Sigma_2 = (\text{Gen}, \text{P}, \text{V})$ .

**Theorem 3.** *The Protocol  $\Sigma_2$  in Figure 4.1 is a PoSpace.*

*Proof.* *Completeness* and *Succinctness* follow from the fact that  $\Sigma_4$  is a PoSpace as proved in Theorem 2. To prove *Soundness*, we reduce an adversary  $\mathcal{A}$  for the scheme  $\Sigma_2$  to an adversary  $\mathcal{B}$  for the scheme  $\Sigma_4$  (of Figure 4.1). Suppose, by contradiction, there exists an adversary  $\mathcal{A}$ , a hint  $h \leftarrow \mathcal{A}(\text{pk}')$  and a value  $S$  such that for  $k$  large enough:

$$\Pr \left[ \text{Snd}_{\Sigma_2, S}^{\mathcal{A}, \mathcal{A}'}(k) \right] \geq k^{-c}.$$

Let  $q(k)$  be a polynomial that upper-bounds the number of queries made by  $\mathcal{A}$  to the RO  $\mathcal{L}$ . For each  $0 \leq i < q(k)$ , let  $Q_i$  be the  $i$ -th query made by  $\mathcal{A}$  to the oracle  $\mathcal{L}$ . Without loss of generality, we can assume that all the  $Q_i$ 's are different and  $\mathcal{A}$  verifies her output certificate before sending it: if  $\mathcal{A}$  hasn't found a valid certificate, she outputs the empty string. Thus, either there exists an index  $m$  such that

$(C, \Pi) \leftarrow \mathcal{A}(h, \mathbf{pk}', \alpha; \rho)$ , with  $Q_m = C$  and  $\Sigma_4.V(\mathbf{pk}, 1^S, \alpha, C, \mathcal{L}(C), \Pi) = 1$ , or  $\mathcal{A}(h, \mathbf{pk}', \alpha; \rho)$  outputs the empty string. Moreover:

$$\Pr[(C, \Pi) \leftarrow \mathcal{A}(h, \mathbf{pk}', \alpha) \wedge Q_m = C] \geq k^{-c} \cdot q(x)^{-1}.$$

Notice that the probability above states that the outputs of  $\mathcal{A}$  will be the same value asked by  $\mathcal{A}$  to  $\mathcal{L}$  at the  $m$ -th call. However, the value  $Q_m$  is a random variable that may depend on all the randomness used until the  $m$ -th call. We divide the randomness used by  $\mathcal{A}$  into two parts,  $\rho_1$  and  $\rho_2$ :  $\rho_1$  is the randomness used by  $\mathcal{A}$  before she has made the  $m$ -th query to the oracle  $\mathcal{L}$ ,  $\rho_2$  is the remaining randomness.

By an averaging argument, there exists a setting of  $\rho_1$  such that the following equation holds

$$\Pr_{\rho_2, \mathcal{H}, \mathcal{L}}[(C, \Pi) := \mathcal{A}(h, \mathbf{pk}', \alpha; \rho_1 \| \rho_2) \wedge Q_m = C] \geq k^{-c} \cdot q(x)^{-1} \quad (5)$$

Now the value  $Q_m$  is a random variable that depends only on the randomness of the challenge  $\alpha$ , the ROs  $\mathcal{H}$  and  $\mathcal{L}$ , and  $\rho_1$ . Therefore, by repeating many times the TM  $\mathcal{A}(h, \mathbf{pk}', \alpha; \rho_1 \| \rho_2)$  for a uniformly random  $\rho_2$  using the same oracles  $\mathcal{H}$ ,  $\mathcal{L}$ , we'll get the same value  $Q_m$  as the  $m$ -th query to  $\mathcal{L}$ . We define the PPT TM with advice  $\mathcal{B}'$  that on input  $\mathbf{pk}$  and advice  $\rho_1, m$  outputs the hint  $\tilde{h} := (\mathcal{A}'(\mathbf{pk}, \mathcal{L}), \mathcal{Q}', m, \rho_1)$ , where  $\mathcal{Q}'$  is the set of RO queries made by  $\mathcal{A}'$  to  $\mathcal{H}$  and  $\mathcal{L}$ .

**Define  $\mathcal{B}(\tilde{h}, \mathbf{pk}, 1^S)$ :**

- **Wait** for the message  $\alpha$ ;
- Compute  $\mathcal{A}(\tilde{h}, \mathbf{pk}, \alpha; \rho_1)$  and reply to RO queries from  $\mathcal{A}$  via  $\mathcal{H}$  or  $\mathcal{L}$  respectively;  
  **until** she makes the  $m$ -th RO query to  $\mathcal{L}$ . Let  $Q_m$  be such a query;
- **Pause** the machine  $\mathcal{A}$ ;
- **if**  $Q_m \in \mathcal{Q}'$  **abort**.
- **Send**  $Q_m$  to  $\Sigma_4.V$ ;
- **Wait** for the message  $\mathbf{v}$  from  $\Sigma_4.V$ ;
- choose  $\rho_2$  uniformly at random;
- **Resume**  $\mathcal{A}$  and reply to the query  $Q_m$  with  $\mathbf{v}$  and use as randomness  $\rho_2$ ;
- **If**  $\mathcal{A}$  outputs the empty string, **send** the empty string to  $\Sigma_4.V$ ;  
  **else** let  $(C, \Pi)$  be the output of  $\mathcal{A}$ ;
- **if**  $C = Q_m$  **then send**  $\Pi$  to  $\Sigma_4.V$   
  **else send** the empty string to  $\Sigma_4.V$ .

First, we show that the value  $Q_m$  has high min-entropy. Since  $|\mathcal{Q}'| \in \text{poly}(k, S)$ , this implies that  $\mathcal{B}$  will abort with negligible probability.

By the equation (5), there exists an  $\mathcal{H}_\alpha$ -labeling  $\ell$  with  $m = O(N/\ln k)$  faults such that  $Q_m = \text{MT}_s(\ell)$ . Let  $\mathcal{R}'$  be the set of RO replies made by  $\mathcal{H}$  and  $\mathcal{L}$  to the RO queries of  $\mathcal{A}'$ . As  $\alpha$  is uniformly random in  $\{0, 1\}^k$  then

$$\tilde{\mathbf{H}}_\infty(\ell | \mathcal{R}', \alpha) \geq (|V(G_{N,k})| - m - 1) \cdot k.$$

In fact only with negligible probability in  $k$  the set  $\mathcal{Q}'$  contains queries to  $\mathcal{H}$  with prefix  $\alpha$ . The min-entropy of  $Q_m$  is  $\omega(\log k)$  because  $\text{MT}$  is collision resistant. The probability that  $Q_m$  is in  $\mathcal{Q}'$  is then negligible.

When  $\mathcal{B}$  doesn't abort the simulation of the protocol,  $\Sigma_2$  is perfect. Also notice that  $\mathcal{B}$  needs space  $S' = S + O(\log(\text{poly}(k)))$  where  $S$  is from the simulations of  $\mathcal{A}$  and  $O(\log \text{poly}(k))$  from the auxiliary indexes. Moreover, if the event in equation (5) holds, then  $\mathcal{B}$  will not send the empty string to  $\Sigma_4.V$ . Hence, it follows that the adversary  $\mathcal{B}$ , with access to the hint function  $\tilde{h}$  and using space  $S'$  for  $k$  large enough, is such that:

$$\Pr \left[ \text{Snd}_{\Sigma_4, S'}^{\mathcal{B}, \mathcal{B}'}(k) \right] \geq k^{-c} \cdot q(k)^{-1} - \text{negl}(k).$$

Therefore, we reach the contradiction that  $\mathcal{B}$  is a winning adversary against protocol  $\Sigma_4$  which was proved to be a PoSpace (Theorem 2). ■

## 5 Weak Proof of Space

The concept of proof of space can lead to multiple interpretations. The main interpretation formalized in the **PoSpace** definition requires that the prover can *handle* (i.e., read/write) space. In this section we provide a definition for a weaker alternative of **PoSpace** we call *weak* Proof of Space (**wPoSpace**). This captures the property that the prover can just access space and formalizes what could effectively be achieved by properly adapting previous work on proof of storage [4, 5, 33] and on memory-bound PoW as defined in [1, 13] (where the adversary model contemplates the existence of cache memory). The definition of **wPoSpace** is similar to the Definition of **PoSpace** (Definition 2) (the only change is in the Soundness part). We will provide a protocol which is a **wPoSpace** but not a **PoSpace**, hence **wPoSpace** is a strictly weaker notion than **PoSpace**.

**Definition 4 (wPoSpace).** Consider  $\Sigma = (\text{Gen}, P, V)$  where  $\text{Gen}$  is a PPT TM, and  $P, V$  are interactive PPT TMs. Let  $k \in \mathbb{N}$  be a security parameter and  $\text{pk} \leftarrow \text{Gen}(1^k)$ . Suppose that the following points hold for all  $S \in \mathbb{N}$ ,  $S > k$ :

(Completeness) and (Succinctness) the same as in the **PoSpace** definition.

(weak-Soundness) For any PPT adversary  $\mathcal{A}$ , the following event:

$$\mathbf{wSnd}_{\Sigma, S}^{\mathcal{A}}(k) := \langle \mathcal{A}, V \rangle(\text{pk}, 1^S) = 1 \wedge \text{Space}_{\mathcal{A}}(\langle \mathcal{A}, V \rangle, \text{pk}, 1^S) < S,$$

has negligible probability (as a function of  $k$ ).

Then we say that  $\Sigma$  is a (one-stage) Weak Proof of Space (**wPoSpace**).

In the Soundness of **PoSpace**, the adversary can take advantage of an unbounded space machine which is then unavailable during the protocol execution. In the Soundness of **wPoSpace**, instead, this is disallowed. Notice that a **wPoSpace**'s adversary can perform some precomputation before the execution of the protocol (i.e., before sending/receiving the first message), however, such a precomputation cannot exceed the space bound given.

**Theorem 4.** There exists a protocol which is a **wPoSpace** but not a **PoSpace**.

*Proof.* We start by providing a protocol which is not a **PoSpace**. Consider the protocol  $\Sigma_4$  in Figure 3.1 where the first message  $\alpha$  sent by  $V$  is always the same, say  $0^k$ . We call  $\Sigma_3$  this modified version of  $\Sigma_4$ . The protocol  $\Sigma_3$  is not a **PoSpace**. For any  $k$  and  $S \in \mathbb{N}$ , consider the hint  $h$  which is the  $\mathcal{H}_{0^k}$ -labeling of  $G_{N,k}$  plus the complete Merkle-Tree of that labeling. We define an adversary that sends the commitment  $C$  that is in  $h$  and, for any verifier's second message  $v$ , reads the right answer from  $h$ . That adversary needs to access  $h$  in *read-only* mode, hence without using any additional working space.

The protocol  $\Sigma_3$  is a **wPoSpace**. The structure of the proof is the same as the one of Theorem 2. We provide a particular case of the PoK Lemma (Lemma 2) where the hint  $h$  is the empty string and  $\tilde{\alpha}$  is  $0^k$ . For the sake of clarity, we restate the PoK Lemma for this particular setting.

**Lemma 3.** Consider the protocol  $\Sigma_3 = (\text{Gen}, V, P)$ , a PPT adversary  $\mathcal{A}$ , a space parameter  $S$  and a security parameter  $k$  such that  $\text{pk} \leftarrow \text{Gen}(1^k)$ .

If  $\Pr \left[ \mathbf{wSnd}_{\Sigma_3, S}^{\mathcal{A}}(k) \right]$  is noticeable, then there exists a randomness  $\tilde{\rho}$  such that  $\mathcal{A}$  fed with the randomness  $\tilde{\rho}$  during the round (1) of the protocol  $\Sigma_3$  calculates an  $\mathcal{H}_{0^k}$ -labeling  $\ell$  of  $G_{N,k}$  with  $m$  faults such that the following holds:

- $m = O\left(\frac{N}{\ln k}\right)$  and
- for any well-labeled node  $u$  in  $\ell$  we have  $(\tilde{\alpha} \| u \| \ell(u_1) \| \dots \| \ell(u_d)) \in \mathcal{Q}$ ,

where  $\mathcal{Q}$  is the set of queries to  $\mathcal{H}$  made by  $\mathcal{A}$  when fed with randomness  $\tilde{\rho}$  until the end of round (1) and  $\{u_1, \dots, u_d\} = \Gamma^-(u)$ .

By contradiction, suppose there exists an adversary  $\mathcal{A}$  for the protocol  $\Sigma_3$ . By applying Lemma 3, we extract a partial  $\mathcal{H}_{0^k}$ -labeling, ensuring that  $\mathcal{A}$  computed that labeling during round (1) (hence satisfying the hypotheses of Lemma 1).

This gives us, with overwhelming probability, a pebbling sequence  $\mathcal{P}$  of  $G_{N,k}$  with  $m$  wildcards and with pebbling space  $S' = \frac{S}{k} + d$ ,  $m = O(\frac{N}{\ln k})$ . The pebbling time of  $\mathcal{P}$  is a lower bound for the execution time of  $\mathcal{A}$  before round (2). The rest of the proof follows exactly the same structure of the proof of Theorem 2 and it is therefore omitted.  $\square$

**Acknowledgements.** We are grateful to Krzysztof Pietrzak for his insightful comments and suggestions.

## References

1. Martin Abadi, Mike Burrows, Mark Manasse, and Ted Wobber. Moderately hard, memory-bound functions. *ACM Trans. Internet Technol.*, 5(2):299–327, May 2005.
2. Mehmod Abliz and Taieb Znati. A guided tour puzzle for denial of service prevention. In *ACSAC*, pages 279–288. IEEE Computer Society, 2009.
3. Giuseppe Ateniese, Ilario Bonacina, Antonio Faonio, and Nicola Galesi. Proofs of space: When space is of the essence. Cryptology ePrint Archive, Report 2013/805, 2013. <http://eprint.iacr.org/>.
4. Giuseppe Ateniese, Randal Burns, Reza Curtmola, Joseph Herring, Lea Kissner, Zachary Peterson, and Dawn Song. Provable data possession at untrusted stores. In *Proceedings of the 14th ACM conference on Computer and communications security*, CCS ’07, pages 598–609, New York, NY, USA, 2007. ACM.
5. Giuseppe Ateniese, Seny Kamara, and Jonathan Katz. Proofs of storage from homomorphic identification protocols. In *ASIACRYPT*, pages 319–333, 2009.
6. Adam Back. Hashcash - a denial of service counter-measure. Technical report, 2002.
7. Boaz Barak and Oded Goldreich. Universal arguments and their applications. In *IEEE Conference on Computational Complexity*, pages 194–203. IEEE Computer Society, 2002.
8. R. B. Boppana, J. Hastad, and S. Zachos. Does co-np have short interactive proofs? *Inf. Process. Lett.*, 25(2):127–132, May 1987.
9. Luís Caires, Giuseppe F. Italiano, Luís Monteiro, Catuscia Palamidessi, and Moti Yung, editors. *Automata, Languages and Programming, 32nd International Colloquium, ICALP 2005, Lisbon, Portugal, July 11-15, 2005, Proceedings*, volume 3580 of *Lecture Notes in Computer Science*. Springer, 2005.
10. Ran Canetti, Ben Riva, and Guy N. Rothblum. Refereed delegation of computation. *Information and Computation*, 226(0):16 – 36, 2013.
11. Kai-min Chung, Yael Kalai, and Salil Vadhan. Improved Delegation of Computation Using Fully Homomorphic Encryption. pages 483–501, 2010.
12. Yevgeniy Dodis, Leonid Reyzin, and Adam Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. In Christian Cachin and Jan L. Camenisch, editors, *Advances in Cryptology - EUROCRYPT 2004*, volume 3027 of *Lecture Notes in Computer Science*, pages 523–540. Springer Berlin Heidelberg, 2004.
13. Cynthia Dwork, Andrew Goldberg, and Moni Naor. On memory-bound functions for fighting spam. *Advances in Cryptology-CRYPTO 2003*, 2003.
14. Cynthia Dwork and Moni Naor. Pricing via processing or combatting junk mail. *Advances in Cryptology-CRYPTO’92*, pages 139–147, 1993.
15. Cynthia Dwork, Moni Naor, and Hoeteck Wee. Pebbling and proofs of work. *Advances in Cryptology CRYPTO 2005*, pages 37–54, 2005.
16. Stefan Dziembowski, Sebastian Faust, Vladimir Kolmogorov, and Krzysztof Pietrzak. Proofs of space. Cryptology ePrint Archive, Report 2013/796, 2013. <http://eprint.iacr.org/>.
17. Stefan Dziembowski, Tomasz Kazana, and Daniel Wichs. Key-evolution schemes resilient to space-bounded leakage. In *CRYPTO*, pages 335–353, 2011.
18. Stefan Dziembowski, Tomasz Kazana, and Daniel Wichs. One-time computable self-erasing functions. In *Proceedings of the 8th conference on Theory of cryptography*, TCC’11, pages 125–143, Berlin, Heidelberg, 2011. Springer-Verlag.
19. Stefan Dziembowski, Krzysztof Pietrzak, and Sebastian Faust. Proofs of space and a greener bitcoin. talk presented at *Workshop on Leakage, Tampering and Viruses*, Warsaw 2013, 2013.
20. Dario Fiore and Rosario Gennaro. Publicly verifiable delegation of large polynomials and matrix computations, with applications. *Proceedings of the 2012 ACM conference on Computer and communications security - CCS ’12*, page 501, 2012.



21. Oded Goldreich and Johan Hastad. On the complexity of interactive proofs with bounded communication. *Information Processing Letters*, 1998.
22. Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. Delegating computation: interactive proofs for muggles. In *Proceedings of the 40th annual ACM symposium on Theory of computing*, STOC '08, pages 113–122, New York, NY, USA, 2008. ACM.
23. Ari Juels and John G. Brainard. Client puzzles: A cryptographic countermeasure against connection depletion attacks. In *NDSS*. The Internet Society, 1999.
24. Nikolaos P. Karvelas and Aggelos Kiayias. Efficient proofs of secure erasure. *SCN*, 2014.
25. Joe Kilian. A note on efficient zero-knowledge proofs and arguments (extended abstract). In *Proceedings of the twenty-fourth annual ACM symposium on Theory of computing*, STOC '92, pages 723–732, New York, NY, USA, 1992. ACM.
26. Thomas Lengauer and Robert E. Tarjan. Asymptotically tight bounds on time-space trade-offs in a pebble game. *Journal of the ACM (JACM)*, 29(4):1087–1130, 1982.
27. Silvio Micali. Computationally sound proofs. *SIAM Journal on Computing*, 30(4):1253–1298, 2000.
28. Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. May 2009.
29. Colin Percival. Stronger key derivation via sequential memory-hard functions. presented at BSDCan '09, 2009.
30. Daniele Perito and Gene Tsudik. Secure code update for embedded devices via proofs of secure erasure. In *ESORICS*, pages 643–662, 2010.
31. Nicholas Pippenger. Superconcentrators. *SIAM J. Comput.*, 6(2):298–304, 1977.
32. R.E.A.C.Paley and A.Zygmund. A note on analytic functions in the unit circle. page 266272, 1932.
33. Hovav Shacham and Brent Waters. Compact proofs of retrievability. In *Proceedings of the 14th International Conference on the Theory and Application of Cryptology and Information Security: Advances in Cryptology, ASIACRYPT '08*, pages 90–107, Berlin, Heidelberg, 2008. Springer-Verlag.
34. Adam Smith and Ye Zhang. Near-linear time, leakage-resilient key evolution schemes from expander graphs. *IACR Cryptology ePrint Archive*, 2013:864, 2013.
35. Martin Tompa. Time-space tradeoffs for computing functions, using connectivity properties of their circuits. In *Proceedings of the tenth annual ACM symposium on Theory of computing*, STOC '78, pages 196–204, New York, NY, USA, 1978. ACM.
36. Brent Waters, Ari Juels, J. Alex Halderman, and Edward W. Felten. New client puzzle outsourcing techniques for dos resistance. In *Proceedings of the 11th ACM conference on Computer and communications security*, CCS '04, pages 246–256, New York, NY, USA, 2004. ACM.

## A Proof of the Pebbling Theorem (Theorem 1)

For the sake of convenience we re-write the statement of that theorem here. For the definitions of pebbling games, pebbling space and pebbling time we refer the reader to Section 2.

**Pebbling Theorem (Theorem 1).** *There exists a family of efficiently sampleable directed acyclic graphs  $\{G_{N,k}\}_{k,N \in \mathbb{N}}$  with constant fan-in  $d$ ,  $N$  input nodes,  $N$  output nodes and  $kN \log N$  nodes in total, such that:*

1. *Any pebbling sequence that pebbles  $\Delta$  output nodes has pebbling space at most  $S$  pebbles and  $m$  wildcards, where  $|\Delta| \geq 4S + 2m + 1$ , and pebbling time  $T$  such that*

$$T \geq |\Delta| \left( \frac{N - 2S - m}{2S + m + 1} \right)^k.$$

2. *There exists a pebbling sequence that pebbles the graph  $G_{N,k}$  which has pebbling space  $N + 2$  and needs time  $O(kN \log N)$ .*
3. *For any node  $v \in V(G_{N,k})$ , the incoming nodes of  $v$  and the position of  $v$  in first lexicographic topologically order of  $G_{N,k}$  are computable in  $O(k \log N)$ .*

The family of graphs  $G_{N,k}$  we build is made from DAGs that are  $k$  layered stack of superconcentrators [31] with  $N$  inputs and  $N$  outputs.

**Definition 5 (superconcentrator).** We say that a directed acyclic graph  $G = (V, E)$  is an  $N$ -superconcentrator if and only if

1.  $|\mathbf{S}(G)| = |\mathbf{T}(G)| = N$ ,
2. for each  $S \subseteq \mathbf{S}(G)$  and for each  $T \subseteq \mathbf{T}(G)$  such that  $|S| = |T| = \ell$  there exist  $\ell$  vertex-disjoint paths from  $S$  to  $T$ , i.e. paths such that no vertex in  $V$  is in two of them.




Fig. A.1. An example of 16-superconcentrator: the Butterfly Graph

**Definition 6 (stack of superconcentrators).** Given an  $N$ -superconcentrator  $G$  we define the graph  $G^{\times k}$ , the stack of  $k$  copies of  $G$ , inductively as follows:  $G^{\times 1} = G$  and  $G^{\times k}$  is obtained from  $G^{\times (k-1)}$  and  $G$  by first renaming the vertexes of  $G$  so that they do not appear in the vertexes of  $G^{\times (k-1)}$ , obtaining a graph  $G'$ , and then by identifying  $\mathbf{T}(G^{\times (k-1)})$  with  $\mathbf{S}(G')$ .

Not any  $k$  layered stack of superconcentrators satisfies points 2 and 3 of the Pebbling Theorem. To obtain those properties we use the following well-known superconcentrator: the *Butterfly Graph*  $B_N$  built by putting together two FFT graphs with  $N$  inputs and outputs [35]. An example of such construction (for  $N = 16$ ) is given in Figure A.1.

It is easy to see that the family of graphs  $G_{N,k} := B_N^{\times k}$  satisfies points 2 and 3 of the Pebbling Theorem and it is a  $k$  layered stack of  $N$ -superconcentrators. It remains to prove only point 1.

For superconcentrator graphs we can prove a tradeoff between pebbling space and pebbling time, even for pebbling games with wildcards. This result is based on a generalization of the “*Basic Lower Bound Argument*” (BLBA) by Tompa [35].

**Lemma 4 (BLBA with wildcards).** Consider an  $N$ -superconcentrator  $G$ , a set  $M \subseteq \mathbf{T}(G)$ , and a pebbling sequence  $\mathcal{P} = (P_0, \dots, P_\ell)$  with  $m$  wildcards that pebbles  $M$ . Let  $A$  be the set of all elements of  $\mathbf{S}(G)$  pebbled and unpebbled at some point in  $\mathcal{P}$ . If  $|M| \geq |P_0| + |P_\ell| + m + 1$  then  $|A| \geq N - |P_0| - |P_\ell| - m$ .

*Proof.* By contradiction suppose that  $|A| < N - |P_0| - |P_\ell| - m$ , this means that  $|A^c| \geq |P_0| + |P_\ell| + m + 1$  and every element in  $A^c$ , by definition, is not pebbled and unpebbled. Take any  $B \subseteq A^c$  such that  $|B| = |M|$  then, as  $G$  is an  $N$ -superconcentrator, we have  $\pi_1, \dots, \pi_{|M|}$  vertex disjoint paths from  $B$  to  $M$ . Let  $W$  be a set of  $m$  faults for  $\mathcal{P}$ . By construction  $|M| > |P_0 \cup P_\ell \cup W|$  thus we have some path  $\pi$  from some element  $v$  of  $B$  to some element  $w$  of  $M$  such that  $\pi \cap (P_0 \cup P_\ell \cup W) = \emptyset$ . By definition of  $A^c$ , and hence of  $B$ ,  $v$  is not pebbled and unpebbled and  $v \notin P_0 \cup P_\ell \cup W$ , thus  $v \notin P_i$  for each  $i \in [\ell]$ . As  $\pi \cap (P_0 \cap W) = \emptyset$ , we must have that  $\pi \cap P_i = \emptyset$  for each  $i \in [\ell]$ . But then the vertex  $w$  is not in  $\bigcup_{i \in [\ell]} P_i$  contradicting the fact that  $\mathcal{P}$  pebbles  $M$ .  $\square$

**Theorem 5 (Pebbling Theorem (point 1)).** Let  $G^{\times k}$  be a stack of  $k$  copies of an  $N$ -superconcentrator  $G$ ,  $\Delta \subseteq \mathbf{T}(G^{\times k})$  and  $\mathcal{P}$  a pebbling sequence for  $G^{\times k}$  with  $m$  wildcards that pebbles  $\Delta$ . Let  $S$  be the pebbling space of  $\mathcal{P}$  and  $T$  the pebbling time of  $\mathcal{P}$ . If  $|\Delta| \geq 4S + 2m + 1$  then

$$T \geq |\Delta| \left( \frac{N - 2S - m}{2S + m + 1} \right)^k.$$

*Proof.* Let  $G_1, \dots, G_k$  be the  $k$  copies of  $G$  that form  $G^{\times k}$ . Suppose we want to pebble a set  $M \subseteq \mathbf{T}(G_i)$  such that  $|M| \geq 2S + m + 1$  and let  $A_1, \dots, A_n$  be disjoint subsets of  $M$  such that  $A_1$  are the first  $2S + m + 1$  elements of  $M$  to be pebbled in  $\mathcal{P}$ ,  $A_2$  are the second  $2S + m + 1$  elements of  $M$  to be pebbled in  $\mathcal{P}$  and so on. Clearly we have that the number of these sets is  $n = \left\lfloor \frac{|M|}{2S + m + 1} \right\rfloor$ . By Lemma 4 we have that for each  $A_j$  there exists a set  $\beta(A_j)$  in  $\mathbf{S}(G_i)$ , whose elements are pebbled and unpebbled, of size at least  $N - 2S - m$ . Thus we have that the total number of elements in  $\mathbf{S}(G_i)$  that have been pebbled and unpebbled is at least  $(N - 2S - m) \left\lfloor \frac{|M|}{2S + m + 1} \right\rfloor$ . Notice now that each  $\beta(A_j) \in \mathbf{T}(G_{i-1})$  hence we can reapply the argument above to it provided that for each  $j$ ,  $|\beta(A_j)| \geq N - 2S - m \geq 2S + m + 1$  (and this is implied by the fact that  $N \geq |\Delta| \geq 4S + 2m + 1$ ). Thus starting from  $\mathbf{T}(G_k) = \mathbf{T}(G^{\times k})$  we can prove by induction, going back to  $G_1$ , that at least

$$|\Delta| \left( \frac{N - 2S - m}{2S + m + 1} \right)^k$$

actions of pebbling and unpebbling take place on  $\mathbf{S}(G_1) = \mathbf{S}(G^{\times k})$ . As each action is an elementary step in the pebbling game  $\mathcal{P}$  we have that the result follows.  $\square$


## B Merkle Tree

Consider the complete binary tree  $T$  over  $n$  leaves. Without loss of generality we can assume that  $n$  is a power of 2 by using padding if necessary. Label the  $i$ -th leaf with  $l_i$ , then, use a collision resistant hash (CRH) function  $H_s : \{0, 1\}^{2k} \rightarrow \{0, 1\}^k$  with random seed  $s$  to propagate the labeling  $\mathbf{l}$  of the leaves to a labeling  $\phi$  along the tree  $T$  up to the root according to the following rules:

1.  $\phi(v) := l_i$  if  $v$  is the  $i$ -th leaf of  $T$ ,
2.  $\phi(v) := H_s(\phi(v_1) \parallel \phi(v_2))$  where  $v_1, v_2$  are the two children of  $v$ .

Then the *commitment* for a string  $\mathbf{l}$  is the label  $\phi(r)$  of the root of the tree; we denote it with  $\mathbf{MT}_s(\mathbf{l})$ . The commitment is succinct in fact it has size  $k$  which is independent of  $n$ .

An *opening* for the  $i$ -th element in the sequence  $\mathbf{l}$  is an ordered sequence formed by elements of  $\{0, 1\}^k$  that are intended to be the label  $l_i$  of the  $i$ -th leaf of  $T$  together with the labels  $\phi(v_j)$  for each  $v_j$  that is a sibling of vertices in the path from the root of the tree to the  $i$ -th leaf. An example of opening is shown in Figure B.1.



**Fig. B.1.** An example of opening:  $\mathbf{c} = (a, b, c, d, e)$

We denote with  $\text{Open}_s(\mathbf{l}, i)$  the opening formed by  $l_i$  together with all the  $\phi(v_j)$ . Notice that from  $\text{Open}_s(\mathbf{l}, i)$  it is possible to compute  $\phi(w)$  for each vertex  $w$  in the path from the root to the leaf at the  $i$ -th

position. In particular it is efficient to compute  $\text{MT}_s(\mathbf{l})$ . The opening is succinct in fact it has size  $k \cdot \log n$  which is (almost) independent of  $n$ . In general, given an opening  $\mathbf{c} = (c_1, \dots, c_{\log n})$  for some position  $i$  we can think of it as equivalent to  $\text{Open}_s(\mathbf{l}, i)$  and compute the label of the root according to the given  $\mathbf{c}$  and position  $i$ . If this value is equal to  $\text{MT}_s(\mathbf{l})$ , we say that  $\mathbf{c}$  *opens the  $i$ -th position to  $c_1$* . As  $\mathbf{H}$  is a CRH function, then for any  $i$ , it is guaranteed that  $\text{MT}_s(\mathbf{l})$  can open the  $i$ -th position only to  $l_i$ .