

# Bandwidth-Hard Functions: Reductions and Lower Bounds

Jeremiah Blocki<sup>\*</sup>  
Department of Computer Science,  
Purdue University  
jblocki@purdue.edu

Ling Ren  
CSAIL, MIT  
renling@mit.edu

Samson Zhou<sup>†</sup>  
Department of Computer Science,  
Purdue University  
samsonzhou@gmail.com

## ABSTRACT

Memory Hard Functions (MHFs) have been proposed as an answer to the growing inequality between the computational speed of general purpose CPUs and Application Specific Integrated Circuits (ASICs). MHFs have seen widespread applications including password hashing, key stretching and proofs of work. Several metrics have been proposed to quantify the “memory hardness” of a function. Cumulative memory complexity (CMC) [8] (or amortized Area  $\times$  Time complexity [4]) attempts to quantify the cost to acquire/build the hardware to evaluate the function — after normalizing the time it takes to evaluate the function. By contrast, bandwidth hardness [30] attempts to quantify the amortized energy costs of evaluating this function on hardware — which in turn is largely dominated by the number of cache misses. Ideally, a good MHF would be both bandwidth hard and have high cumulative memory complexity. While the cumulative memory complexity of leading MHF candidates is well understood, little is known about the *bandwidth hardness* of many prominent MHF candidates.

Our contributions are as follows: First, we provide the first reduction proving that, in the parallel random oracle model, the bandwidth hardness of a Data-Independent Memory Hard Function (iMHF) is described by the red-blue pebbling cost of the directed acyclic graph (DAG) associated with that iMHF. Second, we show that the goals of designing an MHF with high CMC/bandwidth hardness are well aligned. In particular, we prove that *any* function with high CMC also has relatively high energy costs. This result leads to the first *unconditional* lower bound on the energy cost of script in the parallel random oracle model. Third, we analyze the bandwidth hardness of several prominent iMHF candidates such as Argon2i [11], winner of the password hashing competition, aATSample and DRSample [4] — the first practical iMHF with essentially asymptotically optimal CMC. We show Argon2i, aATSample and DRSample are maximally bandwidth hard under appropriate cache size. Finally, we show that the problem of finding a red-blue pebbling with minimum energy cost is NP-hard.

<sup>\*</sup>Research supported in part by NSF Award #1704587.

<sup>†</sup>Research supported in part by NSF Award #1649515.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

CCS '18, October 15–19, 2018, Toronto, ON, Canada

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-5693-0/18/10...\$15.00

<https://doi.org/10.1145/3243734.3243773>

## KEYWORDS

memory-hard functions, energy cost, bandwidth hardness, graph pebbling

## ACM Reference Format:

Jeremiah Blocki, Ling Ren, and Samson Zhou. 2018. Bandwidth-Hard Functions: Reductions and Lower Bounds. In *2018 ACM SIGSAC Conference on Computer and Communications Security (CCS '18)*, October 15–19, 2018, Toronto, ON, Canada. ACM, New York, NY, USA, 20 pages. <https://doi.org/10.1145/3243734.3243773>

## 1 INTRODUCTION

Memory Hard Functions (MHFs) [1, 28] are a crucial building block in the design of password hashing functions, moderately hard key-derivation functions and egalitarian proofs of work [9, 19]. For example, in password hashing it is critically important to ensure that it is prohibitively expensive for an offline attacker to evaluate the function millions or billions of times to check each password in a large cracking dictionary. The development of improved Application Specific Integrated Circuits (ASICs) or Field Programmable Gate Arrays (FPGAs) for computing cryptographic hash functions such as SHA256 makes this goal increasingly challenging. For example, the Antminer S9, an ASIC Bitcoin [27] miner, is able to compute SHA256 hashes at a rate of 13.6 trillion hashes per second using just 1274 Joules of energy per second (Watts). By contrast, the energy needed to compute SHA256 13.6 trillion times on a standard CPU would be about six orders of magnitude higher! In fact, Blocki *et al.* [14] recently argued that non-memory hard key derivation functions (e.g., PBKDF2-SHA256 and BCRYPT are based on hash iteration) *cannot* provide sufficient protection against a rational (economically motivated) offline attacker without introducing unacceptably long authentication delays.

MHFs are based on the observation that memory costs (e.g., latency, bandwidth, energy consumption) tend to be equitable across different architectures. Thus, to develop an “egalitarian” function we want to design a function where evaluation costs are dominated by memory costs. Two of the most prominent approaches to measure the “evaluation cost” of MHFs are memory hardness [8, 28] and bandwidth hardness [30]. Memory hardness [28] seeks to quantify construction costs i.e., the cost to build/obtain the hardware necessary to compute the MHF. By contrast, bandwidth hardness [30] seeks to quantify the energy costs per evaluation i.e., the cost of running the hardware. Ideally, one would hope to design an MHF that is *both* bandwidth hard and memory hard.

Broadly speaking there are two types of MHFs: data-dependent memory hard functions (dMHFs) and data-independent memory hard functions (iMHFs). As the name suggests an iMHF induces

a memory access pattern that is *independent* of the sensitive input (e.g., password), which makes them naturally resistant to certain side channel attacks e.g., cache-timing [10]. Meanwhile, while dMHFs with high memory/bandwidth hardness are potentially easier to construct [2, 6], they are also more vulnerable to side channel attacks. Argon2 [12], winner of the recently completed Password Hashing Competition [29], includes a data-independent mode of operation (Argon2i), a data-dependent mode (Argon2d) and a hybrid mode (Argon2id).

To a large extent, most of the recent cryptanalysis of MHF candidates has focused on memory hardness. In particular, *cumulative memory complexity* (CMC) [8] and the closely related metric *amortized area-time complexity* (aAT) [2, 4] aim to approximate the cost of constructing enough chips to evaluate the function  $T$  times per year. For example, if evaluating the function one time requires us to lock up 1GB of DRAM for 1 second then, *at minimum*, an attacker would need to buy roughly 32 (1GB) DRAM chips to evaluate the function a billion times per year. Alwen *et al.* [6] showed that the dMHF scrypt [28] has maximal CMC  $\Omega(n^2)$  i.e., the amortized area-time complexity scales quadratically with the running time  $n$ . By contrast, Alwen and Blocki [2, 3] showed that *any* iMHF has cumulative memory complexity *at most*  $O\left(\frac{n^2 \log \log n}{\log n}\right)$  and they exhibited even stronger amortization attacks against Password Hashing Competition [29] (PHC) winner Argon2i [12] along with other candidate MHFs such as balloon hashing [17]. Blocki and Zhou [15] showed that Argon2i has CMC at most  $O(n^{1.767})$  and at least  $\Omega(n^{1.75})$ . Alwen *et al.* [5] also gave a theoretical construction of an iMHF with CMC *at least*  $\Omega\left(\frac{n^2}{\log n}\right)$ , which is essentially optimal in an asymptotic sense. More recently, Alwen *et al.* [4] designed two *practical* iMHFs called DRSample and aATSample with the same asymptotic complexity.

By contrast, the notion of *bandwidth-hardness* was only introduced recently [30] with the intention of lower bounding the energy required to evaluate the function. Ren and Devadas [30] observed that metrics such as CMC or aAT do not provide an accurate picture of energy consumption. For example, certain types of memory consume very little energy when idle, but cache misses are costly because we must *retrieve* data from RAM. Memory Bound Functions [1] are functions whose computation *always requires* a large number of cache-misses regardless of computation time. Bandwidth hardness [30] relaxes this notion by requiring that any attacker who evaluates the function must either 1) incur a large number of expensive cache misses, or 2) must perform a larger (e.g., super-linear) amount of computation.

Ren and Devadas proposed to cryptanalyze an iMHF using a variant of the red-blue pebbling game in which red-moves (representing computation performed using data in cache) have a smaller cost  $c_r$  than blue-moves  $c_b$  (representing data movements to/from memory) [30]. Ren and Devadas also proved that the bit reversal graph [25], which forms the core of iMHF candidate Catena-BRG [22], is maximally bandwidth hard in the sense that *any* red-blue pebbling has cost  $\Omega(n \cdot c_b)$ . However, Ren and Devedas [30] did not prove that any attacker in the parallel random oracle model (pROM) can be viewed as a red-blue pebbling so it was not clear whether or not a graph (e.g., Catena-BRG [22]) with high red-blue pebbling cost is *necessarily* bandwidth hard in the pROM model.

Similarly, Ren and Devedas [30] showed that scrypt is bandwidth hard under a restrictive assumption about the cache-architecture adopted by the attacker e.g., they need to assume data from RAM can only be retrieved in large chunks. Prior to our work nothing was known about the bandwidth hardness of key MHF candidates such as PHC winner Argon2i [12, 29], DRSample and aATSample [4].

**Our Contributions.** We formalize the notion of bandwidth hardness in the parallel random oracle model and show that bandwidth hardness is indeed captured by the red-blue pebbling game. This does for bandwidth hardness what Alwen and Serbinenko [8] did for CMC when they showed that CMC is captured by the parallel black pebbling game.

Second, we demonstrate that CMC *lower bounds* can be used to *directly* lower bound energy costs. In particular, we show that energy costs are at least  $\Omega\left(\sqrt{c_b c_r \text{CMC}} - c_b \cdot m\right)$ . Intuitively, an attacker running in time  $t$  will pay computation costs at least  $t \cdot c_r$  and energy cost *at least*  $\left(\frac{\text{CMC}}{t \cdot w} - m\right) \cdot c_b$  where  $m$  denotes the number of  $w$ -bit words that can be stored in cache. Alwen *et al.* [6] show that scrypt has CMC at least  $\Omega(n^2 \cdot w)$ . Combined with our result this implies that scrypt has energy cost at least  $\Omega(\sqrt{c_b c_r} n)$ . While we would ideally hope for a lower bound of  $\Omega(n \cdot c_b)$ , we stress that this is the first lower bound for scrypt that *does not* require any restrictive assumptions about cache-architecture [30]. The result also demonstrates that the goals of high CMC and high bandwidth hardness are well aligned.

Third, we introduce a new technique to lower-bound the red-blue pebbling cost of a DAG and we use this new technical hammer to lower-bound the red-blue pebbling cost of several important iMHF candidates including: Argon2iB (the current version of PHC winner Argon2i [12]), Argon2iA (an older version of Argon2, which is similar to balloon hashing [17]), DRSample and aATSample. In particular, we show that if  $m = O(n^{2/3-\epsilon})$  then any pROM attacker with cache-size  $m \cdot w$  bits *must* incur cost  $\min\{\Omega(n \cdot c_b), \omega(n \cdot c_r)\}$ . In particular, Argon2iB is maximally bandwidth hard whenever the attacker's cache size is sufficiently small. Argon2iB uses a round function with word size  $w = 2^{10}$  Bytes. Assuming that we set our memory hardness parameter  $n = 2^{20}$  (filling  $nw = 1\text{GB}$  of RAM in about 1 second according to Argon2 benchmarks [12]) then  $n^{2/3} \cdot w$  corresponds to a cache-size of 10MB— our lower bounds would not apply if the attacker's cache size is larger. We prove even stronger lower bounds for DRSample and aATSample. In particular, we prove that these functions are *maximally* bandwidth hard as long as  $m = O(n^{1-\epsilon})$ . Interestingly, DRSample and aATSample have asymptotically higher CMC as well, which is consistent with our observation that the goals of designing a MHF with high CMC is well aligned with the goal of designing an maximally bandwidth hard function.

While we prove that DRSample, aATSample and Argon2iB are all maximally bandwidth hard in an asymptotic sense, it would be nice to gain a more precise understanding of the constant factors in these bounds. To this end it would be useful to develop an efficient algorithm to determine the minimum cost red-blue pebbling of a DAG  $G$ . However, our final result is a negative one. We show that

it is NP-Hard to compute the minimum cost red-blue pebbling of a DAG  $G$ .

### 1.1 Graph Pebbling and iMHFs

An iMHF  $f_{G,H}$  is defined by a labeling game over a DAG  $G$  and a random oracle  $H : \{0, 1\}^* \rightarrow \{0, 1\}^w$ . In particular, the label  $\ell_v$  of an intermediate node  $v$  is computed as  $\ell_v = H(v, \ell_{v_1}, \dots, \ell_{v_{\text{indeg}}})$  where  $v_1, \dots, v_{\text{indeg}}$  are the parents of node  $v$  in  $G$ . The output of the function is the label of the final sink node. Before we provide an overview of our technical results it is necessary to first (informally) introduce the black pebbling game and the red-blue pebbling game.

**Black Pebbling.** Given a directed acyclic graph (DAG)  $G = (V, E)$ , the goal of the (parallel) black pebbling game is to place pebbles on all sink nodes of  $G$  (not necessarily simultaneously). The game is played in rounds and we use  $P_i \subseteq V$  to denote the set of currently pebbled nodes on round  $i$ . Initially all nodes are unpebbled,  $P_0 = \emptyset$ , and in each round  $i \geq 1$  we may only include  $v \in P_i$  if all of  $v$ 's parents were pebbled in the previous configuration ( $\text{parents}(v) \subseteq P_{i-1}$ ) or if  $v$  was already pebbled in the last round ( $v \in P_{i-1}$ ). In the sequential pebbling game we can place at most one new pebble on the graph in any round (i.e.,  $|P_i \setminus P_{i-1}| \leq 1$ ), but in the parallel pebbling game no such restriction applies. The space cost of the pebbling is defined to be  $\max_i |P_i|$ , which intuitively corresponds to minimizing the maximum space required during computation of the associated function, and relates to the space-complexity of the black-pebbling game. Gilbert *et al.* [23] studied the space-complexity of the black-pebbling game and showed that this problem is PSPACE – Complete by reducing from the truly quantified boolean formula (TQBF) problem. Given a (parallel) legal black pebbling  $P_1, \dots, P_t$  of a DAG  $G$ , we define the cumulative cost to be  $|P_1| + \dots + |P_t|$ . Then we define  $\Pi_{cc}(G)$  (resp.  $\Pi_{cc}^{\parallel}(G)$ ) as the minimum cumulative cost of any legal sequential (resp. parallel) black pebbling of  $G$ .

**Pebbling Reduction in the pROM Model.** Alwen and Serbinenko [8] show that under the parallel random oracle model (pROM) of computation, the cryptanalysis of an iMHF, under the amortized time-space metric, can be approximately reduced to the cumulative cost of a pebbling strategy. The result is significant in that it allows future cryptanalysis of iMHF candidates to focus on understanding the (parallel) black pebbling costs of the underlying DAG. In particular, a lower bound on the aAT complexity of the best pebbling for a DAG  $G$  immediately yields a lower bound on the aAT complexity of any pROM attacker evaluating the function  $f_{G,H}$ . Intuitively, this means that if  $G$  has sufficiently high (parallel) black pebbling cost then it will be expensive for the attacker to obtain enough hardware to compute the function  $f_{G,H}$  millions/billions of times per year e.g., an offline password cracking adversary.

**Red-Blue Pebbling.** Given a DAG  $G = (V, E)$ , the goal of the red-blue pebbling game [24] is again to place pebbles on all sink nodes of  $G$  (not necessarily simultaneously) from a starting configuration that contains no pebbles on any nodes. The game is again played in rounds, with each node possibly containing a blue pebble or a red pebble at each time step. Informally, at each time step, for any

node  $v$  we can swap between a red pebble at  $v$  and a blue pebble at  $v$  (and vice versa). Each swap is called a blue move, and while there is no limit to the number of blue moves at a single time step, they each have an associated cost  $c_b$ . Simultaneously, we may place a red pebble at a node  $v$  if all of  $v$ 's parents contained red pebbles in the previous configuration. This manner of placing a new red pebble is a red move and each occurrence incurs cost  $c_r$ . We are allowed to have at most  $m$  (cache-size) red-pebbles on the graph at any point in time. In a sequential red-blue pebbling we are allowed to place at most one new red pebble on the graph during each round, while no such constraint applies to a parallel red-blue pebbling. Finally, there is a parameter  $m$  that denotes a threshold on the number of nodes that can contain red pebbles at each time step. The total cost of the red-blue pebbling is the sum of the costs induced by the blue moves and the red moves. We define  $\text{rbpeb}^{\parallel}(G, m)$  (resp.  $\text{rbpeb}(G, m)$ ) to be the minimum cost of any legal parallel (resp. sequential) red-blue pebbling of  $G$  that places at most  $m$  red-pebbles on the graph at any point in time.

### 1.2 Overview of Our Results

**Proving that the Red-Blue Pebbling Game Captures Bandwidth Hardness.** We consider the variant of the red-blue pebble game proposed by Ren and Devadas [30] in which red moves have cost  $c_r$  and blue moves have cost  $c_b$  – note that if  $c_r = 0$  then we recover the traditional goal of minimizing the number of cache misses. Ren and Devadas [30] proposed the adoption of red-blue pebbling to model the bandwidth-complexity of iMHFs, with the idea that red moves correspond to hash computations and blue moves correspond to (more expensive) swaps between cache and memory. However, they did not *prove* any connection between red-blue pebbling costs and the *actual* bandwidth-costs of a pROM attacker.

Our contributions are two-fold. First, we formalize the notion of energy cost of a function  $f_{G,H}$  in the parallel random oracle model. Second, we prove that  $\text{ecost}(f_{G,H})$  the energy cost of  $f_{G,H}$  is closely related to red-blue pebbling costs. In particular, we prove that any pROM machine computing  $f_{G,H}$  with cache-size  $mw$ -bits has energy costs  $\Omega(\text{rbpeb}^{\parallel}(G, 8m))$ . This resolves an open question of [30], and shows that future cryptanalysis of the *bandwidth hardness* of iMHF candidates can focus on the red-blue pebbling cost of the underlying DAG  $G$ .

**THEOREM 1.1.** (Informal, see Theorem 3.3.)  $f_{G,H}$  has energy cost at least  $\text{ecost}(f_{G,H}, mw) \in \Omega(\text{rbpeb}^{\parallel}(G, 8m))$ .

While Theorem 3.3 is similar to a result of Alwen and Serbinenko who showed that the cumulative memory complexity of  $f_{G,H}$  is captured by the black pebbling game [8], we stress that there are several unique challenges in our reduction. In particular, it is easier to extract a black pebbling from the execution trace of a pROM attacker since each new pebble that is placed on the graph during round  $i$  corresponds directly to a random oracle query that was made during the previous round. However, in the red-blue pebbling model only red moves correspond to random oracle queries. Intuitively, we expect that blue moves correspond to labels that are transferred to/from memory, but an attacker may encode each of these labels in an unexpected way (e.g., encryption). Thus, even if

we can observe the data values being transferred to/from memory we stress that we *cannot directly infer* which labels are being transferred making it difficult to extract a legal red-blue pebbling from the execution trace.

We overcome this difficulty by allowing the red-blue pebbling to use a little bit of extra memory (e.g., if the pROM attacker has  $m \cdot w$  bits of cache then the red-blue pebbling is allowed to use  $8m$  red-pebbles) and by introducing the notion of a *red-blue extension pebbling* of a legal black pebbling  $P = (P_1, \dots, P_t)$ . Given a legal black pebbling extracted from the execution trace of the pROM attacker running in time  $t$  we can partition time into intervals  $[t_0 = 1, t_1), [t_1, t_2), \dots, [t_{k-1}, t_k = t]$  such that 1) during each interval  $[t_i, t_{i+1})$  the pROM attacker transfers *at least*  $mw$  bits from memory (at cost  $m \cdot c_b$ ), and 2) there is a red-blue extension pebbling that makes at most  $\mathcal{O}(m)$  blue moves during each interval  $[t_i, t_{i+1})$ .

To partition time into intervals we introduce a set  $\text{QueryFirst}(x, y)$  that intuitively corresponds to the data-labels that *appear first as input* to a random oracle query during the time interval  $[x, y)$  *before* the label appears as the output of some random oracle query during the same interval. We then define  $t_1$  to be the minimum pebbling round such that there exists  $1 \leq j_1 < t_1$  such that  $\text{QueryFirst}(j_1, t_1)$  has size at least  $3m$ . Similarly, once  $t_1 < \dots < t_{i-1}$  have been defined we can define  $t_i > t_{i-1}$  to be the minimum pebbling round such that there exists  $t_i > j_i \geq t_{i-1}$  s.t.  $\text{QueryFirst}(j_i, t_i)$  has size at least  $3m$ . At the beginning of each interval  $[t_i, t_{i+1})$  our red-blue extension pebbling will place red pebbles on all nodes in the set  $\text{QueryFirst}(j_i, t_i)$  (e.g., to “load” these values into cache). We can accomplish this legally since the extension pebbling is allowed to use up to  $8m$  red-pebbles. Once we have red pebbles placed on all of these nodes the extension pebbling is able to finish this interval *without* changing any other blue nodes into red-nodes (i.e., zero cache misses).

To prove that the pROM attacker must transfer *at least*  $mw$  bits from memory during each interval we rely on an extractor argument. In particular, let  $\gamma_i$  encode the messages transferred to/from cache during the interval  $[t_i, t_{i+1})$ . Our extractor will extract  $3m$  labels (without querying the random oracle at these points) by simulating the pROM attacker starting with a hint. The labels we will extract correspond to the nodes in the set  $\text{QueryFirst}(t_i, t_{i+1})$ . The hint consists of  $\gamma_i$  along with other information such as the current state of the cache (at most  $mw$  bits), indices of the  $4m$  labels that we want to extract ( $4m \log n$  bits to encode), and the index of the first query in which each label appears as input to a random oracle query ( $4m \log q$  bits to encode where  $q$  is an upper bound on number of queries made by the attacker). Since a random oracle is incompressible, the extractor’s hint must have length at least  $4mw$  if we expect the extractor to output  $4m$  labels (i.e.,  $4m$  distinct random oracle outputs of length  $w$  assuming there are no hash collisions) without querying the random oracle at these points so it follows that  $|\gamma_i| \geq m \cdot w$ .

**On the Relationship between Bandwidth Complexity and Cumulative Memory Complexity.** We show that bandwidth complexity and cumulative memory complexity are intricately related concepts. If  $\text{rbpeb}^\parallel(G, m)$  is the minimum energy cost<sup>1</sup> of *any*

legal *parallel* reb-blue pebbling of  $G$  with cache size  $m$  and  $\Pi_{cc}$  is the cumulative complexity of *sequential* black pebbling, then

THEOREM 1.2.

$$\text{rbpeb}^\parallel(G, m) \geq 2c_b \left( \frac{\Pi_{cc}(G)}{t} - 2m \right) + c_r t \in \Omega \left( \sqrt{c_b \cdot c_r \cdot \Pi_{cc}(G)} \right),$$

where  $m$  is the cache size,  $t$  is the number of steps in the pebbling,  $c_b$  is the cost of a blue move and  $c_r$  is the cost of a red move.

Theorem 1.2 demonstrates that the goals of designing an MHF with high cumulative complexity and high bandwidth complexity are well aligned. In fact, we use Theorem 1.2 to show that a family  $\{G_n\}_{n=1}^\infty$  of constant indegree DAGs constructed by Schnitger [31] has high energy costs *because*  $\Pi_{cc}(G_n) \in \Omega(n^2)$  [7]. As an intermediate step to proving Theorem 1.2 we show that  $\text{rbpeb}^\parallel(G, m) \geq \text{rbpeb}(G, 2m)$ . This result is interesting as it suggests that an attacker will not be able to dramatically decrease energy costs by exploiting parallelism. By contrast, for *any* constant indegree DAG  $G$  it is known that the parallel cumulative pebbling cost is at most  $\Pi_{cc}^\parallel(G) \in \mathcal{O} \left( \frac{n^2 \log \log n}{\log n} \right)$  [2] while it is known that  $\Pi_{cc}(G_n) \in \Omega(n^2)$  for the constant indegree DAGs constructed by Schnitger [31].

We also prove a similar theorem that directly relates *ecost* and *cmc*. In particular, we show that

$$\text{ecost}(f_{G,H}) \in \Omega \left( \sqrt{c_b c_r \text{cmc}(f_{G,H})} - c_b m \right).$$

Crucially, this bound applies to *any* MHF not just for iMHFs. For iMHFs we could use our pebbling reduction to relate  $\text{ecost}(f_{G,H})$  to  $\text{rbpeb}^\parallel(G)$  and we could use [8] to relate  $\text{cmc}(f_{G,H})$  to  $\Pi_{cc}(G)$ , but no such pebbling reduction is known for dMHFs. Combining our result with a result of Alwen et al. [6] we obtain the following lower bound for *scrypt*:  $\text{ecost}(\text{scrypt}) \in \Omega(\sqrt{c_b c_r} n)$ . While Ren and Devadas [30] previously proved that  $\text{ecost}(\text{scrypt}) \in \Omega(n \cdot c_b)$  their result assumes that the attacker *must* transfer data to/from cache in blocks of size  $w$ . While our lower bound is weaker, it is the first *unconditional* lower bound on the bandwidth hardness of *scrypt*. We conjecture that our results could be extended to show that  $\text{ecost}(\text{scrypt}) \in \Omega(n \cdot c_b)$  without condition.

#### On the Bandwidth Hardness of Important iMHF Candidates.

In Section 5, we provide lower bounds on the bandwidth hardness of several important iMHF candidates including Argon2iA, Argon2iB [12], aATSample and DRSample [4]. We use Argon2iA to refer to v1.1 and we use Argon2iB to refer to versions v1.2+<sup>2</sup>. Thus, Argon2iB (the current version of Argon2i) is particularly important to cryptanalyze as it won the password hashing competition and is being considered for standardization by the Cryptography Form Research Group (CFRG) of the IRTF [13].

<sup>1</sup>In contrast to [3], we use energy cost to refer to bandwidth cost.

<sup>2</sup>The specification of Argon2i has changed several times, but the only changes that affect are analysis are changes that affect the underlying DAG  $G$ . A change to the edge distribution was introduced in v1.2 where a non-uniform indexing was introduced. We use Argon2iB to refer to the version that is currently being considered for standardization by the Cryptography Form Research Group (CFRG) of the IRTF [13].

aATSample and DRSample are important to study as they are the first *practical* iMHF candidate with nearly asymptotically optimal cmc<sup>3</sup>. For the families of graphs generated by aATSample and DRSample [4] we show the following:

**THEOREM 1.3.** *Let  $G$  be a graph generated by aATSample. Then there exists a constant  $C > 0$  so that for all  $m \leq \frac{Cn}{\log n}$ ,*

$$\text{rbpeb}^{\parallel}(G, m) \geq \min(\Omega(n)c_b, \Omega(n \log n)c_r),$$

*holds with high probability.*

**THEOREM 1.4.** *Let  $G$  be a graph generated by DRSample and  $0 < \rho < 1$ . Then there exists a constant  $C > 0$  so that for all  $m \leq Cn^{\rho}$ , with high probability,*

$$\text{rbpeb}^{\parallel}(G, m) \geq \min\left(\Omega(n)c_b, \Omega(n^{3/2-\rho/2})c_r\right).$$

We provide lower bounds on the bandwidth hardness on the family of graphs generated by Argon2iB. The bounds are slightly weaker for Argon2iB in that they only hold if the attacker has cache size  $m \leq Cn^{2/3-\epsilon}$ .

**THEOREM 1.5.** *Let  $G$  be a graph generated by Argon2iB. Then there exists a constant  $C > 0$  so that for any  $0 < \epsilon < 2/3$  and for all  $m \leq Cn^{\frac{2}{3}-\epsilon}$ , with high probability,*

$$\text{rbpeb}^{\parallel}(G, m) \geq \min(\Omega(n)c_b, \Omega(n^{5/3})c_r).$$

At a technical level our template to establish each of these lower bounds is similar. First, we show that the underlying DAG is “block-depth robust.” Second, we show that the graph is “well dispersed.” Essentially, if our block size is  $b$ , then we show that for every interval  $I = [i, j] \subseteq [n/2, n]$  of  $\Omega(n/b)$  nodes in the second half and almost every block  $B$  of  $b$  consecutive nodes in  $[n/2]$  there is an edge from  $B$  to  $I$ . We then consider the pebbling interval  $[t_i, t_j]$  beginning at the time  $t_i$  during which a pebble is first placed on node  $i$  and ending at the time  $t_j$  during which a pebble is first placed on node  $j$ . Because the graph is “well dispersed” we will need to place a red pebble on *at least one node* from almost every block. If the cache size is  $m \in o(n/b)$  then most of these  $\Omega(n/b)$  blocks will begin with no pebbles in cache. Thus, it is either the case that 1) we make  $\Omega(n/b)$  blue moves, or 2) we must re-pebble *almost every* block at some point during the interval  $[t_i, t_j]$ . Because the DAG is block-depth robust this second step will be prohibitively expensive.

**On the Computational Complexity of Minimum Cost Red-Blue Pebbling.** While we can establish asymptotic lower bounds on the energy cost of important iMHF candidates, one would ideally want to find the precise energy cost for each function. In particular, given a graph  $G$  and a cache parameter  $m$  we would like to compute  $\text{rbpeb}^{\parallel}(G, m)$  precisely. However, we show in Appendix D that, unfortunately, exactly computing the red-blue pebbling cost of a DAG  $G$  is NP – Hard, even under realistic assumptions about  $c_b$  and  $c_r$ :

**THEOREM 1.6 (INFORMAL).** *Even for  $c_b \leq 3nc_r$ , the problem of determining the red-blue pebbling cost of a directed acyclic graph  $G$  is NP – Hard.*

<sup>3</sup>In particular,  $\text{cmc}(\text{DRSample}) \in \Omega\left(\frac{n^2 \cdot w}{\log n}\right)$  and  $\text{cmc}(\text{aATSample}) \in \Omega\left(\frac{n^2 \cdot w}{\log n}\right)$  [4] while any iMHF  $f_{G,H}$  has cmc at most  $\text{cmc}(f_{G,H}) \in O\left(\frac{n^2 \cdot w \cdot \log \log n}{\log n}\right)$ .

A result of Demaine and Liu [18, 26] implies that it is PSPACE hard to compute  $\text{rbpeb}^{\parallel}(G, m)$  when  $c_r = 0$  (computation is free)<sup>4</sup>. However, we stress that *in practice* we have  $c_r > 0$  (computation may be *cheap*, but it is not *free*). Furthermore, the decision problem  $\text{rbpeb}^{\parallel} = \text{“is } \text{rbpeb}^{\parallel}(G, m) \leq k \text{”}$  is in NP<sup>5</sup> so the decision problem is fundamentally different when we require  $c_r > 0$ . While the decision problem  $\text{rbpeb}^{\parallel}$  is important for the cryptanalysis of MHFs to the best of our knowledge nothing was known about the complexity of this problem prior to our paper.

Gilbert *et al.* [23] previously showed that the following decision problem was PSPACE complete: Given a DAG  $G$  decide if there is a legal black pebbling with space complexity at most  $m$  i.e., during every pebbling round there are at most  $m$  pebbles on the graph. Gilbert *et al.* showed that the minimum space black pebbling problem was PSPACE – Hard by reduction from the Truly Quantified Boolean Formula (TQBF) problem. Observing that any 3 – SAT instance  $\phi$  with  $n$  variables is also a TQBF instance (albeit with no  $\forall$  quantifiers) their reduction allows us to transform  $\phi$  into a DAG  $G_{\phi}$ . The graph  $G_{\phi}$  has the property that it can be pebbled with at most  $m = 3n + 3$  black pebbles if and only if  $\phi$  is satisfiable. In we detail a gadget to append to  $G_{\phi}$  to create a graph  $H_{\phi}$  so that  $\text{rbpeb}^{\parallel}(H) = x_1$  if  $\phi$  is a satisfiable assignment, but  $\text{rbpeb}^{\parallel}(H_{\phi}) > x_1$  if  $\phi$  is not a satisfiable assignment.

## 2 PRELIMINARIES

We use  $[n]$  to denote the set  $\{1, 2, \dots, n\}$  and  $[a, b] = \{a, a + 1, \dots, b\}$  where  $a, b \in \mathbb{N}$  with  $a \leq b$ . Similarly, we use  $(a, b]$  to denote the set  $[a, b] - \{a\}$ .

We assume a given directed acyclic graph (DAG)  $G = (V, E)$  is labeled in topological order and when  $G$  has size  $n$  we will use  $V = [n]$  to denote the set of vertices. We say a node  $v \in V$  has indegree  $\delta = \text{indeg}(v)$  if there exist  $\delta$  incoming edges  $\delta = |(V \times \{v\}) \cap E|$ . We say that  $G$  has indegree  $\delta = \text{indeg}(G)$  if the maximum indegree of any node of  $G$  is  $\delta$ . A node with indegree 0 is called a source node and a node with no outgoing edges is called a sink. We use  $\text{parents}_G(v) = \{u \in V : (u, v) \in E\}$  to denote the parents of a node  $v \in V$  and similarly for a set  $S \subseteq V$ , we define  $\text{parents}_G(S) = \{u \in V : (u, v) \in E, v \in S\}$ . In general, we use  $\text{ancestors}_G(v) = \bigcup_{i \geq 1} \text{parents}_G^i(v)$  to denote the set of all ancestors of  $v$  – here,  $\text{parents}_G^2(v) = \text{parents}_G(\text{parents}_G(v))$  denotes the grandparents of  $v$  and  $\text{parents}_G^{i+1}(v) = \text{parents}_G(\text{parents}_G^i(v))$ . When  $G$  is clear from context we will simply write  $\text{parents}$  (resp.  $\text{ancestors}$ ). We denote the set of all sinks of  $G$  with  $\text{sinks}(G) = \{v \in V : \nexists (v, u) \in E\}$ , the nodes with no incoming edges.

We often consider the set of all DAGs of equal size  $\mathbb{G}_n = \{G = (V, E) : |V| = n\}$  and often will bound the maximum indegree  $\mathbb{G}_{n,\delta} = \{G \in \mathbb{G}_n : \text{indeg}(G) \leq \delta\}$ . For directed path  $p = (v_1, v_2, \dots, v_z)$  in  $G$ , its length is the number of nodes it traverses,  $\text{length}(p) := z$  (as opposed to the number of edges). We say the

<sup>4</sup>In particular,  $\text{rbpeb}^{\parallel}(G, m) = 0$  if and only if there is a legal black pebbling of  $G$  using at most  $m$  black pebbles where the latter decision problem is PSPACE complete [23].

<sup>5</sup>In particular, for a DAG with  $\text{indeg}(G) = O(1)$  there is always a red-blue pebbling with cost at most  $O(n \cdot c_r + n \cdot c_b)$  – pebble nodes in topological order and never remove a pebble from a node. Thus, the optimal red-blue pebbling runs in time at most  $t \in O\left(n + \frac{n \cdot c_b}{c_r}\right)$  since any red-blue pebbling running in time  $t$  in the parallel random oracle model has cost at least  $t \cdot c_r$ .

depth  $d = \text{depth}(G)$  of DAG  $G$  is the length of the longest directed path in  $G$ .

An iMHF can be specified by a DAG  $G$  and a random oracle  $H$  as in the next definition.

**Definition 2.1.** Given a directed acyclic graph  $G = (V = [n], E)$  with a set of sink nodes  $\text{sinks}(G)$  and a random oracle function  $H : \Sigma^* \rightarrow \Sigma^\ell$  over an alphabet  $\Sigma$ , we define the labeling of graph  $G$  as  $\text{lab}_{G,H} : \Sigma^* \rightarrow \Sigma^*$ . We omit the subscripts  $G, H$  when the dependency on the graph  $G$  and hash function  $H$  is clear. In particular, given an input  $x$  the  $(H, x)$  labeling of  $G$  is defined recursively by

$$\text{lab}_{H,x}(v) = \begin{cases} H(v, x), & \text{indeg}(v) = 0 \\ H(v, \text{lab}_{H,x}(v_1), \dots, \text{lab}_{H,x}(v_d)), & \text{indeg}(v) > 0, \end{cases}$$

where  $v_1, \dots, v_d$  are the parents of  $v$  in  $G$ , according to some pre-determined lexicographical order. We define

$$f_{G,H}(x) = \{\text{lab}_{H,x}(s)\}_{s \in \text{sinks}(G)}.$$

If there is a single sink node  $s_G$  then  $f_{G,H}(x) = \text{lab}_{H,x}(s_G)$ .

We will often consider graphs obtained from other graphs by removing subsets of nodes. Thus if  $S \subset V$ , then let  $G - S$  be the DAG obtained from  $G$  by removing nodes  $S$  and incident edges.

Given a directed acyclic graph (DAG)  $G = (V, E)$  the goal of the red-blue pebbling game is to place pebbles on all sink nodes of  $G$  (not necessarily simultaneously).

Let  $\mathcal{RB} = ((B_0, R_0), (B_1, R_1), \dots, (B_t, R_t))$  (resp.  $\mathcal{RB}^\parallel$ ) denote the set of all sequential (resp. parallel) red-blue pebbblings of a DAG  $G$ . The game is played in rounds and we use  $B_i \subseteq V$  (resp.  $R_i \subseteq V$ ) to denote the set of nodes with blue pebbles (resp. red pebbles) in round  $i$ . Initially, no nodes contain pebbles, so that  $B_0 \cup R_0 = \emptyset$ . The goal is to eventually place a red-pebble on every node in  $V$  (not-necessarily simultaneously) so we require that  $V \subseteq \bigcup_i R_i$ . We also require that in every round  $i > 0$  we have (1) parents  $(R_i \setminus (R_{i-1} \cup B_{i-1})) \subseteq R_{i-1}$ , (2)  $B_i \setminus B_{i-1} \subseteq R_{i-1}$  and (3)  $|R_i| \leq m$  during every time step  $i$ .

We let  $\mathcal{RB}^\parallel(G, m)$  be the set of all valid parallel red-blue pebbblings of  $G$  with a cache-size of  $m$  pebbles. Intuitively, in each round  $i \geq 1$  we may place a red pebble on a node  $v \in V$  if either parents  $(v) \subseteq R_{i-1}$  all of  $v$ 's parents contain red pebbles in the previous configuration (called a *red move*) or  $v$  contained a blue pebble in the previous round (called a *blue move*). On the other hand, we may place a blue pebble at  $v \in P_i$  (also called a *blue move*) if  $v$  contained a red pebble in the previous round. Blue moves represent data transfer to/from memory and are more expensive than red-moves (computation).

We say that a pebbling  $((B_0, R_0), (B_1, R_1), \dots, (B_t, R_t))$  is sequential if  $|R_i \setminus R_{i-1}| \leq 1$  for all  $0 < i \leq t$ , while for a parallel pebbling we make no such restriction. Note that  $\mathcal{RB} \subseteq \mathcal{RB}^\parallel$  since any sequential pebbling is a legal parallel pebbling.

Let  $\#BM_i$  and  $\#RM_i$  denote the number of blue moves and the number of red moves, respectively, during round  $i$ .<sup>6</sup> Formally,

$$\begin{aligned} \#BM_i &= |\{v \in R_i \setminus R_{i-1} : \text{parents}(v) \not\subseteq R_{i-1}\}| + |B_i \setminus B_{i-1}| \\ \#RM_i &= |R_i \setminus R_{i-1}| - |\{v \in R_i \setminus R_{i-1} : \text{parents}(v) \not\subseteq R_{i-1}\}| \end{aligned}$$

<sup>6</sup>In some cases we may have  $v \in B_{i-1}$  and  $\text{parents}(v) \subset R_{i-1}$  so that we could place a pebble on node  $v$  using either a red move or a blue move. In such cases we will assume that this is accomplished by a red move, since blue moves will be more expensive.

Given cost parameters  $c_r$  and  $c_b$ , we define the energy cost of a red-blue pebbling  $(R, B) = ((R_1, B_1), \dots, (R_t, B_t))$  to be

$$\text{rbpeb}^\parallel((R, B)) = \sum_{i=1}^t c_b \#BM_i + c_r \#RM_i.$$

Generally, we assume  $c_b$  is much larger than  $c_r$ . Finally, we define

$$\text{rbpeb}^\parallel(G, m) = \min_{(R, B) \in \mathcal{RB}^\parallel(G, m)} \text{rbpeb}^\parallel((R, B))$$

to be the cost of the optimal red-blue pebbling of  $G$  with maximum cache-size of  $m$  red pebbles.

## 2.1 Depth-Robustness

**Definition 2.2 (Block Depth-Robustness).** Given a node  $v$ , let  $N(v, b) = \{v - b + 1, \dots, v\}$  denote a segment of  $b$  consecutive nodes ending at  $v$ . Similarly, given a set  $S \subseteq V$ , let  $N(S, b) = \bigcup_{v \in S} N(v, b)$ . We say that a DAG  $G$  is  $(e, d, b)$ -block-depth-robust if for every set  $S \subseteq V$  of size  $|S| \leq e$ , we have  $\text{depth}(G - N(S, b)) \geq d$ . If  $b = 1$ , we simply say  $G$  is  $(e, d)$ -depth-robust and if  $G$  is not  $(e, d)$ -depth-robust, we say that  $G$  is  $(e, d)$ -depth-reducible.

Note that when  $b > 1$ ,  $(e, d, b)$ -block-depth robustness is a strictly stronger notion than  $(e, d)$ -depth-robustness since for any set  $S$  with  $|S| \leq e$ , it follows that  $N(S, 1) \subset N(S, b)$ . Hence,  $(e, d, b \geq 1)$ -block depth robustness implies  $(e, d)$ -depth robustness. On the other hand,  $(e, d)$ -depth robustness only implies  $(e/b, d, b)$ -block depth robustness.

The cumulative memory complexity of an iMHF is very closely related to the notion of depth-robustness [2, 4, 5, 15]. In particular, we know that  $\Pi_{cc}^\parallel(G) \geq ed$  for a depth-robust DAG and that  $\Pi_{cc}^\parallel(G) \in \mathcal{O}\left(en + n \cdot \sqrt{dn}\right)$ . We will show that  $\Pi_{cc}^\parallel(G)$  can be used to lower bound  $\text{rbpeb}^\parallel(G, m)$ , thus depth-robustness can also be a useful tool in bandwidth hardness.

## 2.2 Metagraphs

We will also frequently use the notion of a metagraph in our analysis. For a fixed integer  $m \in [n]$ , let  $n' = \lfloor n/m \rfloor$ . For all  $i \in [n']$ , let  $M_i = [(i-1)m + 1, im] \subseteq V$ . Moreover, we denote the first and last thirds respectively of  $M_i$  with

$$M_i^F = \left[ (i-1)m + 1, (i-1)m + \left\lfloor \frac{m}{3} \right\rfloor \right] \subseteq M_i,$$

$$M_i^L = \left[ (i-1)m + \left\lceil \frac{2m}{3} \right\rceil + 1, im \right] \subseteq M_i.$$

Given a DAG  $G$ , we call a DAG  $G_m = (V_m, E_m)$  with the following properties a *metagraph* of  $G$ .

- **Nodes:**  $V_m$  contains one node  $v_i$  per set  $M_i$ , i.e.,  $V_m = \{v_i : i \in [n']\}$ . We call  $v_i$  the *simple node* and  $M_i$  its *meta-node*.
- **Edges:** If the end of a meta-node  $M_i^L$  is connected to the beginning  $M_j^F$  of another meta-node, we connect their simple nodes, i.e.,  $E_m = \{(v_i, v_j) : E \cap (M_i^L \times M_j^F) \neq \emptyset\}$ .

### 3 MODELING ENERGY COMPLEXITY AS RED-BLUE PEBBLING

In this section we show that the energy cost of the function  $f_{G,H}$  is characterized by the red-blue pebbling cost  $\text{rbpeb}^{\parallel}(G, m)$  in the parallel random oracle model just as Alwen and Serbinenko [8] showed that cumulative memory complexity can be characterized by the black pebbling game. Similar to [8] our reduction uses Lemma 3.1 as a core building block. In particular, if the energy cost is significantly smaller than  $\text{rbpeb}^{\parallel}(G, 8m)$  for a pROM attacker with  $m \cdot w$  bits of cache then we can build an extractor that receives a small hint and predicts the random oracle output on a larger set of indices contradicting Lemma 3.1. One of the unique challenges we face when designing our extractor is that it is not obvious how to relate messages between cache and main memory to specific blue pebbling moves. By contrast, a black pebbling move always corresponds to a specific random oracle query.

**LEMMA 3.1.** [21] *Let  $H$  be a set of hints that can be given,  $B$  be a series of random bits and  $\mathcal{A}$  be an algorithm that receives as input some hint  $h \in H$  and can adaptively query  $B$  at specific indices. Let  $\text{WIN}_{\mathcal{A},h}$  denote the event that  $\mathcal{A}$ , given  $h \in H$  as input, eventually outputs a subset of  $k$  indices  $i_1, \dots, i_k$  that were not previously queried as well as the corresponding values  $B[i_1], \dots, B[i_k]$  of each bit then*

$$\Pr [\exists h \in H. \text{WIN}_{\mathcal{A},h}] \leq \frac{|H|}{2^k},$$

where the randomness is taken over the selection of  $B$ .

#### 3.1 Memory and Cache in the Parallel Random Oracle Model

Before we present our reduction it is first necessary to give a formal definition of energy costs in the pROM model.

We define a state of an algorithm  $\mathcal{A}^{H(\cdot)}$  to be the tuple  $(\sigma, \xi)$ , where  $\sigma$  contains the contents of the cache and has size at most  $mw$  bits, and  $\xi$  contains the contents of the memory. We consider a pROM attacker  $\mathcal{A}^{H(\cdot)}$  with cache size  $m \cdot w$  who is given oracle access to a random oracle  $H : \{0, 1\}^* \rightarrow \{0, 1\}^w$ . In particular, the cache is large enough to store  $m$  labels. An execution of  $\mathcal{A}^{H(\cdot)}$  on input  $x$  proceeds in rounds as follows. Initially, the state at time 0 is  $(\sigma_0, \xi_0)$  where  $\xi_0$  is empty and  $\sigma_0$  encodes the initial input  $x$ . At the beginning of round  $i$  the attacker is given the initial state  $(\sigma_{i-1}, \xi_{i-1})$  as well as the answers  $A_{i-1}$  to any random oracle queries that were asked at the end of the last round. The algorithm  $\mathcal{A}^{H(\cdot)}$  may then perform arbitrary computation and/or transfer data between memory and cache. The round ends when the attacker outputs a new state  $(\sigma_i, \xi_i)$  along with a batch of queries  $Q_i = \{q_{i,1}^1, q_{i,2}^1, \dots, q_{i,k_i}^1\}$ . Since the attacker only has cache-size  $m \cdot w$  we only allow the attacker to make at most  $|Q_i| \leq m$  queries during a single step (otherwise the attacker won't even have room to store all of the responses).

We allow the attacker to specify arbitrary functions  $F_1, F_2, F_3$  and  $F_4$  for communication between cache and memory during each round so long as the specification of each function is independent of the random oracle  $H$  (e.g., we cannot query the random oracle in between rounds). In particular, the function  $F_1(\sigma_{i-1}, A_{i-1}) = r_i^1$  is used to specify the first message we will send to memory during

round  $i$  — in the event that we don't send any message to memory we define  $F_1(\sigma_{i-1}) = \perp$ . Similarly, the function  $F_2(\xi_{i-1}, r_i^1) = s_i^1$  specifies the response from memory (or  $\perp$  if there is no response). Once  $r_i^1, s_i^1, \dots, r_i^{j-1}, s_i^{j-1}$  have been defined we set

$$r_i^j = F_1(\sigma_{i-1}, A_{i-1}, r_i^1, s_i^1, \dots, r_i^{j-1}, s_i^{j-1}),$$

$$s_i^j = F_2(\xi_{i-1}, r_i^1, s_i^1, \dots, r_i^{j-1}, s_i^{j-1}, r_i^j).$$

We terminate when  $r_i^j = \perp$  or when  $s_i^j = \perp$ .

We let  $R_i = \{r_i^1, r_i^2, \dots, r_i^{\ell_i}\}$  denote the sequence of messages sent from cache to memory during round  $i$  and we let

$S_i = \{s_i^1, s_i^2, \dots, s_i^{\ell_i}\}$  denote the responses sent from memory back to the cache. Finally, the round ends when the attacker uses the function  $F_3(\xi_{i-1}, R_i, S_i) = \xi_i$  to output a new state  $\xi_i$  for memory and  $F_4(\sigma_{i-1}, R_i, S_i)$  to output a new state  $\sigma_i$  for cache and a new batch  $Q_i$  of at most  $m$  random oracle queries. At this point  $\mathcal{A}^{H(\cdot)}$  outputs the next state  $(\sigma_i, \xi_i)$  along with the next batch of queries  $Q_i$ .

Crucially, the functions  $F_2$  and  $F_3$ , which are used to generate response from main memory and update the state of main memory at the end of the round, do not have access to  $\sigma_{i-1}$  (the state of cache) or  $A_{i-1}$  (the answers to random oracle queries). In particular, any information about  $\sigma_{i-1}$  (cache-state) and  $A_{i-1}$  (most recent answers to random oracle queries) that main memory receives must be communicated through one of the messages in the set  $R_i$ . Similarly, the functions  $F_1$  and  $F_4$  are used to generate the requests sent from cache to main memory, to update the state of cache  $\sigma_i$  at the end of the round and to output the next batch  $Q_i$  of random oracle queries. Crucially these functions do not have access to  $\xi_{i-1}$  (the state of memory). Thus, any information about  $\xi_{i-1}$  must be communicated through one of the responses in the set  $S_i$ .

Dziembowski *et al.* [20] also addresses communication between two parties,  $A_{\text{small}}$  (e.g., a space-bounded virus) and  $A_{\text{big}}$ , over a bounded channel. However, both parties in this model can query the random oracle. This is a crucial difference, since one of the parties in our model, the main memory, is strictly forbidden from querying the random oracle to avoid trivialization of the problem (e.g., the attacker can perform all computation in RAM with no blue moves).

**Execution Trace.** We define the execution trace of the algorithm  $\mathcal{A}^{H(\cdot)}$  by the sequence of cache states, memory states, messages passed between cache and memory, and queries made to the random oracle  $H$ . Formally, the execution trace is  $\text{Trace}_{\mathcal{A},R,H}(x) = \{(\sigma_i, \xi_i, R_i, S_i, Q_i)\}_{i=1}^t$ , where the trace  $\text{Trace}_{\mathcal{A},R,H}(x)$  is dependent on the algorithm  $\mathcal{A}^{H(\cdot)}$ , random oracle  $H$ , internal randomness  $R$ , and input value  $x$ . Then we say the cost of the execution trace is

$$\text{cost}(\text{Trace}_{\mathcal{A},R,H}(x)) = \sum_{i=1}^t \left( c_r k_i + \sum_{j=1}^{\ell_i} \frac{c_b}{w} (|r_i^j| + |s_i^j|) \right).$$

Intuitively, the  $c_r$  term is the cost of all of the queries we make to the random oracle  $H$  and the  $c_b$  terms result from the messages passed between cache and memory — here  $c_b$  denotes the cost of transferring  $w$  bits between cache and memory.



We now formally define the energy cost of computing a function based on its execution trace.

**Definition 3.2.** Given constants  $c_b$  and  $c_r$ , the energy cost  $\text{ecost}$  of a function  $f_{G,H}$  is defined by

$$\text{ecost}_{q,\epsilon}(f_{G,H}, m \cdot w) = \min_{\mathcal{A}, x} \mathbb{E} [\text{cost}(\text{Trace}_{\mathcal{A},R,H}(x))],$$

where the minimum of the expected cost is taken over all valid inputs  $x$  and all algorithms  $\mathcal{A}$  with cache size  $m \cdot w$  bits making at most  $q$  queries that compute  $f_{G,H}(x)$  correctly with probability at least  $\epsilon$ .

### 3.2 Red-Blue Extension Pebbling

We are now ready to prove our main result in this section. Theorem 3.3 lower bounds the energy cost  $\text{ecost}_{q,\epsilon}(f_{G,H}, m \cdot w)$  of the function  $f_{G,H}$  with cache size  $m \cdot w$  using  $\text{rbpeb}^\parallel(G, 8m)$  the red-blue pebbling cost of the DAG  $G$  with  $8m$  red pebbles.

**THEOREM 3.3.** For any DAG  $G$  with  $n$  nodes and any  $\mathcal{A}_{mw}^{H(\cdot)}$  making at most  $q < 2^{w/20}$  queries, then for  $4 \log n < w$ ,

$$\text{ecost}_{q,\epsilon}(f_{G,H}, m \cdot w) \geq \frac{\epsilon}{16} \text{rbpeb}^\parallel(G, 8m).$$

Given a DAG  $G$  and a legal black pebbling  $P = (P_1, \dots, P_t) \in \mathcal{P}^\parallel(G)$  with  $|P_{i+1} \setminus P_i| \leq m$  we say that a (legal) red-blue pebbling  $((B_1, R_1), \dots, (B_t, R_t)) \in \mathcal{RB}^\parallel(G, m)$  is a  $k$ -extension of a black pebbling  $P$  if for all  $i \in [t]$  we can find a small set  $E_i \subseteq V(G)$  such that  $|E_i| \leq k$ ,  $P_i \subseteq B_i \cup R_i$  and  $R_i \cup B_i = P_i \cup E_i$ . We let  $\text{RBExt}(P, m, k) \subseteq \mathcal{RB}^\parallel(G, m)$  denote the set of all possible  $k$ -extensions of  $P$ . To prove Theorem 3.3 we extract a legal black pebbling  $P = (P_1, \dots, P_t) \in \mathcal{P}^\parallel(G)$  from the execution trace of  $\mathcal{A}^{H(\cdot)}$ , and then use  $P$  to build a legal extension pebbling  $((B_1, R_1), \dots, (B_t, R_t)) \in \text{RBExt}(P, m, 7m)$ . Our extension pebbling may use up to  $8m = (m + 7m)$  red-pebbles, but the pROM attacker  $\mathcal{A}$  has a cache size  $mw$  bits — enough to store exactly  $m$  labels. We then show how to *upper bound* the cost of the extension pebbling and *lower bound* the energy cost of of the attacker  $\mathcal{A}$  in the random oracle model.

**Step 1:** Since the  $\epsilon$  term is needed to account for algorithms that fail cheaply, we focus on the instances in which  $\mathcal{A}_{mw}^{H(\cdot)}$  correctly computes  $f_{G,H}$ . We start by using  $\mathcal{A}_{mw}^{H(\cdot)}$  to extract a legal black pebbling following Alwen and Serbinenko [8]. Given an execution trace  $\text{Trace}_{\mathcal{A},R,H}(x)$ , the corresponding black pebbling  $\text{BlackPebble}^H(\text{Trace}_{\mathcal{A},R,H}(x)) = P_0, \dots, P_t$  is defined by setting  $P_0 = \emptyset$  and the pebbles  $P_i$  at each subsequent time step  $i$  based on the corresponding batch of queries  $q_i$  and the following rules:

- For each query  $q$  in batch  $q_i$ , if the query has the form  $v, \text{lab}_{H,x}(v_1), \dots, \text{lab}_{H,x}(v_d)$  for some vertex  $v$  and its parents  $v_1, \dots, v_d$ , then we add a pebble to node  $v$  in  $P_i$ .
- If  $v$  is never used as input for a query in the remaining execution or if there exists another query for  $v$  before  $v$  is used, then  $v$  is deleted from  $P_i$ .

Intuitively, at each time  $j$ ,  $P_j$  contains all nodes  $v$  whose label will appear as input to a future random oracle query *before* the label appears as the output of a random oracle query. Alwen and Serbinenko showed that the black pebbling constructed this way is legal with overwhelming probability.

**THEOREM 3.4.** [8] The pebbling extracted from an execution trace,

$$\text{BlackPebble}^H(\text{Trace}_{\mathcal{A},R,H}(x)) \in \mathcal{P}^\parallel(G),$$

is a legal black pebbling with probability at least  $1 - \frac{q}{2^w}$ , where  $w$  is the label size and  $q$  is the number of queries made by  $\text{Trace}_{\mathcal{A},R,H}$ .

Once our black pebbling has been defined we can define the set  $\text{QueryFirst}(t_1, t_2)$  to be the vertices whose data-labels that will appear as input to a random oracle query during rounds  $[t_1, t_2]$  before the data-label appears as the *output* of some random oracle query made during the same interval. Formally, given  $P$  and an interval  $[t_1, t_2]$  we let

$$\text{QueryFirst}(t_1, t_2) = \bigcup_{i=t_1}^{t_2} \left( \text{parents}(P_i \setminus P_{i-1}) \setminus \left( \bigcup_{j=t_1}^{i-1} (P_j \setminus P_{j-1}) \right) \right).$$

We present a few properties about  $\text{QueryFirst}$  that we will use in the rest of the proof.

**LEMMA 3.5.**  $\forall 0 \leq x < y < z \leq t$ ,

$$\text{QueryFirst}(y, z) \setminus \text{QueryFirst}(x, z) \subseteq \bigcup_{i=x}^y (P_i \setminus P_{i-1}).$$

By definition, any node in the set  $\text{QueryFirst}(y, z)$  but not in the set  $\text{QueryFirst}(x, z)$  for  $y > x$  must appear as input to a random oracle query during rounds  $[x, y]$ . Hence, the left-hand side is a subset of the nodes that are newly pebbled between rounds  $[x, y]$ .

**Step 2:** We partition the pebbling rounds  $[t]$  into sub time-intervals  $(t_0 = 0, t_1], (t_1, t_2], \dots$  recursively as follows. Let  $t_1$  be the minimum pebbling round such that there exists  $j < t_1$  such that  $|\text{QueryFirst}(j, t_1)| \geq 3m$ . As a special case, if  $|\text{QueryFirst}(i, j)| \leq 3m$  for all  $i < j \leq t$  (i.e., no such intervals exist), then set  $t_1 = t$  and output  $(t_0, t_1]$ . In this case, there is a red-blue extension pebbling in  $\text{RBExt}(P, 8m, k)$  that requires 0 blue moves at at most  $\sum_i |P_i \setminus P_{i-1}|$  red-moves.

Once  $t_1 < \dots < t_{i-1}$  have been defined we inductively define  $t_i > t_{i-1}$  to be the minimum round such that there exists  $t_{i-1} \leq j < t_i$  such that  $|\text{QueryFirst}(j, t_i)| \geq 3m$  — if no such  $t_i$  exists then we set  $t_i = t$ .

**Step 3:** We will show that there is an extension pebbling that makes at most  $4m$  blue moves during each interval (except for the first one where it needs 0). In particular, we set  $k = 7m$  and we will define an extension pebbling  $(B^*, R^*) \in \text{RBExt}(P, 8m, k)$  by dividing the cache into two sets of size  $4m$  denoted as  $R_j^{\text{inter}}$  and  $R_j^{\text{legal}}$ , respectively. We will show that  $R_i = R_i^{\text{legal}} \cup R_i^{\text{inter}}$  and  $B_i \supset P_i$  give a legal red-blue pebbling and then bound its cost.

We set  $R_{t_i}^{\text{inter}} = \{\}$  at the start of each time interval  $(t_i, t_{i+1}]$  and for each  $j \in (t_i, t_{i+1}]$  we have

$$R_j^{\text{inter}} = (R_{j-1}^{\text{inter}} \cup (P_j \setminus P_{j-1})) \cap \text{QueryFirst}(j, t_{i+1}).$$

Intuitively,  $R_j^{\text{inter}}$  stores all of the red-pebbles we have computed during the interval  $(t_i, j]$  that are later needed in the interval  $[j + 1, t_{i+1})$ . Thus, any node that is pebbled during rounds  $(t_i, j]$  and subsequently needed in round  $(j + 1, t_{i+1})$  must be in  $R_j^{\text{inter}}$ , which we will keep in cache. This yields the following invariant.



INVARIANT 1. For any  $j \in (t_i, t_{i+1})$ ,

$$\text{QueryFirst}(j+1, t_{i+1}) \cap \bigcup_{i=t_i}^j (P_i \setminus P_{i-1}) \subseteq R_j^{\text{inter}}$$

To maintain legality across all time steps, we add a few rules about blue moves:

- (1) We convert a pebbled node  $v$  from blue to red if node  $v$  is in  $\text{QueryFirst}(t_i, t_{i+1})$ . That is for any  $j \in [t_i, t_{i+1})$ , we define  $R_j^{\text{legal}} = \text{QueryFirst}(t_i, t_{i+1})$ .
- (2) We convert a pebbled node  $v \in R_j$  from red to blue at time  $j$  if node  $v$  is in  $\text{QueryFirst}(t_i, t_{i+1})$  for some later interval  $(t_i, t_{i+1})$  with  $j < t_i$  and if  $v \notin B_j$  is not already stored in memory. In this case it will be helpful to “charge” the cost  $c_b$  of this blue move to the future interval  $(t_i, t_{i+1})$ .

We show the following bound on the size of  $\text{QueryFirst}(j, t_{i+1})$ .

LEMMA 3.6.  $\forall j \in (t_i, t_{i+1}), |\text{QueryFirst}(j, t_{i+1})| \leq 4m$ .

PROOF. By the definition of  $\text{QueryFirst}$ ,  $\text{QueryFirst}(j, t_{i+1}) \subseteq \text{QueryFirst}(j, t_{i+1}-1) \cup \text{parents}(P_{t_{i+1}} \setminus P_{t_{i+1}-1})$ . Due to our choice of  $t_{i+1}$ ,  $\text{QueryFirst}(j, t_{i+1}-1) \leq 3m$ . Since parallelism is bounded by cache size,  $\text{parents}(P_{t_{i+1}} \setminus P_{t_{i+1}-1}) \leq m$ . The lemma then follows.  $\square$

LEMMA 3.7.  $\{R_i\} = \{R_i^{\text{legal}} \cup R_i^{\text{inter}}\}$  is a legal pebbling.

PROOF. First observe  $\text{parents}(P_{j+1} \setminus P_j) \subseteq \text{QueryFirst}(j, t_{i+1})$ . We now prove  $\text{QueryFirst}(j, t_{i+1}) \subseteq R_j$ . Note that any node in  $\text{QueryFirst}(j, t_{i+1})$  must either be in  $\text{QueryFirst}(t_i, t_{i+1})$  or have been pebbled at some point during time steps  $(t_i, j)$ . In the former case, the node would be in  $R_j^{\text{legal}}$ , and in the latter case, the node would be in  $R_j^{\text{inter}}$ . Thus,  $\text{parents}(P_{j+1} \setminus P_j) \subseteq \text{QueryFirst}(j, t_{i+1}) \subseteq R_j$ , and the pebbling is legal.  $\square$

LEMMA 3.8.  $|R_j^{\text{inter}}| \leq 4m$ .

PROOF. Observe that  $R_j^{\text{inter}} \subseteq \text{QueryFirst}(j+1, t_{i+1})$  since elements are only kept in  $R_j^{\text{inter}}$  if they are needed for some later pebbling round.  $|\text{QueryFirst}(j+1, t_{i+1})| \leq 4m$  by Lemma 3.6.  $\square$

Also note that  $|R_{t_i}^{\text{legal}}| = |\text{QueryFirst}(t_i, t_{i+1})| \leq 4m$ . So the extension red-blue pebbling we constructed stores at most  $8m$  labels in cache at any time.

We now bound the cost of the above extension pebbling. Since we never discard necessary red pebbles from  $R_{t_i}^{\text{inter}}$ , the only cache-misses we incur come from  $R_{t_i}^{\text{legal}}$ , at most  $4m$ . We “charge” double for every cache-miss to account for the previous blue move that initially placed a blue pebble on a node. This way, we can also charge the cost of placing new blue pebbles to future rounds. Therefore, the above extension pebbling has cost at most

$$8mc_b + \sum_{j \in (t_i, t_{i+1}]} c_r |P_j \setminus P_{j-1}|.$$

**Step 4:** To complete the proof, we show that during each interval any algorithm  $\mathcal{A}$  must pay red-blue cost at least  $mc_b$  +

$\sum_{j \in (t_i, t_{i+1}]} c_r |P_j \setminus P_{j-1}|$ . Roughly speaking, we will set up an extractor that extracts  $3m$  random oracle labels (i.e.,  $3mw$  truly random bits) by simulating  $\mathcal{A}$  during this time interval. The extractor needs a hint of size  $mw + w(\#words_i)$  where  $\#words_i$  is the total amount of data  $\mathcal{A}$  transfers to/from cache. If  $\#words_i \leq m$  then we will arrive at a contradiction as we compressed a random string of length  $3mw$  – contradicting Lemma 3.1. Thus,  $\mathcal{A}$  must pay blue cost at least  $mc_b$  during each interval, and by construction of  $P = \text{BlackPebble}^H(\text{Trace}_{\mathcal{A}, R, H}(x))$  the red cost is at least  $\sum_{j \in (t_i, t_{i+1}]} c_r |P_j \setminus P_{j-1}|$ . We detail this step in the next section.

### 3.3 Extractor

We now show that if an attacking strategy does not yield a corresponding legal red-blue pebbling, then it can be modified to form an extractor for the labels of a subset of nodes. That is, an extractor with access to the attacking strategy, the state of the cache, and a few select hints can successfully predict a large number of random bits, contradicting Lemma 3.1. The hints we give the extractor will dictate the location of the random bits, and ensure these bits remain “random” (that is, not queried by the extractor). Figure 1 illustrates this setup. In particular, the extractor will use a hint to simulate  $\mathcal{A}^{H(\cdot)}$  but this hint *does not* include the current state of memory  $\xi_i$ . Instead, the hint will encode the messages that the attacker expects to receive from main memory which allows us to *simulate* the attacker without storing the entire (large) state  $\xi_i$ .




Fig. 1: Using the attacker to create an extractor that tries to predict  $3m$  distinct outputs of random oracle  $H(\cdot)$ .

LEMMA 3.9. Given a randomly sampled execution trace

$$\text{Trace}_{\mathcal{A}, R, H}(x) = \{(\sigma_i, \xi_i, R_i, S_i, Q_i)\}_{i=1}^t,$$

let  $\text{BlackPebble}^H(\text{Trace}_{\mathcal{A}, R, H}(x))$  be the extracted black pebbling and  $(t_i, t_{i+1})$  be defined as above. If  $q < 2^{w/20}$  and  $4 \log n < w$ , then with probability at least  $1 - \frac{2}{2^{w/2}}$ , the following holds for all  $i$ :

$$\sum_{i=t_i}^{t_{i+1}} \left( c_r k_i + \sum_{j=1}^{\ell_i} \frac{c_b}{w} (|r_i^j| + |s_i^j|) \right) \geq m,$$

where  $q$  is the total number of random oracle queries made in the execution trace,  $n$  is the number of nodes in the underlying DAG, and the probability is taken over the random coins of  $\mathcal{A}$  and the selection of the random oracle  $H$ .

PROOF. Suppose, by way of contradiction, that for interval  $(t_i, t_{i+1})$ , an attacker transfers less than  $mw$  bits between cache and memory. We define an extractor that can predict  $3m$  labels given access to

the attacker's algorithm, the random oracle, and a small set of hints to help the extractor. Recall that

$$\text{lab}_{H,x}(v) = H(v, \text{lab}_{H,x}(v_1), \dots, \text{lab}_{H,x}(v_d))$$

for all nodes after the first node. Thus for nodes  $y \neq z$ , the values of  $\text{lab}_{H,x}(y)$  and  $\text{lab}_{H,x}(z)$  correspond to different inputs to  $H$ . That is, there are no *input* collisions and so the adversary must separately determine the hash outputs for each of the  $3m$  inputs, which correspond to  $3mw$  truly random bits in total.

The hint given to help the extractor consists of five components:

- (1) The set  $\text{QueryFirst}(t_i, t_{i+1})$  is given as a hint to denote the indices that form the string that the extractor will ultimately predict. Since  $|\text{QueryFirst}(t_i, t_{i+1})| \leq 4m$ , this component of the hint is at most  $4m \log n$  bits.
- (2) For each  $v \in \text{QueryFirst}(t_i, t_{i+1})$ , the index of the first query that appears in which  $\text{lab}(v)$  is needed as input. This component of the hint tells the extractor the queries that require the prediction of random strings, and has size at most  $4m \log q$  bits, where  $q$  is the total number of queries made by the attacker.
- (3) For each  $v \in \text{QueryFirst}(t_i, t_{i+1})$ , the index of the first query when  $\text{lab}(v)$  might be compromised. Observe that if the extractor successfully predicts a random string at a location  $v$ , but then  $\text{lab}(v)$  is later queried by the attacker, we cannot distinguish this case at the end from the case that the extractor simply read  $\text{lab}(v)$  after making the query. In the later case, the extractor is not predicting a random string at all! To avoid this, we give the extractor a hint of the queries that would compromise the randomness of the desired locations. Formally, the hint is the minimal index  $i$  such that  $q_i^j = v$ , which returns the query  $H(q_i^j) = \text{lab}(v)$ . This component of the hint tells the extractor the locations of the random strings to be predicted, and has size at most  $4m \log q$  bits.
- (4) The cache state at  $t_i$  is given as a hint to the extractor to simulate the attacker beginning at time step  $t_i$ . Since the cache has size  $m$ , each containing  $w$ -bit words, the size of this component of the hint is at most  $mw$  bits.
- (5) Messages between the cache and memory during time steps  $(t_i, t_{i+1})$  are also given as a hint to the extractor to simulate the attacker beginning at time step  $t_i$ . By assumption, the attacker transfers less than  $mw$  bits between cache and memory, so the size of this component of the hint is at most  $mw$  bits in total.

Since  $q < 2^{w/20}$  and  $4 \log n < w$ , then the total size, in bits, of the hint is at most

$$4m \log n + 4m \log q + 4m \log q + mw + mw \leq \frac{13}{5}mw.$$

However,  $|\text{QueryFirst}(t_i, t_{i+1})| \geq 3m$ , so the extractor must successfully predict the output of  $3m$  hash outputs, each of size  $w$ , given a hint of size at most  $\frac{13}{5}mw$  bits. Thus, the extractor must predict  $3mw - \frac{13}{5}mw = \frac{2}{5}mw$  random bits. The probability that the extractor predicts  $\frac{2}{5}mw$  random bits is  $2^{-2mw/5}$ , so the probability that the extractor does not predict  $\frac{2}{5}mw$  random bits over any set of time intervals  $(t_i, t_{i+1})$  is at least  $1 - \frac{n^2}{2^{2mw/5}} > 1 - \frac{1}{2^{mw/5}}$ , for sufficiently large  $m$ , by a union bound over the intervals. Since the

probability that the execution does not yield a legal black pebbling is  $\frac{q}{2^w}$  by Theorem 3.4, then the probability that an attacker must transfer  $mw$  bits between cache and memory for each time interval  $(t_i, t_{i+1})$  is at least  $1 - \frac{1}{2^{mw/5}} - \frac{q}{2^w} > 1 - \frac{2}{2^{w/2}}$ .  $\square$

We now justify the correctness of Theorem 3.3.

**Proof of Theorem 3.3:** Recall that  $\text{ecost}_{q,\epsilon}(f_{G,H}, m \cdot w)$  is taken over algorithms that compute  $f_{G,H}(x)$  correctly with probability at least  $\epsilon$ . Thus,  $\text{ecost}_{q,\epsilon}(f_{G,H}, m \cdot w)$  is at least  $\epsilon$  times the expected cost of an execution trace that correctly computes  $f_{G,H}(x)$ . Lemma 3.9 implies that an execution trace that correctly computes  $f_{G,H}(x)$  must transfer at least  $m$  words between memory and cache for each interval  $(t_i, t_{i+1})$  with probability at least  $(1 - \frac{2}{2^{w/2}}) > \frac{1}{2}$ . Recall that by construction, the red-cost for each interval is at least  $\sum_{j \in (t_i, t_{i+1})} c_r |P_j \setminus P_{j-1}|$ . Therefore with probability at least  $\frac{1}{2}$ , the energy cost of an execution trace that correctly computes  $f_{G,H}(x)$  is at least  $mc_b + \sum_{j \in (t_i, t_{i+1})} c_r |P_j \setminus P_{j-1}|$  for each interval  $(t_i, t_{i+1})$ . On the other hand, recall that

$$\text{rbpeb}^\parallel(B^*, R^*) \leq \sum_i \left( 8mc_b + \sum_{j \in (t_i, t_{i+1})} c_r |P_j \setminus P_{j-1}| \right).$$

Hence,

$$\text{ecost}_{q,\epsilon}(f_{G,H}, m \cdot w) > \frac{\epsilon}{16} \text{rbpeb}^\parallel(B^*, R^*).$$

$\square$

## 4 RELATING MEMORY HARDNESS AND BANDWIDTH HARDNESS

In this section, we show that any function with high cumulative memory complexity also has high energy costs. Namely,

**Reminder of Theorem 1.2.**

$$\text{rbpeb}^\parallel(G, m) \geq 2c_b \left( \frac{\Pi_{cc}(G)}{t} - 2m \right) + c_r t \in \Omega \left( \sqrt{c_b \cdot c_r \cdot \Pi_{cc}(G)} \right),$$

where  $m$  is the cache size,  $t$  is the number of steps in the pebbling,  $c_b$  is the cost of a blue move and  $c_r$  is the cost of a red move.

We also show that this connection can be exploited to design a maximally bandwidth hard iMHF. Thus, the goals of designing an MHF with high cumulative memory complexity/bandwidth hardness are well aligned.

$$\text{LEMMA 4.1. } \text{rbpeb}^\parallel(G, m) \geq \min_t \left( 2c_b \left( \frac{\Pi_{cc}(G)}{t} - m \right), tc_r \right).$$

**PROOF.** For any red-blue pebbling  $P$  of DAG  $G$ , let  $R_i$  be the set of red pebbles at time step  $i$  and let  $B_i$  be the set of blue pebbles at time step  $i$ . Setting  $D_i = B_i \cup R_i$  we remark that  $(D_1, \dots, D_t)$  is a valid black pebbling of  $G$ . Thus, by the optimality of  $\Pi_{cc}^\parallel(G)$ ,

$$\Pi_{cc}^\parallel(G) \leq \sum_{i=1}^t |R_i \cup B_i| \leq \sum_{i=1}^t |R_i| + \sum_{i=1}^t |B_i| \leq t \max |B_i| + tm$$

Rearranging terms we have

$$\max_i |B_i| \geq \frac{\Pi_{cc}^\parallel(G)}{t} - m.$$

In the optimal red-blue pebbling, each blue pebble must eventually be converted back to a red pebble, or else it should be discarded. Additionally, without loss of generality, we can assume that during each step we make at least one red move. Otherwise, we could combine consecutive steps into one single step. Thus,

$$\begin{aligned} \text{rbpeb}^\parallel(G, m) &\geq 2 \left| \bigcup_{i=1}^t B_i \right| c_b + t c_r \\ &\geq 2 \max_i |B_i| c_b + t c_r \\ &\geq 2 \left( \frac{\Pi_{cc}^\parallel(G)}{t} - m \right) c_b + t c_r \\ &\geq \min_t \left( 2 \left( \frac{\Pi_{cc}^\parallel(G)}{t} - m \right) c_b, t c_r \right) \end{aligned}$$

□

COROLLARY 4.2. For an  $(e, d)$ -depth robust graph  $G$ ,

$$\text{rbpeb}^\parallel(G, m) \geq \min_t \left( 2 \left( \frac{ed}{t} - m \right) c_b, t c_r \right).$$

PROOF. An  $(e, d)$ -depth robust DAG  $G$  has  $ed \leq \Pi_{cc}^\parallel(G)$  [5]. □

We show that there exists a similar relationship between sequential black pebbling cost and sequential red-blue pebbling cost.

THEOREM 4.3.

$$\text{rbpeb}(G) \geq 2c_b \left( \frac{\Pi_{cc}(G)}{t} - m \right) + c_r t,$$

where  $m$  is the cache size,  $t$  is the number of steps in the pebbling,  $c_b$  is the cost of a blue move and  $c_r$  is the cost of a red move.

PROOF. Given a sequential black pebbling  $P_1, \dots, P_t$  of  $G$ , let  $B_i$  be the set of blue pebbles at time step  $i$ . Then

$$\max_i |B_i| \geq \max_i (|P_i| - m) \geq \left( \frac{\Pi_{cc}(G)}{t} - m \right),$$

where the last step results from a simple averaging argument over all  $t$  steps. Finally, each item in  $B_i$  requires cost  $c_b$  to load into cache and another cost  $c_b$  to be stored in memory (if the item is not ever retrieved from memory, it would not be in  $B_i$  for an optimal pebbling). □

Theorem 1.2 can also be related to parallel pebbling through the following lemma:

LEMMA 4.4.  $\text{rbpeb}(G, 2m) \leq \text{rbpeb}^\parallel(G, m) \leq \text{rbpeb}(G, m)$ .

PROOF.  $\text{rbpeb}^\parallel(G, m) \leq \text{rbpeb}(G, m)$  follows immediately from definition.<sup>7</sup> Now consider  $\text{rbpeb}(G, 2m)$  and  $\text{rbpeb}^\parallel(G, m)$ . Any parallel pebbling with cache size  $m$  can be performed by a sequential pebbling with cache size  $2m$ . Note that at any step, a parallel pebbling with cache size  $m$  can have at most  $m$  labels stored and  $m$  new pebbles placed in each step. Thus, a sequential pebbling with

<sup>7</sup>To see that  $\text{rbpeb}^\parallel(G, m)$  and  $\text{rbpeb}(G, m)$  are not identically equivalent quantities, consider the complete directed bipartite graph  $K_{m,m}$  with  $m$  sources  $A$  and  $m$  sink nodes  $B$  ( $m$  is also the cache size). In the parallel model we can finish in two steps with zero blue moves:  $R_0 = \emptyset$ ,  $R_1 = A$ ,  $R_2 = B$ . In the sequential pebble game we would have to keep pebbles on  $A$  while we begin placing pebbles on  $B$  one by one. Each time we place a red-pebble on a node  $y \in B$  we need to evict some node  $x \in A$  by converting  $x$  into a blue node (and then bring it back into the cache-later).

cache size  $2m$  can emulate this by retaining the stored labels while adding the new pebbles one by one. □

Combining Theorem 4.3 and Lemma 4.4 yields Theorem 1.2.

Alwen and Blocki [2] show  $\Pi_{cc}^\parallel(G) = O\left(\frac{n^2 \log \log n}{\log n}\right)$  for any graph  $G$  with constant indegree. Moreover, there exists a family of DAGs  $\{G_n\}_{n=1}^\infty$  with constant indegree with  $\Pi_{cc}(G_n) \in \Omega(n^2)$  [7, 31].

We now show a relationship similar to Theorem 1.2 between the energy cost and cumulative memory cost [8] of an execution trace, where the cumulative memory cost of an execution trace is defined as:

$$\text{cmc}(\text{Trace}_{\mathcal{A}, R, H}(x)) = \sum |\alpha_i|,$$

where  $\alpha_i$  encodes the state of the attacker at round  $i$ . Similarly,  $\text{cmc}_{q, \epsilon}(f_{G, H}) = \min_{\mathcal{A}, R, x} \text{cmc}(\text{Trace}_{\mathcal{A}, R, H}(x))$  where the minimum is taken over all  $\mathcal{A}$  making at most  $q$  random oracle queries that compute  $f_{G, H}$  correctly with probability at least  $\epsilon$ . While there is no notion of a cache in their pROM model, we could trivially set  $\alpha_i = (\sigma_i, \xi_i)$ . We note that for  $\text{ecost}_{q, \epsilon}(f_{G, H})$  minimum is taken over all  $\mathcal{A}$  making at most  $q$  random oracle queries that compute  $f_{G, H}$  correctly with probability at least  $\epsilon$  and having cache size at most  $mw$  bits, which means that the set of attackers we consider is even more restrictive. We emphasize that  $\mathcal{A}$  can be an arbitrary pROM algorithm, so that the following result also applies to dMHFs such as `script`.

THEOREM 4.5. For any execution trace  $\text{Trace}_{\mathcal{A}, R, H}(x)$  of an algorithm  $\mathcal{A}$  with cache size  $mw$  bits

$$\text{ecost}(\text{Trace}_{\mathcal{A}, R, H}(x)) \geq \left( \frac{\text{cmc}(\text{Trace}_{\mathcal{A}, R, H}(x))}{tw} - m \right) c_b + t c_r,$$

where  $m$  is the cache size,  $t$  is the number of steps,  $c_b$  is the cost of a blue move and  $c_r$  is the cost of a red move.

PROOF. Recall that the energy cost of an execution trace  $\text{Trace}_{\mathcal{A}, R, H}(x) = \{(\sigma_i, \xi_i, R_i, S_i, Q_i)\}_{i=1}^t$  is defined as

$$\begin{aligned} \text{ecost}(\text{Trace}_{\mathcal{A}, R, H}(x)) &= \sum_{i=1}^t \left( c_r |Q_i| + \frac{c_b}{w} (|R_i| + |S_i|) \right) \\ &\geq \max_i \frac{|\xi_i|}{w} c_b + t c_r \geq \left( \frac{\text{cmc}(\text{Trace}_{\mathcal{A}, R, H}(x))}{tw} - m \right) c_b + t c_r \end{aligned}$$

The second step above follows from the observation that for all  $j$  we have  $|\xi_j| \leq \sum_{i=1}^j (|R_i| + |S_i|)$ , and the third step follows from the observation that

$$\text{cmc}(\text{Trace}_{\mathcal{A}, R, H}(x)) - mtw = \sum_{i=1}^t (|\sigma_i| + |\xi_i|) - mtw \leq t \max_i |\xi_i|.$$

□

In particular, by minimizing over all  $t$  it follows that for any trace  $\text{Trace}_{\mathcal{A}, R, H}(x)$  we have

$$\text{ecost}(\text{Trace}_{\mathcal{A}, R, H}(x)) \in \Omega \left( \sqrt{\frac{\text{cmc}(\text{Trace}_{\mathcal{A}, R, H}(x)) \cdot c_b \cdot c_r}{w}} - mc_b \right)$$

Since Alwen *et al.* [6] show that  $\text{cmc}_{q, \epsilon}(\text{script}) \in \Omega(\epsilon n^2 \cdot w)$  for any  $q > 0$  and  $\epsilon > 2^{-w/2} + 2^{-n/20+1}$  it follows that

COROLLARY 4.6. For any  $q > 0$  and  $\epsilon > 2^{-w/2} + 2^{-n/20+1}$ ,

$$\text{ecost}_{q,\epsilon}(\text{script}) \in \Omega(n\sqrt{c_b \cdot c_r}).$$

## 5 BANDWIDTH HARDNESS OF CANDIDATE IMHFS

In this section, we provide lower bounds on the bandwidth hardness on the family of graphs generated by Argon2i [11], aATSample, and DRSample [4]. Given a DAG  $G = ([n], E)$ , a target set  $T \subseteq [n]$  and red/blue subsets  $B, R \subseteq [n]$  with  $|R| \leq m$  we let  $\text{rbpeb}^\parallel(G, m, T, B, R)$  denote the red-blue cost to place red pebbles on a target set  $T$  starting from an initial red-blue pebbling configuration  $B, R$ .

### 5.1 Analysis Framework

We follow a similar strategy for each candidate construction by defining a target set  $T_i = ((i-1)c\ell, ic\ell]$ , and analyzing the structure of the DAG to lower bound the following quantity for that DAG:

$$\min_{R, B' \subseteq [(i-1)c\ell, ic\ell]: |R| \leq m} (|B'|c_b + |\text{ancestors}_{G-R-B'}(T_i)|c_r)$$

We show in Theorem 5.2 that this quantity suffices to lower bound the bandwidth hardness.

LEMMA 5.1.  $\forall T, B, R \subseteq [n]$  such that  $|R| \leq m$  we have

$$\text{rbpeb}^\parallel(G, m, T, B, R) \geq \min_{B' \subseteq B} (|B'|c_b + |\text{ancestors}_{G-R-B'}(T)|c_r),$$

where  $c_b$  is the cost of a blue move and  $c_r$  is the cost of a red move.

PROOF. Let  $P = (B_0, R_0), (B_1, R_1) \dots, (B_t, R_t)$  denote a legal red-blue pebbling sequence given starting configuration  $B_0 = \{v \in B : \exists j \leq t. v \in R_j\}$  (e.g.,  $B_0$  is the subset of all blue pebbles in  $B$  that we will use at some point during the pebbling) and  $R_0 = R$ . By construction of  $B_0$  the pebbling contains at least  $|B_0|$  blue moves at cost  $|B_0|c_b$ . Similarly, we remark that we must place a red-pebble on all of the nodes in  $\text{ancestors}_{G-R-B'}(T)$  at some point. Thus, we have at least  $|\text{ancestors}_{G-R-B'}(T)|c_r$  red-moves. It follows that

$$\text{rbpeb}^\parallel(G, m, T, B, R) \geq \min_{B' \subseteq B} (|B'|c_b + |\text{ancestors}_{G-R-B'}(T)|c_r).$$

□

THEOREM 5.2. Let  $G = ([n], E)$  be any DAG such that  $(i, i+1) \in E$  for each  $i < n$ , let  $c$  be a positive integer and  $T_i = ((i-1)c\ell + 1, ic\ell]$ ,

$$\text{rbpeb}^\parallel(G, m) \geq \sum_{i=1}^{\lfloor \frac{n}{c\ell} \rfloor} \min_{R, B' \subseteq [(i-1)c\ell, ic\ell]: |R| \leq m} (|B'|c_b + |\text{ancestors}_{G-R-B'}(T_i)|c_r).$$

To prove Theorem 5.2, consider an optimal red-blue pebbling and let  $t_i$  denote the first time we place a pebble on node  $ic\ell$ . For each  $i$ , we use Lemma 5.1 to lower bound the red-blue cost incurred between steps  $t_{i-1} + 1$  and  $t_i$ . See Appendix A for more details.

As expected, if  $m = n$  then we have red-blue cost at most  $\text{rbpeb}^\parallel(G, m) \leq nc_r$  for any graph  $G$ . Thus, we require some upper bound on  $m$  to establish lower-bounds for red-blue pebbling cost.

### 5.2 Underlying DAGs

We now describe each of the underlying DAGs whose energy complexity we analyze.

The underlying graph for Argon2iB [12] has a directed path of length  $n$  nodes. Each node  $i$  has parents  $i-1$  and  $r(i) = i(1 - \frac{x^2}{N^2})$ , where  $N \gg n$  (in the implementation of  $N = 2^{32}$ ) and  $x$  is chosen uniformly at random from  $[N]$ . See Algorithm 3 in Appendix C for a more formal description.

While Argon2iA (v1.1) is an outdated version of the password hash function it is still worthwhile to study for several reasons. First, the uniform edge distribution is a natural one which has been adopted by other iMHF constructions [17]. Second, it is possible that this older version of Argon2i may have seen some adoption. Each node  $i$  in Argon2iA has two parents:  $i-1$  and  $r(i) = i(1 - \frac{x}{N})$ , where  $N = 2^{32}$  and  $x$  is chosen uniformly at random from  $[N]$ . Thus, the parents in Argon2iA are slightly less biased towards closer nodes than in Argon2iB. See Algorithm 4 in Appendix C for a more formal description.

DRSample is a family of graphs  $\mathbb{G}_n$  with  $\Pi_{cc}^\parallel(G) \in \Omega(\frac{n^2}{\log n})$  with high probability for any  $G \in \mathbb{G}_n$ . Like Argon2i and Argon2iB, the underlying graph for DRSample has a directed path of length  $n$  nodes. Each node  $i$  has parents  $i-1$  and  $r(i)$ , but the distribution for  $r(i)$  differs greatly from Argon2i and Argon2iB. Roughly speaking, DRSample samples an index  $j$  uniformly at random from  $[1, \log i]$ , an index  $k$  uniformly at random from  $[1, 2^j]$ , and sets  $r(i) = i - k$ . See Algorithm 1 in Appendix C for a more formal description.

A close relative to DRSample, aATSample [4] is also a family of graphs  $\mathbb{G}_n$  with  $\Pi_{cc}^\parallel(G) \in \Omega(\frac{n^2}{\log n})$  with high probability for any  $G \in \mathbb{G}_n$ . aATSample modifies DRSample by appending another directed path with  $n$  nodes that strategically connects to the first half of the graph so that the resulting computational complexity is high. See Algorithm 2 in Appendix C for a more formal description.

### 5.3 Argon2iB

Let  $G$  be a graph generated by Argon2iB and  $G_k$  be the metagraph with  $\frac{n}{k}$  nodes, so that each meta-node in  $G_k$  represents  $k$  nodes in  $G$ . Again, we connect two meta-nodes  $i < j$  in  $G_k$  if there exists a node in the last  $k/3$  nodes of  $i$  to a node in the first  $k/3$  nodes of  $j$ . Then  $G_k$  has the following property:

LEMMA 5.3. [15] For any two meta-nodes  $x < y$  of  $G_k$ , the last third of  $x$  is connected to the first third of  $y$  with probability at least  $\frac{k\sqrt{k}}{k\sqrt{k} + 36\sqrt{n(y-x+1)}}$ .

LEMMA 5.4. Suppose there exists a  $C > 0$  and  $0 < \epsilon < 2/3$  so that  $m = Cn^{\frac{2}{3}-\epsilon}$ . Let  $i > \frac{n}{2}$  and let  $T = [i, i + \ell - 1]$  be an interval of length  $\ell = O(n^{\frac{2}{3}-\epsilon})$ . Then a graph generated by Argon2iB satisfies the following with high probability:

$$\min_{R, B' \subseteq [i-1]: |R| \leq m} (|B'|c_b + |\text{ancestors}_{G-R-B'}(T)|c_r) \geq \min \left( \Omega(n^{\frac{1}{3}-\epsilon})c_b, \Omega(n)c_r \right).$$

PROOF. Consider the metagraph  $G_k$  with  $\frac{n}{k}$  nodes for some constant  $k$  that we shall define. Let  $B$  be the set of nodes in  $G$  that have

a blue pebble at some point, and let  $B_m$  be the set of meta-nodes that contain some node in  $B$ . Partition the second half of graph  $G$  into intervals of size  $\ell = \frac{n}{k} \cdot [i, i + \ell - 1]$ . By Lemma 5.3, there exists a constant  $\alpha$  such that for  $k = \alpha n^{\frac{2}{3} + \epsilon}$ , the meta-nodes in  $[i, i + \ell - 1]$  are connected to  $\Omega\left(\frac{n}{k}\right)$  meta-nodes in  $G_k$  with high probability. Thus, pebbling the interval  $[i, i + \ell - 1]$  requires pebbling at least  $\beta n/k - m - |B_k|$  meta-nodes in  $G_k$  for some constant  $\beta$ . Noting that  $m = O\left(n^{\frac{2}{3} - \epsilon}\right)$  and that the middle  $k/3$  nodes of a meta-node must be pebbled for two meta-nodes that are connected in  $G_k$ , it follows that at least  $\left(\frac{c_1 n}{k} - |B_k|\right) \frac{k}{3} = \Omega(n)$  nodes must be pebbled. Thus, the cost of pebbling  $[i, i + \ell - 1]$  is at least  $\min(\Omega(n/k)c_b, \Omega(n)c_r)$ .  $\square$

**Reminder of Theorem 1.5.** *Let  $G$  be a graph generated by Argon2iB. Then there exists a constant  $C > 0$  so that for any  $0 < \epsilon < 2/3$  and for all  $m \leq Cn^{\frac{2}{3} - \epsilon}$ , with high probability,*

$$\text{rbpeb}^{\parallel}(G, m) \geq \min(\Omega(n)c_b, \Omega(n^{5/3})c_r).$$

**Proof of Theorem 1.5:** Applying Lemma 5.4 to each of the disjoint  $\frac{n}{\ell} = k = O\left(n^{\frac{2}{3} + \epsilon}\right)$  intervals in the second half of graph  $G$ , the theorem follows from Theorem 5.2.  $\square$

## 5.4 DRSample

A DRSample graph can be viewed as a metagraph  $G_b$  with  $\Omega(n/b)$  meta-nodes  $M_i$ . Recall that we connect two meta-nodes  $i < j$  in  $G_b$  if there exists a node in the last  $b/3$  nodes of  $i$  to a node in the first  $b/3$  nodes of  $j$ .

**Lemma 5.5.** *Suppose  $m = O(n^\rho)$  for some constant  $0 < \rho < 1$  and  $i > \frac{n}{2}$ . Let  $T = [i, i + \ell - 1]$  be an interval of length  $\ell = O(n^\rho)$ . Then a graph generated by DRSample satisfies the following with high probability:*

$$\min_{R \subseteq [i-1]: |R| \leq m} \min_{B' \subseteq [i-1]} (|B'|c_b + |\text{ancestors}_{G-R-B'}(T)|c_r) \geq \min\left(\Omega(n^\rho)c_b, \Omega(n^{1/2+\rho/2})c_r\right).$$

Using Lemma 5.5, whose proof appears in Appendix A.2, we have:

**Reminder of Theorem 1.4.** *Let  $G$  be a graph generated by DRSample and  $0 < \rho < 1$ . Then there exists a constant  $C > 0$  so that for all  $m \leq Cn^\rho$ , with high probability,*

$$\text{rbpeb}^{\parallel}(G, m) \geq \min\left(\Omega(n)c_b, \Omega(n^{3/2-\rho/2})c_r\right).$$

**Proof of Theorem 1.4:** Applying Lemma 5.5 to each of the disjoint  $\frac{n}{\ell}$  intervals in the second half of graph  $G$  and observing that  $\ell = O(n^\rho)$ , it follows from Theorem 5.2 that

$$\text{rbpeb}^{\parallel}(G, m) \geq \min\left(\Omega(n)c_b, \Omega(n^{3/2-\rho/2})c_r\right).$$

$\square$

We also give a stronger bound for DRSample when the cache has size  $O(n^\rho/\log n)$  for any  $0 < \rho < 1$  in Appendix A.2.

## 5.5 aATSample

Alwen *et al.* showed the first half of the nodes of aATSample is a  $(e, d, b)$ -block depth robust graph with  $e = \Omega(n/\log n)$ ,  $d = \Omega(n)$  and  $b = \Omega(\log n)$  [4]. This graph can be viewed as a metagraph  $G_k$  with  $\Omega(n/b)$  meta-nodes  $M_i$ . Recall we connect two meta-nodes  $i < j$  in  $G_k$  if there exists a node in the last  $k/3$  nodes of  $i$  connected to a node in the first  $k/3$  nodes of  $j$ . The second half of the nodes of aATSample is a chain of  $n/2$  nodes so that  $\text{parents}([v, v + \ell - 1])$  contains some node in the last  $k/3$  nodes of  $M_i$  for each meta-node  $M_i$  and each interval  $[v, v + \ell - 1]$  of length  $\ell = \Omega(n/b)$  in the second half of the graph. To move a pebble from node  $v$  to node  $v + \ell - 1$ , either the starting configuration must have at least  $e/2$  pebbles on the graph, or a significant fraction of the block depth robust graph in the first half of the graph must be re-pebbled.

**Lemma 5.6.** *Let  $i > \frac{n}{2}$  and  $T = [i, i + \ell - 1]$  be an interval of length  $\ell = \frac{n}{\log n}$ . Then there exists a constant  $C > 0$  so that for  $m \leq \frac{Cn}{\log n}$ , a graph generated by aATSample satisfies the following with high probability:*

$$\min_{R, B' \subseteq [i-1]: |R| \leq m} (|B'|c_b + |\text{ancestors}_{G-R-B'}(T)|c_r) \geq \min\left(\Omega\left(\frac{n}{\log n}\right)c_b, \Omega(n)c_r\right).$$

We now use Lemma 5.6, whose proof appears in Appendix A.1.

**Reminder of Theorem 1.3.** *Let  $G$  be a graph generated by aATSample.*

*Then there exists a constant  $C > 0$  so that for all  $m \leq \frac{Cn}{\log n}$ ,*

$$\text{rbpeb}^{\parallel}(G, m) \geq \min(\Omega(n)c_b, \Omega(n \log n)c_r),$$

*holds with high probability.*

**Proof of Theorem 1.3:** Applying Lemma 5.6 to each of the disjoint  $\log n$  intervals in the second half of graph  $G$ , the theorem follows from Theorem 5.2.  $\square$

## ACKNOWLEDGEMENTS

The authors would like to thank Daniel Wichs for helpful discussion and anonymous reviewers for important comments that improved the presentation of the paper.

## REFERENCES

- [1] Martin Abadi, Michael Burrows, Mark S. Manasse, and Ted Wobber. 2005. Moderately hard, memory-bound functions. *ACM Trans. Internet Techn.* 5, 2 (2005), 299–327.
- [2] Joël Alwen and Jeremiah Blocki. 2016. Efficiently Computing Data-Independent Memory-Hard Functions. In *CRYPTO 2016, Part II (LNCS)*, Matthew Robshaw and Jonathan Katz (Eds.), Vol. 9815. Springer, Heidelberg, 241–271. [https://doi.org/10.1007/978-3-662-53008-5\\_9](https://doi.org/10.1007/978-3-662-53008-5_9)
- [3] Joël Alwen and Jeremiah Blocki. 2017. Towards practical attacks on argon2i and balloon hashing. In *Security and Privacy (EuroS&P), 2017 IEEE European Symposium on*. IEEE, 142–157.
- [4] Joël Alwen, Jeremiah Blocki, and Ben Harsha. 2017. Practical Graphs for Optimal Side-Channel Resistant Memory-Hard Functions. In *ACM CCS 17*, Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu (Eds.). ACM Press, 1001–1017.
- [5] Joël Alwen, Jeremiah Blocki, and Krzysztof Pietrzak. 2017. Depth-Robust Graphs and Their Cumulative Memory Complexity. In *EUROCRYPT 2017, Part III (LNCS)*, Jean-Sébastien Coron and Jesper Buus Nielsen (Eds.), Vol. 10212. Springer, Heidelberg, 3–32.
- [6] Joël Alwen, Binyi Chen, Krzysztof Pietrzak, Leonid Reyzin, and Stefano Tessaro. 2017. Scrypt Is Maximally Memory-Hard. In *EUROCRYPT 2017, Part III (LNCS)*, Jean-Sébastien Coron and Jesper Buus Nielsen (Eds.), Vol. 10212. Springer, Heidelberg, 33–62.

- [7] Joël Alwen, Susanna F. de Rezende, Jakob Nordström, and Marc Vinyals. 2016. Cumulative Space in Black-White Pebbling and Resolution. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science, 9-11 January 2017, Berkeley, California USA*.
- [8] Joël Alwen and Vladimir Serbinenko. 2015. High Parallel Complexity Graphs and Memory-Hard Functions. In *47th ACM STOC*, Rocco A. Servedio and Ronitt Rubinfeld (Eds.). ACM Press, 595–603.
- [9] Adam Back. 2002. Hashcash-a denial of service counter-measure.
- [10] Daniel J. Bernstein. 2005. Cache-Timing Attacks on AES.
- [11] Alex Biryukov, Daniel Dinu, and Dmitry Khovratovich. 2015. Fast and Tradeoff-Resilient Memory-Hard Functions for Cryptocurrencies and Password Hashing. *IACR Cryptology ePrint Archive* 2015 (2015), 430. <http://eprint.iacr.org/2015/430>
- [12] Alex Biryukov, Daniel Dinu, and Dmitry Khovratovich. 2016. Argon2: New Generation of Memory-Hard Functions for Password Hashing and Other Applications. In *IEEE European Symposium on Security and Privacy, EuroS&P 2016, Saarbrücken, Germany, March 21-24, 2016*. 292–302. <https://doi.org/10.1109/EuroSP.2016.31>
- [13] Alex Biryukov, Daniel Dinu, Dmitry Khovratovich, and Simon Josefsson. 2016. The memory-hard Argon2 password hash and proof-of-work function.
- [14] Jeremiah Blocki, Ben Harsha, and Samson Zhou. 2018. On the Economics of Offline Password Cracking. In *IEEE Symposium on Security and Privacy, SP*. 35–53.
- [15] Jeremiah Blocki and Samson Zhou. 2017. On the Depth-Robustness and Cumulative Pebbling Cost of Argon2i. In *TCC 2017, Part I (LNCS)*, Yael Kalai and Leonid Reyzin (Eds.), Vol. 10677. Springer, Heidelberg, 445–465.
- [16] Jeremiah Blocki and Samson Zhou. 2018. On the Computational Complexity of Minimal Cumulative Cost Graph Pebbling. *Financial Cryptography and Data Security, 22nd International Conference, FC (2018)*. (to appear).
- [17] Dan Boneh, Henry Corrigan-Gibbs, and Stuart E. Schechter. 2016. Balloon Hashing: A Memory-Hard Function Providing Provable Protection Against Sequential Attacks. In *ASIACRYPT 2016, Part I (LNCS)*, Jung Hee Cheon and Tsuyoshi Takagi (Eds.), Vol. 10031. Springer, Heidelberg, 220–248. [https://doi.org/10.1007/978-3-662-53887-6\\_8](https://doi.org/10.1007/978-3-662-53887-6_8)
- [18] Erik D. Demaine and Quanquan C. Liu. 2017. Inapproximability of the Standard Pebble Game and Hard to Pebble Graphs. In *Algorithms and Data Structures - 15th International Symposium, WADS 2017, St. John's, NL, Canada, July 31 - August 2, 2017, Proceedings*. 313–324.
- [19] Cynthia Dwork and Moni Naor. 1992. Pricing via Processing or Combatting Junk Mail. In *Advances in Cryptology - CRYPTO, 12th Annual International Cryptology Conference, Proceedings*. 139–147.
- [20] Stefan Dziembowski, Tomasz Kazana, and Daniel Wichs. 2011. Key-Evolution Schemes Resilient to Space-Bounded Leakage. In *Advances in Cryptology - CRYPTO - 31st Annual Cryptology Conference, Proceedings*. 335–353.
- [21] Stefan Dziembowski, Tomasz Kazana, and Daniel Wichs. 2011. One-Time Computable Self-erasing Functions. In *Theory of Cryptography - 8th Theory of Cryptography Conference, TCC Proceedings*. 125–143.
- [22] Christian Forler, Stefan Lucks, and Jakob Wenzel. 2013. Catena: A Memory-Consuming Password Scrambler. *Cryptology ePrint Archive*, Report 2013/525. <http://eprint.iacr.org/2013/525>.
- [23] John R. Gilbert, Thomas Lengauer, and Robert Endre Tarjan. 1979. The Pebbling Problem is Complete in Polynomial Space. In *Proceedings of the 11th Annual ACM Symposium on Theory of Computing (STOC)*. 237–248.
- [24] Jia-Wei Hong and H. T. Kung. 1981. I/O Complexity: The Red-Blue Pebble Game. In *Proceedings of the 13th Annual ACM Symposium on Theory of Computing, May 11-13, 1981, Milwaukee, Wisconsin, USA*. 326–333.
- [25] Thomas Lengauer and Robert E. Tarjan. 1982. Asymptotically Tight Bounds on Time-space Trade-offs in a Pebble Game. *J. ACM* 29, 4 (Oct. 1982), 1087–1130. <https://doi.org/10.1145/322344.322354>
- [26] Quanquan Liu. 2017. *Red-Blue and Standard Pebble Games: Complexity and Applications in the Sequential and Parallel Models*. Master's thesis. Massachusetts Institute of Technology. <http://erikdemaine.org/theses/qliuM.pdf>
- [27] Satoshi Nakamoto. 2008. Bitcoin: A peer-to-peer electronic cash system.
- [28] Colin Percival. 2009. Stronger key derivation via sequential memory-hard functions. *BSDCan* (2009).
- [29] PHC 2013–2015. Password Hashing Competition. <https://password-hashing.net/>
- [30] Ling Ren and Srinivas Devadas. 2017. Bandwidth Hard Functions for ASIC Resistance. In *TCC 2017, Part I (LNCS)*, Yael Kalai and Leonid Reyzin (Eds.), Vol. 10677. Springer, Heidelberg, 466–492.
- [31] Georg Schnitger. 1983. On Depth-Reduction and Grates. In *24th Annual Symposium on Foundations of Computer Science*. 323–328.

## A MISSING PROOFS

**Reminder of Theorem 5.2.** Let  $G = ([n], E)$  be any DAG such

that  $(i, i + 1) \in E$  for each  $i < n$ , let  $c$  be a positive integer and

$$T_i = ((i - 1)c\ell + 1, ic\ell],$$

$$\text{rbpeb}^{\parallel}(G, m) \geq \sum_{i=1}^{\lfloor \frac{n}{c\ell} \rfloor} \min_{R, B' \subseteq [(i-1)c\ell]: |R| \leq m} (|B'|c_b + |\text{ancestors}_{G-R-B'}(T_i)|c_r).$$

**Proof of Theorem 5.2:** (Sketch) Repeatedly invoke the previous lemma. Consider an optimal red-blue pebbling and let  $t_i$  denote the first time we place a pebble on node  $ic\ell$ . For each  $i$  the red-blue cost incurred between steps  $t_{i-1} + 1$  and  $t_i$  starting from some red-blue configuration  $B_{t_{i-1}}, R_{t_{i-1}}$  is at least

$$\begin{aligned} \text{rbpeb}^{\parallel}(G, m, T_i, B_{t_{i-1}}, R_{t_{i-1}}) &\geq \min_{B' \subseteq [(i-1)c\ell]} (|B'|c_b + |\text{ancestors}_{G-R-B'}(T_i)|c_r) \\ &\geq \min_{R, B' \subseteq [(i-1)c\ell]: |R| \leq m} (|B'|c_b + |\text{ancestors}_{G-R-B'}(T_i)|c_r). \end{aligned}$$

To complete the proof we observe that

$$\text{rbpeb}^{\parallel}(G, m) \geq \sum_{i=1}^{\lfloor \frac{n}{c\ell} \rfloor} \text{rbpeb}^{\parallel}(G, m, T_i, B_{t_{i-1}}, R_{t_{i-1}}).$$

□

### A.1 aATSample

**Reminder of Lemma 5.6.** Let  $i > \frac{n}{2}$  and  $T = [i, i + \ell - 1]$  be an interval of length  $\ell = \frac{n}{\log n}$ . Then there exists a constant  $C > 0$  so that for  $m \leq \frac{Cn}{\log n}$ , a graph generated by aATSample satisfies the following with high probability:

$$\min_{R, B' \subseteq [i-1]: |R| \leq m} (|B'|c_b + |\text{ancestors}_{G-R-B'}(T)|c_r) \geq \min \left( \Omega \left( \frac{n}{\log n} \right) c_b, \Omega(n)c_r \right).$$

**Proof of Lemma 5.6:** Consider the metagraph  $G_k$  with  $\frac{n}{\log n}$  nodes and recall that  $G_k$  is  $(\Omega(n/\log n), \Omega(n/\log n))$  depth robust. Let  $B$  be the set of nodes in  $G$  that have a blue pebble at some point, and let  $B_k$  be the set of meta-nodes that contain some node in  $B$ . Then by depth-robustness, there exists a path of length at least  $c_1 n / \log n - |B_k|$  in  $G_k$ .

Now, partition the second half of graph  $G$  into intervals of size  $\ell = n / \log n: [i, i + \ell - 1]$ . Since  $\text{parents}([i, i + \ell - 1]) \cap M_i \neq \emptyset$ , then pebbling the interval  $[i, i + \ell - 1]$  requires pebbling at least  $c_1 n / \log n - m - |B_k|$  meta-nodes in  $G_k$ . Noting that  $m = O\left(\frac{n}{\log n}\right)$  and the middle  $\log n / 3$  nodes of a meta-node must be pebbled for two meta-nodes that are connected in  $G_k$ , it follows that at least  $\left(\frac{c_1 n}{\log n} - |B_k|\right) \frac{\log n}{3} = \Omega(n)$  nodes must be pebbled. Thus, the cost of pebbling  $[i, i + \ell - 1]$  is at least  $\min \left( \Omega \left( \frac{n}{\log n} \right) c_b, \Omega(n)c_r \right)$  nodes. □

## A.2 DRSample

**Reminder of Lemma 5.5.** Suppose  $m = O(n^\rho)$  for some constant  $0 < \rho < 1$  and  $i > \frac{n}{2}$ . Let  $T = [i, i + \ell - 1]$  be an interval of length  $\ell = O(n^\rho)$ . Then a graph generated by DRSample satisfies the following with high probability:

$$\min_{R \subseteq [i-1]: |R| \leq m} \min_{B' \subseteq [i-1]} (|B'| c_b + |\text{ancestors}_{G-R-B'}(T)| c_r) \geq \min \left( \Omega(n^\rho) c_b, \Omega(n^{1/2+\rho/2}) c_r \right).$$

**Proof of Lemma 5.5:** Let  $T = [i, i + \ell]$  where  $\ell = \Omega(n^\rho)$  for some constant  $\frac{1}{2} < \rho < 1$  and let  $r(j)$  denote the predecessor of a node  $j$  in the graph (besides  $j - 1$ ). Let  $X_j$  be an indicator random variable for the event  $\text{far}(j)$ , which we define to be the event that  $|r(j) - r(k)| > b = \sqrt{\frac{n}{16\ell}}$  for all  $k \in [1, j - 1]$ . Observe that

$$\Pr[\text{far}(j)] \geq$$

$$\Pr[\text{far}(j) | r(i) = i - b, r(i + 1) = i - 2b, \dots, r(j - 1) = i - (j - i + 1)b]$$

since the conditioned event has the maximum number of invalid nodes for  $r(j)$ , and with the highest probability of hitting each of these invalid nodes, since they are the closest to  $i$ . Thus,

$$\begin{aligned} \Pr[\text{far}(j)] &\geq \Pr[r(j) < i - (j - i + 1)b] \\ &\geq \Pr[j - r(j) > \ell + (j - i + 1)b] \\ &\geq \Pr[j - r(j) > \ell + (\ell + 1)b] \\ &\geq \Pr\left[j - r(j) > \sqrt{\frac{n\ell}{2}}\right] \end{aligned}$$

since  $j \leq i + \ell$  and  $b = \sqrt{\frac{n}{16\ell}}$ . Hence,

$$\begin{aligned} \Pr[\text{far}(j)] &\geq \frac{\log(j) - \log \sqrt{n\ell}}{\log(j)} \\ &\geq 1 - \left(\frac{1}{2} - \frac{\rho}{2}\right) \left(\frac{\log(n)}{\log(n) - 1}\right) \\ &\geq \frac{1}{2} - \frac{\rho}{2} - o(1) = \Omega(1). \end{aligned}$$

Let  $X = \sum_{k=i}^{i+\ell} X_k$ . With high probability,  $X > c\ell$  for some constant  $c$  that depends on  $\rho$ . Picking  $\ell$  to satisfy  $m < \frac{c\ell}{4}$ , then with high probability, the number of ancestors of  $T$  in  $G - R - B'$  is at most

$$\begin{aligned} (X - |R| - |B'|)b &\geq (X - m - |B'|)b \\ &\geq \left(\frac{c\ell}{4}\right) \sqrt{\frac{n}{16\ell}}. \end{aligned}$$

Thus to pebble  $T$ , either  $\frac{c\ell}{2}$  blue moves are required, or at least  $\left(\frac{c\ell}{4}\right) \sqrt{\frac{n}{16\ell}}$  red moves are required, to pebble the ancestors of  $T$  in  $G - R - B'$ . Hence, the cost is at least  $\min\left(\frac{c\ell}{2} c_b, \left(\frac{c\ell}{4}\right) \sqrt{\frac{n}{16\ell}} c_r\right)$ . Since  $\ell = \Omega(n^\rho)$ , then the cost is at least

$$\min \left( \Omega(n^\rho) c_b, \Omega(n^{1/2+\rho/2}) c_r \right).$$

□

We now give a stronger bound for DRSample when the cache has size  $O(n^\rho / \log n)$  for any  $0 < \rho < 1$ .

**LEMMA A.1.** For each  $x, y \in G_b$  with  $y > x$  and node  $i$  in meta-node  $y$ , there exists an edge from the last third of meta-node  $x$  to node  $i$  with probability at least  $\frac{1}{6|y-x|\log y}$ .

**PROOF.** Recall that for node  $i$ , DRSample creates an edge from  $i$  to parent node  $r(i)$  by first sampling  $j$  from  $[1, \lfloor \log i \rfloor]$ . Then DRSample sets  $r(i) = i - k$  by randomly sampling  $k$  from  $(2^{j-1}, 2^j]$ . Thus, for nodes  $i, j \in G$  with  $i > j$ , there exists an edge from node  $j$  to  $i$  with probability at least  $\frac{1}{2b|y-x|\log i}$ . Taking the union bound over all  $\frac{b}{3}$  nodes in the last third of meta-node  $x$  and observing that  $i < y$  yields the desired result. □

**LEMMA A.2.** For any two meta-nodes  $x, y \in G_b$  with  $x < y$ , the last third of  $x$  is connected to the first third of  $y$  with probability at least  $\frac{b}{b+18|y-x|\log n}$ .

**PROOF.** Let  $p$  be the probability that the final third of  $x$  is connected to the first third of  $y$ . Let  $E_i$  be the event that the  $i^{\text{th}}$  node of meta-node  $y$  is the first node in  $y$  to which there exists an edge from the last third of meta-node  $x$ , so that by Lemma A.1,  $\Pr[E_i] \geq \frac{1}{6|y-x|\log y}$ . Note that furthermore,  $\Pr[E_i]$  is the probability that there exists an edge from the last third of meta-node  $x$  to the  $i^{\text{th}}$  node of meta-node  $y$  and no previous meta-node of  $y$ . Hence,  $\Pr[E_i] \geq \frac{1}{6|y-x|\log y} (1 - p)$ . Thus,

$$\begin{aligned} p &= \Pr[E_1] + \Pr[E_2] + \dots + \Pr[E_{b/3}] \\ &\geq \left(\frac{b}{3}\right) \frac{1}{6|y-x|\log y} (1 - p). \end{aligned}$$

Setting  $\alpha = \left(\frac{b}{3}\right) \frac{1}{6|y-x|\log y}$ , then it follows that  $p + \alpha p \geq \alpha$ , so that  $p \geq \frac{\alpha}{1+\alpha}$ . Since  $y \leq \frac{n}{b}$ ,

$$p \geq \frac{b}{b + 18|y-x|\log y} \geq \frac{b}{b + 18|y-x|\log n}.$$

□

**LEMMA A.3.** Suppose  $m = O(n^\rho / \log n)$  and  $i > \frac{n}{2}$ . Let  $T = [i, i + \ell - 1]$  be an interval of length  $\ell = O(n^\rho)$  for some  $0 < \rho < 1$ . Then a graph generated by DRSample satisfies the following with high probability:

$$\min_{R \subseteq [i-1]: |R| \leq m} \min_{B' \subseteq [i-1]} (|B'| c_b + |\text{ancestors}_{G-R-B'}(T)| c_r) \geq \min(\tilde{\Omega}(n^\rho) c_b, \tilde{\Omega}(n) c_r).$$

**PROOF.** Let  $T$  be an interval of length  $\ell$ . Consider a metagraph  $G_{n/\ell}$  with  $\ell$  nodes. Then by Lemma A.2, a meta-node  $y$  in  $G_\ell$  is connected to a previous meta-node  $x$  with probability  $\frac{\ell}{\ell+18|y-x|\log n} = \Omega(1)$  for  $|y-x| = O\left(\frac{n^\rho}{\log n}\right)$ , since  $\ell = \Omega(n^\rho)$ . Thus in expectation,  $y$  is connected to  $\frac{Cn^\rho}{\log n}$  previous meta-nodes for some constant  $y$ .

Now, partition the second half of graph  $G$  into intervals of size  $n/\ell$ . For each interval of length  $n/\ell$ , to pebble a meta-node, either there already exists a blue pebble in the previous meta-node or we must spend  $\frac{2n}{3\ell}$  red moves to repebble the previous meta-node. Applying the same argument for  $\frac{Cn^\rho}{\log n} - m$  of the previous meta-nodes, and noting that  $\ell = O(n^\rho)$  and  $m = O(n^\rho / \log n)$ , then



the cost of pebbling  $[i, i + \ell - 1]$  is at least  $\tilde{\Omega}(n^\rho) \min\left(c_b, \frac{2n}{3\ell} c_r\right) = \min(\tilde{\Omega}(n^\rho) c_b, \tilde{\Omega}(n) c_r)$ .  $\square$

**THEOREM A.4.** *Let  $G$  be a graph generated by DRSample and  $0 < \rho < 1$ . Then there exists a constant  $C > 0$  so that for all  $m \leq Cn^\rho / \log n$ , it follows that*

$$\text{rbpeb}^\parallel(G, m) \geq \min(\tilde{\Omega}(n) c_b, \tilde{\Omega}(n^{2-\rho}) c_r)$$

with high probability.

**PROOF.** Applying Lemma A.3 to each of the disjoint  $\frac{n}{\ell}$  intervals in the second half of graph  $G$  and observing that  $\ell = O(n^\rho)$ , it follows from Theorem 5.2 that

$$\text{rbpeb}^\parallel(G, m) \geq \min(\tilde{\Omega}(n) c_b, \tilde{\Omega}(n^{2-\rho}) c_r).$$

$\square$

### A.3 Argon2iA

We now consider the family of graphs generated by Argon2iA (Algorithm 4) [11]. Notably, the same underlying graph is also used in Balloon Hashing (Boneh et al. [17]) Let  $G$  be a graph generated by Argon2i and  $G_b$  be the metagraph with  $\frac{n}{b}$  nodes, so that each meta-node in  $G_b$  represents  $b$  nodes in  $G$ . Again, we connect two meta-nodes  $i < j$  in  $G_b$  if there exists a node in the last  $b/3$  nodes of  $i$  to a node in the first  $b/3$  nodes of  $j$ .

**LEMMA A.5.** *For each  $x, y \in G_b$  with  $y > x$  and node  $i$  in meta-node  $y$ , there exists an edge from the last third of meta-node  $x$  to node  $i$  with probability at least  $\frac{1}{6y}$ .*

**PROOF.** Recall that for node  $i$ , Argon2i creates an edge from  $i$  to parent node  $i \left(1 - \frac{k}{N}\right)$ , where  $k \in [N]$  is picked uniformly at random. Thus, for nodes  $i, j \in G$  with  $i > j$ , there exists an edge from node  $j$  to  $i$  with probability at least

$$\begin{aligned} & \Pr\left[(x-1)m+1 \leq i \left(1 - \frac{k}{N}\right) \leq \left(x-1 + \frac{1}{3}\right)m\right] \\ &= \Pr\left[\left(x-1 + \frac{1}{6}\right)m \leq ym \left(1 - \frac{k}{N}\right) \leq \left(x-1 + \frac{1}{3}\right)m\right] \\ &\geq \Pr\left[\frac{y-x+\frac{5}{6}}{y} \geq \frac{k}{N} \geq \frac{y-x+\frac{2}{3}}{y}\right] \\ &\geq \left(\frac{y-x+\frac{5}{6}}{y}\right) - \left(\frac{y-x+\frac{2}{3}}{y}\right) \\ &\geq \frac{1}{6y} \end{aligned}$$

$\square$

**LEMMA A.6.** *For any two meta-nodes  $x, y \in G_b$  with  $x < y$ , the last third of  $x$  is connected to the first third of  $y$  with probability at least  $\frac{b^2}{9y \log y + b^2}$ .*

**PROOF.** Let  $p$  be the probability that the final third of  $x$  is connected to the first third of  $y$ . Let  $E_i$  be the event that the  $i^{\text{th}}$  node of meta-node  $y$  is the first node in  $y$  to which there exists an edge from the last third of meta-node  $x$ , so that by Lemma A.5,  $\Pr[E_1] \geq \frac{1}{6y}$ . Note that furthermore,  $\Pr[E_i]$  is the probability that there exists an

edge from the last third of meta-node  $x$  to the  $i^{\text{th}}$  node of meta-node  $y$  and no previous meta-node of  $y$ . Hence,  $\Pr[E_i] \geq \frac{1}{6y}(1-p)$ . Thus,

$$\begin{aligned} p &= \Pr[E_1] + \Pr[E_2] + \dots + \Pr[E_{b/3}] \\ &\geq \left(\frac{b}{3}\right) \frac{1}{6y}(1-p). \end{aligned}$$

Setting  $\alpha = \left(\frac{b}{3}\right) \frac{1}{6y}$ , then it follows that  $p + \alpha p \geq \alpha$ , so that  $p \geq \frac{\alpha}{1+\alpha}$ . Since  $y \leq \frac{n}{b}$ , then  $p \geq \frac{b}{18y+b} \geq \frac{b^2}{18n+b^2}$ .  $\square$

**LEMMA A.7.** *Let  $i > \frac{n}{2}$  and  $T = [i, i + \ell - 1]$  be an interval of length  $\ell = O(n^{1/2})$ . There exists a constant  $C > 0$  so that for  $0 < \epsilon < 1/2$  and  $m \leq Cn^{1/2}$ , a graph generated by Argon2i satisfies the following with high probability:*

$$\min_{R \subseteq [i-1]: |R| \leq m} \min_{B' \subseteq [i-1]} (|B'| c_b + |\text{ancestors}_{G-R-B'}(T)| c_r) \geq \min\left(\Omega(n^{1/2}) c_b, \Omega(n) c_r\right).$$

**PROOF.** Consider the metagraph  $G_k$  with  $\frac{n}{k}$  nodes for some constant  $k$  that we shall define. Let  $B$  be the set of nodes in  $G$  that have a blue pebble at some point, and let  $B_m$  be the set of meta-nodes that contain some node in  $B$ . Partition the second half of graph  $G$  into intervals of size  $\ell = \frac{n}{k}: [i, i + \ell - 1]$ . By Lemma A.6, there exists a constant  $\alpha$  such that for  $k = \alpha n^{1/2}$ , the meta-nodes in  $[i, i + \ell - 1]$  are connected to  $\Omega\left(\frac{n}{k}\right)$  meta-nodes in  $G_k$  with high probability. Thus, pebbling the interval  $[i, i + \ell - 1]$  requires pebbling at least  $\beta n/k - m - |B_k|$  meta-nodes in  $G_k$  for some constant  $\beta$ . Noting that  $m = O(n/k)$  and that the middle  $k/3$  nodes of a meta-node must be pebbled for two meta-nodes that are connected in  $G_k$ , it follows that at least  $\left(\frac{c_1 n}{k} - |B_k|\right) \frac{k}{3} = \Omega(n)$  nodes must be pebbled. Thus, the cost of pebbling  $[i, i + \ell - 1]$  is at least  $\min(\Omega(n/k) c_b, \Omega(n) c_r)$  nodes.  $\square$

**THEOREM A.8.** *Let  $G$  be a graph generated by Argon2i. Then there exists a constant  $C > 0$  so that for all  $m \leq Cn^{1/2}$ , it follows that*

$$\text{rbpeb}^\parallel(G, m) \geq \min(\Omega(n) c_b, \Omega(n^{3/2}) c_r),$$

with high probability.

**Proof of Theorem A.8:** Applying Lemma A.7 to each of the disjoint  $\frac{n}{\ell} = k = O(n^{1/2})$  intervals in the second half of graph  $G$ , it follows from Theorem 5.2 that

$$\text{rbpeb}^\parallel(G, m) \geq \min(\Omega(n) c_b, \Omega(n^{1/2}) c_r).$$

$\square$

## B BACKGROUND ON THE GILBERT ET AL. BLACK PEBBLING REDUCTION

Recall that Gilbert *et al.* [23] showed that the minimum space black pebbling problem was PSPACE – Hard by reduction from the Truly Quantified Boolean Formula (TQBF) problem. Gilbert *et al.* [23] provide a construction from any instance of TQBF to a DAG  $G_{TQBF}$  with pebbling number  $3n + 3$  if and only if the instance is satisfiable. Here, the pebbling number of a DAG  $G$  is  $\min_{P=(P_1, \dots, P_t) \in \mathcal{P}} \max_{i \leq t} |P_i|$  is the number of pebbles necessary

to pebble  $G$ . An important gadget in their construction is the so-called pyramid DAG. We use both  $\Delta_k$  and a triangle with the number  $k$  inside to denote a  $k$ -pyramid (see Figure 2 for an example of a 3-pyramid). The key property of these DAGs is that *any* legal pebbling  $P = (P_0, \dots, P_t) \in \mathcal{P}^\parallel(\Delta_k)$  of a  $k$ -pyramid requires at least  $\min_i |P_i| \geq k$  pebbles on the DAG at some point in time.



Fig. 2: A 3-Pyramid.

Another gadget, which appears in Figure 3, is the existential quantifier gadget, which requires that  $s_i$ ,  $s_i - 1$ , and  $s_i - 2$  pebbles must be placed in each of the pyramids to ultimately pebble  $q_i$ .




Fig. 3: An existential quantifier, with  $x_i$  set to true in the left figure and  $x_i$  set to false in the right figure.

Any instance of TQBF in which each quantifier is an existential quantifier requires at most a quadratic number of pebbling moves. Specifically, we look at instances of 3-SAT, such as in Figure 4. In such a graph representing an instance of 3-SAT, the sink node to be pebbled is  $q_n$ . By design of the construction, any true statement requires exactly three pebbles for each pyramid representing a clause. On the other hand, a false clause requires four pebbles, so that false statements require more pebbles. Thus, by providing extraneous additions to the construction which force the number of pebbling moves to be a known constant, we can extract the pebbling number, given the space-time complexity. For more details, see the full description in [23].

LEMMA B.1. [23] *The quantified Boolean formula*

$$Q_1 x_1 Q_2 x_2 \cdots Q_n x_n F_n$$

is true if and only if the corresponding DAG  $G_{TQBF}$  has pebbling number  $3n + 3$ .

LEMMA B.2. [16] *Suppose that we have a satisfiable TQBF formula  $Q_1 x_1 Q_2 x_2 \cdots Q_n x_n F_n$  with  $Q_i = \exists$  for all  $i \leq n$ . Then there is a*

*legal sequential pebbling  $P = (P_0, \dots, P_t) \in \mathcal{P}(G_{TQBF})$  of the corresponding DAG  $G_{TQBF}$  from [23] with  $t \leq 6n^2 + 33n$  pebbling moves and  $\max_{i \leq t} |P_i| \leq 3n + 3$ .*

PROOF. We repeat the proof of [16] for completion. We describe the pebbling strategy of Gilbert *et al.* [23], and analyze the pebbling time of their strategy. Let  $T(i)$  be the time it takes to pebble  $q_i$  in the proposed construction for any instance with  $i$  variables,  $i$  clauses, and only existential quantifiers.

Suppose that  $x_i$  is allowed to be true for the existential quantifier  $Q_i = \exists$ . Then vertex  $x'_i$  is pebbled using  $s_i$  moves, where  $s_i = 3n - 3i + 6$ . Similarly, vertices  $d_i$  and  $f_i$  are pebbled using  $s_i - 1$  and  $s_i - 2$  moves respectively. Additionally,  $f_i$  is moved to  $\bar{x}'_i$  and then  $x_i$  is moved to  $x'_i$  in the following step, for a total of two more moves. We then pebble  $q_{i+1}$  using  $T(i + 1)$  moves and finish by placing a pebble on  $\bar{x}_i$  and moving it to  $c_i, b_i, a_i$ , and  $q_i$ , for five more moves. Finally, we use six more moves to pebble an additional clause. Thus, in this case,

$$T_{\text{true}}(i) = s_i + (s_i - 1) + (s_i - 2) + 13 + T(i + 1).$$

On the other hand, if  $x_i$  is allowed to be false for the existential quantifier  $Q_i = \exists$ , then first we pebble  $x'_i, d_i$ , and  $f_i$  sequentially, using  $s_i, s_i - 1$ , and  $s_i - 2$  moves respectively. We then move the pebble from  $f_i$  to  $\bar{x}'_i$  and then to  $\bar{x}_i$ , for a total of two more moves. We then pebble  $q_{i+1}$  using  $T(i + 1)$  moves. The pebble on  $q_{i+1}$  is subsequently moved to  $c_i$  and then  $b_i$ , using two more moves. Picking up all pebbles except those on  $b_i$  and  $x'_i$ , and using them to pebble  $f_i$  takes  $s_i - 2$  more moves. Additionally, the pebble on  $f_i$  is moved to  $\bar{x}'_i$  and then  $a_i$ , while the pebble on  $x'_i$  is moved to  $x_i$  and then  $q_i$ , for four more moves. Finally, we use six more moves to pebble an additional clause. In total,

$$T_{\text{false}}(i) = s_i + (s_i - 1) + (s_i - 2) + (s_i - 2) + 14 + T(i + 1).$$

Therefore,

$$T(i) \leq 4s_i + 10 + T(i + 1),$$

where  $s_i = 3n - 3i + 6$ . Thus,

$$T(i) \leq 12(n - i) + 34 + T(i + 1).$$

Writing  $R(i) = T(n - i)$  then gives

$$R(i) \leq 12i + 34 + R(i - 1),$$

so that  $R(n) \leq \sum_{i=1}^n (12i + 34) = 6n^2 + 40n$ . Hence, it takes at most  $6n^2 + 40n$  moves to pebble the given construction for any instance of TQBF which only includes existential quantifiers.  $\square$

## C SPECIFICATION OF CANDIDATE IMHFS

In this section we provide detailed descriptions of the iMHFs analyzed in the main body of the paper. DRSample is described in Algorithm 1, aATSample is described in Algorithm 2, Argon2iB is described in Algorithm 3 and Argon2iA is described in Algorithm 4.

## D NP-HARDNESS OF THE RED-BLUE PEBBLING COST

In this section, we consider the computational complexity of computing  $\text{rpeb}^\parallel(G, m)$ , defining a decision version below and showing it is NP-Hard.

---

**Algorithm 1:** An algorithm for sampling depth-robust graphs. [4]

---

**Function** DRSample( $n \in \mathbb{N}_{\geq 2}$ ):

```

  V := [v]
  E := {(1, 2)}
  for v ∈ [3, n] and i ∈ [2] do      // Populate edges
    | E := E ∪ {(v, GetParent(v, i))} // Get ith parent
  end
  return G := (V, E).

```

**Function** GetParent( $v, i$ ):

```

  if i = 1 then
    | u := i - 1
  else
    | g' ← [1, ⌊log2(v)⌋ + 1] // Get random range size.
    | g := min(v, 2g') // Don't make edges too long.
    | r ← [max(g/2, 2), g] // Get random edge length.
  end
  return v - r

```

---



---

**Algorithm 2:** An algorithm for sampling a high aAT graph. [4]

---

**Function** aATSample( $H = (\bar{V} = [n], \bar{E}), c \in (0, 1)$ ):

```

  V := [2n]
  E :=  $\bar{E} \cup \{(i, i+1) : i \in [2n-1]\}$ 
  for v ∈ [n+1, 2n] and i ∈ [2] do // Populate new edges of graph.
    | E := E ∪ {(v, GetParentc(v, i))} // Get ith parent of node v
  end
  return G := (V, E).

```

**Function** GetParent<sup>c</sup>( $v, i$ ):

```

  if i = 1 then
    | u := i - 1
  end
  else if v ≤ n then
    | u := GetParentH(v, i) // DRSample
  end
  else
    | m := ⌊c log(n)⌋
    | b := (v - n) mod ⌊n/m⌋
    | u := bm
  end
  return u

```

---

The decision problem  $\text{rbpeb}^\parallel$  is defined as follows:

**Input:** a DAG  $G$  on  $n$  nodes, parameter  $c_b, c_r$ , and integers  $m, d > 0$ .

**Output:** Yes, if  $\text{rbpeb}^\parallel(G, m) \leq d$ ; otherwise No.

---

**Algorithm 3:** An algorithm for sampling depth-robust graphs. [11]

---

**Function** Argon2iB( $n \in \mathbb{N}_{\geq 2}$ ):

```

  V := [v]
  E := {(1, 2)}
  for v ∈ [3, n] and i ∈ [2] do      // Populate edges
    | E := E ∪ {(v, GetParent(v, i))} // Get ith parent
  end
  return G := (V, E).

```

**Function** GetParent( $v, i$ ):

```

  if i = 1 then
    | u := i - 1
  else
    | N := 232 // Set sample range.
    | g ← [1, N] // Get random range length.
    | r := ⌊ $\frac{g^2}{N^2}v$ ⌋ // Set quadratic dependency.
  end
  return v - r

```

---



---

**Algorithm 4:** An algorithm for sampling depth-robust graphs. [11]

---

**Function** Argon2iA( $n \in \mathbb{N}_{\geq 2}$ ):

```

  V := [v]
  E := {(1, 2)}
  for v ∈ [3, n] and i ∈ [2] do      // Populate edges
    | E := E ∪ {(v, GetParent(v, i))} // Get ith parent
  end
  return G := (V, E).

```

**Function** GetParent( $v, i$ ):

```

  if i = 1 then
    | u := i - 1
  else
    | N := 232 // Set sample range.
    | g ← [1, N] // Get random range length.
    | r := ⌊ $\frac{g}{N}v$ ⌋ // Set linear dependency.
  end
  return v - r


```

---

We now show that it is NP – Hard to compute  $\text{rbpeb}^\parallel(G, m)$ . Quanquan Liu [26] observed that when  $c_r = 0$  the problem is PSPACE – Hard via a straightforward reduction from minimum space black pebbling. As we observed previously, when  $c_b/c_r \in \mathcal{O}(\text{poly}(n))$  the decision problem is in NP and has a fundamentally different structure. We show that even when the cost of red moves is significant, the problem remains NP – Hard.

**THEOREM D.1.** For  $c_b \leq 3nc_r$ , the problem  $\text{rbpeb}^\parallel$  is NP – Hard.

Gilbert *et al.* showed that the minimum space black pebbling problem was PSPACE – Hard by reduction from the Truly Quantified Boolean Formula (TQBF) problem. For more details about the Gilbert *et al.* [23] reduction, we refer an interested reader to



**Fig. 4: Graph  $G_{TQBF}$  for  $\exists x_1, x_2, x_3, x_4 \text{ s.t. } (x_1 \vee x_2 \vee x_4) \wedge (x_2 \vee x_3 \vee \bar{x}_4)$ .**


Appendix B. We note that an instance  $\phi$  of 3-SAT with  $n$  variables is still a TQBF instance (albeit with no  $\forall$  quantifiers). Thus, given an instance  $\phi$  of 3-SAT with  $n$  variables, we can create the corresponding DAG  $G_\phi$ , as described in the reduction of Gilbert *et al.* [23]. The graph  $G_\phi$  has the property that it can be pebbled with at most  $3n + 3$  black pebbles if and only if  $\phi$  is satisfiable.

We detail a gadget to append to  $G_\phi$  to create a graph  $H_\phi$  so that  $\text{rbpeb}^{\parallel}(H) = x_1$  if  $\phi$  is a satisfiable assignment, but  $\text{rbpeb}^{\parallel}(H_\phi) > x_1$  if  $\phi$  is not a satisfiable assignment. The key goal of the additional gadget is to ensure that we cannot *significantly* reduce the number of red moves (computation costs) by including a few blue move.

Since an instance  $\phi$  of 3-SAT with  $n$  variables is still a TQBF instance (albeit with no  $\forall$  quantifiers), we can create the corresponding DAG  $G_\phi$ , as described in the reduction of Gilbert *et al.* [23]. For DAG  $G_\phi$  with  $t$  vertices, there exist unique pyramid gadgets with  $3n+3, 3n+2, 3n+1, \dots, 1$  vertices in the bottom layer. Let  $\Delta_i$  be the pyramid gadget with  $\phi$  vertices in the bottom layer. Additionally, let  $\alpha_i$  be the vertex *above* the apex of pyramid  $\Delta_i$ . Create a directed path  $P_1$  with  $3n+3$  vertices so that for each  $1 \leq i \leq 3n+3$ , connect an edge to vertex  $3i-2$  of  $P_1$  from the top vertex of  $\Delta_{3n+4-i}$ .

We then connect the final vertex of  $P_1$  to a directed path  $P_2$  with

$$\left( \frac{(3n+1)(3n)}{2} + 1 \right) + \left( \frac{(3n-2)(3n-3)}{2} + 1 \right) + \dots + (21+1) + (6+1) = \frac{3}{2}n(n+1)^2 + n$$



**Fig. 5: Path  $P$  for  $H_\phi$ .**

vertices. Moreover, the first  $\frac{(3n+1)(3n)}{2}$  vertices of  $P_2$  each have an edge from separate vertices of  $\Delta_{3n+1}$ , starting with the vertices in the bottom layer and moving upwards. We also create an edge to the following vertex from the vertex  $\alpha_{3n+1}$ . The next  $\frac{(3n-2)(3n-3)}{2}$  vertices of  $P_2$  each have an edge from separate vertices of  $\Delta_{3n-2}$ , starting with the vertices in the bottom layer and moving upwards. We also create an edge to the following vertex from the vertex  $\alpha_{3n-2}$ . We continue this process until all vertices from all pyramids of the form  $\Delta_{3i+1}$  are connected to  $P_2$ , as well as the vertices  $\alpha_{3i+1}$ . Finally, we connect  $P_2$  to a sink node. Then by setting  $P$  to be the path  $P_1$  concatenated with  $P_2$ , we have the following result:

LEMMA D.2.  $P$  contains exactly  $3n+3 + \sum_{i=1}^n \left( \frac{(3i+1)(3i)}{2} + 1 \right) = 4n+3 + \frac{3}{2}n(n+1)^2 = \frac{3n^3+6n^2+11n+6}{2}$  vertices.

See Figure 5 for an example. Let  $H_\phi = G_\phi \cup P$ . We claim that  $H_\phi$  with capacity  $3n+4$  will have a certain pebbling cost if and only if  $\phi$  is satisfiable. Thus, if  $\phi$  is satisfiable, the optimal pebbling will correspond to the minimum space black pebbling and will require 0 blue moves.

LEMMA D.3. If  $\phi$  is satisfiable, then there exists a pebbling strategy of  $H_\phi$  with capacity  $3n+4$  and cost exactly

$$\left( \frac{3n^3 + 15n^2 + 40n + 14 + 12c}{2} \right) c_r.$$

PROOF. The total number of nodes in  $G_\phi$  corresponding to variable assignments from the GLT construction is

$$6n + \sum_{i=4}^{3n+3} i = \frac{9n^2 + 33n + 12}{2}.$$

This can be visualized in Figure 4 by the nodes on the left hand side, excluding the nodes  $q_i$ . Additionally, there are  $n$  nodes  $q_i$ , six nodes for each of the  $c$  clauses  $p_i$  for  $1 \leq i \leq c$ , and an additional node for  $p_0$ . Thus,

$$6c + 1 + 7n + \sum_{i=4}^{3n+3} i = \frac{9n^2 + 35n + 14}{2} + 6c$$

nodes must be pebbled in  $G_\phi$ .

By Lemma D.2, the number of nodes in the additional path  $P$  is  $\frac{3n^3 + 6n^2 + 21n + 18}{2}$ . However, pebbling  $P$  requires that each of the pyramids  $\Delta_{3i+1}$  are pebbled a second time, as well as each  $\alpha_{3i+1}$ , requiring an additional

$$\sum_{i=1}^n \left( \frac{(3i+1)(3i)}{2} + 1 \right) = n + \frac{3}{2}n(n+1)^2 = \frac{3n^3 + 6n^2 + 5n}{2}$$

steps.

Thus, the total number of steps required to pebble  $H_\phi$  is

$$\frac{9n^2 + 35n + 14}{2} + 6c + \frac{3n^3 + 6n^2 + 5n}{2} = \frac{3n^3 + 15n^2 + 40n + 14 + 12c}{2}.$$

The GLT construction has pebbling number  $3n + 3$ . Since the nodes in  $P$  are ordered corresponding to the natural pebbling order in  $G_\phi$ , a single additional pebble suffices for  $P$ . Thus, if the capacity of  $G_\phi$  is  $3n + 4$ , then all pebbling moves can be achieved with red moves, so there exists a pebbling strategy with total cost is  $\left( \frac{3n^3 + 15n^2 + 40n + 14 + 12c}{2} \right) c_r$ .

By construction, the pebbling strategy of Lemma D.3 is the optimal pebbling with only red moves. Thus, it remains to show that no strategy containing any blue pebbles has better cost.

Since blue pebbles are more expensive than red pebbles, the only place the above strategy can be possibly improved would be using blue pebbles on nodes in  $H_\phi$  that are pebbled multiple times in the strategy of Lemma D.3. As it turns out, the only nodes that are pebbled multiple times are the pyramids  $\Delta_{3i+1}$ , which are each pebbled twice, as well as the vertices  $\alpha_{3i+1}$ . However, all of the vertices in  $\Delta_{3i+1}$  are parents of vertices in  $P$ . Thus, a blue pebble on any of these nodes must be returned to red, and then subsequently discarded. That is, a blue pebble will replace at most one red move in the pebbling strategy of Lemma D.3. The number of the other red moves is unaffected so that the overall cost cannot be cheaper. Therefore, no pebbling strategy has a better cost than  $\left( \frac{3n^3 + 15n^2 + 40n + 14 + 12c}{2} \right) c_r$ .  $\square$

LEMMA D.4. If  $\phi$  is unsatisfiable, then the pebbling cost of  $H_\phi$  with capacity  $3n + 4$  is greater than

$$\left( \frac{3n^3 + 15n^2 + 40n + 14 + 12c}{2} \right) c_r.$$

PROOF. By the construction of the DAG  $H_\phi$ , if  $\phi$  is unsatisfiable, then  $H_\phi$  has pebbling number at least  $3n+5$ . Thus, if  $H_\phi$  has capacity  $3n+4$ , any pebbling strategy must have a blue pebble at some point.

As in the proof of Lemma D.3, any optimal strategy only places blue pebbles on nodes that the strategy of Lemma D.3 pebbles twice. Again, the only nodes that are pebbled multiple times are the pyramids  $\Delta_{3i+1}$ , which are each pebbled twice, as well as the vertices  $\alpha_{3i+1}$ . Since each vertex in  $\Delta_{3i+1}$  is a parent of a vertex in  $P$ , a blue pebble on any of these nodes must be returned to red, and then subsequently discarded. Thus, a blue pebble will replace at most one red move in the pebbling strategy of Lemma D.3. The remaining red moves are unaffected, so the overall cost is more expensive, under the assumption that blue pebbles are more expensive than red moves. Therefore, any pebbling strategy has a cost greater than  $\left( \frac{3n^3 + 15n^2 + 40n + 14 + 12c}{2} \right) c_r$ .  $\square$

Together, Lemma D.3 and Lemma D.4 imply Theorem D.1.

**Reminder of Theorem D.1.** For  $c_b \leq 3nc_r$ , the problem  $\text{rbpeb}^\parallel$  is NP – Hard.

**Proof of Theorem D.1:** First, we remark that given a DAG  $H_\phi$  with some capacity  $m$ , as well as a complete pebbling strategy as the certificate, the certificate can be verified in polynomial time by checking the validity of each step in the pebbling strategy. Thus, the computation of  $\text{rbpeb}^\parallel(H_\phi)$  is in NP.

We now reduce 3 – SAT to the computation of  $\text{rbpeb}^\parallel(H_\phi)$ . Now, given an instance  $\phi$  of 3 – SAT with  $n$  variables, we construct the above DAG  $H_\phi$ . This procedure clearly takes polynomial time. Moreover, by Lemma D.3, if  $\phi$  is satisfiable, then the optimal pebbling cost of  $H_\phi$  with capacity  $3n + 4$  is exactly

$$\left( \frac{3n^3 + 15n^2 + 40n + 14 + 12c}{2} \right) c_r.$$

On the other hand, by Lemma D.4, if  $\phi$  is unsatisfiable, then the pebbling cost of  $H_\phi$  with capacity  $3n + 4$  is greater than

$$\left( \frac{3n^3 + 15n^2 + 40n + 14 + 12c}{2} \right) c_r.$$

Thus, the computation of  $\text{rbpeb}^\parallel(H_\phi)$  distinguishes whether  $\phi$  is satisfiable or not. Since 3 – SAT is NP – Hard, it follows that the computation of  $\text{rbpeb}^\parallel(H_\phi)$  is NP – Hard.  $\square$