

Depth-Robust Graphs and Their Cumulative Memory Complexity

Joël Alwen¹, Jeremiah Blocki², and Krzysztof Pietrzak¹

¹ IST Austria

² Purdue University

Abstract. Data-independent Memory Hard Functions (iMHFS) are finding a growing number of applications in security; especially in the domain of password hashing. An important property of a concrete iMHF is specified by fixing a directed acyclic graph (DAG) G_n on n nodes. The quality of that iMHF is then captured by the following two pebbling complexities of G_n :

- The parallel cumulative pebbling complexity $\Pi_{cc}^{\parallel}(G_n)$ must be as high as possible (to ensure that the amortized cost of computing the function on dedicated hardware is dominated by the cost of memory).
- The sequential space-time pebbling complexity $\Pi_{st}(G_n)$ should be as close as possible to $\Pi_{cc}^{\parallel}(G_n)$ (to ensure that using many cores in parallel and amortizing over many instances does not give much of an advantage).

In this paper we construct a family of DAGs with best possible parameters in an asymptotic sense, i.e., where $\Pi_{cc}^{\parallel}(G_n) = \Omega(n^2 / \log(n))$ (which matches a known upper bound) and $\Pi_{st}(G_n)$ is within a constant factor of $\Pi_{cc}^{\parallel}(G_n)$.

Our analysis relies on a new connection between the pebbling complexity of a DAG and its depth-robustness (DR) – a well studied combinatorial property. We show that high DR is *sufficient* for high Π_{cc}^{\parallel} . Alwen and Blocki (CRYPTO’16) showed that high DR is *necessary* and so, together, these results fully characterize DAGs with high Π_{cc}^{\parallel} in terms of DR.

Complementing these results, we provide new upper and lower bounds on the Π_{cc}^{\parallel} of several important candidate iMHFs from the literature. We give the first lower bounds on the memory hardness of the Catena and Balloon Hashing functions in a parallel model of computation and we give the first lower bounds of any kind for (a version) of Argon2i.

Finally we describe a new class of pebbling attacks improving on those of Alwen and Blocki (CRYPTO’16). By instantiating these attacks we upperbound the Π_{cc}^{\parallel} of the Password Hashing Competition winner Argon2i and one of the Balloon Hashing functions by $O(n^{1.71})$. We also show an upper bound of $O(n^{1.625})$ for the Catena functions and the two remaining Balloon Hashing functions.

1 Introduction

Moderately hard functions. Functions which are “moderately” hard to compute have found a variety of practical applications including password hashing, key-derivation and for proofs of work. In the context of password hashing, the goal is

to minimize the damage done by a security breach where an adversary learns the password file; Instead of storing $(login, password)$ tuples in the clear, one picks a random salt and stores a tuple $(login, f(password, salt), salt)$, where $f(\cdot)$ is a moderately hard function $f(\cdot)$. This comes at a price, the server verifying a password must evaluate $f(\cdot)$, which thus cannot be too hard. On the other hand, if a tuple $(login, y, salt)$ is leaked, an adversary who tries to find the password by a dictionary attack must evaluate $f(\cdot)$ for every attempt. A popular moderately hard function is PBKDF2 (Password Based Key Derivation Function 2) [?], which basically just iterates a cryptographic hash function H several times (1024 is a typical value).

Unfortunately a moderately hard function like PBKDF2 offers much less protection against adversaries who can build customized hardware to evaluate the underlying hash function than one would hope for. The reason is that the cost of computing a hash function H like SHA256 or MD5 on an ASIC (Application Specific Integrated Circuit) is orders of magnitude smaller than the cost of computing H on traditional hardware [?,?].

Memory-Bound and Memory-Hard Functions. [?] recognized that cache-misses are more egalitarian than computation, in the sense that they cost about the same on different architectures. They propose “memory-bound” functions, which are functions that will incur many expensive cache-misses. This idea was further developed by [?].

Along similar lines, Percival [?] observes that unlike computation, memory costs tend to be relatively stable across different architectures, and suggests to use memory-hard functions (MHF) for password hashing. [?] also introduced the `scrypt` MHF which has found a variety of applications in practice. Very recently it has been proven to indeed offer optimal time/space trade-offs in the random oracle model [?,?].

MHF’s come in two flavours, data-dependent MHF’s (dMHF) such as `scrypt`, and data independent MHF’s (iMHF). The former are potentially easier to construct and allow for more extreme memory-hardness [?,?], but they leave open the possibility of side-channel attacks [?], thus iMHF’s are preferable when the inputs are sensitive, as in the case of password hashing. We shortly discuss the state of the art for dMHF’s at the end of this section.

iMHF as Graphs. An iMHF comes with an algorithm that computes the function using a fixed memory access pattern. In particular the pattern is independent of the input. Such functions can thus be described by a directed acyclic graph (DAG) G , where each node v of the graph corresponds to some intermediate value ℓ_v that appears during the computation of the function, and the edges capture the computation: if ℓ_v is a function of previously computed values $\ell_{i_1}, \dots, \ell_{i_\delta}$, then the nodes i_1, \dots, i_δ are parents of v in G . For an iMHF F , we’ll denote with $G(F)$ the underlying graph. For example $G(\text{PBKDF2})$ is simply a path.

Graph Labeling Functions. Not only can an iMHF be captured by a graph as just outlined, we will actually construct iMHF’s by first specifying a graph, and then defining a “labeling function” on top of it: Given a graph G with node set $V = [n] = \{1, 2, \dots, n\}$, a hash function $H : \{0, 1\}^* \rightarrow \{0, 1\}^w$ and some input x ,

define the labeling of the nodes of G as follows: a source (a node v with indegree 0) has label $\ell_v(x) = H(v, x)$, a node v with parents $v_1 < v_2 < \dots < v_\delta$ has label $\ell_v(x) = H(v, \ell_{v_1}(x), \dots, \ell_{v_\delta}(x))$. For a DAG G with a unique sink s we define the labeling function of G as $f_G(x) = \ell_s(x)$. Note that using the convention from the previous paragraph, we have $G(f_G) = G$.

The Black Pebbling Game One of the main techniques for analyzing iMHF is to use pebbling games played on graphs. First introduced by Hewitt and Paterson [?] and Cook [?] the (sequential) black pebbling game (and its relatives) have been used to great effect in theoretical computer science. Some early applications include space/time trade-offs for various computational tasks such as matrix multiplication [?], the FFT [?,?], integer multiplication [?] and solving linear recursions [?,?]. More recently, pebbling games have been used for various cryptographic applications including proofs of space [?,?], proofs of work [?,?], leakage-resilient cryptography [?], garbled circuits [?], one-time computable functions [?], adaptive security proofs [?,?] and memory-hard functions [?,?,?]. It's also an active research topic in proof complexity (cf. the survey on <http://www.csc.kth.se/~jakobn/research/PebblingSurveyTMP.pdf>).

The black pebbling game is played over a fixed directed acyclic graph (DAG) $G = (V, E)$ in rounds. The goal of the game is to pebble all sink nodes of G (not necessarily simultaneously). Each round $i \geq 1$ is characterized by its pebbling configuration $P_i \subseteq V$ which denotes the set of currently pebbled nodes. Initially $P_0 = \emptyset$, i.e., all nodes are unpebbled. P_i is derived from the previous configuration P_{i-1} according to two simple rules. (1) A node v may be pebbled (added to P_i) if, in the previous configuration all of its parents were pebbled, i.e., $\text{parents}(v) \subseteq P_{i-1}$. (2) A pebble can always be removed from P_i . In the sequential version rule (1) may be applied at most once per round while in the parallel version no such restriction applies. A sequence of configurations $P = (P_0, P_1, \dots)$ is a (sequential) pebbling of G if it adheres to these rules and each sink node of G is contained in at least one configuration.

From a technical perspective, in this paper we investigate upper and lower bounds on various pebbling complexities of graphs, as they can be related to the cost of evaluating the “labeling function” f_G (to be defined below) in various computational models. In particular, let \mathcal{P}_G and \mathcal{P}_G^\parallel denote all valid sequential and parallel peblings of G , respectively. We are interested in the *parallel* cumulative pebbling complexity of G , denoted $\Pi_{cc}^\parallel(G)$, and the sequential space-time complexity of G , denoted $\Pi_{st}(G)$, which are defined as

$$\Pi_{cc}^\parallel(G) = \min_{(P_1, \dots, P_t) \in \mathcal{P}_G^\parallel} \sum_{i=1}^t |P_i| \quad \Pi_{st}(G) = \min_{(P_1, \dots, P_t) \in \mathcal{P}_G} t \cdot \max_i (|P_i|).$$

A main technical result of this paper is a family of graphs with high (in fact, as we'll discuss below, maximum possible) Π_{cc}^\parallel complexity, and where the Π_{st}

complexity is not much higher than the Π_{cc}^{\parallel} complexity.³ Throughout, we'll denote with \mathbb{G}_n the set of all DAGs on n nodes and with $\mathbb{G}_{n,d} \subseteq \mathbb{G}_n$ the DAGs where each node has indegree at most d .

Theorem 1. *There exists a family of DAGs $\{G_n \in \mathbb{G}_{n,2}\}_{n \in \mathbb{N}}$ where*

1. *the parallel cumulative pebbling complexity is*

$$\Pi_{cc}^{\parallel}(G_n) \in \Omega(n^2 / \log(n))$$

2. *and where the sequential space-time complexity matches the parallel cumulative pebbling complexity up to a constant*

$$\Pi_{st}(G_n) \in O(n^2 / \log(n)).$$

The lower bound on Π_{cc}^{\parallel} in item 1. above is basically optimal due to the following bound from [?].⁴

Theorem 2 ([?, Thm. 8]). *For any constant $\epsilon > 0$ and sequence of DAGs $\{G_n \in \mathbb{G}_{n,\delta_n}\}_{n \in \mathbb{N}}$ it holds that*

$$\Pi_{cc}^{\parallel}(G_n) = o\left(\frac{\delta_n n^2}{\log^{1-\epsilon}}\right).$$

In particular if $\delta_n = O(\log^{1-\epsilon})$ then $\Pi_{cc}^{\parallel}(G_n) = o(n^2)$, and

$$\text{if } \delta_n = \Theta(1) \text{ then } \Pi_{cc}^{\parallel}(G_n) = o(n^2 / \log^{1-\epsilon}(n)). \quad (1)$$

Pebbling vs. Memory-hardness. The reason to focus on the graph $G = G(F)$ underlying an iMHF F is that clean combinatorial properties of G – i.e., bounds on the pebbling complexities – imply both upper and lower bounds on the cost of evaluating F in various computational models. For upper bounds (i.e., attacks), no further assumption on F are required to make the transition from pebbling to computation cost. For lower bounds, we have to assume that there's no “shortcut” in computing F , and the only way is to follow the evaluation sequence as given by G . Given the current state of complexity theory, where not even superlinear lower bounds on evaluating any function in \mathcal{NP} are known, we cannot hope to exclude such shortcuts unconditionally. Instead, we assume that the underlying hash function H is a random oracle and circuits are charged unit cost for queries to the random oracle.

For our lower bounds, we must insist on G having constant indegree. The reason is that in reality H must be instantiated with some cryptographic hash

³ Note that $\Pi_{cc}^{\parallel}(G) \leq \Pi_{st}(G)$ as parallelism can only help, and space-time complexity (i.e., number of rounds times the size of the largest state) is always higher than cumulative complexity (the sum of the sizes of all states).

⁴ The statement below is obtained from the result in [?] by treating the core-memory ratio as a constant and observing that, trivially, at most n pebbles are on G during a balloon phase and at most n pebbles are placed in one step during a balloon phase.

function like SHA1, and the indegree corresponds to the input length on which H is invoked. To evaluate H on long inputs, one would typically use some iterated construction like Merkle-Damgard, and the assumption that H behaves like a black-box that can only be queried once the entire input is known would be simply wrong in this case.

As advocated in [?], bounds on $\Pi_{cc}^{\parallel}(G)$ are a reasonable approximation for the cost of evaluating f_G in dedicated hardware, whereas a bound on $\Pi_{st}(G)$ gives an upper bound on the cost of evaluating f_G on a single processor machine. The reason [?] consider cumulative complexity for lower and space-time complexity for the upper bound is that when lower bounding the cost of evaluating f_G we do want to allow for amortization of the cost over arbitrary many instances,⁵ whereas for our upper bound we don't want to make such an assumption. The reason we consider parallel complexity for the lower and only sequential for the upper bound is due to the fact that an adversary can put many (cheap) cores computing H on dedicated hardware, whereas for the upper bound we only want to consider a single processor machine.

If $\Pi_{cc}^{\parallel}(G)$ is sufficiently larger than $|V(G)|$ (in Theorem 1 it's almost quadratic), then the cost of evaluating f_G in dedicated hardware is dominated by the memory cost. As memory costs about the same on dedicated hardware and general purpose processors, if our G additionally satisfies $\Pi_{cc}^{\parallel}(G) \approx \Pi_{st}(G)$, then we get a function f_G whose evaluation on dedicated hardware is not much cheaper than evaluating it on an off the shelf machine (like a single core x86 architecture). This is exactly what the family from Theorem 1 achieves. We elaborate on these computational models and how they are related to pebbling in the full version.

On the positive side, previous to this work, the construction with the best asymptotic bounds was due to [?] and achieved $\Pi_{cc}^{\parallel}(G_n) \in \Omega(n^2/\log^{10}(n))$. However the exponent 10 (and the complexity of the construction) makes this construction uninteresting for practical purposes.

On the negative side [?,?,?] have broken many popular iMHFs in a rather strong asymptotic sense. For example, in [?], the graph underlying Argon2i-A [?], the winner of the recent Password Hashing Competition⁶, was shown to have Π_{cc}^{\parallel} complexity $\tilde{O}(n^{1.75})$. For Catena [?] the upper bound $O(n^{5/3})$ is shown in [?]. In [?] these results were extended to show that Argon2i-B [?] has $\Pi_{cc}^{\parallel}(G) = O(n^{1.8})$. Moreover [?] show that for random instances of these

⁵ Π_{cc}^{\parallel} satisfies a direct product property: pebbling k copies of G cost k times as much as pebbling G , i.e., $\Pi_{cc}^{\parallel}(G^k) = k \cdot \Pi_{cc}^{\parallel}(G)$, but this is not true for Π_{st} complexity.

⁶ The Argon2 specification [?] has undergone several revisions all of which are regularly referred to as "Argon2." To avoid confusion we follow [?] and use Argon2i-A [?] to denote the version of Argon2i from the password hashing competition [?] and we use Argon2i-B [?] to refer the version of Argon2 that is currently being considered for standardization by the Cryptography Form Research Group (CFRG) of the IRTF. We conjecture that the techniques introduced in this paper could also be used to establish tighter bounds for Argon2i-B. However, we leave this as an open challenge for future work.

functions (which is how they are supposed to be used in practice) the attacks actually have far lower Π_{cc}^{\parallel} than these asymptotic analyses indicate.

A New Generic Attack and Its Applications. In this work we improve on the attacks of [?, ?, ?] (Section 6). We give a new parallel pebbling strategy for pebbling DAGs which lack a generalization of depth-robustness. Next we investigate this property for the case of Argon2i-A, the three Balloon-Hashing variants and both Catena variants to obtain new upper bounds on their respective Π_{cc}^{\parallel} . For example, we further improve the upper bound on Π_{cc}^{\parallel} for Argon2i-A and the Single Buffer variant of Balloon-Hashing from $\tilde{O}(n^{1.75})$ to $O(n^{1.708})$.

New Security Proofs. Complementing these results, in Section 5, we give the first security proofs for a variety of iMHFs. Hitherto the only MHF with a full security proof in a parallel computational model was [?] which employed relatively construction specific techniques. When restricted to sequential computation the results of [?, ?] show that Catena has Π_{st} complexity $\Omega(n^2)$. Similar results are also shown for Argon2i-A and Balloon Hashing in [?].

In this work we introduce two new techniques for proving security of iMHFs. In the case of Argon2i-A and Argon2i-B we analyze its depth-robustness to show that its Π_{cc}^{\parallel} is at least $\tilde{\Omega}(n^{5/3})$. The second technique involves a new combinatorial property called dispersion which we show to imply lower bounds on the Π_{cc}^{\parallel} of a graph. We investigate the dispersion properties of the Catena and Balloon Hashing variants to show their Π_{cc}^{\parallel} to be $\tilde{\Omega}(n^{1.5})$. Previously no (non-trivial) lower bounds on Π_{cc}^{\parallel} were known for Argon2i-A, Catena or Balloon Hashing. Interestingly, our results show that Argon2i-A and Argon2i-B have better asymptotic security guarantees than Catena since $\Pi_{cc}^{\parallel} = \Omega(n^{5/3})$ for Argon2i-A and $\Pi_{cc}^{\parallel} = O(n^{13/8})$ for Catena.

While these lower bounds are significantly worse than what we might ideally hope for in a secure iMHF (e.g., $\Pi_{cc}^{\parallel} \geq \Omega(n^2/\log(n))$), we observe that, in light of our new attacks in Section 6, they are nearly tight. Unfortunately, together with the bounds on the sequential complexity of these algorithms our results do highlight a large asymptotic gap between the memory needed when computing the functions on parallel vs. sequential computational devices.

A table summarizing the asymptotic cumulative complexity of various iMHFs can be found in Table 1.

Depth-Robust Graphs. The results in this work rely on a new connection between the depth-robustness of a DAG and its Π_{cc}^{\parallel} complexity. A DAG G is (e, d) -depth-robust if, after removing any subset of at most e nodes there remains a directed path of length at least d . First investigated by Erdős, Graham and Szemerédi [?], several such graphs enjoying low indegree and increasingly extreme depth-robustness have been constructed in the past [?, ?, ?, ?, ?] mainly in the context of proving lower-bounds on circuit complexity and Turing machine time. Depth-robustness has been used as a key tool in the construction of cryptographic objects like proofs of sequential work [?]. In fact depth-robust graphs were already used as a building block in the construction of a high Π_{cc}^{\parallel} graph in [?].

| Algorithm | Lowerbound | Upperbound | Appearing In |
|--|--|--|--------------|
| Argon2i-A | | $\tilde{O}(n^{1.75})$ | [?] |
| Argon2i-A | $\tilde{\Omega}(n^{1.6})$ | $\tilde{O}(n^{1.708})$ | This Work |
| Argon2i-B | | $O(n^{1.8})$ | [?] |
| Argon2i-B | $\tilde{\Omega}(n^{1.6})$ | | This Work |
| Balloon-Hashing: Linear and Double Buffer (DB) | | $O(n^{1.67})$ | [?] |
| Balloon-Hashing: Linear and Double Buffer(DB) | $\tilde{\Omega}(n^{1.5})$ | $\tilde{O}(n^{1.625})$ | This Work |
| Balloon-Hashing: Single Buffer (SB) | | $\tilde{O}(n^{1.75})$ | [?] |
| Balloon-Hashing: Single Buffer (SB) | $\tilde{\Omega}(n^{1.6})$ | $\tilde{O}(n^{1.708})$ | This Work |
| Catena: Dragonfly | | $O(n^{1.67})$ | [?] |
| Catena: Dragonfly | $\tilde{\Omega}(n^{1.5})$ | $\tilde{O}(n^{1.625})$ | This Work |
| Catena: Butterfly | | $O(n^{1.67})$ | [?] |
| Catena: Butterfly | $\tilde{\Omega}(n^{1.5})$ | $o(n^{1.625})$ | This Work |
| [?] | $\Omega\left(\frac{n^2}{\log^{10} n}\right)$ | | [?] |
| Theorem 1 | $\Omega\left(\frac{n^2}{\log n}\right)$ | | This Work |
| Arbitrary iMHF | | $O\left(\frac{n^2 \log \log n}{\log n}\right)$ | [?] |

Table 1: Overview of the asymptotic cumulative complexity of various iMHF.

Depth-Robustness and Π_{cc}^{\parallel} . While the flavour of the results in this work are related to those of [?] the techniques are rather different. As mentioned above already, they stem from a new tight connection between depth-robustness and Π_{cc}^{\parallel} . A special case of this connection shows that if G is (e, d) -depth-robust, then its Π_{cc}^{\parallel} can be lower bounded as

$$\Pi_{cc}^{\parallel}(G) \geq e \cdot d .$$

This complements a result from [?], which gives a pebbling strategy that is efficient for graphs of low depth-robustness. Thus a DAG has high Π_{cc}^{\parallel} if and only if it is very depth-robust.

Moreover, we give a new tool for reducing the indegree of a DAG while not reducing the Π_{cc}^{\parallel} of the resulting graph (in terms of its size). Together these results directly have some interesting consequences

- The family of DAGs $\{G_n \in \mathbb{G}_{n, \log(n)}\}_{n \in \mathbb{N}}$ from Erdős et al. [?] have optimally high $\Pi_{cc}^{\parallel}(G_n) \in \Omega(n^2)$.
- Using our indegree reduction we can turn the above family of $\log(n)$ indegree into a family of indegree 2 DAGs $\{G'_n \in \mathbb{G}_{(n, 2)}\}_{n \in \mathbb{N}}$ with $\Pi_{cc}^{\parallel}(G'_n) \in \Omega(n^2 / \log(n))$, which by Theorem 2 is optimal for constant indegree graphs.

Data-Dependent MHFs. One can naturally extend the Π_{cc}^{\parallel} notion also to “dynamic” graphs – where some edges are only revealed as some nodes are pebbled – in order to analyse data-dependent MHFs (dMHF) like **script**. In this model, [?] show that $\Pi_{cc}^{\parallel}(\mathbf{script}) = \Omega(n^2/\log^2(n))$. Unfortunately unlike for iMHFs, for dMHFs we do not have a proof that a lower bound on Π_{cc}^{\parallel} implies roughly the same lower bound on the cumulative memory complexity in the random oracle model.⁷ Recently a “direct” proof (i.e., avoiding pebbling arguments) – showing that **script** has optimal cumulative memory complexity $\Omega(n^2)$ – has been announced, note that this bound is better than what we can hope to achieve for iMHFs (as stated in Theorem 2). Unfortunately, the techniques that have now been developed to analyse dMHFs seem not to be useful for the iMHF setting.

2 Pebbling Complexities and Depth-Robustness of Graphs

We begin by fixing some common notation. We use the sets $\mathbb{N} = \{0, 1, 2, \dots\}$, $\mathbb{N}^+ = \{1, 2, \dots\}$, and $\mathbb{N}_{\geq c} = \{c, c+1, c+2, \dots\}$ for $c \in \mathbb{N}$. Further, we also use the sets $[c] := \{1, 2, \dots, c\}$ and $[b, c] = \{b, b+1, \dots, c\}$ where $b \in \mathbb{N}$ with $b \leq c$. For a set of sets $A = \{B_1, B_2, \dots, B_z\}$ we use the notation $\|A\| := \sum_i |B_i|$.

2.1 Depth-Robust Graphs

We say that a directed acyclic graph (DAG) $G = (V, E)$ has *size* n if $|V| = n$. A node $v \in V$ has indegree $\delta = \text{indeg}(v)$ if there exist δ incoming edges $\delta = |(V \times \{v\}) \cap E|$. More generally, we say that G has indegree $\delta = \text{indeg}(G)$ if the maximum indegree of any node of G is δ . A node with indegree 0 is called a source node and one 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 $v \in V$. 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 **parents** (**ancestors**). We denote the set of all sinks of G with $\text{sinks}(G) = \{v \in V : \nexists (v, u) \in E\}$ — note that $\text{ancestors}(\text{sinks}(G)) = V$. 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 $\text{length}(p) := z$. The depth $d = \text{depth}(G)$ of DAG G is the length of the longest directed path in G .

We will often consider graphs obtained from other graphs by removing subsets of nodes. Therefore if $S \subset V$ then we denote by $G - S$ the DAG obtained

⁷ [?] introduces a combinatorial conjecture, which if true, means that lower bounds on Π_{cc}^{\parallel} translate to cumulative memory complexity. At this point a strong variant of the conjecture has already been refuted, the state of the conjecture is updated in the eprint version [?] of the paper.

from G by removing nodes S and incident edges. The following is a central definition to our work.

Definition 1 (Depth-Robustness). For $n \in \mathbb{N}$ and $e, d \in [n]$ a DAG $G = (V, E)$ is (e, d) -depth-robust if

$$\forall S \subset V \quad |S| \leq e \Rightarrow \text{depth}(G - S) \geq d.$$

We will make use of the following lemma due to Erdős, Graham and Szemerédi [?], who showed how to construct a family of log indegree DAGs with extreme depth-robustness.

Theorem 3 ([?]). For some fixed constants $c_1, c_2, c_3 > 0$ there exists an infinite family of DAGs $\{G_n \in \mathbb{G}_{n, c_3 \log(n)}\}_{n=1}^{\infty}$ such that G_n is $(c_1 n, c_2 n)$ -depth-robust.

2.2 Graph Pebbling

We fix our notation for the parallel graph pebbling game following [?].

Definition 2 (Parallel/Sequential Graph Pebbling). Let $G = (V, E)$ be a DAG and let $T \subseteq V$ be a target set of nodes to be pebbled. A pebbling configuration (of G) is a subset $P_i \subseteq V$. A legal parallel pebbling of T is a sequence $P = (P_0, \dots, P_t)$ of pebbling configurations of G where $P_0 = \emptyset$ and which satisfies conditions 1 & 2 below. A sequential pebbling additionally must satisfy condition 3.

1. At some step every target node is pebbled (though not necessarily simultaneously).

$$\forall x \in T \exists z \leq t \quad : \quad x \in P_z.$$

2. Pebbles are added only when their predecessors already have a pebble at the end of the previous step.

$$\forall i \in [t] \quad : \quad x \in (P_i \setminus P_{i-1}) \Rightarrow \text{parents}(x) \subseteq P_{i-1}.$$

3. At most one pebble placed per step.

$$\forall i \in [t] \quad : \quad |P_i \setminus P_{i-1}| \leq 1.$$

We denote with $\mathcal{P}_{G,T}$ and $\mathcal{P}_{G,T}^{\parallel}$ the set of all legal sequential and parallel pebbles of G with target set T , respectively. Note that $\mathcal{P}_{G,T} \subseteq \mathcal{P}_{G,T}^{\parallel}$. We will be mostly interested in the case where $T = \text{sinks}(G)$ and then will simply write \mathcal{P}_G and $\mathcal{P}_G^{\parallel}$.

Definition 3 (Time/Space/Cumulative Pebbling Complexity). The time, space, space-time and cumulative complexity of a pebbling $P = \{P_0, \dots, P_t\} \in \mathcal{P}_G^{\parallel}$ are defined to be:

$$\Pi_t(P) = t \quad \Pi_s(P) = \max_{i \in [t]} |P_i| \quad \Pi_{st}(P) = \Pi_t(P) \cdot \Pi_s(P) \quad \Pi_{cc}(P) = \sum_{i \in [t]} |P_i|.$$

For $\alpha \in \{s, t, st, cc\}$ and a target set $T \subseteq V$, the sequential and parallel pebbling complexities of G are defined as

$$\Pi_\alpha(G, T) = \min_{P \in \mathcal{P}_{G, T}} \Pi_\alpha(P) \quad \text{and} \quad \Pi_\alpha^\parallel(G, T) = \min_{P \in \mathcal{P}_{G, T}^\parallel} \Pi_\alpha(P).$$

When $T = \text{sinks}(G)$ we simplify notation and write $\Pi_\alpha(G)$ and $\Pi_\alpha^\parallel(G)$.

It follows from the definition that for $\alpha \in \{s, t, st, cc\}$ and any G the parallel pebbling complexity is always at most as high as the sequential, i.e., $\Pi_\alpha(G) \geq \Pi_\alpha^\parallel(G)$, and cumulative complexity is at most as high as space-time complexity, i.e., $\Pi_{st}(G) \geq \Pi_{cc}(G)$ and $\Pi_{st}^\parallel(G) \geq \Pi_{cc}^\parallel(G)$.

In this work we will consider constant in-degree DAGs $\{G_n \in \mathbb{G}_{n, \theta(1)}\}_{n \in \mathbb{N}}$, and will be interested in the complexities $\Pi_{st}(G_n)$ and $\Pi_{cc}^\parallel(G_n)$ as these will capture the cost of evaluating the labelling function derived from G_n on a single processor machine (e.g. a x86 processor on password server) and amortized AT complexity (which is a good measure for the cost of evaluating the function on dedicated hardware), respectively.

Before we state our main theorem let us observe some simple facts. Every n -node graph can be pebbled in n steps, and we cannot have more than n pebbles on an n node graph, thus

$$\forall G_n \in \mathbb{G}_n : \Pi_{cc}^\parallel(G_n) \leq \Pi_{st}(G_n) \leq n^2.$$

This upper bound is basically matched for the complete graph $K_n = (V = [n], E = \{(i, j) : 1 \leq i < j \leq n\})$ as

$$n(n-1)/2 \leq \Pi_{cc}^\parallel(K_n) \leq \Pi_{st}(K_n) \leq n^2.$$

Graph K_n has the desirable properties that its Π_{st} is within a constant factor to its Π_{cc}^\parallel complexity and that its Π_{cc}^\parallel complexity is maximally high. Unfortunately, K_n has very high indegree, which makes it useless for our purpose to construct memory-hard functions. The path $Q_n = (V = [n], E = \{(i, i+1) : 1 \leq i \leq n-1\})$ on the other hand has indegree 1 and its Π_{st} is even exactly as large as its Π_{cc}^\parallel complexity. Unfortunately it has very low pebbling complexity

$$\Pi_{cc}^\parallel(Q_n) = \Pi_{st}(Q_n) = n$$

which means that in the labelling function we get from Q_n (which is basically PBKDF2 discussed in the introduction) the evaluation cost will not be dominated by the memory cost even for large n . As stated in Theorem 1, in this paper we construct a family of graphs $\{G_n \in \mathbb{G}_{n, 2}\}_{n \in \mathbb{N}}$ which satisfies all three properties at once: (1) the graphs have indegree 2 (2) the parallel cumulative pebbling complexity is $\Pi_{cc}^\parallel(G_n) \in \Omega(n^2 / \log(n))$, which by Theorem 2 is optimal for constant indegree graphs, and (3) $\Pi_{st}(G_n)$ is within a constant factor of $\Pi_{cc}^\parallel(G_n)$.

3 Depth-Robustness Implies High Π_{cc}^{\parallel}

In this section we state and prove a theorem which lowerbounds the Π_{cc}^{\parallel} of a given DAG G in terms of its depth robustness.

Theorem 4. *Let G be an (e, d) -depth-robust DAG, then $\Pi_{cc}^{\parallel}(G) > ed$.*

Proof. Let (P_1, \dots, P_m) be a parallel pebbling of minimum complexity, i.e., $\sum_{i=1}^m |P_i| = \Pi_{cc}^{\parallel}(G)$. For any d , we'll show that there exists a set B of size $|B| \leq \Pi_{cc}^{\parallel}(G)/d$ such that there's no path of length d in $G - B$, or equivalently, G is not $(\Pi_{cc}^{\parallel}(G)/d, d)$ -depth-robust, note that this implies the theorem.

For $i \in [d]$ define $B_i = P_i \cup P_{i+d} \cup P_{i+2d} \dots$. We observe that by construction $\sum_{i=0}^{d-1} |B_i| \leq \sum_{i=1}^m |P_i| = \Pi_{cc}^{\parallel}(G)$, so the size of the B_i 's is $\leq \Pi_{cc}^{\parallel}(G)/d$ on average, and the smallest B_i has size at most this. Let B be the smallest B_i , as just outlined $|B| \leq \Pi_{cc}^{\parallel}(G)/d$.

It remains to show that $G - B$ has no path of length d . For this consider any path v_1, \dots, v_d of length d in G . Let j be minimal such that $v_d \in P_j$ (so v_d is pebbled for the first time in round j of the pebbling). It then must be the case that $v_{d-1} \in P_{j-1}$ (as to pebble v_d in round j there must have been a pebble on v_{d-1} in round $j-1$). In round $j-2$ either the pebble on v_{d-1} was already there, or there was a pebble on v_{d-2} . This argument shows that each of the pebbling configurations $\{P_{j-d+1}, \dots, P_j\}$ must contain at least one node from v_1, \dots, v_d . As B contains each d th pebbling configuration, B contains at least one of these pebbling configurations $\{P_{j-d+1}, \dots, P_j\}$. Specifically we can find $j-d+1 \leq k \leq j$ s.t $P_k \subseteq B$, thus the path v_1, \dots, v_d is not contained entirely in $G - B$. \square

An immediate implication of Theorem 4 and Theorem 3 is that there is an infinite family of DAGs with maximal $\Pi_{cc}^{\parallel}(G) = \Omega(n^2)$ whose indegree scales with $\log n$. Note that this means that allowing indegree as small as $O(\log(n))$ is sufficient to get DAGs whose Π_{cc}^{\parallel} is within a constant factor of the n^2 upper bound on Π_{cc}^{\parallel} for any n node DAG. In the next section we will show how to reduce the indegree to $O(1)$ while only reducing $\Pi_{cc}^{\parallel}(G)$ by a factor of $O(\log(n))$.

Corollary 1 (of Theorem 4 and Theorem 3). *For some constants $c_1, c_2 > 0$ there exists an infinite family of DAGs $\{G_{n,\delta} \in \mathbb{G}_{n,\delta}\}_{n=1}^{\infty}$ with $\delta \leq c_1 \log(n)$ and $\Pi_{cc}^{\parallel}(G) \geq c_2 n^2$. This is optimal in the sense that for any family $\{\delta_n \in [n]\}_{n=1}^{\infty}$ and $\{J_n \in \mathbb{G}_{n,\delta_n}\}_{n=1}^{\infty}$ it holds that $\Pi_{cc}^{\parallel}(J_n) \in O(n^2)$. Moreover if $\delta_n = o(\log(n)/\log \log(n))$ then $\Pi_{cc}^{\parallel}(J_n) = o(n^2) = o(\Pi_{cc}^{\parallel}(G_n))$.*

Corollary 2 lower bounds the cost of pebbling a target set T given a starting pebbling configuration S . In particular, if the ancestors of T in $G - S$ induce an (e, d) -depth-robust DAG then the pebbling cost is at least $\Pi_{cc}^{\parallel}(G - S, T) \geq ed$. We will use Corollary 2 to lower bound Π_{cc}^{\parallel} for iMHFs like Argon2i and SB.

Corollary 2 (of Theorem 4). *Given a DAG $G = (V, E)$ and subsets $S, T \subset V$ such that $S \cap T = \emptyset$ let $G' = G - (V \setminus \text{ancestors}_{G-S}(T))$. If G' is (e, d) -depth robust then the cost of pebbling $G - S$ with target set T is $\Pi_{cc}^{\parallel}(G - S, T) > ed$.*

Proof. Note that $\Pi_{cc}^{\parallel}(G - S, T) \geq \Pi_{cc}^{\parallel}(G')$ since we will need to pebble every node in the set $\text{ancestors}_{G-S}(T) = V(G')$ to reach the target set T in $G - S$. By Theorem 4, we have $\Pi_{cc}^{\parallel}(G') > ed$. \square

Corollary 3 states that it remains expensive to pebble any large enough set of remaining nodes in a depth-robust graph even if we are permitted to first remove an arbitrary node set of limited size. An application of Corollary 3 might involve analysing the cost of pebbling stacks of depth-robust graphs. For example if there are not enough pebbles on the graph at some point in time then there must be some layers with few pebbles. If we can then show that many of the nodes on those layers will eventually need to be (re)pebbled then we can use this lemma to show that the remaining pebbling cost incurred by these layers is large.

Corollary 3 (of Theorem 4). *Let DAG $G = (V, E)$ be (e, d) -depth-robust and let $S, T \subset V$ such that*

$$|S| \leq e \quad \text{and} \quad T \cap S = \emptyset.$$

Then the cost of pebbling $G - S$ with target set T is $\Pi_{cc}^{\parallel}(G - S, T) > (e - |S|)(d - |\text{ancestors}_{G-S}(T)|)$.

Proof. Let $G' = G - (V - \text{ancestors}_{G-S}(T))$ and observe that G' is, at minimum, $(e - |S|, d - |\text{ancestors}_{G-S}(T)|)$ -depth robust. By Corollary 2 we have $\Pi_{cc}^{\parallel}(G - S, T) \geq \Pi_{cc}^{\parallel}(G') > (e - |S|)(d - |\text{ancestors}_{G-S}(T)|)$. \square

We remark that Theorem 4 is a special case of Corollary 3 by setting $S = \emptyset$ letting $T = \text{sinks}(G)$. Recall that $\Pi_{cc}^{\parallel}(G)$ is the parallel pebbling of minimal cumulative cost when pebbling all sinks of G , this requires pebbling all nodes of G at least once.

4 Indegree Reduction: Constant Indegree with Maximal Π_{cc}^{\parallel}

In this section we use the result from the previous section to show a new, more efficient, degree-reduction lemma. We remark that Lemma 1 is similar to [?, Lemma 9] in that both reductions replace high indegree nodes v in G with a path. However, we stress two key differences between the two results. First, our focus is on reducing the indegree while preserving depth-robustness. By contrast, [?, Lemma 9] focuses directly on preserving Π_{cc}^{\parallel} . Second, we note that the guarantee of [?, Lemma 9] is weaker in that it yields a reduced indegree graph G' whose size grows by a factor of indeg ($n' \leq n \times \text{indeg}$) while Π_{cc}^{\parallel} can drop by a factor of indeg — [?, Lemma 9] shows that $\Pi_{cc}^{\parallel}(G') \geq \frac{\Pi_{cc}^{\parallel}(G)}{\text{indeg}-1}$.

By contrast, setting $\gamma = \text{indeg}$ in Lemma 1 yields a reduced indegree graph G' whose size grows by a factor of $2 \times \text{indeg}$ ($n' \leq 2n \times \text{indeg}$) and better depth-robustness $(e', d') = (e, d \times \text{indeg})$. In particular, when we apply Theorem 4 the lower-bound $\Pi_{cc}^{\parallel}(G') \geq ed \times \text{indeg}$ improves by a factor of indeg when compared with the original graph G .

Lemma 1. *Let G be a (e, d) -depth-robust DAG. For $\gamma \in \mathbb{Z}_{\geq 0}$ there exists a $(e, d\gamma)$ -depth-robust DAG G' with*

$$\text{size}(G') \leq (\text{indeg}(G) + \gamma) \cdot \text{size}(G), \quad \text{indeg}(G') = 2 \quad \text{and} \quad \Pi_{st}(G') \leq \frac{\text{size}(G')^2}{\gamma}.$$

Proof. Fix a $\gamma \in \mathbb{Z}_{\geq 0}$ and let $\delta = \text{indeg}(G)$. We identify each node in V' with an element of the set $V \times [\delta + \gamma]$ and we write $\langle v, j \rangle \in V'$. For every node $v \in V$ with $\alpha_v := \text{indeg}(v) \in [0, \delta]$ we add the path $p_v = (\langle v, 1 \rangle, \langle v, 2 \rangle, \dots, \langle v, \alpha_v + \gamma \rangle)$ of length $\alpha_v + \gamma$. We call v the *genesis node* and p_v its *metanode*. In particular $V' = \cup_{v \in V} p_v$. Thus G has size at most $(\delta + \gamma)n$.

Next we add the remaining edges. Intuitively, for the i^{th} incoming edge (u, v) of v we add an edge to G' connecting the end of the metanode of u to the i^{th} node in the metanode of v . More precisely, for every $v \in V$, $i \in [\text{indeg}(v)]$ and edge $(u_i, v) \in E$ we add edge $(\langle u_i, \text{indeg}(u_i) + \gamma \rangle, \langle v, i \rangle)$ to E' . It follows immediately that G' has indegree (at most) 2.

Fix any node set $S \subset V'$ of size $|S| \leq e$. Then at most e metanodes can share a node with S . For each such metanode remove its genesis node in G . As G is (e, d) -depth-robust we are still left with a path p of length (at least) d in G . But that means that after removing S from G' there must remain a path p' in G' running through all the metanodes of p and $|p'| \geq |p|\gamma \geq d\gamma$. In other words G' is $(e, d\gamma)$ -depth-robust.

To see that $\Pi_{st}(G') \leq \text{size}(G')^2/\gamma$ we simply pebble G' in topological order. We note that we never need to keep more than one pebble on any metanode $p_v = (\langle v, 1 \rangle, \langle v, 2 \rangle, \dots, \langle v, \alpha_v + \gamma \rangle)$ with $\alpha_v = \text{indeg}(v)$. Once we pebble the last node $\langle v, \alpha_v + \gamma \rangle$ we can permanently discard any pebbles on the rest of p_v since $\langle v, \alpha_v + \gamma \rangle$ is the only node with outgoing edges. \square

Proof of Theorem 1. Theorem 1 follows by applying Lemma 1 to the family from Theorem 3 with $\gamma = \text{indeg} = \log n$. We get that for some fixed constants $c_1, c_2 > 0$ there exists an infinite family of indegree 2 DAGs $\{G_n \in \mathbb{G}_{n,2}\}_{n=1}^{\infty}$ where G_n is $(c_1 n / \log n, c_2 n)$ -depth robust and $\Pi_{st}(G_n) \leq O(n^2 / \log(n))$. By Theorem 4 then $\Pi_{cc}^{\parallel}(G_n) > (c_1 c_2) n^2 / \log(n)$, which is basically optimal for constant indegree DAGs by Theorem 2. \square

5 Security Proofs of Candidate iMHFs

On the surface, in this and the next section we give both security proofs and nearly optimal attacks for several of the most prominent iMHF proposals. That is we show both lower and (relatively tight) upperbounds on their asymptotic memory-hardness in the PROM. However, more conceptually, we also introduce

two new proof techniques for analysing the depth-robustness of DAGs as well as a new very memory-efficient class of algorithms for pebbling a DAG improving on the techniques used in [?]. Indeed for all candidates considered the attack in the next section is almost optimal in light of the accompanying security proofs in this section.

More specifically, in the first subsection we prove bounds for a class of random graphs which generalize the Argon2i-A construction [?] and the Single Buffer (SB) variant of Balloon Hashing [?]. To prove the lowerbound we use a simple and clean new technique for bounding the depth-robustness of a random DAG. In particular, we show that a random DAG is almost certainly $(e, \tilde{\Omega}(n^2/e^2))$ -depth robust for any $e > \sqrt{n}$. Combined with Theorem 4 we could immediately obtain a lower bound of $\tilde{\Omega}(n^{1.5})$. We can improve the lower bound to $\tilde{\Omega}(n^{5/3})$ by introducing a stronger notion of depth-robustness that we call block depth-robustness.

In the second subsection we prove bounds for a family of layered graphs which generalize both of the Catena constructions [?] as well as Linear (Lin) and Double Buffer (DB) variants of Balloon Hashing [?]. In particular, we introduce a new technique for proving lowerbounds on the cumulative pebbling complexity of a graph without going through the notion of depth-robustness. For example the (single layer) version of the Catena Dragonfly graph has the worst possible depth-robustness of any graph of linear depth. This shows that (in the lower but still non-trivial regimes of) cumulative complexity alternative combinatorial structures exist besides depth-robustness that can also confer some degree of pebbling complexity.

5.1 Lowerbounding the CC of Random DAGs.

We begin by defining a (n, δ, w) -random DAG, the underlying DAGs upon which Argon2i-A and SB are based. The memory window parameter w specifies the intended memory usage and throughput of the iMHF — the cost of the naïve pebbling algorithm is $\Pi_{cc}^{\parallel}(\mathcal{N}) = wn$. In particular, a t -pass Argon2i-A iMHF is based on a $(n, 2, n/t)$ -random DAG. Similarly, a t -pass Single-Buffer (SB) iMHF [?] is based on a $(n, 20, n/t)$ -random DAG. In this section we focus on the $t = 1$ -pass variants of the Argon2i-A and [?] iMHFs.

Definition 4 ((n, δ, w) -random DAG). *Let $n \in \mathbb{N}$, $1 < \delta < n$, and $1 \leq w \leq n$ such that w divides n . An (n, δ, w) -random DAG is a randomly generated directed acyclic (multi)graph with n nodes v_1, \dots, v_n (which we identify with the set $[n]$ according to their topological order) and with maximum in-degree δ for each node. The graph has directed edges (v_i, v_{i+1}) for $1 \leq i < n$ and random forward edges $(v_{r(i,1)}, v_i), \dots, (v_{r(i,\delta-1)}, v_i)$ for each node v_i . Here, $r(i, j)$ is independently chosen uniformly at random from the set $[\max\{0, i - w\}, i - 1]$.*

Theorem 5 states that for a (n, δ, n) -random DAG G such as Argon2i-A or SB we almost certainly have $\Pi_{cc}^{\parallel}(G) = \tilde{\Omega}(n^{5/3})$.

Theorem 5. *Let G be a (n, δ, n) -random DAG then, except with probability $o(n^{-7})$, we have*

$$\Pi_{cc}^{\parallel}(G) = \tilde{\Omega}\left(n^{5/3}\right).$$

Security Lower Bound. To prove the lower bound we rely on a slightly stricter notion of 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 and given a set $S \subseteq V(G)$ 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 $S \subseteq V(G)$ of size $|S| \leq e$ we have $\text{depth}(G - N(S, b)) \geq d$. Notice that when $b = 1$ (e, d, b) -block-depth robustness is equivalent to (e, d) -depth-robustness. However, when $b > 1$ (e, d, b) -block-depth robustness is a strictly stronger notion since the set $N(S, b)$ may have size as large as $|N(S, b)| = eb$.⁸

The proof of Theorem 5 relies on Lemma 2, which states that for any $e \geq \sqrt{n}$, with high probability, a $(n, 2, n)$ -random DAG G will be (e, d, b) -block depth-robust with $d = \frac{n^2}{e^2 \text{polylog}(n)}$ and $b = n/(20e)$. By contrast Lemma 9 states that G will be (e, d) -reducible with $d = \tilde{O}\left(n^2/e^2\right)$.

Lemma 2. *For any $e \geq \sqrt{n}$ any any $\delta \geq 2$ a (n, δ, n) -random DAG will be $\left(e, \Omega\left(\frac{n^2}{e^2 \log(n)}\right), \frac{n}{20e}\right)$ -block depth robust except with negligible probability in n .*

Setting $e = \sqrt{n}$ in Lemma 2 and applying Theorem 4 already implies that $\Pi_{cc}^{\parallel}(G) = \tilde{\Omega}\left(n^{1.5}\right)$. To obtain the stronger bound in Theorem 5 we rely on Corollary 2 combined with a more sophisticated argument exploiting block depth-robustness.

In more detail, let G be an (n, δ, n) -random DAG and let t_j denote the first time we place a pebble on node j . Observe that, since G contains all edges of the form $(j, j + 1)$ it must be that $t_{j+i} - t_j \geq i$ in any legal pebbling of G . We will show that for any $j > n/2$ a legal pebbling must (almost certainly) incur a cost of $\tilde{\Omega}\left(n^{4/3}\right)$ between pebbling steps t_j and t_{j+2k} where $k = \tilde{\Theta}\left(n^{2/3}\right)$. That is $\sum_{t=t_j}^{t_{j+2k}} |P_t| = \tilde{\Omega}\left(n^{4/3}\right)$ for any legal pebbling of G . Thus, $\sum_{t=t_{n/2+1}}^t |P_t| = \tilde{\Omega}\left(n^{4/3} \frac{n/2}{k}\right) = \tilde{\Omega}\left(n^{5/3}\right)$. In the remaining discussion we set $e = \tilde{\Omega}\left(n^{2/3}\right)$, $d = \tilde{\Omega}\left(n^{2/3}\right)$ and $b = \tilde{\Omega}\left(n^{1/3}\right)$.

To show that $\sum_{t=t_j}^{t_{j+2k}} |P_t| = \tilde{\Omega}\left(n^{4/3}\right)$ we consider two cases: we either have $|P_t| \geq e/2 = \tilde{\Omega}\left(n^{2/3}\right)$ pebbles on the DAG during each round $t_j \leq t \leq t_{j+k}$, or we do not. In the first case we trivially have $\sum_{t=t_j}^{t_{j+2k}} |P_t| \geq ke/2 = \tilde{\Omega}\left(n^{4/3}\right)$.

The second case is the trickier one to handle. To address it we essentially show that if, at some moment t' , few pebbles are left on G then between t' and t_{j+2k} it must be that (in particular) a depth-robust sub-graph of G was pebbled which we know requires a high pebbling cost. In more detail, suppose at some moment $t' \in [t_j, t_{j+k}]$ only $|P_{t'}| < e/2$ pebbles remain on G . Then we consider the

⁸ In particular, $(e, d, b \geq 1)$ -block depth robustness implies (e, d) -depth robustness. However, (e, d) -depth robustness only implies $(e/b, d, b)$ -block depth robustness.

sub-graph H induced by the node set $\text{ancestors}_{G_1 - N(P_{t'}, b)}([j + k + 1, j + 2k])$. We observe that, on the one hand, H must be fully pebbled during the interval $[t', t_{j+2k}]$. On the other hand, we observe that $G_1 = G - \{n/2 + 1, \dots, n\}$ is a $(n/2, \delta, n/2)$ -random DAG and, hence, by Lemma 2, G_1 is (almost certainly) (e, d, b) -block depth robust with $e = \tilde{\Omega}(n^{2/3})$, $d = \Omega\left(\frac{n^{2/3}}{\log(n)}\right)$ and $b = \tilde{\Omega}(n^{1/3})$. By exploiting the block depth robustness of G_1 we can show that H must itself be $(\tilde{\Omega}(n^{2/3}), \tilde{\Omega}(n^{2/3}))$ -depth robust. But then by Corollary 2 we get that H has cumulative complexity $\tilde{\Omega}(n^{4/3})$ and so have

$$\sum_{t=t_{j+k+1}}^{t_{j+2k}} |P_t| \geq \Pi_{cc}^{\parallel}(G_1 - P_{t'}, [j + k + 1, j + 2k]) \geq \tilde{\Omega}(n^{4/3}).$$

The proofs of Lemma 2 and Theorem 5 can be found in the full version. We now make a couple of observations about Lemma 2 and Theorem 5.

1. The lower bounds from Lemma 2 and Theorem 5 also apply to Argon2i-B. An Argon2i-B DAG G is similar to an (n, δ, n) -random DAG except that the randomly chosen forward edge $(r(i), i)$ for each node i is not chosen from the uniform distribution. However, these edges are still chosen independently and for each pair $j < i$ we still have $\Pr[r(i) = j] = \Omega(1/i)$. These are the only properties we used in the proofs of Lemma 2 and Theorem 5. Thus, essentially the same analysis shows that (whp) an Argon2i-B DAG G is $(e, \Omega(n^2/e^2), \frac{n}{20e})$ -block depth robust and that $\Pi_{cc}^{\parallel}(G) = \tilde{\Omega}(n^{5/3})$.
2. The lower bound from Lemma 2 is tight up to polylogarithmic factors. In particular, a generalization of an argument of Alwen and Blocki [?] shows that a (n, δ, n) -random DAG is $(e, \tilde{\Omega}\left(\frac{n^2}{e^2}\right))$ -reducible — see Lemma 9. However, this particular upper bound does not extend to Argon2i-B.
3. The lower bound from Theorem 5 might be tight. Alwen and Blocki [?] gave an attack \mathcal{A} such that $\Pi_{cc}^{\parallel}(\mathcal{A}) = O(n^{1.75}\delta \log n)$ for a (n, δ, t) -random DAG. In the following section we reduce the gap of $\tilde{O}(n^{1/12})$ further by developing an improved recursive version of the attack of Alwen and Blocki [?]. In particular, we show that for any $\epsilon > 0$ we have $\Pi_{cc}^{\parallel}(\mathcal{A}) = o(n^{1+\sqrt{1/2+\epsilon}}) = o(n^{1.708})$. Our modified attack also improves the upper bound for other iMHF candidates like Catena [?].
4. Theorem 4 alone will not yield any meaningful lower bounds on the Π_{cc}^{\parallel} of the Catena iMHFs [?]. In particular, the results from Alwen and Blocki [?] imply that for any t -pass variant of Catena the corresponding DAG is (e, d) -reducible for $ed \geq nt$ (typically, $t = O(\text{polylog}(n))$). However, in the remainder of this section, we use an alternative techniques to prove that $\Pi_{cc}^{\parallel}(G) = \Omega(n^{1.5})$ for the both Catena iMHFs and the Linear and DB iMHFs of [?].

5.2 Lowerbounding Dispersed Graphs

In this section we define dispersed graphs and prove a lowerbound on their CC. Next we show that several of the iMHF constructions from the literature are based on such graphs. Thus we obtain proofs of security for each of these constructions (albeit for limited levels of security). In the subsequent section we give an upperbound on the CC of these constructions showing that the lowerbounds in this section are relatively tight.

Generic Dispersed Graphs. Intuitively a (g, k) -dispersed DAG is a DAG ending with a path ϕ of length k which has widely dispersed dependencies. The following definitions make this concept precise.

Definition 5 (Dependencies). Let $G = (V, E)$ be a DAG and $L \subseteq V$. We say that L has a (z, g) -dependency if there exist node disjoint paths p_1, \dots, p_z each ending in L and with length (at least) g .

We are interested in graphs with long paths with many sets of such dependencies.

Definition 6 (Dispersed Graph). Let $g \leq k$ be positive integers. A DAG G is called (g, k) -dispersed if there exists a topological ordering of its nodes such that the following holds. Let $[k]$ denote the final k nodes in the ordering of G and let $L_j = [jg, (j+1)g - 1]$ be the j^{th} subinterval. Then $\forall j \in \llbracket k/g \rrbracket$ the interval L_j has a (g, g) -dependency.

More generally, let $\epsilon \in (0, 1]$. If each interval L_j only has an $(\epsilon g, g)$ -dependency then G is called (ϵ, g, k) -dispersed.

We show that many graphs in the literature consist of a *stack* of dispersed graphs. Our lowerbound on the CC of a dispersed graph grows in the height of this stack. The next definition precisely captures such stacks.

Definition 7 (Stacked Dispersed Graphs). A DAG $G = (V, E)$ is called $(\lambda, \epsilon, g, k)$ -dispersed if there exist $\lambda \in \mathbb{N}^+$ disjoint subsets of nodes $\{L_i \subseteq V\}$, each of size k with following two properties.

1. For each L_i there is a path running through all nodes of L_i .
2. Fix any topological ordering of G . For each $i \in [\lambda]$ let G_i be the sub-graph of G containing all nodes of G up to the last node of L_i . Then G_i is an (ϵ, g, k) -dispersed graph.

We denote the set of $(\lambda, \epsilon, g, k)$ -dispersed graphs by $\mathbb{D}_{\epsilon, g}^{\lambda, k}$.

We are now ready to state and prove the lowerbound on the CC of stacks of dispersed graphs.

Theorem 6.

$$G \in \mathbb{D}_{\epsilon, g}^{\lambda, k} \Rightarrow \Pi_{cc}^{\parallel}(G) \geq \epsilon \lambda g \left(\frac{k}{2} - g \right).$$

Intuitively we sum the CC of pebbling the last k nodes L_i of each sub-graph G_i . For this we consider any adjacent intervals A of $2g$ nodes in L_i . Let p be a path in the $(\epsilon g, g)$ -dependency of the second half of A . Either at least one pebble is always kept on p while pebbling the first half of A (which takes time at least g since a path runs through L_i) or p must be fully pebbled in order to finish pebbling interval A (which also takes time at least g). Either way pebbling A requires an additional CC of g per path in the $(\epsilon g, g)$ -dependency of the second half of A . Since there are $k/2g$ such interval pairs each with ϵg incoming paths in their dependencies we get a total cost for that layer of $k g \epsilon / 2$. So the cost for all layer of G is at least $\lambda k g \epsilon / 2$. The details (for the more general case when g doesn't divide n) can be found in the full version.

The Graphs of iMHFs. We apply Theorem 6 to some important iMHFs from the literature. For this we first describe the particular DAGs (or at least their salient properties) underlying the iMHF candidates for which we prove lowerbounds in this section. Then we state a theorem summarizing our lowerbounds for these graphs. Finally we prove the theorem via a sequence of lemma; one per iMHF being considered.

Catena Dragonfly. We begin with the Catena Dragonfly graph. We briefly recall the properties of the DFG_λ^n construction, relevant to our proof, summarized in the following lemma which follows easily from the definition of DFG_λ^n in [?, Def. 8 & 9].

For this we describe the “bit-reversal” function (from which the underlying bit-reversal graph derives its name). Let $k \in \mathbb{N}^+$ such that $c = \log_2 k$ is an integer. On input $x \in [k]$ the *bit-reversal function* $\mathbf{br}(\cdot) : [k] \rightarrow [k]$ returns $y + 1$ such that the binary representation of $x - 1$ using c bits is the reverse of the binary representation of y using c bits.

Lemma 3 (Catena Dragonfly). *Let $\lambda, n \in \mathbb{N}^+$ be such that $k = n/(\lambda + 1)$ is a power of 2. Let $G = \text{DFG}_\lambda^n$ be the Catena Bit Reversal graph. Then the following holds:*

1. G has n nodes.
2. Number them in topological order with the set $[n]$ and $\forall i \in [0, \lambda]$ let node set $L_i = [1 + ik, (i + 1)k]$. A path runs through all nodes in each set L_i .
3. Node $ki + x \in L_i$ has an incoming edge from $k(i - 1) + \mathbf{br}(x) \in L_{i-1}$.

Catena Butterfly. Next describe the graph underlying the Catena Butterfly graph. We summarize its key properties relevant to our proof in the following lemma (which follows immediately by inspection of the Catena Butterfly definition [?, Def. 10 & 11]).

Lemma 4 (Catena Butterfly Graph). *Let $\lambda, n \in \mathbb{N}^+$ such that $n = \bar{n}(\lambda(2c - 1) + 1)$ where $\bar{n} = 2^c$ for some $c \in \mathbb{N}^+$. Then the Catena Butterfly Graph BFG_λ^n consists of a stack of λ sub-graphs such that the following holds.*

1. The graph BFG_λ^n has n nodes in total.

2. The graph BFG_λ^n is built as a stack of λ sub-graphs $\{G_i\}_{i \in [\lambda]}$ each of which is a superconcentrator⁹. In the unique topological ordering of BFG_λ^n denote the first and final \bar{n} nodes of each G_i as $L_{i,0}$ and $L_{i,1}$ respectively. Then there is a path running through all nodes in each $L_{i,1}$.
3. Moreover, for any $i \in [\lambda]$ and subsets $S \subset L_{i,0}$ and $T \subset L_{i,1}$ with $|S| = |T| = h \leq \bar{n}$ there exist h node disjoint paths p_1, \dots, p_h of length $2c$ from S to T .

Balloon Hashing Linear. Finally we describe the graph underlying both the Linear and DB construction [?]. The graph $G = \text{Lin}_\tau^\sigma$ is a pseudo-randomly constructed τ -layered graph with $\text{indeg}(G) = 21$. It is defined as follows:

- $G = (V, E)$ has $n = \sigma\tau$ nodes $V = [n]$, and G contains a path $1, 2, \dots, n$ running through V .
- For $i \in [0, \tau - 1]$ let $L_i = [i\sigma + 1, (i + 1)\sigma]$ denote the i 'th layer. For each node $x \in L_i$, with $i > 0$, we select 20 nodes $y_1, \dots, y_{20} \in L_{i-1}$ (uniformly at random) and add the directed edges $(y_1, x), \dots, (y_{20}, x)$ to E .

5.3 The Lowerbounds.

Now that we have our lowerbound for stacks of dispersed graphs it remains to analyse for which parameters each of the above three graphs can be viewed as being dispersed graphs. The results of this analysis are summarized in the theorem bellow.

Theorem 7. [*iMHF Constructions Based on Dispersed Graphs*]

- If $\lambda, n \in \mathbb{N}^+$ such that $n = \bar{n}(\lambda(2c - 1) + 1)$ where $\bar{n} = 2^c$ for some $c \in \mathbb{N}^+$ then it holds that

$$\text{BFG}_\lambda^n \in \mathbb{G}_{n,3} \quad \text{BFG}_\lambda^n \in \mathbb{D}_{1, \lceil \sqrt{\bar{n}} \rceil}^{\lambda, \bar{n}} \quad \Pi_{cc}^\parallel(\text{BFG}_\lambda^n) = \Omega\left(\frac{n^{1.5}}{c\sqrt{c\lambda}}\right).$$

- If $\lambda, n \in \mathbb{N}^+$ such that $k = n/(\lambda + 1)$ is a power of 2 then it holds that

$$\text{DFG}_\lambda^n \in \mathbb{G}_{n,2} \quad \text{DFG}_\lambda^n \in \mathbb{D}_{1, \lceil \sqrt{k} \rceil}^{\lambda, k} \quad \Pi_{cc}^\parallel(\text{DFG}_\lambda^n) = \Omega\left(\frac{n^{1.5}}{\sqrt{\lambda}}\right).$$

- If $\sigma, \tau \in \mathbb{N}^+$ such that $n = \sigma * \tau$ then with high probability it holds that

$$\text{Lin}_\tau^\sigma \in \mathbb{G}_{n,21} \quad \text{Lin}_\tau^\sigma \in \mathbb{D}_{0.25, \sqrt{\sigma}/2}^{\tau-1, \sigma} \quad \Pi_{cc}^\parallel(\text{Lin}_\tau^\sigma) = \Omega\left(\frac{n^{1.5}}{\sqrt{\tau}}\right).$$

The theorem is proven in the following three lemma bellow (one lemma per graph). We begin with the graph for Catena Dragonfly.

⁹ A superconcentrator is a DAG with m inputs and outputs such that any subset of $s \in [m]$ inputs and outputs are connected by s node disjoint paths.

Lemma 5. *It holds that $\text{DFG}_\lambda^n \in \mathbb{D}_{1, \sqrt{k}}^{\lambda, k}$ where $k = \frac{n}{(\lambda+1)}$ and $\Pi_{cc}^\parallel(\text{DFG}_\lambda^n) = \Omega\left(\frac{n^{1.5}}{\sqrt{\lambda}}\right)$.*

Proof of Lemma 5. Let $G = \text{DFG}_\lambda^n$ and set $k = n/(\lambda+1)$, $c = \log_2 k$ and $g = \sqrt{k}$. By construction c is an integer. For simplicity assume c is even and so $g \in \mathbb{N}^+$.¹⁰ Number the nodes of G according to (the unique) topological order with the set $[0, n-1]$. It suffices to show that for all $i \in [\lambda]$ the sub-graph G_i consisting of nodes $[(i+1)k-1]$ is (g, k) -dispersed (with probability $\epsilon = 1$). If this holds then Theorem 6 immediately implies that $\Pi_{cc}^\parallel(\text{DFG}_\lambda^n) = \Omega\left(\frac{n^{1.5}}{\sqrt{\lambda}}\right)$.

Recall that G consists of layers L_i of length k . For each $j \in [k/2g]$ let $L_{i,j}$ be the j^{th} interval of $2g$ nodes of L_i . Let $R_{i,j}$ be the second half of $L_{i,j}$. We will show that there are g node-disjoint paths each terminating in $R_{i,j}$ whose remaining nodes are all in layer L_{i-1} . Let node set $S_x = [s+y-(g-2), s+y]$ where $y = \mathbf{br}(x)$ and $s = (i-1)k$. The next three properties follow immediately from Lemma 3 and they imply the lemma.

- $\forall x \in R$ it holds that $S_x \subset L_{i-1}$.
- $\forall x \in R$ there is a path of length g going through the nodes of S_x and ending in x .
- \forall distinct $x, x' \in R$ sets S_x and $S_{x'}$ are disjoint. □

Next we turn to the Catena Dragonfly graph.

Lemma 6. *Let $\lambda, n \in \mathbb{N}^+$ such that $n = \bar{n}(\lambda(2c-1) + 1)$ with $\bar{n} = 2^c$ for some $c \in \mathbb{N}^+$. It holds that $\text{BFG}_\lambda^n \in \mathbb{D}_{1, g}^{\lambda, \bar{n}}$ for $g = \lceil \sqrt{\bar{n}} \rceil$ and $\Pi_{cc}^\parallel(\text{BFG}_\lambda^n) = O\left(\frac{n^{1.5}}{c\sqrt{c\lambda}}\right)$.*

Proof of Lemma 6. Let $G = \text{BFG}_\lambda^n$ and let $G_1, G_2, \dots, G_\lambda$ be the sub-graphs of G described in Lemma 4. We will show that each G_i is (g, \bar{n}) -dispersed for $g = \lceil \sqrt{\bar{n}} \rceil$. Fix arbitrary $i \in [\lambda]$ and L_1 be the last \bar{n} nodes in the (the unique) topological ordering of G_i . We identify the nodes in L_1 with the set $\{1\} \times [\bar{n}]$ such that the second component follows their topological ordering. Let $\bar{g} = \lfloor \bar{n}/g \rfloor$ and for each $j \in [\bar{g}]$ let $L_{1,j} = \{(1, jg+x) : x \in [0, g-1]\}$. We will show that $L_{1,j}$ has a (g, g) -dependency.

Let L_0 be the first \bar{n} nodes of G_i which we identify with the set $\{0\} \times [\bar{n}]$ (again with the second component respecting their topological ordering). Notice that for $n > 1$ and $g = \lceil \sqrt{\bar{n}} \rceil$ it holds that $g(g-2c+1) \leq n$. Thus the set $S = \{(0, i(g-2c+1)) : i \in [g]\}$ is fully contained in L_0 . Property (3) of Lemma 4 implies there exist g node disjoint paths from S to $L_{1,j}$ of length $2c$. In particular $L_{1,j}$ has a $(g, 2c)$ -dependency.

We extend this to a (g, g) -dependency. Let path p , beginning at node $\langle 0, v \rangle \in S$, be a path in the $(g, 2c)$ -dependency of $L_{1,j}$. Prepend to p the path traversing

$$\langle \langle 0, v - (g - 2c - 1) \rangle, \langle 0, v - (g - 2c - 2) \rangle, \dots, \langle 0, v \rangle \rangle$$

¹⁰ The odd case is identical but with messy but inconsequential rounding terms.

to obtain a new path p^+ of length g . As this is a subinterval of L_0 property (2) of Lemma 4 implies this prefix path always exists. Moreover since any paths $p \neq q$ in a $(g, 2c)$ -dependency of $L_{1,i}$ are node disjoint they must, in particular, also begin at distinct nodes $\langle 0, v_p \rangle \neq \langle 0, v_q \rangle$ in S . But by construction of S any such pair of nodes is separated by $g - 2c$ nodes. In particular paths p^+ and q^+ are also node disjoint and so by extending all paths in a $(2c, g)$ -dependency we obtain a (g, g) -dependency for $L_{1,i}$. This concludes the first part of the lemma.

It remains to lowerbound $\Pi_{cc}^{\parallel}(\text{BFG}_{\lambda}^n)$ using Theorem 6.

$$\begin{aligned} \Pi_{cc}^{\parallel}(\text{BFG}_{\lambda}^n) &\geq \lambda g \left(\frac{k}{2} - g \right) \geq \lambda \left\lfloor \sqrt{\bar{n}} \right\rfloor \left(\frac{\bar{n}}{2} - \left\lfloor \sqrt{\bar{n}} \right\rfloor \right) \\ &= \lambda \sqrt{\bar{n}} \left(\frac{\bar{n}}{2} - \sqrt{\bar{n}} \right) - O(\bar{n}) = \Omega(\lambda \bar{n}^{1.5}) \\ &= \Omega\left(\frac{n^{1.5}}{c\sqrt{c\lambda}} \right). \end{aligned}$$

□

Finally we prove a lowerbound for the Linear and DB variants of Balloon Hashing.

Lemma 7. *If $\sigma, \tau \in \mathbb{N}^+$ such that $n = \sigma\tau$ then with high probability it holds that*

$$\text{Lin}_{\tau}^{\sigma} \in \mathbb{G}_{n,21} \quad \text{Lin}_{\tau}^{\sigma} \in \mathbb{D}_{0.25, \sqrt{\sigma}/2}^{\tau-1, \sigma} \quad \Pi_{cc}^{\parallel}(\text{Lin}_{\tau}^{\sigma}) = \Omega\left(\frac{n^{1.5}}{\sqrt{\tau}} \right).$$

Proof. (sketch) It suffices to show that $\text{Lin}_{\tau}^{\sigma} \in \mathbb{D}_{0.25, \sqrt{\sigma}/2}^{\tau-1, \sigma}$. By Theorem 6 it immediately follows that

$$\Pi_{cc}^{\parallel}(\text{Lin}_{\tau}^{\sigma}) \geq \frac{(\tau-1)\sqrt{\sigma}/2}{4} (\sigma/2 - \sqrt{\sigma}/2) = \Omega\left(\frac{n^{1.5}}{\sqrt{\tau}} \right).$$

Fix any $i \in [0, \tau-1]$. Consider layer L_i and given set $S_x = [x, x + \sqrt{\sigma}/2 - 1] \subset L_i$ denoting an interval of $\sqrt{\sigma}/2$ nodes in L_i beginning at node x . Without loss of generality we suppose that each node in S_x only has one randomly chosen parent in L_{i-1} — adding additional edges can only improve dispersity. We partition L_{i-1} into $\sqrt{\sigma}$ intervals of length $\sqrt{\sigma}$. We say that an interval $[u, u + \sqrt{\sigma} - 1] \subset L_{i-1}$ is *covered by* S_x if there is exists edge (y, v) from the second half of the interval to a node in S_x ; that is if $y \in [u + \sqrt{\sigma}/2, u + \sqrt{\sigma} - 1]$ and $v \in S_x$. In this case the path $(u, u+1, \dots, y, v)$ has length $\geq \sqrt{\sigma}/2$ and this path will not intersect the corresponding paths from any of the other (disjoint) intervals in L_{i-1} (recall that we are assuming that $v \in S_x$ only has one parent in L_{i-1}). The probability that an interval $[u, u + \sqrt{\sigma} - 1] \subset L_{i-1}$ is covered by S_x is at least

$$1 - \left(1 - \frac{\sqrt{\sigma}/2}{\sigma} \right)^{\sqrt{\sigma}} \approx 1 - \sqrt{1/e}.$$

Thus, in expectation we will have at least $\mu = \sqrt{\sigma} \left(1 - \sqrt{1/e}\right) \geq 0.39 \times \sqrt{\sigma}$ node disjoint paths of length $\sqrt{\sigma}/2$ ending in S_x . Standard concentration bounds imply that we will have at least $\sqrt{\sigma}/4$ such paths with high probability. \square

6 New Memory-Efficient Evaluation Algorithm and Applications

In this section we introduce a new generic parametrized pebbling algorithm for DAGs (i.e. an evaluation algorithm for an arbitrary iMHF). We upperbound the pebbling strategy’s cumulative pebbling complexity in terms of its parameters. In particular we see that for graphs which are not depth-robust there exist parameter settings for which the algorithm results in low CC pebbling strategies. Next we instantiate the parameters to obtain attacks on the random graphs defined in the previous section. By “attack” we mean that, for Argon2i-A and SB, the algorithm has significantly less asymptotic memory-hardness in the PROM than both that of their naïve algorithms, and even that of the attack in [?].

Review of [?]. In order to describe the results in this section we first review the generic pebbling algorithm PGenPeb of [?] which produces a pebbling P_1, P_2, \dots, P_n of G as follows. PGenPeb takes as input a node set $S \subset V$ of size $|S| = e$ such that removing S reduces the depth of the DAG $\text{depth}(G - S) \leq d$. Intuitively, keeping pebbles on S compresses G in the sense that G can now quickly be entirely (re)pebbled within d (parallel) steps. This is because when S is already pebbled then no remaining unpebbled path has length greater than d . Algorithm PGenPeb never removes pebbles from nodes in S and its goal is to always pebble node i at time i so as to finish in $n = \text{size}(G)$ steps.¹¹ To ensure that parents of node i are all pebbled at time i algorithm PGenPeb sorts nodes in topological order and partitions them into consecutive intervals of g nodes (where $g \in [d, n]$ is another input parameter). Nodes in interval $I_c = [(c-1)g + 1, cg] \cap [n] \subset V$ are pebbled during “light phase” A_c which runs for g time steps. To ensure that the result is a legal pebbling, PGenPeb guarantees the following invariant \mathcal{I} : just before light phase A_c begins (i.e. at time $(c-1)g$) we have $X_{cg} = \text{parents}(I_c) \cap [(c-1)g] \subset P_{(c-1)g}$ so that we begin A_c with all of the necessary pebbles. Now, in phase A_c algorithm PGenPeb simply places a pebble on node $i \in I_c$ at time i .

Notice that for $c = 1$, $X_1 = \emptyset$ and so \mathcal{I} is trivially satisfied. Let $c > 1$. Partition X_{cg} into $X_{cg}^- = X_{cg} \cap [(c-1)g - d + 1]$ and $X_{cg}^+ = X_{cg} \setminus X_{cg}^-$. Since X_{cg}^+ is pebbled in the final d steps of light phase A_{c-1} , PGenPeb can simply not remove those until time step $(c-1)g$. In order to ensure that X_{cg}^- is pebbled at that time PGenPeb also runs a “balloon phase” B_{c-1} in parallel with the final d steps of A_{c-1} .¹² Intuitively in phase B_{c-1} all nodes in $[(c-1)g] \subseteq V$ are quickly “decompressed” by greedily re-pebbling everything possible in parallel. Recall that pebbles are never removed from nodes in S . So at time j all of $S \cap [j]$ is

¹¹ Formally, $i \in P_i$ and $S \cap [i] \subseteq P_i$ for each $i \leq n$.

¹² Recall that $g \geq d$ so A_{c-1} lasts long enough to accommodate B_{c-1} .

already pebbled. Therefore, at time $(c-1)g$, there is no unpebbled path longer than d nodes within the first $[(c-1)g]$ nodes and so B_{c-1} can indeed entirely (and legally) repebble those nodes (and so in particular X_{cg}^-). Thus, together with the nodes in X_{cg}^+ pebbled in the final d steps of \mathcal{A}_{c-1} it follows that \mathcal{I} also holds for \mathcal{A}_c .

The runtime of PGenPeb is n . Thus the cost is at most $\Pi_{cc}^{\parallel}(\text{PGenPeb}) \leq en + \delta gn + \lceil n/g \rceil (dn)$ where $\delta = \text{indeg}(G)$. The en term upper bounds the cost of always keeping pebbles on S , δgn bounds the cost of all light phases, and the third term upper bounds the cost of all balloon phases — each balloon phase costs at most dn and at most $\lceil n/g \rceil$ balloon phases are run.

Notice that (for constant δ) we would like to set $g \leq e$ so that the second term doesn't dominate the first. Conversely, to keep the number of (expensive) balloon phases at a minimum we also want g to be large. Therefore, as long as $e \geq d$, the asymptotically minimal complexity is obtained when $g = e$.

Recursive Attack: Intuition. Our new algorithm relies on the following key insight. Algorithm PGenPeb can actually pebble, not just the sink with the above complexity, but instead any target set $T \subseteq V$ simultaneously.¹³ This more general view allows us to recast the task of the balloon phase as such a pebbling problem. The graph being pebbled is $G' = G - (S \cup [(c-1)g - d + 1])$ and the target set is X_c^- . So instead of implementing balloon phases with an expensive greedy pebbling strategy as in PGenPeb we can apply the same strategy as (the generalized version of) PGenPeb recursively. This is the approach of the new algorithm RGenPeb. (The complete pseudocode of RGenPeb can be found in the full version. For this approach to work we need that not only is G (e, d) -reducible via some set S but that there is also a set S' of size $e' > e$ such that $\text{depth}(G - S') = d' < d$. Only when these conditions can no longer be met do we have to resort to greedy pebbling for the balloon phases. As we show below, it turns out that RGenPeb leads to improved attacks compared to PGenPeb for the DAGs underlying key iMHFs like Argon2i, Catena and Balloon Hashing.

Outline. The remainder of this section has the following structure. First, in Lemma 8 we generalize the results of [?] to upperbound the CC of a graph by the cost of pebbling all light phases plus the CC of the pebblings problems solved by balloon phases. Next we define a generalization of (e, d) -reducible graphs called f -reducible graphs; namely graphs which are $(f(d), d)$ -reducible for all $d \in [n]$. This allows us to state the main theorem of this section. It considers a certain class of functions f and upper bounds the complexity of RGenPeb on such f -reducible graphs using any number k levels of recursion. To apply the theorem to the iMHFs from the literature we prove Lemma 9 which describes the f -reducibility of their underlying DAGs. Thus we obtain the final corollary of the section describing new upperbounds on those iMHFs. At the end of this section we give a more detailed description of RGenPeb and the proof of Lemma 8.

¹³ For example, in the final d steps of the execution one last balloon phase can be run to (re)pebble all of G including T at no added cost to the asymptotic complexity.

Generalizing [?]. In order to derive the new pebbling strategy we first generalize the results of [?]. Given a DAG $G = (V, E)$, node set $T \subseteq V$ and integer t we define $\mathcal{P}_{G,T,t}^{\parallel} \subseteq \mathcal{P}_{G,T}^{\parallel}$ to be the set of all parallel pebbblings (P_1, \dots, P_z) of G such that $z \leq t$. Analogously we let $\Pi_{cc}^{\parallel}(G, T, t) = \min_{P \in \mathcal{P}_{G,T,t}^{\parallel}} \Pi_{cc}^{\parallel}(P)$.

We remark that if $\text{depth}(G) = d$ then $\Pi_{cc}^{\parallel}(G, T, d) \leq dn$ since we can greedily pebble G in topological order in time $\text{depth}(G)$. Lemma 8 provides an alternative upper bound on $\Pi_{cc}^{\parallel}(G, T, 2d)$.

Lemma 8. *Let $G = (V, E)$ be a DAG of size n , indegree δ and $\text{depth}(G) \leq d_0$. If G is (e_1, d_1) -reducible with parameters e_1, d_1 such that $2d_1n \leq e_1d_0$ and $d_1 \leq d_0$ then for any target set $T \subseteq V$ we have*

$$\Pi_{cc}^{\parallel}(G, T, 2d_0) \leq (4\delta + 4) e_1 d_0 + \frac{n}{e_1} \left(\max_{\substack{T' \subseteq V - S_1 \\ |T'| \leq \delta \cdot e_1}} \Pi_{cc}^{\parallel}(G - S_1, T', 2d_1) \right),$$

where $S_1 \subseteq V$ has size $|S_1| \leq e_1$ such that $\text{depth}(G - S_1) \leq d_1$.

To prove the lemma we first define the RGenPeb algorithm and argue the legality of the pebbling it produces at the end of this section. Armed with this, it remains only to upperbound the complexity of a call to RGenPeb in terms of the complexity of the recursive call it makes. This involves a relatively straightforward (but somewhat tedious) counting of the pebbles placed by RGenPeb, the details of which can be found in the full version.

We observe that Lemma 8 generalizes the main result of [?] as that work only considered the special case where balloon phases are implemented with a greedy pebbling strategy. The advantage of the above formulation (and the more general RGenPeb) is that now we can apply the lemma (and algorithm) recursively.

In order to apply this lemma repeatedly we will need graphs which are reducible for a sequence of points parameters (e, d) satisfying the conditions laid out in Lemma 8 relating consecutive parameters. To help characterize such graphs we generalize the notion of reducibility as follows.

Definition 8. *Let $G = (V, E)$ be a DAG with n nodes and let $f : \mathbb{N} \rightarrow \mathbb{N}$ be a function. We say that G is f -reducible if for every positive integer $n \geq d > 0$ there exists a set $S \subseteq V$ of $|S| = f(d)$ nodes such that $\text{depth}(G - S) \leq d$.*

Next we state the main theorem of this section which, for a certain class of natural functions f , upperbounds the CC of any f -reducible graph.

Theorem 8. *Let G be a f -reducible DAG on n nodes then if $f(d) = \tilde{O}\left(\frac{n}{d^b}\right)$ for some constant $0 < b \leq 2/3$ and let $a = \frac{1-2b+\sqrt{1+4b^2}}{2}$. Then for any constant $\epsilon > 0$*

$$\Pi_{cc}^{\parallel}(G) \leq O(n^{1+a+\epsilon}).$$

The proof of Theorem 8 can be found in the full version. We briefly sketch the intuition here. We define a sequence e_1, e_2, \dots and d_1, d_2, \dots such that G is (e_i, d_i) -reducible for each i , $e_i = n^{a_i + \epsilon/3}$ and $d_i = n^{\frac{1-a_i}{b}}$ with

$$a_{i+1} = 1 + \frac{(a-1)(1-a_i)}{b}, \text{ where } a_1 = a = \frac{1-2b + \sqrt{1+4b^2}}{2}.$$

If $b \leq a$ we have $e_{i+1}d_i \geq nd_{i+1}$ for every i so we can repeatedly invoke Lemma 9 as many times as we desire. By exploiting several key properties of the sequence $\{a_i\}_{i=1}^\infty$ we can show that unrolling the recurrence k times yields a pebbling with cost at most $k(4\delta + 2)n^{1+a+\epsilon/3} + n^{1+a+\epsilon/3}d_k$. For any $\epsilon > 0$ we can select the constant k sufficiently large that $d_k \leq n^{\epsilon/3}$. Thus, the pebbling cost is $o(n^{1+a+\epsilon})$.

Analysing Existing iMHFs. We can now turn to applying Theorem 8 to iMHFs from the literature. Lemma 9 below states that an (n, δ, n) -random DAGs and λ -layered DAGs are f -reducible. In particular these are the types of DAGs underlying all of the iMHFs considered in the previous section.

Lemma 9. *Let $f_b(d) = \tilde{O}\left(\frac{n}{d^b}\right)$ then*

1. *Let $\delta = O(\text{polylog}(n))$ then a (n, δ, n) -random DAG is $f_{0.5}$ -reducible with high probability.*
2. *The Catena DAGs DFG_λ^n and BFG_λ^n are both f_1 -reducible for $\lambda = O(\text{polylog}(n))$.*
3. *The Balloon Hashing Linear (and the DB) graph Lin_τ^σ is f_1 -reducible for $\tau = O(\text{polylog}(n))$.*

The proof generalizes the arguments used in [?] to first establish a particular pair (e, d) for which the graphs are reducible. It can be found in the full version.

Together with Theorem 8 and Lemma 9 we now obtain the main application of RGenPeb which is described in the following corollary upperbounding the memory-hardness of each of the considered iMHFs.

Corollary 4. *Let $\epsilon > 0$ be any constant*

1. *Let $\delta = O(\text{polylog}(n))$ then an (n, δ, n) -random DAG G has $\Pi_{cc}^\parallel(G) = O(n^{1+\sqrt{1/2+\epsilon}}) \approx O(n^{1.707+\epsilon})$.*
2. *Both $\Pi_{cc}^\parallel(\text{DFG}_\lambda^n)$ and $\Pi_{cc}^\parallel(\text{BFG}_\lambda^n)$ are in $\tilde{O}\left(n^{\frac{13}{8}}\right) = \tilde{O}(n^{1.625})$.*
3. *$\Pi_{cc}^\parallel(\text{Lin}_\tau^\sigma) = \tilde{O}\left(n^{\frac{13}{8}}\right) = \tilde{O}(n^{1.625})$, where Lin_τ^σ has $n = \tau\sigma$ nodes.*

We remark that Theorem 8 does not yield tighter bounds for Catena iMHFs DFG_λ^n or BFG_λ^n or for Lin_τ^σ . Each DAG is indeed f_b reducible for any $b \leq 2/3$ (even for $b \leq 1$), but for $b \leq 2/3$ it follows that $a = \frac{1-2b+\sqrt{1+4b^2}}{2} \geq 2/3$. Thus, Theorem 8 yields an attack with cost $O\left(n^{\frac{5}{3}+\epsilon}\right)$, which does not improve on the non-recursive PGenPeb attack in [?] as that has cost $O\left(n^{\frac{5}{3}}\right)$. However, we can set $e_1 = n^{5/8}, e_2 = n^{7/8}$ and exploit the fact that the DAGs are (e_i, d_i) -reducible with $d_i = \tilde{O}(n/e_i)$. Applying Lemma 8 twice we have $\Pi_{cc}^\parallel(G) =$

$$O\left(e_1 n + \frac{n}{e_1} \left(e_2 d_1 + \frac{n}{e_2} n d_2\right)\right) = O\left(n^{13/8} + n^{3/8+7/8} d_1 + n^{3/8+1/8+1} d_2\right) = \tilde{O}\left(n^{\frac{13}{8}}\right).$$

Note that $e_2 d_1 = \tilde{O}(n^{10/8}) > \tilde{O}(n^{9/8}) = n d_2$ so it is legal to invoke Lemma 8 for sufficiently large n .

The RGenPeb Algorithm. In the remainder of this section we sketch RGenPeb algorithm and justify that it produces a legal pebbling. The analysis of its complexity in terms of the complexity of its recursive call is contained in the proof of Lemma 8. The final complexity of an execution requires unravelling the recursive statement of Lemma 8 which is done in the proof of Theorem 8.

In the following we will ignore rounding errors here as they are inconsequential for the asymptotic behaviour while adding needless complexity to the exposition. For completeness we observe that if RGenPeb finishes a light phase and there is not enough steps left to complete a full light phase then it can simply run the next light phase as far as it can (and completely omits any further balloon phases). This affects neither the legality of the resulting pebbling nor its asymptotic complexity.

Algorithm RGenPeb takes input a DAG $G = (V, E)$, sets $S_1 \subseteq S_2 \subseteq \dots \subseteq S_k \subseteq V$, integers d_1, d_2, \dots, d_k and a target set T such that $\forall i \in [k] : e_i d_{i-1} \geq n d_i$ and $d_i \geq \text{depth}(G - S_i)$ where $n = |V|$, $e_i = |S_i|$ and $d_0 = \text{depth}(G)$. For this RGenPeb makes use of an arbitrary partition of the nodes of G into $2d_0$ into sets $D_1, D_2, \dots, D_{2d_0}$ such that the following properties hold:¹⁴

TOPOLOGICALLY ORDERED: $\forall j \in [2d_0 - 1] \text{ parents}(D_{j+1}) \subseteq \bigcup_{y \in [j]} D_y$,

MAXIMUM SIZE: $\forall j \leq 2d_0 \ |D_j| \leq \frac{n}{d_0}$.

Intuitively, the set D_j is the set of nodes that will be pebbled by a light phase in the j^{th} step. So for PGenPeb we would simply have $D_j = \{j\}$.

At the top level of the recursion RGenPeb looks relatively similar to PGenPeb. The goal is to pebble $G_0 = G$ with the target set $T_0 = \text{sinks}(G_0)$ in at most $2d_0$ steps which is done by executing a sequence of light phases lasting $m = e_2 d_0 / n$ steps and balloon phases lasting $2d_1$ steps. The requirement that $e_2 d_0 \geq 2d_1 n$ ensures that $m \geq 2d_1$ so that we can complete each balloon phase in time for the upcoming light phase. For $t \in [2d_0]$ let $U_t = \bigcup_{j \in [t]} D_j$ be all nodes pebbled by light phases up to step t . Then for $c \in [2d_0/m]$ the light phase A_c runs during time interval $I_c = [(c-1)m + 1, cm]$ during which it will pebble nodes $U_{cm} \setminus U_{(c-1)m}$. It never removes pebbles from S_1 and, at each time step t it keeps pebbles on $\text{parents}(U_{cm} \setminus U_t)$ as it will still need those to finish the light phase.

As for PGenPeb the light phase A_1 is trivially a legal pebbling. Let $X_{cm} = \text{parents}(U_{cm+m} \setminus U_{cm+1}) \cap U_{cm}$ and $X_{cm}^- = (X_{cm} \cap U_{cm-2d_1}) \setminus S_1$ and $X_{cm}^+ = X_{cm} \setminus X_{cm}^-$. To ensure that all pebbles placed by during light phase A_{c+1} are done so legally it suffices for RGenPeb to ensure that X_{cm} is fully pebbled at time cm . This is done by balloon phase running in parallel to the final $2d_1$ steps of A_c ; that is during the interval $[cm - 2d_1, cm]$. The pebbling for the balloon phase may be

¹⁴ For example we can sort the nodes in topological order and divided them up into the partition. Whenever a set is larger than n/d_0 we insert a new set into the partition with the overflow.

obtained by a recursive call to `RGenPeb` for the graph $G' = G - S - (V \setminus U_{cm-2d_1})$ (G' is the DAG induced by nodes $U_{cm-2d_1} - S$) with target set X_{cm}^- as well as parameters $S_2 \subseteq S_3 \subseteq \dots \subseteq S_k$ and d_2, d_3, \dots, d_k (both lists now have length $k - 1$ and clearly still satisfy the conditions on parameters stated above). If `RGenPeb` is ever called with empty lists $\bar{S} = \emptyset$ and $\bar{d} = \emptyset$ (i.e., $k = 0$) then it simply greedily pebbles G . The result of the recursive call is added to the final $2d_1$ steps of light phase Λ_c . Finally the pebbling is modified to never remove pebbles from X_{cm} during the those final steps of Λ_{c-1} . Notice that each node in X_{cm}^+ is either in S or is pebbled at some point during the final $2d_1$ steps of Λ_{c+1} . Thus we are guaranteed that $X_{cm} \subseteq P_{(c-1)m}$ as desired.

To see why `RGenPeb` produces a legal pebbling it suffices to observe that pebbles placed during light phases always have their parents already pebbled. So if the recursive call returns a legal pebbling for the balloon phase then the final result is also legal. But at the deepest level of the recursion `RGenPeb` resorts to a greedy pebbling which is trivially legal. Thus, by induction, so is the pebbling at the highest level of the recursion.

7 Open Questions

We conclude with several open questions for future research.

- We showed that for some constant $c \geq 0$ we can find a DAG G on n nodes with $\Pi_{cc}^{\parallel}(G) \geq cn^2 / \log(n)$ and $\text{indeg}(G) = 2$. While this result is asymptotically optimal the constant terms are relevant for practical applications to iMHFs. How big can this constant c be? Can we find explicit constructions of constant-indegree, $(c_1n / \log(n), c_2n)$ -depth robust DAGs that match these bounds?
- Provide tighter upper and lower bounds on $\Pi_{cc}^{\parallel}(G)$ for Argon2i-B [?], the most recent version of Argon2i which was submitted to IRTF for standardization.
- Another interesting direction concerns understanding the cumulative pebbling complexity of generic graphs. Given a graph G is it computationally tractable to (approximately) compute $\Pi_{cc}^{\parallel}(G)$? An efficient approximation algorithm for $\Pi_{cc}^{\parallel}(G)$ would allow us to quickly analyze candidate iMHF constructions. Conversely, as many existing iMHF constructions are based on fixed random graphs, [?,?] showing that approximating such a graphs complexity is hard would provide evidence that an adversary will likely not be able to leverage properties of the concrete instance to improve their evaluation strategy for the iMHF. Indeed, it may turn out that the most effective way to construct depth-robust graphs with good constants is via a randomized construction.

Acknowledgments

The authors would like to thank Pierrick Gaudry for his careful reading and many helpful suggestions. The first and third authors were supported by the European Research Council, ERC consolidator grant (682815 - TOCNeT).