On the Depth-Robustness and Cumulative Pebbling Cost of Argon2i

Jeremiah Blocki[*]   Samson Zhou[†]

June 30, 2017

Abstract

Argon2i is a data-independent memory hard function that won the password hashing competition. The password hashing algorithm has already been incorporated into several open source crypto libraries such as libsodium. In this paper we analyze the cumulative memory cost of computing Argon2i. On the positive side we provide a lower bound for Argon2i. On the negative side we exhibit an improved attack against Argon2i which demonstrates that our lower bound is nearly tight. In particular, we show that

(1) An Argon2i DAG is $\left(e, O\left(n^3/e^3\right)\right)$-reducible.

(2) The cumulative pebbling cost for Argon2i is at most $O\left(n^{1.768}\right)$. This improves upon the previous best upper bound of $O\left(n^{1.8}\right)$ [AB17].

(3) Argon2i DAG is $\left(e, \tilde{\Omega}\left(n^3/e^3\right)\right)$-depth robust. By contrast, analysis of [ABP17a] only established that Argon2i was $\left(e, \tilde{\Omega}\left(n^2/e^2\right)\right)$-depth robust.

(4) The cumulative pebbling complexity of Argon2i is at least $\tilde{\Omega}\left(n^{1.75}\right)$. This improves on the previous best bound of $\Omega\left(n^{1.66}\right)$ [ABP17a] and demonstrates that Argon2i has higher cumulative memory cost than competing proposals such as Catena or Balloon Hashing.

We also show that Argon2i has high *fractional* depth-robustness which strongly suggests that data-dependent modes of Argon2 are resistant to space-time tradeoff attacks.

1 Introduction

Memory-hard functions (MHFs) are a promising primitive to help protect low entropy user passwords against offline attacks. MHFs can generally be divided into two categories: data-dependent (dMHF) and data-independent (iMHF). A data-independent MHF (iMHF) is characterized by the property that the memory-access pattern induced by an honest evaluation algorithm is not dependent on the input to the function (e.g., the password). In contexts such as password hashing, iMHFs are useful for their resistance to side-channel attacks such as cache-timing [Ber][1].

---

[1]Unfortunately, this resistance to side-channel attacks has a price; we now know that the dMHFs `scrypt` enjoys strictly greater memory-hardness [ACP+17] than can possibly be achieved for a very broad class of iMHFs [AB16].

Both in theory and in practice, iMHFs (e.g.,[BDK16, CGBS16, CJMS14, Cox14, Wu15, Pin14, AABSJ14]) can be viewed as a directed acyclic graph (DAG) which describes how inputs and outputs of various calls to an underlying compression function are related. That is, the function $f_{G,h}$ can be fully specified in terms of a DAG G and a round function h. The input to the function is the label of the source node(s) and the output of the function is the label of the sink node(s). The label of node v is computed by applying the round function h to the labels of v's parents.

The goal of a MHF is to ensure that it is cost prohibitive for an attacker to evaluate $f_{G,t}$ millions or billions of times even if the attacker can use customized hardware (e.g., FPGAs, ASICs). Thus, we wish to lower bound the "cumulative memory complexity" or "amortized area-time complexity" of any algorithm that computes $f_{G,h}$.

1.1 iMHFs, Graph Pebbling and Depth-Robustness

In the parallel random oracle model, the memory hardness of the iMHF $f_{G,h}$ can be characterized using the parallel black pebbling game on the graph G [AS15, CGBS16, FLW13]. In particular, the "cumulative memory complexity" or "amortized area-time complexity" of $f_{G,h}$ is (essentially) equivalent to the cumulative cost of any legal black pebbling of G in the parallel Random Oracle Model (pROM) [AS15]. 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 ($\mathsf{parents}(v) \subseteq P_{i-1}$) or if v was already pebbled in the last round ($v \in P_{i-1}$). The cumulative cost of the pebbling is defined to be $|P_1| + \ldots + |P_t|$.

Graph pebbling is a particularly useful as a tool to analyze the security of an iMHF [AS15]. A pebbling of G naturally corresponds to an algorithm to compute the iMHF. Alwen and Serbinenko [AS15] proved that in the parallel random oracle model (pROM) of computation, *any* algorithm evaluating such an iMHF could be reduced to a pebbling strategy with (approximately) the same cumulative memory cost.

Recently it has been shown that for a DAG G to have high "amortized area-time complexity" it is both necessary [ABP17a] and sufficient [AB16] for G to be very depth-robust, where an (e, d, b)-block depth robust DAG G has the property that after removing any subset $S \subseteq V(G)$ of up to e blocks of b-consecutive nodes (and adjacent edges) there remains a directed path of length d in $G - S$ (when $b = 1$ we simply say that G is (e, d)-depth robust). It is particularly important to understand the depth-robustness and cumulative pebbling cost of iMHF candidates.

1.2 Argon2i

Of particular importance is the iMHF candidate Argon2i [BDK15], winner of the password hashing competition. Argon2i is being considered for standardization by the Cryptography Form Research Group (CFRG) of the IRTF [BDKJ16] [2].

---

[2]The specification of Argon2i has changed several times. Older versions of the specification constructed G by sampling edges uniformly at random, while this distribution has been modified to a non-uniform distribution in newer versions. Following [AB17] we use Argon2i-A to refer to *all (older) versions* of the algorithm that used a uniform edge distribution. we use Argon2i-B to refer to all versions of the algorithm that use the new non-uniform edge distribution (including the current version that is being considered for standardization by the Cryptography Form Research Group

While significant progress has been made in the last two years in understanding the depth-robustness and cumulative pebbling complexity of candidate iMHFs (e.g., see Table 1) there is still a large gap in the lower and upper bounds for Argon2i, which is arguably the most important iMHF candidate to understand. A table summarizing the asymptotic cumulative complexity of various iMHFs can be found in Table 1.

Algorithm	Lowerbound	Upperbound	Appearing In
Argon2i-A	$\tilde{\Omega}\left(n^{1.\bar{6}}\right)$	$\tilde{O}\left(n^{1.708}\right)$	[ABP17a]
Argon2i-B		$O\left(n^{1.8}\right)$	[AB17]
Argon2i-B	$\tilde{\Omega}\left(n^{1.\bar{6}}\right)$		[ABP17a]
Argon2i-B	$\tilde{\Omega}\left(n^{1.75}\right)$	$O\left(n^{1.767}\right)$	This Work
Balloon-Hashing	$\tilde{\Omega}\left(n^{1.5}\right)$	$\tilde{O}\left(n^{1.625}\right)$	[ABP17a]
Balloon-Hashing: Single Buffer (SB)	$\tilde{\Omega}\left(n^{1.\bar{6}}\right)$	$\tilde{O}\left(n^{1.708}\right)$	[ABP17a]
Catena	$\tilde{\Omega}\left(n^{1.5}\right)$	$\tilde{O}\left(n^{1.625}\right)$	[ABP17a]
	$\Omega\left(\frac{n^2}{\log n}\right)$		[ABP17a]
Arbitrary iMHF		$O\left(\frac{n^2 \log \log n}{\log n}\right)$	[AB16]

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

1.3 Results

We first completely characterize the depth-robustness of Argon2i in Theorem 1, and then apply our bounds to develop (nearly tight) upper and lower bounds on the cumulative pebbling cost of Argon2i — see Theorem 2 and Theorem 3. For comparison, the previous best known upper bound for Argon2i was $O\left(n^{1.8}\right)$ and the best known lower bound was $\Omega\left(n^{5/3}\right)$. Our new bounds are $O\left(n^{1.7676}\right)$ and $\tilde{\Omega}\left(n^{7/4}\right)$ respectively.

Interestingly, Theorem 1 shows that Argon2i is more depth-robust than Argon2i-A as well as other competing iMHFs such as Catena [FLW13] or Balloon Hashing [CGBS16][3]. Furthermore, Theorem 2 in combination with attacks of Alwen et al. [ABP17a] show that Argon2i enjoys strictly greater cumulative memory complexity than Catena [FLW13] or Balloon Hashing [CGBS16] as well as the earlier version Argon2i-A.

Theorem 1 *Argon2i is* $\left(e, \tilde{\Omega}(n^3/e^3), \Omega(n/e)\right)$*-block depth robust with high probability.*

Theorem 2 *For any $\epsilon > 0$ the cumulative pebbling cost of a random Argon2i DAG G is at most* $\Pi_{cc}^{\parallel}(G) = O(n^{1+a+\epsilon})$ *with high probability, where $a = \frac{1/3+\sqrt{1+4/9}}{2} \approx 0.7676$.*

Theorem 3 *With high probability, the cumulative pebbling cost of a random Argon2i DAG G is at least* $\Pi_{cc}^{\parallel}(G) = \tilde{\Omega}\left(n^{7/4}\right)$ *with high probability.*

Theorem 4 *If G is (e,d,b)-block depth robust, then G is $\left(\frac{e}{2}, d, \frac{eb}{2n}\right)$-fractional depth robust.*

(CFRG) of the IRTF[BDKJ16]). Since we are primarily interested in analyzing the current version of the algorithm we will sometime simply write Argon2i instead of Argon2i-B. By contrast, we will always write Argon2i-A whenever we refer to the earlier version.

[3]Argon2i is not as depth-robust as the theoretically optimal constructions of [ABP17a], but at the moment this construction is purely theoretical while Argon2i has been deployed in crypto libraries such as libsodium

Techniques To upper bound the depth-robustness of Argon2i we use the layered attack of [AB16]. Once we know that Argon2i is depth-reducible for multiple different points (e_i, d_i) along a curve, then we can apply a recursive pebbling attack of Alwen et al. [ABP17a] to obtain the upper bounds on cumulative pebbling complexity from Theorem 2.

Lower bounding the depth-robustness of Argon2i is significantly more challenging. We adapt machinery from Erdos et al.[EGS75] to reason about the depth-robustness of meta-graph G_m of an Argon2i DAG G (essentially, the meta-graph is formed by compressing each group of m sequential nodes in G into a single point to obtain a new graph with $n' = n/m$ nodes). We prove that for appropriate choice of m and r^* that the meta-graph is a local expander meaning that for every $r \geq r^*$ every node $x \leq (n/m) + 1 - 2r$ the sets $[x, x + r - 1]$ and $[x + r, x + 2r - 1]$ are connected by an expander graph. We then use local expansion to lower bound the depth-robustness of G_m. Finally, we can apply a result of Alwen et al. [ABP17a] to translate this bound to a lower bound on the block depth robustness of G_m.

Finally, we extend ideas from [ABP17a] to lower bound the cumulative pebbling complexity of an Argon2i DAG. Essentially, we show that any pebbling strategy must either keep $\tilde{\Omega}\left(n^{0.75}\right)$ pebbles on the graph during most pebbling rounds or repebble a $\left(\tilde{\Omega}\left(n^{0.75}\right), \tilde{\Omega}\left(n^{0.75}\right)\right)$-depth robust graph $\tilde{\Omega}\left(n^{0.25}\right)$ times. In the first case the cumulative cost is at least $\Omega\left(n \times n^{0.75}\right)$ since we have at least n pebbling rounds and in the second case we also have that cumulative cost is at least $\Omega\left(n^{0.25} \times n^{1.5}\right)$ since the cost to repebble a $\left(e = \tilde{\Omega}\left(n^{0.75}\right), d = \tilde{\Omega}\left(n^{0.75}\right)\right)$-depth robust graph is at least ed [ABP17a].

2 Related Work

[ABW03] noticed that that cache-misses are more egalitarian than computation and therefore proposed the use of functions which maximize the number of expensive cache misses, "memory-bound" functions. Percival [Per09] observed that memory costs seemed to be more stable across different architectures and proposed the use of memory-hard functions (MHFs) for password hashing. Since the cost of computing the function is primarily memory related (storing/retrieving data values) and cannot be significantly reduced by constructing an ASIC, there presently seems to be a consensus that memory hard functions are the "right tool" for constructing moderately expensive functions. In fact, all entrants in the password hashing competition claimed some form of memory hardness [PHC]. Percival [Per09] introduced a candidate memory hard function called `scrypt`, which has subsequently been shown to be vulnerable to side-channel attacks as its computation yields a memory access pattern that is data-dependent (i.e., depends on the secret input/password). On the positive side this function has been shown to require maximum possible cumulative memory complexity to evaluate [ACP+17].

Alwen and Blocki [AB16] gave an attack on Argon2i-A (an earlier version of Argon2i) with cumulative memory complexity $O(n^{1.75} \log n)$ as well as several other iMHF candidates. They later extended the attack to Argon2i-B (the current version) showing that the function has complexity $O(n^{1.8})$ [AB17]. Alwen and Blocki [AB16] also showed that any iMHF has cumulative memory complexity at most $O\left(\frac{n^2 \log \log n}{\log n}\right)$, and Alwen et al. [ABP17a] later constructed a graph with cumulative pebbling complexity at least $\Omega\left(\frac{n^2 \log \log n}{\log n}\right)$. Alwen et al. [ABP17a] also found a "recursive version" of the [AB16] attack which further reduced the cumulative memory complexity of Argon2i-A

4

to $\tilde{O}\left(n^{1.708}\right)$. At the same time they established a lower bound of $\tilde{\Omega}\left(n^{1.\bar{6}}\right)$ for Argon2i-A and Argon2i-B.

Depth-robust graphs have found several applications in theoretical computer science e.g., proving lowerbounds on circuit complexity and Turing machine time [Val77, PR80, Sch82, Sch83]. [MMV13] constructed proofs of sequential work using depth-robust graph and more recently depth-robust graphs were used to prove lower bounds in the domain of proof complexity [ARN16]. Recent results [AB16, ABP17a] demonstrate that depth-robustness is a necessary and sufficient property for a secure iMHF. Several constructions of graphs with low indegree exhibiting this asymptotically optimally depth-robustness are given in the literature [EGS75, PR80, Sch82, Sch83, MMV13, ABP17b] but none of these constructions are suitable for practical deployment.

3 Preliminaries

Let $\mathbb{N}$ denote the set $\{0, 1, \ldots\}$ and $\mathbb{N}^+ = \{1, 2, \ldots\}$. Let $\mathbb{N}_{\geq c} = \{c, c+1, c+2, \ldots\}$ for $c \in \mathbb{N}$. Define $[n]$ to be the set $\{1, 2, \ldots, n\}$ and $[a, b] = \{a, a+1, \ldots, b\}$ where $a, b \in \mathbb{N}$ with $a \leq b$.

We say that a directed acyclic graph (DAG) $G = (V, E)$ has *size* n if $|V| = n$. A node $v \in V$ has indegree $\delta = \mathsf{indeg}(v)$ if there exist δ incoming edges $\delta = |(V \times \{v\}) \cap E|$. More generally, we say that G has indegree $\delta = \mathsf{indeg}(G)$ if the maximum indegree of any node of G is δ. A node with indegree 0 is called a source node and a node with no outgoing edges is called a sink. We use $\mathsf{parents}_G(v) = \{u \in V : (u, v) \in E\}$ to denote the parents of a node $v \in V$. In general, we use $\mathsf{ancestors}_G(v) = \bigcup_{i \geq 1} \mathsf{parents}_G^i(v)$ to denote the set of all ancestors of v — here, $\mathsf{parents}_G^2(v) = \mathsf{parents}_G(\mathsf{parents}_G(v))$ denotes the grandparents of v and $\mathsf{parents}_G^{i+1}(v) = \mathsf{parents}_G(\mathsf{parents}_G^i(v))$. When G is clear from context we will simply write $\mathsf{parents}$ ($\mathsf{ancestors}$). We denote the set of all sinks of G with $\mathsf{sinks}(G) = \{v \in V : \nexists(v, u) \in E\}$ — note that $\mathsf{ancestors}(\mathsf{sinks}(G)) = V$. 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 $\mathbb{G}_{n, \delta} = \{G \in \mathbb{G}_n : \mathsf{indeg}(G) \leq \delta\}$. For directed path $p = (v_1, v_2, \ldots, v_z)$ in G, its length is the number of nodes it traverses, $\mathsf{length}(p) := z$. The depth $d = \mathsf{depth}(G)$ 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 from G by removing nodes S and incident edges. The following is a central definition to our work.

Definition 5 (Block Depth-Robustness) *Given a node v, let $N(v, b) = \{v - b + 1, \ldots, v\}$ denote a segment of b consecutive nodes ending at v. Similarly, given a set $S \subseteq V$, let $N(S, b) = \cup_{v \in S} N(v, b)$. 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 $\mathsf{depth}(G - N(s, b)) \geq d$. 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.*

Observe when $b > 1$ (e, d, b)-block-depth robustness is a strictly stronger notion that (e, d)-depth-robustness since the set $N(S, b)$ of nodes that we remove may have size as large as $|N(S, b)| = eb$.

We fix our notation for the parallel graph pebbling game following [AS15].

Definition 6 (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 $P = (P_0, \ldots, P_t)$ of pebbling configurations of G where $P_0 = \emptyset$ and which satisfies conditions 1 & 2 below.*

5

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

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

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

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

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

We remark that in the sequential black pebbling game, we face the additional restriction that at most one pebble is place in each step ($\forall i \in [t] \;\; : \;\; |P_i \setminus P_{i-1}| \leq 1$), while in the parallel black pebbling game there is no such restriction. The cumulative complexity of a pebbling $P = \{P_0, \ldots, P_t\} \in \mathcal{P}_G^{\parallel}$ is defined to be $\Pi_{cc}(P) = \sum_{i \in [t]} |P_i|$. The cumulative cost of pebbling a graph G a target set $T \subseteq V$ is defined to be

$$\Pi_{cc}^{\parallel}(G, T) = \min_{P \in \mathcal{P}_{G,T}^{\parallel}} \Pi_{cc}(P) \;.$$

When $T = \mathsf{sinks}(G)$, we simplify notation and write $\Pi_{cc}^{\parallel}(G) = \min_{P \in \mathcal{P}_G^{\parallel}} \Pi_{cc}(P)$.

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.

3.1 Edge Distribution of Argon2i-B

Definition 7 gives the edge distribution for the single-lane/single-pass version of Argon2i-B. The definition also captures the core of the Argon2i-B edge distribution for multiple lane/multiple-pass variants of Argon2i-B. While we focus on the single-lane/single-pass variant for ease of exposition, we stress that all of our results can be extended to multiple-lane/multiple-pass versions of Argon2i-B provided that the parameters $\tau, \ell = O(1)$ are constants. Here, ℓ is the number of lanes and τ is the number of passes and in practice these parameters ℓ and τ will be *always* be constants.

Definition 7 *The Argon2i-B is a graph $G = (V = [n], E)$, where $E = \{(i, i+1) \;\; : \;\; i \in [n-1]\} \cup \{(r(i), i)\}$, where $r(i)$ is a random value distributed as follows:*

$$\Pr[r(i) = j] = \Pr_{x \in [N]} \left[i \left(1 - \frac{x^2}{N^2} \right) \in (j-1, j] \right],$$

since $i \left(1 - \frac{x^2}{N^2} \right)$ is not always an integer. Note that we assume $N \ll n$, and in the Argon2i-B implementation we consider, $N = 2^{32}$.

3.2 Metagraphs

We will use the notion of a metagraph in our analysis. Fix an arbitrary integer $m \in [n]$ and set $n' = \lfloor n/m \rfloor$. Given a DAG G, we will define a DAG G_m called the metagraph of G. For this, we

use the following sets. 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

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

and

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

We define the metagraph $G_m = (V_m, E_m)$ as follows:

Nodes: V_m contains one node v_i per set M_i. 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.

$$V_m = \{v_i : i \in [n']\} \qquad E_m = \{(v_i, v_j) : E \cap (M_i^L \times M_j^F) \neq \emptyset\}.$$

Claim 8 is a simple extension of a result from [ABP17a], which will be useful in our analysis.

Claim 8 ([ABP17a], Claim 1) *If G_m is (e, d)-depth robust, then G is $\left(\frac{e}{2}, \frac{dm}{3}, m\right)$-block depth robust.*

4 Depth-Reducibility of Argon2iB

In this section, we show that the Argon2i-B is depth reducible with high probability. Then, using results from previous layered attacks (such as [AB16], [ABP17a]), we show an upper bound on the computational complexity of Argon2i-B.

Theorem 9 *With high probability, the Argon2i-B graph is $\left(e, \Omega\left(\left(\frac{n}{e}\right)^3\right)\right)$-depth reducible.*

Proof : Recall that for node i, Argon2i-B creates an edge from i to parent node $i\left(1 - \frac{x^2}{N^2}\right)$, where $x \in [N^2]$ is picked uniformly at random. Suppose we have gap size g and L layers, each of size $\frac{n}{L}$. Let i be in layer α, so that $i \in \left[(\alpha-1)\frac{n}{L}, \alpha\frac{n}{L}\right]$. Then the probability that the parent of i is also in layer α, for $\alpha > 1$, is

$$
\begin{aligned}
\mathbf{Pr}\left[(\alpha-1)\frac{n}{L} \leq i\left(1 - \frac{j^2}{N^2}\right)\right] &\leq \mathbf{Pr}\left[(\alpha-1)\frac{n}{iL} \leq \left(1 - \frac{j^2}{N^2}\right)\right] \\
&= \mathbf{Pr}\left[\left(\frac{x^2}{N^2}\right) \leq \frac{iL - (\alpha-1)n}{iL}\right] \\
&\leq \mathbf{Pr}\left[\left(\frac{x^2}{N^2}\right) \leq \frac{\alpha n - (\alpha-1)n}{iL}\right] \\
&\leq \mathbf{Pr}\left[\left(\frac{x^2}{N^2}\right) \leq \frac{n}{(\alpha-1)n}\right] \\
&\leq \frac{1}{\sqrt{\alpha-1}}
\end{aligned}
$$

Thus, the expected number of in-layer edges is

$$\frac{n}{L}\left(1 + \frac{1}{\sqrt{1}} + \frac{1}{\sqrt{2}} + \frac{1}{\sqrt{3}} + \ldots\right) < \frac{n}{L}\left(2\int_1^L \frac{1}{\sqrt{\alpha - 1}}\, d\alpha\right) = 4\frac{n}{\sqrt{L}}.$$

Hence, if we remove an node between every g nodes, as well as all in-layer edges, we have $e = \frac{n}{g} + \frac{4n}{\sqrt{L}}$. We can apply standard concentration bounds to show that e is tightly concentrated around the As a result, the depth is at most g nodes each gap over all L layers, $d = gL$. Therefore, Argon2i-B is $\left(\frac{n}{g} + \frac{4n}{\sqrt{L}}, gL\right)$ depth reducible. Setting $g = \sqrt{L}$ shows $\left(\frac{5n}{\sqrt{L}}, L^{3/2}\right)$ depth reducibility. Consequently, for $e = \frac{5n}{\sqrt{L}}$, then $L^{3/2} = \left(\frac{5n}{e}\right)^3$, and the result follows. $\qquad\square$

Given function f, we say that G is f-reducible if G is $(f(d), d)$-reducible for each value $d \in [n]$. Theorem 10, due to Alwen et al. [ABP17a], upper bounds $\Pi_{cc}^{\parallel}(G)$ for any f-reducible DAG.

Theorem 10 ([ABP17a], 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 $0 < b \leq \frac{2}{3}$ then for any constant $\epsilon > 0$, the cumulative pebbling cost of G is at most $\Pi_{cc}^{\parallel}(G) = O(n^{1+a+\epsilon})$, where $a = \frac{1 - 2b + \sqrt{1 + 4b^2}}{2}$.*

Reminder of Theorem 2. *For any $\epsilon > 0$ the cumulative pebbling cost of a random Argon2i DAG G is at most $\Pi_{cc}^{\parallel}(G) = O(n^{1+a+\epsilon})$ with high probability, where $a = \frac{1/3 + \sqrt{1 + 4/9}}{2} \approx 0.7676$.*

Proof of Theorem 2: By Theorem 9, the Argon2i-B graph is f-reducible for $b = \frac{1}{3}$ with high probability, and the result follows. $\qquad\square$

5 Depth-Robustness for Argon2iB

In this section we show the general block-depth robustness curve of a random Argon2i-B DAG. We will ultimately use these results to lower bound the cumulative pebbling of an Argon2i-B DAG in Section 6. Interestingly, our lower bound from Theorem 1 matches the upper bound from Theorem 9 in the last section up to logarithmic factors. Thus, both results are essentially tight.

Reminder of Theorem 1. *Argon2i is $\left(e, \tilde{\Omega}(n^3/e^3), \Omega(n/e)\right)$-block depth robust with high probability.*

The notion of a (δ, r^*)-local expander will be useful in our proofs. Definition 11 extends the basic notion of a δ-local expander from [EGS75]. [EGS75] showed that for a sufficiently small constant δ, any δ-local expander is $(\Omega(n), \Omega(n))$-depth robust.

Definition 11 *A directed acyclic graph G (with n nodes) is a (δ, r^*)-local expander if for all $r \geq r^*$ and for all $x \leq n - 2r + 1$ and all $A \subseteq \{x, \ldots, x + r - 1\}$, $B \subseteq \{x + r, \ldots, x + 2r - 1\}$ such that $|A|, |B| \geq \delta r$, we have $E(G) \cap A \times B \neq \emptyset$. That is, there exists an edge from some node in X to some node in X'. If $r^* = 1$, then we say G is a δ-local expander.*

Proof Overview: We set $m = \Omega(n/e)$ and construct a metagraph G_m for a random Argon2i-B graph, and bound the probability that two metanodes in G_m are connected, using Claim 12 and Claim 13. Using these bounds, we show that the metagraph G_m for a random Argon2i-B graph is a (δ, r^*)-local expander with high probability for $r^* = \tilde{\Omega}(e^3/n^2)$ and some suitably small constant $\delta > 0$. We then

divide the metagraph into several layers. With respect to a set S, we call a layer "good" if S does not remove too many elements from the layer. We then show that there exists a long path between these layers, which indicates that the remaining graph has high depth.

We now show that the Argon2i-B class of graphs is a (δ, r^*)-local expander with high probability. Given a directed acyclic graph G with n nodes, let G_m be the graph with the metanodes of G, where each metanode has size $m = 6n^{1/3} \log n$, so that G_m has $\frac{n^{2/3}}{6 \log n}$ nodes. First, given two metanodes $x, y \in G_m$ with $x < y$, we bound the probability that for node i in metanode y, there exists an edge from x to i.

Claim 12 *For each $x, y \in G_n$ with $y > x$ and node i in metanode y, there exists an edge from the last third of metanode x to node i with probability at least $\frac{1}{12\sqrt{y}\sqrt{y-x+1}}$.*

Claim 13 *For any two metanodes $x, y \in G_m$ with $x < y$, the last third of x is connected to the first third of y with probability at least $\frac{m\sqrt{m}}{m\sqrt{m}+36\sqrt{n(y-x+1)}}$.*

This allows us to show that the probability there exist subsets $A \subseteq [x, x+r-1]$ and $B \subseteq [x+r, +2r-1]$ of size δr such that A has no edge to B is at most $e^{-\delta r \log(1+\sqrt{\log n})} \binom{r}{\delta r}^2$. We then use Stirling's approximation to show this term is negligible, and then apply the union bound over all vertices x and all $r \geq r^*$, which shows that the metagraph G_m (for Argon2i) is a (δ, r^*)-local expander with high probability.

Lemma 14 *Let $m = n/(20000e)$ then for $r^* = \tilde{\Omega}(e^3/n^2) = \tilde{\Omega}(n/m^3)$ the metagraph G_m (for Argon2i) is a (δ, r^*)-local expander with high probability.*

We now divide G_m into layers $L_1, L_2, \ldots L_{n/(mr^*)}$ of size r^* each. Say that a layer L_i is c-good with respect to a subset $S \subseteq V(G_m)$ if for all $t \geq 0$ we have

$$\left| S \cap \left(\bigcup_{j=i}^{i+t-1} L_j \right) \right| \leq c \left| \left(\bigcup_{j=i}^{i+t-1} L_j \right) \right| \quad , \text{ and } \left| S \cap \left(\bigcup_{j=i-t+1}^{i} L_j \right) \right| \leq c \left| \left(\bigcup_{j=i-t+1}^{i} L_j \right) \right|,$$

We ultimately want to argue that $G_m - S$ has a path through these good layers.

Claim 15 *If $|S| < n/(10000m)$ then at least half of the layers $L_1, L_2, \ldots L_{n/(mr^*)}$ are $(1/1000)$-good with respect to S.*

Fixing a set S let $Y_{1,S}, Y_{2,S}, \ldots,$ denote the c-good layers and let $R_{1,S} = Y_{1,S} - S$ and let $R_{i+1,S} = \{x \in Y_{i+1,S} \mid x$ can be reached from some $y \in R_{i,S}$ in $G_m - S\}$.

Lemma 16 *Suppose that for any S with $|S| \leq e$ and $i \leq n/(2mr^*)$, the set $R_{i,S} \neq \emptyset$. Then G_m is $(e = n/(10000m), n/(2mr^*))$-depth robust and G is $(e = n/(20000m), n/(6r^*), m)$-block depth robust.*

Proof : Removing any $e = n/(10000m)$ nodes from G_m, there is still a path passing through each good layer since $R_{i,S} \neq \emptyset$ and there are at least $n/(2mr^*)$. Removing $e = n/(20000m)$ blocks of nodes of size m from G can affect at most $n/(10000m)$ metanodes. Thus, there is a path of length $(m/3)n/(2mr^*) = n/(6r^*)$ through G. $\qquad \square$

9

We now show that the number of nodes in each reachable good layer $R_{i,S}$ is relatively high, which allows us to construct a path through the nodes in each of these layers. We first show that if two good layers $Y_{i,S}$ and $Y_{i+1,S}$ are close to each other, then no intermediate layer contains too many nodes in S, so we can use expansion to inductively argue that each intermediate layer has many reachable nodes from $R_{i,S}$, and it follows that $R_{i+1,S}$ is large. On the other hand, if $Y_{i,S}$ and $Y_{i+1,S}$ have a large number of intermediate layers in between, then the argument becomes slightly more involved. However, we can use local expansion to argue that most of the intermediate layers have the property that most of the nodes in that layer are reachable. We then use a careful argument to show that as we move close to layer $Y_{i+1,S}$, the density of layers with this property increases. It then follows that $R_{i+1,S}$ is large.

Lemma 17 *Suppose that G_m is a (δ, r^*)-local expander with $\delta = 1/16$ and let $S \subseteq V(G_m)$ be given such that $|S| \leq n/(10000m)$. Then, $R_{i,S} \geq 7r^*/8$.*

Proof of Theorem 1:   Let $m = 20000n/e$ and let G be a random Argon2i DAG. Lemma 14 shows that the metagraph G_m of a random Argon2i DAG G is a (δ, r^*)-local expander with high probability for $r^* = \tilde{\Omega}\left(e^3/n^2\right)$. Now fix any set $S \subseteq G_m$ of size $|S| \leq e$. Claim 15 now implies we have at least $n/(2mr^*)$ good layers $Y_{1,S}, \ldots, Y_{n/(2mr^*)}$. Theorem 1 now follows by applying Lemma 17 and Lemma 16. □

6 Cumulative Pebbling Cost of Argon2iB

We now use the depth-robust results to show a lower bound on the computational complexity of Argon2iB. Given a pebbling of G, we show in Theorem 18 that if at any point the number of pebbles on G is low, then we must completely re-pebble a depth-robust graph. We then appeal to a result which provides a lower bound on the cost of pebbling a depth-robust graph.

Theorem 18 *Suppose G is a DAG that has an edge from $[i, i+b-1]$ to $\left[j, j+\frac{128n\log n}{b}\right]$ for all $\frac{n}{2} + j \leq n - \frac{128n\log n}{b}$ and $1 \leq i \leq \frac{n}{2} - b + 1$. If the subgraph induced by nodes $\left[1, \frac{n}{2}\right]$ is (e, d, b)-block depth robust, then the cost to pebble G is at least $\min\left(\frac{en}{8}, \frac{edb}{8\log n}\right)$.*

First, we exhibit a property which occurs if the number of pebbles on G is low:

Lemma 19 *Suppose G is a DAG that has an edge from $[i, i+b-1]$ to $\left[j, j+\frac{128n\log n}{b}\right]$ for all $\frac{n}{2} + j \leq n - \frac{128n\log n}{b}$ and $1 \leq i \leq \frac{n}{2} - b + 1$. Suppose also that the subgraph induced by nodes $\left[1, \frac{n}{2}\right]$ is (e, d, b)-block depth robust. If $|S| < \frac{e}{2}$, then $H = \mathsf{ancestors}_{G-S}\left(\left[j, j+\frac{128n\log n}{b}\right]\right)$ is $\left(\frac{e}{2}, d\right)$-depth robust.*

Proof :   Let G_1 denote the subgraph induced by first $\frac{n}{2}$ nodes. Note that H contains the graph $W = G_1 - \bigcup_{x in S}[x - b + 1, x]$ since there exists an edge from each interval $[x - b + 1, x]$. Moreover, W is $\left(\frac{e}{2}, d, b\right)$-block depth robust since G_1 is (e, d, b)-block depth robust contains only $\frac{e}{2}$ additional blocks. Finally, since W is a subgraph of H, then H is $\left(\frac{e}{2}, d\right)$-depth robust. □

Lemma 20 ([ABP17a], Corollary 2) *Given a DAG $G = (V, E)$ and subsets $S, T \subset V$ such that $S \cap T = \emptyset$, let $G' = G - (V/\mathsf{ancestors}_{G-S}(T))$. If G' is (e, d)-depth robust, then the cost of pebbling $G - S$ with target set T is $\Pi_{cc}^{\parallel}(G - S, T) > ed$.*

We now prove Theorem 18.

Proof of Theorem 18:   For each interval of length $\frac{2n \log n}{b}$, let t_1 denote the first time we pebble the first node, let t_2 denote the first time we pebble the middle node of the interval, and let t_3 denote the first time we pebble the last node of the interval. We show $\sum_{t \in [t_1, t_3]} |P_t| \geq \min\{en log(n)/(2b), ed/2\}$. Then a pebbling must either:

(1) Keep $\frac{e}{2}$ pebbles on G for at least $\frac{128n \log n}{b}$ steps (i.e., during the entire interval $[t_1, t_2]$)

(2) Pay $\left(\frac{e}{2}\right) d$ to repebble a $(e/2, d)$-depth robust DAG during before round t_3. (Lemma 19)

In the first case, $|P_t| \geq \frac{e}{2}$ for each $t \in [t_1, t_2]$, which is at least $\frac{128n \log n}{b}$ time steps. In the second case, there exists $t \in [t_1, t_2]$ such that $|P_t| < \frac{e}{2}$. Then by Lemma 19 and Lemma 20, $\sum_{t \in [t_1, t_3]} \geq \frac{ed}{2}$. The cost of the first case is $\frac{en \log n}{2b}$ and the cost of the second case is $\frac{ed}{2}$. Since the last $n/2$ nodes can be partitioned into $(n/2)/(2(n/b) \log n) = b/(4 \log n)$ such intervals, then the cost is at least $\left(\frac{b}{4 \log n}\right) \min \left(\frac{en \log n}{2b}, \frac{ed}{2}\right)$, and the result follows.   □

We now provide a lower bound on the probability that there exists an edge between two nodes in the Argon2iB graph.

Claim 21 *Let $i, j \in [n]$ be given $(i \neq j)$ and let G be a random Argon2iB DAG on n nodes. There exists an edge from node j to i in G with probability at least $\frac{1}{4n}$.*

Using the bound on the probability of two nodes being connected, we can also lower bound the probability that two intervals are connected in the Argon2iB graph.

Lemma 22 *Let $b \geq 1$ be a constant. Then with high probability, an Argon2iB DAG has the property that for all pairs i, j such that $\frac{n}{2} \leq j \leq n - \frac{128n \log n}{b}$ and $1 \leq i \leq \frac{n}{2} - b + 1$ there is an edge from $[i, i + b - 1]$ to $\left[j, j + \frac{128n \log n}{b}\right]$.*

Proof :   By Claim 21, the probability that there exists an edge from a specific node $y \in [i, i+b-1]$ to a specific node $x \in \left[j, j + \frac{128n \log n}{b}\right]$ is at least $\frac{1}{4n}$. Then the expected number of edges from $[i, i + b - 1]$ to $\left[j, j + \frac{128n \log n}{b}\right]$ is at least $\frac{1}{4n}(128n \log n) = 32 \log n$. By Chernoff bounds, the probability that there exists no edge from $[i, i+b-1]$ to $\left[j, j + \frac{128n \log n}{b}\right]$ is at most $\frac{1}{n^4}$. Taking a union bound over all possible intervals, the graph of Argon2iB is a DAG that has an edge from $[i, i+b-1]$ to $\left[j, j + \frac{128n \log n}{b}\right]$ and all $\frac{n}{2} + j \leq n - \frac{128n \log n}{b}$ and $1 \leq i \leq \frac{n}{2} - b + 1$ with probability at least $1 - \frac{1}{n^2}$.   □

We now have all the tools to lower bound the computational complexity of Argon2iB.

Reminder of Theorem 3.   *With high probability, the cumulative pebbling cost of a random Argon2i DAG G is at least $\Pi_{cc}^{\parallel}(G) = \tilde{\Omega}\left(n^{7/4}\right)$ with high probability.*

Proof of Theorem 3:   The result follows Theorem 18, Lemma 22, and setting $e = d = n^{3/4}$ and $b = n^{1/4}$.   □

7 Fractional Depth-Robustness

Thus far, our analysis has focused on Argon2i, the data-independent mode of operation for Argon2. In this section, we argue that our analysis of the depth-robustness of Argon2i has important security implications for both data-dependent modes of operation: Argon2 and Argon2id. In particular, we prove a generic relationship between block-depth robustness and fractional depth-robustness of any block-depth robust DAG such as Argon2i. Intuitively, fractional depth-robustness says that even if we delete e vertices from the DAG that a large fraction of the remaining vertices have depth $\geq d$ in the remaining graph.

In the context of a dMHF fractional depth-robustness is a significant metric because the attacker will be repeatedly challenged for a random data-label. Intuitively, if the attacker reduces memory usage and only stores e data labels, then there is a good chance that the attacker will need time $\geq d$ to respond to each challenge. It is known that SCRYPT has cumulative memory complexity $\Omega(n^2)$. However, SCRYPT allows for dramatic space-time trade-off attacks (e.g., attackers could evaluate SCRYPT with memory $O(1)$ if they are willing to run in time $O(n^2)$). Our results are compelling evidence for the hypothesis that similar time space-trade offs are not possible for Argon2 or Argon2id without incurring a dramatic increase in cumulative memory complexity (We believe that providing a formal proof of this claim could be a fruitful avenue of future research). In particular, our results provide strong evidence that any evaluation algorithm either (1) requires space $\Omega\left(n^{0.99}\right)$ for at least n steps, or (2) has cumulative memory complexity $\omega\left(n^2\right)$ since it should take time $\tilde{\Omega}\left(n^3/e^3\right) = \tilde{\Omega}\left(n^{2\epsilon} \times \frac{n}{e}\right)$ on average to respond to a random challenge on with any configuration with space $e = O(n^{1-\epsilon})$. By contrast for SCRYPT, it may only take time $\Omega(n/e)$ to respond to a random challenge starting from a configuration with space e — while this is sufficient to ensure cumulative memory complexity $\Omega(n^2)$, it does not prevent space-time trade-off attacks.

Definition 23 *Recall that the depth of a specific vertex v in graph G, denoted $\mathsf{depth}(v, G)$ is the length of the longest path to v in G. We say that a DAG $G = (V, E)$ is (e, d, f)-fractionally depth robust if for all $S \subseteq V$ with $|S| \leq e$, we have*

$$|\{v \in V : \mathsf{depth}(v, G - S) \geq d\}| \geq f \cdot n.$$

Then we have the following theorem which relates fractional depth-robustness and block depth-robustness.

Reminder of Theorem 4.   *If G is (e, d, b)-block depth robust, then G is $\left(\frac{e}{2}, d, \frac{eb}{2n}\right)$-fractional depth robust.*

Proof of Theorem 4:   Suppose, by way of contradiction, that G is not fractional depth robust. For each node $u \in V$ with $\mathsf{depth}(u, G - S) \geq d$ we have $\left[u, \min_{v \in S \cap [v+1,n]} v - 1\right] \subseteq \{v \in V : \mathsf{depth}(v, G - S) \geq d\}$. Thus, we can cover the entire interval $\left[u, \min_{v \in S \cap [v+1,n]} v - 1\right]$ using at most $\left\lceil \frac{\min_{v \in S \cap [v+1,n]} v - 1 - u}{b} \right\rceil$ blocks of size b. Removing all such intervals requires removing at most $|S| + \frac{eb}{2b} \leq e$ blocks of size b.

But then, the longest path in the resulting graph is at most $d - 1$, which contradicts that G is (e, d, b)-block depth robust. $\square$

Corollary 24 *Argon2i is $\left(e, \tilde{\Omega}(n^3/e^3), \Omega(1)\right)$-fractional depth robust with high probability.*

References

[AABSJ14] Leonardo C Almeida, Ewerton R Andrade, Paulo SLM Barreto, and Marcos A Simplicio Jr. Lyra: Password-based key derivation with tunable memory and processing costs. *Journal of Cryptographic Engineering*, 4(2):75–89, 2014.

[AB16] Joël Alwen and Jeremiah Blocki. Efficiently computing data-independent memory-hard functions. In *Advances in Cryptology CRYPTO'16*. Springer, 2016.

[AB17] Joël Alwen and Jeremiah Blocki. Towards practical attacks on argon2i and balloon hashing. In *2nd IEEE European Symposium on Security and Privacy (EuroS&P 2017)*, 2017.

[ABP17a] Joël Alwen, Jeremiah Blocki, and Krzysztof Pietrzak. Depth-robust graphs and their cumulative memory complexity. In *Advances in Cryptology - EUROCRYPT - 36th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Proceedings, Part III*, pages 3–32, 2017.

[ABP17b] Joel Alwen, Jeremiah Blocki, and Krzysztof Pietrzak. Sustained space complexity. *arXiv preprint arXiv:1705.05313*, 2017.

[ABW03] Martín Abadi, Michael Burrows, and Ted Wobber. Moderately hard, memory-bound functions. In *Proceedings of the Network and Distributed System Security Symposium, NDSS 2003, San Diego, California, USA*, 2003.

[ACP+17] Joël Alwen, Binyi Chen, Krzysztof Pietrzak, Leonid Reyzin, and Stefano Tessaro. `scrypt` is Maximally Memory-Hard. In *Advances in Cryptology-EUROCRYPT 2017*, page (to appear). Springer, 2017. http://eprint.iacr.org/2016/989.

[ARN16] 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*, 2016.

[AS15] Joël Alwen and Vladimir Serbinenko. High Parallel Complexity Graphs and Memory-Hard Functions. In *Proceedings of the Eleventh Annual ACM Symposium on Theory of Computing*, STOC '15, 2015. http://eprint.iacr.org/2014/238.

[BDK15] Alex Biryukov, Daniel Dinu, and Dmitry Khovratovich. Fast and tradeoff-resilient memory-hard functions for cryptocurrencies and password hashing. Cryptology ePrint Archive, Report 2015/430, 2015. http://eprint.iacr.org/2015/430.

[BDK16] Alex Biryukov, Daniel Dinu, and Dmitry Khovratovich. Argon2 password hash. Version 1.3, 2016. https://www.cryptolux.org/images/0/0d/Argon2.pdf.

[BDKJ16] Alex Biryukov, Daniel Dinu, Dmitry Khovratovich, and Simon Josefsson. The memory-hard Argon2 password hash and proof-of-work function. Internet-Draft draft-irtf-cfrg-argon2-00, Internet Engineering Task Force, March 2016.

[Ber] Daniel J. Bernstein. Cache-Timing Attacks on AES.

[CGBS16] Henry Corrigan-Gibbs, Dan Boneh, and Stuart Schechter. Balloon hashing: Provably space-hard hash functions with data-independent access patterns. Cryptology ePrint Archive, Report 2016/027, Version: 20160601:225540, 2016. `http://eprint.iacr.org/`.

[CJMS14] Donghoon Chang, Arpan Jati, Sweta Mishra, and Somitra Kumar Sanadhya. Rig: A simple, secure and flexible design for password hashing version 2.0. 2014.

[Cox14] Bill Cox. Twocats (and skinnycat): A compute time and sequential memory hard password hashing scheme. *Password Hashing Competition. v0 edn.*, 2014.

[EGS75] Paul Erdös, Ronald L. Graham, and Endre Szemeredi. On sparse graphs with dense long paths. Technical report, Stanford, CA, USA, 1975.

[FLW13] Christian Forler, Stefan Lucks, and Jakob Wenzel. Catena: A memory-consuming password scrambler. *IACR Cryptology ePrint Archive*, 2013:525, 2013.

[MMV13] Mohammad Mahmoody, Tal Moran, and Salil P. Vadhan. Publicly verifiable proofs of sequential work. In Robert D. Kleinberg, editor, *Innovations in Theoretical Computer Science, ITCS '13, Berkeley, CA, USA, January 9-12, 2013*, pages 373–388. ACM, 2013.

[Per09] C. Percival. Stronger key derivation via sequential memory-hard functions. In *BSDCan 2009*, 2009.

[PHC] Password hashing competition. `https://password-hashing.net/`.

[Pin14] Krisztián Pintér. Gambit – A sponge based, memory hard key derivation function. Submission to Password Hashing Competition (PHC), 2014.

[PR80] Wolfgang J. Paul and Rüdiger Reischuk. On alternation II. A graph theoretic approach to determinism versus nondeterminism. *Acta Inf.*, 14:391–403, 1980.

[Sch82] Georg Schnitger. A family of graphs with expensive depth reduction. *Theor. Comput. Sci.*, 18:89–93, 1982.

[Sch83] Georg Schnitger. On depth-reduction and grates. In *24th Annual Symposium on Foundations of Computer Science, Tucson, Arizona, USA, 7-9 November 1983*, pages 323–328. IEEE Computer Society, 1983.

[Val77] Leslie G. Valiant. Graph-theoretic arguments in low-level complexity. In Jozef Gruska, editor, *Mathematical Foundations of Computer Science 1977, 6th Symposium, Tatranska Lomnica, Czechoslovakia, September 5-9, 1977, Proceedings*, volume 53 of *Lecture Notes in Computer Science*, pages 162–176. Springer, 1977.

[Wu15] Hongjun Wu. POMELO – A Password Hashing Algorithm, 2015.

A   Missing Proofs

Reminder of Claim 12.   *For each $x, y \in G_n$ with $y > x$ and node i in metanode y, there exists an edge from the last third of metanode x to node i with probability at least $\frac{1}{12\sqrt{y}\sqrt{y-x+1}}$.*

Proof of Claim 12:   Recall that for node i, Argon2iB creates an edge from i to parent node $i\left(1 - \frac{k^2}{N^2}\right)$, where $k \in [N]$ is picked uniformly at random. Thus, for $i, j \in G$ with $i > j$, there exists an edge from node j to i with probability at least

$$
\begin{aligned}
\mathbf{Pr}\left[(x-1)m + 1 \le i\left(1 - \frac{k^2}{N^2}\right) \le \left(x - 1 + \frac{1}{3}\right)m\right] &= \mathbf{Pr}\left[\left(x - 1 + \frac{1}{6}\right)m \le ym\left(1 - \frac{k^2}{N^2}\right) \le \left(x - 1 + \frac{1}{3}\right)m\right] \\
&\ge \mathbf{Pr}\left[\frac{y - x + \frac{2}{3}}{y} \ge \frac{k^2}{N^2} \ge \frac{y - x + \frac{5}{6}}{y}\right] \\
&\ge \sqrt{\frac{y - x + \frac{2}{3}}{y}} - \sqrt{\frac{y - x + \frac{5}{6}}{y}} \\
&\ge \frac{1}{6\sqrt{y}(2\sqrt{y - x + 1})} = \frac{1}{12\sqrt{y}\sqrt{y - x + 1}}.
\end{aligned}
$$

$\square$

Reminder of Claim 13.   *For any two metanodes $x, y \in G_m$ with $x < y$, the last third of x is connected to the first third of y with probability at least $\frac{m\sqrt{m}}{m\sqrt{m} + 36\sqrt{n(y-x+1)}}$.*

Proof of Claim 13:   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^{th} node of metanode y is the first node in y to which there exists an edge from the last third of metanode x, so that by Claim 12, $\mathbf{Pr}\left[E_1\right] \ge \frac{1}{12\sqrt{y}\sqrt{y-x+1}}$. Note that furthermore, $\mathbf{Pr}\left[E_i\right] \ge \frac{1}{12\sqrt{y}\sqrt{y-x+1}}(1 - p)$. Thus,

$$
\begin{aligned}
p &= \mathbf{Pr}\left[E_1\right] + \mathbf{Pr}\left[E_2\right] + \ldots + \mathbf{Pr}\left[E_{m/3}\right] \\
&\ge \left(\frac{m}{3}\right)\frac{1}{12\sqrt{y}\sqrt{y - x + 1}}(1 - p).
\end{aligned}
$$

Setting $\alpha = \left(\frac{m}{3}\right)\frac{1}{12\sqrt{y}\sqrt{y-x+1}}$, then it follows that $p + \alpha p \ge \alpha$, so that $p \ge \frac{\alpha}{1+\alpha}$. Since $y \le \frac{n}{m}$,

$$
p \ge \frac{m/36}{\sqrt{y(y - x + 1)} + m/36} \ge \frac{m\sqrt{m}}{m\sqrt{m} + 36\sqrt{n(y - x + 1)}}
$$

$\square$

Reminder of Lemma 14.   *Let $m = n/(20000e)$ then for $r^* = \tilde{\Omega}(e^3/n^2) = \tilde{\Omega}(n/m^3)$ the metagraph G_m (for Argon2i) is a (δ, r^*)-local expander with high probability.*

Proof of Lemma 14:   Let $r \ge r^*$ and $A \subseteq \{x, \ldots, x + r - 1\}$, $B \subseteq \{x + r, \ldots, x + 2r - 1\}$ be subsets of size δr, for some $x \le n - 2r + 1$. By Stirling's approximation,

$$
\sqrt{2\pi}r^{r+1/2}e^{-r} \le r! \le er^{r+1/2}e^{-r}.
$$

15

Then it follows that

$$
\begin{aligned}
\binom{r}{\delta r} &\leq \frac{e r^{r+1/2} e^{-r}}{2\pi (\delta r)^{\delta r+1/2} (r-\delta r)^{r-\delta r+1/2} e^{-r}} \\
&\leq \frac{e}{2\pi \delta^{\delta r+1/2}(1-\delta)^{r-\delta r+1/2}\sqrt{r}} \\
&= \frac{e^{1+\delta r \log \frac{1}{\delta}+(r-\delta r)\log \frac{1}{1-\delta}}}{2\pi\sqrt{r\delta(1-\delta)}}
\end{aligned}
$$

For two specific metanodes in A and B, the probability the pair is connected is at least $\frac{m\sqrt{m}}{m\sqrt{m}+36\sqrt{nr}}$ by Claim 13. For $36\sqrt{nr} \geq m\sqrt{m}$, the probability is at least $\frac{m\sqrt{m}}{72\sqrt{nr}}$ (otherwise, for $36\sqrt{nr} < m\sqrt{m}$, the probability is at least $\frac{1}{2} > \frac{m\sqrt{m}}{72\sqrt{nr}}$). Now, let p be the probability that there exists an edge from A to a specific metanode in B. Furthermore, let E_i be the event that the i^{th} metanode of A is the first node from which there exists an edge from a specific metanode of B, so that, $\mathbf{Pr}[E_1] \geq \frac{m\sqrt{m}}{72\sqrt{nr}}$. Note that furthermore, $\mathbf{Pr}[E_i] \geq \frac{m\sqrt{m}}{72\sqrt{nr}}(1-p)$. Thus,

$$
\begin{aligned}
p &= \mathbf{Pr}[E_1] + \mathbf{Pr}[E_2] + \ldots + \mathbf{Pr}\left[E_{|A|}\right] \\
&\geq (\delta r)\frac{m\sqrt{m}}{72\sqrt{nr}}(1-p).
\end{aligned}
$$

Since $r \geq r^* = \tilde{\Omega}(n/m^3)$, it follows for an appropriate choice of r' that $p \geq \sqrt{\log n}(1-p)$. Thus, $p \geq \frac{\sqrt{\log n}}{1+\sqrt{\log n}}$ is the probability that there exists an edge from A to a specific metanode in B.

Now, taking the probability over all δr metanodes in B, the probability that A and B are not connected is at most

$$
\begin{aligned}
(1-p)^{\delta r} &= \left(\frac{1}{1+\sqrt{\log n}}\right)^{\delta r} \\
&= e^{-\delta r \log(1+\sqrt{\log n})}
\end{aligned}
$$

Since there $\binom{r}{\delta r}^2$ such sets A and B, the probability that there exists A and B in the above intervals which are not connected by an edge is at most

$$
e^{-\delta r \log(1+\sqrt{\log n})}\binom{r}{\delta r}^2
$$

by a simple union bound. Then from the above Stirling approximation, the probability is at most

$$
\exp\left(2 + 2\delta r \log \frac{1}{\delta} + 2(r-\delta r)\log \frac{1}{1-\delta} - \delta r \log(1+\sqrt{\log n})\right)\frac{1}{4\pi^2 r\delta(1-\delta)} \ ,
$$

where $-\delta r \log(1+\sqrt{\log n})$ is the dominant term in the exponent. Again taking $r \geq r^* = \Omega\left(\frac{n\log n}{m^3}\right)$, the probability that G_m is not a δ-local expander is at most

$$
\begin{aligned}
\mathbf{Pr}[\exists r \geq r^*, x, A, B \text{ with no edge}] &\leq n \sum_{r \geq r^*} \frac{e^{-\Omega(r \log \log n)}}{4\pi^2 r \delta(1-\delta)} \\
&= o\left(\frac{1}{n}\right).
\end{aligned}
$$

16

Thus, G_m is a δ-local expander with high probability. $\qquad\square$

Reminder of Claim 15. *If $|S| < n/(10000m)$ then at least half of the layers $L_1, L_2, \ldots L_{n/(mr^*)}$ are $(1/1000)$-good with respect to S.*

Proof of Claim 15: Let i_1 be the index of the first layer L_{i_1} such that for some $x_1 > 0$ we have $\left|S \cap \left(\bigcup_{t=i_1}^{i_1+x_1-1} L_t\right)\right| \geq c\left|\left(\bigcup_{t=i}^{i_1+x_1-1} L_t\right)\right|$. Once $i_1 < \ldots < i_{j-1}$ and $x_1, \ldots, x_{j-1}$ have been defined we let i_j be the least layer such that $i_j > i_{j-1} + x_{j-1}$ and there exists $x_j > 0$ such that $\left|S \cap \left(\bigcup_{t=i_j}^{i_j+x_j-1} L_t\right)\right| \geq c\left|\left(\bigcup_{t=i_j}^{i_j+x_j-1} L_t\right)\right|$ (assuming that such a pair i_j, x_j exists). Let $i_1 + x_1 < i_2$, $i_2 + x_2 < i_3, \ldots i_{k-1} + x_{k-1} < i_k$ denote a maximal such sequence and let

$$F = \bigcup_{t=1}^{k} [i_t, x_t - 1] \ .$$

Observe that by construction of F we have $|S| \geq c\left|\bigcup_{j \in F} L_j\right| = c|F|r^*$, which means that $|F| \leq |S|/(cr^*) = n/(10000cmr^*)$. Similarly, we can define a maximal sequence $i_1^* > \ldots > i_{k^*}^*$ such that $i_j^* - x_j^* > i_{j+1}^*$ and $\left|S \cap \left(\bigcup_{t=i_j^*-x_j^*+1}^{i_j^*} L_t\right)\right| \geq c\left|\left(\bigcup_{t=i_j^*-x_j^*+1}^{i_j^*} L_t\right)\right|$ for each j. A similar argument shows that $|B| \leq |S|/(cr^*) = n/(10000cmr^*)$, where $B = \bigcup_{t=1}^{k} [i_t^* - x_t^* + 1, i_t^*] \ .$ Finally, we note that if L_i is not c-good then $i \in F \cup B$. Thus, at most $n/(5000cmr^*)$ layers are not c-good, which means that the number of $c = (1/1000)$-good layers is at least

$$\frac{n}{mr^*} - \frac{n}{5mr^*} \geq \frac{n}{2mr^*} \ .$$

$\qquad\square$

Reminder of Lemma 17. *Suppose that G_m is a (δ, r^*)-local expander with $\delta = 1/16$ and let $S \subseteq V(G_m)$ be given such that $|S| \leq n/(10000m)$. Then, $R_{i,S} \geq 7r^*/8$.*

Proof of Lemma 17: We prove by induction. For the base case, we set $R_1 = H_{1,S} - S$. Thus, $|R_1| = |H_{1,S} - S| \geq r^* - (1/1000)r^*$, since $H_{1,S}$ is $(1/1000)$-good with respect to S.

Now, suppose that $|R_j| \geq 7r^*/8$ for each $j \leq i$. If layers $H_{i,S}$ and $H_{i+1,S}$ are within 100 intermediate layers, then since $H_{i,S}$ is $(1/1000)$-good with respect to S, it follows that at most $100/1000 = 1/10$ of the nodes in $H_{i+1,S}$ are also in S. Moreover, since G_m is a (δ, r^*)-local expander with $\delta = 1/16$, then at most δr^* additional nodes in $H_{i+1,S}$ are not reachable from $H_{i,S}$. Therefore,

$$|R_{i+1,S}| \geq |H_{i+1,S} - S| - \delta r^* \geq (1 - 1/1000 - 1/16)r^* \geq (7/8)r^* \ .$$

Otherwise, suppose more than 100 intermediate layers separate layers $H_{i,S}$ and $H_{i+1,S}$. Let $Y_1, \ldots, Y_k$ denote the intermediate layers between $H_{i,S}$ and $H_{i+1,S}$, so that $k > 100$. Let j be the integer such that $2^j \leq k < 2^{j+1}$. Since $R_{i,S}$ is $(1/1000)$-good with respect to S, at most $2^{j+1}r^*/1000$ nodes in $Y_1 \cup \ldots \cup Y_k$ can be in S. Thus, at least $(1/8)$-fraction of the nodes in $Y_{k-2^{j-1}}, \ldots, Y_{k-2^{j-2}+1}$ are reachable from R_i. We now show this is sufficient.

Suppose that at least $(1/8)$-fraction of the nodes in $Y_{k-2u}, \ldots, Y_{k-u-1}$ are reachable from R_i. Then at least $(7/8)$-fraction of nodes in $Y_{k-u}, \ldots, Y_{k-u/2}$ are reachable from R_i, since layer H_{i+1} is both $(1/1000)$-good and a (δ, r^*)-local expander with $\delta = 1/16$. (Note: we are now using layer H_{i+1}, *not layer H_i*). It follows that at least $(7/8)$-fraction of the nodes in Y_k are reachable from R_i. Again,

$$|R_{i+1,S}| \geq |H_{i+1,S} - S| - \delta r^* \geq (1 - 1/1000 - 1/16)r^* \geq (7/8)r^* \ .$$

Thus, at least $(7/8)$-fraction of the nodes in H_{i+1} are reachable, and so $|R_{i+1,S}| \geq (7/8)r^*$. $\quad\square$

Reminder of Claim 21.   *Let $i, j \in [n]$ be given $(i \neq j)$ and let G be a random Argon2iB DAG on n nodes. There exists an edge from node j to i in G with probability at least $\frac{1}{4n}$.*

Proof of Claim 21:   Recall that for node i, Argon2iB creates an edge from i to parent node $i\left(1 - \frac{x^2}{2^{64}}\right)$, where $x \in \left[2^{64}\right]$ is picked uniformly at random. Thus, for $i, j \in G$ with $i > j$, there exists an edge from node j to i with probability at least

$$\mathbf{Pr}\left[j \leq i\left(1 - \frac{x^2}{2^{64}}\right) \leq j + \frac{1}{2}\right] = \mathbf{Pr}\left[\frac{i-j}{i} \geq \frac{x^2}{2^{64}} \geq \frac{i-j-\frac{1}{2}}{i}\right]$$

$$\geq \mathbf{Pr}\left[1 \geq \frac{x^2}{2^{64}} \geq 1 - \frac{1}{2n}\right]$$

$$\geq \frac{1}{4n}.$$

$\square$