

Data-Independent Memory Hard Functions: New Attacks and Stronger Constructions

Jeremiah Blocki¹, Ben Harsha¹, Siteng Kang²,
Seunghoon Lee¹, Lu Xing¹, and Samson Zhou³

¹ Purdue University

² Penn State University

³ Indiana University

Abstract. Memory-hard functions (MHFs) are a key cryptographic primitive underlying the design of moderately expensive password hashing algorithms and egalitarian proofs of work. Over the past few years several increasingly stringent goals for an MHF have been proposed including the requirement that the MHF have high sequential space-time (ST) complexity, parallel space-time complexity, amortized area-time (aAT) complexity and sustained space complexity. Data-Independent Memory Hard Functions (iMHFs) are of special interest in the context of password hashing as they naturally resist side-channel attacks. iMHFs can be specified using a directed acyclic graph (DAG) G with $N = 2^n$ nodes and low indegree and the complexity of the iMHF can be analyzed using a pebbling game. Recently, Alwen et al. [ABH17] constructed a DAG called DRSSample that has aAT complexity at least $\Omega(N^2/\log N)$. Asymptotically DRSSample outperformed all prior iMHF constructions including Argon2i, winner of the password hashing competition (aAT cost $\mathcal{O}(N^{1.767})$), though the constants in these bounds are poorly understood. We show that the greedy pebbling strategy of Boneh et al. [BCS16] is particularly effective against DRSSample e.g., the aAT cost is $\mathcal{O}(N^2/\log N)$. In fact, our empirical analysis *reverses* the prior conclusion of Alwen et al. that DRSSample provides stronger resistance to known pebbling attacks for practical values of $N \leq 2^{24}$. We construct a new iMHF candidate (DRSSample+BRG) by using the bit-reversal graph to extend DRSSample. We then prove that the construction is asymptotically optimal under every MHF criteria, and we empirically demonstrate that our iMHF provides the best resistance to *known* pebbling attacks. For example, we show that any parallel pebbling attack either has aAT cost $\omega(N^2)$ or requires at least $\Omega(N)$ steps with $\Omega(N/\log N)$ pebbles on the DAG. This makes our construction the first practical iMHF with a strong sustained space-complexity guarantee and immediately implies that any parallel pebbling has aAT complexity $\Omega(N^2/\log N)$. We also prove that any sequential pebbling (including the greedy pebbling attack) has aAT cost $\Omega(N^2)$ and, if a plausible conjecture holds, any parallel pebbling has aAT cost $\Omega(N^2 \log \log N / \log N)$ — the best possible bound for an iMHF. We implement our new iMHF and demonstrate that it is just as fast as Argon2. Along the way we propose a simple modification to the Argon2 round function that increases an attacker’s aAT cost by nearly an order of magnitude without increasing running time on a CPU. Finally, we give a pebbling reduction that proves that in the parallel random oracle model (PROM) the cost of evaluating an iMHF like Argon2i or DRSSample+BRG is given by the pebbling cost of the underlying DAG. Prior pebbling reductions assumed that the iMHF round function concatenates input labels before hashing and did not apply to practical iMHFs such as Argon2i, DRSSample or DRSSample+BRG where input labels are instead XORed together.

1 Introduction

Memory Hard Functions (MHFs) are a key cryptographic primitive in the design of password hashing, algorithms and egalitarian proof of work puzzles [Lee11]. In the context of password hashing we want to ensure that the function can be computed reasonably quickly on standard hardware, but that it is prohibitively expensive to evaluate the function millions or billions of times. The first property ensures that legitimate users can authenticate reasonably quickly, while the purpose of the latter goal is to protect low-entropy secrets (e.g., passwords, PINs, biometrics) against brute-force offline guessing attacks. One of the challenges is that the attacker might attempt to reduce computation costs by employing customized hardware such as a Field Programmable Gate Array (FPGA) or an Application Specific Integrated Circuit (ASIC). MHFs were of particular interest in the 2015 Password Hashing Competition [PHC16], where the winner, Argon2 [BDK16], and all but one finalists [FLW14, SAA⁺15, Pes14] claimed some form of memory hardness.

Wiener [Wie04] defined the full cost of an algorithm’s execution to be the number of hardware components multiplied by the duration of their usage e.g., if the algorithm needs to allocate $\Omega(N)$ blocks of memory for $\Omega(N)$ time steps then full evaluation costs would scale quadratically. At an intuitive level, a strong MHF $f(\cdot)$ should have the property that the *full cost* [Wie04] of evaluation grows as fast as possible in the running time parameter N . Towards this end, a number of increasingly stringent security criteria have been proposed for a MHF including sequential space-time complexity, parallel space-time complexity, amortized area-time complexity (aAT) and sustained space-complexity. The sequential (resp. parallel) space-time complexity of a function $f(\cdot)$ measures the space-time cost of the best sequential (resp. parallel) algorithm evaluating $f(\cdot)$ i.e., if a computation runs in time t and requires space s then the space-time cost is given by the product st . The requirement that a hash function has high space-time complexity rules out traditional hash iteration based key-derivation functions like PBKDF2 and bcrypt as both of these functions be computed in linear time $\mathcal{O}(N)$ and constant space $\mathcal{O}(1)$. Blocki et al. [BHZ18] recently presented an economic argument that algorithms with low space-time complexity such as bcrypt and PBKDF2 are no longer suitable to protect low-entropy secrets like passwords i.e., one cannot provide meaningful protection against a rational attacker with customized hardware (FPGA, ASIC) without introducing an unacceptably long authentication delay. By contrast, they argued that MHFs with true cost $\Omega(N^2)$ can ensure that a rational attacker will quickly give up since marginal guessing costs are substantially higher.

The Catena-Bit Reversal MHF [FLW14] has provably optimal sequential space-time complexity $\Omega(N^2)$ — the space-time complexity of any sequential algorithm running in time N is at most $\mathcal{O}(N^2)$ since at most N blocks of memory can be allocated in time N . However, Alwen and Serbinenko [AS15] showed that the *parallel space-time* complexity of this MHF is just $\mathcal{O}(N^{1.5})$. Even parallel space-time complexity has limitations in that it does not amortize nicely. The stronger notion of Amortized Area-Time (aAT) complexity (and the asymptotically equivalent notion of cumulative memory complexity (cmc)) measures the amortized cost of any parallel algorithm evaluating the function $f(\cdot)$ on m distinct inputs. Alwen and Serbinenko [AS15] gave a theoretical example of a function $f(\cdot)$ with the property that the amortized space-time cost of evaluating the function on $m = \sqrt{N}$ distinct inputs is approximately m times cheaper than the parallel space-time cost i.e., evaluating the function on the last $m-1$ inputs is essentially free. This is problematic in the context of password hashing where the attacker wants to

compute the function $f(\cdot)$ multiple times i.e., on each password in a cracking dictionary. The amortization issue is not merely theoretical. Indeed, the aAT complexity of many MHF candidates is significantly lower than $\mathcal{O}(N^2)$ e.g., the aAT complexity of Balloon Hash [BCS16] is just $\mathcal{O}(N^{5/3})$ [AB16, ABP17] and for password hashing competition winner Argon2i [BDK16] the aAT cost is at most $\mathcal{O}(N^{1.767})$ [AB16, AB17, ABP17, BZ17].

The script MHF, introduced by Percival in 2009 [Per09], was proven to have cmc/aAT complexity $\Omega(N^2)$ in the random oracle model [ACP⁺17]. However, it is possible for an script attacker to achieve any space-time trade-off subject to the constraint that $st = \Omega(N^2)$ without penalty e.g., an attacker could evaluate script in time $t = \Omega(N^2)$ with space $s = \mathcal{O}(1)$. Alwen et al. [ABP18] argued that this flexibility potentially makes it easier to develop ASICs for script, and proposed the even stricter MHF requirement of *sustained space complexity*, which demands that any (parallel) algorithm evaluating the function $f(\cdot)$ requires at least t time steps in which the space usage is $\geq s$ — this implies that $\text{aAT} \geq st$. Alwen et al. [ABP18] provided a theoretical construction of a MHF with maximal sustained space complexity i.e., evaluation requires space $s = \Omega(N/\log N)$ for time $t = \Omega(N)$. However, there are no practical constructions of MHFs that provide strong guarantees with respect to sustained space complexity.

Data-Independent vs Data-Dependent Memory Hard Functions. Memory Hard Functions can be divided into two categories: Data-Independent Memory Hard Functions (iMHFs) and Data-Dependent Memory Hard Functions (dMHFs). Examples of dMHFs include script [Per09], Argon2d [BDK16] and Boyen’s halting puzzles [Boy07]. Examples of iMHFs include Password Hashing Competition (PHC) [PHC16] winner Argon2i [BDK16], Balloon Hashing [BCS16] and DRSSample [ABH17]. In this work we primarily focus on the design and analysis of secure iMHFs. iMHFs are designed to resist certain side-channel attacks e.g., cache timing [Ber05] by requiring that the induced memory access pattern does not depend on the (sensitive) input e.g., the user’s password. By contrast, the induced memory access for a dMHFs is allowed to depend on the function input.

Alwen and Blocki [AB16] proved that *any* iMHF has aAT complexity at most $\mathcal{O}(N^2 \log \log N / \log N)$, while the dMHF script provably has aAT complexity $\Omega(N^2)$ in the random oracle model — a result which cannot be matched by any iMHF. However, the aAT complexity of a dMHF may be greatly reduced after a side-channel attack. If a brute-force attacker is trying to find $x \leq m$ s.t. $f(x) = y$ and the attacker also has learned the correct memory access pattern induced by the real input x^* (e.g., via a side-channel attack) then the attacker can quit evaluation $f(x)$ immediately once it is clear that the induced memory access pattern on input $x \neq x^*$ is different. For example, the aAT complexity of script (resp. [Boy07]) after a side-channel attack is just $\mathcal{O}(N)$ (resp. $\mathcal{O}(1)$).

Hybrid Modes. Alwen and Blocki [AB16, AB17] showed that the aAT complexity of most iMHFs was significantly lower than one would hope, but their techniques do not extend to MHFs. In response, the Argon2 spec [KDBJ17] was updated to list Argon2id as the recommended mode of operation for password hashing instead of the purely data-independent mode Argon2i. Hybrid independent-dependent (id) modes, such as Argon2id [KDBJ17], balance side-channel resistance with high aAT complexity by running the MHF in data-independent mode for $N/2$ steps before switching to data-dependent mode for the final $N/2$ steps. If there is a side-channel attack then security reduces to that of the underlying iMHF (e.g., Argon2i), and if there is no side-channel attack then the function is expected to have optimal aAT complexity $\Omega(N^2)$. We remark that, even for a hybrid mode, it is important to ensure that the

underlying iMHF is as strong as possible a side-channel attack on a hybrid “id” mode of operation will reduce security to that of the underlying iMHF.

1.1 Related Work

MHF Goals. Dwork et al. and Abadi et al. [DGN03,ABMW05] introduced the notion of a memory-bound function where we require that *any* evaluation algorithm results in a large number of cache-misses. Ren and Devadas recently introduced a refinement to this notion called bandwidth-hardness [RD17]. To the best of our knowledge Percival was the first to propose the goal that a MHF should have high space-time complexity [Per09] though Boyen’s dMHF construction appears to achieve this goal [Boy07] and the notion of space-time complexity is closely related to the notion of “full cost” proposed by Wiener [Wie04]. Metrics like space-time complexity and Amortized Area-Time Complexity [AS15,ABH17] aim to capture the cost of the hardware (e.g., DRAM chips) the attacker must purchase to compute an MHF — amortized by the number of MHF instances computed over the lifetime of the hardware components. By contrast, bandwidth hardness [RD17] aims to capture the *energy cost* of the electricity required to compute the MHF once. If the attacker uses an ASIC to compute the function then the *energy* expended during computation will typically be small in comparison with the energy expended during a cache-miss. Thus, a bandwidth hard function aims to ensure that *any* evaluation strategy either results in $\Omega(N)$ cache-misses or $\omega(N)$ evaluations of the hash function.

In the full version [BHK⁺18] we argue that, in the context of password hashing, aAT complexity is more relevant than bandwidth hardness because the “full cost” [Wie04] can scale quadratically in the running time parameter N . However, one would ideally want to design a MHF that has high aAT complexity and is also maximally bandwidth hard. Blocki et al. [BRZ18] recently showed that any MHF with high aAT complexity is at least somewhat bandwidth hard. Furthermore, all practical iMHFs (including Catena-Bit Reversal [FLW14], Argon2i and DRSSample) are maximally bandwidth hard [RD17,BRZ18], including our new construction DRS+BRG.

Graph Pebbling and iMHFs. An iMHF $f_{G,H}$ can be viewed as a mode of operation over a directed acyclic graph (DAG) $G = (V = [N], E)$ that encodes data-dependencies (because the DAG is static the memory access pattern will be identical for all inputs) and a compression function $H(\cdot)$. Alwen and Serbinenko [AS15] defined $f_{G,H}(x) = \text{lab}_{G,H,x}(N)$ to be the label of the last node in the graph G on input x . Here, the label of the first node $\text{lab}_{G,H,x}(1) = H(1,x)$ is computed using the input x and for each internal node v with $\text{parents}(v) = v_1, \dots, v_\delta$ we have

$$\text{lab}_{G,H,x}(v) = H(v, \text{lab}_{G,H,x}(v_1), \dots, \text{lab}_{G,H,x}(v_\delta)) .$$

In practice, one requires that the maximum indegree is constant $\delta = \mathcal{O}(1)$ so that the function $f_{G,H}$ can be evaluated in sequential time $\mathcal{O}(N)$. Alwen and Serbinenko [AS15] proved that the cmc complexity (asymptotically equivalent to aAT complexity) of the function $f_{G,H}$ can be fully described in terms of the black pebbling game — defined later in Section 2.2. The result is significant in that it reduces the complex task of building an iMHF with high aAT complexity to the (potentially easier) task of constructing a DAG with maximum pebbling cost. In particular, Alwen and Serbinenko showed that any algorithm evaluating the function $f_{G,H}$ in the parallel random oracle model *must* have cumulative memory cost at least $\Omega\left(w \times \Pi_{cc}^{\parallel}(G)\right)$, where $\Pi_{cc}^{\parallel}(G)$ is the cumulative

pebbling cost of G (defined in Section 2.2), $H: \{0,1\}^* \rightarrow \{0,1\}^w$ is modeled as a random oracle and $w = |H(z)|$ is the number of output bits in a single hash value. Similar, pebbling reductions have been given for bandwidth hardness [BRZ18] and sustained space complexity [ABP18] using the same labeling rule.

While these pebbling reductions are useful in theory, *practical* iMHF implementations do not use the labeling rule proposed in [AS15]. In particular, Argon2i, DRSample and our own iMHF implementation (DRSample+BRG) all use the following labeling rule

$$\text{lab}_{G,H,x}(v) = H(\text{lab}_{G,H,x}(v_1) \oplus \dots \oplus \text{lab}_{G,H,x}(v_\delta)) ,$$

where $v_1, \dots, v_\delta = \text{parents}(v)$ and the DAGs have indegree $\delta=2$. The XOR labeling rule allows one to work with a faster round function $H: \{0,1\}^w \rightarrow \{0,1\}^w$ e.g., Argon2i builds $H: \{0,1\}^{8192} \rightarrow \{0,1\}^{8192}$ using the Blake2b permutation function and DRSample(+BRG) uses the same labeling rule as Argon2i. When we define $f_{G,H}$ using the above, the pebbling reduction of [AS15] no longer applies. Thus, while we know that the pebbling cost of DRSample (resp. Argon2i) is $\Omega(N^2/\log N)$ [ABH17] (resp. $\tilde{\Omega}(N^{1.75})$ [BZ17]), technically it had never been proven that DRSample (resp. Argon2i) has aAT complexity $\Omega(wN^2/\log N)$ (resp. $\tilde{\Omega}(wN^{1.75})$) in the parallel random oracle model.

Argon2i and DRSample. Arguably, two of the most significant iMHFs candidates are Argon2i [BDK16] and DRSample [ABH17]. Argon2i was the winner of the recently completed password hashing competition [PHC16] and DRSample [ABH17] was the first *practical* construction of an iMHF with aAT complexity proven to be *at least* $\Omega(N^2/\log N)$ in the random oracle model. In an asymptotic sense this upper bound almost matches the general upper bound $\mathcal{O}(N^2 \log \log N / \log N)$ on the aAT cost of any iMHF established by Alwen and Blocki [AB16]. A recent line of research [AB16, AB17, ABP17, BZ17] has developed theoretical depth-reducing attacks on Argon2i showing that the iMHF has aAT complexity *at most* $\mathcal{O}(N^{1.767})^4$. The DRSample [ABH17] iMHF modifies the edge distribution of the Argon2i graph to ensure that the underlying directed acyclic graph (DAG) satisfies a combinatorial property called depth-robustness, which is known to be *necessary* [AB16] and *sufficient* [ABP17] for developing an MHF with high aAT complexity.

While the aAT complexity of DRSample is at least $c_1 N^2 / \log N$ for some constant c_1 , the *constant* c in this lower bound is poorly understood — Alwen et al. [ABH17] only proved the lower bound when $c_1 \approx 7 \times 10^{-6}$. Similarly, Argon2i has aAT complexity at least $c_2 N^{1.75} / \log N$ [BZ17] though the constants from this lower bound are also poorly understood⁵. On the negative side the asymptotic lower bounds do not *absolutely* rule out the possibility of an attack that reduces aAT complexity by several orders of magnitude. Alwen et al. [ABH17] also presented an empirical analysis of the aAT cost of DRSample and Argon2i by measuring the aAT cost of these functions against a wide battery of pebbling attacks [AB16, ABP17, AB17]. The results of this empirical analysis were quite positive for DRSample and indicated that DRSample was not only stronger in an asymptotic sense, but that it also provided greater resistance to other pebbling attacks than other iMHF candidates like Argon2i in practice.

Boneh et al. [BCS16] previously presented a greedy pebbling attack that reduced the pebbling cost of Argon2i by a moderate constant factor of 4 to 5. The greedy pebbling attack does not appear to have been included in the empirical analysis of Alwen et

⁴ This latest attack almost matches the *lower bound* of $\tilde{\Omega}(N^{1.75})$ on the aAT complexity of Argon2i.

⁵ Blocki and Zhou did not explicitly work out the constants in their lower bound, but it appears that $c_2 \approx 5 \times 10^{-7}$ [ABH17].

al. [ABH17]. In a strict asymptotic sense the depth-reducing attacks of Alwen and Blocki [AB16, AB17] achieved more substantial $\Omega(N^{0.2+})$ -factor reductions in pebbling cost, which may help to explain the omission of the greedy algorithm in [ABH17]. Nevertheless, it is worth noting that the greedy pebbling strategy is a simple sequential pebbling strategy that would be easy to implement in practice. By contrast, there has been debate about the *practical feasibility* of implementing the more complicated pebbling attacks of Alwen and Blocki [AB16] (Alwen and Blocki [AB17] argued that the attacks do not require unrealistic parallelism or memory bandwidth, but to the best of our knowledge the attacks have yet to be implemented on an ASIC).

1.2 Contributions

Stronger Attacks. We present a theoretical and empirical analysis of the greedy pebbling attack [BCS16] finding that DRSample has aAT complexity at most $\lesssim N^2/\log N$. The greedy pebbling attack that achieves this bound is *sequential*, easy to implement and achieves high attack quality even for practical values of N . In fact, for *practical* values of $N \leq 2^{24}$ we show that DRSample is *more* vulnerable to *known pebbling attacks* than Argon2i, which *reverses* previous conclusions about the *practical* security of Argon2i and DRSample [ABH17]. We next consider a defense proposed by Biryukov et al. [BDK16] against the greedy pebbling attack, which we call the XOR-extension gadget. While this defense defeats the *original* greedy pebbling attack [BCS16], we found a simple generalization of the greedy pebbling attack that thwarts this defense. We also use the greedy pebbling attack to prove that *any* DAG with indegree two has a sequential pebbling with aAT cost $\lesssim \frac{N^2}{4}$.

We also develop a *novel* greedy algorithm for constructing depth-reducing sets, which is the critical first step in the parallel pebbling attacks of Alwen and Blocki [AB16, AB17]. Empirical analysis demonstrates that this greedy algorithm constructs *significantly smaller* depth-reducing sets than previous state of the art techniques [AB16, AB17, ABH17], which leads to higher quality attacks [AB16] and leaving us in an uncomfortable situation where there high quality pebbling attacks against all iMHF candidates e.g., DRSample is susceptible to the greedy pebbling attack while Argon2i is susceptible to depth-reducing attacks [AB16, AB17, ABH17].

New iMHF Candidate with Optimal Security. We next develop a new iMHF candidate DRSample+BRG by overlaying a bit-reversal graph [LT82, FLW14] on top of DRSample, and analyze the new DAG empirically and theoretically. Interestingly, while *neither* DAG (DRSample or BRG) is known to have strong sustained space complexity guarantees, we can prove that *any* parallel pebbling either has maximal sustained space complexity (meaning that there are at least $\Omega(N)$ steps with $\Omega(N/\log N)$ pebbles on the DAG) or has aAT cost at least $\omega(N^2)$. This makes our construction the first practical construction with strong guarantees on the sustained space-complexity — prior constructions of Alwen et al. [ABP18] were theoretical. DRSample+BRG is asymptotically optimal with respect to all proposed MHF metrics including bandwidth hardness (*both* BRG and DRSample are bandwidth hard [RD17, BRZ18]) and aAT complexity (inherited from DRSample [ABH17]). We also show that our construction optimally resists the greedy attack and *any* extensions. In particular, we prove sequential pebbling of the bit-reversal graph has cumulative memory cost (cmc) and aAT cost at least $\Omega(N^2)$. This result generalizes a well-known result that the bit-reversal graph has sequential space-time cost

$\Omega(N^2)$ and may be of independent interest e.g., it demonstrates that Password Hashing Competition Finalist Catena-BRG [FLW14] is secure against *all* sequential attacks.

Our empirical analysis indicates that DRSample+BRG offers strong resistance to *all known* attacks, including the greedy pebbling attack, depth-reducing attacks and several other novel attacks introduced in this paper. In particular, even for very large $N=2^{24}$ (2^{24} 1KB blocks = 16GB) the *best* attack had aAT cost over $\frac{N^2}{11}$ — for comparison *any* DAG with indegree two has aAT cost $\lesssim \frac{N^2}{4}$.

We also show that the aAT/cmc of DRSample+BRG is *at least* $\Omega(N^2 \log \log N / \log N)$ under a plausible conjecture about the depth-robustness of DRSample. As evidence for our conjecture we analyze three state-of-the-art approaches for constructing a depth-reducing set, including the layered attack [AB16], Valiant’s Lemma [AB16, Val77] and the reduction of Alwen et al. [ABP17], which can transform any pebbling with low aAT cost (e.g., the Greedy Pebbling Attack) into a depth-reducing set. We show that each attack fails to refute our conjecture. Thus, even if the conjecture is false we would require significant improvements to state-of-the art to refute it.

Black Pebbling Reduction for XOR Labeling Rule. While Alwen and Serbinenko showed that any algorithm evaluating the graph labeling function $f_{G,H}$ in the parallel random oracle model *must* have cumulative memory cost at least $\Omega(w \times \Pi_{cc}^{\parallel}(G))$, their proof made the restrictive assumption that labels are computed using the concatenation rule $\text{lab}_{G,H,x}(v) = H(v, \text{lab}_{G,H,x}(v_1), \dots, \text{lab}_{G,H,x}(v_\delta))$. However, most *practical* iMHF implementations (e.g., Argon2i and DRSample(+BRG)) all follow the more efficient XOR labeling rule $\text{lab}_{G,H,x}(v) = H(\text{lab}_{G,H,x}(v_1) \oplus \dots \oplus \text{lab}_{G,H,x}(v_\delta))$ where $v_1, \dots, v_\delta = \text{parents}(v)$ and the DAGs have indegree $\delta = \mathcal{O}(1)$. The XOR labeling rule allows one to work with a faster round function $H: \{0,1\}^w \rightarrow \{0,1\}^w$, e.g., Argon2i builds $H: \{0,1\}^{8192} \rightarrow \{0,1\}^{8192}$, to speed up computation so that we fill more memory.

We extend the results of Alwen and Serbinenko to show that, for suitable DAGs, $f_{G,H}$ has cumulative memory cost at least $\Omega(w \times \Pi_{cc}^{\parallel}(G) / \delta)$ when using the XOR labeling rule. The loss of δ is necessary as the pebbling complexity of the complete DAG K_N is $\Pi_{cc}^{\parallel}(K_N) = \Omega(N^2)$, but $f_{K_N,H}$ has cmc/aAT cost at most $\mathcal{O}(N)$ when defined using the XOR labeling rule. In practice, all of the graphs we consider have $\delta = \mathcal{O}(1)$ so this loss is not significant.

One challenge we face in the reduction is that it is more difficult to extract labels from the random oracle query $\text{lab}_{G,H,x}(v_1) \oplus \dots \oplus \text{lab}_{G,H,x}(v_\delta)$ than from the query $\text{lab}_{G,H,x}(v_1), \dots, \text{lab}_{G,H,x}(v_\delta)$. Another challenge we face is that the labeling function $H'(x,y) = H(x \oplus y)$ is not even collision resistant e.g., $H'(y,x) = H'(x,y)$. In fact, one can exploit this property to find graphs G on N nodes where the function $f_{G,H}$ is a constant function: Suppose we start with a DAG $G' = (V' = [N-3], E')$ on $N-3$ nodes that has high pebbling cost $\Pi_{cc}^{\parallel}(G')$ and define $G = (V = [N], E = E' \cup \{(N-3, N-2), (N-3, N-1), (N-4, N-2), (N-4, N-1), (N-2, N), (N-1, N)\})$ by adding directed edges from node $N-3$ and $N-4$ to nodes $N-2, N-1$ and then adding directed edges from $N-2$ and $N-1$ to node N . Note that for any input x we have $\text{lab}_{G,H,x}(N-2) = H(\text{lab}_{G,H,x}(N-3) \oplus \text{lab}_{G,H,x}(N-3)) = \text{lab}_{G,H,x}(N-1)$. It follows that

$$f_{G,H}(x) = \text{lab}_{G,H,x}(N) = H(\text{lab}_{G,H,x}(N-2) \oplus \text{lab}_{G,H,x}(N-1)) = H(0^w)$$
is a constant function. Thus, the claim that $f_{G,H}$ has cumulative memory cost at least $\Omega(w \times \Pi_{cc}^{\parallel}(G) / \delta)$ cannot hold for arbitrary graphs.

The above example exploited the absence of the explicit term v in $\text{lab}_{G,H,x}(v)$ to produce two nodes that always have the same label. However, we can prove that if the DAG $G = (V = [N], E)$ contains all edges of the form $(i, i+1)$ for $i < N$ then *any* algorithm evaluating the function $f_{G,H}$ in the parallel random oracle model *must* have cumulative memory cost at least $\Omega\left(w \times \Pi_{cc}^{\parallel}(G)/\delta\right)$. Furthermore, the cumulative memory cost of an algorithm computing $f_{G,H}$ on m distinct inputs must be at least $\Omega\left(mw \times \Pi_{cc}^{\parallel}(G)\right)$. We stress that all of the practical iMHFs we consider, including Argon2i and DRSample(+BRG), satisfy this condition.

Sequential Round Function. We show how a parallel attacker could reduce aAT costs by nearly an order of magnitude by computation of the Argon2i round function in parallel. For example, the first step to evaluate the Argon2 round function $H(X, Y)$ is to divide the input $R = X \oplus Y \in \{0, 1\}^{8192}$ into 64 groups of 16-byte values $R_0, \dots, R_{63} \in \{0, 1\}^{128}$ and then compute $(Q_0, Q_1, \dots, Q_7) \leftarrow \mathcal{BP}(R_0, \dots, R_7), \dots, (Q_{56}, Q_{57}, \dots, Q_{63}) \leftarrow \mathcal{BP}(R_{56}, \dots, R_{63})$. Each call to the Blake2b permutation \mathcal{BP} can be trivially evaluated in parallel, which means that the attacker can easily reduce the depth of the circuit evaluating Argon2 by a factor of 8 *without* increasing the area of the circuit i.e., memory usage remains constant. The issue affects *all* Argon2 modes of operation (including data-dependent modes like Argon2d and Argon2id) and could potentially be used in combination with other pebbling attacks [AB16, AB17] for an even more dramatic decrease in aAT complexity. We also stress that this gain is independent of any other optimizations that an ASIC attacker might make to speed up computation of \mathcal{BP} e.g., if the attacker can evaluate \mathcal{BP} four-times faster than the honest party then the attacker will be able to evaluate the round function H $8 \times 4 = 32$ -times faster than the honest party. We propose a simple modification to the Argon2 round function by injecting a few additional data-dependencies to ensure that evaluation is inherently sequential. While the modification is simple we show it increases a parallel attacker's aAT costs by nearly an order of magnitude. Furthermore, empirical analysis indicates that our modifications have *negligible* impact on the running time on a CPU.

Implementation of our iMHF. We develop an implementation of our new iMHF candidate DRSample+BRG, which also uses the improved sequential Argon2 round function. The source code is available on Github at <https://github.com/antiparallel-drsbrg-argon/Antiparallel-DRS-BRG>. Empirical tests indicate that the running time of DRSample+BRG is equivalent to that of Argon2 for the honest party, while our prior analysis indicates the aAT costs, energy costs and sustained space complexity are all higher for DRSample+BRG.

2 Preliminaries

In this section we will lay out notation and important definitions required for the following sections.

2.1 Graph Notation and Definitions

We use $G = (V, E)$ to denote a directed acyclic graph and we use $N = 2^n$ to denote the number of nodes in $V = \{1, \dots, N\}$. Given a node $v \in V$, we use $\text{parents}(v) = \{u : (u, v) \in E\}$ to denote the *immediate parents* of node v in G . 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 use $\text{indeg}(G) = \max_v |\text{parents}(v)|$ to denote the maximum indegree of any node in G . All of the practical graphs we consider will contain each of the edges $(i, i+1)$ for $i < N$. Thus, there is a single source node 1 and a single sink node N . Most of the graphs we consider will have $\text{indeg}(G) = 2$ and in this case we will use $r(i) < i$ to denote the *other* parent of node i besides $i-1$. Given a subset of nodes $S \subseteq V$ we use $G-S$ to refer to the graph with all nodes in S deleted and we use $G[S] = G - (V \setminus S)$ to refer to the graph obtained by deleting all nodes except S . Finally, we use $G_{\leq k} = G[\{1, \dots, k\}]$ to refer to the graph induced by the first k nodes.

Block depth-robustness. Block depth-robustness is a stronger variant of depth-robustness. First, we define $N(v, b) = \{v-b+1, v-b+2, \dots, v\}$ to be the set of b contiguous nodes ending at node v . For a set of vertices $S \subseteq V$, we also define $N(S, b) = \bigcup_{v \in S} N(v, b)$. We say that a graph is (e, d, b) block depth-robust if, for every set $S \subseteq V$ of size $|S| \leq e$, $\text{depth}(G - N(S, b)) \geq d$. When $b=1$ we simply say that the graph is (e, d) depth-robust. It is known that highly depth-robust DAGs G have high pebbling complexity, and can be used to construct strong iMHFs with high aAT complexity in the random oracle model [ABP17]. In certain cases, block depth-robustness can be used to establish even *stronger* lower bounds on the pebbling complexity of a graph [ABH17, BZ17]. Alwen et al. gave an algorithm DRSample that (whp) outputs a DAG G that is (e, d, b) block depth-robust with $e = \Omega(N/\log N)$, $d = \Omega(N)$ and $b = \Omega(\log N)$ [ABH17].

Graph labeling functions. As mentioned in the introduction, an iMHF $f_{G,H}$ can be described as a mode of operation over a directed acyclic graph using a round function H . Intuitively, the graph represents data dependencies between the memory blocks that are generated as computation progresses and each vertex represents a value being computed based on some dependencies. The function $f_{G,H}(x)$ can typically be defined as a labeling function i.e., given a set of vertices $V = [N] = \{1, 2, 3, \dots, N\}$, a compression function $H = \{0, 1\}^* \rightarrow \{0, 1\}^m$ (often modeled as a Random Oracle in security analysis), and an input x , we “label” the nodes in V as follows. All source vertices (those with no parents) are labeled as $\ell_v(x) = H(v, x)$ and all other nodes with parents $v_1, v_2, \dots, v_\delta$ are labeled $\ell_v(x) = F_{v,H}(\ell_{v_1}(x), \ell_{v_2}(x), \dots, \ell_{v_\delta}(x))$ for a function $F_{v,H}(\cdot)$ that depends on $H(\cdot)$. The output $f_{G,H}(x)$ is then defined to be the label(s) of the sink node(s) in G .

In theoretical constructions (e.g., [AS15]) we often have $F_{v,H}(\ell_{v_1}(x), \ell_{v_2}(x), \dots, \ell_{v_\delta}(x)) = H(v, \ell_{v_1}(x), \ell_{v_2}(x), \dots, \ell_{v_\delta}(x))$ while in most real world constructions (e.g., Argon2i [BDK16]) we have $F_{v,H}(\ell_{v_1}(x), \ell_{v_2}(x), \dots, \ell_{v_\delta}(x)) = H(\ell_{v_1}(x) \oplus \ell_{v_2}(x) \dots \oplus \ell_{v_\delta}(x))$. To ensure that the function $f_{G,H}$ can be computed in $\mathcal{O}(N)$ steps, we require that G is an N -node DAG with constant indegree δ .

2.2 iMHFs and the Parallel Black Pebbling Game

Alwen and Serbinenko [AS15] and Alwen and Tackmann [AT17] provided reductions proving that in the parallel random oracle model (PROM) the amortized area time complexity of the function $f_{G,H}$ is completely captured by the (parallel) black pebbling game on the DAG G when we instantiate the round function as $F_{v,H}(\ell_{v_1}(x), \ell_{v_2}(x), \dots, \ell_{v_\delta}(x)) = H(v, \ell_{v_1}(x), \ell_{v_2}(x), \dots, \ell_{v_\delta}(x))$. However, *practical constructions* such as Argon2i use a different round function $F_{v,H}^\oplus(\ell_{v_1}(x), \ell_{v_2}(x), \dots, \ell_{v_\delta}(x)) = H\left(\bigoplus_{j=1}^\delta \ell_{v_j}(x)\right)$. In Section 6 we extend prior pebbling reductions to handle the round function $F_{v,H}^\oplus$, which justifies the use of pebbling games to analyze *practical constructions* of iMHFs such as Argon2i or DRSample.

Intuitively, placing a pebble on a node represents computing the corresponding memory block and storing it in memory. The rules of the black pebbling game state that we cannot place a pebble on a node v until we have pebbles on the parents of node v i.e., we cannot compute a new memory block until we have access to all of the memory blocks on which the computation depends. More formally, in the black pebbling game on a directed graph $G=(V,E)$, we place pebbles on certain vertices of G over a series of t rounds. A valid pebbling P is a sequence P_0, P_1, \dots, P_t of sets of vertices satisfying the following properties: (1) $P_0 = \emptyset$, (2) $\forall v \in P_i \setminus P_{i-1}$ we have $\text{parents}(v) \subseteq P_{i-1}$, and (3) $\forall v \in V, \exists i$ s.t. $v \in P_i$.

Intuitively, P_i denotes the *subset* of data-labels stored in memory at time i and $P_i \setminus P_{i-1}$ denotes the new data-labels that are computed during round i — the second constraint states that we can only compute these new data-labels if all of the necessary dependent data values were already in memory. The final constraint says that we must eventually pebble all nodes (otherwise we would never compute the output labels for $f_{G,H}$). We say that a pebbling is *sequential* if $\forall i > 0$ we have $|P_i \setminus P_{i-1}| \leq 1$ i.e., in every round at most *one* new pebble is placed on the graph. We use $\mathcal{P}^\parallel(G)$ (resp. $\mathcal{P}(G)$) to denote the set of all valid parallel (resp. sequential) black pebbblings of the DAG G . We define the space-time cost of a pebbling $P = (P_1, \dots, P_t) \in \mathcal{P}_G^\parallel$ to be $\text{st}(P) = t \times \max_{1 \leq i \leq t} |P_i|$ and the sequential space-time pebbling cost, denoted $\Pi_{st}(G) = \min_{P \in \mathcal{P}_G} \text{st}(P)$, to be the space-time cost of the best legal pebbling of G .

There are many other pebbling games one can define on a DAG including the red-blue pebbling game [JWK81] and the black-white pebbling game [Len81]. Red-blue pebbling games can be used to analyze the bandwidth-hardness of an iMHF [RD17, BRZ18]. In this work, we primarily focus on the (parallel) black pebbling game to analyze the amortized Area-Time complexity and the sustained space complexity of a memory-hard function.

Definition 1 (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 $\alpha \in \{s, t, st, cc\}$ the sequential and parallel pebbling complexities of G are defined as

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

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 $\Pi_{st}^\parallel(G) \geq \Pi_{cc}^\parallel(G)$. Thus, we have $\Pi_{st}(G) \geq \Pi_{cc}(G) \geq \Pi_{cc}^\parallel(G)$ and $\Pi_{st}(G) \geq \Pi_{st}^\parallel(G) \geq \Pi_{cc}^\parallel(G)$. However, the relationship between $\Pi_{st}^\parallel(G)$ and $\Pi_{cc}(G)$ is less clear. It is easy to provide examples of graphs for which $\Pi_{cc}(G) \ll \Pi_{st}^\parallel(G)$ ⁶. Alwen and Serbinenko showed that for the bit-reversal graph $G = \text{BRG}_n$ with $\mathcal{O}(N = 2^n)$ nodes we have

⁶ One such graph G would be to start with the pyramid graph Δ_k , which has $\mathcal{O}(k^2)$ nodes, a single sink node t and append a path W of length k^3 starting at this sink node t . The pyramid graph requires $\Pi_s^\parallel(\Delta_k) = \Theta(k)$ space to pebble and has $\Pi_{cc}(\Delta_k) \leq \Pi_{st}(\Delta_k) \leq k^3$. Similarly, the path W requires at least $\Pi_t^\parallel(W) = \Pi_t(W) = k^3$ steps to pebble the path (even in parallel). Thus, $\Pi_{st}^\parallel(G) \geq k^4$. By contrast, we have $\Pi_{cc}(G) \leq \Pi_{cc}(\Delta_k) + k^3 \leq k^3 + k^3 \ll k^4$.

$\Pi_{st}^{\parallel}(G) = \mathcal{O}(n\sqrt{n})$. In Section 4.2 we show that $\Pi_{cc}(G) = \Omega(N^2)$. Thus, for some DAGs we have $\Pi_{cc}(G) \gg \Pi_{st}^{\parallel}(G)$.

Definition 2 (Sustained Space Complexity [ABP18]). For $s \in \mathbb{N}$ the s -sustained-space (s -ss) complexity of a pebbling $P = \{P_0, \dots, P_t\} \in \mathcal{P}_G^{\parallel}$ is: $\Pi_{ss}(P, s) = |\{i \in [t] : |P_i| \geq s\}|$. More generally, the sequential and parallel s -sustained space complexities of G are defined as

$$\Pi_{ss}(G, s) = \min_{P \in \mathcal{P}_G} \Pi_{ss}(P, s) \quad \text{and} \quad \Pi_{ss}^{\parallel}(G, s) = \min_{P \in \mathcal{P}_G^{\parallel}} \Pi_{ss}(P, s) .$$

We remark that for any s we have $\Pi_{cc}(G) \geq \Pi_{ss}(G, s) \times s$ and $\Pi_{cc}^{\parallel}(G) \geq \Pi_{ss}^{\parallel}(G, s) \times s$.

2.3 Amortized Area-Time Cost (aAT)

Amortized Area-Time (aAT) cost is a way of viewing the cost to compute an iMHF, and it is closely related to the cost of pebbling a graph. Essentially, aAT cost represents the cost to keep pebbles in memory and adds in a factor representing the cost to compute the pebble. Here we require an additional factor, the core-memory ratio R , a multiplicative factor representing the ratio between computation cost vs memory cost. In this paper we are mainly focused on analysis of Argon2, which has previous calculations showing $R=3000$ [BK15]. It can be assumed that this value is being used for R unless otherwise specified. The formal definition of the aAT complexity of a pebbling $P = (P_0, \dots, P_T)$ of the graph G is as follows:

$$\text{aAT}_R(P) = \sum_{i=1}^T |P_i| + R \sum_{i=1}^T |P_i \setminus P_{i-1}|$$

The (sequential) aAT complexity of a graph G is defined to be the aAT complexity of the optimal (sequential) pebbling strategy. Formally,

$$\text{aAT}_R(G) = \min_{P \in \mathcal{P}(G)} \text{aAT}_R(P) , \text{ and } \text{aAT}_R^{\parallel}(G) = \min_{P \in \mathcal{P}^{\parallel}(G)} \text{aAT}_R(P) .$$

One of the nice properties of aAT^{\parallel} and Π_{cc}^{\parallel} complexity is that both cost metrics amortize nicely i.e., if G^m consists of m independent copies of the DAG G then $\text{aAT}_R^{\parallel}(G^m) = m \times \text{aAT}_R^{\parallel}(G)$. We remark that $\text{aAT}_R^{\parallel}(G) \geq \Pi_{cc}^{\parallel}(G)$, but that in most cases we will have $\text{aAT}_R^{\parallel}(G) \approx \Pi_{cc}^{\parallel}(G)$ since the number of queries to the random oracle is typically $o(\Pi_{cc}^{\parallel}(G))$. We will work with $\Pi_{cc}^{\parallel}(G)$ when conducting theoretical analysis and we will use $\text{aAT}_R^{\parallel}(G)$ when conducting empirical experiments, as the constant factor R is important in practice. This also makes it easier to compare our empirical results with prior work [AB17, ABH17].

2.4 Attack Quality

In many cases we will care about how efficient certain pebbling strategies are compared to others. When we work with an iMHF, we have a naïve sequential algorithm \mathcal{N} for evaluation e.g., the algorithm described in the Argon2 specifications [BDK16]. Typically, the naïve algorithm \mathcal{N} is relatively expensive e.g., $\text{aAT}_R(\mathcal{N}) = N^2/2 + RN$. We say

since we can place a pebble on node t with cost $\Pi_{cc}(\Delta_k)$, discard all other pebbles from the graph, and then walk this pebble across the path.

that an attacker \mathcal{A} is *successful* at reducing evaluation costs if $\mathbf{aAT}_R(\mathcal{A}) < \mathbf{aAT}_R(\mathcal{N})$. Following [AB16] we define the quality of the attack as

$$\text{AT-quality}(\mathcal{A}) = \frac{\mathbf{aAT}_R(\mathcal{N})}{\mathbf{aAT}_R(\mathcal{A})},$$

which describes how much more efficiently \mathcal{A} evaluates the function compared to \mathcal{N} .

3 Analysis of the Greedy Pebbling Algorithm

In this section we present a theoretical and empirical analysis of the greedy pebbling attack [BCS16] that *reverses* previous conclusions about the *practical* security of Argon2i vs DRSample [ABH17]. We prove two main results using the greedy algorithm. First, we show that for any N node DAG G with indegree 2 and a unique topological ordering, we have $\mathbf{aAT}_R(G) \leq \frac{N^2+2N}{4} + RN$ — see Theorem 1. Second, we prove that for any constant $\eta > 0$ and a random DRSample DAG G on N nodes, we have $\Pi_{st}(G) \leq (1+\eta)2N^2/\log N$ with high probability — see Theorem 2. We stress that in both cases the bounds are *explicit* not *asymptotic*, and that the pebbling attacks are simple and sequential.

Alwen and Blocki [AB16] previously had shown that any DAG G with constant indegree has $\mathbf{aAT}_R(G) \in \mathcal{O}(N^2 \log \log N / \log N)$, but the constants from this bound were not well understood and did not rule out the existence of an N node DAG G with $\mathbf{aAT}_R(G) \geq N^2/2 + RN$ for *practical* values of N e.g., unless we use more than 16GB of RAM we have $N \leq 2^{24}$ for Argon2i or DRSample⁷. By contrast, Theorem 1 immediately implies that $\mathbf{aAT}_R(G) \leq \frac{N^2+2N}{4} + RN$. Similarly, Alwen et al. [ABH17] previously showed that with high probability a DRSample DAG G has $\mathbf{aAT}_R(G) \in \Omega(N^2/\log N)$, but the constants in this lower bound were not well understood. On a theoretical side, our analysis shows that this bound is tight i.e., $\mathbf{aAT}_R(G) \in \Theta(N^2/\log N)$. It also proves that DRSample does not quite match the generic upper bound of Alwen and Blocki [AB16].

Extension of the Greedy Pebbling Attack. Our analysis leaves us in an uncomfortable position where *every practical* iMHF candidate has high-quality pebbling attacks i.e., greedy pebble for DRSample and depth-reducing attacks for Argon2i. We would like to develop a practical iMHF candidate that provides strong resistance against all known pebbling attacks for all practical values of $N \leq 2^{24}$. We first consider a defense proposed by Biryukov et al. [BDK16] against the greedy pebbling attack. While this defense provides optimal protection against the greedy pebbling attack, we introduce an extension of the greedy pebbling attack that we call the *staggered* greedy pebbling attack and show that the trick of Biryukov et al. [BDK16] fails to protect against the extended attack.

3.1 The Greedy Pebbling Algorithm

We first review the greedy pebbling algorithm. We first introduce some notation. **gc(v):** For each node $v < N$ we let $\mathbf{gc}(v) = \max\{w \mid (v, w) \in E\}$ denote the maximum child of node v — if $v < N$ then the set $\{w \mid (v, w) \in E\}$ is non-empty as it contains the node $v+1$. If node v has no children then set $\mathbf{gc}(v) := v$.

⁷ In Argon2, the block-size is 1KB so when we use $N = 2^{24}$ nodes the honest party would require 16GB ($= N \times \text{KB}$) of RAM to evaluate the MHF. Thus, we view 2^{24} as a reasonable upper bound on the number of blocks that would be used in practical applications.

$\chi(i)$: This represents what we call the crossing set of the i th node. It is defined as $\chi(i) = \{v \mid v \leq i \wedge \text{gc}(v) > i\}$. Intuitively this represents the set of nodes $v \leq i$ incident to a directed edge (v, u) that “crosses over” node i i.e. $u > i$.

Greedy Pebbling Strategy: Set $\text{GP}(G) = P = (P_1, \dots, P_N)$ where $P_i = \chi(i)$ for each $i \leq N$. Intuitively, the pebbling strategy can be described follows: In round i we place a pebble on node i and we then discard *any* pebbles on nodes v that are no longer needed in any future round i.e., for all future nodes $w > i$ we have $v \notin \text{parents}(w)$ (equivalently, the greatest-child of node v is $\text{gc}(v) \leq i$). We refer the reader to the full version [BHK⁺18] for a formal algorithmic description.

We first prove the following *general* lower bound for *any* N node DAG with $\text{indeg}(G) \leq 2$ that has a unique topological ordering i.e., G contains each of the edges $(i, i+1)$. In particular, Theorem 1 shows that for any such DAG G we have $\Pi_{st}(G) \lesssim \frac{N^2}{2}$ and $\Pi_{cc}(G) \lesssim N^2/4$. We stress that this is *twice* as efficient as the naive pebbling algorithm \mathcal{N} , which set $P_i = \{1, \dots, i\}$ for each $i \leq N$ and has cumulative cost $\Pi_{cc}^\parallel(\mathcal{N}) = \frac{N^2}{2}$. Previously, the gold standard was to find constructions of DAGs G with N nodes such that $\Pi_{cc}^\parallel(G) \gtrsim \frac{N^2}{2}$ for *practical* values of N — asymptotic results did not rule out this possibility even for $N \leq 2^{40}$. Theorem 1 demonstrates that the best we could hope for is to ensure $\Pi_{cc}^\parallel(G) \gtrsim \frac{N^2}{4}$ for *practical* values of N .

Theorem 1. *Let $r : \mathbb{N}_{>0} \rightarrow \mathbb{N}$ be any function with the property that $r(i) < i - 1$ for all $i \in \mathbb{N}_{>0}$. Then the DAG $G = (V, E)$ with N nodes $V = \{1, \dots, N\}$ and edges $E = \{(i-1, i) : 1 < i \leq N\} \cup \{(r(i), i) : 2 < i \leq N\}$ has $\Pi_{st}(G) \leq \frac{N^2+2N}{2}$ and $\Pi_{cc}(G) \leq \frac{N^2+2N}{4}$ and $\text{aAT}_R(G) \leq \frac{N^2+2N}{4} + RN$.*

The full proof of Theorem 1 is in the full version [BHK⁺18] of this paper. Intuitively, Theorem 1 follows from the observation that in any pebbling we have $|P_i| \leq i$, and in the greedy pebbling we also have $|P_i| \leq N - i$ since there can be at most $N - i$ nodes w such that $w = r(v)$ for some $v > i$ and other pebbles on any other node would have been discarded by the greedy pebbling algorithm.

3.2 Analysis of the Greedy Pebble Attack on DRSample

We now turn our attention to the specific case of the iMHF DRSample. A DAG G sampled from this distribution has edges of the form $(i, i+1)$ and $(r(i), i)$ where each $r(i) < i$ is independently selected from some distribution. It is not necessary to understand all of the details of this distribution to follow our analysis in this section as the crucial property that we require is given in Claim 1. The proof of Claim 1 (along with a description of DR-Sample [ABH17]) is found in the full version [BHK⁺18] of this paper. Intuitively, Claim 1 follows because we have $\Pr[r(j) = i] \sim \frac{1}{\log j} \times \frac{1}{|j-i|}$ for each node $i < j$ in DRSample.

Claim 1 *Let G be a randomly sampled DRSample DAG with N nodes and let $Y_{i,j}$ be an indicator random variable for the event that $r(j) < i$ for nodes $i < j \leq N$. Then we have $\mathbf{E}[Y_{i,j}] = \Pr[r(j) < i] \leq 1 - \frac{\log(j-i-1)}{\log j}$.*

If $P = (P_1, \dots, P_N) = \text{GP}(G)$, then we remark that $\chi(i)$ can be viewed as an alternate characterization of the set $P_i = \chi(i)$ of pebbles on the graph at time i . Lemma 1 now implies that with high probability, we will have $|P_i| \leq (1+\delta)N/n$ during *all* pebbling rounds.

Lemma 1. *Given a DAG G on $N=2^n$ nodes sampled using the randomized DRSample algorithm for any $\eta>0$, we have*

$$\Pr \left[\max_i |\chi(i)| > (1+\eta) \left(\frac{2N}{n} \right) \right] \leq \exp \left(\frac{-2\eta^2 N}{3n} + n \ln 2 \right).$$

Lemma 1, which bounds the size of $\max_i |\chi(i)|$, is proved in the full version [BHK⁺18]. Intuitively, the proof uses the observation that $|\chi(i)| \leq \sum_{j=i+1}^N Y_{i,j}$ where $Y_{i,j}$ is an indicator random variable for the event that $r(j) \leq i$. This is because $|\chi(i)|$ is upper bounded by the number of edges that “cross” over the node i . We can then use Claim 1 and standard concentration bounds to obtain Lemma 1.

Theorem 2, our main result in this section, now follows immediately from Lemma 1. Theorem 2 states that, except with negligibly small probability, the sequential pebbling cost of a DRSample DAG is at most $(1+\eta) \left(\frac{2N^2}{n} \right) + RN$.

Theorem 2. *Let G be a randomly sampled DRSample DAG with $N=2^n$ nodes. Then for all $\eta>0$ we have*

$$\Pr \left[\Pi_{st}(\text{GP}(G)) > (1+\eta) \left(\frac{2N^2}{n} \right) \right] \leq \exp \left(\frac{-2\eta^2 N}{3n} + n \ln 2 \right).$$

Proof. Fix $\eta>0$ and consider a randomly sampled N -node DRSample DAG G . Recall that $|P_i| = \chi(i)$ where $P = \text{GP}(G)$. It follows that $\Pi_{st}(\text{GP}(G)) \leq N \max_{i \in [N]} |\chi(i)|$. By Lemma 1, except with probability $\exp \left(\frac{-\eta^2 N/n}{3} + n \ln 2 \right)$, we have

$$\Pi_{st}(\text{GP}(G)) \leq N \times \max_{i \in [N]} |\chi(i)| \leq (1+\eta) \left(\frac{2N^2}{n} \right).$$

□

Discussion. Theorem 2 implies that the (sequential) aAT complexity of DRSample is $\text{aAT}_R(G) \lesssim 2N^2/\log N \in \mathcal{O}(N^2/\log N)$, which asymptotically matches the lower bound of $\Omega(N^2/\log N)$ [ABH17]. More significant from a practical standpoint is that the constant factors in the upper bound are given explicitly. Theorem 2 implies attack quality at least $\gtrsim \frac{\log N}{4}$ since the cost of the naïve pebbling algorithm is $N^2/2$. Thus, for *practical* values of $N \leq 2^{24}$ we will get high-quality attacks and our empirical analysis suggests that attack quality actually scales with $\log N$. On a positive note, the pebbling attack is *sequential*, which means that we could adjust the naïve (honest) evaluation algorithm to simply use \mathcal{N} to use $\text{GP}(G)$ instead because the greedy pebbling strategy is *sequential*. While this would lead to an *egalitarian* function, the outcome is still undesirable from the standpoint of password hashing where we want to ensure that the attacker’s absolute aAT costs are as high as possible given a fixed running time N .

3.3 Empirical Analysis of the GP Attack

We ran the greedy pebbling attack against several iMHF DAGs including Argon2i, DRSample and our new construction DRSample+BRG (see Section 4) and compare the attack quality of the greedy pebbling attack with prior depth-reducing attacks. The results, seen in Figure 2 (left), show that the GP attack was especially effective against the DRSample DAG, improving attack quality by a factor of up to 7 (at $n=24$) when

compared to previous state-of-the-art depth-reducing attacks (Valiant, Layered, and various hybrid approaches) [Val77, AB16, ABH17].

The most important observation about Figure 2 (left) is simply how effective the greedy pebbling attack is against DRSample. We remark that attack quality for DRSample with $N = 2^n$ nodes seems to be approximately n — slightly better than the theoretical guarantees from Theorem 2. While DRSample may have the strongest asymptotic guarantees (i.e. $\text{aAT}^\parallel(G) = \Omega(N^2/\log N)$ for DRSample vs. $\text{aAT}^\parallel(G) = \mathcal{O}(N^{1.767})$ for Argon2i) Argon2i seems to provide better resistance to known pebbling attacks for *practical* parameter ranges.

Our tests found that while the Greedy Pebbling attack does sometimes outperform depth-reducing attacks at smaller values of n , the depth-reducing attacks appear to be superior once we reach graph sizes that would likely be used in practice. As an example, when $n = 20$ we find that the attack quality of the greedy pebbling attack is just 2.99, while the best depth-reducing attack achieved attack quality 6.25 [ABH17].

3.4 Defense Against Greedy Pebbling Attack: Attempt 1 XOR extension

Biryukov et al. [BDK16] introduced a simple defense against the greedy pebbling attack of Boneh et al. [BCS16] for iMHFs that make two passes over memory. Normally during computation the block $B_{i+N/2}$ would be stored at memory location i overwriting block B_i . The idea of the defense is to XOR the two blocks $B_{i+N/2}$ and B_i before overwriting block B_i in memory. Biryukov et al. [BDK16] observed that this defense does not *significantly* slow down computation because block B_i would have been loaded into cache before it is overwritten in either case. The effect of performing this extra computation is effectively to add each edge of the form $(i - \frac{N}{2}, i)$ to the DAG G . In particular, this means that the greedy pebbling algorithm will not discard the pebble on node $i - \frac{N}{2}$ until round i , which is when the honest pebbling algorithm would have discarded the pebble anyway. Given a graph $G = (V, E)$ we use $G^\oplus = (V, E^\oplus)$ to denote the XOR-extension graph of G where $E^\oplus = E \cup \{(i - \frac{N}{2}, i) \mid i > \frac{N}{2}\}$. It is easy to see that $\Pi_{cc}^\parallel(\text{GP}(G^\oplus)) \geq \frac{N^2 + 2N}{4}$, which would make it tempting to conclude that the XOR-extension defeats the greedy pebbling attack. **Greedy Pebble Extension:** Given a graph G on N nodes, let $P = (P_1, \dots, P_N) = \text{GP}(G)$ and let $Q = (Q_1, \dots, Q_{N/2}) = \text{GP}(G_{\leq N/2})$. Define $\text{GPE}(G^\oplus) = (P_1^\oplus, \dots, P_N^\oplus)$ where $P_{i+N/2-1}^\oplus = Q_i \cup P_{i+N/2-1}$ and $P_i^\oplus = P_i$ for $i < N/2$. Intuitively, the attack exploits the fact that we *always* ensure that we have a pebble on the extra node $v \in \text{parents}(N/2 + v)$ at time $N/2 + v - 1$ by using the greedy pebble algorithm to synchronously re-pebble the nodes $1, \dots, N/2$ a second time.

Theorem 3 demonstrates that the new generalized greedy pebble algorithm is effective against the XOR-extension gadget. In particular, Corollary 2 states that we still obtain high-quality attacks against DRSample^\oplus so the XOR-gadget does not significantly improve the aAT cost of DRSample.

Theorem 3. *Let $r : \mathbb{N}_{>0} \rightarrow \mathbb{N}$ be any function with the property that $r(i) < i$ for all $i \in \mathbb{N}_{>0}$ and let $G = (V, E)$ be a graph with N nodes $V = \{1, \dots, N\}$ and directed edges $E = \{(i, i+1) \mid i < N\} \cup \{r(i), i \mid 1 < i \leq N\}$. If $P = \text{GP}(G) \in \mathcal{P}(G)$ and $Q \in \mathcal{P}(G_{\leq N/2})$ then the XOR-extension graph G^\oplus of G has amortized Area-Time complexity at most*

$$\text{aAT}_R^\parallel(G^\oplus) \leq \sum_{i=1}^{N/2} |P_i| + \sum_{i=1}^N |Q_i| + \frac{3RN}{2}.$$

Corollary 1. Let $r: \mathbb{N}_{>0} \rightarrow \mathbb{N}$ be any function with the property that $r(i) < i$ for all $i \in \mathbb{N}_{>0}$ and let $G = (V, E)$ be a graph with N nodes $V = \{1, \dots, N\}$ and directed edges $E = \{(i, i+1) \mid i < N\} \cup \{r(i), i \mid 1 < i \leq N\}$. Then for the XOR-extension graph G^\oplus we have $\mathbf{aAT}^\parallel_R(G^\oplus) \leq \frac{5N^2 + 12N}{16} + \frac{3RN}{2}$.

The proof of Theorem 3 can be found in the full version [BHK⁺18]. One consequence of Theorem 3 is that the XOR-extension gadget does not rescue DRSample from the greedy pebble attack — see Corollary 2.

Corollary 2. Fix $\eta > 0$ be a fixed constant and let $G = (V, E)$ be a randomly sampled DRSample DAG with $N = 2^n$ nodes $V = \{1, \dots, N\}$ and directed edges $E = \{(i, i+1) \mid i < N\} \cup \{r(i), i \mid 1 < i \leq N\}$. Then

$$\Pr \left[\mathbf{aAT}^\parallel_R(G^\oplus) > (1+\eta) \left(\frac{3N^2}{n} - \frac{N^2}{n(n-1)} \right) + \frac{3RN}{2} \right] \leq \exp \left(\frac{-\eta^2 N}{3(n-1)} + 1 + n \ln 2 \right).$$

Proof. Fix $\eta > 0$ and let $P = \mathbf{GP}(G)$ where G is a randomly sampled DRSample DAG. By Lemma 1, except with probability $\exp \left(\frac{-2\eta^2 N}{3n} + n \ln 2 \right)$, we have $\max_i |P_i| = \max_i |\chi(i)| \leq (1+\eta) \frac{2N}{n}$, which means that $\sum_{i=1}^N |P_i| \leq (1+\eta) \frac{2N^2}{n}$. Similarly, let $Q = \mathbf{GP}(G_{\leq N/2})$ be a greedy pebbling of the subgraph formed by the first $N/2$ nodes in G . We remark that $G_{\leq N/2}$ can be viewed as a randomly DRSample DAG with $N/2 = 2^{n-1}$ nodes. Thus except with probability $\exp \left(\frac{-\eta^2 N}{3(n-1)} + (n-1) \ln 2 \right)$, we have $\max_{i \leq N/2} |Q_i| = \max_i |\chi(i)| \leq (1+\eta) \frac{N}{n-1}$ since the first $N/2$ nodes of G form a random DRSample DAG with $N/2 = 2^{n-1}$ nodes. This would imply that $\sum_{i=1}^{N/2} |Q_i| \leq (1+\eta) \frac{N}{n-1}$. Putting both bounds together Theorem 3 implies that $\mathbf{aAT}^\parallel(G^\oplus) \leq (1+\eta) \left(\frac{3N^2}{n} - \frac{N^2}{n(n-1)} \right) + \frac{3RN}{2}$. \square

4 New iMHF Construction with Optimal Security

In this section, we introduce a new iMHF construction called DRSample+BRG. The new construction is obtained by overlaying a bit-reversal graph \mathbf{BRG}_n [LT82] on top of a random DRSample DAG. If G denotes a random DRSample DAG with $N/2$ nodes then we will use $\mathbf{BRG}(G)$ to denote the bit-reversal overlay with N nodes. Intuitively, the result is a graph that resists both the greedy pebble attack (which is effective against DRSample alone) and depth-reducing attacks (which DRSample was designed to resist). An even more exciting result is that we can show that DRSample+BRG is the first practical construction to provide strong sustained space complexity guarantees. Interestingly, neither graph (DRSample or BRG) is individually known to provide strong sustained space guarantees. Instead, several of our proofs exploit the synergistic properties of both graphs. We elaborate on the desirable properties of DRSample+BRG below.

First, our new construction inherits desirable properties from *both* the bit-reversal graph and DRSample. For example, $\Pi_{cc}^\parallel(\mathbf{BRG}(G)) \geq \Pi_{cc}^\parallel(G) = \Omega(N^2/\log N)$. Similarly, it immediately follows that $\mathbf{BRG}(G)$ is maximally bandwidth hard. In particular, Ren and Devadas [RD17] showed that \mathbf{BRG}_n is maximally bandwidth hard, and Blocki et al. [BRZ18] showed that DRSample is maximally bandwidth hard.

Second, $\mathbf{BRG}(G)$ provides optimal resistance to the greedy pebbling attack — $\Pi_{cc}^\parallel(\mathbf{GP}(\mathbf{BRG}(G))) \approx N^2/4$. Furthermore, we can show that *any* c -parallel pebbling

attack $P = (P_1, \dots, P_t)$ in which $|P_{t+1} \setminus P_i| \leq c$ has cost $\Pi_{cc}(P) = \Omega(N^2)$. This rules out any *extension* of the greedy pebble attack e.g., GPE is 2-parallel. In fact, we prove that this property already holds for any c -parallel pebbling of the bit reversal graph BRG_n . Our proof that $\Pi_{cc}(\text{BRG}_n) = \Omega(N^2)$ generalizes the well-known result that $\Pi_{st}(\text{BRG}_n) = \Omega(N^2)$ and may be of independent interest.

Third, we can show that *any* parallel pebbling P of $\text{BRG}(G)$ either has $\Pi_{cc}(P) = \Omega(N^2)$ or has maximal sustained space complexity $\Pi_{ss}(P, s) = \Omega(N)$ for space $s = \Omega(N/\log N)$ i.e., there are at least $\Omega(N)$ steps with at least $\Omega(N/\log N)$ pebbles on the graph. To prove this last property we must rely on properties of *both* graphs G and BRG_n i.e., the fact that DRSample is highly block depth-robust and the fact that edges BRG_n are evenly distributed over every interval. This makes $\text{BRG}(G)$ the first *practical* construction of a DAG with provably strong sustained space complexity guarantees.

Finally, we can show that $\Pi_{cc}^{\parallel}(G) = \Omega(N^2 \log \log N / \log N)$, matching the general upper bound of Alwen and Blocki [AB16], under a plausible conjecture about the block-depth-robustness of G . In particular, we conjecture that G is (e, d, b) -block depth-robust for $e = \Omega\left(\frac{N \log \log N}{\log N}\right)$, $d = \Omega\left(\frac{N \log \log N}{\log N}\right)$ and $b = \Omega\left(\frac{\log N}{\log \log N}\right)$. In the full version [BHK⁺18], we also show how to construct a constant indegree DAG G' with $\Pi_{cc}^{\parallel}(G') = \Omega(N^2 \log \log N / \log N)$ from *any* (e, d) -depth robust graph by overlaying a superconcentrator on top of G [Pip77]. However, the resulting construction is not *practically efficient*. Thus we show the bit reversal overlay $G' = \text{BRG}(G)$ satisfies the same complexity bounds under the slightly stronger assumption that G is *block*-depth-robust. As evidence for the conjecture we show that *known* attacks require the removal of a set S of $e = \Omega\left(\frac{N \log \log N}{\log N}\right)$ to achieve $\text{depth}(G - S) \leq \frac{N}{\sqrt{\log N}}$. Thus, we would need to find *substantially* improved depth-reducing attacks to refute the conjectures.

Bit-Reversal Graph Background. The bit reversal graph was originally proposed by Lenguer and Tarjan [LT82] who showed that any sequential pebbling has maximal space-time complexity. Forler et al. [FLW14] previously incorporated this graph into the design of their iMHF candidate Catena, which received special recognition at the password hashing competition [PHC16]. While we are not focused on sequential space-time complexity, the bit reversal graph has several other useful properties that we exploit in our analysis (see Lemma 2).

Local Samplable. We note that one benefit of $\text{DRS} + \text{BRG}$ is that it is locally samplable, a notion mentioned as desirable in [ABH17]. Specifically, we want to be able to compute the parent blocks with time and space $\mathcal{O}(\log |V|)$ with small constants. $\text{DRS} + \text{BRG}$ meets this requirement. Edges sampled from DRSample were shown to be locally navigable in [ABH17], and each bit-reversal edge a simple operation called requires one bit reversal operation, which can easily be computed in time $\mathcal{O}(\log |V|)$. The formal description of the bit-reversal overlay graph $\text{BRG}(G)$ is presented in Definition 4.

The Bit-Reversal DAG. Given a sequence of bits $X = x_1 \circ x_2 \circ \dots \circ x_n$, let $\text{ReverseBits}(X) = x_n \circ x_{n-1} \circ \dots \circ x_1$. Let $\text{integer}(X)$ be the integer representation of bit-string X starting at 1 so that $\text{integer}(\{0, 1\}^n) = [2^n]$ i.e., $\text{integer}(0^n) = 1$ and $\text{integer}(1^n) = 2^n$. Similarly, let $\text{bits}(v, n)$ be the length n binary encoding of $(v-1) \bmod 2^n$ e.g., $\text{bits}(1, n) = 0^n$ and $\text{bits}(2^n, n) = 1^n$ so that for all $v \in [2^n]$ we have $\text{integer}(\text{bits}(v, n)) = v$.

Definition 3. We use the notation BRG_n to denote the bit reversal graph with 2^{n+1} nodes. In particular, $\text{BRG}_n = (V = [2^{n+1}], E = E_1 \cup E_2)$ where $E_1 := \{(i, i+1) : 1 \leq$

$i < 2^{n+1}$ and $E_2 := \{(x, 2^n + y) : x = \text{integer}(\text{ReverseBits}(\text{bits}(y, n)))\}$. That is, E_2 contains an edge from node $x \leq 2^n$ to node $2^n + y$ in BRG_n if and only if $x = \text{integer}(\text{ReverseBits}(\text{bits}(y, n)))$.

Claim 2 states that the cumulative memory cost of the greedy pebbling strategy $\text{GP}(\text{BRG}_n)$ is at least $N^2 + N$.

Claim 2 $\Pi_{cc}(\text{GP}(\text{BRG}_n)) \geq N^2 + N$

Proof. Let $P = (P_1, \dots, P_{2N}) = \text{GP}(\text{BRG}_n)$. We first note that for all $i \leq N$ we have $P_i = \{1, \dots, i\}$ since $\text{gc}(i) > N$ — every node on the bottom layer $[N]$ has an edge to some node on the top layer $[N+1, 2N]$. Second, observe that for any round $i > N$ we have $|P_i \setminus P_{i+1} \cap [N]| \leq 1$ since the only pebble in $[N]$ that might be discarded is the (unique) parent of node i . Thus,

$$\sum_{i=1}^{2N} |P_i| \geq \sum_{i=1}^N i + \sum_{i=1}^N (N - i + 1) = N(N+1). \quad \square$$

Thus, we now define the bit-reversal overlay of the bit reversal graph on a graph G_1 . If the graph G_1 has N nodes then $\text{BRG}(G_1)$ has $2N$ nodes, and the subgraph induced by the first N nodes of $\text{BRG}(G_1)$ is simply G_1 .

Definition 4. Let $G_1 = (V_1 = [N], E_1)$ be a fixed DAG with $N = 2^n$ nodes and $\text{BRG}_n = (V = [2N], E)$ denote the bit-reversal graph. Then we use $\text{BRG}(G_1) = (V, E \cup E_1)$ to denote the bit-reversal overlay of G_1 .

In our analysis, we will rely heavily on the following key-property of the bit-reversal graph from Lemma 2.

Lemma 2. Let $G = \text{BRG}_n$ and $N = 2^n$ so that G has $2N$ nodes. For a given b , partition $[N]$ into $\frac{N}{2^{n-b}} = 2^b$ intervals $I_k = [(k-1)2^{n-b}, k2^{n-b} - 1]$, each having length 2^{n-b} , for $1 \leq k \leq 2^b$. Then for any interval I of length 2^{b+1} , with $I \subseteq [N+1, 2N]$, there exists an edge from each I_k to I , for $1 \leq k \leq 2^b$.

Proof of Lemma 2. Let I be any interval of length 2^b , with $I \subseteq [N+1, 2N]$. Note that every 2^b length bitstring appears as a suffix in I . Thus, there exists an edge from each interval containing a unique 2^b length bitstring as a prefix. It follows that there exists an edge from each I_k to I , for $1 \leq k \leq 2^b$. \square

As we will see, the consequences of Lemma 2 will have powerful implications for the pebbling complexity of $G = \text{BRG}(G_1)$ whenever the underlying DAG G_1 is (e, d, b) -block-depth-robust. In particular, Lemma 3 states that if we start with pebbles on a set $|P_i| < e/2$ then for any initially empty interval I of $\mathcal{O}(N/b)$ consecutive nodes in the top-half of G we have the property that $H := G - \bigcup_{x \in P_i} [x - b + 1, x]$ is an $(e/2, d, b)$ -block-depth-robust graph that will need to be *completely re-pebbled* (at cost at least $\Pi_{cc}^{\parallel}(H) \geq ed/2$) just to advance a pebble across the interval I . See the full version [BHK⁺18] for the proof of Lemma 3.

Lemma 3. Let $G_1 = (V_1 = [N], E)$ be a (e, d, b) -block depth-robust graph with $N = 2^n$ nodes and let $G = \text{BRG}(G_1)$ denote the bit-reversal extension of G_1 with $2N$ nodes $V(G) = [2N]$. For any interval $I = [N + i + 1, N + i + 1 + \frac{4N}{b}] \subseteq [2N]$ and any $S \subseteq [1, N + i]$ with $|S| < \frac{e}{2}$, $\text{ancestors}_{G-S}(I)$ is $(\frac{e}{2}, d, b)$ -block depth-robust.

Lemma 4. Let G be a (e, d, b) -block depth-robust DAG with $N = 2^n$ and let $G' = \text{BRG}(G)$ be the bit reversal overlay of G . Let $P \in \mathcal{P}^\parallel(G')$ be a legal pebbling of G' and let t_v be the first time where $v \in P_{t_v}$. Then for all $v \geq 1$ such that $e' := |P_{t_v+N}| \leq \frac{e}{4}$ and $v \leq N - \frac{32Ne'}{be}$, we have

$$\sum_{j=t_v+N}^{t_{v+N} + \frac{32Ne'}{be} - 1} |P_j| \geq \frac{ed}{2}.$$

Proof of Lemma 4. Let $v \leq N - \frac{32Ne'}{be}$ be given such that the set $S = P_{t_v+N}$ has size at most $e' = |S| \leq e/4$ and set $b' = \frac{eb}{4e'}$. Consider the ancestors of the interval $I = [N+v+1, N+v+\frac{8N}{b'}]$ in the graph $G' - S$. Note that $I \cap S = \emptyset$ since v is the maximum node that has been pebbled at time t_{N+v} . We have

$$H := G - \bigcup_{x \in S} [x - b' + 1, x] \subseteq \text{ancestor}_{G'-S}(I)$$

because for any node $u \in V(G)$ if $u \notin \bigcup_{x \in S} [x - b' + 1, x]$ then $[u, u + b' - 1] \cap S = \emptyset$ which implies that there exists an “ S -free path” from u to I by Lemma 2. Thus, H will have to be repebbled completely at some point during the time interval $[t_{v+N}, t_{v+N+\frac{32Ne'}{be}-1}]$ since $\frac{32Ne'}{be} \geq \frac{8N}{b}$.

Since $b' = \frac{eb}{4e'} \geq b$ we note that the e' intervals of length b' we are removing can be covered by at most $\lceil b'/b \rceil e' = \lceil e/(4e') \rceil e' \leq (e/4) + e' \leq e/2$ intervals of length e . Hence, Lemma 3 implies that H is still $(e/2, d, b)$ -block depth-robust and, consequently, we have that $\Pi_{cc}^\parallel(H) \geq ed/2$ by [ABP17]. We can conclude that

$$\sum_{j=t_v+N}^{t_{v+N} + \frac{32Ne'}{be} - 1} |P_j| \geq \Pi_{cc}^\parallel(H) \geq ed/2. \quad \square$$

4.1 Sustained Space Complexity (Tradeoff Theorem)

We prove that for any parameter $e = \mathcal{O}\left(\frac{N}{\log N}\right)$, either the cumulative pebbling cost of any parallel (legal) pebbling P is at least $\Pi(P) = \Omega(N^3/(e \log N))$, or there are at least $\Omega(N)$ steps with at least e pebbles on the graph i.e., $\Pi_{ss,e}(P) = \Omega(N)$. Note that the cumulative pebbling cost rapidly increases as e decreases e.g., if $e = \sqrt{N}/\log N$ then any pebbling P for which $\Pi_{ss}(P, e) = o(N)$ must have $\Pi(P) = \Omega(N^{2.5})$.

To begin we start with the known result that (with high probability) a randomly sampled DRSample DAG G is (e, d, b) -block depth-robust with $e = \Omega(N/\log N)$, $b = \Omega(\log N)$, and $d = \Omega(N)$ [ABH17]. Lemma 5 now implies that the DAG is also (e', d, b') -block depth-robust for any suitable parameters e' and b' . Intuitively, if we delete e' intervals of length $b' > b$ then we can cover these deleted intervals with *at most* $e' \left(\frac{b'}{b} + 1\right)$ intervals of length b , as illustrated in Figure 1. The formal proof of Lemma 5 is in the full version [BHK⁺18].

Lemma 5. Suppose that a DAG G is (e, d, b) -block depth-robust and that parameters e' and b' satisfy the condition that $e' \left(\frac{b'}{b}\right) + e' \leq \frac{e}{2}$. Then G is (e', d, b') -block depth-robust, and for all S with size $|S| \leq e'$ the graph $H = G - \bigcup_{x \in S} [x - b' + 1, x]$ is $(\frac{e}{2}, d, b)$ -block depth-robust.




Fig. 1: Intervals $\bigcup_{x \in S} [x - b' + 1, x]$ and $\bigcup_{x \in S'} [x - b + 1, x]$ when $b' = 10$ and $b = 3$. Observe that $\bigcup_{x \in S'} [x - b + 1, x] \supset \bigcup_{x \in S} [x - b' + 1, x]$ over the integers.

Together Lemma 4 and Lemma 5 imply that we must incur pebbling cost $\Omega(ed)$ to pebble *any* interval of $\Omega\left(\frac{Ne'}{be}\right)$ consecutive nodes in the top half of $\text{BRG}(G)$, starting from *any* configuration with at most $e' \leq e/4$ pebbles on the graph.

Theorem 4, our main result in this subsection, now follows because for any pebbling $P \in \Pi^\parallel(\text{BRG}(G))$ and any interval I of $\Omega\left(\frac{Ne'}{be}\right)$ nodes in the top-half of G we must either (1) keep at least e' pebbles on the graph while we walk a pebble across the first half of the interval I , or (2) pay cost $\Omega(ed)$ to re-pebble a depth-robust graph. Since there are $\Omega\left(\frac{eb}{e'}\right)$ such disjoint intervals we must either keep $|P_i| \geq e'$ pebbles on the graph for $\Omega(N)$ rounds, or pay cost $\Pi_{cc}^\parallel(P) \geq \frac{e^2 db}{64e'}$.

Theorem 4. *Let G be any (e, d, b) -block depth-robust DAG on $N = 2^n$ nodes, and $G' = \text{BRG}(G)$ be the bit reversal overlay of G . Then for any pebbling $P \in \Pi^\parallel(G)$ and all $e' \leq \frac{e}{4}$, we have either $\Pi_{cc}^\parallel(P) \geq \frac{e^2 db}{64e'}$, or $\Pi_{ss}(P, e') \geq \frac{N}{4} - o(N)$ i.e., at least $\frac{N}{4} - o(N)$ rounds i in which $|P_i| \geq e'$.*

Corollary 3 follows immediately from Theorem 4.

Corollary 3. *Let G be any $\left(\frac{c_1 N}{\log N}, c_2 N, c_3 \log N\right)$ -block depth-robust DAG on $N = 2^n$ nodes for some constants $c_1, c_2, c_3 > 0$ and let $G' = \text{BRG}(G)$ be the bit reversal overlay of G . Then for any $e' < \frac{c_1 N}{4 \log N}$ and any pebbling $P \in \mathcal{P}^\parallel(G')$ we have either $\Pi_{cc}^\parallel(P) \geq \frac{c_1^2 c_2 c_3 N^3}{64e' \log N}$, or $\Pi_{ss}(P, e') \geq \frac{N}{4} - o(N)$ i.e., there are at least $\frac{N}{4} - o(N)$ rounds j in which $|P_j| \geq e'$.*

Remark 1. Alwen et al. previously proved that for constants $c_1 = 2.4 \times 10^{-4}$, $c_2 = 0.03$ and $c_3 = 160$, a randomly sampled DAG G from DRSample will be $\left(\frac{c_1 N}{\log N}, c_2 N, c_3 \log N\right)$ -block depth-robust except with negligible probability [ABH17]. Thus, with high probability Corollary 3 can be applied to the bit reversal overlay $\text{BRG}(G)$. Notice also that as e' decreases, the lower bound on $\Pi_{cc}^\parallel(P)$ increases rapidly e.g., if a pebbling does not have at least $\Omega(N)$ steps with at least $e' = \Omega(\sqrt{N})$ pebbles on the graph, then $\Pi_{cc}^\parallel(P) = \tilde{\Omega}(N^{2.5})$.

A Conjectured (Tight) Lower Bound on $\Pi_{cc}^\parallel(\text{BRG}(G))$. The idea behind the proof of Theorem 5 in the full version [BHK⁺18] is very similar to the proof of Theorem 4 — an attacker must either keep $e/2$ pebbles on the graph most of the time or the attacker must pay $\Omega(edb)$ to re-pebble an (e, d) -depth $\Omega(b)$ times. In fact, a slightly weaker version (worse constants) of Theorem 5 follows as a corollary of

Theorem 4 since $\Pi_{cc}^{\parallel}(P) \geq \epsilon' \times \Pi_{ss}(P, \epsilon')$. Under our conjecture that DRSample DAGs are $(c_1 N \log \log N / \log N, c_2 N \log \log N / \log N, c_3 \log N / \log \log N)$ -block depth-robust graph, Theorem 5 implies that $\Pi_{cc}^{\parallel}(\text{BRG}(G)) = \Omega(N^2 \log \log N / \log N)$. In fact, any pebbling must either keep $\Omega(N \log \log N / \log N)$ pebbles on the graph for $\approx N/4$ steps or the pebbling has cost $\Omega(N^2 \log \log N)$.

Theorem 5. *Let G_1 be an (e, d, b) -block depth-robust graph with $N = 2^n$ nodes. Then $\Pi_{cc}^{\parallel}(\text{BRG}(G_1)) \geq \min\left(\frac{eN}{2}, \frac{edb}{32}\right)$.*

Evidence for Conjecture. In the full version [BHK⁺18] we present evidence for our conjecture on the (block) depth-robustness of DRSample. We show that *all* known techniques for constructing depth-reducing sets *fail* to refute our conjecture. Along the way we introduce a general technique for bounding the size of a set S produced by Valiant’s Lemma⁸. In this attack we partition the edges into sets E_1, \dots, E_n where E_i contains the set of all edges (u, v) such that the most significant different bit of (the binary encoding of) u and v is i . By deleting j of these edge sets (e.g., by removing one node incident to each edge) we can reduce the depth of the graph to $N/2^j$. In the full version [BHK⁺18] we show that for any edge distribution function $r(v) < v$ we have

$$\mathbb{E}[|E_i|] = \frac{N}{2^i} + \sum_{j=0}^{\frac{N}{2^i}-1} \sum_{m=0}^{2^{i-1}-1} \Pr[2^{i-1} + m \geq v - r(v) > m]$$

where the value of the random variable $|E_i|$ will be tightly concentrated around its mean since for each node v the edge distribution function $r(v)$ is independent.

4.2 (Nearly) Sequential Pebblings of BRG_n have Maximum Cost

In this section, we show that for *any* constant $c \geq 1$ *any* c -parallel pebbling P of BRG_n must have cost $\Pi_{cc}(P) = \Omega(N^2)$. A pebbling $P = (P_1, \dots, P_t)$ is said to be c -parallel if we have $|P_{i+1} \setminus P_i| \leq c$ for all round $i < t$. We remark that this rules out *any* natural extension of the greedy pebbling attack e.g., the extension from the previous section that defeated the XOR extension graph G^{\oplus} was a $c = 2$ -parallel pebbling. We also remark that our proof generalizes a well-known result of [LT82] that implied that $\Pi_{st}(\text{BRG}_n) = \Omega(N^2)$ for *any* sequential pebbling. For parallel pebbings it is known that $\Pi_{st}^{\parallel} = \mathcal{O}(N^{1.5})$ [AS15] though this pebbling attack requires parallelism $c = \sqrt{N}$.

⁸ In the full version [BHK⁺18] we also analyze the performance of Valiant’s Lemma attack against Argon2i. Previously, the best known upper bound was that Valiant’s Lemma yields a depth-reducing set of size $e = \mathcal{O}\left(\frac{N \log(N/d)}{\log N}\right)$ for any DAG G with constant indegree. For the specific case of Argon2i this upper bound on e was significantly larger than the upper bound — $e = \tilde{\mathcal{O}}\left(\frac{N}{d^{1/3}}\right)$ — obtained by running the layered attack [AB17, BZ17]. Nevertheless, empirical analysis of both attacks surprisingly indicated that Valiant’s Lemma yields *smaller* depth-reducing sets than the layered attack for Argon2i. We show how to customize the analysis of Valiant’s Lemma attack to a *specific* DAG such as DRSample or Argon2i. Our theoretical analysis of Valiant’s Lemma explains these surprising empirical results. By focusing on Argon2i specifically we can show that, for a target depth d , the attacker yields a depth-reducing set of size $e = \tilde{\mathcal{O}}\left(\frac{N}{d^{1/3}}\right) \ll \mathcal{O}\left(\frac{N \log(N/d)}{\log N}\right)$, which is optimal and *matches* the performance of the layered attack [BZ17].

It is easy to show (e.g., from Lemma 2) that starting from a configuration with $|P_i| \leq e$ pebbles on the graph, it will take $\Omega(N)$ steps to advance a pebble $\mathcal{O}(e)$ steps on the top of the graph. It follows that $\Pi_{st}(\text{BRG}_n) = \Omega(N^2)$. The challenge in lower bounding $\Pi_{cc}(G)$ as in Theorem 6 is that space usage might not remain constant throughout the pebbling. Once we have proved that $\Pi_{cc}(G) = \Omega(N^2)$ we then note that any c -parallel pebbling P can be transformed into a sequential pebbling Q s.t. $\Pi_{cc}(Q) \leq c \times \Pi_{cc}(P)$ by dividing each transition $P_i \rightarrow P_{i+1}$ into c transitions to ensure that $|Q_j \setminus Q_{j-1}| \leq 1$. Thus, it follows that $\Pi_{cc}(P) = \Omega(N^2)$ for any c -parallel pebbling.

Theorem 6. *Let $G = \text{BRG}_n$ and $N = 2^n$. Then $\Pi_{cc}(G) = \Omega(N^2)$.*

The full proof of Theorem 6 can be found in the full version [BHK⁺18]. Briefly, we introduce a potential function Φ and then argue that, beginning with a configuration with at most $\mathcal{O}(e)$ pebbles on the graph, advancing the pebble e steps on the top of the graph either costs $\Omega(Ne)$ (i.e., we keep $\Omega(e)$ pebbles on the graph for the $\Omega(N)$ steps required to advance the pebble e steps) or increases the potential function by $\Omega(Ne)$ i.e., we *significantly* reduce the number of pebbles on the graph during the interval. Note that the cost $\Omega(Ne)$ to advance a pebble e steps on the top of the graph corresponds to an average cost of $\Omega(N)$ per node on the top of the graph. Thus, the total cost is $\Omega(N^2)$. Lemma 6, which states that it is expensive to transition from a configuration with *few* pebbles on the graph to a configuration with *many* well-spread pebbles on the graph, is a core piece of the potential function argument.

Lemma 6. *Let $G = \text{BRG}_n$ for some integer $n > 0$ and $N = 2^n$. Let $P = (P_1, \dots, P_t) \in \mathcal{P}(G)$ be some legal sequential pebbling of G . For a given b , partition $[N]$ into $\frac{N}{2^b} = 2^{n-b}$ intervals $I_x = [(x-1)2^b + 1, x \times 2^b]$, each having length 2^b , for $1 \leq x \leq 2^{n-b}$. Suppose that at time i , at most $\frac{N}{2^{b'+3}}$ of the intervals contain a pebble with $b' \geq b$ and at time j , at least $\frac{N}{2^{b'+1}}$ of the intervals contain a pebble. Then $|P_i| + \dots + |P_j| \geq \frac{N^2}{2^{b'+5}}$ and $(j-i) \geq \frac{2^{b-b'}N}{4}$.*

5 Empirical Analysis

We empirically analyze the quality of DRS+BRG by subjecting it to a variety of known depth-reducing pebbling attacks [AB16, AB17] as well as the “new” greedy pebbling attack. We additionally present a *new heuristic* algorithm for constructing *smaller* depth-reducing sets, which we call greedy depth reduce. We extend the pebbling attack library of Alwen et al. [ABH17] to include the greedy pebbling algorithm [BCS16] as well as our new heuristic algorithm. The source code is available on Github at <https://github.com/NewAttacksAndStrongerConstructions/PebblingAndDepthReductionAttacks>.

5.1 Greedy Depth Reduce

We introduce a novel greedy algorithm for constructing a depth-reducing set S such that $\text{depth}(G-S) \leq d_{tgt}$. Intuitively, the idea is to repeatedly find the node $v \in V(G) \setminus S$ that is incident to the largest number of paths of length d_{tgt} in $G-S$ and add v to S until $\text{depth}(G-S) \leq d_{tgt}$. While we can compute $\text{incident}(v, d_{tgt})$, the number of length d_{tgt} paths incident to v , in polynomial time using dynamic programming, it will

take $\mathcal{O}(Nd_{tgt})$ time and space to fill in the dynamic programming table. Thus, a naïve implementation would run in total time $\mathcal{O}(Nd_{tgt}e)$ since we would need to recompute the array after each iteration. This proves not to be feasible in many instances we encountered e.g. $N=2^{24}$, $d_{tgt}=2^{16}$ and $e \approx 6.4 \times 10^5$ and we would need to run the algorithm multiple times in our experiments. Thus, we adopt two key heuristics to reduce the running time. The first heuristic is to fix some parameter $d' \leq d_{tgt}$ (we used $d'=16$ whenever $d_{tgt} \geq 16$) and repeatedly delete nodes incident to the largest number of paths of length d' until $\text{depth}(G-S) \leq d_{tgt}$. The second heuristic is to select a larger set $T \subseteq V(G) \setminus S$ of k nodes (we set $k = 400 \times 2^{(18-n)/2}$ in our experiments) to delete in each round so that we can reduce the number of times we need to re-compute $\text{incident}(v, d_{tgt})$. We select T in a greedy fashion: repeatedly select a node v (with maximum value $\text{incident}(v, d')$) subject to the constraint $\text{dist}(v, T) \leq r$ for some radius r (we used $r=8$ in our experiments) until $|T| \geq k$ or there are no nodes left to add — here $\text{dist}(v, T)$ denotes the length of the *shortest directed path* connecting v to T in $G-S$. In our experiments we also minimized the number of times we need to run the greedy heuristic algorithm for each DAG G by *first* identifying the target depth value $d_{tgt}^* = 2^j$ with $j \in [n]$ which resulted in the highest quality attack against G when using *other* algorithms (Valiant’s Lemma/Layered Attack) to build the depth-reducing set S . For each DAG G we then ran our heuristic algorithm with target depths $d_{tgt} = 2^j \times d_{tgt}^*$ for each $j \in \{-1, 0, 1\}$. A more formal description of the *heuristic algorithm* can be found in the full version [BHK⁺18].

Figure 3 explicitly compares the performance of our greedy heuristic algorithm with prior state-of-the-art algorithms for constructing depth-reducing sets. Given a DAG G (either Argon2i, DRSample or DRS+BRG) on $N=2^n$ nodes and a target depth d_{tgt} we run each algorithm to find a (small) set S such that $\text{depth}(G-S) \leq d_{tgt}$. The figure on the left (resp. right) plots the size of the depth-reducing set $e=|S|$ vs. the size of the graph N (logscale) when the target depth $d_{tgt}=8$ (resp. $d_{tgt}=16$). Our analysis indicates that our greedy heuristic algorithm outperforms all prior state-of-the-art algorithms for constructing depth-reducing sets including Valiant’s Lemma [Val77] and the layered attack [AB16]. In particular, the greedy algorithm consistently outputs a depth-reducing that is 2.5 to 5 times smaller than the best depth-reducing set found by any other approach — the improvement is strongest for the DRSample graph.

5.2 Comparing Attack Quality

We ran each DAG G (either Argon2i, DRSample or DRS+BRG) with $N=2^n$ nodes against a battery of pebbling attacks including both depth-reducing attacks [AB16, AB17] and the greedy pebble attack. In our analysis we focused on graphs of size $N=2^n$ with n ranging from $n \in [14, 24]$, representing memory ranging from 16MB to 16GB. Our results are shown in Figure 2. While DRSample provided strong resistance to depth-reducing attacks (right), the greedy pebbling attack (left) yields a *very* high-quality attack (for $n \geq 20$ the attack quality is $\approx n$) against DRSample. Similarly, as we can see in Figure 2, Argon2i provides reasonably strong resistance to the *greedy pebble* attack (left), but is vulnerable to depth-reducing attacks (right). DRS+BRG strikes a *healthy* middle ground as it provides good resistance to both attacks. In particular, even if we use our new greedy heuristic algorithm to construct the depth-reducing sets (right), the attack quality never exceeds 6 for DRS+BRG. In summary, DRS+BRG provides the strongest resistance to *known* pebbling attacks for *practical* parameter ranges $n \in [14, 24]$.

As Figure 2 (right) demonstrates attack quality almost always improves when we use the new greedy algorithm to construct depth-reducing sets. The one exception was that for larger Argon2i DAGs prior techniques (i.e., Valiant’s Lemma) outperform greedy. We conjecture that this is because we had to select the parameter $d' \ll d_{tgt}^*$ for efficiency reasons. For DRSample and DRS+BRG the value d_{tgt}^* was reasonably small i.e., for DRSample we always had $d_{tgt}^* \leq 16$ allowing us to set $d' = d_{tgt}^*$. We believe that the greedy heuristic algorithm would outperform prior techniques if we were able to set $d' \sim d_{tgt}^*$ and that this would lead to even higher quality attacks against Argon2i. However, the time to pre-compute the depth-reducing set will increase linearly with d' .

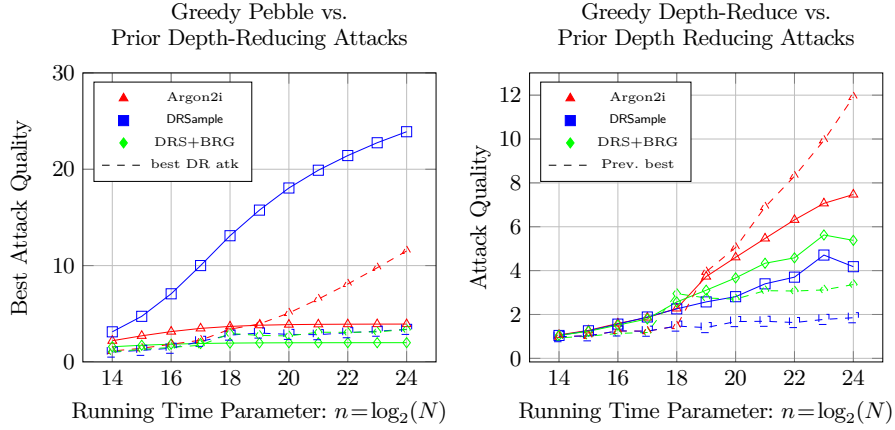


Fig. 2: Attack Quality for Greedy Pebble and Greedy Depth Reduce

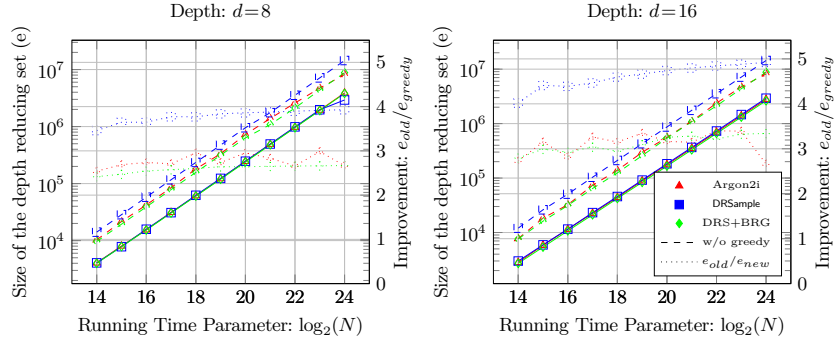


Fig. 3: Greedy Depth-Reduce vs Prior State of the Art

6 Pebbling Reduction

Alwen and Serbinenko [AS15] previously showed that, in the parallel random oracle model, the cumulative memory complexity (cmc) of an iMHFs $f_{G,H}$ can be

characterized by the black pebbling cost $\Pi_{cc}^{\parallel}(G)$ of the underlying DAG. However, their reduction assumed that the output of $f_{G,H}(x) := \text{lab}_{G,H,x}(N)$ is the label of the last node N of G where labels are defined recursively using the concatenation rule $\text{lab}_{G,H,x}(v) := H(v, \text{lab}_{G,H,x}(v_1), \dots, \text{lab}_{G,H,x}(v_\delta))$ where $v_1, \dots, v_\delta = \text{parents}_G(v)$. To improve performance, real world implementations of iMHFs such as Argon2i, DRSSample and our own implementation of BRG(DRSSample) use the XOR labeling rule $\text{lab}_{G,H,x}(v) := H(\text{lab}_{G,H,x}(v_1) \oplus \text{lab}_{G,H,x}(v_2) \oplus \dots \oplus \text{lab}_{G,H,x}(v_\delta))$ so that we can avoid Merkle-Damgard and work with a *faster* round function $H: \{0,1\}^w \rightarrow \{0,1\}^w$ instead of requiring $H: \{0,1\}^{(\delta+1)w} \rightarrow \{0,1\}^w$.

We prove that in the parallel random oracle model, the cumulative memory complexity of $f_{G,H}$ is still captured by $\Pi_{cc}^{\parallel}(G)$ when using the XOR labeling rule (under certain restrictions discussed below that will hold for all of the iMHF constructions we consider in this paper). We postpone a fully formal definition of cumulative memory complexity cmc to the full version [BHK⁺18] as it is identical to [AS15]. Intuitively, one can consider the *execution trace* $\text{Trace}_{\mathcal{A},R,H}(x) = \{(\sigma_i, Q_i)\}_{i=1}^t$ of an attacker $\mathcal{A}^{H(\cdot)}(x; R)$ on input value x with internal randomness R . Here, Q_i denotes the set of random oracle queries made in *parallel* during round i and σ_i denotes the state of the attacker immediately before the queries Q_i are answered. In this case, $\text{cmc}(\text{Trace}_{\mathcal{A},R,H}(x)) := \sum_i |\sigma_i|$ sums the memory required during each round in the parallel random oracle model⁹. For a list of distinct inputs $X = (x_1, x_2, \dots, x_m)$, let $f_{G,H}^{\times m}(X)$ be the ordered tuple $f_{G,H}^{\times m}(X) = (f_{G,H}(x_1), f_{G,H}(x_2), \dots, f_{G,H}(x_m))$. Then the memory cost of a $f_{G,H}^{\times m}$ is defined by

$$\text{cmc}_{q,\epsilon}(f_{G,H}^{\times m}) = \min_{\mathcal{A},x} \mathbb{E}[\text{cmc}(\text{Trace}_{\mathcal{A},R,H}(x))],$$

where the expectation is taken over the selection of the random oracle $H(\cdot)$ as well as the internal randomness R of the algorithm \mathcal{A} . The minimum is taken over all valid inputs $X = (x_1, x_2, \dots, x_m)$ with $x_i \neq x_j$ for $i < j$ and all algorithms $\mathcal{A}^{H(\cdot)}$ that compute $f_{G,H}^{\times m}(X)$ correctly with probability at least ϵ and make at most q queries for each computation of $f_{G,H}(x_i)$. Let $G^{\times m}$ be a DAG with mN nodes, including m sources and m sinks.

Theorem 7, our main result, states that $\text{cmc}_{q,\epsilon}(f_{G,H}^{\times m}) \geq \frac{\epsilon w m}{8\delta} \cdot \Pi_{cc}^{\parallel}(G)$. Thus, the cost of computing $f_{G,H}$ on m distinct inputs and constant indegree graphs G is at least $\Omega\left(m \times w \times \Pi_{cc}^{\parallel}(G)\right)$ — here, we assume that $H: \{0,1\}^w \rightarrow \{0,1\}^w$. We remark that for practical iMHF constructions we will have indegree $\delta \in \{2, 3\}$ so that $\text{cmc}_{q,\epsilon}(f_{G,H}^{\times m}) = \Omega\left(\Pi_{cc}^{\parallel}(G)\right)$. The δ -factor loss is necessary. For example, the complete DAG K_N has maximum pebbling cost $\Pi_{cc}^{\parallel}(K_N) \geq N(N-1)/2$, but $\text{cmc}_{q,\epsilon}(f_{K_N,H}^{\times m}) = \mathcal{O}(Nw)$ when we use the XOR labeling rule¹⁰

Theorem 7. *Let G be a DAG with N nodes, indegree $\delta \geq 2$, and $\text{parents}(u) \neq \text{parents}(v)$ for all pairs $u \neq v \in V$, and let $f_{G,H}$ be a function that follows the XOR labeling rule, with label size w . Let \mathcal{H} be a family of random oracle functions with outputs of label*

⁹ Given a constant R that represents the core/memory area ratio we can define $\text{aAT}^{\parallel}_R(\text{Trace}_{\mathcal{A},R,H}(x)) = \text{cmc}(\text{Trace}_{\mathcal{A},R,H}(x)) + R \sum_i |Q_i|$. We will focus on lower bounds on cmc since the notions are asymptotically equivalent and lower bounds on aAT complexity.

¹⁰ In particular, if we let $L_v = \text{lab}_{K_N,H,x}(v) = H(L_{v-1} \oplus \dots \oplus L_1)$ denotes the label of node v given input x then the prelabel of node v is $Y_v = \text{prelab}_{K_N,H,x}(v) = L_{i-1} \oplus \dots \oplus L_1$. Given only Y_v we can obtain $L_v = H(Y_v)$ and $Y_{v+1} = Y_v \oplus L_v$. Thus, $\text{cmc}_{q,\epsilon}(f_{K_N,H}) = \mathcal{O}(Nw)$ since we can compute $f_{K_N,H}(x) = L_N$ in linear time with space $\mathcal{O}(w)$.

length w and $H = (H_1, H_2)$, where $H_1, H_2 \in \mathcal{H}$. Let m be a number of parallel instances such that $mN < 2^{w/32}$, $q < 2^{w/32}$ be the maximum number of queries to a random oracle, and let $\frac{\epsilon}{4} > 2^{-w/2+2} > \frac{qmN+1}{2^w - m^2N^2 - mN} + \frac{2m^2N^2}{2^w - mN}$. Then $\text{cmc}_{q,\epsilon}(f_{G,H}^{\times m}) \geq \frac{\epsilon mw}{8\delta} \cdot \Pi_{cc}^{\parallel}(G)$.

As in [AS15] the pebbling reduction relies on an extractor argument to show that we can find a black pebbling $P = (P_1, \dots, P_t)$ s.t. $|P_i| = \mathcal{O}(|\sigma_i|/w)$. The extractor takes a hint h of length $|h| = |\sigma_i| + h_2$ and then extracts ℓ distinct random oracle pairs $(x_1, H(x_1)), \dots, (x_\ell, H(x_\ell))$ by simulating the attacker. Here, one can show that $\ell \geq h_2/w + \Omega(|P_i|)$, which implies that $|\sigma_i| = \Omega(w|P_i|)$ since a random oracle cannot be compressed.

There are several additional challenges we must handle when using the XOR labeling rule. First, in [AS15] we effectively use an *independent* random oracle $H_v(\cdot) = H(v, \cdot)$ to compute the label of each node v — a property that does not hold for the XOR labeling rule we consider. Second, when we use the XOR labeling it is more challenging for the extractor to extract the value of labels from random oracle queries made by the (simulated) attacker. For example, the random oracle query the attacker must submit to compute $\text{lab}_{G,H,x}(v)$ is now $\bigoplus_{i=1}^{\delta} \text{lab}_{G,H,x}(v_i)$ instead of $(v, \text{lab}_{G,H,x}(v_1), \dots, \text{lab}_{G,H,x}(v_\delta))$ — in the latter case it is trivial to read each of the labels for nodes v_1, \dots, v_δ . Third, even if H is a random oracle the XOR labeling rule uses a round function $F(x, y) = H(x \oplus y)$ that is not even collision resistant e.g., $F(x, y) = F(y, x)$. Because of this, we will not be able to prove a pebbling reduction for *arbitrary* DAGs G .

In fact, one can easily find examples of DAGs G where $\text{cmc}(f_{G,H}) \ll \Pi_{cc}^{\parallel}(G)$ i.e., the cumulative memory complexity is much less than the cumulative pebbling cost by exploiting the fact that $\text{lab}_{G,H,x}(u) = \text{lab}_{G,H,x}(v)$ whenever $\text{parents}(u) = \text{parents}(v)$. For example, observe that if $\text{parents}(N) = \{u, v\}$ and $\text{parents}(u) = \text{parents}(v)$ then

$$f_{G,H}(x) = \text{lab}_{G,H,x}(N) = H(\text{lab}_{G,H,x}(u) \oplus \text{lab}_{G,H,x}(v)) = H(0^w),$$

so that $f_{G,H}(x)$ becomes a constant function and any attempt to extract a pebbling from an execution trace computing $f_{G,H}$ would be a fruitless exercise!

For this reason, we only prove that $\text{cmc}(f_{G,H}) = \Omega(\Pi_{cc}^{\parallel}(G) \times w)$ when $G = (V = [N], E)$ satisfies the *unique parents* property i.e., for any pair of vertices $u \neq v$ we have $\text{parents}(v) \neq \text{parents}(u)$. We remark that any DAG that contains all edges of the form $(i, i+1)$ with $i < N$ will satisfy this property since $v-1 \notin \text{parents}(u)$. Thus, Argon2i, DRSample and DRSample+BRG all satisfy the unique parents property.

Extractor: We argue that, except with negligible probability, a successful execution trace must have the property that $|\sigma_i| = \Omega(w|P_i|)$ for each round of some legal pebbling P . Our extractor takes a hint, which include σ_i (to simulate the attacker), the set P_i and some (short) additional information e.g., to identify the index of the next random oracle query q_v where the label for node v will appear as input. To address the challenge that the query $q_v = \text{lab}_{G,H,x}(v) \oplus \text{lab}_{G,H,x}(u)$ we increase both the size of the hint and the number of labels being extracted e.g., our hint might additionally include the pair $(u, \text{lab}_{G,H,x}(u))$, which allows us to extract both $\text{lab}_{G,H,x}(v)$ and $\text{lab}_{G,H,x}(u)$ from q_v . Our extractor will attempt to extract labels for each node $v \in P_i$ as well as for a few extra sibling nodes such as u , which means that we must take care to ensure that we never ruin the extracted label $\text{lab}_{G,H,x}(u)$ by submitting the random oracle query $\bigoplus_{i=1}^{\delta} u_i$ to $H(\cdot)$. If G satisfies the *unique parents* property then we can prove that with high probability our *extractor* will be successful. It follows that $|\sigma_i| = \Omega(w|P_i|)$ since the hint must be long enough to encode all of the labels that we extract.

7 An Improved Argon2 Round Function

In this section we show how a parallel attacker could reduce aAT costs by nearly an order of magnitude by computing the Argon2i round function in parallel. We then present a tweaked round function to ensure that the function must be computed *sequentially*. Empirical analysis indicates that our modifications have *negligible* impact on the running time performance of Argon2 for the honest party (sequential), while the modifications will *increase* the attackers aAT costs by nearly an order of magnitude.

Review of the Argon2 Compression Function. We begin by briefly reviewing the Argon2 round function $\mathcal{G}: \{0,1\}^{8192} \rightarrow \{0,1\}^{8192}$, which takes two 1KB blocks X and Y as input and outputs the next block $\mathcal{G}(X,Y)$. \mathcal{G} builds upon a second function $\mathcal{BP}: \{0,1\}^{1024} \rightarrow \{0,1\}^{1024}$, which is the Blake2b round function [SAA⁺15]. In our analysis we treat \mathcal{BP} as a blackbox. For a more detailed explanation including the specific definition of \mathcal{BP} , we refer the readers to the Argon2 specification [BDK16].

To begin, \mathcal{G} takes the intermediate block $R = X \oplus Y$ (which is being treated as an 8x8 array of 16-byte values R_0, \dots, R_{63}), and runs \mathcal{BP} on each row to create a second intermediate stage Q . We then apply \mathcal{BP} to Q column-wise to obtain one more intermediate value Z : Specifically:

$$\begin{aligned} (Q_0, Q_1, \dots, Q_7) &\leftarrow \mathcal{BP}(R_0, R_1, \dots, R_7) & (Z_0, Z_8, \dots, Z_{56}) &\leftarrow \mathcal{BP}(Q_0, Q_8, \dots, Q_{56}) \\ (Q_8, Q_9, \dots, Q_{15}) &\leftarrow \mathcal{BP}(R_8, R_9, \dots, R_{15}) & (Z_1, Z_9, \dots, Z_{57}) &\leftarrow \mathcal{BP}(Q_1, Q_9, \dots, Q_{57}) \\ &\dots & & \\ (Q_{56}, Q_{57}, \dots, Q_{63}) &\leftarrow \mathcal{BP}(R_{56}, R_{57}, \dots, R_{63}) & (Z_7, Z_{15}, \dots, Z_{63}) &\leftarrow \mathcal{BP}(Q_7, Q_{15}, \dots, Q_{63}) \end{aligned}$$

To finish, we have one last XOR, giving the result $\mathcal{G}(X,Y) = R \oplus Z$. **ASIC vs CPU AT cost.** From the above description, it is clear that computation of the round function can be parallelized. In particular, the first (resp. last) eight calls to the permutation \mathcal{BP} are all independent and could easily be evaluated in parallel i.e., compute $\mathcal{BP}(R_0, R_1, \dots, R_7), \dots, \mathcal{BP}(R_{56}, R_{57}, \dots, R_{63})$ then compute $\mathcal{BP}(Q_0, Q_8, \dots, Q_{56}), \dots, \mathcal{BP}(Q_7, Q_{15}, \dots, Q_{63})$ in parallel. Similarly, XORing the 1KB blocks in the first ($R = X \oplus Y$) and last ($\mathcal{G}(X,Y) = R \oplus Z$) steps can be done in parallel. Thus if we let $t_{\mathcal{BP}}^{ASIC}$ (resp. $t_{\mathcal{BP}}^{CPU}$) denote the time to compute \mathcal{BP} on an ASIC (resp. CPU) we have $t_{\mathcal{G}}^{ASIC} \approx 2t_{\mathcal{BP}}^{ASIC}$ whereas $t_{\mathcal{G}}^{CPU} \approx 16 \times t_{\mathcal{BP}}^{CPU}$ since the honest party (CPU) must evaluate each call to \mathcal{BP} sequentially. Suppose that the MHF uses the round function \mathcal{G} to fill N blocks of size 1KB e.g., $N = 2^{20}$ is 1GB. Then the total area-time product on an ASIC (resp. CPU) would approximately be $(A_{mem}^{ASIC} N) \times (t_{\mathcal{G}}^{ASIC} N) \approx 2N^2 \times A_{mem}^{ASIC} t_{\mathcal{BP}}^{ASIC}$ (resp. $(A_{mem}^{CPU} N) \times (16t_{\mathcal{BP}}^{CPU} N)$ where A_{mem}^{ASIC} (resp. A_{mem}^{CPU}) is the area required to store a 1KB block in memory on an ASIC (resp. CPU). Since memory is egalitarian we have $A_{mem}^{ASIC} \approx A_{mem}^{CPU}$ whereas we may have $t_{\mathcal{BP}}^{ASIC} \ll t_{\mathcal{BP}}^{CPU}$. If we can make \mathcal{G} inherently sequential then we have $t_{\mathcal{G}}^{ASIC} \approx 16t_{\mathcal{BP}}^{ASIC}$, which means that the new AT cost on an ASIC is $16N^2 \times A_{mem}^{ASIC} t_{\mathcal{BP}}^{ASIC}$ which is eight times higher than before. We remark that the change would not necessarily increase the running time $N \times t_{\mathcal{G}}^{CPU}$ on a CPU since evaluation is already sequential. We stress that the improvement (resp. attack) applies to *all* modes of Argon2 both data-dependent (Argon2d, Argon2id) and data-independent (Argon2i), and that the attack could potentially be combined with other pebbling attacks [AB16, BCS16].

Remark 2. We remark that the implementation of \mathcal{BP} in Argon2 is heavily optimized using SIMD instructions so that the function \mathcal{BP} would be computed in parallel on *most* computer architectures. Thus, we avoid trying to make \mathcal{BP} sequential as this would slow down *both* the attacker *and* the honest party i.e., both $t_{\mathcal{BP}}^{CPU}$ and $t_{\mathcal{BP}}^{ASIC}$ would increase.

Inherently Sequential Round Function. We present a small modification to the Argon2 compression function that prevents the above attack. The idea is simply to inject extra data-dependencies between calls to \mathcal{BP} to ensure that an attacker must evaluate each call to \mathcal{BP} sequentially just like the honest party would. In short, we require the first output byte from the $i-1^{th}$ call to \mathcal{BP} to be XORed with the i^{th} input byte for the current (i^{th}) call.

In particular, we now compute $\mathcal{G}(X, Y)$ as:

$$\begin{aligned} (Q_0, Q_1, \dots, Q_7) &\leftarrow \mathcal{BP}(R_0, R_1, \dots, R_7) & (Z_0, Z_8, \dots, Z_{56}) &\leftarrow \mathcal{BP}(Q_0, Q_8, \dots, Q_{56}) \\ (Q_8, Q_9, \dots, Q_{15}) &\leftarrow \mathcal{BP}(R_8, R_9 \oplus Q_0, \dots, R_{15}) & (Z_1, Z_9, \dots, Z_{57}) &\leftarrow \mathcal{BP}(Q_1, Q_9 \oplus Z_0, \dots, Q_{57}) \\ &\dots & \dots & \\ (Q_{56}, Q_{57}, \dots, Q_{63}) &\leftarrow \mathcal{BP}(R_{56}, R_{57}, \dots, R_{64} \oplus Q_{48}) & (Z_7, Z_{15}, \dots, Z_{63}) &\leftarrow \mathcal{BP}(Q_7, Q_{15}, \dots, Q_{63} \oplus Z_6) \end{aligned}$$

where, as before, $R = X \oplus Y$ and the output is $\mathcal{G}(X, Y) = Z \oplus R$.

We welcome cryptanalysis of both this round function and the original Argon2 round function. We stress that the primary threat to passwords is brute-force attacks (not hash inversions/collisions etc...) so increasing evaluation costs is arguably the primary goal.

Implementation and Empirical Evaluation. To determine the performance impact this would have on Argon2, we modified the publicly available code to include this new compression function. The source code is available on Github at <https://github.com/antiparallel-drsbrg-argon/Antiparallel-DRS-BRG>. We then ran experiments using both the Argon2 and DRS+BRG edge distributions, and further split these groupings to include/exclude the new round function for a total of four conditions. For each condition, we evaluated 1000 instances of the memory hard function in single-pass mode with memory parameter $N = 2^{20}$ blocks (i.e., $1\text{GB} = N \times 1\text{KB}$). In our experiments, we interleave instances from different conditions to ensure that any incidental interference from system processes affects each condition equally. The experiments were run on a desktop with an Intel Core i5-6600K CPU capable of running at 3.5GHz with 4 cores. After 1000 runs of each instance, we observed only small differences in runtimes, (3%) at most. The exact results can be seen in Table 1 along with 99% confidence intervals. The evidence suggests that there is no large difference between any of these versions and that the anti-parallel modification would not cause a large increase in running time for legitimate users.

Table 1: Anti-parallel runtimes with 99% confidence

	Argon2i	DRS+BRG
Current	1405.541 \pm 1.036 ms	1445.275 \pm 1.076 ms
Anti-parallel	1405.278 \pm 1.121 ms	1445.017 \pm 0.895 ms

Acknowledgments

Seunghoon Lee was supported in part by NSF award CNS #1755708. Ben Harsha was supported by a Rolls-Royce Doctoral Fellowship. The views expressed in this paper are those of the authors and do not necessarily reflect the views of National Science Foundation or Rolls-Royce.

References

- AB16. Joël Alwen and Jeremiah Blocki. Efficiently computing data-independent memory-hard functions. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part II*, volume 9815 of *LNCS*, pages 241–271. Springer, Heidelberg, August 2016.
- AB17. Joël Alwen and Jeremiah Blocki. Towards practical attacks on argon2i and balloon hashing. In *Security and Privacy (EuroS&P), 2017 IEEE European Symposium on*, pages 142–157. IEEE, 2017.
- ABH17. Joël Alwen, Jeremiah Blocki, and Ben Harsha. Practical graphs for optimal side-channel resistant memory-hard functions. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017*, pages 1001–1017. ACM Press, October / November 2017.
- ABMW05. Martin Abadi, Mike Burrows, Mark Manasse, and Ted Wobber. Moderately Hard, Memory-bound Functions. *ACM Trans. Internet Technol.*, 5(2):299–327, May 2005.
- ABP17. Joël Alwen, Jeremiah Blocki, and Krzysztof Pietrzak. Depth-robust graphs and their cumulative memory complexity. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part III*, volume 10212 of *LNCS*, pages 3–32. Springer, Heidelberg, April / May 2017.
- ABP18. Joël Alwen, Jeremiah Blocki, and Krzysztof Pietrzak. Sustained space complexity. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part II*, volume 10821 of *LNCS*, pages 99–130. Springer, Heidelberg, April / May 2018.
- ACP⁺17. Joël Alwen, Binyi Chen, Krzysztof Pietrzak, Leonid Reyzin, and Stefano Tessaro. Script is maximally memory-hard. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part III*, volume 10212 of *LNCS*, pages 33–62. Springer, Heidelberg, April / May 2017.
- AS15. Joël Alwen and Vladimir Serbinenko. High parallel complexity graphs and memory-hard functions. In Rocco A. Servedio and Ronitt Rubinfeld, editors, *47th ACM STOC*, pages 595–603. ACM Press, June 2015.
- AT17. Joël Alwen and Björn Tackmann. Moderately hard functions: Definition, instantiations, and applications. In Yael Kalai and Leonid Reyzin, editors, *TCC 2017, Part I*, volume 10677 of *LNCS*, pages 493–526. Springer, Heidelberg, November 2017.
- BCS16. Dan Boneh, Henry Corrigan-Gibbs, and Stuart E. Schechter. Balloon hashing: A memory-hard function providing provable protection against sequential attacks. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *ASIACRYPT 2016, Part I*, volume 10031 of *LNCS*, pages 220–248. Springer, Heidelberg, December 2016.
- BDK16. Alex Biryukov, Daniel Dinu, and Dmitry Khovratovich. Argon2: new generation of memory-hard functions for password hashing and other applications. In *Security and Privacy (EuroS&P), 2016 IEEE European Symposium on*, pages 292–302. IEEE, 2016.
- Ber05. Daniel J Bernstein. Cache-timing attacks on aes. 2005.
- BHK⁺18. Jeremiah Blocki, Ben Harsha, Siteng Kang, Seunghoon Lee, Lu Xing, and Samson Zhou. Data-independent memory hard functions: New attacks and stronger constructions (full version). Cryptology ePrint Archive, Report 2018/944, 2018. <https://eprint.iacr.org/2018/944>.
- BHZ18. Jeremiah Blocki, Benjamin Harsha, and Samson Zhou. On the economics of offline password cracking. In *2018 IEEE Symposium on Security and Privacy*, pages 853–871. IEEE Computer Society Press, May 2018.
- BK15. Alex Biryukov and Dmitry Khovratovich. Tradeoff cryptanalysis of memory-hard functions. In Tetsu Iwata and Jung Hee Cheon, editors, *ASIACRYPT 2015, Part II*, volume 9453 of *LNCS*, pages 633–657. Springer, Heidelberg, November / December 2015.

- Boy07. Xavier Boyen. Halting password puzzles – hard-to-break encryption from human-memorable keys. In *16th USENIX Security Symposium—SECURITY 2007*, pages 119–134. Berkeley: The USENIX Association, 2007. Available at <http://www.cs.stanford.edu/~xb/security07/>.
- BRZ18. Jeremiah Blocki, Ling Ren, and Samson Zhou. Bandwidth-hard functions: Reductions and lower bounds. In David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang, editors, *ACM CCS 2018*, pages 1820–1836. ACM Press, October 2018.
- BZ17. Jeremiah Blocki and Samson Zhou. On the depth-robustness and cumulative pebbling cost of Argon2i. In Yael Kalai and Leonid Reyzin, editors, *TCC 2017, Part I*, volume 10677 of *LNCS*, pages 445–465. Springer, Heidelberg, November 2017.
- DGN03. Cynthia Dwork, Andrew Goldberg, and Moni Naor. On memory-bound functions for fighting spam. In *Advances in Cryptology - CRYPTO 2003*, volume 2729 of *Lecture Notes in Computer Science*, pages 426–444. Springer, 2003.
- FLW14. Christian Forler, Stefan Lucks, and Jakob Wenzel. Memory-demanding password scrambling. In Palash Sarkar and Tetsu Iwata, editors, *ASIACRYPT 2014, Part II*, volume 8874 of *LNCS*, pages 289–305. Springer, Heidelberg, December 2014.
- JWK81. Hong Jia-Wei and H. T. Kung. I/o complexity: The red-blue pebble game. In *Proceedings of the Thirteenth Annual ACM Symposium on Theory of Computing*, STOC '81, pages 326–333, New York, NY, USA, 1981. ACM.
- KDBJ17. Dmitry Khovratovich, Daniel Dinu, Alex Biryukov, and Simon Josefsson. The memory-hard argon2 password hash and proof-of-work function. *memory*, 2017.
- Lee11. Charles Lee. Litecoin, 2011.
- Len81. Thomas Lengauer. Black-white pebbles and graph separation. *Acta Informatica*, 16(4):465–475, 1981.
- LT82. Thomas Lengauer and Robert E. Tarjan. Asymptotically tight bounds on time-space trade-offs in a pebble game. *J. ACM*, 29(4):1087–1130, October 1982.
- Per09. Colin Percival. Stronger key derivation via sequential memory-hard functions. 2009. 2009.
- Pes14. Alexander Peslyak. yescrypt: password hashing scalable beyond bcrypt and scrypt, 2014.
- PHC16. Password hashing competition, 2016. <https://password-hashing.net/>.
- Pip77. N. Pippenger. Superconcentrators. *SIAM Journal on Computing*, 6(2):298–304, 1977.
- RD17. Ling Ren and Srinivas Devadas. Bandwidth hard functions for ASIC resistance. In Yael Kalai and Leonid Reyzin, editors, *TCC 2017, Part I*, volume 10677 of *LNCS*, pages 466–492. Springer, Heidelberg, November 2017.
- SAA⁺15. Marcos A. Simplicio Jr., Leonardo C. Almeida, Ewerton R. Andrade, Paulo C. F. dos Santos, and Paulo S. L. M. Barreto. Lyra2: Password hashing scheme with improved security against time-memory trade-offs. Cryptology ePrint Archive, Report 2015/136, 2015. <http://eprint.iacr.org/2015/136>.
- Val77. Leslie G Valiant. Graph-theoretic arguments in low-level complexity. In *International Symposium on Mathematical Foundations of Computer Science*, pages 162–176. Springer, 1977.
- Wie04. Michael J. Wiener. The full cost of cryptanalytic attacks. *Journal of Cryptology*, 17(2):105–124, March 2004.