

Efficiently Computing Data-Independent Memory-Hard Functions

Joël Alwen
IST Austria

Jeremiah Blocki
Microsoft Research

February 19, 2016

Abstract

A memory-hard function (MHF) f is equipped with a *space cost* σ and *time cost* τ parameter such that repeatedly computing $f_{\sigma,\tau}$ on an application specific integrated circuit (ASIC) is not economically advantageous relative to a general purpose computer. Technically we would like that any (generalized) circuit for evaluating an iMHF $f_{\sigma,\tau}$ has area \times time (AT) complexity at $\Theta(\sigma^2 * \tau)$. A data-independent MHF (iMHF) has the added property that it can be computed with almost optimal memory and time complexity by an algorithm which accesses memory in a pattern independent of the input value. Such functions can be specified by fixing a directed acyclic graph (DAG) G on $n = \Theta(\sigma * \tau)$ nodes representing its computation graph.

In this work we develop new tools for analyzing iMHFs. First we define and motivate a new complexity measure capturing the amount of *energy* (i.e. electricity) required to compute a function. We argue that, in practice, this measure is at least as important as the more traditional AT-complexity. Next we describe an algorithm \mathcal{A} for repeatedly evaluating an iMHF based on an arbitrary DAG G . We upperbound both its energy and AT complexities per instance evaluated in terms of a certain combinatorial property of G .

Next we instantiate our attack for several general classes of DAGs which include those underlying many of the most important iMHF candidates in the literature. In particular, we obtain the following results which hold for all choices of parameters σ and τ (and thread-count) such that $n = \sigma * \tau$.

- The Catena-Dragonfly function of [FLW13] has AT and energy complexities $O(n^{1.67})$.
- The Catena-Butterfly function of [FLW13] has complexities is $O(n^{1.67})$.
- The Double-Buffer and the Linear functions of [CGBS16] both have complexities in $O(n^{1.67})$.
- The Argon2i function of [BDK15] (winner of the Password Hashing Competition [PHC]) has complexities $O(n^{7/4} \log(n))$.
- The Single-Buffer function of [CGBS16] has complexities $O(n^{7/4} \log(n))$.
- *Any* iMHF can be computed by an algorithm with complexities $O(n^2 / \log^{1-\epsilon}(n))$ for all $\epsilon > 0$. In particular when $\tau = 1$ this shows that the goal of constructing an iMHF with AT-complexity $\Theta(\sigma^2 * \tau)$ is unachievable.

Along the way we prove a lemma upper-bounding the depth-robustness of any DAG which may prove to be of independent interest.

1 Introduction

Moderately hard to compute functions have proven to be useful security primitives. In this work we focus on “memory-hard functions” (MHF) introduced in [Per09]. These aim to serve as password hashing algorithms for storing passwords in a login system, as Key Derivation Functions (also called “key stretching” functions) for password-based cryptography and for building Proof-of-Effort protocols (in particular for use in cryptocurrencies such as Litecoin [Cha11], [Bil13] and others). In each case the main security property we would like to achieve for the MHF is that brute-force attacks (i.e. evaluating the MHF on many inputs) using an application-specific integrated circuit (ASIC) should not be economically viable.

1.1 Memory-Hard Functions and Their Complexity

We interpret this intuitive goal in two ways. Either the cost of *building* the ASIC should be prohibitively expensive (in terms of say USD) or the cost of *running* the ASIC should be prohibitively expensive. In fact, given that the former is a one-time cost which can be amortized over the life-time of the device while the later is a recurring cost, it may often be the case that the later is the most interesting goal to achieve.

The cost of building a circuit is often approximated by its *AT-complexity* [Tho79, BL13, BK15, AS15]; that is the product of the area of the chip and the time it takes the chip to produce the output. In this work we consider MHFs built as modes of operation over an underlying compression function H . Thus, we measure time in units of *tocks*; namely the time it takes to evaluate one instance of H from start to finish.¹ We measure area in units of “memory-area” (MAr); namely the area required to store one output of H (called a *block*). Finally we parametrize our AT-complexity notion AT_R with the *core-memory area ratio* [BK15] $R > 0$, a positive real denoting the number of MAr required to implement one copy of H .²

To estimate the cost of running the chip we will use a new notion which we call the “energy-complexity” or *E-complexity* of the circuit. Intuitively, it approximates the energy (say in kilo-Watt-hours) used in an execution of the chip. More precisely the unit of measure for E-complexity is a “memory-Watt-tock” (MWt) – the number of kWh it takes to store one block for one tock. We also parametrize the complexity notion $E_{\bar{R}}$ with the *core-memory energy ratio*, a positive real $\bar{R} > 0$ which is the number of MWt required to evaluate one instance of H .

To see why this is an interesting measure for achieving our stated security goal consider the case of password hashing. (The case for KDFs follows essentially the same reasoning.) Suppose an attacker manages to pilfer the credentials file from a login-server and now executes an off-line brute-force attack \mathcal{A} implemented in an ASIC with core-memory energy ratio \bar{R} using $E_{\bar{R}}(\mathcal{A})$ MWt per password guess. We model the monetary income from such an attack as being proportional to the number of password guesses made which we denote by $\#eval$.³ Conversely, we model the running cost as being proportional to the electricity consumed by the ASIC while executing the attack, namely its $E_{\bar{R}}$ -complexity times $\#eval$. The attacker can always increase income (i.e. increase $\#eval$) simply by adding more implementations of \mathcal{A} to the ASIC or running the ASIC for more time. Therefore, the attack is profitable (in this model) if and only if the USD cost c of one MWt and the income i per password guess are such that $i > c * E_{\bar{R}}(\mathcal{A})$. Thus we can use $E_{\bar{R}}(\mathcal{A})$ as a key indicator for which ranges of (c, i) an attack is economically viable.

When using an MHF to construct a Proof-of-Work for a cryptocurrency the application of energy complexity follows a slightly different reasoning. In this context our security goal is to make mining (the practice of generating new coins by repeatedly evaluating the MHF until certain type of input/output is found) economically uninteresting on an ASIC. In contrast to the previous case, it now matters to a miner how much time is needed for the evaluations. Indeed, in practice mining ASICs are usually specified by manufacturers in terms of their throughput ($\#eval/s$) [BP1, BF1, Spo]. Thus we model the rate of profit of a mining rig as being proportional to

$$\frac{profit}{second} \approx \frac{\#eval}{second} - \frac{energy}{second} = \frac{\#eval - energy}{second}.$$

A miner can always increase $\#eval/second$ by implementing more copies of the MHF on the ASIC. Therefore, to keep the profit down we want the energy consumption to increase accordingly. How much it increases per increment of $\#eval$ is given precisely by the amortized energy complexity of the (most efficient circuit implementing the) MHF. Thus, demonstrating an MHF with a high enough energy-complexity results in a Proof-of-Work protocol for which running any mining ASIC costs as much as the income it generates from freshly minted coins.⁴

¹I.e. without considering pipelining and other such amortized optimizations.

²This allows our analysis to be applied regardless of the particular VLSI technology employed and the particular implementation of H used when constructing the ASIC.

³Intuitively, the more passwords guesses made the higher the expected number of password (equivalents) recovered by the adversary which can then be monetized.

⁴This calculation may also help explain why, in practice, mining ASIC vendors tend to specify their devices in terms of $\#eval/second$ (which approximates rate of income) [BP1, BF1, Spo], while third-party sites comparing devices across vendors tend focus equally on the energy-complexity of the devices (which is just as important when estimating rate of profit) [MR1, Coi].

Quality of an Attack. A candidate MHF F is specified via an algorithm which evaluates it. (E.g. [Per09, FLW13, BDK15].) We refer to this algorithm as the *naïve* algorithm \mathcal{N} for F and it is understood to be the algorithm used by the honest party. \mathcal{N} is intended to be an algorithm that can be evaluated efficiently on *typical* (i.e. general purpose) computer architectures — where we may not be able to evaluate H multiple times in parallel. As usual, we are interested in what advantage an adversarial evaluation algorithm can have over the honest party. Therefore, one measure of the quality of a given algorithm \mathcal{A} for evaluating (multiple instances of) an MHF F is to compare its complexity to that of \mathcal{N} . In particular for given core-memory ratios R and \bar{R} the *AT-quality* and *energy-quality* of \mathcal{A} are given by

$$\text{AT-quality}_R(\mathcal{A}) = \frac{AT_R(\mathcal{N})}{AT_R(\mathcal{A})} \quad \text{and} \quad \text{E-quality}_{\bar{R}}(\mathcal{A}) = \frac{E_{\bar{R}}(\mathcal{N})}{E_{\bar{R}}(\mathcal{A})}.$$

Here, $AT_R(\mathcal{A})$ (resp. $E_{\bar{R}}(\mathcal{A})$) measures the *amortized* AT complexity (resp. *amortized* energy complexity).⁵ That is $AT_R(\mathcal{A})$ is smallest $AT_{\mathbb{R}}$ complexity of a chip implementing \mathcal{A} divided by $\#inst(\mathcal{A})$ — the number of instances of F computed in an execution of \mathcal{A} . We consider \mathcal{A} an “attack” if either one of these quality measures is greater than 1. (However we remark that all attacks in this work have both qualities *simultaneously* tending towards infinity as $\#inst$ grows.)

Data-Independent and Ideal MHFs. An *data-independent* memory-hard function (iMHF) is a function f for which the associated naïve algorithm \mathcal{N} , on input x , computes $f(x)$ using a memory access pattern that is independent of x . These take on special importance in applications where the MHF is to be evaluated on secret input in an (at least somewhat) hostile environment. This is because (in contrast to their siblings *data-dependent* MHFs) it is much easier to implement an iMHF in such a way that it avoids information leakage via certain side-channel attacks such as Timing attacks. In these attacks, the variation in the time taken to perform certain operations is used to deduce information about the inputs upon which the MHF is being evaluated. Similar attacks have in the past been mounted by local adversarial processes [BM06], adversarial virtual machines in a cloud environment [RTSS09] or even completely remotely [Ber, ASK07]. Therefore, in the context of both KDFs and password hashing data-independence is a desirable property. All MHFs considered in this work are of this form.

In general, an iMHF f can be described via a fixed DAG G representing its computation graph. Each node represents an intermediary value, which is computed via some deterministic round function, using the values represented by the parent nodes in G (e.g. via a single call to H). The source node of G represents the input x while $f(x)$, the output of the computation, is the value represented by the sink node.

Let f be an iMHF given by some DAG G of size n with constant in-degree. There exists a trivial algorithm triv which can always compute f with AT and energy complexities $\Theta(n^2)$.⁶ Given a constant $c > 1$ we consider f to be a *c-ideal* iMHF if, when we take the naïve algorithm to be triv , there exist no attack \mathcal{A} on f with better quality than c (i.e. $\forall \mathcal{A} \text{ E-quality}_{\bar{R}}(\mathcal{A}) \leq c$ and $\text{AT-quality}_R(\mathcal{A}) \leq c$). A primary goal of research in this field is to find an ideal iMHF.⁷

1.2 MHF Candidates

Due to the growing interest in MHFs there are a number of candidate functions. For example in the recently completed Password Hashing Competition [PHC] most entrants claimed some form of memory-hardness. The goal of the PHC was to select a winning algorithm to act as a new standard for password hashing.

⁵Generally, unless explicitly specified otherwise, we are only interested in the *amortized* AT and energy complexities per instance of the MHF computed.

⁶Simply compute each intermediary value in topological order, one value at a time, storing all results in memory until the computation is complete.

⁷Hopefully one permuting as simple as possible an explicit description and naïve implementation and as lightweight as possible round-function.

Catena. To the best of our knowledge the earliest candidate iMHF is the PHC finalist Catena [FLW13]. It received special recognition for its agile framework and its resistance to side-channel attacks. In [FLW13] the authors proposed two different DAGs giving rise to two separate functions. The first, called Catena Bit Reversal, is based on an λ -layered graph BRG_λ^n with n nodes. The second is called Catena Double Butterfly and is based on a different $O(\lambda \log n)$ -layered graph DBG_λ^n . The Catena designers recommended choosing $\lambda \in \{1, 2, 3, 4\}$ [FLW13].

Argon2. One of the most important MHF candidates is Argon2 [BDK15]. Notably, it is the winner of the Password Hashing Competition [PHC]. Argon2 is equipped with a data-dependent mode of operation and an independent mode which is called Argon2i.

Balloon Hashing Most recently, three new candidate iMHFs are proposed in [CGBS16]. These are called the Single-Buffer (SB), Double-Buffer and Linear constructions respectively and are jointly referred to as the Balloon Hashing constructions. The authors provide strong evidence for the memory-hardness of all three candidates albeit assuming the absence of parallelism.

In general iMHF candidates are equipped with a *space-cost* parameter σ (in which the memory required per evaluation is intended to scale) and a *time-cost* parameter τ (in which the time required for an evaluation is intended to scale)⁸. Additionally, Argon2i, the Double-Buffer and the Linear functions are also equipped with a *parallelism* parameter ϕ the property that the naïve algorithm can make efficient use of (up to) ϕ concurrent threads. Viewing these functions as DAGs gives rise to a graph on $n = \sigma * \tau$ nodes. The naïve algorithm \mathcal{N} will run in time n/ϕ and will use space σ at each step. Thus, the energy cost will be $\Theta(\sigma^2 * \tau/\phi)$. The hope is that for all settings of (σ, τ, ϕ) the AT and energy complexity lie in $\Theta(\sigma^2 * \tau/\phi)$.

1.3 Our Contributions

In this work we introduce and motivate the notion of (amortized) energy complexity. Next we give a generic evaluation algorithm PGenPeb for data-independent iMHFs based on arbitrary DAG G . We analyze PGenPeb’s energy and AT complexities in terms of a combinatorial property of G . In particular, we obtain an attack against any iMHF for which G is not *depth-robust*. Informally, a DAG G is not depth-robust if there is a relatively small set S of nodes such that after removing S from G the resulting graph (denoted $G - S$) has low depth (i.e. contains only short paths).

We instantiate the attack for various classes of DAGs. In particular, we exhibit a “depth-reducing” node set S for the Argon2i DAG, both types of Catena DAGs and all three Balloon Hashing DAGs. For example, for any parameters $(\sigma, \tau, \phi = 1)$ with $n = \sigma * \tau$ we obtain an attack on both Catena, the Double-Buffer and the Linear iMHFs with quality $\Omega(n^{1/3})$. Similarly we demonstrate an attack on Argon2i and the Single-Buffer iMHF with quality $\Omega\left(\frac{n^{1/4}}{\ln n}\right)$.⁹

In fact we demonstrate that no DAG with constant indegree is sufficiently depth-robust to completely resist the attack. More precisely, we show that any iMHF is at best c -ideal for $c = \Omega(\log^{1-\epsilon} n)$ and any $\epsilon > 0$. In particular this means that ideal iMHFs, as described above, do not exist.

General Attack on Non-Depth Robust DAGs. We first present in Section 3, a generic evaluation algorithm GenPeb which takes as inputs a node subset S . Because G is not depth-robust there exists a small set S of nodes such that $d = \text{depth}(G - S)$ is relatively small. The basic idea behind our attack is to divide computation steps into two phases: balloon phases and light phases. Each light phase lasts roughly $g \gg d$ time steps. During light phases we discard most of the values that we have computed from memory keeping only values corresponding to nodes in S , the highest node i whose value has been computed and the parents

⁸ τ corresponds to the number of passes through memory. For Catena DAGs this corresponds to the parameter λ , the number of layers in the DAG.

⁹For the cases when $\phi > 1$ PGenPeb maintains the same complexities but the resulting quality decreases somewhat as the complexity of the naïve algorithm improves for Argon2i, the Double-Buffer and the Linear functions. In other words quality decreases not because memory-hardness increases but because the honest algorithm becomes more efficient.

of the nodes whose values we plan to compute in the next g time steps. As the name suggests, light phases are cheap. Our memory usage is low during these light phases and we will compute one instance of the round function (e.g. call to H) during each time step. During a Balloon Phase we quickly restore all of the discarded values to memory so that we can complete the next light phase. Unlike light phases, the balloon phases are more expensive because we are storing up to $O(n)$ values in memory and because we will often make multiple calls to the round function in parallel. However, the key observation is that we will not incur these higher cost in too many time steps. In particular, because the graph $G - S$ has small depth $d \ll g$ and we never discard values for nodes in S the Balloon Phase can be completed very quickly (i.e., in at most $d \ll g$ times steps) by making parallel calls to the round function.

While for any non-depth-robust graph the GenPeb algorithm has good energy complexity, obtaining an evaluation algorithm with low AT-complexity requires a bit more work. Notice that during a light phase most of the memory capacity and round function implementations needed for a balloon phase are no longer being used. Moreover light phases run for significantly more time than the balloon phases. These observations give rise to the low AT-complexity parallel algorithm PGenPeb which evaluates g/d instances of the iMHF concurrently such that at any given time only a single instance is in a balloon phase while all other instances are in light phases. Intuitively this results in more efficient use of available hardware while technically we get that the energy complexity of the algorithm is approximately equal to the AT complexity (Theorem 3.5).

Stacked Sandwich Graphs. In Section 4 we focus on two classes of DAGs called (strict) stacked sandwich graphs. Informally, a DAG G is a λ -stacked sandwich DAG if the nodes can be partitioned into $\lambda + 1$ layers such that, with the possible exception of node i , all of the parents of node $i + 1$ are from previous layers. These classes include the DAGs implicit to both Catena iMHFs as well as the Double-Buffer and Linear iMHFs. We prove that no λ -stacked sandwich graph is depth-robust (Lemma 4.2). For any $t > 1$ there is a set S of n/t nodes such that $\text{depth}(G - S) \leq (\lambda + 1)t$.

(n, δ, w) -Random Graphs. In Section 5 we turn to a class of random graphs called (n, δ, w) -random DAGs. We remark that the graphs implicit to Argon2i and the Single-Buffer iMHF (for a randomly chosen salt) fall into this category of random DAGs. We show (in Lemma 5.3) that, with high probability, by removing just a few nodes these graphs can be transformed into stacked sandwich graphs and are thus not depth-robust.

Attack on any iMHF. In Section 6 we prove that no DAG with constant indegree is sufficiently depth-robust to resist at least some form of attack (Theorem 6.3). In our proof, we rely on a result due to Valiant [Val77] which states that for any DAG G with m edges and depth d there is a set S of $m/\log d$ edges s.t. by deleting them we obtain a graph of depth at most $d/2$ (see Lemma 6.1). Given $\epsilon > 0$ we can repeatedly apply this result obtain a set S of $o\left(\frac{\delta n}{\log^{1-\epsilon} n}\right)$ nodes s.t $\text{depth}(G - S) \leq \frac{n}{\log^2 n}$. Thus if we let the naïve algorithm be (any algorithm complexity comparable to) triv then we have a generic attack \mathcal{A} with quality $\text{AT-quality}_R(\mathcal{A}) = \Omega(\delta^{-1} \log^{1-\epsilon} n)$ and $E_{\bar{R}}(\mathcal{A}) = \Omega(\delta^{-1} \log^{1-\epsilon} n)$.

Exact Security Analysis. Finally we present exact bounds for the energy and AT complexities of all of our attacks. Our analysis demonstrate that our attacks have high quality for practical values of n and \bar{R} — not just as $n \rightarrow \infty$. For example setting $n = 2^{18}$ we already have an attack \mathcal{A} against Argon2i with $\text{AT-quality}_R(\mathcal{A}), \text{E-quality}_{\bar{R}}(\mathcal{A}) > 1$ — using a realistic value $\bar{R} = 3,000$. In general, $\text{E-quality}_{\bar{R}}(\mathcal{A})$ will increase as n increases or as \bar{R} decreases.

1.4 Related Work

The intuitive goal of constructing functions for which VLSI implementations are prohibitively expensive was first laid out by Percival in [Per09]. This property was formalized by asking that evaluating such a function on a PRAM requires large ST-complexity. In particular evaluation algorithms with low amortized complexity such as those in this work were not considered. Percival also introduced the first, and currently most widely

deployed, candidate MHF called `script`. A full proof of security under a strong security definition remains a central open problem in the area. However recently significant progress has been made in this direction in [ACK⁺16]. It is interesting to note though that despite `script` being data-*dependent* the (conditional) lower bound in [ACK⁺16] still does not exceed the upper-bound of Section 6 on the best possible quality of an iMHF.

Catena. In [FLW13] the authors of Catena restricted their analysis of its security to a sequential setting. That is they restrict an adversary to only being able to evaluate one instance of the underlying function H at a time. In this setting and for the case when $\lambda = 1$ the results of [LT82] show that, in a simplified computational model, BRG_1^n has ST-complexity $\Omega(n^2)$. Here ST-complexity denotes the product of the space and time required by any algorithm which evaluates Catena Bit Reversal. The intuition being that large ST-complexity implies large AT-complexity of any implementation in a custom chip.

Argon2. Argon2 [BDK15] was the winner of the Password Hashing Competition [PHC]. Argon2 is equipped with a data-dependent mode of operation and an independent mode which is called Argon2i. The authors recommend using Argon2i for password hashing due to its resistance to side channel attacks [BDK15]. Our attacks only apply to Argon2i, the data independent mode. Recently, Gibbs et al. [CGBS16] gave an attack on Argon2i which reduces the cost of computing Argon2i by a factor of 4.

Balloon Hashing. In [CGBS16] the authors also proposed three iMHFs which resist their attack on Argon2i. These are called Single-Buffer (SB), Double-Buffer and Linear and collectively referred to as the Balloon Hashing iMHFs.¹⁰ Our attacks reduce the cost of computing both Argon2i and SB by a factor of $\tilde{\Omega}(n^{1/4})$.

A Provably Secure MHF. Currently, the only candidate MHF equipped with a full proof of security is the one in [AS15]. There, the authors show an iMHF F for which the energy-complexity of the required storage alone (i.e. disregarding the cost of evaluating the round function) is within a polylogarithmic factor in n of the energy-complexity of the trivial algorithm `triv`. Moreover `triv` uses only a single instance of H (i.e. it is sequential) which implies that, roughly speaking, any evaluation algorithm for F can have E-quality = $O(\text{polylog}(n))$. The results in Section 6 show that this is optimal for any iMHF up to the exponent in the polylogarithmic factor.

Attacking MHFs. The Catena Dragonfly iMHF has been attacked previously [BK15, AS15]. In particular, [AS15] demonstrated an attack on Catena Dragonfly $\text{BRG}_{\lambda=1}^n$ which has energy quality E-quality = $O(\sqrt{n})$. The attack from [BK15] has slightly worse quality $O(n^{1/5})$, but it applies even for Dragonfly variants in which $\lambda > 1$. At a high level the ideas behind both of these attacks is to divide memory into segments, store the leading block in each segment and then recompute the remaining blocks as needed. These attacks only work because the underlying Catena Dragonfly DAG BRG_λ^n allows for quick re-computation of the remaining blocks. In this work we observe that this key idea can be generalized to attack *any* non depth-robust iMHF. In particular, our techniques can be used to attack other iMHFs like Catena Butterfly, Argon2i [BDK15] and SB [CGBS16]. In fact, our attacks can be extended to *any* iMHF because no DAG is sufficiently depth-robust to resist at least some form of attack.

Memory-Bound Functions. An important precursor to memory-hard functions are memory-*bound* functions. First introduced in [ABMW05] here the complexity measure of interest is the number of cache misses required to evaluate the function. On the highest level the motivation is the same as that of memory-hard functions; namely to build moderately hard functions which are more equally hard across different computational devices (compared to the rather unbalanced notion of plain computational complexity). In

¹⁰Gibbs et al. [CGBS16] use “Balloon Hashing” as a title for their iMHF however this similarity with the balloon phase in our evaluation algorithm is a slightly unfortunate coincidence.

particular it was observed that while computational speeds may vary greatly between different devices the same is not as true for memory latency speeds [DGN03]. In contrast memory-hard functions aim to achieve egalitarian hardness by making the cost of custom hardware prohibitively large [Per09]. The first provably secure memory-bound function was (implicitly) given in [DGN03] where it was used to construct a protocol for fighting SPAM email. The construction was later improved in [DNW05] which was also the first result in cryptography to make use of a version of the pebbling model of computation; a technique later adapted in [AS15].

Password Storage. Recent high-profile security breaches (e.g., RockYou, Sony, LinkedIN, Ashley Madison¹¹) highlight the importance of proper password storage practices like salting [Ale04] and key stretching [MT79]¹². However, hash iteration, the technique used by password hash functions like PBKDF2 [Kal00] and `bcrypt` [PM], is typically an insufficient defense against an adversary who could build customized hardware to evaluate the underlying hash function. In particular, the cost of computing a hash function H like SHA256 or MD5 on an ASIC is orders of magnitude smaller than the cost of computing H on traditional hardware [DGN03, NB⁺15]. By contrast, memory costs tend to be relatively stable across different architectures [DGN03], which motivates the use of memory-hard functions for password hashing [Per09].

Several orthogonal lines of research have explored defenses such as: distributing the storage and/or computation of a password hash across multiple servers (e.g., [BJKS03, CLN12]), storing fake password hashes on the server (e.g., [JR13, BBBB10]), the inclusion of secret salt values (e.g., “pepper”) in password hashes [Man96] and the inclusion of the solution(s) to hard AI challenges in password hashes [CHS06, DC08, BBD13].

2 Preliminaries

We begin with some notation. Given a directed acyclic graph (DAG) $G = (V, E)$ of size $|V| = n$ and a subset $S \subseteq V$ we use $G - S$ to denote the resulting DAG after removing all nodes in S . We denote by $\text{depth}(G)$ the length of the longest (directed) path in G and we denote by $\text{indeg}(G)$ the maximum number of directed edges entering a single node. For integers $a \leq b$ we write $[a, b]$ as shorthand for the set $\{a, a + 1, \dots, b\}$ and we write $[a]$ for the set $[1, a]$.

We use $H_\lambda = \sum_{i=1}^\lambda \frac{1}{i}$ to denote the λ 'th harmonic number. In particular H_λ can be approximated by the natural logarithm $H_\lambda \approx \ln \lambda$.

2.1 Complexity and Quality of Attacks

We consider algorithms in the parallel random oracle model (pROM) [AS15] of computation.¹³ That is an algorithm is repeatedly invoked. At invocation $i \in \{1, 2, \dots\}$ the algorithm is given the state (bit-string) σ_{i-1} it produced at the end of the previous invocation. Next \mathcal{A} can make a batch of calls $\mathbf{q}_i = (q_{1,i}, q_{2,i}, \dots)$ to the underlying round function H (modeled as a random oracle (RO)). Then it receives the response from H and can perform arbitrary computation before finally outputting an updated state σ_i . The initial state σ_0 contains the input to the computation which terminates once a special final state is produced by \mathcal{A} . Apart from the explicit states σ the algorithm may keep no other state between invocations. For an input x and coins r we denote by $\mathcal{A}(x; r; H)$ the corresponding (deterministic) execution of \mathcal{A} .

We define the runtime $\text{time}(\mathcal{A})$ to be the maximum running time of \mathcal{A} in any execution (over all choices of x , r and H). Then the *cumulative memory complexity* (CMC) and *cumulative RO complexity* are defined as

¹¹See <http://www.privacyrights.org/data-breach/> (Retrieved 9/1/2015).

¹²Users routinely select lower entropy password [Bon12], which are especially vulnerable to an offline attacker when the underlying password hash function is inexpensive to compute. Furthermore, stricter password restrictions (e.g., requiring a mix of numbers and upper/lower case letters) [SS09] have not been found to greatly improve the entropy of the resulting passwords [KSK⁺11, BKPS13]. In fact, sometime these policies reduced the entropy of user selected passwords [KSK⁺11]. These policies are often associated with high usability costs [FH10].

¹³Alternatively the results in this work also apply to the random access machine model of computation.

$$\text{cmc}(\mathcal{A}) = \max_{x,r,H} \sum_{i \in [T-1]} |\sigma_i| \quad \text{crc}(\mathcal{A}) = \max_{x,r,H} \sum_{i \in [T]} |\mathbf{q}_i|$$

where $|\sigma|$ is the bit-length of state σ , $|\mathbf{q}|$ is the dimension of the vector \mathbf{q} and $\max_{x,r,H}$ denotes the maximum over all possible executions of \mathcal{A} . Similarly the *absolute memory complexity* (AMC) and *absolute RO complexity* are defined to be (ARC)

$$\text{amc}(\mathcal{A}) = \max_{x,r,H} \max_{i \in [T-1]} |\sigma_i| \quad \text{arc}(\mathcal{A}) = \max_{x,r,H} \max_{i \in [T]} |\mathbf{q}_i|.$$

We remark that these complexity measures are stricter than is common, especially with respect to maximizing over all random oracles H . However we use them to upper-bound the complexity of our attacks so this strictness can only serve to strengthen the results.

Using these tools we can now define the complexity of an algorithm as follows.

Definition 2.1 (AT and Energy Complexities) *Let \mathcal{A} be a pROM algorithm which computes $\#inst(\mathcal{A})$ instances of an iMHF in parallel. Then for any core-memory area ratio $R > 0$ and any core-memory energy ratio $\bar{R} > 0$ the (amortized) AT-complexity and the (amortized) energy-complexity of \mathcal{A} are defined to be*

$$AT_R(\mathcal{A}) = [\text{amc}(\mathcal{A}) + R \cdot \text{arc}(\mathcal{A})] \times \frac{\text{time}(\mathcal{A})}{\#inst(\mathcal{A})} \quad E_{\bar{R}}(\mathcal{A}) = \frac{\text{cmc}(\mathcal{A}) + \bar{R} \cdot \text{crc}(\mathcal{A})}{\#inst(\mathcal{A})}.$$

Finally we can define the quality of an attack in terms of how much (if at all) it improves on the naïve algorithm

Definition 2.2 (Attack Quality) *Let f be an MHF with naïve algorithm \mathbb{N} and let \mathcal{A} be a pROM algorithm for evaluating $\#inst(\mathcal{A})$ instance(s) of f . Then for any core-memory area ratio $R > 0$ and any core-memory energy ratio $\bar{R} > 0$ the AT-quality and energy-quality of \mathcal{A} is defined to be*

$$\text{AT-quality}_R(\mathcal{A}) = \frac{AT_R(\mathcal{N})}{AT_R(\mathcal{A})} \quad \text{E-quality}_{\bar{R}}(\mathcal{A}) = \frac{E_{\bar{R}}(\mathcal{N})}{E_{\bar{R}}(\mathcal{A})}.$$

In particular if either quantity is less than 1 then we call \mathcal{A} an attack on f .

Let f be an iMHF based on some DAG G of size n with constant in-degree. Observe f can always be evaluated by computing one intermediate value at a time in topological order while never deleting a computed value. Clearly this always results in correctly computing f and it corresponds to a well defined pROM algorithm triv for evaluating f . Moreover $AT_R(\text{triv}) = \Theta(n(n+R))$ and $E_{\bar{R}}(\text{triv}) = \Theta(n(n+\bar{R}))$. Given a constant $c > 0$ we say that f is a c -ideal iMHF if, when $\text{triv} = \mathcal{N}$ is the naïve, for any attack \mathcal{A} we have $\text{AT-quality}_R(\mathcal{A}) \geq c$ and $\text{E-quality}_{\bar{R}}(\mathcal{A}) \geq c$. This is motivated by the observation that for any iMHF algorithm triv is always a possible way to evaluate it. An ideal iMHF captures the property that triv is (approximately) the best evaluation strategy possible.

Unfortunately we will later show that c -ideal iMHFs do not exist for any constant $c > 0$. As $n \rightarrow \infty$ we will have $\text{AT-quality}_R(\mathcal{A}) = \omega(1)$ and $\text{E-quality}_{\bar{R}}(\mathcal{A}) = \omega(1)$.

2.2 Pebbling and Graph Theory

We provide some shorthand for describing algorithms and give some useful graph theoretic definitions and lemmas.

Graph Pebbling. To simplify exposition, our attacks are often described in the language of parallel graph pebbling [AS15]. However, unlike in [AS15], we merely think of this as shorthand for describing an evaluation strategy of an iMHF rather than describing an algorithm in a distinct model of computation.

In particular any iMHF f which we consider is based on some fixed underlying DAG G with (a single source and sink node) which describes which values are used as inputs to which calls to the round function. To compute f on some input x each node of G is assigned a value (bit-string). The source receives the value

x . The value of any other node v is defined to be the output of the round function applied to the values of the parent nodes of v . Finally $f(x)$ is defined to be the value of the sink node.¹⁴

With this in mind, each round of pebbling corresponds to one invocation in an execution. Placing a pebble on a node v in some round is shorthand for computing the value of v by computing the round function on the values of v 's parents. Clearly this can only be done if (x_1, \dots, x_z) are stored in memory and so, if an algorithm places a pebble on a node whose parents do not all contain a pebble then we call such a move *illegal*. Thus we will always show that our pebbling strategies only produce legal pebbings in order to ensure that they correspond to a feasible pROM algorithm for evaluating iMHF. Finally having a pebble on a node at the end of a round corresponds to storing the value of that node in the state σ for that invocation.

Graph Theory. The key insight behind our attacks is that if a graph is not depth-robust enough then it can be efficiently pebbled.

Definition 2.3 (Depth Robust and Depth Reducible DAGs) For $e, d \in \mathbb{N}$ a DAG $G = (V, E)$ is called (e, d) -depth-robust if

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

If G is not (e, d) -depth-robust then we say that G is (e, d) -reducible. For $0 \leq \beta \leq \alpha \leq 1$ a DAG of size n is (α, β) -fractionally depth-robust if it is $(\alpha n, \beta n)$ -depth-robust.

In order to prove the generic attack on any iMHF we rely on a lemma, originally due to Valiant [Val77], to show that no graph is depth-robust enough not to permit at least some sort of attack.

3 Generic Attack

In this section we describe a general pebbling attack **GenPeb** against any (e, d) -reducible graph. **GenPeb** (G, S, g, d) takes as input a DAG $G = (V, D)$ and a set $S \subseteq V$ of size e such that $\text{depth}(G - S) \leq d$ and a parameter $g \geq d$ which we will define below. In every round **GenPeb** makes progress (i.e., places a pebble on node i in the i 'th round). Thus, $\text{time}(\text{GenPeb}) = n$ as the algorithm will place a pebble on the final node n in the n 'th rounds. Intuitively, **GenPeb** is divided into two types of phases: Balloon Phases and a Light Phases. During light phases we throw out most of the pebbles on the graph keeping only pebbles on nodes in S , the highest pebbled node i and the parents of the nodes $[i, i + g]$ that we plan to pebble in the next g rounds. Every g rounds we execute a balloon phase to ensure that we will always have pebbles placed on the parents of the nodes that we plan to pebble in the next g rounds. Because we never remove pebbles on nodes in S and the DAG $G - S$ has $\text{depth} \leq d$ we will be able to accomplish this goal in at most d rounds. During light phases we keep at most $\delta g + e$ pebbles on the graph and we place at most one new pebble on G in every round. Thus the total cost during all light phases is at most $n(\delta g + e + \bar{R})$. While we may incur higher costs during a balloon phase we are only in the balloon phase for at most $\frac{dn}{g}$ rounds.

We analyze the energy complexity **GenPeb** in terms of the depth reduction parameters e and d . These results are summarized in Theorem 3.4. While **GenPeb** will lead to attacks with good energy-quality **E-quality** $_{\bar{R}}$ the attack may not necessarily have good AT-quality **AT-quality** $_{\bar{R}}$. This is because **GenPeb** may still have high absolute memory and RO complexity due to the balloon phase. However, we can easily circumvent this problem by pebbling multiple copies of the DAG G in parallel, which corresponds to evaluating multiple independent instances of the iMHF. In particular, **PGenPeb** pebbles $\lfloor g/d \rfloor$ instances of G in parallel. We stagger evaluation of the different iMHF instances so that at most one of the $\lfloor g/d \rfloor$ pebbling instances is in a balloon phase at any point in time. To accomplish this **PGenPeb** simply waits $(i - 1)d$ steps to begin pebbling the i 'th instance of G . Thus, **PGenPeb** takes at most $n + \lfloor g/d \rfloor d \leq 2n$ steps to complete. The absolute memory and RO complexity of **PGenPeb** is essentially just the cost of the balloon phase for a single iMHF instance. Thus, **PGenPeb** leads to attacks with good AT-quality **AT-quality** $_{\bar{R}}$ because the cost of the balloon phase can be amortized among the $\lfloor g/d \rfloor$ iMHF instances we compute. Theorem 3.5 states both

¹⁴For concreteness, though not relevant to this work, in most cases the round function is simply the compression function H (with the exception of the Linear iMHF of [CGBS16]).

the energy and AT complexity of PGenPeb. The energy complexity of PGenPeb is roughly equivalent to the energy complexity of GenPeb, and the AT-complexity of PGenPeb is roughly twice the energy complexity of PGenPeb.

In the rest of the paper we will consider several specific families of DAGs like the underlying DAGs in the Catena and Argon2i iMHFs. For Catena, we can find a set $S \subseteq V$ of size $e = n/t$ such that $\text{depth}(G - S) = O(t)$ for every $t > 1$. For Argon2i we can find a set S with expected size $O(n/t + (n \ln \lambda)/\lambda)$ such that $\text{depth}(G - S) \leq t \cdot \lambda$. Combined with Theorem 3.5 we will obtain an attack on Catena with quality $\Omega(n^{1/3})$ and an attack on Argon2i with quality $\Omega(n^{1/4}/\ln n)$.

GenPeb makes use of two subroutines `need` and `keep`. In our complexity analysis we omit the cost of computing these functions. However we stress that in all our attacks they are either trivial (constant) or very easy to compute. By “easy to compute” we mean that the sets returned by these subroutines will have a short description size (e.g., “all nodes” or $[i, j]$) and that it will be trivial to decide whether a given node v is in these sets.

We begin with some useful notation. Fix a DAG G of size n and number its nodes in (arbitrary) topological order from 1 to n . For $i \in [n]$ and $j \geq i$ we write $\text{parents}(i, j)$ for the set of nodes v with an edge (v, u) for some $u \in [i, \min\{j, n\}]$. Next we fix the class of functions from which `need` and `keep` must be chosen in order to prove that GenPeb produces a legal pebbling (and thus defines a pROM evaluation algorithm).¹⁵

Algorithm 1: GenPeb (G, S, g, d)

Arguments : $G = (V, E)$, $S \subseteq V$, $g \in [\text{depth}(G - S), |V|]$, $d \geq \text{depth}(G - S)$

Local Variables: $n = |V|$

```

1 for  $i = 1$  to  $n$  do
2   Pebble node  $i$ .
3    $l \leftarrow \lfloor i/g \rfloor * g + d + 1$ 
4   if  $i \bmod g \in [d]$  then // Balloon Phase
5      $d' \leftarrow d - (i \bmod g) + 1$ 
6      $N \leftarrow \text{need}(l, l + g, d')$ 
7     Pebble every  $v \in N$  which has all parents pebbled.
8     Remove pebble from any  $v \notin K$  where  $K \leftarrow S \cup \text{keep}(i, i + g) \cup \{n\}$ .
9   else // Light Phase
10     $K \leftarrow S \cup \text{parents}(i, i + g) \cup \{n\}$ 
11    Remove pebbles from all  $v \notin K$ .
12  end
13 end
```

Definition 3.1 (Needed Pebbles) Fix a subset of target nodes $T \subseteq V$ and a pebbling configuration $C \subseteq V$ of G .¹⁶ Then a node $v \in V$ is needed for T within d' steps if there exists a completely unpebbled path P ¹⁷ of length $\geq d'$ from v to some node in T . We use $N_{C,T,d'}$ to denote the set of all such nodes. We use $K_{C,T}$ to denote the set of all nodes $v \in C$ such that $v \in T$ or v has a child v' such that $v' \in \bigcup_{i=0}^n N_{C,T,i}$.

Definition 3.2 (Valid need and keep) We say that the pair of functions `need` and `keep` is valid for GenPeb(G, S, d, g) if we always have $\text{need}(i, j, d') \supseteq N_{C,[i,j],d'}$ and $\text{keep}(i, j) \supseteq K_{C,[i,j]}$ whenever GenPeb($G, S, \text{depth}(G - S), g$) queries `need` or `keep`.

In our generic iMHF attack we use the trivial functions `need` and `keep` which always output V (e.g., during the balloon phase we pebble every node we can during each round and we never discard any pebbles). The following fact is easy to see:

¹⁵Later on we instantiate `need` and `keep` in several ways but will always prove that they are valid for the inputs we use them for.

¹⁶That is fix a set C of nodes of V which currently have a pebble on them.

¹⁷That is $P \cap C = \emptyset$.

Fact 3.3 (Generic Valid Subroutine) Fix a DAG $G = (V, E)$ and let `need` and `keep` be the constant function returning V . Then the pair `need` and `keep` is valid for $\text{GenPeb}(G, S, d, g)$ for any set $S \subseteq V$ and any parameters $g \geq d \geq \text{depth}(G - S)$.

While we would already obtain high quality attacks on Catena and Argon2i by using the generic `need` and `keep` subroutines, we show how our attacks can be optimized further by defining the subroutines `need` and `keep` more carefully.

We remark that by leaving `need` and `keep` undefined for now we leave some flexibility in the implementation of the balloon phase in $\text{GenPeb}(G, S, g, d)$. During each round of a balloon phase we may pebble any $v \in V$ which has all parents pebbled, but we are only required to add pebbles to these nodes once it becomes absolutely necessary to finish the balloon phase in time (e.g., there are only d' rounds left in the balloon phase and the vertex v is part of an completely unpebbled path to T of length $\geq d'$). Similarly, we are allowed to remove pebbles provided that they are no longer needed for the balloon phase (e.g., every path to T from that node has an intermediate pebble).

The easiest way to satisfy these conditions is to simply pebble every $v \in V$ which has all parents pebbled, and to never remove pebbles during the balloon phase (Fact 3.3). Indeed this is exactly what we do in our general attack on iMHFs. However, we demonstrate that further optimizations are possible against the Catena and Argon2i iMHFs (e.g., each of the new pebbles we add during a Catena balloon phase does not need to remain on the DAG very long). In each case the subroutines `need` and `keep` will have very simple instantiations — we will not need to perform complicated computations like breadth first search to find these sets.

Fix any G, S, g and d and let $M(G, S, g, d)$ be the largest number of pebbles simultaneously on $G - S - \text{parents}(i, i + g)$ during any round i which is in a Balloon phase of $\text{GenPeb}(G, S, g, d)$.¹⁸ Similarly let $C(G, S, g, d)$ be the largest number of pebbles placed on G during any single round in a Balloon Phase. In the following we prove that GenPeb always produces a legal pebbling. Thus it describes a well formed pROM algorithm \mathcal{A} for evaluating an iMHF based on G . We also show how to use $M(G, S, g, d)$ and $C(G, S, g, d)$ to upper-bound the energy-complexity of \mathcal{A} with hardcoded inputs (G, S, g, d) .

Theorem 3.4 (Energy Complexity of GenPeb) Let $G = (V, E)$ be a DAG, with $\text{indeg}(G) = \delta$. Further let $S \subseteq V$ with $|S| = e$ and $d \geq \text{depth}(G - S)$ and let integer $g \in [d, n]$. Fix any valid pair of subroutines `need` and `keep` and let \mathcal{A} be the pROM algorithm described by GenPeb with hardcoded inputs (G, S, g, d) . Then \mathcal{A} produces a valid pebbling and for any core-memory energy ratio \bar{R} and $M = M(G, S, g, d)$ and $C = C(G, S, g, d)$ it holds that:

$$\begin{aligned} \text{cmc}(\mathcal{A}) &\leq n \left(\frac{d \cdot M}{g} + \delta g + e \right) & \text{crc}(\mathcal{A}) &\leq n \left(\frac{\min\{d \cdot C, n\}}{g} + 1 \right) \\ E_{\bar{R}}(\mathcal{A}) &\leq n \left(\frac{d \cdot M + \min\{dC, n\} \cdot \bar{R}}{g} + \delta g + e + \bar{R} \right). \end{aligned}$$

PROOF. We first prove that $\text{GenPeb}(G, S, g, d)$ produces a legal pebbling to ensure that \mathcal{A} is a well defined algorithm. Then we upper-bound its energy complexity.

Recall that pebbles can be removed at will and by definition, in Step 7, GenPeb only places a pebble if it is legal to do so. Thus the only illegal move could come due to Step 2. Assume no illegal pebble has been placed up to node i . To show that i is then also pebbled legally it suffices to show that each of its parents $P \subseteq V$ are have a pebble at the beginning of round i . The most recent Balloon Phase to have completed before round i (if any) consisted of rounds $B = [i', i' + d]$ where $i - (i' + d) \leq g$. Consider the partition $P_1 = P \cap [i' + d + 1, i - 1]$, $P_2 = P \cap [i', i' + d]$ and $P_3 \cap [1, i' - 1]$ of the set of parents P . By assumption all $v \in P_1$ were pebbled (legally) in the previous g rounds using Step 2 and so were not removed (by definition of K in Step 10). Moreover by assumption all nodes in P_2 where pebbled by Step 2 during B and so were not removed (by definition of K in Step 8 and the validity of the subroutine `keep` the pebble is not removed during the balloon phase $B = [i', i' + d]$ and by definition of K in Step 10 the pebble was not removed during

¹⁸Recall that there are n/g Balloon Phases and the j^{th} Balloon Phase consists of rounds $\{jg + 1, \dots, jg + d\}$

rounds $[i' + d + 1, i - 1]$). Thus it suffices to prove that all $v \in P_3$ contained a pebble at some point during B since then by definition of K in Steps 2 and 10 they too will not be removed.

Let P_4 be the subset of P_3 which don't contain a pebble at the start of B . (If it is empty we are done.) Otherwise, for a given round j let \mathbf{p}_j be all paths which end with a node in P_4 and are unpebbled at the beginning of the round. Let l_j be the length of the longest path in \mathbf{p}_j . We argue that $\forall j \in [i', i' + d - 1]$ then $l_j \leq d - (j - i')$. If this is the case then we are done. Entering the final round $i' + d - 1$ the length of the longest unpebbled path is $l_{i'+d-1} \leq 1$ so by the end of the final round of B all nodes P_4 – the end points of paths \mathbf{p}_j – are pebbled.

We argue that $l_j \leq d - (j - i')$ by induction. Clearly, this is true when $j = i'$ as $l_j \leq \text{depth}(G - S) \leq d$. Now assume that $l_j \leq d - (j - i')$ for some $j \in [i', i' + d - 1]$, let $p \in \mathbf{p}_j$ denote a longest path and let v denote the starting node of p . We first observe that either the starting node v of p has no parents or they are all pebbled.¹⁹ Second, we observe that, because **need** is valid, in round j we must either have $v \in N$ or we must have $l_j < d - (j - i')$. In the latter case we have $l_{j+1} \leq l_j \leq d - (j + 1 - i')$ — because **keep** is valid we are not allowed to remove pebbles from any of the parents of v . In the former case we have $l_{j+1} \leq l_j - 1 \leq d - (j + 1 - i')$ because $v \in \text{need}(i', i' + g, d' = d - (j + 1 - i'))$ by the validity of **need**. Thus in Step 7 of round j node v is pebbled. Finally it remains there till the end of the round since $v \in \text{keep}(j, j + g)$ because there is a completely unpebbled path from v 's children in p to $[i', i' + g + d]$. This completes the proof that **GenPeb** produces a legal pebbling.

Recall that the energy-complexity of \mathcal{A} can be computed as $E_{\bar{R}}(\mathcal{A}) = \text{cmc}(\mathcal{A}) + \bar{R} \cdot \text{crc}(\mathcal{A})$. To upper-bound $\text{cmc}(\mathcal{A})$ we can sum upper-bounds on the cmc of the Balloon phases and the cmc of the Light phases. To compute the Balloon phase term notice that **GenPeb** is in a Balloon Phase for nd/g steps and during each round i of a balloon phase there are, by definition, at most $M(G, S, g, d)$ extra pebbles on $G - S - \text{parents}(i, i + g)$. On the other hand, there are clearly at most n Light phase steps and at the start of each round i of a light phase there are no pebbles on $G - S - \text{parents}(i, i + g)$. Finally, during each round i we pay cumulative memory cost at most e to keep pebbles on nodes in S and at most δg to keep pebbles on nodes in the set $\text{parents}(i, i + g)$, which can be of size at most δg . Adding these three terms and factoring out an n term we get that $\text{cmc}(\mathcal{A}) \leq n \left(\frac{d \cdot M(G, S, g, d)}{g} + \delta g + e \right)$.

Placing a pebble on G corresponds to making a call to H . To upper-bound $\text{crc}(\mathcal{A})$ we observe that in any round of a Light phase only one pebble is ever placed on G (namely in Step 2). During each balloon phase we place at most $C(G, S, g, d)$ pebbles on the graph in each round, and at most n pebbles on the graph in total. Thus we can write $\text{crc}(\mathcal{A}) \leq n \left(\frac{\min\{n, d \cdot C(G, S, g, d)\}}{g} + 1 \right)$. Combing this with the bound on cmc and rearranging terms we obtain the theorem. \square

The following Theorem 3.5 upper-bounds the complexity of **PGenPeb**. The proof closely follows the analysis of **GenPeb** in Theorem 3.4. The key difference is that we evaluate multiple instances, and at that at most one of these instances is in a balloon phase at any point in time. Thus, we get a much tighter bound on AT -complexity because the worst-case memory usage M is approximately the same as the average memory usage of **PGenPeb**.

Theorem 3.5 (Complexity of PGenPeb) *Let $G = (V, E)$ be a DAG, with $\text{indeg}(G) = \delta$. Further let $S \subseteq V$ with $|S| = e$ and $d \geq \text{depth}(G - S)$ and let integer $g \in [d, n]$. Fix any valid pair of subroutines **need** and **keep** and let \mathcal{A} be the pROM algorithm described by **PGenPeb** with hardcoded inputs $(G, S, g, d, \lfloor \frac{g}{d} \rfloor)$. Then for any core-memory area and energy ratios $R > 0$ and $\bar{R} > 0$ and $M = M(G, S, g, d)$ and $C = C(G, S, g, d)$ it holds that:*

$$\begin{aligned} AT_R(\mathcal{A}) &\leq 2n \left[\frac{d(M + RC)}{g} + \delta g + e + R \right] \\ E_{\bar{R}}(\mathcal{A}) &\leq n \left[\frac{dM + \min\{d\bar{R}C, n\bar{R}\}}{g} + \delta g + e + \bar{R} + 1 \right] \end{aligned}$$

¹⁹as otherwise it wouldn't be a longest path in \mathbf{p}_j

Algorithm 2: PGenPeb (G, S, g, d, k)

Arguments : $G, S \subseteq V, g \in [\text{depth}(G - S), |V|]$ $d \geq \text{depth}(G - S), k \leq \lfloor \frac{g}{d} \rfloor$
Local Variables: $n = |V|$, copies $G_1, \dots, G_k = G, S_1, \dots, S_k = S$

```
1 for  $t = 1$  to  $n + kd$  do
2   Parallel for  $j = \max\{1, \frac{t-n}{d}\}$  to  $\min\{k, \frac{t-1}{d}\}$  do
3      $i \leftarrow t - jd$ 
4     Pebble node  $i$  in  $G_j$ .
5     if  $i = n$  then
6       Remove pebbles from all  $v \notin \{n\}$  in  $G_j$ 
7       Break
8     end
9      $l \leftarrow \lfloor i/g \rfloor * g + d + 1$ 
10    if  $i \bmod g \in [d]$  then // Balloon Phase
11       $d' \leftarrow d - (i \bmod g) + 1$ 
12       $N_j \leftarrow \text{need}_j(l, l + g, d')$ 
13      Pebble any  $v \in N_j$  which has all parents pebbled.
14      Remove pebble from any  $v \notin K_j$  where  $K_j \leftarrow S_j \cup \text{keep}_j(i, i + g) \cup \{n\}$ .
15    else // Light Phase
16       $K_j \leftarrow S_j \cup \text{parents}_j(i, i + g) \cup \{n\}$ 
17      Remove pebbles from all  $v \notin K_j$ .
18    end
19  end
20 end
```

PROOF. Algorithm \mathcal{A} evaluates $\#inst(\mathcal{A}) = k = \lfloor \frac{g}{d} \rfloor$ instances of our iMHF G_1, \dots, G_k in parallel. The validity of PGenPeb follows from the validity of GenPeb (Theorem 3.4). Thus \mathcal{A} is a well-defined pROM algorithm and we can focus on the AT and energy-complexities of \mathcal{A} .

We begin with AT-complexity. We claim $\text{amc}(\mathcal{A}) \leq k(e + \delta g) + M$ where $M = M(G, S, g, d)$. To see this observe that during any time step t we are in the balloon phase for at most one iMHF instance G_j ($1 \leq j \leq k$). By definition M upper bounds the number of pebbles on nodes in $V_j - S_j$ during the balloon phase, δg upper bounds the number of pebbles on nodes in $V_{j'} - S_{j'}$ during the light phase for other iMHF instances $G_{j'}$ at any point in time. Finally, at any point in time we can have at most $k \cdot e$ pebbles on nodes in $S_1 \cup \dots \cup S_k$. Similarly we observe that $\text{arc}(\mathcal{A}) \leq C + (k - 1)$ (where $C = C(G, S, g, d)$) because each iMHF instance makes exactly one query to H during a light phase and at most C queries to H during a balloon phase. Thus we can write:

$$AT_R(\mathcal{A}) \leq [(k - 1)(e + \delta g) + M + R \cdot C + R(k - 1)] \times \frac{(n + kd)}{k}$$

because \mathcal{A} runs for $n + kd$ time steps. Rearranging terms and using the facts that $(k - 1)/k < 1$ and $kd \leq n$ we get the first inequality in the theorem.

Now we turn to energy-complexity. For each of the k iMHF instances we use cumulative memory $n \left(\frac{d \cdot M}{g} + \delta g + e \right)$ to compute the instance (Theorem 3.4) plus an extra cost of up to kd to store a pebble on node n for up to kd steps once the iMHF instance finishes. Thus, thus $\text{cmc}(\mathcal{A}) \leq kn \left(\frac{d \cdot M}{g} + \delta g + e \right) + k^2 d$. Moreover, for each of the k iMHF instances we place pebbles on the graph at most $n \left(\frac{d \cdot C}{g} + 1 \right)$ times. Thus, $\text{crc}(\mathcal{A}) \leq kn \left(\frac{d \cdot C}{g} + 1 \right)$ which, (together with the fact that $kd \leq n$) implies the second inequality of the theorem. \square

4 Sandwich Graphs

In this section we focus on the two Catena hash functions [FLW13] as well as the second two Balloon Hashing constructions of [CGBS16]. The first Catena iMHF is given by the *Catena Bit Reversal Graph* (which we denote BRG_λ^n); an n node DAG which consists of a stack of $\lambda \in \mathbb{N}_{\geq 1}$ bit-reversal graphs [LT82]. Each node in a layer is associated with a $\log_2 \left(\frac{n}{\lambda+1} \right)$ bit string and edges between layers correspond to the bit reversal operation²⁰. The Catena designers recommended choosing $\lambda \in \{1, 2, 3, 4\}$ [FLW13]. The second Catena hash function is an iMHF based on the *Catena Double Butterfly Graph*, denoted DBG_λ^n . It is an n node DAG with $O(\lambda \log n)$ layers of nodes.

The “Double-Buffer” and “Linear” iMHFs of [CGBS16] consist of τ layers of σ nodes for a total of $n = \tau * \sigma$ nodes. Each layer is a path with its origin connected to the final node in the path of the previous layer. Moreover all nodes at layers $\tau \geq i \geq 1$ have 20 incoming edges from nodes selected uniformly and independently in the previous layer. In the “Double-Buffer” construction the hash of a node is given by hashing the concatenation of all parent node labels while in the “Linear” construction the parent node labels are first XORed together before being hashed for greater throughput. However this difference will not affect the results in this work.²¹

In this section we demonstrate that all of these DAGs can be computed with lower then hoped for energy and AT complexities (simultaneously) regardless of the random choices made when constructing the graphs. In particular, the iMHF corresponding to both BRG_λ^n and DBG_λ^n can be evaluated with amortized AT complexity $AT_R(\mathcal{A}) = O(n^{5/3} + Rn^{4/3})$ and energy complexity $E_{\bar{R}}(\mathcal{A}) = O(n^{5/3} + Rn^{4/3})$ for any value of λ . In fact, our attacks hold for a more general class of graphs characterized by Definition 4.1 below. Thus, to understand our attacks it is not critical to know the exact specification of these DAGs just that both DAGs are strict sandwich graphs. We refer an interested reader to Appendix A for the actual definitions of the Catena DAGs BRG_λ^n and DBG_λ^n .

Definition 4.1 ((Strict) λ -Stacked Sandwich Graphs) *Let $n, \lambda \in \mathbb{N}_{\geq 1}$ be a integers such that $\lambda + 1$ divides n and let $k = n/(1 + \lambda)$ and let G be a DAG with n nodes. We say that G is a λ -stacked sandwich DAG if G contains a directed path of n nodes (v_1, \dots, v_n) with arbitrary additional edges connecting nodes from lower layers $L_j \doteq \{v_{jk+1}, \dots, v_{j(k+\lambda)}\}$ with $j \leq i$ to the $i + 1^{\text{st}}$ layer L_{i+1} . If the DAG has no edges of the form (u, v) with $u \in L_j$ and $v \in L_{j+2+i}$ for $i \geq 0$ then we say it is a strict λ -stacked sandwich DAG.*

In particular, the Catena bit reversal graph BRG_λ^n is a strict λ -stacked sandwich DAG with n nodes and maximum indegree $\text{indeg} = 2$ — see Fact A.1 in the appendix. The Catena double butterfly graph DBG_λ^n is a strict $(\lambda(2x - 1) + 1)$ -stacked sandwich DAG with n nodes, where $x \leq \log n$ is the integer such that $n = 2^x \cdot (\lambda(2x - 1) + 1)$ — see Fact A.2 in the appendix. Finally, for any parameters t and s , a randomly chosen DAG for the Double-Buffer and Linear iMHFs is a strict t -stacked sandwich graph on $n = ts$ nodes with probability 1.

Summary of The Results in this Section. Lemma 4.2 upper-bounds the depth-robustness of any λ -stacked sandwich DAG G — any λ -stacked sandwich DAG is $(n/t, \lambda t + t)$ -reducible. Thus we can apply the generic attack (Theorem 3.5) to get an upper-bound on the energy and AT complexities of such graphs (Theorem 4.5) and so, in particular, also for the 4 constructions mentioned above. Theorem 4.5 states that there is an attack \mathcal{A} with $AT_R(\mathcal{A})$ and energy complexity $E_{\bar{R}}(\mathcal{A}) = O((\lambda + \delta)n^{5/3} + \bar{R}n^{4/3})$, where δ denotes the maximum indegree of the DAG.

While these results will also be useful in the next section focused on Argon2i for the 4 constructions above the results can be improved somewhat by observing that the constructions are actually a *strict* stacked

²⁰The parameter λ in Catena is related to the parameter τ in Argon2i. The intended space complexity of Catena is $\sigma = 2n/(\lambda + 1)$ and the intended computation time is n . Thus, the intended energy complexity is $2n^2/(\lambda + 1)$.

²¹We remark that we have assumed that the thread count parameter $p = 1$. However we observe that for $p > 1$ the resulting DAG has an almost identical distribution except that every s/p^{th} edge along the path forming a layer is removed. This can only make the job easier of an evaluation algorithm. In particular the complexity of our attacks for the case $p = 1$ are an upper-bound on the complexities of these constructions for $p > 1$.

sandwich DAGs. In particular, we can further decrease the resulting complexity if we first define more targeted **need** and **keep** functions and prove that they are valid for PGenPeb when G is a *strict* λ -stacked sandwich DAG (Lemma 4.3). Theorem 4.6 says that there is an attack \mathcal{A} with $AT_R(\mathcal{A})$ and energy complexity $E_{\bar{R}}(\mathcal{A}) = O(\delta n^{5/3} + \bar{R}n^{4/3})$.

The following Lemma upper-bounds the depth-robustness of any λ -stacked sandwich DAG G . By combining this observation with the generic attack from the previous section we can obtain strong attacks on any λ -stacked sandwich DAG G .

Lemma 4.2 (Sandwich Graphs are Reducible) *Let G be a λ -stacked sandwich DAG then for any integer $t \geq 1$ G is $(n/t, \lambda t + t - \lambda - 1)$ -reducible.*

PROOF. Let $S = \{v_{it} \mid 1 \leq i \leq n/t\}$. We claim that $\text{depth}(G - S) \leq \lambda t + t - \lambda - 1$. Consider any path P in $G - S$. For each layer L_j the path P can contain at most $t - 1$ nodes from layer L_j because any sequence of t consecutive nodes $v_i, v_{i+1}, \dots, v_{i+t}$ must contain at least one node in S . Thus,

$$|P| \leq \sum_{i=0}^{\lambda} |P \cap L_j| \leq (\lambda + 1)(t - 1).$$

□

Algorithm 3: Function: $\text{need}(x, y, d')$

Arguments: $x, y \geq x, d' \geq 0$

Constants : Pebbling round i, g, t .

1 $j \leftarrow (i \bmod g)$ // Current Layer is $L_{\lfloor j/t \rfloor}$
2 **Return** $L_{\lfloor j/t \rfloor} \cap \{it + j \mid i \leq \frac{n}{t}\}$

Algorithm 4: Function: $\text{keep}(x, y)$

Arguments: $x, y \geq x$

Constants : Pebbling round i, g, t .

1 $j \leftarrow (i \bmod g)$
2 $\ell \leftarrow \lfloor j/t \rfloor$ // Current Layer
3 **Return** $L_{\geq \ell-1}$

Lemma 4.3 (Valid need and keep for Strict Sandwich Graphs) *Let G be a strict λ -stacked sandwich DAG on n nodes, let $S = \{it \mid i \leq \frac{n}{\lambda+1}\}$, $d = (\lambda + 1)t$ and $g \geq d$ then the functions **need** and **keep** from Algorithm 3 and Algorithm 3 are valid for GenPeb(G, S, g, d).*

PROOF. Given a round i let $l_i = \lfloor i/g \rfloor * g + d + 1$, let $j_i = i \bmod g$, $k_i = \lfloor j_i/t \rfloor$ and let $d'_i = d - j_i + 1$. Let $B = [i', i' + d]$ denote the rounds of a balloon phase. We first claim that $\forall i \in B$ the following is true at the beginning of the round i .

1. If $k_i \geq 1$ then L_{k_i-1} is pebbled.
2. $L_{k_i} \cap \{it + q \mid i \leq \frac{n}{t} \wedge 0 \leq q < j_i\}$ is pebbled.
3. Any set $L_{\geq k_i} \setminus (S \cup \{it + q \mid i \leq \frac{n}{t} \wedge 0 \leq q < j_i\})$ is unpebbled.

Suppose for now that claims (1)–(3) hold for every round $i \in B$. Then at the beginning of round i there is no unpebbled path from any node $v \in L_{<k_i}$ to $[l_i, l_i + g]$ as any such path must pass through layer L_{k_i-1} . Similarly, there is no completely unpebbled path from any node $v \in L_{k_i} \cap \{it + q \mid i \leq \frac{n}{t} \wedge 0 \leq q < j_i\}$ to $[l_i, l_i + g]$ as every such v is pebbled already. Finally, the maximum length of a path to $[l_i, l_i + g]$ starting from any node in $L_{\geq k_i} \setminus (L_{k_i} \cap \{it + j_i \mid i \leq \frac{n}{t}\})$ is less than d'_i . Thus, any unpebbled path of length $\geq d'_i$ must start in the set $L_{\lfloor j_i/t \rfloor} \cap \{it + j_i \mid i \leq \frac{n}{t}\}$ returned by $\text{need}_{l_i, l_i + g, d'_i}$ in round i . Similarly, any pebbled node with a child v who has an unpebbled path to $[l_i, l_i + g]$ in $L_{\lfloor j_i/t \rfloor} \cap \{it + j_i \mid i \leq \frac{n}{t}\}$ cannot be in a layer $L_{<\lfloor j_i/t \rfloor - 1}$.

Claims (1)–(3) trivially hold at the beginning of the balloon phase $i = i'$. If we assume that these claims hold at the start of some round i then the query $\text{need}_{l_i, l_i + g, d'_i}$ returns the set $L_{\lfloor j_i/t \rfloor} \cap \{it + j_i \mid i \leq \frac{n}{t}\}$. Each $v \in L_{\lfloor j_i/t \rfloor} \cap \{it + j_i \mid i \leq \frac{n}{t}\}$ has all of its parents pebbled as any parent is either in layer $L_{\lfloor j_i/t \rfloor}$ which is completely pebbled or in $L_{\lfloor j_i/t \rfloor} \cap \{it + j_i - 1 \mid i \leq \frac{n}{t}\}$ which is also completely pebbled. Thus, v will be pebbled in Step 7 of **GenPeb** and this pebble will not be removed in Step 8 as $v \in \text{keep}(i, l_i + g)$. Thus, Claim (2) will hold at the start of round $i+1$. If $k_i = k_{i+1}$ then Claim (1) will clearly hold since we start round i with pebbles on L_{k_i-1} and $L_{k_i-1} \subseteq \text{keep}(i, l_i + g)$ so these pebbles won't be removed. Otherwise if $k_{i+1} = k_i + 1$ then we note that $j_i + 1 = t$ so we will have pebbles on every node in $L_{k_i} \cap \{it + q \mid i \leq \frac{n}{t} \wedge 0 \leq q \leq t - 1\} = L_{k_i}$ by the end of round i . Hence, Claim (1) will still hold. Claim (3) will still hold because we do not start with any pebbles on nodes $N_i = L_{\geq k_{i+1}} \setminus (S \cup \{it + q \mid i \leq \frac{n}{t} \wedge 0 \leq q < j_{i+1}\})$ in round i , $N_{i+1} \subseteq N_i$ and $\text{need}(l_i, l_i + g, d'_i)$ does not contain any nodes in N_{i+1} . \square

Lemma 4.4 (Valid need and keep for Sandwich Graphs) *Let G be a λ -stacked sandwich DAG on n nodes, let $S = \{it \mid i \leq \frac{n}{\lambda+1}\}$, $d = (\lambda+1)t$ and $g \geq d$. Further, let keep be the constant function that returns V and let need be the function from Algorithm 3. Then the pair need and keep are valid for $\text{GenPeb}(G, S, g, d)$.*

PROOF. (Sketch) Essentially follows the proof of Lemma 4.3. The only change we make is to claim (1). We will have pebbles on all nodes in $L_{\leq k_i-1}$ at the start of round i instead of pebbles on just L_{k_i-1} . Observe that $V \supset V'$ for any set V' returned by the version of keep defined in Algorithm 3. Thus, claims (1)–(3) follow as before. \square

Theorem 4.5 follows easily from Theorem 3.5, Lemma 4.4 and Lemma 4.2 by setting $g = n^{2/3}$ and $t = n^{1/3}$.

Theorem 4.5 (Complexity of Sandwich Graph) *Let F be an i MHF based on DAG G ; a λ -stacked sandwich DAG on n nodes with $\lambda < n^{1/3}$ and maximum indegree $\text{indeg}(G) = \delta$. Then for any core-memory area and energy ratios R and \bar{R} there exists an evaluation algorithm \mathcal{A} with*

$$\begin{aligned} AT_R(\mathcal{A}) &\leq 2n^{5/3} \left[(1 + \delta) + (\lambda + 1) + \frac{R}{n^{1/3}} + \frac{2R + (\lambda + 1)R}{n^{2/3}} \right] \\ E_{\bar{R}}(\mathcal{A}) &\leq n^{5/3} \left[(\lambda + 1) + (\delta + 1) + \frac{3\bar{R} + 1}{n^{2/3}} + \frac{\bar{R}}{n^{1/3}} \right] \end{aligned}$$

PROOF. First, assume that $\lambda < n^{1/3}$. Let $S = \{it \mid i \leq n/t\}$ denote the set S from the proof of Lemma 4.2 with size n/t and $\text{depth}(G-S) \leq \lambda t + t - \lambda - 1$. \mathcal{A} runs $\text{PGenPeb}(G, S, g, d = \lambda t + t)$ with the function need from Algorithm 3 and the constant function $\text{keep}(x, y) = V$. By Lemma 4.4 and Theorem 3.4 \mathcal{A} produces a valid pebbling. Now we note that trivially $M = M(G, S, g, d) \leq n$. We also have $C = C(G, S, g, d) \leq \left\lceil \frac{n}{(\lambda+1)t} \right\rceil + 1$ because the sets returned by need have size at most $|L_k \cap \{it + j \mid i \leq n/t\}| \leq \left\lceil \frac{n}{(\lambda+1)t} \right\rceil$. Plugging these bounds into Theorem 3.5 we get

$$\begin{aligned} E_{\bar{R}}(\mathcal{A}) &\leq n \left[\frac{d(M + \bar{R}C)}{g} + \delta g + e + \bar{R} + 1 \right] \\ &\leq n \left[\frac{(\lambda + 1)t \left(n + \bar{R} \left\lceil \frac{n}{(\lambda+1)t} \right\rceil + \bar{R} \right)}{g} + \delta g + \frac{n}{t} + \bar{R} + 1 \right]. \end{aligned}$$

In particular we can set $t = n^{1/3}$ and $g = n^{2/3}$ to obtain the final result.

$$\begin{aligned}
E_{\bar{R}}(\mathcal{A}) &\leq n \left[\frac{(\lambda + 1)t \left(n + \bar{R} \left\lceil \frac{n}{(\lambda + 1)t} \right\rceil + \bar{R} \right)}{g} + \delta g + \frac{n}{t} + \bar{R} + 1 \right] \\
&\leq n \left[\frac{(\lambda + 1)t \left(n + \bar{R} \frac{n}{(\lambda + 1)t} + 2\bar{R} \right)}{n^{2/3}} + n^{2/3}(\delta + 1) + \bar{R} + 1 \right] \\
&\leq n \left[(\lambda + 1) \left(n^{2/3} + \frac{2\bar{R}}{n^{1/3}} \right) + \bar{R}n^{1/3} + (\delta + 1)n^{2/3} + \bar{R} + 1 \right] \\
&\leq n^{5/3} \left[(\lambda + 1) + (\delta + 1) + \frac{2(\lambda + 1)\bar{R}}{n} + \frac{\bar{R}}{n^{1/3}} + \frac{\bar{R} + 1}{n^{2/3}} \right] \\
&\leq n^{5/3} \left[(\lambda + 1) + (\delta + 1) + \frac{3\bar{R} + 1}{n^{2/3}} + \frac{\bar{R}}{n^{1/3}} \right].
\end{aligned}$$

Plugging the bounds on $\mathbf{C}(G, S, g, d)$ and $\mathbf{M}(G, S, g, d)$ into Theorem 3.5 for $\text{PGenPeb}(G, s, g, d, k = g/d)$ with $g = n^{2/3}$, $d = (\lambda + 1)n^{1/3}$ and $k = g/d = n^{1/3}/(\lambda + 1)$ we get that

$$\begin{aligned}
AT_R(\mathcal{A}) &\leq \frac{2n}{k} [ke + k\delta g + \mathbf{M}(G, S, g, d) + R \cdot \mathbf{C}(G, S, g, d) + Rk] \\
&\leq \frac{2n(\lambda + 1)}{n^{1/3}} \left[\frac{nk}{t} + \delta kn^{2/3} + n + R \cdot \mathbf{C}(G, S, g, d) + Rk \right] \\
&\leq \frac{2n(\lambda + 1)}{n^{1/3}} \left[\frac{1 + \delta}{\lambda + 1} n + n + R \cdot \frac{n^{2/3}}{(\lambda + 1)} + R(k + 2) \right] \\
&\leq 2n^{5/3}(\lambda + 1) \left[\frac{1 + \delta}{\lambda + 1} + 1 + R \cdot \frac{n^{-1/3}}{(\lambda + 1)} + \frac{2R + kR}{n} \right] \\
&\leq 2n^{5/3} \left[(1 + \delta) + (\lambda + 1) + \frac{R}{n^{1/3}} + \frac{2R + (\lambda + 1)R}{n^{2/3}} \right]
\end{aligned}$$

□

The following Theorem 4.6 for *strict* λ -stacked sandwich DAGs improves on the attack in Theorem 4.5 which applies to the more general class of λ -stacked sandwich DAGs. The only difference is that we use the keep function from , which was proved to be valid in Lemma 4.3, instead of the constant function $\text{keep}(\cdot) = V$ that returns all vertices.

Theorem 4.6 (Complexity of Strict Sandwich Graph) *Let G be a strict λ -stacked sandwich DAG on n nodes with $\lambda < n^{1/3}$ and maximum indegree $\text{indeg}(G) = \delta$ then for any core-memory area and energy ratios R and \bar{R} there exists an evaluation algorithm \mathcal{A} for the corresponding iMHF with*

$$\begin{aligned}
AT_R(\mathcal{A}) &\leq 2n^{5/3} \times \left[3 + \delta + \frac{R}{n^{1/3}} + \frac{3R}{n^{2/3}} \right], \text{ and} \\
E_{\bar{R}}(\mathcal{A}) &\leq n^{5/3} \times \left[3 + \delta + \frac{\bar{R}}{n^{1/3}} + \frac{\bar{R} + 1}{n^{2/3}} \right].
\end{aligned}$$

Theorem 4.6 follows easily from Theorem 3.5, Lemma 4.3 and Lemma 4.2 by setting $t = n^{1/3}$ and $g = n^{2/3}$.

Proof of Theorem 4.6. First, assume that $\lambda < n^{1/3}$. Let $S = \{it \mid i \leq n/t\}$ denote the set S from the proof of Lemma 4.2 with size n/t and $\text{depth}(G - S) \leq \lambda t + t - \lambda - 1$. \mathcal{A} runs $\text{PGenPeb}(G, S, g, d = \lambda t + t)$ with the

subroutine need from Algorithm 3 and the subroutine `keep` from Algorithm 3. By Lemma 4.3 \mathcal{A} produces a valid pebbling. Now we note that $M := M(G, S, g, d) \leq \frac{2n}{\lambda+1}$ because we only need to keep pebbles on at most 2 layers in G at any given time (apart from pebbles on the set S and on the parents of the next g nodes, but these pebbles do not count towards the value M). Plugging these bounds into Theorem 3.5 we get

$$\begin{aligned} E_{\bar{R}}(\mathcal{A}) &\leq n \left[\frac{dM + \min\{dC, n\}\bar{R}}{g} + \delta g + e + \bar{R} + 1 \right] \\ &\leq n \left[\frac{(\lambda+1)t \frac{2n}{\lambda+1} + n\bar{R}}{g} + \delta g + e + \bar{R} + 1 \right] \\ &\leq n \left[\frac{2nt + n\bar{R}}{g} + \delta g + e + \bar{R} + 1 \right]. \end{aligned}$$

In particular we can set $t = n^{1/3}$ and $g = n^{2/3}$ to obtain the final results.

$$\begin{aligned} E_{\bar{R}}(\mathcal{A}) &\leq n \left[\frac{2nt + n\bar{R}}{g} + \delta g + e + \bar{R} + 1 \right] \\ &\leq n \left[2n^{2/3} + n^{1/3}\bar{R} + \delta n^{2/3} + n^{2/3} + \bar{R} + 1 \right] \\ &\leq n^{5/3} \left[3 + \frac{\bar{R}}{n^{1/3}} + \delta + \frac{\bar{R} + 1}{n^{2/3}} \right]. \end{aligned}$$

For the bound on AT_R we observe that $C := C(G, S, g, d) \leq \left\lceil \frac{n}{(\lambda+1)t} \right\rceil + 1$ because the sets returned by need have size at most $|L_k \cap \{it + j \mid i \leq n/t\}| \leq \left\lceil \frac{n}{(\lambda+1)t} \right\rceil$. As before we will use $\text{PGenPeb}(G, S, g, d, k)$ with $g = n^{2/3}$, $d = (\lambda+1)n^{1/3}$ and $k = n^{1/3}/(\lambda+1)$. From Theorem 3.5 we get that

$$\begin{aligned} AT_R(\mathcal{A}) &\leq 2n \left[\frac{d(M + RC)}{g} + \delta g + e + R \right] \\ &\leq 2n \left[\frac{(\lambda+1)(M + RC)}{n^{1/3}} + \delta n^{2/3} + n^{2/3} + R \right] \\ &\leq 2n \left[\frac{(\lambda+1)t \left(\frac{2n}{\lambda+1} + RC \right)}{n^{2/3}} + \delta n^{2/3} + n^{2/3} + R \right] \\ &\leq 2n \left[\frac{2nt + (\lambda+1)t(RC)}{n^{2/3}} + \delta n^{2/3} + n^{2/3} + R \right] \\ &\leq 2n \left[2n^{2/3} + \frac{(\lambda+1)(RC)}{n^{1/3}} + \delta n^{2/3} + n^{2/3} + R \right] \\ &\leq 2n^{5/3} \left[3 + \frac{(\lambda+1)(RC)}{n} + \delta + \frac{R}{n^{2/3}} \right] \\ &\leq 2n^{5/3} \left[3 + \frac{(\lambda+1) \left(R \frac{n}{(\lambda+1)t} + 2R \right)}{n} + \delta + \frac{R}{n^{2/3}} \right] \\ &\leq 2n^{5/3} \left[3 + \frac{(Rn^{2/3} + 2(\lambda+1)R)}{n} + \delta + \frac{R}{n^{2/3}} \right] \\ &\leq 2n^{5/3} \left[3 + \frac{R}{n^{1/3}} + \delta + \frac{3R}{n^{2/3}} \right]. \end{aligned}$$

□

While Theorem 4.5 and Theorem 4.6 only hold for $\lambda < n^{1/3}$ there is a trivial algorithm with complexity $O(n^{5/3})$ when $\lambda > n^{1/3}$. See Theorem A.5 in Appendix A.

We remark that for special cases (e.g., $\lambda = 1$) an alternative implementation of `GenPeb` yields a better upper bound on total memory. In particular, we will use `need` from Algorithm 3 and we will redefine `keep` to simply return the exact same set as `need` in each round of `GenPeb` (so that we don't immediately throw out pebbles in step 14 of the same round). It is easy to show that the pair `keep` and `need` is valid whenever we run `GenPeb` on $\lambda = 1$ -stacked sandwich DAGs (see Lemma A.3 in Appendix A). This observation allows us to immediately re-derive a result of [AS15] stating that for any $\lambda = 1$ -sandwich DAG G with maximum `indeg` = 2 there is an algorithm \mathcal{A} with $\text{cmc}(\mathcal{A}) = O(n^{1.5})$. In fact, Corollary A.4 in Appendix A generalizes the result of [AS15]. We do not require that G has maximum `indeg` = 2. Furthermore, \mathcal{A} also has $\text{amc}(\mathcal{A}) = O(\sqrt{n})$, $\text{arc}(\mathcal{A}) = \sqrt{n}$ and $\text{crc}(\mathcal{A}) = O(n^{1.5})$.

5 Random Graphs

In this section we demonstrate how to extend our attacks to two recent iMHF proposals. The first is the Argon2i iMHF — the variant of Argon2 in which data access patterns are data independent [BDK15]. The authors recommended using Argon2i for password hashing applications to avoid potential side-channel leakage through data dependent memory access patterns [BDK15]. The basic Argon2i DAG is a (pseudo) randomly generated DAG with maximum indegree `indeg` = 2. Thus, we view the Argon2i DAG as a distribution over n node DAGs — see Definition 5.1. The second iMHF considered is the Single-Buffer (SB) construction of [CGBS16] (we considered the Double-Buffer and Linear iMHFs from [CGBS16] in Section 4).

To ease exposition we have opted to analyze the single thread cases of Argon2i and the first construction of [CGBS16]. However (as we briefly describe below) extending these results to multi-thread variants and the remaining two constructions of [CGBS16] is straightforward.

We begin with the following definition which fixes a class of random graphs key to our analysis.

Definition 5.1 ((n, δ, w) -random DAG) *Let $n \in \mathbb{N}$, $1 < \delta < n$, and $1 \leq w \leq n$ such that w divides n . An (n, δ, w) -random DAG is a randomly generated directed acyclic (multi)graph with n nodes v_1, \dots, v_n and with maximum in-degree δ for each vertex. The graph has directed edges (v_i, v_{i+1}) for $1 \leq i < n$ and random forward edges $(v_{r(i,1)}, v_i), \dots, (v_{r(i,\delta-1)}, v_i)$ for each vertex v_i . Here, $r(i, j)$ is independently chosen uniformly at random from the set $[\max\{0, i - w\}, i - 1]$.*

We observe that a t -pass instance of Argon2i iMHF²² is an $(n, 2, n/t)$ -random DAG. Similarly the r -pass “Single-Buffer” construction of [CGBS16] is based on an $(n, 21, n/r)$ -random DAG.

We cannot directly apply our results from the previous section because a (n, δ, w) -random DAG will not be λ -stacked sandwich DAG with high probability. However, we can show that these graphs are ‘close’ to λ -stacked sandwich DAGs in the sense that, with high probability, there is a ‘small’ set S such that we can turn G into a λ -stacked sandwich DAGs just by removing edges incident to vertices in S . Definition 5.2 formalizes this intuition.

Definition 5.2 ((m, λ) -Layered Graph) *Let G be a DAG with $n = k\lambda$ nodes v_1, \dots, v_n with directed edges (v_i, v_{i+1}) for $1 \leq i < n$. Given a set S of nodes we use G_S to denote the resulting DAG if we removed all directed edges (v, v') that are incident to nodes in S (e.g., $v \in S$ or $v' \in S$) except for edges of the form (v_i, v_{i+1}) . We say that G is a (m, λ) -layered DAG if we can find a set S with at most m nodes such that G_S is a λ -stacked sandwich DAG.*

Demonstrating that a graph is (m, λ) -layered is useful because Theorem 5.4 upper-bounds the AT and energy complexities of an iMHF based on such a graph. In particular, Theorem 5.4 relies on Lemma 5.3

²²In the notation of [BDK15] the case when $t = 1$ corresponds to a single pass.

which states that any layered graph is also depth reducible. In Lemma 5.5, for any given probability $\gamma > 0$ we upper-bound the size m when viewing an Argon2i graph as an (m, λ) -layered graph. Thus we get Corollary 5.6 which describes an evaluation algorithm for Argon2i and the Balloon Hashing algorithms and bound their AT and energy complexities. In particular it states both their expected values and upper-bounds holding with a given probability γ .

Lemma 5.3 (Layered Graphs are Reducible) *Let G be a (m, λ) -layered DAG then for any integer $t \geq 1$ G is $(n/t + m, \lambda t + t - \lambda - 1)$ -reducible.*

PROOF. By definition there is a set S_1 of m nodes such that G_{S_1} is a λ -stacked sandwich DAG. Now by Lemma 4.2 we can find a set $S_2 \subseteq V(G)$ of size n/t such that $G_{S_1} - S_2$ has depth at most $\text{depth}(G_{S_1} - S_2) \leq \lambda t + t - \lambda - 1$. Now we set $S = S_1 \cup S_2$ we have $|S| \leq m + n/t$ and $\text{depth}(G - S) \leq \text{depth}(G_{S_1} - S_2) \leq \lambda t - \lambda - 1$. \square

The following theorem upper-bounds the complexity of a layered graph.

Theorem 5.4 (Complexity of Layered Graph) *Let G be a (m, λ) -layered DAG on n nodes with $\lambda < n^{1/3}$ and maximum indegree $\text{indeg}(G) = \delta$ and fix any $t > 0$, $g \geq t(\lambda + 1)$ then there exists an attack \mathcal{A} on the corresponding iMHF such that for any core-memory ratio $R > 0$ and $\bar{R} > 0$ the energy and AT complexities of \mathcal{A} are at most*

$$\begin{aligned} AT_R(\mathcal{A}) &\leq 2n \left[\frac{n}{t} + m + \delta g + R + \frac{(\lambda + 1)t \cdot (n + 2R) + R \cdot n}{g} \right] \\ E_{\bar{R}}(\mathcal{A}) &\leq n \left[\frac{(\lambda + 1)t(n + 2\bar{R}) + \bar{R} \cdot n}{g} + \delta g + \frac{n}{t} + m + \bar{R} \right]. \end{aligned}$$

Theorem 5.4 follows from Lemma 4.4, Lemma 5.3 and Theorem 3.4.

Proof of Theorem 5.4. First, assume that $\lambda < n^{1/3}$. Let $S = S_1 \cup S_2$ denote the set S from the proof of Lemma 5.3 with size $|S_1| = |\{it \mid i \leq n/t\}| = n/t$, $|S_2| = m$ and $\text{depth}(G - S) \leq \lambda t + t - \lambda - 1$. \mathcal{A} runs $\text{GenPeb}(G, S, g, d = \lambda t + t)$ with the function need from Algorithm 3 and the constant function $\text{keep}(x, y) = V$. By Lemma 4.4 \mathcal{A} produces a valid pebbling as G_{S_2} is a λ -stacked sandwich graph. Now we note that trivially $M(G, S, g, d) \leq n$. We also have $C(G, S, g, d) \leq \left\lceil \frac{n}{(\lambda + 1)t} \right\rceil + 1$ because the sets returned by need have size at most $|L_k \cap \{it + j \mid i \leq n/t\}| \leq \left\lceil \frac{n}{(\lambda + 1)t} \right\rceil$. Plugging these bounds into Theorem 3.4 we get $E_{\bar{R}}(\mathcal{A}) \leq$

$$\begin{aligned} &\leq n \left(\frac{d(M(G, S, g, d) + \bar{R} \cdot C(G, S, g, d))}{g} + \delta g + e + \bar{R} \right) \\ &\leq n \left(\frac{(\lambda + 1)t(n + 2\bar{R}) + \bar{R} \cdot n}{g} + \delta g + \frac{n}{t} + m + \bar{R} \right) \end{aligned}$$

Plugging the bounds on $C(G, S, g, d)$ and $M(G, S, g, d)$ into Theorem 3.5 for $\text{PGenPeb}(G, s, g, d, k = g/d)$ we get $AT_R(\mathcal{A}) \leq$

$$\begin{aligned} &\leq \frac{2n}{k} \times [ke + k\delta g + M(G, S, g, d) + R \cdot C(G, S, g, d) + Rk] \\ &\leq 2n \times \left[\frac{n}{t} + m + \delta g + \frac{M(G, S, g, d) + R \cdot C(G, S, g, d)}{k} + R \right] \\ &\leq 2n \times \left[\frac{n}{t} + m + \delta g + \frac{dn + d \cdot R \cdot C(G, S, g, d)}{g} + R \right] \\ &\leq 2n \times \left[\frac{n}{t} + m + \delta g + R + \frac{(\lambda + 1)t \cdot (n + 2R) + R \cdot n}{g} \right]. \end{aligned}$$

\square

The following Lemma 5.5, is the key technical result in this section. It states that, in particular, for any $m \leq n$ and any constants k and δ a $(n, \delta, w = n/k)$ -random DAG will be a $(O(m \log n), O(\frac{n}{m}))$ -layered DAG with high probability.

Recall that by H_λ we denote the λ^{th} harmonic number.

Lemma 5.5 ((n, δ, w)-random DAGs are Layered) *Fix any $\lambda \geq 1$ and let $q = \lceil \frac{n}{\lambda+1} \rceil$. Then a $(n, \delta, w = n/k)$ -random DAG is a (m, λ) -layered DAG, where the random variable m has expected value $\mathbb{E}[m] = k(\delta - 1)q + q + \frac{(\delta-1)q \cdot H_\lambda}{2}$. Furthermore, for any $\gamma > e^{-3(\mathbb{E}[m]-q)/4}$ we have $m \leq \mathbb{E}[m] + \sqrt{3(\mathbb{E}[m]-q) \ln \gamma^{-1}}$ except with probability γ .*

To prove the lemma we show how to construct a set S of (expected) size $O(m \log n)$. We note that our construction is computationally efficient. Intuitively, we partition the nodes of G into equal sized layers of consecutive nodes L_0, \dots, L_λ . We add a node $v \in L_i$ to our set S if *any* of v 's parents are also in layer L_i . In the single pass case ($w = n$) we will add a vertex $v \in L_i$ to S with probability at most $(\delta - 1)/i$. Thus, in expectation we will add at most $(\delta - 1) |L_i| / i$ nodes from layer L_i to S . In total we add at most $\frac{n}{\lambda+1} \sum_{i=1}^\lambda \frac{1}{i} = \frac{nH_\lambda}{\lambda+1}$ nodes from layers $L_{\geq 1}$ in expectation.

Proof of Lemma 5.5. Let G be a $(n, \delta, w = n/k)$ -random DAG with nodes v_1, \dots, v_n . For simplicity assume that $\lambda = n/q - 1$ is an integer so that we can divide G into λ layers L_0, \dots, L_λ of equal size q (If λ is not an integer then the first $\lceil \lambda \rceil - 1$ layers will contain $q + 1$ nodes each and the last layer will contain $\leq q$ nodes). Layer L_i contains nodes $L_i = \{v_{iq+1}, \dots, v_{(i+1)q}\}$. We construct $S \subseteq V(G)$ as follows:

$$S = \bigcup_i \{v_j \in L_i \mid \exists s \leq d - 1. v_{r(j,s)} \in L_i\} .$$

That is if any of v 's parents are in the same layer as v we add v to S .

Given $v_j \in L_i$ we let x_j denote the indicator random variable that is 1 if and only if $v_{r(j,t)} \in L_i$ for some index $t \leq \delta - 1$. We note that by linearity of expectation we have

$$\mathbb{E}[|S|] \leq \sum_{t=0}^{\lambda} \sum_{j=1}^q \mathbb{E}[x_{iq+j}] .$$

Let $\lambda' = \lfloor \frac{w}{q} \rfloor - 1$. Suppose first that $i \leq \lambda'$ and $0 < j \leq q$ so that $iq+j \leq w$ and $v_{iq+j} \in L_i$ ($i > 0$) then the probability that we add v_{iq+j} to S because one of its parents is also in layer L_i is at most $(\delta - 1)(j - 1)/(iq)$. Thus, $\mathbb{E}[x_{iq+j}] \leq (\delta - 1)/i$ for each $v_j \in L_i$. Now suppose that $iq + j > w$ then the probability that we add $v_{r(iq+j)}$ to S because one of its parents is in the same layer is at most $(\delta - 1)(j - 1)/(w) = (\delta - 1)(j - 1)k/n$. Thus, $\mathbb{E}[x_{iq+j}] \leq (\delta - 1)(q - 1)k/n$. Thus, in expectation we have

$$\begin{aligned} \mathbb{E}[|S|] &\leq q + (\delta - 1) \sum_{i=1}^{\lambda} \sum_{j=1}^q \frac{j-1}{im} \\ &\leq q + (\delta - 1) \sum_{i=1}^{\lambda'} \frac{(q-1)}{2i} + q(\delta - 1) \sum_{i=\lambda'+1}^{\lambda} \frac{(q-1)k}{n} \\ &\leq q + (\delta - 1) \sum_{i=1}^{\lambda'} \frac{(q-1)}{2i} + q(\delta - 1)(\lambda - \lambda') \frac{(q-1)k}{n} \\ &\leq q + (\delta - 1) \frac{(q-1)H_{\lambda'}}{2} + q(\delta - 1)k , \end{aligned}$$

where $H_\lambda \approx \ln \lambda$ denotes the λ^{th} harmonic number and the last inequality follows because $(q - 1)\lambda \leq n$. Observe that the DAG G_S with all directed edges originating in S deleted (i.e., delete edges of the form (v, v')

with $v \in S$) will be a λ -stacked sandwich graph because there are no forward edges within each layer apart from the chain (v_i, v_{i+1}) . Let $X = \sum_{j=q+1}^n x_j$. Because the random variables $x_j \in \{0, 1\}$ are independent standard concentration bounds imply that $\Pr[X \geq \mathbb{E}[X] + \tau] \leq \exp\left(\frac{-\tau^2}{2\text{Var}(X) + 2\tau/3}\right)$. In our case we have

$$\begin{aligned} \text{Var}(X) &= \sum_{i=q+1}^n \sum_{j=q+1}^n \mathbb{E}[x_i x_j] - \mathbb{E}[x_i]\mathbb{E}[x_j] = \sum_{i=q+1}^n \mathbb{E}[x_i] - \mathbb{E}[x_i]^2 \\ &\leq \mathbb{E}[X] - (\delta - 1) \frac{(q-1)H_{\lambda'}}{2} + q(\delta - 1)k. \end{aligned}$$

Thus, for any $\gamma > \exp(-3\mathbb{E}[X]/4)$ we can set $\tau = \sqrt{3\mathbb{E}[X] \ln \gamma^{-1}}$ to obtain

$$\Pr[X \geq \mathbb{E}[X] + \tau] \leq \exp\left(\frac{-\tau^2}{\mathbb{E}[X] + 2\tau/3}\right) \leq \exp\left(\frac{3\mathbb{E}[X] \ln \gamma}{3\mathbb{E}[X]}\right) \leq \gamma.$$

Where the second inequality follows from the observation that $2\tau/3 < \mathbb{E}[X]$ whenever $\gamma > \exp(-3\mathbb{E}[X]/4)$. As $|S| = X + q$ we have

$$|S| \leq q + (\delta - 1) \frac{(q-1)H_{\lambda'}}{2} + q(\delta - 1)k + \sqrt{3 \left((\delta - 1) \frac{(q-1)H_{\lambda'}}{2} + q(\delta - 1)k \right) \ln \gamma^{-1}}$$

except with probability γ^{23} . □

As an immediate consequence of Lemma 5.5 we get an attack on the static mode of operation of the Password Hashing Competition winner Argon2. Specifically we can attack the Argon2i variant in which memory accesses patterns are not input dependent — a desirable property to prevent side-channel attacks based on cache timing. As another immediate consequence we also get an attack on the Single-Buffer construction of [CGBS16].

Corollary 5.6 (Complexities of (n, δ, w) -random DAGs) *Let G be a $(n, \delta, n/k)$ -random DAG on n nodes. Then there exists an evaluation algorithm \mathcal{A} for the corresponding iMHF such that for any core-memory ratios $R > 0$ and $\bar{R} > 0$ the expected AT and energy complexities of \mathcal{A} are at most*

$$\begin{aligned} \mathbb{E}[AT_R(\mathcal{A})] &\leq 2n^{7/4} \left[3 + \delta + \frac{(\delta - 1)(H_{n^{1/4}} + 2k)}{2} + \frac{R}{n^{3/4}} + \frac{R}{\sqrt{n}} + \frac{2R}{n} \right] \\ \mathbb{E}[E_{\bar{R}}(\mathcal{A})] &\leq n^{7/4} \left[3 + \delta + \frac{(\delta - 1)(H_{n^{1/4}} + 2k)}{2} + \frac{2\bar{R}}{n} + \frac{2\bar{R}}{\sqrt{n}} \right]. \end{aligned}$$

Furthermore, except with probability $\gamma > e^{-n^{3/4}(\delta-1)}$, we have

$$\begin{aligned} AT_R(\mathcal{A}) &\leq 2n^{7/4} \left[3 + \delta + \left(\frac{(\delta - 1)(H_{n^{1/4}} + 2k)}{2} + \sqrt{\frac{3(\delta - 1)(H_{n^{1/4}} + 2k) \ln \gamma^{-1}}{2n^{3/4}}} \right) + \frac{R}{n^{3/4}} + \frac{R}{\sqrt{n}} + \frac{2R}{n} \right] \\ E_{\bar{R}}(\mathcal{A}) &\leq n^{7/4} \left[3 + \delta + \left(\frac{(\delta - 1)(H_{n^{1/4}} + 2k)}{2} + \sqrt{\frac{3(\delta - 1)(H_{n^{1/4}} + 2k) \ln \gamma^{-1}}{2n^{3/4}}} \right) + \frac{2\bar{R}}{n} + \frac{2\bar{R}}{\sqrt{n}} \right]. \end{aligned}$$

²³If $\gamma < \exp(-3\mathbb{E}[X]/4)$ then we can set $\tau = \sqrt{3\mathbb{E}[X] \ln \gamma^{-1}}$ to obtain a slightly weaker concentration bound. However, the term $\exp(-3\mathbb{E}[X]/4)$ is already negligibly small in all of our applications.

In a nutshell Corollary 5.6 is shown by setting $q = n^{3/4}$ in Lemma 5.5 and $\lambda = n^{1/4} - 1$, $g = n^{3/4}$ and $t = n^{1/4}$ in Theorem 5.4.

Proof of Corollary 5.6. Set $\lambda + 1 = n^{1/4}$, $q = g = n^{3/4}$ and $t = n^{1/4}$. By Lemma 5.5 G is a (m, λ) -layered DAG with $\mathbb{E}[m] \leq n^{3/4} + (\delta - 1)n^{3/4}H_\lambda/2 + k(\delta - 1)n^{3/4}$. Thus, by Theorem 5.4 there exists algorithm \mathcal{A} with

$$\begin{aligned} \mathbb{E}[AT_R(\mathcal{A})] &\leq 2n \left[\frac{n}{t} + \mathbb{E}[m] + \delta g + R + \frac{(\lambda + 1)t \cdot (n + 2R) + R \cdot n}{g} \right] \\ &\leq 2n^{7/4} \left[3 + \delta + (\delta - 1) \left(\frac{H_\lambda}{2} + k \right) + \frac{R}{n^{3/4}} + \frac{R}{\sqrt{n}} + \frac{2R}{n} \right]. \end{aligned}$$

Similarly, substituting in terms for the expected energy complexity we get that

$$\begin{aligned} \mathbb{E}[E_{\bar{R}}(\mathcal{A})] &\leq n \left[\frac{(\lambda + 1)t(n + 2\bar{R}) + \bar{R} \cdot n}{g} + \delta g + \frac{n}{t} + \mathbb{E}[m] + \bar{R} \right] \\ &\leq n \left[(3 + \delta + (\delta - 1)(k + H_\lambda/2))n^{3/4} + 2n^{-1/4}\bar{R} + \bar{R} \cdot n^{1/4} + \bar{R} \right] \\ &\leq n^{7/4} \left[3 + \delta + \frac{(\delta - 1)(H_\lambda + 2k)}{2} + \frac{2\bar{R}}{n} + \frac{2\bar{R}}{\sqrt{n}} \right]. \end{aligned}$$

By Lemma 5.5 for any $\gamma > e^{-3\mathbb{E}[m]/4}$ we have $m \leq \mathbb{E}[m] + \sqrt{3\mathbb{E}[m] \ln \gamma^{-1}}$ except with probability γ . Thus the remainder of the theorem follows from Theorem 5.4 by substituting in the definitions of λ , t and g and the upper-bound on m . \square

6 Ideal iMHFs Don't Exist

In this section we show that ideal iMHFs do not exist. More specifically we show that for every graph G there exists node set S and positive integer $g \geq \text{depth}(G - S)$ such that the iMHF evaluation algorithm $\mathcal{A} = \text{PGenPeb}(G, S, g, d, \lfloor g/d \rfloor)$ has AT and energy-complexity $o(n^2 / \log^{1-\epsilon} n)$ for any constant $\epsilon > 0$. In particular, if we take the naïve algorithm to be $\mathcal{N} = \text{triv}$ then \mathcal{A} is an attack with energy-quality $\omega(\log^{1-\epsilon} n)$. We first prove (Lemma 6.2) that all DAGs are reducible provided that the maximum indegree δ is sufficiently small (e.g., $\delta \leq \log^{0.999} n$). The proof of Lemma 6.2 follows from a result of Valiant [Val77]. Once we have established that all DAGs are reducible we can use PGenPeb to obtain a high quality attack on any iMHF.

Lemma 6.1 ([Val77] Extension) *Given a DAG G with m edges and depth $\text{depth}(G) \leq d = 2^i$ there is a set of m/i edges s.t. by deleting them we obtain a graph of depth at most $d/2$.*

The statement of Lemma 6.1 is slightly different from the statement in [Val77]. Thus, while the proof of Lemma 6.1 follows the argument of [Val77] exactly, we include it here for completeness.

Proof of L. et G be given. A proper labeling of G is a labeling of the nodes with nonnegative numbers such that if (u, v) is an edge, then the label of node u is smaller than the label of node v . Because G has depth at most $d = 2^i$ it is easy to find a proper labeling of G using labels from $[0, d - 1]$. Now we can partition the edges of G into color classes E_1, \dots, E_i s.t. an edge $(u, v) \in E_j$ if and only if the most significant bit where the labels of u and v differ in their binary representation is the j 'th from the left. Because sets E_1, \dots, E_i partition the edges of G so we can always find some index j^* for which $|E_{j^*}| \leq m/i$. Now we claim that $\text{depth}(G - E_{j^*}) \leq d/2$. Consider the labeling of $G - E_{j^*}$ obtained by deleting bit j^* . The resulting labeling using labels from $[0, d/2 - 1]$ and it is easy to verify that this is a proper labeling of $G - E_{j^*}$. Suppose for contradiction that there is a directed edge (u, v) in $G - E_{j^*}$ such that the label of node u is not smaller than the label of node v . Consider the labels for nodes u and v in the proper labeling of G . Because $(u, v) \notin E_{j^*}$ the labels for u and v in the proper labeling of G must either have had the same value in bit j^* or there

must have been some other more significant bit $j < j^*$ in which the labels differed. In either case we arrive at an immediate contradiction: in our original labeling of G the label of u was not smaller than the label of v so this labeling was not proper (Contradiction!). Thus, $G - E_{j^*}$ can have depth at most $d/2$. \square

Lemma 6.2 (All DAGs are Reducible) *Let $G = (V, E)$ be an arbitrary DAG of size $|V| = n = 2^k$ with $\text{indeg}(G) = \delta$. Then for every integer $t \geq 1$ there is a set $S \subseteq V$ of size $|S| \leq \frac{t\delta n}{\log(n)-t}$ such that $\text{depth}(G - S) \leq 2^{k-t}$. Furthermore, there is an efficient algorithm to find S .*

Lemma 6.2 follows by repeatedly invoking Lemma 6.1.

Proof of Lemma 6.2. Let $G_0 = G$ and let $d_0 = 2^k \geq \text{depth}(G_0)$. To construct S we will start with $S = \emptyset$ and apply Lemma 6.1 t times. In particular, let E_0 denote the set of at most m/k edges such that $G_0 - E_0$ has depth at most 2^{k-1} . We note that the set E_0 can be computed in polynomial time. Now let S_0 denote the set of origin vertices in E_0 and observe that $|S_0| \leq \delta|E_0| \leq \delta n / \log d_0$. Let $G_1 = G_0 - S_0$ and let $d_1 = d_0/2$. In general, let $G_{i+1} = G_i - S_i$, where S_i denotes the origin vertices in E_i — the set from applying Lemma 6.1 to G_i . In particular, note that $d_{i+1} = d_i/2 \geq \text{depth}(G_{i+1})$ so after t applications of Lemma 6.1 we have $\text{depth}(G_t) \leq d_t = 2^{k-t}$. Now let $S = S_0 \cup \dots \cup S_{t-1}$ then in the i^{th} application of Lemma 6.1 the size of S grows by at most

$$\frac{m_i}{\log(d_i)} \leq \frac{\delta n}{\log(n/2^i)} = \frac{\delta n}{\log(n) - i}.$$

the final size of S is

$$|S| \leq \sum_{i \in [t]} \frac{\delta n}{\log(n) - i} \leq \frac{t\delta n}{\log(n) - t}.$$

\square

Theorem 6.3 now follows from Lemma 6.2 and Theorem 3.5.

Theorem 6.3 (Complexities of any iMHF) *Let F be an iMHF based on arbitrary DAG $G = (V, E)$ of size $|V| = n$ with in-degree $\text{indeg}(G) = \delta$. Then for every constant $\epsilon > 0$ and fixed ratios $R > 0$ and $\bar{R} > 0$ there exists an evaluation algorithm \mathcal{A} such that*

$$E_{\bar{R}}(\mathcal{A}) = o\left(n \left(\frac{\delta n}{\log^{1-\epsilon}} + \hat{R} \right)\right) = AT_R(\mathcal{A})$$

where $\hat{R} = \max\{R, \bar{R}\}$.

In particular if we let the naïve algorithm for F be $\mathcal{N} = \text{triv}$ then algorithm \mathcal{A} is an attack with qualities

$$\text{E-quality}(\mathcal{A}) = \Omega\left(\frac{\hat{R} + n}{\bar{R} + \delta n / \log^{1-\epsilon} n}\right) = \text{AT-quality}(\mathcal{A})$$

or, for constant R and \bar{R} , simply $\Omega(\delta^{-1} \log^{1-\epsilon} n)$.

In a nutshell, in the proof we set $t = O(\log \log n)$ in Lemma 6.2 and then we set $g = n / \log^{1+\epsilon}(n)$ in Theorem 3.5. The details follow.

Proof of Theorem 6.3. Let constant $\hat{R} = \max(R, \bar{R})$ a $\delta = \text{indeg}(G)$. In Lemma 6.2 set t to be

$$t = \left\lceil \log \left(\frac{n \log^2 n}{n - \hat{R} \log^2 n} \right) \right\rceil.$$

Thus for any constant $\hat{\epsilon} > 0$ we obtain a set S of size

$$|S| \leq \frac{t\delta n}{\log(n) - t} = o\left(\frac{\delta n}{\log^{1+\hat{\epsilon}} n}\right).$$

Moreover, the resulting depth d of $G - S$ is

$$d \leq n/2^t = n \left(\frac{n \log^2 n}{n - \hat{R} \log^2 n} \right)^{-1} = \frac{n - \hat{R} \log^2 n}{\log^2 n}.$$

Now we choose the parameter g in $\text{PGenPeb}(G, S, g, d, g/d)$ and compute the resulting power-complexity and quality of \mathcal{A} .

We now fix $g = n/\log^{1+\epsilon}(n)$ where we require that $1 > \hat{\epsilon} > \epsilon$. We observe that $g \in [d, n]$ as required. Trivially, during any given round of a Balloon phase there can be no more than $M(G, S, g) \leq n$ pebbles on G , nor are more than $C(G, M, g) \leq n$ pebbles ever placed on G during a single balloon phase. Thus we can use Theorem 3.5 to bound

$$\begin{aligned} E_{\bar{R}}(\mathcal{A}) &\leq n \left[\frac{n(\bar{R} + d)}{g} + \delta g + |S| + \bar{R} + 1 \right] \\ &= n \left[(\hat{R} + d) \log^{1+\epsilon}(n) + \delta g + |S| + \bar{R} + 1 \right] \\ &\leq n \left[\hat{R} \log^{1+\epsilon}(n) + \frac{n - \hat{R} \log^2 n}{\log^{1-\epsilon} n} + \delta g + |S| + \bar{R} + 1 \right] \\ &= \frac{n^2}{\log^{1-\epsilon} n} + \frac{\delta n^2}{\log^{1+\epsilon} n} + o\left(\frac{\delta n^2}{\log^{1+\hat{\epsilon}} n}\right) + n\bar{R} + n \\ &= o\left(n \left(\frac{\delta n}{\log^{1-\epsilon} n} + \bar{R} \right)\right) \end{aligned}$$

for any constants \bar{R} and \hat{R} . A similar calculation for AT-complexity implies the theorem.

$$\begin{aligned} AT_{\bar{R}}(\mathcal{A}) &\leq 2n \left[\frac{dn(R + 1)}{g} + \delta g + |S| + R \right] \\ &= 2n \left[d(R + 1) \log^{1+\epsilon}(n) + \delta g + |S| + R \right] \\ &\leq 2n \left[\frac{(R + 1)(n - \hat{R} \log^2 n)}{\log^{1-\epsilon} n} + \delta g + |S| + R \right] \\ &\leq O\left(\frac{n^2}{\log^{1-\epsilon} n}\right) + \frac{\delta n^2}{\log^{1+\epsilon} n} + o\left(\frac{\delta n^2}{\log^{1+\hat{\epsilon}} n}\right) + 2nR \\ &= o\left(n \left(\frac{\delta n}{\log^{1-\epsilon} n} + R \right)\right) \end{aligned}$$

To see why this implies that $\text{E-quality}(\mathcal{A}) = \Omega(\delta^{-1} \log^\epsilon n) = \text{AT-quality}(\mathcal{A})$ note that the the naïve algorithm triv has $E_{\bar{R}}(\mathcal{N}) = n(n + \bar{R})$ and $AT_{\bar{R}}(\mathcal{N}) = n(n + \bar{R})$. \square

Remarks. We remark that for any constants δ and $\epsilon < 1$ Theorem 6.3 yields an attack with quality $\text{E-quality}(\mathcal{A}) = \Omega(\log^\epsilon n) = \text{AT-quality}(\mathcal{A})$. Furthermore, provided that $\delta \leq \log^{\epsilon'} n$ where $\epsilon' < 1$ Theorem 6.3 yields an attack with quality $\text{E-quality}(\mathcal{A}) = \omega(1) = \text{AT-quality}(\mathcal{A})$. If $\bar{R} \neq R$ then we can obtain separate attacks \mathcal{A}_1 and \mathcal{A}_2 in which attack \mathcal{A}_1 has optimal energy quality

$$\text{E-quality}(\mathcal{A}_1) = \Omega\left(\frac{\bar{R} + n}{\bar{R} + \delta^2 n / \log^\epsilon n}\right)$$

and attack \mathcal{A}_2 has optimal AT quality

$$\text{AT-quality}(\mathcal{A}) = \Omega\left(\frac{R + n}{R + \delta^2 n / \log^\epsilon n}\right).$$

Erdoes et al. [EGS75] constructed a graph G of any size n with $\text{indeg}(G) = O(\log(n))$ which is (α, β) -fractionally node robust for some constants $0 < \beta < \alpha < 1$.²⁴ This would imply that our bounds in Theorem 6.3 and Lemma 6.2 are essentially tight. Alwen and Serbinenko [AS15] used the depth robust DAGs from [MMV13] as a building block to construct a family of DAGs with provably high pebbling complexity $\tilde{\Omega}(n^2)$. Thus, our general attack in Theorem 6.3 is optimal up to polylogarithmic factors.

7 Practical Considerations

In this section and Appendix B we investigate the exact (as opposed to asymptotic) complexity of the PGenPeb for a variety of iMHFs and parameter settings.

To get an overview of when PGenPeb is in fact an attack we include Table 1 which shows the smallest power of two value of n for which we can obtain attacks with quality > 1 for the case when $R = \bar{R} = 3,000$ arbitrary space-cost σ and time-cost τ subject to $n = \sigma * \tau$.²⁵ For example, we have an attack on Argon2i, the winner of the Password Hashing Competition, when n is only 2^{18} .

iMHF	Equality(\mathcal{A}) > 1	ATquality(\mathcal{A}) > 1
Catena-BRG ₁ ⁿ	2^{15}	2^{16}
Catena-BRG ₃ ⁿ	2^{18}	2^{19}
Catena-BRG ₅ ⁿ	2^{20}	2^{21}
Catena-DBG ₁ ⁿ	2^{25}	2^{26}
Catena-DBG ₃ ⁿ	2^{29}	2^{30}
Catena-DBG ₅ ⁿ	2^{31}	2^{32}
Argon2i _{σ, τ} : (Expected)	2^{18}	2^{18}
Argon2i _{σ, τ} : $\gamma = 0.1$	2^{18}	2^{18}
Argon2i _{σ, τ} : $\gamma = 0.01$	2^{18}	2^{18}
SB _{σ, τ} : $\gamma = 0.01$	2^{25}	2^{25}
Ideal iMHF: $\text{indeg}(G) = 2$	2^{55}	2^{55}
Ideal iMHF: $\text{indeg}(G) = 3$	2^{85}	2^{85}
Ideal iMHF : $\text{indeg}(G) = 4$	2^{114}	2^{114}

Table 1: Minimum $n = 2^i$ with attack quality > 1 . $R = \bar{R} = 3000$. $n = \sigma * \tau$.

Next, we provide several graphs to show how various parameters affect the complexity of PGenPeb for each iMHF. First, in Figure 1a, for Argon2i and the Single-Buffer (SB) constructions, we show the effect of increasing the time-cost parameter τ while holding all other parameters fixed (in particular for a fixed space-cost σ). We remark that AT-quality decreases as τ grows since the amount of memory needed and number of RO implementations both stay the same. Thus as τ grows area stays constant and time begins to dominate the AT-complexity of both the naïve algorithm and PGenPeb. However since both algorithms already run in minimal time the effect is to decrease the advantage of PGenPeb which has better area (esp. memory) complexity. In other words, the quality is highest for the parameter which makes these iMHFs the most memory-hard.

Next, in Figure 1b we show what effect increasing the indegree δ of the underlying DAG has on the complexity of PGenPeb evaluating an arbitrary iMHF. An interesting observation is that although asymptotically the complexity is subquadratic, the exact security even $\delta > 3$, could still be good enough for practical values of n .

Additionally, we include a few extra plots in Appendix B. Figure 3a and Figure 3b show the effect on PGenPeb’s complexity when increasing λ for Catena Dragonfly (BRG _{λ} ⁿ) and Catena Butterfly (DBG _{λ} ⁿ)

²⁴In [MMV13] the authors give an explicit construction of a DAG which has $\text{indeg}(G) = \log^2 n$ which is depth robust for any α and β arbitrarily close to 1.

²⁵Recall that the complexity of PGenPeb is the same for all values of space-cost σ and time-cost τ where $n = \sigma * \tau$.

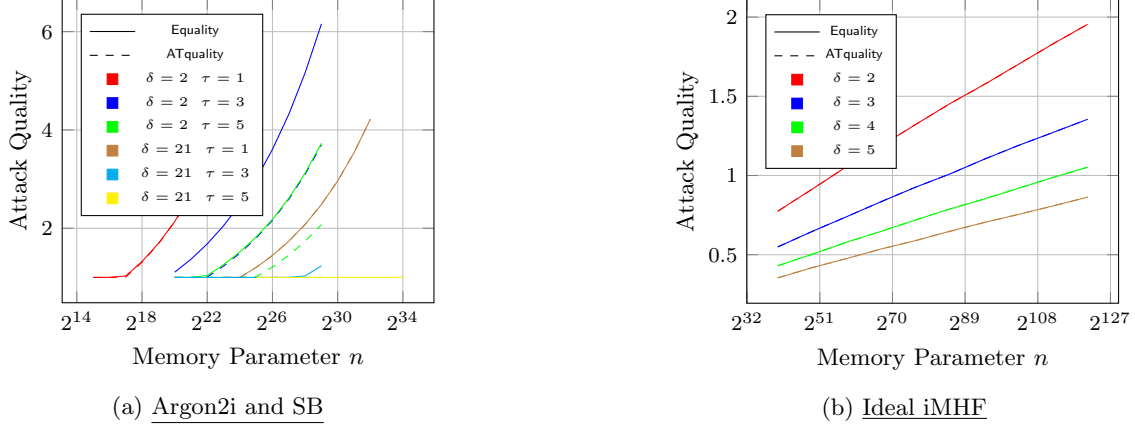


Figure 1: Complexity of PGenPeb for increasing time-cost τ and indegree δ . ($R = \bar{R} = 3000$)

respectively²⁶. Moreover for Argon2i and SB Figure 4a and Figure 4b show the effect (for the case $\tau = 1$ and $R = 3000$) of decreasing the error probability γ which tells us how likely we are to have an attack with at least the plotted quality. Finally Figure 5 shows the effect that different values of the core memory and energy ratios R and \bar{R} have on the complexity of PGenPeb when evaluating Argon2i. Intuitively, the plot shows that attack quality increases as R, \bar{R} decrease, but as we would expect the effect is less significant for larger values of n .

Parameter Optimization. We remark that we optimized the parameters of our attack for each specific value of n in our plots. For example, we showed that any λ -stacked sandwich DAG is $(n/t, t(\lambda+1))$ -reducible for any $t \geq 1$. For each different value of n we ran a script to find the optimal values of t and $g \geq t(\lambda+1)$ which minimize the energy complexity (resp. AT-complexity) of $\text{PGenPeb}(G, S, g, d = t(\lambda+1), k = g/d)$. In our general iMHF attack we used a script to find the optimal value of t in Lemma 6.2 and the optimal value of g .

Naïve Algorithms. The naïve algorithm \mathcal{N} for the Catena Butterfly iMHF has absolute memory complexity $\text{amc}(\mathcal{N}) = n/(\lambda \log n)$ and energy complexity $E_{\bar{R}}(\mathcal{N}) = n(\text{amc}(\mathcal{N}) + \bar{R})$. Similarly, the naïve algorithm \mathcal{N} for Catena Dragonfly has absolute memory complexity $\text{amc}(\mathcal{N}) = n/(\lambda+1)$ and the naïve k -pass algorithm for Argon2i and SB has absolute memory complexity $\text{amc}(\mathcal{N}) = n/k$. Thus, our attack quality decreases with λ or k . We stress that this is not because our attacks becomes less efficient as λ and k increases, but because the \mathcal{N} algorithm requires less and less memory (thus, as λ, k increase the iMHFs become increasingly less ideal). By contrast, the naïve algorithm $\mathcal{N} = \text{triv}$ for our general iMHF (and for Argon2i) has $E_{\bar{R}}(\mathcal{N}) = n(n + \bar{R})$.

Customized Attack Architecture. We have outline efficient attacks on Catena, Argon2i and the Balloon Hashing iMHFs in the theoretical Parallel Random Oracle Machine (pROM) model of computation. Because pROM is a theoretical model of computation it is not obvious a priori that our attacks translate to practically efficient attacks that could be implemented in real hardware because it can be difficult to dynamically reallocate memory between processes in an ASIC (the amount of memory used during each round of a balloon phase is significantly greater than the amount of memory used during each round of a light phase).

²⁶In fact these plots actually show attack quality against *any* strict λ -stacked sandwich DAG. Thus, we do not include separate plots for the Double-Buffer and the Linear functions of [CGBS16] because, like Catena, these iMHFs are based on strict λ -stacked sandwich DAGs. Thus, our attack quality against these functions will be the same as our attack quality against Catena.

In Appendix 8 we argue that this architecture challenge would not be a fundamental barrier to an adversary. In particular, we outline an architecture for our algorithm PGenPeb using Argon2i as an example.

Briefly, we execute $n^{1/4}$ instances of the iMHF in parallel. Our architecture includes $n^{1/4}$ “light phase” chips and a single “Balloon Phase” chip which is responsible for executing all of the balloon phases in a round robin fashion. Each light phase chip only needs $O(n^{3/4} \ln n)$ memory and a single instance of the compression function H . The central balloon phase chip needs to have $O(n \ln n)$ memory and \sqrt{n} instances of the compression functions H .

8 Towards A Practical Attack Architecture

In previous sections we demonstrated efficient attacks on the Catena, Argon2i and the Balloon Hashing iMHFs in the Parallel Random Oracle Machine (pROM) model of computation. Because pROM is a theoretical model of computation it is not obvious a priori that our attacks translate to practically efficient attacks that could be implemented in real hardware. In particular, the amount of memory used during each round of a balloon phase is significantly greater than the amount of memory used during each round of a light phase. However, in an ASIC it can be difficult to dynamically reallocate memory between processes.

In this section we argue that this challenge would not be fundamental barrier to an adversary. Our attacks on these iMHFs will lead to practical attacks when implemented in real hardware. In the following discussion we will focus on our attack PGenPeb(G, S, g, d, k) on Argon2i (i.e., $|S| = O(n^{3/4} \ln n)$, $g = n^{3/4}$, $d = \sqrt{n}$ and $k = n^{1/4}$). However, the architecture we describe could easily be adapted for the Catena or Balloon Hashing iMHFs.

Figure 2 contains a diagram of our proposed attack architecture. Specifically, we will evaluate multiple iMHF instances in parallel. For each iMHF instance we will have a small dedicated chip to evaluate the light phases of that iMHF instance. We will use a single (larger) central unit which is dedicated to execute the balloon phases for each iMHF instance in a round-robin fashion.

In particular, we evaluate $k = O(n^{1/4})$ iMHF instances in parallel so we have k smaller dedicated chips. Each of these smaller dedicated processing unit contains a dedicated memory of size $O(n^{3/4} \log n)$ and a single instance of the underlying random oracle H . Thus, each dedicated processing unit contains sufficient resources to execute a light phases, but not to execute a balloon phase in PGenPeb.

Our central chip will contain \sqrt{n} instances of the random oracle H and $O(n \ln n)$ memory. We will use the central chip to execute balloon phases for each iMHF instance in a round robin fashion. The central chip will store $O(n^{3/4} \ln n)$ random oracle values corresponding the the set S_i for each iMHF instance $i \leq k$ — $O(n \ln n)$ total memory. In addition, for each iMHF instance, the central chip will also store the $g = O(n^{3/4})$ values associated with the parents of the next g nodes for that instance.

Each light phase unit will need to send the central chip the values in the set S_i — we can send these values as they are computed. After executing a balloon phase for instance i the central chip will need to send up to g values (corresponding to the parents of the next g nodes to be pebbled) to the i^{th} light phase chip. Thus, to implement the attack we will also need to have dedicated communication channels between the central chip and each dedicated light phase chip. However, each dedicated memory channel only needs to have constant bandwidth. In particular, we don’t need to send all g values to the i^{th} immediately. Instead, during each pebbling round the central chip will send each dedicated light phase chip the values corresponding to the parent(s) that are required to make progress in the next round. Similarly, every time a dedicated light phase chip for iMHF instance i computes the value for a node $v \in S_i$ we will simply send that value to the central chip.

9 Conclusions and Open Questions

We introduced the new measure of energy complexity. The AT and energy complexity of each of the attacks in this work behave asymptotically the same as the number of evaluations $\#inst$ grows. However when evaluating a single instance the energy complexity can be significantly lower then the AT complexity.




Figure 2: Argon2i: Attack Architecture

Intuitively, this occurs when an evaluation algorithm does not allow for making optimal use of available hardware – that is significant amounts of the circuitry are not used during some of the computation. In particular, especially for MHFs based on (collision resistant) compression functions, it seems that while total (i.e. non-amortized) AT complexity need not scale linearly in $\#inst$ their energy complexities do tend to. One reason for this is that while implementations of H can be reused in parallel across concurrently running evaluations of the MHF increasing neither time nor area of the chip (and so making AT complexity scale badly) the domains/range subsets of H relevant to different inputs to the MHF remain disjoint. In other words while evaluating an MHF may, during some periods, require many implementations of H in order to achieve optimal AT efficiency it may be possible to reuse most of those H implementations for concurrent evaluations during other periods. However it seems much easier to ensure that relevant input/output pairs for H needed in one evaluation of the MHF are of no use when evaluating it on any different input.²⁷

We gave a generic attack on any data independent iMHF with quality $\Omega(\log^{0.999} n)$, and we gave even higher quality attacks against the Catena and Argon2i DAGs — the winner of the password hashing competition. Furthermore, we demonstrated that our attacks on Catena and Argon2i could be highly effective for realistic values of n . Thus, we would strongly advise against the use of the Argon2i and Catena iMHFs for password hashing.

Currently, the only candidate iMHF equipped with a full proof of security is the one in [AS15]. Our results suggest that, up to polylogarithmic factors, this iMHF is optimal. However, this iMHF has several downsides for practical applications. First, the underlying DAGs do not known to have a practically succinct description. Ideally, we would like a DAG $G = (V = [n], E)$ with the property that for any vertex $v \in [n]$ there is an efficient algorithm to compute the set $\text{parents}(v)$ so that G does not need to be stored in memory. Second, the polylogarithmic factors are too large for the security bounds to be meaningful in practice. In particular, the energy complexity is proved to be $\Omega(n^2/(\log^{10} n(\log \log n)^2))$, but for $n < 2^{23}$ the denominator is still larger than the numerator. Could we find an iMHF based on a DAG G with a succinct description such that for any PROM attacker A with energy complexity $n^2/\log n$ ²⁸? The $\log n$ factor might be acceptable for practical values of n (e.g., $n \leq 2^{40}$) especially considering that our attacks typically compute H more times than the naive algorithm and the cost of storing a hash value in memory for 1 time step is about $\bar{R} = 3,000$

²⁷This observation is analogous to the distinction between amortized CC pebbling complexity and amortized ST pebbling complexity in the parallel pebbling game of [AS15].

²⁸It is still possible that a tighter analysis of the construction from [AS15] might show that their DAG satisfies this property.

cheaper than computing H .

Our results imply that depth-robustness is a necessary condition for a good iMHF, but we have not proved that this property is sufficient. There exists a DAG G with $\text{indeg}(G) = O(\log n)$ such that G is $(\alpha n, \beta n)$ -depth robust for some constants $\alpha, \beta > 0$ [EGS75]. Could we prove that for any such $(\alpha n, \beta n)$ -depth robust DAG G the corresponding iMHF has high energy complexity (e.g., say $\Omega(n^2/\log n)$)? Alternatively, could we find an $(\alpha n, \beta n)$ -depth robust DAG G which has lower energy complexity (e.g., say $O(n^2/\log^2 n)$)? Another interesting complexity theoretic question is to resolve the complexity of finding an (approximately) optimal pebbling strategy \mathcal{A} for a particular DAG G . If there were an efficient algorithm to find an (approximately) optimal pebbling strategy in terms of energy or AT -complexity then we could use the algorithm to quickly rule out bad iMHFs.²⁹

Our attacks only apply to data independent iMHFs. Could we build ideal data dependent MHFs? For example, `scrypt` [Per09] is a widely deployed data dependent MHF that provably has high ST-complexity. However, the security proof does not rule out the possibility of attacks with low amortized AT or energy complexity. What is the amortized energy complexity of `scrypt`? Can we develop data dependent MHFs with (provably) high (e.g., $\Omega(n^2)$) energy and AT complexity? If so then it may be worthwhile to explore alternative defenses (e.g., moving target defenses [JGS⁺11] or oblivious RAM [Gol87],[SCSL11]) to mitigate the risks of side-channel attacks on data dependent MHFs. Another defense against side-channel attacks would be to ensure that all memory accesses only depend on the cryptographic hash $H(pwd||salt)$ where $salt$ is a sufficiently long bit string. In this case an attacker who observes the memory access pattern will not be able to infer the password as long as the value $salt$ remains hidden³⁰.

References

- [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.
- [ACK⁺16] Joël Alwen, Binyi Chen, Chethan Kamath, Vladimir Kolmogorov, Krzysztof Pietrzak, and Stefano Tessaro. On the complexity of `scrypt` and proofs of space in the parallel random oracle model. Cryptology ePrint Archive, Report 2016/100, 2016. <http://eprint.iacr.org/>.
- [Ale04] S. Alexander. Password protection for modern operating systems. *login*, June 2004.
- [AS15] Joël Alwen and Vladimir Serbinenko. High Parallel Complexity Graphs and Memory-Hard Functions. In *Proceedings of the Eleventh Annual ACM Symposium on Theory of Computing*, STOC '15, 2015. <http://eprint.iacr.org/2014/238>.
- [ASK07] Onur Aciçmez, Werner Schindler, and Çetin Kaya Koç. Cache based remote timing attack on the AES. In Masayuki Abe, editor, *Topics in Cryptology - CT-RSA 2007, The Cryptographers' Track at the RSA Conference 2007, San Francisco, CA, USA, February 5-9, 2007, Proceedings*, volume 4377 of *Lecture Notes in Computer Science*, pages 271–286. Springer, 2007.
- [BBBB10] Hristo Bojinov, Elie Bursztein, Xavier Boyen, and Dan Boneh. Kamouflage: Loss-resistant password management. In *Computer Security—ESORICS 2010*, pages 286–302. Springer, 2010.
- [BBD13] Jeremiah Blocki, Manuel Blum, and Anupam Datta. Gotcha password hackers! In *Proceedings of the 2013 ACM workshop on Artificial intelligence and security*, pages 25–34. ACM, 2013.

²⁹On the other side if finding a good pebbling strategy is intractable, then it might be possible to build an iMHF for which it is intractable to find a high quality attack (e.g., say $\text{Equality}(\mathcal{A}) = \Omega(\log n)$) even when such an attack exists.

³⁰However, if the adversary observes $salt$ as well as the memory access pattern while the iMHF was being computed with users password pwd as input then his task suddenly becomes much easier — he can eliminate password guesses as soon as memory access pattern does not match.

- [BDK15] Alex Biryukov, Daniel Dinu, and Dmitry Khovratovich. Fast and tradeoff-resilient memory-hard functions for cryptocurrencies and password hashing. Cryptology ePrint Archive, Report 2015/430, 2015. <http://eprint.iacr.org/>.
- [Ber] Daniel J. Bernstein. Cache-Timing Attacks on AES.
- [BF1] Butterfly Labs Products. Retrieved 8, Nov. 2015.
- [Bil13] Billy Markus. Dogecoin, 2013.
- [BJKS03] John G Brainard, Ari Juels, Burt Kaliski, and Michael Szydlo. A new two-server approach for authentication with short secrets. In *USENIX Security*, volume 3, pages 201–214, 2003.
- [BK15] Alex Biryukov and Dmitry Khovratovich. Tradeoff cryptanalysis of memory-hard functions. Cryptology ePrint Archive, Report 2015/227, 2015. <http://eprint.iacr.org/>.
- [BKPS13] Jeremiah Blocki, Saranga Komanduri, Ariel Procaccia, and Or Sheffet. Optimizing password composition policies. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, pages 105–122. ACM, 2013.
- [BL13] Daniel J. Bernstein and Tanja Lange. Non-uniform cracks in the concrete: The power of free precomputation. In Kazue Sako and Palash Sarkar, editors, *Advances in Cryptology - ASIACRYPT 2013 - 19th International Conference on the Theory and Application of Cryptology and Information Security, Bengaluru, India, December 1-5, 2013, Proceedings, Part II*, volume 8270 of *Lecture Notes in Computer Science*, pages 321–340. Springer, 2013.
- [BM06] Joseph Bonneau and Ilya Mironov. Cache-collision timing attacks against aes. In *Cryptographic Hardware and Embedded Systems - CHES 2006, 8th International Workshop*, volume 4249 of *Lecture Notes in Computer Science*, pages 201–215. Springer, 2006.
- [Bon12] J. Bonneau. The science of guessing: analyzing an anonymized corpus of 70 million passwords. In *Security and Privacy (SP), 2012 IEEE Symposium on*, pages 538–552. IEEE, 2012.
- [BP1] Bitmain Products. Retrieved 8, Nov. 2015.
- [CGBS16] Henry Corrigan-Gibbs, Dan Boneh, and Stuart Schechter. Balloon hashing: Provably space-hard hash functions with data-independent access patterns. Cryptology ePrint Archive, Report 2016/027, 2016. <http://eprint.iacr.org/>.
- [Cha11] Charles Lee. Litecoin, 2011.
- [CHS06] Ran Canetti, Shai Halevi, and Michael Steiner. Mitigating dictionary attacks on password-protected local storage. In *Advances in Cryptology-CRYPTO 2006*, pages 160–179. Springer, 2006.
- [CLN12] Jan Camenisch, Anna Lysyanskaya, and Gregory Neven. Practical yet universally composable two-server password-authenticated secret sharing. In *Proceedings of the 2012 ACM conference on Computer and Communications Security*, pages 525–536. ACM, 2012.
- [Coi] Coinplorer: Bitcoin mining hardware comparison. Retrieved 8, Nov. 2015.
- [DC08] Waseem Daher and Ran Canetti. Posh: A generalized captcha with security applications. In *Proceedings of the 1st ACM workshop on Workshop on AISec*, pages 1–10. ACM, 2008.
- [DGN03] Cynthia Dwork, Andrew Goldberg, and Moni Naor. On memory-bound functions for fighting spam. In *Advances in Cryptology - CRYPTO 2003, 23rd Annual International Cryptology Conference, Santa Barbara, California, USA, August 17-21, 2003, Proceedings*, volume 2729 of *Lecture Notes in Computer Science*, pages 426–444. Springer, 2003.

- [DNW05] Cynthia Dwork, Moni Naor, and Hoeteck Wee. Pebbling and proofs of work. In *Advances in Cryptology - CRYPTO 2005: 25th Annual International Cryptology Conference, Santa Barbara, California, USA, August 14-18, 2005, Proceedings*, volume 3621 of *Lecture Notes in Computer Science*, pages 37–54. Springer, 2005.
- [EGS75] Paul Erdős, Ronald L. Graham, and Endre Szemerédi. On sparse graphs with dense long paths. Technical report, Stanford, CA, USA, 1975.
- [FH10] D. Florêncio and C. Herley. Where do security policies come from. In *Proc. of SOUPS*, page 10, 2010.
- [FLW13] Christian Forler, Stefan Lucks, and Jakob Wenzel. Catena: A memory-consuming password scrambler. *IACR Cryptology ePrint Archive*, 2013:525, 2013.
- [Gol87] Oded Goldreich. Towards a theory of software protection and simulation by oblivious RAMs. In *Proceedings of the nineteenth annual ACM symposium on Theory of computing*, pages 182–194. ACM, 1987.
- [JGS⁺11] Sushil Jajodia, Anup K Ghosh, Vipin Swarup, Cliff Wang, and X Sean Wang. *Moving target defense: creating asymmetric uncertainty for cyber threats*, volume 54. Springer Science & Business Media, 2011.
- [JR13] Ari Juels and Ronald L. Rivest. Honeywords: Making password-cracking detectable. In *Proceedings of the 2012 ACM conference on Computer and communications security*. ACM, 2013.
- [Kal00] Burt Kaliski. Pkcs# 5: Password-based cryptography specification version 2.0. 2000.
- [KSK⁺11] S. Komanduri, R. Shay, P.G. Kelley, M.L. Mazurek, L. Bauer, N. Christin, L.F. Cranor, and S. Egelman. Of passwords and people: measuring the effect of password-composition policies. In *Proceedings of the 2011 annual conference on Human factors in computing systems*, pages 2595–2604. ACM, 2011.
- [LT82] Thomas Lengauer and Robert E Tarjan. Asymptotically tight bounds on time-space trade-offs in a pebble game. *Journal of the ACM (JACM)*, 29(4):1087–1130, 1982.
- [Man96] Udi Manber. A simple scheme to make passwords based on one-way functions much harder to crack. *Computers & Security*, 15(2):171–176, 1996.
- [MMV13] Mohammad Mahmoody, Tal Moran, and Salil P. Vadhan. Publicly verifiable proofs of sequential work. In Robert D. Kleinberg, editor, *Innovations in Theoretical Computer Science, ITCS '13, Berkeley, CA, USA, January 9-12, 2013*, pages 373–388. ACM, 2013.
- [MR1] Mining Hardware Comparison. Retrieved 8, Nov. 2015.
- [MT79] Robert Morris and Ken Thompson. Password security: A case history. *Communications of the ACM*, 22(11):594–597, 1979.
- [NB⁺15] Arvind Narayanan, Joseph Bonneau, , Edward W Felten, Andrew Miller, and Steven Goldfeder. *Bitcoin and Cryptocurrency Technology (manuscript)*. 2015. Retrieved 8/6/2015.
- [Per09] C. Percival. Stronger key derivation via sequential memory-hard functions. In *BSDCan 2009*, 2009.
- [PHC] Password hashing competition. <https://password-hashing.net/>.
- [PM] N. Provos and D. Mazieres. Bcrypt algorithm.

- [RTSS09] Thomas Ristenpart, Eran Tromer, Hovav Shacham, and Stefan Savage. Hey, you, get off of my cloud: exploring information leakage in third-party compute clouds. In Ehab Al-Shaer, Somesh Jha, and Angelos D. Keromytis, editors, *Proceedings of the 2009 ACM Conference on Computer and Communications Security, CCS 2009, Chicago, Illinois, USA, November 9-13, 2009*, pages 199–212. ACM, 2009.
- [SCSL11] Elaine Shi, T-H Hubert Chan, Emil Stefanov, and Mingfei Li. Oblivious ram with $o((\log n)^3)$ worst-case cost. In *Advances in Cryptology–ASIACRYPT 2011*, pages 197–214. Springer, 2011.
- [Spo] Spondoolies Tech Products. Retrieved 8, Nov. 2015.
- [SS09] Karen Scarfone and Murugiah. Souppaya. Nist special publication 800-118: Guide to enterprise password management (draft), April 2009.
- [Tho79] Clark D. Thompson. Area-time complexity for VLSI. In Michael J. Fischer, Richard A. DeMillo, Nancy A. Lynch, Walter A. Burkhard, and Alfred V. Aho, editors, *Proceedings of the 11th Annual ACM Symposium on Theory of Computing, April 30 - May 2, 1979, Atlanta, Georgia, USA*, pages 81–88. ACM, 1979.
- [Val77] Leslie G. Valiant. Graph-theoretic arguments in low-level complexity. In Jozef Gruska, editor, *Mathematical Foundations of Computer Science 1977, 6th Symposium, Tatranska Lomnica, Czechoslovakia, September 5-9, 1977, Proceedings*, volume 53 of *Lecture Notes in Computer Science*, pages 162–176. Springer, 1977.

Appendix

A Catena DAGs

The Catena Dragonfly iMHF [FLW13] is based on a stack of λ bit-reversal graphs: BRG_λ^n . More formally, BRG_λ^n is an $n = 2^g(\lambda + 1)$ node DAG G with nodes $V = \{1, \dots, n\}$. These nodes are partitioned into $\lambda + 1$ layers L_0, \dots, L_λ of size 2^g . In particular, $L_i = [i \cdot 2^g + 1, (i + 1)2^g]$. The DAG G includes directed edges $(i, i + 1)$ for $1 \leq i < n$. Additionally, given $x \in [0, 2^g - 1]$ there is a directed edge from node $u = i \cdot 2^g + x + 1$ in layer L_i to $(i + 1)2^g + \mathbf{BR}(x) + 1$, where \mathbf{BR} denotes the binary bit reversal function. For example, if $g = 3$ and $x = 6$ then the binary representation of x is 110 so $\mathbf{BR}(110) = 011$, which is 3 in decimal representation.

From the definition it is easy to verify that this DAG is a strict λ -stacked sandwich DAG.

Fact A.1 *Let $n = 2^g(\lambda + 1)$ then BRG_λ^n is a strict λ -stacked sandwich DAG with maximum indegree $\text{indeg}(\text{BRG}_\lambda^n) = 2$.*

The Catena Butterfly iMHF [FLW13] is based on a stack of λ double-butterfly DAGs: DBG_λ^n . In turn, each double-butterfly DAG is based on the Cooley-Turkey Fast Fourier Transformation algorithm and consists of $O(\log n)$ layers. Formally, a DBG_λ^n consists of $n = 2^g \cdot (\lambda(2g - 1) + 1)$ nodes $V = \{1, \dots, n\}$ and included directed edges $(i, i + 1)$ for each $1 \leq i < n$. Given $1 \leq i \leq \lambda(2g - 1) + 1$ we will write $L_i = \{(i - 1)2^g + 1, \dots, i \cdot 2^g\}$. Note that the layers $L_1, \dots, L_{\lambda(2g-1)+1}$ partition V and each have size 2^g . Given parameters $1 \leq k \leq \lambda$ and $0 \leq j \leq 2^g - 1$ and $1 \leq i \leq 2g - 1$ we will write $v_{i,j}^k$ to denote the node $(k - 1)2^g\lambda(2g - 1) + i2^g\lambda + j$ and we will write $L_i^k = L_{(k-1)\lambda(2g-1)+i} = [v_{i,0}^k, v_{i,2^g-1}^k]$ for notational convenience. For each $1 \leq k \leq \lambda$, $1 \leq i \leq 2g - 2$ and $0 \leq j \leq 2^g - 1$ we will add vertical edges $(v_{i,j}^k, v_{i+1,j}^k)$ between consecutive layers L_i^k and L_{i+1}^k . For each $1 \leq k \leq \lambda$ and $0 \leq j \leq 2^g - 1$ we will also add diagonal edges $(v_{i,j}^k, v_{i+1,j \oplus 2^{g-1}-i}^k)$ whenever $1 \leq i \leq g - 1$ and we will add diagonal edges $(v_{i,j}^k, v_{i+1,j \oplus 2^{i-g+1}}^k)$ whenever $g \leq i \leq 2g - 2$. In both cases the edges always go between consecutive layers L_i^k and L_{i+1}^k . Here, $x \oplus y$ denotes the bit wise XOR operation. For example, if $x = 6$ (i.e., $x = 110$ in binary) and $y = 3$ (i.e., $x = 011$ in binary) then $x \oplus y = 5$ (i.e., 101 in binary).

It is easy to verify that DBG_λ^n is a strict $\lambda(2g-1)$ -stacked sandwich DAG and that the maximum indegree is $\text{indeg}(\text{DBG}_\lambda^n) = 3$.

Fact A.2 *Let $n = 2^g \cdot (\lambda(2g-1) + 1)$ then DBG_λ^n is a strict $\lambda(2g-1)$ -stacked sandwich DAG with maximum indegree $\text{indeg}(\text{BRG}_\lambda^n) = 3$.*

A.1 Improved Attacks for Special Cases

For special cases (e.g., $\lambda = 1$) an alternative implementation of **GenPeb** yields a better upper bound on total memory. In particular, we will use **need** from Algorithm 3 and we will define **keep** to simply return the exact same set as **need** in each round of **GenPeb** so that we don't immediately throw out pebbles in step 14 of the same round. Lemma A.3 states that the pair **keep** and **need** is valid whenever we run **GenPeb** on $\lambda = 1$ -stacked sandwich DAGs.

Lemma A.3 *Let G be a $\lambda = 1$ -stacked sandwich DAG on n nodes and maximum indegree $\text{indeg}(G) = \delta$ and let $S = \{it \mid i \leq n/t\}$ denote the set of size n/t such that $\text{depth}(G - S) \leq 2t$. Let **need** denote the subroutine from Algorithm 3 and let **keep** denote the subroutine that returns the same set as **need** in every pebbling round j of **GenPeb**. Then the pair **need** and **keep** is valid for **GenPeb**(G, S, g, d) whenever $g \geq d \geq 2t$. Furthermore, **GenPeb**(G, S, g, d) keeps at most $n/t + 1 + g(\delta - 1)$ pebbles on G at any point in time.*

PROOF. (sketch) Consider round i of a balloon phase and let $j = i \bmod g$ and let N_i denote the set returned by both **need** and **keep** in round i . Any pebbled node $v \in L_0$ during round i is either (1) the parent of a node in $[i, i+g]$, (2) in the set N_i or (3) not in $N_i \cup \text{parents}([i, i+g])$. In the first case $v \in \text{parents}([i, i+g])$ we do not need to include v in the set **keep**($i, i+g$) as **GenPeb** already keeps these nodes by definition. In the second case $v \in N_i$ we do keep v by definition. In the third case we observe that there is no unpebbled path from any child of v to any node in $[i, i+g]$ — any such path must either travel through the pebbled node $i-1$ in layer L_1 or through a pebbled node in $N_i \cup S$ in layer L_0 . This completes the proof that **need** and **keep** are valid. We note that **GenPeb** discards any pebble that is not in the set $S \cup N_i \cup \text{parents}(i, i+g)$. Thus, we keep at most $1 + |S| + |N_i| + |\text{parents}(i, i+g)| \leq 1 + n/t + n/(2t) + (\delta - 1)g$ pebbles on the graph at any point in time. \square

Corollary A.4 (Complexity of Catena) *Let G be a 1-stacked sandwich DAG on n nodes and maximum indegree $\text{indeg}(G) = \delta$ then there is an algorithm \mathcal{A} with $\text{amc}(\mathcal{A}) = O(\delta\sqrt{n})$, $\text{arc}(\mathcal{A}) = O(\sqrt{n})$, $\text{crc}(\mathcal{A}) = O(n^{1.5})$ and $\text{cmc}(\mathcal{A}) = O(\delta n^{1.5})$. In particular for any core-memory area and energy ratios R and \bar{R} there exists an evaluation algorithm \mathcal{A} for the corresponding iMHF with*

$$\text{E-quality}_{\bar{R}}(\mathcal{A}) = O\left(\frac{\sqrt{n}}{\bar{R} \cdot \delta}\right), \text{ and } \text{AT-quality}_R(\mathcal{A}) = O\left(\frac{\sqrt{n}}{R \cdot \delta}\right).$$

PROOF. \mathcal{A} runs **GenPeb**($G, S, g = \sqrt{n}, d = \sqrt{n}$) with the pair **need** and **keep** from Lemma A.3. Lemma A.3 implies that we only need to keep $O(g\delta)$ pebbles on the graph at any point in time. Thus, $\text{amc}(\mathcal{A}) = O(\delta\sqrt{n})$ and $\text{cmc}(\mathcal{A}) = O(\delta n^{1.5})$. At any point in time \mathcal{A} places $O(n/d)$ new pebbles on the graph. Thus, $\text{arc}(\mathcal{A}) = O(\sqrt{n})$ and $\text{crc}(\mathcal{A}) = O(n^{1.5})$. \square

While Theorem 4.6 and Theorem 4.5 only hold for $\lambda < n^{1/3}$ there is a trivial algorithm with complexity $O(n^{5/3})$ when $\lambda > n^{1/3}$.

Theorem A.5 (Complexity of Strict Sandwich Graph with Large λ .) *Let G be a strict λ -stacked sandwich DAG on n nodes with $\lambda > n^{1/3}$ and maximum indegree $\text{indeg}(G) = \delta$ then for any core-memory area and energy ratios $R > 0$ and $\bar{R} > 0$ there is an attack \mathcal{A} on the corresponding iMHF with*

$$E_{\bar{R}}(\mathcal{A}) \leq \frac{2n^2}{\lambda} + \bar{R}n, \text{ and } \text{AT}_R(\mathcal{A}) \leq \frac{2n^2}{\lambda} + Rn.$$

PROOF. Let \mathcal{A} be the algorithm that runs the naïve pebbling algorithm \mathcal{N} with one simple optimization: when \mathcal{A} first place a pebble on some node $i \in L_k$ in layer k then \mathcal{A} throws out pebbles on nodes in layer $L_{\leq k-1}$. The pebbling is valid for strict λ -stacked sandwich DAGs. In each round we place at most 1 new pebble on the graph and the pebbling stops after n rounds. Thus, $\text{arc}(\mathcal{A}) \leq 1$ and $\text{crc}(\mathcal{A}) \leq n$. We also have $\text{amc}(\mathcal{A}) \leq \frac{2n}{\lambda+1}$ and $\text{cmc}(\mathcal{A}) \leq \frac{2n^2}{\lambda+1}$. The result now follows immediately from the definitions of $AT_{\bar{R}}(\mathcal{A})$ and $E_{\bar{R}}(\mathcal{A})$. \square \square

B Extra Plots

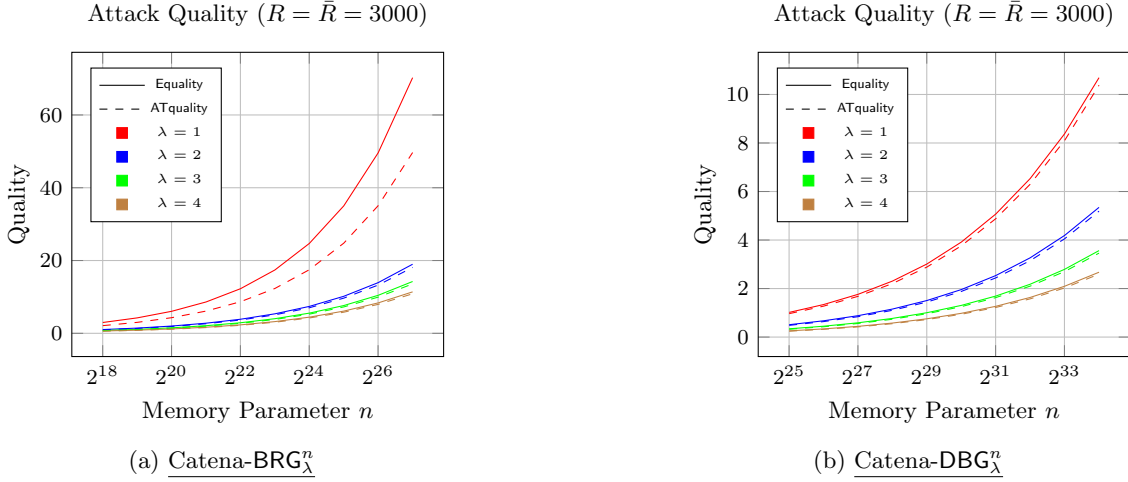


Figure 3: Effect of varying λ on complexity of PGenPeb for Catena iMHFs.




Figure 4: Effect of varying error probability γ on complexity of PGenPeb.

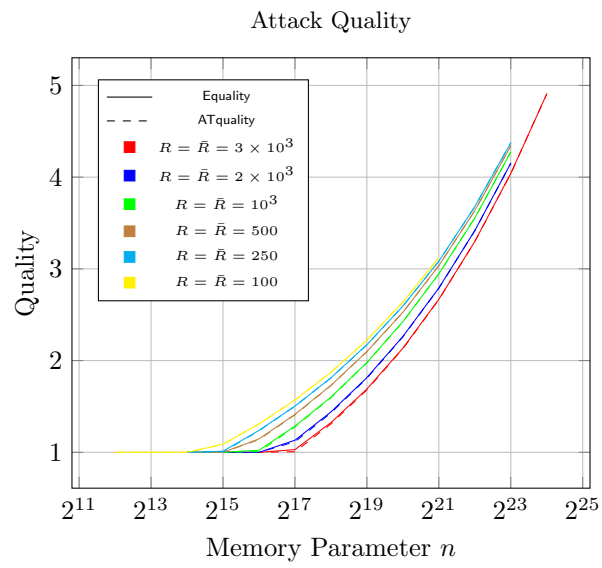


Figure 5: Effect of varying core memory & energy ratios for complexity of PGenPeb evaluating Argon2i ($\tau = 1, \sigma = n, \gamma = 10^{-4}$)