# Trapdoor Memory-Hard Functions 

Benedikt Auerbach©<br>benedikt.auerbach@ista.ac.at

ISTA

Christoph U. Günther©<br>cguenthe@ista.ac.at

ISTA
Krzysztof Pietrzak
pietrzak@ista.ac.at
ISTA
February 23, 2024


#### Abstract

Memory-hard functions (MHF) are functions whose evaluation provably requires a lot of memory. While MHFs are an unkeyed primitive, it is natural to consider the notion of trapdoor MHFs (TMHFs). A TMHF is like an MHF, but when sampling the public parameters one also samples a trapdoor which allows evaluating the function much cheaper.

Biryukov and Perrin (Asiacrypt'17) were the first to consider TMHFs and put forth a candidate TMHF construction called Diodon that is based on the Scrypt MHF (Percival, BSDCan'09). To allow for a trapdoor, Scrypt's initial hash chain is replaced by a sequence of squares in a group of unknown order where the order of the group is the trapdoor. For a length $n$ sequence of squares and a group of order $N$, DIODON's cumulative memory complexity (CMC) is $O\left(n^{2} \log N\right)$ without the trapdoor and $O\left(n \log (n) \log (N)^{2}\right)$ with knowledge of it.

While Scrypt is proven to be optimally memory-hard in the random oracle model (Alwen et al., Eurocrypt'17), Diodon's memory-hardness has not been proven so far. In this work, we fill this gap by rigorously analyzing a specific instantiation of DIOdon. We show that its CMC is lower bounded by $\Omega\left(\frac{n^{2}}{\log n} \log N\right)$ which almost matches the upper bound. Our proof is based Alwen et al.'s lower bound on SCRYPT's CMC but requires non-trivial modifications due to the algebraic structure of Diodon. Most importantly, our analysis involves a more elaborate compression argument and a solvability criterion for certain systems of Diophantine equations.


## Contents

1 Introduction ..... 3
1.1 Memory-Hard Functions ..... 3
1.2 Trapdoor MHFs ..... 4
1.3 The Diodon TMHF ..... 5
1.4 Contributions and Technical Overview ..... 5
1.5 Open Problems ..... 7
2 Preliminaries ..... 7
2.1 Notation ..... 7
2.2 Algebraic Setting ..... 8
2.3 Generic Group Model ..... 9
2.4 Machine Model and Complexity Measure ..... 9
3 A Trapdoor Memory-Hard Function from Factoring ..... 10
3.1 Trapdoor Memory-Hard Functions ..... 10
3.2 Description of TDSCRYPT ..... 11
4 Overview of the Lower Bound Proof ..... 13
5 Single-challenge Time-Memory Trade-Off ..... 14
5.1 Reasoning About $\mathcal{A}_{1}$ 's Queries Algebraically ..... 15
5.2 Proof Skeleton ..... 18
5.3 Analyzing the Behavior of $A \boldsymbol{x}=\boldsymbol{b}$ ..... 19
5.4 Combinatorial Proof of the $\operatorname{rank}(A)$ Lower Bound ..... 20
5.5 Incompressibility Argument ..... 23
6 Multi-challenge Memory Complexity ..... 25
A Detailed Cumulative Memory Complexity Analysis ..... 32
B Bounding the Entries of $A$ ..... 34
C Proof of Claim 5 ..... 35
C. 1 Notation and Data Structures ..... 36
C. 2 Encoder ..... 37
C. 3 Encoding Size ..... 40
C. 4 Success Probability ..... 41

## 1 Introduction

Moderately-hard functions have many applications, the most prominent one being password hashing. Early constructions of such functions aimed to be moderately hard in terms of computation. For example, PBKDF2 [Kal00] is essentially a regular hash function repeated sufficiently many times.

Unfortunately, computationally expensive functions are not egalitarian. Attackers can use specialized hardware (like FPGAs or ASICs) to evaluate some specific computationally expensive function several orders of magnitude more efficiently (in terms of energy and hardware cost) than general-purpose hardware. This creates an asymmetry between the cost for honest users and attackers. Specifically with password hashing in mind, Percival [Per09] introduced the notion of memory-hard functions (MHFs).

### 1.1 Memory-Hard Functions

Informally, a function is memory-hard if its evaluation cost on general-purpose hardware is dominated by the memory (rather than CPU) cost. Since memory costs are roughly the same for specialized- and general-purpose hardware, MHFs are more egalitarian than computationally-hard functions. The first MHF construction was SCRyPt [Per09] followed by many others (e.g., [BDK16, BCS16, ABH17, $\left.\mathrm{BHK}^{+} 19, \mathrm{BH} 22\right]$ ). These constructions primarily differ in their notion of memory-hardness and side-channel resistance.

Memory-Hardness. A popular way to measure the memory-hardness of a function is the cumulative memory complexity (CMC) [AS15]. It basically sums the memory cost over all steps ${ }^{1}$ of an evaluation. A secure MHF not only ensures that the CMC of the honest evaluation algorithm is high, but also that any other, adversarial evaluation algorithm has a CMC not much lower than the honest evaluation algorithm. While the honest evaluation algorithm uses little to no parallelism, the adversarial algorithm is allowed arbitrarily many parallel queries to the random oracle.

Unfortunately, the definition of CMC does not exclude time-memory trade-offs. In particular, an MHF may be evaluated in more steps where each step requires less memory [RD16]. A stronger notion is sustained space complexity [ABP18] which only counts steps where the memory usage is sufficiently high.

Finally, let us mention the related notion of bandwidth-hardness [RD17, BRZ18] capturing the number of cache-misses rather than the memory usage. This captures the energy cost of evaluating a function more accurately, whereas CMC is a better measure for the hardware cost.

Side-channel Resistance. MHFs come in two flavors, data-dependent and dataindependent MHFs (dMHF and iMHF, respectively), which classify the side-channel resistance of MHFs. The memory-access patterns during the evaluation of an iMHF do not depend on the input but are fixed. In contrast, dMHFs allow the memory-access patterns to depend on the input. While dMHFs are easier to construct and can provably

[^0]achieve higher evaluation cost [ABP17], their security can be compromised by side-channel attacks which leak memory access patterns. A notion aiming to combine the advantages of iMHFs and dMHFs using computational assumptions was proposed in [ABZ20].

### 1.2 Trapdoor MHFs

An MHF is an unkeyed primitive, but like for other unkeyed primitives, say, one-way permutations or collision-resistant hash functions, it is natural to consider a keyed version. Biryukov and Perrin [BP17] introduced asymmetrically memory-hard functions which are essentially MHFs admitting a trapdoor. So similarly to an MHF, evaluating the function is guaranteed to cost a lot of memory in general. However, in contrast to an MHF, when sampling the parameters one also generates a secret trapdoor. Knowledge of this trapdoor allows evaluating the function much cheaper using less memory. Therefore, we call such functions trapdoor MHFs (TMHFs) and note that the PURED framework [BLP23]-an effort to classify all types of resource-hard functions - also uses the word "trapdoor" to describe such functions.

In this work, we focus on a data-dependent TMHF that provably achieves high CMC. To precisely quantify the gap between the honest evaluation algorithms (with and without knowledge of the trapdoor) we use the notation $\left(c_{\text {hon }}, c_{\mathrm{td}}\right)$-TMHF. This means that the honest evaluation algorithms without and with knowledge of the trapdoor have a CMC of $c_{\text {hon }}$ and $c_{\mathrm{td}}$, respectively.

Applications. One potential application of TMHFs are proofs of CMC that benefit from efficient private verification. For example, consider an e-mail server that wants to combat junk mail [DN93] by requiring e-mail senders to solve a TMHF. Thanks to the trapdoor, the server can verify the sender's response with fewer resources than the sender. So the server can choose larger parameters that would otherwise be too costly, especially if the gap between $c_{\text {hon }}$ and $c_{\mathrm{td}}$ is large. ${ }^{2}$

Another application is password hashing, albeit a less convincing one. In principle, the trapdoor allows a server to verify logins more cheaply compared to using an MHF. However, the trapdoor needs to be stored securely (e.g., in an HSM) and then it is conceivable to simply encrypt the passwords instead of hashing them. One downside of encryption is that a compromise of the key (e.g., the HSM is broken) reveals all passwords, a catastrophical failure. Since TMHFs would still require some bruteforcing in this case, they might still be preferable.

Related Primitives. Memory-hard puzzles [ABB22] are closely related, yet different. Solving a puzzle is memory-hard, but it is easy to sample a puzzle that will evaluate to a specific solution. While the first property is comparable to evaluating a TMHF without knowledge of the trapdoor, the second property can roughly be seen as the opposite of the trapdoor evaluation.

[^1]
### 1.3 The Diodon TMHF

Diodon [BP17] is the first construction aiming to be TMHF. It is based on Scrypt [Per09], a well-known data-dependent MHF that has provably high memory-hardness $\left[\mathrm{ACK}^{+} 16\right.$, $\left.\mathrm{ACP}^{+} 17\right]$. DIODON as stated in [BP17] offers multiple parameters to fine-tune the security and performance of the function. In this paper, we stick to the natural choice of parameters that most closely resembles SCRYPT ${ }^{3}$ and also work in a slightly different algebraic setting. To avoid confusion, we call this specific instantiation TDScrypt.

TdScrypt essentially replaces the sequential hashing done in Scrypt's initial phase with sequential squaring in a group of unknown order. More precisely, TdScrypt is defined with respect to a group of unknown order where the trapdoor is the group order $N$. Given a parameter $n$ which basically specifies the memory requirement of the evaluation and input group element $W=: W_{0}$, first define $W_{i}:=W_{i-1}^{2}$ for $0<i<n$. Then set $S_{0}:=\mathrm{h}\left(W_{n}, 0 \cdots 0\right)$, and, for $0<i \leq n$, define $S_{i}:=\mathrm{h}\left(W_{j}, S_{i-1}\right)$ where $j:=S_{i-1} \bmod n$ and h is a hash function. Finally, $S_{n}$ is the output of TdScrypt on input $W$.

Evaluation Algorithms. The honest evaluation algorithm without knowledge of the trapdoor repeatedly squares $W$, and stores all intermediate values (i.e., all $W_{i}$ ) in memory. Then it computes all $S_{i}$ in sequence, looking up $W_{j}$ on demand. Adding up the memory consumed over all steps, we get a CMC of $\Theta\left(n^{2} \log N\right)$ because the algorithm stores $n$ group elements of size roughly $\log N$ bits while computing the values from $S_{0}$ to $S_{n}$ within $n$ steps. ${ }^{4}$

In comparison, the trapdoor evaluation algorithm first computes $W_{n}=W^{2^{n}}$ directly by first reducing $2^{n} \bmod N$. Then, it computes the $S_{i}$ sequentially while computing the $W_{j}$ on demand similarly to $W_{n}$. This has a cost of roughly $\Theta\left(n \log (N)^{2} \log (n)\right)$ and we defer the details to Section 3.2. In summary, TdSCRyPT is approximately an $\left(n^{2} \log (N), n \log (n) \log (N)^{2}\right)$-TMHF.

### 1.4 Contributions and Technical Overview

Our main contribution is a rigorous proof of the following lower bound. It bounds the CMC required by any TDScrypt evaluation algorithm without knowledge of the trapdoor.

Theorem 1 (Informal). In the random oracle and generic group model, assuming that factoring is hard, any algorithm $\mathcal{A}$ evaluating TDSCRypt has a cumulative memory complexity lower bounded by $\Omega\left(\frac{n^{2}}{\log n} \log N\right)$.

Recall that the honest evaluation algorithm (without trapdoor) has a cumulativememory complexity of $\Theta\left(n^{2} \log N\right)$, so our lower bound is a factor $1 / \log (n) \operatorname{loose}$. We do not know of any evaluation algorithm achieving $O\left(\frac{n^{2}}{\log n} \log N\right)$ and believe that the loss in tightness is an artifact of our proof.

Proof Strategy. We follow the proof of Alwen et al. [ACP $\left.{ }^{+} 17\right]$ who proved that Scrypt has a CMC of $\Omega\left(n^{2} \omega_{\mathrm{h}}\right)$-which is tight-in the random oracle model (ROM). ${ }^{5}$

[^2]Naturally, we also work in the ROM, but also need to consider Shoup's generic-group model (GGM) [Sho97]. Their proof first considers a single-challenge time-memory trade-off which is then generalized to get a multi-challenge memory complexity lower bound. We will elaborate how these two concepts are related to CMC in the following paragraphs.

We remark that our single-challenge time-memory trade-off proof is more involved and differs considerably from Alwen et al. $\left[\mathrm{ACP}^{+} 17\right]$ where this part of the proof was fairly simple. The generalization to the multi-challenge memory complexity lower bound-by far the most complicated part in $\left[\mathrm{ACP}^{+} 17\right]$ - is fortunately essentially identical. Thus, most of our proof focuses on the single-challenge case and we only sketch the multi-challenge argument while referring to $\left[\mathrm{ACP}^{+} 17\right]$ for details.

Single-challenge Time-Memory Trade-Off. Consider the following single-challenge game. Like when evaluating TdScrypt, the adversary $\mathcal{A}$ receives an input group element $W$ but does not know the group order $N$. It is then given the challenge $j \stackrel{\&}{\leftarrow}\{0, \ldots, n-1\}$ and needs to output $W_{j}=W^{2^{j}}$ as quickly as possible. Before being challenged, $\mathcal{A}$ is allowed to perform precomputation and to store the resulting advice string. Clearly, if the advice string is large enough, $\mathcal{A}$ can store $W_{0}, \ldots, W_{n-1}$ and answer every challenge instantly. Inspired by this, we are interested in a time-memory trade-off: How fast (on average) can $\mathcal{A}$ answer a challenge in relation to the size of the advice string?

To this end, we first show that if $\mathcal{A}$ answers a large fraction of the challenges quickly, it must have a lot of group elements stored, else it could factor. Our approach draws inspiration from proofs showing that repeated squaring (i.e., on input $W$ computing $W^{2^{j}}$ ) [RSW96] is sequential in generic models if factoring is hard [KLX20, Rot22, RS20]. On a high level, they first lower bound the number of queries required by algorithms oblivious to the group order. Using this lower bound, they show that if an adversarial algorithm is faster, its query behavior reveals a non-trivial factor of $N$. While the lower bound for sequential squaring trivially equals the prescribed number of iterations, figuring out a tight enough bound for the single-challenge game with respect to TdScrypt is substantially more complex. Without delving into the details, we prove a bound by analyzing the solvability of certain systems of Diophantine equations. ${ }^{6}$ We use powerful mathematical tools such as a lemma due to van der Waerden [Laz96] and a generalization of the famous distinct subset sums problem due to Erdős [Guy94, C8].

Second, again assuming hardness of factoring, we show that storing $k$ group elements roughly requires a memory of $k \log (N)$ as otherwise we could encode a random injection (the GGM's labeling function) more efficiently than information-theoretically possible. Our proof is inspired by Corrigan-Gibbs and Kogan [CK18] who analyze how helpful preprocessing is for computing discrete logarithms by using an incompressibility argument [DTT10]. In contrast to their work, our argument is more involved. On the one hand, the single-challenge game is more complex than the discrete logarithm problem. On the other, the group order is unknown which complicates bookkeeping in the encoding routine.

Combining the two results above, we get that if an adversary answers challenges quickly on average, then the advice string must be large - assuming factoring is hard.

[^3]Multi-challenge Time-memory Trade-Off. When an algorithm evaluates TdScrypt and outputs $S_{n}$, it almost surely must have computed $S_{1}, \ldots, S_{n}$ in sequence because h is a modeled as a random oracle. Computing $S_{i}$ given $S_{i-1}$ requires $W_{j}$ by definition where is $j:=S_{i-1} \bmod n$ (almost) uniformly random because, again, h is a random oracle. It follows that evaluating TdScrypt requires playing $n$ single-challenge games in sequence. So, a lower bound on the memory complexity of solving multiple challenges implies a lower bound on the CMC of TdScrypt.

### 1.5 Open Problems

First, TdScrypt has a CMC of $O\left(n^{2} \log N\right)$ whereas the lower bound is only $\Omega\left(\frac{n^{2}}{\log n} \log N\right)$. Ideally, this bound should be tight. Looking ahead, one possible way of achieving this is strengthening Lemma 2, a solvability criterion for certain systems of Diophantine equations. Essentially, one would need to show $\operatorname{rank}(A) \geq \ell /(c t)$ for some constant $c$.

Second, TdScrypt's drop in CMC when using the trapdoor is only due to the much lower memory requirement of the trapdoor evaluation. The computation actually increases from $n$ to $n \log (N)$, as in the 2 nd phase of the evaluation (where we compute the $S_{i}$ 's) the normal evaluation just makes $n$ group operations (modular multiplications), while the trapdoor evaluation needs to do $n$ exponentiations. An open problem is constructing a TMHF where the trapdoor evaluation not only improves CMC, but strictly improves on memory usage and computation individually (or at least improves on one of them without decreasing the other).

Finally, coming up with a TMHF that fulfills different notions of memory hardness would be interesting. For example, a data-independent TMHF or a TMHF that ensures high sustained space complexity.

## 2 Preliminaries

### 2.1 Notation

$\mathbb{N}, \mathbb{Z}, \mathbb{Z}_{N}, \mathbb{Q}$, and $\mathbb{R}$ are the sets of natural numbers including 0 , integers, integers modulo $N$, rational, and real numbers, respectively. For these sets, the superscript ${ }^{+}$ denotes the strictly positive subset (e.g., $\mathbb{N}^{+}=\{1,2, \ldots\}$ ). $[a, b]$ denotes the set $\{a, \ldots, b\}$, $[a, b)=\{a, \ldots, b-1\}$, and $[n]$ is a shorthand for $[1, n]$. Furthermore, $\operatorname{Inj}(A, B)$ is the set of all injections from set $A$ to set $B$ and $x \stackrel{\$}{\stackrel{\$}{\leftarrow} X \text { samples an element from the set } X}$ uniformly at random. Vectors are written in boldface (e.g., $\boldsymbol{x}$ ) and matrices in upper case. $|x|$ denotes the absolute value or the length of $x$ depending on whether $x$ is a number or a list, vector, etc., and $\|x\|$ is the number of bits required to encode $x$. Algorithms are usually typeset sans-serif (e.g., Alg), $x:=y$ or $x:=\operatorname{Alg}(\ldots)$ denote assignment or the output of a deterministic algorithm, and $a \leftarrow \operatorname{Alg}(\ldots)$ the output of a probabilistic one.
$\lambda$ always denotes the security parameter and $\operatorname{negl}(\lambda)($ resp. poly $(\lambda))$ are the set of all functions that are negligible (resp. polynomial) in $\lambda$. Furthermore, we use standard Big O notation such as $O, \Omega, \Theta$, and $\omega$. When working with groups, group elements are written upper case, and their exponents with respect to the group generator lower case. Lastly, log denotes the binary logarithm.

### 2.2 Algebraic Setting

To allow for a trapdoor, we require some algebraic structure: the group of quadratic residues with respect to RSA moduli. As a consequence, the memory-hardness of TdScrypt is based on factoring assumptions. In the following, we define modulus generation, related factoring assumptions, and the group of quadratic residues.

Definition 1 (Safe-prime Generator). Let GenSP be an algorithm that samples two distinct safe primes of the same bit length uniformly at random. More precisely, define

$$
\left(p^{\prime}, q^{\prime}\right) \leftarrow \operatorname{GenSP}(\lambda)
$$

and let $p:=\left(p^{\prime}-1\right) / 2$ and $q:=\left(q^{\prime}-1\right) / 2$. It holds that $p^{\prime}, p, q^{\prime}$ and $q$ are prime, $p \neq q$, and $\|p\|=\|q\|=k(\lambda)$ with $k(\lambda) \in \operatorname{poly}(\lambda)$.

In terms of notation, let $N:=p q$ and $N^{\prime}:=p^{\prime} q^{\prime}$ in the rest of the paper. ${ }^{7}$. Next, we give two hardness assumptions stating that factoring $N^{\prime}$ as well as $N$ is hard.

Definition 2 (Factoring $N^{\prime}$ ). The game $\operatorname{Fac}_{\text {GenSP } \mathcal{A}}^{\prime}(\lambda)$ is defined with respect to GenSP and a probabilistic polynomial-time (PPT) adversary $\mathcal{A}$. On input of a security parameter $\lambda$ it runs $\left(p^{\prime}, q^{\prime}\right) \leftarrow \operatorname{GenSP}(\lambda)$, sets $N^{\prime}:=p^{\prime} q^{\prime}$, and invokes the adversary, yielding $\left(p^{*}, q^{*}\right) \leftarrow$ $\mathcal{A}\left(N^{\prime}\right)$. The game returns 1 if $\left\{p^{*}, q^{*}\right\}=\left\{p^{\prime}, q^{\prime}\right\}$ and 0 otherwise.

We say that factoring $N^{\prime}$ is hard if, for every PPT algorithm $\mathcal{A}$, the advantage is negligible, i.e.,

$$
\operatorname{Adv}_{\operatorname{GenSP}, \mathcal{A}}^{\mathrm{Fac}^{\prime}}(\lambda):=\operatorname{Pr}\left[1 \leftarrow \operatorname{Fac}_{\operatorname{GenSP}, \mathcal{A}}^{\prime}(\lambda)\right] \in \operatorname{negl}(\lambda)
$$

where the probability is taken over the randomness of GenSP and $\mathcal{A}$.
Definition 3 (Factoring $N$ ). The game $\operatorname{Fac}_{G e n S P, \mathcal{A}}(\lambda)$ is defined similarly to $\operatorname{Fac}_{G e n S P, \mathcal{A}}^{\prime}(\lambda)$ but with $N^{\prime}$ replaced by $N$. So the game invokes $\left(p^{*}, q^{*}\right) \leftarrow \mathcal{A}(N)$ and returns 1 if and only if $\left\{p^{*}, q^{*}\right\}=\{p, q\}$. We say that factoring $N$ is hard if, for every PPT algorithm $\mathcal{A}$,

$$
\operatorname{Adv}_{\operatorname{GenSP}, \mathcal{A}}^{\mathrm{Fac}}(\lambda):=\operatorname{Pr}\left[1 \leftarrow \operatorname{Fac}_{G e n S P}, \mathcal{A}(\lambda)\right] \in \operatorname{negl}(\lambda)
$$

where the probability is taken over the randomness of GenSP and $\mathcal{A}$.
Let us briefly analyze both assumptions. It is conjectured that the density of safe primes $p^{\prime}$ in the interval $\left[2^{k-1}, 2^{k}\right]$ is of order $1 / k^{2}$ (e.g., [vzGS13]). Note that the map $p^{\prime} \mapsto\left(p^{\prime}-1\right) / 2$ is injective, so this statement implies the same regarding the density of $\left(p^{\prime}-1\right) / 2=p$ in $\left[2^{k-2}, 2^{k-1}\right]$. Thus, assuming the conjecture holds, the standard factoring assumption (i.e., factoring the product of two uniformly sampled arbitrary $k$-bit primes is hard) implies that factoring $N^{\prime}$ as well as $N$ is hard. For simplicity, we say that factoring is hard with respect to GenSP if factoring $N$ as well as factoring $N^{\prime}$ are hard with respect to GenSP.

Equipped with these definitions, we finally define the group of quadratic residues modulo $N^{\prime}$.

[^4]Definition 4 (Group of Quadratic Residues). For $\left(p^{\prime}, q^{\prime}\right) \leftarrow \operatorname{GenSP}(\lambda)$, let $p:=\left(p^{\prime}-1\right) / 2$, $q:=\left(q^{\prime}-1\right) / 2, N^{\prime}:=p^{\prime} q^{\prime}$, and $N:=p q$. The group of quadratic residues modulo $N^{\prime}$ is a subgroup of $\mathbb{Z}_{N^{\prime}}^{*}$ defined as $\mathbb{Q} \mathbb{R}_{N^{\prime}}:=\left\{X^{2} \mid X \in \mathbb{Z}_{N^{\prime}}^{*}\right\}$.

Since $p^{\prime}$ and $q^{\prime}$ are safe primes (and thereby Blum integers), $\mathbb{Q} \mathbb{R}_{N^{\prime}}$ is a cyclic group of order $\left|\mathbb{Q R}_{N^{\prime}}\right|=N$ and a uniformly sampled element $X \leftarrow \mathbb{Q R}_{N^{\prime}}$ is a generator with overwhelming probability $\varphi(N) / N=1-\frac{q+p-1}{N} \in(1-\operatorname{negl}(\lambda))$ where $\varphi$ denotes Euler's totient function.

### 2.3 Generic Group Model

In the proof of our main result we model the cyclic group $\mathbb{Q R}_{N^{\prime}}=\langle g\rangle$ of order $N:=p q$ as a generic group in the style of Shoup's generic group model (GGM) [Sho97]. In the GGM algorithms do not get direct access to the group (that we may identify with the group $\left(\mathbb{Z}_{N},+\right)$ by using the isomorphism $\left.\mathbb{Z}_{N} \rightarrow \mathbb{Q} \mathbb{R}_{N^{\prime}} ; z \mapsto g^{z}\right)$. Instead, access to group elements is provided via abstract labels $\sigma \in \mathcal{L}$ where $\mathcal{L}:=\{0,1\}^{\omega_{\mathcal{L}}}\left(\right.$ with $\left.\omega_{\mathcal{L}} \geq\lceil\log (N)\rceil\right)$ and group operations are performed using an oracle $\mathcal{G}$ that allows for multiplication, division, and inversion.

More precisely, let $\operatorname{Inj}\left(\mathbb{Z}_{N}, \mathcal{L}\right)$ be the set of all injections from $\mathbb{Z}_{N}$ to $\mathcal{L}$ and let $\Sigma \stackrel{\$}{\leftarrow} \operatorname{Inj}\left(\mathbb{Z}_{N}, \mathcal{L}\right)$ be one such injection chosen uniformly at random. $\Sigma$ is called the labeling function, and it defines the oracle $\mathcal{G}:\{+,-$, inv $\} \times \mathcal{L} \times \mathcal{L} \rightarrow \mathcal{L} \cup\{\perp\}$. $\mathcal{G}$ answers a query $\left(\circ, \sigma_{1}, \sigma_{2}\right)$ as follows. If $\sigma_{1}$ and $\sigma_{2}$ are not in the image of $\Sigma$, it returns $\perp$. Otherwise, if $\circ \in\{+,-\}$, it returns $\Sigma\left(\Sigma^{-1}\left(\sigma_{1}\right) \circ \Sigma^{-1}\left(\sigma_{2}\right) \bmod N\right)$, and if $\circ=$ inv, it returns $\Sigma\left(-\Sigma^{-1}\left(\sigma_{1}\right) \bmod N\right)$.

Abusing notation, we will almost always write $\Sigma$ instead of $\mathcal{G}$ when talking about a concrete instantiation of the generic group oracle.

### 2.4 Machine Model and Complexity Measure

Parallel Oracle Model. To define a memory complexity measure, we first require a machine model. Prior work uses the parallel random oracle model [AS15, ACP $\left.{ }^{+} 17\right]$ which allows algorithms to perform an unlimited number of random oracle queries in parallel. Similarly, we define the parallel oracle model which additionally allows for generic group queries. That is, a polynomial-time (PT) algorithm $\mathcal{A}$ has access to two oracles: $\mathcal{G}$, a group-operation oracle, and h , a random oracle. We will sometimes explicitly state the oracles in superscript, i.e., $\mathcal{A}^{\mathcal{G}, \mathrm{h}}$, and omit them when they are clear from context.

The random oracle $\mathrm{h}:\{0,1\}^{*} \rightarrow\{0,1\}^{\omega_{h}}$ maps inputs to bit strings of length $\omega_{\mathrm{h}}$ and we require that $\omega_{\mathrm{h}} \in \Theta(\log N)$. To ensure that the set of random oracles is finite, we assume some sufficiently large, finite bound $*$ on the inputs.

Algorithm $\mathcal{A}$ 's execution proceeds in rounds starting with round 1 . Within each round, $\mathcal{A}$ performs local computation and submits oracle queries at the end of it. Then, at the beginning of the next round, $\mathcal{A}$ receives the response to its queries. Formally, states capture $\mathcal{A}$ 's progress throughout the rounds and its queries to both oracles. At the end of round $i, \mathcal{A}$ produces an output state $\overline{\operatorname{st}}_{i}:=\left(\tau_{i}, \mathbf{q r s}{\underset{i}{\mathcal{G}}}_{i}, \mathbf{q r s}_{i}^{\mathrm{h}}\right)$ where $\tau$ is a bit string, and $\mathbf{q r s}^{\mathcal{G}}$ and $\mathbf{q r s}^{\mathrm{h}}$ are vectors containing queries to $\mathcal{G}$ and h , respectively. Consequently, in round $i+1$, it receives the input state $\mathbf{s t}_{i}:=\left(\tau_{i}, \operatorname{res}_{i}^{\mathcal{G}}, \boldsymbol{r e s}_{i}^{\mathrm{h}}\right)$ where $\mathbf{r e s}_{i}^{\mathcal{G}}$ and $\mathbf{r e s}_{i}^{\mathrm{h}}$ are vectors containing the results of the queries $\mathbf{q r s}_{i}^{\mathcal{G}}$ and $\mathbf{q r s}{ }_{i}^{h}$.

We will only consider deterministic algorithms in the interest of keeping proofs readable and concise. This is essentially without loss of generality in our setting since we only care about algorithms that correctly evaluate TDSCRYPT with sufficiently high, i.e., noticeable, success probability. Such algorithms can be derandomized by trying out a few choices of the randomness, checking each randomness on a polynomial number of random inputs, and then fix the randomness on which the algorithm performed best. This will with overwhelming probability result in a deterministic algorithm whose success probability is close to the randomized one.

Complexity Measure. To evaluate the memory complexity of an algorithm, we use a notion called cumulative memory complexity (CMC) [AS15]. Essentially, CMC is the sum of an algorithm's memory consumption at every point in time, i.e., the area under the memory usage curve.

More formally, consider an algorithm $\mathcal{A}$ running in the parallel oracle model on some input $x$. Its execution results in a sequence of input states $\boldsymbol{s t}_{i}=\left(\tau_{i}, \boldsymbol{r e s}_{i}^{\mathcal{G}}, \boldsymbol{r e s}_{i}^{\mathrm{h}}\right)$. Then, its CMC is given by

$$
\mathrm{cc}_{\mathrm{mem}}\left(\mathcal{A}^{\mathcal{G}, \mathrm{h}}(x)\right):=\sum_{i}\left\|\mathrm{st}_{i}\right\|
$$

where the bit length of the input state st ${ }_{i}$ is defined as $\left\|\mathrm{st}_{i}\right\|:=\left\|\tau_{i}\right\|+\left|\operatorname{res}_{i}^{\mathcal{G}}\right| \omega_{\mathcal{L}}+\left|\operatorname{res}_{i}^{\mathrm{h}}\right| \omega_{\mathrm{h}}$. Without loss of generality, an algorithm submits at least one query per round, so $\left\|s t_{i}\right\| \geq$ $\min \left\{\omega_{\mathcal{L}}, \omega_{\mathrm{h}}\right\} \in \Theta(\log N)$ by definition.

Preprocessing Algorithms. Instead of working with a single deterministic PT algorithm $\mathcal{A}$, we often view it as a pair of deterministic PT algorithms $\mathcal{A}=\left(\mathcal{A}_{0}, \mathcal{A}_{1}\right)$. Intuitively, $\mathcal{A}_{0}$ performs preprocessing and outputs some advice which $\mathcal{A}_{1}$ receives as input. It follows that we do not charge $\mathcal{A}_{0}$ for any memory usage and that we do not count its number of rounds (but emphasize that, unlike in other works, $\mathcal{A}_{0}$ 's computation is not completely unbounded as we cannot allow it to factor $N$ with noticeable probability).

In the terms of the parallel oracle model, given public parameters pp and online input $x$, to evaluate $\mathcal{A}^{\mathcal{G}, \mathrm{h}}(\mathrm{pp}, x)$, first execute $\mathcal{A}_{0}^{\mathcal{G}, \mathrm{h}}(\mathrm{pp})$ resulting in an input state st ${ }_{0}$. Then, starting in round 1 , run $\mathcal{A}_{1}^{\mathcal{G}, \mathrm{h}}\left(\operatorname{st}_{0}, x\right)$ with the online input $x$ yielding the output of $\mathcal{A}$. The CMC of $\mathcal{A}$ is defined as $\mathrm{cc}_{\text {mem }}\left(\mathcal{A}^{\mathcal{G}, \mathrm{h}}(\mathrm{pp}, x)\right):=\operatorname{cc}_{\text {mem }}\left(\mathcal{A}_{1}^{\mathcal{G}, \mathrm{h}}\left(\operatorname{st}_{0}, x\right)\right)$.

## 3 A Trapdoor Memory-Hard Function from Factoring

In this section we define the syntax for trapdoor memory-hard functions (Section 3.1) and then describe TdScrypt, the instantiation of Diodon [BP17] we will use throughout the paper (Section 3.2).

### 3.1 Trapdoor Memory-Hard Functions

A trapdoor memory-hard function (TMHF) is defined by a triple of polynomial-time algorithms (Setup, Eval, TDEval) as follows (with the security parameter $\lambda$ left implicit).

- Setup ()$\rightarrow(\mathrm{pp}, \mathrm{td})$. The probabilistic setup algorithm samples public parameters pp and corresponding trapdoor td. The public parameters implicitly determine the domain $\operatorname{Dom}(\mathrm{pp})$ and range Ran(pp) of the TMHF.
- $\operatorname{Eval}(\mathrm{pp}, w)=: y$. The deterministic evaluation algorithm takes public parameters pp and $w \in \operatorname{Dom}(\mathrm{pp})$ as inputs and returns a $y \in \operatorname{Ran}(\mathrm{pp})$.
- TDEval $(\mathrm{pp}, \mathrm{td}, w)=: y$. The deterministic trapdoor evaluation algorithm takes public parameters pp , trapdoor td , and $w \in \operatorname{Dom}(\mathrm{pp})$ as inputs and returns $y \in \operatorname{Ran}(\mathrm{pp})$.

We require correctness, i.e., for all $(\mathrm{pp}, \mathrm{td}) \leftarrow \operatorname{Setup}()$ and all $w \in \operatorname{Dom}(\mathrm{pp})$, it holds that

$$
\operatorname{Eval}(\mathrm{pp}, w)=\operatorname{TDEval}(\mathrm{pp}, \mathrm{td}, w)
$$

To quantify the quality of a TMHF we have to analyze the cumulative memory required to evaluate it with and without access to the trapdoor. Accordingly, if the required CMC is given by functions $c_{\text {hon }}$ and $c_{\mathrm{td}}$, i.e.,

$$
c_{\mathrm{hon}}(\mathrm{pp})=\mathrm{cc}_{\text {mem }}(\operatorname{Eval}(\mathrm{pp}, w)) \text { and } c_{\mathrm{td}}(\mathrm{pp})=\mathrm{cc}_{\text {mem }}(\operatorname{TDEval}(\mathrm{pp}, \mathrm{td}, w))
$$

the TMHF is referred to as a ( $c_{\text {hon }}, c_{\mathrm{td}}$ )-TMHF.
Naturally, trapdoor evaluations must have a lower CMC than the standard evaluation algorithm. That is, for all $(\mathrm{pp}, \mathrm{td}) \leftarrow$ Setup, there exists some $0<\Delta(\mathrm{pp})<1$ (ideally $\Delta(\mathrm{pp}) \in o(1))$, such that, for all inputs $w \in \operatorname{Dom}(\mathrm{pp})$, we have

$$
c_{\mathrm{td}}(\mathrm{pp})<\Delta(\mathrm{pp}) \cdot c_{\mathrm{hon}}(\mathrm{pp}) .
$$

Moreover, we want our function to be a good MHF when ignoring the trapdoor. This means the CMC of $\operatorname{Eval}(\mathrm{pp}, w)$ should be high by construction, but no adversarial evaluation algorithm should exist that can evaluate the function with much lower CMC. This must hold even when the other adversarial algorithm is allowed to make many parallel queries to the oracles and when it is given some advice that was computed (by any polynomial time preprocessing) dependent on pp (but of course not the input $w$ or the trapdoor).

### 3.2 Description of TDSCRYPT

We will analyze TdScrypt which can be viewed as a concrete instantiation of Diodon [BP17]. Specifically, using notation from [BP17], we set $M:=n, L:=n$, and $\eta=1$. Furthermore, TDSCRYPT is defined over the group of quadratic residues $\mathbb{Q} \mathbb{R}_{N^{\prime}}$ instead of $\mathbb{Z}_{N^{\prime}}$ due to technicalities.

Construction. TdScrypt is defined with respect to integer $n \in \mathbb{N}$ corresponding to the number of iterated steps (i.e., repeated squarings). It relies on computations in the group $\mathbb{Q} \mathbb{R}_{N^{\prime}}$ of quadratic residues modulo $N^{\prime}$ as defined in Definition 4. Accordingly, its public parameters pp consist of an RSA modulus $N^{\prime}:=p^{\prime} q^{\prime}$, where $p^{\prime}$ and $q^{\prime}$ are safe primes generated using safe prime generator GenSP (see Definition 1), a hash function h: $\{0,1\}^{*} \rightarrow$ $\{0,1\}^{\omega_{\mathrm{h}}}$, and $n$. The corresponding trapdoor is $\mathrm{td}:=N=\left(p^{\prime}-1\right)\left(q^{\prime}-1\right) / 4=\left|\mathbb{Q} \mathbb{R}_{N^{\prime}}\right|$. TdSCRYPT's formal description is in Figure 1.

| Setup() | $\operatorname{Eval}(\mathrm{pp}, W)$ | TDEval (pp, td, $W$ ) |
| :---: | :---: | :---: |
| $00\left(p^{\prime}, q^{\prime}\right) \leftarrow \operatorname{GenSP}()$ | $07 W_{0}:=W$ | $16 W_{0}:=W$ |
| $01 N^{\prime}:=p^{\prime} q^{\prime}$ | 08 For $i:=1, \ldots, n-1$ : | $17 \mathrm{~m}:=2^{n} \bmod N$ |
| $02 N:=\left(p^{\prime}-1\right)\left(q^{\prime}-1\right) / 4$ | $09 \quad W_{i}:=W_{i-1}^{2}$ | $18 W_{n}:=W_{0}^{m}$ |
| 03 pick h | 10 Store $W_{i}$ | $19 S_{0}:=\mathrm{h}\left(W_{n}, 0^{\omega_{\mathrm{h}}}\right)$ |
| $04 \mathrm{pp}:=\left(N^{\prime}, \mathrm{h}, n\right)$ | $11 S_{0}:=\mathrm{h}\left(W_{n}, 0^{\omega_{\mathrm{h}}}\right)$ | 20 For $i:=1, \ldots, n$ : |
| $05 \mathrm{td}:=N$ | 12 For $i:=1, \ldots, n$ : | $21 \quad j_{i}:=S_{i-1} \bmod n$ |
| 06 Return (pp,td) | $13 \quad j_{i}:=S_{i-1} \bmod n$ | $22 \quad m_{i}:=2^{j_{i}} \bmod N$ |
|  | $14 \quad S_{i}:=\mathrm{h}\left(W_{j_{i}}, S_{i-1}\right)$ | $23 W_{j_{i}}:=W_{0}^{m_{i}}$ |
|  | 15 Return $S_{n}$ | $24 \quad S_{i}:=\mathrm{h}\left(W_{j_{i}}, S_{i-1}\right)$ |
|  |  | 25 Return $S_{n}$ |

Figure 1: Trapdoor memory-hard function TDSCRyPT ${ }_{n}^{h}$ defined with respect to number of steps $n \in \mathbb{N}$. The TMHF uses a hash function $\mathrm{h}:\{0,1\}^{*} \rightarrow\{0,1\}^{\omega_{\mathrm{h}}}$, domain $\operatorname{Dom}(\mathrm{pp})=$ $\mathbb{Q} \mathbb{R}_{N^{\prime}}$ and range $\{0,1\}^{\omega_{\mathrm{h}}}$. The exponentiations $W_{i}^{2}, W_{0}^{m}$, and $W_{0}^{m_{i}}$ are computed in $\mathbb{Q} \mathbb{R}_{N^{\prime}}$.

In the following, we will assume that the time of evaluating h approximately matches the time of evaluating a group operation, i.e., computing a multiplication in modulo $N^{\prime}$. This could for example be implemented by setting $\mathrm{h}(s):=\mathrm{h}^{\prime \prime}\left(\mathrm{h}^{\prime}(s) \cdot \mathrm{h}^{\prime}(s) \bmod N^{\prime}\right.$ ), where $h^{\prime}:\{0,1\}^{*} \rightarrow \mathbb{Q}_{N^{\prime}}$ and $h^{\prime \prime}: \mathbb{Q R}_{N^{\prime}} \rightarrow\{0,1\}^{\omega_{h}}$ are cryptographic hash functions.

To evaluate the TMHF on input $W \in \operatorname{Dom}(\mathrm{pp}):=\mathbb{Q R}_{N^{\prime}}$ algorithm Eval sets $W_{0}:=W$ and computes and stores the group elements $W_{i}:=W_{i-1}^{2}=W^{2^{i}}$ for $i \in[n-1]$. Then it sets $S_{0}:=\mathrm{h}\left(W_{n}, 0^{\omega_{\mathrm{h}}}\right)$, and, for $i:=[n]$, computes $S_{i}:=\mathrm{h}\left(W_{j_{i}}, S_{i-1}\right)$ where the index $j_{i}:=S_{i-1} \bmod n$. The output of Eval is $S_{n} \in \operatorname{Ran}(\mathrm{pp}):=\{0,1\}^{\omega_{\mathrm{h}}}$.

TDEval computes the values $S_{i}$ without storing the group elements $W_{j_{i}}$, but instead recomputes them efficiently using its knowledge of $N=\left|\mathbb{Q} \mathbb{R}_{N^{\prime}}\right|$. It sets $W_{0}:=W$, computes $W_{n}:=W_{0}^{m}$, and $S_{0}:=\mathrm{h}\left(W_{n}, 0^{\omega_{\mathrm{h}}}\right)$, where $m:=2^{n} \bmod N$. Then for $i \in[n]$ it computes, in order, $j_{i}:=S_{i-1} \bmod n, m_{i}:=2^{j_{i}} \bmod N, W_{i}=W_{0}^{m_{i}}$, and $S_{i}:=\mathrm{h}\left(W_{j_{i}}, S_{i-1}\right)$. Its output is $S_{n}$.

Cumulative Memory Complexity. Recall that the CMC is the memory usage summed over all steps. In the following analysis of TdScrypt we will give a brief estimate of Eval and TDEval's CMC where we define a "step" to constitute the time taken by one group operation. For details on this choice and an in-depth CMC analysis we refer to Appendix A.

For Eval, note that its CMC is dominated by the second loop (Lines 12-14). During the loop, Eval keeps $n$ group elements $W_{0}, \ldots, W_{n}$ in memory, amounting to $\Theta\left(n \log \left(N^{\prime}\right)\right)$ bits. The loop requires $\Theta(n)$ group operations in total as every evaluation of $h$ takes the time of one group operation by definition. We thus obtain

$$
\mathrm{cc}_{\mathrm{mem}}(\mathrm{Eval}(\mathrm{pp}, x)) \in \Theta\left(n \log \left(N^{\prime}\right) \cdot n\right)=\Theta\left(n^{2} \log \left(N^{\prime}\right)\right) .
$$

TDEval's CMC is also dominated by the loop (Lines 20-24). In contrast to Eval, the memory usage during the loop is low because it does not store $n$ group elements. However, every loop iteration requires more steps since TDEval reduces $n$-bit integers of the form $2^{j}$
with $j \in[0, n)$ modulo $N$ (Line 22) and performs exponentiations in $\mathbb{Q R}_{N^{\prime}}$ (Line 23). In general, the former operations requires $\Theta\left(n \log \left(N^{\prime}\right)\right)$ group operations which can be reduced to $\log (n)$ group operations by using a lookup table of size $\Theta\left(\log (n) \log \left(N^{\prime}\right)\right)$. The exponentation requires $\Theta\left(\log N^{\prime}\right)$ group operations using square-and-multiply. So during the loop, TDEval stores the lookup table using $\Theta\left(\log (n) \log \left(N^{\prime}\right)\right)$ bits, and every iteration takes $\Theta\left(\log (n)+\log \left(N^{\prime}\right)\right)$ group operations, resulting in

$$
\begin{aligned}
\operatorname{cc}_{\text {mem }}(\text { TDEval }(\mathrm{pp}, \mathrm{td}, x)) & \in \Theta\left(\log (n) \log \left(N^{\prime}\right) \cdot n\left(\log \left(N^{\prime}\right)+\log (n)\right)\right) \\
& =\Theta\left(n \log \left(N^{\prime}\right)^{2} \log (n)\right) .
\end{aligned}
$$

Summing up, TdScrypt is a $\left(n^{2} \log \left(N^{\prime}\right), n \log (n) \log \left(N^{\prime}\right)^{2}\right)$-TMHF.

## 4 Overview of the Lower Bound Proof

In the previous section, we established that Eval has a CMC of $\Theta\left(n^{2} \log N\right)$ and TDEval one of $\Theta(n \log n \log N)$. These analyses only hold for the specific algorithms stated in Figure 1. However, an adversary (without knowledge of the trapdoor) need not follow Eval. The rest of this paper is devoted to showing that no evaluation algorithm is meaningfully faster than Eval.

Theorem 1 (Restated). Let $n \in \operatorname{poly}(\lambda)$ with $n \geq 8$ and let $\mathcal{A}$ be a deterministic parallel oracle machine that evaluates TDSCRYPT correctly with probability $\chi(\lambda)$ over the choice of the parameters and input $W$. Then, assuming that factoring is hard, in the GGM and $R O M$ with probability at least $\chi(\lambda)-\epsilon(\lambda)$,

$$
\mathrm{cc}_{\mathrm{mem}}\left(\mathcal{A}^{\mathrm{pp}}(W)\right) \in \Omega\left(\frac{n^{2}}{\log n} \log N\right)
$$

where pp are the public parameters, $\epsilon(\lambda) \in \operatorname{negl}(\lambda)$, and the probability is taken over the choice of parameters.

So we show a lower bound on the CMC of evaluating TdScrypt that almost matches Eval and is only a factor of $1 / \log n$ loose. Note that $\log n \in O(\log \lambda)$ by definition of $n$, so the loss in tightness is small asymptotically.

To prove our lower bound, we assume that factoring is hard and work in generic models. This is in line with lower bounds on the memory complexity of (non-trapdoor) memory-hard functions which are situated in the ROM. The rest of the paper is divided into two sections similar to Alwen et al. $\left[\mathrm{ACP}^{+} 17\right]$

Single-challenge Time-Memory Trade-Off (Section 5). Before considering the whole TdScrypt execution, we focus on the hardness of computing $W_{j}:=W^{2^{j}}$ for $j \stackrel{\&}{\leftarrow}[n]$ given $M$ bits of precomputed advice. This is closely related to the second phase of TdSCRYPT which, for all $i \in[n]$, requires computing $S_{i}:=\mathrm{h}\left(W^{2^{j_{i}}}, S_{i-1}\right)$ given challenge $j_{i}:=S_{i-1} \bmod n$. Note that $j_{i}$ is chosen (almost) uniformly from $[0, n)$ when modeling $h$ as a random oracle. Here, the precomputed advice can be thought of as the state of the algorithm before learning challenge $j_{i}$. Our goal is to show that if the advice is short, computing $W_{j}$ takes a long time on average. Stated differently, if $\mathcal{A}$ computes $W_{j}$ quickly on average, the advice must be large.

Multi-challenge Memory Complexity Lower-Bound (Section 6). We abstract the whole evaluation of TDSCRYPT as a multi-challenge game to get the lower bound on $\mathrm{cc}_{\text {mem }}(\mathcal{A})$. To this end, the single-challenge trade-off is applied to every challenge, i.e., to the moment in time before it is known and every point in time before the preceding challenge has been issued. This idea closely follows the corresponding proof in Alwen et al. [ $\left.\mathrm{ACP}^{+} 17\right]$, and we refer interested readers to the paper for details.

## 5 Single-challenge Time-Memory Trade-Off

We consider a pair of deterministic parallel oracle machines $\mathcal{A}=\left(\mathcal{A}_{0}, \mathcal{A}_{1}\right)$ where $\mathcal{A}_{0}$ receives the input $W$ and performs preprocessing, resulting in an advice string. In the context of TdScrypt, this is best thought of as the computation performed by the (adversarial) evaluation algorithm up to learning challenge $j . \mathcal{A}_{1}$, on input of challenge $j \stackrel{\$}{\leftarrow}[0, n)$, uses the advice string to query $W^{2^{j}}$ to $h$ in as few rounds as possible. For now, we explicitly pass the challenge $j$ as an input. Our goal is to relate the advice string size to the number of rounds required by $\mathcal{A}_{1}$ on average across all possible choices of $j$.

To formalize the preceding description, we first specify the set of parameters which determine an execution of TDScrypt and then the game sketched above.

Definition 5 (Parameters). For security parameter $\lambda \in \mathbb{N}$, let params $(\lambda)$ be the set of all possible parameters. It contains all quadruples $(N, \Sigma, \mathrm{~h}, w)$ where $\left(p^{\prime}, q^{\prime}\right) \leftarrow \operatorname{GenSP}(\lambda)$ with $p:=\left(p^{\prime}-1\right) / 2$ and $q:=\left(q^{\prime}-1\right) / 2$ defines the group order $N:=p q$, the labeling function $\Sigma$ is an injection $\operatorname{Inj}\left(\mathbb{Z}_{N}, \mathcal{L}\right), \mathrm{h}$ is a random oracle, and the input $w$ is non-zero, i.e., $w \in \mathbb{Z}_{N} \backslash\{0\}$. Here, $w$ is the discrete logarithm of $W$, i.e., $\Sigma^{-1}(W)=w$ in the GGM. Furthermore, recall that we require $\omega_{\mathcal{L}}$, the length of labels in $\mathcal{L}$, and $\omega_{\mathrm{h}}$, the output length of h, to be larger than $\log N$ and of order $\log N$ respectively.

To avoid clutter, we sometimes leave $\lambda$ implicit and write params instead of params $(\lambda)$.
Definition 6 (Single-challenge Game). Let $(N, \Sigma, \mathrm{~h}, w) \in$ params and $\mathcal{A}=\left(\mathcal{A}_{0}, \mathcal{A}_{1}\right)$ be deterministic parallel oracle machines . To the challenge $j \in[0, n)$ and state st ${ }_{0}:=$ $\mathcal{A}_{0}^{\Sigma, \mathrm{h}}(\Sigma(1), \Sigma(w))$ we associate

$$
\operatorname{TimeSC}_{\mathcal{A}_{1}}^{\Sigma, \mathrm{h}}\left(\mathrm{st}_{0}, j\right)= \begin{cases}\min \left\{i: \Sigma\left(w^{2^{j}}\right) \in \mathbf{q r s}_{i}^{\mathrm{h}}\right\}, & \text { if the minimum exists } \\ \infty, & \text { otherwise }\end{cases}
$$

So TimeSC is the earliest round $i \in \mathbb{N}^{+}$in which $\mathcal{A}_{1}^{\mathcal{G}, \mathrm{h}}\left(\mathrm{st}_{0}, j\right)$ queries $\Sigma\left(w^{2 j}\right)$ to h or $\infty$ if it never does.

Next, we answer the following question: How small can TimeSC be relative to the size of $\mathrm{st}_{0}$ ? That is, we are interested in a time-memory trade-off when playing TimeSC.

A preliminary observation is that if $\left\|s t_{0}\right\| \geq n \omega_{\mathcal{L}}$, then there exists a simple strategy $\mathcal{S S}$ achieving TimeSC $_{\mathcal{S S}}=1$. Indeed, $\mathcal{S S}$ simply stores the answers to every possible challenge in $\mathrm{st}_{0}$. Generalizing to any $\left\|\mathrm{st}_{0}\right\|=M$, we get the following strategy $\mathcal{S S}: \mathcal{S S}_{0}$ encodes $\rho:=\left\lfloor M / \omega_{\mathcal{L}}\right\rfloor$ group elements in st ${ }_{0}$ where the elements are of the form $W^{2^{i}}$ with the $i$ equidistantly spaced across $0, \ldots, n-1$. $\mathcal{S S}_{1}$, on input $j$, picks the maximum $i \leq j$ such that $W^{2^{i}}$ is stored in st ${ }_{0}$ and, if necessary, repeatedly squares it until reaching $W^{2^{j}}$. So TimeSC $\mathcal{S S}>n /(2 \rho)$ for at least half of the challenges.

We want to show that if factoring is hard, no algorithm $\mathcal{A}$ can meaningfully beat $\mathcal{S S}$ for most choices of parameters. That is, there exists a partition of params into sets good and bad such that bad contains a negligible fraction of params and $\mathcal{A}$ can only beat the strategy when the chosen parameters are in bad. Note that we cannot give guarantees for all possible parameter choices since, e.g., sometimes $N$ might be easy to factor. This leads us to Lemma 1 that characterizes the set bad.

Lemma 1 (Single-challenge Trade-Off). For every pair of deterministic parallel oracle machines $\mathcal{A}=\left(\mathcal{A}_{0}, \mathcal{A}_{1}\right)$, all $c \in \mathbb{N}$ with security parameter $\lambda$ large enough, and all $n, M, Q \in \operatorname{poly}(\lambda)$ with $n \geq 8$ and subsets $\operatorname{bad} \subseteq \operatorname{params}(\lambda)$, if, for every $(N, \Sigma, \mathrm{~h}, w) \in \operatorname{bad}$, $\mathcal{A}_{1}$ makes at most $Q$ queries, $\mathrm{st}_{0}:=\mathcal{A}_{0}^{\Sigma, \mathrm{h}}(\Sigma(1), \Sigma(w))$ is of size $\left|\mathrm{st}_{0}\right| \leq M$ and

$$
\underset{j, \underset{\&}{\&}[0, n)}{\operatorname{Pr}}\left[\operatorname{TimeSC}_{\mathcal{A}_{1}}^{\Sigma, \mathrm{h}}\left(\mathrm{st}_{0}, j\right) \leq \frac{n}{6 \rho \log (n / 2)}\right] \geq 1 / 2
$$

where $\rho=(M+\log n+\log Q+c \log \lambda+1) /(\log (N)-3(\log Q-\log n)-3)$, then $|\mathbf{b a d}|<$ $\lambda^{-c}|\operatorname{params}(\lambda)|$.

In other words, bad is the set of parameters where, for every choice of parameters contained therein, $\mathcal{A}_{1}$ answers more quickly than expected. It is also possible to define the single-challenge trade-off the other way around. So there exists a large set good, and the trade-off holds for any parameters in good. Corollary 1 captures this view and also simplifies Lemma 1 by not stating all constants explicitly.

Corollary 1. For every pair of deterministic parallel oracle machines $\mathcal{A}=\left(\mathcal{A}_{0}, \mathcal{A}_{1}\right)$, there exists a negligible function $\epsilon(\lambda) \in \operatorname{negl}(\lambda)$ such that, for all $n, M, Q \in \operatorname{poly}(\lambda)$ with $n \geq 8$, there exists a subset good $\subseteq$ params $(\lambda)$ such that, for every $(N, \Sigma, \mathrm{~h}, w) \in \operatorname{good}, \mathcal{A}_{1}$ makes at most $Q$ queries, $\mathrm{st}_{0}:=\mathcal{A}_{0}^{\Sigma, \mathrm{h}}(\Sigma(w))$ is of size $\left|\mathrm{st}_{0}\right| \leq M$,

$$
\operatorname{Pr}_{j \leftarrow[0, n)}\left[\operatorname{TimeSC}_{\mathcal{A}_{1}}^{\Sigma, \mathrm{h}}\left(\mathrm{st}_{0}, j\right)>\frac{n}{6 \rho \log (n / 2)}\right] \geq 1 / 2
$$

where $\rho \in \Theta((M+\log n+\log Q+\log \lambda) /(\log (N)-(\log Q-\log n)))$ and $|\operatorname{good}| \geq(1-$ $\epsilon(\lambda))|\operatorname{params}(\lambda)|$.

Let us now compare Corollary 1's guarantees to the simple strategy $\mathcal{S S}$. $\rho$ is asymptotically close to $\lfloor M / \log N\rfloor \leq\left\lfloor M / \log \omega_{\mathcal{L}}\right\rfloor$ as in $\mathcal{S S}$. However, $n /(6 \rho \log (n / 2))$ is roughly a $1 / \log (n)$ factor looser than $n /(2 \rho)$. This seemingly is an artifact of our proof (looking ahead, Lemma 2), and we leave possible improvements to future work. The remainder of the section is devoted to the proof and structured as follows: First, we explain how to reason about $\mathcal{A}_{1}$ 's query behavior algebraically and state some algebraic tools related to that. Second, we give the high-level structure of the proof and highlight three central claims. Finally, we prove each claim to complete the proof.

### 5.1 Reasoning About $\mathcal{A}_{1}$ 's Queries Algebraically

Algebraic Representation of Labels. Algorithm $\mathcal{A}_{1}$ receives st ${ }_{0}$ (the output of preprocessing algorithm $\mathcal{A}_{0}(\Sigma(1), \Sigma(w))$ ) and challenge $j$ as input. It then uses the GGM

Table 1: Associating labels in the GGM to linear representations over $\mathbb{Z}_{N}$. Column "Query" lists the queries made by $\mathcal{A}_{1}$ and column "Response" the label returned as response to the query. Labels queried out of the blue, i.e., not previously returned to $\mathcal{A}_{1}$ as answer to a query, are collected in column "O.o.t.B. Inputs". The representations assigned to responses and o.o.t.b. input labels are given in column "Representation".

| Round | Query | O.o.t.B. Inputs | Response | Representation |
| :--- | ---: | ---: | ---: | ---: |
| 1 | $\mathcal{G}\left(+, \sigma_{1}, \sigma_{2}\right)$ |  |  |  |
|  |  | $\sigma_{1}$ |  | $x_{1}$ |
|  | $\sigma_{2}$ | $\sigma_{3}$ | $x_{1}+x_{2}$ |  |
| 2 | $\mathcal{G}\left(+, \sigma_{3}, \sigma_{3}\right)$ |  |  |  |
|  |  |  | $\sigma_{4}$ | $2 x_{1}+2 x_{2}$ |
| 3 | $\mathcal{G}\left(-, \sigma_{3}, \sigma_{5}\right)$ | $\sigma_{5}$ |  | $x_{3}$ |
|  |  |  | $\sigma_{6}$ | $x_{1}+x_{2}-x_{3}$ |
|  |  | $\vdots$ |  |  |

oracle to generate the label $\Sigma\left(w 2^{j}\right)$. As is typical in the GGM (e.g., [CK18, Sho97]), the reduction in our proof exploits $\mathcal{A}_{1}$ 's query behavior. To this end, the reduction associates an algebraic representation to every label that $\mathcal{A}_{1}$ queries to or receives from the oracle. More precisely, the algebraic representations are linear terms over $\mathbb{Z}_{N}$ in several indeterminates $x_{1}, \ldots, x_{m}$. The reduction stores the mapping of algebraic representations to labels in an (initially empty) table $T$ to ensure consistency between queries.

Whenever $\mathcal{A}_{1}$ makes a query $\mathcal{G}\left(+, \sigma_{1}, \sigma_{2}\right)$, the reduction does the following: ${ }^{8}$ First, it checks whether $T$ already contains the label $\sigma_{1}$. If not, it represents $\sigma_{1}$ by a new indeterminate $x_{i}$ and stores this mapping in $T$. Then, $\sigma_{2}$ is processed analogously before the reduction computes the label $\sigma_{3}=\mathcal{G}\left(+, \sigma_{1}, \sigma_{2}\right)$. If $T$ does not contain a representation for $\sigma_{3}$, its representation is set to the sum of the ones of $\sigma_{1}$ and $\sigma_{2}$. Queries for the operations $\circ \in\{-$, inv $\}$ are handled in the analogous way. Table 1 illustrates this explanation by example.

We stress that $\mathcal{A}_{1}$ only receives the (bit string) st ${ }_{0}$ and the challenge $j$ as input but not any labels. As a consequence, $T$ is initially empty. Thus, all algebraic representations will not contain any constant terms. This is in contrast to similar approaches (e.g., [Sho97]) where constant terms arise as a consequence of the adversary explicitly receiving the group generator as input at the beginning of the game. Intuitively, in our scenario $\mathcal{A}_{0}$ must store some labels (e.g., $\Sigma(w))$ in $\mathrm{st}_{0}$ in order for the hint to be useful.

System of Equations. In the above explanation, the reduction executed $\mathcal{A}_{1}$ with some arbitrary challenge $j \in[n]$. In the actual proof, the reduction cares about specific challenges, in particular, the subset $J \subseteq[0, n)$ of challenges that are answered within in at most $t$ rounds ( $t$ will be set later). For every $j \in J$, the reduction starts $\mathcal{A}_{1}\left(\mathrm{st}_{0}, j\right)$. Then,

[^5]it executes all instances of $\mathcal{A}_{1}$ in parallel in lockstep while only keeping track of a single table $T$. In more detail: First, the reduction initializes the empty table $T$. Second, for every $j \in J$, it runs the $\mathcal{A}_{1}\left(\mathrm{st}_{0}, j\right)$ until the first batch of parallel queries to the GGM oracle is made. Iterating over all $j \in J$, the reduction responds to the queries and adds labels to $T$ as described above. Then, it resumes the execution of all $\mathcal{A}_{1}$ instances until they query the GGM a second time. The reduction repeats this procedure for $t$ rounds.

This parallel and lockstep execution gives rise to a system of linear equations corresponding to the challenges $j \in J$ as follows. By definition of $J$, all challenges $j$ lead to $\mathcal{A}_{1}\left(\mathrm{st}_{0}, j\right)$ querying the label $\Sigma\left(w 2^{j}\right)$ within $t$ rounds. In turn, the table must contain this label with a corresponding representation of the form $a_{j, 1} x_{1}+\cdots+a_{j, m} x_{m}$ (here in, say, $m$ indeterminates). By correctness, it must satisfy

$$
\begin{equation*}
a_{j, 1} x_{1}+\cdots+a_{j, m} x_{m}=w 2^{j} \bmod N . \tag{1}
\end{equation*}
$$

Collecting these equations for all $j \in J$, we obtain a system of equations over $\mathbb{Z}_{N}$ of the form $A \boldsymbol{x}=\boldsymbol{b}$, where $A=\left(a_{j i}\right)_{j \in J, i \in[m]}$ and $\boldsymbol{b}=\left(2^{j} w \bmod N \mid j \in J\right)^{\top}$. Since Equation (1) holds for every equation, the system must have a solution $\boldsymbol{x} \in \mathbb{Z}_{m}$. Intuitively, the solution $\boldsymbol{x}$ corresponds to the discrete logarithm of group elements queried by $\mathcal{A}_{1}$ out of the blue (represented by indeterminates) of which there are $m$ in total. In practice, $\mathcal{A}_{1}$ must have (mostly) stored these group elements in $\mathrm{st}_{0}$ as a randomly guessed group element will most likely not be useful in computing a challenge.

Groups of Unknown Order. The above characterization of $A$ requires knowledge of the group order $N$. However, assuming that factoring is hard, the reduction and, more importantly, $\mathcal{A}_{1}$ do not know $N$. Therefore, we will argue that it is almost always valid to consider the system $A \boldsymbol{x}=\boldsymbol{b}$ as being defined over $\mathbb{Z}$ instead of $\mathbb{Z}_{N}$. That is, $A \in \mathbb{Z}^{\ell \times m}$, $\boldsymbol{x} \in \mathbb{Z}^{m}$, and $\boldsymbol{b} \in \mathbb{Z}^{\ell}$ with $\boldsymbol{b}=\left(2^{j} w \mid j \in J\right)^{\top}$ Since $\mathcal{A}_{1}$ is executed for $t$ rounds, for every entry $a_{j i} \in A$, we have $\left|a_{j i}\right| \leq 2^{t}$ as a consequence. ${ }^{9}$ Furthermore, $A \boldsymbol{x}=\boldsymbol{b}$ constitutes a system of Diophantine equations since $A, \boldsymbol{x}$, and $\boldsymbol{b}$ only have integer components. ${ }^{10}$ Since we are working with integers, this system might not have a solution $\boldsymbol{x}$. However, looking ahead, it almost always will, assuming that factoring is hard.

Solving a system of equations $A \boldsymbol{x}=\boldsymbol{b}$ over $\mathbb{Z}$ works similarly to solving one over a field. Recall that in a field one commonly uses Gaussian elimination to transform $A$ into a row- or column-reduced echelon form (i.e., a triangular matrix with pivot entries equal to 1) because then the solution can easily be found algorithmically. The integer analogue to the column reduced echelon form is the well-known column-style Hermite normal form (HNF) which we define as in [MW01].

Definition 7 (Hermite Normal Form). A matrix $H \in \mathbb{Z}^{\ell \times m}$ is in Hermite normal form if
(i) there exists a sequence of integers $1 \leq i_{1}<\cdots<i_{r} \leq \ell$ such that $h_{i j}=0$ for all $i<i_{j}$ where $r$ is the number of non-zero columns; and
(ii) $0 \leq h_{i_{j} k}<h_{i_{j} j}$ for all $0 \leq k<j \leq m$.

[^6]So by Item (i), $H$ is a lower-triangular matrix and all its zero columns are to the right. Furthermore, taking Item (ii) into account, for every matrix $A$, there exists a unique HNF $H$ such that $H=A U$ where $U \in \mathbb{Z}^{m \times m}$ is unimodular (i.e., invertible over $\mathbb{Z}$ ). Moreover, $H$ has $r:=\operatorname{rank}(A)$ non-zero columns [Her09]. In the rest of the paper, we denote the $\operatorname{HNF}$ of $A$ by $\operatorname{HNF}(A)=(H, U)$.

Note that the definition of $\operatorname{rank}(\cdot)$ might differ to one's own intuitive definition. This is due to $\mathbb{Z}^{\ell \times m}$ being a module and not a vector space. In contrast to linear algebra, different definitions of rank are not equivalent, so we define $\operatorname{rank}(\cdot)$ below.

Definition 8 (Rank). $\operatorname{rank}(A)$ is the size of largest subset of $A$ 's columns that are linearly independent. A set of columns $\left\{\boldsymbol{c}_{1}, \ldots, \boldsymbol{c}_{k}\right\}$ is linearly dependent if there exist $\alpha_{1}, \ldots, \alpha_{k} \in \mathbb{Z}$ with at least one $\alpha_{i} \neq 0$ such that $\sum_{i=1}^{k} \alpha_{i} \boldsymbol{c}_{i}=\mathbf{0}$.

### 5.2 Proof Skeleton

We first give a high-level overview of the proof. Assume that Lemma 1 does not hold. So we have a large set of parameters bad where the trade-off does not hold for at least half of the challenges. That is, we have an upper bound on how many rounds the adversary needs to answer them. Thus, we can derive a system of equations $A \boldsymbol{x}=\boldsymbol{b}$ over $\mathbb{Z}$ where the $a_{i j} \in A$ are bounded in magnitude due to the trade-off not holding. Then, we prove the following claims about this system:

- Assuming factoring is hard, $A \boldsymbol{x}=\boldsymbol{b}$ behaves identically over $\mathbb{Z}$ and $\mathbb{Z}_{N}$ almost always. For example, $A \boldsymbol{x}=\boldsymbol{b}$ having a solution over $\mathbb{Z}_{N}$ does not necessarily imply the same holds over $\mathbb{Z}$, yet the implication will almost always hold under the factoring assumption.
- Assuming that $A \boldsymbol{x}=\boldsymbol{b}$ over $\mathbb{Z}$ has a solution, $\operatorname{rank}(A) \geq \rho$ for every choice of parameters in bad.
- $\operatorname{rank}(A)<\rho$ for some choice of parameters in bad. Otherwise, a random labeling function $\Sigma$ could be compressed more efficiently than information-theoretically possible.

Notice that the last two claims contradict each other, so Lemma 1 must hold. We now give the detailed version of the above sketch.

Proof of Lemma 1. We distinguish two cases depending on $\rho$, the first one being trivial. Case $\rho>n /(6 \log (n / 2))$ : Then $n /(6 \rho \log (n / 2))<1$ and so the trade-off holds for any $j \in[0, n)$ since TimeSC $\geq 1$ by Definition 6 .

Case $\rho \leq n /(6 \log (n / 2))$ : Towards contradiction, assume that Lemma 1 does not hold. That is, there exists a pair of deterministic parallel oracle machines $\mathcal{A}=\left(\mathcal{A}_{0}, \mathcal{A}_{1}\right), c \in \mathbb{N}$ with security parameter $\lambda$ large enough and, for infinitely many security parameters $\lambda$, there exist $n, M, Q \in \operatorname{poly}(\lambda)$ with $n \geq 8$ and a subset $\operatorname{bad} \subseteq \operatorname{params}(\lambda)$, such that, for every $(N, \Sigma, \mathrm{~h}, w) \in \operatorname{bad}, \mathcal{A}_{1}$ makes at most $Q$ queries, $\mathrm{st}_{0}:=\mathcal{A}_{0}^{\Sigma, \mathrm{h}}(\Sigma(1), \Sigma(w))$ is of size $\left|\mathrm{st}_{0}\right| \leq M$ and

$$
\underset{j \underset{\sim}{\ddagger}[0, n)}{\operatorname{Pr}}\left[\operatorname{TimeSC} \mathcal{A}_{1}^{\Sigma, \mathrm{h}}\left(\mathrm{st}_{0}, j\right) \leq \frac{n}{6 \rho \log (n / 2)}\right] \geq 1 / 2
$$

where $\rho=(M+\log n+\log Q+c \log \lambda+1) /(\log (N)-3(\log Q-\log n)-3)$ and $|\operatorname{bad}| \geq$ $\lambda^{-c}|\operatorname{params}(\lambda)|$.

So for parameters in bad, it holds that $\mathcal{A}_{1}$ will answer at least $n / 2$ challenges correctly by the end of round $\frac{n}{6 \rho \log (n / 2)}$. By the definition of rounds (cf. Section 2.4), to compute these answers, it can make and receive the results of $t:=\frac{n}{6 \rho \log (n / 2)}-1$ rounds of queries.

As in Section 5.1, consider the system of Diophantine equations $A \boldsymbol{x}=\boldsymbol{b}$ where $A \in \mathbb{Z}^{(n / 2) \times m}$ for some $m \in \mathbb{N}^{+}$arising from $\mathcal{A}_{1}$ 's query behavior. In the case that $\mathcal{A}_{1}$ answers more than $n / 2$ challenges, we avoid ambiguities by selecting $n / 2$ challenges in some deterministic order (e.g., low challenge numbers first). Note that every matrix element $\left|a_{i j}\right| \leq 2^{t}$ and that the vector $\boldsymbol{b}$ does not contain 0 , its elements are pairwise distinct, and of the form $b_{i}=2^{j} w$ for $j \in[0, n)$.

First, assuming factoring is hard, $A \boldsymbol{x}=\boldsymbol{b}$ behaves similarly to $A \boldsymbol{x}=\boldsymbol{b} \bmod N$ almost always. It is proven in Section 5.3.

Claim 1. If Lemma 1 does not hold, there exists a negligible function $\epsilon(\lambda) \in \operatorname{negl}(\lambda)$ such that
(i) $A \boldsymbol{x}=\boldsymbol{b}$ has a solution;
(ii) $\operatorname{HNF}(A)$ and $\operatorname{HNF}(A) \bmod N$ have the same shape (i.e., $\operatorname{HNF}(A) \bmod N$ has no additional zero entries); and
(iii) the diagonal entries of $\operatorname{HNF}(A) \bmod N$ are invertible over $\mathbb{Z}_{N}$
except for probability $\epsilon(\lambda)$ where the probability is taken over the choice of $(N, \Sigma, \mathrm{~h}, w) \in$ params $(\lambda)$.

Second, since $\mathcal{A}$ answers a lot of challenges quickly, its rank must be large as proven in Section 5.4.

Claim 2. If Lemma 1 does not hold and $\rho \leq n /(6 \log n)$, for all $(N, \Sigma, \mathrm{~h}, w) \in \operatorname{bad}$, $\operatorname{rank}(A) \geq \rho$.

Last, $\operatorname{rank}(A)$ cannot be too large by an incompressibility argument presented in Section 5.5.

Claim 3. If Lemma 1 does not hold, there exists a $(N, \Sigma, \mathrm{~h}, w) \in \operatorname{bad}$, such that $\operatorname{rank}(A)<$ $\rho$.

The latter two claims contradict each other, completing the proof.

### 5.3 Analyzing the Behavior of $A \boldsymbol{x}=\boldsymbol{b}$

The proof of Claim 1 follows from the assumed hardness of factoring. In particular, Item (i) and Item (ii) reduce to factoring $N^{\prime}$, and Item (iii) to factoring $N$.

Proof of Claim 1. Let $E_{i}$ denote the event that item $i$ does not hold. We will show that all events happen with negligible probability.

- There exists a negligible function $\epsilon(\lambda) \in \operatorname{negl}(\lambda)$ such that $\operatorname{Pr}_{\text {params }}\left[E_{\text {(iii) }}\right] \leq \epsilon(\lambda)$. To see this, consider the algorithm $\mathcal{B}$ playing the game $\operatorname{Fac}_{G e n S P, \mathcal{B}}(\lambda)$ (cf. Definition 3). It uses $\mathcal{A}$ (guaranteed by the assumption that Lemma 1 does not hold) as a subroutine. On input $N, \mathcal{B}$ picks $w \leftarrow \mathbb{Z}_{N} \backslash\{0\}$, and lazily samples h and $\Sigma$ to respond to $\mathcal{A}$ 's queries. Note that $\mathcal{B}$ can simulate $\Sigma$ and h perfectly because it knows $N$. Assuming that $N, w$ and what has been sampled of $\Sigma$ and h so far is consistent with bad, $\mathcal{B}$ is able to derive the system $A \boldsymbol{x}=\boldsymbol{b}$. Then, it computes $\operatorname{HNF}(A) \bmod N$ and checks whether any diagonal element $a$ cannot be inverted. If so, $\mathcal{B}$ outputs $\left(p^{*}, q^{*}\right)$ where $p^{*}=\operatorname{gcd}(a, N)$ and $q^{*}=N / p^{*}$.
$\mathcal{B}$ is a PPT algorithm and $(N, \Sigma, \mathrm{~h}, w) \in$ bad with probability at least $\lambda^{-c}$. Hence,

$$
\operatorname{Pr}\left[E_{(\mathrm{iii})}\right] \leq \lambda^{c} \operatorname{Adv}_{\mathrm{GensP}, \mathcal{B}}^{\mathrm{Fac}}(\lambda)
$$

which is negligible as factoring $N$ is assumed to be hard.

- There exists a negligible function $\epsilon(\lambda) \in \operatorname{negl}(\lambda)$ such that $\operatorname{Pr}_{\text {params }}\left[E_{(\mathrm{i})}\right]+\operatorname{Pr}_{\text {params }}\left[E_{(\mathrm{ii})}\right] \leq$ $\epsilon(\lambda)$.

The reduction is similar to the previous item except for two differences. First, $\mathcal{B}$ only knows $N^{\prime}$ but not $N$, so it cannot simulate $\Sigma$ perfectly. However, this is not an issue as $\mathcal{B}$ can approximate lazy sampling $\Sigma$ by using $N^{\prime} / 4$ instead of $N$. Since $N^{\prime} / 4=N+O(\sqrt{N})$, answering a polynomial number of $\Sigma$ queries using lazy sampling with $N^{\prime} / 4$ is statistically close to lazy sampling with $N$.
Second, $\mathcal{B}$ factors $N^{\prime}$ instead of $N$. If $A \boldsymbol{x}=\boldsymbol{b}$ has a solution over $\mathbb{Z}$ but not $\mathbb{Z}_{N}$ or if a matrix entry is non-zero over $\mathbb{Z}$ but zero over $\mathbb{Z}_{N}, \mathcal{B}$ learns a multiple of the group order $N$. Such a multiple can be used to factor $N^{\prime}$ with probability at least 1/2 [KL14, Thm. 8.50].

Combining the two items above completes the proof.

### 5.4 Combinatorial Proof of the $\operatorname{rank}(A)$ Lower Bound

The key result of this section is the following number-theoretical lemma. Given a system of Diophantine equations $A \boldsymbol{x}=\boldsymbol{b}$, it bounds $\operatorname{rank}(A)$ from below as a function of $A$ 's number of rows and the magnitude of its entries.

Lemma 2. Let $\ell \in \mathbb{N}$ with $\ell \geq 4, t \in \mathbb{N}$, and $m, w \in \mathbb{N}^{+}$. For any matrix $A=$ $\left(a_{i, j}\right)_{i \in[\ell], j \in[m]} \in \mathbb{Z}^{\ell \times m}$ with $\left|a_{i, j}\right| \leq 2^{t}$ for all $i, j$ and vector $\boldsymbol{b}=\left(2^{j_{1}} w, \ldots, 2^{j_{\ell}} w\right)^{\top}$ with pairwise distinct $j_{1}, \ldots, j_{\ell} \in \mathbb{N}$, if the system $A \boldsymbol{x}=\boldsymbol{b}$ has a solution $\boldsymbol{x} \in \mathbb{Z}^{m}$, then

$$
\operatorname{rank}(A) \geq \frac{\ell}{3 \max \{t, \log \ell\}}
$$

In the context of the single-challenge trade-off, the lemma states the following: If $\mathcal{A}_{1}$ solves a lot of challenges, either it must take a long time or $\operatorname{rank}(A)$ is high. This makes sense since the rank approximately corresponds to the number of group elements stored in $\mathrm{st}_{0}$ which is at most $M / \log N \approx \rho$. Note that Claim 2 formalizes precisely this intuition, so it follows from Lemma 2 almost directly.

Proof of Claim 2. By the assumption that Lemma 1 does not hold, $\mathcal{A}_{1}$ answers at least $\ell:=n / 2 \geq 4$ challenges within $t:=n /(6 \rho \log (n / 2))-1$ rounds of queries. By Claim 1 Item (i), we may assume that bad only contains parameters where the system $A \boldsymbol{x}=\boldsymbol{b}$ has a solution. So we apply Lemma 2 to get

$$
\operatorname{rank}(A) \geq \frac{n}{6 \max \{t, \log (n / 2)\}}
$$

and we consider two cases depending on max: First, if $\operatorname{rank}(A) \geq n /(6 t)$, it follows that $\operatorname{rank}(A) \geq n /(6(t+1))$ and substituting for $t$, we $\operatorname{arrive}$ at $\operatorname{rank}(A) \geq \rho \log (n / 2) \geq \rho$ as desired since $\log (n / 2) \geq 2$. Second, if $\operatorname{rank}(A) \geq n /(6 \log (n / 2))$, by the assumption that $n /(6 \log (n / 2)) \geq \rho$, we also get $\operatorname{rank}(A) \geq \rho$ completing the proof.

We now turn to the proof of Lemma 2. The first ingredient is an elegant numbertheoretic lemma due to van der Waerden (restated and proven in [Laz96]) which precisely describes when a system of Diophantine equations has a solution.

Lemma 3 (van der Waerden). Consider the Diophantine system of linear equations $A \boldsymbol{x}=\boldsymbol{b}$ with $A \in \mathbb{Z}^{\ell \times m}$ and $\boldsymbol{b} \in \mathbb{Z}^{\ell}$. The system has a solution $\boldsymbol{x} \in \mathbb{Z}^{m}$ if and only if, for every rational vector $\boldsymbol{v} \in \mathbb{Q}^{\ell}$ such that $\boldsymbol{v} A \in \mathbb{Z}^{m}$, it also holds that $\boldsymbol{v} \boldsymbol{b} \in \mathbb{Z}$.

We will show that if $\operatorname{rank}(A)$ is too small, such a vector $\boldsymbol{v}$ always exists. To this end, we introduce the second ingredient which is a combinatorial claim. It is a vector version of the famous distinct subset sums problem and associated conjecture ${ }^{11}$ due to Erdős [Guy94, C8].

Claim 4. For $t, m, \ell \in \mathbb{N}$ with $\ell \geq 4$, let $R \subset \mathbb{Z}^{m}$ with $|R|=\ell$ and every $\boldsymbol{r} \in R$ satisfying $\|\boldsymbol{r}\|_{\infty} \leq 2^{t} .{ }^{12}$ If $m<\ell /(3 \max \{t, \log \ell\})$, two subsets $R_{1}, R_{2} \subseteq R$ exist such that at least one subset is non-empty $\left(R_{1} \cup R_{2} \neq \emptyset\right)$, they are disjoint ( $R_{1} \cap R_{2}=\emptyset$ ), and their sums are equal $\left(\sum_{\boldsymbol{r}_{1} \in R_{1}} \boldsymbol{r}_{1}=\sum_{\boldsymbol{r}_{2} \in R_{2}} \boldsymbol{r}_{2}\right)$.

Note that the above claim is slightly more general than the vector version considered by Costa et al. [CDD23, Prop. 2.1], but their proof is easily adapted to our setting.

Proof. First, for any subset's sum, notice that the largest absolute value of any resulting vector coordinate is bounded by $2^{t} \ell$, which follows directly from the triangle inequality and the fact that subsets contain at most $\ell$ elements.

Second, there exist $2^{\ell}$ distinct (not necessarily disjoint) subsets. Assuming that all subsets' sums are distinct, by the pigeonhole principle, it must hold that

$$
\begin{aligned}
\left|\left\{-2^{t} \ell,-2^{t} \ell+1, \ldots, 2^{t} \ell-1,2^{t} \ell\right\}\right|^{m} & \geq 2^{\ell} \\
\left(2^{t+2} \ell\right)^{m} & \geq 2^{\ell} \\
m(t+2+\log \ell) & \geq \ell \\
m & \geq \ell /(3 \max \{t, \log \ell\})
\end{aligned}
$$

[^7]where the second line follows by overapproximating the size of the set, the third by applying $\log$ to both sides, and the fourth by $2 \leq \log \ell \leq \max \{t, \log \ell\}$ since $\ell \geq 4$ by assumption. As a consequence, if $m<\ell /(3 \max \{t, \log \ell\})$, then there exist two distinct, but not necessarily disjoint subsets $R_{1}^{\prime}$ and $R_{2}^{\prime}$ with equal subset-sum.

To get disjoint sets $R_{1}$ and $R_{2}$, we remove all common elements by defining $R_{i}:=$ $R_{i} \backslash\left(R_{1}^{\prime} \cap R_{2}^{\prime}\right)$ for $i \in\{1,2\}$. Notice that at least one set is non-empty since $R_{1}^{\prime}$ and $R_{2}^{\prime}$ are distinct. This completes the proof.

Equipped with these tools, the proof of Lemma 2 is straightforward.
Proof of Lemma 2. Towards contradiction, assume that there exist $A$ as well as $\boldsymbol{b}$ satisfying the constraints in the lemma's statement and that the system $A \boldsymbol{x}=\boldsymbol{b}$ has a solution, but

$$
\operatorname{rank}(A)<\frac{\ell}{3 \max \{t, \log \ell\}}
$$

We will show that this implies the existence of a vector $\boldsymbol{v} \in\{-1,0,1\}^{\ell}$ with $\boldsymbol{v} \neq \mathbf{0}$ and $\boldsymbol{v} A=\mathbf{0}$.

For finding such a $\boldsymbol{v}$, we may assume that $\operatorname{rank}(A)=m$ (i.e., $A$ is full-rank). To see this, assume that $A \in \mathbb{Z}^{\ell \times m}$ is not full rank. We will see that it is sufficient to consider $A^{\prime} \in \mathbb{Z}^{\ell \times \operatorname{rank}(A)}$, the submatrix containing a linearly independent subset of $A$ 's columns that is of maximum size, i.e., $\operatorname{rank}(A)$.

Let the columns of $A$ be $\boldsymbol{c}_{1}, \ldots, \boldsymbol{c}_{m}$ and, without loss of generality, assume that $\boldsymbol{c}_{1}, \ldots, \boldsymbol{c}_{\mathrm{rank}(A)}$ are linearly independent, so $A^{\prime}=\left(\boldsymbol{c}_{1}, \ldots, \boldsymbol{c}_{\mathrm{rank}(A)}\right)$. By definition of linear dependence (cf. Definition 8), for any column $j$ with $\operatorname{rank}(A)<j \leq m$, there exist coefficients $\alpha_{i}$ (with at least one $\alpha_{i} \neq 0$ ) such that

$$
\alpha_{1} \boldsymbol{c}_{1}+\cdots+\alpha_{\mathrm{rank}(A)} \boldsymbol{c}_{\mathrm{rank}(A)}+\alpha_{j} \boldsymbol{c}_{j}=\mathbf{0}
$$

Multiplying by $\boldsymbol{v}$, we get

$$
\alpha_{1} \boldsymbol{v} \boldsymbol{c}_{1}+\cdots+\alpha_{\mathrm{rank}(A)} \boldsymbol{v} \boldsymbol{c}_{\mathrm{rank}(A)}+\alpha_{j} \boldsymbol{v} \boldsymbol{c}_{j}=\boldsymbol{v} \mathbf{0}
$$

which simplifies to $\alpha_{j} \boldsymbol{v} \boldsymbol{c}_{j}=\mathbf{0}$ since $\boldsymbol{v} A^{\prime}=\mathbf{0}$ by assumption. Note that $\alpha_{j} \neq 0$ as otherwise some $\alpha_{i}$ where $1 \leq i \leq \operatorname{rank}(A)$ would need to be non-zero which would imply that $\boldsymbol{c}_{1} \ldots, \boldsymbol{c}_{\mathrm{rank}(A)}$ are not linearly independent. Dividing by $\alpha_{j}$, it follows that $\boldsymbol{v} \boldsymbol{c}_{j}=\mathbf{0}$ for all $\operatorname{rank}(A)<j \leq m$ which implies $\boldsymbol{v} A=0$. It follows that we can restrict our attention to the full-rank case.

Now we use Claim 4 to construct $\boldsymbol{v}$. Whenever $A$ contains two identical rows, it is straightforward to define an appropriate $\boldsymbol{v}$. Thus, we may assume that the rows of $A$ are distinct and form a set $R \subset \mathbb{Z}^{m}$ of size $|R|=\ell$. Combining this with the assumption $m=\operatorname{rank}(A)<\ell /(3 \max \{t, \log \ell\})$, Claim 4 applies, so there exist two disjoint subsets $R_{1}$ and $R_{2}$ with equal sum with one of them being non-empty. Hence, we initialize $\boldsymbol{v}:=\mathbf{0}$, and set $v_{i}:=1$ if $\boldsymbol{r}_{i} \in R_{1}$ and $v_{j}:=-1$ if $\boldsymbol{r}_{j} \in R_{2}$. This is unambiguous as the subsets are disjoint. Further, notice that $\boldsymbol{v} \neq \mathbf{0}$ since one subset is non-empty, and, since the subsets' sums are equal, $\boldsymbol{v} A=\mathbf{0}$ as required.

Last, we apply Lemma 3 to $A \boldsymbol{x}=\boldsymbol{b}$ and $\boldsymbol{v}$. By definition, $\boldsymbol{b} \in \mathbb{Z}^{\ell}$ contains no element $b_{i}=0$, all $b_{i}$ are pairwise distinct and exponentially far apart. Thus, $\boldsymbol{v} \boldsymbol{b}=\beta \in \mathbb{Z}$ with $\beta \neq 0$ since $\boldsymbol{v} \in\{-1,0,1\}$. Setting $\boldsymbol{v}^{\prime}:=\frac{1}{|\beta|+1} \boldsymbol{v}$, we arrive at the conclusion that the system $A \boldsymbol{x}=\boldsymbol{b}$ does not have a solution by van der Waerden (Lemma 3). This is a contradiction, completing the proof.

### 5.5 Incompressibility Argument

Recall that Claim 3 claims the existence of parameters in bad such that $\operatorname{rank}(A)<\rho$. To this end, we will show that rank $_{\min }<\rho$ where rank $_{\min }$ is the minimum rank over all possible parameters in bad, that is,

$$
\begin{equation*}
\operatorname{rank}_{\min }:=\min _{\mathrm{bad}}(\operatorname{rank}(A)) . \tag{2}
\end{equation*}
$$

Our proof use an incompressibility argument (e.g., [Yao90, DTT10, ACP ${ }^{+}$17, AAC ${ }^{+} 17$, CK18]) to bound rank min in terms of $\left|s t_{0}\right|=M$ and since $\rho$ depends on $M$, we get the desired bound. Formally, we follow the framework of De, Trevisan, and Tulsiani [DTT10] stated below.

Lemma 4 (De, Trevisan, and Tulsiani [DTT10]). Let Enc: $\mathcal{S} \times\{0,1\}^{\mu} \rightarrow\{0,1\}^{r}$ and Dec: $\{0,1\}^{\mu} \times\{0,1\}^{r} \rightarrow \mathcal{S}$ be randomized encoding and decoding procedures such that, for every $s \in \mathcal{S}$,

$$
\operatorname{Pr}_{r \leftarrow\{0,1\}^{r}}[\operatorname{Dec}(\operatorname{Enc}(s ; r) ; r)=s] \geq \delta .
$$

Then $\mu \geq \log |\mathcal{S}|-\log 1 / \delta$.
Intuitively, the analysis of preprocessing algorithms $\mathcal{A}=\left(\mathcal{A}_{0}, \mathcal{A}_{1}\right)$ using the framework proceeds as follows. $\mathcal{S}$ is usually a set of functions, injections, etc. where $\mathcal{A}$ has oracle access to one $s \in \mathcal{S}$. Enc, knowing $s \in \mathcal{S}$, runs $\mathcal{A}_{0}^{s}$ to get st ${ }_{0}$ and then $\mathcal{A}_{1}\left(\mathrm{st}_{0}\right)^{s}$ such that $\mathcal{A}_{1}$ 's query behavior implicitly allows one to reconstruct parts of $s$. It outputs an encoding comprised of $\mathrm{st}_{0}$ and an additional hint. Dec receives this encoding and runs $\mathcal{A}_{1}\left(\mathrm{st}_{0}\right)^{\text {hint }}$ where the hint is used to simulate the oracle. Given $\mathcal{A}_{1}$ 's query behavior, Dec recovers $s$ at the cost of $\left\|s t_{0}\right\|+\|$ hint $\|$ bits where the hint is carefully designed to be smaller than $\log |\mathcal{S}|$ bits. Since encoding $s$ using less than $\log |\mathcal{S}|$ bits would constitute information-theoretically impossible compression, $\mathrm{st}_{0}$ must be large enough to contain the remaining information.

In our case, we work with $\mathcal{S}:=$ bad, but effectively we only efficiently encode the labeling function $\Sigma \in \operatorname{Inj}\left(\mathbb{Z}_{N}, \mathcal{L}\right)$ while using a naive encoding for $N$, h, and $w$. Our argument then follows the strategy outlined in the previous paragraph. As in Section 5.1, $\mathcal{A}_{1}$ induces the system of Diophantine equations $A \boldsymbol{x}=\boldsymbol{b}$ where $A$ and $\boldsymbol{b}$ are known. Every element of $\boldsymbol{x} \bmod N$ corresponds to the discrete logarithm of some label in $\mathcal{L}$ and Dec can solve the system to extract $\boldsymbol{x}$, effectively learning some discrete logarithms, and, in turn, the labels for the corresponding group elements without them being encoded in the hint. Given usual linear algebra intuition, the solvability of this system and the uniqueness of its solution are related to the rank of $A$. Here, this is rank min which allows us to connect $\left\|\mathrm{st}_{0}\right\|=M$ to rank ${ }_{\text {min }}$.

Equipped with this high-level explanation, we now state Claim 5. It posits the existence of (Enc, Dec) that efficiently encode the labeling function $\Sigma \in \operatorname{Inj}\left(\mathbb{Z}_{N}, \mathcal{L}\right)$ where the encoding size depends on $M$ and rank $_{\text {min }}$.
Claim 5. If Lemma 1 does not hold, for $(N, \Sigma, \mathrm{~h}, w) \in$ bad, there exist (Enc, Dec) which, given auxiliary inputs $N$, h, and $w$, encodes the labeling function $\Sigma$ using

$$
\log \binom{|\mathcal{L}|}{N}+\log (N!)+M+\log n+\log Q-\text { rank }_{\text {min }} \cdot(\log N-3(\log n+\log Q)-3)
$$

bits with probability 1.

Before constructing such (Enc, Dec), we show how their existence implies Claim 3.
Proof of Claim 3. In general, optimally encoding an element of params requires $\log \mid$ params $\mid=$ $\mu^{\prime}+\log \binom{|\mathcal{L}|}{N}+\log (N!)$ bits where $\mu^{\prime}$ is the length of an optimal encoding of $N$, h , and $w$. Since bad contains at least a $1 / \lambda^{c}$ fraction of params, it follows from elementary logarithmic identities that an optimal encoding of a tuple in bad requires

$$
\mu^{\prime}+\log \binom{|\mathcal{L}|}{N}+\log (N!)+\log \left(\lambda^{-c}\right)=\mu^{\prime}+\log \binom{|\mathcal{L}|}{N}+\log (N!)-c \log \lambda
$$

bits. Now, by Claim 5 and applying Lemma 4 while canceling common terms (implicitly accounting for $\mu^{\prime}$ ), it must hold that

$$
-c \log \lambda \leq M+\log n+\log Q-\text { rank }_{\min } \cdot(\log N-3(\log n+\log Q)-3)
$$

which implies

$$
\text { rank }_{\text {min }} \leq \frac{M+\log n+\log Q+c \log \lambda}{\log N-3(\log Q-\log n)-3}
$$

By the definition of $\rho$ in Lemma 1, it follows that rank $_{\min }<\rho$, so there exists a $(N, \Sigma, \mathrm{~h}, w) \in$ bad such that $\operatorname{rank}(A)<\rho$ as desired.

The only thing left is to show is that (Enc, Dec) for $\Sigma$ as in Claim 5 exist. This requires careful handling of various technicalities, so we will only sketch the proof here and defer the complete proof to Appendix C.

Proof of Claim 5 (Sketch). The high-level idea is that Enc and Dec both run $\mathcal{A}_{1}$ as a subroutine. $\mathcal{A}_{1}$ induces a system of equations $A \boldsymbol{x}=\boldsymbol{b}$ with $\operatorname{rank}(A) \geq \operatorname{rank}_{\text {min }}$. This allows Enc to encode the discrete logarithm of rank min many group element labels using a short hint instead of the naive encoding that costs roughly $\log N$ bits.

Given input $\Sigma$ and auxiliary inputs $N$, h and $w$, Enc computes $\mathrm{st}_{0}$ by running $\mathcal{A}_{0}^{\Sigma, \mathrm{h}}(\Sigma(1), \Sigma(w))$ and then executes $\mathcal{A}_{1}^{\Sigma, \mathrm{h}}\left(\mathrm{st}_{0}\right)$ on every challenge $j \in[0, n)$ for at most $t:=\frac{n}{6 \rho \log (n / 2)}-1$ rounds in parallel in lockstep. During this, Enc successively constructs an encoding of $\Sigma$. This encoding contains the information necessary for Dec to also run $\mathcal{A}_{1}$ without knowing the oracle $\Sigma$.

In particular, Enc (and also Dec) track the algebraic representation of queries as in Section 5.1. The algebraic representation of a label might include some indeterminates which correspond to labels that $\mathcal{A}_{1}$ input to $\Sigma$ out of the blue (i.e., either it guessed the label or it extracted it from $\mathrm{st}_{0}$ ).

On the one hand, Enc needs to ensure that Dec can answer oracle queries to the labeling function $\Sigma$ correctly. To this end, it first encodes the image of $\Sigma$ using $\log \binom{|\mathcal{L}|}{N}$ bits, and, for some queries, it adds a hint of circa $\log N$ bits which specify the label to be returned. Note that this hint is only required for some queries since the output of a query is sometimes uniquely determined by the algebraic representations of the preceding queries (e.g., when $\mathcal{A}_{1}$ repeats a query). Since Dec tracks those representations, it can answer such queries without additional help from Enc.

On the other hand, Enc also needs to tell Dec when the label $\Sigma\left(w^{2^{j}}\right)$ is queried to h . This requires a hint of $\log n+\log Q$ bits; $\log n$ bits to identify the instance of $\mathcal{A}_{1}$ and $\log Q$ to specify the query within that instance. Given rank ${ }_{\min }$ many hints, Dec derives
a system of equations $A \boldsymbol{x}=\boldsymbol{b}$. The vector $\boldsymbol{b} \in \mathbb{Z}^{\text {rank min }^{m}}$ contains the $2^{i} w$ specified by the hints and the rows of $A$ the corresponding algebraic representations. So $\boldsymbol{x} \in \mathbb{Z}^{m}$, when taken modulo $N$, corresponds to the discrete logarithms of some indeterminates within these algebraic representations.

Computing $\operatorname{HNF}(A)=(H, U)$ yields a lower triangular matrix $H$ such that $A \boldsymbol{x}=\boldsymbol{b}$ has a solution if and only if $H \boldsymbol{y}=\boldsymbol{b}$ has a solution where $\boldsymbol{x}=U \boldsymbol{y}$. Recall that $H$ has rank $_{\min }$ many non-zero columns and, due to the triangular shape of $H$, every non-zero column constrains one entry of $\boldsymbol{x}$. For example, consider

$$
\left(\begin{array}{ll}
h_{11} & \\
h_{21} & h_{22} \\
h_{31} & h_{32}
\end{array}\right)\left(\begin{array}{l}
y_{1} \\
y_{2} \\
y_{3}
\end{array}\right)=\left(\begin{array}{l}
2^{i} w \\
2^{j} w \\
2^{k} w
\end{array}\right)
$$

with rank ${ }_{\text {min }}=2$ and pivot elements $h_{11}$ and $h_{22}$.
Next, reduce every matrix and vector modulo $N$. Notice that Claim 1 Items (ii) and (iii) guarantee that $H \bmod N$ still has the same shape and that the pivot elements are invertible. This implies that every non-zero column of $H \bmod N$ uniquely constrains one entry in $\boldsymbol{y} \bmod N$ and thus effectively one indeterminate. So given $\mathcal{A}_{1}$ 's query behavior, Dec derives the discrete logarithm of rank $_{\text {min }}$ many indeterminates with a hint of only rank $_{\text {min }}(\log n+\log Q)$ bits.

Last, we analyze the encoding size relative to the naive encoding of $\log \binom{|\mathcal{L}|}{N}+\log (N!)$ bits. Answering queries in general requires no more bits than the naive encoding: $\log \binom{|\mathcal{L}|}{N}$ bits for the image and roughly $\log N$ bits for every query. For every indeterminate recovered from the system of equations, the hint is only $\log n+\log Q$ bits which is smaller than $\log (N / 2)$ for $\lambda$ large enough. Thus, the total encoding size is

$$
\underbrace{\log \binom{|\mathcal{L}|}{N}+\log (N!)}_{\text {Cost of DLogs in general }}+\underbrace{M}_{\mid \text {sto } \mid}-\underbrace{\operatorname{rank}_{\min } \cdot(\log N-1-\log n-\log Q)}_{\text {Savings due to DLogs from } A \boldsymbol{x}=b} .
$$

Note that the savings in the claim's actual statement are lower since Enc has to handle colliding $\Sigma$ queries which introduce additional overhead.

## 6 Multi-challenge Memory Complexity

We now consider the whole evaluation of TdScrypt, and we want to prove that it requires $\Omega\left(\frac{n^{2}}{\log n} \log N\right)$ memory almost always - no matter the evaluation strategy. This is formalized in the theorem below which is in the spirit of Alwen et al. [ACP ${ }^{+} 17$, Thm. 1].

Theorem 1. Let $n \in \operatorname{poly}(\lambda)$ with $n \geq 8$ and let $\mathcal{A}$ be a deterministic parallel oracle machine that evaluates TDScrypt correctly with probability $\chi(\lambda)$ over the choice of the parameters and input $W$. Then, assuming that factoring is hard, in the GGM and ROM with probability at least $\chi(\lambda)-\epsilon(\lambda)$,

$$
\mathrm{cc}_{\text {mem }}\left(\mathcal{A}^{\Sigma, \mathrm{h}}(\Sigma(1), \Sigma(w))\right) \in \Omega\left(\frac{n^{2}}{\log n} \log N\right)
$$

where $\epsilon(\lambda) \in \operatorname{negl}(\lambda)$, and the probability is taken over $(N, \Sigma, \mathrm{~h}, w) \stackrel{\$}{\leftarrow} \operatorname{params}(\lambda)$.
As stated before, our proof follows Alwen et al.'s [ $\mathrm{ACP}^{+} 17$ ] line of reasoning closely, that is, considering the single-challenge game and then generalizing it to the multi-challenge setting which roughly equals any actual evaluation. While our single-challenge trade-off proof differed considerably, the multi-challenge proof is almost identical. On a high level, the single-challenge trade-offs can simply be swapped out, and this only affects constants. Alwen et al. $\left[\mathrm{ACP}^{+} 17\right]$ took great care to state these explicitly. In contrast, we stick to asymptotics since our primary focus is showing that a TMHF exists.

TdScrypt Definition. Recall the definition of TdScrypt. Given input $\Sigma(w)$, compute the powers $\Sigma\left(2^{i} w\right), 0 \leq i<n$. Then, define $S_{0}=\mathrm{h}\left(\Sigma\left(2^{n} w\right), 0^{\ell}\right)$ and, for $1<i \leq n$, $S_{i}=\mathrm{h}\left(\Sigma\left(2^{j_{i}} w\right), S_{i-1}\right)$ where $j_{i}=S_{i-1} \bmod n$ is the ith challenge. For $k \in[n]$, we define $s_{k}$ to be the round in which the $k$ th challenge is issued, i.e., when the value $S_{k-1}$ is the result of a query.

High-level Proof Strategy. On a high-level, the proof strategy is as follows. We modify Corollary 1 to work with a single adversary $\mathcal{A}$ that can be thought of as a (potentially dishonest) evaluation strategy for TdScrypt. For some $k \in[n]$, we split the execution of $\mathcal{A}$ into two parts, up to round $s_{k}$ when the $k$ th challenge $j_{k}=S_{k-1} \bmod n$ is issued and afterward. These parts map to $\mathcal{A}_{0}$ and $\mathcal{A}_{1}$ as in Corollary 1 where the explicit input $j$ to $\mathcal{A}_{1}$ is replaced by slightly altering (i.e., programming) the random oracle such that $S_{k-1} \bmod n=j$ (but $\left\lfloor S_{k-1} / n\right\rfloor$ is unchanged) and the advice $\mathrm{st}_{0}$ is now the input state $\mathrm{st}_{s_{k}}$. So we get a time-memory trade-off in the style of Corollary 1 that intuitively states: For at least half of the $n$ possible values $j_{k}$ may take, either $\mathcal{A}$ needs a lot of rounds to answer the challenge, or all the input state st $r_{r}$ with $r \leq s_{k}$ have a size $\left\|\mathrm{st}_{r}\right\|$ that is increasing with $r$ and reasonably large when $r$ is closer to $s_{k}$. Then, we can argue that this trade-off must hold for roughly half of all $n$ challenges by Hoeffding's inequality, and we can thus add the input state sizes $\|$ st $_{r} \|$ for all $n$ challenges to get the desired memory complexity.

Proof Sketch. As a first step, we characterize parameter combinations that are not amenable to our proof strategy.

Definition 9 (Collisions). For given $N, \Sigma$, and $w$, the set colliding ${ }_{k}$ contains all random oracles h that cause a collision amongst the preceding $\left\{S_{0}, \ldots, S_{k-1}\right\}$ during the honest $\operatorname{Eval}(\mathrm{pp}, \Sigma(w))$.

Definition 10 (Rounding Impossible). For given $N, \Sigma$, and $w$, the set $\begin{aligned} \\ \text { rounding }\end{aligned}{ }_{k}(n)$ contains all random oracles $h$ that are not amenable to programming some challenge $j \in[0, n)$. That is, there exists $0 \leq i \leq k$ such that $S_{i}>\left\lfloor 2^{\ell} / n\right\rfloor n-1$.

Intuitively, the above two definitions characterize random oracles where we cannot program $j_{k}=S_{k-1} \bmod n$ uniformly. Either the value is not independent of previously issued challenges due to collisions, or $S_{k-1}$ might overflow for some choices of $j_{k}$. The next definition then tells us that any fixed adversary $\mathcal{A}$ almost always learns the values $\left(S_{i}\right)_{0 \leq i \leq n}$ in sequential order.

Definition 11 (Wrong Evaluation Order). For given $N, \Sigma, w$, and $\mathcal{A}$, the set wrongOrder $_{k}$ contains all random oracles h where during the execution of $\mathcal{A}^{\Sigma, \mathrm{h}}(\Sigma(w))$ a query to h of the form $\left(\cdot, S_{j}\right)$ occurs before one of the form $\left(\cdot, S_{i}\right)$ for $0 \leq i<j<k$.

With the above definitions in place, we can now define the time-memory trade-off. Essentially, this is a modified version of the single-challenge game (cf. Definition 6) merged with the time-memory trade-off Lemma 1 (resp. Corollary 1).

Definition 12 (Hard Challenge). Let $\lambda, n \in \mathbb{N}$ with $n \geq 8, k, j \in[0, n),(N, \Sigma, \mathrm{~h}, w) \in$ params $(\lambda)$ with $\mathrm{h} \notin$ colliding $_{k-1} \cup$ rounding $_{k-1}(n)$, and $\mathcal{A}$ a deterministic oracle machine making at most $Q \in \operatorname{poly}(\lambda)$ queries. Consider the execution $\mathcal{A}^{\Sigma, \mathrm{h}}(\Sigma(1), \Sigma(w))$ up to round $s_{k} \in \mathbb{N}^{+}$, which is the earliest round where the query $\left(W^{2^{j_{k-1}}}, S_{k-2}\right) \in$ qrs $^{\mathrm{h}}$ with $j_{k-1}=S_{k-2} \bmod n$ occurs (in the case $k=1$, the query is $\left(W^{2^{n}}, 0^{\ell}\right)$ ). For any round $r$ with $0 \leq r \leq s_{k}$, let $M_{r}=\left\|\mathrm{st}_{r}\right\|$, and associate a fixed value $\rho_{r}$ to it similarly to Lemma 1. ${ }^{13}$ It follows that

$$
\rho_{r} \in \Theta\left(M_{r}+\log n+\log Q+\log \lambda\right) /(\log (N)-\log Q-\log n) .
$$

Define $\mathrm{h}^{\prime}=\mathrm{h}$ except for $S_{k-1}$ which will be changed to $S_{k-1}^{\prime}$ with $\left\lfloor S_{k-1}^{\prime} / n\right\rfloor=\left\lfloor S_{k-1} / n\right\rfloor$ but $S_{k-1}^{\prime} \bmod n=j$, and consider the execution $\mathcal{A}^{\Sigma, \mathrm{h}^{\prime}}(\Sigma(1), \Sigma(x))$. Let $t_{k, j, r} \in \mathbb{N}^{+}$be minimal such that the query $\left(W^{2^{j}}, S_{k-1}\right) \in$ qrs $^{h^{\prime}}$ occurs after $s_{k}$ in round $r+t_{k, j, r}>s_{k}$ where $t_{k, j, r}=\infty$ if it never does. Then, let

$$
r_{k}=\underset{0 \leq r \leq s_{k}}{\arg \max }\left\{\frac{n}{6 \rho_{r} \log (n / 2)}-\left(s_{k}-r\right)\right\},
$$

and define

$$
\operatorname{Hard}_{\mathcal{A}(w)}^{\Sigma, \mathrm{h}, n}(k, j)=1 \text { if and only if } t_{k, j, r_{k}}>\frac{n}{6 \rho_{r_{k}} \log (n / 2)}
$$

We will re-use the notation introduced in these definitions and use it to state the modified single-challenge trade-off. Essentially, it says that most parameters do not fall into colliding, rounding or wrongOrder, and that the time-memory trade-off holds for them.

Lemma 5. For every deterministic parallel oracle machine $\mathcal{A}$, there exist negligible functions $\epsilon(\lambda) \in \operatorname{negl}(\lambda)$ such that, for all $n \in \operatorname{poly}(\lambda)$ with $n \geq 8$, there exists a subset $\operatorname{good} \subseteq \operatorname{params}(\lambda)$ such that, for every $(N, \Sigma, \mathrm{~h}, w) \in \operatorname{good}$ and $k \in \mathbb{N}$ with $1 \leq k \leq n$, $|\operatorname{good}| \geq(1-\epsilon(\lambda))|\operatorname{params}(\lambda)|$ and the following holds:

- $\mathrm{h} \notin$ colliding $_{k-1}$;
- $\mathrm{h} \notin$ rounding $_{k-1}(n)$;
- $\mathrm{h} \notin$ wrongOrder $_{k-1}$; and
- $\operatorname{Pr}_{j,[0, n)}\left[\operatorname{Hard}_{\mathcal{A}(w)}^{\Sigma, \mathrm{h}, n}(k, j)\right] \geq 1 / 2$.

[^8]Proof (Sketch). As in the proof of Lemma 1, towards contradiction, we assume that for some $n$ and $k$, bad is large, i.e., $|\operatorname{bad}| \geq \lambda^{-c} \mid \operatorname{params}(\lambda)$, and we consider four cases. We will show that every case leads to a contradiction.

- $\mathrm{h} \in$ colliding $_{k-1}$ : View $N, \Sigma$ and $w$ as fixed and count how many choices of h cause a collision. There are at most $k^{2} \leq n^{2}$ possible colliding pairs $S_{i}$ and $S_{j}$ and, for every pair, a collision happens for at most a $2^{-\log N}$ fraction of random oracles as h's output length is at least $\log N$. So, at most an $n^{2} / 2^{\log N}$ fraction of random oracles might cause a collision which is negligible and therefore contradicts the assumption that bad covers a polynomial fraction of params (cf. [ACP ${ }^{+} 17$, Clm. 15]).
- $\mathrm{h} \in$ rounding $_{k-1}$ and $\mathrm{h} \in$ wrongOrder $_{k-1}$ : Both analyses are analogous the previous case with the fraction of random oracles being in the order of $1 / N$ as well (cf. $\left[\mathrm{ACP}^{+} 17\right.$, Clms. $\left.16 \& 18\right]$ ).
- $\operatorname{Pr}_{j \leftarrow[n]}\left[\operatorname{Hard}_{\mathcal{A}(w)}^{\Sigma, \mathrm{h}, n}(k, j)\right]<1 / 2$ : This case follows the proof of Lemma 1 with some modifications. Instead of considering explicit preprocessing- and online algorithms $\left(\mathcal{A}_{0}, \mathcal{A}_{1}\right)$, we split the execution of $\mathcal{A}$ into two parts-up to round $r$ (preprocessing) and everything afterward (online). In this setting, we cannot pass the challenge $j$ to the online algorithm explicitly anymore. Instead, we will program the random oracle h to include the challenge implicitly, i.e., setting $S_{k-1} \bmod n=j$. Programming is always possible as long as $\mathrm{h} \notin$ colliding $_{k-1} \cup$ rounding $_{k-1}(n)$, and we may assume that this is the case by the preceding case analyses. For Dec to know when to program $\mathrm{h}, \log Q$ additional bits of hint are required to recognize the query $\left(W^{2^{j_{k-1}}}, S_{k-2}\right)$ (cf. Definition 12). Apart from these modifications, the proof is identical to the one of Lemma 1.

This discussion completes the proof sketch.
We now use Lemma 5's trade-off to lower bound the sum $\sum_{r=1}^{s_{n}+1} \rho_{r}$ where we note that $\rho_{r}$ is a quantity related to the size of the input state $\left\|\mathrm{st}_{r}\right\|$. Note that the sum covers all rounds up to round $s_{n}+1$, i.e., the round in which $\mathcal{A}$ first computes $S_{n}$, the output of TdScrypt. This strategy is taken from [ACP ${ }^{+} 17$, Sec. 5] so we refer interested readers there for more details.

Claim 6 (Lower Bound on $\left.\sum \rho_{r}\right)$. If $(N, \Sigma, \mathrm{~h}, w) \in \operatorname{good}$ and $\mathcal{A}^{\Sigma, \mathrm{h}}(\Sigma(1), \Sigma(w))$ queries $S_{n}$, then $\sum_{r=1}^{s_{n}+1} \rho_{r} \in \Omega\left(n^{2} / \log n\right)$.
Proof (Sketch). First, note that if the parameters are in good, whenever $\mathcal{A}$ queries $S_{n}$, it must have received all $\left\{S_{i}\right\}_{0 \leq i \leq n}$ during its execution, and it must receive them in order. Second, if $\operatorname{Hard}_{\mathcal{A}(w)}^{\Sigma, \mathrm{h}, n}(k, j)=1$, then, for any $0 \leq r \leq s_{k}$,

$$
\begin{equation*}
t_{k, j, r}>\frac{n}{6 \rho_{r} \log (n / 2)} \tag{3}
\end{equation*}
$$

by the choice of $r_{k}\left(\right.$ cf. $\left[\mathrm{ACP}^{+} 17\right.$, Clms. $\left.\left.5 \& 10\right]\right)$. Third, by a generalization of Hoeffding's inequality (cf. [ACP $\left.{ }^{+} 17, \mathrm{Clm} .7\right]$ ), we argue that for any fixing of challenges, at least $n(1 / 2-\epsilon)$ challenges are Hard where $\epsilon>0$ (cf. [ACP ${ }^{+} 17$, Clm. 19]).

Given the above, we apply $\left[\mathrm{ACP}^{+} 17, \mathrm{Clm} .8\right]$ to $\rho_{r}>\frac{n}{6 t_{k, j, r} \log (n / 2)}$, which is Equation (3) rearranged. This gets us $\sum_{r=1}^{s_{n}+1} \rho_{r} \in \Omega\left(n^{2} / \log n\right)$ as desired (cf. [ACP ${ }^{+} 17$, Clm. 11]).

Finally, we prove Theorem 1 by converting the quantity $\sum \rho_{r}$ to memory complexity.
Proof of Theorem 1 (Sketch). By Lemma 5, an overwhelming fraction of parameters is in good, and $\mathcal{A}$ queries $S_{n}$ with probability $\chi(\lambda)$. Applying Claim 6, we get a lower bound of $\sum_{r=1}^{s_{n}+1} \rho_{r} \in \Omega\left(n^{2} / \log n\right)$. Since $M_{r} \in \Omega\left(\rho_{r} \log N\right)$ for all $r \in \mathbb{N}$ by definition, it follows that $\sum_{r=1}^{s_{n}+1} M_{r} \in \Omega\left(\frac{n^{2}}{\log n} \log N\right)$. Note that this sum is equivalent to the $\mathrm{cc}_{\text {mem }}$ of $\mathcal{A}$ which completes the proof sketch.

Acknowledgements. We thank the Eurocrypt reviewers for their thorough review and for pointing out related works. This research was funded in whole or in part by the Austrian Science Fund (FWF) 10.55776/F85.

## References

$\left[\mathrm{AAC}^{+} 17\right]$ Hamza Abusalah, Joël Alwen, Bram Cohen, Danylo Khilko, Krzysztof Pietrzak, and Leonid Reyzin. Beyond hellman's time-memory trade-offs with applications to proofs of space. In Tsuyoshi Takagi and Thomas Peyrin, editors, ASIACRYPT 2017, Part II, volume 10625 of LNCS, pages 357-379. Springer, Heidelberg, December 2017.
[ABB22] Mohammad Hassan Ameri, Alexander R. Block, and Jeremiah Blocki. Memoryhard puzzles in the standard model with applications to memory-hard functions and resource-bounded locally decodable codes. In Clemente Galdi and Stanislaw Jarecki, editors, Security and Cryptography for Networks, pages 45-68, Cham, 2022. Springer International Publishing.
[ABH17] Joël Alwen, Jeremiah Blocki, and Ben Harsha. Practical graphs for optimal side-channel resistant memory-hard functions. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, ACM CCS 2017, pages 1001-1017. ACM Press, October / November 2017.
[ABP17] Joël Alwen, Jeremiah Blocki, and Krzysztof Pietrzak. Depth-robust graphs and their cumulative memory complexity. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, EUROCRYPT 2017, Part III, volume 10212 of LNCS, pages 3-32. Springer, Heidelberg, April / May 2017.
[ABP18] Joël Alwen, Jeremiah Blocki, and Krzysztof Pietrzak. Sustained space complexity. In Jesper Buus Nielsen and Vincent Rijmen, editors, EUROCRYPT 2018, Part II, volume 10821 of LNCS, pages 99-130. Springer, Heidelberg, April / May 2018.
[ABZ20] Mohammad Hassan Ameri, Jeremiah Blocki, and Samson Zhou. Computationally data-independent memory hard functions. In Thomas Vidick, editor, ITCS 2020, volume 151, pages 36:1-36:28. LIPIcs, January 2020.
$\left[\mathrm{ACK}^{+} 16\right]$ 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. In Marc Fischlin and Jean-Sébastien

Coron, editors, EUROCRYPT 2016, Part II, volume 9666 of $L N C S$, pages 358-387. Springer, Heidelberg, May 2016.
$\left[\mathrm{ACP}^{+} 17\right]$ Joël Alwen, Binyi Chen, Krzysztof Pietrzak, Leonid Reyzin, and Stefano Tessaro. Scrypt is maximally memory-hard. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, EUROCRYPT 2017, Part III, volume 10212 of LNCS, pages 33-62. Springer, Heidelberg, April / May 2017.
[AS15] Joël Alwen and Vladimir Serbinenko. High parallel complexity graphs and memory-hard functions. In Rocco A. Servedio and Ronitt Rubinfeld, editors, $4^{7}$ th ACM STOC, pages 595-603. ACM Press, June 2015.
[BCS16] Dan Boneh, Henry Corrigan-Gibbs, and Stuart E. Schechter. Balloon hashing: A memory-hard function providing provable protection against sequential attacks. In Jung Hee Cheon and Tsuyoshi Takagi, editors, ASIACRYPT 2016, Part I, volume 10031 of LNCS, pages 220-248. Springer, Heidelberg, December 2016.
[BDK16] Alex Biryukov, Daniel Dinu, and Dmitry Khovratovich. Argon2: New generation of memory-hard functions for password hashing and other applications. In 2016 IEEE European Symposium on Security and Privacy (EuroS\&P), pages 292-302, 2016.
[BH22] Jeremiah Blocki and Blake Holman. Sustained space and cumulative complexity trade-offs for data-dependent memory-hard functions. In Yevgeniy Dodis and Thomas Shrimpton, editors, CRYPTO 2022, Part III, volume 13509 of LNCS, pages 222-251. Springer, Heidelberg, August 2022.
[ $\left.\mathrm{BHK}^{+} 19\right]$ Jeremiah Blocki, Benjamin Harsha, Siteng Kang, Seunghoon Lee, Lu Xing, and Samson Zhou. Data-independent memory hard functions: New attacks and stronger constructions. In Alexandra Boldyreva and Daniele Micciancio, editors, CRYPTO 2019, Part II, volume 11693 of LNCS, pages 573-607. Springer, Heidelberg, August 2019.
[BLP23] Alex Biryukov and Marius Lombard-Platet. Pured: A unified framework for resource-hard functions. Cryptology ePrint Archive, Paper 2023/1809, 2023. https://eprint.iacr.org/2023/1809.
[BP17] Alex Biryukov and Léo Perrin. Symmetrically and asymmetrically hard cryptography. In Tsuyoshi Takagi and Thomas Peyrin, editors, ASIACRYPT 2017, Part III, volume 10626 of LNCS, pages 417-445. Springer, Heidelberg, December 2017.
[BRZ18] Jeremiah Blocki, Ling Ren, and Samson Zhou. Bandwidth-hard functions: Reductions and lower bounds. In David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang, editors, ACM CCS 2018, pages 1820-1836. ACM Press, October 2018.
[CDD23] Simone Costa, Marco Dalai, and Stefano Della Fiore. Variations on the Erdős distinct-sums problem. Discrete Applied Mathematics, 325:172-185, 2023.
[CK18] Henry Corrigan-Gibbs and Dmitry Kogan. The discrete-logarithm problem with preprocessing. In Jesper Buus Nielsen and Vincent Rijmen, editors, EUROCRYPT 2018, Part II, volume 10821 of LNCS, pages 415-447. Springer, Heidelberg, April / May 2018.
[DN93] Cynthia Dwork and Moni Naor. Pricing via processing or combatting junk mail. In Ernest F. Brickell, editor, CRYPTO'92, volume 740 of LNCS, pages 139-147. Springer, Heidelberg, August 1993.
[DTT10] Anindya De, Luca Trevisan, and Madhur Tulsiani. Time space tradeoffs for attacks against one-way functions and PRGs. In Tal Rabin, editor, CRYPTO 2010, volume 6223 of $L N C S$, pages 649-665. Springer, Heidelberg, August 2010.
[Guy94] Richard K. Guy. Unsolved Problems in Number Theory. Problem Books in Mathematics. Springer, New York, NY, 2 edition, 1994.
[Her09] Charles Hermite. Sur l'introduction des variables continues dans la théorie des nombres, volume 1 of Cambridge Library Collection - Mathematics, page 164-192. Cambridge University Press, 2009.
[Kal00] Burt Kaliski. Pkcs\# 5: Password-based cryptography specification version 2.0. Request for Comments 2898, Internet Engineering Task Force, September 2000.
[KL14] Jonathan Katz and Yehuda Lindell. Introduction to Modern Cryptography, Second Edition. Chapman \& Hall/CRC, 2nd edition, 2014.
[KLX20] Jonathan Katz, Julian Loss, and Jiayu Xu. On the security of time-lock puzzles and timed commitments. In Rafael Pass and Krzysztof Pietrzak, editors, TCC 2020, Part III, volume 12552 of $L N C S$, pages 390-413. Springer, Heidelberg, November 2020.
[Laz96] Felix Lazebnik. On systems of linear diophantine equations. Mathematics Magazine, 69(4):261-266, 1996.
[MVOV97] Alfred J. Menezes, Paul C. Van Oorschot, and Scott A. Vanstone. Handbook of applied cryptography. CRC press, 1997.
[MW01] Daniele Micciancio and Bogdan Warinschi. A linear space algorithm for computing the hermite normal form. In Proceedings of the 2001 International Symposium on Symbolic and Algebraic Computation, ISSAC '01, page 231-236, New York, NY, USA, 2001. Association for Computing Machinery.
[Per09] Colin Percival. Stronger key derivation via sequential memory-hard functions. In BSDCan 2009, 2009.
[RD16] Ling Ren and Srinivas Devadas. Proof of space from stacked expanders. In Martin Hirt and Adam D. Smith, editors, TCC 2016-B, Part I, volume 9985 of LNCS, pages 262-285. Springer, Heidelberg, October / November 2016.
[RD17] Ling Ren and Srinivas Devadas. Bandwidth hard functions for ASIC resistance. In Yael Kalai and Leonid Reyzin, editors, TCC 2017, Part I, volume 10677 of LNCS, pages 466-492. Springer, Heidelberg, November 2017.
[Rot22] Lior Rotem. Revisiting the uber assumption in the algebraic group model: Finegrained bounds in hidden-order groups and improved reductions in bilinear groups. In Dana Dachman-Soled, editor, 3rd Conference on InformationTheoretic Cryptography (ITC 2022), volume 230 of Leibniz International Proceedings in Informatics (LIPIcs), pages 13:1-13:13, Dagstuhl, Germany, 2022. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
[RS20] Lior Rotem and Gil Segev. Generically speeding-up repeated squaring is equivalent to factoring: Sharp thresholds for all generic-ring delay functions. In Daniele Micciancio and Thomas Ristenpart, editors, CRYPTO 2020, Part III, volume 12172 of $L N C S$, pages 481-509. Springer, Heidelberg, August 2020.
[RSW96] Ronald L. Rivest, Adi Shamir, and David A. Wagner. Time-lock puzzles and timed-release crypto. Technical report, USA, 1996.
[Sho97] Victor Shoup. Lower bounds for discrete logarithms and related problems. In Walter Fumy, editor, EUROCRYPT'97, volume 1233 of LNCS, pages 256-266. Springer, Heidelberg, May 1997.
[vzGS13] Joachim von zur Gathen and Igor E. Shparlinski. Generating safe primes. Journal of Mathematical Cryptology, 7(4):333-365, 2013.
[Yao90] Andrew Chi-Chih Yao. Coherent functions and program checkers (extended abstract). In 22nd ACM STOC, pages 84-94. ACM Press, May 1990.

## A Detailed Cumulative Memory Complexity Analysis

To analyze the cumulative memory complexity (CMC) of TDSCRYPT's Eval and TDEval, we first discuss the basic operations performed by them before describing their CMC.

Basic Operations. The relevant parameters are the RSA modulus $N^{\prime}$, the group size $N$, the output length $\omega_{h}$ of the hash function $h$, and the number of iterations $n$. Note that since $N \approx N^{\prime} / 4$ it follows that $N$ and $N^{\prime}$ have essentially the same bit length. As stated in Section 2, we further have that $\omega_{\mathrm{h}} \in \Theta(\log (N))$ and $n \ll N$ since $n \in \operatorname{poly}(\lambda)$ (as otherwise Eval would not run in poly $(\lambda)$ ).

We first give an overview of the operations used in Eval and TDEval, listing the amount of bit operations, i.e., additions and multiplications in the field $\mathbb{F}_{2}$, and the scratch space in terms of bits used in their computation. The results are summarized in Table 2.

Algorithm Eval performs group operations in $\mathbb{Q R}_{N^{\prime}}($ Line 09$)$, evaluates hash function $h$ (Lines 11 and 14), and reduces outputs of $h$ modulo $n$ (Line 13). Recall that, by definition, evaluating $h$ requires the same computational effort as computing a group operation, i.e., $\Theta\left(\log \left(N^{\prime}\right)^{2}\right)$ bit operations while using $\Theta\left(\log \left(N^{\prime}\right)\right)$ bits of scratch space (e.g., [MVOV97, Sec. 14.3.3]). The computation of $j_{i}:=S_{i-1} \bmod n$, on the other hand,

| Operation | Bit Operations | Scratch Space |
| :--- | ---: | ---: |
| Multiplication in $\mathbb{Q R}_{N^{\prime}}$ | $\Theta\left(\log \left(N^{\prime}\right)^{2}\right)$ | $\Theta\left(\log \left(N^{\prime}\right)\right)$ |
| Exponentiation in $\mathbb{Q R}_{N^{\prime}}$ | $\Theta\left(\log (N) \log \left(N^{\prime}\right)^{2}\right)$ | $\Theta\left(\log \left(N^{\prime}\right)\right)$ |
| Multiplication in $\mathbb{Z}_{N}$ | $\Theta\left(\log (N)^{2}\right)$ | $\Theta(\log (N))$ |
| Evaluation of h | $\Theta\left(\log \left(N^{\prime}\right)^{2}\right)$ | $\Theta\left(\log \left(N^{\prime}\right)\right)$ |
| $S_{i-1} \bmod N$ | $\Theta(\log (n) \log (N))$ | $\Theta(\log (N))$ |
| $2^{j} \bmod n$ | $\Theta\left(\log (n) \log (N)^{2}\right)$ | $\Theta(\log (N))$ |

Table 2: Time and memory usage of operations used in Eval and TDEval. Note that $\log (N) \approx \log \left(N^{\prime}\right)>\log (n)$.
requires $\Theta\left(\log (n) \log \left(N^{\prime}\right)\right)$ bit operations and scratch space of $\log (N)$ bits (e.g., [MVOV97, Sec. 14.2.5]), where we used that $\omega_{\mathrm{h}} \in \Theta(\log (N))$.

TDEval also computes $h$ (Lines 19 and 24) and also reduces them modulo $n$ (Line 21), but it additionally performs exponentiations in $\mathbb{Q} \mathbb{R}_{N^{\prime}}$ (Lines 18 and 23) and reduces $n$-bit integers of the form $2^{j}$ with $j \in[0, n)$ modulo $N$ (Lines 17 and 22). Exponentiations can be computed using square-and-multiply at the cost of $\Theta(\log (N))$ group operations. The operation $m:=2^{j} \bmod n$ can be implemented using a lookup table that, for $i \in$ $[0,\lceil\log (n)\rceil]$, stores the values $2^{2^{i}} \bmod N$. Then, by writing $j:=\sum_{i=0}^{\lceil\log (n)\rceil} b_{i} 2^{i}$, one can compute $m$ using at most $\lceil\log (n)\rceil$ group operations in $\mathbb{Z}_{N}$ as

$$
\prod_{i=0}^{\lceil\log (n)\rceil} b_{i} 2^{2^{i}} \bmod N=2^{\sum_{i=0}^{[\log (n)]} b_{i} 2^{i}} \bmod N=2^{j} \bmod N
$$

In turn, one computation of the form $m:=2^{j} \bmod N$ can be done using $\Theta\left(\log (n) \log (N)^{2}\right)$ bit operations while using $\Theta(\log (N))$ bits of scratch space (not accounting for the storage of the lookup table).

Note that the only operation with a time requirement that does not contain a term of the form $\log \left(N^{\prime}\right)^{2}$ or $\log (N)^{2}$ is the computation of the index $j_{i}:=S_{i-1} \bmod n$ in Lines 13 and 21, but that the number of bit operations required for this operation is actually smaller since $\log (n)<\log (N) \approx \log \left(N^{\prime}\right)$. Thus, in the following computation of Eval and TDEval's cumulative memory complexity we will use $\log (N)^{2} \approx \log \left(N^{\prime}\right)^{2}$, the number of bit operations required for one group operation in $\mathbb{Q R}_{N^{\prime}}$ and $\mathbb{Z}_{N}$, as our unit of time. Accordingly, group operations and hash evaluations require one time unit, exponentiation $\log (N)$ units, and computing $2^{j} \bmod N$ requires $\log (n)$ units.

Cumulative Memory Complexity of Eval. As algorithm Eval's scratch space requirements are of order $\log \left(N^{\prime}\right)$ (see Table 2), its memory usage is dominated by storing the group elements $W_{0}, \ldots, W_{n}$ amounting to $\Theta\left(n \log \left(N^{\prime}\right)\right)$ bits of memory sustained throughout the whole second phase, i.e., Lines $12-14$. In the first phase of the evaluation the algorithm computes $n$ group operations, in the second $n$ hash evaluations and $n$ reductions modulo $n$. As a consequence, its cumulative memory complexity in terms of bit storage and group operations is of order

$$
\mathrm{cc}_{\text {mem }}(\operatorname{Eval}(\mathrm{pp}, x)) \in \Theta\left(n \log \left(N^{\prime}\right) \cdot n\right)=\Theta\left(n^{2} \log \left(N^{\prime}\right)\right)
$$

Cumulative Memory Complexity of TDEval. In order to compute the $W_{i}$ at any point in time, TDEval only has to keep track of at most two group elements and one exponent $m_{i}$. Further, it only requires scratch space of order $\log \left(N^{\prime}\right)$. Thus, its memory is dominated by the lookup table used to perform reductions modulo $N$ which has a size of order $\Theta\left(\log (n) \log \left(N^{\prime}\right)\right)$ bits. In the first phase of the evaluation, the algorithm performs one reduction modulo $N$ as well as one exponentiation in $\mathbb{Q} \mathbb{R}_{N^{\prime}}$. In each iteration in the of the second phase of the evaluation TDEval has to perform the same operations as well as an additional reduction modulo $n$ (line 21), thus requiring essentially $\log (N)+\log (n)$ group operations. Thus, its cumulative memory complexity in terms of bit storage and group operations amounts to

$$
\begin{aligned}
\mathrm{cc}_{\text {mem }}(\text { TDEval }(\mathrm{pp}, \mathrm{td}, x)) & \in \Theta\left(\log (n) \log \left(N^{\prime}\right) \cdot n\left(\log \left(N^{\prime}\right)+\log (n)\right)\right) \\
& =\Theta\left(n \log \left(N^{\prime}\right)^{2} \log (n)\right)
\end{aligned}
$$

Summing up, TDSCRYPT is a $\left(n^{2} \log \left(N^{\prime}\right), n \log (n) \log \left(N^{\prime}\right)^{2}\right)$-TMHF.

## B Bounding the Entries of $A$

In this section we analyze the maximal size that the entries of matrix $A \in \mathbb{Z}^{\ell \times m}$ induced by the query behavior of $\mathcal{A}_{1}$ to the group-operation oracle $\mathcal{G}$ can take. Recall that (as discussed in more detail in Section 5.1 and made formal in Appendix C.2) our encoder assigns all labels appearing as input to or output of $\mathcal{G}$ an algebraic representation of the form $\sum_{i=1}^{m} \boldsymbol{a}_{i} \boldsymbol{x}_{\boldsymbol{i}}$ with $\boldsymbol{a} \in \mathbb{Z}^{m}$ which is stored in table $T$. Here, the indeterminates $\boldsymbol{x}$ correspond to the labels being queried out of the blue. Matrix $A$ collects the representations of labels corresponding to challenges that are successfully answered within $t$ rounds of parallel queries.

In the following we characterize the set of representations that possibly can be added to the table $T$. More precisely, using a simple abstraction we inductively define sets $R_{t}$ that for $t \in \mathbb{N}^{+}$contain all $\boldsymbol{a} \in \mathbb{Z}^{m}$ that can be added to $T$ within $t$ rounds of queries to $\mathcal{G}$ with unbounded parallelism.

As a special case, $R_{0}$ is to be understood as all representations corresponding to labels that are used as input to $\mathcal{G}$ within the first round of queries. Note that all of these labels must be out of the blue. The $i$ th label queried out of the blue corresponds to indeterminate $\boldsymbol{x}_{i}$. Thus, its representation is the $i$ th unit vector $\boldsymbol{e}_{i} \in \mathbb{Z}^{m}$. We define

$$
R_{0}:=\left\{ \pm \boldsymbol{e}_{i} \mid i \in\{1, \ldots, m\}\right\} \cup\{\mathbf{0}\} .
$$

Note that $R_{0}$ contains the representations of all labels appearing out of the blue (irrespective of the actual round in which they are first queried) as well the ones of their inverses and the neutral element. This choice of $R_{0}$ overestimates the capabilities of $\mathcal{A}_{1}$ but will allow us to derive a clean characterization of the sequence of sets $R_{t}$.

We now turn to the definition of set $R_{t}$ for $t \geq 1$. The representation of the label output in response to a query to $\mathcal{G}$ is given by the sum, difference, or inverse of the label(s) used as input, depending on whether the used operation was,+- , or inv. Accordingly, for $t \geq 1$ we define

$$
R_{t}:=\left\{\boldsymbol{a}_{\mathbf{1}} \pm \boldsymbol{a}_{\mathbf{2}} \mid \boldsymbol{a}_{\mathbf{1}}, \boldsymbol{a}_{\mathbf{2}} \in R_{t-1}\right\} \cup R_{t-1},
$$

i.e., we add to $R_{t-1}$ all vectors that can be obtained by adding or subtracting two vectors that are reachable within $t-1$ rounds. Note that this in particular captures inversion as $\mathbf{0} \in R_{t-1}$. Further, any upper bound on the elements of $R_{t}$ also applies to the entries of $A$, since the inclusion of additional elements in $R_{0}$ only improves on $\mathcal{A}_{1}$ 's computational capabilities.

To bound the elements of $R_{t}$ note that with every additional query to $\mathcal{G}$ every vector entry can at most double. This immediately implies $\|\boldsymbol{a}\|_{\infty} \leq 2^{t}$ for all $\boldsymbol{a} \in R_{t}$. While this observation is already sufficient for the proof of our lower bound, a more careful analysis yields the following exact characterization of $R_{t}$ in terms of the one-norm.
Lemma 6. Let $m, t \in \mathbb{N}$ and $R_{t}$ be defined as above. Then for $\boldsymbol{a} \in \mathbb{Z}^{m}$ we have

$$
\boldsymbol{a} \in R_{t} \quad \Longleftrightarrow \quad\|\boldsymbol{a}\|_{1} \leq 2^{t} .
$$

Proof. We prove the result by induction on $t$. The statement holds for $t=0$ as we have $R_{0}=\left\{ \pm \boldsymbol{e}_{\boldsymbol{i}} \mid i \in\{1, \ldots, m\}\right\} \cup\{\mathbf{0}\}$, the set of integer vectors with one-norm bounded by 1.

Case $\Longrightarrow$ : Assume the statement is true for all values up to $t$. First, consider $\boldsymbol{a} \in \mathbb{Z}^{m}$ with $\|\boldsymbol{a}\|_{1}>2^{t+1}$ and consider arbitrary $\boldsymbol{b}, \boldsymbol{c} \in \mathbb{Z}^{m}$ satisfying $\boldsymbol{a}=\boldsymbol{b} \pm \boldsymbol{c}$. We have $2^{t+1}<\|\boldsymbol{a}\|_{1} \leq\|\boldsymbol{b}\|_{1}+\|\boldsymbol{c}\|_{1}$ implying that $\|\boldsymbol{b}\|_{1}>2^{t}$ or $\|\boldsymbol{c}\|_{1}>2^{t}$. By the induction hypothesis, one of the two vectors cannot have been computed within $t$ group operations with unbounded parallelism. In turn, $\boldsymbol{a}$ cannot have been computed within $t+1$ group operations.

Case $\Longleftarrow$ : Consider $\boldsymbol{a} \in \mathbb{Z}^{m}$ with $\|\boldsymbol{a}\|_{1} \leq 2^{t+1}$ and let $S_{1}, S_{2} \subseteq\{1, \ldots, m\}$ be sets of (almost) equal size jointly containing all odd entries of $\boldsymbol{a}$. More precisely, we require $\left|S_{2}\right| \leq\left|S_{1}\right| \leq\left|S_{2}\right|+1$ and that $\boldsymbol{a}_{i}$ is odd exactly if $i \in S_{1} \cup S_{2}$. We define vectors $\boldsymbol{b}, \boldsymbol{c} \in \mathbb{N}^{m}$ as

$$
\boldsymbol{b}_{i}=\left\{\begin{array}{ll}
\operatorname{sgn}\left(\boldsymbol{a}_{i}\right) \cdot\left\lceil\left|\boldsymbol{a}_{i}\right| / 2\right\rceil & \text { if } i \in S_{1} \\
\operatorname{sgn}\left(\boldsymbol{a}_{i}\right) \cdot\left\lfloor\left|\boldsymbol{a}_{i}\right| / 2\right\rfloor & \text { if } i \in S_{2} \\
\boldsymbol{a}_{i} / 2 & \text { else }
\end{array} \quad \text { and } \quad \boldsymbol{c}_{i}= \begin{cases}\operatorname{sgn}\left(\boldsymbol{a}_{i}\right) \cdot\left\lfloor\left|\boldsymbol{a}_{i}\right| / 2\right\rfloor & \text { if } i \in S_{1} \\
\operatorname{sgn}\left(\boldsymbol{a}_{i}\right) \cdot\left\lceil\left|\boldsymbol{a}_{i}\right| / 2\right\rceil & \text { if } i \in S_{2}, \\
\boldsymbol{a}_{i} / 2 & \text { else }\end{cases}\right.
$$

where sgn denotes the sign of an integer. Note that by definition of $\boldsymbol{b}$ and $\boldsymbol{c}$ it holds that $\boldsymbol{a}=\boldsymbol{b}+\boldsymbol{c}$. Further, we have that

$$
\begin{aligned}
\|\boldsymbol{b}\|_{1}=\sum_{i=1}^{m}\left|\boldsymbol{b}_{i}\right| & =1 / 2 \cdot\left(\left|S_{1}\right|-\left|S_{2}\right|\right)+\sum_{i=1}^{m}\left|\boldsymbol{a}_{i}\right| / 2 \\
& =1 / 2 \cdot\left(\left(\left|S_{1}\right|-\left|S_{2}\right|\right)+\|\boldsymbol{a}\|_{1}\right)
\end{aligned}
$$

If $\|\boldsymbol{a}\|_{1}$ is even we have $\left|S_{1}\right|=\left|S_{2}\right|$ and thus $2^{t} \geq\|\boldsymbol{a}\|_{1} / 2=\|\boldsymbol{b}\|_{1}=\|\boldsymbol{c}\|_{1}$. On the other hand, if $\|\boldsymbol{a}\|_{1}$ is odd, we have $2^{t+1} \geq\|\boldsymbol{a}\|_{1}+1$ and $\left|S_{1}\right|=\left|S_{2}\right|+1$ which implies $2^{t} \geq 1 / 2\left(\|a\|_{1}+1\right)=\|\boldsymbol{b}\|_{1}=\|\boldsymbol{c}\|_{1}+1$. Thus, by the induction hypothesis we have $\boldsymbol{b}, \boldsymbol{c} \in R_{t}$ and in turn $\boldsymbol{a} \in R_{t+1}$.

## C Proof of Claim 5

For the full proof, we first state Enc's pseudocode, explain it, and finally analyze its encoding size and success probability. We omit Dec's pseudocode since it follows naturally from Enc.

## C. 1 Notation and Data Structures

In the following, let $\operatorname{Im}(f)$ denote the image of the function $f$, and let $\mathfrak{x}_{1}, \mathfrak{x}_{2}, \ldots$ be indeterminates ${ }^{14}$ in $\mathbb{Z}$. Further, we use the shorthand $\hat{o}:=\circ \bmod N$ where $\circ$ might be a vector, matrix, indeterminate, etc. To signify when and why Enc aborts, we write abort ${ }_{\text {reason }}$.

During Enc's execution, it constructs five datastructures: $I$, a list of hints; $\Xi$, a first-in-first-out queue; $T$, a table; $S$, a map of substitutions; and $C$, a list of collision pointers. We use the notation $x \rightarrow Y$ to denote that element $x$ is appended or pushed to datastructure $Y$.

List of Hints $I$. Dec needs a hint to recognize when $\mathcal{A}_{1}\left(\mathrm{st}_{0}, j\right)$ queries the label corresponding to $2^{j} w$ to $h$ for the first time. We are guaranteed that at least $n / 2$ instances of $\mathcal{A}_{1}$ perform this query within $t$ steps and that matrix $A \in \mathbb{Z}^{(n / 2) \times m}$ induced by $\mathcal{A}_{1}$ 's queries has rank ${ }_{\min } \leq m$. In the remainder of this proof, we will not work with $A$ but with a different matrix $A^{\prime} \in \mathbb{Z}^{\text {rank }_{\text {min }} \times m}$ because encoding $A$ would be too costly for the compression argument. $A^{\prime}$ consists of the rows of $A$ that contain a pivot element in $H$, the Hermite normal form (cf. Definition 7) of $A$. Note that $\operatorname{rank}\left(A^{\prime}\right)=\operatorname{rank}(A)=\operatorname{rank}_{\text {min }}$ which follows from the HNF. Indeed, the HNF has rank min non-zero columns and hence rank $_{\text {min }}$ pivot elements. Simple, inefficient HNF algorithms iterate over the rows once, top to bottom, while only touching rows that will contain a pivot in the end. ${ }^{15}$

So for Dec to reconstruct $A^{\prime}$ from $\mathcal{A}_{1}$ 's queries, it needs rank ${ }_{\text {min }}$ many hints. In the remainder of this proof, we will not distinguish between $A$ and $A^{\prime}$ but simply call it $A \in \mathbb{Z}^{\mathrm{rank}_{\min } \times m}$.

Queue $\Xi$. The queue contains labels that Dec needs to, e.g., correctly respond to $\mathcal{A}_{1}$ 's queries. Enc fills $\Xi$ with this information during its execution and, in the final step, appends $\Xi$ to the encoding. The elements pushed to the queue vary in length, but this does not lead to ambiguities during decoding. In particular, we ensure that Dec always knows this length in advance before popping an element from the queue.

Table $T$. Intuitively, $T$ helps Dec answer repeat queries or queries where previous responses uniquely determine the answer. It does so by mapping labels to algebraic representations over $\mathbb{Z}_{N}$, the domain of $\Sigma$, analogously to Section 5.1. In addition, $T$ also stores an algebraic representation over $\mathbb{Z}$, enabling us to connect the compression argument to the query behavior of $\mathcal{A}_{1}$. This is necessary since $\mathcal{A}_{1}$ does not know to the group order $N$ and therefore its queries must be interpreted over $\mathbb{Z}$. To summarize, $T$ contains three columns: The rightmost column contains labels $\sigma_{i} \in \operatorname{Im}(\Sigma)$; the leftmost

[^9]column contains linear terms in the indeterminates $\hat{\mathfrak{x}}_{j}$ over $\mathbb{Z}_{N}$ that represent a unique algebraic representation of $\sigma_{i}$; and the middle column contains an algebraic (but not necessarily unique) representation of $\sigma_{i}$ in the form of linear terms in the indeterminates $\mathfrak{x}_{j}$ over $\mathbb{Z}$ (see Table 3 for an example).

| Representation over $\mathbb{Z}_{N}$ | Representation over $\mathbb{Z}$ | Label |
| :---: | :---: | :---: |
| $2 \hat{\mathfrak{x}}_{10}$ | $\cdots$ |  |
| $4 \hat{\mathfrak{x}}_{5}$ | $2 \mathfrak{x}_{10}$ | $\sigma_{20}$ |
| $\cdots$ |  |  |
|  | $\mathfrak{x}_{5}+\mathfrak{x}_{7}$ | $\sigma_{32}$ |

Table 3: Exemplary table $T$. Note that the second populated row shows the case where a collision has been resolved with $\hat{\mathfrak{x}}_{7}=3 \hat{\mathfrak{x}}_{5}$.

We will ensure that the leftmost (i.e., the representation over $\mathbb{Z}_{N}$ ) and the rightmost (i.e., the label) columns will not contain duplicates. This allows for convenient notation regarding $T$. We write $\sigma_{i} \in T$ to check if there exists a row containing label $\sigma_{i}$ and a row can be selected by either $T(\sigma)$ or $T(x)$ where $x$ is an algebraic representation over $\mathbb{Z}_{N}$ (which is unique and thus unambiguous). A particular column is denoted by the subscripts $\mathbb{Z}_{N}, \mathbb{Z}$, and $\mathcal{L}$. For example, in Table $3, T_{\mathbb{Z}_{N}}\left(\sigma_{32}\right)$ returns $4 \hat{\mathfrak{x}}_{5}$.

Mapping of Substitutions $S$. Recall that $T$ must not contain any duplicate labels. However, $\mathcal{A}_{1}$ might perform queries such that it arrives at the same label in two different ways. So there exist two distinct algebraic representations over $\mathbb{Z}_{N}$ for the same labela collision. Instead of adding a duplicate row, Enc uses the collision to express one indeterminate $\hat{\mathfrak{x}}_{i}$ in terms of its preceding indeterminates $\hat{\mathfrak{x}}_{j}, j<i$. That is, $\hat{\mathfrak{x}}_{i}=\sum_{j<i} c_{j} \hat{\mathfrak{x}}_{j}$ with $c_{j} \in \mathbb{Z}_{N}$. Subsequently, it replaces $\hat{\mathfrak{x}}_{i}\left(\right.$ over $\left.\mathbb{Z}_{N}\right)$ in the leftmost column of $T$. Note that it cannot replace the corresponding indeterminate $\mathfrak{x}_{i}$ (over $\mathbb{Z}$ ) in the middle column. To handle this, $S$ maps $\hat{\mathfrak{x}}_{i}$ to $\sum_{j<i} c_{j} \hat{\mathfrak{j}}_{j}$ keeping track of all substiutions. We write $\hat{\mathfrak{x}}_{i} \in S$ to check whether a mapping exists.

List of Collision Pointers $C$. Enc must tell Dec when the aforementioned collisions take place since Dec cannot detect them on its own. Thus, it stores $C$, the list of collision pointers, in the encoding. To this end, Enc encodes $C$ 's length $|C| \leq n Q$ using $\log n+\log Q$ bits followed by $|C|$ pointers which indicate when a collision occurs and how to resolve it.

## C. 2 Encoder

Pseudocode. The pseudocode of Enc is given in Figure 2.

Explanation. First, Enc runs $\mathcal{A}_{0}$ to get the state st ${ }_{0}$. Recall that Enc knows $\Sigma$ and h, so it can simulate $\mathcal{G}$ and h for $\mathcal{A}_{0}$ perfectly. Note that Enc does not perform any bookkeeping regarding the oracles yet.

Then, it runs $\mathcal{A}_{1}$ with state st ${ }_{0}$ on every challenge $j \in[0, n)$ for at most $t$ rounds. It executes all instances of $\mathcal{A}_{1}$ in parallel in lockstep. That is, Enc waits until all instances

Figure 2: Enc given input $\Sigma$ and auxiliary inputs $N$, h, and $\mathcal{W}$.

```
01 Compute st \({ }_{0}:=\mathcal{A}_{0}^{\Sigma, \mathrm{h}}(\Sigma(1), \Sigma(w))\).
02 Initialize empty data structures \(I, \Xi, T, S\), and \(C\).
ьз For all \(j \in[0, n)\), run \(\mathcal{A}_{1}\left(\mathrm{st}_{0}, j\right)\) in lockstep for \(t\) rounds while simulating \(\mathcal{G}\) (Figure 3)
    and \(h\) (Figure 4).
04 Using the equations \(T_{\mathbb{Z}}\left(\tilde{\sigma}_{j}\right)=w^{2^{j}}\), define a system of equations \(A \boldsymbol{x}=\boldsymbol{b}\) where
    \(A \in \mathbb{Z}^{\text {rank }_{\text {min }} \times m}\) and function \(f\) that maps every column \(i\) to the indeterminate \(\mathfrak{x}_{f(i)}\).
    ( \(m\) and \(f\) are described in the explanation below).
05 Compute the Hermite normal form of \(A\) by \(\operatorname{HNF}(A)=(H, U)\).
    If any \(\hat{h}_{i j}=0\) but \(h_{i j} \neq 0\), aborthnf.
    Solve the system \(\hat{A} \hat{\boldsymbol{x}}=\hat{\boldsymbol{b}}\) such that \(\hat{x}_{i}=\Sigma^{-1}\left(T_{\mathcal{L}}\left(\hat{\mathfrak{x}}_{f(i)}\right)\right)\). During this, if any element
    of \(\mathbb{Z}_{N}\) cannot be inverted, abort \(_{\text {sol }}\) :
08 Consider the equivalent system \(\hat{H} \hat{\boldsymbol{y}}=\hat{\boldsymbol{b}}\) where \(\hat{\boldsymbol{y}}=\hat{U}^{-1} \hat{\boldsymbol{x}}\).
09 For \(i \leq\) rank \(_{\text {min }}\), fix the entries \(\hat{y}_{i} \in \hat{\boldsymbol{y}}\) by the triangular shape of \(\hat{H}\).
10 For every \(i \leq f\left(m-\right.\) rank \(\left._{\text {min }}\right)\), if \(\hat{\mathfrak{x}}_{i} \notin S\), push \(\Sigma^{-1}\left(T_{\mathcal{L}}\left(\hat{\mathfrak{x}}_{i}\right)\right) \rightarrow \Xi\). This fixes the first
    \(m-\) rank \(_{\text {min }}\) values of \(\hat{\boldsymbol{x}}_{i}\).
    Given these constraints on \(\hat{\boldsymbol{y}}\) and \(\hat{\boldsymbol{x}}\), solve \(\hat{U} \boldsymbol{y}=\boldsymbol{x}\).
    Substitute every \(\hat{x}_{i}\) for \(\hat{\mathfrak{x}}_{f(i)}\) in \(T\).
    For every remaining \(\hat{\mathfrak{x}}_{i} \in T_{\mathbb{Z}_{N}}\), push \(\Sigma^{-1}\left(T_{\mathcal{L}}\left(\hat{\mathfrak{x}}_{i}\right)\right) \rightarrow \Xi\).
    Iterate over all \(\sigma \in \operatorname{Im}(\Sigma)\) in lexicographical order and for any \(\sigma_{i} \notin T\), push
    \(\Sigma^{-1}\left(\sigma_{i}\right) \rightarrow \Xi\).
15 Output the encoding comprised of \(\operatorname{Im}(\Sigma), I, C\), \(\mathrm{st}_{0}\), and \(\Xi\).
```

have made queries in the current round before answering them all at the same time. Now Enc performs additional bookkeeping in order to aid Dec in simulating $\mathcal{G}$ and h correctly.

Enc simulates $h$ (Figure 4) by forwarding the queries to $h$. It does not need to store the response in $\Xi$ since Dec also gets $h$ as an auxiliary input. However, Dec cannot recognize when $\mathcal{A}_{1}\left(\mathrm{st}_{0}, j\right)$ queries the label $\sigma$ corresponding to the exponent $2^{j} w$ for the first time. Thus, it adds $(j, i)$ to the list of hints $I$ and also sets the variable $\tilde{\sigma}_{j}=\sigma$ for later use in Figure 2.

The bookkeeping for $\mathcal{G}$ is more elaborate. Enc responds to queries using $\Sigma$ as before, but also populates the table $T$. For every query input $\sigma_{i} \notin T, i \in\{1,2\}$, it adds a row containing the label and fresh indeterminate $\mathfrak{x}_{j}$. For the query output $\sigma_{3}$, Enc pushes it to $\Xi$ if it is a new label, so that Dec can correctly simulate $\mathcal{G}$. Then, in any case, Enc computes the algebraic representations of $\sigma_{3}$ over $\mathbb{Z}_{N}$ and $\mathbb{Z}$. The representations are derived by adding/subtracting the representations of the query inputs $\sigma_{1}$ and $\sigma_{2}$. Here, Enc might encounter a collision, i.e., $\sigma_{3} \in T$ and the computed algebraic representation over $\mathbb{Z}_{N}$ differs from the one assigned to $\sigma_{3}$ in the table.

A collision is beneficial for the compression since it gives a non-trivial relation between the indeterminates. In particular, Enc attempts to eliminate one indeterminate in $T$. To this end, it tries to express $\hat{\mathfrak{x}}_{k}$ (where $k$ is the maximal index within the colliding representations) in terms of the other indeterminates. This might fail since not every element in $\mathbb{Z}_{N}$ has an inverse so Enc aborts in such cases. Enc replaces $\hat{\mathfrak{x}}_{k}$ in $T$ and

Figure 3：Answering $\mathcal{A}_{1}\left(\mathrm{st}_{0}, j\right)$＇s $i$ th query $\left(\circ \in\{+,-\right.$ ，inv $\left.\}, \sigma_{1}, \sigma_{2}\right)$ to $\mathcal{G}$ ．

```
01 If ( }\mp@subsup{\sigma}{1}{},\mp@subsup{\sigma}{2}{})\not\in\operatorname{Im}(\Sigma)\times\operatorname{Im}(\Sigma),\mathrm{ respond with }\perp\mathrm{ .
02 Else:
    If }0\in{+,-}
04 Respond with }\mp@subsup{\sigma}{3}{}=\Sigma(\mp@subsup{\Sigma}{}{-1}(\mp@subsup{\sigma}{1}{})\circ\mp@subsup{\Sigma}{}{-1}(\mp@subsup{\sigma}{2}{}))
05 For i\in{1,2}, if }\mp@subsup{\sigma}{i}{}\not\inT\mathrm{ , choose a fresh indeterminate }\mp@subsup{\mathfrak{x}}{j}{}\mathrm{ and append ( }\mp@subsup{\hat{\mathfrak{x}}}{i}{},\mp@subsup{\mathfrak{x}}{i}{},\mp@subsup{\sigma}{i}{}) \(T\) ．
06 If }\mp@subsup{\sigma}{3}{}\not\inT\mathrm{ , add the }(\mp@subsup{T}{\mp@subsup{\mathbb{Z}}{N}{}}{}(\mp@subsup{\sigma}{1}{})\circ\mp@subsup{T}{\mp@subsup{\mathbb{Z}}{N}{}}{}(\mp@subsup{\sigma}{2}{})\operatorname{mod}N,\mp@subsup{T}{\mathbb{Z}}{}(\mp@subsup{\sigma}{1}{})\circ\mp@subsup{T}{\mathbb{Z}}{}(\mp@subsup{\sigma}{2}{}),\mp@subsup{\sigma}{3}{})->T\mathrm{ and push
        \sigma3}->\Xi\mathrm{ .
07 Else, if }\mp@subsup{\sigma}{3}{}\inT\mathrm{ but }\mp@subsup{T}{\mp@subsup{\mathbb{Z}}{N}{}}{}(\mp@subsup{\sigma}{1}{})\circ\mp@subsup{T}{\mp@subsup{\mathbb{Z}}{N}{}}{}(\mp@subsup{\sigma}{2}{})\not=\mp@subsup{T}{\mp@subsup{\mathbb{Z}}{N}{}}{}(\mp@subsup{\sigma}{3}{})\operatorname{mod}N\mathrm{ :
08 Append the collision (j,i,T(\mp@subsup{\sigma}{3}{}))->C\mathrm{ .}
09 Rearrange }\mp@subsup{T}{\mp@subsup{\mathbb{Z}}{N}{}}{}(\mp@subsup{\sigma}{1}{})\circ\mp@subsup{T}{\mp@subsup{\mathbb{Z}}{N}{}}{}(\mp@subsup{\sigma}{2}{})=\mp@subsup{T}{\mp@subsup{\mathbb{Z}}{N}{}}{}(\mp@subsup{\sigma}{3}{})\operatorname{mod}N\mathrm{ to be of the form }\mp@subsup{\hat{\mathfrak{q}}}{k}{}=\mp@subsup{\Sigma}{l<k}{}c\mp@subsup{\hat{\mathfrak{r}}}{l}{
        where }k\mathrm{ is the largest index within the equation and }\mp@subsup{c}{l}{}\in\mp@subsup{\mathbb{Z}}{N}{}\mathrm{ . If rearranging
        fails (i.e., an element of }\mp@subsup{\mathbb{Z}}{N}{}\mathrm{ cannot be inverted), abort col.
        Add ( (\hat{\mathfrak{x}}}
            Substitute 灄俰㖋}
        Else, ○= inv, so respond to an inversion query in the natural manner analogously
        to the case }\circ\in{+,-}
```

Figure 4：Answering $\mathcal{A}_{1}\left(\mathrm{st}_{0}, j\right)$＇s $i$ th query $(\sigma, \ldots)$ to h ．
01 Respond with $\mathrm{h}(\sigma, \ldots) \rightarrow \Xi$ ．
02 If $\Sigma\left(w^{2^{j}}\right)=\sigma$ ，set $\tilde{\sigma}_{j}:=\sigma$ and append $(j, i) \rightarrow I$ ．
also keeps track of this replacement by adding the appropriate mapping to the list of substitutions $S$ ．

After executing all instances of $\mathcal{A}_{1}$ for at most $t$ rounds，Enc（and also Dec）holds the table $T$ ，substitutions $S$ ，and rank $_{\min }$ many labels $\tilde{\sigma}_{j}=\Sigma\left(w^{2^{j}}\right)$ ．It chooses the first rank min of these labels，and creates a system of equations using the column $T_{\mathbb{Z}_{N}}$ ，i．e．，

$$
\begin{aligned}
a_{11} \mathfrak{x}_{f(1)}+a_{12} \mathfrak{x}_{f(2)}+\cdots+a_{1 m} \mathfrak{x}_{f(m)} & =2^{j_{1}} w \\
a_{21} \mathfrak{x}_{f(1)}+a_{22} \mathfrak{x}_{f(2)}+\cdots+a_{2 m} \mathfrak{x}_{f(m)} & =2^{j_{2}} w \\
& \vdots \\
a_{\text {rank }_{\text {min }} 1} \mathfrak{x}_{f(1)}+a_{\text {rank }_{\text {min }} 2} \mathfrak{x}_{f(2)}+\cdots+a_{r m} \mathfrak{x}_{f(m)} & =2^{j_{\text {rank }}}{ }_{\text {min }} w
\end{aligned}
$$

which can be viewed as a matrix $A \in \mathbb{Z}^{\text {rank }_{\min \times m}}$ and vector $\boldsymbol{b} \in \mathbb{Z}^{m}$ ．Here，$m$ is the number of indeterminates $\hat{\mathfrak{x}}_{i}$ that occur in the system of equations．Note that these indeterminates might only be a subset of all indeterminates，so the function $f$ translates between column indices and indeterminate subscripts．So $i \leq f(i)$ and $f$ might skip some indices．

Next，Enc computes the Hermite normal form of $A$ resulting in a lower－triangular matrix $H$ and unimodular（i．e．，invertible over $\mathbb{Z}$ ）matrix $U$ such that $A U=H$ ．Recall
that $H$ is a lower triangular matrix with rank min non-zero columns on the left and the height of these columns is strictly decreasing. Enc aborts when $H$ and $\hat{H}$ differ in their shape, that is, $\hat{H}$ has additional zero entries because an entry of $H$ was a multiple of $N$. On a high level, this step is necessary to connect our compression argument, which mainly works in $\mathbb{Z}_{N}$, with the query behavior of $\mathcal{A}_{1}$, which must be interpreted over $\mathbb{Z}$ in hidden order groups.

Enc solves the system $\hat{A} \hat{\boldsymbol{x}}=\hat{\boldsymbol{b}}$ by solving $\hat{H} \hat{\boldsymbol{y}}=\hat{\boldsymbol{b}}$ where $\hat{\boldsymbol{x}}=\hat{U} \hat{\boldsymbol{y}}$. The first rank min entries of $\boldsymbol{y}$ are fixed because of the triangular shape of $\hat{H}$ and because every element can be inverted (otherwise, Enc aborts). If rank $\min =m$, then this leads to a unique solution. Otherwise, $m$ - rank $\min$ entries in $\hat{\boldsymbol{y}}$ are not uniquely defined and there might be multiple solutions for $\hat{\boldsymbol{x}}$. Hence, Enc needs to add additional constraints. It does so by fixing the first $m$ - rank min entries of $\hat{\boldsymbol{x}}$. To this end, Enc pushes the discrete logarithm of all $\hat{\mathfrak{x}}_{i}$ with $i \leq f\left(m-\right.$ rank $\left._{\text {min }}\right)$ to $\Xi$ whenever $\hat{\mathfrak{x}}_{i}$ is not in $S$. Because if this is the case, then Dec can compute $\hat{\mathfrak{x}}_{i}$ from the preceding $\hat{\mathfrak{x}}_{j}, j<i$.

Now that the system $\hat{U} \hat{\boldsymbol{y}}=\hat{\boldsymbol{x}}$ is uniquely determined and can be solved to yield $\hat{\boldsymbol{x}}$. As a result, up to rank ${ }_{\text {min }}$ many indeterminates in $T$ might be replaced. Note that only "up to" because some might have already been eliminated due to collisions.

Still, $T$ might contain some indeterminates $\hat{\mathfrak{x}}_{i}$, so Enc pushes the remaining $\hat{\mathfrak{x}}_{i}$ in $T$ to $\Xi$. Finally, the column $T_{\mathbb{Z}_{N}}$ only contains constants which correspond to the discrete logarithm of the labels. $T$ does not contain all labels so Enc pushes the remaining discrete logarithms to $\Xi$ in lexicorgraphical order of the labels.

The encoding output by Enc contains the image $\operatorname{Im}(\Sigma)$, the list of hints $I$ and of collisions $C$, the state st ${ }_{0}$ output by $\mathcal{A}_{0}$, and $\Xi$. Dec follows from the description of Enc and this explanation so it is omitted.

## C. 3 Encoding Size

Let us analyze the size of Enc's encoding.

- $\operatorname{Im}(\Sigma)$ which requires $\log \binom{|\mathcal{L}|}{N}$.
- $I$ is a list of fixed length rank $\min ^{\min }$ and every element is a pointer to a query. Identifying a query requires $\log n$ bits to specify the challenge $j \in[0, n)$ and $\log Q$ bits to indicate the query's index within the execution $\mathcal{A}_{1}\left(\mathrm{st}_{0}, j\right)$. Thus, $\|I\|=$ rank $_{\text {min }}(\log n+\log Q)$.
- $C$ is a list of variable length and every element contains a pointer to a query where a collision happens as well as a row of $T$. Every row of $T$ can be identified by the query in which it was added plus $\lceil\log 3\rceil=2$ bits as every query adds at most three rows. So $\|C\|=(\log n+\log Q)+|C|(2(\log n+\log Q)+2)$ where $(\log n+\log Q)$ encodes the variable length $|C| \leq n Q$. We do not have any non-trivial bound on $|C|$, but this will not be an issue for our analysis.
- $\left\|s \mathrm{t}_{0}\right\|=M$ by definition.
- $\|\Xi\|$ is not bounded as well, but, again, this is not problematic for our analysis.

Now, to handle the variable length of $C$ and $\Xi$, we compare the size of this encoding to the naive encoding of $\Sigma$. The naive encoding also encodes $\operatorname{Im}(\Sigma)$ with $\log \binom{|\mathcal{L}|}{N}$ bits and
then uses $\log (N!)$ bits to assign discrete logarithms to all labels in the image. Hence, we can ignore the size of the image and will focus on how much Enc pays (in terms of bits) for every discrete logarithm. We will show that any element contained in $C$ and $\Xi$ leads to one discrete logarithm and that it is at most as expensive as in the naive encoding.

This is directly related to how much is paid for every row of $T$ as every row of $T$ yields one discrete logarithm in the end.

- Unexpected query inputs require 0 bits initially as they are assigned an indeterminate $\hat{\mathfrak{x}}_{i}$.
- Responses without collision (stored in $\Xi$ ) cost $\log (N-r)$ bits where $r$ is the number of rows in $T$ (and thus the number of labels already assigned) at the time of the query.
- Responses with collision (stored in $C)$ cost $2(\log n+\log Q)+2$ bits but at the same time eliminate one indeterminate $\hat{\mathfrak{x}}_{i}$ in $T$.

Thus, responses without collisions cost as much as in the naive encoding and collisions are cheaper because $2(\log n+\log Q)+2<\log (N)-1$ for $\lambda$ large enough. Indeed, $n$ and $Q$ are polynomial, but $N$ is superpolynomial in $\lambda$. So far, our encoding is at most as expensive as the naive one without accounting for the indeterminates and additional bookkeeping yet.

When solving the system of equations, we get rank ${ }_{\min }$ many indeterminates for the cost of a hint of $\log n+\log Q$ bits. However, we might have already resolved all of those indeterminates by collisions, so in the worst case we pay $3(\log n+\log Q)+2$ bits for rank $_{\text {min }}$ many indeterminates. By the same reasoning as above, $3(\log n+\log Q)+2<\log (N)-1$, so we have a net profit of at least $0<\log (N)-1-(3(\log n+\log Q)+2)$ bits for rank $_{\text {min }}$ many indeterminates.

In the course of solving the system of equations, we might have to fix some values of $\hat{x}_{i}$ by fixing some indeterminates (stored in $\Xi$ ). Since we are fixing these indeterminates topdown according to the rows of $T$, we pay at most $\log (N-r)$ bits per indeterminant - the same as the naive algorithm. Note that we account for collisions in this step by using the mapping $S$. This ensures that we do not overpay for any indeterminate that has already been resolved by a collision. The remaining indeterminants and discrete logarithms (both stored in $\Xi$ ) cost at most $\log (N-r)$ bits again.

To summarize, we pay fewer bits for rank ${ }_{\min }$ discrete logarithms and, accounting for the additional bookkeeping data, pay at most

$$
\underbrace{\log \binom{|\mathcal{L}|}{N}+\log (N!)}_{\text {Naive encoding of } \Sigma}+\underbrace{\log n+\log Q}_{\text {Encoding of }|C|}+\underbrace{M}_{\| \text {sto } \|}-\underbrace{\operatorname{rank}_{\min }(\log (N)-3(\log n+\log Q)-3)}_{\text {Profit over the naive encoding }}
$$

bits in total relative to the naive algorithm.

## C. 4 Success Probability

abort $_{\text {hnf }}$ and abort $_{\text {sol }}$ do not happen for any element in bad as guaranteed by Claim 1. However, the lemma does not cover abort ${ }_{\text {col }}$. Noting that abort ${ }_{c o l}$ may be used to factor
$N$, we modify Claim 1 to account for it as well. So Enc never aborts, so it is successful with probability 1.

The discussion in this section gives a proof of Claim 5.


[^0]:    ${ }^{1}$ MHFs are typically constructed from hash functions. In security proofs, these are modelled as random oracles. So a "step" may be thought of as a query (or many independent queries in parallel) to the random oracle.

[^1]:    ${ }^{2}$ In a similar vein, some proof of work blockchains use MHFs, but the parameters cannot be too large as otherwise the verification of blocks becomes too expensive. Sadly, TMHF are not suited for blockchains since it is unclear who would possess the trapdoor.

[^2]:    ${ }^{3}$ Using notation from [BP17], $M:=n, L:=n$, and $\eta:=1$.
    ${ }^{4}$ We assume that a hash- and group operations take the same amount of time.
    ${ }^{5}$ Note that $\omega_{\mathrm{h}}$ is the output length of the random oracle which corresponds to $\log N$ in our bounds.

[^3]:    ${ }^{6}$ This means that the systems have integer coefficients. Intuitively, if an algorithm is oblivious to the group order, its query behavior might as well be analyzed over $\mathbb{Z}$ instead of $\mathbb{Z}_{N}$.

[^4]:    ${ }^{7}$ Note that usually $N$ and $N^{\prime}$ are defined the other way around. However, most parts of the paper are only concerned with $N$, so we chose this notation to avoid clutter.

[^5]:    ${ }^{8}$ The following explanation is high-level and thus omits technicalities such as collision handling.

[^6]:    ${ }^{9}$ In fact, for every row $\boldsymbol{a}_{j}$, one can even show that $\left\|\boldsymbol{a}_{j}\right\|_{1} \leq 2^{t}$ (cf. Appendix B), but the weaker statement $\left\|\boldsymbol{a}_{j}\right\|_{\infty} \leq 2^{t}$ is sufficient for our proof.
    ${ }^{10}$ A precise definition of $A, \boldsymbol{x}$, and $\boldsymbol{b}$ appears in Appendix C. 2 as part of the proof contained therein.

[^7]:    ${ }^{11}$ Consider a subset $S \subseteq\left\{1, \ldots, 2^{t}\right\}$ of size $|S|=m$. What is the maximum size $m$ so that all subsets of $S$ have distinct sums? The best-known lower bound is $t+2 \leq m$ and upper bounds are roughly $m<t+\log t$ (ignoring constants). Erdős conjectured that $m=t+O(1)$ and offered $\$ 500$ for proof or refutation [Guy94, C8].
    ${ }^{12}\|\boldsymbol{r}\|_{\infty}$ is the infinity norm of the vector $\boldsymbol{r}=\left(r_{1}, \ldots, r_{m}\right)$ and defined as $\max _{i}\left|r_{i}\right|$.

[^8]:    ${ }^{13}$ To be precise, $\rho=(M+\log n+2 \log Q+c \log \lambda+1) /(\log (N)-3(\log Q-\log n)-3)$ where $c$ is as in Lemma 1.

[^9]:    ${ }^{14}$ As we will explain later, in contrast to the proof sketch in Section 5, there is a minor technical difference between the indeterminates $\mathfrak{x}_{i}$ and the entries $x_{i}$ belonging to the vector $\boldsymbol{x}$ of the system $A \boldsymbol{x}=\boldsymbol{b}$.
    ${ }^{15}$ The algorithm starts with $A=A_{0}$ and iteratively computes $A_{i}=\left(\begin{array}{cc}* & 0 \\ * & A_{i+1}\end{array}\right)$ where 0 is an all-zero and * is an arbitrary submatrix. In every iteration, it computes the GCD of $A_{i+1}$ 's topmost row, swaps the GCD to $A_{i+1}$ 's leftmost column, and zeros out all elements to the right of it. This can be realized using elementary column operations and results in a lower-triangular matrix $H$. Then the algorithm performs further computation to ensure other HNF conditions. However, this does not change the shape of $H$, so it is not relevant for our argument.

