# On the Memory-Tightness of Hashed ElGamal

Ashrujit Ghoshal and Stefano Tessaro

Paul G. Allen School of Computer Science & Engineering
University of Washington, Seattle, USA
{ashrujit,tessaro}@cs.washington.edu

**Abstract.** We study the memory-tightness of security reductions in public-key cryptography, focusing in particular on Hashed ElGamal. We prove that any *straightline* (i.e., without rewinding) black-box reduction needs memory which grows linearly with the number of queries of the adversary it has access to, as long as this reduction treats the underlying group generically. This makes progress towards proving a conjecture by Auerbach *et al.* (CRYPTO 2017), and is also the first lower bound on memory-tightness for a concrete cryptographic scheme (as opposed to generalized reductions across security notions). Our proof relies on compression arguments in the generic group model.

**Keywords:** Public-key cryptography, memory-tightness, lower bounds, generic group model, foundations, compression arguments

## 1 Introduction

Security proofs rely on *reductions*, i.e., they show how to transform an adversary $\mathcal{A}$ breaking a scheme into an adversary $\mathcal{B}$ solving some underlying assumed-to-be-hard problem. Generally, the reduction ought to be *tight* – the resources used by $\mathcal{B}$, as well as the attained advantage, should be as close as possible to those of $\mathcal{A}$. Indeed, the more resources $\mathcal{B}$ needs, or the smaller its advantage, the weaker the reduction becomes.

Auerbach *et al.* [**?**] were the first to explicitly point out that *memory* resources have been ignored in reductions, and that this leads to a loss of quality in security results. Indeed, it is conceivable that $\mathcal{A}$'s memory is naturally bounded (say, at most $2^{64}$ bits), and the underlying problem is very sensitive to memory. For example, the best-known algorithm for discrete logarithms in a 4096-bit prime field runs in time (roughly) $2^{156}$ using memory $2^{80}$. With less memory, the best algorithm is the generic one, requiring time $\Theta(\sqrt{p}) \approx 2^{2048}$. Therefore, if $\mathcal{B}$ also uses memory at most $2^{64}$, we can infer a larger lower bound on the necessary time complexity for $\mathcal{A}$ to break the scheme, compared to the case where $\mathcal{B}$ uses $2^{100}$ bits instead.

What can be memory-tight? One should therefore target reductions that are *memory-tight*, i.e., the memory usage of $\mathcal{B}$ is similar to that of $\mathcal{A}$.[1] The work

[1] Generally, $\mathcal{B} = \mathcal{R}^{\mathcal{A}}$ for a black-box reduction $\mathcal{R}$, and one imposes the slightly stronger requirement that $\mathcal{R}$ uses small memory, independent of that of $\mathcal{A}$.

of Auerbach *et al.* [**?**], and its follow-up by Wang *et al.* [**?**], pioneered the study of memory-tight reductions. In particular, and most relevant to this work, they show *negative* results (i.e., that certain reductions cannot be memory tight) using *streaming lower bounds*.

Still, these lower bounds are tailored at general notions (e.g., single- to multi-challenge reductions), and lower bounds follow from a natural connection with classical frequency problems on streams. This paper tackles the more ambitious question of proving impossibility of memory-tight reductions for concrete *schemes*, especially those based on algebraic structures. This was left as an open problem by prior works.

Hashed ElGamal. Motivated by a concrete open question posed in [**?**], we consider here the CCA-security of Hashed ElGamal. In its KEM variant, the scheme is based on a cyclic group $G = \langle g \rangle$ – the secret key $\mathsf{sk}$ is a random element from $\mathbb{Z}_{|G|}$, whereas the public key is $\mathsf{pk} = g^{\mathsf{sk}}$. Then, encapsulation produces a ciphertext-key pair

$$C \leftarrow g^r \ , \quad K \leftarrow \mathsf{H}(\mathsf{pk}^r) \ .$$

for $r \leftarrow \mathbb{Z}_{|G|}$ and a hash function $\mathsf{H} : G \rightarrow \{0,1\}^\ell$. Decapsulation occurs by computing $K \leftarrow \mathsf{H}(C^{\mathsf{sk}})$.

The CCA-security of Hashed ElGamal in the random-oracle model was proved by Abdalla, Bellare, and Rogaway [**?**] based on the *Strong Diffie-Hellman* (SDH) assumption (also often called GapDH), and we briefly review the proof.[2] First, recall that in the SDH assumption, the attacker is asked to compute $g^{uv}$ from $g^u$ and $g^v$, given additionally access to a *decision* oracle $\mathsf{O}_\mathsf{v}$ which on input $h, y \in G$, tells us whether $h^v = y$.

The reduction sets the Hashed ElGamal public-key to $\mathsf{pk} = g^v$ (setting implicitly $\mathsf{sk} = v$), the challenge ciphertext to be $C^* = g^u$, and the corresponding key $K^*$ to be a random string. Then, it simulates both the random oracle and the decapsulation oracle to the adversary $\mathcal{A}$ (which is run on inputs $\mathsf{pk}, C^*$ and $K^*$), until a random-oracle query for $g^{uv}$ is made (this can be detected using the $\mathsf{O}_\mathsf{v}$ oracle). The challenge is to simulate both oracles consistently: As the reduction cannot compute discrete logarithms, it uses the oracle $\mathsf{O}_\mathsf{v}$ to detect whether a random-oracle query $X$ and a decapsulation query $C_i$ satisfy $\mathsf{O}_\mathsf{v}(C_i, X) = \mathsf{true}$, and, if this is the case, answers them with the same value.

This reduction requires memory to store all prior decapsulation and random-oracle queries. Unlike other reductions, the problem here is not to store the random-oracle output values (which could be compressed using a PRF), but the actual *inputs* to these queries, which are under adversarial control. This motivates the conjecture that a reduction using little memory does not exist, but the main challenge is of course to prove this is indeed the case.

[2] Abdalla et al. [**?**] do not phrase their paper in terms of the KEM/DEM paradigm [**?**,**?**], which was introduced concurrently – instead, they prove that an intermediate assumption, called Oracle Diffie-Hellman (ODH), follows from SDH in the ROM. However, the ODH assumption is structurally equivalent to the CCA security of Hashed ElGamal KEM for one challenge ciphertext.

Our result, in summary. We provide a *memory* lower bound for reductions that are *generic* with respect to the underlying group $G$. Specifically, we show the existence of an (inefficient) adversary $\mathcal{A}$ in *the generic group model (GGM)* which breaks the CCA security of Hashed ElGamal via $O(k)$ random oracle/decapsulation queries, but such that no reduction using less than $k \cdot \lambda$ bits of memory can break the SDH assumption *even* with access to $\mathcal{A}$, where $\lambda$ is the bit-size of the underlying group elements.

Our lower bound is strong in that it shows we do not even have a trade-off between advantage and memory, i.e., if the memory is smaller than $k \cdot \lambda$, then the advantage is very small, as long as the reduction makes a polynomial number of queries to $\mathsf{O}_\mathsf{v}$ and to the generic group oracle. It is however also important to discuss two limitations of our lower bound. The first one is that the reduction – which receives $g, g^v$ in the SDH game – uses $\mathsf{pk} = g^v$ as the public key to the Hashed ElGamal adversary. The second one is that the reduction is straightline, i.e., it does not perform any rewinding.

We believe that our impossibility result would extend even when the reduction is not straightline. However, allowing for rewinding appears to be out of reach of our techniques. Nonetheless, we *do* conjecture a lower bound on the memory of $\Omega(k \log k)$ bits, and discuss the reasoning behind our conjecture in detail in the full version.

We stress that our result applies to reductions in the GGM, but treats the adversary as a black box. This captures reductions which are black-box in their usage of the group and the adversary. (In particular, the reduction cannot see generic group queries made by the adversary, as in a GGM security proofs.) Looking at the GGM reduces the scope of our result. However, it is uncommon for reductions to depend on the specifics of the group, although our result can be bypassed for specific groups, e.g., if the group has a pairing.

Concurrent related work. Concurrently to our work, Bhattacharyya [**?**] provides memory-tight reductions of KEM-CCA security for variants of Hashed ElGamal. At first glance, the results seem to contradict ours. However, they are entirely complementary – for example, a first result shows a memory tight reduction for the KEM-CCA security of the "Cramer-Shoup" variant of Hashed ElGamal – this variant differs from the (classical) Hashed ElGamal we consider here and is less efficient. The second result shows a memory-tight reduction for the version considered in this paper, but assumes that the underlying group has a pairing. This is a good example showing our result can be bypassed for specific groups i.e. groups with pairings, but we also note that typical instantiations of the scheme are on elliptic curves for which no pairing exists.

### 1.1 Our Techniques

We give a high-level overview of our techniques here. We believe some of these to be novel and of broader interest in providing other impossibility results.

THE SHUFFLING GAME. Our adversary against Hashed ElGamal[3] $\mathcal{A}$ first attempts to detect whether the reduction is using a sufficient amount of memory. The adversary $\mathcal{A}$ is given as input the public key $g^v$, as well as $g^u$, as well as a string $C^* \in \{0,1\}^\ell$, which is either a real encapsulation or a random string. It first samples $k$ values $i_1, \ldots, i_k$ from $\mathbb{Z}_p$. It then:

**(1)** Asks for decapsulation queries for $C_j \leftarrow g^{i_j}$, obtaining values $K_j$, for $j \in [k]$
**(2)** Picks a random permutation $\pi : [k] \to [k]$.
**(3)** Asks for RO queries for $H_j \leftarrow \mathsf{H}(V_j)$ for $j \in [k]$, where $V_j \leftarrow g^{v \cdot i_{\pi(j)}}$.

After this, the adversary checks whether $K_j = H_{\pi(j)}$ for all $j \in [k]$, and if so, it continues its execution, breaking the ODH assumption (inefficiently). If not, it just outputs a random guess.

The intuition here is that no reduction using substantially less than $k \cdot \log p$ bits succeeds in passing the above test – in particular, because the inputs $C_j$ and $V_j$ are (pseudo-)random, and thus incompressible. If the test does not pass, the adversary $\mathcal{A}$ is rendered useless, and thus not helpful to break SDH.

*Remark 1.* The adversary here is described in a way that requires secret randomness, not known to the reduction, and it is easier to think of $\mathcal{A}$ in this way. We will address in the body how to generically make the adversary deterministic.

*Remark 2.* We stress that this adversary requires memory – it needs to remember the answers $C_1, \ldots, C_k$. However, recall that we adopt a black-box approach to memory-tightness, where our requirement is that the reduction itself uses little memory, regardless of the memory used by the adversary. We also argue this is somewhat necessary – it is not clear how to design a reduction which adapts its memory usage to the adversary, even if given this information in a non-black-box manner. Also, we conjecture different (and much harder to analyze) memoryless adversaries exist enabling a separation. An example is multi-round variant, where each round omits **(2)**, and **(3)** only asks a single query $\mathsf{H}(V_j)$ for a random $j \leftarrow_\$ [k]$, and checks consistency. Intuitively, the chance of passing each round is roughly $k \log p/s$, but we do not know how to make this formal.

INTRODUCING THE GGM. Our intuition is however *false* for an arbitrary group. For instance, if the discrete logarithm (DL) problem is easy in the group, then the reduction can simply simulate the random oracle via a PRF, as suggested in [**?**]. Ideally, we could prove that if the DL problem is hard in $G$, then any PPT reduction given access to $\mathcal{A}$ and with less than $k \cdot \log p$ bits of memory fails to break SDH.[4] Unfortunately, it will be hard to capture a single hardness property of $G$ sufficient for our proof to go through. Instead, we will model the

[3] The paper will in fact use the cleaner formalization of the ODH assumption, so we stick to Hashed ElGamal only in the introduction.

[4] This statement is somewhat confusing, so note that in general, the existence of a reduction is *not* a contradiction with the hardness of DL, as the reduction is meant to break SDH only given access to an adversary breaking the scheme, and this does not imply the ability to break SDH *without* access to the adversary.

group via the *generic group model* (GGM) [?,?]: We model a group of prime order $p$ defined via a random injection $\sigma : \mathbb{Z}_p \to \mathcal{L}$. An algorithm in the model typically has access to $\sigma(1)$ (in lieu of $g$) and an evaluation oracle which on input $\mathbf{a}, \mathbf{b} \in \mathcal{L}$ returns $\sigma(\sigma^{-1}(\mathbf{a}) + \sigma^{-1}(\mathbf{b}))$. (We will keep writing $g^i$ instead of $\sigma(i)$ in the introduction, for better legibility.)

THE PERMUTATION GAME. In order to fool $\mathcal{A}$, the reduction can learn information about $\pi$ via the $\mathsf{O}_\mathsf{v}$ oracle. For example, it can try to input $C_j = g^{i_j}$ and $V_{j'} = g^{v i_{\pi(j')}}$ (both obtained from $\mathcal{A}$'s queries), and $\mathsf{O}_\mathsf{v}(C_j, V_{j'}) = \mathsf{true}$ if and only if $\pi(j') = j$. More generally, the reduction can compute, for any $\vec{a} = (a_1, \ldots, a_k)$ and $\vec{b} = (b_1, \ldots, b_k)$,

$$C^* = g^{\sum_{j=1}^k a_j i_j} = \prod_{j=1}^k C_j^{a_j} \ , \ \ V^* = g^{\sum_{j=1}^k b_j v \cdot i_{\pi(j)}} = \prod_{j=1}^k V_j^{b_j} \ ,$$

and the query $\mathsf{O}_\mathsf{v}(C^*, V^*)$ returns true iff $b_j = a_{\pi(j)}$ for all $j \in [k]$, which we write as $\vec{b} = \pi(\vec{a})$. We abstract this specific strategy in terms of an information-theoretic game – which we refer to as the *permutation game* – which gives the adversary access to an oracle $O$ which takes as inputs pairs of vectors $(\vec{a}, \vec{b})$ from $\mathbb{Z}_p^k$, and returns $\mathsf{true}$ iff $\vec{b} = \pi(\vec{a})$ for a secret permutation $\pi$. The goal of the adversary is to recover $\pi$.

Clearly, a strategy can win with $O(k^2)$ oracle queries $(\vec{e}_i, \vec{e}_j)$ for all $i, j$, where $\vec{e}_i \in \mathbb{Z}_p^k$ is the unit vector with a 1 in the $i$-th coordinate, and 0 elsewhere. This strategy requires in particular querying, in its first component, vectors which have rank $k$. Our first result will prove that this is necessary – namely, assume that an adversary makes a sequence of $q$ queries $(\vec{x}_1, \vec{y}_1), \ldots, (\vec{x}_q, \vec{y}_q)$ such that the rank of $\vec{x}_1, \ldots, \vec{x}_p$ is at most $\ell$, then the probability to win the permutation game is of the order $O(q^\ell / k!)$. We will prove this via a compression argument.

Note that roughly, this bound tells us that to win with probability $\epsilon$ and $q$ queries to the oracle, the attacker needs

$$\ell = \Omega\left(\frac{k \log k - \log(1/\epsilon)}{\log(q)}\right) \ .$$

A REDUCTION TO THE PERMUTATION GAME. We will think of the execution of the reduction against our adversary as consisting of two stages – we refer to them as $\mathcal{R}_1$ and $\mathcal{R}_2$. The former learns the decapsulation queries $g^{i_1}, \ldots, g^{i_k}$, whereas the latter learns the RO queries $g^{i_{\pi(1)} v}, \ldots, g^{i_{\pi(k)} v}$, and (without loss of generality) attempts to guess the permutation $\pi$. We will lower bound the size of the state $\phi$ that $\mathcal{R}_1$ passes on to $\mathcal{R}_2$. Both stages can issue $\mathsf{O}_\mathsf{v}$ and $\mathsf{Eval}$ queries.

Note that non-trivial $\mathsf{O}_\mathsf{v}$ queries (i.e., those revealing some information about the permutation), are (except with very small probability) issued by $\mathcal{R}_2$, since no information about $\pi$ is ever revealed to $\mathcal{R}_1$. As one of our two key steps, we will provide a reduction from the execution of $\mathcal{R}_1, \mathcal{R}_2$ against $\mathcal{A}$ in the GGM to the permutation game – i.e., we build an adversary $\mathcal{D}$ for the latter game simulating the interaction between $\mathcal{R}_1, \mathcal{R}_2$ and $\mathcal{A}$, and such that $\mathcal{R}_1, \mathcal{R}_2$ "fooling" $\mathcal{A}$ results in $\mathcal{D}$ guessing the permutation.

MEMORY VS. RANK. The main question, however, is to understand the complexity of $\mathcal{D}$ in the permutation game, and in particular, the *rank* $\ell$ of the first component of its queries – as we have seen above, this affects its chance of winning the game.

To do this, we will take a slight detour, and specifically consider a set $\mathcal{Z} \subseteq \mathcal{L}$ of labels (i.e., outputs of $\sigma$) that the reduction $\mathcal{R}_2$ comes up with (as inputs to either of $\mathsf{Eval}$ or $\mathsf{O_v}$) on its own (in the original execution), i.e., no earlier $\mathsf{Eval}$ query of $\mathcal{R}_2$ returned them, and that have been previously learnt by $\mathcal{R}_1$ as an output of its $\mathsf{Eval}$ queries. (The actual definition of $\mathcal{Z}$ is more subtle, and this is due to the ability of the adversary to come up with labels *without* knowing the corresponding pre-image.)

Then, we will show two statements about $\mathcal{Z}$:

- **(i)** On the one hand, we show that the rank $\ell$ of the oracle queries of the adversary $\mathcal{D}$ is upper bound by $|\mathcal{Z}|$ in its own simulation of the execution of $\mathcal{R}_1, \mathcal{R}_2$ with $\mathcal{A}$.
- **(ii)** On the other hand, via a compression argument, we prove that the size of $\mathcal{Z}$ is related to the length of $\phi$, and this will give us our final upper bound.

This latter statement is by itself not very surprising – one can look at the execution of $\mathcal{R}_2$, and clearly every label in $\mathcal{Z}$ that appears "magically" in the execution must be the result of storing them into the state $\phi$. What makes this different from more standard compression arguments is the handling of the generic group model oracle (which admits non-trivial operations). In particular, our compression argument will compress the underlying map $\sigma$, and we will need to be able to figure out the pre-images of these labels in $\mathcal{Z}$. We give a very detailed technical overview in the body explaining the main ideas.

MEMORY-TIGHT AGM REDUCTION. The Algebraic Group Model (AGM) was introduced in [**?**]. AGM reductions make strong extractability assumptions, and the question of their memory-tightness is an interesting one. In the full version we construct a reduction to the discrete logarithm problem that runs an adversary against the KEM-CCA security of Hashed ElGamal in the AGM such that the reduction is memory-tight but not tight with respect to advantage. We note that the model of our reduction is different than a (full-fledged) GGM reduction which is not black-box, in that it can observe the GGM queries made by the adversary. Our result does not imply any impossibility for these. In turn, AGM reductions are weaker, but our results do not imply anything about them, either.

## 2 Preliminaries

In this section, we review the formal definition of the generic group model. We also state ODH and SDH as introduced in [**?**] in the generic group model.

NOTATION. Let $\mathbb{N} = \{0, 1, 2, \cdots\}$ and, for $k \in \mathbb{N}$, let $[k] = \{1, 2, \cdots, k\}$. We denote by $\mathsf{InjFunc}(S_1, S_2)$ the set of all injective function from $S_1$ to $S_2$.

We also let $*$ denote a wildcard element. For example $\exists t : (t, *) \in T$ is true if the set $T$ contains an ordered pair whose first element is $t$ (the type of the

wildcard element shall be clear from the context). Let $\mathcal{S}_k$ denote the set of all permutations on $[k]$. We use $f : \mathsf{D} \to \mathsf{R} \cup \{\perp\}$ to denote a partial function, where $f(x) = \perp$ indicates the value of $f(x)$ is undefined. Define in particular $D(f) = \{d \in \mathsf{D} : f(d) \neq \perp\}$ and $R(f) = \{r \in \mathsf{R} : \exists d \in \mathsf{D} : \sigma(d) = r\}$. Moreover, we let $\overline{D(f)} = \mathsf{D} \backslash D(f)$ and $\overline{R(f)} = \mathsf{R} \backslash R(f)$.

We shall use pseudocode descriptions of games inspired by the code-based framework of [**?**]. The output of a game is denoted using the symbol $\Rightarrow$. In all games we assume the flag $\mathsf{bad}$ is set to $\mathsf{false}$ initially. In pseudocode, we denote random sampling using $\leftarrow_{\$}$, assignment using $\leftarrow$ and equality check using $=$. In games that output boolean values, we use the term "winning" the game to mean that the output of the game is $\mathsf{true}$.

We also introduce some linear-algebra notation. Let $S$ be a set vectors with equal number of coordinates. We denote the rank and the linear span of the vectors by $\mathsf{rank}(S)$ and $\mathsf{span}(S)$ respectively. Let $\vec{x}, \vec{y}$ be vectors of dimension $k$. We denote $\vec{z}$ of dimension $2k$ which is the concatenation of $\vec{x}, \vec{y}$ as $\vec{z} = (\vec{x}, \vec{y})$. We denote the element at index $i$ of a vector $\vec{x}$ as $\vec{x}[i]$.

### 2.1 Generic Group Model

The *generic group model* [**?**] captures algorithms that do not use any special property of the encoding of the group elements, other than assuming every element of the group has a unique representation, and that the basic group operations are allowed. This model is useful in proving lower bounds for some problems, but we use it here to capture reductions that are not specific to the underlying group.

More formally, let the order of the group be a large prime $p$. Let $\mathbb{Z}_p = \{0, 1, 2, \cdots, p-1\}$. Let $\mathcal{L} \subset \{0,1\}^*$ be a set of size $p$, called the set of *labels*. Let $\sigma$ be a random injective mapping from $\mathbb{Z}_p$ to $\mathcal{L}$. The idea is that now every group element in $\mathbb{Z}_p$ is represented by a label in $\mathcal{L}$. An algorithm in this model takes as input $\sigma(1), \sigma(x_1), \sigma(x_2), \cdots, \sigma(x_n)$ for some $x_1, \cdots, x_n \in \mathbb{Z}_p$ (and possibly other inputs which are not group elements). The algorithm also has access to an oracle named $\mathsf{Eval}$ which takes as input two labels $\mathbf{a}, \mathbf{b} \in \mathcal{L}$ and returns $\mathbf{c} = \sigma(\sigma^{-1}(\mathbf{a}) + \sigma^{-1}(\mathbf{b}))$. Note that for any $d$, given $\sigma(x_i)$, $\sigma(d \cdot x_i)$ can be computed using $O(\log d)$ queries to $\mathsf{Eval}$. We denote this operation as $\mathsf{Exp}(\sigma(x_i), d)$. We assume that all labels queried by algorithms in the generic group model are valid i.e. all labels queried by algorithms in the generic group model are in $\mathcal{L}$.[5]

Oracle Diffie-Hellman assumption (ODH). We first formalize the Oracle Diffie-Hellman Assumption (ODH) [**?**], which we are going to use in lieu of the CCA security of Hashed ElGamal. Suppose, a group has generator $g$ and order $p$. The domain of hash function $\mathsf{H}$ is all finite strings and range is $\{0,1\}^{\mathsf{hLen}}$. The assumption roughly states for $u, v \leftarrow_{\$} \mathbb{Z}_p, W \leftarrow_{\$} \{0,1\}^{\mathsf{hLen}}$, the distributions $(g^u, g^v, \mathsf{H}(g^{uv}))$ and $(g^u, g^v, W)$ are indistinguishable to an adversary who has

[5] We stress that we assume a strong version of the model where the adversary *knows* $\mathcal{L}$.

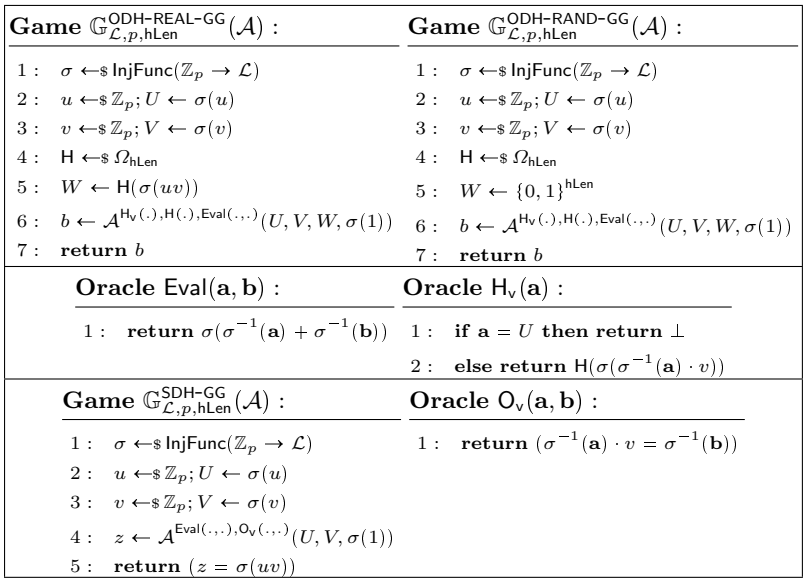


**Fig. 1.** Games for ODH and SDH assumptions

access to the oracle $\mathsf{H_v}$ where $\mathsf{H_v}(g^x)$ returns $\mathsf{H}(g^{xv})$ with the restriction that it is not queried on $g^u$.

We give a formalization of this assumption in the random-oracle and generic group models. For a fixed $\mathsf{hLen} \in \mathbb{N}$, let $\Omega_{\mathsf{hLen}}$ be the set of hash functions mapping $\{0,1\}^*$ to $\{0,1\}^{\mathsf{hLen}}$. In Figure **??**, we formally define the Games $\mathbb{G}^{\mathsf{ODH\text{-}REAL\text{-}GG}}_{\mathcal{L},p,\mathsf{hLen}}$, $\mathbb{G}^{\mathsf{ODH\text{-}RAND\text{-}GG}}_{\mathcal{L},p,\mathsf{hLen}}$. The advantage of violating ODH is defined as

$$\mathsf{Adv}^{\mathsf{ODH\text{-}GG}}_{\mathcal{L},p,\mathsf{hLen}}(\mathcal{A}) = \left|\Pr\left[\mathbb{G}^{\mathsf{ODH\text{-}REAL\text{-}GG}}_{\mathcal{L},p,\mathsf{hLen}}(\mathcal{A}) \Rightarrow 1\right] - \Pr\left[\mathbb{G}^{\mathsf{ODH\text{-}RAND\text{-}GG}}_{\mathcal{L},p,\mathsf{hLen}}(\mathcal{A}) \Rightarrow 1\right]\right| .$$

Strong Diffie-Hellman Assumption (SDH). This is a stronger version of the classical CDH assumption. This assumption roughly states that CDH is hard even in the presence of a DDH-oracle $\mathsf{O_v}$ where $\mathsf{O_v}(g^x, g^y)$ is $\mathsf{true}$ if and only if $x \cdot v = y$.

We formally define the game $\mathbb{G}^{\mathsf{SDH\text{-}GG}}$ in the generic group model in Figure **??**. The advantage of violating SDH is defined as

$$\mathsf{Adv}^{\mathsf{SDH\text{-}GG}}_{\mathcal{L},p,\mathsf{hLen}}(\mathcal{A}) = \left|\Pr\left[\mathbb{G}^{\mathsf{SDH\text{-}GG}}_{\mathcal{L},p,\mathsf{hLen}}(\mathcal{A}) \Rightarrow \mathsf{true}\right]\right| .$$

Note in particular that one *can* upper bound this advantage unconditionally. We shall drop the $\mathcal{L}$ from the subscript of advantages and games henceforth since the set of labels $\mathcal{L}$ remains the same throughout our paper.

Black Box reductions in the GGM. We consider black-box reductions in the generic group model. We will limit ourselves to an informal description, but this can easily be formalized within existing formal frameworks for reductions (see e.g. [**?**]). We let the reduction $\mathcal{R}$ access an adversary $\mathcal{A}$, and denote by $\mathcal{R}^{\mathcal{A}}$

the resulting algorithm – understood here is that $\mathcal{R}$ supplies inputs, answers queries, etc. In addition, we let $\mathcal{R}$ and $\mathcal{A}$ access the $\mathsf{Eval}$ oracle available in the GGM. We stress that the GGM oracle is not under the reduction's control here – typically, the reduction itself will break a (hard) problem in the GGM with help of $\mathcal{A}$. We will allow (for simplicity) $\mathcal{A}$ to be run depending on some secret private coins[6] not accessible by $\mathcal{R}$. Reductions can run $\mathcal{A}$ several times (with fresh private coins). We call a reduction *straigthline* if it only runs $\mathcal{A}$ *once*.

In our case, the reduction $\mathcal{R}$ will be playing $\mathbb{G}_{p,\mathsf{hLen}}^{\mathsf{SDH\text{-}GG}}$. It receives as inputs $\sigma(1)$, $U = \sigma(u)$, $V = \sigma(v)$, and has access to the $\mathsf{Eval}$, $\mathsf{O_v}$ oracles, as well as an adversary $\mathcal{A}$ for $\mathbb{G}_{p,\mathsf{hLen}}^{\mathsf{ODH\text{-}REAL\text{-}GG}}$ or $\mathbb{G}_{p,\mathsf{hLen}}^{\mathsf{ODH\text{-}RAND\text{-}GG}}$. The reduction needs therefore to supply inputs $(\sigma(1), U', V', W)$ to $\mathcal{A}$, and to answer its queries to the oracles $\mathsf{H_v}$, as well as queries to $\mathsf{H}$. We will call such a reduction *restricted* if it is straightline *and* $V' = V$.

### 2.2 Compression Lemma

In our lower bound proof we use the compression lemma that was formalized in [**?**] which roughly means that it is impossible to compress every element in a set with cardinality $c$ to a string less than $\log c$ bits long, even relative to a random string. We state the compression lemma here as a proposition.

**Proposition 1.** *Suppose, there is a (not necessarily efficient) procedure* $\mathsf{Encode} : \mathcal{X} \times \mathcal{R} \to \mathcal{Y}$ *and a (not necessarily efficient) decoding procedure* $\mathsf{Decode} : \mathcal{Y} \times \mathcal{R} \to \mathcal{X}$ *such that*

$$\Pr_{x \in \mathcal{X}, r \in \mathcal{R}} \left[\mathsf{Decode}(\mathsf{Encode}(x, r), r) = x\right] \geqslant \epsilon \,,$$

*then* $\log |\mathcal{Y}| \geqslant \log |\mathcal{X}| - \log(1/\epsilon)$.

### 2.3 Polynomials

Let $\mathsf{p}(X_1, \cdots, X_n)$ be a $n$ variate polynomial. We denote by $\mathsf{p}(x_1, \cdots, x_n)$ the evaluation of $\mathsf{p}$ at the point $(x_1, \cdots, x_n)$ throughout the paper. The polynomial ring in variables $X_1, \cdots, X_n$ over the field $\mathbb{Z}_p$ is denoted by $\mathbb{Z}_p[X_1, \cdots, X_n]$.

### 2.4 Key Encapsulation Mechanism (KEM)

A key-encapsulation mechanism (KEM) consists of three probabilistic polynomial time (PPT) algorithms $\mathsf{Gen}, \mathsf{Encap}, \mathsf{Decap}$. The key generation algorithm $\mathsf{Gen}$ is probabilistic and outputs a key-pair $(\mathsf{pk}, \mathsf{sk})$. The encapsulation algorithm $\mathsf{Encap}$ is a probabilistic algorithm that takes $\mathsf{pk}$ as input and outputs a ciphertext $c$ and a key $K$ where $K \in \mathcal{K}$ for some non-empty set $\mathcal{K}$. The decapsulation algorithm $\mathsf{Decap}$ is a deterministic algorithm that takes as input the secret key $\mathsf{sk}$ and a ciphertext $c$ outputs a key $K \in \mathcal{K}$ if $(\mathsf{sk}, c)$ is a valid secret key-ciphertext pair and $\perp$ otherwise. For correctness, it is required that for all pairs

[6] If we want to allow the reduction to control random bits, we model them explicitly as an additional input.

$(\mathsf{pk}, \mathsf{sk})$ output by $\mathsf{Gen}$, if $(K, c)$ is output by $\mathsf{Encap}(\mathsf{pk})$ then $K$ is the output of $\mathsf{Decap}(\mathsf{sk}, c)$.

Single challenge KEM-CCA security. The single challenge CCA security of a KEM is defined by a pair of games called $\mathbb{G}^{\mathsf{KEM\text{-}CCA\text{-}REAL}}, \mathbb{G}^{\mathsf{KEM\text{-}CCA\text{-}RAND}}$. In both games a $(\mathsf{pk}, \mathsf{sk})$ pair is generated by $\mathsf{Gen}$, and $(c, K)$ is output by the encapsulation algorithm $\mathsf{Encap}$ on input $\mathsf{pk}$. The adversary is provided with $(\mathsf{pk}, c, K)$ in $\mathbb{G}^{\mathsf{KEM\text{-}CCA\text{-}REAL}}$ and with $(\mathsf{pk}, c, K')$ in $\mathbb{G}^{\mathsf{KEM\text{-}CCA\text{-}RAND}}$ where $K'$ is a randomly sampled element of $\mathcal{K}$. The adversary has access to the decapsulation oracle with $\mathsf{sk}$ as the secret key and it can make decapsulation queries on any ciphertext except the ciphertext $c$ and has to output a bit. We define the advantage of violating single challenge KEM-CCA security is defined as the absolute value of the difference of probabilities of the adversary outputting 1 in the two games. A KEM is single challenge CCA-secure if for all non-uniform PPT adversaries the advantage of violating single challenge KEM-CCA security is negligible.

Single challenge KEM-CCA of Hashed ElGamal. We describe the KEM for Hashed ElGamal in a group with order $p$ and generator $g$ and a hash function $\mathsf{H}$. The function $\mathsf{Gen}$ samples $v$ at random from $\mathbb{Z}_p$, and returns $(g^v, v)$ as the $(\mathsf{pk}, \mathsf{sk})$ pair. The function $\mathsf{Encap}$ on input $v$, samples $u$ at random from $\mathbb{Z}_p$ and returns $g^u$ as the ciphertext and $\mathsf{H}(g^{uv})$ as the key $K$. The function $\mathsf{Decap}$ on input $c$, returns $\mathsf{H}(c^v)$. Note that $\mathsf{Decap}$ in KEM of Hashed ElGamal is identical to the $\mathsf{H}_\mathsf{v}$ function as defined in the ODH assumption. It follows that the single challenge KEM-CCA security of Hashed ElGamal is equivalent to the ODH assumption. In particular, in the generic group model when $\mathsf{H}$ is modeled as a random oracle, the single challenge KEM-CCA security of Hashed ElGamal is equivalent to the ODH assumption in the random oracle and generic group model.

## 3 Memory Lower Bound on the ODH-SDH Reduction

### 3.1 Result and Proof Outline

In this section, we prove a memory lower bound for restricted black-box reductions from ODH to SDH. We stress that the restricted reduction has access only to the $\mathsf{H}, \mathsf{H}_\mathsf{v}$ queries of the adversary. As discussed above, the ODH assumption is equivalent to the single-challenge KEM-CCA security of Hashed ElGamal, this proves a memory lower-bound for (restricted) black-box reductions of single challenge KEM-CCA security of Hashed ElGamal to the SDH assumption.

**Theorem 1 (Main Theorem).** *In the generic group model, with group order $p$, there exists an ODH adversary $\mathcal{A}$ that makes $k$ $\mathsf{H}$ queries and $k$ $\mathsf{H}_\mathsf{v}$ queries (where $k$ is a polynomial in $\log p$), a function $\epsilon_1(p, \mathsf{hLen})$ which is negligible in $\log p, \mathsf{hLen}$, and a function $\epsilon_2(p)$ which is negligible in $\log p$, such that,*

1. $\mathsf{Adv}^{\mathsf{ODH\text{-}GG}}_{p,\mathsf{hLen}}(\mathcal{A}) = 1 - \epsilon_1(p, \mathsf{hLen})$.

2. *For all restricted black-box reductions $\mathcal{R}$, with $s$ bits of memory and making a total of $q$ (assuming $q \geqslant k$) queries to $\mathsf{O}_{\mathsf{v}}$, $\mathsf{Eval}$,*

$$\mathsf{Adv}^{\mathsf{SDH\text{-}GG}}_{p,\mathsf{hLen}}(\mathcal{R}^{\mathcal{A}}) \leqslant 2 \cdot 2^{\frac{s}{2}} \left(\frac{48q^3}{p}\right)^{\frac{k}{8c}} \left(1 + \frac{6q}{p}\right)^{q} + \frac{4q^2 \log p + 13q^2 + 5q}{p} + \epsilon_2(p) \ ,$$

*where* $c = 4\lceil \frac{\log q}{\log k} \rceil$.

This result implies that if $\mathsf{Adv}^{\mathsf{SDH\text{-}GG}}_{p,\mathsf{hLen}}(\mathcal{R}^{\mathcal{A}})$ is non-negligible for a reduction $\mathcal{R}$ making $q$ queries where $q$ is a polynomial in $\log p$, then $s = \Omega(k \log p)$ i.e. the memory required by any restricted black-box reduction grows with the number of queries by $\mathcal{A}$. Hence, there does not exist any efficient restricted black-box reduction from ODH to SDH that is memory-tight.

In the full version, we discuss how rewinding can slightly improve the memory complexity to (roughly) $O(k \log k)$, with heavy computational cost (essentially, one rewinding per oracle query of the adversary). We conjecture this to be optimal, but a proof seems to evade current techniques.

De-randomization. Before we turn to the proof – which also connects several technical lemmas presented across the next sections, let us discuss some aspects of the results. As explained above, our model allows for the adversary $\mathcal{A}$ to be run with randomness unknown to $\mathcal{R}$. This aspect *may* be controversial, but we note that there is a generic way for $\mathcal{A}$ to be made deterministic. Recall that $\mathcal{A}$ must be inefficient for the separation to even hold true. For example, $\mathcal{A}$ can use the injection $\sigma$ from the generic group model to generate its random coin – say, using $\sigma^{-1}(\mathbf{a}_i)$ as coins a priori fixed labels $\mathbf{a}_1, \mathbf{a}_2, \ldots$. It is a standard – albeit tedious and omitted – argument to show that unless the reduction ends up querying the pre-images (which happens with negligible probability only), the $\sigma^{-1}(\mathbf{a}_i)$'s are good random coins.

Strengthening beyond SDH. We would like to note that our result can be strengthened without much effort to a reduction between ODH and a more general version of SDH. Informally, we can extend our result to every problem which is hard in the generic group model in presence of an $\mathsf{O}_{\mathsf{v}}$ oracle. For example, this could be a problem where given $g$, $g^u$, and $g^v$, the attacker needs to output $g^{f(u,v)}$, where $f$ is (a fixed) two-variate polynomial with degree at least 2. We do not include the proof for the strengthened version for simplicity. However, it appears much harder to extend our result to different types of oracles than $\mathsf{O}_{\mathsf{v}}$, as our proof is tailored at this oracle.

*Proof.* Here, we give the overall structure, the key lemmas, and how they are combined – quantitatively – to obtain the final result.

First off, Lemma **??** establishes that there exists an adversary $\mathcal{A}$ such that $\mathsf{Adv}^{\mathsf{ODH\text{-}GG}}_{p,\mathsf{hLen}}(\mathcal{A})$ is close to 1, which we will fix (i.e., when we refer to $\mathcal{A}$, we refer to the one guaranteed to exist by the lemma). The proof of Lemma **??** is in Section **??** and the proof of Lemma **??** is in Section **??**.

**Lemma 1.** *There exists an adversary $\mathcal{A}$ and a function $\epsilon_1(p, \mathsf{hLen})$ such that is negligible in $\log p, \mathsf{hLen}$, and*

$$\mathsf{Adv}^{\mathsf{ODH\text{-}GG}}_{p,\mathsf{hLen}}(\mathcal{A}) = 1 - \epsilon_1(p, \mathsf{hLen}) \ .$$

After that, we introduce a game, called $\mathbb{G}_1$ and described in Figure **??** in Section **??**. Very informally, this is a game played by a two-stage adversary $\mathcal{R}_1, \mathcal{R}_2$ which can pass a state to each other of size $s$ bits and have access to the $\mathsf{Eval}, \mathsf{O_v}$ oracles. The game captures the essence of the reduction $\mathcal{R}$ the adversary $\mathcal{A}$ of having a sufficient amount of memory. This is made formal in Lemma **??**, where we show that the probability of the reduction $\mathcal{R}$ winning the SDH-GG game while running $\mathcal{A}$ is bounded by the probability of winning $\mathbb{G}_1$.

**Lemma 2.** *For every restricted black box reduction $\mathcal{R}$ to SDH-GG that runs $\mathcal{A}$, there exist adversaries $\mathcal{R}_1, \mathcal{R}_2$ playing $\mathbb{G}_1$, such that the number of queries made by $\mathcal{R}_1, \mathcal{R}_2$ to $\mathsf{Eval}, \mathsf{O_v}$ is same as the number of queries made by $\mathcal{R}$ to $\mathsf{Eval}, \mathsf{O_v}$, the state passed from $\mathcal{R}_1$ to $\mathcal{R}_2$ is upper bounded by the memory used by $\mathcal{R}$ and,*

$$\mathsf{Adv}^{\mathsf{SDH\text{-}GG}}_{p,\mathsf{hLen}}(\mathcal{R}^{\mathcal{A}}) \leqslant \Pr\left[\mathbb{G}_1 \Rightarrow \mathsf{true}\right] + \frac{4k^2(\log p)^2}{p} + \frac{4qk\log p + q^2}{p} \ .$$

We introduce Games $\mathbb{G}_2, \mathbb{G}_3$ in Figure **??** in Section **??**. These games are identical to $\mathbb{G}_1$ except for the condition to output $\mathsf{true}$. The condition to output $\mathsf{true}$ in these games are disjoint and the disjunction of the two conditions is equivalent to the condition to output $\mathsf{true}$ in $\mathbb{G}_1$. A little more specifically, both games depend on a parameter $l$, which can be set arbitrarily, and in $\mathbb{G}_3$ and $\mathbb{G}_2$ the winning condition of $\mathbb{G}_1$ is strengthened by additional ensuring that a certain set defined during the execution of the game is smaller or larger than $l$, respectively. Therefore, tautologically,

$$\Pr\left[\mathbb{G}_1 \Rightarrow \mathsf{true}\right] = \Pr\left[\mathbb{G}_2 \Rightarrow \mathsf{true}\right] + \Pr\left[\mathbb{G}_3 \Rightarrow \mathsf{true}\right] \ . \tag{1}$$

We now prove the following two lemmas below, in Sections **??** and **??**,

**Lemma 3.** *For the game $\mathbb{G}_2$,*

$$\Pr\left[\mathbb{G}_2 \Rightarrow \mathsf{true}\right] \leqslant \frac{q^l}{k!} + \frac{2q(2k+3q+2)}{p} + \frac{5q}{p} + \frac{k^2+k+2}{p} \ .$$

**Lemma 4.** *If the size of the state $\phi$ output by $\mathcal{R}_1$ is $s$ bits and $(\mathcal{R}_1, \mathcal{R}_2)$ make $q$ queries in total in $\mathbb{G}_3$, then*

$$\Pr\left[\mathbb{G}_3 \Rightarrow \mathsf{true}\right] \leqslant 2 \cdot 2^{\frac{s}{2}} \left(\frac{8q^2(2k+2+3q)}{p}\right)^{\frac{l}{2}} \left(1 + \frac{6q}{p}\right)^{\frac{2q-l}{2}} + \frac{k^2+k+2}{p} \ .$$

Combining **??** and the result of Lemmas **????** we get,

$$\Pr\left[\mathbb{G}_1 \Rightarrow \mathsf{true}\right] \leqslant 2 \cdot 2^{\frac{s}{2}} \left(\frac{8q^2(2k+2+3q)}{p}\right)^{\frac{l}{2}} \left(1 + \frac{6q}{p}\right)^{\frac{2q-l}{2}} + \frac{2(k^2+k+2)}{p} + \frac{q^l}{k!} + \frac{2q(2k+3q+2)}{p} + \frac{5q}{p} \ . \tag{2}$$

Since $\left(1+\frac{6q}{p}\right)^{\frac{2q-l}{2}} \leqslant \left(1+\frac{6q}{p}\right)^{q}$, combining Lemma **??**, **??** we get,

$$\mathsf{Adv}^{\mathsf{SDH\text{-}GG}}_{p,\mathsf{hLen}}(\mathcal{R}^{\mathcal{A}}) \leqslant 2\cdot 2^{\frac{s}{2}}\left(\frac{8q^2(2k+2+3q)}{p}\right)^{\frac{l}{2}}\left(1+\frac{6q}{p}\right)^{q}+\frac{2(k^2+k+2)}{p}+ \frac{2q(2k+3q+2)}{p}+\frac{5q}{p}+\frac{4k^2(\log p)^2}{p}+\frac{4qk\log p+q^2}{p}+\frac{q^l}{k!}\,.$$

We let,

$$\epsilon_2(p)=\frac{q^l}{k!}+\frac{2(k^2+k+2)}{p}+\frac{4k^2(\log p)^2}{p}\,.$$

Setting $c=\lceil\frac{\log q}{\log k}\rceil$ and $l=\frac{k}{4c}$, $\frac{q^l}{k!}\leqslant\frac{k^{k/4}}{k!}$. By Sterling's approximation $k!\geqslant k^{k+1/2}e^{-k}$. Therefore,

$$\frac{k^{k/4}}{k!}=\frac{k^{k/4}}{k^{k/4}}\frac{e^k}{k^{k/4}}\frac{1}{k^{k/2+1/2}}\,.$$

For $k>e^4$ (we can set $k>e^4$), $\frac{q^l}{k!}\leqslant\frac{1}{k^{k/2+1/2}}$ i.e. $\frac{q^l}{k!}$ is negligible in $\log p$ for $k$ polynomial in $\log p$. Also, $\frac{2(k^2+k+2)}{p}+\frac{4k^2(\log p)^2}{p}$ is negligible in $\log p$ (since $k$ is a polynomial in $\log p$). So, $\epsilon_2(p)$ is negligible in $\log p$. We have that,

$$\mathsf{Adv}^{\mathsf{SDH\text{-}GG}}_{p,\mathsf{hLen}}(\mathcal{R}^{\mathcal{A}}) \leqslant 2\cdot 2^{\frac{s}{2}}\left(\frac{8q^2(2k+2+3q)}{p}\right)^{\frac{k}{8c}}\left(1+\frac{6q}{p}\right)^{q}+ \frac{2q(2k+3q+2)}{p}+\frac{5q}{p}+\frac{4qk\log p+q^2}{p}+\epsilon_2(p)\,.$$

where $c=4\lceil\frac{\log q}{\log k}\rceil$. Assuming $q\geqslant k$ (and thus $q>e^4>2$), we get,

$$\mathsf{Adv}^{\mathsf{SDH\text{-}GG}}_{p,\mathsf{hLen}}(\mathcal{R}^{\mathcal{A}}) \leqslant 2\cdot 2^{\frac{s}{2}}\left(\frac{48q^3}{p}\right)^{\frac{k}{8c}}\left(1+\frac{6q}{p}\right)^{q}+\frac{4q^2\log p+13q^2+5q}{p}+\epsilon_2(p)\,.$$

□

# 4 Proof of Theorem

## 4.1 Adversary $\mathcal{A}$ against ODH

In this section, we construct the ODH adversary $\mathcal{A}$ needed for the proof.

**Lemma ??.** *There exists an adversary $\mathcal{A}$ and a function $\epsilon_1(p,\mathsf{hLen})$ such that is negligible in $\log p, \mathsf{hLen}$, and*

$$\mathsf{Adv}^{\mathsf{ODH\text{-}GG}}_{p,\mathsf{hLen}}(\mathcal{A})=1-\epsilon_1(p,\mathsf{hLen})\,.$$

The adversary $\mathcal{A}$ is formally defined in Figure **??**. The proof of Lemma **??** itself is deferred to the full version. Adversary $\mathcal{A}$ samples $i_1,\cdots,i_k$ from $\mathbb{Z}_p$, and computes $\sigma(i_j),\sigma(i_j\cdot v)$ for all $j$ in $[k]$. It then makes $\mathsf{H_v}$ queries on $\sigma(i_j)$'s for

**Adversary** $\mathcal{A}^{\mathsf{H_v}(.),\mathsf{H}(.),\mathsf{Eval}(.,.)}(U,V,W,\sigma(1))$ :

$i_1,\cdots,i_k \leftarrow\!\!\$\ \mathbb{Z}_p$
**foreach** $j \in [k]$ **do**
  $Q_1[j] \leftarrow \mathsf{Exp}(\sigma(1), i_j); Q_2[j] \leftarrow \mathsf{Exp}(V, i_j)$
$\mathsf{honest} \leftarrow 1$
**foreach** $j \in [k]$ **do**
  $\mathsf{ans}_1[j] \leftarrow \mathsf{H_v}(Q_1[j])$
$\pi \leftarrow\!\!\$\ \mathcal{S}_k$
**foreach** $j \in [k]$ **do**
  $\mathsf{ans}_2[\pi(j)] \leftarrow \mathsf{H}(Q_2[\pi(j)])$
**if** $\exists j, l \in [k], j \neq l : (\mathsf{ans}_1[j] = \mathsf{ans}_1[l] \vee \mathsf{ans}_2[j] = \mathsf{ans}_2[l])$ **then** $\mathsf{honest} \leftarrow 0$
**if** $\exists j \in [k] : \mathsf{ans}_1[j] \neq \mathsf{ans}_2[j]$ **then** $\mathsf{honest} \leftarrow 0$
**if** $\mathsf{honest} = 1$ **then**
  $\mathsf{temp} \leftarrow \sigma(1); v \leftarrow 1$
  **while** $(\mathsf{temp} \neq V)$
    $\mathsf{temp} \leftarrow \mathsf{Eval}(\mathsf{temp}, \sigma(1)); v \leftarrow v + 1$
  $\mathsf{inp} \leftarrow \mathsf{Exp}(U, v); W' \leftarrow \mathsf{H}(\mathsf{inp}); b \leftarrow (W' = W)$
**else** $b \leftarrow\!\!\$\ \{0,1\}$
**return** $b$

**Fig. 2.** The adversary $\mathcal{A}$

all $j$ in $[k]$. Adversary $\mathcal{A}$ then samples a permutation $\pi$ on $[k] \to [k]$, and then makes $\mathsf{H}$ queries on $\sigma(i_{\pi(j)} \cdot v)$'s for all $j$ in $[k]$. If answers of all the $\mathsf{H}$ queries are distinct and the answers of all the $\mathsf{H_v}$ queries are distinct and for all $j$ in $[k]$, $\mathsf{H_v}(\sigma(i_j)) = \mathsf{H}(\sigma(i_j \cdot v))$, $\mathcal{A}$ computes the discrete logarithm of $V$ outputs the correct answer. Otherwise it returns a bit uniformly at random. Note that $\mathcal{A}$ is inefficient, but only if it is satisfied from the responses it gets from the reduction using it.

### 4.2 The Shuffling Games

The Game $\mathbb{G}_1$. We first introduce the two-stage game $\mathbb{G}_1$ played by a pair of adversaries $\mathcal{R}_1$ and $\mathcal{R}_2$. (With some foresight, these are going to be two stages of the reduction.) It is formally described in Figure **??**. Game $\mathbb{G}_1$ involves sampling $\sigma, i_1, \cdots, i_k, v$ from $\mathbb{Z}_p$, then running $\mathcal{R}_1$, followed by sampling permutation $\pi$ from $\mathcal{S}_k$ and then running $\mathcal{R}_2$. The first stage $\mathcal{R}_1$ has inputs $\sigma(i_1), \cdots, \sigma(i_k)$ and it outputs a state $\phi$ of $s$ bits along with $k$ strings in $\{0,1\}^{\mathsf{hLen}}$. The second stage $\mathcal{R}_2$ has inputs $\phi, \sigma(i_{\pi(1)} \cdot v), \cdots, \sigma(i_{\pi(k)} \cdot v)$ and it outputs $k$ strings in $\{0,1\}^{\mathsf{hLen}}$. Both the stages $\mathcal{R}_1, \mathcal{R}_2$ have access to oracles $\mathsf{Eval}, \mathsf{O_v}$. Game $\mathbb{G}_1$ outputs $\mathsf{true}$ if all the $k$ strings output by $\mathcal{R}_1$ are distinct, and if all the $k$ strings output by $\mathcal{R}_2$ are distinct, and if for all $j \in [k]$, the $j^{\text{th}}$ string output by $\mathcal{R}_2$ is identical to the $\pi(j)^{\text{th}}$ string output by $\mathcal{R}_1$. Additionally, $\mathbb{G}_1$ involves some bookkeeping. The $\mathsf{Eval}, \mathsf{O_v}$ oracles in $\mathbb{G}_1$ take an extra parameter named $\mathsf{from}$ as input which indicates whether the query was from $\mathcal{R}_1$ or $\mathcal{R}_2$.

**Game** $\mathbb{G}_1$ :

$\sigma \leftarrow\$ \mathsf{InjFunc}(\mathbb{Z}_p, \mathcal{L}); i_1, \cdots, i_k, v \leftarrow\$ \mathbb{Z}_p$

$\mathcal{X} \leftarrow \{\sigma(1), \sigma(v), \sigma(i_1), \cdots, \sigma(i_k)\}; \mathcal{Y}_1 \leftarrow \{\sigma(1), \sigma(v), \sigma(i_1), \cdots, \sigma(i_k)\}$

$\phi, s_1, \cdots, s_k \leftarrow \mathcal{R}_1^{\mathsf{Eval}(.,.,1), \mathsf{O_v}(.,.,1)}(\sigma(1), \sigma(v), \sigma(i_1), \cdots, \sigma(i_k))$

$\pi \leftarrow\$ \mathcal{S}_k; \mathcal{Y}_2 \leftarrow \{\sigma(1), \sigma(v), \sigma(i_1 \cdot v), \cdots, \sigma(i_k \cdot v)\}; \mathcal{Z} \leftarrow \varnothing$

$s'_1, s'_2, \cdots, s'_k \leftarrow \mathcal{R}_2^{\mathsf{Eval}(.,.,2), \mathsf{O_v}(.,.,2)}(\phi, \sigma(1), \sigma(v), \sigma(i_{\pi(1)} \cdot v), \cdots, \sigma(i_{\pi(k)} \cdot v))$

$\mathsf{win} \leftarrow (\forall j \in [k] : s_{\pi(j)} = s'_j) \wedge (\forall j, l \in [k] : j \neq l \implies s_j \neq s_l \wedge s'_j \neq s'_l)$

**return** win

**Oracle** $\mathsf{Eval}(\mathbf{a}, \mathbf{b}, \mathsf{from})$ :

$\mathbf{c} \leftarrow \sigma(\sigma^{-1}(\mathbf{a}) + \sigma^{-1}(\mathbf{b}))$

**if** from $= 1$ **then**

  **if** $\mathbf{c} \notin \mathcal{Y}_1$ **then** $\mathcal{X} \xleftarrow{\cup} \{\mathbf{c}\}$

  $\mathcal{Y}_1 \xleftarrow{\cup} \{\mathbf{a}, \mathbf{b}, \mathbf{c}\}$

**if** from $= 2$ **then**

  **if** $\mathbf{a} \in \mathcal{X} \backslash \mathcal{Y}_2$ **then** $\mathcal{Z} \xleftarrow{\cup} \{\mathbf{a}\}$

  **if** $\mathbf{b} \in \mathcal{X} \backslash \mathcal{Y}_2$ **then** $\mathcal{Z} \xleftarrow{\cup} \{\mathbf{b}\}$

  $\mathcal{Y}_2 \xleftarrow{\cup} \{\mathbf{a}, \mathbf{b}, \mathbf{c}\}$

**return** $\mathbf{c}$

**Oracle** $\mathsf{O_v}(\mathbf{a}, \mathbf{b}, \mathsf{from})$ :

**if** from $= 1$ **then** $\mathcal{Y}_1 \xleftarrow{\cup} \{\mathbf{a}, \mathbf{b}\}$

**if** from $= 2$ **then**

  **if** $\mathbf{a} \in \mathcal{X} \backslash \mathcal{Y}_2$ **then** $\mathcal{Z} \xleftarrow{\cup} \{\mathbf{a}\}$

  **if** $\mathbf{b} \in \mathcal{X} \backslash \mathcal{Y}_2$ **then** $\mathcal{Z} \xleftarrow{\cup} \{\mathbf{b}\}$

  $\mathcal{Y}_2 \xleftarrow{\cup} \{\mathbf{a}, \mathbf{b}\}$

**return** $(v \cdot \sigma^{-1}(\mathbf{a}) = \sigma^{-1}(\mathbf{b}))$

**Fig. 3.** Game $\mathbb{G}_1$. We use the phrase $\mathcal{R}_1, \mathcal{R}_2$ win $\mathbb{G}_1$ to mean $\mathbb{G}_1 \Rightarrow \mathsf{true}$. We shall use this convention for all games in the paper that output boolean values.

We introduce the phrase "seen by" before describing the bookkeeping. A label has been "seen by" $\mathcal{R}_1$ if it was an input to $\mathcal{R}_1$, queried by $\mathcal{R}_1$ or an answer to a previously made $\mathsf{Eval}(.,.,1)$ query. A label has been "seen by" $\mathcal{R}_2$ if it was an input to $\mathcal{R}_2$, queried by $\mathcal{R}_2$ or an answer to a previously made $\mathsf{Eval}(.,.,2)$ query. We describe the sets $\mathcal{X}, \mathcal{Y}_1, \mathcal{Y}_2, \mathcal{Z}$ which are used for bookkeeping in $\mathbb{G}_1$.

- The labels in $\mathcal{X}$ are answers to $\mathsf{Eval}(.,.,1)$ queries such that it has not yet been "seen by" $\mathcal{R}_1$ before the query.
- $\mathcal{Y}_1$ contains all the labels that are input to $\mathcal{R}_1$, queried by $\mathcal{R}_1$ or answers to $\mathsf{Eval}(.,.,1)$ queries i.e. it is the set of labels "seen by" $\mathcal{R}_1$.
- $\mathcal{Y}_2$ contains all the labels that are input to $\mathcal{R}_2$, queried by $\mathcal{R}_1$ or answers to $\mathsf{Eval}(.,.,2)$ queries i.e. it is the set of labels "seen by" $\mathcal{R}_2$.
- All labels in $\mathcal{Z}$ are queried by $\mathcal{R}_2$ and have not been "seen by" $\mathcal{R}_2$ before the query and are in $\mathcal{X}$

The following lemma tells us that we can (somewhat straightforwardly) take a reduction as in the theorem statement, and transform it into an equivalent pair $\mathcal{R}_1, \mathcal{R}_2$ of adversaries for $\mathbb{G}_1$. The point here is that the reduction is very unlikely to succeed in breaking the SDH assumption without doing an effort equivalent to winning $\mathbb{G}_1$ to get $\mathcal{A}$'s help – otherwise, it is left with breaking SDH directly in the generic group model, which is hard. The proof is deferred to the full version.

**Game** $\boxed{\mathbb{G}_2}, \boxed{\mathbb{G}_3}$ :

1 : $\sigma \leftarrow\$ \mathsf{InjFunc}(\mathbb{Z}_p, \mathcal{L}); i_1, \cdots, i_k, v \leftarrow\$ \mathbb{Z}_p$

2 : $\mathcal{X} \leftarrow \{\sigma(1), \sigma(v), \sigma(i_1), \cdots, \sigma(i_k)\}; \mathcal{Y}_1 \leftarrow \{\sigma(1), \sigma(v), \sigma(i_1), \cdots, \sigma(i_k)\}$

3 : $\phi, s_1, \cdots, s_k \leftarrow \mathcal{R}_1^{\mathsf{Eval}(.,.,1), \mathsf{O_v}(.,.,1)}(\sigma(1), \sigma(v), \sigma(i_1), \cdots, \sigma(i_k))$

4 : $\pi \leftarrow\$ \mathcal{S}_k; \mathcal{Y}_2 \leftarrow \{\sigma(1), \sigma(v), \sigma(i_1 \cdot v), \cdots, \sigma(i_k \cdot v)\}; \mathcal{Z} \leftarrow \varnothing$

5 : $s'_1, s'_2, \cdots, s'_k \leftarrow \mathcal{R}_2^{\mathsf{Eval}(.,.,2), \mathsf{O_v}(.,.,2)}(\phi, \sigma(1), \sigma(v), \sigma(i_{\pi(1)} \cdot v), \cdots, \sigma(i_{\pi(k)} \cdot v))$

6 : $\mathsf{win} \leftarrow (\forall j \in [k] : s_{\pi(j)} = s'_j) \wedge (\forall j, l \in [k] : j \neq l \implies s_j \neq s_l \wedge s'_j \neq s'_l)$

7 : $\boxed{\textbf{return } (\mathsf{win} \wedge |\mathcal{Z}| < l)}$ $\boxed{\textbf{return } (\mathsf{win} \wedge |\mathcal{Z}| \geqslant l)}$

**Fig. 4.** Games $\mathbb{G}_2, \mathbb{G}_3$. The $\mathsf{Eval}, \mathsf{O_v}$ oracles in $\mathbb{G}_2, \mathbb{G}_3$ are identical to those in $\mathbb{G}_1$ and hence we do not rewrite it here. The newly introduced changes compared to $\mathbb{G}_1$ are highlighted. The statement within the thinner box is present only in $\mathbb{G}_3$ and the statement within the thicker box is present only in $\mathbb{G}_2$.

**Lemma ??.** *For every restricted black box reduction $\mathcal{R}$ to* SDH-GG *that runs $\mathcal{A}$, there exist adversaries $\mathcal{R}_1, \mathcal{R}_2$ playing $\mathbb{G}_1$, such that the number of queries made by $\mathcal{R}_1, \mathcal{R}_2$ to $\mathsf{Eval}, \mathsf{O_v}$ is same as the number of queries made by $\mathcal{R}$ to $\mathsf{Eval}, \mathsf{O_v}$, the state passed from $\mathcal{R}_1$ to $\mathcal{R}_2$ is upper bounded by the memory used by $\mathcal{R}$ and,*

$$\mathsf{Adv}^{\mathsf{SDH\text{-}GG}}_{p,\mathsf{hLen}}(\mathcal{R}^{\mathcal{A}}) \leqslant \Pr\left[\mathbb{G}_1 \Rightarrow \mathsf{true}\right] + \frac{4k^2(\log p)^2}{p} + \frac{4qk \log p + q^2}{p} .$$

The Games $\mathbb{G}_2$ and $\mathbb{G}_3$. In Figure **??** we define $\mathbb{G}_2, \mathbb{G}_3$ which have an added check on the cardinality of $\mathcal{Z}$ to output $\mathsf{true}$. Everything else remains unchanged (in particular $\mathsf{Eval}, \mathsf{O_v}$ are unchanged from $\mathbb{G}_1$ and we do not specify them again here). The statement within the thinner box is present only in $\mathbb{G}_3$ and statement within the thicker box is present only in $\mathbb{G}_2$. The changes from $\mathbb{G}_1$ have been highlighted. We shall follow these conventions of using boxes and highlighting throughout the paper.

The games $\mathbb{G}_2, \mathbb{G}_3$ are identical to $\mathbb{G}_1$ except for the condition to output $\mathsf{true}$. Since this disjunction of the conditions to output $\mathsf{true}$ in $\mathbb{G}_2, \mathbb{G}_3$ is equivalent to the condition to output $\mathsf{true}$ in $\mathbb{G}_1$, and the conditions to output $\mathsf{true}$ in $\mathbb{G}_2, \mathbb{G}_3$ are disjoint, we have,

$$\Pr\left[\mathbb{G}_1 \Rightarrow \mathsf{true}\right] = \Pr\left[\mathbb{G}_2 \Rightarrow \mathsf{true}\right] + \Pr\left[\mathbb{G}_3 \Rightarrow \mathsf{true}\right] .$$

### 4.3 Proof of Lemma ??

Recall we are going to prove the following lemma.

**Lemma ??.** *For the game $\mathbb{G}_2$,*

$$\Pr\left[\mathbb{G}_2 \Rightarrow \mathsf{true}\right] \leqslant \frac{q^l}{k!} + \frac{2q(2k + 3q + 2)}{p} + \frac{5q}{p} + \frac{k^2 + k + 2}{p} .$$

| **Game** $\mathbb{PG}(\mathcal{A})$ : | **Oracle** $O(\vec{x}, \vec{y})$ : // $\vec{x} \in \mathbb{Z}_p^k, \vec{y} \in \mathbb{Z}_p^k$ |
|---|---|
| 1 : $\pi \leftarrow\$ \mathcal{S}_k$ | 1 : **return** $(\forall i \in [k] : \vec{x}[\pi(i)] = \vec{y}[i])$ |
| 2 : $\pi' \leftarrow \mathcal{A}^{O(\cdot,\cdot)}$ | |
| 3 : **return** $(\pi = \pi')$ | |

**Fig. 5.** The permutation game $\mathbb{PG}$ being played by adversary $\mathcal{A}$ is denoted by $\mathbb{PG}(\mathcal{A})$

We introduce a new game – called the *permutation game* and denoted $\mathbb{PG}$ – in order to upper bound $\Pr[\mathbb{G}_2 \Rightarrow \mathsf{true}]$. In the rest of this proof, we are going to first define the game, and upper bound the winning probability of an adversary. Then, we are going to reduce an adversary for $\mathbb{G}_2$ to one for $\mathbb{PG}$.

The Permutation Game. In Game $\mathbb{PG}$, an adversary has to guess a randomly sampled permutation $\pi$ over $[k]$. The adversary has access to an oracle that takes as input two vectors of length $k$ and returns $\mathsf{true}$ if the elements of the first vector, when permuted using $\pi$, results in the second vector and $\mathsf{false}$ otherwise. Figure **??** formally describes the game $\mathbb{PG}$.

In the following, we say an adversary playing $\mathbb{PG}$ is a $(q, l)$-query adversary if it makes at most $q$ queries to $O$, and the rank of the vectors that were the first argument to the $O$ queries returning $\mathsf{true}$ is at most $l$.

The following lemma – which we prove via a compression argument – yields an upper bound on the probability of winning the game for a $(q, l)$-query adversary.

**Lemma 5.** *For a $(q, l)$-query adversary $\mathcal{A}$ playing $\mathbb{PG}$ the following is true.*

$$\Pr[\mathbb{PG}(\mathcal{A}) \Rightarrow \mathsf{true}] \leqslant \frac{q^l}{k!} .$$

*Proof.* We construct an encoding of $\pi$ by running adversary $\mathcal{A}$. In order to run $\mathcal{A}$, all the $O$ queries need to be correctly answered. This can be naively done by storing the sequence number of queries whose answers are $\mathsf{true}$. In fact, of all such queries, we need to just store the sequence number of just those whose first argument is not in the linear span of vectors which were the first argument of previous such queries i.e. we store the sequence number of only those $O$ queries returning $\mathsf{true}$ whose first argument form a basis of the first argument of all $O$ queries returning $\mathsf{true}$. This approach works because for every vector $\vec{x}$, there is only a unique vector $\vec{y}$ such that $O(\vec{x}, \vec{y}) = 1$. The random tape of the adversary can be derived using the common randomness of $\mathsf{Encode}, \mathsf{Decode}$ and hence the adversary produces identical queries and output. For simplicity, we do not specify this explicitly in the algorithms and treat $\mathcal{A}$ as deterministic. The formal description of the algorithms $\mathsf{Encode}, \mathsf{Decode}$ are in Figure **??**.

Observe that $S$ is a basis of vectors $\vec{x}$ such that $O(\vec{x}, \vec{y}) = \mathsf{true}$. Note that for an $O(\vec{x}, \vec{y})$ query returning $\mathsf{true}$, if $\vec{x} \in S$ then the sequence number of the query is stored in $\mathsf{enc}$. Therefore, $(\vec{x}, \vec{y}) \in S'$ in $\mathsf{Decode}$. Again, for an $O(\vec{x}, \vec{y})$ query returning $\mathsf{true}$, if $\vec{x} \notin S$ then the sequence number of the query is not stored in $\mathsf{enc}$ and therefore $(\vec{x}, \vec{y}) \notin S'$. So, for an $O(\vec{x}, \vec{y})$ query returning $\mathsf{true}$, $(\vec{x}, \vec{y}) \in S'$ iff $\vec{x} \in S$. Since, for all $(\vec{x}, \vec{y})$ such that $O(\vec{x}, \vec{y}) = \mathsf{true}$ we have that

| **Procedure** $\mathsf{Encode}(\pi)$ : | **Oracle** $O(\vec{x}, \vec{y})$ : |
|---|---|
| 1 : $c \leftarrow 0$ | 1 : $c \leftarrow c + 1$ |
| 2 : $S \leftarrow \varnothing$ | 2 : **if** $(\exists i \in [k] : \vec{x}[\pi(i)] \neq \vec{y}[i])$ **then** |
| 3 : $\mathsf{enc} \leftarrow \varnothing$ | 3 : **return** false |
| 4 : $\pi' \leftarrow \mathcal{A}^{O(.,.)}$ | 4 : **else** |
| 5 : **return** enc | 5 : **if** $\vec{x} \notin \mathsf{span}(S)$ **then** |
| | 6 : $S \leftarrow S \cup \{\vec{x}\}$ |
| | 7 : $\mathsf{enc} \leftarrow \mathsf{enc} \cup \{c\}$ |
| | 8 : **return** true |
| **Procedure** $\mathsf{Decode}(\mathsf{enc})$ : | **Oracle** $O(\vec{x}, \vec{y})$ : |
| 1 : $c \leftarrow 0$ | 1 : $c \leftarrow c + 1$ |
| 2 : $S' \leftarrow \varnothing$ | 2 : **if** $c \in \mathsf{enc}$ **then** |
| 3 : $\pi' \leftarrow \mathcal{A}^{O(.,.)}$ | 3 : $S' \leftarrow S' \cup \{(\vec{x}, \vec{y})\}$ |
| 4 : **return** $\pi'$ | 4 : **return** true |
| | 5 : **return** $((\vec{x}, \vec{y}) \in \mathsf{span}(S'))$ |

**Fig. 6.** Encoding and decoding $\pi$ using $\mathcal{A}$

for all $i \in [k]$, $\vec{y}[i] = \vec{x}[\pi^{-1}(i)]$, it follows that $S'$ forms a basis of vectors $(\vec{x}, \vec{y})$ such that $O(\vec{x}, \vec{y}) = \mathsf{true}$.

In $\mathsf{Decode}(\mathsf{enc})$, the simulation of $O(\vec{x}, \vec{y})$ is perfect because

- If $c$ is in $\mathsf{enc}$, then $\vec{x} \in S$ in $\mathsf{Encode}$. From the definition of $S$ in $\mathsf{Encode}$, it follows that $O(\vec{x}, \vec{y})$ should return $\mathsf{true}$.
- Otherwise we check if $(\vec{x}, \vec{y}) \in \mathsf{span}(S')$ and return $\mathsf{true}$ if the check succeeds, $\mathsf{false}$ otherwise. This is correct since in $S'$ is a basis of vectors $(\vec{x}, \vec{y})$ such that $O(\vec{x}, \vec{y}) = \mathsf{true}$.

The encoding is a set of $|S|$ query sequence numbers. Since there are at most $q$ queries, the encoding space is at most $\binom{q}{|S|}$. Using $\mathcal{X}$ to be the set $S_k$, $\mathcal{Y}$ to be the set of all possible encodings, $\mathcal{R}$ to be the set of random tapes of $\mathcal{A}$, it follows from Proposition **??** that,

$$\Pr\left[\text{Decoding is sucessful}\right] \leqslant \frac{\binom{q}{|S|}}{k!} .$$

Since the simulation of $O(\vec{x}, \vec{y})$ is perfect in $\mathsf{Decode}$, decoding is successful if $\mathbb{PG}(\mathcal{A}) \Rightarrow \mathsf{true}$. Therefore,

$$\Pr\left[\mathbb{PG}(\mathcal{A}) \Rightarrow \mathsf{true}\right] \leqslant \frac{\binom{q}{|S|}}{k!} \leqslant \frac{q^{|S|}}{k!} .$$

Since $\mathcal{A}$ is a $(q, l)$-query adversary, $|S| \leqslant l$. Thus, we have,

$$\Pr\left[\mathbb{PG}(\mathcal{A}) \Rightarrow \mathsf{true}\right] \leqslant \frac{q^l}{k!} \tag{3}$$

□

| **Procedure** $\mathsf{PopulateSetsEval}(\mathbf{a},\mathbf{b},\mathbf{c},\mathsf{from})$ : | **Procedure** $\mathsf{PopulateSetsO_v}(\mathbf{a},\mathbf{b},\mathsf{from})$ : |
|---|---|
| 1 : **if** $\mathsf{from} = 1$ **then** | 1 : **if** $\mathsf{from} = 1$ **then** $\mathcal{Y}_1 \xleftarrow{\cup} \{\mathbf{a},\mathbf{b}\}$ |
| 2 : **if** $\mathbf{c} \notin \mathcal{Y}_1$ **then** $\mathcal{X} \xleftarrow{\cup} \{\mathbf{c}\}$ | 2 : **if** $\mathsf{from} = 2$ **then** |
| 3 : $\mathcal{Y}_1 \xleftarrow{\cup} \{\mathbf{a},\mathbf{b},\mathbf{c}\}$ | 3 : **if** $\mathbf{a} \in \mathcal{X}\backslash\mathcal{Y}_2$ **then** $\mathcal{Z} \xleftarrow{\cup} \{\mathbf{a}\}$ |
| 4 : **if** $\mathsf{from} = 2$ **then** | 4 : **if** $\mathbf{b} \in \mathcal{X}\backslash\mathcal{Y}_2$ **then** $\mathcal{Z} \xleftarrow{\cup} \{\mathbf{b}\}$ |
| 5 : **if** $\mathbf{a} \in \mathcal{X}\backslash\ \mathcal{Y}_2$ **then** $\mathcal{Z} \xleftarrow{\cup} \{\mathbf{a}\}$ | 5 : $\mathcal{Y}_2 \xleftarrow{\cup} \{\mathbf{a},\mathbf{b}\}$ |
| 6 : **if** $\mathbf{b} \in \mathcal{X}\backslash\mathcal{Y}_2$ **then** $\mathcal{Z} \xleftarrow{\cup} \{\mathbf{b}\}$ | |
| 7 : $\mathcal{Y}_2 \xleftarrow{\cup} \{\mathbf{a},\mathbf{b},\mathbf{c}\}$ | |

**Fig. 7.** Subroutines $\mathsf{PopulateSetsEval}, \mathsf{PopulateSetsO_v}$

Reduction to $\mathbb{PG}$. We next show that the $\Pr\left[\mathbb{G}_2 \Rightarrow \mathsf{true}\right]$ is upper bounded in terms of the probability of a $(q,l)$-query adversary winning the game $\mathbb{PG}$.

**Lemma 6.** *There exists a $(q,l)$-query adversary $\mathcal{D}$ against the permutation game $\mathbb{PG}$ such that*

$$\Pr\left[\mathbb{G}_2 \Rightarrow \mathsf{true}\right] \leqslant \Pr\left[\mathbb{PG}(\mathcal{D}) \Rightarrow \mathsf{true}\right] + \frac{2q(2k+3q+2)}{p} + \frac{5q}{p} + \frac{k^2+k+2}{p}\ .$$

*Proof.* We transform $\mathcal{R}_1, \mathcal{R}_2$ playing $\mathbb{G}_2$ to an adversary $\mathcal{D}$ playing the game $\mathbb{PG}$ through a sequence of intermediate games and use the upper bound on the probability of winning the game $\mathbb{PG}$ established previously to prove an upper bound on $\Pr\left[\mathbb{G}_2 \Rightarrow \mathsf{true}\right]$. In order to make the pseudocode for subsequent games compact we define the two subroutines $\mathsf{PopulateSetsEval}, \mathsf{PopulateSetsO_v}$ and invoke them from $\mathsf{Eval}, \mathsf{O_v}$. The subroutines $\mathsf{PopulateSetsEval}, \mathsf{PopulateSetsO_v}$ are formally described in Figure **??**.

The game $\mathbb{G}_4$. We next describe game $\mathbb{G}_4$ where we introduce some additional bookkeeping. In $\mathbb{G}_4$, every valid label that is an input to $\mathcal{R}_1, \mathcal{R}_2$ or queried by $\mathcal{R}_1, \mathcal{R}_2$ or an answer to a query of $\mathcal{R}_1, \mathcal{R}_2$, is mapped to a polynomial in $\mathbb{Z}_p[I_1, \cdots, I_k, V, T_1, \cdots, T_{2q}]$ where $q$ is the total number of $\mathsf{Eval}, \mathsf{O_v}$ queries made by $\mathcal{R}_1, \mathcal{R}_2$. The polynomial associated with label $\mathbf{a}$ is denoted by $\mathsf{p}_{\mathbf{a}}$. Similarly, we define $\Lambda$ to be a mapping from polynomials to labels. For all labels $\mathbf{a} \in \mathcal{L}$, $\Lambda(\mathsf{p}_{\mathbf{a}}) = \mathbf{a}$. The mapping from labels to polynomials is done such that for every label $\mathbf{a}$ mapped to $\mathsf{p}_{\mathbf{a}}$,

$$\sigma^{-1}(\mathbf{a}) = \mathsf{p}_{\mathbf{a}}(i_1, \cdots, i_k, v, t_1, \cdots, t_{2q})\ .$$

For compactness, let us denote $(i_1, \cdots, i_k, v, t_1, \cdots, t_{2q})$ by $\vec{i}$. Before running $\mathcal{R}_1$, $\mathsf{p}_{\sigma(1)}, \mathsf{p}_{\sigma(v)}, \mathsf{p}_{\sigma(i_1)}, \cdots, \mathsf{p}_{\sigma(i_k)}, \mathsf{p}_{\sigma(i_1 \cdot v)}, \cdots, \mathsf{p}_{\sigma(i_k \cdot v)}$ are assigned polynomials $1, V, I_1, \cdots, I_k, I_1V, \cdots, I_kV$ respectively and for all other labels $\mathbf{a} \in \mathcal{L}$, $\mathsf{p}_{\mathbf{a}} = \perp$. The function $\Lambda$ is defined accordingly. For labels $\mathbf{a}$ queried by $\mathcal{R}_1, \mathcal{R}_2$ that have not been previously mapped to any polynomial (i.e. $\mathsf{p}_{\mathbf{a}} = \perp$), $\mathsf{p}_{\mathbf{a}}$ is assigned $T_{\mathsf{new}}$ ($\mathsf{new}$ starting from 1 and being incremented for every such label queried), the variable $t_{\mathsf{new}}$ is assigned the pre-image of the label and $\Lambda(T_{\mathsf{new}})$ is assigned $\mathbf{a}$. Since there are $q$ queries (each with two inputs), there can be at most $2q$ labels

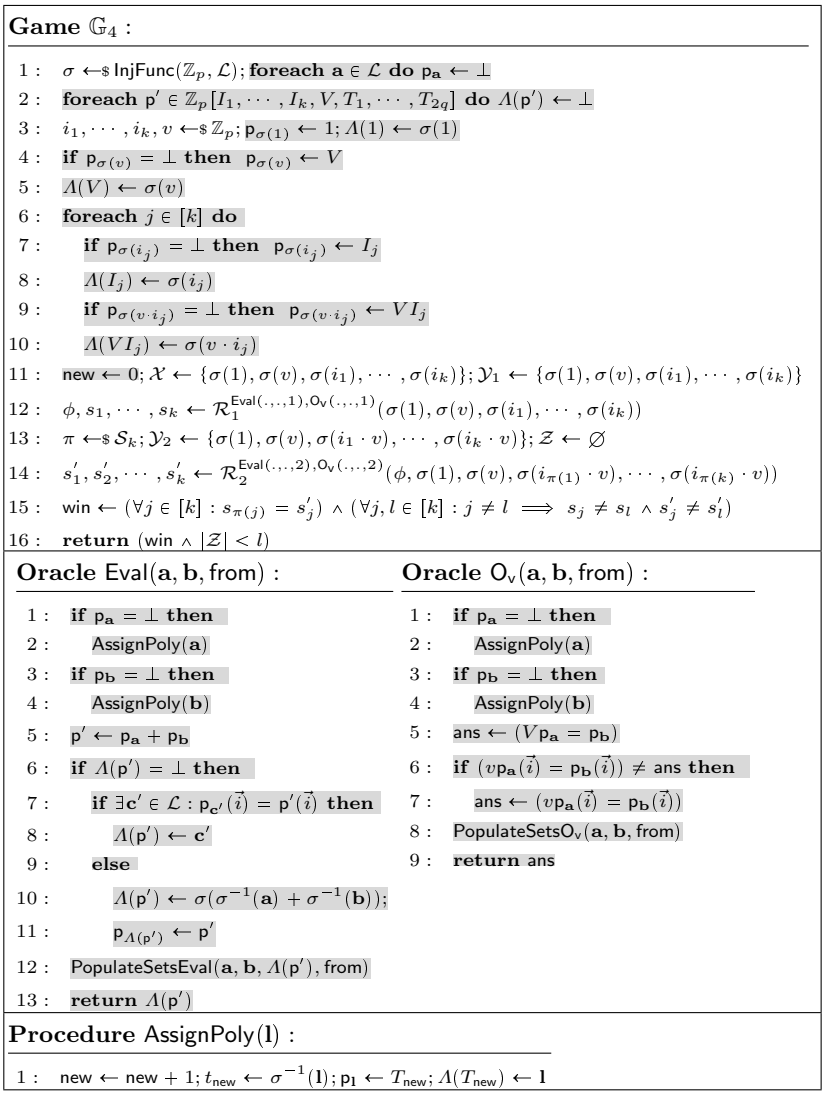


**Fig. 8.** $\mathbb{G}_4$ introduces additional bookkeeping. The newly introduced changes compared to $\mathbb{G}_2$ are highlighted.

that had not previously been mapped to any polynomial. Hence, the polynomials have variables $I_1, \cdots, I_k, V, T_1, \cdots, T_{2q}$.

For an $\mathsf{Eval}(\mathbf{a}, \mathbf{b}, .)$ query where $\mathbf{c} = \sigma(\sigma^{-1}(\mathbf{a}) + \sigma^{-1}(\mathbf{b}))$, let $\mathsf{p}' = \mathsf{p}_{\mathbf{a}} + \mathsf{p}_{\mathbf{b}}$. From the definition of $\mathsf{p}$, we have that $\mathsf{p}'(\vec{i}) = \sigma^{-1}(\mathbf{a}) + \sigma^{-1}(\mathbf{b})$. If $\Lambda(\mathsf{p}') \neq \perp$, then by definition of $\Lambda$, we have $\Lambda(\mathsf{p}') = \mathbf{c}$. If $\Lambda(\mathsf{p}') = \perp$, then exactly one of the following two must be true.

1. The label $\mathbf{c}$ has been mapped to a polynomial which is different from $\mathsf{p}'$. In this case $\mathsf{p}_{\mathbf{c}}(\vec{i}) = \mathsf{p}'(\vec{i})$ and $\Lambda(\mathsf{p}')$ is assigned $\mathbf{c}$.
2. The label $\mathbf{c}$ has not been mapped to any polynomial. In this case, $\mathsf{p}_{\mathbf{c}}$ is assigned $\mathsf{p}'$ and $\Lambda(\mathsf{p}')$ is assigned $\mathbf{c}$.

The label $\Lambda(\mathsf{p}')$ is returned as the answer of the $\mathsf{Eval}$ query. Note that the output of $\mathsf{Eval}$ is $\mathbf{c} = \sigma(\sigma^{-1}(\mathbf{a}) + \sigma^{-1}(\mathbf{b}))$ in all cases, i.e. it is the same as the output of $\mathsf{Eval}$ in $\mathbb{G}_2$.

For an $\mathsf{O_v}(\mathbf{a}, \mathbf{b}, .)$ query, we first assign the boolean value $V\mathsf{p}_{\mathbf{a}} = \mathsf{p}_{\mathbf{b}}$ to $\mathsf{ans}$. Note that if $\mathsf{ans}$ is $\mathsf{true}$, then $v \cdot \sigma^{-1}(\mathbf{a}) = \sigma^{-1}(\mathbf{b})$. However, we might have that $v \cdot \sigma^{-1}(\mathbf{a}) = \sigma^{-1}(\mathbf{b})$ and $V\mathsf{p}_{\mathbf{a}} \neq \mathsf{p}_{\mathbf{b}}$. When this happens, the boolean value $v(\mathsf{p}_{\mathbf{a}}(\vec{i}) = \mathsf{p}_{\mathbf{b}}(\vec{i}))$ is assigned to $\mathsf{ans}$. Oracle $\mathsf{O_v}$ returns $\mathsf{ans}$. From the definition of $\mathsf{p}$, it follows that the value returned by $\mathsf{O_v}$ in $\mathbb{G}_4$ is $(v \cdot \sigma^{-1}(\mathbf{a}) = \sigma^{-1}(\mathbf{b}))$ i.e. it is the same as the output of $\mathsf{O_v}$ in $\mathbb{G}_2$.

Figure **??** formally describes $\mathbb{G}_4$. The changes in $\mathbb{G}_4$ compared to $\mathbb{G}_2$ have been highlighted. We have already pointed out that the outputs of $\mathsf{O_v}, \mathsf{Eval}$ in $\mathbb{G}_4$ are identical to those in $\mathbb{G}_2$. Since the other changes involve only additional bookkeeping, the outputs of $\mathbb{G}_2, \mathbb{G}_4$ are identical. Therefore

$$\Pr\left[\mathbb{G}_4 \Rightarrow \mathsf{true}\right] = \Pr\left[\mathbb{G}_2 \Rightarrow \mathsf{true}\right] . \tag{4}$$

THE GAME $\mathbb{G}_{11}$. We introduce a new game named $\mathbb{G}_{11}$ in Figure **??**. Initially, for all polynomials $\mathsf{p}$, $\Lambda(\mathsf{p}) = \perp$. In this game $\Lambda(1), \Lambda(V)$, $\Lambda(I_j)$'s, and $\Lambda(VI_j)$'s are assigned distinct labels sampled from $\mathcal{L}$. Adversary $\mathcal{R}_1$ is run with input labels $\Lambda(1), \Lambda(V), \Lambda(I_1), \cdots, \Lambda(I_k)$ and $\mathcal{R}_2$ has input labels $\Lambda(1), \Lambda(V), \Lambda(I_{\pi(1)} \cdot V), \cdots, \Lambda(I_{\pi(k)} \cdot V)$. The bookkeeping is identical to that in $\mathbb{G}_4$. Observe from the pseudocode that the mapping $\Lambda$ is injective in this game and hence $\Lambda^{-1}$ is well defined.

For every $\mathsf{Eval}$ or $\mathsf{O_v}$ query, if for the input label $\mathbf{l}$, $\Lambda^{-1}(\mathbf{l})$ is $\perp$, then $\mathbf{l}$ is assigned to $\Lambda(T_{\mathsf{new}})$. For every such input label, $\mathsf{new}$ is incremented. For an $\mathsf{Eval}(\mathbf{a}, \mathbf{b}, .)$ query, if $\Lambda(\Lambda^{-1}(\mathbf{a}) + \Lambda^{-1}(\mathbf{b}))$ is not defined, then it is assigned a random label in $\overline{R(\Lambda)}$. The label $\Lambda(\Lambda^{-1}(\mathbf{a}) + \Lambda^{-1}(\mathbf{b}))$ is returned as answer. For $\mathsf{O_v}(\mathbf{a}, \mathbf{b}, .)$, query $\mathsf{true}$ is returned iff $V\Lambda^{-1}(\mathbf{a})$ and $\Lambda^{-1}(\mathbf{b})$ are the same polynomials.

We next upper bound $\Pr\left[\mathbb{G}_4 \Rightarrow \mathsf{true}\right]$ in terms of $\Pr\left[\mathbb{G}_{11} \Rightarrow \mathsf{true}\right]$ in Lemma **??**.

**Lemma 7.** *For the games $\mathbb{G}_4, \mathbb{G}_{11}$, we have,*

$$\Pr\left[\mathbb{G}_4 \Rightarrow \mathsf{true}\right] \leqslant \Pr\left[\mathbb{G}_{11} \Rightarrow \mathsf{true}\right] + \frac{2q(2k+3q+2)}{p} + \frac{5q}{p} + \frac{k^2+k+2}{p} .$$

The proof of Lemma **??** has been deferred to the full version.

THE ADVERSARY $\mathcal{D}$. Next, we construct the adversary $\mathcal{D}$ that plays $\mathbb{PG}$ by simulating $\mathbb{G}_{11}$ to $\mathcal{R}_1, \mathcal{R}_2$, where the permutation $\pi$ is the secret permutation from $\mathbb{PG}$. As we will discuss below, the core of the adversary $\mathcal{D}$ will boil down

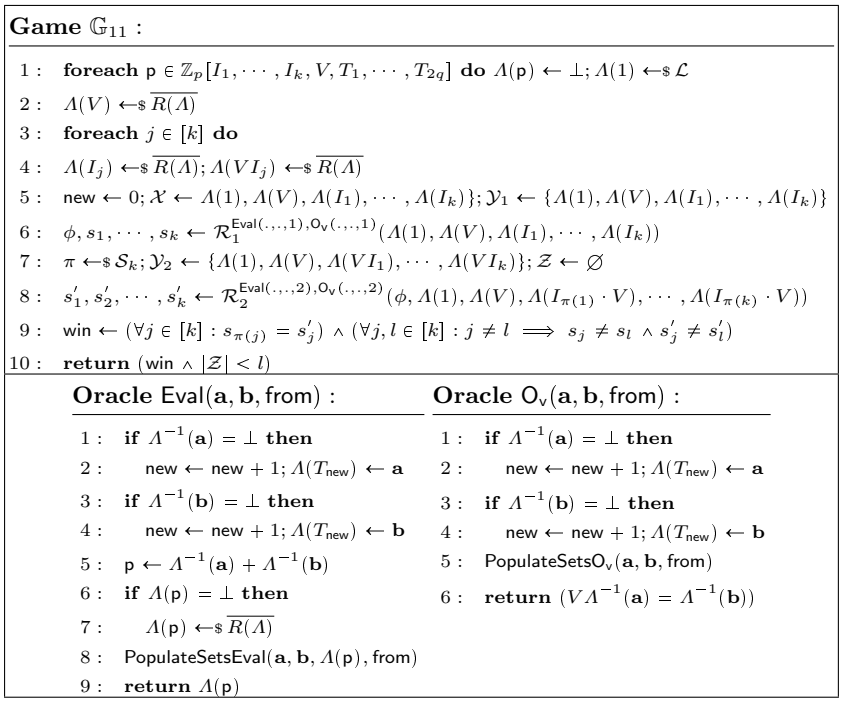


**Fig. 9.** Game $\mathbb{G}_{11}$

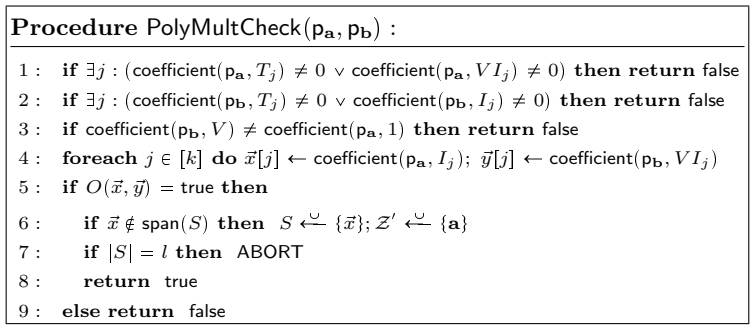


**Fig. 10. Subroutine** $\mathsf{PolyMultCheck}$ **for simulating** $\mathsf{O_v}$**.** In particular, $\mathsf{coefficient}(\mathsf{p}, M)$ returns the coefficient of the monomial $M$ in the polynomial $\mathsf{p}$. The sets $S$ and $\mathcal{Z}'$ have no effect on the behavior, and are only used in the analysis of $\mathcal{D}$. The symbol ABORT indicates that $\mathcal{D}$ aborts and outputs $\perp$.

to properly simulating the $\mathsf{O_v}$ oracle using the $O$ oracle from $\mathbb{PG}$ *and* simulating the labels $\sigma(i_{\pi(j)})$ (and the associated polynomials) correctly without knowing $\pi$. After a correct simulation, $\mathcal{D}$ will simply extract the permutation $\pi$.

To see how this can be done, let us first have a closer look at $\mathbb{G}_{11}$. Let us introduce the shorthand $K_j = VI_{\pi(j)}$ for $j \in [k]$. With this notation, every

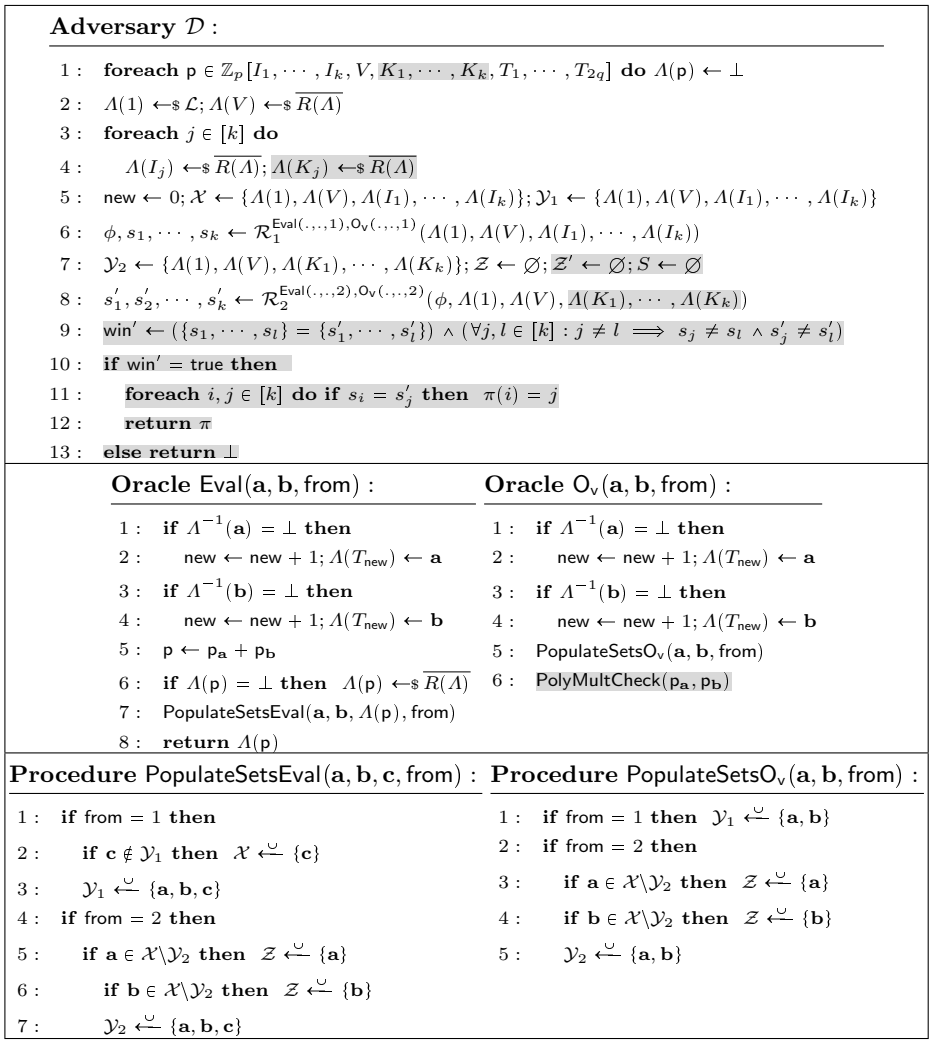


**Fig. 11.** Adversary $\mathcal{D}$ which plays the permutation game $\mathbb{PG}$. The changes in $\mathcal{D}$ compared to $\mathbb{G}_{11}$ have been highlighted.

polynomial input to or output from Eval is a linear combination of the monomials $1, I_1, \ldots, I_k, V, K_1, \ldots, K_k, T_1, T_2, \ldots$ . Now, it is convenient to slightly rethink the check of whether $V\mathsf{p_a} = \mathsf{p_b}$ within $\mathsf{O_v}$ with this notation. First off, we observe that if either of the polynomial contains a monomial of the form $T_i$, the check fails. In fact, it is immediately clear that the check can only possibly succeed is if $\mathsf{p_a}$ is a linear combination of 1 and the $I_j$'s and $\mathsf{p_b}$ is a linear combination of

$V$ and the $K_j$'s. Now, assume that

$$\mathsf{p_a}(I_1, \ldots, I_k) = a_0 + \sum_{j=1}^{k} \vec{x}[j] \cdot I_j \ ,$$

$$\mathsf{p_b}(V, K_1, \ldots K_k) = b_0 \cdot V + \sum_{j=1}^{k} \vec{y}[j] \cdot K_j \ .$$

Then, $V \cdot \mathsf{p}_a = \mathsf{p}_b$ if and only if $a_0 = b_0$ *and* $\vec{y}[j] = \vec{x}[\pi(j)]$ for all $j \in [k]$. If we are now in Game $\mathbb{PG}$, and $\pi$ is the chosen permutation, then this is equivalent to $O(\vec{x}, \vec{y}) = \mathsf{true}$ and $a_0 = b_0$.

This leads naturally to the adversary $\mathcal{D}$, which we formally describe in Figure **??**. The adversary will simply sample labels $\mathbf{f}_1, \ldots, \mathbf{f}_k$ for $\sigma(v \cdot i_{\pi(1)}), \ldots, \sigma(v \cdot i_{\pi(k)})$, and associate with them polynomials in the variables $K_1, \ldots, K_j$. Other than that, it simulates the game $\mathbb{G}_{11}$, with the exception that the check $V \cdot \mathsf{p_a} = \mathsf{p_b}$ is not implemented using the above approach – summarized in Figure **??**. Note that $\mathcal{D}$ aborts when $|S| = l$ and makes at most $q$ queries to $O$. Thus $\mathcal{D}$ is a -query adversary against $\mathbb{PG}$. If $\mathcal{D}$ does not abort, then its simulation of $\mathbb{G}_{11}$ is perfect. If $\mathbb{G}_{11} \Rightarrow \mathsf{true}$ and $\mathcal{D}$ does not abort, then $\mathsf{win}'$ shall be $\mathsf{true}$ and $\mathcal{D}$ will output the correct $\pi$.

The rest of the proof will now require proving that whenever $\mathbb{G}_{11}$ outputs $\mathsf{true}$ our adversary $\mathcal{D}$ will never abort due to the check $|S| = l$. Since $\mathbb{G}_{11} \Rightarrow \mathsf{true}$ only if $|\mathcal{Z}| < l$, the following lemma implies that $\mathcal{D}$ does not abort if $\mathbb{G}_{11} \Rightarrow \mathsf{true}$.

**Lemma 8.** *Let $(\vec{x}_1, \vec{y}_1), \cdots, (\vec{x}_u, \vec{y}_u)$ be the queries made by $\mathcal{D}$ to $O$ which return* $\mathsf{true}$*. Then,*

$$\mathsf{rank}(\vec{x}_1, \cdots, \vec{x}_u) \leqslant |\mathcal{Z}| \ .$$

The proof of Lemma **??** has been deferred to the full version.

We have established that if $\mathbb{G}_{11}$ outputs $\mathsf{true}$, then $\mathcal{D}$ will not abort and hence $\mathcal{D}$ simulates $\mathbb{G}_{11}$ to $\mathcal{R}_1, \mathcal{R}_2$ perfectly. If $\mathsf{win} = \mathsf{true}$ in $\mathbb{G}_{11}$, the checks by $\mathcal{D}$ succeed and $\mathcal{D}$ outputs the correct permutation and wins $\mathbb{PG}$. Therefore, $\mathcal{D}$ is a $(q, l)$-query adversary such that $\mathbb{PG}(\mathcal{D}) \Rightarrow \mathsf{true}$ if $\mathbb{G}_{11} \Rightarrow \mathsf{true}$. Hence,

$$\Pr\left[\mathbb{G}_{11} \Rightarrow \mathsf{true}\right] \leqslant \Pr\left[\mathbb{PG}(\mathcal{D}) \Rightarrow \mathsf{true}\right] \ . \tag{5}$$

Combining Lemma **??** and **??**,**??** we get,

$$\Pr\left[\mathbb{G}_2 \Rightarrow \mathsf{true}\right] \leqslant \Pr\left[\mathbb{PG}(\mathcal{D}) \Rightarrow \mathsf{true}\right] + \frac{2q(2k+3q+2)}{p} + \frac{5q}{p} + \frac{k^2+k+2}{p} \ . \tag{6}$$

□

Combining **????**, we get,

$$\Pr\left[\mathbb{G}_2 \Rightarrow \mathsf{true}\right] \leqslant \frac{q^l}{k!} + \frac{2q(2k+3q+2)}{p} + \frac{5q}{p} + \frac{k^2+k+2}{p} \ .$$

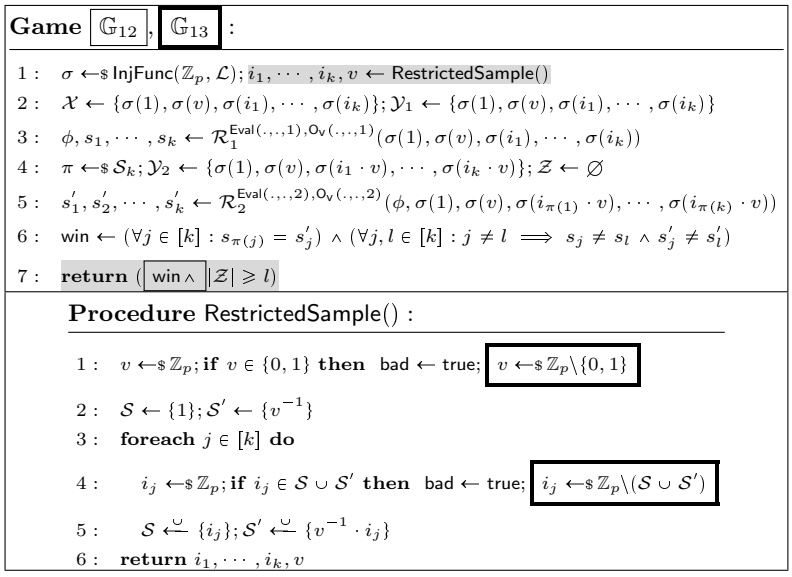


**Fig. 12.** Games $\mathbb{G}_{12}, \mathbb{G}_{13}$. The $\mathsf{Eval}, \mathsf{O_v}$ oracles in $\mathbb{G}_{12}, \mathbb{G}_{13}$ are identical to those in $\mathbb{G}_3$ and hence we do not rewrite it here. The statement within the thinner box is present only in $\mathbb{G}_{12}$ and the statement within the thicker box is present only in $\mathbb{G}_{13}$. The newly introduced changes compared to $\mathbb{G}_3$ are highlighted.

### 4.4 Memory Lower Bound when $|\mathcal{Z}| \geqslant l$ (Proof of Lemma ??)

Recall that we need to prove the following lemma, which we do by using a compression argument.

**Lemma ??.** *If the size of the state $\phi$ output by $\mathcal{R}_1$ is $s$ bits and $(\mathcal{R}_1, \mathcal{R}_2)$ make $q$ queries in total in $\mathbb{G}_3$, then*

$$\Pr\left[\mathbb{G}_3 \Rightarrow \mathsf{true}\right] \leqslant 2 \cdot 2^{\frac{s}{2}} \left(\frac{8q^2(2k+2+3q)}{p}\right)^{\frac{l}{2}} \left(1 + \frac{6q}{p}\right)^{\frac{2q-l}{2}} + \frac{k^2+k+2}{p} .$$

*Proof.* Our proof does initial game hopping, with easy transitions. It first introduces a new game, $\mathbb{G}_{12}$ whose minor difference from game $\mathbb{G}_3$ is that it samples $i_1, \cdots, i_k, v$ using $\mathsf{RestrictedSample}$ which was previously used in game $\mathbb{G}_{11}$. It adds a $\mathsf{bad}$ flag while sampling $i_1, \cdots, i_k, v$ which is set to $\mathsf{true}$ if $v$ is in $\{0,1\}$ or if $|1, v, i_1, \cdots, i_k, i_1 \cdot v, \cdots, i_k \cdot v| < 2k+2$. The $\mathsf{bad}$ event does not affect the output of $\mathbb{G}_{12}$ in any way. Observe that even though the sampling of $i_1, \cdots, i_k, v$ is written in a different manner in $\mathbb{G}_{12}$, it is identical to that in $\mathbb{G}_3$. In all other respects these two games are identical.

$$\Pr\left[\mathbb{G}_3 \Rightarrow \mathsf{true}\right] = \Pr\left[\mathbb{G}_{12} \Rightarrow \mathsf{true}\right] . \tag{7}$$

Games $\mathbb{G}_{12}, \mathbb{G}_{13}$ differ in the procedure $\mathsf{RestrictedSample}$ and the condition to return $\mathsf{true}$. Note that the conditions of $\mathsf{bad}$ being set to $\mathsf{true}$ is identical in

$\mathbb{G}_{12}, \mathbb{G}_{13}$ and given that $\mathsf{bad}$ is not set to $\mathsf{true}$, $\mathbb{G}_{13}$ returns $\mathsf{true}$ whenever $\mathbb{G}_{12}$ returns $\mathsf{true}$. Therefore,

$$\Pr\left[\mathbb{G}_{12} \Rightarrow \mathsf{true}\right] \leqslant \Pr\left[\mathbb{G}_{13} \Rightarrow \mathsf{true}\right] + \Pr\left[\mathsf{bad} = \mathsf{true} \text{ in } \mathbb{G}_{13}\right] .$$

It is not hard to show (details in the full version) that the probability of $\mathsf{bad}$ being set to $\mathsf{true}$ in $\mathsf{RestrictedSample}$ is at most $\frac{k^2+k+2}{p}$. Since in $\mathbb{G}_{13}$ $\mathsf{bad}$ is set only in $\mathsf{RestrictedSample}$, the probability of $\mathsf{bad}$ being set to true is the same. Hence, we get,

$$\Pr\left[\mathbb{G}_{12} \Rightarrow \mathsf{true}\right] \leqslant \Pr\left[\mathbb{G}_{13} \Rightarrow \mathsf{true}\right] + \frac{k^2+k+2}{p} . \tag{8}$$

The compression argument. We assume $\Pr\left[\mathbb{G}_{13} \Rightarrow \mathsf{true}\right] = 2\epsilon$. We say a $\sigma$ is "good" in $\mathbb{G}_{13}$ if

$$\Pr\left[\mathbb{G}_{13} \Rightarrow \mathsf{true} \,\middle|\, \sigma \text{ was sampled in } \mathbb{G}_{13}\right] \geqslant \epsilon .$$

It follows from Markov's inequality that at least $\epsilon$ fraction of $\sigma$'s are "good".

The following lemma captures the essence of our compression argument.

**Lemma 9.** *If the state output by $\mathcal{R}_1$ has size $s$ bits, all the "good" $\sigma$'s can be encoded in an encoding space of size at most*

$$2^s p! \left(1 + \frac{6q}{p}\right)^{(2q-l)} \left(\frac{p}{8q^2(2k+2+3q)}\right)^{-l} ,$$

*and decoded correctly with probability $\epsilon$.*

We next give some intuition regarding how we achieve compression and defer the formal proof of Lemma **??** to the full version.

Intuition regarding compression. Observe in $\mathbb{G}_{13}$, the labels in $\mathcal{Z}$ were queried by $\mathcal{R}_2$ (these labels were not seen by $\mathcal{R}_2$ before they were queried) and were answers to $\mathcal{R}_1$ and were not seen by $\mathcal{R}_1$ before the query. The core idea is that for all $\mathbf{a} \in \mathcal{L} \backslash \mathcal{Z}$, we store exactly one of $\mathbf{a}$ or its pre-image in the encoding and for all labels in $\mathcal{Z}$, we store neither the label nor its pre-image. Since $\mathcal{R}_2$ queries all the labels in $\mathcal{Z}$, these labels can be found by running $\mathcal{R}_2$ while decoding. Since all the labels in $\mathcal{Z}$ are answers to queries of $\mathcal{R}_1$ and were not seen by $\mathcal{R}_1$ before the query, their pre-images can be figured out while running $\mathcal{R}_1$.

High level outlines of $\mathsf{Encode}$, $\mathsf{Decode}$. In $\mathsf{Encode}$, we simulate the steps of $\mathbb{G}_{13}$ to $\mathcal{R}_1, \mathcal{R}_2$, including bookkeeping and then run $\mathcal{R}_1$ again assuming the particular $\sigma$ we are compressing is sampled in $\mathbb{G}_{13}$. In $\mathsf{Decode}$, we run $\mathcal{R}_2$ and then $\mathcal{R}_1$ to recover $\sigma$. We treat the values $i_1, \cdots, i_k, v, \pi$ as part of the common randomness provided to $\mathsf{Encode}, \mathsf{Decode}$ (we assume they are sampled from the same distribution they are sampled from in $\mathbb{G}_{13}$). The random tapes of $\mathcal{R}_1, \mathcal{R}_2$ can also be derived from the common randomness of $\mathsf{Encode}, \mathsf{Decode}$. For simplicity, we do not specify this explicitly in the algorithms and treat $\mathcal{R}_1, \mathcal{R}_2$ as deterministic.

RUNNING $\mathcal{R}_2$. First off, we assume that $\mathcal{R}_1$ queries labels that it has "seen" before and $\mathcal{R}_2$ queries labels that $\mathcal{R}_1$ has "seen" or it has "seen" before. We shall relax this assumption later. Ideally, we would want to just store only $\phi$, the inputs labels to $\mathcal{R}_2$ and the labels that are answers to $\mathcal{R}_2$'s queries. We append the input labels of $\mathcal{R}_2$ and labels that are answers to its $\mathsf{Eval}$ queries that it has not "seen" before to a list named $\mathsf{Labels}$. However, it is easy to see that this information is not enough to answer $\mathsf{O_v}$ queries during decoding, as answering $\mathsf{O_v}$ queries inherently requires knowledge about pre-images of $\mathcal{R}_2$. This naturally leads to the idea of maintaining a mapping of all the labels "seen by" $\mathcal{R}_2$ to their pre-images.

THE MAPPING $\mathsf{T}$ OF LABELS TO PRE-IMAGE EXPRESSIONS. The pre-images of input labels and the labels that were results of sequence of $\mathsf{Eval}$ queries on its input labels by $\mathcal{R}_2$, are known. However, $\mathcal{R}_2$ might query labels which were neither an input to it nor an answer to one of its $\mathsf{Eval}$ queries. Such a label is in $\mathcal{Z}$ since we have assumed that all labels queried by $\mathcal{R}_2$ were "seen by" $\mathcal{R}_1$ or "seen by" $\mathcal{R}_2$ before. We represent the pre-images of labels in $\mathcal{Z}$ using a placeholder variable $X_n$ where $n$ is incremented for every such label. Note that the pre-image of every label seen by $\mathcal{R}_2$ can be expressed as a linear polynomial in the $X_n$'s (these linear polynomials are referred to as pre-image expressions from hereon). Therefore we maintain a mapping of all labels "seen by" and their pre-image expressions in a list of tuples named $\mathsf{T}$. Our approach is inspired by a similar technique used by Corrigan-Gibbs and Kogan in [**?**]. Like in [**?**], we *stress* that the mapping $\mathsf{T}$ is not a part of the encoding.

For $\mathsf{Eval}$ queries, we can check if there is a tuple in $\mathsf{T}$ whose pre-image expression is the sum of the pre-image expressions of the input labels. If that is the case, we return the label of such a tuple. Otherwise, we append the answer label to $\mathsf{Labels}$. For $\mathsf{O_v}$ queries, we can return $\mathsf{true}$ if the pre-image expression of the first input label multiplied by $v$ gives the pre-image expression of the second input label. Otherwise we return $\mathsf{false}$.

SURPRISES. There is a caveat, however. There might arise a situation that the label which is the answer to the $\mathsf{Eval}$ query is present in $\mathsf{T}$ but its pre-image expression is not the sum of the pre-image expressions of the input labels. We call such a situation a "surprise" and we call the answer label in that case a "surprise label". For $\mathsf{O_v}$ queries, there might be a surprise when the answer of the $\mathsf{O_v}$ query is $\mathsf{true}$ but the pre-image expression of the first input label multiplied by $v$ is different pre-image expression of the second input label. In this case we call the second input label the surprise label. We assign a sequence number to each query made by $\mathcal{R}_2$, starting from 1 and an index to each tuple in $\mathsf{T}$, with the indices being assigned to tuples in the order they were appended to $\mathsf{T}$. To detect the query where the surprise happens, we maintain a set named $\mathsf{Srps}_1$ that contains tuples of query sequence numbers and indices of the surprise label in $\mathsf{T}$. This set $\mathsf{Srps}_1$ is a part of the encoding. Note that whenever there is a surprise, it means that two different pre-image expressions evaluate to the same value. Since these two pre-image expressions are linear polynomials, at least one variable can be eliminated from $\mathsf{T}$ by equating the two pre-image expressions.

RUNNING $\mathcal{R}_1$. Now that we have enough information in the encoding to run $\mathcal{R}_2$, we consider the information we need to add to the encoding to run $\mathcal{R}_1$ after $\mathcal{R}_2$ is run. First, we need to provide $\mathcal{R}_1$ its input labels. Our initial attempt would be to append the input labels of $\mathcal{R}_1$ (except $\sigma(1), \sigma(v)$, which are already present) to $\mathsf{Labels}$. However, some of these input labels to $\mathcal{R}_1$ might have already been "seen by" $\mathcal{R}_2$. Since all labels "seen by" $\mathcal{R}_2$ are in $\mathsf{T}$, we need a way to figure out which of $\sigma(i_j)$'s are in $\mathsf{T}$. Note that such a label was either queried by $\mathcal{R}_2$ or an answer to a query of $\mathcal{R}_2$ (cannot have been an input to $\mathcal{R}_2$ given the restrictions on $i_1, \cdots, i_k, v$). Suppose $q$ was the sequence number of the query in which $\sigma(i_j)$ was queried or an answer. The tuple $(q, b, j)$ is added to the set $\mathsf{Inputs}$ where $b$ can take values $\{1, 2, 3\}$ depending on whether $\sigma(i_j)$ was the first input label, the second input label or the answer label respectively. This set $\mathsf{Inputs}$ is a part of the encoding. The rest of the labels $\sigma(i_j)$, which do not appear in $\mathsf{T}$, are added to $\mathsf{T}$ with their pre-images and the labels are appended to $\mathsf{Labels}$. Note that for all queries of $\mathcal{R}_1$, it follows from our assumption that the input labels will be in $\mathsf{T}$. For every surprise, we add a tuple of sequence number and an index in $\mathsf{T}$ to the set $\mathsf{Srps}_2$.

RELAXING THE ASSUMPTION. When we allow $\mathcal{R}_2$ to query labels it has not seen before or $\mathcal{R}_1$ has not seen, there are two issues. First, we need to add a tuple for the label in $\mathsf{T}$ (since $\mathsf{T}$, by definition contains a tuple for all labels queried by $\mathcal{R}_2$). We solve this issue by adding the tuple made of the label and its pre-image. We have no hope of recovering the pre-image later, hence, we append the pre-image to a list named $\mathsf{Vals}$. This list needs to be a part of the encoding since the pre-image of the label needs to be figured out to be added to $\mathsf{T}$ during decoding. For queries of $\mathcal{R}_1$, if the input label is not present in $\mathsf{T}$, we do the same thing. The second issue that comes up when we relax the assumption is that we need to distinguish whether an input label was in $\mathcal{Z}$ or not. We solve this issue by maintaining a set of tuples named $\mathsf{Free}$. For all labels in $\mathcal{Z}$ that are not an input label to $\mathcal{R}_1$, we add the tuple consisting of the sequence number of the query of $\mathcal{R}_2$ and $b$ to $\mathsf{Free}$ where $b$ set to 1 indicates it was the first input label and $b$ set to 2 indicates it was the second input label.

THE FINAL STEPS. The labels the are absent in $\mathsf{T}$ are appended to a list named $\mathsf{RLabels}$. If $|\mathcal{Z}| < l$, a fixed encoding $\mathsf{D}$ (the output of $\mathsf{Encode}$ for some fixed $\sigma$ when $|\mathcal{Z}| \geqslant l$) is returned. Otherwise the encoding of $\sigma$ consisting of $\mathsf{Labels}$, $\mathsf{RLabels}$, $\mathsf{Vals}$, $\mathsf{Inputs}$, $\mathsf{Srps}_1$, $\mathsf{Srps}_2$, $\mathsf{Free}$, $\phi$ is returned.

WRAPPING UP. The set of all "good" $\sigma$'s has size at least $\epsilon p!$ (where we have used that the total number of injective functions from $\mathbb{Z}_p \to \mathcal{L}$ is $p!$). Using $\mathcal{X}$ to be the set of the "good" $\sigma$'s, $\mathcal{Y}$ to be the set of encodings, $\mathcal{R}$ to be the set of cartesian product of the domains of $i_1, \cdots, i_k, v, \pi$, the set of all random tapes of $\mathcal{R}_1$ the set of all random tapes of $\mathcal{R}_2$ and $\mathcal{L}$, it follows from Lemma **??** and

Proposition **??** that

$$\log\left(\Pr\left[\text{Decoding is correct}\right]\right) \leqslant s + (2q-l)\log\left(1+\frac{6q}{p}\right) - l\log\left(\frac{p}{8q^2(2k+2+3q)}\right) - \log\epsilon\ .$$

We have from Lemma **??** that $\Pr\left[\text{Decoding is correct}\right] \leqslant \epsilon$. Therefore,

$$2\log\epsilon \leqslant s + (2q-l)\log\left(1+\frac{6q}{p}\right) - l\log\left(\frac{p}{8q^2(2k+2+3q)}\right)\ .$$

Since $\Pr\left[\mathbb{G}_{13}\right] = 2\epsilon$, using **????** we have,

$$\Pr\left[\mathbb{G}_3 \Rightarrow \mathsf{true}\right] \leqslant 2\cdot 2^{\frac{s}{2}}\left(\frac{8q^2(2k+2+3q)}{p}\right)^{\frac{l}{2}}\left(1+\frac{6q}{p}\right)^{\frac{2q-l}{2}} + \frac{k^2+k+2}{p}\ .$$

□

## 5 Conclusions

Despite a clear restriction of our result to straightline reductions, we believe the main contribution of this work is the introduction of novel techniques for proving lower bounds on the memory of reductions that will find wider applicability. In particular, we clearly departed from the framework of prior works [**?**,**?**] tailored at the usage of lower bounds for streaming algorithms, and provided the first lower bound for "algebraic" proofs in the public-key domain. The idea of a problem-specific proof of memory could be helpful elsewhere.

Of course, there are several open problems. It seems very hard to study the role of rewinding for such reductions. In particular, the natural approach is to resort to techniques from communication complexity (and their incarnation as streaming lower bounds), as they are amenable to the multi-pass case. The simple combinatorial nature of these lower bounds however is at odds with the heavily structured oracles we encounter in the generic group model. Another problem we failed to solve is to give an adversary $\mathcal{A}$ in our proof which uses little memory – we discuss a candidate in the body, but analyzing it seems to give us difficulties similar to those of rewinding.

This latter point makes a clear distinction, not discussed by prior works, between the *way* in which we prove memory-tightness (via reductions using small memory), and its most general interpretation, as defined in [**?**], which would allow the reduction to adapt its memory usage to that of $\mathcal{A}$.

## Acknowledgements

We thank the anonymous reviewers of EUROCRYPT 2020 for helpful comments. This work was partially supported by NSF grants CNS-1553758 (CAREER), CNS-1719146, and by a Sloan Research Fellowship.

## References


1. Michel Abdalla, Mihir Bellare, and Phillip Rogaway. The oracle Diffie-Hellman assumptions and an analysis of DHIES. In David Naccache, editor, *CT-RSA 2001*, volume 2020 of *LNCS*, pages 143–158. Springer, Heidelberg, April 2001.
2. Benedikt Auerbach, David Cash, Manuel Fersch, and Eike Kiltz. Memory-tight reductions. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part I*, volume 10401 of *LNCS*, pages 101–132. Springer, Heidelberg, August 2017.
3. Mihir Bellare and Phillip Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In Serge Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 409–426. Springer, Heidelberg, May / June 2006.
4. Rishiraj Bhattacharyya. Memory-tight reductions for practical key encapsulation mechanisms. In *PKC 2020*.
5. 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.
6. Ronald Cramer and Victor Shoup. Design and analysis of practical public-key encryption schemes secure against adaptive chosen ciphertext attack. *SIAM Journal on Computing*, 33(1):167–226, 2003.
7. 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 *LNCS*, pages 649–665. Springer, Heidelberg, August 2010.
8. Georg Fuchsbauer, Eike Kiltz, and Julian Loss. The algebraic group model and its applications. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part II*, volume 10992 of *LNCS*, pages 33–62. Springer, Heidelberg, August 2018.
9. Ueli M. Maurer. Abstract models of computation in cryptography (invited paper). In Nigel P. Smart, editor, *10th IMA International Conference on Cryptography and Coding*, volume 3796 of *LNCS*, pages 1–12. Springer, Heidelberg, December 2005.
10. Omer Reingold, Luca Trevisan, and Salil P. Vadhan. Notions of reducibility between cryptographic primitives. In Moni Naor, editor, *TCC 2004*, volume 2951 of *LNCS*, pages 1–20. Springer, Heidelberg, February 2004.
11. 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.
12. Victor Shoup. A proposal for an ISO standard for public key encryption. Cryptology ePrint Archive, Report 2001/112, 2001. http://eprint.iacr.org/2001/112.
13. Yuyu Wang, Takahiro Matsuda, Goichiro Hanaoka, and Keisuke Tanaka. Memory lower bounds of reductions revisited. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part I*, volume 10820 of *LNCS*, pages 61–90. Springer, Heidelberg, April / May 2018.