Permuted Puzzles and Cryptographic Hardness

Elette Boyle
IDC Herzliya

Justin Holmgren
Princeton University

Mor Weiss
IDC Herzliya

Abstract

A *permuted puzzle* problem is defined by a pair of distributions $\mathcal{D}_0, \mathcal{D}_1$ over Σ^n. The problem is to distinguish samples from $\mathcal{D}_0, \mathcal{D}_1$, where the symbols of each sample are *permuted* by a single secret permutation π of $[n]$.

The conjectured hardness of specific instances of permuted puzzle problems was recently used to obtain the first candidate constructions of Doubly Efficient Private Information Retrieval (DE-PIR) (Boyle et al. & Canetti et al., TCC'17). Roughly, in these works the distributions $\mathcal{D}_0, \mathcal{D}_1$ over $\mathbb{F}^n$ are evaluations of either a moderately low-degree polynomial or a random function. This new conjecture seems to be quite powerful, and is the foundation for the first DE-PIR candidates, almost two decades after the question was first posed by Beimel et al. (CRYPTO'00). While permuted puzzles are a natural and general class of problems, their hardness is still poorly understood.

We initiate a formal investigation of the cryptographic hardness of permuted puzzle problems. Our contributions lie in three main directions:

- **Rigorous formalization.** We formalize a notion of permuted puzzle distinguishing problems, extending and generalizing the proposed permuted puzzle framework of Boyle et al. (TCC'17).
- **Identifying hard permuted puzzles.** We identify natural examples in which a one-time permutation *provably* creates cryptographic hardness, based on "standard" assumptions. In these examples, the original distributions $\mathcal{D}_0, \mathcal{D}_1$ are easily distinguishable, but the permuted puzzle distinguishing problem is computationally hard. We provide such constructions in the random oracle model, and in the plain model under the Decisional Diffie-Hellman (DDH) assumption. We additionally observe that the Learning Parity with Noise (LPN) assumption itself can be cast as a permuted puzzle.
- **Partial lower bound for the DE-PIR problem.** We make progress towards better understanding the permuted puzzles underlying the DE-PIR constructions, by showing that a toy version of the problem, introduced by Boyle et al. (TCC'17), withstands a rich class of attacks, namely those that distinguish solely via statistical queries.

Contents

1 Introduction

Computational hardness assumptions are the foundation of modern cryptography. The approach of building cryptographic systems whose security follows from well-defined computational assumptions has enabled us to obtain fantastical primitives and functionality, pushing far beyond the limitations of information theoretic security. But, in turn, the resulting systems are only as secure as the computational assumptions lying beneath them. As cryptographic constructions increasingly evolve toward usable systems, gaining a deeper understanding of the true hardness of these problems—and the relationship between assumptions—is an important task.

To date, a relatively select cluster of structured problems have withstood the test of time (and intense scrutiny), to the point that assuming their hardness is now broadly accepted as "standard." These problems include flavors of factoring [RSA78, Rab79] and computing discrete logarithms [DH76], as well as certain computational tasks in high-dimensional lattices and learning theory [GKL88, BFKL93, Ajt96, BKW00, Ale03, Reg05]. A central goal in the foundational study of cryptography is constructing cryptographic schemes whose security provably follows from these (or weaker) assumptions.

In some cases, however, it may be beneficial – even necessary – to introduce and study *new* assumptions (indeed, every assumption that is "standard" today was at some point freshly conceived). There are several important cryptographic primitives (notable examples include indistinguishability obfuscation (IO) [BGI+01, GGH+13] and SNARKs [BCC+17]) that we do not currently know how to construct based on standard assumptions. Past experience has shown that acheiving new functionalities from novel assumptions, especially *falsifiable* assumptions [Nao03, GW11, GK16], can be a stepping stone towards attaining the same functionality from standard assumptions. This was the case for fully homomorphic encryption [RAD78, Gen09, BV11], as well as many recent primitives that were first built from IO and later (following a long line of works) based on more conservative assumptions (notably, non-interactive zero-knowledge protocols for NP based on LWE [KRR17, CCRR18, HL18, CCH+19, PS19], and the cryptographic hardness of finding a Nash equilibrium based on the security of the Fiat-Shamir heuristic [BPR15, HY17, CHK+19]). Finally, cryptographic primitives that can be based on diverse assumptions are less likely to "go extinct" in the event of a devastating new algorithmic discovery.

Of course, new assumptions should be introduced with care. We should strive to extract some intuitive reasoning justifying them, and some evidence for their hardness. A natural approach is to analyze the connection between the new assumption and known (standard) assumptions, with the ultimate goal of showing that the new assumption is, in fact, implied by a standard assumption. However, coming up with such a reduction usually requires deep understanding of the new assumption, which can only be obtained through a systematic study of it.

DE-PIR and permuted polynomials. A recent example is the new computational assumption underlying the construction of *Doubly Efficient Private Information Retrieval* (DE-PIR) [BIPW17, CHR17], related to pseudorandomness of permuted low-degree curves.

Private Information Retrieval (PIR) [CGKS95, KO97] schemes are protocols that enable a client to access entries of a database stored on a remote server (or multiple servers), while hiding from the server(s) which items are retrieved. If no preprocessing of the database takes place, the security guarantee inherently requires the server-side computation to be linear in the size of the database for each incoming query [BIM00]. Database preprocessing was shown to yield computational savings in the multi-server setting [BIM00], but the goal of single-server PIR protocols with sublinear-time computation was a longstanding open question, with no negative results or (even heuristic) candidate solutions. Such a primitive is sometimes referred to as *Doubly Efficient (DE)* PIR.[1]

Recently, two independent works [BIPW17, CHR17] provided the first candidate constructions of single-server DE-PIR schemes, based on a new conjecture regarding the hardness of distinguishing *permuted* local-decoding queries (for a Reed-Muller code [Ree54, Mul54] with suitable parameters) from a uniformly random set of points. Specifically, although given the queries $\{z_1, \ldots, z_k\} \subseteq [N]$ of the local decoder it is possible to

[1] Namely, computationally efficient for both client and server.

guess (with a non-trivial advantage) the index i which is being locally decoded, the conjectures of [BIPW17, CHR17] very roughly assert that adding a secret permutation can computationally hide i. More precisely, if an adversary instead sees (many) samples of sets of *permuted* queries $\{\pi(z_1), \ldots, \pi(z_k)\}$, where $\pi : [N] \to [N]$ is a secret fixed permutation (the same for all samples), then the adversary cannot distinguish these from independent uniformly random size-k subsets of $[N]$.

This new assumption (which we will refer to as PermRM, see Conjecture 1 in Section 6.2) allowed for exciting progress forward in the DE-PIR domain. But what do we really know about its soundness? Although [BIPW17, CHR17] provide some discussion and cryptanalysis of the assumption, our understanding of it is still far from satisfactory.

Permuted puzzles. The PermRM assumption can be cast as a special case in a broader family of hardness assumptions: as observed in [BIPW17], it can be thought of as an example of an instance where a secret random permutation seems to make an (easy) "distinguishing problem" hard, namely the permutation is the only sources of computational hardness. It should be intuitively clear that such permutations may indeed create hardness. For example, while one can easily distinguish a picture of a cat from that of a dog, this task becomes much more challenging when the pixels are permuted. There are also other instances in which random secret permutations were used to introduce hardness (see Section 1.2 below). Therefore, using permutations as a source of cryptographic hardness seems to be a promising direction for research, and raises the following natural question:

> *Under which circumstances can a secret random permutation be a source of cryptographic hardness?*

1.1 Our Results

We initiate a formal investigation of the cryptographic hardness of permuted puzzle problems. More concretely, our contributions can be summarized within the following three directions.

Rigorous formalization. We formalize a notion of *permuted puzzle distinguishing problems*, which extends and generalizes the proposed framework of [BIPW17]. Roughly, a permuted puzzle distinguishing problem is associated with a pair of distributions $\mathcal{D}_0, \mathcal{D}_1$ over strings in Σ^n, together with a random permutation π over $[n]$. The permuted puzzle consists of the distributions $\mathcal{D}_{0,\pi}, \mathcal{D}_{1,\pi}$ which are defined by sampling a string s according to $\mathcal{D}_0, \mathcal{D}_1$ (respectively), and permuting the entries of s according to π. A permuted puzzle is *computationally hard* if no efficient adversary can distinguish between a sample from $\mathcal{D}_{0,\pi}$ or $\mathcal{D}_{1,\pi}$, even given arbitrarily many samples of its choice from either of the distributions. We also briefly explore related hardness notions, showing that a weaker and simpler variant (which is similar to the one considered in [BIPW17]) is implied by our notion of hardness, and that in some useful cases the weaker hardness notion implies our hardness notion. Our motivation for studying the stronger (and perhaps less natural) hardness notion is that the weaker variant is insufficient for the DE-PIR application.

Identifying Hard Permuted Puzzles. We identify natural examples in which a one-time permutation *provably* introduces cryptographic hardness, based on standard assumptions. In these examples, the distributions $\mathcal{D}_0, \mathcal{D}_1$ are efficiently distinguishable, but the permuted puzzle distinguishing problem is computationally hard. We provide such constructions in the random oracle model, and in the plain model under the Decisional Diffie-Hellman (DDH) assumption [DH76]. We additionally observe that the Learning Parity with Noise (LPN) assumption [BKW00, Ale03] itself can be cast as a permuted puzzle. This is formalized in the following theorem (see Proposition 4.3, Proposition 5.12, and Proposition 5.5 for the formal statements).

Informal Theorem 1.1 (Hard Permuted Puzzles)**.** *There exists a computationally-hard permuted puzzle distinguishing problem:*

- *In the random oracle model.*

- *If the DDH assumption holds.*
- *If the LPN assumption holds.*

Statistical Query Lower Bound for DE-PIR Toy Problem. We make progress towards better understanding the PermRM assumption underlying the DE-PIR constructions of [BIPW17, CHR17]. Specifically, we show that a toy version of the problem, which was introduced in [BIPW17], provably withstands a rich class of learning algorithms known as *Statistical Query (SQ) algorithms*.

Roughly, the toy problem is to distinguish randomly permuted graphs of random univariate polynomials of relatively low degree from randomly permuted graphs of random functions. More formally, for a function $f : X \to Y$, we define its 2-dimensional graph $\mathsf{Graph}(f) : X \times Y \to \{0,1\}$ where $\mathsf{Graph}(f)(x,y) = 1 \Leftrightarrow y = f(x)$. For a security parameter λ and a field $\mathbb{F}$, the distributions $\mathcal{D}_0, \mathcal{D}_1$ in the toy problem are over $\{0,1\}^n$ for $n = |\mathbb{F}|^2$, and output a sample $\mathsf{Graph}(\gamma)$ where $\gamma : \mathbb{F} \to \mathbb{F}$ is a uniformly random degree-λ polynomial in $\mathcal{D}_0$, and a uniformly random function in $\mathcal{D}_1$.

We analyze the security of the toy problem against SQ learning algorithms. Our motivation for focusing on learning algorithms in general is that permuted puzzles are a special example of a learning task. Indeed, the adversary's goal is to classify a challenge sample, given many labeled samples. Thus, it is natural to explore approaches from learning theory as potential solvers for (equivalently, attacks on) the permuted puzzle. Roughly speaking, most known learning algorithms can be categorized within two broad categories. The first category leverages linearity, by identifying correlations with subspaces and using algorithms based on Gaussian elimination to identify these. The second category, which is our focus in this work, is SQ algorithms. Informally, an SQ algorithm obtains no labeled samples. Instead, it can make *statistical queries* that are defined by a boolean-valued function f, and the algorithm then obtains the outcome of applying f to a random sample. A statistical query algorithm is an SQ algorithm that makes polynomially many such queries. We show that the toy problem is hard for SQ algorithms (see Theorem 6.3):

Informal Theorem 1.2. *The BIPW toy problem is hard for statistical query algorithms.*

We contrast this statistical-query lower bound with the bounded-query statistical indistinguishability lower bound of [CHR17]. That result showed that there is some fixed polynomial B such that no adversary can distinguish B DE-PIR queries from random, even if computationally unbounded. In contrast, our result proves a lower bound for adversaries (also computationally unbounded), that have no a-priori polynomial bound on the number of queries that they can make – in fact, they can make up to $2^{\epsilon\lambda}$ queries where λ is the security parameter and ϵ is a small positive constant. However, they are restricted in that they cannot see the result of any individual query in its entirety; instead, adversaries can only see the result of applying bounded (up to $\epsilon\lambda$-bit) output functions separately to each query.

1.2 Other Instances of Hardness from Random Permutations

There are other instances in which random secret permutations were used to obtain computational hardness. The *Permuted Kernel Problem (PKP)* is an example in the context of a search problem. Roughly, the input in PKP consists of a matrix $A \in \mathbb{Z}_p^{m \times n}$ and a vector $\vec{v} \in \mathbb{Z}_p^n$, where p is a large prime. A solution is a permutation π on $[n]$ such that the vector $\vec{v}'$ obtained by applying π to the entries of $\vec{v}$ is in the kernel of A. PKP is known to be NP-complete in the worst-case [GJ02], and conjectured to be hard on average [Sha89], for sufficiently large $n - m$ and p. It is the underlying assumption in Shamir's identification scheme [Sha89], and has lately seen renewed interest due to its applicability to post-quantum cryptography (e.g., [LP12, FKM+18, KMP19]). Despite being studied for 3 decades, the best known algorithms to date run in exponential time; see [KMP19] and the references therein.

1.3 Techniques

We now proceed to discuss our results and techniques in greater detail.

1.3.1 Defining Permuted Puzzles

We generalize and extend the intuitive puzzle framework proposed in [BIPW17], by formally defining the notions of (permuted) puzzle distinguishing problems.

We formalize a *puzzle distinguishing problem* as a pair of distributions $\mathcal{D}_0, \mathcal{D}_1$ over Σ^n, for some alphabet Σ and some input length n. Very roughly, hardness of a puzzle distinguishing problem means one cannot distinguish a single sample from $\mathcal{D}_0$ or $\mathcal{D}_1$, even given oracle access to $\mathcal{D}_0$ *and* $\mathcal{D}_1$. We say that a puzzle problem is (s, ϵ)-*hard* if any size-s adversary distinguishes $\mathcal{D}_0$ from $\mathcal{D}_1$ with advantage at most ϵ. This concrete hardness notion naturally extends to computational hardness of an *ensemble* of puzzles, in which case we allow the distributions to be *keyed* (by both public and secret key information) and require that they be efficiently sampleable given the key.

With this notion of puzzle distinguishing problems, we turn to defining a *permuted puzzle* which, informally, is obtained by sampling a random permutation π once and for all as part of the secret key, and permutating all samples according to π. Hardness of a permuted puzzle is defined identically to hardness of (standard) puzzle distinguishing problems.

We also consider a simpler hardness definition, in which the adversary is given oracle access *only* to a randomly selected $\mathcal{D}_b$ (but not to $\mathcal{D}_{1-b}$), and attempts to guess b. We say that a puzzle distinguishing problem is *weak computationally hard* if every adversary of polynomial size obtains a negligible advantage in this modified distinguishing game. Weak computational hardness captures the security notion considered in [BIPW17], but is too weak for certain applications, as it allows for trivial permuted puzzles, e.g., $\mathcal{D}_0 = \left\{0^{n/2}1^{n/2}\right\}, \mathcal{D}_1 = \left\{1^{n/2}0^{n/2}\right\}$. More generally, and as discussed in Remark 3.10 (Section 3), weak computational hardness is generally weaker than the definition discussed above (which is more in line with the DE-PIR application). Concretely, we show that the definition discussed above implies the weaker definition, and that in certain cases (e.g., when $\mathcal{D}_1$ is the uniform distribution), the weaker definition implies the stronger one. This last observation will be particularly useful in proving security of our permuted puzzle constructions.

1.3.2 Hard Permuted Puzzle in the Random Oracle (RO) Model

Our first permuted puzzle is in the random oracle model. Recall that a permuted puzzle is defined as the permuted version of a puzzle distinguishing problem. For our RO-based permuted puzzle, the underlying puzzle distinguishing problem is defined as follows. There is no key, but both the sampling algorithm and the adversary have access to the random oracle H. The sampling algorithm samples a uniformly random input x_0 for H, and uniformly random seeds $s_1, \ldots, s_n$, where $n = \lambda$, and computes x_n sequentially as follows. For every $1 \leq i \leq n$, $x_i \stackrel{\mathsf{def}}{=} H(s_i, x_{i-1})$. The sample is then $(x_0, x'_n, s_1, \ldots, s_n)$ where $x'_n \stackrel{\mathsf{def}}{=} x_n$ in $\mathcal{D}_0$, and x'_n is uniformly random in $\mathcal{D}_1$. Notice that in this (unpermuted) puzzle distinguishing problem one can easily distinguish samples from $\mathcal{D}_0$ and $\mathcal{D}_1$, by sequentially applying the oracle to x_0 and the seeds, and checking whether the output is x'_n. This will hold with probability 1 for samples from $\mathcal{D}_0$, and only with negligible probability for samples from $\mathcal{D}_1$ (assuming H has sufficiently long outputs). The corresponding permuted puzzle is obtained by applying a fixed random permutation π^* to the seeds $(s_1, \ldots, s_n)$.[2]

Hardness of the Permuted Puzzle. We focus on a simpler case in which the adversary receives only the challenge sample (and does not request any additional samples from its challenger). This will allow us to present the main ideas of the analysis, and (as we show in Section 4), the argument easily extends to the general case.

At a very high level, we show that the hardness of the permuted puzzle stems from the fact that to successfully guess b, the adversary has to guess the underlying random permutation π^*, *even though it has*

[2] We note that syntactically, this is not a permuted puzzle since the permutation should be applied to the *entire* sample. However, this simplified view of the permuted puzzle captures the fact that in our construction, the permutation essentially operates only over the seeds. In the actual construction, this is achieved by tagging the different parts of the sample (with either "input", "output", or "seed") such that any permutation over the entire sample uniquely determines a permutation over the seeds; see Section 4.

oracle access to H.

We first introduce some terminology. For a random oracle H, input x_0 and seeds $s'_1, \ldots, s'_n$, each permutation π over the seeds uniquely defines a corresponding "output" x^π_n through a length-$(n+1)$ "path" P_π defined as follows. Let $x^\pi_0 \stackrel{\mathsf{def}}{=} x_0$, and for every $1 \leq i \leq n$, let $s''_i \stackrel{\mathsf{def}}{=} s'_{\pi^{-1}(i)}$ and $x^\pi_i \stackrel{\mathsf{def}}{=} H\left(s''_i, x^\pi_{i-1}\right)$. Then the label of the i'th node on the path P_π is x^π_i. We say that a node v with label x on some path P_π is *reachable* if x was the oracle answer to one of the adversary's queries in the distinguishing game. We note that when $s'_i = s_{\pi^*(i)}$, i.e., the seeds are permuted with the permutation used in the permuted puzzle, then $x^{\pi^*}_i = x_i$ for every $1 \leq i \leq n$. We call P_{π^*} the *special path*.

We will show that with overwhelming probability, unless the adversary queries H on all the x_i's on the special path (i.e., on $x^{\pi^*}_0, x^{\pi^*}_1, \ldots, x^{\pi^*}_n = x_n$), then he obtains only a negligible advantage in guessing b. Hardness of the permuted puzzle then follows because there are $n!$ possible paths, and the adversary has a negligible chance of guessing the special path (because π^* is a secret random permutation).

We would first like to prove that all node labels, over all paths P_π, are unique. This, however, is clearly false, because the paths are not disjoint: for example, the label of node 0 in all of them is x_0. More generally, if $\pi \neq \pi'$ have the same length-k prefix for some $0 \leq k < \lambda$, then for every $0 \leq i \leq k$, the i'th nodes on $\mathsf{P}_\pi, \mathsf{P}_{\pi'}$ have the same label. In this case, we say that the i'th nodes *correspond to the same node*. Let Unique denote the event that across all paths there do not exist two nodes that (1) do *not* correspond to the same node, but (2) have the same label. Our first observation is that Unique happens with overwhelming probability. Indeed, this holds when H's output is sufficiently large (e.g., of the order of $3\lambda \cdot \log \lambda$), because there are only $\lambda \cdot \lambda!$ different nodes (so the number of pairs is roughly of the order of $2^{2\lambda \cdot \log \lambda}$).

Let $\mathcal{E}$ denote the event that the adversary queries H on the label of an unreachable node, and let $\mathsf{ReachQ} = \bar{\mathcal{E}}$ denote its complement. Our next observation is that conditioned on Unique, ReachQ happens with overwhelming probability. Indeed, conditioned on Unique, the label of an unreachable node is uniformly random, even given the entire adversarial view (including previous oracle answers). Thus, querying H on an unreachable node corresponds to guessing the random node label. When H's output length is sufficiently large (on the order of $3\lambda \cdot \log \lambda$ as discussed above) this happens only with negligible probability.

Consequently, it suffices to analyze the adversarial advantage in the distinguishing game conditioned on $\mathsf{Unique} \wedge \mathsf{ReachQ}$. Notice that in this case, the only *potential* difference between the adversarial views when $b = 0$ and when $b = 1$ is in the label of the endpoint v_{end} of the special path P_{π^*}, which is x'_n when $b = 0$, and independent of x'_n when $b = 1$. Indeed, conditioned on Unique, the label of v_{end} appears nowhere else (i.e., is not the label of any other node on any path). Therefore, conditioned on $\mathsf{ReachQ} \wedge \mathsf{Unique}$, the label of v_{end} appears as one of the oracle answers only if v_{end} is reachable, i.e., only if the adversary queried H on all the node labels on the special path.

1.3.3 Hard Permuted Puzzles in the Plain Model

Our second permuted puzzle is based on the Decisional Diffi-Helman (DDH) assumption. The underlying puzzle distinguishing problem is defined over a multiplicative cyclic group G of prime order p with generator g. The public key consists of G, g and a uniformly random vector $\vec{u} \leftarrow \left(\mathbb{Z}^*_p\right)^n$. A sample from $\mathcal{D}_0, \mathcal{D}_1$ is of the form $(g^{x_1}, \ldots, g^{x_n})$, where in $\mathcal{D}_0$ $(x_1, \ldots, x_n)$ is chosen as a uniformly random vector that is orthogonal to $\vec{u}$, whereas in $\mathcal{D}_1$ $(x_1, \ldots, x_n)$ is uniformly random. As discussed below, in this (unpermuted) puzzle distinguishing problem one can easily distinguish samples from $\mathcal{D}_0$ and $\mathcal{D}_1$. The corresponding permuted puzzle is obtained by applying a fixed random permutation to the samples $(g^{x_1}, \ldots, g^{x_n})$.

Why are both DDH and a permutation needed? The computational hardness of the permuted puzzles stems from the *combination* of the DDH assumption and the permutation, as we now explain. To see why the DDH assumption is needed, notice that in $\mathcal{D}_0$, all sampled $(x_1, \ldots, x_n)$ belong to an $(n-1)$-dimensional subspace of $\mathbb{Z}^n_p$, whereas in $\mathcal{D}_1$ this happens only with negligible probability, because each sample is uniformly and independently sampled. Consider a simpler version in which $\mathcal{D}_0, \mathcal{D}_1$ simply output the vector $(x_1, \ldots, x_n)$. In this case, one can obtain an overwhelming distinguishing advantage by (efficiently) checking whether all samples $(x_1, \ldots, x_n)$ lie within an $(n-1)$-dimensional subspace, and if so guess that

the underlying distribution is $\mathcal{D}_0$. This "attack" can be executed even if the samples are permuted (as is the case in a permuted puzzle), because applying a permutation to the $(x_1, \ldots, x_n)$ is a linear operation, and therefore preserves the dimension of the subspace. Therefore, a permutation on its own is insufficient to get computational hardness, and we need to rely on the DDH assumption.

To see why the permutation is needed, notice that even if the DDH assumption holds in G, given $(g^{x_1}, \ldots, g^{x_n})$ one can efficiently test whether the underlying exponents $(x_1, \ldots, x_n)$ are orthogonal to a known vector $\vec{u}$, by only computing exponentiations and multiplications in G. Notice that for a sufficiently large p, the exponents of a sample from $\mathcal{D}_1$ will be orthogonal to $\vec{u}$ only with negligible probability, so this "attack" succeeds with overwhelming probability.

Hardness of the permuted puzzle. We now show that the *combination* of the DDH assumption, and permuted samples, gives computational hardness. Notice that it suffices to prove that the permuted puzzle is weak computationally hard, because $\mathcal{D}_1$ is random over G^n (see Section 1.3.1). In this case, the adversarial view $\mathsf{V}_b, b \in \{0,1\}$ consists of the public key $(G, g, \vec{u})$, and a polynomial number of permuted samples of the form $(g^{x_1}, \ldots, g^{x_n})$ which were all sampled according to $\mathcal{D}_b$ and permuted using the same random permutation π.

Our first observation is that V_b is computationally indistinguishable from the distribution $\mathcal{H}_b$ in which the public key is $(G, g, \pi'(\vec{u}))$ for $\pi' \stackrel{\mathsf{def}}{=} (\pi)^{-1}$, and the samples from $\mathcal{D}_b$ are *unpermuted*.

Our second observation is that the DDH assumption implies that $\mathcal{H}_b$ is computationally indistinguishable from the distribution $\mathcal{H}'_b$ in which the $(x_1, \ldots, x_n)$ additionally lie in a random 1-dimensional subspace $L_{b,\vec{v}}$. That is, $(x_1, \ldots, x_n)$ are chosen at random from $L_{b,\vec{v}}$, where in $\mathcal{H}'_0$ $\vec{v}$ is random subject to $\vec{v} \cdot \vec{u} = 0$, and in $\mathcal{H}'_1$ $\vec{v}$ is uniformly random. Specifically, we show that the problem of distinguishing between $\mathcal{H}_b, \mathcal{H}'_b$ can be efficiently reduced to the task of distinguishing between a polynomial number of length-$(n-1)$ vectors of the form $(g^{y_1}, \ldots, g^{y_{n-1}})$, where the $(y_1, \ldots, y_{n-1})$ are all sampled from a random 1-dimensional subspace of $\mathbb{Z}_p^{n-1}$ or all sampled from the full space $\mathbb{Z}_p^{n-1}$. If the DDH assumption holds in G then a polynomial-sized adversary cannot efficiently distinguish between these distributions [BHHO08]. Consequently, it suffices to show that $\mathcal{H}'_0, \mathcal{H}'_1$ are computationally close.

The final step is to show that $\mathcal{H}'_0, \mathcal{H}'_1$ are computationally (in fact, statistically) close. The only difference between the two distributions is in the choice of $\vec{v}$ (which is orthogonal to $\vec{u}$ in $\mathcal{H}'_0$, and random in $\mathcal{H}'_1$), where all other sampled values are either identical or deterministically determined by the choice of $\vec{v}$. Notice that in $\mathcal{H}'_1$, $(\pi(\vec{u}), \vec{v})$ is uniformly random in $\mathbb{Z}_p^n \times \mathbb{Z}_p^n$. Thus, to show that $\mathcal{H}'_0, \mathcal{H}'_1$ are statistically close and conclude the proof, it suffices to prove that $(\pi(\vec{u}), \vec{v})$ in $\mathcal{H}'_0$ is statistically close to uniform over $\mathbb{Z}_p^n \times \mathbb{Z}_p^n$. Very roughly, this follows from the leftover hash lemma due to the following observations. First, $\pi(\vec{u})$ has high min entropy even conditioned on $\vec{u}$ (because π is random). Second, the family of inner product functions with respect to a fixed vector (i.e., $h_{\vec{v}}(\vec{v}') = \vec{v} \cdot \vec{v}'$) is a pair-wise independent hash function.

Permuted Puzzles and the Learning Parity with Noise (LPN) Assumption. The argument used in the DDH-based permuted puzzle can be generalized to other situations in which it is hard to distinguish between the uniform distribution and a hidden permuted kernel (but easy to distinguish when the kernel is *not* permuted). This more general view allows us to cast the LPN assumption as a permuted puzzle, see Section 5.1.

1.3.4 Statistical-Query Lower Bound

We show that SQ algorithms that make polynomially many queries obtain only a negligible advantage in distinguishing the distributions $\mathcal{D}_0, \mathcal{D}_1$ in the toy problem presented in Section 1.1. Recall that a sample in the toy problem is a permuted $\mathsf{Graph}(\gamma)$ where γ is either a uniformly random degree-λ polynomial (in $\mathcal{D}_0$), or a uniformly random function (in $\mathcal{D}_1$), and that the SQ algorithm obtains the outputs of boolean-valued functions f of its choice on random samples. Very roughly, we will show that the outcome of f on (permutation of) a random sample $x \leftarrow \mathcal{D}_b$ is independent of the challenge bit b and the permutation π.

Notice that every permutation π over $\mathsf{Graph}(\gamma)$ defines a partition $\Phi \stackrel{\mathsf{def}}{=} \{\pi(\{i\} \times \mathbb{F})\}_{i \in \mathbb{F}}$ of $\mathbb{F} \times \mathbb{F}$, where each set in the partition corresponds to a single x value. We say that π *respects* the partition Φ. Notice also that each set contains a single non-0 entry (which is $\pi(i, \gamma(i))$, where i is the value of x that corresponds to the set). Thus, an SQ algorithm can compute this partition, so we cannot hope to hide it. Instead, we show indistinguishability even when the adversary is given the partition.

Our main observation is that for every partition Φ, and any boolean-valued function f, there exists $p_{f,\Phi} \in [0,1]$ such that for every $b \in \{0,1\}$, with overwhelming probability over the choice of random permutation π that respects the partition Φ, the expectation $\mathbb{E}_{x \leftarrow \mathcal{D}_b}[f(\pi(x))]$ is very close to $p_{f,\Phi}$, where $\pi(x)$ denote that the entries of x are permuted according to π. Crucially, $p_{f,\Phi}$ is independent of the challenge bit b, any particular sample x, and the permutation (other than the partition).

We prove this observation in two steps. First, we show that in expectation over the choice of the permutation, $\mathbb{E}_{x \leftarrow \mathcal{D}_0}[f(\pi(x))]$ and $\mathbb{E}_{x \leftarrow \mathcal{D}_1}[f(\pi(x))]$ have the same value. To see this, we write the expectations over $x \leftarrow \mathcal{D}_b$ as a weighted sum $\sum_x P_b(x) f(\pi(x))$, and apply linearity of the expectation over π. To show that this is independent of b, we observe that for any fixed x, the distribution of $\pi(x)$ is the same (i.e. does not depend on x).

Next, we show that for any distribution $\mathcal{D}$, the variance (over the choice of the permutation π) of $\mathbb{E}_{x \leftarrow \mathcal{D}_b}[f(\pi(x))]$ is small. The variance is by definition the difference between

$$\mathop{\mathbb{E}}_{\pi}\left[\mathop{\mathbb{E}}_{x \leftarrow \mathcal{D}_b}[f(\pi(x))]^2\right] \tag{1}$$

and

$$\mathop{\mathbb{E}}_{\pi}\left[\mathop{\mathbb{E}}_{x \leftarrow \mathcal{D}_b}[f(\pi(x))]\right]^2. \tag{2}$$

We show that both Eq. (1) and Eq. (2) can be expressed as an expectation (over some distribution of g, g') of $\mathbb{E}_{\pi}\left[\left(f(\pi(\mathsf{Graph}(g))), f(\pi(\mathsf{Graph}(g')))\right)\right]$. We observe that this depends only on the Hamming distance between g and g'. Finally, we observe that the distribution of (g, g') is uniform in Eq. (2) and two independent samples from $\mathcal{D}_b$ in Eq. (1). To complete the bound on the variance, we show that when g, g' are sampled independently from $\mathcal{D}_b$ (specifically, the interesting case is when they are sampled from $\mathcal{D}_0$), then the distribution of the Hamming distance between g and g' is nearly the same as when g and g' are independent uniformly random functions.

To prove this, we prove a lemma (Lemma A.4) stating that when t-wise independent random variables $(X_1, \ldots, X_n)$ satisfy $\Pr[X_i \neq \star_i] = p_i$ for some values of $\star_i$ and p_i such that $\sum_{i \in [n]} p_i \leq \frac{t}{4} \geq \omega(\log \lambda)$, then $(X_1, \ldots, X_n)$ are statistically $\mathrm{negl}(\lambda)$-close to mutually independent. We apply this with X_i being the indicator random variable for the event that $g(i) \neq g'(i)$. This lemma quantitatively strengthens a lemma of [CHR17].

1.3.5 Open Problems and Future Research Directions

The broad goal of basing DE-PIR on standard assumptions was a motivating starting point for this work, in which we put forth the framework of permuted puzzles. In describing hard permuted puzzles, we take a "bottom-up" approach by describing such constructions based on standard cryptographic assumptions. Since these permuted puzzles are still not known to imply DE-PIR, we try to close the gap between the permuted puzzle on which DE-PIR security is based, and provably hard permuted puzzles, by taking a "top down" approach, and analyzing the security of a toy version of the DE-PIR permuted puzzle, against a wide class of possible attacks.

Our work still leaves open a fascinating array of questions, we discuss some of them below. First, it would be very interesting to construct a hard permuted puzzle based only on the existence of one-way functions, as well as to provide "public key" hard permuted puzzles, namely ones in which the key generation algorithm needs no secret key, based on standard assumptions. In the context of DE-PIR and its related permuted puzzle, it would be interesting to construct DE-PIR based on other (and more standard) assumptions, as well as to analyze the security of its underlying permuted puzzle (and its toy version) against a wider class of attacks.

2 Preliminaries

For a set X, we write $x \leftarrow X$ to denote that x is sampled uniformly at random from X. For a distribution $\mathcal{D}$, we use $\mathrm{Supp}(\mathcal{D})$ to denote its support. The min entropy of $\mathcal{D}$ is $\mathsf{H}_\infty(\mathcal{D}) \stackrel{\mathsf{def}}{=} \min_{x \in \mathrm{Supp}(\mathcal{D})} \log \frac{1}{\Pr[x]}$. For a pair X, Y of random variables, we denote their statistical distance by $d_{\mathsf{TV}}(X, Y)$. We use $\cdot$ to denote inner product, i.e., for a pair $\vec{x} = (x_1, \ldots, x_n), \vec{y} = (y_1, \ldots, y_n)$ of vectors, $\vec{x} \cdot \vec{y} \stackrel{\mathsf{def}}{=} \sum_{i=1}^{n} x_i y_i$. We use $[n]$ to denote the set $\{1, \ldots, n\}$, and S_n to denote the group of permutations of $[n]$.

Notation 2.1 (Permutation of a vector)**.** For a vector $\vec{x} = (x_1, \ldots, x_n)$, and a permutation $\pi \in S_n$, we denote:

$$\pi(\vec{x}) \stackrel{\mathsf{def}}{=} \left(x_{\pi^{-1}(1)}, \ldots, x_{\pi^{-1}(n)}\right).$$

3 Distinguishing Problems and Permuted Puzzles

In this section, we formally define (permuted) puzzle problems which are, roughly, a (special case) of ensembles of keyed "string-distinguishing" problems.

We begin in Section 3.1 by developing terminology for general string-distinguishing and puzzle problems. In Section 3.2 we present the formal distinguishing challenge and define hardness. Then, in Section 3.3, we discuss the case of *permuted* puzzles, and present an alternative indistinguishability notion that is equivalent in certain cases.

3.1 String-Distinguishing Problems

At the core, we consider string-distinguishing problems, defined by a pair of distributions over n-element strings. We begin by defining a finite instance.

Definition 3.1 (String-Distinguishing Problems)**.** A string-distinguishing problem is a tuple $\Pi = (n, \Sigma, \mathcal{D}_0, \mathcal{D}_1)$, where n is a positive integer, Σ is a non-empty finite set, and each $\mathcal{D}_b$ is a distribution on Σ^n. We call n the string length, and Σ the string alphabet.

More generally, an *oracle-dependent* string-distinguishing problem is a function $\Pi^{(\cdot)}$ that maps an oracle $O : \{0,1\}^* \to \{0,1\}$ to a string-distinguishing problem Π^O.

For example, we will consider permuted puzzle string-distinguishing problems relative to a random oracle in Section 4. Note that oracle-dependent string-distinguishing problems are strictly more general than string-distinguishing problems, as the distributions can simply ignore the oracle.

Remark 3.2 (**Oracle Outputs**). In the above, we modeled the oracle as outputting a single bit for simplicity. However, any (deterministic) oracle with multi-bit output can be emulated given a corresponding single-bit-output oracle, at the cost of making more oracle queries.

We will be interested in distinguishing problems where the distributions $\mathcal{D}_0$ and $\mathcal{D}_1$ may depend on common sampled "key" information. Parts of this key may be publicly available, or hidden from a distinguishing adversary (discussed in Definition 3.7); these parts are denoted pk, sk, respectively.

Definition 3.3 (Keyed Families)**.** A *keyed family* of (oracle-dependent) string-distinguishing problems is a tuple $(\mathcal{K}, \{\Pi_k\}_{k \in \mathcal{K}})$, where $\mathcal{K}$ is a distribution on a non-empty finite set of pairs $(\mathsf{pk}, \mathsf{sk})$ and each Π_k is an (oracle-dependent) string-distinguishing problem. We refer to the support of $\mathcal{K}$ as the key space, and also denote it by $\mathcal{K}$.

Note that any string-distinguishing problem can trivially be viewed as a keyed family by letting $\mathcal{K}$ be a singleton set.

Example 3.4 (Keyed Family: Dimension-t Subspaces). For a finite field $\mathbb{F}$, and $n \in \mathbb{N}$, consider an example keyed family of string-distinguishing problems $(\mathcal{K}, \{\Pi_k\}_{k \in \mathcal{K}})$ as follows:

- $\mathcal{K}$ samples a random $t \leftarrow \{1, \ldots, n-1\}$, and a random subspace $L \subseteq \mathbb{F}^n$ of dimension t, sets $\mathsf{pk} = t$ and $\mathsf{sk} = L$, and outputs $(\mathsf{pk}, \mathsf{sk})$.
- For a key $k = (t, L)$, the corresponding string-distinguishing problem is $\Pi_k = (n, \mathbb{F}, \mathcal{D}_0, \mathcal{D}_1)$ where $\mathcal{D}_0$ outputs a uniformly random $\vec{v} \in L$, and $\mathcal{D}_1$ outputs a uniformly random $\vec{v} \in \mathbb{F}^n$.

Note that in this example, it will be computationally easy to distinguish between the distributions $\mathcal{D}_0, \mathcal{D}_1$ given sufficiently many samples.

We next define a puzzle problem which, informally, is an *efficiently sampleable* ensemble of keyed families of string-distinguishing problems.

Definition 3.5 (Puzzle problem)**.** A puzzle problem is an ensemble $\{(\mathcal{K}_\lambda, \{\Pi_k^{(\cdot)}\}_{k \in \mathcal{K}_\lambda})\}_{\lambda \in \mathbb{Z}^+}$ of keyed families of (oracle-dependent) string-distinguishing problems associated with probabilistic polynomial-time algorithms KeyGen and Samp such that:

- For any $\lambda \in \mathbb{Z}^+$, $\mathsf{KeyGen}(1^\lambda)$ outputs a sample from $\mathcal{K}_\lambda$.
- For any $k \in \mathcal{K}_\lambda$, any $b \in \{0,1\}$, and any oracle $O : \{0,1\}^* \to \{0,1\}$, $\mathsf{Samp}^O(k, b)$ outputs a sample from $\mathcal{D}_b$, where $\Pi_k^O = (n, \Sigma, \mathcal{D}_0, \mathcal{D}_1)$.

Remark 3.6 (**Abbreviated terminology**). Somewhat abusing notation, we will also refer to a *single* keyed family of string-distinguishing problems as a puzzle problem.

3.2 Distinguishing Games and Hardness

We will focus on puzzle problems where it is computationally hard to distinguish between the pair of distributions. This notion of hardness is formalized through the following distinguishing game. Roughly, the distinguishing adversary is given a challenge sample x from a randomly selected $\mathcal{D}_b$, and query access to both distributions (denoted by choices β below), and must identify from which $\mathcal{D}_b$ the x was sampled.

Definition 3.7 (Distinguishing Game)**.** Let $\mathcal{P} = (\mathcal{K}, \{\Pi_k\}_{k \in \mathcal{K}})$ be a puzzle problem, and let $\mathcal{O}$ be a distribution of oracles. The distinguishing game $\mathcal{G}^{\mathcal{O}}_{\mathsf{dist}}[\mathcal{P}]$ is run between an "adversary" $\mathcal{A}$ and a fixed "challenger" $\mathcal{C}$, and is defined as follows:

1. $\mathcal{C}$ samples a key $k = (\mathsf{pk}, \mathsf{sk})$ from $\mathcal{K}$, and $O \leftarrow \mathcal{O}$, and denote $\Pi_k^O = (n, \Sigma, \mathcal{D}_0, \mathcal{D}_1)$. $\mathcal{C}$ sends pk to $\mathcal{A}$, who is also given oracle access to O throughout the game.
2. $\mathcal{C}$ samples a random bit $b \leftarrow \{0,1\}$, samples $x \leftarrow \mathcal{D}_b$, and sends x to $\mathcal{A}$.
3. The following is repeated an arbitrary number of times: $\mathcal{A}$ sends a bit β to $\mathcal{C}$, who samples $x' \leftarrow \mathcal{D}_\beta$ and sends x' to $\mathcal{A}$.
4. $\mathcal{A}$ outputs a "guess" bit $b' \in \{0,1\}$.

$\mathcal{A}$ is said to win the game if $b' = b$. $\mathcal{A}$'s advantage is $\mathsf{Adv}_{\mathcal{A}}(\mathcal{G}^{\mathcal{O}}_{\mathsf{dist}}[\mathcal{P}]) \stackrel{\mathsf{def}}{=} 2 \cdot \left| \Pr[b' = b] - \frac{1}{2} \right|$.

Informally, a permuted puzzle is computationally hard if any polynomial-time adversary wins the distinguishing game of Definition 3.7 with negligible advantage. We first formalize the notion of *concrete* hardness.

Definition 3.8 (Concrete Hardness)**.** A puzzle problem $\mathcal{P} = (\mathcal{K}, \{\Pi_k\}_{k \in \mathcal{K}})$ is said to be (s, ϵ)-hard (with respect to oracle distribution $\mathcal{O}$) if in the game $\mathcal{G}^{\mathcal{O}}_{\mathsf{dist}}[\mathcal{P}]$, all adversaries $\mathcal{A}$ of size at most s have advantage at most ϵ.

We say a puzzle problem $\left\{(\mathcal{K}_\lambda, \{\Pi_k^{(\cdot)}\}_{k \in \mathcal{K}_\lambda})\right\}_{\lambda \in \mathbb{Z}^+}$ is $\big(s(\cdot), \epsilon(\cdot)\big)$-hard (with respect to an ensemble $\{\mathcal{O}_\lambda\}$ of oracle distributions) if each $(\mathcal{K}_\lambda, \{\Pi_k^{(\cdot)}\}_{k \in \mathcal{K}_\lambda})$ is $\big(s(\lambda), \epsilon(\lambda)\big)$-hard with respect to $\mathcal{O}_\lambda$.

Definition 3.9 (Asymptotic Hardness)**.** As usual, we say simply that $\mathcal{P}$ is (computationally) hard if for every $s(\lambda) \leq \lambda^{O(1)}$, there exists $\epsilon(\lambda) \leq \lambda^{-\omega(1)}$ such that for every $\lambda \in \mathbb{Z}^+$, $\mathcal{P}$ is $(s(\cdot), \epsilon(\cdot))$-hard.

$\mathcal{P}$ is *statistically* hard if for some $\epsilon(\lambda) \leq \lambda^{-\omega(1)}$, $\mathcal{P}$ is $(\infty, \epsilon(\cdot))$-hard against adversaries that are restricted to making a polynomial number of queries to their oracle and challenger in the distinguishing game of Definition 3.7.

Remark 3.10 (**Discussion on Definition**). A slightly simpler and more natural definition would be to give the adversary access to (polynomially-many samples from) *only* a randomly selected $\mathcal{D}_b$, where the adversary must identify b.

For keyed puzzles, these definitions are in general *not* equivalent. Consider, for example, a modified version of Example 3.4, where both $\mathcal{D}_0$ and $\mathcal{D}_1$ are defined by random dimension-t subspaces, L_0 and L_1. Then over the choice of the key (including L_0, L_1), the distributions $\mathcal{D}_0$ and $\mathcal{D}_1$ on their own are *identical*: that is, even an unbounded adversary with arbitrarily many queries would have 0 advantage in the simplified challenge. However, given t samples from *both* distributions, as in Definition 3.7, $\mathcal{D}_0$ and $\mathcal{D}_1$ are trivially separated, and a sample x can be correctly labeled with noticeable advantage. On the other hand, hardness with respect to our definition implies hardness with respect to the simplified notion, by a hybrid argument over the number of queries (see Lemma 3.13).

Since our motivation for studying puzzles come from applications where correlated samples from the corresponding distributions can be revealed (e.g., correlated PIR queries on different indices i), we thus maintain the more complex, stronger definition.

The definitional separation in the example above stems from the fact that given access to only one distribution $\mathcal{D}_b$, one cannot necessarily simulate consistent samples from $\mathcal{D}_0$ and $\mathcal{D}_1$. However, in certain instances, this issue does not arise; for example, if one of the two is simply the uniform distribution over strings. We formally address this connection in the following section: presenting the simplified indistinguishability notion in Definition 3.12, and proving equivalence for certain special cases in Lemma 3.16.

3.3 Permuted Puzzles and a Related Indistinguishability Notion

In this work we will focus on *permuted* puzzles. This is a special case of puzzle problems, as we now define. Here, the key includes an additional secret random *permutation* on the indices of the n-element strings, and strings output by the distributions $\mathcal{D}_0, \mathcal{D}_1$ will be permuted as dictated by π.

Definition 3.11 (Permuted Puzzle Problems)**.** For a puzzle problem $\mathcal{P} = \{(\mathcal{K}_\lambda, \{\Pi_k^{(\cdot)}\}_{k \in \mathcal{K}_\lambda})\}_{\lambda \in \mathbb{Z}^+}$, we define the associated permuted puzzle problem $\mathsf{Perm}(\mathcal{P}) \stackrel{\text{def}}{=} \{(\mathcal{K}'_\lambda, \{\Pi'^{(\cdot)}_{k'}\}_{k' \in \mathcal{K}'_\lambda})\}_{\lambda \in \mathbb{Z}^+}$, where:

- A sample from $\mathcal{K}'_\lambda$ is $\big(\mathsf{pk}, (\mathsf{sk}, \pi)\big)$, where:
  - $(\mathsf{pk}, \mathsf{sk})$ is sampled from $\mathcal{K}_\lambda$, and
  - If $\Pi_k = (n, \Sigma, \mathcal{D}_0, \mathcal{D}_1)$, then π is sampled uniformly at random from the symmetric group S_n.
- For any key $k' = (\mathsf{pk}, (\mathsf{sk}, \pi))$, if $\Pi_{(\mathsf{pk},\mathsf{sk})} = (n, \Sigma, \mathcal{D}_0, \mathcal{D}_1)$ then $\Pi'_{k'} = (n, \Sigma, \mathcal{D}'_0, \mathcal{D}'_1)$, where a sample from $\mathcal{D}'_b$ is $\pi(x)$ for $x \leftarrow \mathcal{D}_b$.

Recall (Notation 2.1) for vector $x \in \Sigma^n$ and $\pi \in S_n$, that $\pi(x)$ denotes the index-permuted vector.

As discussed in Remark 3.10, we now present a simplified notion of indistinguishability, and show that in certain special cases, this definition aligns with Definition 3.9. In such cases, it will be more convenient to work with the simplified version.

Definition 3.12 (Weak Hardness of Puzzle Problems)**.** Let $\mathcal{P} = (\mathcal{K}, \{\Pi_k\}_{k \in \mathcal{K}})$ and $\mathcal{O}$ be as in Definition 3.7. The simplified distinguishing game $\mathcal{G}^{\mathcal{O}}_{\mathsf{dist},s}[\mathcal{P}]$ is defined similarly to $\mathcal{G}^{\mathcal{O}}_{\mathsf{dist}}[\mathcal{P}]$, except that in Step 3, $\mathcal{C}$ samples $x' \leftarrow \mathcal{D}_b$ (instead of $x' \leftarrow \mathcal{D}_\beta$).

A puzzle problem $\mathcal{P} = (\mathcal{K}, \{\Pi_k\}_{k \in \mathcal{K}})$ is weak (s, ϵ)-hard if $\mathsf{Adv}_{\mathcal{A}}(\mathcal{G}^{\mathcal{O}}_{\mathsf{dist},s}[\mathcal{P}]) \leq \epsilon$ for any size-s adversary $\mathcal{A}$. Weak computational hardness is defined similarly to Definition 3.9.

Note that *weak* computational (statistical) hardness (with respect to Definition 3.12) is implied by hardness with respect to Definition 3.7:

Lemma 3.13 (Standard ⇒ Weak)**.** *Let $\mathcal{P} = \{(\mathcal{K}_\lambda, \{\Pi_k^{(\cdot)}\}_{k\in\mathcal{K}_\lambda})\}_{\lambda\in\mathbb{Z}^+}$ be a puzzle problem. If $\mathcal{P}$ is computationally (statistically, respectively) hard in the standard sense (Definition 3.9) then it is weak computationally (statistically, respectively) hard (Definition 3.12).*

Proof. Assume towards negation that $\mathcal{P}$ is not weak computationally hard, and let $\mathcal{A} = \{\mathcal{A}_\lambda\}$ be a (non-uniform) polynomial-time adversary that obtains a non-negligible distinguishing advantage $\epsilon = \epsilon(\lambda)$ in the simplified distinguishing game of Definition 3.12. Let q denote a bound on the number of samples which $\mathcal{A}$ obtains from his challenger throughout the game. We define a sequence of hybrids $\mathcal{H}_0, \ldots \mathcal{H}_q$ where $\mathcal{H}_i$ consists of q samples, the first i ones sampled from $\mathcal{D}_0$, and the rest from $\mathcal{D}_1$. Notice that $\mathcal{A}$ has advantage ϵ in distinguishing $\mathcal{H}_0$ from $\mathcal{H}_q$, and so there exists some $i^* \in [q]$ such that $\mathcal{A}$ distinguishes between $\mathcal{H}_{i^*}$ and $\mathcal{H}_{i^*-1}$ with non-negligible advantage ϵ/q.

We now describe an adversary $\mathcal{A}'$ that obtains advantage ϵ/q in the distinguishing game of Definition 3.7. $\mathcal{A}'$ obtains a challenge sample x sampled from $\mathcal{D}_b$ (where b is the challenge bit chosen by the challenger in the distinguishing game). $\mathcal{A}'$ then requests i^* samples from $\mathcal{D}_0$ and $q - i^* - 1$ samples from $\mathcal{D}_1$, and obtains samples $x_1, \ldots, x_{i^*}, x_{i^*+2}, \ldots, x_q$. It then runs $\mathcal{A}$ on input $x_1, \ldots, x_{i^*}, x, x_{i^*+2}, \ldots, x_q$. Notice that if x is sampled from $\mathcal{D}_0$ then $\mathcal{A}$ is run with a sample from $\mathcal{H}_{i^*+1}$, otherwise it is run with a sample from $\mathcal{H}_{i^*}$, so the distinguishing advantage of $\mathcal{A}'$ is ϵ/q, contradicting computational hardness. □

The more interesting direction is that weak hardness implies (standard) hardness in the case that one of the two distributions $\mathcal{D}_0$ or $\mathcal{D}_1$ is efficiently sampleable and *permutation-invariant*, in the following sense.

Definition 3.14 (Permutation-Invariant Distributions)**.** Let $n \in \mathbb{N}$, let Σ be a non-empty set, and let $\mathcal{D}$ be a distribution over Σ^n. For a permutation $\pi \in S_n$, let $\mathcal{D}_\pi$ be the distribution induced by sampling $x \leftarrow \mathcal{D}$ and outputting $\pi(x)$. We say that $\mathcal{D}$ is $\mathsf{permutation\text{-}invariant}$ if for a uniformly random $\pi \in S_n$, the joint distribution $\mathcal{D}_\pi \times \mathcal{D}_\pi$ is identical to $\mathcal{D} \times \mathcal{D}_\pi$.

Remark 3.15. One example of a permutation-invariant distribution $\mathcal{D}$ particularly useful in this work is the uniform distribution over Σ^n.

Lemma 3.16 (In certain cases Weak ⇒ Standard)**.** *Let $\mathcal{P} = \{(\mathcal{K}_\lambda, \{\Pi_k^{(\cdot)}\}_{k\in\mathcal{K}_\lambda})\}_{\lambda\in\mathbb{Z}^+}$ be a puzzle problem. If:*

- *The corresponding permuted puzzle $\mathsf{Perm}(\mathcal{P})$ is* weak *computationally hard (Definition 3.12).*
- *For every λ, every $k = (\mathsf{pk}, \mathsf{sk}) \in \mathrm{Supp}(\mathcal{K}_\lambda)$, and every $\Pi_k = (n, \Sigma, \mathcal{D}_0, \mathcal{D}_1)$:*
  - *$\mathcal{D}_1$ is permutation-invariant.*
  - *One can efficiently sample from $\mathcal{D}_1$ without sk.*

Then $\mathsf{Perm}(\mathcal{P})$ is computationally hard in the standard sense (Definition 3.9).

Proof. Assume towards negation that $\mathsf{Perm}(\mathcal{P})$ is not computationally hard, and let $\mathcal{A} = \{\mathcal{A}_\lambda\}$ be a (non-uniform) polynomial-time adversary that obtains a non-negligible distinguishing advantage in the distinguishing game of Definition 3.7. Let $\vec{\beta}$ denote the bits which $\mathcal{A}$ sent to its challenger in Step 3 of the game. We first show that without loss of generality, $\vec{\beta} = \vec{0}\vec{1}$ or $\vec{1}\vec{0}$ (i.e., all queries with $\beta = 0$ are made first, and all queries with $\beta = 1$ are made last, or vice-versa). Specifically, we show that the existence of $\mathcal{A}$ implies the existence of a (non-uniform) polynomial-time $\mathcal{A}' = \{\mathcal{A}'_\lambda\}$ which obtains the same distinguishing advantage, and whose queries are of this form.

$\mathcal{A}'_\lambda$ interacts with its challenger $\mathcal{C}'$, and emulates the challenger $\mathcal{C}$ for $\mathcal{A}_\lambda$. Let $T = T(\lambda)$ be a bound on the number of queries which $\mathcal{A}_\lambda$ makes (e.g., its runtime). $\mathcal{A}'_\lambda$ obtains a challenge x from $\mathcal{C}'$, which it forwards to $\mathcal{A}_\lambda$ as the challenge. Then, $\mathcal{A}'_\lambda$ makes T queries with $\beta = 0$, followed by T queries with $\beta = 1$. Let $x'_1, \ldots, x'_{2T}$ denote the challenger answers to the queries. Then, $\mathcal{A}'_\lambda$ enters Step 3 of the distinguishing

game with $\mathcal{A}_\lambda$, answering the i'th query with $\beta = 0$ ($\beta = 1$, respectively) with x'_i (x'_{T+i}, respectively). When $\mathcal{A}$ makes a guess b', $\mathcal{A}'_\lambda$ outputs b' as its own guess. Then $\mathcal{A}'$ is polynomial-time and obtains the same advantage as $\mathcal{A}$.

Consequently, we now assume that the queries of $\mathcal{A}$ are of the form $\vec{\beta} = 0^k 1^l$ for some $k = k(\lambda), l = l(\lambda) : \mathbb{N} \to \mathbb{N}$ (the case that $\vec{\beta} = \vec{10}$ is symmetric). Let $\mathsf{V}_{\lambda,b}$ denote the view of $\mathcal{A}_\lambda$ in the distinguishing game with bit b. That is, $\mathsf{V}_{\lambda,b} = \left(y^b, y'_1, \ldots, y'_k, y''_1, \ldots, y''_l\right)$ such that $y^b = \pi_\lambda\left(x^b\right)$, $y'_i = \pi_\lambda(x'_i), i = 1, \ldots, k$, and $y''_i = \pi_\lambda(x''_i), i = 1, \ldots, l$, where $\pi_\lambda \leftarrow S_{n_\lambda}$, $x^b \leftarrow \mathcal{D}_{\lambda,b}$, $x'_1, \ldots, x'_k \leftarrow \mathcal{D}_{\lambda,0}$, and $x''_1, \ldots, x''_l \leftarrow \mathcal{D}_{\lambda,1}$. Let $\mathcal{H}_\lambda$ denote the view of $\mathcal{A}_\lambda$ when the challenge, and all query answers, are sampled from $\mathcal{D}_{\lambda,1}$. That is, $\mathcal{H}_\lambda = \left(y, y'_1, \ldots, y'_{k+l}\right)$ such that $y = \pi_\lambda(x)$, and $y'_i = \pi_\lambda(x'_i), i = 1, \ldots, k+l$, where $\pi_\lambda \leftarrow S_{n_\lambda}$, $x, x'_1, \ldots, x'_{k+l} \leftarrow \mathcal{D}_{\lambda,1}$. For $b = 0, 1$ let $\mathsf{V}_b = \{\mathsf{V}_{\lambda,b}\}$, and let $\mathcal{H} = \{\mathcal{H}_\lambda\}$. We use the weak computational indistinguishability of $\mathsf{Perm}(\mathcal{P})$, and the properties of $\mathcal{D}_{1,\lambda}$, to show that $\mathcal{H}$ is computationally close to both V_0 and V_1.

$\mathcal{H}$ is computationally close to V_0. Assume towards negation that there exists a polynomial-sized family circuit $\mathcal{A}^0 = \{\mathcal{A}^0_\lambda\}$ which obtains a non-negligible advantage $\epsilon(\lambda)$ in distinguishing $\mathcal{H}$ and V_0. We construct an efficient adversary $\mathcal{A}' = \{\mathcal{A}'_\lambda\}$ which obtains advantage $\epsilon(\lambda)$ in the weak security game.

$\mathcal{A}'_\lambda$ requests $k+1$ samples from its challenger, obtaining the answers $y, y'_1, \ldots, y'_k$. Then, it samples $y''_1, \ldots, y''_l \leftarrow \mathcal{D}_{\lambda,1}$, and runs $\mathcal{A}^0_\lambda$ with $(y, y'_1, \ldots, y'_k, y''_1, \ldots, y''_l)$, outputting the guess which $\mathcal{A}^0_\lambda$ makes. Then $\mathcal{A}'$ is polynomial-time because one can efficiently sample from $\mathcal{D}_{\lambda,1}$ without the secret key. Moreover, when $b = 0$ in $\mathcal{A}'$'s simplified distinguishing game then $\mathcal{A}^0$'s input is distributed as in V_0, whereas when $b = 1$ its input is distributed as in $\mathcal{H}$. Indeed, the only difference between the input which $\mathcal{A}'$ provides to $\mathcal{A}^0$, and $\mathsf{V}_0, \mathcal{H}$, is the samples $y''_1, \ldots y''_l$, which are permuted in $\mathsf{V}_0, \mathcal{H}$ but not in the input provided by $\mathcal{A}'$. However, since $y, y'_1, \ldots, y'_k$ were generated using a random secret permutation π, then $\mathcal{A}^0$'s inputs from $\mathcal{A}'$ are distributed identically to V_0 ($\mathcal{H}$, respectively) when $b = 0$ ($b = 1$, respectively) because $\mathcal{D}$ is permutation-invariant.

$\mathcal{H}$ is computationally close to V_1. The proof is similar to the previous case. Assume towards negation that there exists a polynomial-sized family circuit $\mathcal{A}^1 = \{\mathcal{A}^1_\lambda\}$ which obtains a non-negligible advantage in distinguishing $\mathcal{H}$ and V_1. We again construct an adversary $\mathcal{A}' = \{\mathcal{A}'_\lambda\}$ breaking weak hardness. $\mathcal{A}'_\lambda$ requests k samples from its challenger, obtaining the answers $y'_1, \ldots, y'_k$, and runs $\mathcal{A}^1_\lambda$ with $(y, y'_1, \ldots, y'_k, y''_1, \ldots, y''_l)$, where $y, y''_1, \ldots, y''_l \leftarrow \mathcal{D}_{\lambda,1}$, and outputs $\mathcal{A}^1_\lambda$'s guess. Then similarly to the previous case, $\mathcal{A}'$ is polynomial-time and obtains the same distinguishing advantage as $\mathcal{A}^1$ because $\mathcal{A}^1$'s inputs from $\mathcal{A}'$ are distributed identically to V_1 ($\mathcal{H}$, respectively) when $b = 0$ ($b = 1$, respectively). □

Finally, we show that the existence of hard permuted puzzles for which the original distributions $\mathcal{D}_0, \mathcal{D}_1$ are *statistically far* implies the existence of OWFs. This follows from the fact that if $\mathcal{P}$ is not statistically hard, then $\mathsf{Perm}(\mathcal{P})$ is not statistically hard either.

Lemma 3.17. *If $\mathcal{P}$ is a puzzle problem that is not statistically hard, then $\mathsf{Perm}(\mathcal{P})$ is not statistically hard.*

Proof. Let $\mathcal{P} = \{(\mathcal{K}_\lambda, \{\Pi_k\}_{k \in \mathcal{K}_\lambda})\}_{\lambda \in \mathbb{Z}^+}$ be a given puzzle problem, with $\Pi_k = (n_k, \Sigma_k, \mathcal{D}_{k,0}, \mathcal{D}_{k,1})$. We first prove the following lemma.

Lemma 3.18. *If $\mathcal{P}$ is not statistically hard, then there is some $q(\lambda) \leq \lambda^{O(1)}$ and $\epsilon(\lambda) \geq \lambda^{-\Omega(1)}$ such that for infinitely many $\lambda \in \mathbb{Z}^+$, the random variables $M^{(0)}$ and $M^{(1)}$, sampled by the following process, are statistically $\epsilon(\lambda)$-far.*

1. *Sample $K \leftarrow \mathcal{K}_\lambda$.*
2. *For each $i \in [q(\lambda)]$ and $j \in \{0,1\}$, independently sample $m^{(0)}_{i,j} \leftarrow \mathcal{D}_{K,0}$ and $m^{(1)}_{i,j} \leftarrow \mathcal{D}_{K,j}$.*
3. *Output $M^{(b)} = \left(m^{(b)}_{i,j}\right)_{i \in [q(\lambda)], j \in \{0,1\}}$.*

Proof. Since $\mathcal{P}$ is not statistically hard then there exists an adversary $\mathcal{A}'$ in the distinguishing game of Definition 3.7 that makes a polynomial number q' of queries to its challenger, and obtains a non-negligible advantage $2\epsilon'$ in distinguishing between $\mathcal{D}_0, \mathcal{D}_1$. In particular, if $D_{b,k}$ denotes a random variable whose distribution is that of $\mathcal{D}_b$ conditioned on key $K = k$, then by Markov's inequality for infinitely many $\lambda \in \mathbb{Z}^+$,

with probability at least ϵ' over the choice of $K \leftarrow \mathcal{K}_\lambda$, $\mathcal{A}'$ obtains advantage ϵ' in the distinguishing game of Definition 3.7, conditioned on the chosen key being K. We say that K is *good* if it satisfies the above.

We set $q = (q'+1) \cdot \frac{4\lambda}{(\epsilon')^2}$ and $\epsilon = \epsilon' - e^{-2\lambda} \cdot (1+\epsilon')$. We construct an adversary $\mathcal{A}$ against $M^{(0)}, M^{(1)}$, which operates as follows. $\mathcal{A}$ obtains q pairs of samples of the form $\left(m_{i,0}^{(b)}, m_{i,1}^{(b)}\right)$, where $m_{i,0}^{(b)} \in \mathcal{D}_{K,0}$ and $m_{i,1}^{(b)} \in \mathcal{D}_{K,b}$ (for some key K). We divide these samples into two columns: a "left" column containing all samples of the form $m_{i,0}^{(b)}$, and a "right" column containing all samples of the form $m_{i,1}^{(b)}$. $\mathcal{A}$ divides its q pairs into subsets of $q'+1$ pairs, and emulates $\mathcal{A}'$ $4\lambda/(\epsilon')^2$ times as follows, using the l'th subset of samples in the l'th emulation. It picks a random bit $b_l \leftarrow \{0,1\}$, if $b_l = 0$ ($b_l = 1$) then it gives the first sample from the left (right) column to $\mathcal{A}'$ as the challenge sample. Then, it answers $\mathcal{A}'$'s queries to $\mathcal{D}_0, \mathcal{D}_1$ using the first unused sample from the left and right columns, respectively, and records $\mathcal{A}'$'s guess for b_l. For each iteration l, let X_l be indicator of the event that $\mathcal{A}'$ guessed b_l correctly, and let $X = \frac{(\epsilon')^2}{4\lambda} \sum_l X_l$. If $X \geq 1/2 + \epsilon'/2$ then $\mathcal{A}$ guesses that $b = 1$, otherwise it guesses that $b = 0$.

We now analyze the distinguishing advantage of $\mathcal{A}$, conditioned on the key K. Notice first that if $b = 0$ then the left and right columns are identically distributed (they are distributed according to $\mathcal{D}_{0,K}$) and so $\mathcal{A}'$ obtains no distinguishing advantage, i.e., $E[X_l] = E[X] = 1/2$, so by Hoeffding's inequality,

$$\Pr\left[X \geq 1/2 + \epsilon'/2\right] \leq e^{-2(\epsilon'/2)^2 \cdot \frac{4\lambda}{(\epsilon')^2}} = e^{-2\lambda}.$$

On the other hand, if $b = 1$ then conditioned on the key K being good, each iteration exactly emulates the distinguishing game for $\mathcal{A}'$, and so if K is good then $E[X_l] \geq 1/2 + \epsilon'$ so $E[X] \geq 1/2 + \epsilon'$. Therefore, by Hoeffding's inequality:

$$\Pr\left[X < 1/2 + \epsilon'/2\right] \leq \Pr\left[X \leq (1/2 + \epsilon') - \epsilon'/2\right] \leq e^{-2(\epsilon'/2)^2 \cdot \frac{4\lambda}{(\epsilon')^2}} = e^{-2\lambda}.$$

Thus, $\mathcal{A}$'s distinguishing advantage is:

$$\left|\Pr_K\left[\mathcal{A}\left(M_K^{(1)}\right) = 1\right] - \Pr_K\left[\mathcal{A}\left(M_K^{(0)}\right) = 1\right]\right|$$

$$\geq \left|\Pr_K\left[\mathcal{A}\left(M_K^{(1)}\right) = 1 | K \text{ is good}\right] \cdot \Pr_K\left[K \text{ is good}\right] + \Pr_K\left[\mathcal{A}\left(M_K^{(1)}\right) = 1 | K \text{ is bad}\right] \cdot \Pr_K\left[K \text{ is bad}\right] - \Pr_K\left[\mathcal{A}\left(M_K^{(0)}\right) = 1\right]\right|$$

$$\geq \left|\epsilon' \cdot \left(1 - e^{-2\lambda}\right) + 0 - e^{-2\lambda}\right| \geq \epsilon' - e^{-2\lambda} \cdot (1+\epsilon').$$

□

In particular, if we denote by $M_k^{(b)}$ a random variable whose distribution is that of $M^{(b)}$ conditioned on $K = k$, then there is a statistical distinguishing algorithm $\mathcal{A}$ such that for infinitely many $\lambda \in \mathbb{Z}^+$ and with probability at least $\epsilon(\lambda)/2$ over the choice of $K \leftarrow \mathcal{K}_\lambda$, we have

$$\left|\Pr\left[\mathcal{A}(M_K^{(0)}) = 1\right] - \Pr\left[\mathcal{A}(M_K^{(1)}) = 1\right]\right| \geq \frac{\epsilon(\lambda)}{2}.$$

(This follows from Markov's inequality.)

Next, we define random variables $\tilde{M}^{(0)}$ and $\tilde{M}^{(1)}$, sampled by the following process. Let $N : \mathbb{Z}^+ \to \mathbb{Z}^+$ be a polynomially bounded function such that for every $k \in \mathcal{K}_\lambda$ it holds that $n_k \leq N(\lambda)$. Let $\tilde{q}(\lambda) = 3q(\lambda) \cdot \frac{N(\lambda)^2}{\epsilon(\lambda)^2}$.

1. Sample $K \leftarrow \mathcal{K}_\lambda$ and $\pi \leftarrow S_{n_K}$.
2. For each $i \in [\tilde{q}(\lambda)]$ and $j \in \{0,1\}$, independently sample $m_{i,j}^{(0)} \leftarrow \mathcal{D}_{K,0}$ and $m_{i,j}^{(1)} \leftarrow \mathcal{D}_{K,j}$. Define $\tilde{m}_{i,j}^{(b)} = \pi\left(m_{i,j}^{(b)}\right)$.
3. Output $\tilde{M}^{(b)} = \left(\tilde{m}_{i,j}^{(b)}\right)_{i \in [\tilde{q}(\lambda)], j \in \{0,1\}}$.

To complete the proof, we show that $\tilde{M}^{(0)}$ and $\tilde{M}^{(1)}$ are statistically distinguishable, and use this to show that $\mathsf{Perm}(\mathcal{P})$ is not statistically hard.

Lemma 3.19. *$\tilde{M}^{(0)}$ and $\tilde{M}^{(1)}$ are statistically distinguishable.*

Proof. Let $\mathcal{A}$ be a distinguisher for $M^{(0)}$ vs. $M^{(1)}$, and recall that $\mathcal{A}$ makes at most q queries. That is, suppose that for infinitely many $\lambda \in \mathbb{Z}^+$, $\left|\Pr[\mathcal{A}(M^{(1)}) = 1] - \Pr[\mathcal{A}(M^{(0)}) = 1]\right| \geq \epsilon(\lambda)$. To show that $\tilde{M}^{(0)}$ and $\tilde{M}^{(1)}$ are statistically distinguishable, we show how they can be mapped to $\frac{N(\lambda)^2}{4\epsilon(\lambda)^2}$ independent samples that are (depending on whether we started with $\tilde{M}^{(0)}$ or $\tilde{M}^{(1)}$) either statistically indistinguishable or (with non-negligible probability) distinguishable by $\mathcal{A} \circ \pi$ for a suitable permutation π. Thus $\tilde{M}^{(0)}$ and $\tilde{M}^{(1)}$ can be statistically distinguished by trying all permutations π, and checking whether any $\mathcal{A} \circ \pi$ distinguishes these samples (correctness of this distinguisher is proven with Hoeffding's bound and a union bound over all possible $N(\lambda)!$ permutations).

More specifically, we interpret the samples from $\tilde{M}^{(b)}$ as consisting of left and right columns (for $j = 0$ and $j = 1$, resp.), containing lists of samples from $\mathcal{D}_{K,0}$ and $\mathcal{D}_{K,b}$ (resp.) for some key K, permuted according to the same permutation π'. The idea is to divide these lists into subsets of size $3q$, and use these subsets to construct "tests" for $\mathcal{A}$. Each such "test" will contain, in its left column, the first $2q$ samples from the left column, and the right column will consist of the remaining q samples from the left column, concatenated with the first q samples from the right column. Notice that if $b = 0$ then the two columns are identically distributed, otherwise these are samples according to $M^{(0)}$ and $M^{(1)}$, respectively, permuted according to π'. Therefore, $\mathcal{A} \circ (\pi')^{-1}$ can distinguish between these distributions. The idea is to run $\mathcal{A} \circ (\pi')^{-1}$ with each of these "tests" and check whether it distinguishes (we need multiple tests to use Chernoff's bound). Of course, π' is unknown, so instead we run $\mathcal{A}$ with *all possible permutations* and check whether there *exists some permutation* for which it distinguishes. This will be the case for $(\pi')^{-1}$ in case $b = 1$, but no such permutation exists when $b = 0$. We proceed to formalize this argument.

The key is how to map $\tilde{M}^{(0)}$ and $\tilde{M}^{(1)}$ to such samples, that (depending on b) are either statistically indistinguishable or (with non-negligible probability) distinguishable by $\mathcal{A} \circ \pi$ for some π. As before, let $\tilde{M}_k^{(b)}$ denote a random variable distributed like $\tilde{M}^{(b)}$ conditioned on $K = k$. We define a function ϕ such that:

- For any $k \in \mathcal{K}_\lambda$ and any $\pi \in S_{n_k}$, $\phi\left(\pi, \tilde{M}_k^{(0)}\right)$ is $N(\lambda)^2/4\epsilon(\lambda)^2$ independent pairs of i.i.d. random variables.
- With probability at least $\epsilon(\lambda)/2$ over a choice of $K \leftarrow \mathcal{K}_\lambda$, there exists $\pi \in S_{n_K}$ such that $\phi\left(\pi, \tilde{M}_k^{(1)}\right)$ is $N(\lambda)^2/4\epsilon(\lambda)^2$ independent samples that are $\epsilon(\lambda)/2$-distinguishable by $\mathcal{A}$.

Sketch of ϕ. When the input to ϕ is $\tilde{M}_k^{(b)}$, regardless of whether b is 0 or 1, ϕ has access to $\tilde{q}(\lambda)$ independent samples that are distributed like (permuted) samples from $\mathcal{D}_{k,0}$ (specifically the entries $\tilde{m}_{i,0}^{(b)}$ for $i \in [\tilde{q}(\lambda)]$). ϕ also has access to $\tilde{q}(\lambda)$ independent (permuted) samples from $\mathcal{D}_{k,b}$ (specifically the entries $\tilde{m}_{i,1}^{(b)}$). So, let the i^{th} output of $\phi(\pi, \tilde{M}^{(b)})$ be

$$\left(\begin{bmatrix} \pi(\tilde{m}^{(b)}_{(4i)\cdot q(\lambda),0}) & \pi(\tilde{m}^{(b)}_{(4i+1)\cdot q(\lambda),0}) \\ \vdots & \vdots \\ \pi(\tilde{m}^{(b)}_{(4i+1)\cdot q(\lambda)-1,0}) & \pi(\tilde{m}^{(b)}_{(4i+2)\cdot q(\lambda)-1,0}) \end{bmatrix}, \begin{bmatrix} \pi(\tilde{m}^{(b)}_{(4i+2)\cdot q(\lambda),0}) & \pi(\tilde{m}^{(b)}_{(4i+3)\cdot q(\lambda),1}) \\ \vdots & \vdots \\ \pi(\tilde{m}^{(b)}_{(4i+3)\cdot q(\lambda)-1,0}) & \pi(\tilde{m}^{(b)}_{(4i+4)\cdot q(\lambda)-1,1}) \end{bmatrix}\right)$$

If $b = 0$, this will be a pair of i.i.d. random variables, whereas if $b = 1$ this pair of random variables will be distinguishable by $\mathcal{A}$ (for the right choice of π). □

This completes the proof of Lemma 3.17. □

Lemma 3.20. *If $\mathcal{P}$ is a puzzle problem that is not statistically hard, but $\mathsf{Perm}(\mathcal{P})$ is computationally hard, then there exists a one-way function.*

Proof. By a result of [Gol90], it suffices to prove that there exist a pair of distribution ensembles that are efficiently sampleable, statistically far, and computationally indistinguishable. By Lemma 3.17, the pair $\tilde{M}^{(0)}$ and $\tilde{M}^{(1)}$ as defined in the proof of Lemma 3.17 are efficiently sampleable and statistically far. We now use the fact that $\mathsf{Perm}(\mathcal{P})$ is computationally hard to prove that they are computationally indistinguishable.

Assume towards negation that $\tilde{M}^{(0)}$ and $\tilde{M}^{(1)}$ are computationally distinguishable by an efficient adversary $\tilde{\mathcal{A}}$ that obtains a non-negligible distinguishing advantage $\tilde{\epsilon}$. We construct an efficient adversary $\mathcal{A}$ that obtains advantage $\epsilon = \tilde{\epsilon}/\tilde{q}$ in the distinguishing game of Definition 3.7, in contradiction to $\mathsf{Perm}(\mathcal{P})$ being computationally hard. We use a hybrid argument, defining, for $0 \leq i \leq \tilde{q}$, a hybrid $\mathcal{H}_i$ which contains the first i samples from $\tilde{M}^{(0)}$ and the rest of the samples are from $\tilde{M}^{(1)}$. Then

$$\tilde{\epsilon} \leq \left|\Pr\left[\tilde{\mathcal{A}}\left(\tilde{M}^{(0)}\right)=1\right]-\Pr\left[\tilde{\mathcal{A}}\left(\tilde{M}^{(0)}\right)=1\right]\right| = \left|\sum_{i=1}^{\tilde{q}}\left(\Pr\left[\tilde{\mathcal{A}}\left(\mathcal{H}_i\right)=1\right]-\Pr\left[\tilde{\mathcal{A}}\left(\mathcal{H}_{i-1}\right)=1\right]\right)\right|$$

$\mathcal{A}$ picks a random $i \in [\tilde{q}]$, and asks for i samples from $\mathcal{D}_0$ and $\tilde{q}-i-1$ samples from $\mathcal{D}_1$. It uses the samples from $\mathcal{D}_0$ as the first i samples, the samples from $\mathcal{D}_1$ as the last $\tilde{q}-i-1$ samples, and its challenge sample as the sample $i+1$. Then, it runs $\tilde{\mathcal{A}}$ on the resultant list of samples and outputs whatever $\tilde{\mathcal{A}}$ outputs. Notice that if the challenge bit in $\mathcal{A}$'s distinguishing game is $b=0$ ($b=1$, resp.) then $\tilde{\mathcal{A}}$ is emulated with input $\mathcal{H}_{i+1}$ ($\mathcal{H}_i$, resp.). Therefore, $\mathcal{A}$'s distinguishing advantage is:

$$\left|\frac{1}{\tilde{q}}\sum_{i=1}^{\tilde{q}}\left(\Pr\left[\tilde{\mathcal{A}}\left(\mathcal{H}_i\right)=1\right]-\Pr\left[\tilde{\mathcal{A}}\left(\mathcal{H}_{i-1}\right)=1\right]\right)\right| \geq \frac{\tilde{\epsilon}}{\tilde{q}}.$$

□

4 Hard Permuted Puzzles in the Random Oracle Model

We show that there exist computationally hard permuted puzzles in the random oracle model. We first formally define the notion of a random oracle.

Definition 4.1 (Random Oracle)**.** We use the term random oracle to refer to the uniform distribution on functions mapping $\{0,1\}^* \to \{0,1\}$.

Construction 4.2 (Permuted puzzles in the ROM)**.** *Let H be a random oracle. For a security parameter λ, we interpret H as a function $H_\lambda : \{0,1\}^{m_\lambda+\lambda} \to \{0,1\}^{m_\lambda}$ for $m_\lambda = 2(\lambda+1)\log\lambda$ (also see Remark 3.2). We define a puzzle problem $\mathcal{P} = \left\{(\mathcal{K}_\lambda, \{\Pi_k\}_{k\in\mathcal{K}_\lambda})\right\}$ by the following* KeyGen *and* Samp *algorithms:*

- $\mathsf{KeyGen}\left(1^\lambda\right)$ *outputs 1^λ as the public key (the secret key is empty).*[3]

  We note that for any λ, the corresponding string distinguishing problem $\Pi_\lambda = \left(n, \Sigma, \mathcal{D}_0^{(\cdot)}, \mathcal{D}_1^{(\cdot)}\right)$ has $n = \lambda+2$ and $\Sigma = \{0,1\}^{m_\lambda} \times \{\texttt{INPUT}, \texttt{OUTPUT}, \texttt{SEED}\}$.

- $\mathsf{Samp}\,(k,b)$ *where $k = 1^\lambda$ outputs a sample from $\mathcal{D}_{\lambda,b}^{H_\lambda}$ for $H_\lambda : \{0,1\}^{m_\lambda+\lambda} \to \{0,1\}^{m_\lambda}$ as defined above, where $\mathcal{D}_{\lambda,b}^{H_\lambda}$ is defined as follows.*

  - *A sample from $\mathcal{D}_{\lambda,0}^{H_\lambda}$ is of the form $(\sigma_1, \ldots, \sigma_{\lambda+2})$, where:*
    - *For $i \in [\lambda]$, $\sigma_i = (s_i, \texttt{SEED})$ for uniformly random and independent $s_1, \ldots, s_\lambda$ in $\{0,1\}^{m_\lambda}$.*
    - *$\sigma_{\lambda+1} = (x_0, \texttt{INPUT})$, where x_0 is uniformly random in $\{0,1\}^{m_\lambda}$.*

[3] We note that in this permuted puzzle construction the key generation stage is obsolete.

* $\sigma_{\lambda+2} = (x_\lambda, \texttt{OUTPUT})$, *where for each* $i \in [\lambda]$, $x_i = H_\lambda(s'_i, x_{i-1})$, *where* s'_i *is the length-*λ *prefix of* s_i. *(That is, the random oracle uses length-*λ *seeds, and the rest of the bits in the seed are ignored.)*

– $\mathcal{D}^{H_\lambda}_{\lambda,1}$ *is defined identically to* $\mathcal{D}^{H_\lambda}_{\lambda,0}$, *except that* x_λ *is uniformly random in* $\{0,1\}^{m_\lambda}$, *independent of* x_0, H_λ, *and* $s_1, \ldots, s_\lambda$.

Proposition 4.3. *The puzzle problem* $\mathcal{P}$ *of Construction 4.2 is computationally easy, and the corresponding permuted puzzle problem* $\mathsf{Perm}(\mathcal{P})$ *is statistically hard, with respect to a random oracle.*

We note that $\mathcal{P}$ is computationally easy in an extremely strong sense: a polynomial-sized adversary can obtain advantage $1 - \mathrm{negl}(\lambda)$ in the distinguishing game.

Before proving Proposition 4.3, we first set some notation. For a permutation $\pi \in S_n$, we define the permutation π' which, informally, is the permutation over the seeds which is induced by π. That is:

- For $i \in [\lambda]$, if $\pi(\sigma_i) = \sigma_j$ for $j \in [\lambda]$ then $\pi'(\sigma_i) = \pi(\sigma_i)$.
- For $i \in [\lambda]$, if $\pi(\sigma_i) = \sigma_j$ for $j \in \{\lambda+1, \lambda+2\}$, and $\pi(\sigma_j) = \sigma_l$ for $l \notin \{\lambda+1, \lambda+2\}$, then $\pi'(\sigma_i) = \pi(\pi(\sigma_i)) = \sigma_l$.
- For $i \in [\lambda]$, if $\pi(\sigma_i) = \sigma_j$ for $j \in \{\lambda+1, \lambda+2\}$, and $\pi(\sigma_j) = \sigma_l$ for $l \in \{\lambda+1, \lambda+2\}$, then $\pi'(\sigma_i) = \pi(\pi(\pi(\sigma_i))) = \pi(\sigma_l)$.

Next, we define the notion of a "permutation tree". Intuitively, the root node is labeled by the input x_0. Edge labels on a path from x_0 correspond to a partial permutation on the indices of the seeds for the random oracle. Node labels correspond to the output of the random oracle when recursively applied to x_0, with the seeds permuted according to the edge labels on the path.

Definition 4.4 (Permutation tree)**.** For an oracle H, and a sample $((s_1, \texttt{SEED}), \ldots (s_\lambda, \texttt{SEED}), (x_0, \texttt{INPUT}), (x_\lambda, \texttt{OUTPUT}))$ from $\mathcal{D}^H_{\lambda,0}$, we define a tree $\mathsf{T}_{x_0,H}$ of depth λ with $\lambda!$ leaves as follows. The tree is leveled from level 0 (the leaves) to level λ (the root). For every $1 \le l \le \lambda$, each node in level l has l edges leaving it. The edges are labeled recursively as follows:

- The edges leaving the root are labeled from left to right by $1, 2, \ldots, \lambda$. (We note that the edges are directed from level i to level $i-1$. If an edge is directed from a node v in level i to a node u in level $i-1$, then we say that v is u's parent.)
- For a node v in level $1 \le l < \lambda$, let L_v denote the labels of the edges on the path from the root to v. Then the edges leaving v are labeled by the labels in $[\lambda] \setminus L_v$, in increasing order from left to right.

The nodes of the tree are likewise labeled recursively, as follows:

- The label of the root is x_0.
- For every level $0 \le l < \lambda$, the label of a node v in level l is computed as follows. Let z denote the label of v's parent in level $l+1$, and let i denote the label of the edge leading from v's parent to v. Then v's label is set to $H(s'_i, z)$, where s'_i is the length-λ prefix of s_i.

For a $k \le \lambda$, a path $\mathsf{P} = (e_1, e_2, \ldots, e_k)$ is defined by the edge labels $e_1, \ldots, e_k$. It follows the edge labeled e_1 from the root to a node v_1 in level $\lambda - 1$, then follows the edge labeled e_2 from v_1 to a node v_2 in level $\lambda - 2$ and so on, until it follows the edge labeled e_k to a node in level $\lambda - k$. Let $\pi \in S_n$ denote the permutation chosen as part of the secret key in the permuted puzzle, then we call the path $(\pi'(1), \pi'(2), \ldots, \pi'(\lambda))$ the $\mathsf{special\ path}$, and we call the leaf ℓ at the end of the special path the $\mathsf{special\ leaf}$.

We first consider an $\mathcal{A}$ that receives a *single* sample from $\mathcal{D}^{H_\lambda}_{\lambda,0}$. Let $\mathcal{Q}$ denote the set of queries which $\mathcal{A}$ makes to H_λ, where a query $q \in \mathcal{Q}$ is of the form $q = (i, v)$ and denotes that $\mathcal{A}$ queried H_λ on v with seed

s'_i.[4] Let $\mathcal{Q}_t$ denote the restriction of $\mathcal{Q}$ to the first t queries which $\mathcal{A}$ makes. For a path $(e_1, \ldots, e_k)$ in the tree, where $v_0, v_1, \ldots, v_k$ are the nodes on the path (in particular, v_0 is the root), and $y_0 = x_0, y_1, \ldots, y_k$ are the corresponding node labels, we say that v_k is *reachable* if for every $1 \leq i \leq k$, $(e_i, y_{i-1}) \in \mathcal{Q}$. We use $\mathsf{Labels}\,(T_{x_0,H}, \mathcal{Q})$ to denote the set of labels of reachable nodes in $T_{x_0,H}$.

We define the following events.

1. $\mathsf{Collide}$**:** this event happens if there exist two nodes v, v' in the tree $\mathsf{T}_{x_0,H_\lambda}$ that have the same label.
2. PrevQ**:** this event occurs when the adversary queries the oracle about a node label of an unreachable node. Formally, PrevQ is the event that there exists a node v in the tree $\mathsf{T}_{x_0,H_\lambda}$ with label y such that v is not reachable but $(i, y) \in \mathcal{Q}$ for some $i \in [\lambda]$.
3. Reach**:** this is the event that the special leaf is reachable at the end of the distinguishing game of Definition 3.7.

We will prove Proposition 4.3 in four steps. First, we bound the probability that events $\mathsf{Collide}, \mathsf{PrevQ}$ occur. Then, we show that if $\mathcal{A}$ has a non-negligible advantage in the distinguishing game of Definition 3.7 then Reach happens with non-negligible probability. Third, we derive a contradiction by showing that conditioned on $\neg\mathsf{Collide}$ and $\neg\mathsf{PrevQ}$, Reach can only occur with negligible probability, since if it occurs then the adversary can guess a uniformly random permutation. In the final step, we generalize the argument to hold even when the adversary receives multiple samples.

Analysis of events $\mathsf{Collide}, \mathsf{PrevQ}$**.** We show that for our choice of m_λ, each of the events $\mathsf{Collide}$, and $\mathsf{PrevQ}|\neg\mathsf{Collide}$, happen only with $\mathrm{negl}\,(\lambda)$ probability, where in the following the probability is over the choice of H_λ and the randomness of $\mathcal{A}$.

Lemma 4.5. $\Pr[\mathsf{Collide}] = \mathrm{negl}\,(\lambda)$.

We prove the lemma by bounding, using induction, the collision probability of nodes in different levels of the tree. The inductive argument allows us to argue about the collision probability of two intermediate nodes v, v', since by the induction hypothesis we can bound the probability that the labels of their parent nodes collide. We note that such label collisions in parent nodes are the reason the lemma does not follow directly by a standard union bound over all nodes.

Proof of Lemma 4.5. First, notice that it suffices to prove the lemma conditioned on the event $\mathcal{E}$ that there are no seed collisions, because $\Pr[\neg\mathcal{E}] = \Pr[\exists i \neq i', s'_i = s'_{i'}] \leq \lambda^2 \cdot 2^{-\lambda} = \mathrm{negl}\,(\lambda)$.

For a distance $0 \leq d \leq \lambda$ from the root, let k_d denote the number of nodes in levels $\lambda, \lambda-1, \ldots, \lambda-d$ of the tree. Let $\mathcal{E}_d$ denote the event that the label of some node in one of the levels $\lambda, \lambda-1, \ldots, \lambda-d$ collides with the label of some other node in the tree $\mathsf{T}_{x_0,H_\lambda}$. We prove by induction on d that $\Pr[\mathcal{E}_d] \leq k_d \cdot 2\lambda! \cdot 2^{-m_\lambda}$. The lemma then follows from this claim because for $d = \lambda$ we get, using the fact that $k_\lambda \leq 2\lambda!$, that

$$\Pr[\mathsf{Collide}|\mathcal{E}] \leq 2\lambda! \cdot 2\lambda! \cdot 2^{-m_\lambda} \leq^{(*)} 8\pi e\lambda \left(\frac{\lambda}{e}\right)^{2\lambda} \cdot \lambda^{-2(\lambda+1)} = \mathrm{negl}\,(\lambda)$$

where the inequality denoted $(*)$ holds because $n! < \sqrt{2\pi e n}\left(\frac{n}{e}\right)^n$ and $m = 2(\lambda+1)\log\lambda$. Therefore, $\Pr[\mathsf{Collide}] \leq \Pr[\mathsf{Collide}|\mathcal{E}] + \Pr[\neg\mathcal{E}] = \mathrm{negl}\,(\lambda)$.

We now prove the claim. The basis is for $d = 0$, in which case we need to bound the probability that any node other than the root has the root label x_0. Since H is a random function, any single node label collides with x_0 only with 2^{-m_λ} probability, so by a union bound over the (at most) $2\lambda!$ nodes in the tree, $\Pr[\mathcal{E}_0] \leq 2\lambda! \cdot 2^{-m_\lambda}$ which proves the base case because $k_0 = 1$.

[4] We note that $\mathcal{A}$ can also query H_λ on seeds which are not part of the sample, i.e., with $s \notin \{s_1, \ldots, s_\lambda\}$. However, the oracle answers to such queries are uniformly random and independent of $\mathcal{D}^{H_\lambda}_{\lambda,0}$, so they can be trivially simulated with random values and we therefore disregard them.

For the step, assume the claim holds up to $d-1$ and we prove it holds for d. Notice that

$$\Pr[\mathcal{E}_d] = \Pr[\mathcal{E}_d \wedge \neg\mathcal{E}_{d-1}] + \Pr[\mathcal{E}_{d-1}] \leq \Pr[\mathcal{E}_d|\neg\mathcal{E}_{d-1}] + k_{d-1} \cdot 2\lambda! \cdot 2^{-m_\lambda}$$

where the right inequality follows from the law of conditional probability and the induction hypothesis. We now analyze $\Pr[\mathcal{E}_d|\neg\mathcal{E}_{d-1}]$. Let v be a node in level $\lambda - d < \lambda$, let y denote the label of its parent node u, and let i denote the label of the edge $u \to v$. Let $v' \neq v$ be any other node in the tree. If v' is the root then we have already proven in the base case that its label collides with the label of v only with probability 2^{-m_λ}. Otherwise, let y' be the label of the parent node u' of v', and let i' denote the label of the edge $u' \to v'$. Then the labels of v and v' collide only if $H(s'_i, y) = H(s'_{i'}, y')$. Since we have conditioned on $\neg\mathcal{E}_{d-1}$, $y' \neq y$. Moreover, since we have conditioned on $\neg\mathcal{E}$, $s'_i \neq s'_{i'}$ so $\Pr[H(s'_i, y) = H(s'_{i'}, y')] \leq 2^{-m_\lambda}$ because H is a random function. Therefore, the probability that the label of any node $v' \neq v$ (including the root) collides with the label of v is at most 2^{-m_λ}. Taking the union bound over the (at most) $2\lambda!$ nodes in the tree, the probability that the label of v collides with any other node in the tree is at most $2\lambda! \cdot 2^{-m_\lambda}$. Taking a union bound over the $k_d - k_{d-1}$ nodes in level $\lambda - d$, $\Pr[\mathcal{E}_d|\neg\mathcal{E}_{d-1}] \leq (k_d - k_{d-1}) \cdot 2\lambda! \cdot 2^{-m_\lambda}$, which proves the induction step. □

Remark 4.6 (**Collision probability for multiple samples**). The collision probability of node labels remains negligible even when $\mathcal{A}$ is given $t = \text{poly}(\lambda)$ samples from its challenger. This will be useful later when proving Proposition 4.3. To see why this holds, notice that in this case there are t permutation trees (one per samples). Let $\mathsf{Collide}^*$ be the event that there exist two nodes v, v' in two permutation trees $\mathsf{T}_{x,H_\lambda}, \mathsf{T}_{x',H_\lambda}$ respectively (possibly with $x = x'$) that have the same label. Then except with probability $t^2 \cdot \lambda^2 \cdot 2^{-\lambda} = \text{negl}(\lambda)$ no pair of the seeds collide (across all t trees). Conditioned on this event, let $\mathcal{E}^*_d$ be the event that the label of some node in levels $\lambda, \lambda-1, \ldots, \lambda-d$ in *one of the* permutation trees collides with the label of some other node in *one of the trees*. Then similarly to the proof of Lemma 4.5, $\Pr[\mathcal{E}_d] \leq t^2 \cdot k_d \cdot 2\lambda! \cdot 2^{-m_\lambda}$, and so $\Pr[\mathsf{Collide}^*] = \text{negl}(\lambda)$ because $t = \text{poly}(\lambda)$. Indeed, the proof is by induction on d, and we only describe the changes from the proof described above. In the base case, we have t root nodes, and for each of them we take a union bound over the (at most) $t \cdot 2\lambda!$ nodes in *all* trees, so $\Pr[\mathcal{E}^*_0] \leq t^2 \cdot k_0 \cdot 2\lambda! \cdot 2^{-m_\lambda}$. As for the step, condition on $\neg\mathcal{E}^*_{d-1}$, and consider a node v in level $\lambda - d < \lambda$ and some other node $v' \neq v$ in some tree. Then the labels of v and v' collide only with 2^{-m_λ} probability as in the proof of Lemma 4.5. Taking a union bound over the (at most) $t \cdot 2\lambda!$ nodes in *all* trees, and then another union bound over the $t \cdot (k_d - k_{d-1})$ nodes in level $\lambda - d$ in *all trees*, we get that $\Pr[\mathcal{E}^*_d] \leq t^2 \cdot k_d \cdot 2\lambda! \cdot 2^{-m_\lambda}$.

The following lemma will be used to bound the probability that PrevQ occurs.

Lemma 4.7. *Conditioned on* $\neg\mathsf{Collide}$*, if a node* v *in* $\mathsf{T}_{x_0,H_\lambda}$ *is unreachable then its label is uniformly random in* $\{0,1\}^{m_\lambda} \setminus \mathsf{Labels}(T_{x_0,H}, \mathcal{Q})$*, even conditioned on the entire adversarial view.*

Proof. Let $(e_1, \ldots, e_k)$ be the path from the root v_0 to node v in the tree, where $v_0, v_1, \ldots, v_k = v$ are the nodes on the path, and $y_0 = x_0, y_1, \ldots, y_{k-1}, y_k = y$ are the corresponding node labels. Since we have conditioned on $\neg\mathsf{Collide}$, these labels are unique and appear nowhere else in the tree. Moreover, v is unreachable so there exists an $1 \leq l \leq k$ such that $(e_l, y_{l-1}) \notin \mathcal{Q}$. Let l^* be the largest such l. Then y_{l^*} is uniformly distributed in $\{0,1\}^{m_\lambda} \setminus \mathsf{Labels}(T_{x_0,H}, \mathcal{Q})$, even conditioned on the entire adversarial view, because y_{l^*} is unique. Moreover, it uniquely determines the label y, so y is also uniformly distributed in $\{0,1\}^{m_\lambda} \setminus \mathsf{Labels}(T_{x_0,H}, \mathcal{Q})$ because H_λ is a random function and we have conditioned on $\neg\mathsf{Collide}$. □

Lemma 4.8. $\Pr[\mathsf{PrevQ}|\neg\mathsf{Collide}] = \text{negl}(\lambda)$.

Proof. Let v be an unreachable node in the tree, and let y be its label. Since we have conditioned on $\neg\mathsf{Collide}$ then all node labels in the tree $\mathsf{T}_{x_0,H_\lambda}$ are unique, so y appears nowhere else in the tree. Therefore, because v is unreachable then Lemma 4.7 guarantees that y is uniformly distributed in $\{0,1\}^{m_\lambda} \setminus \mathsf{Labels}(T_{x_0,H}, \mathcal{Q})$ even when conditioned on the entire adversarial view. Therefore, the probability that a single query of $\mathcal{A}$ is of the form (i, y) for some i is $\frac{\lambda}{2^{m_\lambda} - 2\lambda!}$, and using the union bound the probability that one of $\mathcal{A}$'s queries

is of the form (i, y) is at most $\frac{|\mathcal{Q}|\cdot\lambda}{2^{m_\lambda}-2\lambda!}$. Taking a union bound over the (at most) $2\lambda!$ nodes in the tree,

$$\Pr\left[\mathsf{PrevQ}|\neg\mathsf{Collide}\right] \leq 2\lambda! \cdot \frac{|\mathcal{Q}|\cdot\lambda}{2^{m_\lambda}-2\lambda!} \leq 2\sqrt{2\pi e\lambda}\cdot\left(\frac{\lambda}{e}\right)^{\lambda}\cdot|\mathcal{Q}|\cdot\lambda\cdot\frac{1}{\lambda^{2(\lambda+1)}-2\lambda!} = \mathrm{negl}\left(\lambda\right)$$

where against we used the fact that $n! < \sqrt{2\pi en}\left(\frac{n}{e}\right)^n$, and the rightmost equality holds because $|\mathcal{Q}| = \mathrm{poly}\left(\lambda\right)$.

□

Indistinguishability of adversarial views conditioned on the events. Let $\mathsf{V}_0, \mathsf{V}_1$ denote the view of $\mathcal{A}$ when the challenger in the distinguishing game of Definition 3.7 chooses $b = 0, b = 1$, respectively. We show that the views are indistinguishable, conditioned on none of the events discussed above occurring.

Lemma 4.9. *Conditioned on* $(\neg\mathsf{Collide} \wedge \neg\mathsf{PrevQ} \wedge \neg\mathsf{Reach})$*, then* $d_{\mathsf{TV}}\left(\mathsf{V}_0, \mathsf{V}_1\right) = \mathrm{negl}\left(\lambda\right)$.

Proof. The adversarial view $\mathsf{V}_\star$ for $\star \in \{0,1\}$ consists of three parts. First, it includes the sample. Second, it includes the restriction of H_λ's "truth" table to the queries in $\mathcal{Q}$. Third, it contains a partial view of the tree $\mathsf{T}_{x_0,H_\lambda}$, which includes the tree structure and the edge labels, as well as the node labels which correspond to queries $\mathcal{A}$ made to H_λ. That is, if node u in $\mathsf{T}_{x_0,H_\lambda}$ has label y, the edge $u \to v$ has label i, and $(i, y) \in \mathcal{Q}$, then v's label also appears in the partial view of the tree. We denote this tree with partial node labels by T'_{x,H_λ}. Notice that if x_λ is removed from $\mathsf{V}_0, \mathsf{V}_1$ then the views are identical.

Conditioned on $\neg\mathsf{Collide} \wedge \neg\mathsf{PrevQ}$, only reachable nodes have labels in $\mathsf{T}'_{x_0,H_\lambda}$. Conditioned on $\neg\mathsf{Reach}$, the special leaf is unreachable so its label does not appear in $\mathsf{T}'_{x_0,H_\lambda}$. Conditioned on $\neg\mathsf{Collide}$, this label appears nowhere else in the tree. Therefore, x_λ is uniformly random in V_1, and uniformly distributed over a set of size $2^{m_\lambda} - 2\lambda!$ in V_0, so $d_{\mathsf{TV}}\left(\mathsf{V}_0, \mathsf{V}_1\right) = \mathrm{negl}\left(\lambda\right)$. □

Bounding the adversarial advantage through permutation guessing. Assume that $\mathcal{A}$ has a noticeable advantage in guessing b in the distinguishing game of Definition 3.7. Then Lemma 4.9 guarantees that at least one of the events $\mathsf{Collide}, \mathsf{PrevQ}, \mathsf{Reach}$ occur with noticeable probability. Combining this with Lemmas 4.5 and 4.8, if $\mathcal{A}$ has a noticeable advantage then Reach happens with some noticeable probability p (both when $b = 0$ and when $b = 1$). We now derive a contradiction by showing that $\mathcal{A}$ can be used to guess a uniformly random permutation with probability p, which is impossible. We proceed to formalize this intuition. We define a *permutation guessing game* in which, roughly, the adversarial goal is to guess a random permutation on which it has no information.

Definition 4.10 (Permutation Guessing Game)**.** The $\mathsf{permutation\ guessing\ game}$ is parameterized by a security parameter λ, and run between an adversary $\mathcal{A}_p$ and a challenger $\mathcal{C}_p$:

- $\mathcal{C}_p$ picks a random permutation $\pi \in S_\lambda$, and sends "Init" to $\mathcal{A}_p$.
- The game proceeds for $\mathrm{poly}\left(\lambda\right)$ rounds, where in each round $\mathcal{A}_p$ sends a permutation $\pi'' \in S_\lambda$ to $\mathcal{C}_p$.
- $\mathcal{A}_p$ wins if in one of the rounds $\pi'' = \pi$.

Clearly, any (even computationally unbounded) adversary wins the permutation guessing same only with $\mathrm{negl}\left(\lambda\right)$ probability.

Lemma 4.11. *Conditioned on* $\neg\mathsf{Collide} \wedge \neg\mathsf{PrevQ}$*, there exists an adversary* $\mathcal{A}_p$ *that wins the permutation guessing game with probability at least* $\Pr\left[\mathsf{Reach}\right]$.

Proof. We describe the adversary $\mathcal{A}_p$, which runs $\mathcal{A}$ as a sub-routine. $\mathcal{A}_p$ chooses $s_1, \ldots, s_\lambda, x_0, x_\lambda \in_R \{0,1\}^{m_\lambda}$. It sends $\left(\left(s_1, \mathtt{SEED}\right), \ldots, \left(s_\lambda, \mathtt{SEED}\right), \left(x_0, \mathtt{INPUT}\right), \left(x_\lambda, \mathtt{OUTPUT}\right)\right)$ to $\mathcal{A}$ as the sample. It answers all of $\mathcal{A}$'s queries to H_λ randomly but consistently. That is, $\mathcal{A}_p$ records a table of past queries of $\mathcal{A}$ to the oracle. Given a query (i, y), if it already appears in the table then $\mathcal{A}_p$ returns the value written in that entry. Otherwise, $\mathcal{A}_p$ picks the answer uniformly at random, and records the query in the table. Additionally, $\mathcal{A}_p$

maintains the partial permutation tree T'_{x_0,H^*} (here, we use H^* to denote the oracle which $\mathcal{A}_p$ simulates for $\mathcal{A}$). Whenever a leaf v in the tree becomes reachable, $\mathcal{A}_p$ asks its challenger about the permutation π'' defined by the path $(e_1,\ldots,e_\lambda)$ to v, namely $\pi''(i)=e_i$ for every $i\in[\lambda]$. Notice that if Reach happens then $\mathcal{A}_p$ wins (here, we think of the random permutation π as being both the permutation chosen by the challenger, and the permutation underlying $\mathcal{D}^{H_\lambda}_{\lambda,0}$; this is possible since both are uniformly random in S_λ). Conditioned on $\neg\mathsf{Collide}\wedge\neg\mathsf{PrevQ}$, $\mathcal{A}$'s view in the reduction is identical to V_1, so Reach occurs with the same probability as in V_1. □

The following is a direct corollary of Lemma 4.11 because the permutation guessing game can be won only with negligible probability.

Corollary 4.12. $\Pr[\mathsf{Reach}|\neg\mathsf{Collide}\wedge\neg\mathsf{PrevQ}]=\mathrm{negl}(\lambda)$.

The following is a direct corollary of Corollary 4.12 and Lemmas 4.5, 4.8 and 4.9.

Corollary 4.13. *Any adversary $\mathcal{A}$ in the distinguishing game of Definition 3.7 that obtains a single sample from its challenger has only a negligible advantage in guessing b.*

Proof sketch for Proposition 4.3. We now prove Proposition 4.3.

Proof sketch for Proposition 4.3. Notice first that $\mathcal{P}$ is computationally easy, even given a single sample. Indeed, given a sample $((s_1,\mathtt{SEED}),\ldots,(s_\lambda,\mathtt{SEED}),(x_0,\mathtt{INPUT}),(x_\lambda,\mathtt{OUTPUT}))$ an adversary can sequentially call the oracle to compute $x_i=H_\lambda\left(s'_{i-1},x_{i-1}\right)$ and check whether the outcome of the λ'th call is x_λ. This will hold with probability 1 for a sample from $\mathcal{D}_0$, but will only hold with probability $2^{-m_\lambda}\leq 2^{-\lambda}$ for samples from $\mathcal{D}_1$.

To show statistical hardness of $\mathsf{Perm}(\mathcal{P})$, we describe how to extend the argument described above to hold for an adversary who receives $t=\mathrm{poly}(\lambda)$ samples from its challenger. In this case, we have t permutation trees - one for every sample received from the challenger, each with its own special leaf.

- Recall that $\mathsf{Collide}^*$ is the event that there exist two nodes v,v' in two permutation trees $\mathsf{T}_{x,H_\lambda},\mathsf{T}_{x',H_\lambda}$ respectively (possibly with $x=x'$) that have the same label. Then as discussed in Remark 4.6, $\Pr[\mathsf{Collide}^*]=\mathrm{negl}(\lambda)$.
- Lemma 4.7 now holds (with a similar proof) when conditioning on $\neg\mathsf{Collide}^*$ and the entire adversarial view (which includes all samples).
- Let PrevQ^* be the event that in *some* tree T_{x,H_λ} there exists a node v with label y such that v is not reachable, but $\mathcal{A}$ queried H_λ about (i,y) for some $i\in[\lambda]$. A similar proof to the proof of Lemma 4.8 shows that $\Pr[\mathsf{PrevQ}^*|\neg\mathsf{Collide}^*]=\mathrm{negl}(\lambda)$. (The only difference is that the bound now increases by a multiplicative factor of t due to the larger number of nodes over all permutation trees; since $t=\mathrm{poly}(\lambda)$ the upper bound is still $\mathrm{negl}(\lambda)$.)
- Let Reach^* be the event that the special leaf of *some* tree T_{x,H_λ} is reachable at the end of the game. We can now prove that conditioned on $(\neg\mathsf{Collide}^*\wedge\neg\mathsf{PrevQ}^*\wedge\neg\mathsf{Reach}^*)$, it holds that $d_{\mathsf{TV}}(\mathsf{V}_0,\mathsf{V}_1)=\mathrm{negl}(\lambda)$. The proof is similar to the proof of Lemma 4.9, where the adversarial view now includes t samples and t partial permutation trees, and we note that if the special leaves in all of these permutation trees are removed then the views are identical. Similarly to the proof of Lemma 4.9, conditioning on $(\neg\mathsf{Collide}^*\wedge\neg\mathsf{PrevQ}^*\wedge\neg\mathsf{Reach}^*)$ guarantees that none of these special leaves appear in the views.
- Finally, we can prove, similarly to the proof of Lemma 4.11, that conditioned on $\neg\mathsf{Collide}^*\wedge\neg\mathsf{PrevQ}^*$, there exists an $\mathcal{A}_p$ that wins the permutation guessing game with probability at least Reach^*, which means $\Pr[\mathsf{Reach}^*|\neg\mathsf{Collide}^*\wedge\neg\mathsf{PrevQ}^*]=\mathrm{negl}(\lambda)$. The adversary $\mathcal{A}_p$ answers all sample queries of $\mathcal{A}$ as in the proof of Lemma 4.11, maintains a permutation tree for each sample, and whenever a leaf in *any* of the trees becomes reachable it asks its challenger about the corresponding permutation. We conclude from the discussion above that $\Pr[\neg\mathsf{Collide}^*\wedge\neg\mathsf{PrevQ}^*\wedge\neg\mathsf{Reach}^*]\geq 1-\mathrm{negl}(\lambda)$, so $\mathsf{V}_0,\mathsf{V}_1$ are statistically close. □

5 Hard Permuted Puzzles in the Plain Model

In this section we discuss permuted puzzle problems based on hidden permuted kernels. At a high level, these puzzles have the following structure. First, the distributions $\mathcal{D}_0, \mathcal{D}_1$ are associated with a group G with generator g, and a uniformly random public "constraint vector" $\vec{c}$. Samples from $\mathcal{D}_0$ and $\mathcal{D}_1$ are vectors in G^m, of the form $g^{\vec{x}}$. Specifically, $\mathcal{D}_1$ samples a uniformly random vector in G^m, whereas $\mathcal{D}_0$ samples a vector $\vec{x}$ that is uniformly random subject to being orthogonal to $\vec{c}$. Intuitively, since $\mathcal{D}_1$ is uniformly random, weak computational hardness of the permuted puzzle problem implies computational hardness by Lemma 3.16.

Remark 5.1 (**An alternative formulation of the problem**). In the high-level blueprint of a permuted puzzle problem described above, the constraint vector $\vec{c}$ is given "in the clear" (namely, we assume it is public, and indistinguishability does not rely on the secrecy of $\vec{c}$), and the samples $\vec{x}$ are permuted according to a random permutation $\pi \in S_n$, namely, the adversary obtains $\pi(\vec{x})$ (recall that $\pi(\vec{x}) = \left(x_{\pi^{-1}(1)}, \ldots, x_{\pi^{-1}(n)}\right)$ Let $\mathcal{C}$ denote the set of "good" vectors $\vec{c}$, i.e., vectors that satisfy the requirement, and let G^n denote the domain over which $\mathcal{D}_0, \mathcal{D}_1$ are defined. Let $\mathcal{D}'_b \stackrel{\text{def}}{=} \left(\vec{c}, (\pi(\vec{x}_i))_{i\in[q]}\right)_{\vec{c}\leftarrow\mathcal{C}, \pi\leftarrow S_n, \vec{x}_i\leftarrow\mathcal{D}_b}$ denote the distribution over the adversary's view in the simplified distinguishing game of Definition 3.12, where b is the challenge bit, and q is the number of samples the adversary receives from the challenger. Denote $\mathcal{D}''_b \stackrel{\text{def}}{=} \left(\pi(\vec{c}), (\vec{x}_i)_{i\in[q]}\right)_{\vec{c}\leftarrow\mathcal{C}, \pi\leftarrow S_n, \vec{x}_i\leftarrow\mathcal{D}_b}$. The permuted puzzle problems described in this section will have the property that $\mathcal{D}'_b \approx \mathcal{D}''_b$ for $b \in \{0,1\}$, which will be used in the security proofs.

5.1 Permuted Puzzles and the Learning Parity With Noise (LPN) Assumption

We now describe how to cast the Learning Parity with Noise (LPN) assumption as a permuted puzzle.

Notation. For $n \in \mathbb{N}$, let $\mathcal{R}_n$ denote the distribution that outputs a uniformly random $\vec{x} \leftarrow \mathbb{F}_2^n$. For a fixed $\vec{s} \in \mathbb{F}_2^n$, and $\gamma \in (0,1)$, let $\mathcal{D}_{\mathsf{LPN},\vec{s},\gamma}$ denote the distribution over $\mathbb{F}_2^n$ that with probability γ outputs a uniformly random $\vec{x} \leftarrow \mathbb{F}_2^n$, and otherwise (with probability $1-\gamma$) outputs a uniformly random element of the set $\{\vec{x} \in \mathbb{F}_2^n \ : \ \vec{x} \cdot \vec{s} = 0\}$.

Assumption 5.2 (Learning Parity with Noise (LPN))**.** Let $\gamma \in (0,1)$. The γ-Learning Parity with Noise (γ-LPN) assumption conjectures that for every polynomial-sized oracle circuit ensemble $\mathcal{A} = \{\mathcal{A}_\lambda\}_\lambda$ there exists a negligible function $\epsilon(\lambda)$ such that for every λ,

$$\mathsf{Adv}^{\mathsf{LPN}}_{\mathcal{A}}(\lambda) \stackrel{\text{def}}{=} \left| \Pr_{\vec{s}\leftarrow\mathbb{F}_2^\lambda}\left[\mathcal{A}^{\mathcal{D}_{\mathsf{LPN},\vec{s},\gamma}}(1^\lambda) = 1\right] - \Pr\left[\mathcal{A}^{\mathcal{R}_\lambda}(1^\lambda) = 1\right] \right| \leq \epsilon(\lambda).$$

Remark 5.3 (**Equivalence to standard LPN formulation**). Recall that the standard LPN assumption states that for all $0 < \gamma < \frac{1}{2}$, any polynomial-time adversary obtains only a negligible advantage in distinguishing between (polynomially many samples from) the following distributions:

- $(\vec{a}_i, \langle \vec{a}_i, \vec{s}\rangle + e_i)_{i=1}^m$, where for every i, $\vec{a}_i \leftarrow \mathbb{F}_2^n$ and e_i is sampled from a Bernoulli distribution with $\Pr[e_i = 1] = \gamma$; vs.
- $(\vec{a}_i, u_i)_{i=1}^m$, where each $(\vec{a}_i, u_i)$ is sampled uniformly at random from $\mathbb{F}_2^{n+1}$.

We now show that if the standard LPN assumption holds with parameters $\lambda - 1, \gamma/2$, then Assumption 5.2 holds with parameters λ, γ, where the distinguishing advantage increases by at most $2^{-\lambda}$. To see why, notice that in Assumption 5.2 if $\vec{s} = 0$ then $\mathcal{D}_{\mathsf{LPN},\vec{s},\gamma}$ and $\mathcal{R}_\lambda$ are identically distributed, whereas in the standard LPN formulation they might be distinguishable (with some advantage ≤ 1). Conditioned on $\vec{s} \neq 0$, in 5.2 there exists at least one coordinate $i \in [\lambda]$ such that the i'th coordinate of a sample from $\mathcal{D}_{\mathsf{LPN},\vec{s},\gamma}$ is a noisy linear function of the other coordinates. Moreover, since $\vec{s}$ is uniformly random over non-zero vectors, i is uniformly distributed in $[\lambda]$. In contrast, in the standard LPN formulation this "special" coordinate is the last one. Now, assume Assumption 5.2 does not hold, and let D be the corresponding distinguisher. We use

D to break the standard LPN assumption with parameters $\lambda - 1, \gamma/2$. The distinguisher D' for LPN picks a random $i \leftarrow [\lambda]$, and then emulates D. Whenever D asks for a sample, D' obtains a sample $\vec{a}, b$ from its oracle, and forwards $a_1, \ldots, a_{i-1}, b, a_i, \ldots, a_{\lambda-1}$ to D. If the oracle of D' is the uniform distribution, then so is the oracle it simulates for D. Otherwise, the location of the "special" coordinate in samples provided to D is distributed as in samples from $\mathcal{D}_{\mathsf{LPN},\vec{s},\gamma}$. Moreover, for every sample, with probability $1-\gamma$, the error term $e = 0$, in which case the sample provided to D is orthogonal to $\vec{s}$; and with the remaining γ probability, e is uniformly random, in which case the sample provided to D is uniformly random. Therefore, D' perfectly simulates $\mathcal{D}_{\mathsf{LPN},\vec{s},\gamma}$ for D, and obtains the same distinguishing advantage.

We now describe how to view Assumption 5.2 as assuming that $\mathsf{Perm}(\mathcal{P})$ is computationally hard for some computationally easy puzzle $\mathcal{P}$. For $i \in [n]$, we denote by $\vec{\mathsf{v}}_{n,i}$ the string $1^i \cdot 0^{n-i}$ (i.e., a canonical n-bit string of Hamming weight i).

Construction 5.4 (Permuted puzzle problem from LPN)**.** *For a noise parameter $\gamma \in (0, 1/2)$, we define a puzzle problem $\mathcal{P} = \left\{(\mathcal{K}_\lambda, \{\Pi_k\}_{k \in \mathcal{K}_\lambda})\right\}$ by the following* KeyGen *and* Samp *algorithms:*

- $\mathsf{KeyGen}\left(1^\lambda\right)$ *samples a weight w according to the binomial distribution over $[n]$. It outputs w as the secret key (there is no public key).*

  For a key k generated by $\mathsf{KeyGen}\left(1^\lambda\right)$, the corresponding string-distinguishing problem $\Pi_k = \left(n, \Sigma, \mathcal{D}_0, \mathcal{D}_1\right)$ has string length $n = \lambda$ and alphabet $\Sigma = \mathbb{F}_2$.

- $\mathsf{Samp}\left(w, b\right)$ *outputs a sample from $\mathcal{D}_{\lambda,b}$, where $\mathcal{D}_{\lambda,0} = \mathcal{D}_{\mathit{LPN},\vec{\mathsf{v}}_{\lambda,w},\gamma}$, and $\mathcal{D}_{\lambda,1} = \mathcal{R}_\lambda$.*

Proposition 5.5. *For any constant $\gamma \in (0, 1/2)$, the γ-LPN assumption is equivalent to the computational hardness of the permuted puzzle problem* $\mathsf{Perm}\left(\mathcal{P}_\gamma\right)$ *of Construction 5.4.*

Proof. Now observe that the permuted distribution $\mathcal{D}'_{\lambda,0}$ of the permuted puzzle is exactly $\mathcal{D}_{\mathsf{LPN},\vec{s},\gamma}$, where $\vec{s} = \pi\left(\vec{\mathsf{v}}_{\lambda,\mathsf{w}}\right)$ for a uniformly random $\pi \in S_\lambda$, and a weight $\mathsf{w} \in [\lambda]$ which was sampled according to the binomial distribution, so $\vec{s}$ is uniformly random in $\mathbb{F}_2^n$. Therefore, the distinguishing advantage in the distinguishing game of the permuted puzzle corresponds exactly to the γ-LPN assumption (because additionally $\mathcal{D}'_{\lambda,1} = \mathcal{R}_\lambda$). □

Remark 5.6 (**(Unpermuted) puzzle problem is computationally easy**)*.* We note that the (unpermuted) puzzle problem of Construction 5.4 is computationally easy. Indeed, in the unpermuted puzzle problem there are only λ possible "secret" vectors (i.e., $\vec{\mathsf{v}}_{\lambda,1}, \ldots, \vec{\mathsf{v}}_{\lambda,\lambda}$). Given a polynomial number of samples from $\mathcal{D}_{\lambda,0}$ the adversary can determine, with overwhelming probability, which of these is the secret vector used in $\mathcal{D}_{\lambda,0}$, and can then determine (with constant advantage) whether the challenge sample is from $\mathcal{D}_{\lambda,0}$ or $\mathcal{D}_{\lambda,1}$.

5.2 Permuted Puzzles Based on DDH

In this section we describe a permuted puzzle problem based on the DDH assumption. We first recall the standard DDH assumption, and describe an equivalent formulation which we use.

Definition 5.7 (Group Samplers)**.** A group sampler is a probabilistic polynomial-time algorithm $\mathcal{G}$ that on input 1^λ outputs a pair (G, g), where G is a multiplicative cyclic group of order $p = \Theta(2^\lambda)$, and g is a generator of G. We assume that p is included in the group description G, and that there exists an efficient algorithm that given G and descriptions of group elements g_1, g_2 outputs a description of $g_1 \cdot g_2$.

Definition 5.8 (DDH assumption)**.** For any cyclic group G of order p with generator g, define the following distributions:

- $\mathcal{D}_{\mathsf{DDH}}(G, g)$ is uniform over the set $\left\{(g^x, g^y, g^{xy}) : x, y \in \mathbb{Z}_p\right\}$.
- $\mathcal{R}_{\mathsf{DDH}}(G, g)$ is uniform over G^3.

For a group sampler $\mathcal{G}$, the DDH assumption over $\mathcal{G}$ conjectures that for any polynomial-sized circuit family $\mathcal{A} = \{\mathcal{A}_\lambda\}_\lambda$ there exists a negligible function $\epsilon(\lambda)$ such that for every λ:

$$\mathsf{Adv}_{\mathcal{A}}^{\mathsf{DDH}(\mathcal{G})}(\lambda) \stackrel{\mathsf{def}}{=} \left| \Pr_{\substack{(G,g)\leftarrow\mathcal{G}(1^\lambda)\\ v\leftarrow\mathcal{D}_{\mathsf{DDH}}(G,g)}}[\mathcal{A}_\lambda(v)=1] - \Pr_{\substack{(G,g)\leftarrow\mathcal{G}(1^\lambda)\\ v\leftarrow\mathcal{R}_{\mathsf{DDH}}(G,g)}}[\mathcal{A}_\lambda(v)=1] \right| \leq \epsilon(\lambda).$$

We will use the matrix version of DDH, defined next. Informally, in matrix DDH the adversary is given many vectors of the form $(g^{x_1}, \ldots, g^{x_n})$, and the conjecture is that no polynomial-time adversary can distinguish between the case that the $(x_1, \ldots x_n)$ are sampled uniformly from $\mathbb{Z}_p^n$, and the case that $(x_1, \ldots, x_n)$ are sampled from a random 1-dimensional subspace of $\mathbb{Z}_p^n$.

Definition 5.9 (Matrix DDH assumption)**.** For a cyclic group G of order p, and $n, q \in \mathbb{N}$, define

$$\mathsf{Rk}_i\left(G^{q\times n}\right) = \left\{ g^A = \left(g^{a_{ij}}\right)_{i\in[q],j\in[n]} \;:\; A \in \mathbb{Z}_p^{q\times n}, \mathsf{rank}(A) = i \right\}.$$

Let $\mathcal{G}$ be as in Definition 5.8, and let $n = n(\lambda), q = q(\lambda)$ be polynomials such that $q(\lambda) \geq n(\lambda)$ for every λ. The matrix DDH assumption over $\mathcal{G}$ conjectures that for any polynomial-sized circuit family $\mathcal{A} = \{\mathcal{A}_\lambda\}_\lambda$ there exists a negligible function $\epsilon(\lambda)$ such that for every λ:

$$\mathsf{Adv}_{\mathcal{A}}^{\mathsf{M\text{-}DDH}(\mathcal{G})}(\lambda) \stackrel{\mathsf{def}}{=} \left| \Pr_{\substack{(G,g)\leftarrow\mathcal{G}(1^\lambda)\\ v\leftarrow\mathsf{Rk}_n\left(G^{q\times n}\right)}}[\mathcal{A}_\lambda(v)=1] - \Pr_{\substack{(G,g)\leftarrow\mathcal{G}(1^\lambda)\\ v\leftarrow\mathsf{Rk}_1\left(G^{q\times n}\right)}}[\mathcal{A}_\lambda(v)=1] \right| \leq \epsilon(\lambda).$$

Boneh et al. proved [BHHO08, Lemma 1] that the DDH assumption over $\mathcal{G}$ implies the matrix DDH assumption over $\mathcal{G}$:

Imported Theorem 5.10 (DDH implies matrix-DDH [BHHO08])**.** *Let λ be a security parameter, let $\mathcal{G}$ be as in Definition 5.8, and let $n = n(\lambda), q = q(\lambda)$ be polynomials. Then for any polynomial-sized adversary circuit $\mathcal{A}_{M\text{-}DDH}$ there exists an adversary $\mathcal{A}_{DDH}$ of size $|\mathcal{A}_{M\text{-}DDH}| + \mathrm{poly}(q,n)$ such that $\mathsf{Adv}_{\mathcal{A}_{M\text{-}DDH}}^{M\text{-}DDH(\mathcal{G})}(\lambda) \leq (n-1)\cdot \mathsf{Adv}_{\mathcal{A}_{DDH}}^{DDH(\mathcal{G})}(\lambda)$.*

We are now ready to define the permuted puzzle problem based on DDH.

Construction 5.11 (Permuted puzzle problem from DDH)**.** *Let $\mathcal{G}$ be as in Definition 5.8. We define a puzzle problem $\mathcal{P} = \left\{(\mathcal{K}_\lambda, \{\Pi_k\}_{k\in\mathcal{K}_\lambda})\right\}$ by the following* KeyGen *and* Samp *algorithms:*

- KeyGen *on input 1^λ samples $(G,g) \leftarrow \mathcal{G}(1^\lambda)$, where $\mathcal{G}$ is the group sampling algorithm of Definition 5.8. Let p denote the order of G. Then,* KeyGen *samples a uniformly random vector $\vec{u} \in \mathbb{Z}_p^n$ for $n = \lambda^2$ and outputs $(G, g, \vec{u})$ as a public key (there is no secret key).*

  We note that for any $k = (G, g, \vec{u})$, the corresponding string distinguishing problem $\Pi_k = (n, \Sigma, \mathcal{D}_0, \mathcal{D}_1)$ has alphabet $\Sigma = G$.

- $\mathsf{Samp}(k, b)$ *for $k = (n, \Sigma, \mathcal{D}_0, \mathcal{D}_1)$ outputs a sample from $\mathcal{D}_b$, where:*
  - *$\mathcal{D}_0$ is uniform over $\left\{g^{\vec{x}} \in G^n : \vec{x}\cdot\vec{u} = 0\right\}$.*
  - *$\mathcal{D}_1$ is uniform over G^n.*

Proposition 5.12. *The puzzle problem $\mathcal{P}$ of Construction 5.11 is computationally easy. Moreover, if $\mathcal{G}$ is an ensemble of groups in which the matrix DDH assumption of Definition 5.9 holds, then the corresponding permuted puzzle problem $\mathsf{Perm}(\mathcal{P})$ is computationally hard.*

We note that $\mathcal{P}$ is computationally easy in an extremely strong sense: a polynomial-sized adversary can obtain advantage $1 - \mathrm{negl}(\lambda)$ in the distinguishing game.

Proof. We first show that $\mathcal{P}$ is computationally easy, even given only the challenge sample $(g^{x_1}, \ldots, g^{x_n})$. Indeed, for every $i \in [n]$ the adversary computes $h_i = (g^{x_i})^{u_i}$, and outputs 1 if and only if $\prod_{i=1}^{n} h_i = \mathrm{id}$, where id is the identity of the group. If $(g^{x_1}, \ldots, g^{x_n})$ was sampled from $\mathcal{D}_0$ then the adversary outputs 1 with probability 1, otherwise he outputs 1 with probability $1/p = 2^{-\Omega(\lambda)}$ (he outputs 1 only if $x_n = -u_n^{-1} \cdot \sum_{i=1}^{n-1} x_i u_i$, and x_n is random in $\mathbb{Z}_p$). Therefore, the adversary's advantage in the distinguishing game of Definition 3.7 is $1 - 2^{-\Omega(\lambda)}$.

We now prove that $\mathsf{Perm}(\mathcal{P})$ is computationally hard. For any key k, the distribution $\mathcal{D}_1$ is permutation-invariant (because it is random) and can be sampled from without the secret key (because there is no secret key). Therefore, it suffices to prove weak computational hardness because by Lemma 3.16 this implies computational hardness. We proceed to prove weak computational hardness.

To prove weak computational hardness, we need to show that the view of an adversary in the simplified distinguishing game of Definition 3.12 when $b = 0$ is computationally indistinguishable from the view when $b = 1$. Without loss of generality, suppose that there is a polynomial $q(\lambda) \geq n$ so that on every input, the adversary makes exactly $q(\lambda)$ queries. In this case, the adversary's view is $(G, g, g^{\pi(\vec{x}_1)}, \ldots, g^{\pi(\vec{x}_q)})$, where $g^{\vec{x}_1}, \ldots, g^{\vec{x}_q}$ are sampled independently from $\mathcal{D}_b$ and π is a uniformly random permutation of $[n]$.

We will prove the indistinguishability of the cases $b = 0$ and $b = 1$ with a hybrid argument. We define ensembles of hybrid distributions $\{\mathcal{H}^i_\lambda\}_{\lambda \in \mathbb{Z}^+}$ for each $i \in [4]$. To obtain a sample from $\mathcal{H}^i_\lambda$:

1. Sample $(G, g) \leftarrow \mathcal{G}(1^\lambda)$. Let p denote the order of G.
2. Sample a vector $\vec{u}$ uniformly at random from $\mathbb{Z}_p^n$, and sample $\pi \leftarrow S_n$ where $n = \lambda^2$. Let $\vec{c}$ denote $\pi(\vec{u})$.
3. Choose a subspace $V \subseteq Z_p^n$ in a way that depends on i (described in Table 1 below).
4. Sample $\vec{x}_1, \ldots, \vec{x}_q$ independently and uniformly from V.
5. Output $(G, g, \vec{c}, g^{\vec{x}_1}, \ldots, g^{\vec{x}_q})$.

Table 1: The subspaces from which the group exponents $(\vec{x}_j)_{j \in [q]}$ are sampled in each hybrid distribution.

Hybrid	V
$\mathcal{H}^1_\lambda$	$\{\vec{x} : \vec{x} \cdot \vec{u} = 0\}$
$\mathcal{H}^2_\lambda$	$\{\alpha \vec{v} : \alpha \in \mathbb{Z}_p\}$ for random $\vec{v} \in \mathbb{Z}_p^n$ s.t. $\vec{u} \cdot \vec{v} = 0$
$\mathcal{H}^3_\lambda$	$\{\alpha \vec{v} : \alpha \in \mathbb{Z}_p\}$ for random $\vec{v} \in \mathbb{Z}_p^n$
$\mathcal{H}^4_\lambda$	$\mathbb{Z}_p^n$

We will show in Claims 5.13 and 5.14 below that $\mathcal{H}^1_\lambda$ is identically distributed to the view of an adversary in the simplified distinguishing game of Definition 3.12 when $b = 0$, and similarly $\mathcal{H}^4_\lambda$ is identically distributed to the view when $b = 1$. In Claim 5.15, we show that if the DDH assumption on $\mathcal{G}$ holds, then the ensembles $\{\mathcal{H}^1_\lambda\}$ and $\{\mathcal{H}^2_\lambda\}$ are computationally indistinguishable, and so too are $\{\mathcal{H}^3_\lambda\}$ and $\{\mathcal{H}^4_\lambda\}$. Finally, in Claim 5.16, we show that $\{\mathcal{H}^2_\lambda\}$ and $\{\mathcal{H}^3_\lambda\}$ are statistically close.

Claim 5.13. *For any $\lambda \in \mathbb{N}$, $\mathcal{H}^1_\lambda$ is identically distributed to the view of an adversary in the simplified distinguishing game of Definition 3.12 when $b = 0$.*

Proof. The distributions that we need to show equivalent are those of $(G, g, \vec{u}, g^{\pi(\vec{x}_1)}, \ldots, g^{\pi(\vec{x}_q)})$ and $\left(G, g, \pi(\vec{u}), g^{\vec{x}_1}, \ldots, g^{\vec{x}_q}\right)$, when sampling

$$\begin{aligned} &(G, g) \leftarrow \mathcal{G}(1^\lambda) \\ &\vec{u} \leftarrow \mathbb{Z}_p^n, \text{ where } p = |G| \\ &\pi \leftarrow S_n \\ &\text{For } i = 1, \ldots, q\text{:} \\ &\quad \vec{x}_i \leftarrow \left\{\vec{x} : \vec{x} \cdot \vec{u} = 0\right\}. \end{aligned} \tag{3}$$

This follows from the following observations.

- Conditioned on any G and g, the distributions of $(\vec{u}, \pi)$ and of $\big(\pi(\vec{u}), \pi\big)$ are identical. This follows from the fact that $\vec{u}$ is uniformly random, so for *any* permutation π, the distribution of $\pi(\vec{u})$ is uniform on $\mathbb{Z}_p^n$, even conditioned on π.
- Consequently:
  - Conditioned on any $(G, g, \vec{u})$, the distribution of $\big(\pi(\vec{x}_1), \ldots, \pi(\vec{x}_q)\big)$ can be sampled as follows:

$$\begin{array}{l}\text{Sample } \pi \leftarrow S_n \text{ (this is identical to the distribution of } \pi \text{ in (3) conditioned on } (G, g, \vec{u}))\\ \text{For } i = 1, \ldots, q\text{:}\\ \quad \text{Sample } \vec{x}_i \leftarrow \big\{\vec{x} : \vec{x} \cdot \vec{u} = 0\big\}.\\ \text{Output } \big(\pi(\vec{x}_1), \ldots, \pi(\vec{x}_q)\big).\end{array} \tag{4}$$

  - Conditioned on any $\big(G, g, \pi(\vec{u})\big)$, the distribution of $\vec{x}_1, \ldots, \vec{x}_q$ can be sampled as follows:

$$\begin{array}{l}\text{Sample } \bar{\pi} \leftarrow S_n \text{ (this is identical to the distribution of } \pi^{-1} \text{ in (3) conditioned on } \big(G, g, \pi(\vec{u})\big))\\ \text{For } i = 1, \ldots, q\text{:}\\ \quad \text{Sample } \vec{x}_i \leftarrow \big\{\vec{x} : \vec{x} \cdot \bar{\pi}\big(\pi(\vec{u})\big) = 0\big\}.\\ \text{Output } (\vec{x}_1, \ldots, \vec{x}_q).\end{array} \tag{5}$$

- The inner product operation is permutation-symmetric, i.e. for any permutation $\pi \in S_n$ and any vectors $\vec{u}$ and $\vec{v}$, it holds that $\vec{u} \cdot \vec{v} = \pi(\vec{u}) \cdot \pi(\vec{v})$. Applying this to (5) above (with $\bar{\pi}^{-1}$, $\vec{x}$, and $\bar{\pi}(\pi(\vec{u}))$, we obtain that the sampling procedure from (5) is equivalent to the following procedure:

$$\begin{array}{l}\text{Sample } \bar{\pi} \leftarrow S_n\\ \text{For } i = 1, \ldots, q\text{:}\\ \quad \text{Sample } \vec{x}_i \leftarrow \big\{\vec{x} : \bar{\pi}^{-1}(\vec{x}) \cdot \pi(\vec{u}) = 0\big\}.\\ \text{Output } (\vec{x}_1, \ldots, \vec{x}_q).\end{array} \tag{6}$$

  But this is equivalent to:

$$\begin{array}{l}\text{Sample } \bar{\pi} \leftarrow S_n\\ \text{For } i = 1, \ldots, q\text{:}\\ \quad \text{Sample } \vec{x}_i \leftarrow \big\{\vec{x} : \vec{x} \cdot \pi(\vec{u}) = 0\big\}.\\ \text{Output } \big(\bar{\pi}(\vec{x}_1), \ldots, \bar{\pi}(\vec{x}_q)\big),\end{array} \tag{7}$$

  which is clearly equivalent to the sampling procedure in (4) above (for the vector $\pi(\vec{u})$). □

The following claim can be proved similarly to Claim 5.13.

Claim 5.14. *For any $\lambda \in \mathbb{N}$, $\mathcal{H}_\lambda^4$ is identically distributed to the view of an adversary in the simplified distinguishing game of Definition 3.12 when $b = 1$.*

Claim 5.15. *Assume that the DDH assumption holds in $\mathcal{G}$. Then:*

1. *The distribution ensembles $\{\mathcal{H}_\lambda^1\}$ and $\{\mathcal{H}_\lambda^2\}$ are computationally indistinguishable*
2. *The distribution ensembles $\{\mathcal{H}_\lambda^3\}$ and $\{\mathcal{H}_\lambda^4\}$ are computationally indistinguishable.*

Proof. We show that (1) holds by reduction to the matrix-DDH problem over matrices of dimension $q \times (n-1)$. The argument for (2) is similar.

Suppose for contradiction that there is a non-uniform polynomial time algorithm $\mathcal{A}$ that makes q oracle queries and obtains a non-negligible advantage $\epsilon = \epsilon(\lambda)$ in distinguishing samples from $\mathcal{H}_\lambda^1$ vs. $\mathcal{H}_\lambda^2$, for every λ. We construct a polynomial time adversary $\mathcal{A}_{\mathsf{M\text{-}DDH}}$ such that $\mathsf{Adv}^{\mathsf{M\text{-}DDH}}_{\mathcal{A}_{\mathsf{M\text{-}DDH}}}(\lambda) = \epsilon(\lambda)$, which is a contradiction to the DDH assumption by Imported Theorem 5.10.

$\mathcal{A}_{\mathsf{M\text{-}DDH}}$ on input a security parameter 1^λ, a group G with generator g and order p, and matrix $A \in G^{q \times (n-1)}$ (either from $\mathsf{Rk}_{n-1}\big(G^{q\times(n-1)}\big)$ or $\mathsf{Rk}_1\big(G^{q\times(n-1)}\big)$) operates as follows:

1. Samples $\vec{u} \in \mathbb{Z}_p^n$ exactly as it is chosen in $\mathcal{H}_\lambda^1$, and samples a random permutation $\pi \leftarrow S_\lambda$.
2. Constructs the $q \times n$ matrix A' which is obtained from A by appending to each row $(g^{x_1}, \ldots, g^{x_{n-1}})$ the group element $g^{-u_n^{-1} \cdot \sum_{i=1}^{n-1} x_i u_i}$. (Notice that this gives $\vec{x} = (x_1, \ldots, x_n)$, for $x_n \stackrel{\text{def}}{=} -u_n^{-1} \cdot \sum_{i=1}^{n-1} x_i u_i$, which satisfies $\vec{x} \cdot \vec{u} = 0$.) Let $g^{\vec{x}_i}$ denote the i^{th} row of A'.
3. Invokes $\mathcal{A}$ on input $\left(G, g, \pi(\vec{u}), g^{\vec{x}_1}, \ldots, g^{\vec{x}_q}\right)$.

$\mathcal{A}_{\mathsf{M\text{-}DDH}}$ runs in polynomial time because $\mathcal{A}$ does, because generating A' from A only requires a polynomial number of exponentiations and multiplications in G as well as computing the multiplicative inverse of a polynomial number of field elements. Moreover, if $A \in \mathsf{Rk}_{n-1}\left(G^{q \times (n-1)}\right)$ then the samples provided to $\mathcal{A}$ are distributed as in $\mathcal{H}_\lambda^1$, whereas if $A \in \mathsf{Rk}_1\left(G^{q \times (n-1)}\right)$ then the samples are distributed as in $\mathcal{H}_\lambda^2$, so $\mathsf{Adv}^{\mathsf{M\text{-}DDH}}_{\mathcal{A}_{\mathsf{M\text{-}DDH}}}(\lambda) = \epsilon(\lambda)$. □

Claim 5.16. *The distribution ensembles $\{\mathcal{H}_\lambda^2\}$ and $\{\mathcal{H}_\lambda^3\}$ are statistically $e^{-\Omega(\lambda)}$-close.*

Proof. Recall the differences between $\mathcal{H}_\lambda^2$ and $\mathcal{H}_\lambda^3$. They are defined as the output of the following procedure (with variations in step 3 depending on whether we are in $\mathcal{H}_\lambda^2$ or $\mathcal{H}_\lambda^3$):

1. Sample $(G, g) \leftarrow \mathcal{G}(1^\lambda)$. Let p denote the order of G.
2. Sample a vector $\vec{u}$ uniformly at random from $\mathbb{Z}_p^n$, and sample $\pi \leftarrow S_n$ where $n = \lambda^2$.
3. In $\mathcal{H}_\lambda^2$, let $\vec{v}$ be uniformly random in $\mathbb{Z}_p^n$ conditioned on $\vec{u} \cdot \vec{v} = 0$.
   In $\mathcal{H}_\lambda^3$, let $\vec{v}$ be uniformly random in $\mathbb{Z}_p^n$.
4. Sample $\alpha_1, \ldots, \alpha_q$ independently from $\mathbb{Z}_p$.
5. Output $(G, g, \pi(\vec{u}), g^{\alpha_1 \vec{v}}, \ldots, g^{\alpha_q \vec{v}})$.

Conditioned on any (G, g) with $|G| = p$, it clearly holds in $\mathcal{H}_\lambda^3$ that the distribution of $\left(\pi(\vec{u}), \vec{v}\right)$ is uniform on $\mathbb{Z}_p^n \times \mathbb{Z}_p^n$. We will show that in $\mathcal{H}_\lambda^2$, the distribution of $\left(\pi(\vec{u}), \vec{v}\right)$ is statistically close to uniform on $\mathbb{Z}_p^n \times \mathbb{Z}_p^n$. This will imply Claim 5.16 because the distribution of $(G, g, \pi(\vec{u}), g^{\alpha_1 \vec{v}}, \ldots, g^{\alpha_q \vec{v}})$ conditioned on any $(G, g, \pi(\vec{u}), \vec{v})$ is the same in both $\mathcal{H}_\lambda^2$ and $\mathcal{H}_\lambda^3$.

To see the statistical closeness of $\left(\pi(\vec{u}), \vec{v}\right)$ to uniform, we will consider the following equivalent rejection sampling procedure for sampling $\left(\pi(\vec{u}), \vec{v}\right)$:

1. A candidate $\vec{u}'$ for $\pi(\vec{u})$ is sampled uniformly at random from $\mathbb{Z}_p^n$.
2. A candidate $\vec{v}$ is sampled uniformly at random from $\mathbb{Z}_p^n$.
3. A candidate permutation π is sampled uniformly at random from S_n.
4. If $\pi^{-1}(\vec{u}') \cdot \vec{v} = 0$, then we output $(\vec{u}', \vec{v})$. Otherwise, we start over from step 1.

Claim 5.17. *With all but $e^{-\Omega(\lambda^2)}$ probability over the choice of $\vec{u}'$, the distribution of $(\vec{v}, \pi^{-1}(\vec{u}') \cdot \vec{v})$ conditioned on $\vec{u}'$ is $e^{-\Omega(\lambda^2)}$-close to uniform over $\mathbb{Z}_p^n \times \mathbb{Z}_p$.*

Proof. We show that the distributions are close when $\pi(\vec{u}')$ has high min entropy even conditioned on $\vec{u}'$, and show that the min entropy is high with overwhelming probability.

We claim first that for almost all $\vec{u}'$, the distribution of $\pi(\vec{u}')$ conditioned on $\vec{u}'$ has high min-entropy. Indeed, with all but $e^{-\Omega(\lambda^2)}$ probability, $u'_{2i} \neq u'_{2i+1}$ for at least $n/3$ values of $i \in [n/2]$, in which case the min entropy of $\pi(\vec{u}')$ is at least $n/3$ even conditioned on $\vec{u}'$ (this min-entropy lower bound is easiest to see when conditioning on the unordered sets $\{\pi(2i), \pi(2i+1)\}$ for all such i).

View $\vec{v}$ as a seed for the pairwise-independent hash family $\{h_{\vec{v}}\}$, where $h_{\vec{v}}(\vec{u}) = \vec{u} \cdot \vec{v}$. (That is, for every $\vec{u}' \neq \vec{u}$, $\Pr_{\vec{v} \leftarrow \mathbb{Z}_p^n}\left[h_{\vec{v}}(\vec{u}') = h_{\vec{v}}(\vec{u})\right] = 1/p$.) The leftover hash lemma (see Imported Theorem 5.18 below) states that for these $\vec{u}'$, the distribution of $(\vec{v}, \pi^{-1}(\vec{u}') \cdot \vec{v})$ conditioned on $\vec{u}'$ is ϵ-close to uniform for $\epsilon = 2^{\Omega(\log p - n)} \leq e^{-\Omega(\lambda^2)}$, which implies the claim. □

Imported Theorem 5.18 (Leftover hash lemma). *Let $\mathcal{H} = \left\{h_\alpha : \mathbb{Z}_p^n \to \mathbb{Z}_p\right\}_\alpha$ be a family of pairwise independent hash function (i.e., for every $x, x' \in \mathbb{Z}_p^n$, $\Pr_{h_\alpha \leftarrow \mathcal{H}}\left[h_\alpha(x) = h_\alpha(x')\right] = 1/p$). Let X be distributed over $\mathbb{Z}_p^n$ with min entropy k. Then for a uniformly random α, $d_{\mathsf{TV}}\left((\alpha, h_\alpha(X)), (\alpha, U)\right) \leq 2^{-\Omega(k - \log p)}$, where U is the uniform distribution over $\mathbb{Z}_p$.*

Claim 5.19. *In the rejection sampling procedure, the number of trials is at most $2p\lambda^2$ with all but $e^{-\Omega(\lambda^2)}$ probability.*

Proof. Each trial is the last one with probability at least $\frac{1}{p} - e^{-\Omega(\lambda^2)} \geq \frac{1}{2p}$. Thus, the probability that there are more than $2p\lambda^2$ trials is bounded by $\left(1 - \frac{1}{2p}\right)^{2p\lambda^2} \leq e^{-\lambda^2}$. □

Thus, the rejection sampling procedure described is statistically $2p\lambda^2 \cdot e^{-\Omega(\lambda^2)}$-close to the following:

1. A candidate $\vec{u}'$ for $\pi(\vec{u})$ is sampled uniformly at random from $\mathbb{Z}_p^n$.
2. A candidate $\vec{v}$ is sampled uniformly at random from $\mathbb{Z}_p^n$.
3. A candidate permutation π is sampled uniformly at random from S_n.
4. With probability $\frac{1}{p}$, output $(\vec{u}', \vec{v})$. Otherwise, we start over from step 1, up to $2p\lambda^2$ times in total. If the repetition limit is reached, output $\perp$.

But clearly the output of this procedure is $e^{-\Omega(\lambda^2)}$-close to uniform on $\mathbb{Z}_p^n \times \mathbb{Z}_p$. Claim 5.16 follows. □

Proposition 5.12 now follows from the definition of the hybrids, Lemma 3.16, and Claims 5.13 to 5.16 using a standard hybrid argument. □

6 Statistical Query Lower Bound

In this section we discuss a specific permuted puzzle toy problem introduced by [BIPW17], and study its hardness against a large class of potential adversarial algorithms called *statistical-query algorithms*. We first define this class of algorithms in Section 6.1, then present the toy problem in Section 6.2 and prove it is secure against such algorithms. We prove some useful relevant lemmas in Appendix A.

6.1 Statistical Query Algorithms

Definition 6.1 (Statistical Query Algorithms). Let $\mathcal{P} = (\mathcal{K}, \{\Pi_k\}_{k \in \mathcal{K}})$ be a puzzle problem. A statistical q-query algorithm for $\mathcal{G}_{\mathsf{dist},s}[\mathcal{P}]$ is a stateful adversary $\mathcal{A}$ using an "inner adversary" $\mathcal{A}_{\mathsf{SQ}}$ as follows.

1. Upon receiving the public key pk, $\mathcal{A}$ forwards it to $\mathcal{A}_{\mathsf{SQ}}$.
   Recall that pk is part of the key k, and denote $\Pi_k = (n, \Sigma, \mathcal{D}_0, \mathcal{D}_1)$.
2. The following is repeated q times:
   (a) $\mathcal{A}_{\mathsf{SQ}}$ outputs a boolean-valued function f.[5]
   (b) $\mathcal{A}$ requests a sample $x \leftarrow \mathcal{D}_b$ from the challenger (where $b \in \{0,1\}$ is the challenger's secret bit), computes $f(x)$ (this is a single bit), and forwards $f(x)$ to $\mathcal{A}_{\mathsf{SQ}}$.
3. When $\mathcal{A}_{\mathsf{SQ}}$ outputs a "guess" bit b', $\mathcal{A}$ forwards b' to the challenger.

Remark 6.2. We consider only statistical query algorithms for the simplified distinguishing game $\mathcal{G}_{\mathsf{dist},s}$ of Definition 3.12 because our lower bounds (proven in Section 6.2) hold for puzzle problems in which weak computational hardness (i.e., hardness of $\mathcal{G}_{\mathsf{dist},s}$) is equivalent to computational hardness (i.e., hardness of the more standard distinguishing game $\mathcal{G}_{\mathsf{dist}}$ of Definition 3.7) by Lemma 3.16.

[5] We do not assume any bound on the description size or complexity of f, which will not matter for our lower bounds.

Statistical Query (SQ) algorithms constitute a broad class of distinguishing algorithms, that is incomparable in power to polynomial-time algorithms. For example, an SQ algorithm can distinguish between a PRG output and a uniformly random string with a single query. On the other hand, SQ algorithms cannot distinguish between a distribution that is uniform on $\{0,1\}^n$ and one that is uniform on a random high-dimensional subspace of $\{0,1\}^n$. These distributions can be distinguished (given many samples) in polynomial time by a simple rank computation.

Still, in the context of distinguishing problems, SQ algorithms seem to be a powerful class of adversarial algorithms. In fact, except for the aforementioned examples of algorithms which exploit algebraic structure, we are not aware of any natural distinguishing algorithms that cannot be simulated by statistical query algorithms. A challenging and important open problem, which we leave for future work, is to formalize a class of algorithms that use algebraic structure (or even only linear algebra), possibly together with statistical queries, and to prove lower bounds against this class.

6.2 The Toy Problem and Lower Bound

The works [CHR17, BIPW17] base the security of their DE-PIR schemes on the PermRM conjecture, for which they also discuss different variants (e.g., noisy versions). Boyle et al. [BIPW17] also put forth a toy version of the problem, for which we will prove a lower bound against SQ algorithms. We first recall the PermRM conjecture and its toy version.

Conjecture 1 (PermRM, Conjecture 4.2 in [BIPW17])**.** *Let $m \in \mathbb{N}$ be a dimension parameter, let $\lambda \in \mathbb{N}$ be a security parameter, let $d = d_m(n)$ be the minimal integer such that $n \geq \binom{m+d}{d}$, and let $\mathbb{F}$ be a finite field satisfying $|\mathbb{F}| > d\lambda + 1$. Define a probabilistic algorithm $\mathsf{Samp}(b, \pi, v)$ that operates as follows:*

- *If $b = 0$:*
  1. *Select m random degree-λ polynomial $p_1, \ldots, p_m \leftarrow \mathbb{F}[X]$ such that for every $1 \leq i \leq \lambda$, $p_i(0) = v$. Notice that these polynomials determine a curve $\gamma(t)$ in $\mathbb{F}^m$, given by $\{(p_1(t), \ldots, p_m(t)) \ : \ t \in \mathbb{F}\}$.*
  2. *Sample $d\lambda + 1$ distinct points on the curve $\gamma(t)$, determined by non-zero parameters $t_0, \ldots, t_{d\lambda} \leftarrow \mathbb{F}$.*
  3. *Output the points, in order, where each point is permuted according to $\pi : \mathbb{F}^m \to \mathbb{F}^m$, namely output*
$$\left(\pi\left(p_1(t_i), \ldots, p_m(t_i)\right)\right)_{i=0}^{d\lambda} \in \left(\mathbb{F}^m\right)^{d\lambda+1}.$$
- *If $b = 1$: sample $d\lambda + 1$ random points in $\mathbb{F}^m$ $(w_0, \ldots, w_{d\lambda}) \leftarrow \left(\mathbb{F}^m\right)^{d\lambda+1}$, and output $(w_0, \ldots, w_{d\lambda})$.*

The PermRM conjecture is that for every efficient non-uniform $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ there exists a negligible function $\mu(\lambda) = \mathrm{negl}(\lambda)$ such that:

$$\Pr\left[\begin{array}{c} \left(1^n, 1^{|\mathbb{F}|}, \mathsf{aux}\right) \leftarrow \mathcal{A}_1\left(1^\lambda\right) \\ \pi \leftarrow S_{(\mathbb{F}^m)}; b \leftarrow \{0,1\} \\ b' \leftarrow \mathcal{A}_2^{\mathsf{Samp}(b,\pi,\cdot)}\left(1^n, \mathsf{aux}\right) \end{array} \ : \ b' = b\right] \leq 1/2 + \mu(\lambda)$$

Let $\mathbb{F} = \{\mathbb{F}_\lambda\}_{\lambda \in \mathbb{Z}^+}$ denote an ensemble of finite fields with $|\mathbb{F}_\lambda| = \Theta(\lambda^2)$. Let $q = q_\lambda$ denote $|\mathbb{F}_\lambda|$.
For a function $f : X \to Y$, we define $\mathsf{Graph}(f) : X \times Y \to \{0,1\}$ such that

$$\mathsf{Graph}(f)(x,y) = \begin{cases} 1 & \text{if } y = f(x) \\ 0 & \text{otherwise.} \end{cases}$$

Define the puzzle problem $\Pi_\lambda = (n, \{0,1\}, \mathcal{D}_0, \mathcal{D}_1)$, where $n = q^2$, and $\mathcal{D}_0$ and $\mathcal{D}_1$ are defined as follows.

- A sample from $\mathcal{D}_0$ is $\mathsf{Graph}(\gamma)$, where $\gamma : \mathbb{F} \to \mathbb{F}$ is a uniformly random degree-λ polynomial.
- A sample from $\mathcal{D}_1$ is $\mathsf{Graph}(U)$, where $U : \mathbb{F} \to \mathbb{F}$ is a uniformly random function.

Conjecture 2 ([BIPW17]). *The permuted puzzle problem* $\mathcal{P} \stackrel{\text{def}}{=} \mathsf{Perm}(\{\Pi_\lambda\}_{\lambda\in\mathbb{Z}^+})$ *is computationally hard.*

Theorem 6.3. *The simplified distinguishing game* $\mathcal{G}_{\mathsf{dist},s}[\mathcal{P}]$ *is hard for statistical-query algorithms. That is, for all polynomially bounded* $q(\cdot)$*, the advantage of any statistical* $q(\lambda)$*-query adversary in* $\mathcal{G}_{\mathsf{dist},s}[\mathcal{P}]$ *is at most* $e^{-\Omega(\lambda)}$*.*

Proof. We will show that even if we give the statistical query adversary additional information about π, it cannot distinguish permuted samples from $\mathcal{D}_0$ from permuted samples from $\mathcal{D}_1$. Specifically, we will give the adversary (for free) the unordered partition $\Phi_1 \cup \cdots \cup \Phi_q$ of $\mathbb{F} \times \mathbb{F}$, where $\Phi_i = \pi(\{i\} \times \mathbb{F})$. (Intuitively, Φ_i is the image under π of all points in which the X coordinate equals i. In particular, $\pi(\mathsf{Graph}(f))$ takes value "1" at exactly one coordinate in Φ_i.) Note that it is indeed possible for a statistical query adversary to learn $\Phi \stackrel{\text{def}}{=} \{\Phi_1, \ldots, \Phi_q\}$: if (x, y) and (x', y') belong to the same Φ_i, then for a random sample $z \leftarrow \mathcal{D}_b$, it is never the case that $\pi(z)_{(x,y)} = \pi(z)_{(x',y')} = 1$. However, if (x, y) and (x', y') do *not* belong to the same Φ_i, then $\pi(z)_{(x,y)} = \pi(z)_{(x',y')} = 1$ with probability at least $\frac{1}{q^2}$.

We say that a permutation π $\mathsf{respects}$ a partition $\Phi = \{\Phi_1, \ldots, \Phi_q\}$ if $\{\pi(\{i\}\times\mathbb{F})\}_i = \Phi$. For any partition Φ, we will write $\Pr_\Phi$ to denote the probability space in which a permutation π is sampled uniformly at random from the set of permutations that respect Φ. Similarly, we will write $\mathbb{E}_\Phi$ to denote expectations in $\Pr_\Phi$, and we write Var_Φ to denote variances in $\Pr_\Phi$.

We will show that there is some negligible function $\nu : \mathbb{Z}^+ \to \mathbb{R}$ such that for any function $f : \{0,1\}^n \to \{0,1\}$ and any partition Φ, there exists some $p_{f,\Phi} \in [0,1]$ such that for every $b \in \{0,1\}$, it holds that

$$\Pr_\Phi\left[\mathop{\mathbb{E}}_{x\leftarrow\mathcal{D}_b}\left[f(\pi(x))\right] - p_{f,\Phi}\right| \geq \nu(\lambda)\Bigg] \leq \nu(\lambda).$$

Crucially, $p_{f,\Phi}$ is independent of the challenge bit b, the specific sample x, and the secret permutation π (except for its dependence on Φ). Thus, the answer to a query f can be simulated by computing $p_{f,\Phi}$.

The following two observations are at the core of our proof. Recall that Δ denotes the Hamming distance. For a pair of functions $g, g' : X \to Y$, we denote $\Delta(g, g') = |\{x \in X \ : \ g(x) \neq g'(x)\}|$.

Claim 6.4. *For any partition* Φ*, any function* $g : \mathbb{F} \to \mathbb{F}$*, and any fixed permutation* π^* *that respects* Φ*, the distribution of* $\pi(\mathsf{Graph}(g))$ *under* $\Pr_\Phi$ *is identical to the distribution of* $\pi^*(\mathsf{Graph}(u))$ *when* $u : \mathbb{F} \to \mathbb{F}$ *is a uniformly random function.*

Proof. To sample a random permutation π conditioned on $\{\pi(\{i\} \times \mathbb{F})\}_i = \Phi \stackrel{\text{def}}{=} \{\Phi_1, \ldots, \Phi_q\}$, one can sample a uniformly random permutation $\sigma : \mathbb{F} \to \mathbb{F}$ and q independent bijections $\pi_i : \mathbb{F} \to \Phi_{\sigma(i)}$, and then define $\pi(j, k) = \pi_j(k)$.

$\pi(\mathsf{Graph}(g))$ is defined by the set of points $\{\pi(j, g(j))\}_{j\in\mathbb{F}} = \{\pi_j(g(j))\}$. It is clear that sampling g uniformly at random corresponds to independently picking each $g(j)$ at random, which produces an identical distribution of $\pi(\mathsf{Graph}(g))$ as picking the bijections $\{\pi_j\}$ independently and uniformly at random. Thus, $\pi^*(\mathsf{Graph}(u))$ for a fixed π^* which respects the partition Φ, and a random u, is distributed identically to $\pi(\mathsf{Graph}(g))$ for a fixed g and a random π that respects Φ. □

Claim 6.5. *For any partition* Φ*, any functions* $g, g' : \mathbb{F} \to \mathbb{F}$*, and any fixed permutation* π^* *that respects* Φ*, the distribution of* $\big(\pi(\mathsf{Graph}(g)), \pi(\mathsf{Graph}(g'))\big)$ *under* $\Pr_\Phi$ *is identical to the distribution of* $\big(\pi^*(\mathsf{Graph}(u)), \pi^*(\mathsf{Graph}(u'))\big)$*, where* $u, u' : \mathbb{F} \to \mathbb{F}$ *are jointly uniformly random conditioned on* $\Delta(u, u') = \Delta(g, g')$*.*

Proof. We first consider the distribution under $\Pr_\Phi$ of $(x, x') = \big(\pi(\mathsf{Graph}(g)), \pi(\mathsf{Graph}(g'))\big)$, where g and g' are fixed. Because g and g' are functions, both x and x' will consist mostly of zeros, but for each $j \in \mathbb{F}$, they will contain a 1 in exactly one position in Φ_j. Recall from the proof of Claim 6.4 that π can be sampled by sampling a uniformly random permutation $\sigma : \mathbb{F} \to \mathbb{F}$ and q independent bijections $\pi_i : \mathbb{F} \to \Phi_{\sigma(i)}$, and defining $\pi(j, k) = \pi_j(k)$. Therefore, for any $j \in \mathbb{F}$ if $g(j) = g'(j)$ then x and x' will agree on the position within $\Phi_{\sigma(j)}$ at which they contain a 1 entry. Otherwise, they will disagree. Other than that, the

positions are uniformly random within $\Phi_{\sigma(j)}$ because π_j is a random bijection. Moreover, since σ is a random permutation, the set of Φ_i's for which x, x' agree on the 1-entry is a random subset of size $\Delta(g, g')$.

Now consider the distribution of $(y, y') = \left(\pi^*(\mathsf{Graph}(u)), \pi^*(\mathsf{Graph}(u'))\right)$ where π^* is fixed and defined by σ^* and $\{\pi_i^*\}_{i\in\mathbb{F}}$. The same arguments show that for every $j \in \mathbb{F}$, y, y' agree on the positions within $\Phi_{\sigma^*(j)}$ at which they contain a 1 if and only if $u(j) = u'(j)$. Since u, u' are random and independent, the positions in $\Phi_{\sigma^*(j)}$ in which y, y' have a 1 are otherwise random because these positions are $\pi_j^*(u(j))$ and $\pi_j^*(u'(j))$, respectively. Additionally, the Φ_i's for which y, y' agree on the position of the 1 entry is a uniformly random subset of size $\Delta(g, g') = \Delta(u, u')$, because this set is $\{\sigma^*(j) : u(j) = u'(j)\}$, and u, u' are random and independent.

□

Claim 6.6. *If $g_0, g_1 : \mathbb{F} \to \mathbb{F}$ are two independent uniformly random degree-λ polynomials, then $\Delta(g_0, g_1)$ is $e^{-\Omega(\lambda)}$-close to $\Delta(g_0', g_1')$ for uniformly random $g_0', g_1' : \mathbb{F} \to \mathbb{F}$.*

Proof. For $i \in \mathbb{F}$, let X_i (respectively, Y_i) be indicator of the event that $g_0(i) = g_1(i)$ (respectively, $g_0'(i) = g_1'(i)$). Then X_i, Y_i are λ-wise independent with $\mathbb{E}[X_i] = \mathbb{E}[Y_i] = |\mathbb{F}|^{-1}$. The claim now follows from Lemma A.4 (Page 37) for $n = |\mathbb{F}|$. □

Now, we will show that $\mathbb{E}_{x\leftarrow\mathcal{D}_0}[f(\pi(x))]$ and $\mathbb{E}_{x\leftarrow\mathcal{D}_1}[f(\pi(x))]$, viewed as random variables that depend on π, have the same expectation and also have very small (negligible) variance.

Claim 6.7. *For any $f : \{0,1\}^n \to \{0,1\}$ and any partition Φ,*

$$\mathop{\mathbb{E}}_{\Phi}\left[\mathop{\mathbb{E}}_{x\leftarrow\mathcal{D}_0}[f(\pi(x))]\right] = \mathop{\mathbb{E}}_{\Phi}\left[\mathop{\mathbb{E}}_{x\leftarrow\mathcal{D}_1}[f(\pi(x))]\right].$$

Proof. Consider any $f : \{0,1\}^n \to \{0,1\}$ and any partition Φ. By Claim 6.4, there is a distribution $\mathcal{U}$ that is equal to the distribution (in $\Pr_\Phi$) of $\pi(\mathsf{Graph}(g))$ for all functions $g : \mathbb{F} \to \mathbb{F}$. Let μ denote $\mathbb{E}_{x'\leftarrow\mathcal{U}}[f(x')]$. Let P_b denote the probability mass function of $\mathcal{D}_b$. Then for any $b \in \{0,1\}$,

$$\begin{aligned}
\mathop{\mathbb{E}}_{\Phi}\left[\mathop{\mathbb{E}}_{x\leftarrow\mathcal{D}_b}[f(\pi(x))]\right] &= \mathop{\mathbb{E}}_{\Phi}\left[\sum_x P_b(x)\cdot f(\pi(x))\right] \\
&= \sum_x P_b(x)\cdot \mathop{\mathbb{E}}_{\Phi}[f(\pi(x))] \\
&= \sum_x P_b(x)\cdot \mu \\
&= \mu,
\end{aligned}$$

which does not depend on b. □

Now we analyze the variance. Recall that our goal is to show that $\mathrm{Var}_\Phi\left[\mathbb{E}_{x\leftarrow\mathcal{D}_b}[f(\pi(x))]\right]$ is negligible for $b \in \{0,1\}$. Because of Claim 6.6, this follows from the following more general claim.

Claim 6.8. *Let $\mathcal{D}$ be any distribution on functions mapping $\mathbb{F}$ to $\mathbb{F}$. Suppose that when g and g' are sampled independently from $\mathcal{D}$ and $u, u' : \mathbb{F} \to \mathbb{F}$ are independent uniformly random functions, the distribution of $\Delta(g, g')$ is statistically ϵ-close to that of $\Delta(u, u')$.*

Then, for any $f : \{0,1\}^n \to \{0,1\}$, any partition Φ,

$$\mathop{\mathrm{Var}}_{\Phi}\left[\mathop{\mathbb{E}}_{g\leftarrow\mathcal{D}}\left[f\left(\pi(\mathsf{Graph}(g))\right)\right]\right] \le \epsilon.$$

Proof. Let P denote the probability mass function of $\mathcal{D}$, and let π^* be an arbitrary permutation in S_n such that $\{\pi^*(\{i\}\times\mathbb{F})\}_i = \Phi$. By the definition of variance,

$$\mathop{\mathrm{Var}}_{\Phi}\left[\mathop{\mathbb{E}}_{g\leftarrow\mathcal{D}}[f(\pi(\mathsf{Graph}(g)))]\right] = \mathop{\mathbb{E}}_{\Phi}\left[\mathop{\mathbb{E}}_{g\leftarrow\mathcal{D}}[f(\pi(\mathsf{Graph}(g)))]^2\right] - \mathop{\mathbb{E}}_{\Phi}\left[\mathop{\mathbb{E}}_{g\leftarrow\mathcal{D}}[f(\pi(\mathsf{Graph}(g)))]\right]^2.$$

For the first term, we have

$$\begin{aligned}\mathop{\mathbb{E}}_{\Phi}[\mathop{\mathbb{E}}_{g\leftarrow\mathcal{D}}[f(\pi(\mathsf{Graph}(g)))]^2] &= \mathop{\mathbb{E}}_{\Phi}\left[\left(\sum_g P(g)\cdot f(\pi(\mathsf{Graph}(g)))\right)^2\right] \\ &= \sum_{g,h} P(g)\cdot P(h)\cdot \mathop{\mathbb{E}}_{\Phi}[f(\pi(\mathsf{Graph}(g)))\cdot f(\pi(\mathsf{Graph}(h)))] && \text{(Claim 6.5)} \\ &= \mathop{\mathbb{E}}_{g,h\leftarrow\mathcal{D}}\left[\mathop{\mathbb{E}}_{\substack{u,v:\mathbb{F}\to\mathbb{F}\\ \Delta(u,v)=\Delta(g,h)}}\left[f(\pi^*(\mathsf{Graph}(u)))\cdot f(\pi^*(\mathsf{Graph}(v)))\right]\right].\end{aligned}$$

For the second term, we have

$$\begin{aligned}&\mathop{\mathbb{E}}_{\Phi}\left[\mathop{\mathbb{E}}_{g\leftarrow\mathcal{D}}[f(\pi(\mathsf{Graph}(g)))]\right]^2 \\ &= \left(\sum_g P(g)\cdot \mathop{\mathbb{E}}_{\Phi}\left[f(\pi(\mathsf{Graph}(g)))\right]\right)^2 \\ &= \left(\sum_g P(g)\cdot \mathop{\mathbb{E}}_{u:\mathbb{F}\to\mathbb{F}}\left[f(\pi^*(\mathsf{Graph}(u)))\right]\right)^2 && \text{(Claim 6.4)} \\ &= \mathop{\mathbb{E}}_{u:\mathbb{F}\to\mathbb{F}}\left[f(\pi^*(\mathsf{Graph}(u)))\right]^2 \\ &= \mathop{\mathbb{E}}_{u,v:\mathbb{F}\to\mathbb{F}}\left[f(\pi^*(\mathsf{Graph}(u)))\cdot f(\pi^*(\mathsf{Graph}(v)))\right] \\ &= \mathop{\mathbb{E}}_{g,h:\mathbb{F}\to\mathbb{F}}\left[\mathop{\mathbb{E}}_{\substack{u,v:\mathbb{F}\to\mathbb{F}\\ \Delta(u,v)=\Delta(g,h)}}\left[f(\pi^*(\mathsf{Graph}(u)))\cdot f(\pi^*(\mathsf{Graph}(v)))\right]\right] && \text{(law of total expectation).}\end{aligned}$$

The difference between these two expressions is only in the distribution of g and h over which the (outer) expectation is taken. Furthermore, the value whose expectation is computed lies in $[0,1]$ and depends only on the Hamming distance between g and h. The claim follows. □

Theorem 6.3 follows from Claims 6.6 to 6.8 and Chebyshev's inquality. □

Acknowledgments

We thank Yuval Ishai for many useful discussions. We thank Fermi Ma for helpful discussions, in particular for pointing out that the blueprint of the DDH-based permuted puzzle extends also to the LPN setting, for simplifying our proof of Claim 5.17 (which we had previously proved using Fourier analysis), and for allowing us to include these observations in the current work. We thank the anonymous TCC reviewers for helpful comments.

This work was supported in part by ISF grant 1861/16, AFOSR Award FA9550-17-1-0069, and the Simons Collaboration on Algorithms and Geometry and National Science Foundation grant No. CCF-1714779.

References

[Ajt96] Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing, Philadelphia, Pennsylvania, USA, May 22-24, 1996*, pages 99–108, 1996.

[Ale03] Michael Alekhnovich. More on average case vs approximation complexity. In *44th Symposium on Foundations of Computer Science (FOCS 2003), 11-14 October 2003, Cambridge, MA, USA, Proceedings*, pages 298–307, 2003.

[BCC+17] Nir Bitansky, Ran Canetti, Alessandro Chiesa, Shafi Goldwasser, Huijia Lin, Aviad Rubinstein, and Eran Tromer. The hunting of the SNARK. *J. Cryptology*, 30(4):989–1066, 2017.

[BFKL93] Avrim Blum, Merrick L. Furst, Michael J. Kearns, and Richard J. Lipton. Cryptographic primitives based on hard learning problems. In *Advances in Cryptology - CRYPTO '93, 13th Annual International Cryptology Conference, Santa Barbara, California, USA, August 22-26, 1993, Proceedings*, pages 278–291, 1993.

[BGI+01] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. In *Advances in Cryptology - CRYPTO 2001, 21st Annual International Cryptology Conference, Santa Barbara, California, USA, August 19-23, 2001, Proceedings*, pages 1–18, 2001.

[BHHO08] Dan Boneh, Shai Halevi, Michael Hamburg, and Rafail Ostrovsky. Circular-secure encryption from decision Diffie-Hellman. In *Advances in Cryptology - CRYPTO 2008, 28th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2008. Proceedings*, pages 108–125, 2008.

[BIM00] Amos Beimel, Yuval Ishai, and Tal Malkin. Reducing the servers computation in private information retrieval: PIR with preprocessing. In *Advances in Cryptology - CRYPTO 2000, 20th Annual International Cryptology Conference, Santa Barbara, California, USA, August 20-24, 2000, Proceedings*, pages 55–73, 2000.

[BIPW17] Elette Boyle, Yuval Ishai, Rafael Pass, and Mary Wootters. Can we access a database both locally and privately? In *TCC (2)*, volume 10678 of *Lecture Notes in Computer Science*, pages 662–693. Springer, 2017.

[BKW00] Avrim Blum, Adam Kalai, and Hal Wasserman. Noise-tolerant learning, the parity problem, and the statistical query model. In *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing, May 21-23, 2000, Portland, OR, USA*, pages 435–440, 2000.

[BPR15] Nir Bitansky, Omer Paneth, and Alon Rosen. On the cryptographic hardness of finding a Nash equilibrium. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 1480–1498, 2015.

[BV11] Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. *ECCC 2011*, 18:109, 2011.

[CCH+19] Ran Canetti, Yilei Chen, Justin Holmgren, Alex Lombardi, Rothblum Guy N, Ron D. Rothblum, and Daniel WIchs. Fiat-Shamir: From practice to theory. In *STOC*, 2019.

[CCRR18] Ran Canetti, Yilei Chen, Leonid Reyzin, and Ron D. Rothblum. Fiat-Shamir and correlation intractability from strong KDM-secure encryption. In *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part I*, pages 91–122, 2018.

[CGKS95] Benny Chor, Oded Goldreich, Eyal Kushilevitz, and Madhu Sudan. Private information retrieval. In *36th Annual Symposium on Foundations of Computer Science, Milwaukee, Wisconsin, USA, 23-25 October 1995*, pages 41–50, 1995.

[CHK+19] Arka Rai Choudhuri, Pavel Hubávcek, Chethan Kamath, Krzysztof Pietrzak, Alon Rosen, and Guy N. Rothblum. Finding a Nash equilibrium is no easier than breaking Fiat-Shamir. *IACR Cryptology ePrint Archive*, 2019:158, 2019.

[CHR17] Ran Canetti, Justin Holmgren, and Silas Richelson. Towards doubly efficient private information retrieval. In *TCC (2)*, volume 10678 of *Lecture Notes in Computer Science*, pages 694–726. Springer, 2017.

[DH76] Whitfield Diffie and Martin E. Hellman. New directions in cryptography. *IEEE Trans. Information Theory*, 22(6):644–654, 1976.

[FKM+18] Jean-Charles Faugère, Eliane Koussa, Gilles Macario-Rat, Jacques Patarin, and Ludovic Perret. PKP-based signature scheme. *IACR Cryptology ePrint Archive*, 2018:714, 2018.

[Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In *STOC 2009, Proceedings*, pages 169–178. ACM, 2009.

[GGH+13] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA*, pages 40–49, 2013.

[GJ02] Michael R Garey and David S Johnson. *Computers and intractability*, volume 29. wh freeman New York, 2002.

[GK16] Shafi Goldwasser and Yael Tauman Kalai. Cryptographic assumptions: A position paper. In *TCC (A1)*, volume 9562 of *Lecture Notes in Computer Science*, pages 505–522. Springer, 2016.

[GKL88] Oded Goldreich, Hugo Krawczyk, and Michael Luby. On the existence of pseudorandom generators (extended abstract). In *29th Annual Symposium on Foundations of Computer Science, White Plains, New York, USA, 24-26 October 1988*, pages 12–24, 1988.

[Gol90] Oded Goldreich. A note on computational indistinguishability. *Inf. Process. Lett.*, 34(6):277–281, 1990.

[GW11] Craig Gentry and Daniel Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In *STOC*, pages 99–108. ACM, 2011.

[HL18] Justin Holmgren and Alex Lombardi. Cryptographic hashing from strong one-way functions (or: One-way product functions and their applications). In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 850–858. IEEE, 2018.

[HY17] Pavel Hubávcek and Eylon Yogev. Hardness of continuous local search: Query complexity and cryptographic lower bounds. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, pages 1352–1371, 2017.

[KMP19] Eliane Koussa, Gilles Macario-Rat, and Jacques Patarin. On the complexity of the permuted kernel problem. *IACR Cryptology ePrint Archive*, 2019:412, 2019.

[KO97] Eyal Kushilevitz and Rafail Ostrovsky. Replication is NOT needed: SINGLE database, computationally-private information retrieval. In *38th Annual Symposium on Foundations of Computer Science, FOCS '97, Miami Beach, Florida, USA, October 19-22, 1997*, pages 364–373, 1997.

[KRR17] Yael Tauman Kalai, Guy N. Rothblum, and Ron D. Rothblum. From obfuscation to the security of Fiat-Shamir for proofs. In *Advances in Cryptology - CRYPTO 2017 - 37th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 20-24, 2017, Proceedings, Part II*, pages 224–251, 2017.

[LL14] Nathan Linial and Zur Luria. Chernoff's inequality - a very elementary proof, 2014.

[LP12] Rodolphe Lampe and Jacques Patarin. Analysis of some natural variants of the PKP algorithm. In *SECRYPT 2012 - Proceedings of the International Conference on Security and Cryptography, Rome, Italy, 24-27 July, 2012, SECRYPT is part of ICETE - The International Joint Conference on e-Business and Telecommunications*, pages 209–214, 2012.

[Mul54] David E. Muller. Application of boolean algebra to switching circuit design and to error detection. *Trans. I.R.E. Prof. Group on Electronic Computers*, 3(3):6–12, 1954.

[Nao03] Moni Naor. On cryptographic assumptions and challenges. In *CRYPTO*, volume 2729 of *Lecture Notes in Computer Science*, pages 96–109. Springer, 2003.

[PS19] Chris Peikert and Sina Shiehian. Noninteractive zero knowledge for NP from (plain) learning with errors. *IACR Cryptology ePrint Archive*, 2019:158, 2019.

[Rab79] Michael O. Rabin. Digitalized signatures and public-key functions as intractable as factorization. *Technical Report, MIT Laboratory for Computer Science*, 1979.

[RAD78] Ronald L Rivest, Len Adleman, and Michael L Dertouzos. On data banks and privacy homomorphisms. *Foundations of secure computation, Academia Press*, 1978.

[Ree54] Irving S. Reed. A class of multiple-error-correcting codes and the decoding scheme. *Trans. of the IRE Professional Group on Information Theory (TIT)*, 4:38–49, 1954.

[Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005*, pages 84–93, 2005.

[RSA78] Ronald L. Rivest, Adi Shamir, and Leonard M. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM*, 21(2):120–126, 1978.

[Sha89] Adi Shamir. An efficient identification scheme based on permuted kernels (extended abstract). In *Advances in Cryptology - CRYPTO '89, 9th Annual International Cryptology Conference, Santa Barbara, California, USA, August 20-24, 1989, Proceedings*, pages 606–609, 1989.

A Useful Lemmas

In this section, we will routinely use the following standard binomial coefficient bounds: For any n, k,

$$\left(\frac{n}{k}\right)^k \leq \binom{n}{k} \leq \left(\frac{en}{k}\right)^k.$$

We will also use the following fact.

Fact A.1. *For all constants $0 < \epsilon < 1$, there exists a constant $\delta > 0$ such that for any n and any $t \leq (1-\epsilon)n$,*

$$\binom{n}{t} \leq (1+\epsilon)^t \cdot \binom{(1-\delta)n}{t}.$$

Proof. For any $0 < \delta < \epsilon$, we have

$$\begin{aligned}
\frac{\binom{n}{t}}{\binom{(1-\delta)n}{t}} &= \prod_{i=0}^{t-1} \frac{n-i}{(1-\delta)n-i} \\
&\leq \left(\frac{n-t}{(1-\delta)n-t}\right)^t \\
&= \left(\frac{1-\frac{t}{n}}{1-\delta-\frac{t}{n}}\right)^t,
\end{aligned}$$

where in the second step we have used the fact that $\frac{n-1}{k-1} \geq \frac{n}{k}$ if $n \geq k > 1$. The last expression is bounded by $(1+\epsilon)^t$ for sufficiently small $\delta > 0$. □

We will rely on the following theorem.

Imported Theorem A.2 ([LL14]). *Let $X_1, \ldots, X_n$ be $\{0,1\}$-valued random variables, let $0 < \beta < 1$, and let $0 < t < \beta n$. Then*

$$\Pr\left[\sum_{i=1}^{n} X_i \geq \beta n\right] \leq \frac{1}{\binom{\beta n}{t}} \cdot \sum_{\substack{A \subseteq [n] \\ |A|=t}} \mathbb{E}\left[\prod_{i \in A} X_i\right].$$

In applying Imported Theorem A.2, we will need to bound quantities of the form $\sum_{\substack{A \subseteq [n] \\ |A|=t}} \prod_{i \in A} p_i$. We will do so with the following claim:

Claim A.3. *For any $p_1, \ldots, p_n \geq 0$, it holds that*

$$\sum_{\substack{A \subseteq [n] \\ |A|=t}} \prod_{i \in A} p_i \leq \frac{(pn)^t}{t!} \leq \left(\frac{epn}{t}\right)^t,$$

where $p \stackrel{\text{def}}{=} \frac{1}{n} \cdot \sum_i p_i$.

Proof. We have

$$\begin{aligned}
\sum_{\substack{A \subseteq [n] \\ |A|=t}} &= \sum_{i_1 < \cdots < i_t \in [n]} \prod_{j \in [t]} p_{i_j} \\
&= \frac{1}{t!} \cdot \sum_{\substack{\text{distinct} \\ i_1, \ldots, i_t \in [n]}} \prod_{j \in [t]} p_{i_j} \\
&\leq \frac{1}{t!} \cdot \sum_{i_1, \ldots, i_t} \prod_{j \in [t]} p_{i_j} && \text{because each } p_i \geq 0 \\
&= \frac{1}{t!} \cdot \left(\sum_j p_j\right)^t \\
&= \frac{(pn)^t}{t!}.
\end{aligned}$$

The second inequality of the claim follows from Stirling's approximation, which states that $\ln(t!) \geq t \ln t - t$, so $\ln(\sqrt[t]{t!}) \geq \ln t - 1$ and thus $t! \geq \left(\frac{t}{e}\right)^t$. □

Lemma A.4. *Let $X = (X_1, \ldots, X_n)$ and $Y = (Y_1, \ldots, Y_n)$ be t-wise independent $\{0,1\}$-valued random variables such that for all $i \in [n]$, $\mathbb{E}[Y_i] = \mathbb{E}[X_i] \stackrel{\text{def}}{=} p_i$, let p denote $\frac{1}{n} \cdot \sum_i p_i$, and suppose that $p < \frac{t}{2en}$. Then the total variation distance $d_{\mathsf{TV}}(X, Y)$ is at most*

$$\left(\frac{2epn}{t}\right)^{t/2} \cdot \left(\frac{n + \frac{t}{t-2epn}}{\prod_{i \in [n]}(1-p_i)} + 2\right)$$

Proof. Without loss of generality, suppose that $Y_1, \ldots, Y_n$ are mutually independent (moving to the general case incurs only a factor of two loss in the obtained bound).

Let $|X|$ denote the number of $i \in [n]$ for which $X_i = 1$. We first establish a high-probability upper bound on X.

Claim A.5.

$$\Pr[|X| \geq \frac{t}{2}] \leq \left(\frac{2epn}{t}\right)^{t/2}.$$

Proof. $|X| \geq \frac{t}{2}$ only if for some $S \subseteq [n]$ with $|S| = \frac{t}{2}$, it holds that $X_S = \vec{1}$. By a union bound and t-wise independence, this holds with probability at most $\sum_{|S|=t/2} \prod_{i \in S} p_i$. The claim then follows from Claim A.3. □

We now bound $\left|\frac{\Pr[X=z]}{\Pr[Y=z]} - 1\right|$. We first consider the case that z is 0^n.

Lemma A.6. *Let $X = (X_1, \ldots, X_n)$ be t-wise independent $\{0,1\}$-valued random variables, and denote $p \stackrel{\text{def}}{=} \frac{1}{n} \cdot \sum_i \mathbb{E}[X_i]$. Suppose that $p < \frac{t}{en}$. Then*

$$\left|\Pr[X = \vec{0}] - \prod_{i=1}^{n} \Pr[X_i = 0]\right| \leq \left(n + \frac{t}{t - epn}\right) \cdot \left(\frac{epn}{t}\right)^t$$

Proof. Let p_j denote $\Pr[X_j = 1]$. Using this notation, we need to bound

$$\left|\Pr[X = \vec{0}] - \prod_{j=1}^{n}(1 - p_j)\right|.$$

By the inclusion-exclusion principle,

$$\Pr[X = \vec{0}] = \sum_{i=0}^{n}(-1)^i \cdot \sum_{\substack{S \subseteq [n] \\ |S|=i}} \Pr[X_S = \vec{1}].$$

Since X is t-wise independent, this is equal to

$$\sum_{i=0}^{t}(-1)^i \cdot \sum_{\substack{S \subseteq [n] \\ |S|=i}} \prod_{j \in S} p_j + \sum_{i=t+1}^{n}(-1)^i \cdot \sum_{\substack{S \subseteq [n] \\ |S|=i}} \Pr[X_S = \vec{1}]$$

$$= \prod_{j=1}^{n}(1 - p_j) - \left(\sum_{\substack{S \subseteq [n] \\ |S| \geq t+1}} (-1)^{|S|} \cdot \prod_{i \in S} p_i\right) + \sum_{i=t+1}^{n}(-1)^i \cdot \sum_{\substack{S \subseteq [n] \\ |S|=i}} \Pr[X_S = \vec{1}].$$

We now bound the two error terms separately. For the first error term,

$$\begin{aligned}
\left|\sum_{\substack{S \subseteq [n] \\ |S| \geq t+1}} (-1)^{|S|} \cdot \prod_{i \in S} p_i\right| &\leq \sum_{i=t+1}^{n} \sum_{\substack{S \subseteq [n] \\ |S|=i}} \prod_{i \in S} p_i \\
&\leq \sum_{i=t+1}^{n} \left(\frac{epn}{i}\right)^i && \text{(By Claim A.3)} \\
&\leq \sum_{i=t}^{n} \left(\frac{epn}{t}\right)^i \\
&\leq \sum_{i=t}^{\infty} \left(\frac{epn}{t}\right)^i \\
&= \frac{t}{t - epn} \cdot \left(\frac{epn}{t}\right)^t && \text{(because } epn < t\text{)}.
\end{aligned}$$

To bound the second term, we first rewrite it (noticing that each outcome where $|X| = w$ is an outcome in which $X_S = \vec{1}$ for exactly $\binom{w}{i}$ choices of $S \subseteq [n]$ with $|S| = i$) as:

$$\sum_{i=t+1}^{n} (-1)^i \cdot \sum_{\substack{S\subseteq[n]\\|S|=i}} \Pr[X_S = \vec{1}] = \sum_{w=t+1}^{n} \Pr[|X| = w] \cdot \sum_{i=t+1}^{w} (-1)^i \cdot \binom{w}{i}.$$

Using Imported Theorem A.2 and t-wise independence,

$$\begin{aligned}\left|\sum_{w=t+1}^{n} \Pr[|X| = w] \cdot \sum_{i=t+1}^{w} (-1)^i \cdot \binom{w}{i}\right| &\leq \sum_{w=t+1}^{n} \left|\Pr[|X| \geq w] \cdot \sum_{i=t+1}^{w} (-1)^i \cdot \binom{w}{i}\right| \\ &\leq \sum_{w=t+1}^{n} \binom{w}{t}^{-1} \cdot \left(\sum_{\substack{A\subseteq[n]\\|A|=t}} \prod_{j\in A} p_j\right) \cdot \left|\sum_{i=t+1}^{w} (-1)^i \cdot \binom{w}{i}\right|. \qquad (8)\end{aligned}$$

By Claim A.3, it holds that $\sum_{\substack{A\subseteq[n]\\|A|=t}} \prod_{j\in A} p_j$ is at most $\left(\frac{epn}{t}\right)^t$.

We prove that

$$\left|\sum_{i=t+1}^{w} (-1)^i \cdot \binom{w}{i}\right| \leq \binom{w}{t}$$

as follows: If $t \geq w/2$, then the summation consists of summands that decrease monotonically in absolutely value, and have alternating signs. Thus the sum is less (in absolute value) than the first term $\binom{w}{t+1}$, which is less than $\binom{w}{t}$. On the other hand, if $t < w/2$ then we can use the identity $0 = (1-1)^w = \sum_{i=0}^{w} (-1)^i \cdot \binom{w}{i}$ to rewrite it as

$$\left|\sum_{i=0}^{t} (-1)^i \cdot \binom{w}{i}\right|,$$

in which case the summation consists of terms that *increase* monotonically in absolute value and alternate in sign; hence the summation is bounded (in absolute value) by the last term $\binom{w}{t}$.

Thus, we can bound Eq. (8) by

$$\sum_{w=t+1}^{n} \left(\frac{epn}{t}\right)^t \leq n \cdot \left(\frac{epn}{t}\right)^t. \qquad \square$$

Finally, we bound the statistical distance between X and Y. Let G denote the set of $z \in \{0,1\}^n$ with $|z| < t/2$.

$$\begin{aligned} d_{\mathsf{TV}}(X,Y) &= \frac{1}{2} \sum_{z\in\{0,1\}^n} \Big|\Pr[X = z] - \Pr[Y = z]\Big| \\ &= \frac{1}{2} \sum_{z\notin G} \Big|\Pr[X = z] - \Pr[Y = z]\Big| + \frac{1}{2} \sum_{z\in G} \Big|\Pr[X = z] - \Pr[Y = z]\Big| \\ &\leq \frac{1}{2} \cdot \big(\Pr[X \notin G] + \Pr[Y \notin G]\big) + \frac{1}{2} \sum_{z\in G} \Big|\Pr[X = z] - \Pr[Y = z]\Big| \end{aligned}$$

Using Claim A.5, this is at most

$$\left(\frac{2epn}{t}\right)^{t/2} + \frac{1}{2}\sum_{z\in G}\left|\Pr[X=z]-\Pr[Y=z]\right|$$
$$= \left(\frac{2epn}{t}\right)^{t/2} + \frac{1}{2}\sum_{z\in G}\Pr[Y=z]\cdot\left|\frac{\Pr[X=z]}{\Pr[Y=z]}-1\right|$$
$$\leq \left(\frac{2epn}{t}\right)^{t/2} + \frac{1}{2}\max_{z\in G}\left|\frac{\Pr[X=z]}{\Pr[Y=z]}-1\right|.$$

Take any $z \in G$, let S denote the set of i's for which $z_i = 1$, and let $\bar{S}$ denote the complement of S. Note that conditioned on $X_S = \vec{1}$, the distribution of $X_{\bar{S}}$ is $t/2$-wise independent, because X is t-wise independent and $|S| \leq t/2$. In particular, for all $i \notin S$, $\Pr\left[X_i = 1 \middle| X_S = \vec{1}\right] = \Pr\left[X_i = 1\right] = p_i$. Therefore,

$$\left|\frac{\Pr[X=z]}{\Pr[Y=z]}-1\right| = \left|1-\frac{\Pr[X_S=\vec{1}]\cdot\Pr\left[X_{\bar{S}}=\vec{0}\middle|X_S=\vec{1}\right]}{\Pr[Y_S=\vec{1}]\cdot\Pr\left[Y_{\bar{S}}=\vec{0}\middle|Y_S=\vec{1}\right]}\right|$$

Using the fact that $\Pr[X_S = \vec{1}] = \Pr[Y_S = \vec{1}]$ (which holds because X, Y are t-wise independent and $\mathbb{E}[Y_i] = \mathbb{E}[X_i]$ for every $i \in [n]$) this is equal to

$$\left|1-\frac{\Pr\left[X_{\bar{S}}=\vec{0}\middle|X_S=\vec{1}\right]}{\Pr\left[Y_{\bar{S}}=\vec{0}\middle|Y_S=\vec{1}\right]}\right|$$

which, since Y is mutually independent, is equal to

$$\left|1-\frac{\Pr\left[X_{\bar{S}}=\vec{0}\middle|X_S=\vec{1}\right]}{\prod_{i\in\bar{S}}\Pr\left[Y_i=0\right]}\right|$$
$$= \left|1-\frac{\Pr\left[X_{\bar{S}}=\vec{0}\middle|X_S=\vec{1}\right]}{\prod_{i\in\bar{S}}(1-p_i)}\right|$$

Using Lemma A.6 for $X_{\bar{S}}\middle|X_S = \vec{1}$ (which is $t/2$-wise independent), this is at most

$$\frac{(n+\frac{t}{t-2epn})\cdot(2epn/t)^{t/2}}{\prod_{i\in\bar{S}}(1-p_i)} \leq \frac{(n+\frac{t}{t-2epn})\cdot(2epn/t)^{t/2}}{\prod_{i\in[n]}(1-p_i)}$$

Thus,

$$d_{\mathsf{TV}}(X,Y) \leq \left(\frac{2epn}{t}\right)^{t/2} + \frac{1}{2}\cdot\frac{(n+\frac{t}{t-2epn})\cdot(2epn/t)^{t/2}}{\prod_{i\in[n]}(1-p_i)}$$
$$= \left(\frac{2epn}{t}\right)^{t/2}\cdot\left(\frac{n+\frac{t}{t-2epn}}{2\cdot\prod_{i\in[n]}(1-p_i)}+1\right).$$

Recall that we assumed $Y_1, \ldots, Y_n$ were mutually independent. To handle general Y, let $U = (U_1, \ldots, U_n)$ denote independent $\{0,1\}$-valued random variables where $\mathbb{E}[U_i] = \mathbb{E}[X_i] (= \mathbb{E}[Y_i])$. Then

$$d_{\mathsf{TV}}(X,Y) \leq d_{\mathsf{TV}}(X,U) + d_{\mathsf{TV}}(U,Y) \leq \left(\frac{2epn}{t}\right)^{t/2}\cdot\left(\frac{n+\frac{t}{t-2epn}}{\prod_{i\in[n]}(1-p_i)}+2\right),$$

which proves the lemma. □