

Indistinguishability Lifting for Keyed Oracles, Compressed Ideal Cipher, and More Applications

Ritam Bhaumik¹  and Yu-Hsuan Huang² 

¹ Technology Innovation Institute (TII), United Arab Emirates

² Max Planck Institute for Security and Privacy (MPI-SP), Germany
ritam.bhaumik@tii.ae, yu-hsuan.huang@mpi-sp.org

Abstract. Cryptographic security proofs often involve an adversary interacting with a larger, *keyed oracle* that consists of (potentially exponentially) many independent instances of a smaller, *base oracle*. Examples include ideal ciphers, which provide access to an independent random permutation for each key, and random oracles, which provide an independent random bit string for each input. However, showing quantum indistinguishability between two such keyed oracles can be tricky, since a single query made by an adversary may involve a superposition that covers all instances of the base oracles simultaneously.

In this paper, we establish a generic *indistinguishability lifting theorem* of the following form: if the two base oracles are indistinguishable under quantum queries, then their corresponding keyed oracles are too, up to an $O(q^2)$ multiplicative loss in distinguishing advantage, where q is the number of queries made by the adversary. Our lifting theorem applies to both statistical and computational settings, and to oracles that are stateful as well. It is also optimal in that it matches the obvious Grover search attack for a certain (contrived) choice of oracles.

As an immediate application, we extend Carolan’s compressed permutation oracle to an efficiently implementable *compressed ideal cipher*, and use it to prove preimage resistance of the Davies-Meyer compression function in the quantum ideal cipher model. Thanks to our lifting theorem, the soundness of our compressed ideal cipher reduces to that of Carolan’s oracle, and any further improvement on the latter would automatically carry over to the former.

As our second application, we give a modular construction that doubles the message length of any quantum-secure strong pseudorandom permutation. Along the way, we show that an existing two-round tweakable Feistel construction is indistinguishable from a random permutation under quantum bidirectional queries. This is done via a dedicated polynomial-method argument, which may be of independent interest.

1 Introduction

Indistinguishability in cryptography. Indistinguishability has played an indispensable role in modern cryptography. From Shannon’s treatment of perfect secrecy [Sha49], through the seminal works of Goldwasser and Micali [GM82, GM84] on computational security of public-key encryption, to Yao’s formulation of pseudorandomness [Yao82], a recurring theme is that information generated from two different sources may nevertheless be indistinguishable to an observer.

Along with the development of cryptography as a field, the use of indistinguishability has further been extended: not only as a property between two pieces of input data, but rather concerning the entire interaction an adversary can have with the cryptographic scheme [GGM86, LR88] at hand. This is often captured in the form of *indistinguishability between oracles*, wherein an adversary is allowed to make a bounded number of queries to one oracle O_0 or another O_1 , and is tasked to distinguish which oracle it is interacting with. This also serves as a useful tool for security analyses. For example, one may perform a sequence of replacements to the oracle at hand, each of which remains indistinguishable, and in the end arrive at the final oracle that is harmless to the considered cryptographic scheme.

Indistinguishability between keyed oracles. In cryptography, it is very common to consider a *keyed oracle*, which consists of (potentially exponentially) many independent instances of a base oracle O_b , and the adversary is allowed to decide which instance of the oracle O_b to interact with in each query, indicated via a key $k \in \mathcal{K}$. For instance, in the random oracle model [BR93], a random oracle consists of a family of independently sampled bit strings, one string per key, while in the ideal cipher model [BRS02], an ideal cipher consists of a family of independently sampled random permutation oracles, one permutation per key.

Naturally, one would hope to argue indistinguishability between two such keyed oracles O_0^K and O_1^K . For this to hold, the indistinguishability of the base oracles O_0 and O_1 is clearly necessary, and establishing it depends on the specific oracles under consideration. This leads to a generic question:

Does indistinguishability of the base oracles O_0 and O_1 imply that of the keyed oracles O_0^K and O_1^K ? If so, at what security loss?

Classically, the answer is yes, and a standard hybrid argument settles the security loss: for each instance of the base oracle O_0 that the adversary has actually queried, we may replace it by the other base oracle O_1 , at an additive security loss ϵ . Having replaced all q instances that are queried by the adversary, we have effectively replaced the entire keyed oracle O_0^K with O_1^K , at a total loss summing up to $q\epsilon$. This is the best one can achieve, since a matching attack exists for a certain (contrived) choice of oracles.³

In the quantum case, however, the situation becomes significantly trickier. Indeed, a quantum adversary may be able to query the keyed oracle in a quantum superposition of keys that, in turn, touches all instances of the base oracle simultaneously, and so the (naive) classical hybrid argument does not work in general. On the other hand, replacing the entire family of instances one by one often causes a security loss that explodes exponentially fast in the length of the keys, which is often unaffordable from the security perspective.

Limitations of prior work. In the literature, there have been several ways to mitigate the aforementioned difficulties, but they all come with a price and limitations, which we discuss below.

The first limitation lies in the obtained distinguishing bound. In the computational setting, Zhandry’s small-range distribution technique [Zha12] (essentially) provides a way to simulate the keyed oracle O_b^K using another smaller keyed oracle $O_b^{[L]}$ (for *stateless* oracles, which we discuss in the next paragraph) whose key space size is much smaller, of size L . An adversary $\mathcal{A}$ distinguishing this smaller keyed oracle $O_b^{[L]}$ can then be transformed into an adversary $\mathcal{B}$ distinguishing the underlying base oracles O_b , via an L -step hybrid argument, which introduces a factor L loss in the distinguishing advantage $\text{Adv}(\mathcal{B})$. However, this small-range simulator is imperfect, with an additive security loss of order $O(q^3/L)$ for adversaries given q quantum queries. To make this additive simulation error negligible, the parameter L then has to be set superpolynomially large, thereby introducing a superpolynomially large multiplicative loss in the final distinguishing bound. In particular, if one optimizes the overall bound via equating $L \cdot \text{Adv}(\mathcal{B})$ and the simulation error $O(q^3/L)$, the achieved bound would be on the order of $q^{1.5} \sqrt{\text{Adv}(\mathcal{B})}$, which is rather non-tight. Further restricting to the statistical setting, where the base oracles O_0 and O_1 are sampled from distributions that are close in a statistical distance ϵ , a dedicated tight query lower bound that matches Grover search can be achieved [Bel19]. However, in terms of concrete advantage, prior work [BDF⁺11, SY17, BM26] only reaches bounds of the form $O(q^c \sqrt{\epsilon})$ (for $c \geq 1$). That is, they all lose a square root in ϵ .

Beyond quantitative loss, the second limitation lies in scope. All existing lifting results concern only *stateless* oracles. Namely, the base oracle O_b is a function (or unitary [AGQY22]) that is sampled up front according to a certain distribution. Many useful oracles are *not* of this kind. For instance, Zhandry’s compressed oracle [Zha19] and the analogous compressed permutation oracles of Unruh [Unr23] and Carolan [Car26] are stateful simulators whose quantum memory evolves with every query. Moreover, their usefulness crucially relies on such memory, since a reduction may use it to embed/observe useful information to/from the adversary interacting with the oracle. Extending, say, a compressed permutation oracle to a compressed *ideal cipher*—which would be useful for analyzing permutation-based constructions—means lifting (quantum) indistinguishability of *stateful* base oracles O_0 and O_1 , to the corresponding keyed oracles O_0^K and O_1^K . To the best of our knowledge, no prior work has established a generic lifting result for this stateful setting.

1.1 Our Contributions

In this paper, we answer the above generic question affirmatively, to the best possible extent. Namely, we show that if (possibly stateful) oracles O_0 and O_1 are indistinguishable in q quantum queries, then so are their corresponding keyed oracles O_0^K and O_1^K , up to a factor $8q^2$ in the distinguishing bound.

³ Concretely, think of the oracle O_b that outputs a pre-selected random bit from $\text{Ber}(b \cdot \epsilon)$. Then O_0 and O_1 are ϵ indistinguishable but O_0^K and O_1^K are $\Omega(q\epsilon)$ distinguishable.

Theorem (Quantum indistinguishability lifting, informal). *Every algorithm distinguishing $\mathcal{O}_0^K$ from $\mathcal{O}_1^K$ with q quantum queries and advantage ϵ can be transformed into a similarly efficient algorithm distinguishing $\mathcal{O}_0$ from $\mathcal{O}_1$ with q quantum queries and advantage at least $\epsilon/(8q^2)$.*

Our lifting theorem holds both statistically and computationally, and side-steps all the above limitations of prior works: One can immediately see that, contrary to prior works, our achieved indistinguishability bound does not have any square-root loss. Not only that, our bound is actually *optimal*: for a certain (contrived) choice of oracles $\mathcal{O}_0$ and $\mathcal{O}_1$ (see Section 3.1), the standard Grover search would have achieved a matching distinguishing advantage. Moreover, our indistinguishability lifting theorem holds for *stateful* oracles, making it applicable to more useful oracles, such as the (family of) compressed permutation oracles proposed by Unruh [Unr23] and Carolan [Car26].

Compressed ideal cipher. As the first and immediate application of our indistinguishability lifting theorem, we show that existing compressed permutation oracle techniques in the literature [Unr23, Car26, HHM⁺26, FLM⁺26] can be extended to the ideal cipher model. We do so via extending the particular oracle cpO considered in [Car26], to its corresponding keyed oracle cpO^K . Thanks to our lifting theorem, we can readily lift the indistinguishability of cpO and a random permutation P , to that of cpO^K and the corresponding ideal cipher. Slightly less obvious though, is whether the new, keyed compressed permutation oracle admits an efficient implementation cIC, which we address using a standard isometry-based argumentation.

In terms of concrete bounds, on a message space of n -bit strings, the oracle cpO is indistinguishable from a random permutation P by [Car26], under a (single-instance) bound of order $O(q^3/2^{n/4})$. Very recently, two works [Kim26, CM26] that are independent and concurrent to ours (see Section 1.2) further improve this bound to $O(q/2^{n/2})$. This would in turn give us a bound for the distinguishing advantage between cIC and the corresponding ideal cipher P^K , under a bound of order $O(q^3/2^{n/2})$. More generally, let $\Delta_{\text{cpO}}(n, q)$ be the maximum distinguishing advantage between cpO and P given q quantum queries, our obtained bound for cIC and P^K is $8q^2 \cdot \Delta_{\text{cpO}}(n, q)$ as spelled out below.

Theorem (Compressed ideal cipher, informal). *There exists an efficiently implementable compressed ideal cipher cIC such that every oracle algorithm making at most q quantum bidirectional queries has distinguishing advantage at most*

$$8q^2 \cdot \Delta_{\text{cpO}}(n, q)$$

between cIC and the ideal cipher P^K .

Given the fast development toward improving single-instance bounds, we would like to highlight the usefulness of our indistinguishability lifting theorem. Indeed, our lifting theorem works generically, irrespective of the specifics of the underlying oracles. Therefore, should there be any further improvement to the underlying compressed permutation bound, we would immediately obtain a corresponding improved bound for the compressed ideal cipher, without having to spend additional efforts redoing the proof from scratch.

As an application of our compressed ideal cipher, we prove preimage resistance of the Davies-Meyer compression function in the quantum ideal cipher model. We do so via mimicking existing proof techniques in the quantum random oracle model [CFHL21], bounding query-by-query the magnitude of the compressed ideal cipher database containing such a preimage.

Domain extension of quantum PRPs. The second application of our lifting theorem is domain extension for quantum-secure strong pseudorandom permutations (PRPs).⁴ We show that every quantum-secure strong PRP on n -bit message space can be extended to another one that operates on $2n$ -bit message space. This is done via a two-step, modular construction, as illustrated in Fig. 1.

⁴ Prior to our work, Bhaumik et al. also [BCF⁺24] studied domain extension of (forward-only) quantum PRP, which requires an additional random-access assumption.




Fig. 1. Domain extension of a PRP from n -bit to $2n$ -bit messages.

In the first step, we construct a tweakable cipher TWK from PRP, via using a pseudorandom function (PRF) PRF that, given any tweak t as input, acts as a key-derivation procedure for PRP:

$$\text{TWK}_K(t, x) := \text{PRP}_{\text{PRF}_K(t)}(x) .$$

Here, we note that the required PRF can easily be obtained from PRP using the PRP-PRF switching lemma [Zha15]. To show that this generic construction of tweakable cipher is indeed quantum-secure, i.e. indistinguishable from an ideal cipher, we first replace queries to $\text{PRF}_K(t)$ with queries to a random oracle $\text{RO}(t)$. Then, the resulting oracle $\text{TWK}'_K(t, x) := \text{PRP}_{\text{RO}(t)}(x)$ can be seen as a keyed oracle PRP^K of PRP with an independent PRP key sampled for each tweak. Thanks to the fact that our indistinguishability lifting theorem also applies in the computational setting, we immediately conclude that it is indistinguishable from an ideal cipher. Prior to our work, this generic construction of tweakable cipher has been studied [LLMM18] in the classical setting, but the quantum security remained untreated.

In the second step, we use the two-round variant of the domain extender proposed in [CDMS10], whose classical strong pseudorandomness is known, to construct the intended $2n$ -bit PRP from the tweakable cipher TWK. We further extend this strong pseudorandomness to the quantum setting using a tailored polynomial-method analysis, which may be of independent interest.

1.2 Concurrent Work

Two very recent works, by Kim [Kim26] and by Carolan and Majenz [CM26], which are independent and concurrent to ours, revisit the soundness of Carolan’s compressed permutation oracle cpO [Car26]. Both improve its single-instance bound from $O(q^3/2^{n/4})$ to $O(q/2^{n/2})$ using different routes. Kim gives a concrete bound of $4q/2^{n/2}$ for $q \leq 2^n/4$, which is also formalized in Lean. As applications, Kim obtains tight query complexities for preimage and collision search in the sponge, and for collision search in the keyless Davies-Meyer construction. Carolan and Majenz construct the compression isometry directly from a POVM on the naive purification of a random permutation, and use representation theory of the symmetric group to control its normalization. As applications, they give explicit-constant bounds for the sponge, as well as preimage and collision bounds for the Davies-Meyer construction.

As discussed in Section 1.1, plugging these improved bounds into our lifting theorem immediately gives a bound of order $O(q^3/2^{n/2})$ between our compressed ideal cipher clC and the ideal cipher. Notably, Carolan and Majenz also treat the ideal cipher model directly: they give each key an independent compressed database, which is the same structure as our clC, and observe that a query under a key k acts only on the k -th database. Their one-query analysis then carries over without any dependence on the number of keys, giving a bound of order $O(q/2^{n/2})$ for the ideal cipher as well. For the specific oracle cpO, their direct analysis is therefore tighter than our generic route by a factor of q^2 . Likewise, their Davies-Meyer preimage bound of order $O(q^2/2^n)$ is tighter than ours in Theorem 3, whose dominant term comes from the lifting. We view the two approaches as complementary. Their extension relies on the structure of their soundness proof for cpO, whereas our lifting theorem treats the base oracles as black boxes, and so applies to arbitrary (possibly stateful) base oracles, including other compressed permutation oracles whose soundness proofs may not have this structure.

1.3 Other Related Work

Quantum security of symmetric cryptography. Our work belongs to the broader study of the quantum security of symmetric constructions. Kaplan et al. [KLLN16] showed that adversaries with superposi-

tion access to the keyed construction can break several constructions that are secure classically. On the positive side, quantum security has been established for message authentication [BZ13, AMRS20], for HMAC and NMAC [HI21], and for authenticated encryption from tweakable block ciphers [BBC⁺21], as well as for the sponge construction in several settings [CBH⁺17, CP24, CPZ24, ACMT25]. For the FX key-length extension [JST21], Even-Mansour cipher [ABKM22], and the Key-Alternating Feistel cipher [BBC⁺25], (various notions of) security has been established against adversaries with classical access to the keyed construction and quantum access to the underlying primitive. Closer to our applications, Zhandry [Zha16] constructed quantum-secure PRPs from quantum-secure PRFs and classical format-preserving encryption, obtaining the first quantum-secure PRP from one-way functions, Hosoyamada and Iwata [HI19] attempted to show that the 4-round Luby-Rackoff construction is a quantum-secure PRP, but a gap in their proof was subsequently identified by Bhaumik et al. [BCEJ24]; eventually, Carolan [Car26] showed that the 7-round Feistel construction is a strong quantum-secure PRP. The quantum security of the Davies-Meyer construction was studied by Hosoyamada and Yasuda [HY18]. For the analysis of quantum queries to random permutations, we also mention the permutation superposition oracles of [MMW25] and the lower bounds for two-sided permutation inversion of [ABPS24].

Quantum security of PRG with advice. An active line of work [HXY19, CGLQ20, Liu23, HL26] has established increasingly tighter bounds on the PRG security of random oracles, against quantum adversaries given L (qu)bits of oracle-dependent advice. As noted in [CGLQ20], in the classical advice setting [HXY19] obtained a bound of order $\tilde{O}(\sqrt[3]{Lq^4/N} + q^4/N)$, where q is the number of (quantum) random oracle queries, and N is the size of the key space. This is then improved to $\tilde{O}(\sqrt[3]{Lq/N} + q^2/N)$ by [CGLQ20], the authors of which further extended the results to the quantum advice setting with a bound $\tilde{O}(\sqrt[10]{L^5q/N} + L^4q^2/N)$. The subsequent work [Liu23] improves the quantum advice bound to $O(\sqrt[3]{Lq/N} + \sqrt{q^2/N})$. Huang and Liao [HL26] further improve it to $O(\sqrt[3]{Lq/N} + q^2/N)$, replacing the square-root term with $O(q^2/N)$. They also establish a multi-challenge bound $O(\sqrt[3]{nL(q+n)/N} + nq^2/N)$, where n is the number of PRG challenges given to the adversary. Notably, the polynomial method is also used in [HL26], though in a different way, to obtain a tighter PRG security bound.

Quantum security of sum of permutations. Two recent papers by Bhaumik, Guo, Guo, and Jha [BGGJ26] and Choi, Hhan, and Jang [CHJ26] study the quantum indistinguishability of the sum (or XOR) of independent random permutations from a random function, a classical way to convert permutations into functions beyond the birthday bound. Choi, Hhan, and Jang show that the XOR of $r \geq 2$ random permutations over n -bit strings is quantum-secure beyond the quantum birthday bound, while Bhaumik et al. treat sums over arbitrary finite abelian groups, together with linear postprocessing such as truncation and Dinur’s LxOP construction. Both works use Fourier-analytic variants of the polynomial method, in which the advantage of a q -query quantum distinguisher is controlled by components of degree at most $2q$, determined by $2q$ input-output pairs of the permutations. This is close in spirit to our analysis of the two-round construction in Section 5.2, where we also write the output probability of a distinguisher as a polynomial of degree at most $2q$ in the input-output indicators of a random permutation, although the constructions and the resulting arguments are different. Their results also concern the first step of our domain extension, where we obtain PRF from PRP via the PRP-PRF switching lemma; a beyond-birthday conversion such as the XOR of permutations could replace this step, although the overall bound of our domain extension also contains an $O(q^4/2^n)$ term from the second step, which such a replacement would not affect.

Other notions of lifting. Our lifting theorem can be regarded as a decision counterpart of existing lifting theorems for search games, which lift security under classical oracle access to security under quantum oracle access. This line of work builds on top of the measure-and-reprogram technique [DFMS19, DFM20], originally developed for applications such as Fiat-Shamir transformations. Subsequent work by Yamakawa and Zhandry [YZ21] shows that it is more broadly applicable for search games in the quantum random oracle model. This family of lifting theorems is further expanded by Cojocaru et al. [CHL⁺25] to the quantum bidirectional random permutation model, as well as the quantum ideal cipher model. In contrast, our lifting theorem concerns *decision* games, with quantum oracle access both before and after the lifting. Interestingly, our lifting theorem can also recover a weaker form of the aforementioned *search* lifting, which we will discuss in Appendix C, whereas the existing search lifting theorems do *not* directly imply our result.

Classical proofs with ideal permutations and ideal ciphers. Classically, ideal permutations and ideal ciphers underlie many security proofs, including those of the Even-Mansour cipher [EM93, EM97] and key-alternating ciphers [BKL⁺12, CS14], and the equivalence between the random oracle and ideal cipher models [CPS08, HKT11]. Tweakable block ciphers were introduced by Liskov, Rivest, and Wagner [LRW02], and Minematsu [Min09] built $2n$ -bit block ciphers from n -bit tweakable block ciphers. In classical proofs in the ideal cipher model, the permutation associated with each key is typically sampled lazily when the key is first queried, so that only the queried keys enter the analysis. This is exactly the step that fails against adversaries querying the key in superposition, and that our lifting theorem replaces.

2 Preliminaries

2.1 Basic Notation

For every integer $n \in \mathbb{Z}_{>0}$ we write $[n] := \{1, \dots, n\}$, with the convention that $[0] := \emptyset$ is the empty set. For a predicate or event P , we write $\mathbb{1}_P$ for the indicator function that evaluates to 1 if it is satisfied, and 0 otherwise. When P is described in a long expression, we may also write $\mathbb{1}[P]$ to make the formula more neat.

We use the falling Pochhammer (or falling-factorial) notation

$$(a)_k := a(a-1)\cdots(a-k+1) \quad (k \geq 1), \quad (a)_0 := 1.$$

We write $\left\{ \begin{smallmatrix} k \\ j \end{smallmatrix} \right\}$ for the Stirling number of the second kind, namely the number of partitions of a k -element set into j nonempty, unlabeled blocks. Equivalently, these numbers are the coefficients in the change of basis

$$x^k = \sum_{j=0}^k \left\{ \begin{smallmatrix} k \\ j \end{smallmatrix} \right\} (x)_j. \quad (1)$$

2.2 Oracle Notation

For an algorithm/procedure $\mathcal{A}$ that produces binary output, we denote by $\Pr[\mathcal{A}]$ the probability that $\mathcal{A}$ outputs 1, and likewise we denote by $SD(\mathcal{A}, \mathcal{B})$ the statistical distance of the two output distributions of the two algorithms, respectively.

Throughout this paper, we write $\mathcal{A}^O$ for an algorithm given *controlled* oracle access to O , meaning that in each query it can decide on the fly whether to disable the query. In case O is a permutation or an invertible cipher, we take it as understood that $\mathcal{A}$ can make bidirectional queries as well. This is then further emphasized with the notation $O^{\leftrightarrow}$.

Moreover, for any PRF construction PRF, we define

$$\mathbf{Adv}_{\text{PRF}}^{\text{RO}}(\mathcal{A}) := SD(\mathcal{A}^{\text{PRF}_K}, \mathcal{A}^{\text{RO}})$$

to be the distinguishing advantage from a random oracle RO over the randomness of key K . Likewise, we denote by $\mathbf{Adv}_{\text{PRP}}^{\text{RP}\leftrightarrow}(\mathcal{A})$ and $\mathbf{Adv}_{\text{TWK}}^{\text{IC}\leftrightarrow}(\mathcal{A})$ the advantage as a PRP, and advantage as a tweakable cipher, given quantum bidirectional queries.

Stateful oracles. A (stateful) oracle O consists of a memory register M , an initial state ρ on M , and a unitary U_O acting on a query register Q and on M . This is without loss of generality: any randomness or intermediate measurement of the oracle can be purified into ρ and M . A controlled query to O applies $|0\rangle\langle 0|_C \otimes \mathbb{I} + |1\rangle\langle 1|_C \otimes U_O$ for a control qubit C held by the algorithm; in particular, a disabled query leaves M untouched. For a bidirectional oracle, Q contains a direction bit d and $U_O = \sum_d |d\rangle\langle d| \otimes U_O^d$, so the direction may be in superposition. An algorithm making q queries applies its own unitaries on its registers, interleaved with q (controlled) queries, and finally measures its output.

For a finite key set $\mathcal{K}$, the keyed oracle $O^{\mathcal{K}}$ has memory $\bigotimes_{k \in \mathcal{K}} M_k$, where each M_k is a copy of M , initial state $\rho^{\otimes \mathcal{K}}$, and query unitary

$$U_{O^{\mathcal{K}}} := \sum_{k \in \mathcal{K}} |k\rangle\langle k|_K \otimes (U_O)_{QM_k},$$

where K is the key register and $(U_O)_{QM_k}$ acts as U_O on QM_k and as the identity on $M_{k'}$ for $k' \neq k$. In other words, the instances start in independent copies of ρ , and each query is routed to the instance named by the key.

2.3 Polynomial Inequalities

We recall here Markov brothers' inequality, which is used multiple times in this paper. Let $f(x)$ be a (real or complex valued) polynomial, and let $[a, b] \subseteq \mathbb{R}$ be a real interval. Then the Markov brothers' inequality would tell us the following

$$|f'(x)| \leq \frac{2(\deg f)^2}{b-a} \cdot \max_{t \in [a,b]} |f(t)| \quad \forall x \in [a, b]. \quad (2)$$

When we are interested in the interior points $x \in (a, b)$ away from the boundary, then Bernstein's inequality would give us a linear bound in the degree of f :

$$|f'(x)| \leq \frac{\deg f}{\sqrt{(x-a)(b-x)}} \cdot \max_{t \in [a,b]} |f(t)| \quad \forall x \in (a, b). \quad (3)$$

A Concrete Variant. We recall a concrete variant of Markov brothers' inequality that does not rely on maximization of t . Let $d \geq \deg f$, and for every $0 \leq j \leq d$, let $x_j := \cos(j\pi/d)$, and

$$L_j(x) := \prod_{i \in \{0, \dots, d\} \setminus \{j\}} \frac{x - x_i}{x_j - x_i}$$

be the j th element of the corresponding Lagrange basis, i.e. $L_j(x_i) = \mathbb{1}_{i=j}$ and $\deg L_j = d$. Then we immediately get

$$f'(x) = \left(\sum_{j \in \{0, \dots, d\}} L_j(x) \cdot f(x_j) \right)' = \sum_{j \in \{0, \dots, d\}} L_j'(x) \cdot f(x_j)$$

Observe that, from the Duffin-Schaeffer inequality [DS41] we have that $\sum_j |L_j'(x)| \leq d^2$ for every $x \in [-1, 1]$, and so in this case we have

$$|f'(x)| \leq d^2 \left| \mathbb{E}_{X \leftarrow \mathcal{D}_x^d} [\text{sgn}_X^x \cdot f(X)] \right| \quad \forall x \in [-1, 1],$$

where $\text{sgn}_X^x \in \{\pm 1\}$ is the sign⁵ of $L_j'(x)$ for the only j such that $x_j = X$, and $\mathcal{D}_x^d$ is the distribution over $\{x_0, \dots, x_d\}$ with

$$\Pr_{X \leftarrow \mathcal{D}_x^d} [X = x_j] := \frac{|L_j'(x)|}{\sum_i |L_i'(x)|} \quad \forall j \in \{0, \dots, d\}.$$

For an $x \in [a, b]$ sitting in the general interval $[a, b]$, the similar bound as follows can be achieved via rescaling the interval to $[-1, 1]$

$$|f'(x)| \leq \frac{2d^2}{b-a} \left| \mathbb{E}_{T \leftarrow \mathcal{D}_{x, [a,b]}^d} [\Delta \text{sgn}_T^x \cdot f(T)] \right| \quad \forall x \in [a, b], \quad (4)$$

where $\Delta \text{sgn}_T^x := \text{sgn}_{\Delta(T)}^{\Delta(x)}$ for $\Delta(x) := (2x - a - b)/(b - a)$, and the distribution $\mathcal{D}_{x, [a,b]}^d$ is defined as follows

$$\Pr_{T \leftarrow \mathcal{D}_{x, [a,b]}^d} [T = t] := \Pr_{X \leftarrow \mathcal{D}_{\Delta(x)}^d} [X = \Delta(t)].$$

Lifting integer-point upper bounds to an interval. An immediate consequence of (2) is that, if the interval length is sufficiently long, specifically if $b - a \geq 4(\deg f)^2 > 0$, then

$$|f(x)| \leq 2 \max_{n \in [a,b] \cap \mathbb{Z}} |f(n)| \quad \forall x \in [a, b], \quad (5)$$

⁵ Whether the sign at zero is 1 or -1 is irrelevant to our discussion.

meaning that one can lift an upper bound on integer points of f to all continuous points, if it applies to sufficiently many points to begin with.

This can be seen by picking $x \in [a, b]$ that maximizes $|f(x)|$, its closest integer point $n \in [a, b]$, and the following chain of inequalities

$$|f(x)| \leq |f(n)| + \max_{t \in [a, b]} |f'(t)| \leq |f(n)| + \frac{2(\deg f)^2}{b-a} |f(x)| \leq |f(n)| + (1/2) \cdot |f(x)|$$

where the first inequality follows from basic calculus and the fact that $|n - x| \leq 1$, the second inequality follows from (2), and the last follows from the assumption that $b - a \geq 4(\deg f)^2$. Moving the rightmost term to the left immediately concludes the proof of (5).

3 Quantum Indistinguishability Lifting

In this section, we formally state and prove the main theorem of this paper. Throughout, let O_0 and O_1 be any two (possibly stateful) oracles, and let O_0^K and O_1^K be their corresponding keyed oracles, as have been defined in Section 2.2.

Theorem 1 (Quantum indistinguishability lifting). *Let $\mathcal{A}$ be a binary-output oracle algorithm given at most q quantum queries to either O_0^K or O_1^K . Then, there exists a binary-output oracle algorithm $\mathcal{B}$ making at most q quantum controlled queries to O_i , O_0^K , and O_1^K respectively, such that*

$$SD(\mathcal{A}^{O_0^K}, \mathcal{A}^{O_1^K}) \leq 8q^2 \cdot SD(\mathcal{B}^{O_0, O_0^K, O_1^K}, \mathcal{B}^{O_1, O_0^K, O_1^K}). \quad (6)$$

Moreover, $\mathcal{B}$ has a runtime similar to that of $\mathcal{A}$.

Remark 1. If indistinguishability of O_0 and O_1 holds statistically, or if the oracles O_0 and O_1 and their initial states ρ_0 and ρ_1 are efficiently implementable, then the queries made by $\mathcal{B}$ to O_0^K and O_1^K can be internalized.

Proof. Our reduction algorithm $\mathcal{B}$ works as follows. First, it samples a parameter $p^* \leftarrow [0, 1]$ uniformly and another parameter $z^* \leftarrow \mathcal{D}_{0, [0, 1]}^{2q}$ following the distribution $\mathcal{D}_{x, [a, b]}^d$ specified in (4) (with $(x, a, b, d) := (0, 0, 1, 2q)$). It then runs a parametrized subprocedure $\mathcal{B}_{p, z}$ with $(p, z) := (p^*, z^*)$, and flips the output bit of $\mathcal{B}_{p, z}$ whenever the sign $\Delta \text{sgn}_{z^*}^0 \in \{\pm 1\}$ (as defined in (4)) is negative. Finally, it outputs the potentially flipped bit.

As for $\mathcal{B}_{p, z}$, it maintains two functions H and G mapping from $\mathcal{K} \rightarrow \{0, 1\}$, which are independently sampled according to Bernoulli distributions $\text{Ber}(p)$ and $\text{Ber}(z)$ at each entry respectively. Then it runs $\mathcal{A}$ while simulating the oracle queries made by $\mathcal{A}$ through an interface $\mathcal{O}$: on input $|k, x\rangle$ (or its superposition), forward the $|x\rangle$ part to O_i if $G(k)$ is one, and otherwise forward it to O_b^K with $b := H(k)$. We summarize the behavior of $\mathcal{B}$ and $\mathcal{B}_{p, z}$ below.

$\mathcal{B}^{O_i, O_0^K, O_1^K}:$ 1: $(p^*, z^*) \leftarrow \mathcal{U}([0, 1]) \otimes \mathcal{D}_{0, [0, 1]}^{2q}$ 2: $s^* := \mathbb{1}[\Delta \text{sgn}_{z^*}^0 = -1]$ 3: return $s^* \oplus \left(b \leftarrow \mathcal{B}_{p^*, z^*}^{O_i, O_0^K, O_1^K} \right)$	$\mathcal{B}_{p, z}^{O_i, O_0^K, O_1^K}:$ 1: for $k \in \mathcal{K}$: 2: $H(k) \leftarrow \text{Ber}(p)$ 3: $G(k) \leftarrow \text{Ber}(z)$ 4: return $b \leftarrow \mathcal{A}^{\mathcal{O}}$	$\mathcal{O}(k, x):$ 1: if $G(k)$ return $O_i(x)$ 2: else if $H(k)$ 3: return $O_1^K(k, x)$ 4: else return $O_0^K(k, x)$
---	---	---

Fig. 2. The (behavior of the) reduction algorithm $\mathcal{B}$ for indistinguishability lifting, which runs a parametrized subprocedure $\mathcal{B}_{p, z}$ with p and z substituted with p^* and z^* sampled at the beginning, and sometimes flips the output dependent on $\Delta \text{sgn}_{z^*}^0 \in \{\pm 1\}$ defined in (4). Here, $\mathcal{U}([0, 1])$ denotes the uniform distribution over $[0, 1]$, and $\mathcal{D}_{0, [0, 1]}^{2q}$ is as defined in (4). Note that, while we only write the oracle interface $\mathcal{O}$ in classical pseudocode, quantum queries are also handled in the obvious way.

Remark 2. While in Fig. 2 the reduction algorithm looks inefficient, it is actually not. This is because the functions H and G are just (biased) random oracles, where each entry is sampled independently from a Bernoulli distribution with a parameter sampled at the beginning. This can be efficiently implemented using existing techniques. We discuss this matter in more detail in Appendix B.

Our proof is performed via a rather non-obvious use of the polynomial method. Let $f(p)$ and $g_p(z)$ be functions defined as follows:

$$f(p) := \Pr \left[\mathcal{B}_{p,0}^{\mathcal{O}_i, \mathcal{O}_0^K, \mathcal{O}_1^K} \right], \quad g_p(z) := \Pr \left[\mathcal{B}_{p,z}^{\mathcal{O}_1, \mathcal{O}_0^K, \mathcal{O}_1^K} \right] - \Pr \left[\mathcal{B}_{p,z}^{\mathcal{O}_0, \mathcal{O}_0^K, \mathcal{O}_1^K} \right], \quad (7)$$

where the choice of $i \in \{0, 1\}$ in the definition of $f(p)$ is irrelevant (since $\mathcal{O}_i$ is never queried).⁶ From a procedural perspective, the absolute difference $|f(1) - f(0)|$ coincides with $\mathcal{A}$'s distinguishing advantage between $\mathcal{O}_0^K$ and $\mathcal{O}_1^K$. Therefore, we have

$$SD \left(\mathcal{A}^{\mathcal{O}_0^K}, \mathcal{A}^{\mathcal{O}_1^K} \right) = |f(1) - f(0)| = \left| \int_0^1 f'(p) dp \right| = \left| \mathbb{E}_{p^*} [f'(p^*)] \right|. \quad (8)$$

Via a careful analysis deferred to Lemma 2, there is a connection between the derivatives of f and g , namely $f'(p) = g'_p(0)$. Furthermore, by Lemma 1, the query complexity of $\mathcal{A}$ implies that $g_p(z)$ is a polynomial in z of degree at most $2q$, and hence so is the polynomial $\bar{g}(z) := \mathbb{E}_{p^*} [g_{p^*}(z)]$. The derivative $g'_p(0)$ can be controlled (on average over $p \sim p^*$) via

$$\begin{aligned} \left| \mathbb{E}_{p^*} [g'_{p^*}(0)] \right| &= \left| \frac{d}{dz} \mathbb{E}_{p^*} [g_{p^*}(z)] \Big|_{z=0} \right| \\ &= |\bar{g}'(0)| \leq 8q^2 \cdot \left| \mathbb{E}_{z^*} [\Delta \text{sgn}_{z^*}^0 \cdot \bar{g}(z^*)] \right| = 8q^2 \cdot \left| \mathbb{E}_{p^*, z^*} [\Delta \text{sgn}_{z^*}^0 \cdot g_{p^*}(z^*)] \right| \\ &= 8q^2 \cdot \left| \Pr \left[s^* \oplus \left(b \leftarrow \mathcal{B}_{p^*, z^*}^{\mathcal{O}_1, \mathcal{O}_0^K, \mathcal{O}_1^K} \right) \right] - \Pr \left[s^* \oplus \left(b \leftarrow \mathcal{B}_{p^*, z^*}^{\mathcal{O}_0, \mathcal{O}_0^K, \mathcal{O}_1^K} \right) \right] \right| \\ &= 8q^2 \cdot SD \left(\mathcal{B}^{\mathcal{O}_0, \mathcal{O}_0^K, \mathcal{O}_1^K}, \mathcal{B}^{\mathcal{O}_1, \mathcal{O}_0^K, \mathcal{O}_1^K} \right), \end{aligned} \quad (9)$$

where the only (non-equality) inequality is via (4), and in the second to last equality $s^* := \mathbb{1}_{\Delta \text{sgn}_{z^*}^0 = -1}$. We now explain the equalities one by one. The first equality is via $g_p(z)$ being a polynomial whose coefficients are continuous in p , so that expectation commutes with derivative. The second equality is via the definition of $\bar{g}(z)$, and so is the third to last equality. The second to last equality is via the fact that $\Pr[1 \oplus b_1] - \Pr[1 \oplus b_0] = -(\Pr[b_1] - \Pr[b_0])$ for any binary random variables b_0 and b_1 . The last equality is via the definition of $\mathcal{B}$.

Collecting the bounds so far, we conclude the proof with

$$SD \left(\mathcal{A}^{\mathcal{O}_0^K}, \mathcal{A}^{\mathcal{O}_1^K} \right) = \left| \mathbb{E}_{p^*} [f'(p^*)] \right| = \left| \mathbb{E}_{p^*} [g'_{p^*}(0)] \right| \leq 8q^2 \cdot SD \left(\mathcal{B}^{\mathcal{O}_0, \mathcal{O}_0^K, \mathcal{O}_1^K}, \mathcal{B}^{\mathcal{O}_1, \mathcal{O}_0^K, \mathcal{O}_1^K} \right)$$

where the first equality follows from (8), the second equality follows from the fact that $f'(p) = g'_p(0)$, and the only (non-equality) inequality follows from (9). $\square$

In the remainder, we state and prove Lemmas 1 and 2 that are used in the above proof.

Lemma 1. *For every $0 < p < 1$, the function $g_p(z)$ defined in (7) is a polynomial in z of degree at most $2q$.*

Proof. For every $G^\circ : \mathcal{K} \rightarrow \{0, 1\}$, let

$$h_{p,z,i}(G^\circ) := \Pr \left[\mathcal{B}_{p,z}^{\mathcal{O}_i, \mathcal{O}_0^K, \mathcal{O}_1^K} \mid G = G^\circ \right] \quad (10)$$

so that $g_p(z) = \mathbb{E}_G [h_{p,z,1}(G) - h_{p,z,0}(G)]$. A standard argument in the polynomial method (see e.g. [BBC⁺98]) implies that the function $h_{p,z,i}(G)$ is a multilinear polynomial in $(G(k))_{k \in \mathcal{K}}$ of degree at most $2q$. Indeed, the oracle G is only referenced in $\mathcal{O}$ while running $\mathcal{B}_{p,z}$, and each of the q queries to $\mathcal{O}$ implements the unitary

$$|k, x, y\rangle \mapsto G(k) |k\rangle \otimes \mathcal{O}_i |x, y\rangle + (1 - G(k)) \mathcal{O}_{H(k)}^K |k, x, y\rangle,$$

⁶ Note that $f(p)$ and $g_p(z)$ are polynomials in p and z respectively, and so they are differentiable.

which would only increase the degree of the amplitudes (as a polynomial in $(G(k))_k$) by one. Rewrite $h_{p,z,i}(G)$ in Fourier expansion (i.e. in terms of monomials):

$$h_{p,z,i}(G) = \sum_{S \subseteq \mathcal{K}: |S| \leq 2q} \widehat{h}_{p,z,i}(S) \prod_{k \in S} G(k) ,$$

and then exploiting the independence of $(G(k))_{k \in \mathcal{K}}$, we have that

$$\begin{aligned} \mathbb{E}_G [h_{p,z,i}(G)] &= \mathbb{E}_G \left[\sum_{S \subseteq \mathcal{K}: |S| \leq 2q} \widehat{h}_{p,z,i}(S) \prod_{k \in S} G(k) \right] \\ &= \sum_{S \subseteq \mathcal{K}: |S| \leq 2q} \widehat{h}_{p,z,i}(S) \prod_{k \in S} \mathbb{E}_G [G(k)] = \sum_{S \subseteq \mathcal{K}: |S| \leq 2q} \widehat{h}_{p,z,i}(S) \cdot z^{|S|} , \end{aligned}$$

is a polynomial in z of degree at most $2q$, and hence

$$\deg_z g_p(z) = \deg_z \left(\mathbb{E}_G [h_{p,z,1}(G)] - \mathbb{E}_G [h_{p,z,0}(G)] \right) \leq 2q .$$

□

Lemma 2. Let $f(p)$ and $g_p(z)$ be defined as in (7). Then, for every $0 < p < 1$, we have

$$f'(p) = g'_p(0) .$$

Proof. For every $S \subseteq \mathcal{K}$ let

$$W_{i,p}(S) := \Pr \left[\mathcal{B}_{p,z}^{\mathcal{O}_i, \mathcal{O}_0^K, \mathcal{O}_1^K} \mid G(k) \equiv \mathbb{1}_{k \in S} \right] , \quad \Delta_p(S) := W_{1,p}(S) - W_{0,p}(S) ,$$

where the choice of z in the above definition is irrelevant (since z is used to sample G , which is fixed by the conditioning on S). We can then rewrite $g_p(z)$ as follows:

$$g_p(z) = \sum_{S \subseteq \mathcal{K}} \Pr [G(k) \equiv \mathbb{1}_{k \in S}] \cdot \Delta_p(S) = \sum_{S \subseteq \mathcal{K}} z^{|S|} (1-z)^{|\mathcal{K}| - |S|} \cdot \Delta_p(S) .$$

Differentiating both sides with respect to z at 0, we get

$$g'_p(0) = \sum_{k \in \mathcal{K}} \Delta_p(\{k\}) - |\mathcal{K}| \Delta_p(\emptyset) = \sum_{k \in \mathcal{K}} \Delta_p(\{k\}) ,$$

where in the second equality we exploit that $\Delta_p(\emptyset)$ vanishes.

Toward calculating $f'(p)$, let

$$V(S) := \Pr \left[\mathcal{B}_{p,0}^{\mathcal{O}_i, \mathcal{O}_0^K, \mathcal{O}_1^K} \mid H(k) \equiv \mathbb{1}_{k \in S} \right] ,$$

where again the choice of i and p in the above definition is irrelevant (since $\mathcal{O}_i$ is never queried, and p is used to sample H which is fixed by S).

Then, by the fact that the random variables $(H(k))_{k \in \mathcal{K}}$ are independent, we get

$$f(p) = \sum_{S \subseteq \mathcal{K}} \Pr [H(k) \equiv \mathbb{1}_{k \in S}] \cdot V(S) = \sum_{S \subseteq \mathcal{K}} V(S) \cdot \prod_{k \in \mathcal{K}} \Pr [H(k) = \mathbb{1}_{k \in S}] .$$

Observe that, taking the derivative of $\Pr [H(k) = \mathbb{1}_{k \in S}]$ with respect to p would give us

$$\frac{d}{dp} \Pr [H(k) = \mathbb{1}_{k \in S}] = \frac{d}{dp} (p \cdot \mathbb{1}_{k \in S} + (1-p) \cdot \mathbb{1}_{k \notin S}) = \mathbb{1}_{k \in S} - \mathbb{1}_{k \notin S} .$$

Therefore, taking the derivative of $f(p)$, we obtain

$$f'(p) = \sum_{k \in \mathcal{K}} \sum_{S \subseteq \mathcal{K}} V(S) \cdot (\mathbb{1}_{k \in S} - \mathbb{1}_{k \notin S}) \prod_{k' \in \mathcal{K} \setminus \{k\}} \Pr [H(k') = \mathbb{1}_{k' \in S}] ,$$

Further observing that

$$\begin{aligned} \sum_{S \subseteq \mathcal{K}} V(S) \cdot \mathbb{1}_{k \in S} \prod_{k' \in \mathcal{K} \setminus \{k\}} \Pr[H(k') = \mathbb{1}_{k' \in S}] &= W_{1,p}(\{k\}) \\ \sum_{S \subseteq \mathcal{K}} V(S) \cdot \mathbb{1}_{k \notin S} \prod_{k' \in \mathcal{K} \setminus \{k\}} \Pr[H(k') = \mathbb{1}_{k' \in S}] &= W_{0,p}(\{k\}) , \end{aligned}$$

we immediately obtain

$$f'(p) = \sum_{k \in \mathcal{K}} (W_{1,p}(\{k\}) - W_{0,p}(\{k\})) = \sum_{k \in \mathcal{K}} \Delta_p(\{k\}) = g'_p(0) ,$$

which concludes the proof. $\square$

3.1 On Black-box Optimality

It is easy to see that our bound in Theorem 1 cannot be further improved, unless relying on certain specifics of the considered oracles. Indeed, consider an oracle O_p parametrized by $p \in [0, 1]$, where every output $O_p(X)$ is set to a bit b that is sampled at the beginning, with probability p being one:

$$O_p(X) := b \quad \forall X, \quad b \xleftarrow{\text{init}} \text{Ber}(p) .$$

Then the distinguishing advantage between O_0 and O_ϵ is at most ϵ , and a simple Grover search would distinguish the keyed oracles O_0^K and O_ϵ^K with advantage up to $\Omega(q^2\epsilon)$.

However, that is not to say our bound is tight for all choices of oracles. For example, the distinguishing advantage between $O_{(1 \pm \epsilon)/2}$ is ϵ , but the distinguishing advantage between $O_{(1 \pm \epsilon)/2}^K$ can be further bounded by

$$SD\left(\mathcal{A}^{O_{(1-\epsilon)/2}^K}, \mathcal{A}^{O_{(1+\epsilon)/2}^K}\right) = |f((1+\epsilon)/2) - f((1-\epsilon)/2)| \leq \int_{(1-\epsilon)/2}^{(1+\epsilon)/2} |f'(p)| dp$$

where $f(p) := \Pr[\mathcal{A}^{O_p^K}] \in \mathbb{R}[p]$ is a polynomial in p of degree at most $2q$. Via Bernstein's inequality (Eq. (3)), we obtain $|f'(p)| \leq 2q/\sqrt{p(1-p)} \leq 8q/\sqrt{3}$ for every $p \in [1/4, 3/4]$. This would then tell us that, whenever $\epsilon \leq 1/4$,

$$SD\left(\mathcal{A}^{O_{(1-\epsilon)/2}^K}, \mathcal{A}^{O_{(1+\epsilon)/2}^K}\right) \leq \int_{(1-\epsilon)/2}^{(1+\epsilon)/2} (8q/\sqrt{3}) dp \leq 8q\epsilon/\sqrt{3} .$$

We summarize the above discussion below.

Proposition 1. *There exists a pair of efficiently implementable oracles O_0 and O_1 that is ϵ -indistinguishable, no matter how many queries are made, but for some oracle algorithm $\mathcal{A}$ making at most q quantum queries, and a sufficiently large key space $\mathcal{K}$, we have that*

$$SD\left(\mathcal{A}^{O_0^K}, \mathcal{A}^{O_1^K}\right) \geq \Omega(q^2\epsilon) .$$

On the other hand, there also exists a pair of efficiently implementable oracles O'_0 and O'_1 that is ϵ -distinguishable in one classical query, but for every oracle algorithm $\mathcal{A}$ making at most q quantum queries, we have that

$$SD\left(\mathcal{A}^{O'_0^K}, \mathcal{A}^{O'_1^K}\right) \leq O(q\epsilon) .$$

4 Compressed Ideal Cipher and an Application

4.1 Carolan's Compressed Permutation Oracle

Before looking into the case of ideal cipher, we first recall Carolan's compressed permutation oracle [Car26]. Essentially, it is a simulator of random permutation which enjoys useful properties for security proofs.

Consider a message space $\mathcal{X} := \{0, 1\}^n$ that consists of all n -bit strings, a random permutation P is sampled uniformly from $\text{Sym}(\{0, 1\}^n)$ the set of all permutations over $\mathcal{X}$, and allows an algorithm $\mathcal{A}$ to make queries in the form of the following unitary

$$O_P : |x, y\rangle \mapsto |x, y \oplus P(x)\rangle ,$$

for every $x, y \in \mathcal{X}$. Likewise, we also allow $\mathcal{A}$ to make backward queries, which are of the form $O_{P^{-1}}$.

The compressed permutation oracle, which we denote by cpO , is a stateful simulation of P that maintains the same bidirectional query access, but differs in the way each query is responded to. The internal state of cpO is maintained in a register D , which we call the *database register*. The state of this register is supported by vectors of the form $|D\rangle$, where $D : \mathcal{X} \rightarrow \mathcal{X} \cup \{\perp\}$, called a *database*, is a partially specified permutation over $\mathcal{X}$, with the special symbol $\perp$ denoting that a specific function value has not been specified. Specifically, for every distinct $x, y \in \mathcal{X}$ such that $D(x)$ and $D(y)$ are non- $\perp$, it must be that $D(x) \neq D(y)$. As a short-hand notation, it is standard to write $|D|$ for the number of non- $\perp$ entries, $\text{Dom}(D)$ for the set of x such that $D(x) \neq \perp$, $\text{Im}(D)$ for the set of such $D(x)$'s, and $(x, y) \in D$ for $D(x) = y \neq \perp$.

Initially, the database register $D := |\perp\rangle$ is set to the constant- $\perp$ database, and each forward query is implemented by the following unitary

$$\text{cpO}^\rightarrow : |x, y, D\rangle \mapsto |x\rangle C_x \cdot P_x \cdot C_x |y, D\rangle .$$

Here, for every $x \in \mathcal{X}$ and database D such that $D(x) = \perp$, one writes

$$|+_{x,D}\rangle := \frac{1}{\sqrt{2^n - |D|}} \sum_{y \notin \text{Im}(D)} |D[x \mapsto y]\rangle$$

for the uniform superposition of all (injective) databases that agree with D outside of x ; the *compression operator* C_x maps $|D\rangle \mapsto |+_{x,D}\rangle$ and $|+_{x,D}\rangle \mapsto |D\rangle$ and fixes the orthogonal complement, and the operator P_x works as follows:

$$P_x : |y, D\rangle \mapsto \begin{cases} |y \oplus D(x), D\rangle & \text{if } D(x) \neq \perp \\ |y, D\rangle & \text{otherwise} . \end{cases}$$

For backward queries an additional flipping operator $F : |D\rangle \mapsto |D^{-1}\rangle$ with $(x, y) \in D$ if and only if $(y, x) \in D^{-1}$ is conjugated to the $\text{cpO}^\rightarrow$ operator. Namely, the following unitary is applied

$$\text{cpO}^\leftarrow : |x, y, D\rangle \mapsto |x\rangle \otimes F \cdot C_x \cdot P_x \cdot C_x \cdot F |y, D\rangle .$$

Indeed, as has been shown in [Car26], and subsequently sharpened by [Kim26, CM26], this simulator is indistinguishable from the random permutation. In the statement below, let

$$\Delta_{\text{cpO}}(n, q) := \sup_{\mathcal{A}} SD(\mathcal{A}^P, \mathcal{A}^{\text{cpO}})$$

where the supremum is taken over all oracle algorithms $\mathcal{A}$ making at most q quantum queries to either cpO or P bidirectionally, and n specifies the bit-length of elements in $\mathcal{X}$.

Lemma 3 (Bounds achieved in [Kim26, CM26]). *For $n, q \in \mathbb{Z}_{>0}$, we have*

$$\Delta_{\text{cpO}}(n, q) \leq O\left(\frac{q}{2^{n/2}}\right) .$$

4.2 Compressed Ideal Cipher

Now, we are ready to address the case of ideal cipher. Consider $P^{\mathcal{K}}(k, x) := P_k(x)$ an oracle where P_k is sampled independently from $\text{Sym}(\mathcal{X})$ for each k . Again, we allow an algorithm to query this oracle bidirectionally, where a forward query takes as input $(k, x) \in \mathcal{K} \times \mathcal{X}$ and then responds $P_k(x)$, while a backward query responds $P_k^{-1}(x)$.

We analogously define the *compressed ideal cipher* oracle cIC , which is a stateful simulation of $P^{\mathcal{K}}$ that maintains bidirectional query access. The internal database register D is now supported by state vectors of the form $|D\rangle$, where $D : \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{X} \cup \{\perp\}$ must be a valid database. Namely, for every

$k \in \mathcal{K}$ and distinct $x, x' \in \mathcal{X}$ such that $D(k, x)$ and $D(k, x')$ are non- $\perp$, it must be that they are distinct, i.e. $D(k, x) \neq D(k, x')$. For notational convenience, we will denote by $D_k(x) := D(k, x)$ the part of the database associated to the key k .

Initially, the database register D is set to the constant- $\perp$ database, and, upon each query, the compressed ideal cipher oracle clC would apply the operators $\text{cpO}^\rightarrow$ (resp. $\text{cpO}^\leftarrow$) on D_k . That is, the forward and backward queries are respectively implemented by the following unitary operators:

$$\begin{aligned} \text{clC}^\rightarrow : |k, x, y, D\rangle &\mapsto |k, x\rangle \mathbf{C}_{k,x} \cdot \mathbf{P}_{k,x} \cdot \mathbf{C}_{k,x} |y, D\rangle \\ \text{clC}^\leftarrow : |k, x, y, D\rangle &\mapsto |k, x\rangle \mathbf{F}_k \cdot \mathbf{C}_{k,x} \cdot \mathbf{P}_{k,x} \cdot \mathbf{C}_{k,x} \cdot \mathbf{F}_k |y, D\rangle . \end{aligned}$$

Here, $\mathbf{P}_{k,x}$ maps $|k, x, y, D\rangle \mapsto |k, x, y \oplus D(k, x), D\rangle$ if $D(k, x) \neq \perp$ and fixes the state otherwise; for every $(k, x) \in \mathcal{K} \times \mathcal{X}$ and D such that $D(k, x) = \perp$, the operator $\mathbf{C}_{k,x}$ swaps the two vectors $|+_{k,x,D}\rangle$ and $|D\rangle$ and fixes the orthogonal complement, where

$$|+_{k,x,D}\rangle := \frac{1}{\sqrt{2^n - |D_k|}} \sum_{y \notin \text{Im}(D_k)} |D[(k, x) \mapsto y]\rangle ;$$

also the flipping operator $\mathbf{F}_k$ maps $|D\rangle \mapsto |D'\rangle$ where $D'_{k'} := D_{k'}$ for every $k' \neq k$ and $D'_k = D_k^{-1}$.

As an application of our indistinguishability lifting theorem, we show that the compressed ideal cipher oracle as described above is indistinguishable from the ideal cipher $P^\mathcal{K}$.

Theorem 2. *Let $\mathcal{A}$ be an oracle algorithm making at most q quantum queries to either the compressed ideal cipher clC or the ideal cipher $P^\mathcal{K}$ bidirectionally, where the message space $\mathcal{X} = \{0, 1\}^n$ consists of all n -bit strings. Then, we have*

$$SD(\mathcal{A}^{\text{clC}}, \mathcal{A}^{P^\mathcal{K}}) \leq 8q^2 \cdot \Delta_{\text{cpO}}(n, q) \leq O\left(\frac{q^3}{2^{n/2}}\right) .$$

Proof. Our lifting theorem (Theorem 1) immediately implies another oracle algorithm $\mathcal{B}$ making at most q quantum queries such that

$$SD(\mathcal{A}^{\text{cpO}^\mathcal{K}}, \mathcal{A}^{P^\mathcal{K}}) \leq 8q^2 \cdot SD(\mathcal{B}^{\text{cpO}}, \mathcal{B}^P) .$$

Hence, it suffices to argue that the keyed compressed permutation oracle $\text{cpO}^\mathcal{K}$, which runs one independent instance of cpO per key k , behaves identically to our previously defined compressed ideal cipher clC . This can be seen via considering the following isometry

$$\mathcal{J} : |D\rangle \mapsto \bigotimes_{k \in \mathcal{K}} |D_k\rangle ,$$

which, by construction, has the easy-to-verify property that

$$\mathcal{J} \cdot \text{clC}^\rightarrow = (\text{cpO}^\mathcal{K})^\rightarrow \cdot \mathcal{J} , \quad \mathcal{J} \cdot \text{clC}^\leftarrow = (\text{cpO}^\mathcal{K})^\leftarrow \cdot \mathcal{J} .$$

This would imply that, for the hybrid game Hyb_i that replaces the first i queries of $\mathcal{A}^{\text{cpO}^\mathcal{K}}$ to querying clC , we have that Hyb_i behaves identically to Hyb_{i+1} for every i . Hence Hyb_0 , which is simply $\mathcal{A}^{\text{cpO}^\mathcal{K}}$, behaves identically to Hyb_q , which is $\mathcal{A}^{\text{clC}}$. We therefore have that

$$SD(\mathcal{A}^{\text{clC}}, \mathcal{A}^{P^\mathcal{K}}) \leq 8q^2 \cdot \Delta_{\text{cpO}}(n, q) \leq O\left(\frac{q^3}{2^{n/2}}\right) ,$$

where the asymptotic bound is via plugging in the results that we have recalled in Lemma 3. $\square$

4.3 Preimage Resistance of Davies-Meyer Construction

As an application of our compressed ideal cipher oracle, we show preimage resistance of the Davies-Meyer construction:

$$\text{DM}(L, R) := P^\mathcal{K}(L, R) \oplus R ,$$

where we model $P^\mathcal{K} : \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{X}$ as a (quantum) ideal cipher.

Theorem 3. Let $\mathcal{A}$ be an oracle algorithm making at most $q < |\mathcal{X}| - 1$ quantum queries to P^K bidirectionally. Then for every $y \in \mathcal{X}$, and for $(L, R) \leftarrow \mathcal{A}^{P^K}(y)$, we have

$$\Pr [\text{DM}(L, R) = y] \leq \frac{32(q+1)^2 + 1}{|\mathcal{X}| - q - 1} + 8(q+1)^2 \Delta_{\text{cpO}}(n, q+1) .$$

Proof. First, we consider a game G_0 . In this game, we run $\mathcal{A}$, but substitute each forward query (k, x) to P^K with a query to $P^K(k, x) \oplus S$, where $S \leftarrow \mathcal{X}$ is sampled at the beginning of the game. Likewise, we replace the backward queries in the obvious way. After the output (L, R) is produced by $\mathcal{A}$, the game aborts if $R = Y_S := y \oplus S$ and otherwise outputs the result of the check whether $P^K(L, R) \oplus R = Y_S$. The substitution to each query has no effect on the output distribution of $\mathcal{A}$, and since S is kept secret, the abort only happens with probability at most $1/|\mathcal{X}|$. Therefore,

$$\Pr_{(L,R) \leftarrow \mathcal{A}^{P^K}} [\text{DM}(L, R) = y] \leq \Pr [G_0] + 1/|\mathcal{X}| . \quad (11)$$

Next, we consider a game G_1 , in which we further substitute all queries to P^K in G_0 with queries to the compressed ideal cipher oracle clC . Due to our prior proven indistinguishability of compressed ideal cipher, we have that

$$\Pr [G_0] \leq \Pr [G_1] + 8(q+1)^2 \Delta_{\text{cpO}}(n, q+1) , \quad (12)$$

where the +1 in query complexity comes from the final verification query. In this game, let $|\phi_0\rangle_{\text{XYZD}}$ be the initial state of $\mathcal{A}$'s query input and output registers X and Y , working register Z , and the database register D of clC , which is just a constant- $\perp$ database. Moreover, let $|\phi_i\rangle$ be the state of XYZD right before the i th query. Then for every fixed choice y_s of Y_S , the projector

$$\Pi^{y_s} := \sum_{k, x: x \neq y_s} |(k, x), y_s \oplus x\rangle \langle (k, x), y_s \oplus x|_{\text{XY}}$$

would capture the final successful event of $\mathcal{A}$ outputting a preimage of y_s , and the game does not abort. Therefore, we have that

$$\Pr [G_1 \mid Y_S = y_s] \leq \|\Pi^{y_s} \cdot \text{clC}^\rightarrow |\phi_{q+1}\rangle_{\text{XYZD}}\|^2 \quad (13)$$

where we note that $|\phi_{q+1}\rangle$ can be written in the form of $\sum_{k,x} \alpha_{k,x} |k, x, 0\rangle |\psi_{k,x}\rangle_{\text{ZD}}$, so that $(|0\rangle \langle 0|)_Y |\phi_{q+1}\rangle = |\phi_{q+1}\rangle$.

Similar to [CFHL21], define the following sets of databases

$$\text{SZ}_i := \{D \mid |D| \leq i\} , \quad \text{PRMG}_{y_s} := \{D \mid \exists ((k, x), w) \in D : w \oplus x = y_s\} ,$$

where we then abuse the notation and treat these sets as projectors onto the databases in these sets respectively.

Throughout the game, since $\mathcal{A}$ makes at most q quantum queries, it follows that $|\phi_i\rangle$ is supported by all states of database size at most q , i.e. $\text{SZ}_q |\phi_i\rangle = |\phi_i\rangle$. For every $(k, x) \in \mathcal{K} \times \mathcal{X}$, let D be a database such that $D(k, x) = \perp$, and let $\text{clC}_{k,x,D}^\rightarrow$ be the restriction of $\text{clC}^\rightarrow$ to the invariant subspace where X is fixed in the computational basis to $|k, x\rangle$, and D is supported by database states $|D'\rangle$ such that $D'(k', x') = D(k', x')$ for every $(k', x') \neq (k, x)$. Given that all entries in D' other than $D'(k, x)$ are fixed, the state space of D is spanned by all possible $|D'(k, x)\rangle$, in which we have:

$$\begin{aligned} \|\text{PRMG}_{y_s} \cdot \text{C}_{k,x} \cdot \text{SZ}_q \cdot (\mathbb{I} - \text{PRMG}_{y_s})\|^2 &\leq \left\| \sum_{\substack{w \in \mathcal{X} \cup \{\perp\} \setminus \{y_s \oplus x\} \\ w \notin \text{Im}(D_k)}} |y_s \oplus x\rangle \langle y_s \oplus x| \text{C}_{k,x} \cdot \text{SZ}_q |w\rangle \langle w| \right\|^2 \\ &= \sum_{\substack{w \in \mathcal{X} \cup \{\perp\} \setminus \{y_s \oplus x\} \\ w \notin \text{Im}(D_k)}} |\langle y_s \oplus x | \text{C}_{k,x} \cdot \text{SZ}_q |w\rangle|^2 \\ &\leq \mathbb{1}_{D \in \text{SZ}_q} \cdot \left(\frac{1}{|\mathcal{X}| - |D_k|} + \frac{|\mathcal{X}| - |D_k|}{(|\mathcal{X}| - |D_k|)^2} \right) \leq \frac{2}{|\mathcal{X}| - q} , \end{aligned} \quad (14)$$

where the last inequality is via $|D_k| \leq |D| \leq q$ for every $D \in \text{SZ}_q$, and the second to last inequality is via observing the fact that

$$\left| \langle y_s \oplus x | C_{k,x} \cdot \text{SZ}_q | \perp \rangle \right| \leq \frac{1}{\sqrt{|\mathcal{X}| - |D_k|}}, \quad \left| \langle y_s \oplus x | C_{k,x} \cdot \text{SZ}_q | w \rangle \right| \leq \frac{1}{|\mathcal{X}| - |D_k|}.$$

Now, observing that $P_{k,x}$ commutes with both SZ_q and PRMG_{y_s} , and that $(\mathbb{I} - \text{SZ}_{q+1}) \cdot C_{k,x} \cdot \text{SZ}_q = 0$, triangle inequality would give us that

$$\begin{aligned} & \left\| \text{PRMG}_{y_s} \cdot \text{clC}_{k,x}^{\rightarrow} \cdot \text{SZ}_q \cdot (\mathbb{I} - \text{PRMG}_{y_s}) \right\| \\ & \leq \left\| \text{PRMG}_{y_s} \cdot C_{k,x} \cdot \text{SZ}_q \cdot (\mathbb{I} - \text{PRMG}_{y_s}) \right\| + \left\| \text{PRMG}_{y_s} \cdot C_{k,x} \cdot \text{SZ}_{q+1} \cdot (\mathbb{I} - \text{PRMG}_{y_s}) \right\| \\ & \leq \sqrt{\frac{2}{|\mathcal{X}| - q}} + \sqrt{\frac{2}{|\mathcal{X}| - q - 1}} \leq \sqrt{\frac{8}{|\mathcal{X}| - q - 1}}. \end{aligned} \quad (15)$$

Likewise, by the symmetry of clC , we can similarly define $\text{clC}_{k,x,D}^{\leftarrow}$ in the obvious way, and have

$$\left\| \text{PRMG}_{y_s} \cdot \text{clC}_{k,x,D}^{\leftarrow} \cdot \text{SZ}_q \cdot (\mathbb{I} - \text{PRMG}_{y_s}) \right\| \leq \sqrt{\frac{8}{|\mathcal{X}| - q - 1}}. \quad (16)$$

Combining (15) and (16) and the fact that $\text{SZ}_q |\phi_i\rangle = |\phi_i\rangle$, we have that

$$\begin{aligned} \left\| \text{PRMG}_{y_s} |\phi_{i+1}\rangle \right\| & \leq \sum_{d \in \{\rightarrow, \leftarrow\}} \left\| \text{PRMG}_{y_s} \cdot (\text{clC}^d \otimes |d\rangle\langle d|) \cdot (\mathbb{I} - \text{PRMG}_{y_s}) |\phi_i\rangle \right\| + \left\| \text{PRMG}_{y_s} |\phi_i\rangle \right\| \\ & \leq \sqrt{\frac{32}{|\mathcal{X}| - q - 1}} + \left\| \text{PRMG}_{y_s} |\phi_i\rangle \right\| \end{aligned}$$

where the projector $|d\rangle\langle d|$ acts on a designated register where $\mathcal{A}$ indicates the direction of a clC query. Applying the inequality recursively would then give us

$$\left\| \text{PRMG}_{y_s} |\phi_{q+1}\rangle \right\| \leq q \sqrt{\frac{32}{|\mathcal{X}| - q - 1}}. \quad (17)$$

Therefore, in the invariant subspace of every fixed k, x again, we obtain

$$\begin{aligned} & \left\| \Pi^{y_s} \cdot \text{clC}^{\rightarrow} \cdot \text{SZ}_q \cdot (\mathbb{I} - \text{PRMG}_{y_s}) \cdot (|0\rangle\langle 0|)_{\mathcal{Y}} \right\| \\ & \leq \left\| \Pi^{y_s} \cdot (C_{k,x} \cdot P_{k,x} \cdot C_{k,x}) \cdot \text{SZ}_q \cdot (\mathbb{I} - \text{PRMG}_{y_s}) \cdot (|0\rangle\langle 0|)_{\mathcal{Y}} \right\| \\ & \leq \left\| \Pi^{y_s} \cdot P_{k,x} \cdot (\mathbb{I} - \text{PRMG}_{y_s}) \cdot \text{SZ}_q \cdot (|0\rangle\langle 0|)_{\mathcal{Y}} \right\| + \sqrt{\frac{8}{|\mathcal{X}| - q - 1}} = \sqrt{\frac{8}{|\mathcal{X}| - q - 1}}, \end{aligned} \quad (18)$$

where the second inequality is via (14) and the fact that Π^{y_s} and $(|0\rangle\langle 0|)_{\mathcal{Y}}$ commute with $C_{k,x}$, and the last equality is via the fact that $\Pi^{y_s} \cdot P_{k,x} \cdot (\mathbb{I} - \text{PRMG}_{y_s}) \cdot (|0\rangle\langle 0|)_{\mathcal{Y}} = 0$. Therefore we get

$$\begin{aligned} \left\| \Pi^{y_s} \cdot \text{clC}^{\rightarrow} |\phi_{q+1}\rangle \right\| & \leq \left\| \Pi^{y_s} \cdot \text{clC}^{\rightarrow} \cdot \text{SZ}_q |\phi_{q+1}\rangle \right\| \\ & \leq \left\| \Pi^{y_s} \cdot \text{clC}^{\rightarrow} \cdot \text{SZ}_q \cdot (\mathbb{I} - \text{PRMG}_{y_s}) |\phi_{q+1}\rangle \right\| + q \sqrt{\frac{32}{|\mathcal{X}| - q - 1}} \leq (2q + 1) \sqrt{\frac{8}{|\mathcal{X}| - q - 1}} \end{aligned}$$

where the first inequality is via $\text{SZ}_q |\phi_{q+1}\rangle = |\phi_{q+1}\rangle$, the second inequality is via (17), and the last inequality is via $(|0\rangle\langle 0|)_{\mathcal{Y}} |\phi_{q+1}\rangle = |\phi_{q+1}\rangle$ and (18). Combining the above with (11), (12) and (13), the proof is concluded. $\square$

5 Domain Extension of Quantum PRP

As an application of our generic lifting theorem, we show that, given a quantum-secure PRP $\text{PRP}_K : \{0, 1\}^n \rightarrow \{0, 1\}^n$ that operates on n -bit messages, where the subscript K denotes a PRP key of length λ , there is a quantum-secure PRP PRP'_K that operates on $2n$ -bit messages.

Theorem 4. Let $\text{PRP}_K : \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a quantum-secure strong PRP of key length λ . Then there is a quantum-secure strong PRP $\text{PRP}'_{K_3} : \{0, 1\}^{2n} \rightarrow \{0, 1\}^{2n}$ of key length $2\lambda \lceil \lambda/n \rceil$, which makes at most $2\lceil \lambda/n \rceil + 2$ queries to PRP , such that the following holds. For every attacker $\mathcal{A}$ making at most q bidirectional quantum queries, there is a similarly efficient attacker $\mathcal{B}$ making at most $6q$ quantum bidirectional queries such that

$$\text{Adv}_{\text{PRP}'}^{\text{RP}\leftrightarrow}(\mathcal{A}) \leq (72q^2 + 2\lceil \lambda/n \rceil) \text{Adv}_{\text{PRP}}^{\text{RP}\leftrightarrow}(\mathcal{B}) + \frac{1 + 11664\lceil \lambda/n \rceil q^3 + 32770q^4}{2^n}.$$

We achieve this via a three-step construction, where the preparation step follows from existing results, and each of the remaining two steps is via a generic transformation that may be of independent interest.

Preparation: from PRP to PRF. First, we derive a quantum-secure PRF $\text{PRF}_{K_1} : \{0, 1\} \times \{0, 1\}^n \rightarrow \{0, 1\}^\lambda$ of key length $\lambda_1 := 2\lambda \lceil \lambda/n \rceil$ bits. This step is trivial: for $\ell := \lceil \lambda/n \rceil$ and 2ℓ independent PRP keys $K_1 := (K_{1,b,i})_{b \in \{0,1\}, i \in [\ell]}$, we simply concatenate ℓ instances of PRP,

$$\text{PRF}_{K_1}(b, X) := \text{PRP}_{K_{1,b,1}}(X) \parallel \dots \parallel \text{PRP}_{K_{1,b,\ell}}(X), \quad (19)$$

and then truncate the output to λ bits. Zhandry's PRP-PRF switching lemma [Zha12] implies this is indeed a secure PRF. For every attacker making q quantum queries, the precise bound of this step is

$$\text{Adv}_{\text{PRF}}^{\text{RO}}(\mathcal{A}) \leq 2\lceil \lambda/n \rceil \text{Adv}_{\text{PRP}}^{\text{RP}\leftrightarrow}(\mathcal{B}) + \frac{54\lceil \lambda/n \rceil q^3}{2^n}, \quad (20)$$

for a similarly efficient attacker $\mathcal{B}$ making at most q quantum queries, where the constant 27 in each of the $2\lceil \lambda/n \rceil$ switching steps comes from the constant $(8\pi^2/3)$ spelled out concretely in [Xag24, Theorem 1], and the factor $2\lceil \lambda/n \rceil$ comes from a hybrid argument over the $2\lceil \lambda/n \rceil$ independent keys.

Step 1: from PRP and PRF to tweakable cipher. Second, we derive a quantum-secure tweakable cipher $\text{TWK}_{K_2} : (\{0, 1\} \times \{0, 1\}^n) \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ with tweak space $\{0, 1\} \times \{0, 1\}^n$ and key length $\lambda_2 := \lambda_1$ bits, from PRP and PRF as follows:

$$\text{TWK}_{K_2}(t, x) := \text{PRP}_{\text{PRF}_{K_1}(t)}(x), \quad (21)$$

where K_2 is simply the PRF key K_1 . It simply uses PRF to generate a key for each tweak t , and the goal here is to show that this construction is indistinguishable from a quantum ideal cipher oracle, which we will argue in Section 5.1 using our indistinguishability lifting theorem Theorem 1.

Step 2: from tweakable cipher to PRP of double message length. Finally, we derive a quantum-secure PRP $\text{PRP}'_{K_3} : \{0, 1\}^{2n} \rightarrow \{0, 1\}^{2n}$ of key length $\lambda_3 := \lambda_2$, from $\text{TWK} : (\{0, 1\} \times \{0, 1\}^n) \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ as follows:

$$\text{PRP}'_{K_3}(L, R) := (\text{TWK}_{K_3}((1, U), L), U), \quad U := \text{TWK}_{K_3}((0, L), R), \quad (22)$$

where K_3 is a single TWK key, and the two rounds are separated by the first bit of the tweak. In Section 5.2, we show via the polynomial method that the construction above is quantum indistinguishable from a random permutation.

Combining (20) with Lemmas 4 and 5 below proves the theorem. Indeed, the attacker against TWK in Lemma 5 makes at most $3q$ queries, so the attacker against PRF in Lemma 4 makes at most $6q$ queries and the attacker against PRP in Lemma 4 makes at most $3q$ queries. This gives

$$\begin{aligned} \text{Adv}_{\text{PRP}'}^{\text{RP}\leftrightarrow}(\mathcal{A}) &\leq 2\lceil \lambda/n \rceil \text{Adv}_{\text{PRP}}^{\text{RP}\leftrightarrow}(\mathcal{B}) + \frac{54\lceil \lambda/n \rceil (6q)^3}{2^n} \\ &\quad + 8(3q)^2 \text{Adv}_{\text{PRP}}^{\text{RP}\leftrightarrow}(\mathcal{B}) + 2^{-n} + \frac{32770q^4}{2^n}, \end{aligned}$$

where $\mathcal{B}$ is the one of the two PRP attackers with the larger advantage, and makes at most $6q$ queries. The key length is $\lambda_3 = \lambda_1 = 2\lambda \lceil \lambda/n \rceil$, and each evaluation of PRP' makes two evaluations of TWK, each consisting of $\lceil \lambda/n \rceil$ calls to PRP for PRF and one further call to PRP.

Remark 3. For simplicity we only consider doubling the message-length. It is possible to further extend our technique toward extending input length to k times. We simply need a PRF of input size of $O(kn)$ bits, which would allow us to construct the tweakable cipher TWK of tweak length $O(kn)$ bits in Step 1, thereby allowing us to perform a recursion of depth $O(\log k)$ in Step 2.

5.1 A Generic Tweakable Cipher Construction

In the statement below, let $\text{PRP}_K : \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a quantum-secure strong PRP of key length λ bits, and $\text{PRF}_{K_1} : \{0, 1\}^m \rightarrow \{0, 1\}^\lambda$ be a quantum secure PRF of key length λ_1 bits. (In our application, $m = n + 1$, where we identify $\{0, 1\} \times \{0, 1\}^n$ with $\{0, 1\}^{n+1}$.) We show that the tweakable cipher construction in (21):

$$\text{TWK}_{K_2}(t, x) := \text{PRP}_{\text{PRF}_{K_1}(t)}(x) ,$$

where $K_2 := K_1 \leftarrow \{0, 1\}^{\lambda_1}$, is also quantum-secure.

Lemma 4. *For every oracle algorithm $\mathcal{A}$ making at most q quantum bidirectional queries, there exist similarly efficient oracle algorithms $\mathcal{B}$ and $\mathcal{C}$ making at most $2q$ and q quantum queries respectively, such that*

$$\text{Adv}_{\text{TWK}}^{\text{IC} \leftrightarrow}(\mathcal{A}) \leq \text{Adv}_{\text{PRF}}^{\text{RO}}(\mathcal{B}) + 8q^2 \cdot \text{Adv}_{\text{PRP}}^{\text{RP} \leftrightarrow}(\mathcal{C}) + 2^{-n} .$$

Here, queries made by $\mathcal{A}$ and $\mathcal{C}$ are allowed to be bidirectional.

Proof. Let $H : \{0, 1\}^m \rightarrow \{0, 1\}^\lambda$ be a random oracle, $P : \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a random permutation, and $P^{\{0, 1\}^m} : \{0, 1\}^m \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be an ideal cipher. The goal of our proof is to show that, for $K_2 := K_1 \leftarrow \{0, 1\}^{\lambda_1}$, we have

$$\text{Adv}_{\text{TWK}}^{\text{IC} \leftrightarrow}(\mathcal{A}) = \text{SD} \left(\mathcal{A}^{\text{TWK}_{K_2}}, \mathcal{A}^{P^{\{0, 1\}^m}} \right)$$

To that end, consider an intermediate hybrid oracle, defined as follows

$$\text{Hyb}(t, x) := \text{PRP}_{H(t)}(x) ,$$

with inverse queries defined in the obvious way. Now, immediately from the quantum security of PRF, we can switch out the PRF underlying TWK and hence obtain

$$\text{SD} \left(\mathcal{A}^{\text{TWK}_{K_2}}, \mathcal{A}^{\text{Hyb}} \right) \leq \text{Adv}_{\text{PRF}}^{\text{RO}}(\mathcal{B}) ,$$

for a similarly efficient oracle algorithm $\mathcal{B}$. Since in each query that $\mathcal{A}$ makes, $\mathcal{B}$ would need to make 2 queries to the underlying PRF/ H oracle, it sums up to $2q$ queries in total.

Next, we note that the oracle Hyb can be regarded as the oracle $\text{PRP}^{\{0, 1\}^m}$ that samples an exponential family of PRP keys $K_t \leftarrow \{0, 1\}^\lambda$, one for each tweak t independently and then responds with $\text{PRP}_{K_t}(x)$ on input (t, x) . Hence, by applying our indistinguishability lifting theorem (Theorem 1), we obtain

$$\begin{aligned} \text{SD} \left(\mathcal{A}^{\text{Hyb}}, \mathcal{A}^{P^{\{0, 1\}^m}} \right) &= \text{SD} \left(\mathcal{A}^{\text{PRP}^{\{0, 1\}^m}}, \mathcal{A}^{P^{\{0, 1\}^m}} \right) \\ &\leq 8q^2 \cdot \text{SD} \left(\mathcal{D}^{\text{PRP}, \text{PRP}^{\{0, 1\}^m}, P^{\{0, 1\}^m}}, \mathcal{D}^{P, \text{PRP}^{\{0, 1\}^m}, P^{\{0, 1\}^m}} \right) \\ &\leq 8q^2 \cdot (\text{SD}(\mathcal{C}^{\text{PRP}}, \mathcal{C}^P) + 2 \cdot 2^{-n} / (16q^2)) = 8q^2 \cdot \text{Adv}_{\text{PRP}}^{\text{RP} \leftrightarrow}(\mathcal{C}) + 2^{-n} , \end{aligned}$$

for a similarly efficient distinguishing algorithm $\mathcal{D}$, and the final reduction algorithm $\mathcal{C}$ that internally simulates the oracle queries to $\text{PRP}^{\{0, 1\}^m}$ and $P^{\{0, 1\}^m}$ up to a small additive error. Concretely, similar to what was pointed out in [HHM⁺26, Appendix A], we can simulate a random permutation up to a δ error within a runtime overhead of order $O(q^2 \cdot n^2 (n + \log(1/\delta))^2)$, counting basic arithmetics as unit runtime. Picking $\delta := 2^{-(n+m)} / (16q^2)$ would make the simulation of the entire family $P^{\{0, 1\}^m}$ of error bounded by $2^{-n} / (16q^2)$, within a running time overhead of order $O(q^2 n^2 (n + m + \log q)^2)$.

Collecting the bounds so far, the proof is concluded. $\square$

5.2 From Tweakable Cipher to Double-Input PRP

In the statement below, let $\text{TWK}_{K_3} : (\{0, 1\} \times \{0, 1\}^n) \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a quantum-secure tweakable cipher with tweak space $\{0, 1\} \times \{0, 1\}^n$ and key length λ_3 bits. We show that the strong PRP construction described in Eq. (22), which we recall below:

$$\text{PRP}'_{K_3}(L, R) := (\text{TWK}_{K_3}((1, U), L), U) , \quad U := \text{TWK}_{K_3}((0, L), R) ,$$

is also quantum-secure. Here, both rounds use the same key K_3 , and are separated by the first bit of the tweak.

Lemma 5. For every oracle algorithm $\mathcal{A}$ making at most q quantum bidirectional queries, there exists a similarly efficient oracle algorithm $\mathcal{B}$ making at most $3q$ quantum bidirectional queries, such that

$$\mathbf{Adv}_{\text{PRP}'}^{\text{RP} \leftrightarrow}(\mathcal{A}) \leq \mathbf{Adv}_{\text{TWK}}^{\text{IC} \leftrightarrow}(\mathcal{B}) + \frac{32770q^4}{2^n}.$$

Proof. Let $\mathcal{B}^E$ run $\mathcal{A}$ and answer each of its queries by evaluating PRP' with TWK_{K_3} replaced by its oracle E . A forward query (L, R) is answered by computing $U := E((0, L), R)$ into an ancilla, XOR-ing $E((1, U), L)$ and U into the answer register, and uncomputing U ; an inverse query (V, U) is answered symmetrically via $L := E^{-1}((1, U), V)$ and $R := E^{-1}((0, L), U)$. Since the direction bit only selects which tweak, input, and direction are loaded into the query registers, $\mathcal{B}$ makes 3 bidirectional queries to E per query of $\mathcal{A}$, and has a similar running time as $\mathcal{A}$ does. Clearly $\mathcal{B}^{\text{TWK}_{K_3}}$ behaves identically to $\mathcal{A}^{\text{PRP}'_{K_3}}$. On the other hand, let $\text{IC} \leftarrow \text{Sym}(\{0, 1\}^n)^{\{0, 1\} \times \{0, 1\}^n}$ be an ideal cipher with tweak space $\{0, 1\} \times \{0, 1\}^n$, and let $G(L, R) := \text{IC}((0, L), R)$ and $H(U, L) := \text{IC}((1, U), L)$. Then G and H are two independent ideal ciphers over the message space $\{0, 1\}^n$ and tweak space $\{0, 1\}^n$, and $\mathcal{B}^{\text{IC}}$ behaves identically to $\mathcal{A}^{F_{G,H}}$, where

$$F_{G,H}(L, R) := (H(G(L, R), L), G(L, R)).$$

Therefore, for a random permutation $\text{RP} \leftarrow \text{Sym}(\{0, 1\}^{2n})$, the triangle inequality gives

$$SD(\mathcal{A}^{\text{PRP}'_{K_3}}, \mathcal{A}^{\text{RP}}) \leq \mathbf{Adv}_{\text{TWK}}^{\text{IC} \leftrightarrow}(\mathcal{B}) + SD(\mathcal{A}^{F_{G,H}}, \mathcal{A}^{\text{RP}}). \quad (23)$$

For the remainder of the proof we focus on bounding the last term of (23). We will show that this term is bounded by $32770q^4/2^n$, which will complete the proof.

Output probability for a fixed permutation. We first observe that there is a multilinear polynomial a of degree at most $2q$ in the variables $(\mathbb{1}_{\pi(x)=y})_{x,y \in \{0,1\}^{2n}}$ such that, for every fixed permutation $\pi \in \text{Sym}(\{0, 1\}^{2n})$,

$$\Pr[\mathcal{A}^\pi] = a\left((\mathbb{1}_{\pi(x)=y})_{x,y}\right).$$

Indeed, each query implements the unitary

$$|x, y\rangle \mapsto |x, y \oplus \pi^{\pm 1}(x)\rangle = \sum_z \mathbb{1}_{\pi^{\pm 1}(x)=y \oplus z} |x, z\rangle,$$

where for inverse queries $\mathbb{1}_{\pi^{-1}(x)=y \oplus z} = \mathbb{1}_{\pi(y \oplus z)=x}$, so the entries of this unitary are linear in these variables. Hence, after q queries, every amplitude is a polynomial of degree at most q , and the output probability, being a sum of squared magnitudes of amplitudes, is a polynomial of degree at most $2q$. Since the variables only take values in $\{0, 1\}$, we may take this polynomial to be multilinear.

Output probability for a rerandomized permutation. Now, consider the following rerandomization procedure $\pi' \leftarrow \text{R}(\pi)$: it independently samples $P, Q \leftarrow \text{Sym}(\{0, 1\}^n)^{\{0, 1\}^n}$ and produces

$$\pi'(L, R) := (Q(U, V), U), \quad (V, U) := \pi(L, P(L, R)).$$

In other words, $\pi' = \tau \circ \pi \circ \sigma$ for the random permutations $\sigma(L, R) := (L, P(L, R))$ and $\tau(V, U) := (Q(U, V), U)$ of $\{0, 1\}^{2n}$. We have

$$\Pr[\mathcal{A}^{\text{RP}}] = \Pr[\mathcal{A}^{\text{R}(\text{RP})}] , \quad \Pr[\mathcal{A}^{F_{G,H}}] = \Pr[\mathcal{A}^{\text{R}(F_{G,H})}]. \quad (24)$$

Indeed, $\tau \circ \text{RP} \circ \sigma$ is again a uniformly random permutation, since σ and τ are independent of RP . Moreover, $\text{R}(F_{G,H}) = F_{G',H'}$, where $G'(L, \cdot) := G(L, \cdot) \circ P(L, \cdot)$ and $H'(U, \cdot) := Q(U, \cdot) \circ H(U, \cdot)$, and G' and H' are again two independent ideal ciphers.

Next, for every fixed permutation $\pi \in \text{Sym}(\{0, 1\}^{2n})$ and for $\pi' \leftarrow \text{R}(\pi)$, we have

$$\begin{aligned} \Pr[\mathcal{A}^{\text{R}(\pi)}] &= \mathbb{E}\left[a\left((\mathbb{1}_{\pi'(x)=y})_{x,y}\right)\right] \\ &= \mathbb{E}\left[\sum_S \widehat{a}(S) \prod_{(x,y) \in S} \mathbb{1}_{\pi'(x)=y}\right] \\ &= \sum_S \widehat{a}(S) \cdot \Pr\left[\forall ((L, R), (V, U)) \in S \right. \\ &\quad \left. \pi'(L, R) = (V, U) \right], \end{aligned} \quad (25)$$

where the first equality applies the previous paragraph to π' , the second expands a into monomials, and the third uses linearity of expectation. Here, the summation runs over all

$$S = \{(x_1, y_1), \dots, (x_k, y_k)\} \subseteq \{0, 1\}^{2n} \times \{0, 1\}^{2n}$$

of size at most $2q$ such that the x_i 's are all distinct and the y_i 's are all distinct,⁷ and $\widehat{a}(S) \in \mathbb{R}$ is the coefficient corresponding to the monomial $\prod_{(x,y) \in S} \mathbb{1}_{\pi'(x)=y}$.

A matrix C^π for counting. Towards calculating (25), for every fixed permutation $\pi \in \text{Sym}(\{0, 1\}^{2n})$, let C^π be a matrix whose rows and columns are indexed by $L \in \{0, 1\}^n$ and $U \in \{0, 1\}^n$, respectively, with each entry $C_{L,U}^\pi$ defined as follows:

$$C_{L,U}^\pi := \left| \{R \in \{0, 1\}^n \mid \pi(L, R) \in \{0, 1\}^n \times \{U\}\} \right|.$$

Now fix a set $S = \{(x_1, y_1), \dots, (x_k, y_k)\}$ as in (25), and write $(L_i, R_i) := x_i$ and $(V_i, U_i) := y_i$ for each i (where $L_i, R_i, V_i, U_i \in \{0, 1\}^n$), and let

$$e_{L,U} := \left| \{i \mid (L, U) = (L_i, U_i)\} \right|, \quad |e_{L,\star}| := \sum_U e_{L,U}, \quad |e_{\star,U}| := \sum_L e_{L,U}.$$

We will show that

$$\Pr_{\pi' \leftarrow \mathbb{R}(\pi)} \left[\forall ((L, R), (V, U)) \in S \quad \pi'(L, R) = (V, U) \right] = \frac{\prod_{L,U \in \{0,1\}^n} (C_{L,U}^\pi)^{e_{L,U}}}{\prod_{L \in \{0,1\}^n} (2^n)^{|e_{L,\star}|} \prod_{U \in \{0,1\}^n} (2^n)^{|e_{\star,U}|}}. \quad (26)$$

For each i , let $\tilde{R}_i := P(L_i, R_i)$, and let $\tilde{V}_i$ denote the first half of $\pi(L_i, \tilde{R}_i)$. By definition of π' , we have $\pi'(x_i) = y_i$ if and only if the second half of $\pi(L_i, \tilde{R}_i)$ equals U_i and $Q(U_i, \tilde{V}_i) = V_i$. First fix $L \in \{0, 1\}^n$. The R_i with $L_i = L$ are distinct and $P(L, \cdot)$ is a uniform permutation, so $(\tilde{R}_i)_{i:L_i=L}$ is a uniformly random sequence of $|e_{L,\star}|$ distinct elements of $\{0, 1\}^n$. Among the $(2^n)^{|e_{L,\star}|}$ such sequences, exactly $\prod_U (C_{L,U}^\pi)^{e_{L,U}}$ satisfy $\pi(L, \tilde{R}_i) \in \{0, 1\}^n \times \{U_i\}$ for all i with $L_i = L$, since for each U the $e_{L,U}$ values $\tilde{R}_i$ with $U_i = U$ must be distinct elements of the set counted by $C_{L,U}^\pi$. Next fix $U \in \{0, 1\}^n$, and condition on P . The pairs $(L_i, \tilde{R}_i)$ are distinct, so the values $\tilde{V}_i$ with $U_i = U$ are distinct since π is injective, and so are the V_i with $U_i = U$ since the y_i are distinct. Since $Q(U, \cdot)$ is a uniform permutation independent of P , the event $Q(U, \tilde{V}_i) = V_i$ for all i with $U_i = U$ has probability $1/(2^n)^{|e_{\star,U}|}$. Multiplying over all L and U , using the independence of the permutations $P(L, \cdot)$ and $Q(U, \cdot)$, gives (26).

Therefore, combined with (25), we have

$$\Pr[\mathcal{A}^{\mathbb{R}(\pi)}] = \sum_S \widehat{a}(S) \cdot \frac{\prod_{L,U \in \{0,1\}^n} (C_{L,U}^\pi)^{e_{L,U}}}{\prod_{L \in \{0,1\}^n} (2^n)^{|e_{L,\star}|} \prod_{U \in \{0,1\}^n} (2^n)^{|e_{\star,U}|}} =: g(C^\pi), \quad (27)$$

which is a polynomial of degree at most $2q$ in the entries of C^π , since $\sum_{L,U} e_{L,U} = |S| \leq 2q$.

Some useful properties of C^π . Before moving ahead, we establish some basic facts about the matrix C^π , whose proofs are deferred to Section 5.3.

Claim 1. The transformation $\pi \mapsto C^\pi$ forms a surjection from

$$\text{Sym}(\{0, 1\}^{2n}) \rightarrow \mathcal{C},$$

where $\mathcal{C}$ is the set of all $2^n \times 2^n$ matrices over $\mathbb{Z}_{\geq 0}$ whose row sums and column sums are all equal to 2^n .

Claim 2. For every $G, H \in \text{Sym}(\{0, 1\}^n)^{\{0,1\}^n}$, the matrix $C^{F_{G,H}}$ coincides with the all-ones matrix J , i.e., $J_{L,U} := 1$ for all $L, U \in \{0, 1\}^n$.

⁷ Note that the probability in (25) would have vanished if S cannot be completed to a permutation.

An auxiliary polynomial $f(r)$. Combining (24), (27), and Claim 2, we obtain

$$SD(\mathcal{A}^{F_{G,H}}, \mathcal{A}^{RP}) = |\mathbb{E}[g(C^{F_{G,H}})] - \mathbb{E}[g(C^{RP})]| = |g(J) - \mathbb{E}[g(C^{RP})]|, \quad (28)$$

where the first equality uses (24) together with $\Pr[\mathcal{A}^{R(\pi)}] = g(C^\pi)$ from (27), and the second uses Claim 2. We recall from (27) that g is a polynomial of degree at most $2q$. Hence, with the monomial notation

$$C^{\diamond w} := \prod_{L,U} C_{L,U}^{w_{L,U}},$$

we can write

$$g(C) = \sum_w \widehat{g}(w) C^{\diamond w},$$

where w ranges over $\mathbb{Z}_{\geq 0}^{2^n \times 2^n}$ with $|w| \leq 2q$ and $\widehat{g}(w) \in \mathbb{R}$.

We use the following notation for matrices $A, w \in \mathbb{Z}_{\geq 0}^{2^n \times 2^n}$ (with rows and columns indexed by $\{0, 1\}^n$): $|A_{L,\star}| := \sum_U A_{L,U}$, $|A_{\star,U}| := \sum_L A_{L,U}$, $|A| := \sum_{L,U} A_{L,U}$, and $A \leq w$ if $A_{L,U} \leq w_{L,U}$ for all L, U . We also write

$$\left\{ \begin{matrix} w \\ A \end{matrix} \right\} := \prod_{L,U} \left\{ \begin{matrix} w_{L,U} \\ A_{L,U} \end{matrix} \right\}.$$

For real numbers $x, r \in \mathbb{R}$ and an integer $k \geq 0$ we write $(x|r)_k := \prod_{0 \leq j < k} (x - jr)$, so that $(x|1)_k = (x)_k$ and $(x|0)_k = x^k$. Writing

$$\begin{aligned} \gamma_A(r) &:= \frac{1}{(2^{2n}|r|_{|A|})} \prod_L (2^n|r)_{|A_{L,\star}|} \prod_U (2^n|r)_{|A_{\star,U}|}, \\ \tau_{w,A}(r) &:= \frac{(2^{2n}|r|_{2q})}{2^{4nq}} \widehat{g}(w) \left\{ \begin{matrix} w \\ A \end{matrix} \right\} r^{|w|-|A|} \gamma_A(r), \end{aligned}$$

we define the function f as

$$f(r) := \sum_w \sum_{A \leq w} \tau_{w,A}(r). \quad (29)$$

We make the following claim about f , which we prove in Section 5.3.

Claim 3. The function f defined in (29) is a polynomial of degree at most $4q$, and it satisfies

$$f(0) = g(J), \quad f(1) = \prod_{0 \leq j < 2q} \left(1 - \frac{j}{2^{2n}}\right) \cdot \mathbb{E}[g(C^{RP})], \quad (30)$$

and $|f(r)| \leq 2$ for every $r \in \{0, \dots, \lfloor 2^n/(256q^2) \rfloor\}$.

Assume $2^{2n} \geq 2q$, since otherwise $32770q^4/2^n > 1$ and the lemma holds trivially. Then every factor $1 - j/2^{2n}$ with $0 \leq j < 2q$ is positive, so we have

$$\begin{aligned} SD(\mathcal{A}^{F_{G,H}}, \mathcal{A}^{RP}) &= \left| f(0) - \frac{f(1)}{\prod_{0 \leq j < 2q} \left(1 - \frac{j}{2^{2n}}\right)} \right| \\ &\leq |f(0) - f(1)| + |f(1)| \cdot \left(\frac{1}{\prod_{0 \leq j < 2q} \left(1 - \frac{j}{2^{2n}}\right)} - 1 \right) \\ &= |f(0) - f(1)| + \mathbb{E}[g(C^{RP})] \cdot \left(1 - \prod_{0 \leq j < 2q} \left(1 - \frac{j}{2^{2n}}\right) \right) \\ &\leq |f(0) - f(1)| + \sum_{0 \leq j < 2q} \frac{j}{2^{2n}} \leq |f(0) - f(1)| + \frac{2q^2}{2^{2n}}. \end{aligned} \quad (31)$$

The first inequality follows from triangle inequality, the second inequality follows from the fact that $0 \leq \mathbb{E}[g(C^{RP})] \leq 1$ together with $1 - \prod_j (1 - a_j) \leq \sum_j a_j$ for $a_j \in [0, 1]$, the last inequality follows from $\sum_{0 \leq j < 2q} j = q(2q - 1) \leq 2q^2$, and the two equalities follow from Eq. (30).

Bounding $|f(0) - f(1)|$. To complete the proof, it suffices to bound $|f(0) - f(1)|$, for which we use the Markov brothers' inequality twice. We will show that $|f(0) - f(1)| \leq 32768q^4/2^n$ whenever $2^n \geq 16384q^4$. Let $T := 2^n/(256q^2)$, so that $T \geq 64q^2 \geq 1$ under this assumption. By the mean value theorem and (2) on the interval $[0, T] \supseteq [0, 1]$, we have

$$|f(0) - f(1)| \leq \max_{r \in [0,1]} |f'(r)| \leq \frac{2(\deg f)^2}{T} \max_{r \in [0,T]} |f(r)| \leq \frac{32q^2}{T} \max_{r \in [0,T]} |f(r)| ,$$

where the last step uses $\deg f \leq 4q$ from Claim 3. To bound the remaining maximum, we apply (5) to f on the interval $[0, T]$, which is possible since $T \geq 64q^2 \geq 4(\deg f)^2$. The integers in $[0, T]$ are exactly $r = 0, \dots, \lfloor T \rfloor$, and at each of them $|f(r)| \leq 2$ by Claim 3. Therefore

$$\max_{r \in [0,T]} |f(r)| \leq 2 \max_{r \in [0,T] \cap \mathbb{Z}} |f(r)| \leq 4 ,$$

and hence $|f(0) - f(1)| \leq 128q^2/T = 32768q^4/2^n$. Combining this with (31), we get

$$SD(\mathcal{A}^{F_{G,H}}, \mathcal{A}^{\text{RP}}) \leq |f(0) - f(1)| + \frac{2q^2}{2^n} \leq \frac{32768q^4}{2^n} + \frac{2q^2}{2^n} \leq \frac{32770q^4}{2^n}$$

whenever $2^n \geq 16384q^4$. Otherwise, the right-hand side exceeds 1, and the bound holds trivially. Plugging this into (23), we conclude that

$$\text{Adv}_{\text{PRP}}^{\text{RP} \leftrightarrow}(\mathcal{A}) \leq \text{Adv}_{\text{TWK}}^{\text{IC} \leftrightarrow}(\mathcal{B}) + \frac{32770q^4}{2^n} ,$$

which completes the proof of Lemma 5. $\square$

5.3 Remaining Proofs of the Claims

Claim 1. The transformation $\pi \mapsto C^\pi$ forms a surjection from

$$\text{Sym}(\{0, 1\}^{2^n}) \rightarrow \mathcal{C} ,$$

where $\mathcal{C}$ is the set of all $2^n \times 2^n$ matrices over $\mathbb{Z}_{\geq 0}$ whose row sums and column sums are all equal to 2^n .

Proof of Claim 1. We first show that $C^\pi \in \mathcal{C}$ for every $\pi \in \text{Sym}(\{0, 1\}^{2^n})$. For every U , the column sum $\sum_L C_{L,U}^\pi$ is the number of pairs (L, R) such that $\pi(L, R) \in \{0, 1\}^n \times \{U\}$, which is 2^n since π is a bijection. Likewise, for every L , the row sum $\sum_U C_{L,U}^\pi$ is the number of R such that $\pi(L, R) \in \{0, 1\}^{2^n}$, which is 2^n .

We now show surjectivity: given any matrix $C \in \mathcal{C}$, we construct $\pi \in \text{Sym}(\{0, 1\}^{2^n})$ such that $C^\pi = C$. First, for every $L \in \{0, 1\}^n$, we pick a partition $\{S_{L,U}\}_{U \in \{0,1\}^n}$ of $\{0, 1\}^n$ such that $|S_{L,U}| = C_{L,U}$, which exists since the entries of row L of C are nonnegative integers summing to 2^n . Likewise, for every $U \in \{0, 1\}^n$, we pick a partition $\{S'_{L,U}\}_{L \in \{0,1\}^n}$ of $\{0, 1\}^n$ such that $|S'_{L,U}| = C_{L,U}$ for every L , which exists since column U of C sums to 2^n . Next, for every $L, U \in \{0, 1\}^n$ we pick any bijection

$$\pi_{L,U} : \{L\} \times S_{L,U} \rightarrow S'_{L,U} \times \{U\} ,$$

which exists since $|S_{L,U}| = C_{L,U} = |S'_{L,U}|$. We then assemble these bijections into π via

$$\pi(L, R) := \pi_{L,U}(L, R) ,$$

where U is the unique string such that $R \in S_{L,U}$; it is unique since $\{S_{L,U}\}_U$ is a partition of $\{0, 1\}^n$. To see that π is a bijection, note that for every U the sets $\{S'_{L,U}\}_L$ form a partition of $\{0, 1\}^n$, so every $(V, U) \in \{0, 1\}^{2^n}$ lies in exactly one of the sets $\text{Im}(\pi_{L,U}) = S'_{L,U} \times \{U\}$, and each $\pi_{L,U}$ is a bijection onto its image. Finally, by construction exactly $|S_{L,U}| = C_{L,U}$ inputs from row L are mapped into column U , and hence $C^\pi = C$. $\square$

Claim 2. For every $G, H \in \text{Sym}(\{0, 1\}^n)^{\{0,1\}^n}$, the matrix $C^{F_{G,H}}$ coincides with the all-ones matrix J , i.e., $J_{L,U} := 1$ for all $L, U \in \{0, 1\}^n$.

Proof of Claim 2. Fix $L, U \in \{0, 1\}^n$. The second half of $F_{G,H}(L, R)$ equals U exactly when $G(L, R) = U$. Since $G(L, \cdot)$ is a permutation, there is exactly one such R , namely $R = G^{-1}(L, U)$, where $G^{-1}(L, \cdot)$ denotes the inverse of $G(L, \cdot)$. Thus $C_{L,U}^{F_{G,H}} = 1$ for every L, U , and hence $C^{F_{G,H}} = J$. $\square$

Claim 3. The function f defined in (29) is a polynomial of degree at most $4q$, and it satisfies

$$f(0) = g(J), \quad f(1) = \prod_{0 \leq j < 2q} \left(1 - \frac{j}{2^{2n}}\right) \cdot \mathbb{E}[g(C^{\text{RP}})], \quad (30)$$

and $|f(r)| \leq 2$ for every $r \in \{0, \dots, \lfloor 2^n/(256q^2) \rfloor\}$.

Proof of Claim 3. In addition to the notation introduced before (29), for matrices $A, B, w \in \mathbb{Z}_{\geq 0}^{2^n \times 2^n}$ and any matrix C we write

$$\begin{pmatrix} w \\ B \end{pmatrix} := \prod_{L,U} \begin{pmatrix} w_{L,U} \\ B_{L,U} \end{pmatrix}, \quad (C)_{\circ A} := \prod_{L,U} (C_{L,U})_{A_{L,U}}.$$

Since $|A| \leq 2q$, each denominator $(2^{2n} |r|)_{|A|}$ in $\gamma_A(r)$ divides $(2^{2n} |r|)_{2q}$, so each $\tau_{w,A}$ is a polynomial in r , and hence so is f . More precisely, using

$$(2^{2n} |r|)_{2q} = (2^{2n} |r|)_{|A|} \prod_{|A| \leq j < 2q} (2^{2n} - jr),$$

we can write

$$\tau_{w,A}(r) = \frac{\widehat{g}(w)}{2^{4nq}} \begin{Bmatrix} w \\ A \end{Bmatrix} r^{|w|-|A|} \prod_L (2^n |r|)_{|A_{L,\star}|} \prod_U (2^n |r|)_{|A_{\star,U}|} \prod_{|A| \leq j < 2q} (2^{2n} - jr). \quad (32)$$

Since $\sum_L |A_{L,\star}| = \sum_U |A_{\star,U}| = |A|$, each $\tau_{w,A}$ has degree at most $(|w| - |A|) + 2|A| + (2q - |A|) = 2q + |w| \leq 4q$ in r , so $\deg f \leq 4q$. Applying (1) entrywise gives

$$C^{\circ w} = \prod_{L,U} \left(\sum_{A_{L,U} \leq w_{L,U}} \begin{Bmatrix} w_{L,U} \\ A_{L,U} \end{Bmatrix} (C_{L,U})_{A_{L,U}} \right) = \sum_{A \leq w} \left(\prod_{L,U} \begin{Bmatrix} w_{L,U} \\ A_{L,U} \end{Bmatrix} (C_{L,U})_{A_{L,U}} \right) = \sum_{A \leq w} \begin{Bmatrix} w \\ A \end{Bmatrix} (C)_{\circ A}. \quad (33)$$

The value $f(0)$. With the convention $0^0 = 1$, $\tau_{w,A}(0) = 0$ unless $|A| = |w|$, and together with $A \leq w$ this forces $A = w$. Since $\{w\} = 1$ and $(2^n |0|)_k = 2^{nk}$,

$$\tau_{w,w}(0) = \frac{\widehat{g}(w)}{2^{4nq}} \cdot 2^{n|w|} \cdot 2^{n|w|} \cdot 2^{2n(2q-|w|)} = \widehat{g}(w).$$

Hence $f(0) = \sum_w \widehat{g}(w) = g(J)$, since every monomial $C^{\circ w}$ equals 1 at $C = J$.

The matrices C_t . For an integer $t \geq 1$, let $\sigma_t : \{0, 1\}^n \times [t] \rightarrow [t] \times \{0, 1\}^n$ be a uniformly random bijection, and let C_t be the $2^n \times 2^n$ matrix with entries

$$C_{t,L,U} := \left| \{i \in [t] \mid \sigma_t(L, i) \in [t] \times \{U\}\} \right|.$$

Every row and column of C_t sums to t . Identifying $[2^n]$ with $\{0, 1\}^n$, σ_t is a uniform element of $\text{Sym}(\{0, 1\}^{2n})$ when $t = 2^n$, so C_{2^n} is distributed as C^{RP} . We show that for every $A \in \mathbb{Z}_{\geq 0}^{2^n \times 2^n}$,

$$(2^n t)_{|A|} \cdot \mathbb{E}[(C_t)_{\circ A}] = \prod_L (t)_{|A_{L,\star}|} \prod_U (t)_{|A_{\star,U}|}. \quad (34)$$

If $|A_{L,\star}| > t$ for some L , both sides vanish: the right-hand side contains $(t)_{|A_{L,\star}|} = 0$, and since row L of C_t sums to $t < |A_{L,\star}|$, some $A_{L,U} > C_{t,L,U}$, so $(C_{t,L,U})_{A_{L,U}} = 0$. The same holds if $|A_{\star,U}| > t$ for some U . Otherwise $|A| \leq 2^n t$. Write $\text{In}_L := [L] \times [t]$ and $\text{Out}_U := [t] \times \{U\}$. The factor $(C_{t,L,U})_{A_{L,U}}$ counts the ordered $A_{L,U}$ -tuples of distinct elements of In_L mapped by σ_t into Out_U . Elements counted for different (L, U) are distinct, since they lie in different In_L or are mapped into different Out_U . Hence $(C_t)_{\circ A}$ is the number of pairs $(\mathbf{x}, \mathbf{y})$ of families $\mathbf{x} = (x_{L,U,i})$ and $\mathbf{y} = (y_{L,U,i})$, indexed by L, U and $i \leq A_{L,U}$, such that the $x_{L,U,i} \in \text{In}_L$ are pairwise distinct, the $y_{L,U,i} \in \text{Out}_U$ are pairwise distinct, and $\sigma_t(x_{L,U,i}) = y_{L,U,i}$ for all indices. There are $\prod_L (t)_{|A_{L,\star}|}$ choices of $\mathbf{x}$ and $\prod_U (t)_{|A_{\star,U}|}$ choices of $\mathbf{y}$, and for each such pair, the uniform bijection σ_t on a set of size $2^n t$ satisfies the $|A|$ prescribed equalities with probability $1/(2^n t)_{|A|}$. Linearity of expectation gives (34).

The value $f(1)$. Fix w and $A \leq w$, so $|A| \leq |w| \leq 2q$. Using $(x|1)_k = (x)_k$, $(2^{2n})_{2q} = (2^{2n})_{|A|} \prod_{|A| \leq j < 2q} (2^{2n} - j)$ and (34) with $t = 2^n$,

$$\tau_{w,A}(1) = \frac{\widehat{g}(w)}{2^{4nq}} \left\{ \begin{matrix} w \\ A \end{matrix} \right\} \prod_L (2^n)_{|A_{L,*}|} \prod_U (2^n)_{|A_{*,U}|} \prod_{|A| \leq j < 2q} (2^{2n} - j) = \frac{(2^{2n})_{2q}}{2^{4nq}} \widehat{g}(w) \left\{ \begin{matrix} w \\ A \end{matrix} \right\} \mathbb{E}[(C_{2^n})_{\circ A}] .$$

We have $(2^{2n})_{2q}/2^{4nq} = \prod_{0 \leq j < 2q} (1 - j/2^{2n})$. Summing over $A \leq w$ using (33), then over w , yields

$$f(1) = \prod_{0 \leq j < 2q} \left(1 - \frac{j}{2^{2n}} \right) \cdot \mathbb{E}[g(C_{2^n})] = \prod_{0 \leq j < 2q} \left(1 - \frac{j}{2^{2n}} \right) \cdot \mathbb{E}[g(C^{\text{RP}})] .$$

Bounding $|f(r)|$: the padded family f_r . For $r = 0$, $|f(0)| = g(J) \leq 1$, since $J = C^{F_{G,H}}$ by Claim 2 and $g(C^\pi) = \Pr[\mathcal{A}^{R(\pi)}] \in [0, 1]$ by (27). Now fix an integer $1 \leq r \leq 2^n/(256q^2)$, and put $t_* := 2^n/r \geq 256q^2$, which need not be an integer. For integers $s \geq 0$, let W_s be the sum of s independent uniformly random $2^n \times 2^n$ permutation matrices ($W_0 := 0$). For every integer $t \in [1, t_*]$, sample C_t and $W_{2^n - rt}$ independently and define

$$f_r(t) := \frac{(2^n t)_{2q}}{(2^n t_*)_{2q}} \cdot \mathbb{E}[g(rC_t + W_{2^n - rt})] . \quad (35)$$

The matrix $rC_t + W_{2^n - rt}$ has nonnegative integer entries and all row and column sums equal to $rt + (2^n - rt) = 2^n$, so by Claim 1 it equals C^π for some π , and g takes a value in $[0, 1]$ on it. Since $2q \leq 2^n \leq 2^n t \leq 2^n t_*$, the ratio in (35) lies in $(0, 1]$. Hence $0 \leq f_r(t) \leq 1$ for every integer $t \in [1, t_*]$.

Moments of W_s . Fix $B \in \mathbb{Z}_{\geq 0}^{2^n \times 2^n}$ and list the positions (L, U) with multiplicity $B_{L,U}$ as $(L_1, U_1), \dots, (L_k, U_k)$, where $k = |B|$; that is, each position (L, U) appears exactly $B_{L,U}$ times in the list. Write $W_s = \sum_{\ell=1}^s M_\ell$ with $M_1, \dots, M_s$ independent uniform permutation matrices, and let M denote a single uniform permutation matrix. As for C_t , we write $W_{s,L,U}$ and $M_{\ell,L,U}$ for the (L, U) -th entries of W_s and M_ℓ . By the choice of $(L_1, U_1), \dots, (L_k, U_k)$, we have

$$W_s^{\circ B} = \prod_{L,U} W_{s,L,U}^{B_{L,U}} = \prod_{m=1}^k W_{s,L_m,U_m} = \prod_{m=1}^k \left(\sum_{\ell=1}^s M_{\ell,L_m,U_m} \right) .$$

When we expand this product, each summand is a product of k entries of the M_ℓ , one from each factor, for a total of s^k summands. With each summand we can associate a function $\varphi : [k] \rightarrow [s]$ by setting $\varphi(m)$ to be the index ℓ such that M_{ℓ,L_m,U_m} is the m -th factor in the product. With this notation we can write

$$W_s^{\circ B} = \sum_{\varphi : [k] \rightarrow [s]} \prod_{m=1}^k M_{\varphi(m), L_m, U_m} .$$

By linearity of expectation, we obtain

$$\mathbb{E}[W_s^{\circ B}] = \sum_{\varphi : [k] \rightarrow [s]} \mathbb{E} \left[\prod_{m=1}^k M_{\varphi(m), L_m, U_m} \right] . \quad (36)$$

The polynomial μ_B . We will show that $\mathbb{E}[W_s^{\circ B}]$ is a polynomial in s of degree at most k . Fix $\varphi : [k] \rightarrow [s]$, and let $\ell_1 < \dots < \ell_i$ be the distinct values in the range of φ . Define $K_b := \varphi^{-1}(\ell_b)$ for $b \in [i]$, so that $\kappa(\varphi) := \{K_1, \dots, K_i\}$ is a partition of $[k]$. Then we can write

$$\prod_{m=1}^k M_{\varphi(m), L_m, U_m} = \prod_{b=1}^i \prod_{m \in K_b} M_{\ell_b, L_m, U_m} .$$

The b th inner product $\prod_{m \in K_b} M_{\ell_b, L_m, U_m}$ depends only on the entries of M_{ℓ_b} , and $M_{\ell_1}, \dots, M_{\ell_i}$ are independent since the ℓ_b are distinct, so the expectation of the product factorizes. Each M_{ℓ_b} is distributed as M , hence

$$\mathbb{E} \left[\prod_{m=1}^k M_{\varphi(m), L_m, U_m} \right] = \prod_{b=1}^i \mathbb{E} \left[\prod_{m \in K_b} M_{L_m, U_m} \right] =: c_{\kappa(\varphi)} .$$

We observe that $c_{\kappa(\varphi)}$ depends on φ only through the unordered partition $\kappa(\varphi)$.

Now fix a partition κ of $[k]$ into nonempty blocks $K_1, \dots, K_i$, indexed so that $\min K_1 < \dots < \min K_i$. A map $\varphi : [k] \rightarrow [s]$ has $\kappa(\varphi) = \kappa$ if and only if φ is constant on each block K_b and takes distinct values on distinct blocks. Thus the number of such maps is exactly the number of ways to choose a sequence of i distinct elements from $[s]$, which is $(s)_i$.

Grouping the terms of (36) by $\kappa(\varphi)$ gives

$$\mathbb{E} [W_s^{\diamond B}] = \sum_{\varphi: [k] \rightarrow [s]} c_{\kappa(\varphi)} = \sum_{\kappa} \sum_{\varphi: \kappa(\varphi) = \kappa} c_{\kappa} = \sum_{\kappa} c_{\kappa} |\{\varphi : \kappa(\varphi) = \kappa\}| = \sum_{\kappa} c_{\kappa} (s)_{|\kappa|} =: \mu_B(s). \quad (37)$$

Since $(s)_i$ is a polynomial in s of degree i and every partition of $[k]$ has at most k blocks, μ_B is a polynomial of degree at most $k = |B|$. If $B = 0$, $k = 0$, and the only partition of $[0] = \emptyset$ is the empty one, with $|\kappa| = 0$ and $c_{\kappa} = 1$ (an empty product), so $\mu_0 = (s)_0 = 1$. If $B \neq 0$, then $k \geq 1$, so every partition of $[k]$ has $|\kappa| \geq 1$ blocks, and $(0)_i = 0$ for $i \geq 1$; hence $\mu_B(0) = 0$.

A polynomial P_r extending f_r . We will show that f_r agrees at every integer $t \in [1, t_*]$ with a polynomial P_r in t , which, unlike f_r , can also be evaluated at the possibly non-integer point t_* . Fix w with $|w| \leq 2q$ and an integer $t \in [1, t_*]$, and let $s := 2^n - rt$. Then s is an integer with $s \geq 2^n - rt_* = 0$, so (37) applies to W_s . Applying the binomial theorem to each entry of $rC_t + W_s$, we have

$$\begin{aligned} (rC_t + W_s)^{\diamond w} &= \prod_{L,U} (rC_{t,L,U} + W_{s,L,U})^{w_{L,U}} = \prod_{L,U} \left(\sum_{B_{L,U} \leq w_{L,U}} \binom{w_{L,U}}{B_{L,U}} r^{w_{L,U} - B_{L,U}} C_{t,L,U}^{w_{L,U} - B_{L,U}} W_{s,L,U}^{B_{L,U}} \right) \\ &= \sum_{B \leq w} \binom{w}{B} r^{|w| - |B|} C_t^{\diamond(w-B)} W_s^{\diamond B}. \end{aligned}$$

In the last step, each term of the expanded product picks one summand $B_{L,U}$ from the factor at every position (L, U) , and we collect these choices into a matrix $B \leq w$; the powers of r then multiply to $r^{\sum_{L,U} (w_{L,U} - B_{L,U})} = r^{|w| - |B|}$.

Next we take expectations. Since C_t and W_s are independent, the expectation of each term factorizes as $\mathbb{E} [C_t^{\diamond(w-B)}] \cdot \mathbb{E} [W_s^{\diamond B}]$, and the second factor equals $\mu_B(s) = \mu_B(2^n - rt)$ by (37). For the first factor, we apply (33) to C_t with $w - B$ in place of w , and obtain $\mathbb{E} [C_t^{\diamond(w-B)}] = \sum_{A \leq w-B} \left\{ \begin{smallmatrix} w-B \\ A \end{smallmatrix} \right\} \mathbb{E} [(C_t)_{\diamond A}]$ by linearity of expectation. Altogether,

$$\mathbb{E} [(rC_t + W_s)^{\diamond w}] = \sum_{B \leq w} \binom{w}{B} r^{|w| - |B|} \mu_B(2^n - rt) \sum_{A \leq w-B} \left\{ \begin{smallmatrix} w-B \\ A \end{smallmatrix} \right\} \mathbb{E} [(C_t)_{\diamond A}]. \quad (38)$$

We now clear the denominators hidden in the moments $\mathbb{E} [(C_t)_{\diamond A}]$. Every A in the inner sum of (38) satisfies $|A| \leq |w - B| \leq |w| \leq 2q$, so we can split the product defining $(2^n t)_{2q}$ at $j = |A|$, and then apply (34):

$$\begin{aligned} (2^n t)_{2q} \cdot \mathbb{E} [(C_t)_{\diamond A}] &= \prod_{|A| \leq j < 2q} (2^n t - j) \cdot (2^n t)_{|A|} \mathbb{E} [(C_t)_{\diamond A}] \\ &= \prod_L (t)_{|A_{L,*}|} \prod_U (t)_{|A_{*,U}|} \prod_{|A| \leq j < 2q} (2^n t - j) =: \rho_A(t). \end{aligned}$$

Since $g(C) = \sum_w \widehat{g}(w) C^{\diamond w}$ is a finite sum, linearity of expectation gives

$$\mathbb{E} [g(rC_t + W_s)] = \sum_w \widehat{g}(w) \mathbb{E} [(rC_t + W_s)^{\diamond w}].$$

Plugging this and (38) into (35), and multiplying each moment $\mathbb{E} [(C_t)_{\diamond A}]$ by the factor $(2^n t)_{2q}$ from (35), we obtain $f_r(t) = P_r(t)$ for every integer $t \in [1, t_*]$, where

$$P_r(t) := \frac{1}{(2^n t_*)_{2q}} \sum_w \widehat{g}(w) \sum_{B \leq w} \binom{w}{B} r^{|w| - |B|} \mu_B(2^n - rt) \sum_{A \leq w-B} \left\{ \begin{smallmatrix} w-B \\ A \end{smallmatrix} \right\} \rho_A(t). \quad (39)$$

The right-hand side of (39) is defined for every real t , and it is a polynomial in t . Indeed, the prefactor does not depend on t , and it is well defined since each factor of $(2^n t_*)_{2q}$ satisfies $2^n t_* - j = 2^{2n}/r - j \geq 256q^2 \cdot 2^n - 2q > 0$ for $j < 2q$. Moreover, $\mu_B(2^n - rt)$ is the polynomial μ_B composed with the affine map $t \mapsto 2^n - rt$, $\rho_A(t)$ is a product of linear polynomials in t , and all sums in (39) are finite.

Degree of P_r . We will show that $\deg P_r \leq 4q$, by bounding the degree in t of each (w, B, A) term of (39). The factors $\widehat{g}(w)$, $\binom{w}{B}$, $r^{|w|-|B|}$, $\{w-B\}$ and the prefactor do not depend on t . Composing a polynomial with an affine map does not increase its degree, so $\mu_B(2^n - rt)$ has degree at most $\deg \mu_B \leq |B|$ in t . The polynomial $\rho_A(t)$ is a product of linear factors in t : each $(t)_k$ contributes k of them, and the last product in $\rho_A(t)$ contributes $2q - |A|$ of them. Since $\sum_L |A_{L,\star}| = \sum_U |A_{\star,U}| = |A|$, we get

$$\deg \rho_A = \sum_L |A_{L,\star}| + \sum_U |A_{\star,U}| + (2q - |A|) = |A| + |A| + (2q - |A|) = 2q + |A|.$$

Finally, $A \leq w - B$ implies $|A| \leq |w| - |B|$, so the degree of the (w, B, A) term is at most

$$|B| + 2q + |A| \leq |B| + 2q + |w| - |B| = 2q + |w| \leq 4q,$$

where the last step uses $|w| \leq 2q$. Since P_r is a finite sum of such terms, $\deg P_r \leq 4q$.

The value $P_r(t_*)$. We will show that

$$f(r) = \prod_{0 \leq j < 2q} \left(1 - \frac{jr}{2^{2n}}\right) \cdot P_r(t_*). \quad (40)$$

At $t = t_*$, the argument of μ_B in (39) is $2^n - rt_* = 0$, and we have shown that $\mu_B(0) = 0$ for $B \neq 0$ and $\mu_0 = 1$. Hence only the terms with $B = 0$ remain, and since $\binom{w}{0} = 1$ and $w - 0 = w$,

$$P_r(t_*) = \sum_w \sum_{A \leq w} \widehat{g}(w) \left\{ \begin{matrix} w \\ A \end{matrix} \right\} \cdot \frac{r^{|w|} \rho_A(t_*)}{(2^n t_*)_{2q}}. \quad (41)$$

To compare this with $f(r)$, we rewrite the last factor in terms of generalized falling factorials. For every integer $k \geq 0$ and real x , we have

$$(x/r)_k = \prod_{0 \leq j < k} \left(\frac{x}{r} - j\right) = \prod_{0 \leq j < k} \frac{x - jr}{r} = r^{-k} (x|r)_k.$$

Recall that $t_* = 2^n/r$ and $2^n t_* = 2^{2n}/r$. We rewrite the three products in $\rho_A(t_*)$ one by one. Applying the identity with $x = 2^n$ to each $(t_*)_{|A_{L,\star}|}$ and each $(t_*)_{|A_{\star,U}|}$, and using $\sum_L |A_{L,\star}| = \sum_U |A_{\star,U}| = |A|$ to collect the powers of r , we get

$$\prod_L (t_*)_{|A_{L,\star}|} = r^{-|A|} \prod_L (2^n|r)_{|A_{L,\star}|}, \quad \prod_U (t_*)_{|A_{\star,U}|} = r^{-|A|} \prod_U (2^n|r)_{|A_{\star,U}|}.$$

Each of the $2q - |A|$ factors of the last product in $\rho_A(t_*)$ satisfies $2^n t_* - j = 2^{2n}/r - j = r^{-1}(2^{2n} - jr)$, so

$$\prod_{|A| \leq j < 2q} (2^n t_* - j) = r^{-(2q-|A|)} \prod_{|A| \leq j < 2q} (2^{2n} - jr).$$

Multiplying these three expressions, and noting that the exponents of r add up to $-|A| - |A| - (2q - |A|) = -(2q + |A|)$, we get

$$\rho_A(t_*) = r^{-(2q+|A|)} \prod_L (2^n|r)_{|A_{L,\star}|} \prod_U (2^n|r)_{|A_{\star,U}|} \prod_{|A| \leq j < 2q} (2^{2n} - jr).$$

For the denominator, applying the identity with $x = 2^{2n}$ and $k = 2q$ gives

$$(2^n t_*)_{2q} = (2^{2n}/r)_{2q} = r^{-2q} (2^{2n}|r)_{2q},$$

which is nonzero, since $jr < 2q \cdot 2^n / (256q^2) < 2^{2n}$ for every $j < 2q$. Dividing $r^{|w|} \rho_A(t_*)$ by this, the powers of r combine to $r^{|w|-(2q+|A|)+2q} = r^{|w|-|A|}$, so

$$\begin{aligned} \frac{r^{|w|} \rho_A(t_*)}{(2^n t_*)_{2q}} &= \frac{r^{|w|-|A|}}{(2^{2n}|r)_{2q}} \prod_L (2^n|r)_{|A_{L,\star}|} \prod_U (2^n|r)_{|A_{\star,U}|} \prod_{|A| \leq j < 2q} (2^{2n} - jr) \\ &= \frac{2^{4nq}}{(2^{2n}|r)_{2q}} \cdot \frac{r^{|w|-|A|}}{2^{4nq}} \prod_L (2^n|r)_{|A_{L,\star}|} \cdot \prod_U (2^n|r)_{|A_{\star,U}|} \prod_{|A| \leq j < 2q} (2^{2n} - jr). \end{aligned}$$

Multiplying by $\widehat{g}^{(w)}\{A\}$ and comparing the last line with (32), we see that

$$\widehat{g}^{(w)}\{A\} \cdot \frac{r^{|w|} \rho_A(t_*)}{(2^n t_*)_{2q}} = \frac{2^{4nq}}{(2^{2n} |r)_{2q}} \cdot \tau_{w,A}(r) .$$

Substituting this into (41) and using $f = \sum_w \sum_{A \leq w} \tau_{w,A}$, we get

$$P_r(t_*) = \frac{2^{4nq}}{(2^{2n} |r)_{2q}} \sum_w \sum_{A \leq w} \tau_{w,A}(r) = \frac{2^{4nq}}{(2^{2n} |r)_{2q}} \cdot f(r) .$$

Finally, since $2^{4nq} = (2^{2n})^{2q}$, we have

$$\frac{(2^{2n} |r)_{2q}}{2^{4nq}} = \prod_{0 \leq j < 2q} \frac{2^{2n} - jr}{2^{2n}} = \prod_{0 \leq j < 2q} \left(1 - \frac{jr}{2^{2n}}\right) ,$$

so multiplying both sides of the previous display by this product gives (40). Moreover, since $0 \leq jr < 2^{2n}$ for every $j < 2q$, each factor of this product lies in $(0, 1]$, and hence so does the product.

Concluding the bound. We now apply (5) to the polynomial P_r on the interval $[a, b] := [1, t_*]$. This requires $b - a \geq 4(\deg P_r)^2$. Since $\deg P_r \leq 4q$, it suffices that $t_* - 1 \geq 64q^2$, which holds because $t_* \geq 256q^2$ and $q \geq 1$. The integers in $[1, t_*]$ are exactly $t = 1, \dots, \lfloor t_* \rfloor$, and at each of them we have $P_r(t) = f_r(t) \in [0, 1]$. Therefore (5) gives

$$|P_r(t_*)| \leq 2 \max_{t \in [1, t_*] \cap \mathbb{Z}} |P_r(t)| = 2 \max_{t \in [1, t_*] \cap \mathbb{Z}} |f_r(t)| \leq 2 .$$

Combining this with (40), and using that the product in (40) lies in $(0, 1]$, we obtain

$$|f(r)| = \prod_{0 \leq j < 2q} \left(1 - \frac{jr}{2^{2n}}\right) \cdot |P_r(t_*)| \leq |P_r(t_*)| \leq 2 .$$

This holds for every integer $1 \leq r \leq 2^n / (256q^2)$. Together with $|f(0)| \leq 1$, which we showed at the beginning of the paragraph on the padded family, we conclude that $|f(r)| \leq 2$ for every $r \in \{0, \dots, \lfloor 2^n / (256q^2) \rfloor\}$, which completes the proof of Claim 3. $\square$

References

- ABKM22. Gorjan Alagic, Chen Bai, Jonathan Katz, and Christian Majenz. Post-quantum security of the Even-Mansour cipher. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part III*, volume 13277 of *LNCS*, pages 458–487. Springer, Cham, May / June 2022. (CITED ON P. 5.)
- ABPS24. Gorjan Alagic, Chen Bai, Alexander Poremba, and Kaiyan Shi. On the two-sided permutation inversion problem. *CiC*, 1(1):27, 2024. (CITED ON P. 5.)
- ACMT25. Gorjan Alagic, Joseph Carolan, Christian Majenz, and Saliha Tokat. The sponge is quantum indifferentiable. In *66th FOCS*, pages 2591–2630. IEEE Computer Society Press, December 2025. (CITED ON P. 5.)
- AGQY22. Prabhanjan Ananth, Aditya Gulati, Luowen Qian, and Henry Yuen. Pseudorandom (function-like) quantum state generators: New definitions and applications. In Eike Kiltz and Vinod Vaikuntanathan, editors, *TCC 2022, Part I*, volume 13747 of *LNCS*, pages 237–265. Springer, Cham, November 2022. (CITED ON P. 2.)
- AMRS20. Gorjan Alagic, Christian Majenz, Alexander Russell, and Fang Song. Quantum-access-secure message authentication via blind-unforgeability. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part III*, volume 12107 of *LNCS*, pages 788–817. Springer, Cham, May 2020. (CITED ON P. 5.)
- BBC⁺98. Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald de Wolf. Quantum lower bounds by polynomials. In *39th FOCS*, pages 352–361. IEEE Computer Society Press, November 1998. (CITED ON P. 9.)
- BBC⁺21. Ritam Bhaumik, Xavier Bonnetain, André Chailloux, Gaëtan Leurent, María Naya-Plasencia, André Schrottenloher, and Yannick Seurin. QCB: Efficient quantum-secure authenticated encryption. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part I*, volume 13090 of *LNCS*, pages 668–698. Springer, Cham, December 2021. (CITED ON P. 5.)

- BBC⁺25. Jyotirmoy Basak, Ritam Bhaumik, Amit Kumar Chauhan, Ravindra Jejurikar, Ashwin Jha, Anandarup Roy, André Schrottenloher, and Suprita Talnikar. Post-quantum security of key-alternating feistel ciphers. In Goichiro Hanaoka and Bo-Yin Yang, editors, *ASIACRYPT 2025, Part I*, volume 16245 of *LNCS*, pages 446–478. Springer, Singapore, December 2025. (CITED ON P. 5.)
- BCEJ24. Ritam Bhaumik, Benoît Cogliati, Jordan Ethan, and Ashwin Jha. Mind the bad norms - revisiting compressed oracle-based quantum indistinguishability proofs. In Kai-Min Chung and Yu Sasaki, editors, *ASIACRYPT 2024, Part IX*, volume 15492 of *LNCS*, pages 215–247. Springer, Singapore, December 2024. (CITED ON P. 5.)
- BCF⁺24. Ritam Bhaumik, André Chailloux, Paul Frixons, Bart Mennink, and María Naya-Plasencia. Block cipher doubling for a post-quantum world. *CiC*, 1(3):4, 2024. (CITED ON P. 3.)
- BDF⁺11. Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. Random oracles in a quantum world. In Dong Hoon Lee and Xiaoyun Wang, editors, *ASIACRYPT 2011*, volume 7073 of *LNCS*, pages 41–69. Springer, Berlin, Heidelberg, December 2011. (CITED ON P. 2.)
- Bel19. Aleksandrs Belovs. Quantum Algorithms for Classical Probability Distributions. In Michael A. Bender, Ola Svensson, and Grzegorz Herman, editors, *27th Annual European Symposium on Algorithms (ESA 2019)*, volume 144 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 16:1–16:11, Dagstuhl, Germany, 2019. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. (CITED ON P. 2.)
- BGGJ26. Ritam Bhaumik, Chun Guo, Xiaoning Guo, and Ashwin Jha. Indistinguishability of sum of permutations: A Fourier analytic route to classical and quantum security. arXiv preprint arXiv:2609.35421, 2026. (CITED ON P. 5.)
- BKL⁺12. Andrey Bogdanov, Lars R. Knudsen, Gregor Leander, François-Xavier Standaert, John P. Steinberger, and Elmar Tischhauser. Key-alternating ciphers in a provable setting: Encryption using a small number of public permutations - (extended abstract). In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 45–62. Springer, Berlin, Heidelberg, April 2012. (CITED ON P. 6.)
- BM26. Marvin Beckmann and Christian Majenz. Quantum oracle distribution switching and its applications to fully anonymous ring signatures. Cryptology ePrint Archive, Paper 2026/293, 2026. (CITED ON PP. 2 AND 30.)
- BR93. Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In Dorothy E. Denning, Raymond Pyle, Ravi Ganesan, Ravi S. Sandhu, and Victoria Ashby, editors, *ACM CCS 93*, pages 62–73. ACM Press, November 1993. (CITED ON P. 1.)
- BRS02. John Black, Phillip Rogaway, and Thomas Shrimpton. Black-box analysis of the block-cipher-based hash-function constructions from PGV. In Moti Yung, editor, *CRYPTO 2002*, volume 2442 of *LNCS*, pages 320–335. Springer, Berlin, Heidelberg, August 2002. (CITED ON P. 1.)
- BZ13. Dan Boneh and Mark Zhandry. Quantum-secure message authentication codes. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 592–608. Springer, Berlin, Heidelberg, May 2013. (CITED ON P. 5.)
- Car26. Joseph Carolan. Compressed permutation oracles. In *Proceedings of the 58th Annual ACM Symposium on Theory of Computing, STOC '26*, page 150–161, New York, NY, USA, 2026. Association for Computing Machinery. (CITED ON PP. 2, 3, 4, 5, 11, AND 12.)
- CBH⁺17. Jan Czapkowski, Leon Groot Bruinderink, Andreas Hülsing, Christian Schaffner, and Dominique Unruh. Post-quantum security of the sponge construction. Cryptology ePrint Archive, Report 2017/771, 2017. (CITED ON P. 5.)
- CDMS10. Jean-Sébastien Coron, Yevgeniy Dodis, Avradip Mandal, and Yannick Seurin. A domain extender for the ideal cipher. In Daniele Micciancio, editor, *TCC 2010*, volume 5978 of *LNCS*, pages 273–289. Springer, Berlin, Heidelberg, February 2010. (CITED ON P. 4.)
- CFHL21. Kai-Min Chung, Serge Fehr, Yu-Hsuan Huang, and Tai-Ning Liao. On the compressed-oracle technique, and post-quantum security of proofs of sequential work. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part II*, volume 12697 of *LNCS*, pages 598–629. Springer, Cham, October 2021. (CITED ON PP. 3, 14, AND 30.)
- CGK⁺23. Alexandru Cojocaru, Juan Garay, Aggelos Kiayias, Fang Song, and Petros Wallden. Quantum Multi-Solution Bernoulli Search with Applications to Bitcoin’s Post-Quantum Security. *Quantum*, 7:944, March 2023. (CITED ON P. 30.)
- CGLQ20. Kai-Min Chung, Siyao Guo, Qipeng Liu, and Luowen Qian. Tight quantum time-space tradeoffs for function inversion. In *61st FOCS*, pages 673–684. IEEE Computer Society Press, November 2020. (CITED ON P. 5.)
- CHJ26. Wonseok Choi, Minki Hhan, and Junyoung Jang. Quantum security of XOR of permutations via Fourier analysis. Cryptology ePrint Archive, Paper 2026/2241, 2026. (CITED ON P. 5.)
- CHL⁺25. Alexandru Cojocaru, Minki Hhan, Qipeng Liu, Takashi Yamakawa, and Aaram Yun. Quantum lifting for invertible permutations and ideal ciphers. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part II*, volume 16001 of *LNCS*, pages 481–512. Springer, Cham, August 2025. (CITED ON P. 5.)
- CM26. Joseph Carolan and Christian Majenz. Compressed permutation oracles revisited. arXiv preprint arXiv:2609.28469, 2026. (CITED ON PP. 3, 4, AND 12.)

- CP24. Joseph Carolan and Alexander Poremba. Quantum one-wayness of the single-round Sponge with invertible permutations. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part VI*, volume 14925 of *LNCS*, pages 218–252. Springer, Cham, August 2024. (CITED ON P. 5.)
- CPS08. Jean-Sébastien Coron, Jacques Patarin, and Yannick Seurin. The random oracle model and the ideal cipher model are equivalent. In David Wagner, editor, *CRYPTO 2008*, volume 5157 of *LNCS*, pages 1–20. Springer, Berlin, Heidelberg, August 2008. (CITED ON P. 6.)
- CPZ24. Joseph Carolan, Alexander Poremba, and Mark Zhandry. (Quantum) indistinguishability and pre-computation. Cryptology ePrint Archive, Report 2024/1727, 2024. (CITED ON P. 5.)
- CS14. Shan Chen and John P. Steinberger. Tight security bounds for key-alternating ciphers. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 327–350. Springer, Berlin, Heidelberg, May 2014. (CITED ON P. 6.)
- DFM20. Jelle Don, Serge Fehr, and Christian Majenz. The measure-and-reprogram technique 2.0: Multi-round Fiat-Shamir and more. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part III*, volume 12172 of *LNCS*, pages 602–631. Springer, Cham, August 2020. (CITED ON PP. 5 AND 30.)
- DFMS19. Jelle Don, Serge Fehr, Christian Majenz, and Christian Schaffner. Security of the Fiat-Shamir transformation in the quantum random-oracle model. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part II*, volume 11693 of *LNCS*, pages 356–383. Springer, Cham, August 2019. (CITED ON PP. 5 AND 30.)
- DFPW19. Ronald DeVore, Simon Foucart, Guergana Petrova, and Przemysław Wojtaszczyk. Computing a quantity of interest from observational data. *Constructive Approximation*, 49(3):461–508, 2019. (CITED ON P. 30.)
- DS41. R. J. Duffin and A. C. Schaeffer. A refinement of an inequality of the brothers markoff. *Transactions of the American Mathematical Society*, 50(3):517–528, 1941. (CITED ON P. 7.)
- EM93. Shimon Even and Yishay Mansour. A construction of a cipher from a single pseudorandom permutation. In Hideki Imai, Ronald L. Rivest, and Tsutomu Matsumoto, editors, *ASIACRYPT’91*, volume 739 of *LNCS*, pages 210–224. Springer, Berlin, Heidelberg, November 1993. (CITED ON P. 6.)
- EM97. Shimon Even and Yishay Mansour. A construction of a cipher from a single pseudorandom permutation. *Journal of Cryptology*, 10(3):151–162, June 1997. (CITED ON P. 6.)
- FGdG⁺26. Pierre-Alain Fouque, Phillip Gajland, Hubert de Groote, Jonas Janneck, and Eike Kiltz. A closer look at falcon. In Joan Daemen and Emmanuel Thomé, editors, *EUROCRYPT 2026, Part IV*, volume 16545 of *LNCS*, pages 3–31. Springer, Cham, May 2026. (CITED ON P. 30.)
- FLM⁺26. Ben Foxman, Alex Lombardi, Fermi Ma, Barak Nehoran, and John Wright. Quantum lazy sampling and path recording for any group. Cryptology ePrint Archive, Paper 2026/1510, 2026. (CITED ON P. 3.)
- GGM86. Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *Journal of the ACM*, 33(4):792–807, October 1986. (CITED ON P. 1.)
- GM82. Shafi Goldwasser and Silvio Micali. Probabilistic encryption and how to play mental poker keeping secret all partial information. In *14th ACM STOC*, pages 365–377. ACM Press, May 1982. (CITED ON P. 1.)
- GM84. Shafi Goldwasser and Silvio Micali. Probabilistic encryption. *Journal of Computer and System Sciences*, 28(2):270–299, 1984. (CITED ON P. 1.)
- HHM⁺26. Yu-Hsuan Huang, Andreas Hülsing, Varun Maram, Silvia Ritsch, and Abishanka Saha. Feistel tools: Reprogramming and query-recording for QRPs. Cryptology ePrint Archive, Paper 2026/146, 2026. (CITED ON PP. 3 AND 17.)
- HI19. Akinori Hosoyamada and Tetsu Iwata. 4-round Luby-Rackoff construction is a qPRP. In Steven D. Galbraith and Shiho Moriai, editors, *ASIACRYPT 2019, Part I*, volume 11921 of *LNCS*, pages 145–174. Springer, Cham, December 2019. (CITED ON P. 5.)
- HI21. Akinori Hosoyamada and Tetsu Iwata. On tight quantum security of HMAC and NMAC in the quantum random oracle model. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part I*, volume 12825 of *LNCS*, pages 585–615, Virtual Event, August 2021. Springer, Cham. (CITED ON P. 5.)
- HKT11. Thomas Holenstein, Robin Künzler, and Stefano Tessaro. The equivalence of the random oracle model and the ideal cipher model, revisited. In Lance Fortnow and Salil P. Vadhan, editors, *43rd ACM STOC*, pages 89–98. ACM Press, June 2011. (CITED ON P. 6.)
- HL26. Yu-Hsuan Huang and Heming Liao. Multi-reprogramming with quantum advice in the quantum random oracle model. Cryptology ePrint Archive, Paper 2026/2140, 2026. (CITED ON P. 5.)
- HXY19. Minki Hhan, Keita Xagawa, and Takashi Yamakawa. Quantum random oracle model with auxiliary input. In Steven D. Galbraith and Shiho Moriai, editors, *ASIACRYPT 2019, Part I*, volume 11921 of *LNCS*, pages 584–614. Springer, Cham, December 2019. (CITED ON P. 5.)
- HY18. Akinori Hosoyamada and Kan Yasuda. Building quantum-one-way functions from block ciphers: Davies-Meyer and Merkle-Damgård constructions. In Thomas Peyrin and Steven Galbraith, editors, *ASIACRYPT 2018, Part I*, volume 11272 of *LNCS*, pages 275–304. Springer, Cham, December 2018. (CITED ON P. 5.)
- JST21. Joseph Jaeger, Fang Song, and Stefano Tessaro. Quantum key-length extension. In Kobbi Nissim and Brent Waters, editors, *TCC 2021, Part I*, volume 13042 of *LNCS*, pages 209–239. Springer, Cham, November 2021. (CITED ON P. 5.)
- Kim26. Sunyeop Kim. Improved soundness for compressed permutation oracles and tight quantum preimage and collision bounds for the sponge. Cryptology ePrint Archive, Paper 2026/2135, 2026. (CITED ON PP. 3, 4, AND 12.)

- KLLN16. Marc Kaplan, Gaëtan Leurent, Anthony Leverrier, and María Naya-Plasencia. Breaking symmetric cryptosystems using quantum period finding. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part II*, volume 9815 of *LNCS*, pages 207–237. Springer, Berlin, Heidelberg, August 2016. (CITED ON P. 4.)
- Liu23. Qipeng Liu. Non-uniformity and quantum advice in the quantum random oracle model. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part I*, volume 14004 of *LNCS*, pages 117–143. Springer, Cham, April 2023. (CITED ON P. 5.)
- LLMM18. Jooyoung Lee, Atul Luykx, Bart Mennink, and Kazuhiko Minematsu. Connecting tweakable and multi-key blockcipher security. *DCC*, 86(3):623–640, 2018. (CITED ON P. 4.)
- LR88. Michael Luby and Charles Rackoff. How to construct pseudorandom permutations from pseudorandom functions. *SIAM Journal on Computing*, 17(2), 1988. (CITED ON P. 1.)
- LRW02. Moses Liskov, Ronald L. Rivest, and David Wagner. Tweakable block ciphers. In Moti Yung, editor, *CRYPTO 2002*, volume 2442 of *LNCS*, pages 31–46. Springer, Berlin, Heidelberg, August 2002. (CITED ON P. 6.)
- Min09. Kazuhiko Minematsu. Beyond-birthday-bound security based on tweakable block cipher. In Orr Dunkelman, editor, *FSE 2009*, volume 5665 of *LNCS*, pages 308–326. Springer, Berlin, Heidelberg, February 2009. (CITED ON P. 6.)
- MMW25. Christian Majenz, Giulio Malavolta, and Michael Walter. Permutation superposition oracles for quantum query lower bounds. In Michal Koucký and Nikhil Bansal, editors, *57th ACM STOC*, pages 1508–1519. ACM Press, June 2025. (CITED ON P. 5.)
- Sha49. Claude E. Shannon. Communication theory of secrecy systems. *Bell Systems Technical Journal*, 28(4):656–715, 1949. (CITED ON P. 1.)
- SY17. Fang Song and Aaram Yun. Quantum security of NMAC and related constructions - PRF domain extension against quantum attacks. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part II*, volume 10402 of *LNCS*, pages 283–309. Springer, Cham, August 2017. (CITED ON P. 2.)
- Unr23. Dominique Unruh. Towards compressed permutation oracles. In Jian Guo and Ron Steinfeld, editors, *ASIACRYPT 2023, Part IV*, volume 14441 of *LNCS*, pages 369–400. Springer, Singapore, December 2023. (CITED ON PP. 2 AND 3.)
- Xag24. Keita Xagawa. On the efficiency of generic, quantum cryptographic constructions. *CiC*, 1(1):8, 2024. (CITED ON P. 16.)
- Yao82. Andrew Chi-Chih Yao. Theory and applications of trapdoor functions (extended abstract). In *23rd FOCS*, pages 80–91. IEEE Computer Society Press, November 1982. (CITED ON P. 1.)
- YZ21. Takashi Yamakawa and Mark Zhandry. Classical vs quantum random oracles. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part II*, volume 12697 of *LNCS*, pages 568–597. Springer, Cham, October 2021. (CITED ON P. 5.)
- Zha12. Mark Zhandry. How to construct quantum random functions. In *53rd FOCS*, pages 679–687. IEEE Computer Society Press, October 2012. (CITED ON PP. 2 AND 16.)
- Zha15. Mark Zhandry. A note on the quantum collision and set equality problems. *Quantum Info. Comput.*, 15(7–8):557–567, May 2015. (CITED ON P. 4.)
- Zha16. Mark Zhandry. A note on quantum-secure PRPs. Cryptology ePrint Archive, Report 2016/1076, 2016. (CITED ON P. 5.)
- Zha19. Mark Zhandry. How to record quantum queries, and applications to quantum indistinguishability. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part II*, volume 11693 of *LNCS*, pages 239–268. Springer, Cham, August 2019. (CITED ON PP. 2 AND 30.)

A AI Disclosure

We use AI in several ways, including generating TikZ code, identifying calculation, spelling, and grammatical errors, and discussing the presentation of the paper. We also use AI for interactive technical discussions, which are intertwined with our own ideas and discussions. Our main result (Theorem 1) has gone through multiple iterations of technical improvements. One of these arose from an insight obtained during a discussion with GPT-5.6 Sol Ultra: using an existential argument based on the polynomial method, one can unnoticeably embed a base oracle into a keyed family of its instances. The proof of Lemma 5 has also undergone a similar iterative process, in which GPT-6 Astra helped us with several counting arguments concerning randomized matrices and contributed to the idea of rerandomizing the permutation oracle, though this idea was buried in a single sentence among several pages of derivations.

B Efficient Simulation of Biased Random Oracles

In Section 3, our reduction samples a parameter p according to a certain distribution $\mathcal{D}$ over the interval $[0, 1]$, and then implements an oracle $H : \mathcal{K} \rightarrow \{0, 1\}$ in which every entry $H(k) \leftarrow \text{Ber}(p)$ is sampled independently according to the biased Bernoulli distribution of parameter p . We discuss here how this can be efficiently simulated, in the following two different cases.

Case 1: $\mathcal{D}$ is efficiently sampleable. In the case when $\mathcal{D}$ is efficiently sampleable, one can again follow an approach similar to the compressed oracle technique [Zha19, CFHL21]. Once the parameter p is sampled and fixed, we can simply use the Bernoulli recording oracle by [CGK⁺23]. On a very high level, consider the purification of H stored in a designated register $\mathcal{D}$, which is initially of the form $\bigotimes_{k \in \mathcal{K}} |\hat{0}_p\rangle$, with

$$|\hat{0}_p\rangle := \sqrt{1-p} |0\rangle + \sqrt{p} |1\rangle, \quad |\hat{1}_p\rangle := \sqrt{p} |0\rangle - \sqrt{1-p} |1\rangle.$$

We can then compress it into a compact representation via applying the compression operator $\mathbf{C} := |\perp\rangle\langle\hat{0}_p| + |\hat{0}_p\rangle\langle\perp| + |\hat{1}_p\rangle\langle\hat{1}_p|$ to each entry of $\mathcal{D}$, which would compress the initial state to an empty database $|\perp\rangle$. In each of the queries, one would simply need to implement the unitary $\mathbf{C} \cdot \mathbf{O} \cdot \mathbf{C}$ for $\mathbf{O} : |k, y, H\rangle \mapsto |k, y \oplus H(k), H\rangle$, which can indeed be made efficient.

Case 2: $\mathcal{D}$ is uniform over $[0, 1]$. In the case when $\mathcal{D}$ is uniform over $[0, 1]$, the resulting parameter $p \in [0, 1]$ would have infinitely large entropy, and so it is not even possible to store the entire p . Nevertheless, there are two possible ways out of this problem.

First, if we simply pick a sufficiently high precision $\lambda \in \mathbb{Z}_{>0}$ and truncate p to a λ -bit approximation $p_\lambda := \lfloor p \cdot 2^\lambda \rfloor \cdot 2^{-\lambda}$, then given that $|p_\lambda - p| \leq 2^{-\lambda}$, the resulting oracle $H_\lambda \leftarrow \text{Ber}(p_\lambda)^{\otimes K}$ would differ by a statistical distance at most $SD(H_\lambda, H) \leq |\mathcal{K}| \cdot 2^{-\lambda}$. Since p_λ is efficiently sampleable, we have fallen back to the first case, in which we know how to simulate H_λ .

On the other hand, if we insist on a perfect simulation (without any loss), the simulation can also be done without sampling the entire p . Observe that for any oracle algorithm $\mathcal{A}$ making at most q quantum queries to H , conditioned on any fixed choice of H , the probability of $\mathcal{A}$ outputting 1 is a multilinear polynomial $a(H)$ of degree at most $2q$ in $(H(k))_k$. Therefore, it suffices to simulate a randomized function $F(k)$ such that for every distinct $k_1, \dots, k_\ell$ with $\ell \leq 2q$, we have

$$\mathbb{E}_F \left[\prod_{i \in [\ell]} F(k_i) \right] = \mathbb{E}_H \left[\prod_{i \in [\ell]} H(k_i) \right] = \int_0^1 p^\ell dp = \frac{1}{\ell + 1}, \quad (42)$$

where the last two equalities follow from straight-forward calculation. Then we would have

$$\Pr[\mathcal{A}^F] = \mathbb{E}_F[a(F)] = \sum_{S \subseteq \mathcal{K}: |S| \leq 2q} \frac{\widehat{a}(S)}{|S| + 1} = \mathbb{E}_H[a(H)] = \Pr[\mathcal{A}^H],$$

where $\widehat{a}(S)$ is the coefficient of the monomial $\prod_{k \in S} H(k)$ in $a(H)$.

To achieve (42), we solve for $w_0, \dots, w_M \in \mathbb{Q}_{\geq 0}$ such that

$$\sum_{0 \leq j \leq M} w_j \cdot (j/M)^\ell = \frac{1}{\ell + 1} \quad \forall \ell \in \{0, \dots, 2q\}$$

with the convention that $0^0 := 1$, since then sampling p' with the distribution such that $\Pr[p' = j/M] := w_j$, and then sampling $F \leftarrow \text{Ber}(p')^{\otimes K}$ would immediately achieve (42). Observe that this is a numerical integration formula in disguise: namely, for $f_\ell(x) := x^\ell$, we are given its evaluation $f_\ell(j/M)$ for each $j \in \{0, \dots, M\}$, and tasked to find a set of coefficients $w_0, \dots, w_M$ that recovers $\int_0^1 f_\ell(x) dx = 1/(\ell + 1)$ via the formula $\sum_{0 \leq j \leq M} w_j f_\ell(j/M)$ for all $\ell \in [2q]$ simultaneously. By Equation (5.5) in [DFPW19], such a formula would indeed exist for some $M \leq \max_\ell (\deg f_\ell + 1)^2 = (2q + 1)^2$. Furthermore, as has been pointed out in the same paper, the coefficients $w_0, \dots, w_M$ can be cast as a solution of a linear program, which can be solved efficiently.

C Lifting Hardness of Search Games

In this section, we show that our indistinguishability lifting theorem (Theorem 1) is useful for lifting hardness of search games. Specifically, it can be used to recover a weaker form of the measure-and-reprogram technique [DFMS19, DFM20]. Such a weaker form can be seen in the literature. For example, Falcon uses a multi-target SIS assumption [FGdG⁺26, BM26], which can also be captured by our statement here.

Corollary 1. Let $V : \mathcal{Y} \times \mathcal{Z} \rightarrow \{0, 1\}$ be a predicate that can be evaluated efficiently, and $H : \mathcal{X} \rightarrow \mathcal{Y}$ be a random oracle. Then, for every oracle algorithm $\mathcal{A}$ making at most q quantum queries to H , there exists a similarly efficient algorithm $\mathcal{B}$ (where we count each evaluation of V as unit runtime), such that

$$\Pr \left[V(H(X), Z) \mid (X, Z) \leftarrow \mathcal{A}^H \right] \leq 8(q+1)^2 \cdot \Pr \left[V(Y, Z) \mid \begin{array}{l} Y \leftarrow \mathcal{Y} \\ Z \leftarrow \mathcal{B}(Y) \end{array} \right].$$

Proof. For $b \in \{0, 1\}$, consider the oracle $\mathcal{O}_b$ defined with two modes, with the selected mode indicated in a designated input register that is always measured by the oracle at the beginning of each query. On the read mode, the oracle always outputs Y that is sampled uniformly at the initialization stage. On the challenge mode, the oracle takes an input z , fully measures it and records it in an internal state, and then outputs $b \wedge V(Y, z)$. We summarize this oracle in Fig. 3.

Furthermore, let $\mathcal{O}_b^X$ be the corresponding keyed oracle, which can be efficiently simulated, since querying the read mode is indistinguishable from querying a random oracle, while querying a challenge mode can be implemented coherently because V can be evaluated efficiently.

$$\begin{array}{ll} \mathcal{O}_b(\text{"read"}) := Y \stackrel{\text{init}}{\leftarrow} \mathcal{Y} & \mathcal{O}_b(\text{"challenge"}, z): \\ & 1: \text{measure and record } z \\ & 2: \text{return } b \wedge V(Y, z) \end{array}$$

Fig. 3. The oracle $\mathcal{O}_b$ used for lifting hardness of search games.

Next, we transform the search algorithm $\mathcal{A}$ querying H to a distinguishing algorithm $\mathcal{D}$ querying $\mathcal{O}_b^X$ as follows: the algorithm $\mathcal{D}$ runs $\mathcal{A}$, and forwards each query to H to $\mathcal{O}_b^X$ in the read mode. After obtaining the output (X, Z) of $\mathcal{A}$, it makes one more query $B \leftarrow \mathcal{O}_b^X(X, \text{"challenge"}, Z)$ and outputs the response B . By construction we then have

$$\Pr \left[\mathcal{D}^{\mathcal{O}_b^X} \right] = b \cdot \epsilon, \text{ where } \epsilon := \Pr \left[V(H(X), Z) \mid (X, Z) \leftarrow \mathcal{A}^H \right].$$

This implies that $SD(\mathcal{D}^{\mathcal{O}_0^X}, \mathcal{D}^{\mathcal{O}_1^X}) = \epsilon$, and so applying our lifting theorem (Theorem 1) to $\mathcal{D}$ would give us a similarly efficient distinguisher $\mathcal{E}$ such that

$$SD(\mathcal{D}^{\mathcal{O}_0^X}, \mathcal{D}^{\mathcal{O}_1^X}) \leq 8(q+1)^2 \cdot SD(\mathcal{E}^{\mathcal{O}_0}, \mathcal{E}^{\mathcal{O}_1}).$$

We note that since $\mathcal{O}_0^X$ and $\mathcal{O}_1^X$ are efficiently simulatable, they have been internalized as a part of $\mathcal{E}$.

Finally, let E be the event where $\mathcal{E}^{\mathcal{O}_b}$ has made a challenge query z such that $V(Y, z)$ is satisfied. Conditioned on E not happening, the two games behave identically, and moreover the probability of it happening is independent of the choice of b . Therefore

$$SD(\mathcal{E}^{\mathcal{O}_0}, \mathcal{E}^{\mathcal{O}_1}) \leq \Pr[E].$$

On the other hand, when E happens, a Z such that $V(Y, Z)$ is satisfied can indeed be extracted via scanning through the challenge queries made by $\mathcal{E}$. This would then give us a similarly efficient reduction algorithm $\mathcal{B}$ that, on input Y , simulates the oracle $\mathcal{O}_b$ (the choice of b is irrelevant), and in the end extracts and outputs the aforementioned Z if it exists. The concrete bound obtained here is as follows:

$$\Pr[E] \leq \Pr \left[V(Y, Z) \mid \begin{array}{l} Y \leftarrow \mathcal{Y} \\ Z \leftarrow \mathcal{B}(Y) \end{array} \right].$$

Collecting the bounds so far, we conclude the proof. $\square$