Compressed Permutation Oracles

Joseph Carolan

University of Maryland

Abstract

The analysis of quantum algorithms which query random, invertible permutations has been a long-standing challenge in cryptography. Many techniques which apply to random oracles fail, or are not known to generalize to this setting. As a result, foundational cryptographic constructions involving permutations often lack quantum security proofs. With the aim of closing this gap, we develop and prove soundness of a compressed permutation oracle. Our construction shares many of the attractive features of Zhandry's original compressed function oracle: the purification is a small list of input-output pairs which meaningfully reflect an algorithm's knowledge of the oracle.

We then apply this framework to show that the Feistel construction with seven rounds is a strong quantum PRP, resolving an open question of (Zhandry, 2012). We further re-prove essentially all known quantum query lower bounds in the random permutation model, notably the collision and preimage resistance of both Sponge and Davies-Meyer, hardness of double-sided zero search and sparse predicate search, and give new lower bounds for cycle finding and the one-more problem.

Contents

1 Introduction **3**
- 1.1 Our results . . . 4
- 1.2 Technical summary . . . 5
- 1.3 Related work . . . 8
- 1.4 Reader's guide . . . 8
- 1.5 Acknowledgements . . . 9

2 Preliminaries **9**
- 2.1 Quantum . . . 9
- 2.2 Notation and definitions . . . 10
- 2.3 Compressed oracles . . . 12

3 Compressed permutation oracles **17**
- 3.1 Properties . . . 19

4 Purified permutations **21**
- 4.1 Query analysis . . . 22
- 4.2 Internal Feistel . . . 24
- 4.3 Canonical compression . . . 32
- 4.4 Cromulent twirling . . . 36

5 Soundness **39**
- 5.1 Valid, elegant, and sophisticated subspaces . . . 40
- 5.2 A twirling lemma . . . 43
- 5.3 Sanitized and ideal masked Feistel . . . 45
- 5.4 Preserved subspaces . . . 53
- 5.5 Putting the pieces together . . . 56

6 Feistel security **58**
- 6.1 Twirling . . . 59
- 6.2 Putting the pieces together . . . 64

7 Search bounds **66**
- 7.1 Simple examples . . . 68
- 7.2 A general search lower bound . . . 70
- 7.3 Application to sponge . . . 73
- 7.4 Application to Davies-Meyer . . . 75

1 Introduction

Quantum algorithms which query random permutations and their inverse are ubiquitous in quantum and post-quantum cryptography [Zha16; Ros22; Zha21; Coj+25; Ala+23; Ala+25]. This model is commonly referred to as the quantum random permutation model, or QRPM. Like the related quantum random oracle model (the QROM), this model is an idealization. Where the QROM is an idealization of hash functions, the QRPM is an idealization of cryptographic permutations and ciphers.

Many tools exist for analyzing the QROM [Zha12; Bon+11; Unr14; YZ21]. One of the most powerful is the compressed oracle [Zha19], which analyzes the purification of a random oracle in a specially chosen basis that has an intuitive interpretation. In particular, the purification can be written as a *database* of input-output pairs which meaningfully correspond to points the adversary has knowledge of. Moreover, each query adds at most a single such pair to the database. The introduction of this tool led to a slew of new results in quantum and post-quantum cryptography [Don+21; Chu+21], as well as query complexity [LZ19; HM23].

The QRPM has proven more challenging to analyze. For instance, the compressed oracle technique has resisted many attempts at generalizations to permutation oracles [Unr21; Unr23; Ros22; Cza+19]—or any correlated oracle distribution. Even basic search problems remained open until recently [CP24; HY18; Zha21; Unr21; MMW24], and despite recent progress on permutation based hash functions [CPZ24; Ala+25; Coj+25], many important constructions still lack quantum security proofs.

An emblematic example is the Feistel cipher, also known as the Luby-Rackoff cipher or Feistel network, which is a standard construction of a psuedorandom permutation from a psuedorandom function [LR88]. A number of influential block ciphers, such as DES [Nat77] and Blowfish [Sch93], are based on this construction. For purposes of illustration, suppose that h, k, f are pseudorandom functions from n bits to n bits. Then, the Feistel network constructs a $2n$ bit permutation φ as the following sequence of operations. To compute $\varphi(x)$:

1. Take the first n bits of x, referred to as x_L, and compute $h(x_L)$. XOR this value into the last n bits of x, referred to as x_R, obtaining $x' = x_L \| (x_R \oplus h(x_L))$.

2. Take the last n bits of x', and compute $k(x'_R)$. XOR the result into x'_L, obtaining $x'' = (x'_L \oplus k(x'_R)) \| x'_R$.

3. Take the first n bits of x'', and compute $f(x''_L)$. XOR the result into x''_R, obtaining $y = x''_L \| (x''_R \oplus f(x''_L))$, where $y = \varphi(x)$.

This procedure is depicted in Figure 1. The three functions corresponds to three rounds, and the construction can be extended with more functions in alternating direction for any number of rounds. It is known that the three round construction is indistinguishable from a random permutation with only forwards classical queries, and that four rounds are indistinguishable with forwards and inverse classical queries [LR88]. With quantum queries—the strongest plausible attack model—essentially nothing is known

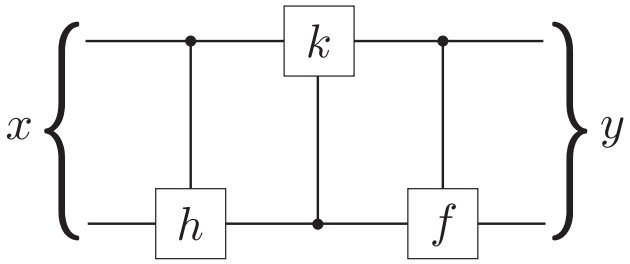


Figure 1: An illustration of three round Feistel.

about security. This question was first raised by Zhandry in 2012 [Zha12], yet remains beyond existing techniques.[1]

Hosoyamada and Iwata [HI19] claimed a proof that four round Feistel is a quantum psuedorandom permutation or qPRP[2], but their proof was found to contain a bug [Bha+24]. The same paper which found the bug proved security against non-adaptive, forwards only queries. On the algorithmic side, it is known that there is a quantum distinguishing attack for three rounds with forwards only queries [KM10], and four rounds with both forwards and backwards queries [Ito+19]. These results indicate that classical proof techniques for Feistel security must break down.

1.1 Our results

Our main result is a compressed oracle technique for analyzing two-way accessible permutations. Our construction is simple to apply, and shares many features of Zhandry's original compressed function oracle [Zha19]. In particular, our construction has a purifying register which is a list of input-output pairs, referred to as the compressed database. We establish that the construction is

1. *Sound*, in that it faithfully represents algorithms querying two-way accessible random permutations.

2. *Bounded*, in that the size of the compressed database grows by at most one with each query.

3. *Meaningful*, in that the points defined by the database correspond to points the algorithm "knows".

We defer to Section 3 for a formal description of both the construction and the above properties—note that the construction itself is only a few pages. At a high level, the construction is similar to the compressed function oracle but with two changes. First, the compression operator explicitly maintains injectivity of the database. Second, to answer inverse queries, we invert the database before and after answering.

[1] For instance, in a 2025 revision to [Zha16], Zhandry says "The quantum security of Feistel networks remains a largely open, very challenging problem."

[2] Without the qualifier of strong, this means forwards-only secure.

The idea of inverting the truth table to answer inverse queries is due to Unruh [Unr23]; our construction is a small modification to Unruh's proposal.[3] The proof that this simple construction is sound constitutes the main technical component of this work.

Theorem (Informal version of Theorem 5.19)**.** *No algorithm can distinguish the compressed permutation oracle described above from a truly random size N permutation unless it makes $\Omega(\sqrt[12]{N})$ bidirectional quantum queries.*

While our bound is exponential in the size of an input to the permutation, it does not seem tight. Once the construction is in place, we apply it to a number of problems in the QRPM. The main application is a proof that the seven round Feistel construction is a strong quantum PRP.

Theorem (Informal version of Theorem 6.5)**.** *No algorithm can distinguish the seven round Feistel construction from a truly random size N permutation unless it makes $\Omega(\sqrt[12]{N})$ bidirectional quantum queries.*

Additionally, we give a generic quantum lower bound for any predicate search problem in the QRPM in terms of the sparsity of the predicate. We defer to Section 7 for a formal statement, but note that this result is analogous to search lower bounds proved in the QROM by standard compressed oracle methods [Zha19; Chu+21]. Given our loose soundness bound, the result is not always tight.

However, essentially all known search lower bounds in the QRPM [HY18; Coj+25; CP24; CPZ24; Ala+25; MMW24] can be seen as a special case of this theorem. We reprove the preimage and collision resistance of the sponge hash function [Ber+07; Ber+11b; Ber+11a], which underlies the current international standard SHA3, and the Davies-Meyer construction [Win84], which underlies SHA1, SHA2, and MD5. Previous proofs are due to [Ala+25] and [Coj+25], both of which required the development of new techniques. For the sponge, our bound is tighter than the one in [Ala+25]. We also prove the first lower bounds for the one-more problem and cycle finding in the QRPM, both of which appear—to the author's knowledge—beyond the scope of existing techniques.

Intriguingly, we observe that the looseness in our bounds is entirely due to the soundness error of our construction. If the soundness analysis were improved, then all of the previously mentioned lower bounds would be tight.[4] This suggests that our compressed permutation oracle accurately captures the difficulty of these problems.

Given the range of applications and simplicity of our construction, we expect it to be useful for many other problems in the QRPM.

1.2 Technical summary

To explain the intuition underlying our security proof, let us begin by considering an approach which does not appear to work. A natural starting point for analyzing quantum algorithms which query invertible permutations is to purify the permutation oracle, mirroring the approach of Zhandry for the compressed function oracle. If we represent a

[3] In [Unr23], however, soundness is only conjectured.

[4] With the possible exception of the bound for the one-more and cycle finding problem.

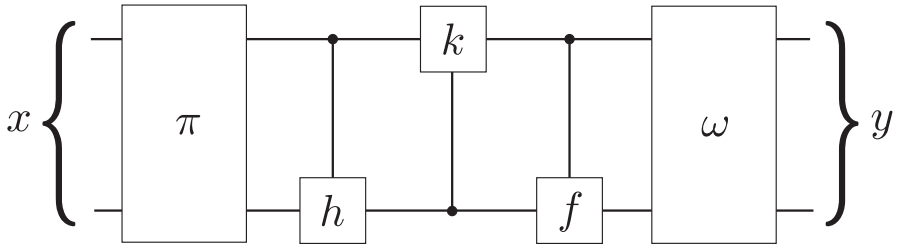


Figure 2: The masked Feistel decomposition.

permutation/function from $[N]$ to $[N]$ by its truth table, we can contrast the two purifications:

$$\underbrace{\frac{1}{\sqrt{N!}} \sum_{\substack{y_1,\ldots,y_N,\\ y_i \neq y_j}} |y_1\rangle \otimes \cdots \otimes |y_N\rangle}_{\text{permutation}}, \qquad \text{vs.} \qquad \underbrace{\frac{1}{\sqrt{N^N}} \sum_{y_1,\ldots,y_N} |y_1\rangle \otimes \cdots \otimes |y_N\rangle}_{\text{function}}. \tag{1}$$

The two states look quite similar, but there is a subtle difference. The purification of random functions, in this representation, is a product state: each entry of the truth table is unentangled, reflecting the fact that the oracle outputs are uncorrelated. Then, a simple indicator that a query algorithm has learned about, say, y_i is to check whether the i-th register above remains in uniform superposition (no knowledge), or is entangled with an external register (knowledge gained). This observation is what enables standard compressed oracles.

In the permutation case, there is weak entanglement between registers i and j, reflecting the condition that $y_i \neq y_j$. While this might appear like a minor difference, it turns out that the intuition above completely breaks. In particular, the i-th register in the purification of a permutation is *maximally* entangled with the remaining registers. For every $j \neq i$, y_j eliminates one of the $[N]$ possibilities for y_i. Taken together, the complement of y_i suffices to fix y_i. To check if y_i is queried, we must do more than check if register i is entangled, because it is maximally entangled before any queries are made.

To get around this, we must devise some other way to understand the entanglement between a query algorithm and the purification. One approach taken by Unruh [Unr21] and Rosmanis [Ros22] is to identify an injective database (a.k.a. an injective partial function) I with the sum of permutations agreeing with I on all well-defined points. Unfortunately, this leads to compressed database states which are non-orthogonal. This significantly complicates the analysis, especially in the case of inverse queries, limiting the scope of applications to simple search problems.

We will take a different approach based on expanding the purifying register. The point of this expansion is that we will be able to define *orthogonal* database states, which nonetheless characterize entanglement between the algorithm and oracle in an intuitive way. In more detail, we will write the full permutation φ as a three round Feistel network defined by functions labeled h, k, f, sandwiched on either side by random permutations π and ω. We refer to this ensemble as *masked Feistel*, and it is depicted in Figure 2.

We can write the purification of this ensemble, using $N = 4^n$ as the permutation domain size, as

$$\frac{1}{\sqrt{(N!)^2(2^n)^{3\cdot 2^n}}} \sum_{\pi,\omega\in\mathbf{S}_N, h,k,f\in\mathbf{F}_{\sqrt{N},\sqrt{N}}} |\pi,\omega,h,k,f\rangle \tag{2}$$

where $\mathbf{F}_{\sqrt{N},\sqrt{N}}$ is the set of functions from $\{0,1\}^n$ to $\{0,1\}^n$ and $\mathbf{S}_N$ is the set of permutations on $\{0,1\}^{2n}$. We denote the permutation defined by this whole ensemble as φ. This φ is a product of a uniform random permutation (say π) and some independent permutation, and is therefore uniform random. At this point, it may appear that we have made backwards progress, given the complexity of this purification. However, let us ignore for a moment the outer π,ω permutations, and consider what it means for φ to satisfy $\varphi(x) = y$. This fixes three input-output pairs in the h,k,f functions, corresponding to the three values reached along the evaluation of $\varphi(x)$—we call these three pairs a chain. In other words, input-output properties of φ correspond in a direct way to input-output properties of the random functions h,k,f.

Now, we apply the standard compressed oracle theory to analyze h,k,f. In this picture, the purifying registers for these functions are databases which begin empty. Each new query to the ensemble adds at most one new input-output pair to these databases, which we will use to reconstruct the input-output pairs of φ. Then, so long as each chain has no collisions (i.e. the same function input-output pair appearing in multiple chains), we can uniquely recover an injective database.

Of course, an adversarial algorithm could easily thwart this plan, by exploiting the periodicity of the three-round Feistel to distinguish it from a random permutation; or even making two queries that collide on the first n bits. This latter problem is the "trivialization of norm" faced by prior attempts to analyze quantum security of Feistel [HI19; Bha+24]. The purpose of the twirling permutations π and ω is to prevent the algorithm from detecting this internal structure. We do not compress these permutation, instead keeping them in uniform superposition. Intuitively, with high probability over the twirl, every new query will lead to a distinct first n bits on the internal Feistel. In turn, with high probability over the selection of the round function values, this will lead to a new chain (i.e. sequence of database values defining an input-output pair of φ) distinct from prior chains. This intuition can be formalized using the standard compressed oracle theory. At a high level, we define an isometry $\mathcal{I}$ which maps a superposition over tuples π,ω,D_h,D_k,D_f, where D_h,D_k,D_f correspond to function databases, to an injective database I . This isometry essentially acts as[5]

$$\mathcal{I}\left(\frac{1}{\sqrt{K}} \sum_{\substack{\pi,\omega,D_h,D_k,D_f,\\ \text{supporting only } I,\\ \text{no internal collisions}}} |\pi,\omega,D_h,D_k,D_f\rangle\right) = |I\rangle\,, \tag{3}$$

[5]The exact definition of an internal collision can be found in Section 4, and of the isometry can be found in Section 5.1.

where supporting I means the databases contain information to fix only the input-output pairs in I (and none else), and K is some normalization factor. The key point is that the states on the left hand side are mutually orthogonal, as they are supported on distinct sets of computational basis states.

Then, we show that this isometry *intertwines* the masked Feistel oracle with our compressed permutation oracle. In other words, if cmfO is the compressed masked Feistel oracle and cpO is the compressed permutation oracle, we show $\mathcal{I} \cdot \mathsf{cmfO} \approx \mathsf{cpO} \cdot \mathcal{I}$. Applying this across each query, we show that the final purified permutation (represented here as the tuple π, ω, h, k, f) is mapped by $\mathcal{I}$ to the purification in the compressed permutation oracle experiment (a partial injective database). Then, by Uhlmann's theorem, the mixed state the algorithm sees after tracing out the purifying register is close in trace distance.

We stress that the masked Feistel ensemble is necessary only as a tool for proving soundness, and is not necessary to understand and apply the construction. Indeed, our applications to search lower bounds make no mention of masked Feistel, instead following essentially the same blueprint as compressed function oracle lower bounds. On the other hand, our proof that the seven round Feistel construction is a strong qPRP is quite similar. In particular, we can consider the seven round Feistel to be an instance of masked (three round) Feistel, where the twirling π, ω are drawn from two round Feistel and two round inverse Feistel respectively. We show that this distribution *also* has the requisite statistical properties to hide the internal three rounds, allowing us to reuse many techniques from the soundness proof.

1.3 Related work

Prior work on constructing compressed permutation oracles has been met with limited success. Constructions were put forth by Unruh [Unr21] and Czajkowski et al [Cza+19], though both were later found to be incorrect. A notable exception is the recent work of Majenz, Malavolta and Walter [MMW24], referred to as the superposition permutation oracle. This construction was successfully employed to show search hardness for sparse relations, and is based on compressed oracle ideas. Unfortunately, this technique differs from standard compressed oracles in that the purification is not a small list of input-output pairs. This added complexity has limited the application of this technique to simple search problems involving only a single input-output pair—in contrast, our search lower bounds apply for predicates on any number of input-output pairs.

In the case where inverse queries are not allowed, a proposal by Rosmanis [Ros22] based on representation theory has been used to re-prove lower bounds for permutation inversion. Furthermore, a number of results from the QROM can be lifted to this setting [Cza+17; Cza+19], as a random one-way permutation is quantum indistinguishable from a random function [Zha15; Yue13].

1.4 Reader's guide

This manuscript is written in order of mathematical dependence. Sections 2 and 3 lay out the core ideas and notation, and should be read first. It may be helpful to read Section 7,

especially Section 7.1, afterwards to build intuition about our construction, as the search applications are relatively simple. The application to Feistel in Section 6 should be read after the soundness proof.

For the soundness proof, Section 5.5 and the proof of the main theorem provides a high level overview of the pieces, and uses only notation from Sections 4.1 and 5.1. Section 4 defines the purification of permutation oracles whose analysis forms the core of this work, though it is quite technical, and many of the definitions and lemmas may seem unmotivated if read immediately after Section 3. Section 5 relates this purification to our compressed permutation oracle.

1.5 Acknowledgements

The author is grateful to Gorjan Alagic, Christian Majenz, and Saliha Tokat for insightful discussions and significant feedback on this manuscript. JC acknowledges support from the U.S. Department of Energy grant DE-SC0020264.

2 Preliminaries

Throughout this document, we will assume $N = 4^n$ is a power of four, and identify $[N]$ with bit strings in the natural way. The choice of 4^n reflects that we are constructing a $2n$ bit permutation from n bit to n bit functions.

2.1 Quantum

Given a pure state $|\psi\rangle \in \mathcal{H}$, we denote the corresponding density matrix by $\rho(|\psi\rangle) = |\psi\rangle\langle\psi|$. We recall the trace norm of a density matrix, $\|\rho\|_1 = \mathrm{Tr}[\sqrt{\rho\rho^\dagger}]$. The corresponding trace distance of two mixed states upper bounds the probability that they can be distinguished by any algorithm.

Lemma 2.1. *Let σ, τ be density matrices over $\mathcal{H}$. Then, given either σ or τ with equal probability, the maximum distinguishing advantage $D(\sigma, \tau)$ is given by*

$$D(\sigma, \tau) = \frac{1}{2} \|\sigma - \tau\|_1 . \tag{4}$$

For a tensor product Hilbert space, we use capital sans serif subscripts to label each factor, e.g. $|\psi\rangle_{\mathsf{AB}} \in \mathcal{H}_{\mathsf{AB}} = \mathcal{H}_{\mathsf{A}} \otimes \mathcal{H}_{\mathsf{B}}$. We also refer to the state as having a register A and a register B. When an operator O_{A} acts on $\mathcal{H}_{\mathsf{A}}$, we can naturally extend it's action to $\mathcal{H}_{\mathsf{AB}}$ as $O_{\mathsf{A}} \otimes \mathbb{1}_{\mathsf{B}}$: we will drop this tensor product for convenience, writing $O_{\mathsf{A}} |\psi\rangle_{\mathsf{AB}}$ in place of $O_{\mathsf{A}} \otimes \mathbb{1}_{\mathsf{B}} |\psi\rangle_{\mathsf{AB}}$.

We will show that two oracles are indistinguishable by showing that the possible views of an adversary making queries to one of the two oracles are described by density matrices which are close in trace distance. Uhlmann's theorem allows us to bound the trace distance of two density matrices by constructing an isometry on their purifications.

Lemma 2.2. *Let $|\psi\rangle \in \mathcal{H}_\mathsf{A} \otimes \mathcal{H}_\mathsf{B}$ and $|\phi\rangle \in \mathcal{H}_\mathsf{A} \otimes \mathcal{H}_\mathsf{C}$ be normalized quantum states, and let $V : \mathcal{H}_\mathsf{C} \to \mathcal{H}_\mathsf{B}$ be an isometry. Let $\sigma_\mathsf{A} = \mathrm{Tr}_\mathsf{B}[\rho(|\psi\rangle_{\mathsf{AB}})]$ and $\tau_\mathsf{A} = \mathrm{Tr}_\mathsf{C}[\rho(|\phi\rangle_{\mathsf{AC}})]$. Then we have*

$$\frac{1}{2}\left\|\sigma_\mathsf{A} - \tau_\mathsf{A}\right\|_1 \leq \left\||\psi\rangle_{\mathsf{AB}} - V_\mathsf{C}\,|\phi\rangle_{\mathsf{AC}}\right\|. \tag{5}$$

Proof. We have

$$\begin{aligned}
\left\|\sigma_\mathsf{A} - \tau_\mathsf{A}\right\|_1 &= \left\|\mathrm{Tr}_\mathsf{B}[\rho(|\psi\rangle_{\mathsf{AB}})] - \mathrm{Tr}_\mathsf{C}[\rho(|\phi\rangle_{\mathsf{AC}})]\right\|_1 && && (6)\\
&= \left\|\mathrm{Tr}_\mathsf{B}[\rho(|\psi\rangle_{\mathsf{AB}})] - \mathrm{Tr}_\mathsf{B}[\rho(V_\mathsf{C}\,|\phi\rangle_{\mathsf{AC}})]\right\|_1 && (V \text{ an isometry}) && (7)\\
&\leq \left\|\rho(|\psi\rangle_{\mathsf{AB}}) - \rho(V_\mathsf{C}\,|\phi\rangle_{\mathsf{AC}})\right\|_1 && (L_1 \text{ trace monotonicity}) && (8)\\
&\leq 2\left\||\psi\rangle_{\mathsf{AB}} - V_\mathsf{C}\,|\phi\rangle_{\mathsf{AC}}\right\| && (L_1 \text{ and } L_2 \text{ relation}) && (9)
\end{aligned}$$

□

Another simple and useful fact is that the norm of an operator given by a direct sum is the maximum of the norm of any of the terms. If an operator $O : H(A) \otimes \mathcal{H}(B) \to \mathcal{H}(A) \otimes \mathcal{H}(B)$ can be written $O_{\mathsf{AB}} = \bigoplus_{a \in A} |a\rangle\langle a|_\mathsf{A} \otimes O_{a,\mathsf{B}}$ for some family of operators $O_a : \mathcal{H}(B) \to \mathcal{H}(B)$, we will sometimes say that it is *controlled* on register A.

Lemma 2.3. *Let $\mathcal{H} = \mathcal{H}_1 \oplus \cdots \oplus \mathcal{H}_k$ be a Hillbert space, and $O : \mathcal{H} \to \mathcal{H}$ be an operator of the form $O = O_1 \oplus \cdots \oplus O_k$ where $O_i : \mathcal{H}_i \to \mathcal{H}_i$. Then we have*

$$\|O\| = \max_{i \in [k]} \|O_i\|. \tag{10}$$

A projector $\Pi : \mathcal{H} \to \mathcal{H}$ is an idempotent operator, where $\Pi^2 = \Pi$. Projectors are orthogonal if they are Hermitian. We will only work with orthogonal projectors. We sometimes say $|\psi\rangle \in \Pi$ to mean that $\Pi\,|\psi\rangle = |\psi\rangle$.

2.2 Notation and definitions

Databases. We use $\mathbf{F}_{M,N}$ to denote the set of functions from $[M]$ to $[N]$. We use $\mathbf{S}_N$ to denote the set of permutations of $[N]$. We use $\mathbf{D}_{M,N}$ to denote the set of partial functions from $[M]$ to $[N]$, or equivalently the set of functions from $[M]$ to $[N] \cup \{\perp\}$. We use $\mathbf{I}_{M,N}$ to denote set of partial *injective* functions from $[M]$ to $[N]$, or equivalently the set of functions from $[M]$ to $[N] \cup \{\perp\}$ which are injective on all elements of the range except $\perp$.[6] We will sometimes drop the subscript when it is clear from context. For our purposes, we will almost always work with $\mathbf{I}_{N,N}$ and $\mathbf{D}_{\sqrt{N},\sqrt{N}}$, where we suppose $N = 4^n$ is a power of four. We take total functions to be a special case of partial functions, meaning

$$\mathbf{F}_{M,N} \subset \mathbf{D}_{M,N} \qquad\qquad \mathbf{S}_N \subset \mathbf{I}_N. \tag{11}$$

[6] As in, there may be multiple pre-images of $\perp$, but at most one preimage of any $y \in [N]$.

We will interchange between the partial and total function representations, and refer to such objects as *databases*, which may or may not be injective. For a $D \in \mathbf{D}_{M,N}$, we use $\mathsf{Dom}(D)$ to denote the domain of D as a partial function, or equivalently the preimages of $[N]$ as a total function. The image $\mathsf{Im}(D)$ refers to the image of D as a partial function, and similarly define $\mathsf{Dom}(I)$ and $\mathsf{Im}(I)$ for $I \in \mathbf{I}_{M,N}$. The size of a database can be defined as $|D| = |\mathsf{Dom}(D)|$, and similarly for injective databases $|I| = |\mathsf{Dom}(I)|$. When M, N are understood from context, we write $\mathbf{D}_t$ to denote the set of $D \in \mathbf{D}$ of size $|D| = t$, and similarly $\mathbf{D}_{\leq t}$ to denote the set of $D \in \mathbf{D}$ of size $|D| \leq t$. We similarly define $\mathbf{I}_t$ and $\mathbf{I}_{\leq t}$.

We use $\perp$ to denote the trivial partial function, or equivalently the total function which maps all inputs to $\perp$. Given a $D \in \mathbf{D}_{M,N}$, we use $D[x \to y]$ to denote the the function

$$D[x \to y](x') = \begin{cases} D(x') & (\text{If } x \neq x') \\ y & (\text{Otherwise}) \end{cases} \tag{12}$$

We similarly define $I[x \to y]$, with the subtlety that we may need to discard a conflicting element from the domain of I to ensure injectivity,

$$I[x \to y](x') = \begin{cases} I(x') & (\text{If } x \neq x', \text{and either } y = \perp \text{ or } I(x') \neq y) \\ \perp & (\text{If } x \neq x', y \neq \perp, \text{and } I(x') = y) \\ y & (\text{If } x = x') \end{cases} \tag{13}$$

We will almost always use this notation when x is not already in the domain, and in the case of injective functions where y is not already in the range. We can identify databases with relations as well, and write $(x, y) \in D$ when $D(x) = y$, and similarly for $(x, y) \in I$.

Given a $D \in \mathbf{D}$, we define $D|^x$ as the set of databases which agree with D on all inputs except x, in particular

$$D|^x := \{D[x \to y] \,:\, y \in [N] \cup \{\perp\}\}. \tag{14}$$

We let $\mathbf{D}|^x$ be the set of databases which are undefined on x, i.e.,

$$\mathbf{D}|^x := \{D \in \mathbf{D} \,:\, x \notin \mathsf{Dom}(D)\}. \tag{15}$$

Note that $\{D|^x \,:\, D \in \mathbf{D}|^x\}$ is a partition of $\mathbf{D}$ for any fixed x. We define similar notions for injective databases $I \in \mathbf{I}$, once again taking care to maintain injectivity in the natural way.

$$I|^x := \{I[x \to y] \,:\, y \in ([N] \setminus \mathsf{Im}(I)) \cup \{I(x)\} \cup \{\perp\}\} \tag{16}$$

$$\mathbf{I}|^x := \{I \in \mathbf{I} \,:\, x \notin \mathsf{Dom}(I)\}, \tag{17}$$

where once again $\{I|^x \,:\, I \in \mathbf{I}|^x\}$ is a partition of $\mathbf{I}$ for any fixed x. Given a set S, we denote a uniform random element of S by $x \sim S$. Given a string $x \in \{0,1\}^{2n}$, we use x_L to refer to the first n bits and x_R to refer to the last n bits.

Spaces. Given a set S, we denote by $\mathcal{H}(S)$ the Hilbert space spanned by an orthogonal basis labeled by elements of S, i.e.,

$$\mathcal{H}(S) \coloneqq \mathsf{span}\{|x\rangle \,:\, x \in S\}, \qquad \langle x|x'\rangle = \delta_{x,x'}. \tag{18}$$

In such a space, for a set $A \subset S$ we denote the uniform superposition of elements in A by

$$|+_A\rangle \coloneqq \frac{1}{\sqrt{|A|}} \sum_{x\in A} |x\rangle\,. \tag{19}$$

Given sets $A_1, \ldots, A_l \subset S$ such that $\bigcup_{i=1}^{l} A_i = S$ and $A_i \cap A_j = \emptyset$ for all $i \neq j$, we say that $\{A_1, \ldots, A_l\}$ is a partition of S. In that case, we may express $\mathcal{H}(S)$ as the direct sum $\mathcal{H}(A_1) \oplus \cdots \oplus \mathcal{H}(A_l)$.

We will often deal with the Hilbert spaces $\mathcal{H}(\mathbf{D})$ and $\mathcal{H}(\mathbf{I}) \leq \mathcal{H}(\mathbf{D})$. We define the projector Π^t and $\Pi^{\leq t}$ on the former space as

$$\Pi^t \coloneqq \sum_{D\in\mathbf{D}_t} |D\rangle\langle D|, \qquad \Pi^{\leq t} \coloneqq \sum_{k=0}^{t} \Pi_k. \tag{20}$$

We can define these operators on $\mathcal{H}(\mathbf{I})$ by their restriction to this subspace.

A norm without a subscript, $\|\cdot\|$, refers to the standard L_2 norm. We use $\|\cdot\|_1$ for the trace norm. For the Hilbert space $\mathcal{H}_{\mathsf{A}} \otimes \mathcal{H}(\mathbf{D})$ or $\mathcal{H}_{\mathsf{A}} \otimes \mathcal{H}(\mathbf{I})$, for any Hilbert space $\mathcal{H}_{\mathsf{A}}$, we use the notation $\|\cdot\|_{\leq t}$ for the standard L_2 norm restricted to databases of size at most t. In other words, for a state $|\psi\rangle$ or operator O on $\mathcal{H}_{\mathsf{AI}} = \mathcal{H}_{\mathsf{A}} \otimes \mathcal{H}(\mathbf{I})$, we define

$$\||\psi\rangle_{\mathsf{AI}}\|_{\leq t} \coloneqq \left\|\Pi_{\mathsf{I}}^{\leq t} \cdot |\psi\rangle_{\mathsf{AI}}\right\|, \qquad \|O_{\mathsf{AI}}\|_{\leq t} \coloneqq \left\|O_{\mathsf{AI}} \cdot \Pi_{\mathsf{I}}^{\leq t}\right\|. \tag{21}$$

We also define consistency projectors for databases, which assert that a given input-output pair appears in the database. For databases over $\mathcal{H}(\mathbf{F}_{M,N})$ and an input-output pair $(x, y) \in [M] \times [N]$, we define

$$\Pi^{(x,y)\in\mathsf{db}} \coloneqq \sum_{D\in\mathbf{D}\,:\,D(x)=y} |D\rangle\langle D|, \tag{22}$$

and similarly for the space of injective databases.

2.3 Compressed oracles

In this section, we briefly recall the standard compressed oracle due to [Zha19]. Our presentation diverges from the original to more closely match the presentation of compressed permutation oracles. Throughout, we assume M and N are powers of two, and identify elements of $[M]$ and $[N]$ with bitstrings by the natural representation.

Let $f \sim \mathbf{F}_{M,N}$ be a uniform random function, and consider the oracle

$$\mathcal{O}_f\,|x, y\rangle_{\mathsf{XY}} \coloneqq |x, y \oplus f(x)\rangle_{\mathsf{XY}}\,. \tag{23}$$

Let $\mathcal{A}^{\mathcal{O}_f}$ be a quantum algorithm which makes q queries to $\mathcal{O}_f$. This algorithm is described by a Hilbert space $\mathcal{H}_{\mathsf{A}}$ for its internal state and query registers, as well as a sequence $A_0, \ldots, A_q$ of unitary operators on $\mathcal{H}_{\mathsf{A}}$ describing the action in between each query. We refer to registers X, Y as the query registers, corresponding to the two registers in Equation (23). For simplicity, we consider these registers to be inside of register A. We can write the final mixed state of the algorithm as

$$\rho_{\mathcal{A}}^{(\mathcal{O}_f)} = \mathbb{E}_{f \sim \mathbf{F}_{M,N}} \left[\rho(A_q \mathcal{O}_f \ldots A_1 \mathcal{O}_f A_0 \left|0\right\rangle) \right]. \tag{24}$$

Alternatively, we could purify this experiment, defining a register $\mathcal{H}_{\mathsf{F}} = \mathcal{H}(\mathbf{F})$ to hold the purified function, and replacing the standard oracle with the purified oracle

$$\mathsf{P}_{\mathsf{XYF}} \left|x, y\right\rangle_{\mathsf{XY}} \left|f\right\rangle_{\mathsf{F}} \coloneqq \left|x, y \oplus f(x)\right\rangle_{\mathsf{XY}} \left|f\right\rangle_{\mathsf{F}}. \tag{25}$$

Considering the purified experiment over $\mathcal{H}_{\mathsf{AF}}$ and tracing out F results in the same view,

$$\rho_{\mathcal{A}}^{(\mathcal{O}_f)} = \mathrm{Tr}_{\mathsf{F}} \left(\rho \left(A_{q,\mathsf{A}} \mathsf{P}_{\mathsf{AF}} \ldots A_{1,\mathsf{A}} \mathsf{P}_{\mathsf{AF}} A_{0,\mathsf{A}} \left|0\right\rangle_{\mathsf{A}} \left|+_{\mathbf{F}}\right\rangle_{\mathsf{F}} \right) \right). \tag{26}$$

The compressed picture. The compressed oracle provides an alternative, powerful picture of quantum accessible random functions. In this picture, the purifying register is a database which begins empty. Each query modifies only the query point in the database, somewhat analogous to classical lazy sampling.

To define this formally, we will first consider the subspace $\mathcal{H}(D|^x) \leq \mathcal{H}(\mathbf{D})$ for some $D \in \mathbf{D} = \mathbf{D}_{M,N}$ and $x \in [M] \setminus \mathsf{Dom}(D)$. This subspace captures modifying the database on only the x-th position. Let us write $\left|+_{x,D}\right\rangle = \frac{1}{\sqrt{N}} \sum_{y \in [N]} \left|D[x \to y]\right\rangle$. We can now define the *function compression* operator fC on this space, indexed by x and D, as

$$\mathsf{fC}_{x,D} \coloneqq \mathbb{1} - \left|+_{x,D}\right\rangle\!\left\langle+_{x,D}\right| - \left|D\right\rangle\!\left\langle D\right| + \left|+_{x,D}\right\rangle\!\left\langle D\right| + \left|D\right\rangle\!\left\langle+_{x,D}\right|. \tag{27}$$

In other words, this operator swaps the uniform superposition of outputs with the undefined output, for the input x. Recalling that $\{D|^x : D \in \mathbf{D}|^x\}$ is a partition of $\mathbf{D}$, we may then define the full compression operator and its controlled analog as

$$\mathsf{fC}_x \coloneqq \bigoplus_{D \in \mathbf{D}|^x} \mathsf{fC}_{x,D}, \qquad\qquad \mathsf{fC}_{\mathsf{XD}} \left|x\right\rangle_{\mathsf{X}} = \left|x\right\rangle_{\mathsf{X}} \otimes \mathsf{fC}_{x,\mathsf{D}}. \tag{28}$$

The final piece we will need is the extended purified oracle P, which acts as

$$\mathsf{P} \left|x, y\right\rangle \left|D\right\rangle \coloneqq \begin{cases} \left|x, y \oplus D(x)\right\rangle \left|D\right\rangle & (\text{If } x \in \mathsf{Dom}(D)) \\ \left|x, y\right\rangle \left|D\right\rangle & (\text{Otherwise}) \end{cases} \tag{29}$$

Now, we define the full compressed function oracle as the purified oracle conjugated by the compression operator, i.e.,

$$\mathsf{cfO}_{\mathsf{XYP}} \coloneqq \mathsf{fC}_{\mathsf{XP}} \cdot \mathsf{P}_{\mathsf{XYP}} \cdot \mathsf{fC}_{\mathsf{XP}}^{\dagger}, \tag{30}$$

noting that fC is an involution so the dagger is strictly unnecessary. The compressed oracle experiment is somewhat similar to the purified experiment, except we allow the purifying register to include partial functions, $\mathcal{H}_{\mathsf{D}} = \mathcal{H}(\mathbf{D})$. We refer to this register as the *compressed oracle*, and it is initialized to $|\perp\rangle_{\mathsf{D}}$. We additionally replace purified oracle calls with the compressed oracle cfO. This corresponds to the final state

$$|\psi_{\mathcal{A}}^{(\mathsf{cfO})}\rangle_{\mathsf{AD}} := A_{q,\mathsf{A}}\mathsf{cfO}_{\mathsf{AD}} \dots A_{1,\mathsf{A}}\mathsf{cfO}_{\mathsf{AD}}A_{0,\mathsf{A}} |0\rangle_{\mathsf{A}} |\perp\rangle_{\mathsf{D}}. \tag{31}$$

The following facts were shown by [Zha19], and are core to the theory. At a high level, compressed oracles are:

1. *Sound*, in that they faithfully represent algorithms querying random functions.
2. *Bounded*, in that the size of the compressed database grows by at most one with each query.
3. *Meaningful*, in that the points defined by the database correspond to points the algorithm "knows".

These key properties are formalized in the following lemmas.

Lemma 2.4 (Soundness)**.** *The view of an algorithm querying the compressed oracle is the same as an algorithm querying a truly random function,*

$$\rho_{\mathcal{A},\mathsf{A}}^{(\mathcal{O}_f)} = \mathrm{Tr}_{\mathsf{D}}\left[|\psi_{\mathcal{A}}^{(\mathsf{cfO})}\rangle\langle\psi_{\mathcal{A}}^{(\mathsf{cfO})}|_{\mathsf{AD}}\right]. \tag{32}$$

Lemma 2.5 (Bounded growth)**.** *The final state of the compressed oracle after q queries is supported entirely on functions with domain size q,*

$$|\psi_{\mathcal{A}}^{(\mathsf{cfO})}\rangle_{\mathsf{AD}} \in \mathcal{H}_{\mathsf{A}} \otimes \mathcal{H}(\mathbf{D}_{\leq q}). \tag{33}$$

To establish the fundamental lemma (that formalizes the third property above), we will interpret the adversary as outputting a tuple of l input-output pairs. We will compare the probabilities that these pairs appear in the database with the probability that they agree with the oracle.[7] To do this, we define the projector

$$\Pi_{\mathsf{AD}}^{(\mathsf{cfO})} := \sum_{\{(x_1,y_1),\dots,(x_l,y_l)\}\in([M]\times[N])^{\times l}} \left(\Pi^{(x_1,y_1)\in\mathsf{db}} \dots \Pi^{(x_l,y_l)\in\mathsf{db}}\right)_{\mathsf{D}} \otimes \tag{34}$$

$$(|x_1,y_1,\dots,x_l,y_l\rangle\langle x_1,y_1,\dots,x_l,y_l| \otimes \mathbb{1})_{\mathsf{A}}, \tag{35}$$

[7] Sometimes, this lemma is formulated in the context of search games with a target predicate. Our version implies this application as a special case, as shown in Lemma 7.1.

which asserts that the output of the adversary corresponds to input-output pairs which agree with the database. We will also require the analogous operator that decompresses before each check,

$$\Pi_{\mathsf{AD}}^{(\mathcal{O}_f)} := \sum_{\{(x_1,y_1),\ldots,(x_l,y_l)\}\in([M]\times[N])^{\times l}} \left(\mathsf{fC}_{x_1}\Pi^{(x_1,y_1)\in\mathsf{db}}\mathsf{fC}_{x_1}^\dagger \ldots \mathsf{fC}_{x_l}\Pi^{(x_l,y_l)\in\mathsf{db}}\mathsf{fC}_{x_l}^\dagger\right)_{\mathsf{D}} \otimes \tag{36}$$

$$(|x_1,y_1,\ldots,x_l,y_l\rangle\langle x_1,y_1,\ldots,x_l,y_l| \otimes \mathbb{1})_{\mathsf{A}}. \tag{37}$$

Because the compressed oracle is indistinguishable from a random oracle, this projector is equivalent to checking if the adversary output is consistent with the oracle.

Lemma 2.6 (Fundamental lemma)**.** *The probability that an adversary finds l input-output pairs of the oracle is exponentially close to the probability that the corresponding pairs appears in the compressed database,*

$$\left\|\Pi_{\mathsf{AD}}^{(\mathcal{O}_f)}\,|\psi_{\mathcal{A}}^{(\mathsf{cfO})}\rangle_{\mathsf{AD}}\right\| \leq \left\|\Pi_{\mathsf{AD}}^{(\mathsf{cfO})}\,|\psi_{\mathcal{A}}^{(\mathsf{cfO})}\rangle_{\mathsf{AD}}\right\| + \sqrt{\frac{l}{N}}. \tag{38}$$

Further properties. Another important property of compressed oracles is that they are *valid*, in that decompressing on any given input will cause the truth table to have a well-defined output. A simple example of a valid database is $|\perp\rangle_{\mathsf{D}}$. If we consider applying decompression $\mathsf{fC}_x^\dagger$ to this state, we obtain

$$\mathsf{fC}_{x,\mathsf{D}}^\dagger\,|\perp\rangle_{\mathsf{D}} = \sum_{y\in[N]} |(x,y)\rangle_{\mathsf{D}}\,, \tag{39}$$

where we use relation notation for the database. It is clear from the above expression that the output of x is well defined. It is known that this property is preserved under compressed oracle queries.

To make this formal, we define the *compressed* and *decompressed* subspaces. The projector onto the former will be denoted by Ξ, and the latter by Γ. We define these operators for a fixed input x as

$$\Gamma_x := \sum_{D\in\mathbf{D}\text{ s.t. }x\in\mathsf{Dom}(D)} |D\rangle\langle D|, \qquad\qquad \Xi_x := \mathsf{fC}_x\Gamma_x\mathsf{fC}_x^\dagger. \tag{40}$$

Note that the compression operator interchanges compressed and decompressed subspaces. We define the general operators as the products of these over all values of x,

$$\Gamma := \prod_{x\in[M]} \Gamma_x, \qquad\qquad \Xi := \prod_{x\in[M]} \Xi_x. \tag{41}$$

Each Γ_x and fC_x commutes with $\Gamma_{x'}$ and $\mathsf{fC}_{x'}$ for $x \neq x'$, so this is a well defined projector. The full decompressed subspace is simply the projector onto all total functions, from which it is clear that the standard purified oracle maintains this subspace. Analogously, the compressed oracle maintains the compressed subspace. This guarantees that, after applying the fC_x operator, the compressed oracle will have a well-defined output for x.

Lemma 2.7. *For any $|\psi\rangle \in \Xi$, we have $\mathsf{cfO}\,|\psi\rangle \in \Xi$.*

The following further properties of compressed oracles are not, to our knowledge, shown in the literature. However, they have straightforward proofs, which we include.

Lemma 2.8. *Fix some input $x \in [M]$. For each $D \in \mathbf{D}|^x$, choose a set $S_{x,D} \subseteq [N]$ such that $|S_{x,D}| \geq N - t$. Let $|S_{x,D}\rangle = \frac{1}{\sqrt{|S_{x,D}|}} \sum_{y \in S_{x,D}} |D[x \to y]\rangle$.*

Define the operator

$$G_{x,D} := \mathbb{1} - |D\rangle\langle D| - |S_{x,D}\rangle\langle S_{x,D}| + |D\rangle\langle S_{x,D}| + |S_{x,D}\rangle\langle D|, \tag{42}$$

$$G_x := \bigoplus_{D \in \mathbf{D}|^x} G_{x,D}. \tag{43}$$

Then we have

$$\|\mathsf{fC}_x - G_x\| = O\left(\sqrt{\frac{t}{N}}\right). \tag{44}$$

Proof. First, let us fix a database D with $x \notin D$ and use the shorthand $|+_{x,D}\rangle$ for the even superposition $\sum_{y \in [N]} \frac{1}{\sqrt{N}} D[x \to y]$. Then we have

$$\begin{aligned} \||+_{x,D}\rangle - |S_{x,D}\rangle\|^2 &= \sum_{y \notin S_{x,D}} \frac{1}{N} + \sum_{y \in S_{x,D}} \left(\frac{1}{\sqrt{N}} - \frac{1}{\sqrt{|S_{x,D}|}}\right)^2 && (45)\\ &\leq \frac{t}{N} + N \cdot \left(\frac{1}{\sqrt{N}} - \frac{1}{\sqrt{N-t}}\right)^2 && (46)\\ &= O\left(\frac{t}{N}\right). && (47) \end{aligned}$$

We may then write

$$\begin{aligned} \mathsf{fC}_{x,D} - G_{x,D} = (|+_{x,D}\rangle\langle+_{x,D}| - |S_{x,D}\rangle\langle S_{x,D}|) + && (48)\\ |D\rangle\left(\langle+_{x,D}| - \langle S_{x,D}|\right) + \left(|+_{x,D}\rangle - |S_{x,D}\rangle\right)\langle D|. && (49) \end{aligned}$$

By inspecting each term and applying the triangle inequality, one can verify that

$$\|\mathsf{fC}_{x,D} - G_{x,D}\| = O(\sqrt{t/N}). \tag{50}$$

The claim now follows from Lemma 2.3. □

Lemma 2.9. *Fix an input $x \in [N]$ for a compressed database of a random function, and let R denote an ancillary register. Let $|\phi_y\rangle_{\mathsf{R}}$ be a family of states indexed by $y \in [N]$ which are pairwise orthogonal, and consider a state of the form*

$$|\psi\rangle_{\mathsf{DR}} = \sum_{D \in \mathbf{D}, D(x) \neq \perp} \alpha_D\, |D\rangle_{\mathsf{D}}\, |\phi_{D(x)}\rangle_{\mathsf{R}}. \tag{51}$$

In other words, there is a record of the output of x outside the database register. Then, for any operator $G_{x,\mathsf{D}}$ of the form described in Lemma 2.8, we have

$$\|(\mathbb{1}_{\mathsf{DR}} - G_{x,\mathsf{D}})\,|\psi\rangle_{\mathsf{DR}}\| = O\left(\sqrt{\frac{1}{N-t}}\right). \tag{52}$$

Proof. We can write

$$\begin{aligned}\|(\mathbb{1}_{\mathsf{DR}} - G_{x,\mathsf{D}})\,|\psi\rangle_{\mathsf{DR}}\| &= \left\| \sum_{D\in\mathbf{D},D(x)\neq\perp} \frac{\alpha_D}{\sqrt{|S_{x,D[x\to\perp]}|}} \left(|D[x\to\perp]\rangle - |S_{x,D[x\to\perp]}\rangle\right)_{\mathsf{D}} |\phi_{D(x)}\rangle_{\mathsf{R}} \right\| \\ &= \sqrt{\sum_{D\in\mathbf{D},D(x)\neq\perp} \frac{|\alpha_D|^2}{|S_{x,D[x\to\perp]}|}} \\ &= O\left(\sqrt{\frac{1}{N-t}}\right), \end{aligned} \tag{53}$$

where the second line follows from the fact that states in the sum corresponding $D \neq D' \in \mathbf{D}$ will be orthogonal. To see this, observe that we could define an isometry which combined every input-output pair except $D(x)$ from register D, and determined $D(x)$ from register R, to construct D in a new register. □

3 Compressed permutation oracles

In this section, we present our compressed permutation oracle framework. The construction itself is similar to the function case, though establishing soundness turns out to be much more difficult. Our construction is similar to that of Unruh [Unr23], though we explicitly and unitarily maintain injectivity of the database. We begin with basic definitions of permutation oracles.

Let $\varphi \sim \mathbf{S}_N$ be a uniform random permutation, and consider the oracle

$$\mathcal{O}_\varphi\,|b,x\rangle_{\mathsf{X}}\,|y\rangle_{\mathsf{Y}} := |b,x\rangle_{\mathsf{X}}\,|y \oplus \varphi^{1-2b}(x)\rangle_{\mathsf{Y}}\,, \tag{54}$$

for a bit b and $x, y \in [N]$. In other words, $\mathcal{O}_\varphi$ is a standard oracle for the permutation φ with direction controlled by the bit b, on query point x. We do not explicitly consider controlled queries (i.e., the possibility of not querying the oracle), but we note that all our proofs apply to this case as well.[8] We will throughout consider algorithms which may query both φ and φ^{-1}, as $\mathcal{O}_\varphi$ exposes both directions.

Let $\mathcal{A}^{\mathcal{O}_\varphi}$ be a quantum algorithm that makes q queries to $\mathcal{O}_\varphi$. This algorithm is described by a Hilbert space $\mathcal{H}_{\mathsf{A}}$ for its internal state, as well as a sequence $A_0, \ldots, A_q$ of

[8] The easiest way to see this is to replace controlled queries with a controlled swap of the output register (holding y) with the uniform superposition $|+_{[N]}\rangle$, before and after the query. A standard oracle acts as identity when the output register is in uniform superposition.

unitary operators on $\mathcal{H}_{\mathsf{A}}$, describing the action in between each query. We can write the final mixed state of the algorithm as

$$\rho_{\mathcal{A}}^{(\mathcal{O}_\varphi)} = \mathbb{E}_{\varphi\sim \mathbf{s}_N}\left[\rho(A_q\mathcal{O}_\varphi \dots A_1\mathcal{O}_\varphi A_0\,|0\rangle)\right]. \tag{55}$$

The compressed picture. Similar to the compressed function oracle, we will take a purifying database which begins empty. A query on input x will affect only the database entries involving x. In contrast to the function case, however, this modification will be *controlled* on the rest of the database to avoid collisions. This results in an operator which cannot be expressed as a direct product over database entries, which reflects the correlations between permutation outputs. Further, compression operators on different input points will no longer commute.

To formalize the action, fix for now some $I \in \mathbf{I}$ and $x \in [N] \setminus \mathsf{Dom}(I)$, and consider the subspace $\mathcal{H}(I|^x) \leq \mathcal{H}(\mathbf{I})$. This subspace reflects modifying only the x-th query point, though in this case we will perform this modification controlled on the rest of the database. Let us define

$$|+_{x,I}\rangle = \frac{1}{\sqrt{N-|I|}} \sum_{y\in[N]\setminus \mathsf{Im}(I)} |I[x\to y]\rangle\ . \tag{56}$$

We can now define the *permutation compression* operator pC on this subspace, indexed by x and I, as

$$\mathsf{pC}_{x,I} := \mathbb{1} - |+_{x,I}\rangle\langle+_{x,I}| - |I\rangle\langle I| + |+_{x,I}\rangle\langle I| + |I\rangle\langle+_{x,I}|. \tag{57}$$

In other words, this operator swaps the uniform superposition of *non-colliding* outputs with the undefined output, for the input x. Recalling that $\{I|^x\ :\ I\in\mathbf{I}|^x\}$ is a partition of $\mathbf{I}$, we may then define the full compression operator and its controlled variant as

$$\mathsf{pC}_x := \bigoplus_{I\in\mathbf{I}|^x} \mathsf{pC}_{x,I}, \qquad\qquad \mathsf{pC}_{\mathsf{XI}}\,|x\rangle_{\mathsf{X}} = |x\rangle_{\mathsf{X}}\otimes \mathsf{pC}_{x,\mathsf{I}}. \tag{58}$$

We recall the extended purified oracle P, which acts as

$$\mathsf{P}\,|x,y\rangle\,|I\rangle := \begin{cases} |x, y\oplus I(x)\rangle\,|I\rangle & (\text{If } x\in\mathsf{Dom}(I)) \\ |x,y\rangle\,|I\rangle & (\text{Otherwise}) \end{cases} \tag{59}$$

The final ingredient we need is the flip operator F, which inverts the database (viewed as a partial function). In other words, $I^{-1}(y) = x \neq \perp$ if and only if $I(x) = y$, else $I^{-1}(y) = \perp$. Here, the database is required to be injective, and so we can define the flip unitary

$$\mathsf{F}\,|I\rangle = |I^{-1}\rangle\,. \tag{60}$$

The full compressed permutation oracle is similar to the compressed function oracle, except we maintain injectivity of the database. This allows us to invert before answering

inverse queries. In particular, we define the full compressed permutation oracle as

$$\mathsf{cpO}\,|b\rangle := |b\rangle \otimes \begin{cases} \mathsf{pC}\cdot\mathsf{P}\cdot\mathsf{pC}^\dagger & (\text{If } b=0) \\ \mathsf{F}\cdot\mathsf{pC}\cdot\mathsf{P}\cdot\mathsf{pC}^\dagger\cdot\mathsf{F}^\dagger & (\text{If } b=1) \end{cases} \tag{61}$$

where again the operators pC and F are involutions, and so the dagger may be suppressed. We take the purifying register to be $\mathcal{H}_\mathsf{I} = \mathcal{H}(\mathbf{I})$. We refer to this register as the *compressed permutation oracle*, and it is initialized to $|\perp\rangle_\mathsf{I}$. We additionally replace oracle calls with the compressed oracle cpO. This corresponds to the final state

$$|\psi^{(\mathsf{cpO})}_{\mathcal{A}}\rangle_{\mathsf{AI}} = A_{q,\mathsf{A}}\mathsf{cpO}_{\mathsf{AI}}\ldots A_{1,\mathsf{A}}\mathsf{cpO}_{\mathsf{AI}}A_{0,\mathsf{A}}\,|0\rangle_\mathsf{A}\,|\perp\rangle_\mathsf{I}\,. \tag{62}$$

Throughout the rest of this document, we will establish that this simple construction satisfies properties analogous to those of the compressed random oracle [Zha19].

3.1 Properties

The key properties we will show are that our compressed oracle is:

1. *Sound*, in that it faithfully represent algorithms querying two-way accessible random permutations.
2. *Bounded*, in that the size of the compressed database grows at most linearly with the number of queries.
3. *Meaningful*, in that the points defined by the database correspond to points the algorithm "knows".

It is fairly straightforward to establish bounded growth and the fundamental lemma, points (2) and (3) above, which we do in this section. For now, we only state soundness; proving it will be quite technically involved, and is the focus of the next two sections. Note also that soundness is only approximate.

Lemma 3.1 (Soundness)**.** *The view of an efficient algorithm querying the compressed permutation oracle is negligibly close to the view of the same algorithm querying a truly random permutation, i.e.,*

$$\rho^{(\mathcal{O}_\varphi)}_{\mathcal{A},\mathsf{A}} \overset{\text{negl}}{\approx} \mathrm{Tr}_\mathsf{I}[\rho(|\psi^{(\mathsf{cpO})}_{\mathcal{A}}\rangle_{\mathsf{AI}})]. \tag{63}$$

Proof. Follows from Theorem 5.19. □

Lemma 3.2 (Bounded growth)**.** *The final state of the compressed permutation oracle after q queries is supported solely on databases of size q,*

$$\Pi_{\leq q,\mathsf{I}}\,|\psi^{(\mathsf{cpO})}_{\mathcal{A}}\rangle_{\mathsf{AI}} = |\psi^{(\mathsf{cpO})}_{\mathcal{A}}\rangle_{\mathsf{AI}}\,. \tag{64}$$

Proof. Suppose that the query register X holds the state $|0,x\rangle_{\mathsf{X}}$. Then, the cpO operator preserves the subspace $\mathcal{H}(I|^x)$ on the purifying I register. It follows that the operator can increase the size of the purifying register by at most 1. A similar argument follows for the case of an inverse query. □

To establish the fundamental lemma, we will once again interpret the adversary as outputting a tuple of l input-output pairs. We will compare the probabilities that these pairs appear in the database with the probability that they agree with the (compressed) oracle. To do this, we define the projector

$$\Pi_{\mathsf{AI}}^{(\mathsf{cpO})} := \left(\sum_{\{(x_1,y_1),\ldots,(x_l,y_l)\}\in[N]^{\times 2l}} \Pi^{(x_1,y_1)\in\mathsf{db}} \ldots \Pi^{(x_l,y_l)\in\mathsf{db}} \right)_{\mathsf{I}} \otimes \tag{65}$$

$$(|x_1,y_1,\ldots,x_l,y_l\rangle\langle x_1,y_1,\ldots,x_l,y_l| \otimes \mathbb{1})_{\mathsf{A}}. \tag{66}$$

We will also require the the analogous operator after decompressing,

$$\Pi_{\mathsf{AI}}^{(\mathcal{O}_\varphi)} := \sum_{\{(x_1,y_1),\ldots,(x_l,y_l)\}\in[N]^{\times 2l}} \left(\mathsf{pC}_{x_1} \Pi^{(x_1,y_1)\in\mathsf{db}} \mathsf{pC}_{x_1}^\dagger \ldots \mathsf{pC}_{x_l} \Pi^{(x_l,y_l)\in\mathsf{db}} \mathsf{pC}_{x_l}^\dagger \right)_{\mathsf{I}} \otimes \tag{67}$$

$$(|x_1,y_1,\ldots,x_l,y_l\rangle\langle x_1,y_1,\ldots,x_l,y_l| \otimes \mathbb{1})_{\mathsf{A}}. \tag{68}$$

Once soundness of the compressed permutation oracle is established, this check corresponds to verifying that the adversary output is consistent with the oracle. We can now state and prove the fundamental lemma. Due to the non-commutativity of compression operators, the dependence on l is quadratically worse than in the function case, though this may not be tight.

Lemma 3.3 (Fundamental lemma)**.** *The probability that an adversary finds some list of l input-output pairs of the oracle after t queries is exponentially close to the probability that the pair appears in the compressed database,*

$$\left\| \Pi_{\mathsf{AI}}^{(\mathcal{O}_\varphi)} |\psi_{\mathcal{A}}^{(\mathsf{cpO})}\rangle_{\mathsf{AI}} \right\| \leq \left\| \Pi_{\mathsf{AI}}^{(\mathsf{cpO})} |\psi_{\mathcal{A}}^{(\mathsf{cpO})}\rangle_{\mathsf{AI}} \right\| + \frac{l}{\sqrt{N-t-l}}. \tag{69}$$

Proof. We have

$$\left\| \Pi_{\mathsf{AI}}^{(\mathcal{O}_\varphi)} |\psi_{\mathcal{A}}^{(\mathsf{cpO})}\rangle_{\mathsf{AI}} \right\| - \left\| \Pi_{\mathsf{AI}}^{(\mathsf{cpO})} |\psi_{\mathcal{A}}^{(\mathsf{cpO})}\rangle_{\mathsf{AI}} \right\| \leq \left\| (\Pi_{\mathsf{AI}}^{(\mathcal{O}_\varphi)} - \Pi_{\mathsf{AI}}^{(\mathsf{cpO})}) |\psi_{\mathcal{A}}^{(\mathsf{cpO})}\rangle_{\mathsf{AI}} \right\| \tag{70}$$

$$\leq \left\| \Pi_{\mathsf{AI}}^{(\mathcal{O}_\varphi)} - \Pi_{\mathsf{AI}}^{(\mathsf{cpO})} \right\|_{\leq t}. \tag{71}$$

Note that both these operators are controlled on A, so by Lemma 2.3 it suffices to analyze the difference for some fixed set $\mathbf{x} = \{(x_1,y_1),\ldots,(x_l,y_l)\}$. We then obtain

$$\left\| \Pi_{\mathbf{x},\mathsf{I}}^{(\mathcal{O}_\varphi)} - \Pi_{\mathbf{x},\mathsf{I}}^{(\mathsf{cpO})} \right\|_{\leq t} \leq \sum_{i=1}^{l} \left\| \Pi_{(x_i,y_i)\in\mathsf{db},\mathsf{I}} - \mathsf{pC}_{x_i,\mathsf{I}} \Pi_{(x_i,y_i)\in\mathsf{db},\mathsf{I}} \mathsf{pC}_{x_i,\mathsf{I}}^\dagger \right\|_{\leq t+l} \tag{72}$$

by repeated applications of triangle inequality. We can analyze each term individually. Note that the operator $\Pi_{(x,y)\in\mathsf{db},\mathsf{I}} - \mathsf{pC}_{x,\mathsf{I}}\Pi_{(x,y)\in\mathsf{db},\mathsf{I}}\mathsf{pC}^{\dagger}_{x,\mathsf{I}}$ preserves the subspace $\mathcal{H}(I|^x)$ for any $I \in \mathbf{I}|^x$. Therefore, it suffices to consider the action on a state

$$|\phi\rangle_{\mathsf{I}} = \sum_{y'\in[N]\setminus\mathsf{Im}(I)\cup\{\perp\}} \alpha_{y'} |I[x \to y']\rangle_{\mathsf{I}}, \tag{73}$$

for some $I \in \mathbf{I}|^x$ of size at most $t + l$. We write $s = \frac{1}{\sqrt{N-|\mathsf{Im}(I)|}} \sum_{y'\in[N]\setminus\mathsf{Im}(I)} \alpha_{y'}$. We can then compute

$$\Pi_{(x,y)\in\mathsf{db},\mathsf{I}} |\phi\rangle_{\mathsf{I}} = \alpha_y |I[x \to y]\rangle \tag{74}$$

$$\mathsf{pC}_{x,\mathsf{I}}\Pi_{(x,y)\in\mathsf{db},\mathsf{I}}\mathsf{pC}^{\dagger}_{x,\mathsf{I}} |\phi\rangle_{\mathsf{I}} = \mathsf{pC}_{x,\mathsf{I}} \left(\alpha_y + \frac{\alpha_\perp - s}{\sqrt{N-|\mathsf{Im}(I)|}}\right) |I[x \to y]\rangle \tag{75}$$

$$= \left(\alpha_y + \frac{\alpha_\perp - s}{\sqrt{N-|\mathsf{Im}(I)|}}\right) \cdot \left(|I[x \to y]\rangle - \frac{1}{\sqrt{N-|\mathsf{Im}(I)|}} |+_{x,I}\rangle + \frac{1}{\sqrt{N-|\mathsf{Im}(I)|}} |I\rangle\right) \tag{76}$$

$$= \alpha_y |I[x \to y]\rangle_{\mathsf{I}} + O\left(\frac{1}{\sqrt{N-|\mathsf{Im}(I)|}}\right). \tag{77}$$

where the last line follows from $|s|, |\alpha_\perp| \leq 1$. Recalling that $N - |\mathsf{Im}(I)| \geq N - l - t$,

$$\left\| \Pi_{(x,y)\in\mathsf{db},\mathsf{I}} |\phi\rangle_{\mathsf{I}} - \mathsf{pC}_{x,\mathsf{I}}\Pi_{(x,y)\in\mathsf{db},\mathsf{I}}\mathsf{pC}^{\dagger}_{x,\mathsf{I}} |\phi\rangle_{\mathsf{I}} \right\| = O\left(\frac{1}{\sqrt{N-l-t}}\right), \tag{78}$$

which combined with Equation (72) proves the claim. □

4 Purified permutations

In this section, we introduce the ensemble of permutations which we will analyze to prove both soundness of our construction and of Feistel, and prove properties about the purified ensemble. We use φ to denote the permutation which is being queried. We will write φ as a product of constituent permutations. Some of the constituent permutations will be highly structured, which will enable us to apply the standard compressed oracle theory.

In particular, we consider an ensemble defined by two random permutations $\pi, \omega \in \mathbf{S}_N$, which we refer to as the outer or twirling permutations, and three random functions $h, k, f : \{0,1\}^n \to \{0,1\}^n$, which we refer to as the internal or inner Feistel. We will always take h, k, f to be uniform random functions, but we do not yet commit to a particular distribution for π and ω. This tuple defines the permutation $\varphi_{(\pi,\omega,h,k,f)}$. The formula for this permutation is $\varphi_{(\pi,\omega,h,k,f)}(x) = y$, where

$$u = \pi(x), \qquad v = \omega^{-1}(y), \tag{79}$$

$$m \oplus u_R = h(u_L), \qquad v_L \oplus u_L = k(m), \qquad m \oplus v_R = f(v_L). \tag{80}$$

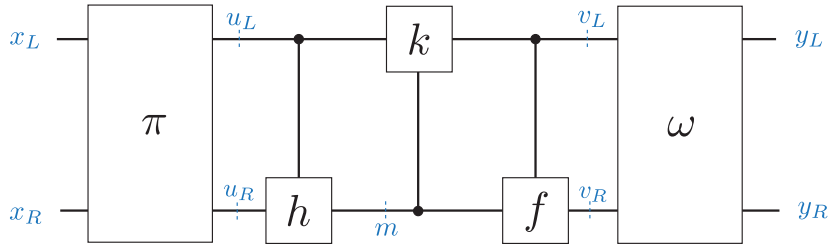


Figure 3: The masked feistel construction, with commonly chosen variable names in blue.

This is depicted in Figure 3, along with the variables above. For a generic value on the left wire, we will usually use w or l; on the right wire, usually z or r.

We call this construction *masked Feistel*, as the outer permutations will hide the structure of the internal Feistel construction. We will analyze this ensemble by purifying. The purification register lives in the space $\mathcal{H}_{\mathsf{P}} = \mathcal{H}(\mathbf{S}_N)^{\otimes 2} \otimes \mathcal{H}(\mathbf{F}_{\sqrt{N},\sqrt{N}})^{\otimes 3}$ for the permutations π, ω and functions h, k, f respectively. We label registers holding the parts of a purification by $\Pi, \Omega, \mathsf{H}, \mathsf{K}, \mathsf{F}$ respectively. For concreteness, it may help to imagine π, ω being chosen at uniform random; this is the case we will analyze in our soundness proof. In this case, the permutation φ is uniform random, as it is the product of a uniform random permutation (say π) and an independently chosen permutation.

4.1 Query analysis

Let us define the purified query operator mfP on this space to act as

$$\mathsf{mfP} \left|x, y\right\rangle_{\mathsf{XY}} \left|\pi, \omega, h, k, f\right\rangle_{\mathsf{P}} := \left|x, y \oplus \varphi_{(\pi,\omega,h,k,f)}\right\rangle_{\mathsf{XY}} \left|\pi, \omega, h, k, f\right\rangle_{\mathsf{P}}. \tag{81}$$

Looking ahead, we will answer inverse queries by conjugating this operator by a flip of the truth table; for now we focus on forwards queries for simplicity. Let $\mathcal{D}$ be the distribution on $\mathbf{S}_N^{\times 2}$ from which the twirling permutations are drawn. We can write the initial purified state

$$\left|\psi_0\right\rangle_{\mathsf{P}} = \sum_{\pi,\omega \in \mathbf{S}_N} \sqrt{p_{\mathcal{D}}(\pi,\omega)} \left|\pi, \omega\right\rangle_{\Pi\Omega} \otimes \frac{1}{\sqrt{(2^n)^{3 \cdot 2^n}}} \sum_{h,k,f \in \mathbf{F}_{2^n,2^n}} \left|h, k, f\right\rangle_{\mathsf{HKF}}. \tag{82}$$

Then, answering queries using the mfP operator and tracing out the purifying register at the end of the experiment produces the same view as randomly selecting φ from the appropriate distribution.

Now, we may use compressed oracle representations for the internal functions HKF. Answering a query on input x in this compressed picture can be thought of as the following out-of-place circuit:

1. Compute $u = \pi(x)$ into a fresh register
2. Call the compressed oracle on H with input u_L, and write the output z to a fresh register.

3. Call the compressed oracle on K with input $u_R \oplus z$, and write the output w to a fresh register.
4. Call the compressed oracle on F with input $u_L \oplus w$, write the output z' to a fresh register.
5. Compute $v = u \oplus (w \| (z \oplus z'))$, and output $\omega(v)$.
6. Uncompute all intermediate values.

However, this circuit has many redundant steps. For instance, at the end of step (2) the database in register H is compressed on input u_L. The next operation which acts on this register is a decompression on u_L, in step (6). These operations are inverses of one another, and so they cancel out. Similarly, step (2) writes z to a fresh wire, only for the next step to be a decompression conditioned on z. It will instead be easier to imagine performing the decompression in steps (2-4) without writing the result to a new register, nor performing the final compression in these steps. Then, the query can be answered in step (5) using information in the databases contain in registers HKF. Finally, we undo by compressing steps (2-4), in the reverse order.

We will call this sequence of decompressions the *masked Feistel decompression*, writing $\mathsf{mfC}^\dagger$, and mfC for the corresponding masked Feistel compression. At an intuitive level, this decompression operator performs the following, where we use D_h, D_k, D_f to denote the databases stored in $\mathsf{H}, \mathsf{K}, \mathsf{F}$ respectively. We introduce some auxiliary variables only for notational convenience; one should not think of these as being stored in any register.

1. Decompress H on input u_L, where $u = \pi(x)$.
2. Decompress K on input $u_R \oplus D_h(u_L)$.
3. Decompress F on input $u_L \oplus D_k(u_R \oplus D_h(u_L))$.

Note that the database entries referenced are guaranteed to exist here, as they were just decompressed, and the compressed oracle maintains validity. After this operator is performed, the value $\varphi(x)$ is guaranteed to be determined, and can be XORed into the target query register Y. Following this, undoing the above steps in reverse order will re-compress the database. This method of answering queries is equivalent to the out-of-place circuit described above.

We will now define the compression operator mfC more formally. Let

$$U^{(H)}_{u,\mathsf{H}} := \mathsf{fC}_{u_L,\mathsf{H}} \tag{83}$$

$$U^{(K)}_{u,\mathsf{HK}} |D_h\rangle_{\mathsf{H}} := |D_h\rangle_{\mathsf{H}} \otimes \begin{cases} \mathbb{1}_{\mathsf{K}} & (\text{If } u_L \notin \mathsf{Dom}(D_h)) \\ \mathsf{fC}_{u_R \oplus D_h(u_L),\mathsf{K}} & (\text{Otherwise}) \end{cases} \tag{84}$$

$$U^{(F)}_{u,\mathsf{HKF}} |D_h, D_k\rangle_{\mathsf{HK}} := |D_h, D_k\rangle_{\mathsf{HK}} \otimes \begin{cases} \mathbb{1}_{\mathsf{F}} & (\text{If } u_L \notin \mathsf{Dom}(D_h)) \\ \mathbb{1}_{\mathsf{F}} & (\text{If } u_R \oplus D_h(u_L) \notin \mathsf{Dom}(D_k)) \\ \mathsf{fC}_{u_L \oplus D_k(u_R \oplus D_h(u_L)),\mathsf{F}} & (\text{Otherwise}) \end{cases} \tag{85}$$

$$U_{u,\mathsf{HKF}} := U^{(H)}_{u,\mathsf{H}} U^{(K)}_{u,\mathsf{HK}} U^{(F)}_{u,\mathsf{HKF}}, \tag{86}$$

where U_u compresses the internal databases on input u. Note that the choice of $\mathbb{1}$ in the above cases is arbitrary; as described above, when these operators are called in the context of the masked Feistel construction the values are guaranteed to exist, and so these cases do not occur.

We can now define mfC by applying the twirl, to obtain

$$\mathsf{mfC}_{x,\Pi\mathsf{HKF}} |\pi\rangle_\Pi := |\pi\rangle_\Pi \otimes U_{\pi(x),\mathsf{HKF}}, \tag{87}$$

$$\mathsf{mfC}_{\mathsf{X}\Pi\mathsf{HKF}} |x\rangle_\mathsf{X} := |x\rangle_\mathsf{X} \otimes \mathsf{mfC}_{x,\Pi\mathsf{HKF}}. \tag{88}$$

We also define an operator analogous to the flip operator, which we denote by mfF, which acts as

$$\mathsf{mfF}_\mathsf{P} |\pi\rangle_\Pi |\omega\rangle_\Omega |D_h\rangle_\mathsf{H} |D_k\rangle_\mathsf{K} |D_f\rangle_\mathsf{F} := |\omega^{-1}\rangle_\Pi |\pi^{-1}\rangle_\Omega |D_f\rangle_\mathsf{H} |D_k\rangle_\mathsf{K} |D_k\rangle_\mathsf{F}, \tag{89}$$

in other words this operator inverts and exchanges both twirl permutations, and interchanges the databases in F and H. This operator is an involution, and if a purification P represents permutation φ_P, then the flipped purification P' where $\mathsf{mfF} |P\rangle = |P'\rangle$ represents the permutation $\varphi_P^{-1} = \varphi_{P'}$. Now, we can define the *compressed masked Feistel oracle* cmfO as

$$\mathsf{cmfO} |b\rangle := |b\rangle \otimes \begin{cases} \mathsf{mfC} \cdot \mathsf{mfP} \cdot \mathsf{mfC}^\dagger & (\text{If } b = 0) \\ \mathsf{mfF} \cdot \mathsf{mfC} \cdot \mathsf{mfP} \cdot \mathsf{mfC}^\dagger \cdot \mathsf{mfF}^\dagger & (\text{If } b = 1) \end{cases} \tag{90}$$

Observe that in this case we have $\mathsf{mfC} \neq \mathsf{mfC}^\dagger$, as the order of compression/decompressions on the HKF databases are flipped. This is in contrast with the compressed function and compressed permutation oracles, where compression is an involution.

We use notation $D = (D_h, D_k, D_f) \in \mathbf{D}_{2^n,2^n}^{\times 3}$ for the compressed databases of functions h, k, f, and think of $P = (\pi, \omega, D)$ as a possible state of the purification registers. We use $\mathbf{P}$ to denote set of possible purifications (i.e. $P \in \mathbf{P}$). The size of a purification is the size of its largest database, $|P| = \max\{|D_h|, |D_k|, |D_f|\}$. We will define a number of desirable properties of both the twirling permutations (π, ω) and the internal Feistel databases (D_h, D_k, D_f), beginning with the latter.

4.2 Internal Feistel

In this section, we will analyze the internal three rounds of Feistel. In particular, our main object of study will be triples $D = (D_h, D_k, D_f)$ of databases. A core concept is a chain, which is a set of database points which suffice to define an input-output pair of the Feistel permutation.

Definition 4.1. *Fix a $D = (D_h, D_k, D_f)$. We say there is a* chain *from u to v in D if $D_h(u_L) \oplus u_R = D_f(v_L) \oplus v_R = m$, and $u_L \oplus v_L = D_k(m)$. The values (u, m, v) define the chain, and (u_L, m, v_L) are the input points along the chain.*

We also say that the *rightward semi-chain* from u is length 0 if $D_h(u_L) = \perp$, is length 1 if $D_k(D_h(u_L) \oplus u_R) = \perp$, is of length 2 if $D_f(u_L \oplus D_k(D_h(u_L) \oplus u_R)) = \perp$, and is of length 3 (and therefore a full chain to some v) otherwise.

We say that D *supports* the pair (u, v) if there is a chain from u to v. We use $\mathsf{Supp}(D)$ to denote the injective database supported by the D tuple, and $\mathsf{Dom}(D)$ and $\mathsf{Im}(D)$ to denote the domain and range of this database (i.e. we write $\mathsf{Dom}(D)$ in place of $\mathsf{Dom}(\mathsf{Supp}(D))$). We write $D(u)$ in place of $\mathsf{Supp}(D)(u)$, and similar for $D^{-1}(v)$. Finally, for $u \in \mathsf{Dom}(D)$ we write $D[u \to \bot]$ to denote D with the chain beginning at u removed.

Two chains *collide* if any of the function inputs along the chain collide.

Definition 4.2. *If $(u, m, v) \neq (u', m', v')$ are distinct chains in D, then they* collide *if $m = m'$, $u_L = u'_L$, or $v_L = v'_L$.*

A database triple $D = (D_h, D_k, D_f)$ is *canonical* if no two chains in D collide, and every input output pair in each database participates in exactly one chain. The set of all canonical database triples is denoted $\mathbf{C} \subset \mathbf{D}^{\times 3}_{2^n,2^n}$, and $\mathbf{C}_t$ denotes the set of canonical database where each D_h, D_k, and D_f are of size t.

Lemma 4.3. *Let $D = (D_h, D_k, D_f) \in \mathbf{C}_t$ be a canonical triple of databases of size t each. Then there are t full chains $u_1 \to v_1, \ldots, u_t \to v_t$ with distinct u_i and distinct v_i. There are $t(t-1)$ length 2 rightward semi-chains, $t \cdot (2^n - t)$ of length 1, and $(2^n - t) \cdot 2^n$ of length 0. There are the same number of leftward semi-chains of the appropriate length.*

Proof. The set of chains being size t is the definition of a canonical database. To compute the number of length 2 semi-chains, we choose any $w \in D_h$, and z such that $z \oplus D_h(w) \in D_k$, but which does not lead to an already queried value in D_f. We know that $w = (u_i)_L, z = (u_i)_R$ satisfy the first two conditions but not the third, but if a different $z' \neq z$ lead to a queried value in D_f then we would have a colliding chain. This would violate canonicity, so we may choose any $z' \neq z$, of which there are $t - 1$ satisfying the remaining constraints.

For the number of semi-chains of length 1, we choose $w \in D_h$ and any z s.t. $D_h(w) \oplus z \notin D_k$. There are t choices for w, and $2^n - t$ choices for z given w. Similarly, for semi-chains of length 0 we may pick any $w \notin D_h$ of which there are $2^n - t$, and any z of which there are 2^n. We can verify that

$$t + t(t-1) + t \cdot (2^n - t) + (2^n - t) \cdot 2^n = t^2 + (t + 2^n) \cdot (2^n - t) \tag{91}$$

$$= 2^{2n}. \tag{92}$$

□

Given some injective database $I \in \mathbf{I}_t$, let $\mathbf{D}(I)$ denote the set of canonical databases which support I. Note that this set is empty if any one of the following are true.

Definition 4.4. *An $I \in \mathbf{I}_t$ is* non-allowable *if any of the following predicates are true.*

1. *There is an* input left collision *in I if there are distinct $x, x' \in \mathsf{Dom}(I)$ such that $x_L = x'_L$.*
2. *There is an* output left collision *in I if there are distinct $y, y' \in \mathsf{Im}(I)$ such that $y_L = y'_L$.*
3. *There is an* internal left collision *in I if there is an $(x, y) \in I, x' \in \mathsf{Dom}(I), y' \in \mathsf{Im}(I)$ with $(x, y) \neq (x', y')$ and $x_L \oplus y_L = x'_L \oplus y'_L$.*

We say I is allowable if none of these events occur, i.e. $\mathbf{D}(I) \neq \emptyset$, and denote by $\mathbf{A}_t \subset \mathbf{I}_t$ the set of allowable databases.[9]

Lemma 4.5. *Definition 4.4 characterizes all non-allowable databases.*

Proof. Suppose that I does not satisfy Definition 4.4. If there is an input or output left collision, then there cannot be disjoint chains for the left colliding inputs or outputs. Suppose then that there are not input or output left collisions, yet there is an internal collision. Let $(u, v) \in I, u' \in \mathsf{Dom}(I), v' \in \mathsf{Im}(I)$ be such a collision. Suppose that D were a canonical database supporting I, such that there is a chain from u to v. We write this chain as

$$\underbrace{(u_L, z_h)}_{\in D_h} \mapsto \underbrace{(u_R \oplus z_h, w_k)}_{\in D_k} \mapsto \underbrace{(u_L \oplus w_k, z_f)}_{\in D_f}, \tag{93}$$

where we observe that $w_k = u_L \oplus v_L$. Now observe that $u'_L \in \mathsf{Dom}(D_h)$ and $v'_L \in \mathsf{Dom}(D_f)$, so let us define

$$z'_h = D_h(u'_L), \qquad z'_f = D_f(v'_L). \tag{94}$$

Consider the inputs $u'' = u'_L \| (z'_h \oplus u_R \oplus z_h)$ and $v'' = v'_L \| (z'_f \oplus u_R \oplus z_h)$. There is then a chain from u'' to v'' of the form

$$\underbrace{(u''_L, z'_h)}_{\in D_h} \mapsto \underbrace{(u_R \oplus z_h, w_k)}_{\in D_k} \mapsto \underbrace{(v''_L, z'_f)}_{\in D_f}, \tag{95}$$

where we use the fact that $u'_L \oplus w_k = v'_L = v''_L$. This chain collides with the chain found from u to v. We know either $u'' \neq u$ (if $u' \neq u$, and assumption of no input left collisions) or $v'' \neq v$ (if $v' \neq v$, and assumption of no output left collisions), so there is a colliding chain in D, a contradiction.

Now let us consider the other direction. Suppose that I does satisfy Definition 4.4, we will construct a database D which supports exactly I and is canonical. Fix some ordering on the input output pairs as $I = (u_1, v_1), (u_2, v_2), \ldots, (u_t, v_t)$. Then, beginning with an empty $D = (D_h, D_k, D_f)$, iteratively construct the database as in Algorithm 1.

Algorithm 1: Database extension

1 For every $j \in [t]$:
2 Choose some z_j such that $(u_j)_R \oplus z_j \notin \mathsf{Dom}(D_k)$, and add $((u_j)_L, z_j)$ to D_h.
3 Add $((u_j)_R \oplus z_j, (u_j)_L \oplus (v_j)_L)$ to D_k.
4 Add $((v_j)_L, (v_j)_R \oplus z_j \oplus (u_j)_R)$ to D_f.

After $j - 1$ iterations, step (2) has $2^n - j + 1$ choices, and hence will succeed (unless $t \geq 2^n$, in which case database states are anyways not defined). If step (2) succeeds, then it will create new length 1 right facing semi-chains beginning at $(u_j)_L \| z_R$ for every $z_R \in \{0,1\}^n$. Of these, $j - 1$ will be extended immediately to length 2, corresponding to

[9] In fact, condition (3) entails conditions (1) and (2), though we explicitly state all three for clarity.

the values already in D_k. On the left wire, these length 2 right facing semi-chains lead to the values $(u_j)_L \oplus (u_i)_L \oplus (v_i)_L$ for $i \in [j-1]$. From the allowability of I, we know that these values will never be extended to full length 3 chains, as $(u_j)_L \oplus (u_i)_L \oplus (v_i)_L \neq (v_p)_L$ for any $p \in [t]$.

Step (3) will add a new pair to D_k if step (2) succeeds, and of the previously described length 1 right facing semi-chains, will extend only the one beginning at $u_j = (u_j)_L \| (u_j)_R$, as desired. However, it will also extend length 1 right facing semi-chains of the form $(u_i)_L \| (z_i \oplus (u_j)_R)$, for every $i \in [j-1]$. These now-length-2 right facing semi-chains will lead to a value $(u_i)_L \oplus (u_j)_L \oplus (v_j)_L$ on the left wire after the application of h and k. However, we know that this value is distinct from all $(v_p)_L$ for any $p \in [t]$, by the fact that I has no internal left collisions. Hence, these chains are never extended to length 3 chains. Therefore, each iterations adds a chain from u_j to v_j, disjoint from all prior chains, and adds no other chains. The databases are clearly minimal at the end of this procedure, so D is a canonical database tuple supporting I. □

Lemma 4.6. *For any $I \in \mathbf{A}_t$, we have the equation $|\mathbf{D}(I)| = \binom{2^n}{t}$.*

Proof. Consider the algorithm described in the proof of Lemma 4.5. In the first iteration of step (1), there are 2^n choices, in the second, $2^n - 1$, and etc. up to $2^n - t + 1$ in the final. Any choice which is not of the form described leads immediately to a chain collision, and hence is not allowable. Steps (2) and (3) do not allow any choices, as the values in D_k and D_f are constrained by the value of v. However, the choice of ordering of input-output pairs is arbitrary, leading to a factor $t!$ redundancy, and therefore a total of $\binom{2^n}{t}$ databases consistent with I. □

Definition 4.7. *Let $A \in \mathbf{A}_t$, and $u \notin \mathsf{Dom}(I)$. We say that A* allows *u if there exists a $v \in \{0,1\}^{2n}$ such that $A[u \to v] \in \mathbf{A}_{t+1}$. We denote this as $A \heartsuit u$, and if $A[u \to v]$ remains allowable then we write $A \heartsuit (u, v)$ and say (u, v) is allowed.*

Remark 4.8. *If A allows u, then u' with $u'_L = u_L$ is allowed as well. There are at least $2^n - t^2 - t$ allowable left substrings, by a direct counting argument.*

Definition 4.9. *Fix some $A \in \mathbf{A}_t$ and $u \in \{0,1\}^n$ which A allows. Let $\mathbf{V}_{u,A}$ denote the set of values $v \in \{0,1\}^{2n}$ such that $A[u \to v]$ is allowable, and let $\mathbf{L}_{u,A}$ denote the set of left substrings in $\mathbf{V}_{u,A}$. Formally:*

$$\mathbf{V}_{u,A} \coloneqq \{v \,:\, A[u \to v] \in \mathbf{A}_{t+1}\}, \tag{96}$$

$$\mathbf{L}_{u,A} \coloneqq \{v_L \,:\, v \in V_{u,A}\}. \tag{97}$$

We will sometimes use a length n bitstring $l \in \{0,1\}^n$ as the index $\mathbf{L}_{l,A}$ or $\mathbf{V}_{l,A}$. This is well defined, as these sets only depend on the left n bits of u.

Lemma 4.10. *For any $A \in \mathbf{A}_t$ and $u \in [N] \setminus \mathsf{Dom}(A)$ which A allows, we have*

$$|\mathbf{L}_{u,A}| \geq 2^n - 2t^2 - 2t, \tag{98}$$

and further for any $l \in \mathbf{L}_{u,A}$ and $r \in \{0,1\}^n$, we have $l \| r \in \mathbf{V}_{x,A}$.

Proof. The second statement is straightforward, as allowability does not constrain the right substring in any way. For the first, observe that an assignment $A[u \to v]$ can cause an internal left collision in one of two ways:

1. There exists $u' \in \mathsf{Dom}(A) \cup \{u\}$, $v' \in \mathsf{Im}(A) \cup \{v\}$ with $(u', v') \neq (u, v)$ such that
$$u'_L \oplus v'_L = u_L \oplus v_L \tag{99}$$
2. There exists $u' \in \mathsf{Dom}(A) \cup \{u\}$, $(u'', v'') \in A$ such that
$$u'_L \oplus v_L = u''_L \oplus v''_L \tag{100}$$

Note that in case (1), when $v = v'$ the equation reduces to the equality $u'_L = u_L$, which is not true by the fact that A allows u. Hence, there are $t^2 + t$ equations for the first ($t + 1$ for u', t for v'), which therefore eliminate at most $t^2 + t$ possible values for v_L. Similarly, there are $t + 1$ choices for u' and t choices for (u'', v'') in case (2), giving at most $2t^2 + 2t$ eliminated choices of v_L. Recall that the lack of internal left collisions implies no output left collisions (in particular, case (1) with $u = u'$ gives $v_L \neq v'_L$ for all $v \in \mathsf{Im}(A)$), so any choice not excluded leads to an allowable database. □

Database extensions. Now we turn to studying subsets of three round databases which can occur as intermediate states between canonical databases. We will refer to these as extensions. At a high level, these sets represent databases which would be canonical, except for one chain which is only partially defined.

Definition 4.11. *We say that a tuple D is a* one-extension *if there exists:*

1. *some canonical $D' = (D'_h, D'_k, D'_f)$,*
2. *some $l \in \{0,1\}^n$ such that D' allows $l \| r$ for some r, yet $l \| r \notin \mathsf{Dom}(D')$ for any r,*
3. *some $z \in \{0,1\}^n$,*

such that $D = (D'_h[l \to z], D'_k, D'_f)$.

We would say that D above is a one-extension of D' with (l, z). We use $\mathbf{E}(\mathbf{C}_t)$ to denote the set of one-extensions of size t canonical databases, and $\mathbf{E}(D)$ to denote the set of one-extensions of D. We define two-extensions similarly.

Definition 4.12. *We say that a tuple D is a* two-extension *if there exists:*

1. *some D' which is a one-extension of D'' with (l, z),*
2. *some $m \in \{0,1\}^n$ such that $m \notin \mathsf{Dom}(D'_k)$,*
3. *some $w \in \{0,1\}^n$ such that $l \oplus w \in \mathbf{L}_{l, A(D'')}$,*

such that $D = (D'_h, D'_k[m \to w], D'_f)$.

We would say that D above is a two-extension of D'' with $(l, z), (m, w)$, and a one-extension of D' with (m, w). We use $\mathbf{E}^2(\mathbf{C}_t)$ to denote the set of two-extensions of size t canonical databases, $\mathbf{E}^2(D)$ to denote the set of two-extensions of a $D \in \mathbf{C}$, and $\mathbf{E}(D')$ to denote the set of two-extensions obtained as one-extensions of $D' \in \mathbf{E}(\mathbf{C})$. Finally, we define a three extension,

Definition 4.13. *We say that a tuple D is a* three-extension *if there exists:*

1. *some D' which is a two-extension of D'' with $(l, z), (m, w)$,*
2. *some $z' \in \{0,1\}^n$,*

such that $D = (D'_h, D'_k, D'_f[l \oplus w \to z'])$.

Similarly, we use $\mathbf{E}^3(\mathbf{C}_t)$ to denote three-extensions of canonical databases, $\mathbf{E}^3(D)$ for $D \in C_t$ to denote three-extensions of D, $\mathbf{E}^2(D)$ for $D \in \mathbf{E}(\mathbf{C}_t)$ to denote three-extensions obtained as two-extensions of a fixed one-extension, and $\mathbf{E}(D)$ for $D \in \mathbf{E}^2(\mathbf{C}_t)$ to denote three-extensions obtained as one-extensions of a fixed two-extension.

Lemma 4.14. *Fix some $u \in [N]$. For any $t \geq 0$, canonical $D \in \mathbf{C}_{t+1}$ with $u \in \mathsf{Dom}(D)$, there exists a unique $D' \in \mathbf{C}_t$ with $u \notin \mathsf{Dom}(D)$ such that $D \in \mathbf{E}^3(D')$. Additionally, all three-extensions of canonical databases are themselves canonical.*

Proof. For the first statement, consider the database D' obtained by removing from D the chain beginning at u. Canonicity is preserved under removing a chain (this cannot add a chain collision, or make the database non-minimal), so $D \in \mathbf{E}^3(D')$. Further, every $D'' \neq D'$ will have either a chain beginning at $u' \in \mathsf{Dom}(D)$ and ending at a $v' \neq D(u')$, or a chain beginning at some $u'' \neq u$ and $u'' \notin \mathsf{Dom}(D)$. In either case, no extension can remove this chain, and such a chain does not appear in D, so $D'' \notin \mathbf{C}^3(D')$.

For the second statement, recall the correctness of Algorithm 1, proved in Lemma 4.5. In particular, an extension can be seen as one iteration of said algorithm, over all possible allowable (u, v) pairs. □

Now we prove some disjointness properties for extension sets. These properties will be useful later on to argue that the span of various extension sets will be orthogonal, and hence certain operators can be written in a direct sum manner. This in turn will be key to analyzing the complicated effect of compression operators on the databases when they are entangled with the other registers. Towards this, we introduce the pipe notation for database triples. We use an upwards arrow $\uparrow$ for subsets that will occur while evaluating forwards and/or adding output points, and a downwards arrow $\downarrow$ for subsets that occur while evaluating backwards and/or removing output points.

Definition 4.15. *Fix some $D \in \mathbf{C}_t$ and $u \in [N] \setminus \mathsf{Dom}(D)$ which D allows. Then we define*

$$D \uparrow_h^u := \{D\} \cup \{(D_h[u_L \to z], D_k, D_f) \,:\, z \in \{0,1\}^n \setminus \{u_R \oplus \mathsf{Dom}(D_k)\}\}. \tag{101}$$

For some $D' \in D \uparrow_h^u \subset \mathbf{E}(D)$ with extension (u_L, z) and $z \in \{0,1\}^n \setminus \{u_R \oplus \mathsf{Dom}(D_k)\}$, define

$$D' \uparrow_k^u := \{D'\} \cup \{(D'_h, D'_k[u_R \oplus z \to w], D'_f) \,:\, w \in \{u_L \oplus \mathbf{L}_{u,\mathsf{Supp}(D)}\}\}. \tag{102}$$

For some $D'' \in D' \uparrow^u_k \subset \mathbf{E}^2(D)$ with extension $(u_L, z), (u_R \oplus z, w)$ as above, we define

$$D'' \uparrow^u_f := \{D''\} \cup \{(D''_h, D''_k, D''_f[u_L \oplus w \to z']) : z \in \{0,1\}^n\}. \tag{103}$$

When this notation is applied to sets of databases, it is the union of the application to each individual one. We will collapse the notation $(D \uparrow^u_h) \uparrow^u_k = D \uparrow^u_{h,k}$, and similarly $((D \uparrow^u_h) \uparrow^u_k) \uparrow^u_f = D \uparrow^u_{h,k,f}$.

Lemma 4.16. *For any $D \neq D' \in \mathbf{C}$ and $u \in [N]$ which is allowed by D and D' but supported by neither, meaning $u \notin \mathsf{Dom}(D) \cup \mathsf{Dom}(D')$, we have*

$$D \uparrow^u_h \cap D' \uparrow^u_h = \emptyset, \qquad D \uparrow^u_{h,k} \cap D' \uparrow^u_{h,k} = \emptyset, \qquad D \uparrow^u_{h,k,f} \cap D' \uparrow^u_{h,k,f} = \emptyset. \tag{104}$$

Proof. There is WLOG a chain in D but not D' that does not involve u, and this chain cannot be removed by any extension, nor can it be added by any extension that does not originate from u. □

Now we turn to defining and analyzing the downwards arrow notation. In contrast to the upwards arrow, the downwards arrow notation should be thought of as changing or removing an existing input-output point.

Definition 4.17. *Let $D \in \mathbf{C}$ and $u \in \mathsf{Dom}(D)$. Then we define*

$$D \downarrow^u_h := \{(D_h[u_L \to z], D_k, D_f) : z \in \{0,1\}^n \cup \{\perp\} \setminus \{u_R \oplus \mathsf{Dom}(D_k)\}\} \tag{105}$$
$$D \downarrow^u_k := \{(D_h, D_k[u_R \oplus D_h(u_L) \to w], D_f) : w \in \{u_L \oplus \mathbf{L}_{u,D[u\to\perp]}\} \cup \{\perp\}\} \tag{106}$$
$$D \downarrow^u_f := \{(D_h, D_k, D_f[D(u)_L \to z]) : z \in \{0,1\}^n \cup \{\perp\}\} \tag{107}$$

Lemma 4.18. *For any $D \neq D' \in \mathbf{C}$ and $u \in \mathsf{Dom}(D) \cap \mathsf{Dom}(D')$, we have the disjointness relations*

$$D \downarrow^u_h \cap D' \downarrow^u_h = \emptyset, \qquad D \downarrow^u_k \cap D' \downarrow^u_k = \emptyset. \tag{108}$$

The identity $D \downarrow^u_f \cap D' \downarrow^u_f = \emptyset$ holds if $D[u \to \perp] \neq D'[u \to \perp]$.

Proof. Let us begin with the first equation. Observe that such a pair $D \neq D'$ has either disjoint chains not involving u, or $D(u) \neq D'(u)$. Let us begin with the first case. In this case, changing the image of u_L in D_h as described (or making it undefined) will modify the chain from u to $D(u)$ and/or u to $D'(u)$, changing the last n bits of the preimage. However, it will not create any further chains, as $D_h(u_L)$ is chosen such that it does not connect to any other pair in D_k. Hence, the disjoint chain not involving u will remain, and the databases will be distinct.

In the other case, consider reassignments of the form $D_h[u_L \to z]$ and $D'_h[u_L \to z']$. In such a case, there may be no chain from u to either $D(u)$ or $D'(u)$ in the redefined D or D' respectively. If one of z or z' is $\perp$, observe that there are exactly two input-output pairs, $(r, w) \in D_k$ and $(u_L \oplus w, z_f) \in D_f$ which do not participate in a chain. We have $u_L \oplus w \| r \oplus z_f = D(u)$. Similarly, under the redefinition $D'_h[u_L \to z]$ there are exactly two input-output pairs, $(r', w') \in D'_k$ and $(u_L \oplus w', z'_f) \in D'_f$ which do not participate in a chain. We have $u_L \oplus w' \| r' \oplus z'_f = D'(u) \neq D(u)$. Therefore, these databases must be distinct. Similarly, if neither z nor z' is $\perp$, there will remain some chain to $D(u)$ in the first set and to $D'(u) \neq D(u)$ in the second.

Now let us turn to the middle statement. Similar to before, the claim is easily handled if there is disjoint chains in D, D' not involving u, or if the reassignment is to $\perp$ in either case. In the other case, we perform the non-trivial reassignment $D_k[u_R \oplus D_h(u_L) \to w'']$ and $D'_k[u_R \oplus D'_h(u_L) \to w''']$, and have $D(u) \neq D'(u)$. In this case, there are exactly two input-output pairs, $(u_L, z) \in D_h$ and $(u_L \oplus w, z_f) \in D_f$ which do not participate in a chain for the former. Observe that these pairs would uniquely be completed to a chain under undoing the reassignment, i.e. $D_k[u_R \oplus z \to w] = D_k$. Under this unique completion, we would naturally obtain endpoint $D(u)$ for the chain beginning at u. Similarly, there is a unique completion of the modified D', whose endpoint would be $D'(u)$. These unique completions lead to distinct chains, and so the databases cannot be equal.

The final case follows similarly to the previous two: by the assumption that $D[u \to \perp] \neq D'[u \to \perp]$ there are distinct chains not involving u in D and D', which will not be modified by the described reassignment in D_f. It is worth observing that if $D(u)_L = D'(u)_L$, on the other hand, then the claim fails to hold. □

With these in place, orthogonality of the spaces spanned by these sets follows. In the following lemma, for some $S \subset \mathbf{D}$ we take $\mathcal{H}(S) \leq \mathcal{H}(\mathbf{D})$ as a subspace.

Corollary 4.19 (of Lemmas 4.16 and 4.18)**.** *For any $D \neq D' \in \mathbf{C}$ and $u \in [N]$ which is allowed by D and D' but supported by neither, we have*

$$\mathcal{H}(D \uparrow_h^u) \perp \mathcal{H}(D' \uparrow_h^u), \quad \mathcal{H}(D \uparrow_{h,k}^u) \perp \mathcal{H}(D' \uparrow_{h,k}^u), \quad \mathcal{H}(D \uparrow_{h,k,f}^u) \perp \mathcal{H}(D' \uparrow_{h,k,f}^u). \tag{109}$$

Further, for any $D \neq D' \in \mathbf{C}$ and $u \in \mathsf{Dom}(D) \cap \mathsf{Dom}(D')$, we have

$$\mathcal{H}(D \downarrow_h^u) \perp \mathcal{H}(D' \downarrow_h^u), \qquad \mathcal{H}(D \downarrow_k^u) \perp \mathcal{H}(D' \downarrow_k^u), \tag{110}$$

and $\mathcal{H}(D \downarrow_f^u) \perp \mathcal{H}(D' \downarrow_f^u)$ when $D[u \to \perp] \neq D'[u \to \perp]$.

We can now introduce the notation of assignment for database triples. For $u \notin \mathsf{Dom}(D)$ and $v \notin \mathsf{Im}(D)$ with (u, v) allowable, we define $D[u \to v]$ as the subset of $\mathbf{E}^3(D)$ which contains a chain from u to v. In other words,

$$D[u \to v] := \{D' \,:\, D'(u) = v, D'[u \to \perp] = D\}. \tag{111}$$

Corollary 4.20 (of Lemma 4.14)**.** *If $D \neq D'$ and both allow (u, v), then $D[u \to v] \cap D'[u \to v] = \emptyset$.*

These extension sets will allow us to define a modified version of masked Feistel compression, which is close in operator norm to the original but easier to analyze.

4.3 Canonical compression

In this section, we describe an alternative to the U_u compression operator, which we will label by G_u and refer to as *canonical compression*. This operator is constructed such that it maintains canonicity of the database, as long as the input u is allowed. In order to define it, we will require notation for the set of ouput values which we can insert without resulting in an internal collision. Let us fix some databases $D = (D_h, D_k, D_f)$ with $D(u) = \perp$, and write the sets

$$H(u, D) := \{z : u_R \oplus z \notin \mathsf{Dom}(D_k)\} \tag{112}$$

$$K(u, D) := \{w : u_L \oplus w \in \mathbf{L}_{u,\mathsf{Supp}(D)}\}, \tag{113}$$

where $\mathbf{L}_{u,\mathsf{Supp}(D)}$ is as defined in Section 4.2. When we use $\mathbf{L}_{u,\mathsf{Supp}(D')}$ for $D' \in \mathbf{E}(D)$, this is equivalent to $\mathbf{L}_{u,\mathsf{Supp}(D)}$. We can define compression operators which swap the uniform superposition over values in this set with the undefined label, rather than the full superposition. To formalize this, fix some $D_h \in \mathbf{D}|^{u_L}$ and some $D_k \in \mathbf{D}$, to define

$$|+_{u_L,H(u,D_k)}\rangle := \frac{1}{\sqrt{|H(u,D_k)|}} \sum_{z \in H(u,D_k)} |D_h[u_L \to z]\rangle \tag{114}$$

$$\begin{aligned} G^{(H)}_{u,D_h,D_k,\mathsf{H}} := (\mathbb{1} - |+_{u_L,H(u,D_k)}\rangle\langle+_{u_L,H(u,D_k)}| - |D_h\rangle\langle D_h| + \\ |+_{u_L,H(u,D_k)}\rangle\langle D_h| + |D_h\rangle\langle+_{u_L,H(u,D_k)}|)_{\mathsf{H}} \end{aligned} \tag{115}$$

$$G^{(H)}_{u,\mathsf{HK}} |D_k\rangle_{\mathsf{K}} := \bigoplus_{D_h \in \mathbf{D}|^{u_L}, D_k \in \mathbf{D}} |D_k\rangle\langle D_k|_{\mathsf{K}} \otimes G^{(H)}_{u,D_h,D_k,\mathsf{H}}. \tag{116}$$

We similarly define a canonical compression operator for the K register, first by fixing some $D = (D_h, D_k, D_f) \in \mathbf{E}(\mathbf{C})$ with $D_k(u_R \oplus D_h(u_L)) = \perp$; we use $m = u_R \oplus D_h(u_L)$ going forward, and defining

$$|+_{z,K(u,D)}\rangle := \frac{1}{\sqrt{|K(u,D)|}} \sum_{w \in K(u,D)} |D_h[z \to w]\rangle \tag{117}$$

$$G^{(K)}_{u,D,\mathsf{K}} := \begin{cases} (\mathbb{1} - |+_{m,K(u,D)}\rangle\langle+_{m,K(u,D)}| - |D_k\rangle\langle D_k| + \\ |+_{m,K(u,D)}\rangle\langle D_k| + |D_k\rangle\langle+_{m,K(u,D)}|)_{\mathsf{K}} & (\text{If } D_h(u_L) \neq \perp) \\ \mathbb{1}_{\mathsf{K}} & (\text{Otherwise}) \end{cases} \tag{118}$$

$$G^{(K)}_{u,\mathsf{HKF}} |D_k\rangle_{\mathsf{K}} := \mathbb{1} \oplus \bigoplus_{\substack{D \in \mathbf{E}(\mathbf{C}) \\ \text{s.t. } D_k(u_R \oplus D_h(u_L)) = \perp}} |D_h, D_f\rangle\langle D_h, D_f|_{\mathsf{HF}} \otimes G^{(K)}_{u,D_h,D_f,\mathsf{K}}. \tag{119}$$

Note that this is a unitary operator, as $K(u, D)$ does not depend on $D_k(m)$, which cannot participate in a chain as it is $\perp$. We take $G^{(F)}_{u,\mathsf{HKF}} = U^{(F)}_{u,\mathsf{HKF}}$, and define the full canonical compression operator

$$G_{u,\mathsf{HKF}} = G^{(H)}_{u,\mathsf{HK}} G^{(K)}_{u,\mathsf{HKF}} G^{(F)}_{u,\mathsf{HKF}}. \tag{120}$$

Remark 4.21. *The canonical compression operators are exponentially close to the standard compression operators after a bounded number of queries. In particular, we have*

$$\left\|G^{(H)}_{u,\mathsf{HK}} - U^{(H)}_{u,\mathsf{H}}\right\|_{\leq t} = O\left(\sqrt{\frac{t}{2^n}}\right), \qquad \left\|G^{(K)}_{u,\mathsf{HKF}} - U^{(K)}_{u,\mathsf{HK}}\right\|_{\leq t} = O\left(\sqrt{\frac{t^2}{2^n}}\right), \tag{121}$$

$$\left\|G^{(F)}_{u,\mathsf{HKF}} - U^{(F)}_{u,\mathsf{HKF}}\right\|_{\leq t} = 0, \qquad \left\|G_{u,\mathsf{HKF}} - U_{u,\mathsf{HKF}}\right\|_{\leq t} = O\left(\sqrt{\frac{t^2}{2^n}}\right). \tag{122}$$

Proof. The two equations on the top line follow from Lemma 2.8, and the bottom left by definition. Then, the bottom right follows from the other three equations by a simple sequence of triangle inequalities. □

Properties. We can work out the action of G_u and its components in some examples. These facts will be useful in proving soundness of our compressed permutation oracle.

Lemma 4.22. *Let $A \in \mathbf{A}_t$ be an allowable, size t injective database and let $u \in [N] \setminus \mathsf{Dom}(A)$ be an input which A allows. Then we have*

$$G_{u,\mathsf{HKF}} \left|+_{\mathbf{D}(A)}\right\rangle_{\mathsf{HKF}} = \frac{1}{\sqrt{|\mathbf{V}_{u,A}|}} \sum_{v \in \mathbf{V}_{u,A}} \left|+_{\mathbf{D}(A[u\to v])}\right\rangle_{\mathsf{HKF}}. \tag{123}$$

Proof. Let us fix some $D = (D_h, D_k, D_f)$ which is canonical with $\mathsf{Supp}(D) = A$, and compute the action on such a state. We have

$$G_{u,\mathsf{HKF}} \left|D\right\rangle_{\mathsf{HKF}} = \frac{1}{\sqrt{(2^n - t)\cdot|\mathbf{L}_{u,\mathsf{Supp}(D)}|\cdot 2^n}} \sum_{\substack{z_1\in\{0,1\}^n\setminus\{\mathsf{Dom}(D_k)\oplus u_R\},\\ w\in\{\mathbf{L}_{u,\mathsf{Supp}(D)}\oplus u_L\},\\ z_2\in\{0,1\}^n}} \left|D_h[u_L \to z_1]\right\rangle_{\mathsf{H}} \otimes \tag{124}$$

$$\left|D_k[u_R \oplus z_1 \to w]\right\rangle_{\mathsf{K}} \left|D_f[u_L \oplus w \to z_2]\right\rangle_{\mathsf{F}}. \tag{125}$$

This state can be seen to be the uniform superposition of canonical databases which support $A[u \to v]$ for some v, and agree with D[10], i.e. the set $D \uparrow^u_{h,k,f}$. Now observe that every canonical D' supporting $A[u \to v]$ for some v will agree with one, and only one, canonical D supporting I, from Lemma 4.14. Hence, when we extend the above calculation to the uniform superposition of all canonical D supporting I, we find

$$G_{u,\mathsf{HKF}} \left|+_{\mathbf{D}(A)}\right\rangle_{\mathsf{HKF}} = \frac{1}{\sqrt{|\mathbf{V}_{u,A}|}} \sum_{v \in \mathbf{V}_{u,A}} \left|+_{\mathbf{D}(A[u\to v])}\right\rangle_{\mathsf{HKF}}. \tag{126}$$

□

[10]Here we say database triples D, D' agree if for every $w \in \mathsf{Dom}(D_h) \cap \mathsf{Dom}(D'_h)$, we have $D_h(w) = D'_h(w)$, and similarly for D_f and D_k.

Lemma 4.23. *Fix some $A \in \mathbf{A}_t$ and $u \in [N] \setminus \mathsf{Dom}(A)$ which A allows, to write the state*

$$|\psi\rangle_{\mathsf{HKF}} \coloneqq \sum_{v \in \mathbf{V}_{u,A}} \alpha_v \, |+_{\mathbf{D}(A[u \to v])}\rangle_{\mathsf{HKF}} \,. \tag{127}$$

Then we have

$$\left\| |\psi\rangle_{\mathsf{HKF}} - G^{(H)}_{u,\mathsf{HK}} \, |\psi\rangle_{\mathsf{HKF}} \right\| = O\left(\frac{1}{\sqrt{2^n - t}} \right). \tag{128}$$

Proof. Let us consider some $D \in \mathbf{D}(A)$, and take $|A| = t$. We can write the component of $|\psi\rangle$ which is an extension of D as

$$\begin{aligned} |\psi_D\rangle_{\mathsf{HKF}} &= \Pi^{\mathcal{H}(\mathbf{E}^3(D))}_{\mathsf{HKF}} \, |\psi\rangle_{\mathsf{HKF}} && (129) \\ &= \frac{1}{\sqrt{|\mathbf{D}(A)| \cdot (2^n - t)}} \sum_{z \in \{0,1\}^n \setminus \{u_R \oplus \mathsf{Dom}(D_k)\}} |D_h[u_L \to z]\rangle_{\mathsf{H}} \otimes && (130) \\ &\quad \sum_{v \in \mathbf{V}_{u,A}} \alpha_v \, |D_k[u_R \oplus z \to u_L \oplus v_L], D_f[v_L \to u_R \oplus z \oplus v_R]\rangle_{\mathsf{KF}} \,. && (131) \end{aligned}$$

Further, recall from Lemma 4.14 we can write

$$|\psi\rangle_{\mathsf{HKF}} = \sum_{D \in \mathbf{D}(A)} |\psi_D\rangle_{\mathsf{HKF}} \,, \tag{132}$$

so by linearity it suffices to analyze each component separately. Let us define

$$\mathcal{H}_D = \bigoplus_{v \in \mathbf{V}_{u,A}} \mathcal{H}(D[u \to v] \downarrow^u_h), \tag{133}$$

where we have $\mathcal{H}_D \perp \mathcal{H}_{D'}$ for $D \neq D' \in \mathbf{D}(A)$ from Corollary 4.20 and Lemma 4.18. Observe that $|\psi_D\rangle \in \mathcal{H}_D$, and furthermore the image under $G^{(H)}_{u,\mathsf{HK}}$ will remain in this space. Therefore, we can write

$$\left\| |\psi\rangle_{\mathsf{HKF}} - G^{(H)}_{u,\mathsf{HK}} \, |\psi\rangle_{\mathsf{HKF}} \right\| = \sqrt{\sum_{D \in \mathbf{D}(A)} \left\| |\psi_D\rangle_{\mathsf{HKF}} - G^{(H)}_{u,\mathsf{HK}} \, |\psi_D\rangle_{\mathsf{HKF}} \right\|^2}. \tag{134}$$

Now, we have $\langle \psi_D | \psi_D \rangle = \frac{1}{|\mathbf{D}(A)|}$. Furthermore, observe that the KF registers, i.e. Equation (131), will have a single set of input-output registers which participate in the chain from u to v. If $D_k[u_L \to z']$ is re-assigned to $z' \neq z$, then this chain will not begin at u, rather at $u_L \| (u_R \oplus z \oplus z')$. Hence, there is a record of z external to $D_k(u_L)$. We can therefore apply Lemma 2.9, obtaining

$$\left\| |\psi_D\rangle_{\mathsf{HKF}} - G^{(H)}_{u,\mathsf{HK}} \, |\psi_D\rangle_{\mathsf{HKF}} \right\|^2 = O\left(\frac{1}{|\mathbf{D}(A)| \cdot (2^n - t)} \right), \tag{135}$$

$$\left\| |\psi\rangle_{\mathsf{HKF}} - G^{(H)}_{u,\mathsf{HK}} \, |\psi\rangle_{\mathsf{HKF}} \right\| = O\left(\frac{1}{\sqrt{2^n - t}} \right). \tag{136}$$

□

Lemma 4.24. *Fix some $A \in \mathbf{A}_t$ and $u \in [N] \setminus \mathsf{Dom}(A)$ which A allows, to write the state*

$$|\psi\rangle_{\mathsf{HKF}} \coloneqq \sum_{v \in \mathbf{V}_{u,A}} \alpha_v \, |+_{\mathbf{D}(A[u \to v])}\rangle_{\mathsf{HKF}} \,. \tag{137}$$

Then we have

$$\left\| |\psi\rangle_{\mathsf{HKF}} - G^{(K)}_{u,\mathsf{HKF}} \, |\psi\rangle_{\mathsf{HKF}} \right\| = O\left(\frac{1}{\sqrt{2^n - t}}\right). \tag{138}$$

Proof. This proof proceeds similarly to that of Lemma 4.23. □

Lemma 4.25. *Fix some $A \in \mathbf{A}_t$ and $u \in [N] \setminus \mathsf{Dom}(A)$ which A allows, to write the state*

$$|\psi\rangle_{\mathsf{HKF}} \coloneqq \sum_{v \in \mathbf{V}_{u,A}} \alpha_v \, |+_{\mathbf{D}(A[u \to v])}\rangle_{\mathsf{HKF}} \,. \tag{139}$$

We take $\alpha_v = 0$ for $v \notin \mathbf{V}_{u,A}$. Then we have

$$\left\| |\psi\rangle_{\mathsf{HKF}} - G^{(F)}_{u,\mathsf{F}} \, |\psi\rangle_{\mathsf{HKF}} \right\| = O\left(\sqrt{\frac{T}{2^n - t}}\right), \tag{140}$$

where $T = \sum_{l \in \{0,1\}^n} T_l$, and $T_l = \left| \sum_{r \in \{0,1\}^n} \alpha_{l \| r} \right|^2$.

Proof. We will use the same setup as the proof for Lemma 4.23. Let us consider some $D \in \mathbf{D}(A)$, and take $|A| = t$. We can write the component of $|\psi\rangle$ which is an extension of this state as

$$|\psi_D\rangle_{\mathsf{HKF}} = \Pi^{\mathcal{H}(\mathbf{E}^3(D))}_{\mathsf{HKF}} \, |\psi\rangle_{\mathsf{HKF}} \tag{141}$$

$$= \frac{1}{\sqrt{|\mathbf{D}(A)| \cdot (2^n - t)}} \sum_{z \in \{0,1\}^n \setminus \{u_R \oplus \mathsf{Dom}(D_k)\}} |D_h[u_L \to z]\rangle_{\mathsf{H}} \otimes \tag{142}$$

$$\sum_{v \in \mathbf{V}_{u,A}} \alpha_v \, |D_k[u_R \oplus z \to u_L \oplus v_L], D_f[v_L \to u_R \oplus z \oplus v_R]\rangle_{\mathsf{KF}} \,. \tag{143}$$

Further, recall from Lemma 4.14 that these states are orthogonal and span the relevant space, so we can write

$$|\psi\rangle_{\mathsf{HKF}} = \sum_{D \in \mathbf{D}(A)} |\psi_D\rangle_{\mathsf{HKF}} \,, \tag{144}$$

so by linearity it suffices to analyze each component separately. Let us define

$$\mathcal{H}_D = \bigcup_{v \in \mathbf{V}_{u,A}} \mathcal{H}(D[u \to v]), \tag{145}$$

where we have $\mathcal{H}_D \perp \mathcal{H}_{D'}$ for $D \neq D' \in \mathbf{D}(A)$ from Corollary 4.20 and Lemma 4.18. Observe that $|\psi_D\rangle \in \mathcal{H}_D$, and furthermore the image under $G^{(F)}_{u,\mathsf{HKF}}$ will remain in this space. Therefore, we can write

$$\left\| |\psi\rangle_{\mathsf{HKF}} - G^{(F)}_{u,\mathsf{HKF}} |\psi\rangle_{\mathsf{HKF}} \right\| = \sqrt{\sum_{D \in \mathbf{D}(A)} \left\| |\psi_D\rangle_{\mathsf{HKF}} - G^{(F)}_{u,\mathsf{HKF}} |\psi_D\rangle_{\mathsf{HKF}} \right\|^2}. \tag{146}$$

We can analyze the action of $G^{(F)}_{u,\mathsf{HKF}}$ on this state as

$$G^{(F)}_{u,\mathsf{HKF}} |\psi_D\rangle_{\mathsf{HKF}} = \frac{1}{\sqrt{|\mathbf{D}(A)| \cdot (2^n - t)}} \sum_{\substack{z \in \{0,1\}^n \setminus \{u_R \oplus \mathsf{Dom}(D_k)\} \\ l \in \mathbf{L}_{u,A}}} |D_h[u_L \to z], D_k[u_R \oplus z \to u_L \oplus l]\rangle_{\mathsf{HK}} \otimes$$
$$\mathsf{fC}_{l,\mathsf{F}} \sum_{r \in \{0,1\}^n} \alpha_{l\|r} |D_f[l \to u_R \oplus z \oplus r]\rangle_{\mathsf{F}}, \tag{147}$$

from which it is clear to see that

$$\left\| |\psi_D\rangle - G^{(F)}_u |\psi_D\rangle \right\| = \frac{1}{\sqrt{|\mathbf{D}(A)| \cdot (2^n - t)}} \cdot \tag{148}$$
$$\sqrt{\sum_{\substack{z \in \{0,1\}^n \setminus \{u_R \oplus \mathsf{Dom}(D_k)\} \\ l \in \mathbf{L}_{u,A}}} \left\| \sum_{r \in \{0,1\}^n} (\mathbb{1}_{\mathsf{F}} - \mathsf{fC}_{l,\mathsf{F}}) \alpha_{l\|r} |D_f[l \to u_R \oplus z \oplus r]\rangle_{\mathsf{F}} \right\|^2}$$
$$= \frac{1}{\sqrt{|\mathbf{D}(A)| \cdot (2^n - t)}} \cdot O\left(\sqrt{\sum_{l \in \mathbf{L}_{u,A}} \left| \sum_{r \in \{0,1\}^n} \alpha_{l\|r} \right|^2} \right) \tag{149}$$
$$= O\left(\sqrt{\frac{T}{|\mathbf{D}(A)| \cdot (2^n - t)}} \right), \tag{150}$$

which combined with Equation (146) shows the claim. □

4.4 Cromulent twirling

We now turn to working out properties of the twirling permutations. A number of statistical properties are required of the distribution for our analysis to hold. First, define the left and right action of permutations on databases as

$$\pi * I := \{(\pi(x), y) : (x, y) \in I\} \tag{151}$$
$$I * \omega := \{(x, \omega^{-1}(y)) : (x, y) \in I\}. \tag{152}$$

This action corresponds to the internal database which is supported; if the full masked Feistel ensemble supports I, then the internal three rounds of Feistel will support $\pi * I * \omega$. In this way, we will require certain properties of π and ω to ensure that this internal database is allowable, as well as to understand the action of a query.

Definition 4.26. *We say π, ω resolve $I \in \mathbf{I}_t$ if $\pi * I * \omega \in \mathbf{A}_t$. We let $\mathbf{R}_I$ be the set of permutation pairs (π, ω) which resolve I. For some fixed $A \in \mathbf{A}_t$, we let $\mathbf{R}_{I \to A}$ be the set of permutation pairs (π, ω) such that $\pi * I * \omega = A$. For some fixed $x \in [N]$, we let $\mathbf{R}_{x,I}$ be the set of π, ω which resolve I and for which $\pi * I * \omega \heartsuit \pi(x)$, meaning the internal database allows $\pi(x)$.*

Let $\mathcal{D}$ be a distribution on $\mathbf{S}_N^{\times 2}$, and $p_{\mathcal{D}}$ be the corresponding probability function. We use the notation $\mathcal{D}_I$ to denote the induced distribution of $\mathcal{D}$ conditioned on resolving I. In other words,

$$p_{\mathcal{D}_I}(\pi, \omega) = p_{\mathcal{D}}(\pi, \omega \,|\, (\pi, \omega) \in \mathbf{R}_I). \tag{153}$$

Similarly, $\mathcal{D}_{I \to A}$ for $A \in \mathbf{A}$ is the distribution conditioned on satisfying $\pi * I * \omega = A$, and $\mathcal{D}_{x,I}$ is the distribution conditioned on the inner distribution allowing $\pi(x)$. Formally,

$$p_{\mathcal{D}_{I \to A}}(\pi, \omega) = p_{\mathcal{D}}(\pi, \omega \,|\, (\pi, \omega) \in \mathbf{R}_{I \to A}), \tag{154}$$
$$p_{\mathcal{D}_{x,I}}(\pi, \omega) = p_{\mathcal{D}}(\pi, \omega \,|\, (\pi, \omega) \in \mathbf{R}_{x,I}). \tag{155}$$

We say that such a distribution is *cromulent* if it satisfies the following three properties. These are statistical properties about the distribution which will suffice to show that the twirl hides the internal three Feistel rounds.

Definition 4.27. *Distribution $\mathcal{D}$ is* cromulent *if, for all $I \in \mathbf{I}_t$, $x \neq x' \in [N] \setminus \mathsf{Dom}(I)$, $l \in \{0,1\}^n$, and $y \neq y' \in [N] \setminus \mathsf{Im}(I)$, we have*

$$\Pr_{(\pi,\omega) \sim \mathcal{D}_I}[(\pi, \omega) \in \mathbf{R}_{x,I}] = 1 - O(t^2/2^n), \tag{156}$$

$$\Pr_{(\pi,\omega) \sim \mathcal{D}_{x,I}}[\omega^{-1}(y)_L = l] = \frac{1}{2^n}, \tag{157}$$

$$\Pr_{(\pi,\omega) \sim \mathcal{D}_{x,I}}[\omega^{-1}(y)_L = \omega^{-1}(y')_L = l] = \frac{1}{2^{2n}} \pm O\left(\frac{t^2}{2^{3n}}\right), \tag{158}$$

and for all $(\pi, \omega) \in \mathbf{R}_{I[x \to y]}$ and $A = \pi * I * \omega$ we have

$$\sqrt{\frac{p_{\mathcal{D}_{I[x \to y]}}(\pi, \omega)}{N - t}} - \sqrt{\frac{p_{\mathcal{D}_{x,I}}(\pi, \omega)}{|\mathbf{V}_{\pi(x),A}|}} = \pm O\left(\sqrt{\frac{t^2}{2^n}} \cdot \sqrt{\frac{p_{\mathcal{D}_{x,I}}(\pi, \omega)}{|\mathbf{V}_{\pi(x),A}|}}\right) \tag{159}$$

Further, all equations hold in the inverse direction, if we replace (π, ω) with (ω^{-1}, π^{-1}), I with I^{-1}, x with y and x' with y'.

The first condition captures the intuition that each new query should, with high probability, lead to a non-colliding left wire at the start of the internal Feistel. The second and third point capture properties necessary for compression to work as desired: compression should only remove a chain when the algorithm uncomputes an input-output pair. For plain three-round Feistel, uncomputing only the last n bits of an input would suffice to remove the last element of its chain. These conditions capture that the last n bits are sufficiently scrambled by ω to render this attack, and anything like it, impossible.

The last equation is a technical requirement related to normalization, and is due to small deviations in the number of allowable databases.

Note that the final remark concerning inverses is implied by the previous equations if $\mathcal{D}$ is invariant under the flip operation, $(\pi, \omega) \mapsto (\omega^{-1}, \pi^{-1})$. In all of the settings we consider, this will be the case. It therefore suffices to show the properties for the listed cases only. The following lemmas establish that the uniform distribution $p_{\mathcal{D}}(\pi, \omega) = 1/(N!)^2$ is cromulent. We will in turn show that, so long as the twirling permutations are cromulent, the masked Feistel ensemble is indistinguishable from our compressed permutation oracle.

Lemma 4.28. *Fix some $I \in \mathbf{I}_t$ and $x \in [N] \setminus \mathsf{Dom}(I)$. Then we have*

$$\Pr_{(\pi,\omega)\sim\mathbf{R}_I}[\pi * I * \omega \text{ allows } \pi(x)] = 1 - O(t^2/2^n). \tag{160}$$

Proof. Consider sampling all values $\pi(x)$ for $x \in \mathsf{Dom}(I)$, and all values of ω such that $\pi * I * \omega \in \mathbf{A}_t$. The remaining values of π may be sampled uniformly at random, and so $\pi(x)$ comes from a set of size $2^{2n} - t$. There are $O(t^2 2^n)$ values which are not allowed (from the $O(t^2)$ left substrings times 2^n right substrings), so we have

$$\Pr_{(\pi,\omega)\sim\mathbf{R}_I}[\pi * I * \omega \text{ allows } \pi(x)] = \frac{2^{2n} - O(t^2 2^n)}{2^{2n} - t} \tag{161}$$

$$= 1 - O\left(\frac{t^2}{2^n}\right). \tag{162}$$

□

Lemma 4.29. *Fix some $I \in \mathbf{I}_t$, some $x \in [N] \setminus \mathsf{Dom}(I)$, and $l \in \{0,1\}^n$. Then, for all $y \neq y' \in [N] \setminus \mathsf{Im}(I)$, we have*

$$\Pr_{\pi,\omega\sim\mathbf{R}_{x,I}}[\omega^{-1}(y)_L = l] = \frac{1}{2^n}, \tag{163}$$

$$\Pr_{\pi,\omega\sim\mathbf{R}_{x,I}}[\omega^{-1}(y)_L = \omega^{-1}(y')_L = l] = \frac{1}{2^{2n}} \pm O\left(\frac{t^2}{2^{3n}}\right). \tag{164}$$

Proof. For the first equation, observe that allowability in this sense is translation invariant on the internal left wire: if $\pi, \omega \in \mathbf{R}_I$, then we can construct $\omega'(v) = \omega(v \oplus (s\|0^n))$ and $\pi'(x) = \pi(x) \oplus (s\|0^n)$. Then, for any $s \in \{0,1\}^n$ the newly constructed permutations will also be in $\mathbf{R}_I$. Therefore, the distribution induced by $\omega^{-1}(y)_L$ is uniform.

For the second equation, let us analyze the probability that that $\omega^{-1}(y)_L = \omega^{-1}(y')_L$. Consider fixing all of the other input-output points of π and ω such that the internal database is allowable. Then, both $\omega^{-1}(y)_L$ and $\omega^{-1}(y')_L$ are chosen uniformly at random from a universe of size $2^{2n} - t$. Therefore, there will be a collision on the left wire with probability $2^{-n} \pm O(t 2^{-2n})$. Combining this with the previous translation invariance,

observe that we can randomly shift the constructed π, ω as described above after this procedure. This gives a 2^{-n} probability that $\omega(y) = l$, independent of the prior event, leaving us with the result

$$\Pr_{\pi,\omega\sim\mathbf{R}_{x,I}}[\omega^{-1}(y)_L = \omega^{-1}(y')_L = l] = \frac{1}{2^{2n}} \pm O\left(\frac{t}{2^{3n}}\right). \tag{165}$$

This is tighter than the bound stated, showing the claim. □

Lemma 4.30. *For any $I \in \mathbf{I}_t$, $x \in [N] \setminus \mathsf{Dom}(I)$, $y \in [N] \setminus \mathsf{Im}(I)$, and $\pi, \omega \in \mathbf{R}_{I[x\to y]}$ with $A = \pi * I * \omega$, we have*

$$\frac{1}{\sqrt{|\mathbf{R}_{I[x\to y]}| \cdot (N-t)}} - \frac{1}{\sqrt{|\mathbf{R}_{x,I}| \cdot |\mathbf{V}_{\pi(x),A}|}} = \pm O\left(\frac{t^2}{2^n\sqrt{|\mathbf{R}_{I[x\to y]}| \cdot (N-t)}}\right). \tag{166}$$

Proof. Observe that we have

$$\frac{|\mathbf{R}_{I[x\to y]}|}{|\mathbf{R}_{x,I}|} = 1 \pm O\left(\frac{t^2}{2^n}\right), \qquad \frac{N-t}{|\mathbf{V}_{\pi(x),\mathbf{A}}|} = 1 \pm O\left(\frac{t^2}{2^n}\right), \tag{167}$$

the first by an argument analogous to Lemma 4.28 and the second by Lemma 4.10. If we let $a = |\mathbf{R}_{I[x\to y]}| \cdot (N-t)$, then we can write

$$\frac{1}{\sqrt{|\mathbf{R}_{x,I}| \cdot |\mathbf{V}_{\pi(x),A}|}} = \frac{1}{\sqrt{a(1 \pm O(t^2/2^n)}} \tag{168}$$

$$= \frac{1}{\sqrt{a}} \pm O\left(\frac{t^2}{2^n \cdot \sqrt{a}}\right), \tag{169}$$

where the last line follows from the taylor expansion of $\frac{1}{\sqrt{a+x}}$ about $x = 0$. □

Going forward, we will say that $P = (\pi, \omega, D)$ supports I if D supports $\pi * I * \omega$.

5 Soundness

In this section, we will establish that our compressed permutation oracle is indistinguishable from a random permutation oracle. We will do this by showing that an oracle for the masked Feistel construction with uniform twirl (which is therefore a uniform permutation oracle) is indistinguishable from our compressed permutation oracle.

In fact, all of the machinery we develop prior to the main theorem will assume only that the masked Feistel has a cromulent twirling distributions. We take $\mathcal{D}$ to be the distribution on $\mathbf{S}_N^{\times 2}$ from which the twirling permutations are drawn. The only part of this section where we assume that $\mathcal{D}$ is uniform is in the main Theorem 5.19. In Section 6, we will use a different cromulent distribution to establish the quantum security of the Feistel construction, re-using much of the machinery developed here.

5.1 Valid, elegant, and sophisticated subspaces

We will begin with the *sophisticated* subspace of $\mathcal{H}_{\mathsf{P}}$, which will capture purifications corresponding to compressed permutation databases. Towards this, we first define the notation $|\mathbf{P}(I)\rangle$ for $I \in \mathbf{I}$, where dependence on $\mathcal{D}$ is suppressed, as

$$|\mathbf{P}(I)\rangle := \sum_{(\pi,\omega)\in\mathbf{R}_I} \sqrt{p_{\mathcal{D}_I}(\pi,\omega)}\,|\pi,\omega\rangle_{\Pi\Omega} \otimes |+_{\mathbf{D}(\pi * I * \omega)}\rangle_{\mathsf{HKF}}\,. \tag{170}$$

Observe that these states satisfy $\langle \mathbf{P}(I)|\mathbf{P}(I')\rangle = \delta_{I,I'}$, as for any $I \neq I'$ the two states have distinct supports. The span of these states form a subspace, which we call the *sophisticated* subspace, and refer to by

$$\mathcal{H}^{\mathsf{soph}} := \mathsf{span}\{|\mathbf{P}(I)\rangle \,:\, I \in \mathbf{I}_{\leq 2^n}\}, \qquad \Pi^{\mathsf{soph}} := \sum_{I\in\mathbf{I}_{\leq 2^n}} |\mathbf{P}(I)\rangle\langle\mathbf{P}(I)|. \tag{171}$$

Then, we can takes these states to define our intertwining isometry.

Definition 5.1. *We define the* intertwiner $\mathcal{I} : \mathcal{H}(\mathbf{P}) \to \mathcal{H}(\mathbf{I}) \oplus \mathcal{H}(\mathbf{P})$ *as the isometry*

$$\mathcal{I} = \mathbb{1} + \sum_{I\in\mathbf{I}_{\leq 2^n}} |I\rangle\langle\mathbf{P}(I)| - |\mathbf{P}(I)\rangle\langle\mathbf{P}(I)|. \tag{172}$$

In other words, $\mathcal{I}$ is an isometry which maps $|\mathbf{P}(I)\rangle$ to $|I\rangle$, and is the identity on all states in $\mathcal{H}^{\mathsf{soph},\perp}$.

As a technical note, we will take the compressed permutation oracle purifying register to be $\mathcal{H}(\mathbf{I}) \oplus \mathcal{H}(\mathbf{P})$ in this section. We will consider the permutation compression operator pC, the purified query operation P, and flip operation F, to act as the identity on the $\mathcal{H}(\mathbf{P})$ component. We will still label this entire Hilbert space as register I, with the expectation that only a small portion of the state will lie in the $\mathcal{H}(\mathbf{P})$ component. Observe that all of the aforementioned operators (and therefore cpO) preserve the two subspaces, in that they can be written as a direct sum $O_1 \oplus O_2$ for operator O_1 on $\mathcal{H}(\mathbf{I})$ and O_2 on $\mathcal{H}(\mathbf{P})$. Therefore, the compressed oracle experiment with this extended purification is perfectly indistinguishable from the standard compressed oracle experiment.[11]

Let us observe the simple fact that $\mathcal{I}$ is an intertwiner between the flip operators F and mfF on the sophisticated subspace, as well as between the purified oracles P and mfP. Observe that the the sophisticated subspace is maintained by the masked Feistel flip operator mfF and masked Feistel purified query operator mfP.

Remark 5.2. *We have*

$$\mathsf{F}_{\mathsf{I}} \cdot \mathcal{I}_{\mathsf{P}} \cdot \Pi_{\mathsf{P}}^{\mathsf{soph}} = \mathcal{I}_{\mathsf{P}} \cdot \mathsf{mfF}_{\mathsf{P}} \cdot \Pi_{\mathsf{P}}^{\mathsf{soph}}, \qquad \mathsf{P}_{\mathsf{XYI}} \cdot \mathcal{I}_{\mathsf{P}} \cdot \Pi_{\mathsf{P}}^{\mathsf{soph}} = \mathcal{I}_{\mathsf{P}} \cdot \mathsf{mfP}_{\mathsf{XYP}} \cdot \Pi_{\mathsf{P}}^{\mathsf{soph}}. \tag{173}$$

[11] The reason we do this instead of constructing the isometry from $\mathcal{H}(\mathbf{I})$ to $\mathcal{H}(\mathbf{P})$ is that it is unclear how to define validity for the compressed permutation oracle. In the isometric subspace, we will see that it is easy to define, which in turn will be used to conclude validity of our construction.

Proof. For the first equation, observe that if $P = (\pi, \omega, D_h, D_k, D_f) \in \mathbf{P}$ supports injective database I, then the flipped $(\omega^{-1}, \pi^{-1}, D_f, D_k, D_h)$ supports I^{-1}. This holds even outside the sophisticated subspace, so it holds on it as well.

For the second equation, recall that $\mathcal{I}_{\mathsf{P}}$ sends $|\mathbf{P}(I)\rangle$ to $|I\rangle$, which determines it's action entirely on the sophisticated subspace. The support of every standard-basis state in the sum defining $|\mathbf{P}(I)\rangle$ is I, and so mfP answers queries identically to P. □

We will define the *valid* subspace of $\mathcal{H}_{\mathsf{P}} = \mathcal{H}(\mathbf{P})$ as the span of states where the HKF registers lie purely in the valid subspace as defined in Section 2.3. We will use the notation $\mathcal{H}^{\mathsf{val}}_{\mathsf{P}}$ to denote this subspace, and Π^{val} to denote the projector onto it, which can be defined as

$$\Pi^{\mathsf{val}}_{\mathsf{P}} \coloneqq \Xi_{\mathsf{H}}\Xi_{\mathsf{K}}\Xi_{\mathsf{F}}. \tag{174}$$

From Lemma 2.7, queries to the internal functions preserves this subspace. It follows that cmfO preserves it as well, meaning cmfO can be written of the form $\mathsf{cmfO} = O_1 \oplus O_2$ where $O_1 : \mathcal{H}^{\mathsf{val}}_{\mathsf{P}} \to \mathcal{H}^{\mathsf{val}}_{\mathsf{P}}$ and $O_2 : \mathcal{H}^{\mathsf{val},\perp}_{\mathsf{P}} \to \mathcal{H}^{\mathsf{val},\perp}_{\mathsf{P}}$ are unitary operators.

A related subspace is the *query valid* subspace $\mathcal{H}^{\mathsf{qval}}_{\mathsf{XP}}$, this time a subspace of $\mathcal{H}_{\mathsf{X}} \otimes \mathcal{H}_{\mathsf{P}}$. The projector onto this subspace is denoted Π^{qval}, and the subspace captures the property that, after decompression, the query point will have a well-defined output. To define this, we will first extend the definition of the decompressed projector Γ_x to $\mathcal{H}(\mathbf{P})$, writing

$$\Gamma_{x,\mathsf{P}} |P\rangle_{\mathsf{P}} \coloneqq \begin{cases} 0 & (\text{If } \mathsf{Supp}(P)(x) = \perp) \\ |P\rangle_{\mathsf{P}} & (\text{Otherwise}) \end{cases} \tag{175}$$

$$\Pi^{\in \mathsf{db}}_{\mathsf{XP}} |x\rangle_{\mathsf{X}} \coloneqq |x\rangle_{\mathsf{X}} \otimes \Gamma_{x,\mathsf{P}}. \tag{176}$$

We also need the notation $\mathsf{ctrl\text{-}mfF}_{\mathsf{XP}}$ denote the $\mathsf{mfF}_{\mathsf{P}}$ operator controlled on the b bit of the X register, which can be written

$$\mathsf{ctrl\text{-}mfF}_{\mathsf{XP}} |b, x\rangle_{\mathsf{X}} = |b, x\rangle_{\mathsf{X}} \otimes \begin{cases} \mathbb{1}_{\mathsf{P}} & (\text{If } b = 0) \\ \mathsf{mfF}_{\mathsf{P}}. & (\text{If } b = 1) \end{cases} \tag{177}$$

With this in hand, we can define the query valid projector as

$$\Pi^{\mathsf{qval}}_{\mathsf{XP}} \coloneqq (\mathsf{ctrl\text{-}mfF} \cdot \mathsf{mfC} \cdot \Pi^{\in \mathsf{db}} \cdot \mathsf{mfC}^{\dagger} \cdot \mathsf{ctrl\text{-}mfF}^{\dagger})_{\mathsf{XP}} \tag{178}$$

Observe that $\mathcal{H}^{\mathsf{qval}} \leq \mathcal{H}_{\mathsf{X}} \otimes \mathcal{H}^{\mathsf{val}}$, or in words, if the purification register is valid then decompressing on any query point will result in a well defined answer. This has the simple corollary that every state reached by an algorithm lies in the query valid subspace.

Corollary 5.3. *Let $\mathcal{A}$ be a quantum algorithm making q queries with intermediate operations $A_0, \ldots, A_q$, and define*

$$|\psi^{(\mathsf{cmfO})}\rangle_{\mathsf{AXYP}} \coloneqq A_{q,\mathsf{AXY}}\mathsf{cmfO}_{\mathsf{XYP}} \ldots A_{1,\mathsf{AXY}}\mathsf{cmfO}_{\mathsf{XYP}} A_{0,\mathsf{AXY}} |0\rangle_{\mathsf{AXY}} |\mathbf{P}(\perp)\rangle_{\mathsf{P}}. \tag{179}$$

Then we have $|\psi^{(\mathsf{cmfO})}\rangle_{\mathsf{AXYP}} \in \Pi^{\mathsf{val}}_{\mathsf{P}}$, and thus $|\psi^{(\mathsf{cmfO})}\rangle_{\mathsf{AXYP}} \in \Pi^{\mathsf{qval}}_{\mathsf{XP}}$.

Note that the projectors onto valid or query valid do not commute with projectors onto the sophisticated subspace. We would like to focus our attention on states which are both query valid and sophisticated, as these are the ones which we will have tools to analyze. Unfortunately, given the preceding fact, it is difficult to understand this intersection space.

Instead, we will work with a subspace of the sophisticated space which we call the *elegant* subspace, denoted $\mathcal{H}^{\mathsf{ele}}$. This space will act as a proxy for the intersection of sophisticated and query valid. It is a subspace of $\mathcal{H}_{\mathsf{X}} \otimes \mathcal{H}^{\mathsf{soph}}$. We write Π^{ele} for the projector onto $\mathcal{H}^{\mathsf{ele}}$, and can define for any $b \in \{0,1\}, x \in [N]$

$$|+_{x,\mathbf{P}(I)}\rangle := \frac{1}{\sqrt{N-|I|}} \sum_{y \in [N] \setminus \mathsf{Im}(I)} |\mathbf{P}(I[x \to y])\rangle\,, \tag{180}$$

$$\Pi^{\mathsf{ele}}_{\mathsf{XP}} |b,x\rangle_{\mathsf{X}} := |b,x\rangle_{\mathsf{X}} \otimes \sum_{I \in \mathbf{I}^x : |I| < 2^m} \left(\Pi^{\mathsf{soph}} - |+_{x,\mathbf{P}(I)}\rangle\langle+_{x,\mathbf{P}(I)}|\right)_{\mathsf{P}} \tag{181}$$

To understand this definition, consider a state of the form

$$|\psi\rangle_{\mathsf{XP}} = |0,x\rangle_{\mathsf{X}} \left(\alpha_\perp |\mathbf{P}(I)\rangle + \sum_{y \in [N] \setminus \mathsf{Im}(I)} \alpha_y |\mathbf{P}(I[x \to y])\rangle\right)_{\mathsf{P}}, \tag{182}$$

for some $x \in [N]$ and $I \in \mathbf{I}^{|x}$ of size less then 2^n. This state will lie in the elegant subspace if and only if $\sum_{y \in [N] \setminus \mathsf{Im}(I)} \alpha_y = 0$. If we replace the $\mathbf{P}(I)$ and $\mathbf{P}(I[x \to y])$ with I and $I[x \to y]$ above (considering a purification register $\mathcal{H}(\mathbf{I})$), then this is precisely the condition that we will obtain a valid output after applying decompression $\mathsf{pC}^\dagger_{\mathsf{XP}}$. In this way, the elegant subspace captures states which, after applying the intertwiner $\mathcal{I}$, will have a query valid answer on decompression $\mathsf{pC}^\dagger_{\mathsf{XP}}$.

Note also that we do not consider the value of b in the X register; this subspace will be most relevant when the decompression operator is applied during the compressed masked Feistel. In this scenario, the database will have already been flipped in the case $b = 1$. As a result, this subspace is defined assuming a forwards query. It will also be worth defining a version of the elegant subspace conjugated by the controlled flip, called *flip elegant*. We will define it by its projector. Intuitively, this projector capture the same structure as the elegant projector; the state before applying $\mathsf{ctrl\text{-}mfF}$ will be flip elegant (i.e. the input state to cmfO), and the state after applying the $\mathsf{ctrl\text{-}mfF}$ will be elegant (i.e. the input state to decompression $\mathsf{mfC}^\dagger$). We write

$$\Pi^{\mathsf{fele}}_{\mathsf{XP}} := (\mathsf{ctrl\text{-}mfF} \cdot \Pi^{\mathsf{ele}} \cdot \mathsf{ctrl\text{-}mfF}^\dagger)_{\mathsf{XP}}. \tag{183}$$

While this subspace is not exactly the same as the intersection of sophisticated and query valid, it can act as a proxy. We can formalize this using the following lemma.

Lemma 5.4. *Let $|\psi\rangle_{\mathsf{XP}} \in \mathcal{H}_{\mathsf{X}} \otimes \mathcal{H}^{\mathsf{soph}}$ be a sophisticated quantum state satisfying $|\psi\rangle_{\mathsf{XP}} \in \Pi^{\mathsf{soph}}_{\mathsf{P}}$, and supported on purifications of size at most t. If we define*

$$|\phi\rangle_{\mathsf{XP}} := \Pi^{\mathsf{qval}}_{\mathsf{XP}} |\psi\rangle_{\mathsf{XP}}\,, \qquad |\mu\rangle_{\mathsf{XP}} := \Pi^{\mathsf{fele}}_{\mathsf{XP}} |\psi\rangle_{\mathsf{XP}}\,, \tag{184}$$

then we have

$$\| |\phi\rangle_{\mathsf{XP}} - |\mu\rangle_{\mathsf{XP}} \| = O\left(\sqrt{\frac{t^2}{2^n}}\right). \tag{185}$$

We defer the proof of this lemma until we have more technical machinery.

5.2 A twirling lemma

The following lemma will be key in analyzing masked Feistel. It allows us to formalize the intuition that the twirling permutations hide the internal Feistel structure. For intuition, consider the lemma statement without the condition $\sum_{y\in[N]} \alpha_y = 0$. Then, we could set $\alpha_y = 2^{-n}$, and we would obtain $\mathbb{E}_{\omega\sim D}[T] = 2^n$. What we show is that, for any orthogonal vector, this quantity is comparatively small. This, in turn, will allow us to show that the masked Feistel compression operator is indeed close to the identity on the subspace where the corresponding compressed permutation oracle acts as the identity.

Lemma 5.5. *Fix some vector $|\alpha\rangle = \sum_{y\in[N]} \alpha_y |y\rangle$ of unit norm, and such that $\sum_{y\in[N]} \alpha_y = 0$. Let $\mathcal{D}$ be a distribution such that for all $y \neq y' \in [N]$ with $\alpha_y \neq 0$ and $\alpha_{y'} \neq 0$, and $l \in \{0,1\}^n$,*

$$\Pr_{(\pi,\omega)\sim\mathcal{D}}[\omega^{-1}(y)_L = l] = 2^{-n}, \quad \Pr_{(\pi,\omega)\sim\mathcal{D}}[\omega^{-1}(y)_L = \omega^{-1}(y')_L = l] = 2^{-2n} \pm O(t^2 2^{-3n}). \tag{186}$$

Then we have

$$\mathbb{E}_{(\pi,\omega)\sim\mathcal{D}}[T] = \mathbb{E}_{(\pi,\omega)\sim\mathcal{D}}\left[\sum_{l\in\{0,1\}^n} \left|\sum_{r\in\{0,1\}^n} \alpha_{\omega(l\|r)}\right|^2\right] = O(t^2). \tag{187}$$

Proof. Let us compute the expectation of T_l, defined as

$$T_l = \left|\sum_{r\in\{0,1\}^n} \alpha_{\omega(l\|r)}\right|^2. \tag{188}$$

We can write I_y to be the indicator variable that is 1 if $\omega^{-1}(y)_L = l$. If $\alpha_y = 0$, then we take I_y to be 1 with probability 2^{-n} independent of the other $I_{y'}$, which we can do without

changing the following equation. Then we can write

$$\mathbb{E}[T_l] = \mathbb{E}\left[\left|\sum_{y\in[N]} I_y\alpha_y\right|^2\right] \tag{189}$$

$$= \mathbb{E}\left[\sum_{y\in[N]} I_y^2\left|\alpha_y\right|^2 + \sum_{y\neq y'\in[N]} I_yI_{y'}\left|\alpha_y\alpha_{y'}\right|\right] \tag{190}$$

$$= \sum_{y\in[N]} \mathbb{E}[I_y]\alpha_y\alpha_y^* + \sum_{y\neq y'\in[N]} \mathbb{E}[I_yI_{y'}]\alpha_y\alpha_{y'}^* \tag{191}$$

$$= 2^{-n}\underbrace{\sum_{y\in[N]} \alpha_y\alpha_y^*}_{S_1} + 2^{-2n}\underbrace{\sum_{y\neq y'\in[N]} \alpha_y\alpha_{y'}^*}_{S_2} + \underbrace{\sum_{y\neq y'\in[N]} \pm O(t^2 2^{-3n})\alpha_y\alpha_{y'}^*}_{S_3}, \tag{192}$$

where the last line follows from cromulence of $\mathcal{D}$. It is clear that $S_1 = 1$. We can compute S_2 as

$$S_2 = \sum_{y\neq y'\in[N]} \alpha_y\alpha_{y'}^* \tag{193}$$

$$= \sum_{y\in[N]} \alpha_y\left(\sum_{y'\in[N]\setminus\{y\}} \alpha_{y'}^*\right) \tag{194}$$

$$= \sum_{y\in[N]} -\alpha_y\alpha_y^* \qquad (\text{By } \sum_{y\in[N]} \alpha_y = 0) \tag{195}$$

$$= -1. \tag{196}$$

It remains to compute S_3. Note that we cannot factor out the $O(t^2 2^{-3n})$, as the constant pre-factor and sign may be different for different values of y and y'. Instead we apply triangle inequality, obtaining

$$S_3 = \sum_{y\neq y'\in[N]} \pm O(t^2 2^{-3n})\alpha_y\alpha_{y'}^* \tag{197}$$

$$\leq O(t^2 2^{-3n}) \sum_{y\neq y'\in[N]} |\alpha_y||\alpha_{y'}^*| \qquad (\text{Triangle inequality}) \tag{198}$$

$$\leq O(t^2 2^{-3n})\left(\sum_{y\in[N]} |\alpha_y|\right)^2 \tag{199}$$

$$\leq O(t^2 2^{-n}). \tag{200}$$

We therefore have $0 \leq \mathbb{E}[T_l] = O(t^2 2^{-n})$ for all $l \in \{0,1\}^n$, and so the claim follows by linearity of expectation, and the identity $T = \sum_{l\in\{0,1\}^n} T_l$. □

5.3 Sanitized and ideal masked Feistel

We will here define two modification to the masked Feistel oracle, called the *ideal* masked Feistel denoted $\overline{\mathsf{mfC}}$, and the *sanitized* masked Feistel denoted $\widetilde{\mathsf{mfC}}$, both of which will be helpful in our analysis. For this section, we will imagine the X register to contain only the query point X, i.e. the forward query case. Recall that we answer inverse queries by flipping the truth table, answering as a forward query, and flipping back. In this way, due to the symmetry of both masked Feistel and our compressed permutation oracle, it will suffice to analyze only forwards queries here.

Let us begin by defining the ideal masked Feistel compression operator $\overline{\mathsf{mfC}}$. This operator is the identity on the non-elegant subspace. On the elegant subspace, we will first define it for some fixed $I \in \mathbf{I}|^x$ of size $|I| < 2^n$, defining $\mathsf{mfC}_{x,I}$ as an operator on $\mathcal{H}(\{\mathbf{P}(I') \,:\, I' \in I|^x\})$.

$$\overline{\mathsf{mfC}}_{x,I} := \mathbb{1} - |\mathbf{P}(I)\rangle\langle\mathbf{P}(I)| - |+_{x,\mathbf{P}(I)}\rangle\langle+_{x,\mathbf{P}(I)}| + |\mathbf{P}(I)\rangle\langle+_{x,\mathbf{P}(I)}| + |+_{x,\mathbf{P}(I)}\rangle\langle\mathbf{P}(I)| \quad (201)$$

The masked Feistel decompression and its coherently controlled counterpart are then naturally defined as

$$\overline{\mathsf{mfC}}_{x,\mathsf{P}} := \bigoplus_{I\in\mathbf{I}|^x,|I|<2^n} \overline{\mathsf{mfC}}_{x,I,\mathsf{P}}, \qquad \overline{\mathsf{mfC}}_{\mathsf{XP}}\,|x\rangle_{\mathsf{X}} := |x\rangle_{\mathsf{X}} \otimes \overline{\mathsf{mfC}}_{x,\mathsf{P}}. \quad (202)$$

Note the clear similarity to the definition of the permutation compression operator pC; the only difference being the use of $\mathbf{P}(I)$ in place of I, reflecting the distinct—but isomorphic—Hilbert spaces on which these operators act. This can be formalized via the following two remarks.

Remark 5.6. *The operator $\mathcal{I}$ intertwines ideal masked Feistel decompression $\overline{\mathsf{mfC}}^\dagger$ on the elegant subspace with decompression $\mathsf{pC}^\dagger$ on the corresponding subspace.*

$$\mathcal{I}_{\mathsf{P}} \cdot \overline{\mathsf{mfC}}^\dagger_{\mathsf{P}} \cdot \Pi^{\mathsf{ele}}_{\mathsf{P}} = \mathsf{pC}^\dagger_{\mathsf{P}} \cdot \mathcal{I}_{\mathsf{P}} \cdot \Pi^{\mathsf{ele}}_{\mathsf{P}} \quad (203)$$

Remark 5.7. *The action of $\overline{\mathsf{mfC}}_{\mathsf{XP}}$ interchanges the elegant space and the intersection of $\Pi^{\in\mathsf{db}}$ and the sophisticated space,*

$$\Pi^{\mathsf{ele}}_{\mathsf{XP}}\overline{\mathsf{mfC}}_{\mathsf{XP}} = \overline{\mathsf{mfC}}_{\mathsf{XP}}(\Pi^{\in\mathsf{db}}\Pi^{\mathsf{soph}})_{\mathsf{XP}}. \quad (204)$$

We use this to define the ideal masked feistel oracle $\overline{\mathsf{cmfO}}$, which acts as

$$\overline{\mathsf{cmfO}}_{\mathsf{XYP}} := \mathsf{ctrl\text{-}mfF}_{\mathsf{XP}} \cdot \overline{\mathsf{mfC}}_{\mathsf{XP}} \cdot \mathsf{mfP}_{\mathsf{XYP}} \cdot \overline{\mathsf{mfC}}^\dagger_{\mathsf{XP}} \cdot \mathsf{ctrl\text{-}mfF}^\dagger_{\mathsf{XP}}. \quad (205)$$

The following remark is based on similar reasoning to the prior.

Remark 5.8. *The ideal compressed mask Feistel operator is exactly intertwined by $\mathcal{I}$ with the compressed permutation oracle, on the flip elegant subspace.*

$$\mathcal{I}_{\mathsf{P}} \cdot \overline{\mathsf{cmfO}}_{\mathsf{XYP}} \cdot \Pi^{\mathsf{fele}}_{\mathsf{P}} = \mathsf{cpO}_{\mathsf{XYP}} \cdot \mathcal{I}_{\mathsf{P}} \cdot \Pi^{\mathsf{fele}}_{\mathsf{P}}. \tag{206}$$

The ideal operator is what we are aiming for, in that its action is precisely analogous to the compressed permutation oracle up to a change of basis. To show that the action of the ideal $\overline{\mathsf{cmfO}}$ is close the actual cmfO, we will define the sanitized compression as an intermediate operation. This operator is similar to standard compression, except it will maintain canonicity of the internal D_h, D_k, D_f tuple of databases. We define

$$\widetilde{\mathsf{mfC}}_{x,\mathsf{P}} |\pi\rangle_{\Pi} := |\pi\rangle_{\Pi} \otimes G_{\pi(x),\mathsf{HKF}}. \tag{207}$$

Observe that this operator is close the masked Feistel compression operator.

Corollary 5.9 (of Lemma 4.22)**.** *The sanitized and unsanitized masked Feistel compression operators are close:* $\left\| \widetilde{\mathsf{mfC}} - \mathsf{mfC} \right\|_{\leq t} = O\left(\sqrt{\frac{t^2}{2^n}}\right)$.

We will therefore be able to replace the unsanitized mfC with its sanitized variant $\widetilde{\mathsf{mfC}}$. It will further be helpful to define the *allowability* projector $\Pi^{\heartsuit}_{x,\mathsf{P}}$. This projects onto the subspace in which the query point x is allowed by the purification, i.e.

$$\Pi^{\heartsuit}_{x,\mathsf{P}} = \sum_{\pi,\omega,D\,:\,D\heartsuit\pi(x)} |\pi,\omega,D\rangle\langle\pi,\omega,D|_{\mathsf{P}}. \tag{208}$$

Naturally, the coherently controlled version is defined as $\Pi^{\heartsuit}_{\mathsf{XP}} |x\rangle_{\mathsf{X}} = |x\rangle_{\mathsf{X}} \otimes \Pi^{\heartsuit}_{x,\mathsf{P}}$. A useful fact about this projector is that it acts gently on the sophisticated subspace when the input point is undefined, which follows from the cromulence of $\mathcal{D}$.

Lemma 5.10. *We have*

$$\left\| (\mathbb{1} - \Pi^{\heartsuit}_{x}) \Pi^{x\in\mathsf{db},\perp} \Pi^{\mathsf{soph}} \right\|_{\leq t} = O\left(\sqrt{\frac{t^2}{2^n}}\right). \tag{209}$$

Proof. First, let us consider some fixed $I \in \mathbf{I}_{\leq 2^n}$ where $x \notin \mathsf{Dom}(I)$ (any state orthogonal

to this family is annihilated by $\Pi^{x\in\mathsf{db},\perp}\Pi^{\mathsf{soph}}$). We can write

$$|\mathbf{P}(I)\rangle_{\mathsf{P}} = \sum_{(\pi,\omega)\in\mathbf{R}_I} \sqrt{p_{\mathcal{D}_I}(\pi,\omega)}\,|\pi,\omega\rangle_{\Pi\Omega}\otimes|\mathbf{D}(\pi*I*\omega)\rangle_{\mathsf{HKF}}\,, \tag{210}$$

$$\left\|\Pi^{\heartsuit}_{x,\mathsf{P}}\,|\mathbf{P}(I)\rangle_{\mathsf{P}} - |\mathbf{P}(I)\rangle_{\mathsf{P}}\right\| = \left\| \sum_{\substack{(\pi,\omega)\in\mathbf{R}_I,\\(\pi,\omega)\notin\mathbf{R}_{x,I}}} \sqrt{p_{\mathcal{D}_I}(\pi,\omega)}\,|\pi,\omega\rangle_{\Pi\Omega}\otimes|\mathbf{D}(\pi*I*\omega)\rangle_{\mathsf{HKF}} \right\| \tag{211}$$

$$= \sqrt{\Pr_{(\pi,\omega)\sim\mathcal{D}_I}[(\pi,\omega)\notin\mathbf{R}_{x,I}]} \tag{212}$$

$$= O\left(\sqrt{\frac{t^2}{2^n}}\right), \qquad \text{(By cromulence)} \tag{213}$$

in particular the property established in Lemma 4.28. Further, observe that the state in Equation (211) is supported entirely on purifications P with the property $\mathsf{Supp}(P) = I$. Therefore, if we had considered a different database I, we would have an orthogonal state. The claim then follows from Lemma 2.3. □

Now, we will show closeness of ideal and real in the case where the purification does not support the input point.

Lemma 5.11. *We have*

$$\left\|(\overline{\mathsf{mfC}}^{\dagger} - \mathsf{mfC}^{\dagger})_{\mathsf{XP}}\cdot(\Pi^{\in\mathsf{db},\perp}\Pi^{\mathsf{ele}})_{\mathsf{XP}}\right\|_{\le t} = O\left(\sqrt{\frac{t^2}{2^n}}\right). \tag{214}$$

Proof. We begin with a simpler claim.

Claim. *Let $x\in[N]$ and $I\in\mathbf{I}|^x$ be of size $t = |I| < 2^n$. Then we have*

$$\left\|\widetilde{\mathsf{mfC}}_x^{\dagger}\cdot\Pi^{\heartsuit}_x\,|\mathbf{P}(I)\rangle - |+_{x,\mathbf{P}(I)}\rangle\right\| = O\left(\sqrt{\frac{t^2}{2^n}}\right). \tag{215}$$

Further, $\widetilde{\mathsf{mfC}}_x\cdot\Pi^{\heartsuit}_x\,|\mathbf{P}(I)\rangle$ is supported entirely on the span of $|\mathbf{P}(I[x\to y])\rangle$ for $y\in[N]\setminus\mathsf{Im}(I)$.

To prove this claim, observe that the state after the projection is

$$\Pi^{\heartsuit}_x\,|\mathbf{P}(I)\rangle = \sum_{(\pi,\omega)\in\mathbf{R}_{x,I}} \sqrt{p_{\mathcal{D}_I}(\pi,\omega)}\,|\mathbf{D}(\pi*I*\omega)\rangle_{\mathsf{HKF}}\otimes|\pi,\omega\rangle_{\Pi\Omega} \tag{216}$$

$$= \left(1 - O\left(\frac{t^2}{2^n}\right)\right)\cdot\sum_{(\pi,\omega)\in\mathbf{R}_{x,I}} \sqrt{p_{\mathcal{D}_{x,I}}(\pi,\omega)}\,|\mathbf{D}(\pi*I*\omega)\rangle_{\mathsf{HKF}}\otimes|\pi,\omega\rangle_{\Pi\Omega}\,, \tag{217}$$

by cromulence. Applying the $\widetilde{\mathsf{mfC}}_x$, we obtain

$$\widetilde{\mathsf{mfC}}_x \cdot \Pi_{\heartsuit x} |\mathbf{P}(I)\rangle = \left(1 - O\left(\frac{t^2}{2^n}\right)\right) \cdot \sum_{(\pi,\omega)\in\mathbf{R}_{x,I}} \sqrt{p_{\mathcal{D}_{x,I}}(\pi,\omega)} \left(G^{\dagger}_{\pi(x)} |\mathbf{D}(\pi * I * \omega)\rangle\right)_{\mathsf{HKF}} \otimes |\pi,\omega\rangle_{\Pi\Omega} \tag{218}$$

$$= \left(1 - O\left(\frac{t^2}{2^n}\right)\right) \cdot \sum_{\substack{(\pi,\omega)\in\mathbf{R}_{x,I},\\ A:=\pi*I*\omega}} \left(\sum_{v\in\mathbf{V}_{\pi(x),A}} \sqrt{\frac{p_{\mathcal{D}_{x,I}}(\pi,\omega)}{|\mathbf{V}_{\pi(x),A}|}} |\mathbf{D}(A[\pi(x)\to v])\rangle\right)_{\mathsf{HKF}} \otimes |\pi,\omega\rangle_{\Pi\Omega} \tag{219}$$

$$= \sum_{\substack{y\in[N]\setminus\mathsf{Dom}(I),\\ (\pi,\omega)\in\mathbf{R}_{I[x\to y]},\\ A:=\pi*I*\omega}} \left(1 \pm O\left(\sqrt{\frac{t^2}{2^n}}\right)\right) \cdot \sqrt{\frac{p_{\mathcal{D}_{I[x\to y]}}(\pi,\omega)}{N-t}} \cdot |\mathbf{D}(A[\pi(x)\to\omega^{-1}(y)])\rangle_{\mathsf{HKF}} \otimes |\pi,\omega\rangle_{\Pi\Omega}\,, \tag{220}$$

where the second line follows from Lemma 4.22, and the last line from cromulence, in particular the property established in Lemma 4.30. Observe that the above state would, without the $1 \pm O\left(\frac{t^2}{2^n}\right)$ term, be precisely $|+_{x,\mathbf{P}(I)}\rangle$. The two claims follow.

With this in hand, recall that any state orthogonal to $|\mathbf{P}(I)\rangle$ for $I \in \mathbf{I}|^x$ will be annihilated by $\Pi^{\in\mathsf{db},\perp}\Pi^{\mathsf{ele}}$. Further, by definition we have

$$\overline{\mathsf{mfC}}_{x,\mathsf{P}} |\mathbf{P}(I)\rangle := |+_{x,\mathbf{P}(I)}\rangle\,. \tag{221}$$

Thus, for any state in the family $|x\rangle_{\mathsf{X}} |\mathbf{P}(I)\rangle_{\mathsf{P}}$, we have that the images of $\widetilde{\mathsf{mfC}}^{\dagger}_x$ and of $\overline{\mathsf{mfC}}^{\dagger}$ lie in the span of $|x\rangle_{\mathsf{X}} |\mathbf{P}(I[x\to y])\rangle_{\mathsf{P}}$. We would obtain an orthogonal subspace for any different choice if $x \in [N]$ and $I \in \mathbf{I}|^x$, so these operators are close for any state in the linear combination. Putting everything together, we have

$$\left\|(\overline{\mathsf{mfC}}^{\dagger} - \mathsf{mfC}^{\dagger})_{\mathsf{X}P} \cdot (\Pi^{\in\mathsf{db},\perp}\Pi^{\mathsf{ele}})_{\mathsf{X}\mathsf{P}}\right\|_{\le t} \le \underbrace{\left\|(\widetilde{\mathsf{mfC}}^{\dagger} - \mathsf{mfC}^{\dagger})_{\mathsf{X}P} \cdot (\Pi^{\in\mathsf{db},\perp}\Pi^{\mathsf{ele}})_{\mathsf{X}\mathsf{P}}\right\|_{\le t}}_{O(\sqrt{t^2/2^n})} +$$

$$\underbrace{\left\|(\widetilde{\mathsf{mfC}}^{\dagger}\Pi^{\heartsuit} - \widetilde{\mathsf{mfC}}^{\dagger})_{\mathsf{X}P} \cdot (\Pi^{\in\mathsf{db},\perp}\Pi^{\mathsf{ele}})_{\mathsf{X}\mathsf{P}}\right\|_{\le t}}_{O(\sqrt{t^2/2^n})} + \underbrace{\left\|(\overline{\mathsf{mfC}}^{\dagger} - \widetilde{\mathsf{mfC}}^{\dagger}\Pi^{\heartsuit})_{\mathsf{X}P} \cdot (\Pi^{\in\mathsf{db},\perp}\Pi^{\mathsf{ele}})_{\mathsf{X}\mathsf{P}}\right\|_{\le t}}_{O(\sqrt{t^2/2^n})} \tag{222}$$

$$= O\left(\sqrt{\frac{t^2}{2^n}}\right). \tag{223}$$

□

Another important case is the portion of the elegant subspace where the query point is already defined prior to decompression.

Lemma 5.12. *We have*

$$\left\|(\overline{\mathsf{mfC}}^{\dagger} - \mathsf{mfC}^{\dagger})_{\mathsf{XP}} \cdot (\Pi^{\in \mathsf{db}} \Pi^{\mathsf{ele}})_{\mathsf{XP}}\right\|_{\leq t} = O\left(\sqrt{\frac{t^2}{2^n}}\right). \tag{224}$$

Proof. Let $|\psi\rangle_{\mathsf{XP}} \in \mathcal{H}_{\mathsf{X}} \otimes \mathcal{H}^{\mathsf{ele}}$ be an elegant quantum state, where the P register is supported on databases of size at most t, and the query point is defined: $|\psi\rangle_{\mathsf{XP}} \in \Pi^{\in \mathsf{db}}_{\mathsf{XP}}$. Note that the entire operation is controlled on the X register, so by Lemma 2.3 it suffices to do the analysis for some fixed $x \in [N]$. Further, we can replace $\mathsf{mfC}^{\dagger}$ by $\widetilde{\mathsf{mfC}}^{\dagger}$ by Corollary 5.9. The action of ideal decompression $\overline{\mathsf{mfC}}^{\dagger}$ on this subspace is the identity, which we would like to show is close the action of $\widetilde{\mathsf{mfC}}^{\dagger}$. We will drop the dagger † from $\widetilde{\mathsf{mfC}}$, observing that if $U^{\dagger}$ is close to $\mathbb{1}$ on some subspace then by linearity so is U. In other words, we will show for any fixed $x \in [N]$ that $\left\|\widetilde{\mathsf{mfC}}_x |\psi\rangle - |\psi\rangle\right\|_{\leq t} = O\left(\sqrt{\frac{t^2}{2^n}}\right)$ for any elegant $|\psi\rangle \in \mathcal{H}^{\mathsf{ele}}$ with $|\psi\rangle \in \Pi^{x \in \mathsf{db}}$.

Let us expand $\widetilde{\mathsf{mfC}}$ into it's constituent parts. We will begin with the $G^{(F)}$ application, considering

$$\left\|\left(\sum_{\pi \in \mathbf{S}_N} |\pi\rangle\langle\pi|_{\Pi} \otimes G^{(F)}_{\pi(x),\mathsf{HKF}} - \mathbb{1}_{\mathsf{P}}\right) |\psi\rangle_{\mathsf{P}}\right\|. \tag{225}$$

Observe that both states (the one after applying the left or the right operator in the above parenthesis) will be supported entirely on purifications P such that $\mathsf{Supp}(P)|^x = I|^x$. This is because a re-assignment to the final element of the chain beginning at $\pi(x)$ (in D_f) can only modify the chain beginning at $\pi(x)$. Let us make then the assumption that the state $|\psi\rangle$ is supported purely on states of the form $|\mathbf{P}(I[x \to y])\rangle$ for $I \in \mathbf{I}|^x$ of size $|I| = t - 1$. The other components of $|\psi\rangle$ will remain orthogonal, so by Lemma 2.3 it suffices to do the analysis on this component only. We will write

$$|\psi\rangle = \sum_{y \in [N] \setminus \mathsf{Im}(I)} \alpha_y |\mathbf{P}(I[x \to y])\rangle_{\mathsf{P}}, \tag{226}$$

where $\sum_{y \in [N] \setminus \mathsf{Im}(I)} \alpha_y = 0$. We can equivalently write

$$|\psi\rangle_{\mathsf{P}} = \sum_{\substack{\pi,\omega \in \mathbf{R}_{x,I}, \\ A := \pi * I * \omega, \\ u := \pi(x)}} \sqrt{p_{\mathcal{D}_{x,I}}(\pi,\omega)}\, |\pi,\omega\rangle_{\Pi\Omega} \otimes \left(\sum_{v \in \mathbf{V}_{u,A}} \alpha_{\omega(v)} |+_{\mathbf{D}(A[u \to v])}\rangle_{\mathsf{HKF}}\right). \tag{227}$$

Applying the $G^{(F)}$ operator, we find

$$\left(\sum_{\pi \in \mathbf{S}_N} |\pi\rangle\langle\pi|_{\Pi} \otimes G^{(F)}_{\pi(x),\mathsf{HKF}}\right) |\psi\rangle_{\mathsf{P}} = \sum_{\substack{\pi,\omega \in \mathbf{R}_{x,I}, \\ A := \pi * I * \omega, \\ u := \pi(x)}} \sqrt{p_{\mathcal{D}_{x,I}}(\pi,\omega)}\, |\pi,\omega\rangle_{\Pi\Omega} \otimes \left(G^{(F)}_{u,\mathsf{HKF}} \sum_{v \in \mathbf{V}_{u,A}} \alpha_{\omega(v)} |+_{\mathbf{D}(A[u \to v])}\rangle_{\mathsf{HKF}}\right), \tag{228}$$

giving us difference

$$\left\| \left(\sum_{\pi \in \mathbf{S}_N} |\pi\rangle\langle\pi|_\Pi \otimes G^{(F)}_{\pi(x),\mathsf{HKF}} - \mathbb{1}_\mathsf{P} \right) |\psi\rangle_\mathsf{P} \right\| = \left\| \sum_{\substack{\pi,\omega \in \mathbf{R}_{x,I}, \\ A := \pi * I * \omega, \\ u := \pi(x)}} \sqrt{p_{\mathcal{D}_{x,I}}(\pi,\omega)}\, |\pi,\omega\rangle_{\Pi\Omega} \otimes \left(\left(G^{(F)}_{u,\mathsf{HKF}} - \mathbb{1}_\mathsf{HKF} \right) \sum_{v \in \mathbf{V}_{u,A}} \alpha_{\omega(v)} \left|+_{\mathbf{D}(A[u \to v])}\right\rangle_\mathsf{HKF} \right) \right\| \tag{229}$$

$$= \sqrt{\sum_{\pi,\omega \in \mathbf{R}_{x,I}} p_{\mathcal{D}_{x,I}}(\pi,\omega) \cdot T_{\omega,\alpha}} \tag{230}$$

where $T_{\omega,\alpha} = \sum_{l \in \{0,1\}^n} T_{l,\omega,\alpha}$ and $T_{l,\omega,\alpha} = \left(\sum_{r \in \{0,1\}^n} \left| \alpha_{\omega(l\|r)} \right| \right)^2$. This last line follows from Lemma 4.25. This can be rewritten as

$$\left\| \left(\sum_{\pi \in \mathbf{S}_N} |\pi\rangle\langle\pi|_\Pi \otimes G^{(F)}_{\pi(x),\mathsf{HKF}} - \mathbb{1}_\mathsf{P} \right) |\psi\rangle_\mathsf{P} \right\| = \sqrt{\mathbb{E}_{(\pi,\omega)\sim\mathcal{D}_{x,I}} \left[\frac{T_{\omega,\alpha}}{2^n - t} \right]} \tag{231}$$

$$= O\left(\sqrt{\frac{t^2}{2^n}} \right), \tag{232}$$

where the last line follows from Lemma 5.5 combined with cromulence, and in particular Lemma 4.29. Therefore, we may replace the application of $G^{(F)}$ in $\widetilde{\mathsf{mfC}}$ with the identity $\mathbb{1}$ at the above cost. A similar argument shows that the application of $G^{(K)}$ and $G^{(H)}$ can be replaced with the identity $\mathbb{1}$ at a cost $O(\sqrt{t^2/2^n})$, using Lemma 4.24 and Lemma 4.23 repsectively in place of Lemma 4.25. Indeed, for these cases the final step corresponding to Equations (231) and (232) can be ommitted, due to the simpler form of Lemmas 4.23 and 4.24. □

Corollary 5.13 (of Lemmas 5.11 and 5.12)**.** *We have*

$$\left\| (\overline{\mathsf{mfC}}^\dagger - \mathsf{mfC}^\dagger)\Pi^{\mathsf{ele}} \right\|_{\leq t} = O\left(\sqrt{\frac{t^2}{2^n}} \right), \tag{233}$$

$$\left\| (\overline{\mathsf{mfC}} - \mathsf{mfC})\Pi^{\in \mathsf{db}}\Pi^{\mathsf{soph}} \right\|_{\leq t} = O\left(\sqrt{\frac{t^2}{2^n}} \right). \tag{234}$$

Proof. The first equation is immediate from the referenced lemmas. To conclude the sec-

ond, by Corollary 5.9 we may substitute $\widetilde{\mathsf{mfC}}$ for mfC. Then, observe that

$$\left\|(\overline{\mathsf{mfC}}-\widetilde{\mathsf{mfC}})\Pi^{\in\mathsf{db}}\Pi^{\mathsf{soph}}\right\|_{\leq t}=\left\|(\overline{\mathsf{mfC}}-\widetilde{\mathsf{mfC}})\overline{\mathsf{mfC}}^{\dagger}\cdot\Pi^{\mathsf{ele}}\cdot\overline{\mathsf{mfC}}\right\|_{\leq t} \tag{235}$$

$$=\left\|\Pi^{\mathsf{ele}}\cdot\overline{\mathsf{mfC}}-\widetilde{\mathsf{mfC}}\widetilde{\mathsf{mfC}}^{\dagger}\cdot\Pi^{\mathsf{ele}}\cdot\overline{\mathsf{mfC}}\right\|_{\leq t}+O\left(\sqrt{\frac{t^2}{2^n}}\right) \tag{236}$$

$$=O\left(\sqrt{\frac{t^2}{2^n}}\right). \tag{237}$$

□

We can now prove the main result for this section, establishing the soundness of our intertwiner on the elegant subspace.

Lemma 5.14. *On the flip elegant subspace, $\mathcal{I}$ is an approximate intertwiner between the compressed permutation oracle and the compressed masked Feistel oracle,*

$$\left\|(\mathcal{I}_{\mathsf{P}}\cdot\mathsf{cmfO}_{\mathsf{XYP}}-\mathsf{cpO}_{\mathsf{XYI}}\cdot\mathcal{I}_{\mathsf{P}})\Pi^{\mathsf{fele}}_{\mathsf{P}}\right\|_{\leq t}=O\left(\sqrt{\frac{t^2}{2^n}}\right). \tag{238}$$

Proof. We can write

$$\mathsf{cmfO}_{\mathsf{XYP}}\Pi^{\mathsf{fele}}_{\mathsf{XP}}=\mathsf{ctrl\text{-}mfF}_{\mathsf{XP}}\cdot\mathsf{mfC}_{\mathsf{XP}}\cdot\mathsf{mfP}_{\mathsf{XYP}}\cdot\mathsf{mfC}^{\dagger}_{\mathsf{XP}}\cdot\mathsf{ctrl\text{-}mfF}^{\dagger}_{\mathsf{XP}}\cdot\Pi^{\mathsf{fele}}_{\mathsf{XP}} \tag{239}$$

$$=\mathsf{ctrl\text{-}mfF}_{\mathsf{XP}}\cdot\mathsf{mfC}_{\mathsf{XP}}\cdot\mathsf{mfP}_{\mathsf{XYP}}\cdot\mathsf{mfC}^{\dagger}_{\mathsf{XP}}\cdot\Pi^{\mathsf{ele}}_{\mathsf{XP}}\cdot\mathsf{ctrl\text{-}mfF}^{\dagger}_{\mathsf{XP}}\cdot\Pi^{\mathsf{fele}}_{\mathsf{XP}} \tag{240}$$

$$=\mathsf{ctrl\text{-}mfF}_{\mathsf{XP}}\cdot\mathsf{mfC}_{\mathsf{XP}}\cdot\mathsf{mfP}_{\mathsf{XYP}}\cdot\overline{\mathsf{mfC}}^{\dagger}_{\mathsf{XP}}\cdot\Pi^{\mathsf{ele}}_{\mathsf{XP}}\cdot\mathsf{ctrl\text{-}mfF}^{\dagger}_{\mathsf{XP}}\cdot\Pi^{\mathsf{fele}}_{\mathsf{XP}}+O\left(\sqrt{\frac{t^2}{2^n}}\right) \quad \text{(Corollary 5.13)}$$

$$=\mathsf{ctrl\text{-}mfF}_{\mathsf{XP}}\cdot\mathsf{mfC}_{\mathsf{XP}}\cdot(\Pi^{\in\mathsf{db}}\Pi^{\mathsf{soph}})_{\mathsf{XP}}\cdot\mathsf{mfP}_{\mathsf{XYP}}\cdot\overline{\mathsf{mfC}}^{\dagger}_{\mathsf{XP}}\cdot\mathsf{ctrl\text{-}mfF}^{\dagger}_{\mathsf{XP}}\cdot\Pi^{\mathsf{fele}}_{\mathsf{XP}}+O\left(\sqrt{\frac{t^2}{2^n}}\right) \quad \text{(Remark 5.7)}$$

$$=\mathsf{ctrl\text{-}mfF}_{\mathsf{XP}}\cdot\overline{\mathsf{mfC}}_{\mathsf{XP}}\cdot(\Pi^{\in\mathsf{db}}\Pi^{\mathsf{soph}})_{\mathsf{XP}}\cdot\mathsf{mfP}_{\mathsf{XYP}}\cdot\overline{\mathsf{mfC}}^{\dagger}_{\mathsf{XP}}\cdot\mathsf{ctrl\text{-}mfF}^{\dagger}_{\mathsf{XP}}\cdot\Pi^{\mathsf{fele}}_{\mathsf{XP}}+O\left(\sqrt{\frac{t^2}{2^n}}\right) \quad \text{(Corollary 5.13)}$$

$$=\mathsf{ctrl\text{-}mfF}_{\mathsf{XP}}\cdot\overline{\mathsf{mfC}}_{\mathsf{XP}}\cdot\mathsf{mfP}_{\mathsf{XYP}}\cdot\overline{\mathsf{mfC}}^{\dagger}_{\mathsf{XP}}\cdot\mathsf{ctrl\text{-}mfF}^{\dagger}_{\mathsf{XP}}\cdot\Pi^{\mathsf{fele}}_{\mathsf{XP}}+O\left(\sqrt{\frac{t^2}{2^n}}\right)$$

$$=\overline{\mathsf{cmfO}}_{\mathsf{XYP}}\Pi^{\mathsf{fele}}_{\mathsf{XP}}+O\left(\sqrt{\frac{t^2}{2^n}}\right). \tag{241}$$

The claim now follows from Remark 5.8. □

The final case is decompression on states which are sophisticated but not elegant. For a technical reason, we will show that the decompression operator results in an invalid output and non-sophisticated purification. This will be useful in bounding the component of the state on this subspace. We use $\mathcal{H}_A/\mathcal{H}_B$ for spaces $\mathcal{H}_B\leq\mathcal{H}_A$ to denote the orthogonal complement of $\mathcal{H}_B$ in $\mathcal{H}_A$.

Lemma 5.15. *Let $|\psi\rangle_{\mathsf{XP}} \in \mathcal{H}_{\mathsf{X}} \otimes (\mathcal{H}^{\mathsf{soph}}/\mathcal{H}^{\mathsf{ele}})$ be a sophisticated, but not elegant, quantum state supported on purifications of size at most t. Then we have*

$$\left\|\Pi^{\mathsf{soph}}_{\mathsf{XP}} \cdot \widetilde{\mathsf{mfC}}^{\dagger}_{\mathsf{XP}} \cdot |\psi\rangle_{\mathsf{XP}}\right\| = O\left(\sqrt{\frac{1}{2^n}}\right), \quad \left\|\Pi^{\in\mathsf{db}}_{\mathsf{XP}} \cdot \widetilde{\mathsf{mfC}}^{\dagger}_{\mathsf{XP}} \cdot |\psi\rangle_{\mathsf{XP}}\right\| = O\left(\sqrt{\frac{1}{2^n}}\right). \quad (242)$$

Proof. Observe that that $|\psi\rangle_{\mathsf{XP}} \in \Pi^{\in\mathsf{db}}_{\mathsf{XP}}$, the orthogonal complement of $\Pi^{\in\mathsf{db}}_{\mathsf{XP}}$ in the sophisticated space $\mathcal{H}^{\mathsf{soph}}$ lies in the elegant subspace $\mathcal{H}^{\mathsf{ele}}$. By definition, we can write the sanitized masked Feistel decompression as

$$\widetilde{\mathsf{mfC}}^{\dagger}_{\mathsf{XP}} = \sum_{\pi\in\mathbf{S}_N} |\pi\rangle\langle\pi|_{\Pi} \otimes \left(G^{(F)}_{\pi(x),\mathsf{HKF}} G^{(K)}_{\pi(x),\mathsf{HKF}} G^{(H)}_{\pi(x),\mathsf{HK}}\right), \quad (243)$$

where we recall that the $G^{(H)}, G^{(K)}, G^{(F)}$ unitaries are involutions. By an argument analogous to that given in the proof of Lemma 5.12, the first two operations applied to the state, $G^{(H)}, G^{(K)}$, perturb the state a distance at most $O(\sqrt{1/2^n})$. By triangle inequality, we may therefore consider only the action of of the final $G^{(F)}$ operator. For now, let us fix some $x \in [N]$ to be the value in the X register, and $I \in \mathbf{I}^{|x|}$ of size $|I| = t-1$, and work out the action on the $|+_{x,\mathbf{P}(I)}\rangle_{\mathsf{P}}$ state. We find

$$\left(\sum_{\pi\in\mathbf{S}_N} |\pi\rangle\langle\pi|_{\Pi} \otimes G^{(F)}_{\pi(x),\mathsf{HKF}}\right)|+_{x,\mathbf{P}(I)}\rangle_{\mathsf{P}} = \sum_{\substack{\pi,\omega\in\mathbf{R}_{x,I},\\ A:=\pi*I*\omega,\\ u:=\pi(x)}} \sqrt{p_{\mathcal{D}_{x,I}}(\pi,\omega)}\,|\pi,\omega\rangle_{\Pi\Omega} \otimes \left(G^{(F)}_{u,\mathsf{HKF}} \sum_{v\in\mathbf{V}_{u,A}} \frac{1}{\sqrt{|\mathbf{V}_{u,A}|}}\,|+_{\mathbf{D}(A[u\to v])}\rangle_{\mathsf{HKF}}\right) \quad (244)$$

$$= \sum_{\substack{\pi,\omega\in\mathbf{R}_{x,I},\\ A:=\pi*I*\omega,\\ u:=\pi(x)}} \sqrt{p_{\mathcal{D}_{x,I}}(\pi,\omega)}\,|\pi,\omega\rangle_{\Pi\Omega} \otimes G^{(F)}_{u,\mathsf{HKF}}\Bigg(\sum_{\substack{D\in\mathbf{D}(A),\\ z\in H(u,D),\\ w\in K(u,D),\\ r\in\{0,1\}^n}} \frac{1}{\sqrt{|\mathbf{D}(A)|\cdot|H(u,D)|\cdot|K(u,D)|\cdot 2^n}}\cdot |D_h[u_L\to z]\rangle_{\mathsf{H}}\,|D_k[u_R\oplus z\to w]\rangle_{\mathsf{K}}\,|D_f[u_L\oplus w\to r]\rangle_{\mathsf{F}}\Bigg) \quad (245)$$

$$= \sum_{\substack{\pi,\omega\in\mathbf{R}_{x,I},\\ A:=\pi*I*\omega,\\ u:=\pi(x)}} \sqrt{p_{\mathcal{D}_{x,I}}(\pi,\omega)}\,|\pi,\omega\rangle_{\Pi\Omega} \otimes$$

$$\left(\sum_{\substack{D\in\mathbf{D}(A),\\ z\in H(u,D),\\ w\in K(u,D)}} \frac{1}{\sqrt{|\mathbf{D}(A)|\cdot|H(u,D)|\cdot|K(u,D)|}}\cdot \right.$$

$$\left. |D_h[u_L \to z]\rangle_{\mathsf{H}} |D_k[u_R \oplus z \to w]\rangle_{\mathsf{K}} |D_f\rangle_{\mathsf{F}} \right). \tag{246}$$

The key properties of this state is that it is fully undefined on input x, as can be seen by the lack of assignment of $u_L \oplus w$ in D_f; it is thus annihilated by $\Pi^{\in\mathsf{db}}_{\mathsf{XP}}$. Further, there is a length two rightward semi-chain beginning at u which is not completed, in every term in the superposition. The databases are non-minimal, and the state is therefore annihilated by $\Pi^{\mathsf{soph}}_{\mathsf{XP}}$. Further, the actual state $|\psi\rangle$ is a linear combination of states of this form, so we have

$$\left\| \Pi^{\in\mathsf{db}}_{\mathsf{XP}} \cdot \left(\sum_{\pi\in\mathbf{S}_N} |\pi\rangle\langle\pi|_{\Pi} \otimes G^{(F)}_{\pi(x),\mathsf{HKF}} \right) |\psi\rangle_{\mathsf{XP}} \right\| = 0, \tag{247}$$

and similarly for the sophisticated projector. The claim now follows from triangle inequality. □

5.4 Preserved subspaces

We are now ready to show that the elegant subspace of sophisticated provides a faithful proxy of the query-valid subspace, proving Lemma 5.4. We restate it below.

Lemma 5.4. *Let $|\psi\rangle_{\mathsf{XP}} \in \mathcal{H}_{\mathsf{X}} \otimes \mathcal{H}^{\mathsf{soph}}$ be a sophisticated quantum state satisfying $|\psi\rangle_{\mathsf{XP}} \in \Pi^{\mathsf{soph}}_{\mathsf{P}}$, and supported on purifications of size at most t. If we define*

$$|\phi\rangle_{\mathsf{XP}} := \Pi^{\mathsf{qval}}_{\mathsf{XP}} |\psi\rangle_{\mathsf{XP}}, \qquad |\mu\rangle_{\mathsf{XP}} := \Pi^{\mathsf{fele}}_{\mathsf{XP}} |\psi\rangle_{\mathsf{XP}}, \tag{184}$$

then we have

$$\| |\phi\rangle_{\mathsf{XP}} - |\mu\rangle_{\mathsf{XP}} \| = O\left(\sqrt{\frac{t^2}{2^n}}\right). \tag{185}$$

Proof. We can write

$$\| |\phi\rangle - |\mu\rangle \| = \left\| \Pi^{\mathsf{qval}} |\mu\rangle + \Pi^{\mathsf{qval}}(|\psi\rangle - |\mu\rangle) - |\mu\rangle \right\| \tag{248}$$

$$\leq \left\| \Pi^{\mathsf{qval}} |\mu\rangle - |\mu\rangle \right\| + \left\| \Pi^{\mathsf{qval}}(|\psi\rangle - |\mu\rangle) \right\| \tag{249}$$

$$= \left\| \Pi^{\mathsf{qval},\perp} |\mu\rangle \right\| + \left\| \Pi^{\mathsf{qval}}(|\psi\rangle - |\mu\rangle) \right\|. \tag{250}$$

Therefore, it suffices to prove that

$$\left\| \Pi^{\mathsf{qval},\perp} |\mu\rangle \right\| = O\left(\sqrt{\frac{t^2}{2^n}}\right), \qquad \left\| \Pi^{\mathsf{qval}}(|\psi\rangle - |\mu\rangle) \right\| = O\left(\sqrt{\frac{t^2}{2^n}}\right). \tag{251}$$

We perform the analysis for the case where the flip bit is $b = 0$, the other case follows similarly. Let us begin with the right hand equation. By definition, we have that $|\psi\rangle - |\mu\rangle \in \Pi^{\mathsf{fele},\perp}$ is non-flip-elegant. Then we have

$$\begin{aligned}
\left\|\Pi^{\mathsf{qval}}(|\psi\rangle - |\mu\rangle)\right\| &= \left\|\mathsf{mfC}_{\mathsf{XP}}\Pi^{\in\mathsf{db}}_{\mathsf{XP}}\mathsf{mfC}^{\dagger}_{\mathsf{XP}}(|\psi\rangle_{\mathsf{XP}} - |\mu\rangle_{\mathsf{XP}})\right\| && (252)\\
&= \left\|\Pi^{\in\mathsf{db}}_{\mathsf{XP}}\widetilde{\mathsf{mfC}}^{\dagger}_{\mathsf{XP}}(|\psi\rangle_{\mathsf{XP}} - |\mu\rangle_{\mathsf{XP}})\right\| + O\left(\sqrt{\frac{t^2}{2^n}}\right) && \text{(By Corollary 5.9)}\\
&= O\left(\sqrt{\frac{t^2}{2^n}}\right). && \text{(By Lemma 5.15)}
\end{aligned}$$

In the other case, where $|\mu\rangle \in \Pi^{\mathsf{fele}}$, we have

$$\begin{aligned}
\left\|\Pi^{\mathsf{qval},\perp}|\mu\rangle\right\| &= \left\|\mathsf{mfC}_{\mathsf{XP}}\Pi^{\in\mathsf{db},\perp}_{\mathsf{XP}}\mathsf{mfC}^{\dagger}_{\mathsf{XP}}|\mu\rangle_{\mathsf{XP}}\right\| && (253)\\
&= \left\|\Pi^{\in\mathsf{db},\perp}_{\mathsf{XP}}\overline{\mathsf{mfC}}^{\dagger}_{\mathsf{XP}}|\mu\rangle_{\mathsf{XP}}\right\| + O\left(\sqrt{\frac{t^2}{2^n}}\right) && \text{(By Corollaries 5.9 and 5.13)}\\
&= O\left(\sqrt{\frac{t^2}{2^n}}\right). && \text{(By Remark 5.7)}
\end{aligned}$$

□

Another important property which we can now prove is that the cmfO operator approximately preserves the elegant subspace.

Lemma 5.16. *Let $|\psi\rangle_{\mathsf{XYP}} \in \mathcal{H}_{\mathsf{XY}} \otimes \mathcal{H}^{\mathsf{fele}}$ be a flip elegant quantum state supported on purifications of size at most t. Then we have*

$$\left\|\Pi^{\mathsf{fele},\perp}_{\mathsf{P}} \cdot \mathsf{cmfO}_{\mathsf{XYP}}|\psi\rangle_{\mathsf{XYP}}\right\| = O\left(\sqrt{\frac{t^2}{2^n}}\right) \tag{254}$$

Proof. Let us analyze the case for forward queries, the proof for inverse queries follows similarly.

$$\begin{aligned}
\left\|\Pi^{\mathsf{ele},\perp}_{\mathsf{P}} \cdot \mathsf{cmfO}_{\mathsf{XYP}}|\psi\rangle_{\mathsf{XYP}}\right\| &= \left\|\Pi^{\mathsf{ele},\perp}_{\mathsf{P}} \cdot \mathsf{mfC}_{\mathsf{XP}} \cdot \mathsf{P}_{\mathsf{XYP}} \cdot \mathsf{mfC}^{\dagger}_{\mathsf{XP}}|\psi\rangle_{\mathsf{XYP}}\right\| && (255)\\
&= \left\|\Pi^{\mathsf{ele},\perp}_{\mathsf{P}} \cdot \overline{\mathsf{mfC}}_{\mathsf{XP}} \cdot \mathsf{P}_{\mathsf{XYP}} \cdot \overline{\mathsf{mfC}}^{\dagger}_{\mathsf{XP}}|\psi\rangle_{\mathsf{XYP}}\right\| + O\left(\sqrt{\frac{t^2}{2^n}}\right),
\end{aligned}$$

by Corollary 5.13, and the fact that P preserves the subspace $\Pi^{\in\mathsf{db}}\Pi^{\mathsf{soph}}$. Applying Remark 5.7, we find that the norm term is 0, showing the claim. □

Using this lemma, we can show that the compressed masked Feistel oracle approximately preserves the sophisticated subspace.

Lemma 5.17. *Let $\mathcal{A}$ be a q query quantum algorithm with intermediate operations $A_0, \ldots, A_q$, and denote the final state by*

$$|\phi_q\rangle_{\mathsf{AP}} = A_{q,\mathsf{A}}\mathsf{cmfO}_{\mathsf{AP}} \ldots A_{1,\mathsf{A}}\mathsf{cmfO}_{\mathsf{AP}}A_{0,\mathsf{A}} |0\rangle_{\mathsf{A}} |\bot\rangle_{\mathsf{P}} . \tag{256}$$

Then we have

$$\left\| \Pi_{\mathsf{P}}^{\mathsf{soph},\bot} |\phi_q\rangle_{\mathsf{AP}} \right\| = O\left(\sqrt{\frac{q^4}{2^n}}\right). \tag{257}$$

Proof. We will prove this statement inductively. Clearly the $A_{t,\mathsf{A}}$ operations cannot break validity nor increase the projector onto the non-sophisticated subspace, as operators on different registers commute. Therefore, we need only analyze the action of each query.

Suppose that the statement holds for the first t queries, and denote the state after these queries as

$$|\phi_t\rangle_{\mathsf{AP}} = A_{t,\mathsf{A}}\mathsf{cmfO}_{\mathsf{AP}} \ldots A_{1,\mathsf{A}}\mathsf{cmfO}_{\mathsf{AP}}A_{0,\mathsf{A}} |0\rangle_{\mathsf{A}} |\bot\rangle_{\mathsf{P}} . \tag{258}$$

We will split the A register into a further XY component for the query registers. We can write this state as

$$|\phi_t\rangle_{\mathsf{AXYP}} = \underbrace{\Pi_{\mathsf{XP}}^{\mathsf{qval}}\Pi_{\mathsf{P}}^{\mathsf{soph}} |\phi_t\rangle_{\mathsf{AXYP}}}_{|\mu\rangle} + \underbrace{\Pi_{\mathsf{XP}}^{\mathsf{qval}}\Pi_{\mathsf{P}}^{\mathsf{soph},\bot} |\phi_t\rangle_{\mathsf{AXYP}}}_{|\delta\rangle} + \underbrace{\Pi_{\mathsf{XP}}^{\mathsf{qval},\bot} |\phi_t\rangle_{\mathsf{AXYP}}}_{=0} . \tag{259}$$

We have $\||\delta\rangle\| \leq \left\| \Pi_{\mathsf{P}}^{\mathsf{soph},\bot} |\phi_t\rangle_{\mathsf{AXYP}} \right\|$. Turning to the first term, we have

$$\mathsf{cmfO}_{\mathsf{XYP}} |\mu\rangle_{\mathsf{XYP}} = \mathsf{cmfO}_{\mathsf{XYP}} \left(\Pi_{\mathsf{XP}}^{\mathsf{fele}} |\phi_t\rangle_{\mathsf{AXYP}} + O\left(\sqrt{\frac{t^2}{2^n}}\right) \right) \quad \text{(By Lemma 5.4)} \tag{260}$$

$$= \mathsf{cmfO}_{\mathsf{XYP}}\Pi_{\mathsf{XP}}^{\mathsf{fele}} |\phi_t\rangle_{\mathsf{AXYP}} + O\left(\sqrt{\frac{t^2}{2^n}}\right). \tag{261}$$

Recalling from Lemma 5.16 that cmfO approximately preserves the flip elegant subspace, and that the flip elegant subspace is a subspace of sophisticed $\mathcal{H}^{\mathsf{ele}} \leq \mathcal{H}_{\mathsf{X}} \otimes \mathcal{H}^{\mathsf{soph}}$, we have $\left\|\Pi^{\mathsf{ele},\bot}\mathsf{cmfO}_{\mathsf{XYP}} |\mu\rangle_{\mathsf{XYP}}\right\| = O(\sqrt{t^2/2^n})$. Plugging into Equation (259) and applying triangle inequality, we find that

$$\left\| \Pi_{\mathsf{P}}^{\mathsf{soph},\bot}\mathsf{cmfO}_{\mathsf{XYP}} |\phi_t\rangle_{\mathsf{AXYP}} \right\| = \left\| \Pi_{\mathsf{P}}^{\mathsf{soph},\bot} |\phi_t\rangle_{\mathsf{AXYP}} \right\| + O\left(\sqrt{\frac{t^2}{2^n}}\right). \tag{262}$$

The claim now follows by induction. □

Finally, we can show that the $\mathcal{I}$ operator acts as an approximate intertwiner between the two compressed permutation oracles, so long as the state is valid and close to sophisticated.

Lemma 5.18. *Let $|\psi\rangle_{\mathsf{AXYP}}$ be a quantum state supported on purifications of size at most t, and which is valid. Let $\delta = \left\|\Pi_{\mathsf{P}}^{\mathsf{soph},\perp} |\psi\rangle_{\mathsf{AXYP}}\right\|$ be the size of the unsophisticated component. Then we have*

$$\left\|\left(\mathcal{I}_{\mathsf{P}} \cdot \mathsf{cmfO}_{\mathsf{XYP}} - \mathsf{cpO}_{\mathsf{XYI}} \cdot \mathcal{I}_{\mathsf{P}}\right) |\psi\rangle_{\mathsf{AXYP}}\right\| = O\left(\delta + \sqrt{\frac{t^2}{2^n}}\right). \tag{263}$$

Proof. We can write $|\psi\rangle_{\mathsf{AXYP}}$ as

$$|\psi\rangle_{\mathsf{AXYP}} = \underbrace{\Pi_{\mathsf{XP}}^{\mathsf{qval}} \Pi_{\mathsf{P}}^{\mathsf{soph}} |\psi\rangle_{\mathsf{AXYP}}}_{|\mu\rangle} + \underbrace{\Pi_{\mathsf{XP}}^{\mathsf{qval}} \Pi_{\mathsf{P}}^{\mathsf{soph},\perp} |\psi\rangle_{\mathsf{AXYP}}}_{|\delta\rangle} + \underbrace{\Pi_{\mathsf{XP}}^{\mathsf{qval},\perp} |\psi\rangle_{\mathsf{AXYP}}}_{=0}. \tag{264}$$

The final term being 0 follows from the fact that the state is valid. Observe that $\||\delta\rangle\| \leq \delta$, and the operator acting on the state is norm at most 2, so this incurs an additive $O(\delta)$. Turning to the remaining component, we can write

$$\begin{aligned}
\mathcal{I}_{\mathsf{P}} \cdot \mathsf{cmfO}_{\mathsf{XYP}} |\mu\rangle_{\mathsf{AXYP}} &= \mathcal{I}_{\mathsf{P}} \cdot \mathsf{cmfO}_{\mathsf{XYP}} \cdot \Pi_{\mathsf{XP}}^{\mathsf{fele}} |\psi\rangle_{\mathsf{AXYP}} + O\left(\sqrt{\frac{t^2}{2^n}}\right) && \text{(By Lemma 5.4)} \\
&= \mathsf{cpO}_{\mathsf{XYI}} \cdot \mathcal{I}_{\mathsf{P}} \cdot \Pi_{\mathsf{XP}}^{\mathsf{fele}} |\psi\rangle_{\mathsf{AXYP}} + O\left(\sqrt{\frac{t^2}{2^n}}\right) && \text{(By Lemma 5.14)}
\end{aligned} \tag{265}$$

Plugging into Equation (264) and applying triangle inequality shows the claim. □

We are now ready to show the main result of this manuscript, that the compressed permutation oracle construction is sound. We remind the reader that up to this point, we assumed only that the twirling distributions were cromulent. In the next section, we will make an explicit choice for these distributions.

5.5 Putting the pieces together

Let $\mathcal{A}$ be a q query quantum algorithm which queries a random permutation oracle. Suppose that the algorithm has a workspace register A, and input/output registers X and Y of dimension 2^{2n+1} and 2^{2n}. The extra input bit indicates the direction of the query. We describe such an algorithm by a sequence of unitary operators $A_0, \ldots, A_q$ acting on the AXY system, interleaved with oracles $\mathcal{O}_{\varphi,\mathsf{XY}}$ which acts $\forall b \in \{0,1\}, x, y \in \{0,1\}^{2n}$ as

$$\mathcal{O}_{\varphi,\mathsf{XY}} |b,x\rangle_{\mathsf{X}} |y\rangle_{\mathsf{Y}} = |b,x\rangle_{\mathsf{X}} |y \oplus \varphi^{1-2b}(x)\rangle_{\mathsf{Y}}. \tag{266}$$

For simplicity, we will treat the registers XY as subregisters of the adversary state A. We can then write the view of such an algorithm as the density matrix

$$\rho_{\mathsf{A}}^{(\mathcal{O})} = \mathop{\mathbb{E}}_{\varphi \sim \mathbf{S}_{2^{2n}}} \left[\rho(A_{q,\mathsf{A}} \ldots \mathcal{O}_{\varphi,\mathsf{A}} A_{0,\mathsf{A}} |0\rangle_{\mathsf{A}})\right]. \tag{267}$$

Let us compare to the compressed oracle experiment, where we answer queries according to our compressed permutation oracle. In this experiment, we have a purifying register I in the Hilbert space $\mathcal{H}(\mathbf{I})$, spanned by all injective partial functions from $\{0,1\}^{2n} \to \{0,1\}^{2n}$. In this case, we answer queries using the unitary cpO, as described in Section 3. We think of the I register as being outside the algorithm's view, so it is traced out. We can write the final state in this experiment as

$$\rho_{\mathsf{A}}^{(\mathsf{cpO})} = \mathrm{Tr}_{\mathsf{I}} \left[\rho(A_{q,\mathsf{A}} \dots \mathsf{cpO}_{\mathsf{AI}} A_{0,\mathsf{A}} \left|0\right\rangle_{\mathsf{A}} \left|\perp\right\rangle_{\mathsf{I}}) \right]. \tag{268}$$

Our main result is that these two views are indistinguishable unless the algorithm has made a number of queries comparable to the twelveth-root of the permutation size.

Theorem 5.19. *Any q query quantum algorithm cannot distinguish the compressed permutation oracle from a uniform random permutation oracle on $[N] \simeq \{0,1\}^{2n}$ except with advantage*

$$\frac{1}{2} \left\| \rho_{\mathsf{A}}^{(\mathcal{O})} - \rho_{\mathsf{A}}^{(\mathsf{cpO})} \right\|_1 = O\left(\frac{q^3}{N^{1/4}} \right). \tag{269}$$

In particular, constant advantage requires $\Omega(\sqrt[12]{N})$ queries.

Proof. We will prove this statement by purifying. The natural purification of a random permutation oracle, which purifies $\rho_{\mathsf{A}}^{(\mathcal{O})}$, is to maintain a uniform superposition of permutation truth tables outside the view of the adversary. We will consider the alternate, more complicated purification described in Section 4, with a uniform twirl. This can be written as follows, where P is the purifying register containing sub-registers $\Pi\Omega\mathsf{HKF}$,

$$\left|\psi^{(\mathcal{O})}\right\rangle_{\mathsf{AXYP}} = A_{q,\mathsf{A}} \dots \mathsf{mfP}_{\mathsf{AP}} A_{0,\mathsf{A}} \left|0\right\rangle_{\mathsf{A}} \otimes \sqrt{\frac{1}{(N!)^2 (2^n)^{3 \cdot 2^n}}} \sum_{\pi,\omega \in \mathbf{S}_N, h,k,f \in \mathbf{F}_{2^n,2^n}} \left|\pi,\omega,h,k,f\right\rangle_{\mathsf{P}}. \tag{270}$$

In other words, we take the twirling distribution $\mathcal{D}$ to be uniform over $\mathbf{S}_N^{\times 2}$. This distribution is cromulent, as shown in Section 4.4. This oracle is perfectly indistinguishable from a random permutation oracle, $\mathrm{Tr}_{\mathsf{P}}[\rho(\left|\psi^{(\mathcal{O})}\right\rangle_{\mathsf{AXYP}})] = \rho_{\mathsf{A}}^{(\mathcal{O})}$, as the permutation is a product of a random permutation (say π) with an independent permutation.

Now, by the standard compressed oracle theory, the state

$$\left|\psi^{(\mathsf{cmfO})}\right\rangle_{\mathsf{AP}} = A_{q,\mathsf{A}} \dots \mathsf{cmfO}_{\mathsf{AP}} A_{0,\mathsf{A}} \left|0\right\rangle_{\mathsf{A}} \otimes \frac{1}{N!} \sum_{\pi,\omega \in \mathbf{S}_N} \left|\pi,\omega,\perp,\perp,\perp\right\rangle_{\mathsf{P}} \tag{271}$$

is perfectly indistinguishable after tracing out the P register, in that we have

$$\rho_{\mathsf{A}}^{(\mathcal{O})} = \mathrm{Tr}_{\mathsf{P}}[\rho(\left|\psi^{(\mathcal{O})}\right\rangle_{\mathsf{AP}})] = \mathrm{Tr}_{\mathsf{P}}[\rho(\left|\psi^{(\mathsf{cmfO})}\right\rangle_{\mathsf{AP}})]. \tag{272}$$

The purification of the compressed permutation oracle is

$$\left|\psi^{(\mathsf{cpO})}\right\rangle_{\mathsf{AI}} = A_{q,\mathsf{A}} \dots \mathsf{cpO}_{\mathsf{AI}} A_{0,\mathsf{A}} \left|0\right\rangle_{\mathsf{A}} \left|\perp\right\rangle_{\mathsf{I}}, \qquad \mathrm{Tr}_{\mathsf{I}}[\rho(\left|\psi^{(\mathsf{cpO})}\right\rangle_{\mathsf{AI}})] = \rho_{\mathsf{A}}^{(\mathsf{cpO})}. \tag{273}$$

Now we can consider the hybrid states,

$$|\psi_0^{(\mathsf{Hyb})}\rangle = A_{q,\mathsf{A}} \dots \mathsf{cpO}_{\mathsf{AI}} A_{0,\mathsf{A}} \mathcal{I}_{\mathsf{P}} |0\rangle_{\mathsf{A}} |\mathbf{P}(\perp)\rangle_{\mathsf{P}} \tag{274}$$

$$|\psi_1^{(\mathsf{Hyb})}\rangle = A_{q,\mathsf{A}} \dots \mathsf{cpO}_{\mathsf{AI}} A_{1,\mathsf{A}} \mathcal{I}_{\mathsf{P}} \mathsf{cmfO}_{\mathsf{AP}} A_{0,\mathsf{A}} |0\rangle_{\mathsf{A}} |\mathbf{P}(\perp)\rangle_{\mathsf{P}} \tag{275}$$

$$\dots$$

$$|\psi_q^{(\mathsf{Hyb})}\rangle = \mathcal{I}_{\mathsf{P}} A_{q,\mathsf{A}} \dots \mathsf{cmfO}_{\mathsf{AP}} A_{0,\mathsf{A}} |0\rangle_{\mathsf{A}} |\mathbf{P}(\perp)\rangle_{\mathsf{P}}, \tag{276}$$

where $|\psi^{(\mathsf{cpO})}\rangle_{\mathsf{AI}} = |\psi_0^{(\mathsf{Hyb})}\rangle$ and $\mathcal{I}_{\mathsf{P}} |\psi^{(\mathsf{cmfO})}\rangle_{\mathsf{AP}} = |\psi_q^{(\mathsf{Hyb})}\rangle$. For $t \leq q$, let us define

$$|\phi_t\rangle_{\mathsf{AP}} = A_{t,\mathsf{A}} \dots \mathsf{cmfO}_{\mathsf{AP}} A_{0,\mathsf{A}} |0\rangle_{\mathsf{A}} |\mathbf{P}(\perp)\rangle_{\mathsf{P}}, \tag{277}$$

observing that $\left\| \Pi_{\mathsf{P}}^{\mathsf{soph},\perp} |\phi_t\rangle_{\mathsf{AP}} \right\| = O(\sqrt{t^4/2^n})$ by Lemma 5.16, and further the $|\phi_t\rangle_{\mathsf{AP}} \in \Pi_{\mathsf{P}}^{\mathsf{val}}$, lying entirely in the valid subspace as cmfO preserves validity. Observing that $\mathcal{I}_{\mathsf{P}}$ clearly commutes with the $A_{t,\mathsf{A}}$ operator, we can then compute the successive differences between hybrid states to be

$$\begin{aligned} \left\| |\psi_{t+1}^{(\mathsf{Hyb})}\rangle - |\psi_t^{(\mathsf{Hyb})}\rangle \right\| &= \| (\mathcal{I}_{\mathsf{P}} \cdot \mathsf{cmfO}_{\mathsf{AP}} - \mathsf{cpO}_{\mathsf{AI}} \cdot \mathcal{I}_{\mathsf{P}}) |\phi_t\rangle_{\mathsf{AP}} \| \\ &= O\left(\sqrt{\frac{t^4}{2^n}}\right). \qquad \text{(By Lemma 5.14)} \end{aligned} \tag{278}$$

Therefore, we have $\left\| \mathcal{I}_{\mathsf{P}} |\psi^{(\mathsf{cmfO})}\rangle_{\mathsf{AP}} - |\psi^{(\mathsf{cpO})}\rangle_{\mathsf{AI}} \right\| = O\left(\sqrt{\frac{q^6}{2^n}}\right)$. This equation combined with Lemma 2.2 proves the claim. □

6 Feistel security

Let $\mathsf{Feist}^{(m)}$ denote the distribution induced on permutations by m rounds of Feistel with truly random round functions. This is query-indistinguishable from the case where the round functions are pseudorandom. The size of all permutations considered is $N = 4^n$. Our convention is that the first round is "left-to-right", in that input x is transformed to $x_L \| (x_R \oplus h(x_L))$ in the first round.

Our proof that the seven round Feistel construction is a qPRP is nearly identical to the soundness proof of our compressed permutation oracle. In particular, we will show that Feistel is indistinguishable from the compressed permutation oracle. To do so, we will apply the machinery we have developed for masked Feistel to the case where the twirling distribution $\mathcal{D}$ draws π uniformly from $\mathsf{Feist}^{(2)}$, and ω uniformly from $(\mathsf{Feist}^{(2)})^{-1}$. In other words, we define

$$p_{\mathcal{D}}(\pi, \omega) = p_{\mathsf{Feist}^{(2)} \times \mathsf{Feist}^{(2)}}(\pi, \omega^{-1}). \tag{279}$$

Note that this distribution is invariant under the transformation $(\pi, \omega) \mapsto (\omega^{-1}, \pi^{-1})$, maintaining the flip symmetry of all our constructions thus far. The main technical component of this section is proving that this distribution is *cromulent*, which amounts to proving purely classical statistical properties. Once this is in place, we can use the results derived in the previous two sections to give a simple proof of security.

6.1 Twirling

We will study the statistical properties of the two round Feistel construction. In particular, we will show that $\mathcal{D} = \mathsf{Feist}^{(2)} \times (\mathsf{Feist}^{(2)})^{-1}$ is a cromulent distribution. From the flip symmetry above, it will suffice to show the properties in Definition 4.27 for the listed forward cases only. We can think of sampling from $\mathcal{D}_I$ as a procedure in which we assign inputs to the constituent functions in stages. We use $D_1^{(\pi)}$ and $D_2^{(\pi)}$ to denote the partial databases representing the two Feistel functions composing π throughout this procedure. We use $D_1^{(\omega)}$ to represent the last Feistel round, and $D_2^{(\omega)}$ to represent the second to last, in ω. Observe that $D_1^{(\cdot)}$ is left-to-right, and $D_2^{(\cdot)}$ is right-to-left, for both π and ω. We stress that this is purely a classical procedure, which is useful as a tool to prove probabilistic properties about this distribution. We use ideas from Definition 4.4 in this description.

1. For each $x \in \mathsf{Dom}(I)$, sample random values for $D_1^{(\pi)}(x_L)$ and $D_2^{(\pi)}(x_R \oplus D_1^{(\pi)}(x_L))$, and for each $y \in \mathsf{Im}(I)$, sample random values for $D_1^{(\omega)}(y_L)$ and $D_2^{(\omega)}(y_R \oplus D_1^{(\omega)}(y_L))$.
2. Observing that $A = \pi * I * \omega$ is now fully defined, if A is non-allowable, restart the procedure, and otherwise continue.
3. Sample the remaining undefined points at uniform random.

Now we observe a number of symmetries in this distribution. First, let us define a *right shift* in π, parameterized by an $s \in \{0,1\}^n$, as the transformation $(D_1^{(\pi)}, D_2^{(\pi)}) \mapsto_s (D_1^{(\pi)'}, D_2^{(\pi)'})$ where

$$D_1^{(\pi)}(w_1) = z_1 \Leftrightarrow D_1^{(\pi)'}(w_1) = z_1 \oplus s, \tag{280}$$
$$D_2^{(\pi)}(z_2) = w_2 \Leftrightarrow D_2^{(\pi)'}(z_2 \oplus s) = w_2. \tag{281}$$

Similarly, a right shift in ω parameterized by an $s \in \{0,1\}^n$ is the transformation $(D_1^{(\omega)}, D_2^{(\omega)}) \mapsto_s (D_1^{(\omega)'}, D_2^{(\omega)'})$ where

$$D_1^{(\omega)}(w_1) = z_1 \Leftrightarrow D_1^{(\omega)'}(w_1) = z_1 \oplus s, \tag{282}$$
$$D_2^{(\omega)}(z_2) = w_2 \Leftrightarrow D_2^{(\omega)'}(z_2 \oplus s) = w_2. \tag{283}$$

We define a left shift as a joint transformation on (π, ω), again parameterized by an $s \in \{0,1\}^n$, which maps $(D_2^{(\pi)}, D_2^{(\omega)}) \mapsto_s (D_2^{(\pi)'}, D_2^{(\omega)'})$ where

$$D_2^{(\pi)}(z_\pi) = w_\pi \Leftrightarrow D_2^{(\pi)'}(z_\pi) = w_\pi \oplus s, \tag{284}$$
$$D_2^{(\omega)}(z_\omega) = w_\omega \Leftrightarrow D_2^{(\omega)'}(z_\omega) = w_\omega \oplus s. \tag{285}$$

We can equivalently think of sampling from $\mathcal{D}_I$ with additional re-randomization, as illustrated by these symmetries. In the following lemma, inside of each step the order does not matter (e.g. the selections in step (1) can be made in any order, as long as they all occur before the selections in step (4)).

Lemma 6.1. *The following procedure is equivalent to sampling (π, ω) from $\mathcal{D}_I$, where π is defined by the two round Feistel $D_1^{(\pi)}, D_2^{(\pi)}$ and ω is defined by the inverse two round Feistel of $D_1^{(\omega)}, D_2^{(\omega)}$.*

1. *For each $x \in \mathsf{Dom}(I)$, sample random values for $D_1^{(\pi)}(x_L)$ and $D_2^{(\pi)}(x_R \oplus D_1^{(\pi)}(x_L))$, and for each $y \in \mathsf{Im}(I)$, sample random values for $D_1^{(\omega)}(y_L)$ and $D_2^{(\omega)}(y_R \oplus D_1^{(\omega)}(y_L))$. Noting that this suffices to define $A = \pi * I * \omega$, do this conditioned on the fact that A is allowable.*
2. *Sample random $s \sim \{0,1\}^n$, perform a right shift by s on π.*
3. *Sample random $s' \sim \{0,1\}^n$, perform a right shift by s' on ω.*
4. *Sample the remaining undefined points of the $D_i^{(\pi)}$ and $D_i^{(\omega)}$ at uniform random.*
5. *Sample random $s'' \sim \{0,1\}^n$, perform a left shift by s'' on (π, ω).*

Proof. Recall that allowability of $A = \pi * I * \omega$ is a condition on only the left substrings of A. A random right shift does not affect the left substrings of A, and so the right shifts in steps (2) and (3) do not change the distribution. Furthermore, the allowability conditions are a set of linear equations over $\mathbb{Z}_2^n$ on the left substrings with 2 or 4 variables in each. Shifting all of the variables of such an equation does not change its truth/falsehood, as every shift s has order 2 with $s \oplus s = 0$, so the last step does not affect allowability. □

Note that we can reason similarly about, e.g., $\mathcal{D}_{x,I}$, simply adding the condition of allowing x in step (1). The same argument implies that the random shifts do not change the truth or falsehood of this condition. With this in place, we are ready to prove that $\mathcal{D}$ satisfies the cromulent conditions.

Lemma 6.2. *Fix some $I \in \mathbf{I}_t$, $x \in [N] \setminus \mathsf{Dom}(I)$. Then we have*

$$\Pr_{(\pi,\omega)\sim\mathcal{D}_I}[(\pi, \omega) \in \mathbf{R}_{x,I}] = 1 - O\left(\frac{t^2}{2^n}\right) \tag{286}$$

when $t = O(2^{n/3})$.

Note the additional condition on t compared to the uniform twirl case; this nonetheless holds in our application.

Proof. We will use the characterization in Lemma 6.1. There are two cases to consider.

1. There is an $x' \in I$ with $x_L = x'_L$. Let $X \subset I$ be the set of such x''s. Let us imagine selecting the value $D_1^{(\pi)}(x_L)$, and the corresponding values in $D_2^{(\pi)}$, last in step (1). Prior to this selection, define $Z = \{x''_R \oplus D_1^{(\pi)}(x''_L) : x'' \in I, x''_L \neq x_L\}$ as the set of values on the right wire which have been determined, where $|Z| < t$. We use $Z' = \{z : \exists x' \in X \cup \{x\} \text{ s.t. } x'_R \oplus z \in Z\}$ as the set of assignments to $D_1^{(\pi)}(x_L)$

which would cause a collision with these values on the right wire. The probability of selecting some $D_1^{(\pi)}(x_L)$ which is not in Z' is uniform, so it only remains to analyze the probability of some selection in Z'. Such a selection may be more likely if it causes some x' to immediately satisfy allowability. In such a case, there are a factor $2^n/(2^n - O(t^2)) = 1 + O(t^2/2^n)$ more ways to choose the database values, as there would have been $2^n - O(t^2)$ allowable values of $D_2^{(\pi)}(x'_R \oplus D_1^{(\pi)}(x'_L))$ otherwise by Lemma 4.10. Over the at most t values in X, this gives a factor at most $(1 + O(t^2/2^n))^t = 1 + O(t^3/2^n) = O(1)$, meaning values in Z' are at most a constant factor more likely than values outside. It follows that values in Z' have probability $O(2^{-n})$. We further have $|Z'| = O(t^2)$ by a simple counting argument.

Over the random choice (conditioned on prior events) of $D_1^{(\pi)}(x_L)$, we then have

$$\Pr_{(\pi,\omega)\sim\mathcal{D}_I}[D_1^{(\pi)}(x_L) \in Z'] = O(t^2/2^n) \tag{287}$$

from the union bound. Supposing that this event does not happen, the value $D_2^{(\pi)}(x_R \oplus D_1^{(\pi)}(x_L))$ will be selected at uniform random in step (4). Once again by Lemma 4.10, this will lead to an allowable value of $\pi(x)_L$ with probability $1 - O(t^2/2^n)$, showing the claim.

2. There is no $x' \in I$ with $x_L = x'_L$. Then, the value $D_1^{(\pi)}(x_L)$ will be selected in step (4), and will result in a value $D_2^{(\pi)}(x_R \oplus D_1^{(\pi)}(x_L))$ which is not already selected with probability $1 - O(t/2^n)$. Similar to the previous case, conditioned on this value being sampled fresh in step (4) there is once again a probability $1 - O(t^2/2^n)$ of x being disallowed, proving the claim.

□

Lemma 6.3. *Fix some $I \in \mathbf{I}_t$, $y \neq y' \in [N] \setminus \mathsf{Im}(I)$, $x \in [N] \setminus \mathsf{Dom}(I)$ and $l \in \{0,1\}^n$. Then we have*

$$\Pr_{(\pi,\omega)\sim\mathcal{D}_{x,I}}[\omega^{-1}(y)_L = l] = \frac{1}{2^n} \tag{288}$$

$$\Pr_{(\pi,\omega)\sim\mathcal{D}_{x,I}}[\omega^{-1}(y)_L = \omega^{-1}(y')_L = l] = \frac{1}{2^{2n}} \pm O\left(\frac{t^2}{2^{3n}}\right). \tag{289}$$

Proof. We will use the characterization in Lemma 6.1. For the first equation, consider performing steps (1-4). Then, step (5) will re-randomize the value of $\omega^{-1}(y)$, such that it will satisfy $\omega^{-1}(y)_L = l$ with probability 2^{-n}.

For the second equation, let us begin by analyzing the probability that $\omega^{-1}(y)_L = \omega^{-1}(y')_L$ after steps (1-4). There are the following cases.

1. $y_L = y'_L$. Then, after the selection of $D_1^{(\omega)}(y_L)$, it is guaranteed that $y_R \oplus D_1^{(\omega)}(y_L) \neq y'_R \oplus D_1^{(\omega)}(y'_L)$. Let us define

$$z = y_R \oplus D_1^{(\omega)}(y_L), \qquad z' = y_R \oplus D_1^{(\omega)}(y_L). \tag{290}$$

Now, consider performing all of the selections which can be performed prior to selecting $D_2^{(\omega)}(z)$. WLOG we can assume that $D_2^{(\omega)}(z')$ is already selected (otherwise, swap y and y'). The selection of $D_2^{(\omega)}(z)$ will occur either at the end of step (1) (in which case there are at least $2^n - O(t^2)$ admissible values) or at the end of step (4) (in which case there are 2^n admissible values), depending on the prior selections. In all cases, we have shown

$$\frac{1}{2^n} \leq \Pr_{(\pi,\omega)\sim\mathcal{D}_{x,I}}[\omega^{-1}(y)_L = \omega^{-1}(y')_L \,|\, y_L = y'_L] \leq \frac{1}{2^n - O(t^2)}. \tag{291}$$

2. $y_L \neq y'_L$, and there is $y'', y''' \in \mathsf{Im}(I)$ with $y''_L = y_L$ and $y'''_L = y'_L$. We will call this case E_2. Consider in step (1) first selecting all of the values relevant to π, all of the relevant values for $D_1^{(\omega)}$ except $D_1^{(\omega)}(y_L)$, and none of the values for $D_2^{(\omega)}$. Let us define $z' = y'_R \oplus D_1^{(\omega)}(y'_L)$. In the selection of $D_1^{(\omega)}(y_L)$, there are at least $2^n - O(t^2)$ admissible values, i.e. values not immediately violating the allowability condition. If we define $z = y_R \oplus D_1^{(\omega)}(y_L)$, then we will have $z = z'$ with probability no more than $\frac{1}{2^n - O(t^2)}$; call this event E_{col}.

   In the case of E_{col}, we are guaranteed that $\omega^{-1}(y)_L \neq \omega^{-1}(y')_L$, as if there is a right collision there cannot be a left collision. In the other case, we can consider once again selecting $D_2^{(\omega)}(z)$ after $D_2^{(\omega)}(z')$, and after all the (other) selections in step (1) are made. Similar to the previous case, depending on the step in which this occurs, the probability of $\omega^{-1}(y)_L = \omega^{-1}(y')_L$ lies between $\frac{1}{2^n}$ and $\frac{1}{2^n - O(t^2)}$. We then have

$$\begin{aligned}
\Pr_{(\pi,\omega)\sim\mathcal{D}_{x,I}}[\omega^{-1}(y)_L = \omega^{-1}(y')_L \,|\, E_2] = &\Pr_{(\pi,\omega)\sim\mathcal{D}_{x,I}}[\omega^{-1}(y)_L = \omega^{-1}(y')_L \,|\, E_2, E_{\mathsf{col}}] + \\
&\underbrace{\Pr_{(\pi,\omega)\sim\mathcal{D}_{x,I}}[\omega^{-1}(y)_L = \omega^{-1}(y')_L \,|\, E_2, \neg E_{\mathsf{col}}]}_{=0} \qquad (292)\\
= &\left(1 - \frac{1}{2^n - O(t^2)}\right) \cdot \left(\frac{1}{2^n \pm O(t^2)}\right) \qquad (293)\\
= &\frac{1}{2^n} \pm O\left(\frac{t^2}{2^{2n}}\right) \qquad (294)
\end{aligned}$$

3. There is no $y'' \in \mathsf{Im}(I)$ with $y''_L = y_L$, or no such y'' with $y''_L = y'_L$. We can WLOG assume the first one. We call this case E_3. Now, we imagine performing all of the selections in step (1), as well as selecting $D_1^{(\omega)}(y'_L)$ to define $z' = y'_R \oplus D_1^{(\omega)}(y'_L)$, and selecting $D_2^{(\omega)}(z')$ to define $w' = y'_L \oplus D_2^{(\omega)}(z')$ (this selection may occur in step (1), or as the first selection in step (4)). Now let $d = y_L \oplus w'$. Let E_d denote the event that d appears in the image of $D_2^{(\omega)}$ at this stage. By the shift invariance of allowability, and the fact that $D_2^{(\omega)}$ is of size at most $t + 1$, we have that E_d occurs with probability $O(t/2^n)$.

Now we consider selecting $D_1^{(\omega)}(y_L)$ and defining $z = y_R \oplus D_1^{(\omega)}(y_L)$. Let E_z denote the event that z is in the domain of $D_2^{(\omega)}$. There are at most $t+1$ such values, so E_z occurs with probability $O(t/2^n)$ over the randomness of the selection. Note that E_z is independent of E_d. If E_z occurs, then the value $w = y_L \oplus D_2^{(\omega)}(z)$ is already defined, and we have that $\omega^{-1}(y)_L = \omega^{-1}(y')_L$ can only occur if $D_2^{(\omega)}(z) = d$, i.e. if E_d occurred.

Turning to the other case, we consider randomly selecting $D_2^{(\omega)}(z)$, which defines $w = y_L \oplus D_2^{(\omega)}(z)$. With probability 2^{-n}, we will have $w = w'$. This occurs regardless of E_d. Assembling all the cases, we can write

$$\begin{aligned}\Pr_{(\pi,\omega)\sim\mathcal{D}_{x,I}}[\omega^{-1}(y)_L = \omega^{-1}(y')_L \mid E_3] = &\underbrace{\Pr_{(\pi,\omega)\sim\mathcal{D}_{x,I}}[\omega^{-1}(y)_L = \omega^{-1}(y')_L \mid E_3, E_d, E_z]}_{O(t^2/2^{2n})} + \\ &\underbrace{\Pr_{(\pi,\omega)\sim\mathcal{D}_{x,I}}[\omega^{-1}(y)_L = \omega^{-1}(y')_L \mid E_3, \neg E_d, E_z]}_{=0} + \\ &\underbrace{\Pr_{(\pi,\omega)\sim\mathcal{D}_{x,I}}[\omega^{-1}(y)_L = \omega^{-1}(y')_L \mid E_3, \neg E_z]}_{(1-O(t/2^n))(2^{-n})} \qquad (295)\end{aligned}$$

$$= \frac{1}{2^n} \pm O\left(\frac{t^2}{2^{2n}}\right). \qquad (296)$$

Together, these cases imply that we have $\omega^{-1}(y)_L = \omega^{-1}(y')_L$ with probability $2^{-n} \pm O(t^2/2^{2n})$, prior to step (5). Over the random shift in step (5), we will have $\omega^{-1}(y) = l$ with probability 2^{-n}, independent of the prior event. Multiplying together gives the claim. □

Lemma 6.4. *For any $I \in \mathbf{I}_t$, $x \in [N] \setminus \mathsf{Dom}(I)$, $y \in [N] \setminus \mathsf{Im}(I)$, and $\pi, \omega \in \mathbf{R}_{I[x\to y]}$ with $A = \pi * I * \omega$, we have*

$$\sqrt{\frac{p_{\mathcal{D}_{I[x\to y]}}(\pi,\omega)}{(N-t)}} - \sqrt{\frac{p_{\mathcal{D}_{x,I}}(\pi,\omega)}{|\mathbf{V}_{\pi(x),A}|}} = \pm O\left(\sqrt{\frac{t^2}{2^n}} \cdot \sqrt{\frac{p_{\mathcal{D}_{I[x\to y]}}(\pi,\omega)}{N-t}}\right). \qquad (297)$$

Proof. Observe that we have

$$\frac{p_{\mathcal{D}_{I[x\to y]}}(\pi,\omega)}{p_{\mathcal{D}_{x,I}}(\pi,\omega)} = 1 \pm O\left(\frac{t^2}{2^n}\right), \qquad \frac{N-t}{\mathbf{V}_{\pi(x),\mathbf{A}}} = 1 \pm O\left(\frac{t^2}{2^n}\right), \qquad (298)$$

the first by an argument analogous to Lemma 6.2 and the second by Lemma 4.10. If we let $a = (N-t)/p_{\mathcal{D}_{I[x\to y]}}(\pi,\omega)$, then we can write

$$\sqrt{\frac{p_{\mathcal{D}_{x,I}}(\pi,\omega)}{|\mathbf{V}_{\pi(x),A}|}} = \frac{1}{\sqrt{a(1\pm O(t^2/2^n)}} \tag{299}$$

$$= \frac{1}{\sqrt{a}} \pm O\left(\frac{t^2}{2^n\cdot\sqrt{a}}\right), \tag{300}$$

where the last line follows from the taylor expansion of $\frac{1}{\sqrt{a+x}}$ about $x=0$. □

6.2 Putting the pieces together

As in Section 5, let $\mathcal{A}$ be a q query quantum algorithm which queries a permutation oracle $\mathcal{O}_\varphi$ for $\varphi \sim \mathsf{Feist}^{(7)}$. Suppose that the algorithm has a workspace register A, and input/output registers X and Y of dimension 2^{2n+1} and 2^{2n}. The extra input bit indicates the direction of the query. We describe such an algorithm by a sequence of unitary operators $A_0,\dots,A_q$ acting on the AXY system, interleaved with oracles $\mathcal{O}_{\varphi,\mathsf{XY}}$ which acts $\forall b\in\{0,1\}, x,y\in\{0,1\}^{2n}$ as

$$\mathcal{O}_{\varphi,\mathsf{XY}}\left|b,x\right\rangle_{\mathsf{X}}\left|y\right\rangle_{\mathsf{Y}} = \left|b,x\right\rangle_{\mathsf{X}}\left|y\oplus\varphi^{1-2b}(x)\right\rangle_{\mathsf{Y}}. \tag{301}$$

For simplicity, we will treat the registers XY as subregisters of the adversary state A. We can then write the view of such an algorithm on querying feistel as the density matrix

$$\rho_{\mathsf{A}}^{(\mathsf{Feist})} = \mathop{\mathbb{E}}_{\varphi\sim\mathsf{Feist}^{(7)}}\left[\rho(A_{q,\mathsf{A}}\dots\mathcal{O}_{\varphi,\mathsf{A}}A_{0,\mathsf{A}}\left|0\right\rangle_{\mathsf{A}})\right]. \tag{302}$$

We can write the view of the same algorithm querying a truly random permutation as

$$\rho_{\mathsf{A}}^{(\mathcal{O})} = \mathop{\mathbb{E}}_{\varphi\sim\mathbf{S}_N}\left[\rho(A_{q,\mathsf{A}}\dots\mathcal{O}_{\varphi,\mathsf{A}}A_{0,\mathsf{A}}\left|0\right\rangle_{\mathsf{A}})\right]. \tag{303}$$

We will show that these two views are indistinguishable unless the algorithm has made a number of queries comparable to the twelveth-root of the permutation size.

Theorem 6.5. *Any q query quantum algorithm cannot the seven round balanced Feistel construction from a uniform random permutation oracle on $[N]$ except with advantage*

$$\left\|\rho_{\mathsf{A}}^{(\mathcal{O})} - \rho_{\mathsf{A}}^{(\mathsf{Feist})}\right\|_1 = O\left(\frac{q^3}{N^{1/4}}\right). \tag{304}$$

In particular, constant advantage requires $\Omega(\sqrt[12]{N})$ bidirectional queries.

Proof. We will prove this statement by purifying. We will use our compressed oracle as an intermediary between the two, where we have a purifying register I. We can write the final state in this experiment as

$$\rho_{\mathsf{A}}^{(\mathsf{cpO})} = \mathrm{Tr}_{\mathsf{I}}\left[\rho(A_{q,\mathsf{A}}\dots\mathsf{cpO}_{\mathsf{AI}}A_{0,\mathsf{A}}\left|0\right\rangle_{\mathsf{A}}\left|\bot\right\rangle_{\mathsf{I}})\right]. \tag{305}$$

By Theorem 5.19, we have

$$\left\|\rho_{\mathsf{A}}^{(\mathsf{cpO})} - \rho_{\mathsf{A}}^{(\mathcal{O})}\right\|_1 = O\left(\frac{q^3}{N^{1/4}}\right). \tag{306}$$

We will therefore analyze the compresed permutation oracle experiment in place of the ideal permutation oracle.

To purify the seven round Feistel construction, we will follow the idea described in Section 4. In particular, we will consider the two leftmost Feistel rounds to define the permutation π, the two rightmost Feistel rounds to define the permutation ω, and the internal three rounds to correspond to the h, k, f functions. This can be written as follows, where P is the purifying register containing sub-registers $\Pi\Omega\mathsf{HKF}$,

$$|\psi^{(\mathcal{O})}\rangle_{\mathsf{AXYP}} = A_{q,\mathsf{A}} \dots \mathsf{mfP}_{\mathsf{AP}} A_{0,\mathsf{A}} |0\rangle_{\mathsf{A}} \otimes \sqrt{\frac{1}{(2^n)^{3\cdot 2^n}}} \sum_{\pi,\omega\in\mathbf{S}_N, h,k,f\in\mathbf{F}_{2^n,2^n}} \sqrt{p_{\mathcal{D}}(\pi,\omega)}\, |\pi,\omega,h,k,f\rangle_{\mathsf{P}}\,. \tag{307}$$

where $\mathcal{D}$ is defined as the distribution $\mathsf{Feist}^{(2)} \times (\mathsf{Feist}^{(2)})^{-1}$. This distribution is cromulent, as shown in Section 6.1, meaning we can apply the results of Section 5. This oracle is perfectly indistinguishable from the seven round Feistel construction, $\mathrm{Tr}_{\mathsf{P}}[\rho(|\psi^{(\mathcal{O})}\rangle_{\mathsf{AXYP}})] = \rho_{\mathsf{A}}^{(\mathsf{Feist})}$. Now, by the standard compressed oracle theory, the state

$$|\psi^{(\mathsf{mfC})}\rangle_{\mathsf{AP}} = A_{q,\mathsf{A}} \dots \mathsf{cmfO}_{\mathsf{AP}} A_{0,\mathsf{A}} |0\rangle_{\mathsf{A}} \otimes \sum_{\pi,\omega\in\mathbf{S}_N} \sqrt{p_{\mathcal{D}}(\pi,\omega)}\, |\pi,\omega,\perp,\perp,\perp\rangle_{\mathsf{P}} \tag{308}$$

is perfectly indistinguishable after tracing out the P register, in that we have

$$\rho^{(\mathcal{O})} = \mathrm{Tr}_{\mathsf{P}}[|\psi^{(\mathcal{O})}\rangle\langle\psi^{(\mathcal{O})}|] \tag{309}$$

$$= \mathrm{Tr}_{\mathsf{P}}[|\psi^{(\mathsf{cmfO})}\rangle\langle\psi^{(\mathsf{cmfO})}|]. \tag{310}$$

For comparison, we will take the following purification of the compressed oracle,

$$|\psi^{(\mathsf{cpO})}\rangle_{\mathsf{AI}} = A_{q,\mathsf{A}} \dots \mathsf{cpO}_{\mathsf{AI}} A_{0,\mathsf{A}} |0\rangle_{\mathsf{A}} |\perp\rangle_{\mathsf{I}}\,, \qquad \mathrm{Tr}_{\mathsf{I}}[\rho(|\psi^{(\mathsf{cpO})}\rangle_{\mathsf{AI}})] = \rho_{\mathsf{A}}^{(\mathsf{cpO})}. \tag{311}$$

Now we can consider the hybrid states,

$$|\psi_0^{(\mathsf{Hyb})}\rangle = A_{q,\mathsf{A}} \dots \mathsf{cpO}_{\mathsf{AI}} A_{0,\mathsf{A}} \mathcal{I}_{\mathsf{P}} |0\rangle_{\mathsf{A}} |\mathbf{P}(\perp)\rangle_{\mathsf{P}} \tag{312}$$

$$|\psi_1^{(\mathsf{Hyb})}\rangle = A_{q,\mathsf{A}} \dots \mathsf{cpO}_{\mathsf{AI}} A_{1,\mathsf{A}} \mathcal{I}_{\mathsf{P}} \mathsf{cmfO}_{\mathsf{AP}} A_{0,\mathsf{A}} |0\rangle_{\mathsf{A}} |\mathbf{P}(\perp)\rangle_{\mathsf{P}} \tag{313}$$

$$\dots$$

$$|\psi_q^{(\mathsf{Hyb})}\rangle = \mathcal{I}_{\mathsf{P}} A_{q,\mathsf{A}} \dots \mathsf{cmfO}_{\mathsf{AP}} A_{0,\mathsf{A}} |0\rangle_{\mathsf{A}} |\mathbf{P}(\perp)\rangle_{\mathsf{P}}\,, \tag{314}$$

where $|\psi^{(\mathsf{cpO})}\rangle_{\mathsf{AI}} = |\psi_0^{(\mathsf{Hyb})}\rangle$ and $\mathcal{I}_{\mathsf{P}} |\psi^{(\mathsf{cmfO})}\rangle_{\mathsf{AP}} = |\psi_q^{(\mathsf{Hyb})}\rangle$. For $t \le q$, let us define

$$|\phi_t\rangle_{\mathsf{AP}} = A_{t,\mathsf{A}} \dots \mathsf{cmfO}_{\mathsf{AP}} A_{0,\mathsf{A}} |0\rangle_{\mathsf{A}} |\mathbf{P}(\perp)\rangle_{\mathsf{P}}\,, \tag{315}$$

observing that $\left\|\Pi_{\mathsf{P}}^{\mathsf{soph},\perp}\left|\phi_t\right\rangle_{\mathsf{AP}}\right\| = O(\sqrt{t^4/2^n})$ by Lemma 5.16, and further the $\left|\phi_t\right\rangle_{\mathsf{AP}} \in \Pi_{\mathsf{P}}^{\mathsf{val}}$, lying entirely in the valid subspace as cmfO preserves validity. Observing that $\mathcal{I}_{\mathsf{P}}$ clearly commutes with the $A_{t,\mathsf{A}}$ operator, we can then compute the successive differences between hybrid states to be

$$\begin{aligned}\left\||\psi_{t+1}^{(\mathsf{Hyb})}\rangle - |\psi_t^{(\mathsf{Hyb})}\rangle\right\| &= \left\|\left(\mathcal{I}_{\mathsf{P}} \cdot \mathsf{cmfO}_{\mathsf{AP}} - \mathsf{cpO}_{\mathsf{AI}} \cdot \mathcal{I}_{\mathsf{P}}\right)\left|\phi_t\right\rangle_{\mathsf{AP}}\right\| \\ &= O\left(\sqrt{\frac{t^4}{2^n}}\right). \qquad \text{(By Lemma 5.14)} \qquad (316)\end{aligned}$$

Therefore, we have $\left\|\mathcal{I}_{\mathsf{P}}\left|\psi^{(\mathsf{cmfO})}\right\rangle_{\mathsf{AP}} - \left|\psi^{(\mathsf{cpO})}\right\rangle_{\mathsf{AI}}\right\| = O\left(\sqrt{\frac{q^6}{2^n}}\right)$. This equation combined with Lemma 2.2 proves the claim. □

7 Search bounds

We can also apply this framework to prove new search type query lower bounds. The most important ingredient in such bounds is the fundamental lemma, which relates the compressed permutation oracle database to points the algorithm "knows". This result can be applied to show lower bounds for various search problems in which the goal is to find a set of input output points having some property. All of our bounds follow the same rough format.

1. Switch to the compressed permutation oracle.
2. Show that the database is unlikely to contain a solution after q queries.
3. Apply Lemma 3.3 to conclude that the adversary is unlikely to have found a solution.

The first step is straightforward, while the second step varies in complexity depending on the problem. Often, it is enough to reason combinatorially about the databases, analogous to compressed function oracles. We will show how to use the fundamental lemma to apply the final step in general.

Throughout this section, our bounds will incur a loss due to the distinguishing advantage against the compressed permutation oracle. We quantify this by the loss term

$$\mathsf{Dist}(q) = O\left(\frac{q^3}{N^{1/4}}\right), \qquad (317)$$

and this term is often the loosest. Improving the soundness error of our construction would improve all the lower bounds derived. Most would be tight, ignoring the $\mathsf{Dist}(q)$ term.

To formalize this procedure, we will capture search problems by a target predicate $\mathcal{R} \subset ([N]^{\times 2})^*$, or a set of finite lists of input-output pairs representing the set of potential

solutions. Then, we consider an adversary $\mathcal{A}$ making q quantum queries to a permutation oracle $\mathcal{O}_\varphi$. The predicate $\mathcal{R}$ is known to the adversary (meaning its intermediate operations may depend on $\mathcal{R}$), although naturally the underlying permutation is not, outside of the information available by querying the oracle. The goal of the adversary is to find some input-output pairs of the permutation which satisfy the predicate. In particular, the game proceeds as follows.

Game 1. *Predicate search.*

1. *A random $\varphi \sim \mathbf{S}_N$ is selected to construct $\mathcal{O}_\varphi$.*
2. *An adversary $\mathcal{A}$ makes q quantum queries to $\mathcal{O}_\varphi$, and outputs a list $\mathbf{x} = [(x_1, y_1), \ldots, (x_l, y_l)]$.*
3. *The adversary wins if $\varphi(x_i) = y_i$ for all i, and $\mathbf{x} \in \mathcal{R}$.*

Alternatively, we could consider a game in which the adversary queries the compressed permutation oracle, and the final condition involves only analyzing the compressed oracle database. We say an $I \in \mathbf{I}_t$ satisfies $\mathcal{R}$, denoted $I \cap \mathcal{R} \neq \emptyset$, if there exists $\mathbf{x} = [(x_1, y_1), \ldots, (x_l, y_l)]$ such that $I(x_i) = y_i$ for all $i \in [l]$ and $\mathbf{x} \in \mathcal{R}$. We define the predicate projector

$$\Pi_{\mathcal{R}} |I\rangle = \begin{cases} |I\rangle & (\text{If } I \cap \mathcal{R} \neq \emptyset) \\ 0 & (\text{Otherwise}) \end{cases} \tag{318}$$

which projects onto databases which contain a set of input-output pairs in the predicate.

Game 2. *Predicate search with compressed oracles.*

1. *The compressed databases is initialized to $|\bot\rangle_{\mathsf{I}}$.*
2. *The adversary $\mathcal{A}$ with workspace A makes q quantum queries to $\mathsf{cpO}_{\mathsf{AI}}$.*
3. *The projective measurement $\mathcal{M}_{\mathcal{R}} = \{\Pi_{\mathcal{R}}, \Pi_{\mathcal{R}}^{\bot}\}$ is applied to register I. The adversary wins if the first outcome is achieved.*

It turns out that the winning probability of the second game provides an upper bound on the first. We can compare the same adversary across the two games by simply ignoring the output of $\mathcal{A}$ in the second game. This is analogous to the result of [Zha19] relating the compressed function oracle database to the winning probability of a search game.

Lemma 7.1. *Suppose that $\mathcal{A}$ wins Game 2 with probability p_2, and wins Game 1 with probability p_1. Then we have the bound*

$$\sqrt{p_1} \leq \sqrt{p_2} + \frac{l}{\sqrt{N-q-l}} + \mathsf{Dist}(q+l). \tag{319}$$

This bound also implies

$$p_1 = O\left(p_2 + \frac{l^2}{N-q-l} + \mathsf{Dist}(q+l)^2\right). \tag{320}$$

Proof. For Game 1, we can suppose WLOG that the adversary outputs some $\mathbf{x} \in \mathcal{R}$. If not, then modify the last step to check whether the output satisfies the predicate, and if needed switch to some output which satisfies the predicate. This cannot decrease the success probability, as the adversary is guaranteed to lose if $\mathbf{x} \notin \mathcal{R}$. Call this modified adversary $\mathcal{A}'$. Now, consider running Game 1 with the compressed permutation oracle and adversary $\mathcal{A}'$, replacing the final check with the projector $\Pi_{\mathsf{AI}}^{(\mathcal{O}_\varphi)}$ from Lemma 3.3. The resultant state is a distance at most $\mathsf{Dist}(q+l)$, by the definition of the distinguishing advantage.

Then, we can replace this projector with $\Pi_{\mathsf{AI}}^{(\mathsf{cpO})}$ by Lemma 3.3, and the resultant state is a distance at most $\frac{l}{\sqrt{N-q-l}}$. Now, this projector can only succeed if there is a tuple of input-output pairs in the database which are also in $\mathcal{R}$, i.e. the state is in $\Pi_{\mathcal{R},\mathsf{I}}$. The magnitude of this component is then at most $\sqrt{p_2}$. □

7.1 Simple examples

The simplest example of a problem for which we can show a lower bound is the one-more problem. In this problem, an adversary makes q queries to a random permutation and is tasked with finding $q+1$ input-output pairs.

Problem 1. *Sample $\varphi \sim \mathbf{S}_N$, to produce two way accessible oracle $\mathcal{O}_\varphi$. $\mathcal{A}$ makes q queries to $\mathcal{O}_\varphi$, and succeeds if it finds $\{(x_1, y_1), \ldots, (x_{q+1}, y_{q+1})\}$ such that $\varphi(x_i) = y_i$ for all $i \in [q+1]$, and $x_i \neq x_{i'}$ for $i \neq i'$.*

This problem does not appear to be in the scope of existing techniques for permutation problems. However, it is a simple corollary of our framework.

Theorem 7.2. *No q query quantum algorithm $\mathcal{A}$ can solve the one more problem, Problem 1, with probability exceeding*

$$p_{\mathcal{A}} = O\left(\frac{q^2}{N} + \mathsf{Dist}(q)^2\right). \tag{321}$$

In particular, exponential queries are required for constant error.

Proof. We can analyze this game in the compressed permutation oracle picture, as in Game 2. In this context, the predicate $\mathcal{R}$ is the set of all length $q+1$ lists of distinct input-output pairs. After q queries, the database is of size at most q by Lemma 3.2. Therefore, no satisfying list can appear in the database, so the success probability of the adversary in this game is zero. The claim now follows from Lemma 7.1. □

Another simple problem outside the scope of existing techniques is the cycle finding problem. In this problem, the adversary is tasked with finding a cycle in the permutation.

Problem 2. *Sample $\varphi \sim \mathbf{S}_N$, to produce two way accessible oracle $\mathcal{O}_\varphi$. $\mathcal{A}$ makes q queries to $\mathcal{O}_\varphi$, and succeeds if it finds $x_1, \dots, x_l$ such that $\varphi(x_i) = x_{i+1 \mod l}$ for all $i \in [l]$.*

Theorem 7.3. *No q query quantum algorithm $\mathcal{A}$ can solve cycle finding, Problem 2, with probability exceeding*

$$p_{\mathcal{A}} = O\left(\frac{q^3}{N} + \mathsf{Dist}(q)^2\right). \tag{322}$$

In particular, exponential queries are required for constant error.

Proof. We use the framework of Game 2, where the predicate $\mathcal{R}$ is the set of all cycles. Let $\mathcal{A}$ be an algorithm with Hilbert space $\mathcal{H}_{\mathsf{A}}$ containing query subregisters XY, and intermediate operations $A_0, \dots, A_q$ on this space. The success probability of this algorithm is given by

$$p_{\mathcal{A}} = \left\| \Pi_{\mathcal{R},\mathsf{I}} A_{q,\mathsf{A}} \mathsf{cpO}_{\mathsf{AI}} \dots A_{1,\mathsf{A}} \mathsf{cpO}_{\mathsf{AI}} A_{0,\mathsf{A}} \left|0\right\rangle_{\mathsf{A}} \left|\perp\right\rangle_{\mathsf{A}} \right\|^2. \tag{323}$$

To bound this quantity, we will show a bound on the commutator norm

$$\left\| [\Pi_{\mathcal{R},\mathsf{I}}, \mathsf{cpO}_{\mathsf{AI}}] \right\|_{\leq t} = O\left(\sqrt{\frac{t}{N}}\right). \tag{324}$$

To do this, we can modify the compression operator such that it cannot add cycles to the database. In particular, fix an $I \in \mathbf{I}_t$ and $x \notin \mathsf{Dom}(I)$. We will first define an operator on $\mathcal{H}(I|^x)$, writing

$$\left|\widetilde{+}_{x,I}\right\rangle_{\mathsf{I}} := \frac{1}{\sqrt{N - |\mathsf{Dom}(I) \cup \mathsf{Im}(I)|}} \sum_{y \in [N] \setminus \mathsf{Dom}(I) \setminus \mathsf{Im}(I)} \left|I[x \to y]\right\rangle \tag{325}$$

$$\widetilde{\mathsf{pC}}_{x,I} := \mathbb{1} - \left|I\right\rangle\!\left\langle I\right| - \left|\widetilde{+}_{x,I}\right\rangle\!\left\langle \widetilde{+}_{x,I}\right| + \left|\widetilde{+}_{x,I}\right\rangle\!\left\langle I\right| + \left|I\right\rangle\!\left\langle \widetilde{+}_{x,I}\right| \tag{326}$$

$$\widetilde{\mathsf{pC}}_{\mathsf{XI}} = \bigoplus_{x \in [N], I \in \mathbf{I}^x} \left|x\right\rangle\!\left\langle x\right|_{\mathsf{X}} \otimes \mathsf{pC}_{x,I,\mathsf{I}}. \tag{327}$$

Note that the above are unitary operators so long as the database is of size less than $N/2$, which is anyways beyond where the bound trivializes. Observe that $\left\| \left|\widetilde{+}_{x,I}\right\rangle - \left|+_{x,I}\right\rangle \right\|_{\leq t} = O(\sqrt{t/N})$, so $\left\| \mathsf{pC} - \widetilde{\mathsf{pC}} \right\|_{\leq t} = O(\sqrt{t/N})$ by Lemma 2.3. Now we have

$$[\Pi_{\mathcal{R},\mathsf{I}}, \widetilde{\mathsf{pC}}_{\mathsf{XI}}] = 0, \qquad [\Pi_{\mathcal{R},\mathsf{I}}, \mathsf{F}_{\mathsf{I}} \widetilde{\mathsf{pC}}_{\mathsf{XI}} \mathsf{F}_{\mathsf{I}}^\dagger] = 0, \tag{328}$$

$$[\Pi_{\mathcal{R},\mathsf{I}}, \mathsf{P}_{\mathsf{XYI}}] = 0, \qquad [\Pi_{\mathcal{R},\mathsf{I}}, \mathsf{F}_{\mathsf{I}} \mathsf{P}_{\mathsf{XYI}} \mathsf{F}_{\mathsf{I}}^\dagger] = 0, \tag{329}$$

as $\widetilde{\mathsf{pC}}$ cannot add a cycle to the database. If we write

$$\widetilde{\mathsf{cpO}}_{\mathsf{XYI}} = \left|0\right\rangle\!\left\langle 0\right|_{\mathsf{X}} \otimes (\widetilde{\mathsf{pC}}_{\mathsf{XI}} \cdot \mathsf{P}_{\mathsf{XYI}} \cdot \widetilde{\mathsf{pC}}_{\mathsf{XI}}^\dagger) + \left|1\right\rangle\!\left\langle 1\right|_{\mathsf{X}} \otimes (\mathsf{F}_{\mathsf{I}} \cdot \widetilde{\mathsf{pC}}_{\mathsf{XI}} \cdot \mathsf{P}_{\mathsf{XYI}} \cdot \widetilde{\mathsf{pC}}_{\mathsf{XI}}^\dagger \cdot \mathsf{F}_{\mathsf{I}}^\dagger) \tag{330}$$

$$= \left|0\right\rangle\!\left\langle 0\right|_{\mathsf{X}} \otimes (\widetilde{\mathsf{pC}}_{\mathsf{XI}} \cdot \mathsf{P}_{\mathsf{XYI}} \cdot \widetilde{\mathsf{pC}}_{\mathsf{XI}}^\dagger) + \left|1\right\rangle\!\left\langle 1\right|_{\mathsf{X}} \otimes (\mathsf{F}_{\mathsf{I}} \cdot \widetilde{\mathsf{pC}}_{\mathsf{XI}} \cdot (\mathsf{F}^\dagger \mathsf{F})_{\mathsf{I}} \cdot \mathsf{P}_{\mathsf{XYI}} \cdot (\mathsf{F}^\dagger \mathsf{F})_{\mathsf{I}} \cdot \widetilde{\mathsf{pC}}_{\mathsf{XI}}^\dagger \cdot \mathsf{F}_{\mathsf{I}}^\dagger), \tag{331}$$

then we have $\left\|\mathsf{cpO} - \widetilde{\mathsf{cpO}}\right\|_{\leq t} = O(\sqrt{t/N})$, and $\widetilde{\mathsf{cpO}}_{\mathsf{XYI}}\Pi_{\mathcal{R},\mathsf{I}} = \Pi_{\mathcal{R},\mathsf{I}}\widetilde{\mathsf{cpO}}_{\mathsf{XYI}}$, so the claimed Equation (324) follows. Now we can write

$$\sqrt{p_{\mathcal{A}}} = \left\|\Pi_{\mathcal{R},\mathsf{I}} A_{q,\mathsf{A}}\mathsf{cpO}_{\mathsf{AI}} \dots A_{1,\mathsf{A}}\mathsf{cpO}_{\mathsf{AI}}A_{0,\mathsf{A}} \left|0\right\rangle_{\mathsf{A}} \left|\perp\right\rangle_{\mathsf{A}}\right\| \tag{332}$$

$$= \left\|A_{q,\mathsf{A}}\mathsf{cpO}_{\mathsf{AI}}\Pi_{\mathcal{R},\mathsf{I}} A_{q-1,\mathsf{A}}\mathsf{cpO}_{\mathsf{AI}} \dots A_{1,\mathsf{A}}\mathsf{cpO}_{\mathsf{AI}}A_{0,\mathsf{A}} \left|0\right\rangle_{\mathsf{A}} \left|\perp\right\rangle_{\mathsf{A}}\right\| + O\left(\sqrt{\frac{q}{N}}\right) \tag{333}$$

$$\dots \tag{334}$$

$$= \left\|A_{q,\mathsf{A}}\mathsf{cpO}_{\mathsf{AI}} \dots A_{1,\mathsf{A}}\mathsf{cpO}_{\mathsf{AI}}A_{0,\mathsf{A}}\Pi_{\mathcal{R},\mathsf{I}} \left|0\right\rangle_{\mathsf{A}} \left|\perp\right\rangle_{\mathsf{A}}\right\| + \sum_{t=1}^{q} O\left(\sqrt{\frac{t}{N}}\right) \tag{335}$$

$$= O\left(\sqrt{\frac{q^3}{N}}\right). \tag{336}$$

□

7.2 A general search lower bound

The proof strategy developed for the cycle finding lower bounding can be applied much more broadly, to give a lower bound for any search problem. This bound is in terms of a quantity we refer to as the *sparsity* of the underlying predicate. Intuitively, this quantity measures the likelihood that each new query results in a solution being added to the database. Recalling that we allow forwards and inverse queries, we define sparsity in a two-sided way that allows both direction of query.

Definition 7.4. *Let $\mathcal{R} \subset ([N]^{\times 2})^*$ be a predicate, and $I \in \mathbf{I}_t$ such that $I \cap \mathcal{R} = \emptyset$. Then we define the I-sparsity s_I and t-sparsity s_t as*

$$s_I(\mathcal{R}) = \max\left\{ \max_{x \in [N]\setminus \mathsf{Dom}(I)} |\{y \notin \mathsf{Im}(I) \,:\, I[x \to y] \cap \mathcal{R} \neq \emptyset\}|, \right.$$
$$\left. \max_{y \in [N]\setminus \mathsf{Im}(I)} |\{x \notin \mathsf{Dom}(I) \,:\, I[x \to y] \cap \mathcal{R} \neq \emptyset\}| \right\}$$
$$s_t(\mathcal{R}) = \max_{I \in \mathbf{I}_{\leq t} \,:\, I \cap \mathcal{R} = \emptyset} s_I(\mathcal{R}). \tag{337}$$

Now, we can give a lower bound for any predicate search problem in terms of this quantity. Note that special cases of this theorem are known for predicates of constant size, meaning subsets of $([N]^{\times 2})^{O(1)}$, due to [MMW24; Coj+25].

Theorem 7.5. *Let $\mathcal{R} \subset ([N]^{\times 2})^*$ be a predicate, and $\mathcal{A}$ be a quantum algorithm making q bi-directional queries to a random permutation of size $[N]$ to output a list $\mathbf{x} = [(x_1, y_1), \dots, (x_l, y_l)]$ of input-output pairs. Then, $\mathcal{A}$ wins Game 1 under predicate $\mathcal{R}$ with probability at most*

$$p_{\mathcal{A}} = O\left(\mathsf{Dist}(q)^2 + \frac{l^2}{N} + \left(\sum_{t=0}^{q-1} \sqrt{\frac{s_t(\mathcal{R})}{N}}\right)^2\right). \tag{338}$$

Proof. We use the framework of Game 2. Let $\mathcal{A}$ act on Hillbert space $\mathcal{H}_{\mathsf{A}}$ containing query registers XY with intermediate operations $A_0, \ldots, A_q$ on this space. The success probability of this algorithm is given by

$$p_{\mathcal{A}} = \left\| \Pi_{\mathcal{R},\mathsf{I}} A_{q,\mathsf{A}} \mathsf{cpO}_{\mathsf{AI}} \ldots A_{1,\mathsf{A}} \mathsf{cpO}_{\mathsf{AI}} A_{0,\mathsf{A}} \left|0\right\rangle_{\mathsf{A}} \left|\perp\right\rangle_{\mathsf{A}} \right\|^2. \tag{339}$$

To bound this quantity, we will show a bound on the commutator norm

$$\left\| [\Pi_{\mathcal{R},\mathsf{I}}, \mathsf{cpO}_{\mathsf{AI}}] \right\|_{\leq t} = O\left(\sqrt{\frac{s_t(\mathcal{R})}{N}} \right). \tag{340}$$

To do this, we can modify the compression operator such that it preserves dissatisfaction of $\mathcal{R}$ perfectly. In particular, fix an $I \in \mathbf{I}_t$ and $x \notin \mathsf{Dom}(I)$. We define the set $\mathcal{R}^{\rightarrow}_{x,I}$ to be the set of images whose (forwards) assignment on x would result in the predicate being satisfied, or the empty set if the predicate is already satisfied. Formally,

$$\mathcal{R}^{\rightarrow}_{x,I} := \begin{cases} \{y \,:\, I[x \to y] \in \mathcal{R}\} & (\text{If } I \cap \mathcal{R} = \emptyset) \\ \emptyset & (\text{Otherwise}) \end{cases} \tag{341}$$

Observe that $|\mathcal{R}^{\rightarrow}_{x,I}| \leq s_I(\mathcal{R})$. Then we write

$$\left|\widetilde{+}^{\rightarrow}_{x,I}\right\rangle_{\mathsf{I}} := \frac{1}{\sqrt{N - |\mathcal{R}^{\rightarrow}_{x,I} \cup \mathsf{Im}(I)|}} \sum_{y \in [N] \setminus \mathsf{Im}(I) \setminus \mathcal{R}^{\rightarrow}_{x,I}} \left|I[x \to y]\right\rangle \tag{342}$$

$$\widetilde{\mathsf{pC}}^{\rightarrow}_{x,I} := \mathbb{1} - \left|I\right\rangle\!\left\langle I\right| - \left|\widetilde{+}^{\rightarrow}_{x,I}\right\rangle\!\left\langle \widetilde{+}^{\rightarrow}_{x,I}\right| + \left|\widetilde{+}^{\rightarrow}_{x,I}\right\rangle\!\left\langle I\right| + \left|I\right\rangle\!\left\langle \widetilde{+}^{\rightarrow}_{x,I}\right| \tag{343}$$

$$\widetilde{\mathsf{pC}}^{\rightarrow}_{\mathsf{XI}} = \bigoplus_{x \in [N], I \in \mathbf{I}^x} \left|x\right\rangle\!\left\langle x\right|_{\mathsf{X}} \otimes \widetilde{\mathsf{pC}}^{\rightarrow}_{x,I}. \tag{344}$$

In the other direction, we define the set $\mathcal{R}^{\leftarrow}_{y,I}$ to be the set of preimages whose (inverse) assignment on y would result in the predicate being satisfied, and otherwise the empty set. Formally,

$$\mathcal{R}^{\leftarrow}_{y,I} := \begin{cases} \{x \,:\, (I[y \to x])^{-1} \in \mathcal{R}\} & (\text{If } I^{-1} \cap \mathcal{R} = \emptyset) \\ \emptyset & (\text{Otherwise}) \end{cases} \tag{345}$$

Observe that $|\mathcal{R}^{\leftarrow}_{y,I}| \leq s_{I^{-1}}(\mathcal{R})$. Then we write

$$\left|\widetilde{+}^{\leftarrow}_{y,I}\right\rangle_{\mathsf{I}} := \frac{1}{\sqrt{N - |\mathcal{R}^{\leftarrow}_{y,I} \cup \mathsf{Im}(I)|}} \sum_{x \in [N] \setminus \mathsf{Im}(I) \setminus \mathcal{R}^{\leftarrow}_{y,I}} \left|I[y \to x]\right\rangle \tag{346}$$

$$\widetilde{\mathsf{pC}}^{\leftarrow}_{y,I} := \mathbb{1} - \left|I\right\rangle\!\left\langle I\right| - \left|\widetilde{+}^{\leftarrow}_{y,I}\right\rangle\!\left\langle \widetilde{+}^{\leftarrow}_{y,I}\right| + \left|\widetilde{+}^{\leftarrow}_{y,I}\right\rangle\!\left\langle I\right| + \left|I\right\rangle\!\left\langle \widetilde{+}^{\leftarrow}_{y,I}\right| \tag{347}$$

$$\widetilde{\mathsf{pC}}^{\leftarrow}_{\mathsf{XI}} = \bigoplus_{y \in [N], I \in \mathbf{I}^y} \left|y\right\rangle\!\left\langle y\right|_{\mathsf{X}} \otimes \widetilde{\mathsf{pC}}^{\leftarrow}_{y,I}. \tag{348}$$

Observe that $\left\||\widetilde{+}^{\rightarrow}_{x,I}\rangle - |+_{x,I}\rangle\right\| = O(\sqrt{|\mathcal{R}^{\rightarrow}_{x,I}|/N})$, and $\left\||\widetilde{+}^{\leftarrow}_{y,I}\rangle - |+_{y,I}\rangle\right\| = O(\sqrt{|\mathcal{R}^{\leftarrow}_{y,I}|/N})$. It follows that $\left\|\mathsf{pC} - \widetilde{\mathsf{pC}}^{\rightarrow}\right\|_{\leq t} = O(\sqrt{s_t(\mathcal{R})/N})$ and $\left\|\mathsf{pC} - \widetilde{\mathsf{pC}}^{\leftarrow}\right\|_{\leq t} = O(\sqrt{s_t(\mathcal{R})/N})$ by Lemma 2.3 and the prior observation. Now we have

$$[\Pi_{\mathcal{R},\mathsf{I}}, \widetilde{\mathsf{pC}}^{\rightarrow}_{\mathsf{XI}}] = 0, \qquad [\Pi_{\mathcal{R},\mathsf{I}}, \mathsf{F}_{\mathsf{I}} \cdot \widetilde{\mathsf{pC}}^{\leftarrow}_{\mathsf{XI}} \cdot \mathsf{F}^{\dagger}_{\mathsf{I}}] = 0, \tag{349}$$

$$[\Pi_{\mathcal{R},\mathsf{I}}, \mathsf{P}_{\mathsf{XYI}}] = 0, \qquad [\Pi_{\mathcal{R},\mathsf{I}}, \mathsf{F}_{\mathsf{I}} \cdot \mathsf{P}_{\mathsf{XYI}} \cdot \mathsf{F}^{\dagger}_{\mathsf{I}}] = 0, \tag{350}$$

as the modified operators cannot add any input-output pair which causes $\mathcal{R}$ to be satisfied. If we write

$$\begin{aligned}\widetilde{\mathsf{cpO}}_{XYI} &= |0\rangle\langle 0|_{\mathsf{X}} \otimes (\widetilde{\mathsf{pC}}^{\rightarrow}_{\mathsf{XI}} \cdot \mathsf{P}_{\mathsf{XYI}} \cdot \widetilde{\mathsf{pC}}^{\rightarrow,\dagger}_{\mathsf{XI}}) + |1\rangle\langle 1|_{\mathsf{X}} \otimes (\mathsf{F}_{\mathsf{I}} \cdot \widetilde{\mathsf{pC}}^{\leftarrow}_{\mathsf{XI}} \cdot \mathsf{P}_{\mathsf{XYI}} \cdot \widetilde{\mathsf{pC}}^{\leftarrow,\dagger}_{\mathsf{XI}} \cdot \mathsf{F}^{\dagger}_{\mathsf{I}}) \qquad (351)\\ &= |0\rangle\langle 0|_{\mathsf{X}} \otimes (\widetilde{\mathsf{pC}}^{\rightarrow}_{\mathsf{XI}} \cdot \mathsf{P}_{\mathsf{XYI}} \cdot \widetilde{\mathsf{pC}}^{\rightarrow,\dagger}_{\mathsf{XI}}) + |1\rangle\langle 1|_{\mathsf{X}} \otimes (\mathsf{F}_{\mathsf{I}} \cdot \widetilde{\mathsf{pC}}^{\leftarrow}_{\mathsf{XI}} \cdot (\mathsf{F}^{\dagger}\mathsf{F})_{\mathsf{I}} \cdot \mathsf{P}_{\mathsf{XYI}} \cdot (\mathsf{F}^{\dagger}\mathsf{F})_{\mathsf{I}} \cdot \widetilde{\mathsf{pC}}^{\leftarrow,\dagger}_{\mathsf{XI}} \cdot \mathsf{F}^{\dagger}_{\mathsf{I}}),\end{aligned}$$

then we have $\left\|\mathsf{cpO} - \widetilde{\mathsf{cpO}}\right\|_{\leq t} = O(\sqrt{s_t(\mathcal{R})/N})$, and $\widetilde{\mathsf{cpO}}_{\mathsf{XYI}}\Pi_{\mathcal{R},\mathsf{I}} = \Pi_{\mathcal{R},\mathsf{I}}\widetilde{\mathsf{cpO}}_{\mathsf{XYI}}$, so the claimed Equation (340) follows. Now we can write

$$\sqrt{p_{\mathcal{A}}} = \left\|\Pi_{\mathcal{R},\mathsf{I}} A_{q,\mathsf{A}} \mathsf{cpO}_{\mathsf{AI}} \dots A_{1,\mathsf{A}} \mathsf{cpO}_{\mathsf{AI}} A_{0,\mathsf{A}} |0\rangle_{\mathsf{A}} |\perp\rangle_{\mathsf{A}}\right\| \tag{352}$$

$$= \left\|A_{q,\mathsf{A}} \mathsf{cpO}_{\mathsf{AI}} \Pi_{\mathcal{R},\mathsf{I}} A_{q-1,\mathsf{A}} \mathsf{cpO}_{\mathsf{AI}} \dots A_{1,\mathsf{A}} \mathsf{cpO}_{\mathsf{AI}} A_{0,\mathsf{A}} |0\rangle_{\mathsf{A}} |\perp\rangle_{\mathsf{A}}\right\| + O\left(\sqrt{\frac{s_{q-1}(\mathcal{R})}{N}}\right) \tag{353}$$

$$\dots \tag{354}$$

$$= \left\|A_{q,\mathsf{A}} \mathsf{cpO}_{\mathsf{AI}} \dots A_{1,\mathsf{A}} \mathsf{cpO}_{\mathsf{AI}} A_{0,\mathsf{A}} \Pi_{\mathcal{R},\mathsf{I}} |0\rangle_{\mathsf{A}} |\perp\rangle_{\mathsf{A}}\right\| + \sum_{t=0}^{q-1} O\left(\sqrt{\frac{s_t(\mathcal{R})}{N}}\right) \tag{355}$$

$$= \sum_{t=0}^{q-1} O\left(\sqrt{\frac{s_t(\mathcal{R})}{N}}\right). \tag{356}$$

This, combined with Lemma 7.1 implies the claim. □

One of the simplest examples from the literature of a predicate search problem is the double-sided zero search problem, due to Unruh [Unr21; Unr23].

Problem 3. *Sample $\varphi \sim \mathbf{S}_N$ for $N = 2^{2n}$, to produce two way accessible oracle $\mathcal{O}_\varphi$. $\mathcal{A}$ makes q queries to $\mathcal{O}_\varphi$, and succeeds if it finds $(x\|0^n, y\|0^n)$ such that $\varphi(x\|0^n) = y\|0^n)$.*

This is a search problem with predicate $\mathcal{R} = \{[(x\|0^n, y\|0^n)] : x, y \in \{0,1\}^n\}$. It is easy to compute $s_t(\mathcal{R}) = O(2^n)$ for $t = o(2^n)$, giving us the following bound.

Corollary 7.6 (of Theorem 7.5)**.** *Any algorithm $\mathcal{A}$ making q quantum queries solves Problem 3 with probability at most*

$$p_{\mathcal{A}} = O\left(\frac{q^2}{2^n} + \mathsf{Dist}(q)^2\right). \tag{357}$$

Lower bounds for this problem are known already [CP24; MMW24; Coj+25], even tight ones. Observe that our bound would be tight, ignoring that $\mathsf{Dist}(q)^2$ term. The problems we will look at next are less well understood, and for many no tight bound is known. The bounds we show as applications of our general theorem will in some cases improve on known results and/or be tight.

Even more intriguingly, all of them would be tight if the $\mathsf{Dist}(q)^2$ term were ignored. This suggests that our construction accurately captures the difficulty of these problems, and improving our soundness analysis provides a direct path towards showing tight quantum query lower bounds for many problems in the random permutation model.

7.3 Application to sponge

The sponge construction lifts an invertible permutation to a hash function with arbitrary domain. It is parameterized by integers $r, c \in \mathbb{N}$ and the permutation $\varphi : \{0,1\}^{r+c} \to \{0,1\}^{r+c}$. To hash a message $m \in (\{0,1\}^r)^*$ of the form $m = m_1 \| m_2 \| \dots \| m_l$, we initialize the internal state to $u_0 = 0^{r+c}$. Then, for every i from 1 up to l, we XOR m_i into the first r bits of u and apply φ, obtaining $u_{i+1} = \varphi(u_i \oplus (m_i \| 0^c))$. The output is then the first r bits of the final state u_l.[12] This construction is due to Bertoni et al [Ber+07; Ber+08], and when instantiated with the Keccak permutation underlies the NIST standardized hash function SHA3 [Ber+11b; Ber+11a].

It is known that this construction is quantum indifferentiable from a random oracle [Ala+25], which in particular implies preimage and collision resistance. The best known bound for these is due to the same work, which shows $\tilde{\Omega}(2^{\min(r,c)/5})$ queries are required for constant success probability at either collision or preimage finding. This result required a long and detailed proof tailored to this specific construction, and is not tight. Using our framework, we can improve these bounds in a direct manner.

We will use the notation $x[a:b]$ to denote the a-th (inclusive) through b-th (exclusive) bits of string x with zero based indexing. Then, we can define the preimage predicate of the sponge for image $w \in \{0,1\}^r$ as

$$\mathcal{R}^{w-\mathsf{pre}} := \left\{ [(x_1, y_1), \dots, (x_l, y_l)] \ : \ \begin{array}{c} x_1[r:r+c]=0^c, \text{ and} \\ y_i[r:r+c]=x_{i+1}[r:r+c], \text{ and} \\ y_l[0:r]=z. \end{array} \right\} \tag{358}$$

Problem 4 (Sponge inversion.)**.** *Given quantum query access to an ideal random permutation $\mathcal{O}_\varphi$, find a sponge pre-image of some fixed $w \in \{0,1\}^r$. Equivalently, solve the search problem defined by predicate $\mathcal{R}^{w-\mathsf{pre}}$.*

Theorem 7.7. *A q query quantum algorithm $\mathcal{A}$ solves sponge inversion, Problem 4, with probability no more than*

$$p_{\mathcal{A}} = O\left(\frac{q^2}{2^r} + \frac{q^3}{2^c} + \mathsf{Dist}(q)^2\right) \tag{359}$$

[12] For simplicity, we consider only one round of squeezing.

Proof. To show this bound, we will change the predicate slightly. In particular, define the internal z-preimage predicate as

$$\mathcal{R}^{z-\mathsf{ipre}} := \left\{ [(x_1, y_1), \ldots, (x_l, y_l)] \ : \begin{array}{c} x_1[r{:}r+c]=0^c, \text{ and} \\ y_i[r{:}r+c]=x_{i+1}[r{:}r+c], \text{ and} \\ y_l[r{:}r+c]=z. \end{array} \right\}, \tag{360}$$

which can be used to construct the internal collision predicate

$$\mathcal{R}^{\mathsf{icol}} := \mathcal{R}^{0^c-\mathsf{ipre}} \cup \bigcup_{z \in \{0,1\}^c \setminus 0^c} \mathcal{R}^{z-\mathsf{ipre}} \times \mathcal{R}^{z-\mathsf{ipre}}. \tag{361}$$

We then let $\mathcal{R}' = \mathcal{R}^{\mathsf{icol}} \cup \mathcal{R}^{w-\mathsf{pre}}$; in other words, we allow the adversary to find either a pre-image of w or an internal collision. Clearly, this task is no harder than finding a sponge preimage. It is known that $s_t(\mathcal{R}') = O\left(2^c + t \cdot 2^r\right)$ [LM22], from which the claim follows. □

We can also define the sponge collision predicate,

$$\mathcal{R}^{\mathsf{col}} := \bigcup_{w \in \{0,1\}^r} \mathcal{R}^{w-\mathsf{pre}} \times \mathcal{R}^{w-\mathsf{pre}}. \tag{362}$$

This gives rise to the following search task.

Problem 5 (Sponge collision.)**.** *Given quantum query access to a random permutation oracle $\mathcal{O}_\varphi$, find a sponge collision. Equivalently, solve the search problem defined by predicate $\mathcal{R}^{\mathsf{col}}$.*

Theorem 7.8. *A q query quantum algorithm $\mathcal{A}$ solves sponge collision, Problem 4, with probability no more than*

$$p_{\mathcal{A}} = O\left(\frac{q^3}{2^{\min(r,c)}} + \mathsf{Dist}(q)^2\right) \tag{363}$$

Proof. To show this bound, we will change the predicate slightly. In particular, we once again allow the adversary to find either a sponge collision or an internal collision, letting $\mathcal{R}' = \mathcal{R}^{\mathsf{icol}} \cup \mathcal{R}^{\mathsf{col}}$. It is known that $s_t(\mathcal{R}') = O\left(t \cdot 2^r + t \cdot 2^c\right)$ [Unr21], from which the claim follows. □

Observe that using quantum collision and claw finding [BHT97], there is an algorithm solving Problem 5 with probability $O(q^3 2^{-\min(r,c)})$. It follows that Theorem 7.8 is tight whenever $\min(r,c)/3 \leq (r+c)/12$, and improves on the prior art [Ala+25] when $\min(r,c)/5 \leq (r+c)/12$. Furthermore, ignoring the $\mathsf{Dist}(q)^2$ term the bound would be tight for any r and c.

Similarly using Grover search [Gro96] or claw finding, there is an algorithm solving Problem 4 with probability $O(q^3/2^r + q^2/2^c)$. Then Theorem 7.7 is tight whenever $\min(r/2, c/3) \leq (r+c)/12$, and improves on the prior art [Ala+25] when $\min(r,c)/5 \leq (r+c)/12$. Again, ignoring the $\mathsf{Dist}(q)^2$ term the bound would be tight for any r and c.

7.4 Application to Davies-Meyer

Davies-Meyer is a construction that turns an ideal cipher into a fixed length compressing hash function [Win84]. It underlies SHA1, SHA2, and MD5. A cipher is a keyed permutation, and the quantum ideal cipher model allows quantum query access across both the key and the resulting permutation. We will state our results for the key-less case for simplicity.[13]

The Davies-Meyer function $\mathsf{DM} : \{0,1\}^n \to \{0,1\}^n$ is defined by a permutation $\varphi : \{0,1\}^n \to \{0,1\}^n$, and is

$$\mathsf{DM}^{\varphi}(x) := \varphi(x) \oplus x. \tag{364}$$

Problem 6 (Davies-Meyer inversion.)**.** *Given bidirectional quantum query access to a random permutation $\varphi : \{0,1\}^n \to \{0,1\}^n$, find an $x \in \{0,1\}^n$ such that $\mathsf{DM}^{\varphi}(x) = 0^n$.*

Theorem 7.9. *Any q query quantum algorithm $\mathcal{A}$ solves Davies-Meyer inversion, Problem 6, with probability bounded by*

$$p_{\mathcal{A}} = O\left(\frac{q^2}{2^n} + \mathsf{Dist}(q)^2\right). \tag{365}$$

Proof. We can define the zero preimage predicate $\mathcal{R} = \{[(x,x)] : x \in \{0,1\}^n\}$, which satisfies $s_t(\mathcal{R}) = 1$. Then the claim follows from Theorem 7.5. □

Problem 7 (Davies-Meyer collision.)**.** *Given bidirectional quantum query access to a random permutation $\varphi : \{0,1\}^n \to \{0,1\}^n$, find a pair $x \neq x' \in \{0,1\}^n$ such that $\mathsf{DM}^{\varphi}(x) = \mathsf{DM}^{\varphi}(x')$.*

Theorem 7.10. *Any q query quantum algorithm $\mathcal{A}$ solves Davies-Meyer collision, Problem 7, with probability bounded by*

$$p_{\mathcal{A}} = O\left(\frac{q^3}{2^n} + \mathsf{Dist}(q)^2\right). \tag{366}$$

Proof. We can define the collision predicate

$$\mathcal{R} = \left\{[(x,y),(x',y')] : \begin{array}{c} x,y,x',y' \in \{0,1\}^n, \text{ and} \\ x \oplus y = x' \oplus y', \text{ and} \\ x \neq x', y \neq y'. \end{array}\right\}, \tag{367}$$

which satisfies $s_t(\mathcal{R}) = t$. Then the claim follows from Theorem 7.5. □

[13] Our framework can be applied to the ideal cipher model, though we leave a detailed exploration to future work.

Note that [HY18; Coj+25] prove bounds for preimage and collision resistance of the Davies-Meyer construction, and these bounds are tighter than the ones given here. However, no tight bound is known for collision resistance: the best known is $\Omega(2^{n/4})$ for constant success probability due to [Coj+25].

Ignoring the $\mathsf{Dist}(q)^2$ term, our bound would be $\Omega(2^{n/3})$, matching quantum collision finding. Tightening our framework then provides a path towards deriving a tight collision bound for the Davies-Meyer construction.

References

[Nat77] U.S. Department of Commerce National Institute of Standards. *Data Encryption Standard (DES)*. Tech. rep. Federal Information Processing Standards Publication 46. Washington, D.C.: National Bureau of Standards, 1977. URL: https://csrc.nist.gov/pubs/fips/46/final.

[Win84] Robert S. Winternitz. "A Secure One-Way Hash Function Built from DES". In: *1984 IEEE Symposium on Security and Privacy*. 1984, pp. 88–88. DOI: 10.1109/SP.1984.10027.

[LR88] Michael Luby and Charles Rackoff. "How to Construct Pseudorandom Permutations from Pseudorandom Functions". In: *SIAM Journal on Computing* 17.2 (1988), pp. 373–386. DOI: 10.1137/0217022. eprint: https://doi.org/10.1137/0217022. URL: https://doi.org/10.1137/0217022.

[Sch93] Bruce Schneier. "Description of a New Variable-Length Key, 64-bit Block Cipher (Blowfish)". In: *Fast Software Encryption, Cambridge Security Workshop, Cambridge, UK, December 9-11, 1993, Proceedings*. Vol. 809. Lecture Notes in Computer Science. Springer, 1993, pp. 191–204. DOI: 10.1007/3-540-58108-1_24.

[Gro96] Lov K. Grover. "A fast quantum mechanical algorithm for database search". In: *Proceedings of the Twenty-Eighth Annual Symposium on the Theory of Computing*. New York: ACM Press, 1996, pp. 212–219.

[BHT97] Gilles Brassard, Peter Høyer, and Alain Tapp. *Quantum Algorithm for the Collision Problem*. Avalable at https://arxiv.org/abs/quant-ph/9705002. 1997.

[Ber+07] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles van Assche. "Sponge functions". In: *ECRYPT Hash Workhsop*. 2007.

[Ber+08] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. "On the Indifferentiability of the Sponge Construction". In: *Advances in Cryptology – EUROCRYPT 2008*. Ed. by Nigel Smart. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, pp. 181–197. ISBN: 978-3-540-78967-3.

[KM10] Hidenori Kuwakado and Masakatu Morii. "Quantum distinguisher between the 3-round Feistel cipher and the random permutation". In: *Proc. IEEE International Symposium on Information Theory*. IEEE, 2010, pp. 2682–2685.

[Ber+11a] G. Bertoni, J. Daemen, M. Peeters, and G. Van Assche. *Cryptographic sponge functions*. Submission to NIST (Round 3). 2011. URL: http://sponge.noekeon.org/CSF-0.1.pdf.

[Ber+11b] G. Bertoni, J. Daemen, M. Peeters, and G. Van Assche. *The Keccak SHA-3 submission*. Submission to NIST (Round 3). 2011. URL: http://keccak.noekeon.org/Keccak-submission-3.pdf.

[Bon+11] Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. "Random Oracles in a Quantum World". In: *Advances in Cryptology – ASIACRYPT 2011*. Ed. by Dong Hoon Lee and Xiaoyun Wang. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 41–69. ISBN: 978-3-642-25385-0.

[Zha12] Mark Zhandry. "How to Construct Quantum Random Functions". In: *Proceedings of the 53rd Annual Symposium on Foundations of Computer Science*. FOCS '12. Washington, DC, USA: IEEE Computer Society, 2012, pp. 679–687. ISBN: 978-0-7695-4874-6. DOI: 10.1109/FOCS.2012.37.

[Yue13] Henry Yuen. *A quantum lower bound for distinguishing random functions from random permutations*. 2013. arXiv: 1310.2885 [cs.CC]. URL: https://arxiv.org/abs/1310.2885.

[Unr14] Dominique Unruh. "Revocable Quantum Timed-Release Encryption". In: *Advances in Cryptology – EUROCRYPT 2014*. Ed. by Phong Q. Nguyen and Elisabeth Oswald. Berlin, Heidelberg: Springer Berlin Heidelberg, 2014, pp. 129–146. ISBN: 978-3-642-55220-5.

[Zha15] Mark Zhandry. "A Note on the Quantum Collision and Set Equality Problems". In: *Quantum Info. Comput.* 15.7-8 (May 2015), pp. 557–567.

[Zha16] Mark Zhandry. *A Note on Quantum-Secure PRPs*. 2016. arXiv: 1611.05564 [cs.CR]. URL: https://arxiv.org/abs/1611.05564.

[Cza+17] Jan Czajkowski, Leon Groot Bruinderink, Andreas Hülsing, Christian Schaffner, and Dominique Unruh. *Post-quantum security of the sponge construction*. Cryptology ePrint Archive, Paper 2017/771. https://eprint.iacr.org/2017/771. 2017. URL: https://eprint.iacr.org/2017/771.

[HY18] Akinori Hosoyamada and Kan Yasuda. *Building Quantum-One-Way Functions from Block Ciphers: Davies-Meyer and Merkle-Damgård Constructions*. Cryptology ePrint Archive, Paper 2018/841. 2018. URL: https://eprint.iacr.org/2018/841.

[Cza+19] Jan Czajkowski, Christian Majenz, Christian Schaffner, and Sebastian Zur. *Quantum Lazy Sampling and Game-Playing Proofs for Quantum Indifferentiability*. Cryptology ePrint Archive, Paper 2019/428. https://eprint.iacr.org/2019/428. 2019. URL: https://eprint.iacr.org/2019/428.

[HI19] Akinori Hosoyamada and Tetsu Iwata. "4-Round Luby-Rackoff Construction is a qPRP". In: *Advances in Cryptology – ASIACRYPT 2019: 25th International Conference on the Theory and Application of Cryptology and Information Security, Kobe, Japan, December 8–12, 2019, Proceedings, Part I*. Kobe, Japan: Springer-Verlag, 2019, pp. 145–174. ISBN: 978-3-030-34577-8. DOI: 10.1007/978-3-030-34578-5_6. URL: https://doi.org/10.1007/978-3-030-34578-5_6.

[Ito+19] Gembu Ito, Akinori Hosoyamada, Ryutaroh Matsumoto, Yu Sasaki, and Tetsu Iwata. "Quantum Chosen-Ciphertext Attacks Against Feistel Ciphers". In: *Topics in Cryptology – CT-RSA 2019: The Cryptographers' Track at the RSA Conference 2019, San Francisco, CA, USA, March 4–8, 2019, Proceedings*. San Francisco, CA, USA: Springer-Verlag, 2019, pp. 391–411. ISBN: 978-3-030-12611-7. DOI: 10.1007/978-3-030-12612-4_20. URL: https://doi.org/10.1007/978-3-030-12612-4_20.

[LZ19] Qipeng Liu and Mark Zhandry. "On Finding Quantum Multi-collisions". In: *Advances in Cryptology – EUROCRYPT 2019*. Ed. by Yuval Ishai and Vincent Rijmen. Cham: Springer International Publishing, 2019, pp. 189–218. ISBN: 978-3-030-17659-4.

[Zha19] Mark Zhandry. "How to Record Quantum Queries, and Applications to Quantum Indifferentiability". In: *Advances in Cryptology – CRYPTO 2019*. Ed. by Alexandra Boldyreva and Daniele Micciancio. Cham: Springer International Publishing, 2019, pp. 239–268. ISBN: 978-3-030-26951-7.

[Chu+21] 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: *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer. 2021, pp. 598–629.

[Don+21] Jelle Don, Serge Fehr, Christian Majenz, and Christian Schaffner. *Online-Extractability in the Quantum Random-Oracle Model*. Cryptology ePrint Archive, Paper 2021/280. 2021. URL: https://eprint.iacr.org/2021/280.

[Unr21] Dominique Unruh. *Compressed Permutation Oracles (And the Collision-Resistance of Sponge/SHA3)*. Cryptology ePrint Archive, Paper 2021/062. https://eprint.iacr.org/2021/062. 2021. URL: https://eprint.iacr.org/2021/062.

[YZ21] Takashi Yamakawa and Mark Zhandry. "Classical vs quantum random oracles". In: *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer. 2021, pp. 568–597.

[Zha21] Mark Zhandry. "Redeeming Reset Indifferentiability and Applications to Post-quantum Security". In: *Advances in Cryptology – ASIACRYPT 2021*. Ed. by Mehdi Tibouchi and Huaxiong Wang. Cham: Springer International Publishing, 2021, pp. 518–548. ISBN: 978-3-030-92062-3.

[LM22] Charlotte Lefevre and Bart Mennink. "Tight Preimage Resistance of the Sponge Construction". In: *Advances in Cryptology – CRYPTO 2022*. Ed. by Yevgeniy Dodis and Thomas Shrimpton. Cham: Springer Nature Switzerland, 2022, pp. 185–204. ISBN: 978-3-031-15985-5.

[Ros22] Ansis Rosmanis. *Tight Bounds for Inverting Permutations via Compressed Oracle Arguments*. 2022. arXiv: 2103.08975 [quant-ph].

[Ala+23] Gorjan Alagic, Chen Bai, Alexander Poremba, and Kaiyan Shi. *On the Two-sided Permutation Inversion Problem*. 2023. arXiv: 2306.13729 [quant-ph].

[HM23] Yassine Hamoudi and Frédéric Magniez. "Quantum Time–Space Tradeoff for Finding Multiple Collision Pairs". In: *ACM Transactions on Computation Theory* 15.1–2 (June 2023), pp. 1–22. ISSN: 1942-3462. DOI: 10.1145/3589986. URL: http://dx.doi.org/10.1145/3589986.

[Unr23] Dominique Unruh. "Towards Compressed Permutation Oracles". In: *Advances in Cryptology – ASIACRYPT 2023: 29th International Conference on the Theory and Application of Cryptology and Information Security, Guangzhou, China, December 4–8, 2023, Proceedings, Part IV*. Guangzhou, China: Springer-Verlag, 2023, pp. 369–400. ISBN: 978-981-99-8729-0. DOI: 10.1007/978-981-99-8730-6_12. URL: https://doi.org/10.1007/978-981-99-8730-6_12.

[Bha+24] Ritam Bhaumik, Benoît Cogliati, Jordan Ethan, and Ashwin Jha. "Mind the Bad Norms: Revisiting Compressed Oracle-Based Quantum Indistinguishability Proofs". In: *Advances in Cryptology – ASIACRYPT 2024: 30th International Conference on the Theory and Application of Cryptology and Information Security, Kolkata, India, December 9–13, 2024, Proceedings, Part IX*. Kolkata, India: Springer-Verlag, 2024, pp. 215–247. DOI: 10.1007/978-981-96-0947-5_8. URL: https://doi.org/10.1007/978-981-96-0947-5_8.

[CP24] Joseph Carolan and Alexander Poremba. "Quantum One-Wayness of the Single-Round Sponge with Invertible Permutations". In: *Advances in Cryptology – CRYPTO 2024*. Ed. by Leonid Reyzin and Douglas Stebila. Cham: Springer Nature Switzerland, 2024, pp. 218–252. ISBN: 978-3-031-68391-6.

[CPZ24] Joseph Carolan, Alexander Poremba, and Mark Zhandry. *(Quantum) Indifferentiability and Pre-Computation*. Cryptology ePrint Archive, Paper 2024/1727. 2024. URL: https://eprint.iacr.org/2024/1727.

[MMW24] Christian Majenz, Giulio Malavolta, and Michael Walter. *Permutation Superposition Oracles for Quantum Query Lower Bounds*. Cryptology ePrint Archive, Paper 2024/1140. 2024. URL: https://eprint.iacr.org/2024/1140.

[Ala+25] Gorjan Alagic, Joseph Carolan, Christian Majenz, and Saliha Tokat. *The Sponge is Quantum Indifferentiable*. Cryptology ePrint Archive, Paper 2025/731. 2025. URL: https://eprint.iacr.org/2025/731.

[Coj+25] Alexandru Cojocaru, Minki Hhan, Qipeng Liu, Takashi Yamakawa, and Aaram Yun. *Quantum Lifting for Invertible Permutations and Ideal Ciphers*. Cryptology ePrint Archive, Paper 2025/738. 2025. URL: https://eprint.iacr.org/2025/738.