

Improved Soundness for Compressed Permutation Oracles and Tight Quantum Preimage and Collision Bounds for the Sponge

Sunyeop Kim^{1,2}

¹ Korea University

`sunyeopkim@korea.ac.kr`

² Nanyang Technological University

Abstract. We give a sharper bound on the error introduced by replacing a random permutation oracle with Carolan’s compressed permutation oracle. Building on Rosmanis’s representation-theoretic approach, we represent exact permutation states in the compressed oracle’s database space, allowing a direct comparison of the two query operations. For a uniform permutation on N points, this comparison bounds the soundness error by $4q/\sqrt{N}$ after $q \leq N/4$ queries. As the main application, we obtain tight query complexity for constant success probability in sponge preimage and collision search by substituting this bound into Carolan’s search bounds. The same analysis improves bounds for the one-more problem and cycle finding, and gives tight query complexity for constant success probability in keyless Davies–Meyer collision search.

Keywords: Post-quantum cryptography · Quantum random permutation model · Compressed oracle · Quantum query complexity

1 Introduction

A compressed oracle records a quantum algorithm’s queries in a superposition of small databases. For random functions, this representation turns many search lower bounds into a count of the database entries that complete a solution [12]. A permutation imposes an additional constraint: distinct inputs have distinct outputs, and an algorithm may query in either direction. Carolan’s compressed permutation oracle maintains these constraints using databases represented by partial injective functions [2].

Carolan proves soundness of his compressed permutation oracle with error $O(q^3/N^{1/4})$ after q queries on a domain of size N . His proof bounds the Euclidean distance between purifications of the two experiments. Transferring the database analysis to a search bound introduces the square of this soundness error, $O(q^6/\sqrt{N})$. We improve the soundness bound to $O(q/\sqrt{N})$ by embedding permutation states isometrically into the database space and comparing the exact and compressed oracle operations using input and output relabeling symmetries.

1.1 Main results

We use $\text{Dist}(q)$ to measure how far the random-permutation and compressed-oracle experiments can differ after at most q queries. It is defined through purifications of the adversary’s final states in section 2.4. We prove

$$\text{Dist}(q) \leq \frac{4q}{\sqrt{N}}, \quad q \leq N/4.$$

The proof constructs an isometric embedding W_N into the database space. After applying W_N to the final purification in the random-permutation experiment, we bound its distance from the final purification in Carolan’s compressed-oracle experiment (theorem 1). Tracing out the purifying registers also bounds the distinguishing advantage by $4q/\sqrt{N}$. A Lean 4 formalization of theorem 1 is available at <https://github.com/SUNYEOP/improved-soundness-compressed-permutation-oracle-lean>.

We substitute this upper bound in Carolan’s search bounds (sections 4 and 5). Table 1 summarizes the resulting bounds and marks new upper bounds and tight query complexities.

Consider the sponge with rate r , capacity c , one output block, and a fixed injective padding map to nonempty block sequences. For an adversary making q queries and returning messages of total padded length at most L blocks, theorem 4 bounds the success probability by $O((q + L)^2/2^r + (q + L)^3/2^c)$ for preimage search and $O((q + L)^3/2^{\min(r,c)})$ for collision search. For the keyless Davies–Meyer construction $\text{DM}^\varphi(x) = \varphi(x) \oplus x$ on n -bit strings, theorem 5 gives collision probability $O((q + 1)^3/2^n)$ and tight $\Theta(2^{n/3})$ query complexity for constant success probability.

1.2 Proof overview

We first place the exact permutation state in the database space without changing inner products. A partial database determines a uniform superposition of consistent permutations, but these superpositions can overlap even when the databases are distinct. Mapping them directly to orthogonal database records would therefore change the experiment. Decomposing by relabeling symmetry and normalizing each component gives the required isometric embedding. It sends the uniform permutation state to the empty database.

The key observation is that fixing $\varphi(x) = y$ leaves a permutation between $N - 1$ inputs and $N - 1$ outputs. We can therefore reuse the same embedding on the remaining labels and then adjoin (x, y) . All possible responses are retained in superposition. This gives an operation to compare with Carolan’s permutation compression at x , which instead swaps a database state undefined at x with the uniform superposition of its extensions at x .

Both operations respect input relabelings fixing x and arbitrary output relabelings. This symmetry reduces their comparison to one overlap per component, but does not determine how close that overlap is to one. The main calculation in lemma 3 expresses these overlaps through the normalization factors for the

Table 1. Bounds obtained here for two-sided permutation queries. The search rows give success-probability upper bounds. For sponge, L bounds the total number of blocks in the padded messages. Tightness concerns query complexity for constant success probability, measured in $q + L$ for sponge with identity padding.

Problem	Bound	New	Tight query complexity
Soundness	$4q/\sqrt{N}$	✓	—
Sponge preimage (one output block)	$O\left(\frac{(q+L)^2}{2^r} + \frac{(q+L)^3}{2^c}\right)$	✓	✓
Sponge collision (one output block)	$O\left(\frac{(q+L)^3}{2^{\min(r,c)}}\right)$	✓	✓
One-more problem	$O((q+1)^2/N)$	✓	—
Cycle finding	$O((q+1)^3/N)$	✓	—
Double-sided zero	$O((q+1)^2/2^n)$	—	✓
Keyless Davies–Meyer inversion	$O((q+1)^2/2^n)$	—	✓
Keyless Davies–Meyer collision	$O((q+1)^3/2^n)$	✓	✓

full and remaining permutations. The resulting one-query error is at most $4/\sqrt{N}$ along an exact q -query computation with $q \leq N/4$. A hybrid argument replaces the queries one at a time, applying the one-query bound to states generated in the random-permutation experiment. Summing these errors gives $4q/\sqrt{N}$.

The applications use Carolan’s sparsity argument [2, proof of Thm. 7.5]. For a database of size at most t containing no solution, s_t bounds the number of new outputs for a fixed input, or new inputs for a fixed output, that cause the database to contain a solution. The increase per query in the amplitude on databases containing a solution is $O(\sqrt{s_t/N})$. We combine this bound with the new soundness theorem in Carolan’s predicate-search reduction.

1.3 Related work

Unruh proposed a compressed permutation oracle with forward and inverse access, but left its soundness as a conjecture [11]. Carolan modified the construction to maintain database injectivity unitarily and proved a quantitative soundness bound [2]. We retain Carolan’s oracle and predicate-search reductions, and improve its soundness bound. Our proof uses Rosmanis’s oracle subspaces $\mathcal{A}_t$, their irreducible decomposition, and the eigenvalues he computed to prove tight forward-query inversion bounds [8]. The new step is the overlap calculation between conditioning on the oracle response and Carolan’s permutation compression.

Exact compressed oracles of Grinko and Yoshida [6] and exact path-recording oracles of Foxman, Lombardi, Ma, Nehoran, and Wright [5] use different representations from Carolan’s database.³ For permutations, Grinko and Yoshida use a Fourier basis; their Appendix E.1 distinguishes nonorthogonal consistent-permutation states from Carolan’s orthonormal database records. Exactness in these representations does not by itself give a distance bound for Carolan’s oracle; using his sparsity analysis would require a comparison with his permutation compression.

Single-round sponge one-wayness and preimage resistance were established in [3,7]. Cojocaru et al. apply a general lifting theorem to multi-round sponge search [4]; in their bounds, the exponent in the number of queries grows with the number of permutation calls per sponge evaluation. Alagic et al. prove quantum indifferentiability of the sponge and give preimage and collision bounds for variable-length messages [1]. Our improved soundness bound sharpens Carolan’s preimage and collision bounds for the sponge with one output block.

Carolan analyzes inversion and collision search for the keyless Davies–Meyer construction [2]. The known tight inversion bound follows from the generalized double-sided search theorem of Cojocaru et al. [4]. Substituting our improved soundness bound into Carolan’s analysis recovers this inversion bound and gives tight query complexity for collision search.

2 Oracle models and notation

For a positive integer N , write $[N] = \{1, \dots, N\}$ and let $\mathbf{S}_N$ be the symmetric group. For a finite set S , define

$$\mathcal{H}(S) := \text{span}\{|s\rangle : s \in S\}, \quad \langle s | t \rangle = \delta_{s,t}.$$

For nonempty S , write $|+_S\rangle = |S|^{-1/2} \sum_{s \in S} |s\rangle$. We use $\|\cdot\|$ for the Euclidean norm of a vector and the operator norm of a linear map.

2.1 Quantum query model

We follow Carolan’s two-sided query notation [2, Sec. 3]. Let $\varphi \sim \mathbf{S}_N$ be uniform. The query register $\mathbf{X}$ contains a direction bit $b \in \{0, 1\}$ and a query point $x \in [N]$, and $\mathbf{Y}$ is the output register. Represent elements of $[N]$ by distinct bit strings of a fixed length and let $\oplus$ denote bitwise XOR. The query acts as

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

Thus $b = 0$ gives a forward query and $b = 1$ an inverse query; both b and x may be in superposition. We also allow controlled queries: a control qubit determines whether to apply the query or the identity. The adversary’s initial state, including any advice, is independent of φ .

³ The hook product in our normalization (3.9) is also used in Foxman et al.’s dimension-ratio calculation [5, Eq. (3.27)].

To purify the random permutation, introduce an oracle register with state space $\mathcal{H}(\mathbf{S}_N)$, initialized to

$$|+\mathbf{s}_N\rangle = \frac{1}{\sqrt{N!}} \sum_{\varphi \in \mathbf{S}_N} |\varphi\rangle.$$

Each query applies $\mathcal{O}_\varphi$ controlled by the permutation in this register. Tracing out the oracle register at the end gives the same final state as choosing φ uniformly at the start of the experiment. By retaining auxiliary registers, we may take the initial state to be pure and defer measurements until the end. We write $\rho_{\mathcal{A}}^{(\mathcal{O}_\varphi)}$ for the adversary's final state after tracing out the oracle register and the auxiliary purifying registers.

For density matrices σ and τ , define their trace distance by

$$D(\sigma, \tau) = \frac{1}{2} \|\sigma - \tau\|_1, \quad \|M\|_1 = \text{Tr} \sqrt{M^\dagger M}.$$

This is the maximum difference between the acceptance probabilities of a measurement applied to the two states. For normalized purifications $|\psi\rangle$ and $|\phi\rangle$ of σ and τ , respectively, in a common space,

$$D(\sigma, \tau) \leq \| |\psi\rangle - |\phi\rangle \|,$$

see Carolan's Lemmas 2.1–2.2 [2]. Thus a bound on the distance between purifications also bounds the distinguishing advantage.

2.2 Carolan's compressed permutation oracle

We recall the oracle using Carolan's definitions and notation [2, Secs. 2.2–3]. Let $\mathbf{I} = \mathbf{I}_{N,N}$ denote the set of partial injective functions from $[N]$ to $[N]$, called *injective databases*. We write $I(x) = \perp$ when I is undefined at x . Define

$$\text{Dom}(I) = \{x \in [N] : I(x) \neq \perp\}, \quad \text{Im}(I) = \{I(x) : x \in \text{Dom}(I)\},$$

and define the size of I by $|I| = |\text{Dom}(I)|$. Let $\perp$ denote the empty database, given by $\perp(x) = \perp$ for every x . We use $\mathbf{I}_t$ and $\mathbf{I}_{\leq t}$ for the subsets of size exactly t and at most t , respectively. We also call an input-output pair (x, y) an *edge*, recorded in I when $I(x) = y$. The database space is

$$\mathcal{H}(\mathbf{I}) = \bigoplus_{t=0}^N \mathcal{H}(\mathbf{I}_t).$$

The update $I[x \rightarrow y]$ sets the value at x to y and, if $y \neq \perp$, deletes any conflicting preimage of y . In particular, $I[x \rightarrow \perp]$ deletes the edge at x . When $x \notin \text{Dom}(I)$ and $y \in [N] \setminus \text{Im}(I)$, the update simply adds (x, y) . The inverse database I^{-1} exchanges inputs and outputs.

For J undefined at x , define

$$J|^x = \{J\} \cup \{J[x \rightarrow y] : y \in [N] \setminus \text{Im}(J)\}, \quad \mathbf{I}^x = \{J \in \mathbf{I} : x \notin \text{Dom}(J)\}.$$

Put $d = N - |J|$ and

$$|+_{x,J}\rangle = \frac{1}{\sqrt{d}} \sum_{y \in [N] \setminus \text{Im}(J)} |J[x \rightarrow y]\rangle.$$

On $\mathcal{H}(J|x)$, define permutation compression by

$$\text{pC}_{x,J}|J\rangle = |+_{x,J}\rangle, \quad \text{pC}_{x,J}|+_{x,J}\rangle = |J\rangle,$$

and let it act as the identity on the orthogonal complement of these two vectors in $\mathcal{H}(J|x)$. The sets $J|x$, indexed by $J \in \mathbf{I}|^x$, partition $\mathbf{I}$, so $\text{pC}_x = \bigoplus_{J \in \mathbf{I}|^x} \text{pC}_{x,J}$ is a well-defined involution.

Let P_x be Carolan's extended purified oracle P with query point fixed at x : it XORs $I(x)$ into the output register when $x \in \text{Dom}(I)$ and acts trivially otherwise. Define $\text{F}|I\rangle = |I^{-1}\rangle$. The forward block of the compressed oracle is

$$\text{cpO}_x = \text{pC}_x \text{P}_x \text{pC}_x,$$

and its inverse block is $\text{FcpO}_x \text{F}^\dagger$. The full direction-controlled oracle is denoted cpO . When $\mathcal{A}$ queries cpO , we denote the final joint pure state by $|\psi_{\mathcal{A}}^{(\text{cpO})}\rangle$, as in Carolan [2, Eq. (62)], with the database initialized to $|\perp\rangle_{\mathbf{I}}$. The adversary's final reduced state is $\rho_{\mathcal{A}}^{(\text{cpO})}$, obtained by tracing out the database and auxiliary purifying registers. Register subscripts are omitted when unambiguous.

Carolan's bounded-growth lemma [2, Lem. 3.2] shows that a query increases database size by at most one. Starting from the empty database, the state after q queries is therefore supported on $\mathcal{H}(\mathbf{I}_{\leq q})$.

2.3 Additional notation for verification queries

The ordinary lookup P_x leaves the output register unchanged both when $I(x)$ is undefined and when its value is the all-zero string. To prevent an undefined entry from passing verification as an all-zero value, we use the encoding $\text{enc}(y) = (1, y)$, called a *marked answer*, and call a query using this encoding a *marked query*. On the exact oracle, this is implemented by one ordinary query and a bit flip on an additional marker qubit.

For a basis value z of the output register, including the marker qubit, define

$$\text{P}_x^{\text{mark}}|z\rangle|I\rangle = \begin{cases} |z \oplus (1, I(x))\rangle|I\rangle, & x \in \text{Dom}(I), \\ |z\rangle|I\rangle, & x \notin \text{Dom}(I). \end{cases}$$

The compressed marked query is $\text{pC}_x \text{P}_x^{\text{mark}} \text{pC}_x$, with its inverse block obtained by conjugating with F . Permutation compression, the initial database, and the final-state notation remain unchanged.

Both pC_x and P_x^{mark} preserve each subspace $\mathcal{H}(J|x)$, whose basis databases have sizes $|J|$ and $|J| + 1$. The compressed marked query therefore changes database size by at most one; database inversion and coherent controls preserve this property. In section 3 we prove soundness for sequences of ordinary and marked queries, including those appended to check the input-output pairs returned by the adversary.

2.4 Distance between the oracle experiments

For a fixed domain size N , consider algorithms in the preceding query model, allowing ordinary or marked queries. For each algorithm $\mathcal{A}$, let $|\Psi\rangle$ and $|\Phi\rangle$ range over normalized purifications of $\rho_{\mathcal{A}}^{(\mathcal{O}_\varphi)}$ and $\rho_{\mathcal{A}}^{(\text{cpO})}$, respectively, in a common extension of the same adversary workspace. Define

$$\text{Dist}(q) := \sup_{\mathcal{A}: \text{at most } q \text{ queries}} \inf_{\Psi, \Phi} \|\Psi - \Phi\|. \quad (2.1)$$

Carolan's Theorem 5.19 states a trace-distance bound for the final reduced states. The predicate-search reduction compares square roots of success probabilities: a trace-distance bound ε gives only a $\sqrt{\varepsilon}$ bound on their difference in general, whereas a purification-distance bound ε gives ε . For ordinary forward and inverse queries, his proof already establishes a purification-distance bound of $O(q^3/N^{1/4})$ [2, proof of Thm. 5.19, following Eq. (278)]. We therefore use $\text{Dist}(q)$ for the distance defined above; Carolan uses this notation for a soundness-loss bound in Eq. (317).

3 Soundness of the compressed permutation oracle

This section proves $\text{Dist}(q) \leq 4q/\sqrt{N}$ for $q \leq N/4$. We recall Rosmanis's permutation-space results in section 3.1, then use them to embed the exact permutation register into the database space, preserving inner products, and construct the conditioning maps in section 3.2. With both oracles acting in the same space, the overlap calculation in section 3.3 bounds the effect of replacing one exact query by cpO on reachable exact states. A hybrid argument sums these differences in section 3.4.

3.1 Permutation spaces and results from Rosmanis

We follow Rosmanis's representation-theoretic notation, retaining Carolan's oracle notation. Rosmanis [8, Secs. 6.1–6.2] reviews the background used here; Serre [10] and Sagan [9] give more detailed accounts of the representation theory of finite groups and symmetric groups, respectively.

Relabeling and decomposition. Input and output relabelings act by

$$V_\pi^\tau |\varphi\rangle = |\tau\varphi\pi^{-1}\rangle, \quad V_\pi = V_\pi^\epsilon, \quad V^\tau = V_\epsilon^\tau,$$

where ϵ is the identity permutation. A partition $\lambda \vdash N$ is a nonincreasing list of row lengths summing to N , represented by a Young diagram. The irreducible representations of $\mathbf{S}_N$ are the Specht modules $\mathcal{S}^\lambda$. An irreducible representation has no nonzero proper invariant subspace. The λ -isotypic subspace is the sum of all subrepresentations isomorphic to $\mathcal{S}^\lambda$. The input and output isotypic subspaces $\mathcal{H}_\lambda$ and $\mathcal{H}^\lambda$ coincide in the permutation space, and their common value is denoted

by $\mathcal{H}^\lambda$. Rosmanis [8, Sec. 7, Eqs. (13)–(15)] gives the joint decomposition, with one copy of each $\mathcal{S}^\lambda \otimes \mathcal{S}^\lambda$:

$$\mathcal{H}(\mathbf{S}_N) = \bigoplus_{\lambda \vdash N} \mathcal{H}^\lambda, \quad \mathcal{H}^\lambda \cong \mathcal{S}^\lambda \otimes \mathcal{S}^\lambda.$$

The first factor carries the input action and the second the output action, using self-duality of the Specht modules. Schur's lemma then says that an operator commuting with both actions preserves each component and is a multiple of the identity there.

Permutation states and oracle subspaces. Identifying Rosmanis's assignment α with a database I , we write his states $|v_\alpha\rangle$ and nested oracle subspaces [8, Sec. 3.1] as

$$|v_I\rangle = \frac{1}{\sqrt{(N-t)!}} \sum_{\varphi \supset I} |\varphi\rangle, \quad \mathcal{A}_t = \text{span}\{|v_I\rangle : I \in \mathbf{I}_t\}, \quad |I| = t.$$

The states are also called consistent-permutation states [6, App. E]. Rosmanis's Claim 13 [8] identifies the irreducible components of these oracle subspaces: for $t < N/2$,

$$\mathcal{A}_t = \bigoplus_{\substack{\lambda \vdash N \\ \lambda_1 \geq N-t}} \mathcal{H}^\lambda. \quad (3.1)$$

After q forward queries, the oracle register is supported on $\mathcal{A}_q$ [8, Thm. 3(1)]. We check the inverse and marked query variants used here in lemma 1.

Eigenvalues used for normalization. For a partition ρ , let $d_\rho = \dim \mathcal{S}^\rho$, let $|\rho|$ be its number of boxes, and let $\rho_j^\top$ be the height of column j . The hook-length formula is

$$h_\rho(i, j) = \rho_i - j + \rho_j^\top - i + 1, \quad h(\rho) = \prod_{(i,j) \in \rho} h_\rho(i, j), \quad d_\rho = \frac{|\rho|!}{h(\rho)}.$$

Following Rosmanis's notation [8, Sec. 6.2 and App. A], we write $h_1(\lambda)$ for the product of the hook lengths in the first row. For $t = N - \lambda_1 < N/2$, his Lemma 19 [8, App. A] gives

$$\Phi_t = (N-t)! \sum_{I \in \mathbf{I}_t} |v_I\rangle \langle v_I|, \quad \Phi_t|_{\mathcal{H}^\lambda} = h_1(\lambda) \mathbf{1} \quad (\lambda_1 = N-t, \ t < N/2). \quad (3.2)$$

Fixing one input. Let $G_x = \{\pi \in \mathbf{S}_N : \pi(x) = x\} \cong \mathbf{S}_{N-1}$. Write $\mu \prec \lambda$ when μ is obtained by removing one corner box from λ , and write Res_H^G for restriction to a subgroup. The branching rule [8, Secs. 6.2 and 7.1] gives one copy of $\mathcal{S}^\mu \otimes \mathcal{S}^\lambda$ for each $\mu \prec \lambda$:

$$\text{Res}_{G_x \times \mathbf{S}_N}^{\mathbf{S}_N \times \mathbf{S}_N} \mathcal{H}^\lambda \cong \bigoplus_{\mu \prec \lambda} \mathcal{S}^\mu \otimes \mathcal{S}^\lambda. \quad (3.3)$$

After fixing both x and its image y , the remaining bijections form an $(N-1)$ -point permutation space. Its input μ -isotypic subspace is $\mathcal{S}^\mu \otimes \mathcal{S}^\mu$, by the same decomposition [8, Sec. 7.1, Eq. (16)].

3.2 Embedding and conditioning of permutation states

We now use these facts to place the exact oracle in Carolan's database space. Define

$$T_{N,t}|I\rangle = |v_I\rangle, \quad I \in \mathbf{I}_t, \quad \Phi_t = (N-t)! T_{N,t} T_{N,t}^\dagger.$$

The database basis states $|I\rangle$ are orthogonal, but their images under $T_{N,t}$ need not be. We therefore normalize $T_{N,t}^\dagger$ separately on each irreducible component.

For $\lambda = (N-t, \theta)$, define $\deg(\lambda) = N - \lambda_1 = t$ and

$$g_{N,\lambda} = \binom{N}{t} \frac{d_\theta}{d_\lambda}. \quad (3.4)$$

Here degree refers to the oracle subspaces: it is the least level t for which $\overset{\lambda}{\mathcal{H}} \subseteq \mathcal{A}_t$, equivalently the number of boxes below the first row. By the hook-length formula, $g_{N,\lambda} = h_1(\lambda)/(N-t)!$, so (3.2) gives the required eigenvalue of $T_{N,t} T_{N,t}^\dagger$ when $t < N/2$. We define the embedding on the full permutation space so that exact queries act unitarily on its image. Appendix A supplies the normalization formula for every irreducible component and extends (3.1) to all $0 \leq t \leq N$.

Now define

$$W_N v = g_{N,\lambda}^{-1/2} T_{N,t}^\dagger v, \quad v \in \overset{\lambda}{\mathcal{H}}, \quad t = \deg(\lambda). \quad (3.5)$$

Here $W_N v$ is a vector in $\mathcal{H}(\mathbf{I}_t)$, the subspace spanned by databases of size t . Set

$$\mathcal{H}_N = W_N \mathcal{H}(\mathbf{S}_N), \quad \mathcal{H}_{N,t} = W_N \left(\bigoplus_{\deg(\lambda)=t} \overset{\lambda}{\mathcal{H}} \right), \quad \mathcal{H}_{N,\leq t} = \bigoplus_{j \leq t} \mathcal{H}_{N,j}.$$

The space $\mathcal{H}_{N,t}$ contains the embedded degree- t components, not every size- t database vector. On databases the relabeling action is $V_\pi^\tau |I\rangle = |\tau I \pi^{-1}\rangle$.

Fix a query point x . We construct D_x so that its database image records the response $y = \varphi(x)$, allowing the exact query to act by database lookup. To condition on this response, let $j_{x,y}$ adjoin (x, y) to a bijection between the remaining labels, and let $\iota_{x,y}$ adjoin that edge to a residual database. Let $W_{N-1}^{x,y}$ be the embedding for those residual label sets. Define

$$J_x = \sum_y \iota_{x,y} W_{N-1}^{x,y} j_{x,y}^\dagger, \quad D_x = J_x W_N^\dagger|_{\mathcal{H}_N}. \quad (3.6)$$

We call D_x the conditioning map. Let $\mathcal{O}_x$ be the exact forward query and put $\hat{\mathcal{O}}_x = W_N \mathcal{O}_x W_N^\dagger$ on $\mathcal{H}_N$, including its action on the output register.

The following lemma checks the properties of these maps that are needed to apply the decompositions in section 3.1.

Lemma 1 (Properties of the embedding and conditioning maps). *The map W_N is an isometry, sends $|+\mathbf{s}_N\rangle$ to $|\perp\rangle$, and respects input and output relabelings. For ordinary and marked queries, in either direction and with coherent controls, an exact state after q queries has oracle-register support in $\mathcal{H}_{N,\leq q}$ after embedding by W_N .*

The maps J_x and D_x are isometries whose images lie in the span of database states defined at x . They and $\mathbf{pC}_x$ respect the $G_x \times \mathbf{S}_N$ relabeling actions, and

$$D_x \widehat{\mathcal{O}}_x = \mathbf{P}_x D_x. \quad (3.7)$$

For a component indexed by (λ, μ) in (3.3), the residual states lie in the input μ -isotypic subspace of the remaining permutation space. Their degree is

$$u = (N - 1) - \mu_1.$$

Thus D_x uses normalization $g_{N-1,\mu}^{-1/2}$ and maps this component to states supported on databases of size $u + 1$: size $t + 1$ if $\mu_1 = \lambda_1 - 1$, and size t otherwise, where $t = \deg(\lambda)$.

Proof. The eigenvalue $g_{N,\lambda}$ makes (3.5) isometric on each $\mathcal{H}^\lambda$. The images of distinct isotypic subspaces of the same degree remain orthogonal, while components of different degrees lie in distinct, mutually orthogonal subspaces $\mathcal{H}(\mathbf{I}_t)$. The degree-zero formula gives $W_N|+\mathbf{s}_N\rangle = |\perp\rangle$. Relabeling bijects extensions and preserves each isotypic subspace, so it commutes with the construction of W_N and its normalization.

We next show that a query increases the degree by at most one. For $I \in \mathbf{I}_t$ with $t < N$,

$$\mathbf{1}[\varphi(x) = y]|v_I\rangle = \begin{cases} |v_I\rangle, & I(x) = y, \\ \frac{|v_{I[x \rightarrow y]}\rangle}{\sqrt{N-t}}, & x \notin \text{Dom}(I), \ y \notin \text{Im}(I), \\ 0, & \text{otherwise.} \end{cases}$$

The inverse indicator $\mathbf{1}[\varphi^{-1}(y) = x]$ imposes the same constraint. Ordinary and marked queries are sums of these indicators tensored with output-register operations, so they map $\mathcal{A}_t$ into $\mathcal{A}_{t+1}$. The nesting of the oracle subspaces also covers identity branches under coherent control. Starting from $\mathcal{A}_0$, induction gives the claimed support, since workspace operations do not change the oracle-register support.

The summands of J_x have orthogonal database ranges because they specify distinct images of x . Since edge adjoining and the residual embeddings are isometric,

$$\begin{aligned} J_x^\dagger J_x &= \sum_y J_{x,y} (W_{N-1}^{x,y})^\dagger W_{N-1}^{x,y} J_{x,y}^\dagger \\ &= \sum_y J_{x,y} J_{x,y}^\dagger = \mathbf{1}. \end{aligned}$$

The last equality holds because each permutation maps x to exactly one y . Thus J_x is isometric, as is $D_x = J_x W_N^\dagger|_{\mathcal{H}_N}$, and their images lie in the span of database states defined at x . On the component $\varphi(x) = y$, both sides of (3.7) apply the same output-register operation, also for marked queries.

A relabeling with $\pi(x) = x$ carries the set $I|^x$ to $(\tau I \pi^{-1})|^x$ and its uniform extension vector to the corresponding uniform vector. It preserves the swap defining pC_x . In J_x , relabeling sends the term indexed by y to the term indexed by $\tau(y)$; edge adjoining and the residual embedding both respect this action. This proves the stated compatibility for J_x and D_x .

Finally, an input relabeling fixing x preserves $\varphi(x) = y$. Thus $j_{x,y}^\dagger$ intertwines the full and residual G_x -actions and preserves input representation types. Writing Π_{μ_x} for the input μ -isotypic projector, we have

$$\begin{aligned} j_{x,y}^\dagger V_\pi &= V_\pi j_{x,y}^\dagger, & \pi \in G_x, \\ j_{x,y}^\dagger \Pi_{\mu_x} &= \Pi_{\mu_x} j_{x,y}^\dagger, \end{aligned}$$

where the operators on the two sides act on the full and residual spaces, respectively. By the decomposition in section 3.1, the residual image lies in the $\mathcal{S}^\mu \otimes \mathcal{S}^\mu$ subspace of the remaining $(N-1)$ -point permutation space. Its embedding has degree $u = (N-1) - \mu_1$ and normalization $g_{N-1,\mu}^{-1/2}$; adjoining (x, y) gives the stated database sizes. This proves the lemma. $\square$

3.3 Comparison with permutation compression

We now compare D_x with pC_x . By lemma 1, $D_x^\dagger \text{pC}_x$ commutes with the $G_x \times \mathbf{S}_N$ action on $\mathcal{H}_N$. The decomposition (3.3) has one copy of each pair (λ, μ) , so Schur's lemma makes this operator a multiple $c_{\lambda,\mu} \mathbb{1}$ of the identity on each component. The distance between the two maps is controlled by how far $c_{\lambda,\mu}$ is from one. We compute these overlaps here and bound the resulting distance in section 3.4.

We first identify how pC_x acts on the embedded states. Orthogonality to the lower-degree oracle subspaces gives the following cancellation identity.

Lemma 2 (Vanishing sums over extensions). *Let $h \in \mathcal{H}_{N,t}$, let $J \in \mathbf{I}_{t-1}$, and let $x \notin \text{Dom}(J)$. In the database basis,*

$$\sum_{y \notin \text{Im}(J)} \langle J[x \rightarrow y] \mid h \rangle = 0.$$

The analogous identity holds after exchanging inputs and outputs.

Proof. For every J as above,

$$\sum_{y \notin \text{Im}(J)} T_{N,t} |J[x \rightarrow y]\rangle = \sqrt{N-t+1} T_{N,t-1} |J\rangle.$$

A degree- t vector in $\mathcal{H}(\mathbf{S}_N)$ is orthogonal to $\mathcal{A}_{t-1}$. Taking its inner product with this identity and using (3.5) proves the claimed vanishing sum. Exchanging inputs and outputs gives the other identity. $\square$

Consequently, when $\mathbf{pC}_x$ acts on a vector in $\mathcal{H}_{N,t}$, the part already defining x is unchanged, whereas the part not defining x is uniformly extended to a fresh output. The cancellation removes the otherwise possible size- $(t-1)$ component. Combining this action with the residual degree and normalization in lemma 1 gives the overlap coefficient.

Lemma 3 (Overlap on irreducible components). *Let $t = N - \lambda_1$ and let $\mu \prec \lambda$. On the irreducible component indexed by (λ, μ) ,*

$$D_x^\dagger \mathbf{pC}_x = c_{\lambda, \mu} \mathbb{1}, \quad c_{\lambda, \mu} = \begin{cases} \sqrt{g_{N, \lambda} / g_{N-1, \mu}}, & \mu_1 = \lambda_1 - 1, \\ \sqrt{g_{N-1, \mu} / g_{N, \lambda}}, & \mu_1 = \lambda_1. \end{cases}$$

Both ratios are positive and at most one.

Proof. Take a vector v in this component of $\mathcal{H}(\mathbf{S}_N)$ and put $h = W_N v$.

First suppose $\mu_1 = \lambda_1$. Then the residual degree is $t-1$, so $D_x h$ is supported on size- t databases defined at x . Write $\Pi^{x \in \text{db}}$ for the projector onto databases defining x . For a size- t database I containing (x, y) ,

$$T_{N,t}|I\rangle = j_{x,y} T_{N-1,t-1}|I[x \rightarrow \perp]\rangle; \quad (3.8)$$

both sides have normalization $1/\sqrt{(N-t)!}$. Comparing the coefficients in (3.5) and (3.6) yields

$$\Pi^{x \in \text{db}} h = \sqrt{\frac{g_{N-1, \mu}}{g_{N, \lambda}}} D_x h.$$

By lemma 2, the size- t part of $\mathbf{pC}_x h$ is the left-hand side; its remaining part is supported on size- $(t+1)$ databases and is orthogonal to $D_x h$. Taking inner products proves the second case.

Now suppose $\mu_1 = \lambda_1 - 1$. The residual degree is t . Every coefficient of h on a size- t database defining x vanishes: after removing (x, y) , (3.8) tests a residual degree- t vector against a consistent-permutation state for a database of size $t-1$. Hence $\mathbf{pC}_x h = E_x h$, where, on size- t databases not defining x ,

$$E_x |I\rangle = \frac{1}{\sqrt{N-t}} \sum_{y \notin \text{Im}(I)} |I[x \rightarrow y]\rangle.$$

The state $D_x h$ is supported on size- $(t+1)$ databases. Using (3.8) at level $t+1$ and the identity

$$\sum_{y \notin \text{Im}(I)} T_{N,t+1}|I[x \rightarrow y]\rangle = \sqrt{N-t} T_{N,t}|I\rangle$$

shows that the factor $\sqrt{N-t}$ cancels the factor $1/\sqrt{N-t}$ from $E_x^\dagger$. The residual embedding contributes $g_{N-1, \mu}^{-1/2}$, giving

$$E_x^\dagger D_x h = g_{N-1, \mu}^{-1/2} T_{N,t}^\dagger v = \sqrt{\frac{g_{N, \lambda}}{g_{N-1, \mu}}} h.$$

Taking the inner product with h proves the first case. $\square$

3.4 Proof of the soundness bound

To bound the overlap coefficients, we apply the hook-length formula to (3.4). For $\lambda = (N - t, \theta)$, cancelling the hook lengths below the first row gives

$$g_{N,\lambda} = \prod_{j=1}^{\theta_1} \left(1 + \frac{\theta_j^\top}{N - t - j + 1} \right). \quad (3.9)$$

Write $\lambda = (a, \theta)$ with $a = N - t$ and put $\ell_j = \theta_j^\top$. In both cases below, assume $t < N/2$. When μ is obtained from λ by removing a box from the first row, cancelling the hook-length products gives

$$c_{\lambda,\mu}^2 = \prod_{j=1}^{\theta_1} \left(1 - \frac{\ell_j}{(a - j + 1)(a - j + \ell_j)} \right),$$

and therefore

$$1 - c_{\lambda,\mu}^2 \leq \frac{t}{(N - 2t + 1)^2}. \quad (3.10)$$

If instead a corner of θ is removed from column c of height $\ell = \theta_c^\top$, cancellation leaves

$$c_{\lambda,\mu}^2 = \frac{a - c + \ell}{a - c + \ell + 1}, \quad 1 - c_{\lambda,\mu}^2 \leq \frac{1}{N - 2t + 2}.$$

Inequality (3.10) follows from $1 - \prod_j (1 - u_j) \leq \sum_j u_j$ and $\sum_j \ell_j = t$. Both right-hand sides are at most $2/N$ when $t \leq N/4$.

Since $\mathbf{pC}_x$ and D_x are isometries on $\mathcal{H}_N$, the identity on each component gives

$$\|(D_x - \mathbf{pC}_x)h\|^2 = 2(1 - c_{\lambda,\mu}) \|h\|^2 \leq 2(1 - c_{\lambda,\mu}^2) \|h\|^2.$$

Taking the maximum over the irreducible components of $\mathcal{H}_{N,\leq s}$ yields

$$\|(D_x - \mathbf{pC}_x)\Pi_{\mathcal{H}_{N,\leq s}}\| \leq \frac{2}{\sqrt{N}}, \quad s \leq N/4.$$

Tensoring the operator with the identity on an auxiliary register does not change its norm.

For a unit vector $h \in \mathcal{H}_{N,\leq t}$, the state $\widehat{\mathcal{O}}_x h$ has degree at most $t + 1$. Using (3.7), the involution $\mathbf{pC}_x$, and unitary invariance,

$$\begin{aligned} \|\mathbf{cpO}_x h - \widehat{\mathcal{O}}_x h\| &= \|\mathbf{P}_x \mathbf{pC}_x h - \mathbf{pC}_x \widehat{\mathcal{O}}_x h\| \\ &\leq \|\mathbf{P}_x (\mathbf{pC}_x - D_x) h\| + \|(D_x - \mathbf{pC}_x) \widehat{\mathcal{O}}_x h\| \\ &\leq \frac{4}{\sqrt{N}}, \quad t + 1 \leq N/4. \end{aligned} \quad (3.11)$$

The maps $T_{N,t}$ and W_N intertwine the inversion operators on their respective domain and codomain spaces. Inversion preserves the label λ and the scalar $g_{N,\lambda}$,

so the same bound holds for inverse queries. With coherent control over the query direction, query point, or whether to make a query, the difference operator is a direct sum of blocks satisfying the same bound.

The calculation uses only that, on databases defined at x , P_x applies the same operation to the output register as the exact query. It therefore also holds for the marked answer encoding introduced in section 2.3. Ordinary and marked queries may be interleaved in the same experiment, with each counting as one query.

The following theorem sums the one-query bound along a hybrid sequence to compare the final states after q queries.

Theorem 1 (Improved soundness). *For $q \leq N/4$, $\text{Dist}(q) \leq 4q/\sqrt{N}$. More precisely, let a quantum algorithm $\mathcal{A}$, with oracle-independent initial advice, make q coherently controlled forward or inverse queries to a uniform permutation on $[N]$, allowing ordinary or marked queries. Let $h_{\mathcal{A}}$ be the image under W_N of its exact final purification. The states $h_{\mathcal{A}}$ and $|\psi_{\mathcal{A}}^{(\text{cpO})}\rangle$ are purifications of $\rho_{\mathcal{A}}^{(\text{O}_\varphi)}$ and $\rho_{\mathcal{A}}^{(\text{cpO})}$, respectively, in a common space, and*

$$\left\| h_{\mathcal{A}} - |\psi_{\mathcal{A}}^{(\text{cpO})}\rangle \right\| \leq \frac{4q}{\sqrt{N}}, \quad q \leq N/4.$$

Consequently, for every final projector G on the adversary workspace,

$$\left| \|Gh_{\mathcal{A}}\| - \|G|\psi_{\mathcal{A}}^{(\text{cpO})}\rangle\| \right| \leq \text{Dist}(q) \leq \frac{4q}{\sqrt{N}}.$$

Proof. We extend the transported exact query arbitrarily to a unitary on the orthogonal complement of $\mathcal{H}_N$. A telescoping argument compares the two q -query evolutions. At the j th difference, the first $j-1$ queries are exact, so the input to query j belongs to $\mathcal{H}_{N, \leq j-1}$, and the last $q-j$ queries are compressed. The operations after query j are unitary and therefore preserve the norm of this difference. Summing (3.11) gives the purification-distance bound. For each algorithm, the infimum in (2.1) is at most the distance of this particular pair. Taking the supremum over algorithms with at most q queries gives the bound on $\text{Dist}(q)$. For the last assertion, the two projected norms depend only on the reduced states. The reverse triangle inequality bounds their difference by the distance of every purification pair. Taking the infimum over pairs and then using (2.1) proves the theorem. $\square$

4 Applying Carolan's search bound

Carolan's search bounds apply with the distance Dist from (2.1). We substitute the upper bound in theorem 1 in his Lemma 7.1 and Theorem 7.5; the database analysis is unchanged [2].

Following Carolan [2, Sec. 7], let $\mathcal{R} \subseteq ([N]^{\times 2})^*$ be a predicate on finite lists of input-output pairs. Write $I \cap \mathcal{R} \neq \emptyset$ when I contains every pair of some list in $\mathcal{R}$, and let $\Pi_{\mathcal{R}}$ be the projector onto the span of those database states.

Let $s_t(\mathcal{R})$ denote the t -sparsity of $\mathcal{R}$, as defined by Carolan [2, Def. 7.4].

A list of at most l input-output pairs can be checked with at most l marked queries (section 2.3). With this notation, Carolan's predicate-search reduction and sparsity argument, together with theorem 1, give the following corollary.

Corollary 1 (Predicate search). *Let an adversary make q queries to a uniform random permutation and return a list of at most l input-output pairs. It wins when every pair (x, y) satisfies $\varphi(x) = y$ and the list belongs to a fixed predicate $\mathcal{R}$. Suppose $q + l \leq N/4$. For the same adversary in the compressed experiment, let $p_{\text{db}} = \left\| \Pi_{\mathcal{R}} |\psi_{\mathcal{A}}^{(\text{cpO})}\rangle \right\|^2$. Then*

$$\sqrt{p_{\mathcal{A}}} \leq \sqrt{p_{\text{db}}} + l \sqrt{\frac{2}{N - q - l}} + \text{Dist}(q + l). \quad (4.1)$$

Using theorem 1, if $\perp \cap \mathcal{R} = \emptyset$, then

$$p_{\mathcal{A}} = O \left(\left(\sum_{t=0}^{q-1} \sqrt{\frac{s_t(\mathcal{R})}{N}} \right)^2 + \frac{(q+l)^2}{N} \right).$$

Proof. Applying the last assertion of theorem 1 to the algorithm followed by its at most l marked verification queries gives the term $\text{Dist}(q + l)$ in the reduction of Carolan's Lemma 7.1. Write $\Pi^{(x,y) \in \text{db}}$ for the consistency projector onto the span of database states satisfying $I(x) = y$. With the output register initialized to zero, the branch of a compressed marked query accepting $\text{enc}(y)$ acts on the database register as $\text{pC}_x \Pi^{(x,y) \in \text{db}} \text{pC}_x$. For the stated verification constant, on a subspace $\mathcal{H}(J|x)$ with $y \notin \text{Im}(J)$ and $d = N - |J|$,

$$\left\| \text{pC}_x \Pi^{(x,y) \in \text{db}} \text{pC}_x - \Pi^{(x,y) \in \text{db}} \right\| = \sqrt{1 - (1 - 1/d)^2} \leq \sqrt{2/d}$$

and the left-hand side is zero on $\mathcal{H}(J|x)$ when $y \in \text{Im}(J)$. We apply the telescoping argument in Carolan's fundamental lemma [2, Lem. 3.3] to the ordered marked checks and consistency projectors. Products of the consistency projectors preserve database size at most q , so $d \geq N - q \geq N - q - l$. Summing over the l checks proves the first inequality.

The database part of Carolan's proof of Theorem 7.5 gives

$$\sqrt{p_{\text{db}}} = O \left(\sum_{t < q} \sqrt{\frac{s_t(\mathcal{R})}{N}} \right)$$

in this query range. The same argument applies to marked lookup, which leaves the database register unchanged and hence commutes with $\Pi_{\mathcal{R}}$. The probability bound follows by substituting $\text{Dist}(q + l) \leq 4(q + l)/\sqrt{N}$, squaring the first inequality, and using $N - q - l \geq 3N/4$. $\square$

5 Applications

We substitute theorem 1 into Carolan’s search bounds [2, Secs. 7.1–7.4], using the number of queries in each Dist term below, including witness construction and verification.

Remark 1 (Verification queries). Carolan’s application statements use $\text{Dist}(q)$, omitting the verification count l included in $\text{Dist}(q+l)$ in Lemma 7.1. The sponge reductions in Theorems 7.7–7.8 do not explicitly count the queries needed to construct lists of input-output pairs from the messages returned by the adversary. For clarity, appendix B specifies the list lengths and the queries needed to construct and verify the lists in each application.

5.1 One-more problem

The adversary returns $q+1$ distinct pairs (x_i, y_i) and wins when $\varphi(x_i) = y_i$ for all i .

Theorem 2 (One-more problem). *Every q -query quantum algorithm satisfies*

$$p_{\mathcal{A}}^{\text{om}} = O\left(\frac{(q+1)^2}{N}\right).$$

Proof. The proof of Carolan’s Theorem 7.2, with $l = q+1$ as in appendix B.1, gives

$$p_{\mathcal{A}}^{\text{om}} = O\left(\frac{(q+1)^2}{N} + \text{Dist}(2q+1)^2\right). \quad (5.1)$$

The claim follows from theorem 1. $\square$

5.2 Cycle finding

A *simple cycle of length l* is a tuple of pairwise distinct points $x_1, \dots, x_l$ such that $\varphi(x_i) = x_{i+1}$, with indices modulo l .

Theorem 3 (Cycle finding). *The probability that a q -query algorithm returns a simple cycle of any length is*

$$p_{\mathcal{A}}^{\text{cyc}} = O\left(\frac{(q+1)^3}{N}\right).$$

Proof. The proof of Carolan’s Theorem 7.3 and the long-cycle reduction in appendix B.2 give

$$p_{\mathcal{A}}^{\text{cyc}} = O\left(\frac{q^3 + (q+1)^2}{N} + \text{Dist}(2q+1)^2\right). \quad (5.2)$$

The claim follows from theorem 1. $\square$

5.3 Double-sided zero search

Let $N = 2^{2n}$. Double-sided zero search asks for $x, y \in \{0, 1\}^n$ such that

$$\varphi(x \| 0^n) = y \| 0^n.$$

Corollary 2 (Double-sided zero). *Every q -query quantum algorithm has*

$$p_{\mathcal{A}}^{\text{dsz}} = O\left(\frac{(q+1)^2}{2^n}\right).$$

Proof. The proof of Carolan's Corollary 7.6, with $l = 1$ as in appendix B.3, gives

$$p_{\mathcal{A}}^{\text{dsz}} = O\left(\frac{q^2}{2^n} + \frac{1}{N} + \text{Dist}(q+1)^2\right). \quad (5.3)$$

The claim follows from theorem 1. $\square$

This gives the known tight $\Theta(2^{n/2})$ query complexity for constant success probability [3,4].

5.4 Sponge with one output block

Let $N = 2^{r+c}$ and let $\varphi : \{0, 1\}^{r+c} \rightarrow \{0, 1\}^{r+c}$ be a uniformly random permutation. Fix an injective padding map from messages to nonempty sequences of r -bit blocks. For $\text{pad}(M) = (m_1, \dots, m_k)$, define the sponge with one r -bit output block by

$$\begin{aligned} u_0 &= 0^{r+c}, & x_i &= u_{i-1} \oplus (m_i \| 0^c), \\ u_i &= \varphi(x_i), & \text{Sp}^\varphi(M) &= u_k[0 : r]. \end{aligned}$$

The permutation calls give the input-output pairs (x_i, u_i) for $1 \leq i \leq k$. The fixed preimage target $w \in \{0, 1\}^r$ is chosen independently of φ . In the collision game, the two padded messages must be distinct.

The slices $z[0 : r]$ and $z[r : r+c]$ denote the rate and capacity parts. Let L bound the total padded block length of the messages returned by the adversary. Counting the queries used to construct and check the list of pairs as in appendix B.4 gives the following bounds.

Theorem 4 (Sponge preimage and collision). *Suppose a q -query adversary returns a preimage whose padded length is at most L . Then*

$$p_{\mathcal{A}}^{\text{pre}} = O\left(\frac{(q+L)^2}{2^r} + \frac{(q+L)^3}{2^c}\right).$$

If it returns a collision whose messages have total padded length at most L , then

$$p_{\mathcal{A}}^{\text{col}} = O\left(\frac{(q+L)^3}{2^{\min(r,c)}}\right).$$

Proof. The proofs of Carolan’s Theorems 7.7–7.8, applied to the reduction making at most $q + L$ queries with $l = L$ as in appendix B.4, give

$$p_{\mathcal{A}}^{\text{pre}} = O\left(\frac{(q + L)^2}{2^r} + \frac{(q + L)^3}{2^c} + \frac{L^2}{N} + \text{Dist}(q + 2L)^2\right), \quad (5.4)$$

$$p_{\mathcal{A}}^{\text{col}} = O\left(\frac{(q + L)^3}{2^{\min(r,c)}} + \frac{L^2}{N} + \text{Dist}(q + 2L)^2\right). \quad (5.5)$$

Both claims follow from theorem 1. $\square$

For nonempty block messages with identity padding, these bounds give tight $\Theta(\min\{2^{r/2}, 2^{c/3}\})$ and $\Theta(2^{\min(r,c)/3})$ query complexities for constant success probability in preimage and collision search, respectively, measured in the sum $q + L$ of the number of queries and the total number of blocks in these messages [1, Sec. II-B].

5.5 Davies–Meyer

As in Carolan [2, Sec. 7.4], consider a single permutation on $N = 2^n$ points and

$$\text{DM}^\varphi(x) = \varphi(x) \oplus x.$$

For a fixed oracle-independent target w , inversion asks for $\varphi(x) = x \oplus w$. Collision finding asks for distinct x, x' with $\text{DM}^\varphi(x) = \text{DM}^\varphi(x')$.

Theorem 5 (Keyless Davies–Meyer). *For a uniformly random permutation on $\{0, 1\}^n$, every q -query algorithm $\mathcal{A}$ satisfies*

$$p_{\mathcal{A}}^{\text{inv}} = O\left(\frac{(q + 1)^2}{2^n}\right), \quad p_{\mathcal{A}}^{\text{col}} = O\left(\frac{(q + 1)^3}{2^n}\right).$$

Proof. The proofs of Carolan’s Theorems 7.9–7.10, with verification counted as in appendix B.5, give

$$\begin{aligned} p_{\mathcal{A}}^{\text{inv}} &= O\left(\frac{q^2 + 1}{N} + \text{Dist}(q + 1)^2\right), \\ p_{\mathcal{A}}^{\text{col}} &= O\left(\frac{q^3 + 1}{N} + \text{Dist}(q + 3)^2\right). \end{aligned} \quad (5.6)$$

Both claims follow from theorem 1. $\square$

These bounds give tight $\Theta(2^{n/2})$ and $\Theta(2^{n/3})$ query complexities for constant success probability in inversion and collision search, respectively, matching Grover search and quantum collision finding [2, Sec. 7.4].

AI use disclosure. Codex, using GPT-6 Astra in Ultra mode, was used extensively to develop proof ideas, carry out technical analyses, produce the Lean 4 formalization, and draft and revise the manuscript. The author takes full responsibility for the paper and the accompanying Lean formalization.

References

1. Alagic, G., Carolan, J., Majenz, C., Tokat, S.: The sponge is quantum indifferentiable. In: 2025 IEEE 66th Annual Symposium on Foundations of Computer Science (FOCS). pp. 2591–2630. IEEE (2025). <https://doi.org/10.1109/FOCS63196.2025.00135>
2. Carolan, J.: Compressed permutation oracles. In: Proceedings of the 58th Annual ACM Symposium on Theory of Computing. pp. 150–161. ACM (2026). <https://doi.org/10.1145/3798129.3800736>, section, statement, and equation numbers refer to arXiv:2509.18586v1
3. Carolan, J., Poremba, A.: Quantum one-wayness of the single-round sponge with invertible permutations. In: Advances in Cryptology – CRYPTO 2024, Part VI. Lecture Notes in Computer Science, vol. 14925, pp. 218–252. Springer Nature Switzerland (2024). https://doi.org/10.1007/978-3-031-68391-6_7
4. Cojocaru, A., Hhan, M., Liu, Q., Yamakawa, T., Yun, A.: Quantum lifting for invertible permutations and ideal ciphers. In: Advances in Cryptology – CRYPTO 2025, Part II. Lecture Notes in Computer Science, vol. 16001, pp. 481–512. Springer Nature Switzerland (2025). https://doi.org/10.1007/978-3-032-01878-6_16, section, statement, and equation numbers refer to arXiv:2504.18188v1
5. Foxman, B., Lombardi, A., Ma, F., Nehoran, B., Wright, J.: Quantum lazy sampling and path recording for any group. arXiv:2606.30281 (2026), <https://arxiv.org/abs/2606.30281v1>, version cited: arXiv:2606.30281v1
6. Grinko, D., Yoshida, S.: Quantum simulation of random unitaries from Clebsch–Gordan transforms. arXiv:2509.26623 (2025), <https://arxiv.org/abs/2509.26623v2>, version cited: arXiv:2509.26623v2
7. Majenz, C., Malavolta, G., Walter, M.: Permutation superposition oracles for quantum query lower bounds. In: Proceedings of the 57th Annual ACM Symposium on Theory of Computing. pp. 1508–1519. ACM (2025). <https://doi.org/10.1145/3717823.3718266>
8. Rosmanis, A.: Tight bounds for inverting permutations via compressed oracle arguments. arXiv:2103.08975 (2021), <https://arxiv.org/abs/2103.08975v2>, version cited: arXiv:2103.08975v2
9. Sagan, B.E.: The Symmetric Group: Representations, Combinatorial Algorithms, and Symmetric Functions, Graduate Texts in Mathematics, vol. 203. Springer New York, 2 edn. (2001). <https://doi.org/10.1007/978-1-4757-6804-6>
10. Serre, J.P.: Linear Representations of Finite Groups, Graduate Texts in Mathematics, vol. 42. Springer New York (1977). <https://doi.org/10.1007/978-1-4684-9458-7>
11. Unruh, D.: Towards compressed permutation oracles. In: Advances in Cryptology – ASIACRYPT 2023, Part IV. Lecture Notes in Computer Science, vol. 14441, pp. 369–400. Springer Nature Singapore (2023). https://doi.org/10.1007/978-981-99-8730-6_12
12. Zhandry, M.: How to record quantum queries, and applications to quantum indistinguishability. In: Advances in Cryptology – CRYPTO 2019, Part II. Lecture Notes in Computer Science, vol. 11693, pp. 239–268. Springer International Publishing (2019). https://doi.org/10.1007/978-3-030-26951-7_9

A Normalization of W_N on the full permutation space

The embedding W_N in section 3.2 is defined on every irreducible component, although the soundness bound only uses low degrees. Rosmanis's Lemma 19 states its eigenvalue formula for $t < N/2$ [8]. We give the all-level version needed for this global definition, following the character calculation in his Appendix A.3 but retaining the general branching multiplicity.

Let $f^{\lambda/(N-t)}$ count standard tableaux of the skew diagram left after removing the leftmost $N - t$ boxes of the first row of λ . Such a tableau fills the remaining boxes with $1, \dots, t$, once each, increasing along rows and columns. Set $m_{\lambda,t} = f^{\lambda/(N-t)}$, with value zero if $\lambda_1 < N - t$. The resulting eigenvalue formula supplies both the global normalization and the decomposition of $\mathcal{A}_t$.

Lemma 4 (Eigenvalues of $T_{N,t}T_{N,t}^\dagger$). *For $0 \leq t \leq N$, on $\mathcal{H}^\lambda$ the operator $T_{N,t}T_{N,t}^\dagger$ is a multiple of the identity, with eigenvalue*

$$\binom{N}{t} \frac{m_{\lambda,t}}{d_\lambda}.$$

This eigenvalue is zero exactly when $\lambda_1 < N - t$. At $t = N - \lambda_1$, it equals $g_{N,\lambda}$.

Proof. Write $\mathfrak{F}(g) = |\{i : g(i) = i\}|$. Counting the size- t databases consistent with both φ and σ gives

$$\langle \varphi | T_{N,t} T_{N,t}^\dagger | \sigma \rangle = \frac{\binom{\mathfrak{F}(\varphi^{-1}\sigma)}{t}}{(N-t)!}. \quad (\text{A.1})$$

The quantity $\mathfrak{F}(\varphi^{-1}\sigma)$ counts the inputs on which φ and σ agree and is invariant under their simultaneous input and output relabeling. Thus this operator commutes with both actions and is a multiple of the identity on each $\mathcal{H}^\lambda$.

Let $\chi_\lambda(g) = \text{Tr}(\mathcal{S}^\lambda(g))$ be the character of $\mathcal{S}^\lambda$. The action on t -permutations of $[N]$, that is, ordered t -tuples of distinct points, has character $X_t(g) = (\mathfrak{F}(g))_t$, where $(a)_t = a(a-1)\cdots(a-t+1)$ and $(a)_0 = 1$. The stabilizer of a tuple is $\mathbf{S}_{N-t}$. By Frobenius reciprocity, the multiplicity of $\mathcal{S}^\lambda$ in this action equals the multiplicity of the trivial representation in its restriction to $\mathbf{S}_{N-t}$. Iterating the branching rule identifies this multiplicity with the number of paths from λ to the one-row partition $(N-t)$, namely $m_{\lambda,t}$. Character orthogonality gives

$$m_{\lambda,t} = \frac{1}{N!} \sum_{g \in \mathbf{S}_N} X_t(g) \chi_\lambda(g^{-1}).$$

Put $k_t(g) = X_t(g)/(t!(N-t)!)$. By (A.1), $T_{N,t}T_{N,t}^\dagger = \sum_g k_t(g) V_g$. On the d_λ -dimensional input factor its eigenvalue is its trace divided by d_λ . Changing variables from g to g^{-1} and using $k_t(g^{-1}) = k_t(g)$ gives the eigenvalue of $T_{N,t}T_{N,t}^\dagger$

on $\overset{\lambda}{\mathcal{H}}$:

$$\begin{aligned} \frac{1}{d_\lambda} \sum_{g \in \mathbf{S}_N} k_t(g) \chi_\lambda(g) &= \frac{1}{d_\lambda} \sum_{g \in \mathbf{S}_N} k_t(g) \chi_\lambda(g^{-1}) \\ &= \binom{N}{t} \frac{m_{\lambda,t}}{d_\lambda}. \end{aligned}$$

The multiplicity is positive exactly when $\lambda_1 \geq N - t$. At $t = N - \lambda_1$, removing the first row leaves $\theta = (\lambda_2, \lambda_3, \dots)$, so $m_{\lambda,t} = d_\theta$. Substituting this identity into the eigenvalue formula gives $g_{N,\lambda}$ by (3.4), proving the last assertion. $\square$

Since the image of $T_{N,t}$ equals the image of $T_{N,t} T_{N,t}^\dagger$, the lemma also gives

$$\mathcal{A}_t = \text{im } T_{N,t} = \bigoplus_{\substack{\lambda \vdash N \\ \lambda_1 \geq N-t}} \overset{\lambda}{\mathcal{H}} \quad (0 \leq t \leq N).$$

This justifies the all-level use of the decomposition in lemmas 1 and 2.

B Counting verification queries

This appendix counts the verification queries in the reductions used in section 5. Squaring (4.1) and using $N - q - l \geq 3N/4$ gives

$$p_{\mathcal{A}} = O\left(p_{\text{db}} + \frac{l^2}{N} + \text{Dist}(q+l)^2\right), \quad q + l \leq N/4. \quad (\text{B.1})$$

B.1 One-more problem

As in the proof of Carolan's Theorem 7.2, a database of size at most q cannot contain $q + 1$ distinct input-output pairs. Thus $p_{\text{db}} = 0$. Taking $l = q + 1$ in (B.1) gives (5.1), provided $2q + 1 \leq N/4$.

B.2 Cycle finding

Split success into cycles of length at most q and cycles of greater length. For the first event, the proof of Carolan's Theorem 7.3 bounds the database probability by $O(q^3/N)$. There are at most q verification queries, giving

$$\Pr[\text{success and } l \leq q] = O\left(\frac{q^3 + q^2}{N} + \text{Dist}(2q)^2\right).$$

For the second event, the reduction retains the first $q + 1$ edges of the simple cycle returned by the adversary. They form a solution to the one-more problem whenever that event occurs, so (5.1) bounds its probability. Adding the two probabilities and using $\text{Dist}(2q) \leq \text{Dist}(2q + 1)$ gives (5.2), under $2q + 1 \leq N/4$.

B.3 Double-sided zero search

Let $N = 2^{2^n}$. Carolan's Corollary 7.6 uses a predicate on one input-output pair with sparsity $O(2^n)$, giving $p_{\text{db}} = O(q^2/2^n)$. The adversary specifies both endpoints, so $l = 1$ in (B.1) gives (5.3), under $q + 1 \leq N/4$.

B.4 Sponge

The reduction pads the messages returned by the adversary and makes the permutation queries prescribed by the sponge computation, using at most L additional queries and recording at most L input-output pairs. Checking these pairs takes at most L further queries, so the database analysis uses $q + L$ queries and the comparison uses $q + 2L$.

We use the predicates in the proofs of Carolan's Theorems 7.7–7.8, augmented by internal collisions. Here a nonempty sponge path is a list $[(x_1, y_1), \dots, (x_k, y_k)]$ with $x_1[r : r + c] = 0^c$ and $y_i[r : r + c] = x_{i+1}[r : r + c]$ for $i < k$. An internal collision is a nonempty path returning to the initial capacity value 0^c , or two paths for distinct block messages ending at the same capacity value. Whenever a predicate involves two message paths, the corresponding block messages must be distinct. We use Carolan's sponge sparsity bounds [2, proofs of Thms. 7.7–7.8] with the collision conditions above. Including the empty database gives, for $t \geq 0$, the respective preimage and collision sparsities

$$O(2^c + (t + 1)2^r), \quad O((t + 1)(2^r + 2^c)).$$

The added 1 accounts for a first edge returning to the initial capacity 0^c .

Nonempty injective padding ensures that each successful preimage or collision yields a witness for these predicates. Carolan's sparsity argument for $q + L$ queries [2, proof of Thm. 7.5], followed by (B.1) with $l = L$, yields (5.4) and (5.5), provided $q + 2L \leq N/4$.

B.5 Davies–Meyer

Let $N = 2^n$. Without loss of generality, take the inversion target to be 0^n . The predicate in Carolan's Theorem 7.9 has sparsity one, so $p_{\text{db}} = O(q^2/N)$. Its single pair gives the inversion bound in (5.6), under $q + 1 \leq N/4$.

For collision finding, the adversary returns only distinct inputs x, x' . The reduction makes one additional query to obtain $y = \varphi(x)$ and forms the list

$$[(x, y), (x', y \oplus x \oplus x')].$$

The list satisfies Carolan's collision predicate and is correct exactly when x, x' collide. The proof of his Theorem 7.10 gives $p_{\text{db}} = O((q + 1)^3/N)$ from sparsity t . Using (B.1) for the reduction making at most $q + 1$ queries, with $l = 2$, gives the collision bound in (5.6), under $q + 3 \leq N/4$.