

Quantum One-Wayness of the Single-Round Sponge with Invertible Permutations

Joseph Carolan*
University of Maryland

Alexander Poremba†
MIT

Abstract

Sponge hashing is a widely used class of cryptographic hash algorithms which underlies the current international hash function standard SHA-3. In a nutshell, a sponge function takes as input a bit-stream of any length and processes it via a simple iterative procedure: it repeatedly feeds each block of the input into a so-called *block function*, and then produces a digest by once again iterating the block function on the final output bits. While much is known about the post-quantum security of the sponge construction when the block function is modeled as a random function or one-way permutation, the case of *invertible* permutations, which more accurately models the construction underlying SHA-3, has so far remained a fundamental open problem.

In this work, we make new progress towards overcoming this barrier and show several results. First, we prove the “double-sided zero-search” conjecture proposed by Unruh (eprint’ 2021) and show that finding zero-pairs in a random $2n$ -bit permutation requires at least $\Omega(2^{n/2})$ many queries—and this is tight due to Grover’s algorithm. At the core of our proof lies a novel “symmetrization argument” which uses insights from the theory of Young subgroups. Second, we consider more general variants of the double-sided search problem and show similar query lower bounds for them. As an application, we prove the quantum one-wayness of the single-round sponge with invertible permutations in the quantum random permutation model.

1 Introduction

Hash functions are one of the most fundamental objects in cryptography. They are used in a multitude of applications, such as integrity checks for data packages (e.g., software updates), password storage on cloud servers, or as important components in digital signature schemes—whether it is to design hash-based signatures or simply to construct signatures schemes for variable input lengths [KL14].

*jcarolan@umd.edu

†poremba@mit.edu

In recent years, the National Institute of Standards and Technology (NIST) announced a new international hash function standard known as SHA-3. Unlike its predecessor SHA-2, which was rooted in the Merkle-Damgård construction [Mer88; Mer90; Dam87], the new hash function standard uses Keccak [Ber+11b]—a family of cryptographic functions based on the idea of *sponge hashing* [Ber+11a]. This particular approach allows for both variable input length and variable output length, which makes it particularly attractive towards the design of cryptographic hash functions. The internal state of a sponge function gets updated through successive applications of a so-called *block function* $\varphi : \{0,1\}^{r+c} \rightarrow \{0,1\}^{r+c}$, where we call the parameters $r \in \mathbb{N}$ the *rate* and $c \in \mathbb{N}$ the *capacity* of the sponge. The evaluation of the sponge function $\text{Sp}^\varphi : \{0,1\}^* \rightarrow \{0,1\}^*$ takes place in two phases:

1. (Absorption phase:) During each round, a new block consisting of r many bits of the input gets “fed” into the block function φ . Following the sponge metaphor, we say that the sponge function “absorbs” the input. More formally, suppose that the input consists of r -bit blocks m_1, m_2, \dots, m_ℓ . During the first round, we compute $\varphi(m_1||0^c) = y_1||z_1$, for some $y_1 \in \{0,1\}^r$ and $z_1 \in \{0,1\}^c$. Then, to absorb the next block m_2 , we compute $\varphi(y_1 \oplus m_2||z_1) = y_2||z_2$, and so on.
2. (Squeezing phase:) During each round, a new r -bit block is produced by essentially running the absorption phase in reverse. Each r -bit digest becomes a block of the final output of the function Sp^φ . Following the sponge metaphor, we say that the sponge function gets “squeezed” to produce fresh random bits. More formally, suppose the last round of the absorption phase results in a pair $y_\ell||z_\ell$. We then let the digest y_ℓ serve as the first block of the output. To produce the second block, we compute $\varphi(y_\ell||z_\ell) = y_{\ell+1}||z_{\ell+1}$, and output $y_{\ell+1}$, and so on.

Several works have since analyzed the security properties behind the sponge construction [Ber+08; FGK22; Aks+23]. In particular, in [Ber+08] it was shown that the sponge hash function enjoys a strong form of security called *indifferentiability* in the case when the underlying block function φ is modeled as an invertible random permutation. Indifferentiability already implies many desired properties of cryptographic hash functions, such as collision-resistance or pseudorandomness.

While much is known about the security of the sponge in a classical world, our understanding changes significantly once we take the threat of large-scale quantum computing into account. The good news is that most hash functions are believed to only be mildly affected, whereas the majority of pre-quantum public-key cryptography faces serious threats due to Shor’s algorithm (and its variants) [Sho97; Reg24]. The reason hash functions are believed to be less susceptible to quantum attacks is due to their inherent lack of structure, which means that generic quantum attacks tend to achieve at most a square-root speed-up relative to their classical counterparts. To this day, however, our understanding of the post-quantum security of the sponge remains somewhat incomplete. In fact, the post-quantum security of the sponge construction is only well-understood in the special case when the block function is modeled as a *non-invertible* random permutation [Cza+17; Cza+19]. The case of invertible random permutations, which more accurately models the construction underlying SHA-3, has so far remained a major open problem.




Figure 1: The single-round sponge.

Single-round sponge hashing. In the special case when there is only a single round of absorption and squeezing, the sponge function $\text{Sp}^\varphi : \{0, 1\}^r \rightarrow \{0, 1\}^r$ has a simple form which is illustrated in Figure 1; namely, on input $x \in \{0, 1\}^r$, the output is given by $y = \text{Sp}^\varphi(x)$, where y corresponds to the first r bits of $\varphi(x||0^c)$. In other words, Sp^φ is defined as the restriction of φ onto the first r bits of its output.

Despite its simplicity, the single-round sponge already features many of the technical difficulties that one encounters when analyzing the security of the (many-round) sponge—especially once we model φ as an invertible random permutation. In the post-quantum setting, this requires us to analyze quantum algorithms $\mathcal{A}^{\varphi, \varphi^{-1}}$ which are allowed to make superposition queries to both the “forward oracle” O_φ , as well as the “backward oracle” $O_{\varphi^{-1}}$. For example, one may ask:

- (One-wayness:) How many queries does a quantum algorithm $\mathcal{A}^{\varphi, \varphi^{-1}}$ need to invert Sp^φ on a randomly chosen input? In other words, given y which corresponds to the first r bits of $\varphi(x||0^c)$, for a random $x \sim \{0, 1\}^r$, how many queries does \mathcal{A} need to find x' such that the first r many bits of $\varphi(x'||0^c)$ equal y ?
- (Collision-resistance:) How many queries does a quantum algorithm $\mathcal{A}^{\varphi, \varphi^{-1}}$ need to find a collision for Sp^φ ? Here, \mathcal{A} needs to find a colliding pair x, x' such that the first r many bits of $\varphi(x||0^c)$ and $\varphi(x'||0^c)$ match.

To the best of our knowledge, there are currently no known (tight) query lower bounds for either of the two above problems—despite the fact that they concern the seemingly simple variant of single-round sponge hashing. While prior work by Zhandry [Zha21] does in fact show the strong notion of reset-indifferentiability of the single-round sponge with invertible permutations (which implicitly implies the hardness of the query problems above), the techniques do not seem to yield tight lower bounds and only apply to restricted parameter ranges (see Section 1.3). The lack of lower bounds is mainly due to the fact that current proof techniques for the non-invertible case [Cza+17; Cza+19] seem to break down once the inverse oracle $O_{\varphi^{-1}}$ enters the picture.

Consider, for example, the basic property of *zero-preimage-resistance*; that is, the hardness of finding a pre-image of the all-zero string. Czajkowski et al. [Cza+17] showed that if a (non-invertible) φ is both collision-resistant (when restricted to the left and right half of its output) and zero-preimage-resistant (when restricted to the right half of its output), then the (many-round) sponge construction is collision-resistant. Unfortunately, this approach fails immediately if an adversary can submit inverse queries to the block function φ . Clearly, the restriction of φ to the right half of its output is no longer zero-preimage-resistant—an adversary can simply query $O_{\varphi^{-1}}$ on inputs of the form $y||0^c$, for any input $y \in \{0, 1\}^r$, in order to find a zero-preimage.

Another difficulty in generalizing existing proof techniques lies in the fact it is currently not known how to construct *compressed permutation oracles* [Unr21; Unr23]. Zhandry [Zha18] previously introduced compressed oracles as a tool to prove query lower bounds for problems involving random oracles; specifically, as a means to “record” quantum queries to a random oracle. Compressed oracles have proven to be extremely useful in analyzing the post-quantum security of hash functions in the quantum random oracle model, particularly in the case of random sponges [Cza+19; Unr21; Unr23]. However, unlike random functions where all the function outputs are sampled independently, this is not true for random permutations. This makes it difficult to construct compressed oracles for permutations, especially for modelling algorithms that can query the permutation in both directions [Unr21; Unr23].

Double-sided zero-search. In light of the difficulty in proving the post-quantum security of the sponge with invertible permutations, Unruh proposed the following simple conjecture which seems beyond the scope of current techniques.

Conjecture 1 (Double-sided zero-search, [Unr21; Unr23]). *Any quantum algorithm $A^{\varphi, \varphi^{-1}}$ which has quantum query-access to a random permutation $\varphi : \{0, 1\}^{2n} \rightarrow \{0, 1\}^{2n}$ and its inverse φ^{-1} must make $\Omega(2^{n/2})$ many queries to find a zero pair (x, y) such that $\varphi(x||0^n) = y||0^n$ with constant success probability.*

Note that the above lower bound is already tight—a simple Grover search allows one to find a zero-pair using $\Theta(2^{n/2})$ many queries, if one exists. Notice that there is a slight subtlety that arises in the statement of Conjecture 1 since not every permutation has a zero pair; for example, the permutation $\varphi(x||y) := (x||y \oplus 1^n)$ has no zero pairs since $\varphi(x||0^n) = x||1^n$, for any $x \in \{0, 1\}^n$, whereas the identity permutation $\varphi := \text{id}$ has precisely 2^n zero pairs. However, for a random permutation there is at least one zero pair with constant probability—a property we show in Lemma 3.5.

While it is not immediately clear how the double-sided zero-search problem from Conjecture 1 relates to the security of the sponge construction, it was suggested by Unruh [Unr21] that it may already provide some evidence for its collision-resistance. The problem has so far also resisted attempts¹ at solving it using standard techniques from quantum query complexity; for example, the adversary method [Amb02] or the polynomial method [Bea+01]. Therefore, a resolution to Conjecture 1 may already offer interesting new insights into proving the post-quantum security of the sponge construction with invertible permutations.

1.1 Our contributions

We now give an overview of our contributions in this work.

Resolving Unruh’s conjecture. We prove the double-sided zero-search conjecture (Conjecture 1) due to Unruh [Unr21; Unr23] and show that finding zero-pairs in

¹This was pointed out by Unruh [Unr21].

a random $2n$ -bit permutation with constant probability requires at least $\Theta(2^{n/2})$ many queries. Specifically, we show the following in Theorem 4.1.

Theorem (Informal). *Any quantum algorithm for DOUBLE-SIDED ZERO-SEARCH that makes T queries to an invertible permutation succeeds with probability at most $O(T^2/2^n)$.*

In other words, any algorithm must make at least $T = \Omega(\sqrt{\epsilon 2^n})$ many queries in order to succeed with probability ϵ . As we observe in Corollary 4.2, this immediately yields a tight lower bound of $\Theta(2^{n/2})$ for constant success probability due to Grover, or more generally a tight bound of $\Theta(\sqrt{\epsilon 2^n})$ for any success probability $\epsilon > 0$.

Our proof takes place in two steps:

- **Reduction from worst-case unstructured search:** Our first insight lies in the fact that we can reduce UNSTRUCTURED SEARCH with K out of 2^n marked elements to a specific instance of DOUBLE-SIDED ZERO-SEARCH without any overhead; namely, one in which the $2n$ -bit permutation φ has exactly K zero pairs. Suppose we have quantum oracle access to a function

$$f : \{0,1\}^n \rightarrow \{0,1\} \quad \text{with} \quad |f^{-1}(1)| = K.$$

We now construct the following permutation $\varphi : \{0,1\}^{2n} \rightarrow \{0,1\}^{2n}$ such that, for any inputs $x, y \in \{0,1\}^n$, it holds that:

$$\varphi(x||y) = \begin{cases} x||y & \text{if } f(x) = 1 \\ x||(y \oplus 1^n) & \text{if } f(x) = 0. \end{cases}$$

Notice that while φ is not random, it has exactly K zero pairs on inputs of the form $x||0^n$ whenever x satisfies $f(x) = 1$. Furthermore, $\varphi = \varphi^{-1}$ and so backwards queries are not helpful—thus worst case hardness follows.

- **Worst-case to average-case reduction via symmetrization:** Our next insight lies in the fact that we can re-randomize any worst-case permutation φ with K zero pairs into an average-case permutation φ^{sym} with K zero pairs:
 1. Select two $2n$ -bit permutations σ, ω such that both preserve the property of ending in 0^n , but are otherwise independently random.²
 2. Symmetrize φ by letting $\varphi^{\text{sym}} = \omega \circ \varphi \circ \sigma$.

A zero pair of φ can be reconstructed from a zero pair of φ^{sym} , but the location of the zero pairs and the permutation behaviour everywhere else is now randomized. We have inverse access to φ , and so we also have inverse access to φ^{sym} . Hardness of the average case is now implied by hardness of the worst case.

Our symmetrization argument (formally shown in Lemma 3.12) uses insights from the theory of *Young subgroups*. We give an extensive treatment of the subject in Section 3.

In Section 4.2, we give an alternative proof of Unruh’s double-sided zero-search conjecture (Conjecture 1) which combines our technique of symmetrization with a more conventional approach rooted in one-way to hiding [AHU18]. The advantage

²More formally, we require that σ and ω map strings of the form $(x||0^n)$ to $(y||0^n)$, for $x, y \in \{0,1\}^n$.

of our alternative proof is that it is more direct—it avoids the two-step template of the previous proof. At a high level, it uses the superposition oracle framework [Amb02] and introduces a *function register* which is outside of the view of the query algorithm, and contains a uniform superposition over permutations. Note that the superposition oracle formalism makes it especially easy to analyze quantum query-algorithms which query permutations both in the forward, as well as in the backward direction.

Once we introduce the function register, we can then “symmetrize it in superposition”. This framework allows us to directly show that no quantum algorithm can distinguish whether it is querying a random invertible permutation with exactly K zero pairs, or a random invertible permutation with no zero pairs—unless it makes a large number of queries. We remark, however, that contrary to the work of Zhandry [Zha18], we do not need to “compress” the superposition oracles and use inefficient representations instead.

Subset pairs and double-sided search. Motivated by the sponge construction, we generalize these techniques to handle search problems which are non-uniform in the sense of (1) different constraints on preimages vs images and (2) non-uniform distribution on permutations. In particular, we define a variant of two-sided search in which one is given query access to an n -bit permutation $\varphi : \{0,1\}^n \rightarrow \{0,1\}^n$ and its inverse, for an integer $n = r + c$, and asked to produce inputs that end in c many zeros and whose outputs begin in r many zeros. We further consider a class of non-uniform distributions on permutations that weights permutations according to their number of solutions. Notably, under this new distribution a solution is always guaranteed to exist, avoiding a subtlety in the original uniform problem. We show that these modifications do not make the problem significantly easier in Theorem 4.7.

Theorem (Informal). *Any quantum algorithm for NON-UNIFORM DOUBLE-SIDED SEARCH that makes at most T queries to an invertible (non-uniform) random permutation and succeeds with probability $\epsilon > 0$ satisfies $\epsilon = O(T^2/2^{\min(r,c)})$.*

The argument mirrors that of the uniform case, though with a slightly more complicated worst-case instance.

Quantum one-wayness of the single-round sponge. As an application of our techniques, we give a reduction from our NON-UNIFORM DOUBLE-SIDED SEARCH problem to the one-wayness (formally defined in Definition 5.1) of the single-round sponge. As a corollary of our bound for NON-UNIFORM DOUBLE-SIDED SEARCH, in Theorem 5.2 we establish the first arbitrary parameter post-quantum security result for sponge hashing in the random permutation model.

Theorem (Informal). *Any T -query quantum algorithm that breaks the quantum one-wayness of the single-round sponge, where the block function is instantiated with an invertible random permutation, has a success probability of at most $\epsilon = O(T^2/2^{\min(r,c)})$.*

By modeling the block function as an invertible random permutation, our theorem shows the post-quantum security of the single-round sponge in the quantum random

oracle model [Bon+11]. To prove our main theorem, we first switch to an alternative but equivalent characterization of one-wayness of the sponge (Lemma 5.3), where the image is chosen uniformly at random before the actual block function is sampled (in this case, φ). Note that this approach, however, results in a non-uniform distribution over φ . To complete the proof, we then give a reduction from the NON-UNIFORM DOUBLE-SIDED SEARCH problem (formally defined in Definition 5.1).

Combinatorics of subset-pairs. Along the way, we work out the combinatorics of the expected number of input-output pairs of any given kind (i.e., for a general *subset-pair*), as well as strong tail bounds in the case where the expected number is small. We show that the number of such pairs decays exponentially in this case for both the uniform distribution in Theorem 3.15 and a certain non-uniform one in Theorem 3.20.

Theorem (Informal). *Whenever the average number of subset pairs of a random permutation is a constant, the probability that there are more than K subset pairs is at most $\exp(-\Omega(K))$.*

In the uniform case, the number of subset pairs of permutations follows a so-called *hypergeometric distribution*—a well studied distribution corresponding to sampling without replacement. We relate our non-uniform distribution, in which permutations are favored proportional to their number of subset pairs, to the uniform distribution, and show how to derive strong tail bounds in this case as well.

1.2 Open questions

This work does not address post-quantum security of the many-round sponge, which is a natural and important next direction. In particular, we do not address whether the many-round sponge when instantiated with an invertible random permutation is one-way. However, we believe that our techniques could potentially open up new directions towards proving the post-quantum security of the many-round sponge, especially when combined with insights from the non-invertible case [Cza+17].

Another natural question is whether one can show stronger security properties of the sponge with invertible permutations, even in the single-round case. For example, whether the single-round sponge is *collision-resistant* (or more generally, *collapsing*) when the block function is instantiated with an invertible random permutation. This is not addressed by our work. However, in this case, it is conceivable that there exist direct reductions from collision-resistance (or collapsing) to natural query problems involving invertible permutations, similar to our reduction for one-wayness. In such a case, our techniques would be well suited to analyzing this type of problem.

1.3 Related work

Czajkowski, Bruinderink, Hülsing, Schaffner and Unruh [Cza+17] showed that the many-round sponge is *collapsing*—a strengthening of collision-resistance—in the case when the block function is modeled as a random function or a random (non-invertible) permutation. In the same model, Czajkowski, Majenz, Schaffner and Zur [Cza+19] proved the quantum indistinguishability of the (many-round) sponge. However, in

contrast with the former work, the latter relies on Zhandry’s compressed oracle technique [Zha18]. Inspired by Zhandry’s compressed oracle method, Rosmanis [Ros22] recently also gave tight bounds for the permutation inversion problem. We remark that the results in the aforementioned works do not apply in the case where the adversary has access to an invertible random permutation. Zhandry [Zha21] showed that the single-round sponge (in the special case when the message length is roughly half the block length of the permutation) is quantumly indifferentiable from a random oracle (even if the adversary has access to the inverse of the permutation). Because indistinguishability implies that the function behaves “just like a random oracle”, this immediately implies many desirable properties for this particular parameter range, such as one-wayness and collision-resistance. Our results, however, are incomparable to those of Zhandry. While [Zha21] obtains a much more general result, the advantage of our techniques is that we do not require any constraints on the message/block length. Moreover, the restriction to half the block length in [Zha21] is necessary, as the strong form of *reset-indifferentiability* which is achieved is not possible if the message exceeds half the block length. Finally, we remark that our techniques yield tight lower bounds for the one-wayness of the single-round sponge, whereas the techniques in [Zha21] seem insufficient for this purpose. Unruh proposed the notion of *compressed permutation oracles* [Unr21; Unr23] which are aimed at dealing with invertible permutations. However, the soundness of the construction is still unresolved and currently relies on a conjecture. Unruh also put forward the double-sided zero-search conjecture [Unr21; Unr23] as a simple and non-trivial query problem in the context of invertible random permutations which may potentially have implications for collision-resistance. Alagic, Bai, Poremba and Shi [Ala+23] showed space-time trade-offs for two-sided permutation inversion—the task of inverting a random but invertible permutation on a random input. While the inverter does in fact receive access to an inverse oracle, it is crucial that such an oracle is *punctured* at the challenge input; otherwise, the inversion task would become trivial. Although this work does not seem to have any immediate implications for the single-round sponge, it does suggest that permutation inversion remains hard in the presence of inverse oracles.

Concurrent work. After completing our work, we became aware of a different manuscript on the single-round sponge by Majenz, Malavolta and Walter [MMW24] which was prepared independently. They develop a generalization of Zhandry’s compressed oracle technique in the setting where the query algorithm has access to a permutation and its inverse, and they also obtain query lower bounds for both sponge inversion and Unruh’s double-sided zero-search. While their framework also applies to permutation inversion with respect to arbitrary relations on (single) input-output pairs, it does not yield tight query lower bounds as in our work.

Acknowledgments. We thank Gorjan Alagic, Chen Bai, Luke Schaeffer and Kaiyan Shi for many useful discussions. JC acknowledges funding from the Department of Defense. AP is supported by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Co-design Center for Quantum Advantage (C2QA) under contract number DE-SC0012704.

1.4 Organization

In Section 2, we review the fundamentals of quantum computing and quantum search lower bounds. In Section 3, we review basic group theory—with a particular focus on the theory of Young subgroups—and state the key group-theoretic lemma for our symmetrization argument. We also introduce the concept of a *zero pair* as well as a *subset pair*, which constitute the main objects in the query problems we study. Moreover, we analyze many general combinatorial properties of these objects.

In Section 4, we prove tight lower bounds for the DOUBLE-SIDED ZERO-SEARCH problem using a worst-case to average-case reduction, as well as a direct proof based on the superposition oracle framework. We also define and show lower bounds for a non-uniform variant of this problem which we call NON-UNIFORM DOUBLE-SIDED SEARCH—a problem which we naturally encounter later in Section 5. In the final section, we give a reduction from NON-UNIFORM DOUBLE-SIDED SEARCH to the one-way game of the single-round sponge with invertible permutations.

2 Preliminaries

Let us first introduce some basic notation and relevant background.

Notation. For $N \in \mathbb{N}$, we use $[N] = \{1, 2, \dots, N\}$ to denote the set of integers up to N . The symmetric group on $[N]$ is denoted by S_N . In slight abuse of notation, we oftentimes identify elements $x \in [N]$ with bit strings $x \in \{0, 1\}^n$ via their binary representation whenever $N = 2^n$ and $n \in \mathbb{N}$. Similarly, we identify permutations $\pi \in S_N$ with permutations $\pi : \{0, 1\}^n \rightarrow \{0, 1\}^n$ over bit strings of length n .

We write $\text{negl}(\cdot)$ to denote any *negligible* function, which is a function f such that, for every constant $c \in \mathbb{N}$, there exists an integer N such that for all $n > N$, $f(n) < n^{-c}$.

Probability theory. The notation $x \sim X$ describes that an element x is drawn uniformly at random from the set X . Similarly, if \mathcal{D} is a distribution, we let $x \sim \mathcal{D}$ denote sampling x according to \mathcal{D} . We denote the expectation value of a random variable X by $\mathbb{E}[X] = \sum_x x \Pr[X = x]$. We make use of the following distribution.

Definition 2.1 (Hypergeometric distribution). *A hypergeometric distribution describes the number of marked elements found by drawing T many objects from a total of N , of which K are marked. More formally, for $X \sim \text{Hypergeometric}(N, K, T)$, we have*

$$\Pr[X = k] = \frac{\binom{K}{k} \binom{N-K}{T-k}}{\binom{N}{T}}.$$

We use the following tail bound for the hypergeometric distribution.

Lemma 2.2 (Hoeffding’s inequality, [Hoe63; Chv79]). *Let $X \sim \text{Hypergeometric}(N, K, T)$ and define $p = K/N$. Then, for any $t \in \mathbb{R}$ with $0 \leq t < 1 - p$, it holds that*

$$\Pr[X \geq (p + t) \cdot T] \leq \exp(-T \cdot D_{\text{KL}}(p + t || p)),$$

where $D_{\text{KL}}(q||p)$ denotes the Kullback-Leibler divergence with

$$D_{\text{KL}}(q||p) = q \cdot \ln \left(\frac{q}{p} \right) + (1 - q) \cdot \ln \left(\frac{1 - q}{1 - p} \right).$$

Quantum computing. A finite-dimensional complex Hilbert space is denoted by \mathcal{H} , and we use subscripts to distinguish between different systems (or registers); for example, we let \mathcal{H}_A be the Hilbert space corresponding to a system A . The tensor product of two Hilbert spaces \mathcal{H}_A and \mathcal{H}_B is another Hilbert space which we denote by $\mathcal{H}_{AB} = \mathcal{H}_A \otimes \mathcal{H}_B$. We let $\mathcal{L}(\mathcal{H})$ denote the set of linear operators over \mathcal{H} . A quantum system over the 2-dimensional Hilbert space $\mathcal{H} = \mathbb{C}^2$ is called a *qubit*. For $n \in \mathbb{N}$, we refer to quantum registers over the Hilbert space $\mathcal{H} = (\mathbb{C}^2)^{\otimes n}$ as n -qubit states. We use the word *quantum state* to refer to both pure states (unit vectors $|\psi\rangle \in \mathcal{H}$) and density matrices $\rho \in \mathcal{D}(\mathcal{H})$, where we use the notation $\mathcal{D}(\mathcal{H})$ to refer to the space of positive semidefinite linear operators of unit trace acting on \mathcal{H} .

The *trace distance* between two density matrices $\rho, \sigma \in \mathcal{D}(\mathcal{H})$ is given by

$$\text{TD}(\rho, \sigma) = \frac{1}{2} \text{Tr} \left[\sqrt{(\rho - \sigma)^\dagger (\rho - \sigma)} \right].$$

Note that the trace distance between two pure states $|\psi\rangle, |\phi\rangle \in (\mathbb{C}^2)^{\otimes n}$ satisfies

$$\text{TD}(|\psi\rangle\langle\psi|, |\phi\rangle\langle\phi|) \leq \| |\psi\rangle - |\phi\rangle \|,$$

where $\| \cdot \|$ denotes the Euclidean distance over the vector space $(\mathbb{C}^2)^{\otimes n}$.

When the dimensions are clear from context, we use $\mathbb{1}$ to denote the identity matrix. A *unitary* $U : \mathcal{L}(\mathcal{H}_A) \rightarrow \mathcal{L}(\mathcal{H}_A)$ is a linear operator such that $U^\dagger U = U U^\dagger = \mathbb{1}_A$. A *projector* Π is a Hermitian operator such that $\Pi^2 = \Pi$. Oftentimes, we use the shorthand notation $\bar{\Pi} = \mathbb{1} - \Pi$.

A quantum algorithm is a uniform family of quantum circuits $\{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$, where each circuit \mathcal{A}_λ is described by a sequence of unitary gates and measurements; moreover, for each $\lambda \in \mathbb{N}$, there exists a deterministic Turing machine that, on input 1^λ , outputs a circuit description of \mathcal{A}_λ . We say that a quantum algorithm \mathcal{A} has oracle access to a classical function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$, denoted by \mathcal{A}^f , if \mathcal{A} is allowed to use a unitary gate O^f at unit cost in time. The unitary O^f acts as follows on the computational basis states of a Hilbert space $\mathcal{H}_X \otimes \mathcal{H}_Y$ of $n + m$ qubits:

$$O^f : |x\rangle_X \otimes |y\rangle_Y \longrightarrow |x\rangle_X \otimes |y \oplus f(x)\rangle_Y,$$

where the operation \oplus denotes bit-wise addition modulo 2. Oracles with quantum query-access have been studied extensively, for example in the context of quantum complexity theory [Ben+97], as well as in cryptography [Bon+11; AHU18; Ala+20].

2.1 Lower bounds for unstructured search

The quantum query complexity of worst-case UNSTRUCTURED SEARCH is one of the most fundamental results in quantum computing [Gro96; Ben+97; Boy+98]. In this

problem, a quantum algorithm is given query access to a function $f : [N] \rightarrow \{0, 1\}$ and is tasked with finding an $x \in [N]$ such that $f(x) = 1$. We will consider this problem subject to the promise that there are exactly K preimages of 1, i.e., $|f^{-1}(1)| = K$. Note also that only ever consider the worst-case complexity of this problem, meaning the success probability is only over the internal randomness of the algorithm.

The exact query complexity of this problem is well-understood [Gro96; Za199; DH09]. We use the following lower bound, which is a direct corollary of the work of Dohotaru and Høyer.

Theorem 2.3 ([DH09], Theorem 8). *Any algorithm which solves UNSTRUCTURED SEARCH with K out of N marked elements with worst-case success probability $\epsilon > 0$ requires at least*

$$T \geq \frac{\sqrt{N/K}}{2\sqrt{2}} \left(1 + \sqrt{\epsilon} - \sqrt{1 - \epsilon} - \frac{2}{\sqrt{N/K}} \right)$$

many queries.

We can rearrange the expression in Theorem 2.3 to obtain a bound on the worst-case success probability of a T -query quantum algorithm.

Corollary 2.4. *Any algorithm which solves UNSTRUCTURED SEARCH with K marked elements out of N with worst-case success probability $\epsilon > 0$, satisfies the inequality*

$$\epsilon \leq \frac{8(T+1)^2 K}{N}.$$

Proof. We have

$$\begin{aligned} T &\geq \frac{\sqrt{N/K}}{2\sqrt{2}} \left(1 + \sqrt{\epsilon} - \sqrt{1 - \epsilon} - \frac{2}{\sqrt{N/K}} \right) && \text{(Theorem 2.3)} \\ &\geq \frac{\sqrt{N/K}}{2\sqrt{2}} \left(\sqrt{\epsilon} - \frac{2}{\sqrt{N/K}} \right). \end{aligned}$$

This implies that $T + 1 \geq \sqrt{N\epsilon/K}/\sqrt{8}$, which yields the desired inequality. \square

This bound is no longer exact, but will suffice for our purposes.

3 Subset Pairs and the Symmetric Group

In this section, we recall some basic facts from the theory of Young subgroups. Due to the nature of the query problems in Section 4, we oftentimes encounter subgroups of the symmetric group S_N consisting of permutations that fix certain subsets of $[N]$; for example, zero pairs in the case of DOUBLE-SIDED ZERO-SEARCH in Problem 1. This forces us to analyze certain subgroups of S_N —called Young subgroups—which are direct products of subgroups on fewer elements that form a partition of $[N]$.

We use the machinery of Young subgroups to prove a key lemma which shows our symmetrization procedure is sound, and to offer an algebraic characterization of the input/output pairs of a permutation satisfying certain constraints.

3.1 Group Theory

We first define relevant notions in group theory, and recall known results about the symmetric group S_N on N elements. For a more complete overview of the subject, we refer the reader to the work of James [Jam84].

Let S_N denote the symmetric group consisting of permutations which act on the set $[N] := \{1, \dots, N\}$. For a subset $A \subset [N]$, let S_A denote the maximal subgroup of S_N which fixes every element in the complement of A , i.e. $[N] \setminus A$. Now let A_1, \dots, A_l be a partition of $[N]$ such that the disjoint union satisfies $\bigsqcup_{i \in [l]} A_i = [N]$.

Definition 3.1 (Young subgroup). *A subgroup H of the symmetric group S_N is a Young subgroup if it can be expressed as $H = S_{A_1} \times \dots \times S_{A_l}$, where \times denotes the internal direct product and the collection of subsets $\{A_i\}_{i \in [l]}$ forms a partition of $[N]$.*

The concept of a double coset, which we review below, will also be relevant.

Definition 3.2 (Double cosets). *Let H, K be subgroups of a group G . The double cosets of G under (H, K) , denoted $H \setminus G / K$, are the sets of elements which are invariant under left multiplication by H and right multiplication by K . In particular,*

$$H \setminus G / K = \{ \{h x k : h \in H, k \in K\} : x \in G \}.$$

It is well-known that G is the disjoint union of its double cosets for any subgroups $H, K \leq G$. We focus on the double cosets of the symmetric group S_N , specifically those generated by Young subgroups. These subgroups admit the following characterization, adapted from Jones [Jon96] and James [Jam84].

Theorem 3.3 ([Jon96], Theorem 2.2). *Let H, K be Young subgroups of S_N , with corresponding partitions A_1, \dots, A_l (for H) and B_1, \dots, B_m (for K). Let $\pi \in S_N$ and $C = H\pi K$ be the corresponding double coset. Any other permutation $\pi' \in S_N$ is in C if and only if for all $i \leq l, j \leq m$ we have $|A_i \cap \pi' B_j| = |A_i \cap \pi B_j|$.*

Intuitively, the above characterization says that the double cosets defined by Young subgroups (H, K) correspond to sets of permutations which look the same if one only considers how they distribute the elements of each B_j among the different A_i 's. Two permutations π, π' are in the same double coset if and only if for every A_i, B_j both π and π' send the same number of elements from B_j to A_i . This characterization will prove useful when showing the random self-reducibility of two-sided search problems. The following lemma and proof is a key component of this reduction, and is based heavily on the approach of Wildon [Wil], Theorem 4.4.

Lemma 3.4. *For any subgroups H, K of a (finite) group G with $x, g \in G$ both in the same (H, K) double coset, there are exactly $|x^{-1}Hx \cap K|$ ways of choosing $h \in H$ and $k \in K$ such that $g = h x k$.*

Proof. Fix some $x \in G$, and consider any $g \in G$ in the same double coset as x . Then, there exist elements $h \in H, k \in K$ such that

$$\begin{aligned} g &= h x k \\ &= x(x^{-1}h x)k. \end{aligned}$$

Any alternative $k' \in K$ can be written as $k' = u_k k$ for $u_k \in K$, and an alternative $h' \in H$ can be written as $x^{-1}h'x = (x^{-1}hx)u_h$ for $u_h \in x^{-1}Hx$. These new elements satisfy

$$\begin{aligned} g' &= h' x k' \\ &= x(x^{-1}hx)u_h u_k k \end{aligned}$$

It is apparent that $g' = g$ holds if and only if $u_h = u_k^{-1}$. This implies that

$$u_k, u_h \in K \cap x^{-1}Hx.$$

Thus, there are $|x^{-1}Hx \cap K|$ alternative ways to choose a satisfying $h \in H$ and $k \in K$. \square

3.2 Subset Pairs

In this section, we introduce the notion of a subset pair, which will be the main object of study. Specifically, we are interested in the complexity of finding input/output pairs of a permutation which lie within certain pairs of subsets.

Zero pairs. Before giving the general definition of a subset pair, we first define a zero pair as a specific and relevant example of a subset pair.

Example 3.1 (Zero pair). Let $Z \subset [2^{2n}]$ be the subset consisting of all $z \in [2^{2n}]$ with binary decomposition $(z_1, z_2, \dots, z_{2n}) \in \{0, 1\}^{2n}$ such that

$$z_{n+1} = z_{n+2} = \dots = z_{2n} = 0.$$

Then, Z is a subset of size 2^n and, when identifying $S_{2^{2n}}$ with permutations over $2n$ -bit strings, $(x||0^n)$ and $(y||0^n)$ form a Z -pair for a permutation $\pi : \{0, 1\}^{2n} \rightarrow \{0, 1\}^{2n}$, if $\pi(x||0^n) = y||0^n$. For brevity, we sometimes call (x, y) a zero pair of the permutation π . We show the following concerning zero pairs in random permutations.

Lemma 3.5. A random permutation $\pi \sim S_{2^{2n}}$ contains at least one zero pair with probability

$$\Pr_{\pi \sim S_{2^{2n}}} [|Z_\pi| \geq 1] = 1 - 1/e + o(1).$$

Here, $|Z_\pi|$ denotes the random variable for the number of zero pairs of the permutation π .

Proof. Let $N = 2^{2n}$ and let $Z \subset [N]$ be the subset from Example 3.1, where $|Z| = \sqrt{N}$. It suffices to estimate the probability that a random permutation $\pi \sim S_N$ has no zero pairs. The total number of permutations over $[N]$ which have no zero pairs is

$$\frac{(N - \sqrt{N})!}{(N - 2\sqrt{N})!} \cdot (N - \sqrt{N})! .$$

This is because we need to match all \sqrt{N} elements in the subset Z exclusively with elements outside of Z . Therefore, we get

$$\begin{aligned} \Pr_{\pi \sim S_N} [|Z_\pi| \geq 1] &= 1 - \Pr_{\pi \sim S_N} [|Z_\pi| = 0] \\ &= 1 - \frac{(N - \sqrt{N})! \cdot (N - \sqrt{N})!}{N! \cdot (N - 2\sqrt{N})!} \\ &= 1 - \binom{N - \sqrt{N}}{\sqrt{N}} / \binom{N}{\sqrt{N}}. \end{aligned}$$

Therefore, the claim follows from the fact that $\binom{N - \sqrt{N}}{\sqrt{N}} / \binom{N}{\sqrt{N}}$ is a monotonically increasing sequence which approaches the limit $1/e$ as $N \rightarrow \infty$. \square

We will use the notation G_Z to denote the subgroup $G_Z \leq S_{2n}$ of permutations which preserve membership in Z .

Definition 3.6 (Subgroup G_Z). *A permutation π is a member of the subgroup $G_Z \leq S_{2n}$ if and only if π fixes the subset Z such that $\pi(Z) = Z$. In other words, for every $x \in \{0, 1\}^n$, there exists $y \in \{0, 1\}^n$ such that $\pi(x||0^n) = y||0^n$.*

Note that the subgroup G_Z decomposes as $G_Z \cong S_Z \times S_{[N] \setminus Z}$. The algebraic structure of the double cosets induced by $G_Z \times G_Z$ will be central in proving the hardness of DOUBLE-SIDED ZERO-SEARCH in Problem 1.

General statements. We now turn to defining a subset pair in full generality, and prove a key “symmetrization lemma” using the theory of Young subgroups.

Definition 3.7 (Subset pair). *Let $N \in \mathbb{N}$, let $X_1, X_2 \subseteq [N]$ be subsets and $X = (X_1, X_2)$. We say that a pair (i, j) is an X -pair for a permutation $\pi \in S_N$ if $i \in X_1, j \in X_2$ and $\pi(i) = j$. We denote the set of all X -pairs of a permutation $\pi \in S_N$ by*

$$X_\pi = \{(i, j) \mid i \in X_1, j \in X_2 \text{ and } \pi(i) = j\}.$$

Remark 3.8. *Whenever $X = (X_1, X_2)$ is a pair of identical subsets with $X_1 = X_2 = Z$, for some $Z \subseteq [N]$, we sometimes refer to the pairs (i, j) in Definition 3.7 as Z -pairs instead. This is the case for the zero pair subsets from Example 3.1.*

For an $X = (X_1, X_2)$, define the subgroups $G_1 \leq S_N$ and $G_2 \leq S_N$ with

$$\begin{aligned} G_1 &:= \{\pi \in S_N : \pi(X_1) = X_1\} \\ G_2 &:= \{\pi \in S_N : \pi(X_2) = X_2\}. \end{aligned}$$

In other words, G_i consists of all permutations in S_N that fix X_i . Moreover, it follows from Definition 3.2 that G_1 and G_2 are Young subgroups of S_N .

Fact 3.9. *The Young subgroups $G_1 \leq S_N$ and $G_2 \leq S_N$ decompose as*

$$G_i \cong S_{X_i} \times S_{[N] \setminus X_i}, \quad \text{for } i \in \{1, 2\}.$$

We introduce the following notation for the double cosets of S_N under (G_1, G_2) .

Definition 3.10. For $\kappa \in \mathbb{N} \cup \{0\}$, and $X = (X_1, X_2)$ with $X_1, X_2 \subset [N]$, define the subset

$$S_N^{\kappa, X} := \{\pi \in S_N : \pi \text{ has exactly } \kappa \text{ many } X\text{-pairs}\}.$$

We oftentimes omit the superscript X when it is clear from context. Note that $S_N = \bigsqcup_{\kappa=0}^{|X_1|} S_N^{\kappa}$, i.e., S_N is the disjoint union over all the S_N^{κ} . Note, however, that whenever we have $\kappa > |X_1|$ or $\kappa > |X_2|$, it must be that $S_N^{\kappa} = \emptyset$. The following corollary comes from the characterization in Theorem 3.3, since G_1 and G_2 are both Young subgroups.

Corollary 3.11. The sets $\{S_N^{\kappa}\}$ are the double cosets of S_N under (G_1, G_2) ; formally,

$$G_1 \backslash S_N / G_2 = \{S_N^{\kappa} \mid \kappa \in \{0, 1, \dots, |X_1|\}\}.$$

Proof. From Theorem 3.3, a permutation φ 's double coset is determined by the data $(|\varphi(X_1) \cap X_2|, |\varphi(X_1) \cap ([N] \setminus X_2)|, |\varphi([N] \setminus X_1) \cap X_2|, |\varphi([N] \setminus X_1) \cap ([N] \setminus X_2)|)$. Note that once the number of X -pairs (equivalently, $|\varphi(X_1) \cap X_2|$) is fixed the remaining quantities are determined. In particular, we have

$$\begin{aligned} |\varphi(X_1) \cap ([N] \setminus X_2)| &= |X_1| - |\varphi(X_1) \cap X_2|, \\ |\varphi([N] \setminus X_1) \cap X_2| &= |X_2| - |\varphi(X_1) \cap X_2|, \\ |\varphi([N] \setminus X_1) \cap ([N] \setminus X_2)| &= N - |X_1| - |X_2| + |\varphi(X_1) \cap X_2|, \end{aligned}$$

and hence the number of X pairs uniquely determines φ 's double coset. \square

We are now ready to prove the main tool for analyzing random double-sided search problems, such as DOUBLE-SIDED ZERO-SEARCH in Problem 1.

Lemma 3.12 (Symmetrization lemma). Consider any fixed permutation $\pi \in S_N^{\kappa}$, for some $\kappa \in \mathbb{N} \cup \{0\}$, and suppose that $\omega \sim G_1, \sigma \sim G_2$ are sampled independently and uniformly at random. Then, $\omega \circ \pi \circ \sigma$ is a uniformly random element of S_N^{κ} .

Proof. From Lemma 3.4, it follows that for any fixed $\pi \in S_N^{\kappa}$ and some $\varphi \in S_N^{\kappa}$ the number of ω, σ pairs with $\omega \in G_1, \sigma \in G_2$ such that $\varphi = \omega \circ \pi \circ \sigma$ is $|\pi G_1 \pi^{-1} \cap G_2|$, which notably is independent of φ . In particular, we have

$$\Pr_{\omega \sim G_1, \sigma \sim G_2} [\omega \circ \pi \circ \sigma = \varphi] = \frac{|\pi G_1 \pi^{-1} \cap G_2|}{\sum_{\varphi' \in S_N^{\kappa}} |\pi G_1 \pi^{-1} \cap G_2|} = \frac{1}{|S_N^{\kappa}|}.$$

The probability is the same for all elements $\varphi \in S_N^{\kappa}$, and so the element $\omega \circ \pi \circ \sigma$ is uniformly random among elements of S_N^{κ} . \square

3.3 Combinatorics of Subset Pairs

In this section, we work out the combinatorics behind the expected number of subset pairs (on average over the choice of permutation), as well derive strong tail bounds in the case where the expected number is small. We show that the number of subset pairs decays exponentially for the uniform distribution (see Theorem 3.15), as well as for a particular choice of non-uniform distribution (see Theorem 3.20).

Uniform case. We begin by considering the case of uniformly random $\pi \sim S_N$ and derive the expected number of X -pairs as well as tail bounds. We show the following:

Theorem 3.13. *Let $N \in \mathbb{N}$ and let $X_1, X_2 \subseteq [N]$ be subsets with $X = (X_1, X_2)$. Then, on average over the uniform choice of $\pi \sim S_N$, the expected number of X -pairs equals*

$$\mathbb{E}_{\pi \sim S_N} [|X_\pi|] = \frac{|X_1||X_2|}{N}.$$

Proof. For a uniform random $\pi \sim S_N$, we let I_x denote the indicator variable which equals 1, if $\pi(x) \in X_1$, and 0 otherwise. Thus, we have

$$\mathbb{E}[I_x] = \frac{|X_2|}{N}, \quad |X_\pi| = \sum_{x \in X_1} I_x.$$

From the above, it then follows that

$$\begin{aligned} \mathbb{E}_{\pi \sim S_N} [|X_\pi|] &= \sum_{x \in X_1} \mathbb{E}[I_x] && \text{(Linearity of Expectation)} \\ &= \frac{|X_1||X_2|}{N}. \end{aligned}$$

□

A simple corollary of this bound is an alternative characterization of the total number of subset pairs over all permutations.

Corollary 3.14. *Let $N \in \mathbb{N}$ and fix arbitrary subsets $X_1, X_2 \subseteq [N]$ with $X = (X_1, X_2)$. Then, there exist precisely*

$$\sum_{\pi \in S_N} |X_\pi| = \frac{|X_1||X_2|}{N} \cdot N!$$

many X -pairs over all permutations $\pi \in S_N$.

Proof. This is easily verified as follows:

$$\begin{aligned} \sum_{\pi \in S_N} |X_\pi| &= \mathbb{E}_{\pi \sim S_N} [|X_\pi|] \cdot |S_N| \\ &= \frac{|X_1||X_2|}{N} \cdot N!. \end{aligned} \quad \text{(Theorem 3.13)}$$

□

We now state and prove tail bounds, which will be sufficiently tight for the case where the expected number of subset pairs is small.

Theorem 3.15. *Let $N \in \mathbb{N}$ and $X_1, X_2 \subseteq [N]$ be subsets with $X = (X_1, X_2)$. Then, for any real number $u \geq 6 \mathbb{E}_{\sigma \sim S_N} [|X_\sigma|]$, it holds that*

$$\Pr_{\pi \sim S_N} \left[|X_\pi| \geq \mathbb{E}_{\sigma \sim S_N} [|X_\sigma|] + u \right] \leq \exp \left(-\frac{3}{4}u \right).$$

Proof. The number of permutations having κ many X -pairs for $0 \leq \kappa \leq |X_1|, |X_2|$ is hypergeometrically distributed (see Definition 2.1). In other words,

$$\Pr_{\pi \sim S_N} [|X_\pi| = \kappa] = \frac{\binom{|X_2|}{\kappa} \binom{N-|X_2|}{|X_1|-\kappa}}{\binom{N}{|X_1|}}.$$

This is because each image of X_1 can be assigned without replacement by a random permutation to one of N images. Next, we want to count how many images there are in X_2 , which we interpret as a “success”. Let us consider the hypergeometric distribution $\text{Hypergeometric}(N, K, T)$ with the following parameters:

$$\begin{aligned} N &:= N && \text{(Total Objects)} \\ K &:= |X_2| && \text{(Success Objects)} \\ T &:= |X_1| && \text{(Trials)} \end{aligned}$$

Let $p = |X_2|/N$ denote the probability that the first draw is a success. Applying Hoeffding’s inequality (see Lemma 2.2), we get

$$\begin{aligned} \Pr_{\pi \sim S_N} [|X_\pi| \geq (p+t)|X_1|] &\leq \exp(-|X_1| D_{\text{KL}}(p+t||p)) \\ &\leq \exp\left(-\frac{3}{4}|X_1|t\right) \quad \text{(From Lemma 3.16).} \end{aligned}$$

By substituting $u = |X_1|t$ and applying Theorem 3.13, we obtain

$$\Pr_{\pi \sim S_N} [|X_\pi| \geq \mathbb{E}_{\sigma \sim S_N} [|X_\sigma|] + u] \leq \exp\left(-\frac{3}{4}u\right) \quad \text{(When } u \geq 6 \mathbb{E}_{\sigma \sim S_N} [|X_\sigma|]).$$

□

We use the following technical lemma.

Lemma 3.16. *Let $p, t \in (0, 1)$ such that $t \geq 6p$ and $p+t < 1$. Then, we can lower bound the Kullback-Leibler divergence as follows:*

$$D_{\text{KL}}(p+t||p) > 3t/4.$$

Proof. Recall that the Kullback-Leibler divergence $D_{\text{KL}}(p+t||p)$ is defined as

$$D_{\text{KL}}(p+t||p) = \underbrace{(p+t) \ln\left(\frac{p+t}{p}\right)}_{=D_1} + \underbrace{(1-p-t) \ln\left(\frac{1-p-t}{1-p}\right)}_{=D_2}.$$

We now bound D_1 and D_2 separately. Beginning with the first term, we get

$$D_1 > t \ln\left(\frac{p+t}{p}\right) > t \ln\left(\frac{t}{p}\right) \geq t \ln(6),$$

where we used that $t \geq 6p$. We now continue with the second term. First, we consider the following first-order derivatives with respect to t :

$$\begin{aligned}\frac{\partial}{\partial t} (D_2) &= -\ln\left(\frac{p+t-1}{p-1}\right) - 1 \\ \frac{\partial}{\partial t} (D_2) \Big|_{t=0} &= -1.\end{aligned}$$

Calculating the second-order derivatives, we also find that

$$\begin{aligned}\frac{\partial^2}{\partial t^2} (D_2) &= -\frac{1}{p+t-1} \\ \frac{\partial^2}{\partial t^2} (D_2) &> 0.\end{aligned}\quad (\text{From } p+t < 1).$$

It follows that D_2 is concave upwards in the relevant regime, and the derivative at $t = 0$ is -1 . We further have $D_2|_{t=0} = 0$, so when $p+t < 1$ and $t \geq 0$ we have $D_2 > -t$. Combining these results, we obtain the desired inequality:

$$D(p+t||p) > t \ln(6) - t > 3t/4.$$

□

Non-uniform case. In this section, we work out the combinatorics for the number of subset pairs with respect to a particular *non-uniform* distribution which assigns more weight to permutations that have a large amount of X -pairs.

Definition 3.17. Let $N \in \mathbb{N}$. The non-uniform distribution \mathcal{D}_X over permutations in S_N is parameterized by a pair of subsets $X = (X_1, X_2)$, and is defined as follows:

$$\Pr_{\Phi \sim \mathcal{D}_X} [\Phi = \varphi] = \frac{|X_\varphi|}{\sum_{\sigma \in S_N} |X_\sigma|}, \quad \text{for } \varphi \in S_N.$$

When X is clear from the context, we sometimes denote \mathcal{D}_X by \mathcal{D} . We begin by bounding the average number of X -pairs.

Theorem 3.18. Let $N \in \mathbb{N}$ and $X_1, X_2 \subset [N]$ be non-empty subsets. Let $|X_\pi|$ denote the number of X -pairs with respect to $\pi \in S_N$. Then, the average number of X -pairs over the choice of permutation from $\pi \sim \mathcal{D}_X$ satisfies

$$1 \leq \mathbb{E}_{\pi \sim \mathcal{D}_X} [|X_\pi|] \leq 1 + \frac{|X_1||X_2|}{N}.$$

Proof. Recall that the subsets $\{S_N^\kappa\}$ form a partition $S_N = \bigsqcup_{\kappa=0}^{|X_1|} S_N^\kappa$, i.e., S_N is the disjoint union over all the S_N^κ , we have $N! = \sum_{\kappa=0}^{|X_1|} |S_N^\kappa|$, and thus

$$\mathbb{E}_{\sigma \sim S_N} [|X_\sigma|] = \sum_{\kappa=0}^{|X_1|} \kappa \cdot \Pr_{\sigma \sim S_N} [\sigma \in S_N^\kappa] = \sum_{\kappa=0}^{|X_1|} \kappa \cdot \frac{|S_N^\kappa|}{N!}.$$

Therefore, we obtain the following identity,

$$\sum_{\sigma \in S_N} |X_\sigma| = \sum_{\kappa=0}^{|X_1|} \sum_{\sigma \in S_N^\kappa} |X_\sigma| = \sum_{\kappa=0}^{|X_1|} \kappa \cdot |S_N^\kappa| = N! \cdot \mathbb{E}_{\sigma \sim S_N} [|X_\sigma|]. \quad (1)$$

We can now write the average number of X -pairs as follows:

$$\begin{aligned} \mathbb{E}_{\pi \sim \mathcal{D}_X} [|X_\pi|] &= \sum_{\kappa=0}^{|X_1|} \kappa \cdot \Pr_{\pi \sim \mathcal{D}_X} [\pi \in S_N^\kappa] \\ &= \sum_{\kappa=0}^{|X_1|} \kappa \cdot \sum_{\pi \in S_N^\kappa} \Pr_{\Pi \sim \mathcal{D}_X} [\Pi = \pi] \\ &= \sum_{\kappa=0}^{|X_1|} \kappa \cdot \frac{\sum_{\pi \in S_N^\kappa} |X_\pi|}{\sum_{\sigma \in S_N} |X_\sigma|} \\ &= \sum_{\kappa=0}^{|X_1|} \kappa \cdot \frac{\kappa \cdot |S_N^\kappa|}{\sum_{\sigma \in S_N} |X_\sigma|} \\ &= \left(\mathbb{E}_{\sigma \sim S_N} [|X_\sigma|] \right)^{-1} \sum_{\kappa=0}^{|X_1|} \kappa^2 \cdot \frac{|S_N^\kappa|}{N!} \quad (\text{By Equation (1)}) \\ &= \frac{\mathbb{E}_{\sigma \sim S_N} [|X_\sigma|^2]}{\mathbb{E}_{\sigma \sim S_N} [|X_\sigma|]}. \end{aligned} \quad (2)$$

In other words, the average number of zero pairs in the non-uniform case is the ratio between the second moment and the first moment in the uniform case. We also have

$$\frac{|X_1||X_2|}{N} \leq \mathbb{E}_{\sigma \sim S_N} [|X_\sigma|^2] \leq \frac{|X_1||X_2|}{N} + \frac{|X_1|^2|X_2|^2}{N^2} \quad (\text{Lemma 3.19})$$

$$\mathbb{E}_{\sigma \sim S_N} [|X_\sigma|] = \frac{|X_1||X_2|}{N}. \quad (\text{Theorem 3.13})$$

Putting everything together, we find that

$$1 \leq \mathbb{E}_{\pi \sim \mathcal{D}} [|X_\pi|] \leq 1 + \frac{|X_1||X_2|}{N}.$$

□

Recall that the previous lemma relied on the following fact:

Lemma 3.19. *Let $N \in \mathbb{N}$ and $X_1, X_2 \subset [N]$ be non-empty subsets. Let $|X_\pi|$ denote the random variable for the number of X -pairs for $\pi \in S_N$. Then, the second moment satisfies*

$$\left(\frac{|X_1||X_2|}{N} \right) \leq \mathbb{E}_{\pi \sim S_N} [|X_\pi|^2] \leq \left(\frac{|X_1||X_2|}{N} + \frac{|X_1|^2|X_2|^2}{N^2} \right).$$

Proof. For a uniform random $\pi \sim S_N$ and some $x \in [N]$, we let I_x be the indicator variable which equals 1, if $\pi(x) \in X_2$, and 0 otherwise. Then, we find that

$$\mathbb{E}[I_x] = \frac{|X_2|}{N}, \quad |X_\pi|^2 = \left(\sum_{x \in X_1} I_x \right)^2$$

Therefore, we can express the second moment as follows:

$$\begin{aligned} \mathbb{E}_{\pi \sim S_N} [|X_\pi|^2] &= \sum_{x, y \in X} \mathbb{E}[I_x I_y] && \text{(Linearity of Expectation)} \\ &= \sum_{x \in X_1} \mathbb{E}[I_x] + \sum_{x \neq y \in X_1} \mathbb{E}[I_x I_y]. \end{aligned}$$

Note that the the first term is equal to

$$\mathbb{E}_{\sigma \sim S_N} [|X_\sigma|] = \frac{|X_1| |X_2|}{N}.$$

The second term is obviously positive, hence the lower bound. The second term has less than $|X_1|^2$ summands, and each summand can be bounded above via

$$\frac{|X_2|(|X_2| - 1)}{N(N - 1)} < \frac{|X_2|^2}{N^2}.$$

Multiplying this through, we can upper bound the term as $\frac{|X_1|^2 |X_2|^2}{N^2}$. This gives the desired inequality. \square

We can now give a tail bound for the number of X -pairs in the non-uniform case. Our bound is tailored to the case where $|X_1| \cdot |X_2| = N$.

Theorem 3.20. *Let $N \in \mathbb{N}$ and let $X = (X_1, X_2)$ be a pair of subsets $X_1, X_2 \subseteq [N]$ such that $|X_1| \cdot |X_2| = N$. Then, for any real number $u \geq 6$ $\mathbb{E}_{\sigma \sim \mathcal{D}_X} [|X_\sigma|]$, it holds that*

$$\Pr_{\pi \sim \mathcal{D}_X} \left[|X_\pi| \geq \mathbb{E}_{\sigma \sim \mathcal{D}_X} [|X_\sigma|] + u \right] \leq 3 \exp(-4u/9).$$

Proof. Note that because $|X_1| \cdot |X_2| = N$ we can write \mathcal{D}_X as

$$\begin{aligned} \Pr_{\Phi \sim \mathcal{D}_X} [\Phi = \varphi] &= \frac{|X_\varphi|}{\sum_{\sigma \in S_N} |X_\sigma|} \\ &= \frac{|X_\varphi|}{N!} && \text{(Corollary 3.14)} \end{aligned}$$

Using the above equation, we can now write the probability distribution on X -pairs as follows:

$$\Pr_{\pi \sim \mathcal{D}} [|X_\pi| = \kappa] = \kappa \cdot \Pr_{\sigma \sim S_N} [|X_\sigma| = \kappa],$$

for any $\kappa \in \{0, 1, \dots, |X_1|\}$. In particular, for $\kappa > 6$, we have the bound

$$\Pr_{\pi \sim S_N} [|X_\pi| \geq \kappa] \leq \exp(-3\kappa/4). \quad (\text{Theorem 3.15})$$

Therefore, we obtain

$$\begin{aligned} \Pr_{\pi \sim \mathcal{D}_X} [|X_\pi| = \kappa] &\leq \kappa \cdot \exp(-3\kappa/4) \\ &\leq \exp\left(-\left(\frac{9-2\ln(6)}{12}\right)\kappa\right). \quad (\text{From } \kappa \geq 6). \end{aligned}$$

Putting everything together and bounding the geometric series, we get that

$$\begin{aligned} \Pr_{\pi \sim \mathcal{D}_X} [|X_\pi| \geq \kappa] &= \sum_{i=\kappa}^{|X_1|} \Pr_{\pi \sim \mathcal{D}} [|X_\pi| = i] \\ &\leq \sum_{i=\kappa}^{\infty} \exp\left(-\left(\frac{9-2\ln(6)}{12}\right)i\right) \\ &\leq \left(1 - \exp\left(-\frac{9-2\ln(6)}{12}\right)\right)^{-1} \exp\left(-\left(\frac{9-2\ln(6)}{12}\right)\kappa\right) \\ &\leq 3 \exp\left(\frac{-4\kappa}{9}\right). \end{aligned}$$

□

4 Query Problem Lower Bounds

Unruh [Unr21; Unr23] proposed the task of “double-sided zero-search”—a simple query problem dealing with invertible permutations which seems to go beyond the scope of current techniques, and may potentially offer new insights into the post-quantum security of the sponge construction. Unruh conjectured that it requires at least $\Omega(2^{n/2})$ many queries to solve the problem with constant success probability—and this is tight due to Grover’s algorithm.

In this section, we prove query-lower bounds for Unruh’s original double-sided zero-search problem, as well as a non-uniform variant that will be useful in proving the one-wayness of the single-round sponge with invertible permutations. We remark that our lower bounds on double-sided zero-search are tight up to a constant factor, and therefore resolve Unruh’s original conjecture.

4.1 Double-Sided Zero-Search

In light of the difficulty in proving the post-quantum security of the sponge with invertible permutations, Unruh [Unr21; Unr23] proposed the following simple query problem which seems beyond the scope of current techniques.

Problem 1 (DOUBLE-SIDED ZERO-SEARCH). *Given quantum query-access to a uniformly random permutation $\varphi : \{0,1\}^{2n} \rightarrow \{0,1\}^{2n}$ as well as its inverse φ^{-1} , output a pair of strings $x, y \in \{0,1\}^n$ such that $\varphi(x||0^n) = y||0^n$.*

Recall that we call $x, y \in \{0,1\}^n$ which satisfy $\varphi(x||0^n) = y||0^n$ a *zero pair* of φ .

We prove that any algorithm must make at least $T = \Omega(\sqrt{\epsilon 2^n})$ many queries in order to find a zero pair for a random $2n$ -bit permutation with probability ϵ . At a high level, the first step of the proof is showing that there is some worst case instances on which the problem is hard for any fixed number of zero pairs, in Lemma 4.3. Then we give a worst to average case reduction to show hardness for a random permutation with a fixed number of zero pairs in Lemma 4.4. Finally, we give a bound for the general problem using the aforementioned bounds as well as tail bounds on the number of zero pairs.

Theorem 4.1. *Any quantum algorithm for DOUBLE-SIDED ZERO-SEARCH that makes T queries to a random $2n$ -bit permutation or its inverse and succeeds with probability $\epsilon > 0$ satisfies*

$$\epsilon \leq \frac{50(T+1)^2}{2^n}.$$

Proof. Let φ be a uniformly random $2n$ -bit permutation and $K = |Z_\varphi| \in \{0, 1, \dots, 2^n\}$ be the random variable corresponding to the number of zero pairs in φ . By the law of total probability, we can write the success probability ϵ of a T -query algorithm as

$$\begin{aligned} \epsilon &= \sum_{\kappa=0}^{2^n} \Pr[K = \kappa] \cdot \Pr[\text{success} | K = \kappa] \\ &= \underbrace{\sum_{\kappa=1}^6 \Pr[K = \kappa] \cdot \Pr[\text{success} | K = \kappa]}_{=P_1} + \underbrace{\sum_{\kappa=7}^{2^n} \Pr[K = \kappa] \cdot \Pr[\text{success} | K = \kappa]}_{=P_2}. \end{aligned}$$

We bound each term separately—note that splitting the sum at $\kappa = 6$ is chosen to give the tightest bound. Beginning with the first, we get

$$\begin{aligned} P_1 &\leq \max_{\kappa \in [6]} \{\Pr[\text{success} | K = \kappa]\} && \text{(By Convexity)} \\ &\leq \frac{48(T+1)^2}{2^n}. && \text{(By Lemma 4.4)} \end{aligned}$$

Continuing with the second, we find that

$$\begin{aligned} P_2 &\leq \sum_{\kappa=7}^{2^n} \Pr[\text{success} | K = \kappa] \cdot \exp(-3\kappa/4) && \text{(By Theorem 3.15)} \\ &\leq \sum_{\kappa=7}^{2^n} \frac{8(T+1)^2 \kappa}{2^n} \cdot \exp(-3\kappa/4) && \text{(By Lemma 4.4)} \\ &\leq \frac{2(T+1)^2}{2^n}. && \text{(Arithmetico-geometric series)} \end{aligned}$$

Therefore, we obtain the following bound on total success probability:

$$\epsilon \leq \frac{50(T+1)^2}{2^n}$$

□

We observe that this bound is tight up to constant factors for any ϵ , as there is at least single zero pair with $\Omega(1)$ probability (Lemma 3.5), which can be found using Grover's algorithm [Gro96] with probability $\Theta(T^2/2^n)$ after T queries. In particular, we have the following corollary.

Corollary 4.2. *Given quantum oracle-access to a random $2n$ -bit permutation and its inverse, the query complexity of outputting a zero pair with probability ϵ is $\Theta(\sqrt{\epsilon 2^n})$.*

Reduction from worst-case unstructured search. In the following lemma, we prove a query-lower bound for a worst-case variant of DOUBLE-SIDED ZERO-SEARCH in Problem 1, where we are given the promise that the underlying permutation has exactly K zero pairs. We show the following.

Lemma 4.3. *Any T -query quantum algorithm for the DOUBLE-SIDED ZERO-SEARCH problem succeeding with probability $\epsilon > 0$ on a $2n$ -bit worst-case permutation with exactly $K > 0$ zero pairs must satisfy the inequality*

$$\epsilon \leq \frac{8(T+1)^2 K}{2^n}.$$

Proof. We reduce the worst-case UNSTRUCTURED SEARCH problem with K out of 2^n marked elements to a specific instance of DOUBLE-SIDED ZERO-SEARCH for a $2n$ -bit permutation φ with exactly K zero pairs without any query overhead.

Suppose we are given quantum query access to a function

$$f : \{0,1\}^n \rightarrow \{0,1\}$$

such that $|f^{-1}(1)| = K$. We now construct a permutation $\varphi : \{0,1\}^{2n} \rightarrow \{0,1\}^{2n}$, where φ is defined as follows for any pair of strings $x, y \in \{0,1\}^n$:

$$\varphi(x||y) = \begin{cases} x||y & \text{if } f(x) = 1 \\ x||(y \oplus 1^n) & \text{if } f(x) = 0. \end{cases}$$

It is apparent that we can implement a query to φ using a single query to f . Further, we have $\varphi = \varphi^{-1}$ since φ is a composition of 1 and 2-cycles; hence, we can also implement queries to φ^{-1} . Now let $x||0^n$ be a zero pair of φ . It follows that $f(x) = 1$, as otherwise $\varphi(x||0^n) = x||1^n$. Similarly, if, for any x' , we have $f(x') = 1$, then $x'||0^n$ is a zero pair of φ . There are therefore K zero pairs in φ and, given a zero pair of φ , it is straight forward to find a preimage of 1 under f , which can be used to solve UNSTRUCTURED SEARCH. The claim now follows from Corollary 2.4. □

Worst-case to average-case reduction. In the following lemma, we prove a query-lower bound for an average-case variant of DOUBLE-SIDED ZERO-SEARCH in Problem 1, where we are given the promise that the underlying permutation has exactly K zero pairs. In this case, average refers to a uniform random permutation over all that have exactly K zero pairs. We show the following.

Lemma 4.4. *Any quantum algorithm for DOUBLE-SIDED ZERO-SEARCH on a uniform random permutation φ , subject to the constraint that φ has exactly $K > 0$ zero pairs, making T queries to φ, φ^{-1} and succeeding with probability $\epsilon > 0$ satisfies the inequality*

$$\epsilon \leq \frac{8(T+1)^2 K}{2^n}$$

Proof. Let $N = 2^{2n}$ and let $Z \subset [N]$ denote the set of elements whose binary decomposition ends in 0^n . Let $G_Z \leq S_N$ be the subgroup of permutations that fix Z , i.e.,

$$G_Z \cong S_Z \times S_{[N] \setminus Z}.$$

We can re-randomize any worst-case permutation $\varphi : \{0, 1\}^{2n} \rightarrow \{0, 1\}^{2n}$ with K zero pairs into an average-case permutation φ^{sym} with K zero pairs as follows.

1. Randomly and independently sample permutations $\omega, \sigma \sim G_Z$.
2. Define the symmetrized permutation $\varphi^{\text{sym}} = \omega \circ \varphi \circ \sigma$.

It follows from Lemma 3.12 that φ^{sym} is uniform random among permutations with K zero pairs. Let $x, y \in Z$ be a zero pair of φ^{sym} such that $\varphi^{\text{sym}}(x) = y$. We have that

$$x' = \sigma(x), \quad y' = \omega^{-1}(y) \tag{3}$$

satisfy $x', y' \in Z$ and $\varphi(x') = y'$, hence a zero pair of φ can be constructed for free from a zero pair of φ^{sym} . We can simulate queries to φ^{sym} as well as its inverse with a single query to φ, φ^{-1} respectively, so the reduction incurs no overhead. Therefore, the claim follows from Lemma 4.3. \square

4.2 Alternative proof in the superposition-oracle framework

In this section, we give an alternative proof for Lemma 4.4, i.e., the query lower bound for DOUBLE-SIDED ZERO-SEARCH. Because this lemma is a key component of Theorem 4.1, it presents an alternative approach for resolving Conjecture 1.

We work with two-way accessible superposition oracles for invertible permutations. The superposition oracle framework is a powerful tool which has been used to prove query lower bounds in a variety of settings, e.g., [Amb+11; Zha18; Ros22; Ala+22]. Zhandry [Zha18] also introduced the notion of compressed oracles as a means to “record” quantum queries to a random oracle. However, contrary to the work of Zhandry [Zha18], we do not need to “compress” the oracles and use inefficient representations instead. First, we introduce some relevant notation.

Superposition oracles. Let $\mathcal{F}_n = \{f : \{0,1\}^n \rightarrow \{0,1\}^n\}$ be the family of n -bit functions. The *function register* for $f \in \mathcal{F}_n$ is a collection $F = \{F_x\}_{x \in \{0,1\}^n}$ with

$$|f\rangle_F = \bigotimes_{x \in \{0,1\}^n} |f(x)\rangle_{F_x}.$$

A query to f in the superposition oracle framework amounts to the operation O with

$$O_{XYF} |x\rangle_X |y\rangle_Y |f\rangle_F = \text{CNOT}_{F_x:Y} |x\rangle_X |y\rangle_Y |f\rangle_F = |x\rangle_X |y \oplus f(x)\rangle_Y |f\rangle_F.$$

We can also model inverse queries in the superposition oracle framework. This amounts to the operation O^{-1} which is defined as the unitary

$$O_{XYF}^{-1} |x\rangle_X |y\rangle_Y |f\rangle_F = |x \oplus (\bigoplus_{x':f(x')=y} x')\rangle_X |y\rangle_Y |f\rangle_F.$$

For example, if f is a permutation, the inverse oracle amounts to the operation

$$O_{XYF}^{-1} |x\rangle_X |y\rangle_Y |f\rangle_F = |x \oplus f^{-1}(y)\rangle_X |y\rangle_Y |f\rangle_F.$$

Symmetrization. Let $N = 2^{2n}$ and consider a random permutation $\varphi \in S_N^\kappa$ which has exactly κ zero pairs. In the superposition oracle framework, we can model queries to such a permutation using a function register of the form

$$|\Phi_\kappa\rangle_F = |S_N^\kappa|^{-\frac{1}{2}} \sum_{\varphi \in S_N^\kappa} |\varphi\rangle_F. \quad (4)$$

We now switch to an alternative characterization of the function registers which uses our technique of symmetrization via subset pairs of the symmetric group. Let $Z \subset [N]$ denote the set of strings whose binary decomposition ends in 0^n many zeroes. Let $G_Z \leq S_N$ be the subgroup of permutations in S_N that fix Z , i.e.,

$$G_Z \cong S_Z \times S_{[N] \setminus Z}.$$

Define the product group $G = G_Z \times G_Z$, where G consists of pairs of permutations (σ, ω) such that each permutation fixes Z , i.e., $\sigma(Z) = Z = \omega(Z)$. Define the unitary representation $U : G \rightarrow \text{GL}(\mathbb{C}^{\mathcal{F}})$ with $(\sigma, \omega) \mapsto U_{\sigma, \omega}$ and

$$U_{\sigma, \omega} |f\rangle_F = |\omega \circ f \circ \sigma\rangle_F, \quad \text{for } f \in \mathcal{F}_{2n}.$$

Using symmetrization, we can now equivalently instantiate the function register F in Equation (4) via the extension

$$|\Phi_\kappa\rangle_F \mapsto |\pi_\kappa^{\text{sym}}\rangle_{\Sigma\Omega F} = \frac{1}{\sqrt{|G|}} \sum_{(\sigma, \omega) \in G} |\sigma\rangle_\Sigma |\omega\rangle_\Omega U_{\sigma, \omega} |\pi_\kappa\rangle_F,$$

where π_κ is an *arbitrary* fixed permutation in S_N^κ with exactly κ zero pairs. To see why this is a valid purification, we can use our symmetrization argument in Lemma 3.12 to argue that—once we trace out registers $\Sigma\Omega$ —the reduced state equals

$$\frac{1}{|G|} \sum_{(\sigma, \omega) \in G} U_{\sigma, \omega} |\pi_\kappa\rangle \langle \pi_\kappa|_F U_{\sigma, \omega}^\dagger = \frac{1}{|S_N^\kappa|} \sum_{\varphi \in S_N^\kappa} |\varphi\rangle \langle \varphi|_F.$$

Double-sided zero-search revisited. To give an alternative proof for Lemma 4.4 in the superposition oracle framework, we first prove a lower bound for the decision problem. We show that no quantum algorithm can distinguish whether it is querying a random invertible permutation with exactly κ zero pairs, or a random invertible permutation with no zero pairs—unless it makes a large number of queries. Our proof is rooted in the hybrid argument [Ben+97] and one-way to hiding [AHU18], and allows the query algorithm to query two oracles: one in the forward direction and one in the backward direction. We spell out the full details on the hybrid argument in terms of our notation for the convenience of the reader.

Theorem 4.5. *Let $N = 2^{2n}$ and let $\kappa \in \{0, 1, \dots, \sqrt{N}\}$. Then, for any quantum algorithm \mathcal{D} which makes at most T many queries, it holds that*

$$\left| \Pr_{\varphi \sim S_N^\kappa} [\mathcal{D}^{\varphi, \varphi^{-1}}(1^n) = 1] - \Pr_{\varphi \sim S_N^0} [\mathcal{D}^{\varphi, \varphi^{-1}}(1^n) = 1] \right| \leq 2T \sqrt{\frac{\kappa}{2^n}}.$$

Proof. Suppose that a distinguisher \mathcal{D} makes a total amount of T oracle queries. In the superposition oracle framework, we can model the adversary/oracle interaction by introducing a function register F which is outside of the view of \mathcal{D} . Suppose the register F is initially in the state $|\Phi^0\rangle_F$. Using symmetrization, we can now equivalently instantiate the register F as the uniform superposition

$$|\Phi^0\rangle_F \mapsto |\pi_0^{\text{sym}}\rangle_{\Sigma\Omega F} = \frac{1}{\sqrt{|G|}} \sum_{(\sigma, \omega) \in G} |\sigma\rangle_\Sigma |\omega\rangle_\Omega U_{\sigma, \omega} |\pi_0\rangle_F,$$

where $\pi_0(x_1||x_2) = (x_1||x_2 \oplus 1^n)$ is a fixed permutation in S_N^0 with no zero pairs.

For the remainder of the proof, we will analyze the effect of swapping out the symmetrization of $\pi_0 \in S_N^0$ with symmetrization of a different permutation which comes from S_N^κ . Specifically, we will choose a particular permutation π_K which is generated as follows: first sample a random subset $K \subset \{0, 1\}^n$ (independently of everything else) of size $|K| = \kappa$, and then let π_K be the permutation

$$\pi_K(x_1||x_2) = \begin{cases} x_1||x_2 & \text{if } x_1 \in K; \\ x_1||x_2 \oplus 1^n & \text{otherwise.} \end{cases}$$

Note that π_K differs from π_0 on exactly $\kappa \cdot 2^n$ many inputs of the form $(x_1||x_2)$, where $x_1 \in K$ and x_2 is arbitrary. Moreover, π_K has exactly κ zero pairs of the form $(x_1||0^n)$ whenever $x_1 \in K$. Thus, $\pi_K \in S_N^\kappa$. Notice that we can map between the symmetrizations $|\pi_0^{\text{sym}}\rangle$ and $|\pi_K^{\text{sym}}\rangle$ via the following unitary

$$\text{SWAP}_{\Sigma\Omega F}^{\pi_K, \pi_0} = U_{\Sigma\Omega: F}(\mathbb{1}_{\Sigma\Omega} \otimes W_F^K) U_{\Sigma\Omega: F}^\dagger,$$

where $U_{\Sigma\Omega: F}$ is a controlled unitary of the form

$$U |\sigma\rangle_\Sigma |\omega\rangle_\Omega |f\rangle_F = |\sigma\rangle_\Sigma |\omega\rangle_\Omega U_{\sigma, \omega} |f\rangle_F, \quad \text{for } f \in \mathcal{F}_{2n},$$

and where W^K is the unitary which flips between $|\pi_K\rangle$ and $|\pi_0\rangle$.³ Therefore,

$$\text{SWAP}_{\Sigma\Omega F}^{\pi_K, \pi_0} |\pi_0^{\text{sym}}\rangle_{\Sigma\Omega F} = |\pi_K^{\text{sym}}\rangle_{\Sigma\Omega F}.$$

Our goal is to show the following property: on average over the choice of K , the query magnitude on symmetrized oracle calls that involve π_0 on $(x_1||x_2)$ with $x_1 \in K$ and arbitrary x_2 will be small. We proceed via a standard hybrid argument which involves a sequence of intermediate states. Note that oracle calls to $|\omega \circ \pi_0 \circ \sigma\rangle_F$, where the permutations σ, ω come from registers Σ and Ω , only involve K in one of two ways:

- a forward query in the X register is made on $\sigma^{-1}(x_1||x_2)$ with $x_1 \in K$; or
- a backward query in the Y register is made on $\omega(x_1||x_2)$ with $x_1 \in K$.

Define the following projectors that exclude each of the two cases:

$$\Pi_{X\Sigma}^K = \sum_{x' \in \{0,1\}^{2n}} |x'\rangle\langle x'|_X \otimes \bigotimes_{\substack{x_1 \in K \\ x_2 \in \{0,1\}^n}} (\mathbb{1} - |x_1||x_2\rangle\langle x_1||x_2|)_{\Sigma_{x'}} \quad (5)$$

$$\Xi_{Y\Omega}^K = \sum_{y' \in \{0,1\}^{2n}} |y'\rangle\langle y'|_Y \otimes \bigotimes_{\substack{x_1 \in K \\ x_2 \in \{0,1\}^n}} (\mathbb{1} - |y'\rangle\langle y'|)_{\Omega_{x_1||x_2}}. \quad (6)$$

If we apply the projectors $\Pi_{X\Sigma}^K$ and $\Xi_{Y\Omega}^K$ to each forward/backward query just before the oracle is applied, then swapping from π_0 to π_K before the oracle evaluation has no effect. Formally, we have the commutation relations⁴

$$\left[\text{SWAP}_{\Sigma\Omega F}^{\pi_K, \pi_0}, O_{XYF} \Pi_{X\Sigma}^K \right] = 0 \quad (7)$$

$$\left[\text{SWAP}_{\Sigma\Omega F}^{\pi_K, \pi_0}, O_{XYF}^{-1} \Xi_{Y\Omega}^K \right] = 0. \quad (8)$$

For $i \in [T]$, we now define the ensemble of projectors $\{P_i^K\}_{i \in [T]}$, where

$$P_i^K = \begin{cases} \Pi_{X\Sigma}^K, & \text{if } O_i = O; \\ \Xi_{Y\Omega}^K, & \text{if } O_i = O^{-1}. \end{cases} \quad (9)$$

Suppose F is initially in the state $|\Phi^0\rangle_F$ and \mathcal{D} starts out with some normalized pure state $|\eta_0\rangle$ on registers XYE . We can model the interaction between \mathcal{D} and the oracle as sequence of oracle queries and unitaries such that for $i \in [T]$,

$$|\psi_i\rangle_{XYE\Sigma\Omega F} = U_i O_i U_{i-1} O_{i-1} \cdots U_1 O_1 |\psi_0\rangle. \quad (10)$$

Here, the unitaries $O_i \in \{O, O^{-1}\}$ represent a call to the forward/backward oracles acting on registers XYF , the U_i represent intermediate unitaries performed by \mathcal{D} which act on XYE , for some workspace register E , and where $|\psi_0\rangle$ is the state

$$|\psi_0\rangle_{XYE\Sigma\Omega F} = |\eta_0\rangle_{XYE} \otimes |\pi_0^{\text{sym}}\rangle_{\Sigma\Omega F}.$$

³ W^K has a well-defined extension to all $|f\rangle_F$ since we just need to XOR the string $(0^n||1^n)$ into the function output whenever the input is of the form $(x_1||y)$, for $x_1 \in K$ and $y \in \{0,1\}^n$

⁴For two linear operators A and B , the commutator is defined as $[A, B] := AB - BA$.

To study the effects of swapping between symmetrizations of $\pi_0 \in S_N^0$ and $\pi_K \in S_N^K$, we now define i -th intermediate state as

$$|\psi_i^K\rangle_{XYE\Sigma\Omega F} = U_i O_i U_{i-1} O_{i-1} \cdots U_1 O_1 |\psi_0^K\rangle,$$

where $|\psi_0^K\rangle$ is the initial state corresponding to a symmetrization of π_K , i.e.,

$$|\psi_0^K\rangle_{XYE\Sigma\Omega F} = |\eta_0\rangle_{XYE} \otimes |\pi_K^{\text{sym}}\rangle_{\Sigma\Omega F}.$$

We now introduce the K -dependent oracle unitaries $\{O_i^K\}_{i \in [T]}$ where

$$O_i^K = (\text{SWAP}_{\Sigma\Omega F}^{\pi_K, \pi_0})^\dagger (O_i)_{XYF} \text{SWAP}_{\Sigma\Omega F}^{\pi_K, \pi_0}$$

acts on registers $XY\Sigma\Omega F$. This yields the following equivalent characterization

$$|\psi_i^K\rangle_{XYE\Sigma\Omega F} = U_i O_i^K U_{i-1} O_{i-1}^K \cdots U_1 O_1^K |\psi_0\rangle. \quad (11)$$

By unitarity, we can now bound the distance between the i -th hybrid states as

$$\begin{aligned} \left\| |\psi_i\rangle - |\psi_i^K\rangle \right\|^2 &= \left\| U_i O_i |\psi_{i-1}\rangle - U_i O_i^K |\psi_{i-1}^K\rangle \right\|^2 \\ &= \left\| O_i |\psi_{i-1}\rangle - O_i^K |\psi_{i-1}^K\rangle \right\|^2 \\ &= \left\| O_i |\psi_{i-1}\rangle - O_i^K |\psi_{i-1}\rangle + O_i^K |\psi_{i-1}\rangle - O_i^K |\psi_{i-1}^K\rangle \right\|^2 \\ &\leq \left\| O_i |\psi_{i-1}\rangle - O_i^K |\psi_{i-1}\rangle \right\|^2 + \left\| O_i^K |\psi_{i-1}\rangle - O_i^K |\psi_{i-1}^K\rangle \right\|^2 \\ &\quad + 2 \cdot \left\| O_i |\psi_{i-1}\rangle - O_i^K |\psi_{i-1}\rangle \right\| \cdot \left\| O_i^K (|\psi_{i-1}\rangle - |\psi_{i-1}^K\rangle) \right\| \\ &= \left\| (O_i - O_i^K) |\psi_{i-1}\rangle \right\|^2 + \left\| |\psi_{i-1}\rangle - |\psi_{i-1}^K\rangle \right\|^2 \\ &\quad + 2 \cdot \left\| (O_i - O_i^K) |\psi_{i-1}\rangle \right\| \cdot \left\| |\psi_{i-1}\rangle - |\psi_{i-1}^K\rangle \right\|, \end{aligned} \quad (12)$$

where the triangle inequality gives the inequality. We now observe that

$$\begin{aligned} \left\| (O_i - O_i^K) |\psi_{i-1}\rangle \right\| &= \left\| (O_i - O_i^K) \bar{P}_i^K |\psi_{i-1}\rangle + (O_i - O_i^K) P_i^K |\psi_{i-1}\rangle \right\| \\ &\leq \left\| (O_i - O_i^K) \bar{P}_i^K |\psi_{i-1}\rangle \right\| + \left\| (O_i - O_i^K) P_i^K |\psi_{i-1}\rangle \right\| \\ &\leq 2 \cdot \left\| \bar{P}_i^K |\psi_{i-1}\rangle \right\| + \left\| (O_i - O_i^K) P_i^K |\psi_{i-1}\rangle \right\| \\ &= 2 \cdot \left\| \bar{P}_i^K |\psi_{i-1}\rangle \right\|. \end{aligned} \quad (13)$$

Here, the second to last line uses that $(O_i - O_i^K)$ has operator norm at most 2. For the last line, we first invoke the fact that $\text{SWAP}_{\Sigma\Omega F}^{\pi_K, \pi_0}$ and P_i^K commute, and then use the commutation relations from Equations (7) and (8) so that

$$\begin{aligned} \left\| (O_i - O_i^K) P_i^K |\psi_{i-1}\rangle \right\| &= \left\| \text{SWAP}_{\Sigma\Omega F}^{\pi_K, \pi_0} (O_i - O_i^K) P_i^K |\psi_{i-1}\rangle \right\| \quad (\text{By unitarity}) \\ &= \left\| (\text{SWAP}_{\Sigma\Omega F}^{\pi_K, \pi_0} O_i P_i^K - O_i \text{SWAP}_{\Sigma\Omega F}^{\pi_K, \pi_0} P_i^K) |\psi_{i-1}\rangle \right\| \\ &= \left\| (\text{SWAP}_{\Sigma\Omega F}^{\pi_K, \pi_0} O_i P_i^K - O_i P_i^K \text{SWAP}_{\Sigma\Omega F}^{\pi_K, \pi_0}) |\psi_{i-1}\rangle \right\| \\ &= \left\| [\text{SWAP}_{\Sigma\Omega F}^{\pi_K, \pi_0}, O_i P_i^K] |\psi_{i-1}\rangle \right\| = 0. \end{aligned}$$

By combining Equation (12) and Equation (13), we get that

$$\begin{aligned} \left\| |\psi_i\rangle - |\psi_i^K\rangle \right\|^2 &\leq 4 \cdot \left\| \bar{P}_i^K |\psi_{i-1}\rangle \right\|^2 + \left\| |\psi_{i-1}\rangle - |\psi_{i-1}^K\rangle \right\|^2 \\ &\quad + 4 \cdot \left\| \bar{P}_i^K |\psi_{i-1}\rangle \right\| \cdot \left\| |\psi_{i-1}\rangle - |\psi_{i-1}^K\rangle \right\| \\ &= \left(\left\| |\psi_{i-1}\rangle - |\psi_{i-1}^K\rangle \right\| + 2 \cdot \left\| \bar{P}_i^K |\psi_{i-1}\rangle \right\| \right)^2, \end{aligned} \quad (14)$$

and thus

$$\left\| |\psi_T\rangle - |\psi_T^K\rangle \right\| \leq 2 \cdot \sum_{i=1}^T \left\| \bar{P}_i^K |\psi_i\rangle \right\|. \quad (15)$$

Therefore, we can bound the average distance as follows:

$$\begin{aligned} &\text{TD}\left(|\psi_T\rangle\langle\psi_T|, \mathbb{E}_{\substack{K \subset \{0,1\}^n \\ |K|=\kappa}} |\psi_T^K\rangle\langle\psi_T^K|\right) \\ &\leq \mathbb{E}_{\substack{K \subset \{0,1\}^n \\ |K|=\kappa}} \text{TD}\left(|\psi_T\rangle\langle\psi_T|, |\psi_T^K\rangle\langle\psi_T^K|\right) \quad (\text{By convexity}) \\ &\leq \mathbb{E}_{\substack{K \subset \{0,1\}^n \\ |K|=\kappa}} \left\| |\psi_T\rangle - |\psi_T^K\rangle \right\| \quad (\text{Norm inequality}) \\ &\leq \mathbb{E}_{\substack{K \subset \{0,1\}^n \\ |K|=\kappa}} 2 \cdot \sum_{i=1}^T \left\| \bar{P}_i^K |\psi_i\rangle \right\| \quad (\text{By Equation (15)}) \\ &= 2T \mathbb{E}_{i \sim [T]} \mathbb{E}_{\substack{K \subset \{0,1\}^n \\ |K|=\kappa}} \left\| \bar{P}_i^K |\psi_i\rangle \right\| \\ &\leq 2T \sqrt{\mathbb{E}_{i \sim [T]} \mathbb{E}_{\substack{K \subset \{0,1\}^n \\ |K|=\kappa}} \left\| \bar{P}_i^K |\psi_i\rangle \right\|^2}. \quad (\text{Jensen's inequality}). \end{aligned}$$

To complete the proof, it suffices to bound the expectation. By the definition of $\{P_i^K\}_{i \in [T]}$ in Equation (9), we have to consider two cases for each $i \in [T]$. First, suppose that the i -th query is a forward query with $O_i = O$. Then,

$$\mathbb{E}_{\substack{K \subset \{0,1\}^n \\ |K|=\kappa}} \left\| \bar{P}_i^K |\psi_i\rangle \right\|^2 = \mathbb{E}_{\substack{K \subset \{0,1\}^n \\ |K|=\kappa}} \left[\left\| (\mathbb{1} - \Pi^K)_{X\Sigma} |\psi_i\rangle_{XY\Sigma\Omega E} \right\|_2^2 \right] \leq \frac{\kappa}{2^n}. \quad (16)$$

Next, suppose that the i -th query is to the inverse oracle $O_i = O^{-1}$. Then, we have

$$\mathbb{E}_{\substack{K \subset \{0,1\}^n \\ |K|=\kappa}} \left\| \bar{P}_i^K |\psi_i\rangle \right\|^2 = \mathbb{E}_{\substack{K \subset \{0,1\}^n \\ |K|=\kappa}} \left[\left\| (\mathbb{1} - \Xi^K)_{Y\Omega} |\psi_i\rangle_{XY\Sigma\Omega E} \right\|_2^2 \right] \leq \frac{\kappa}{2^n}. \quad (17)$$

Putting everything together and using Equation (16) and Equation (17), we arrive that

$$\text{TD}\left(|\psi_T\rangle\langle\psi_T|, \mathbb{E}_{\substack{K \subseteq \{0,1\}^n \\ |K|=\kappa}} |\psi_T^K\rangle\langle\psi_T^K|\right) \leq 2T \cdot \sqrt{\frac{\kappa}{2^n}}.$$

This proves the claim. \square

Alternative to resolving Unruh’s conjecture. We can use Theorem 4.5 to prove an alternative version of Lemma 4.4, which in turn is the key component of Theorem 4.1, and thus in resolving Conjecture 1.

Lemma 4.6 (Alternative to Lemma 4.4). *Any quantum T -query algorithm for the DOUBLE-SIDED ZERO-SEARCH problem with respect to a random $2n$ -bit invertible permutation with exactly $\kappa > 0$ zero pairs, $\pi \sim S_{2n}^\kappa$, that succeeds with probability $\epsilon > 0$ satisfies the inequality*

$$\epsilon \leq 2(T + 1) \sqrt{\frac{\kappa}{2^n}}.$$

Before proceeding with the proof, we observe a difference between Lemma 4.6 and Lemma 4.4. While both are tight up to constant factors for a constant success probability (and thus sufficient to resolve Conjecture 1), the bound in Lemma 4.6 is of the form $\epsilon = O(T\sqrt{\kappa}/2^n)$ whereas in Lemma 4.4 the bound derived is of the form $\epsilon = O(T^2\kappa/2^n)$. The latter bound is quadratically tighter for small $\epsilon \ll 1$, and so we state our main theorem in these terms. We leave tightening the analysis in the superposition-oracle framework as an interesting open problem.

Proof. Suppose that $\mathcal{A}^{\varphi, \varphi^{-1}}$ finds a zero pair with probability $\epsilon > 0$ whenever φ is random $2n$ -bit permutation with exactly κ zero pairs. We can transform such an \mathcal{A} into a distinguisher between permutations with κ zero pairs and no zero pairs by running it on a given permutation, and guessing that φ has κ zero pairs if and only if \mathcal{A} finds a zero pair. Formally, given a permutation φ, φ^{-1} the distinguisher $\mathcal{D}^{\varphi, \varphi^{-1}}$ proceeds as follows:

1. Run $\mathcal{A}^{\varphi, \varphi^{-1}}(1^n)$ to obtain⁵ a pair $(x||0^n, y||0^n)$.
2. If $\varphi(x||0^n) = y||0^n$, then output 1 (guess that φ has κ zero pairs)
3. Otherwise, output 0 (guess that φ has no zero pairs)

If \mathcal{A} makes T queries then \mathcal{D} makes $T + 1$ queries, and the distinguishing advantage of \mathcal{D} is the probability that \mathcal{A} succeeds given that φ has κ zero pairs. We have

$$\left| \Pr_{\varphi \sim S_N^\kappa} [\mathcal{D}^{\varphi, \varphi^{-1}}(1^n) = 1] - \Pr_{\varphi \sim S_N^0} [\mathcal{D}^{\varphi, \varphi^{-1}}(1^n) = 1] \right| = \Pr_{\varphi \sim S_N^\kappa} [\mathcal{A}^{\varphi, \varphi^{-1}}(1^n)].$$

⁵If the output is not of this form, simply output 0 (guess φ has no zero pairs).

Now applying Theorem 4.5, we obtain

$$\Pr_{\varphi \sim S_N^{\kappa}} [\mathcal{A}^{\varphi, \varphi^{-1}}(1^n)] \leq 2(T+1) \sqrt{\frac{\kappa}{2^n}}.$$

□

4.3 Non-Uniform Double-Sided Search

In this section we define a new variant of Unruh’s original double-sided zero-search problem [Unr21; Unr23]. Namely, we consider the DOUBLE-SIDED SEARCH problem in which the distribution on permutations is *non-uniform*, and the constraints on preimages versus images are more flexible. This problem will be useful for proving the one-wayness of the single-round sponge, which is our motivation for studying it.

Notation. For the remainder of this section, we fix the following notation:

- $N = 2^n$ and $r, c > 0$ are integers such that $r + c = n$.
- $X = (X_1, X_2)$ is a pair of subsets, where $X_1 \subset \{0, 1\}^n$ is the set of bitstrings ending in 0^c , $X_2 \subset \{0, 1\}^n$ is the set of bitstrings beginning in 0^r .
- $|X_\varphi|$ denotes the number of X -pairs with respect to a permutation $\varphi \in S_N$.
- G_1 is the subgroup of S_N consisting of permutations which preserve membership in X_1 , whereas G_2 is the subgroup of S_N which preserves membership in X_2 .

Note that $|X_1| \cdot |X_2| = N$, and G_1, G_2 are both Young subgroups. In this section, we consider the distribution \mathcal{D}_X from Definition 3.17 which assigns more weight to permutations $\varphi \in S_N$ that have a large amount of X -pairs. We will occasionally omit the subscript which specifies $X = (X_1, X_2)$, for the pair of subsets X_1, X_2 , when it is clear from context.

Problem 2 (NON-UNIFORM DOUBLE-SIDED SEARCH). *Given quantum query access to a permutation $\varphi : \{0, 1\}^n \rightarrow \{0, 1\}^n$ as well as its inverse φ^{-1} , where φ is sampled according to the non-uniform distribution \mathcal{D}_X from Definition 3.17, output a pair of strings $x \in X_1, y \in X_2$ such that $\varphi(x) = y$.*

In the remainder of this section, we prove that any algorithm must make at least $T = \Omega(\sqrt{\epsilon 2^{\min(r,c)}})$ many queries in order to find an X -pair for an n -bit permutation sampled from \mathcal{D}_X with probability ϵ . At a high level, the first step of the proof is showing that there exist worst-case instances on which the problem is hard for any fixed number of X -pairs (see Lemma 4.9). Then, we give a worst-case to average-case reduction to show hardness for a random permutation with a fixed number of X -pairs in Lemma 4.10. In the below theorem we give a bound for the general problem using the aforementioned results, as well as tail bounds on the number of subset pairs derived in Section 3.3.

Theorem 4.7. *Any quantum algorithm for NON-UNIFORM DOUBLE-SIDED SEARCH that makes T queries to an invertible permutation and succeeds with probability $\epsilon > 0$ satisfies*

$$\epsilon \leq \frac{80(T+1)^2}{2^{\min(r,c)}}.$$

Proof. Let $\varphi \sim \mathcal{D}_X$ be a permutation sampled according to the distribution \mathcal{D}_X from Definition 3.17 and let $K = |X_\varphi| \in \{0, 1, \dots, 2^{\min(r,c)}\}$ be the random variable corresponding to the number of X -pairs in φ . By the law of total probability, we can write the success probability of any T -query algorithm as the sum

$$\begin{aligned} \epsilon &= \sum_{\kappa=0}^{2^{\min(r,c)}} \Pr[K = \kappa] \cdot \Pr[\text{success} | K = \kappa] \\ &= \underbrace{\sum_{\kappa=1}^6 \Pr[K = \kappa] \cdot \Pr[\text{success} | K = \kappa]}_{=P_1} + \underbrace{\sum_{\kappa=7}^{2^{\min(r,c)}} \Pr[K = \kappa] \cdot \Pr[\text{success} | K = \kappa]}_{=P_2}. \end{aligned}$$

We bound each term separately—note that splitting the sum at $\kappa = 6$ is chosen to give the tightest bound. Beginning with the first, we find

$$\begin{aligned} P_1 &\leq \max_{\kappa \in [6]} \{\Pr[\text{success} | K = \kappa]\} && \text{(By Convexity)} \\ &\leq \frac{48(T+1)^2}{2^{\min(r,c)}}. && \text{(By Lemma 4.10)} \end{aligned}$$

Next, continuing with the second term, we get

$$\begin{aligned} P_2 &\leq \sum_{\kappa=7}^{2^{\min(r,c)}} \Pr[\text{success} | K = \kappa] \cdot 3 \exp(-4\kappa/9) && \text{(By Theorem 3.20)} \\ &\leq \sum_{\kappa=7}^{2^{\min(r,c)}} \frac{24(T+1)^2 \kappa}{2^{\min(r,c)}} \cdot \exp(-4\kappa/9) && \text{(By Lemma 4.10)} \\ &\leq \frac{32(T+1)^2}{2^{\min(r,c)}}. && \text{(Arithmetico-geometric series)} \end{aligned}$$

This concludes the proof, as we can bound the success probability as

$$\epsilon \leq \frac{80(T+1)^2}{2^{\min(r,c)}}.$$

□

This bound is tight up to constant factors for any ϵ , as there is at least a single X -pair with $\Omega(1)$ probability (Lemma 3.5), which can be found using Grover's algorithm [Gro96] with probability $\Theta(T^2/2^{\min(r,c)})$ after T queries. In particular, we have the following corollary.

Corollary 4.8. *Given quantum oracle-access to an n -bit permutation chosen from \mathcal{D}_X , the query complexity of outputting an X -pair with probability ϵ is $\Theta(\sqrt{\epsilon 2^{\min(r,c)}})$.*

Reduction from worst-case unstructured search. In the following lemma, we prove a query-lower bound for a worst-case variant of NON-UNIFORM DOUBLE-SIDED SEARCH in Problem 2, where we are given the promise that the underlying permutation has exactly K many X -pairs. We show the following.

Lemma 4.9. *Any T -query quantum algorithm for NON-UNIFORM DOUBLE-SIDED SEARCH on a worst-case permutation with exactly $K > 0$ many X -pairs that succeeds with probability $\epsilon > 0$ must satisfy the inequality*

$$\epsilon \leq \frac{8(T+1)^2 K}{2^{\min(r,c)}}.$$

Proof. We can reduce UNSTRUCTURED SEARCH with K out of $2^{\min(r,c)}$ marked elements to a specific instance of NON-UNIFORM DOUBLE-SIDED SEARCH with respect to a permutation φ on $N = 2^n$ elements and K many X pairs. Moreover, as we will show, our reduction does not incur any query overhead.

Suppose we are given query access to a function

$$f : \{0,1\}^{\min(r,c)} \rightarrow \{0,1\}$$

such that $|f^{-1}(1)| = K$. We can construct a permutation $\varphi : \{0,1\}^n \rightarrow \{0,1\}^n$ for $x \in \{0,1\}^{\min(r,c)}, y \in \{0,1\}^{\max(r,c)}$ as follows, where s^R denotes the reverse of string s :

$$\varphi(x||y) = \begin{cases} (x||y)^R & \text{if } f(x) = 1 \\ (x||y)^R \oplus 1^r || 0^c & \text{if } f(x) = 0. \end{cases}$$

It is apparent that we can implement a query to φ using a single query to f . Moreover, the image $\varphi(x||y)$ always contains x^R in the last $\min(r,c)$ bits, which suffices to determine $f(x)$, and thus to determine $x||y$ by inverting whichever case was applied. Hence, we can implement φ^{-1} with a single query to f . Now let $x \in \{0,1\}^{\min(r,c)}, y \in \{0,1\}^{\max(r,c)}$ such that y ends in 0^c and $\varphi(x||y)$ begins with 0^r . It follows that $f(x) = 1$ and $y = 0^{\max(r,c)}$, as otherwise either y does not end in 0^c or $\varphi(x||y)$ does not begin with 0^r . Similarly, if for any x' we have $f(x') = 1$ then $x' || 0^{\max(r,c)}$ is an X -pair of φ . Therefore, there are K many X -pairs in φ , and given an X -pair of φ it is straight forward to find a preimage of 1 under f , which can be used to solve UNSTRUCTURED SEARCH. The claim now follows from Corollary 2.4. \square

Worst-case to average-case reduction. In the following lemma, we prove a query-lower bound for an average-case variant of NON-UNIFORM DOUBLE-SIDED SEARCH in Problem 2, where we are given the promise that the underlying permutation has exactly K many X -pairs. In this case, average refers to a uniform random permutation subject to this promise. We show the following.

Lemma 4.10. *Any T -query algorithm for NON-UNIFORM DOUBLE-SIDED SEARCH on a random permutation which has exactly $K > 0$ many X -pairs that succeeds with probability $\epsilon > 0$ must satisfy the inequality*

$$\epsilon \leq \frac{8(T+1)^2 K}{2^{\min(r,c)}}.$$

Proof. We can re-randomize any worst-case permutation φ with K many X -pairs toward an average-case permutation φ^{sym} with K many X -pairs as follows.

1. Randomly and independently sample $\omega \sim G_1, \sigma \sim G_2$, where ω preserves membership in X_1 and σ preserves membership in X_2 .
2. Define the symmetrized permutation $\varphi^{\text{sym}} = \omega \circ \varphi \circ \sigma$.

It follows from Lemma 3.12 that φ^{sym} is random among permutations with K many X -pairs. Let $(x, y) \in X_1 \times X_2$ be an X -pair of φ^{sym} such that $\varphi^{\text{sym}}(x) = y$. Then,

$$x' = \sigma(x), \quad y' = \omega^{-1}(y) \quad (18)$$

which satisfy $x' \in X_1, y' \in X_2$ and $\varphi(x') = y'$. Hence, an X -pair of φ can be constructed for free from an X -pair of φ^{sym} . We can simulate queries to φ^{sym} as well as its inverse with a single query to φ, φ^{-1} respectively, so the reduction incurs no overhead. Lemma 2.4 now implies the claim. \square

5 One-Wayness of the Sponge Construction

In this section, we prove our main result; namely, the quantum one-wayness of the single-round sponge when instantiated with invertible random permutations.

The sponge construction. Recall that the sponge construction [Ber+11a] uses two main parameters which we call the *rate* r and the *capacity* c of the sponge. The internal state of the sponge function gets updated through successive applications of a *block function* $\varphi : \{0, 1\}^{r+c} \rightarrow \{0, 1\}^{r+c}$. Let $\sigma^{(0)} = (0^r || \text{IV})$ denote the initial state of the sponge, where $\text{IV} \in \{0, 1\}^c$ is some initialization vector; for simplicity, we consider the choice $\text{IV} = 0^c$. In the following, we use the notation $x_r \in \{0, 1\}^r$ and $x_c \in \{0, 1\}^c$ to denote the first r bits as well as the last c bits of $x \in \{0, 1\}^{r+c}$, respectively.

The sponge hash function $\text{Sp}^\varphi : \{0, 1\}^* \rightarrow \{0, 1\}^r$ is evaluated as follows on a sequence of r -bit blocks $m = (m_1, \dots, m_\ell)$:

- (Absorption phase) for $i \in [\ell]$, update the state by computing

$$\sigma^{(i)} = \varphi((\sigma_r^{(i-1)} \oplus m_i) || \sigma_c^{(i-1)}).$$

- (Squeezing phase) parse the final state $\sigma^{(\ell)}$ as $(\sigma_r^{(\ell)}, \sigma_c^{(\ell)}) \in \{0, 1\}^r \times \{0, 1\}^c$, and output the digest $\text{Sp}^\varphi(m_1, \dots, m_\ell) = \sigma_r^{(\ell)}$ consisting of the first r bits of $\sigma^{(\ell)}$.

We remark that the sponge construction also allows for variable-length output by essentially running the absorption phase in reverse [Ber+11a]; we choose to ignore this option for simplicity as we only focus on single-round sponge hashing.

One-wayness of the single-round sponge. Recall that, in the special case when there is only a single round of absorption, the sponge function $\text{Sp}^\varphi : \{0, 1\}^r \rightarrow \{0, 1\}^r$ takes on a simple form; namely, on input $x \in \{0, 1\}^r$, the output is given by the string $y = \text{Sp}^\varphi(x)$, where y corresponds to the first r bits of $\varphi(x||0^c)$.

We now answer the following question on the one-wayness of the single-round sponge; namely, how many queries does a quantum algorithm $\mathcal{A}^{\varphi, \varphi^{-1}}$ need to invert Sp^φ on a randomly chosen input, when φ is an invertible random permutation?

Let us first define the one-wayness experiment in full generality. Note that the sponge function Sp^φ depends on φ , which in turn is modeled as a uniformly random permutation which we have oracle access to. Therefore, we need to consider a variant of one-wayness which is parameterized by an ensemble of oracle functions.

Definition 5.1 (One-wayness). Let $\lambda \in \mathbb{N}$ be the security parameter and let $\mathcal{O} = \{\mathcal{O}_\lambda\}$ be an ensemble of oracle functions. A function family

$$\mathcal{H}^\mathcal{O} = \left\{ \{h^\mathcal{O} : \{0, 1\}^{m(\lambda)} \rightarrow \{0, 1\}^{n(\lambda)}\}_{\mathcal{O} \in \mathcal{O}_\lambda} \right\}_{\lambda \in \mathbb{N}}$$

is called one-way if (i) there exists a polynomial-time algorithm $\mathcal{M}^\mathcal{O}(1^\lambda)$ that efficiently computes $h^\mathcal{O}$ with oracle access to \mathcal{O} and (ii), for any efficient quantum algorithm \mathcal{A} , it holds

$$\Pr_{\substack{x \sim \{0, 1\}^m \\ \mathcal{O} \sim \mathcal{O}_\lambda}} \left[h^\mathcal{O}(x) = h^\mathcal{O}(x') : x' \leftarrow \mathcal{A}^\mathcal{O}(1^\lambda, h^\mathcal{O}(x)) \right] \leq \text{negl}(\lambda).$$

We now prove a lower bound for single-round sponge hashing, where $\ell = 1$ and the block function $\varphi : \{0, 1\}^{r+c} \rightarrow \{0, 1\}^{r+c}$ is an invertible random permutation. Specifically, we show that any algorithm must make at least $T = \Omega(\sqrt{\epsilon} 2^{\min(r,c)})$ many queries in order to find a pre-image of a randomly chosen image with probability ϵ .

Theorem 5.2. Let $n = r + c$ and $N = 2^n$. Suppose that a quantum algorithm \mathcal{A} breaks the one-wayness of the single-round Sponge $\text{Sp}^\varphi : \{0, 1\}^r \rightarrow \{0, 1\}^r$ with probability

$$\Pr_{\substack{x \sim \{0, 1\}^r \\ \varphi \sim S_N}} \left[\text{Sp}^\varphi(x) = \text{Sp}^\varphi(x') : x' \leftarrow \mathcal{A}^{\varphi, \varphi^{-1}}(1^n, \text{Sp}^\varphi(x)) \right] = \epsilon.$$

Then, if \mathcal{A} makes a total amount of T quantum queries to φ, φ^{-1} , it satisfies the inequality

$$\epsilon \leq \frac{80(T+1)^2}{2^{\min(r,c)}}.$$

Proof. Suppose that a quantum algorithm \mathcal{A} breaks the one-wayness of the single-round sponge hash function $\text{Sp}^\varphi : \{0, 1\}^r \rightarrow \{0, 1\}^r$ with probability $\epsilon > 0$. We reduce from the problem of NON-UNIFORM DOUBLE-SIDED SEARCH in Problem 2. Let $X = (X_1, X_2)$ be a pair of subsets, where $X_1 \subset \{0, 1\}^n$ denotes the set of bitstrings ending in 0^c and $X_2 \subset \{0, 1\}^n$ denotes the set of bitstrings beginning in 0^r .

Consider the following reduction $\mathcal{B}^{\pi, \pi^{-1}}(1^n)$ which receives quantum query access to a random permutation $\pi : \{0, 1\}^n \rightarrow \{0, 1\}^n$ chosen from \mathcal{D}_X (Definition 3.17) as well as its inverse π^{-1} .

1. \mathcal{B} samples a random $y \sim \{0,1\}^r$.
2. \mathcal{B} runs $\mathcal{A}^{\varphi, \varphi^{-1}}(1^n, y)$ to get $x' \in \{0,1\}^r$; whenever \mathcal{A} makes a query to either φ or φ^{-1} , the reduction \mathcal{B} answers with either O_φ or $O_{\varphi^{-1}}$, where φ is the permutation

$$\varphi := \text{XOR}_{(y||0^c)} \circ \pi$$

with $\text{XOR}_{(y||0^c)}(a||b) = (a \oplus y)||b$ for $a \in \{0,1\}^r$ and $b \in \{0,1\}^c$.

3. \mathcal{B} computes $\varphi(x' || 0^c) = y_{x'} || z_{x'}$ using a single query to O_π , and outputs the pair $(x' || 0^c)$ and $(0^r || z_{x'})$.

By Lemma 5.3, we know \mathcal{A} sees an instance (φ, y) drawn from the proper one-way game distribution. Note that whenever $\mathcal{A}^{\varphi, \varphi^{-1}}(1^n, y)$ succeeds at breaking one-wayness with respect to Sp^φ , it holds that $y_{x'} = y$, and thus $\varphi(x' || 0^c) = y || z_{x'}$. But then the output $(x' || 0^c)$ and $(0^r || z_{x'})$ of the reduction \mathcal{B} forms an X -pair of π , since

$$\begin{aligned} \pi(x' || 0^c) &= \text{XOR}_{(y||0^c)}(\varphi(x' || 0^c)) \\ &= \text{XOR}_{(y||0^c)}(y || z_{x'}) \\ &= 0^r || z_{x'}. \end{aligned}$$

The lower bound now follows from Theorem 4.7. □

Next, we prove the following technical lemma which we require for our main result. Here, we denote by \mathcal{D}_X the non-uniform distribution from Definition 3.17.

Lemma 5.3. *Let $n = r + c$ for $r, c \in \mathbb{N}$ and $N = 2^{r+c}$. Let $X = (X_1, X_2)$, where $X_1 \subset \{0,1\}^n$ denotes the set of bitstrings ending in 0^c and $X_2 \subset \{0,1\}^n$ denotes the set of bitstrings beginning in 0^r . Consider the following two distributions:*

<u>Distribution \mathcal{D}_1 :</u>	<u>Distribution \mathcal{D}_2 :</u>
1. Sample $\varphi \sim S_N$;	1. Sample $y \sim \{0,1\}^r$;
2. Sample $x \sim \{0,1\}^r$;	2. Sample $\pi \sim \mathcal{D}_X$ with
3. Let $y = \text{Sp}^\varphi(x)$;	$\Pr_{\Pi \sim \mathcal{D}_X} [\Pi = \pi] = X_\pi / \sum_{\sigma \in S_N} X_\sigma $;
4. Output (φ, y) .	3. Output (φ, y) with $\varphi = \text{XOR}_{(y 0^c)} \circ \pi$.

Then, the distributions \mathcal{D}_1 and \mathcal{D}_2 are identical.

Proof. To show that the distributions \mathcal{D}_1 and \mathcal{D}_2 are identical, it suffices to show that each pair $(\varphi, y) \in S_N \times \{0,1\}^r$ is equally likely to occur under both \mathcal{D}_1 and \mathcal{D}_2 .

First, we consider the probability that (φ, y) occur under \mathcal{D}_1 . Thus,

$$\begin{aligned} \Pr_{(\Phi, Y) \sim \mathcal{D}_1} [(\Phi, Y) = (\varphi, y)] &= \Pr[\Phi = \varphi] \cdot \Pr[Y = y | \Phi = \varphi] \\ &= \frac{1}{N!} \cdot \frac{|\{x \in \{0,1\}^r : \text{Sp}^\varphi(x) = y\}|}{2^r}. \end{aligned}$$

Next, we consider the probability that (φ, y) occurs under \mathcal{D}_2 . We find that

$$\begin{aligned} \Pr_{(\Phi, Y) \sim \mathcal{D}_2} [(\Phi, Y) = (\varphi, y)] &= \Pr[Y = y] \cdot \Pr[\Phi = \varphi | Y = y] \\ &= \frac{1}{2^r} \cdot \frac{|\{x \in \{0, 1\}^r : \text{Sp}^\pi(x) = 0^r \text{ and } \varphi = \text{XOR}_{(y||0^c)} \circ \pi, \text{ for } \pi \in S_N\}|}{\sum_{\sigma \in S_N} |\{x \in \{0, 1\}^r : \text{Sp}^\sigma(x) = 0^r, \text{ for } \sigma \in S_N\}|}. \end{aligned}$$

Note that since $|X_1| \cdot |X_2| = N$, we can invoke Theorem 3.13 to show that there is precisely one X -pair on average over the uniform choice over S_N . But since there are $N!$ such permutations the total number of X pairs must be $N!$. Thus,

$$N! = \sum_{\sigma \in S_N} |\{x \in \{0, 1\}^r : \text{Sp}^\sigma(x) = 0^r, \text{ for } \sigma \in S_N\}|.$$

Moreover, because the permutation π is uniquely defined by the permutation φ and the image y , it is also easy to see that

$$\begin{aligned} &|\{x \in \{0, 1\}^r : \text{Sp}^\pi(x) = 0^r \text{ and } \varphi = \text{XOR}_{(y||0^c)} \circ \pi, \text{ for } \pi \in S_N\}| \\ &= |\{x \in \{0, 1\}^r : \text{Sp}^\varphi(x) = y\}|. \end{aligned}$$

Putting everything together, it follows that

$$\Pr_{(\Phi, Y) \sim \mathcal{D}_2} [(\Phi, Y) = (\varphi, y)] = \frac{1}{N!} \cdot \frac{|\{x \in \{0, 1\}^r : \text{Sp}^\varphi(x) = y\}|}{2^r}.$$

This completes the proof. □

References

- [Hoe63] Wassily Hoeffding. “Probability Inequalities for Sums of Bounded Random Variables”. In: *Journal of the American Statistical Association* 58.301 (1963), pp. 13–30. ISSN: 01621459. URL: <http://www.jstor.org/stable/2282952> (visited on 02/25/2024).
- [Chv79] V. Chvatal. “The tail of the hypergeometric distribution”. In: *Discrete Mathematics* 25.3 (1979), pp. 285–287. ISSN: 0012-365X. DOI: [https://doi.org/10.1016/0012-365X\(79\)90084-0](https://doi.org/10.1016/0012-365X(79)90084-0). URL: <https://www.sciencedirect.com/science/article/pii/0012365X79900840>.
- [Jam84] James. *The Representation Theory of the Symmetric Group*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 1984.
- [Dam87] Ivan Damgård. “Collision Free Hash Functions and Public Key Signature Schemes”. In: *Advances in Cryptology - EUROCRYPT '87, Workshop on the Theory and Application of Cryptographic Techniques, Amsterdam, The Netherlands, April 13-15, 1987, Proceedings*. Vol. 304. Lecture Notes in Computer Science. Springer, 1987, pp. 203–216. DOI: [10.1007/3-540-39118-5_19](https://doi.org/10.1007/3-540-39118-5_19).

- [Mer88] Ralph C. Merkle. "A Digital Signature Based on a Conventional Encryption Function". In: *Advances in Cryptology — CRYPTO '87*. Ed. by Carl Pomerance. Berlin, Heidelberg: Springer Berlin Heidelberg, 1988, pp. 369–378. ISBN: 978-3-540-48184-3.
- [Mer90] Ralph C. Merkle. "A Certified Digital Signature". In: *Advances in Cryptology — CRYPTO' 89 Proceedings*. Ed. by Gilles Brassard. New York, NY: Springer New York, 1990, pp. 218–238. ISBN: 978-0-387-34805-6.
- [Gro96] Lov K. Grover. "A fast quantum mechanical algorithm for database search". In: *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*. STOC '96. Philadelphia, Pennsylvania, USA: Association for Computing Machinery, 1996, pp. 212–219. ISBN: 0897917855. DOI: [10.1145/237814.237866](https://doi.org/10.1145/237814.237866). URL: <https://doi.org/10.1145/237814.237866>.
- [Jon96] Andrew R. Jones. "A Combinatorial Approach to the Double Cosets of the Symmetric Group with respect to Young Subgroups". In: *European Journal of Combinatorics* 17.7 (1996), pp. 647–655. ISSN: 0195-6698. DOI: <https://doi.org/10.1006/eujc.1996.0056>. URL: <https://www.sciencedirect.com/science/article/pii/S0195669896900566>.
- [Ben+97] Charles H. Bennett et al. "Strengths and Weaknesses of Quantum Computing". In: *SIAM Journal on Computing* 26.5 (Oct. 1997), pp. 1510–1523. ISSN: 1095-7111. DOI: [10.1137/S0097539796300933](https://doi.org/10.1137/S0097539796300933). URL: <http://dx.doi.org/10.1137/S0097539796300933>.
- [Sho97] Peter W. Shor. "Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer". In: *SIAM Journal on Computing* 26.5 (Oct. 1997), pp. 1484–1509. ISSN: 1095-7111. DOI: [10.1137/S0097539795293172](https://doi.org/10.1137/S0097539795293172). URL: <http://dx.doi.org/10.1137/S0097539795293172>.
- [Boy+98] Michel Boyer et al. "Tight Bounds on Quantum Searching". In: *Fortschritte der Physik* 46.4–5 (June 1998), pp. 493–505. ISSN: 1521-3978. DOI: [10.1002/\(sici\)1521-3978\(199806\)46:4/5<493::aid-prop493>3.0.co;2-p](https://doi.org/10.1002/(sici)1521-3978(199806)46:4/5<493::aid-prop493>3.0.co;2-p). URL: [http://dx.doi.org/10.1002/\(SICI\)1521-3978\(199806\)46:4/5%3C493::AID-PROP493%3E3.0.CO;2-P](http://dx.doi.org/10.1002/(SICI)1521-3978(199806)46:4/5%3C493::AID-PROP493%3E3.0.CO;2-P).
- [Zal99] Christof Zalka. "Grover's quantum searching algorithm is optimal". In: *Phys. Rev. A* 60 (4 Oct. 1999), pp. 2746–2751. DOI: [10.1103/PhysRevA.60.2746](https://doi.org/10.1103/PhysRevA.60.2746). URL: <https://link.aps.org/doi/10.1103/PhysRevA.60.2746>.

- [Bea+01] Robert Beals et al. “Quantum lower bounds by polynomials”. In: *J. ACM* 48.4 (July 2001), pp. 778–797. ISSN: 0004-5411. DOI: [10.1145/502090.502097](https://doi.org/10.1145/502090.502097). URL: <https://doi.org/10.1145/502090.502097>.
- [Amb02] Andris Ambainis. “Quantum Lower Bounds by Quantum Arguments”. In: *Journal of Computer and System Sciences* 64.4 (2002), pp. 750–767. ISSN: 0022-0000. DOI: <https://doi.org/10.1006/jcss.2002.1826>. URL: <https://www.sciencedirect.com/science/article/pii/S002200000291826X>.
- [Ber+08] Guido Bertoni et al. “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.
- [DH09] Cătălin Dohotaru and Peter Høyer. “Exact Quantum Lower Bound for Grover’s Problem”. In: *Quantum Info. Comput.* 9.5 (May 2009), pp. 533–540. ISSN: 1533-7146.
- [Amb+11] Andris Ambainis et al. “Symmetry-Assisted Adversaries for Quantum State Generation”. In: *2011 IEEE 26th Annual Conference on Computational Complexity*. IEEE, June 2011. DOI: [10.1109/ccc.2011.24](https://doi.org/10.1109/ccc.2011.24). URL: <http://dx.doi.org/10.1109/CCC.2011.24>.
- [Ber+11a] G. Bertoni et al. *Cryptographic sponge functions*. Submission to NIST (Round 3). 2011. URL: <http://sponge.noekeon.org/CSF-0.1.pdf>.
- [Ber+11b] G. Bertoni et al. *The Keccak SHA-3 submission*. Submission to NIST (Round 3). 2011. URL: <http://keccak.noekeon.org/Keccak-submission-3.pdf>.
- [Bon+11] Dan Boneh et al. “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.
- [KL14] Jonathan Katz and Yehuda Lindell. *Introduction to Modern Cryptography, Second Edition*. 2nd. Chapman & Hall/CRC, 2014. ISBN: 1466570261.
- [Cza+17] Jan Czakowski et al. *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>.
- [AHU18] Andris Ambainis, Mike Hamburg, and Dominique Unruh. *Quantum security proofs using semi-classical oracles*. Cryptology ePrint Archive, Paper 2018/904. <https://eprint.iacr.org/2018/904>. 2018. URL: <https://eprint.iacr.org/2018/904>.

- [Zha18] Mark Zhandry. *How to Record Quantum Queries, and Applications to Quantum Indifferentiability*. Cryptology ePrint Archive, Paper 2018/276. <https://eprint.iacr.org/2018/276>. 2018. URL: <https://eprint.iacr.org/2018/276>.
- [Cza+19] Jan Czakowski et al. *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>.
- [Ala+20] Gorjan Alagic et al. “On Quantum Chosen-Ciphertext Attacks and Learning with Errors”. In: *Cryptography* 4.1 (2020). ISSN: 2410-387X. DOI: [10.3390/cryptography4010010](https://doi.org/10.3390/cryptography4010010). URL: <https://www.mdpi.com/2410-387X/4/1/10>.
- [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>.
- [Zha21] Mark Zhandry. “Redeeming Reset Indifferentiability and Applications to Post-quantum Security”. In: *Advances in Cryptology – ASIACRYPT 2021: 27th International Conference on the Theory and Application of Cryptology and Information Security, Singapore, December 6–10, 2021, Proceedings, Part I*. Singapore, Singapore: Springer-Verlag, 2021, pp. 518–548. ISBN: 978-3-030-92061-6. DOI: [10.1007/978-3-030-92062-3_18](https://doi.org/10.1007/978-3-030-92062-3_18). URL: https://doi.org/10.1007/978-3-030-92062-3_18.
- [Ala+22] Gorjan Alagic et al. “Post-Quantum Security of the Even-Mansour Cipher”. In: *Advances in Cryptology – EUROCRYPT 2022*. Ed. by Orr Dunkelman and Stefan Dziembowski. Cham: Springer International Publishing, 2022, pp. 458–487. ISBN: 978-3-031-07082-2.
- [FGK22] Cody Freitag, Ashrujit Ghoshal, and Ilan Komargodski. “Time-Space Tradeoffs for Sponge Hashing: Attacks and Limitations for Short Collisions”. In: *Advances in Cryptology – CRYPTO 2022*. Ed. by Yevgeniy Dodis and Thomas Shrimpton. Cham: Springer Nature Switzerland, 2022, pp. 131–160. ISBN: 978-3-031-15982-4.
- [Ros22] Ansis Rosmanis. *Tight Bounds for Inverting Permutations via Compressed Oracle Arguments*. 2022. arXiv: [2103.08975](https://arxiv.org/abs/2103.08975) [quant-ph].
- [Aks+23] Akshima et al. “On Time-Space Lower Bounds for Finding Short Collisions in Sponge Hash Functions”. In: *Theory of Cryptography: 21st International Conference, TCC 2023, Taipei, Taiwan, November 29–December 2, 2023, Proceedings, Part III*. Taipei, Taiwan: Springer-Verlag, 2023, pp. 237–270. ISBN: 978-3-031-48620-3. DOI: [10.1007/978-3-031-48620-3](https://doi.org/10.1007/978-3-031-48620-3).

78-3-031-48621-0_9. URL: https://doi.org/10.1007/978-3-031-48621-0_9.

- [Ala+23] Gorjan Alagic et al. *On the Two-sided Permutation Inversion Problem*. 2023. arXiv: [2306.13729](https://arxiv.org/abs/2306.13729) [quant-ph].
- [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](https://doi.org/10.1007/978-981-99-8730-6_12). URL: https://doi.org/10.1007/978-981-99-8730-6_12.
- [MMW24] Christian Majenz, Giulio Malavolta, and Michael Walter. *Permutation Superposition Oracles for Quantum Query Lower Bounds*. Cryptology ePrint Archive, Paper 2024/1140. <https://eprint.iacr.org/2024/1140>. 2024. URL: <https://eprint.iacr.org/2024/1140>.
- [Reg24] Oded Regev. *An Efficient Quantum Factoring Algorithm*. 2024. arXiv: [2308.06572](https://arxiv.org/abs/2308.06572) [quant-ph].
- [Wil] Mark Wildon. *A Model for the Double Cosets of Young Subgroups*. Royal Holloway, University of London. URL: <https://www.ma.rhul.ac.uk/~uvah099/Maths/doubleRevised2.pdf>.