Quantum Indifferentiability of SHA-3

Jan Czajkowski[*]

QuSoft, University of Amsterdam

February 22, 2021

Abstract

In this paper we prove quantum indifferentiability of the sponge construction instantiated with random (invertible) permutations. With this result we bring the post-quantum security of the standardized SHA-3 hash function to the level matching its security against classical adversaries. To achieve our result, we generalize the compressed-oracle technique of Zhandry (Crypto'19) by defining and proving correctness of a compressed permutation oracle. We believe our technique will find applications in many more cryptographic constructions.

Contents

---

[*]j.czajkowski@uva.nl

1

1 Introduction

With efforts to build a quantum computer on the rise, cryptography resilient to quantum attackers (i.e. attackers equipped with a large-scale quantum computer) becomes a necessity. *Post-quantum* cryptography [BBD09], in other words quantum-safe (classical) cryptography, is considered to be important enough that the NIST organized the process to standardize post-quantum cryptographic schemes [NIS17].

A prominent primitive, used in many of the submissions to the NIST standardization process, is the SHA-3 hash function [NIS14]. The cryptographic construction used in the design of this hash function, is the sponge construction [Ber+07]. One of the reasons the sponge construction was chosen as the backbone of this standardized hash function, was the strong security guarantees offered by the proof of its *indifferentiability* from a random oracle [Ber+08]. Indifferentiability [MRH04] is a notion similar to the Ransom Oracle Model (ROM) [BR93]—where we claim a hash function is a random oracle—but much better tailored to cryptographic constructions. To prove indifferentiability, we assume the internal function called by the construction to be ideal (a uniformly random function or permutation) but also give the adversary access to it.

For five years now, the research community has worked to rigorously prove post-quantum security of the sponge construction [Cza+18; CHS19; Cza+19; Unr21]. One of the goals of this line of research was to match the security level guaranteed by proofs valid for classical adversaries. In [Cza+19] the authors managed to prove quantum indifferentiability, but with one shortcoming, they only treated uniformly random internal functions. Now is it important to note, that SHA-3 is based on an invertible *permutation*, this design choice is dictated, among other reasons, by the cryptographic practice showing that it is much easier to design secure permutations than one-way functions. In [Unr21] Unruh claims to prove quantum collision-resistance of SHA-3. In our work, we finally manage to close the gap between classical and post-quantum security of SHA-3: We prove quantum indifferentiability of the sponge construction instantiated with random permutations.

The natural proof technique suitable for indifferentiability proofs is that of lazy-sampling. In a recent work, Zhandry developed the compressed-oracle technique [Zha19], which is the quantum version of lazy-sampling. Until now, however, the quantum techniques did not cover lazy-sampling of random (invertible) permutations. Changing that, is the main technical contribution of this paper. Basing on the quantum game-playing framework of [Cza+19], we define a quantum compressed permutation oracle. Our result introduces errors, but they are smaller than the leading terms in the distinguishing advantage of the indifferentiability adversary. Arguably, the total bound on the distinguishing advantage that we prove is tight. Although we do not show an algorithm with a matching query complexity, our bound on the distinguishing advantage[1] of $O(\sqrt{q^3/2^c})$ gives rise to security of up to $\Omega\left(2^{c/3}\right)$ queries. A generic collision finding algorithm for the sponge construction, e.g. from [Cza+18], is expected to run in $O\left(2^{c/3}\right)$ queries to the internal function. In [Unr21], Unruh also proposes a technique for dealing with invertible random permutations in a quantum world, his approach, however, is not suitable for indifferentiability proofs, as the permutation oracle is not efficiently simulatable.

In section 2 we give more details on the methods we use to achieve our results. The compressed permutation oracle is defined and proven indistinguishable from a random permutation in section 3. The main result of this paper, quantum indifferentiability of the sponge construction instantiated with random permutations, is located in section 4. A reader mostly interested in the application of our new technique can go directly to section 4, as we provide a high level introduction to the main ingredients of the proof.

---

[1]By distinguishing advantage, we mean the absolute value of the difference of the probabilities that the adversary outputs 1 when interacting with the construction or a random oracle.

2 Preliminaries

In this section we start by going over notation used in this paper. Later we present some key concepts that we use to achieve our results. A summary of symbols used throughout the paper can be found in the Symbol Index.

In this paper we assume the reader has basic knowledge of quantum computing. If any concept is not familiar though, we refer to [NC10; Wol11]. A quantum state is a unit vector in a Hilbert space $|\psi\rangle \in \mathbb{C}^d$. We often refer to vectors of norm smaller than 1 as states as well. In what follows we denote the Euclidean norm of a vector by $\||\psi\rangle\|$. The trace norm of an operator A on a Hilbert space is defined as the trace of the absolute value of the operator: $\|\mathsf{A}\|_1 := \mathrm{Tr}\sqrt{\mathsf{A}^\dagger \mathsf{A}}$. Subspaces of tensor products of several Hilbert spaces are referred to by superscripts, so if $\mathcal{H}_D = \bigotimes_{i=1}^q \mathcal{H}_{X_i} \otimes \mathcal{H}_{Y_i}$, then D^X refers to $\bigotimes_{i=1}^q \mathcal{H}_{X_i}$, and same for Y. We write A^D to denote that A acts on register D.

We write $[N] := \{0, 1, \ldots, N-1\}$ for the set of size N. Whenever we consider sets of bits $\{0,1\}^n$, we denote the bitwise XOR by $\oplus$.

To denote vectors we write $\vec{x} \in [N]^s$. We also use set operations[2] on these vectors, by identifying them with sets with elements from $\vec{x}$. We consider $\vec{x}$ to be always sorted in a raising fashion and contain no repetitions, so set operations are unambiguous. By $\vec{y}_{\vec{x}}$ we denote $(y_{x_1}, \ldots, y_{x_s})$, where $s = |\vec{x}|$. Vector $\vec{y}_{\vec{x}}$ is sorted according to $\vec{x}$.

The set of outputs of *injective* partial functions is denoted by

$$\mathcal{I}(s \mid \vec{y}) := \{\vec{y}' \in [N]^s : \forall i, j \neq i, y_i' \neq y_j' \wedge y_i' \notin \vec{y}\}, \text{ and its size by} \tag{1}$$

$$(N)_s := |\mathcal{I}(s)| = N(N-1)\cdots(N-s+1), \tag{2}$$

where in the definition of $(N)_s$ we use the notation $\mathcal{I}(s) := \mathcal{I}(s \mid \emptyset)$.

By $x \leftarrow \mathsf{A}$ we denote sampling x from a distribution or getting the output of a randomized algorithm. By square brackets we denote (classical or quantum) oracle access to some algorithm, we also use A^H if the oracle is denoted by a more confined symbol.

An object that we make heavy use of in this paper is the random oracle:

$$\mathsf{R} : \{0,1\}^* \times \mathbb{N} \to \{0,1\}^*. \tag{3}$$

A random oracle grants access to a function sampled from distribution $\mathfrak{R}$ on functions $\{0,1\}^* \times \mathbb{N} \to \{0,1\}^*$, that is defined as follows: To sample a function $h \leftarrow \mathfrak{R}$ we

- choose g uniformly at random from $\{g : \{0,1\}^* \to \{0,1\}^\infty\}$, where by $\{0,1\}^\infty$ we denote the set of infinitely long bitstrings,

- for each $(x, \ell) \in \{0,1\}^* \times \mathbb{N}$ set $h(x, \ell) := \lfloor g(x) \rfloor_\ell$, that is output the first ℓ bits of the output of g.

2.1 Indifferentiability

Whenever, the hash function is constructed out of a public internal function (like in the case of, e.g., SHA-2 [NIS15] and SHA-3 [NIS14]), a security notion that includes the adversary's access to the internal (and publicly specified) function and aims at showing that the hash function behaves as a uniformly random function, is *indifferentiability* [MRH04].

In the setting of indifferentiability the adversary is given access to two *interfaces*, the *public* interface, denoted by "pub" and the *private* interface, denoted by "priv". Indifferentiability of a system C is phrased in terms of a distinguishing game, where the adversary distinguishes between the real and the ideal world. In the real world, the private interface is the construction calling the internal function, denoted as $\mathsf{C}^{\mathrm{priv}}[\mathsf{C}^{\mathrm{pub}}]$, and the public interface is the internal

---

[2]Such as the union $\cup$, intersection $\cap$, or subtraction $\backslash$.

function C^{pub}. In the ideal world, where the construction is often replaced by a random oracle, the private interface is just R_k^{priv}, but the public interface is provided by a *simulator* that replaces the internal function, she is also given access to the ideal system and the interface is $S[R_k^{pub}]$. In case of R being a random oracle both the public and the private interfaces of R are equal and just return outputs of a random function.

Definition 1 (Indifferentiability [MRH04])**.** *A cryptographic system* C *is* (q, ε)*-indifferentiable from* R*, if there is an efficient* (*classical or quantum*) *simulator* S *and a negligible function* ε *such that for any efficient* (*classical or quantum*) *distinguisher* D *with binary output (0 or 1) the advantage*

$$\left| \mathbb{P}\left[1 \leftarrow D[C_k^{priv}[C_k^{pub}], C_k^{pub}] \right] - \mathbb{P}\left[1 \leftarrow D[R_k^{priv}, S[R_k^{pub}]] \right] \right| \leq \varepsilon(k) \,, \tag{4}$$

where k *is the security parameter. The distinguisher makes at most* q (*classical or quantum*) *queries.*

2.2 Compressed-Oracles Technique

The compressed oracle technique was developed by Zhandry [Zha19]. The key idea of this technique is to purify a quantum-accessible random oracle so all the randomness is kept in a quantum register. The standard oracle for a uniformly random function is defined as:

$$\mathsf{StO}|x, y\rangle_{XY} \sum_{f \in \mathcal{F}} \frac{1}{\sqrt{|\mathcal{F}|}} |f\rangle_F = \sum_{f \in \mathcal{F}} \frac{1}{\sqrt{|\mathcal{F}|}} |x, y + f(x)\rangle_{XY} |f\rangle_F, \tag{5}$$

where $x, y \in \mathcal{X}$, for $\mathcal{X} = [N]$, addition is done modulo N, and f is the full truth table of the function from $\mathcal{F} := \{f : \mathcal{X} \to \mathcal{X}\}^3$. The update procedure StO just adds the correct row of f to adversary's Y register. The crucial fact is that an adversary A interacting with StO acting on the oracle register from Eq. (5) has zero probability of distinguishing it from $f \xleftarrow{\$} \mathcal{F}$.

A compressed oracle CStO is an oracle with the oracle register F compressed to register D that just holds the outputs to queries asked by the adversary A. The update procedure in this case is more complicated but allows to keep a superposition of small database of the form $D = ((x_1, y_1), \ldots, (x_q, y_q))$, where q is the bound on the maximal number of quantum queries made by A. In the case A made $s < q$ queries, the last $q - s$ entries hold $\perp$ in the X-type part of the entries. More details on this technique can be found in [Zha19; LZ19; Cza+19; HI19; Chu+20; Unr21]. Similarly to StO, the distinguishing advantage between CStO and $f \xleftarrow{\$} \mathcal{F}$ is 0.

A change of the adversary's basis, that can be done by applying the quantum Fourier transform $\mathsf{QFT}_N |y\rangle := \frac{1}{\sqrt{N}} \sum_{\eta \in [N]} \omega_N^{\eta \cdot y} |\eta\rangle$, where $\omega_N := e^{\frac{2\pi i}{N}}$ is the N-th root of unity, gives rise to the phase oracle:

$$\mathsf{PhO}|x, \eta\rangle_{XY} \sum_{f \in \mathcal{F}} \frac{1}{\sqrt{|\mathcal{F}|}} |f\rangle_F = \sum_{f \in \mathcal{F}} \frac{1}{\sqrt{|\mathcal{F}|}} \omega_N^{\eta \cdot f(x)} |x, \eta\rangle_{XY} |f\rangle_F. \tag{6}$$

Similarly, by changing the adversary's interface to the Fourier basis, we get the compressed phase oracle CPhO. For readers less familiar with quantum accessible oracles, we note that a quantum query does not always increase the number of input-output pairs A has knowledge of. Say A queries (x, η_x) at some point. If A later queries the same input with a different η, then she updates her knowledge of $f(x)$. If $\eta = -\eta_x$ then she might even "forget" the query x, or—if she interacts with a compressed oracle—remove x from the database.

Whenever we talk about CStO and CPhO, we consider them to be oracles for the uniform distribution. Moreover, when providing the subscript $\mathcal{X}$, by CPhO$_{\mathcal{X}}$ we mean that functions

---

[3]In general, the standard and compressed oracle can be defined for any set of functions.

4

are distributed uniformly from $\{f : \mathcal{X} \to \mathcal{X}\}$. More on compressed oracles for different distributions can be found in [Cza+19].

Building upon the compressed oracle technique the authors of [Cza+19] developed a quantum game-playing framework. The key object of this framework, that we make heavy use of, is the *punctured* compressed oracle. Puncturing is defined as performing a binary measurement of the oracle register D of a compressed oracle after every query done by A. As long as the probability of one of the outputs (say 0) is almost 1, the puncturing does not give the adversary a major distinguishing advantage. The measurement, however, allows us to argue about the adversary's behavior conditioned on the database begin of some specific form.

The measurements in punctured oracles are binary projective measurement. We project to databases that fulfill a given *relation*. A relation is a subset of all possible databases of size bounded by q. An important relation, that we use in the reminder of this paper, is the *collision* relation:

$$R_{\text{coll}} := \left\{ D \in \bigcup_{s=1}^{q} (\mathcal{X} \times \mathcal{Y})^s : \exists i, j \neq i, x_i \neq x_j, y_i = y_j \right\}, \tag{7}$$

where $D = ((x_1, y_1), \ldots, (x_s, y_s))$. Measuring a relation on the quantum compressed oracle includes checking the actual size of the database, so it omits the entries starting with $\perp$.

A formal definition of a compressed oracle H punctured on relation R, denoted by $H \setminus R$, is stated below. More details on the above concepts can be found in [Cza+19].

Definition 2 (Punctured compressed oracle $H \setminus R$, Definition 9 in [Cza+19])**.** *Let H be a compressed oracle and R a relation on its database. The punctured compressed oracle $H \setminus R$ is equal to H, except that R is measured after every query. By* Find *we denote the event that R outputs 1 at least once among all queries.*

2.3 Game-Playing Proofs

Game-playing proofs are a common framework used in cryptography. Originally developed by Bellare and Rogaway [BR06] the main concept is that if two games—in other words, interactive protocols—are identical until a "bad" event happens, then they are indistinguishable except for the probability that this event occurs.

Using recently developed techniques, this reasoning can be applied to *quantum* interactive protocols as well. Below, we present the classical and the quantum game-playing framework. We use the statements from below in our indifferentiability proofs in section 4.

2.3.1 Classical Game-Playing Proofs

Two games are considered to be *identical-until-bad* if they are syntactically identical except for code that follows after setting Bad to true. Bad is a special flag, we also write Bad to denote the event "Bad set to true". The main ingredient of the classical game-playing framework is the fundamental game-playing lemma:

Lemma 3 (Fundamental lemma of game-playing, Lemma 2 in [BR06])**.** *Let G and H be identical-until-bad games and let A be an adversary that outputs a bit b. Then*

$$\left| \mathbb{P}[1 \leftarrow A^H] - \mathbb{P}[1 \leftarrow A^G] \right| \leq \mathbb{P}[\text{Bad} = 1 : A^G]. \tag{8}$$

2.3.2 Quantum Game-Playing Proofs

The quantum game-playing framework, developed in [Cza+19], makes use of two powerful proof techniques. One being the compressed oracle technique by Zhandry [Zha19], the other

being the One-way To Hiding (O2H) lemma by Unruh [Unr14; AHU19]. The former (which we described in section 2.2), together with the idea of puncturing oracles, is used to define quantum games. The latter provides the quantum counterpart of the fundamental game playing lemma.

Here we present the version of the O2H lemma from [AHU19] that was later repurposed to treat punctured compressed oracles in [Cza+19].

Lemma 4 (Compressed oracle O2H, Theorem 10 in [Cza+19]). *Let R_1 and R_2 be relations on the database of a quantum oracle* H. *Let* A *be an oracle algorithm making q quantum queries, then*

$$\left| \mathbb{P}[1 \leftarrow A^{H\setminus R_1}] - \mathbb{P}[1 \leftarrow A^{H\setminus R_1 \cup R_2}] \right| \leq \sqrt{(q+1)\mathbb{P}[\text{Find} : A^{H\setminus R_1 \cup R_2}]}, \qquad (9)$$

where Find *is the event that measuring $R_1 \cup R_2$ succeeds.*

An important comment that relates to the main subject of this work is that the above lemma works just as well for *invertible* oracles, such as random permutations. Following the proof, in [Cza+19], we notice that it relies only on measuring the database and the general interface of compressed oracles. These aspects do not change in our definition of the compressed permutation oracle. More discussion on which, can be found in section 3.

The notion of identical-until-bad games can also be brought to the world of quantum compressed oracles.

Definition 5 (Almost identical oracles, Definition 11 in [Cza+19]). *Let* H *and* G *be compressed oracles and R_i, $i = 1, 2$ relations on their databases. We call the oracles* $H \setminus R_1$ *and* $G \setminus R_2$ *almost identical if they are equal conditioned on the events $\neg\text{Find}_1$ and $\neg\text{Find}_2$ respectively, i.e. for any quantum algorithm* A

$$\mathbb{P}[1 \leftarrow A^{H\setminus R_1} \mid \neg\text{Find}_1] = \mathbb{P}[1 \leftarrow A^{G\setminus R_2} \mid \neg\text{Find}_2]. \qquad (10)$$

Distinguishing advantage for almost-identical oracles can be bounded using the following lemma.

Lemma 6 (Distinguishing almost identical punctured oracles, Lemma 12 in [Cza+19]). *If* $H \setminus R_1$ *and* $G \setminus R_2$ *are almost identical according to Def.5 then for any $b \in \{0, 1\}$*

$$\left| \mathbb{P}[1 \leftarrow A^{H\setminus R_1}] - \mathbb{P}[1 \leftarrow A^{G\setminus R_2}] \right| \leq 2\mathbb{P}[\text{Find}_1 : A^{H\setminus R_1}] + 2\mathbb{P}[\text{Find}_2 : A^{G\setminus R_2}]. \qquad (11)$$

2.4 Sponge Construction

The sponge construction SPONGE [Ber+07] is prominently used to design the standardized hash function SHA-3 [NIS14]. It can be instantiated with a function or a permutation $f : \{0,1\}^{r+c} \to \{0,1\}^{r+c}$. The construction is divided in two phases: In the absorbing phase the padded input message is XOR-ed block by block (blocks are r-bits long) into the first r bits of the internal state (consisting of $r + c$ bits) of SPONGE$_f$ and after each block f is applied to the whole state. In the following squeezing phase the first r bits are output and f is applied, this repeats until ℓ bits of output is generated. This distinction into the first r bits and the last c bits calls for special notation. For a state $s \in \{0,1\}^{r+c}$, the first r bits are called the *outer* part and are denoted by $\bar{s}$, the last c bits are called the *inner* part and are denoted by $\hat{s}$. By PAD : $\{0,1\}^* \to (\{0,1\}^r)^*$ (where by $(\{0,1\}^r)^*$ we denote the strings consisting of an arbitrary number of r-bit blocks) we denote a padding function. The requirements on PAD in the sponge construction are that it is an efficiently computable injection such that $|\text{PAD}(m)| \geq r$ and that the last bit of PAD(m) is never 0 (this ensures injectivity for inputs of different lengths). The *rate* of the sponge is r and the *capacity* is c, these key parameters, together with the internal function and the padding function give rise to the hash function SPONGE$_f$[PAD, r, c] : $\{0,1\}^* \times \mathbb{N} \to \{0,1\}^*$, where $\{0,1\}^* := \bigcup_{n=0}^{\infty} \{0,1\}^*$. In

---

Algorithm 1: $\textsc{Sponge}_f[\textsc{pad}, r, c]$

Input : $m \in \{0,1\}^*$, $\ell \geq 0$
Output: $z \in \{0,1\}^\ell$

1   $p := \textsc{pad}(m)$
2   $s := (0,0) \in \{0,1\}^r \times \{0,1\}^c$
3   **for** $i = 1$ to $|p|_r$ **do**   `// Absorbing phase`
4         $s := (\bar{s} \oplus \lfloor p \rfloor_i, \hat{s})$
5         $s := f(s)$
6   $z := \bar{s}$ `// Squeezing phase`
7   **while** $|z| < \ell$ **do**
8         $s := f(s)$
9         $z := z \| \bar{s}$
10   Output z

---

what follows, by $|p|_r$ we denote the number of r-bit blocks in p and by $\lfloor p \rfloor_i$ we denote the i-th r-bit block of p. The formal definition of the sponge construction is presented in Algorithm 1.

It was already noticed in [Ber+07] that one can associate a graph $G = (\mathcal{V}, \mathcal{E})$ with the internal function of a sponge. This graph is called the *sponge graph* and plays an important role in indifferentiability proofs of the sponge construction [Ber+08; Cza+19]. The set of nodes of this graph is $\mathcal{V} := \{0,1\}^{r+c}$, it corresponds to all possible states of the sponge. A directed edge connects any two nodes (s,t) whenever $f(s) = t$, meaning there are 2^{r+c} edges in $\mathcal{E}$. From each node starts exactly one edge. A *supernode* groups the nodes with the same inner-part, thus we have 2^c supernodes and each such supernode consists of 2^r nodes. Edges between nodes are also edges between supernodes.

The initial state of any call of $\textsc{Sponge}$ is the $(0^r, 0^c)$ node. This node is called the *root*. The first r-bit block $\lfloor p \rfloor_1$ of the padded message $p = \textsc{pad}(m)$ is then added to the outer-part of the state and fed to the internal function $f(\lfloor p \rfloor_1 \| 0^c) = s_2$. This evaluation of f defines the edge $(\lfloor p \rfloor_1 \| 0^c, s_2) \in \mathcal{E}$. In the absorbing phase this operation is repeated until we run out of blocks of p. In the squeezing phase we no longer modify the state, in other words just add 0^r to the outer part. Note that knowing just p and G we can get to the last node traversed by $\textsc{Sponge}_f(m)$. This leads to the definition of a *sponge path*.

Definition 7 (Sponge path, Definition 3 in [Ber+08]). *First, the empty string is a sponge path to the node $0^r \| 0^c$. Then, if p is a sponge path to node $s = \bar{s} \| \hat{s}$ and there is an edge $(\bar{s} \oplus a \| \hat{s}, t)$ in the sponge graph G, $p' = p \| a$ is a sponge path to node t.*

Given the above definition, let us say that p forms a sponge path to s, then we define a function

$$\mathsf{SpPath}(s, G) := p. \tag{12}$$

Output of the above function is the input to the construction $\textsc{Sponge}_f(., \ell = r)$ that yields the output $\bar{s}$.

In indifferentiability proofs, the simulator has access to the *simulator graph*. The graph kept by the simulator differs from the sponge graph discussed above by the number of edges in it. As the simulator lazy samples the internal function f the set of edges $\mathcal{E}$ grows by at most one edge per a single adversary's query. Other than that, everything that has been said above holds. We refer to the simulator graph G as just the (sponge) graph whenever it is clear from context.

A supernode is called *rooted* if there is a path (a regular path that is just a set of edges connected by the end-start nodes) leading to it that starts in the root (the 0^c-supernode). The

set $\mathcal{R}$ is the set of all rooted supernodes in G. By $\mathcal{U}$ we denote the set of supernodes with a node with an outgoing edge.

A simulator graph is called *saturated* if $\mathcal{R} \cup \mathcal{U} = \{0,1\}^c$. It means that for every inner state in $\{0,1\}^c$ there is an edge in G that leads to it from 0^c (the root) or leads from it to another node. Saturation will be important in the proof of indifferentiability as the simulator wants to pick outputs of f without colliding inner parts (so not in $\mathcal{R}$) and making the path leading from 0^c to the output longer by just one edge (so not in $\mathcal{U}$).

3 Compressed Oracle for Random Permutations

In this section we define the compressed oracle for *invertible* random permutation. Our approach is rather simple, we use a compressed oracle for the uniform distribution and puncture it on collisions. The collision-free databases can be easily inverted and the same update procedure as for the forward direction provides the backward-direction oracle.

The hard part in this section is proving that our compressed permutation oracle is in fact indistinguishable from a full permutation oracle. The full oracle keeps a superposition of all bijective functions and responds to both forward and backward queries.

Our approach to proving indistinguishability relies on a relatively simple auxiliary *good* state. The good state is generated by the adversary interacting with the compressed oracle for the uniform distribution, after the interaction we project the database to contain only collision-free databases. Such state is easy to handle and we can prove its closeness of the state coming from the full permutation oracle. We then reuse the proof from [Cza+19] to show closeness of the good state and the compressed oracle state.

In the end of this section we also provide bounds on the probability of Find for punctured compressed permutation oracles. These results are important for our indifferentiability proof.

3.1 Definition of the Compressed Permutation Oracle

The core of the compressed permutation oracle is a compressed oracle for the uniform distribution punctured on the collision relation $\mathsf{CPhO}_\mathcal{X} \setminus R_{\text{coll}}$. To allow for inverse access we define a unitary that flips the database, so the outputs become inputs and vice versa. Then queries are handled with the same $\mathsf{CPhO}_\mathcal{X} \setminus R_{\text{coll}}$ procedure.

In more detail, we provide access to the inverse of the permutation by flipping the database: treat the content of D^Y as inputs and D^X as outputs and sort by D^Y, we call the appropriate unitary Flip. Note that Flip^D works as anticipated only if both in D^Y and D^X there are no collisions; The definition of Flip can be easily extended to a full unitary by, e.g., keeping the order unchanged within a tuple with colliding outputs. We thereby obtain a (superposition) database where the entries (y_i, x_i) are sorted in a rising order according to the first value. A more detailed description of Flip is:

1. Controlled on D^Y copy D^X to a fresh D'^Y and arrange in an increasing order according to values of y_i in D^Y.

2. Controlled on D^X and D'^Y copy D^Y to a fresh D'^X in the new order.

3. Controlled on D' erase the old D: Take the smallest x_i in D'^Y and subtract it from the first register of D^X, also subtract the corresponding y_i, and so on.

A formal definition of the compressed permutation oracle $\mathsf{CPerO}_\mathcal{X}$ is:

Definition 8. *The compressed permutation oracle for permutations $f : \mathcal{X} \to \mathcal{X}$ is defined as the following procedure:*

$$\mathsf{CPerO}_{\mathcal{X}} := \mathsf{CPhO}_{\mathcal{X}} \setminus R_{\mathrm{coll}}, \tag{13}$$

$$\mathsf{CPerO}_{\mathcal{X}}^{-1} := \mathsf{Flip}^D \circ (\mathsf{CPhO}_{\mathcal{X}} \setminus R_{\mathrm{coll}}) \circ \mathsf{Flip}^D, \tag{14}$$

where Flip *is defined above and* $\mathsf{CPhO}_{\mathcal{X}}$ *is the compressed phase oracle for the uniform distribution over* $\{f : \mathcal{X} \to \mathcal{X}\}$.

3.2 Indistinguishability of CPerO and PerO

By PerO we denote the full oracle for a uniform distribution over permutations, technically it is defined in the same way as PhO from Eq. (6) but we always use it with the oracle register holding a superposition of permutations:

$$\mathsf{PerO}_{\mathcal{X}} |x, \eta\rangle_{XY} \sum_{f \in \mathcal{P}} \frac{1}{\sqrt{|\mathcal{P}|}} |f\rangle_F = \sum_{f \in \mathcal{P}} \frac{1}{\sqrt{|\mathcal{P}|}} \omega_N^{\eta \cdot f(x)} |x, \eta\rangle_{XY} |f\rangle_F, \tag{15}$$

where $\mathcal{P} := \{f : \mathcal{X} \to \mathcal{X} : f \text{ is a bijection}\}$. The main statement of this section is the correctness of the compressed permutation oracle.

Theorem 9 (Correctness of $\mathsf{CPerO}_{\mathcal{X}}$). *No quantum algorithm* A *making at most q quantum queries can distinguish* $\mathsf{CPerO}_{\mathcal{X}}$ *from the full permutation oracle* $\mathsf{PerO}_{\mathcal{X}}$ *except for a small error:*

$$\left| \mathbb{P} \left[1 \leftarrow \mathsf{A}^{\mathsf{PerO}_{\mathcal{X}}} \right] - \mathbb{P} \left[1 \leftarrow \mathsf{A}^{\mathsf{CPerO}_{\mathcal{X}}} \right] \right| \leq \frac{18 q^{5/2}}{N - q}, \tag{16}$$

where $N = |\mathcal{X}|$.

Before getting into the proof of the main statement we suggest the reader to go over the observation presented in Appendix A. The observation is a theorem stating the distinguishability advantage of any adversary trying to discern a uniformly random function from a random permutation using a single quantum query. The approach to the proof of Theorem 17 is quite related to that from of proof of Theorem 9.

In the following three paragraphs we define and provide some useful notation on the two important classes of quantum states that we will make heavy use of.

The permutation states First we go over the full permutation oracle. The state that results from A interacting with the full permutation oracle PerO is

$$|\Psi^{\mathsf{PerO}}\rangle_{AF} := \sum_{\vec{x}, \vec{\eta}, w} \alpha_{\vec{x}, \vec{\eta}, w} |\psi(\vec{x}, \vec{\eta}, w)\rangle_A$$

$$\underbrace{\sum_{\vec{y}_{\vec{x}} \in \mathcal{I}(s)} \frac{1}{\sqrt{(N)_s}} \omega_N^{\vec{\eta} \cdot \vec{y}_{\vec{x}}}}_{(*): \text{ effective queries}} \sum_{\vec{y}_{\mathcal{X} \setminus \vec{x}} \in \mathcal{I}(N - s | \vec{y}_{\vec{x}})} |y_0, y_1, \dots, y_{N-1}\rangle_F, \tag{17}$$

where register A holds all of the adversary's state, this includes her work register, query register, and any other auxiliary registers she decides to use. We denote it by $\psi(\vec{x}, \vec{\eta}, w)$, where w can be any value of finite size. In the above state $\vec{x} = (x_1, \dots, x_s)$ is the set of *effective* queries. As querying a quantum oracle can result in "un-querying" a value, by $\vec{x}$ we denote only the queried values that are entangled with the oracle register after all q queries of A. For these queries we also have $\eta_{x_i} \neq 0$ for all $x_i \in \vec{x}$. The size of $\vec{x}$ is s and $s \leq q$.

The permutation state is the partial trace of the above state

$$\rho_{\mathsf{PerO}} := \mathrm{Tr}_F |\Psi^{\mathsf{PerO}}\rangle_{AF} \langle \Psi^{\mathsf{PerO}}| = \sum_{\vec{x}, \vec{x}', \vec{x}''} \rho_{\mathsf{PerO}}(\vec{x} \cup \vec{x}', \vec{x} \cup \vec{x}''), \tag{18}$$

where in the above sums we have $\vec{x}' \cap \vec{x} = \emptyset$, $\vec{x}'' \cap \vec{x} = \emptyset$, and $\vec{x}'' \cap \vec{x}' = \emptyset$. Moreover $\vec{x} \cup \vec{x}'$ and $\vec{x} \cup \vec{x}''$ range over $\mathcal{X}^q$. The state is defined as

$$\rho_{\mathsf{PerO}}(\vec{x} \cup \vec{x}', \vec{x} \cup \vec{x}'') := \mathrm{Tr}_F |\Psi^{\mathsf{PerO}}(\vec{x} \cup \vec{x}')\rangle_{AF}\langle\Psi^{\mathsf{PerO}}(\vec{x} \cup \vec{x}'')|, \qquad (19)$$

where $|\Psi^{\mathsf{PerO}}(\vec{x} \cup \vec{x}')\rangle_{AF}$ is the branch of the superposition in $|\Psi^{\mathsf{PerO}}\rangle$ corresponding to *effective* queries in $\vec{x} \cup \vec{x}'$.

Below we define *well-formed* matrices. The intuition behind well-formed matrices is that they are parts of quantum states resulting from the interaction of an adversary with the full permutation oracle. They are natural parts of density matrices corresponding to A interacting with PerO.

Definition 10. *A vector is* well-formed *if it is of the form from Eq. (17). Coefficients $\alpha_{\vec{x},\vec{\eta},w}$ can be such that the overall state is sub-normalized.*

A matrix is well-formed *if it is of the form from Eq. (19) with the vectors being well formed.*

The good states In what follows we write $\vec{x}$ to denote the effective queries. The state $|\Psi^{\mathsf{Good}}\rangle_{AD}$ corresponds to the adversary's state after q queries. The size of the database is $s = |\vec{x}|$. After q queries s can range from 0 to q and the joint state of A and the oracle can be a superposition over different database sizes. We denote the outputs given to A by $\vec{y} := (y_1, \ldots, y_s)$. We define the good state as:

$$|\Psi^{\mathsf{Good}}\rangle_{AD} := \sum_{\vec{x},\vec{\eta},w} \alpha_{\vec{x},\vec{\eta},w} |\psi(\vec{x}, \vec{\eta}, w)\rangle_A \sum_{\vec{y} \in \mathcal{I}(s)} \frac{1}{\sqrt{(N)_s}} \omega_N^{\vec{\eta}\cdot\vec{y}} \sum_{y_{s+1},\ldots,y_q \in [N]} \frac{1}{\sqrt{N^{q-s}}}$$
$$|(x_1, y_1), \ldots, (x_s, y_s), (\bot, y_{s+1}), \ldots, (\bot, y_q)\rangle_D. \qquad (20)$$

The good state is also defined as running the adversary and giving her access to the compressed standard oracle $\mathsf{CPhO}_{\mathcal{X}}$. After q queries we project register D to collision free databases.

Inverse queries In what follows we treat elements of tuples of effective queries $\vec{x}$ as pairs in $\{+, -\} \times \mathcal{X}$, where the first entry signifies the interface (forward or backward) of the permutation. We don't write it explicitly throughout this paper because it is not necessary for the main argument. We just note that a query value always implicitly contains the interface it was made to.

Below we write the part corresponding to effective queries of an adversary that queries x_1 in the forward direction and y_2 in the backward direction:

$$\sum_{f \in \mathcal{P}} \frac{1}{\sqrt{N!}} \omega_N^{\eta_1 f(x_1)} \omega_N^{\eta_2 f^{-1}(y_2)} |f\rangle_F = \frac{1}{\sqrt{N!}} \sum_{z_{x_1} \in [N]} \omega_N^{\eta_1 z_{x_1}} |z_{x_1}\rangle \sum_{s_{y_2} \neq x_1} \omega_N^{\eta_2 s_{y_2}} |s_{y_2}\rangle$$
$$\sum_{z_1 \notin \{z_{x_1}, y_2\}} |z_1\rangle \sum_{\vec{z} \in \mathcal{I}(N-2|\{z_{x_1}, y_2, z_1\})} |z_2\rangle \cdots |z_{N-2}\rangle, \qquad (21)$$

where the second line is correct only if $y_2 \neq z_{x_1}$. The above equation exemplifies that the form we wrote the states from Eqs. (17, 20) is suitable to capture inverse queries.

Proof of Theorem 9. We first simplify the distinguishing advantage to include interaction with the punctured oracle conditioned on ¬Find:

$$\left|\mathbb{P}\left[1 \leftarrow \mathsf{A}^{\mathsf{PerO}}\right] - \mathbb{P}\left[1 \leftarrow \mathsf{A}^{\mathsf{CPerO}}\right]\right|$$
$$\leq \left|\mathbb{P}\left[1 \leftarrow \mathsf{A}^{\mathsf{PerO}}\right] - \mathbb{P}\left[1 \leftarrow \mathsf{A}^{\mathsf{CPerO}} \mid \neg\mathsf{Find}\right]\mathbb{P}\left[\neg\mathsf{Find}\right]\right| + \mathbb{P}\left[\mathsf{Find}\right] \qquad (22)$$
$$\leq \left|\mathbb{P}\left[1 \leftarrow \mathsf{A}^{\mathsf{PerO}}\right] - \mathbb{P}\left[1 \leftarrow \mathsf{A}^{\mathsf{CPerO}} \mid \neg\mathsf{Find}\right]\right| + \mathbb{P}\left[\mathsf{Find}\right]. \qquad (23)$$

In the first step above we write down the right term as $\mathbb{P}[1 \leftarrow \mathsf{A}^{\mathsf{CPerO}}] = \mathbb{P}[1 \leftarrow \mathsf{A}^{\mathsf{CPerO}} \mid \neg\mathsf{Find}]\mathbb{P}\left[\neg\mathsf{Find}\right] + \mathbb{P}[1 \leftarrow \mathsf{A}^{\mathsf{CPerO}} \mid \mathsf{Find}]\mathbb{P}\left[\mathsf{Find}\right]$. Next we use the triangle inequality and bound $\mathbb{P}[1 \leftarrow \mathsf{A}^{\mathsf{CPerO}} \mid \mathsf{Find}] \leq 1$. If an adversary outputs 1 with greater probability when she interacts with PerO we instead consider A' that operates exactly like A with the only difference that she flips the output. The second inequality is a result of this reasoning and bounding $\mathbb{P}\left[\neg\mathsf{Find}\right] \leq 1$.

By $\Phi_{\mathsf{CPerO}}^{AD}$ we denote the joint state of A and the database D conditioned on $\neg\mathsf{Find}$ ($\Phi_{\mathsf{CPerO}}^{A}$ is the same state after a partial trace over D). To generate the final answer, any quantum algorithm has to perform a quantum measurement, let us denote by Q_1 is the measurement corresponding to A outputting 1. Then we have

$$\left|\mathbb{P}\left[1 \leftarrow \mathsf{A}^{\mathsf{PerO}}\right] - \mathbb{P}\left[1 \leftarrow \mathsf{A}^{\mathsf{CPerO}} \mid \neg\mathsf{Find}\right]\right|$$
$$= \left|\mathrm{Tr}\left(\mathsf{Q}_1\rho_{\mathsf{PerO}}^{A}\right) - \mathrm{Tr}\left(\mathsf{Q}_1\Phi_{\mathsf{CPerO}}^{A}\right)\right| \leq \frac{1}{2}\left\|\rho_{\mathsf{PerO}}^{A} - \Phi_{\mathsf{CPerO}}^{A}\right\|_1, \tag{24}$$

where the inequality above comes from the fact that trace distance gives the optimal measurement for distinguishing two states [NC10]. Using the bound from Eq. (24) we bound Eq. (23) as follows:

$$\text{Eq. (23)} \leq \frac{1}{2}\left\|\rho_{\mathsf{PerO}}^{A} - \Phi_{\mathsf{CPerO}}^{A}\right\|_1 + \mathbb{P}\left[\mathsf{Find}\right] \tag{25}$$

$$\leq \frac{1}{2}\left\|\rho_{\mathsf{PerO}}^{A} - \rho_{\mathsf{Good}}^{A}\right\|_1 + \frac{1}{2}\left\|\rho_{\mathsf{Good}}^{AD} - \Phi_{\mathsf{CPerO}}^{AD}\right\|_1 + \mathbb{P}\left[\mathsf{Find}\right] \leq \frac{18q^{5/2}}{N-q}, \tag{26}$$

where in the second inequality we use the triangle inequality and the fact that trace distance is decreasing under partial trace. The final bound is a simplification (using the assumption that $\sqrt{q} \geq \frac{N-q}{N-5q}$ and $q \in O(\sqrt[4]{N})$) of bounds from Lemma 11 and Lemma 13 that are proven in the reminder of this section. $\qquad\square$

In the following lemma we prove a bound on the trace distance of the full permutation state and the good state.

Lemma 11. *The permutation and good states are close in trace distance:*

$$\frac{1}{2}\left\|\rho_{\mathsf{PerO}}^{A} - \rho_{\mathsf{Good}}^{A}\right\|_1 \leq \frac{1}{2}\frac{5q^2}{N-5q}. \tag{27}$$

Proof. To prove the statement we first observe that ρ_{Good} is diagonal in the effective queries. Moreover the diagonal terms of ρ_{PerO} are equal to the good state. The only difference in the two states are the off-diagonal parts of ρ_{PerO}. The main idea of our proof is that one can "reduce" the off-diagonal part to the diagonal part of a state generated by a slightly different adversary. This reduction introduces the factor resulting in the final bound on the distance of the states.

Diagonal terms Let us consider two worlds an adversary A can be in. In the ideal world she interacts with the full permutation oracle, resulting in the joint state from Eq. (17). In the real world she interacts with a compressed oracle for the uniform distribution. In the real world, after A's interaction we project the database register to collision-free databases, resulting in the good state from Eq. (20).

One can see that the outputs in the part of the oracle register corresponding to the effective queries are the same in both cases. However, by saving the inputs in register D^X, the density matrix reduced to the adversarial register will differ. The good density matrix ρ_{Good} will consist solely of terms diagonal in effective queries $\vec{x}$. Note however that these—diagonal in queries— terms in ρ_{PerO} are exactly the same. One can see it by taking the state from Eq. (17), fixing the

branch of superposition corresponding to a single $\vec{x}$, and calculating $\rho_{\mathsf{PerO}}(\vec{x},\vec{x})$. Tracing over the oracle register F results in a sum over $\vec{y}_{\vec{x}} \in \mathcal{I}(s)$. For $|\Psi^{\mathsf{Good}}\rangle$ from Eq. (20) with a fixed $\vec{x}$ the same operations will lead to the same matrix (the coefficients are the same because we take the same A and the sum over $\vec{y}$ is the same for diagonal terms). The sum over $\vec{x}$ is the whole diagonal. As the good state has no off-diagonal terms but gives the same diagonal terms we have

$$\sum_{\vec{x}} \rho_{\mathsf{PerO}}(\vec{x},\vec{x}) = \rho_{\mathsf{Good}}. \tag{28}$$

Reducing the off-diagonals Below we state and prove the claim that a well-formed off-diagonal matrix can be reduced to a well-formed matrix that is "closer" to the diagonal.

Claim 12. *Any* well-formed *matrix (i.e. following Definition 10)* $\rho_{\mathsf{PerO}}(\vec{x} \cup \vec{x}' \cup \{x\}, \vec{x} \cup \vec{x}'')$, *for any* $\vec{x}, \vec{x}', \vec{x}'', \{x\}$ *with empty intersections, can be reduced to a sum of well-formed matrices that do not include the effective query* x:

$$\rho_{\mathsf{PerO}}(\vec{x} \cup \vec{x}' \cup \{x\}, \vec{x} \cup \vec{x}'') = -\frac{s}{N-s}\Big(\widetilde{\rho}_{\mathsf{PerO}}^{\mathsf{UPD}}(\vec{x} \cup \vec{x}', \vec{x} \cup \vec{x}'')$$
$$+ \widetilde{\rho}_{\mathsf{PerO}}^{\mathsf{ADD}}(\widehat{\vec{x}} \cup \vec{x}', \widehat{\vec{x}} \cup \overline{\vec{x}}'') + \widetilde{\rho}_{\mathsf{PerO}}^{\mathsf{REM}}(\overline{\vec{x}} \cup \vec{x}', \overline{\vec{x}} \cup \widehat{\vec{x}}'')$$
$$+ \widetilde{\rho}_{\mathsf{PerO}}^{\mathsf{REM}}(\vec{x} \cup \overline{\vec{x}}', \vec{x} \cup \vec{x}'') + \widetilde{\rho}_{\mathsf{PerO}}^{\mathsf{REM}}(\vec{x} \cup \vec{x}', \vec{x} \cup \overline{\vec{x}}'')\Big), \tag{29}$$

where $s := |\vec{x} \cup \vec{x}' \cup \vec{x}''|$, $\widehat{\vec{x}}$ *is the set* $\vec{x}$ *with a single element added, and* $\overline{\vec{x}}$ *is* $\vec{x}$ *with a single element removed. The same element is added and removed when* $\widehat{\vec{x}}$ *and* $\overline{\vec{x}}$ *appears in both inputs to* $\widetilde{\rho}_{\mathsf{PerO}}$. *Overline and hat notation works the same for other sets. Detailed definition of the new matrices can be found in the proof.*

Proof. Let us inspect $\rho_{\mathsf{PerO}}(\vec{x} \cup \vec{x}' \cup \{x\}, \vec{x} \cup \vec{x}'')$ in more detail. Below we have $t := |\vec{x}|$, $t' := |\vec{x}'|$, and $t'' := |\vec{x}''|$ (and $s = t + t' + t''$):

$$\rho_{\mathsf{PerO}}(\vec{x} \cup \vec{x}' \cup \{x\}, \vec{x} \cup \vec{x}'') = \sum_{\vec{\eta},\vec{\eta}'} \sum_{w,w'} \alpha_{\vec{x} \cup \vec{x}' \cup \{x\},\vec{\eta},w} \overline{\alpha}_{\vec{x} \cup \vec{x}'',\vec{\eta}',w'}$$

$$|\psi(\vec{x} \cup \vec{x}' \cup \{x\}, \vec{\eta}, w)\rangle_A \langle \psi(\vec{x} \cup \vec{x}'', \vec{\eta}', w')|$$

$$\frac{1}{(N)_{t+t'+t''}} \sum_{\vec{z}_{\vec{x}} \in \mathcal{I}(t)} \omega_N^{\vec{\eta}_{\vec{x}} \cdot \vec{y}_{\vec{x}}} \overline{\omega}_N^{\vec{\eta}'_{\vec{x}} \cdot \vec{y}_{\vec{x}}} \sum_{\vec{y}_{\vec{x}'} \in \mathcal{I}(t'|\vec{y}_{\vec{x}})} \omega_N^{\vec{\eta}_{\vec{x}'} \cdot \vec{y}_{\vec{x}'}} \overline{\omega}_N^{\vec{\eta}'_{\vec{x}'} \cdot \vec{y}_{\vec{x}'}}$$

$$\sum_{\vec{y}_{\vec{x}''} \in \mathcal{I}(t''|\vec{y}_{\vec{x}} \cup \vec{y}_{\vec{x}'})} \omega_N^{\vec{\eta}_{\vec{x}''} \cdot \vec{y}_{\vec{x}''}} \overline{\omega}_N^{\vec{\eta}'_{\vec{x}''} \cdot \vec{y}_{\vec{x}''}} \frac{1}{N-t-t'-t''} \sum_{y_x \notin \vec{y}_{\vec{x}} \cup \vec{y}_{\vec{x}'} \cup \vec{y}_{\vec{x}''}} \omega_N^{\eta_x \cdot y_x}. \tag{30}$$

To reduce the effective query of x we evaluate the last sum in the above formula

$$\frac{1}{N-t-t'-t''} \sum_{y_x \notin \vec{y}_{\vec{x}} \cup \vec{y}_{\vec{x}'} \cup \vec{y}_{\vec{x}''}} \omega_N^{\eta_x \cdot y_x}$$

$$= \frac{1}{N-t-t'-t''}\Big(\underbrace{N\delta_{\eta_x,0}}_{=0} - \sum_{i=1}^{t} \omega_N^{\eta_x \cdot y_{x_i}} - \sum_{i'=1}^{t'} \omega_N^{\eta_x \cdot y_{x'_{i'}}} - \sum_{i''=1}^{t''} \omega_N^{\eta_x \cdot y_{x''_{i''}}} \Big). \tag{31}$$

We note that the result of evaluating the sum are changes to other effective queries held by the state.

Given the above observation, we define the matrices from the statement of the lemma. We can write the initial matrix as

$$\rho_{\mathsf{PerO}}(\vec{x} \cup \vec{x}' \cup \{x\}, \vec{x} \cup \vec{x}'')$$
$$= -\frac{s}{N-s}\mathrm{Tr}_{FA\widetilde{W}} |\widetilde{\Psi}^{\mathsf{PerO}}(x; \vec{x} \cup \vec{x}')\rangle\langle\widetilde{\Psi}^{\mathsf{PerO}}(\vec{x} \cup \vec{x}'')|, \tag{32}$$

where $|\widetilde{\Psi}^{\mathsf{PerO}}(x;\vec{x}\cup\vec{x}')\rangle$ and $|\widetilde{\Psi}^{\mathsf{PerO}}(\vec{x}\cup\vec{x}'')\rangle$ are well-formed vectors. The former vector is defined as the branch of the superposition from Eq. (17) corresponding to $\vec{x}\cup\vec{x}'$ with the effective queries (the part marked with $(*)$) changed to

$$\sum_{\tilde{x}\in\vec{x}\cup\vec{x}'\cup\vec{x}''}\frac{1}{\sqrt{s}}|\tilde{x}\rangle_{A\widetilde{W}}\sum_{\vec{y}_{\vec{x}\cup\vec{x}'\cup\vec{x}''}\in\mathcal{I}(s)}\frac{1}{\sqrt{(N)_s}}\omega_N^{\vec{\eta}\cdot\vec{y}_{\vec{x}\cup\vec{x}'\cup\vec{x}''}}\omega_N^{\eta_x\cdot y_{\tilde{x}}}. \tag{33}$$

The latter vector is the original $|\Psi^{\mathsf{PerO}}(\vec{x}\cup\vec{x}'')\rangle$ with the additional register $\sum_{\tilde{x}\in\vec{x}\cup\vec{x}'\cup\vec{x}''}\frac{1}{\sqrt{s}}|\tilde{x}\rangle_{A\widetilde{W}}$. Note that without that the trace would be ill-defined. The numerator s in Eq. (29) comes from the normalization of the new register.

To prove that $|\widetilde{\Psi}^{\mathsf{PerO}}(x;\vec{x}\cup\vec{x}')\rangle$ is well-formed we define a new adversary $\widetilde{A}$ based on the initial A. Our vector is the result of subtracting the states generated by the two adversaries. If a vector is a branch of the superposition generated by a valid adversary, then naturally it is well-formed.

The new adversary $\widetilde{A}$ first prepares a fresh register $\frac{1}{\sqrt{s}}\sum_{\tilde{x}\in\vec{x}\cup\vec{x}'\cup\vec{x}''}|\tilde{x}\rangle_{A\widetilde{W}}$. Next, after every query performed by A, the new $\widetilde{A}$ immediately applies $\mathsf{PerO}^{\dagger}$ controlled on A^X holding x, so uncomputes the query x. Next (immediately after the uncomputation of x), $\widetilde{A}$ applies $\mathsf{PerO}^{A\widetilde{W}Y}$ controlled on A^X being x. To sum up, whenever the old adversary would query x, the new adversary queries a superposition of $\tilde{x}$ instead. It is crucial to note that if there was no effective query to x then $\widetilde{A}$ generates the same state as A.

Say that $\widetilde{A}$ produces the vector $|\widetilde{\Psi}^{\mathsf{PerO}}(\vec{x}\cup\vec{x}')\rangle$ and A the vector $|\Psi^{\mathsf{PerO}}(\vec{x}\cup\vec{x}')\rangle$, then

$$|\widetilde{\Psi}^{\mathsf{PerO}}(x;\vec{x}\cup\vec{x}')\rangle := |\widetilde{\Psi}^{\mathsf{PerO}}(\vec{x}\cup\vec{x}')\rangle - |\Psi^{\mathsf{PerO}}(\vec{x}\cup\vec{x}')\rangle. \tag{34}$$

Now that we showed that $|\widetilde{\Psi}^{\mathsf{PerO}}(x;\vec{x}\cup\vec{x}')\rangle$ is a well-formed state we define the different matrices in the right hand side of Eq. (29). The new vector is a sum of five well-formed vectors, they are well-formed because they are just branches of the superposition constituting a well-formed vector:

$$\begin{aligned}|\widetilde{\Psi}^{\mathsf{PerO}}(x;\vec{x}\cup\vec{x}')\rangle = &|\widetilde{\Psi}^{\mathsf{PerO}}_{\mathsf{UPD}}(x;\vec{x}\cup\vec{x}')\rangle + |\widetilde{\Psi}^{\mathsf{PerO}}_{\mathsf{ADD}}(\widehat{\vec{x}}\cup\vec{x}')\rangle \\ &+ |\widetilde{\Psi}^{\mathsf{PerO}}_{\mathsf{REM}}(x;\overline{\vec{x}}\cup\vec{x}')\rangle + |\widetilde{\Psi}^{\mathsf{PerO}}_{\mathsf{REM}}(x;\vec{x}\cup\overline{\vec{x}}')\rangle + |\widetilde{\Psi}^{\mathsf{PerO}}_{\mathsf{REM}}(\vec{x}\cup\vec{x}')\rangle. \end{aligned} \tag{35}$$

The vectors in the right-hand side of Eq. (35) come from updating one of the other queries, adding a new query to a query from $\vec{x}''$, or removing a query from $\vec{x}$ or $\vec{x}'$. The branch corresponding to $\eta_x \neq -\eta_{\tilde{x}}$ (for any $\tilde{x}\in\vec{x}\cup\vec{x}'$) is denoted by $|\widetilde{\Psi}^{\mathsf{PerO}}_{\mathsf{UPD}}(x;\vec{x}\cup\vec{x}')\rangle$. The branch where $\tilde{x}\in\vec{x}''$ and $\eta_x \neq \eta_{\tilde{x}}$ is marked with the subscript ADD. State $|\widetilde{\Psi}^{\mathsf{PerO}}_{\mathsf{REM}}(\vec{x}\cup\vec{x}')\rangle$ denote the branch where $\tilde{x}\in\vec{x}''$ and $\eta_x = \eta_{\tilde{x}}$. The branches where for some $\tilde{x}\in\vec{x}\cup\vec{x}'$ we have $\eta_x = -\eta_{\tilde{x}}$ are denoted by the states in Eq. (35) with subscript REM. The first two REM states correspond to removing a query to $\vec{x}$ or $\vec{x}'$.

The matrices in the right hand side of Eq. (29) are defined by taking the partial trace over registers $FA^{\widetilde{W}}$ of the states from Eq. (35) multiplied by $\langle\widetilde{\Psi}^{\mathsf{PerO}}(\vec{x}\cup\vec{x}'')|$. $\qquad\square$

By using Claim 12 recursively on the new well-formed matrices from Eq. (29), we reduce every matrix $\rho_{\mathsf{PerO}}(\vec{x}\cup\vec{x}'\cup\{x\},\vec{x}\cup\vec{x}'')$ to a well-formed matrix diagonal in effective queries. We combine the diagonal matrices from Eq. (29) that are a result of the same number d of recursive steps. The diagonal matrices that we get are well-formed. This is because they are a sum of well-formed vectors coming from Claim 12, that are reduced off-diagonal terms of matrices with distinct effective queries. Intuitively speaking, we sum the branches we picked apart in the proof of Claim 12.

More concretely we have

$$\rho_{\mathsf{PerO}} = \sum_{\vec{x},\vec{x}',\vec{x}''} \rho_{\mathsf{PerO}}(\vec{x}\cup\vec{x}',\vec{x}\cup\vec{x}'') \tag{36}$$

$$= \sum_{d=1}^{2q-2} \sum_{t=1}^{t_{\max}(d)} \left(\prod_{s=t}^{t+d-1} \frac{s}{N-s}\right) \sum_{\vec{x}:|\vec{x}|=t} \widetilde{\rho}_{\mathsf{PerO}}^{d}(\vec{x},\vec{x}), \tag{37}$$

where in the first sum $\vec{x}\cap\vec{x}' = \emptyset$, $\vec{x}\cap\vec{x}'' = \emptyset$, and $\vec{x}'\cap\vec{x}'' = \emptyset$. The sum over d goes over the depth of recursion. The bound in the sum over t is just the maximal number of effective queries on the diagonal, this number depends on d but is bounded by q so we do not calculate it strictly. The product over s is the factor coming from applying Claim 12 recursively d times. The matrix

$$\widetilde{\rho}_{\mathsf{PerO}}^{d}(\vec{x},\vec{x}) = \mathrm{Tr}_{FA\widetilde{W}}\left(|\widetilde{\Psi}_{1}^{\mathsf{PerO}}(d;\vec{x})\rangle\langle\widetilde{\Psi}_{2}^{\mathsf{PerO}}(d;\vec{x})| + |\widetilde{\Psi}_{2}^{\mathsf{PerO}}(d;\vec{x})\rangle\langle\widetilde{\Psi}_{1}^{\mathsf{PerO}}(d;\vec{x})|\right) \tag{38}$$

for some $|\widetilde{\Psi}_{1}^{\mathsf{PerO}}(d;\vec{x})\rangle$ and $|\widetilde{\Psi}_{2}^{\mathsf{PerO}}(d;\vec{x})\rangle$ that are well-formed vectors holding a sum of vectors from Claim 12. The parameter d signifies the vectors contributing to Eq. (38), it is not crucial to keep track of it but we supply it for completeness.

Final bound Now we can define

$$|\widetilde{\Psi}^{\mathsf{Good}}(d;\vec{x})\rangle := |\widetilde{\Psi}_{1}^{\mathsf{Good}}(d;\vec{x})\rangle + |\widetilde{\Psi}_{2}^{\mathsf{Good}}(d;\vec{x})\rangle. \tag{39}$$

First of all, $|\widetilde{\Psi}_{i}^{\mathsf{Good}}(d;\vec{x})\rangle$ for $i \in \{1,2\}$ are defined as the same vectors as $|\widetilde{\Psi}_{1}^{\mathsf{PerO}}(d;\vec{x})\rangle$ but generated by A interacting with a compressed oracle for the uniform distribution and projected to injective databases at the end.

Finally we can calculate the norm of the off-diagonal terms of the permutation state. By $\pm$ denote adding and removing the state:

$$\left\| \sum_{d=1}^{2q-2} \sum_{t=1}^{t_{\max}(d)} \left(\prod_{s=t}^{t+d-1} \frac{s}{N-s}\right) \sum_{\vec{x}:|\vec{x}|=t} \left(\widetilde{\rho}_{\mathsf{PerO}}^{d}(\vec{x},\vec{x}) \pm |\widetilde{\Psi}^{\mathsf{Good}}(\vec{x})\rangle\langle\widetilde{\Psi}^{\mathsf{Good}}(\vec{x})|\right) \right\|$$

$$= \left\| \sum_{d=1}^{2q-2} \sum_{t=1}^{t_{\max}(d)} \left(\prod_{s=t}^{t+d-1} \frac{s}{N-s}\right) \sum_{\vec{x}:|\vec{x}|=t} \left(\mathrm{Tr}_{FA\widetilde{W}}|\widetilde{\Psi}_{1}^{\mathsf{PerO}}(d;\vec{x})\rangle\langle\widetilde{\Psi}_{1}^{\mathsf{PerO}}(d;\vec{x})| \right.\right.$$

$$\left.\left. + \mathrm{Tr}_{FA\widetilde{W}}|\widetilde{\Psi}_{2}^{\mathsf{PerO}}(d;\vec{x})\rangle\langle\widetilde{\Psi}_{2}^{\mathsf{PerO}}(d;\vec{x})| + |\widetilde{\Psi}^{\mathsf{Good}}(d;\vec{x})\rangle\langle\widetilde{\Psi}^{\mathsf{Good}}(d;\vec{x})|\right) \right\| \tag{40}$$

$$\leq \sum_{d=1}^{2q-2} \sum_{t=1}^{t_{\max}(d)} \left(\prod_{s=t}^{t+d-1} \frac{s}{N-s}\right) \left\| \sum_{\vec{x}:|\vec{x}|=t} \left(\mathrm{Tr}_{FA\widetilde{W}}|\widetilde{\Psi}_{1}^{\mathsf{PerO}}(d;\vec{x})\rangle\langle\widetilde{\Psi}_{1}^{\mathsf{PerO}}(d;\vec{x})| \right.\right.$$

$$\left.\left. + \mathrm{Tr}_{FA\widetilde{W}}|\widetilde{\Psi}_{2}^{\mathsf{PerO}}(d;\vec{x})\rangle\langle\widetilde{\Psi}_{2}^{\mathsf{PerO}}(d;\vec{x})| + |\widetilde{\Psi}^{\mathsf{Good}}(\vec{x})\rangle\langle\widetilde{\Psi}^{\mathsf{Good}}(\vec{x})|\right) \right\| \tag{41}$$

$$\leq 2 \sum_{d=1}^{2q-2} \sum_{t=1}^{t_{\max}(d)} \left(\prod_{s=t}^{t+d-1} \frac{s}{N-s}\right), \tag{42}$$

where the first equality comes from the fact that the cross terms of $|\widetilde{\Psi}^{\mathsf{Good}}(\vec{x})\rangle\langle\widetilde{\Psi}^{\mathsf{Good}}(\vec{x})|$ equal $\widetilde{\rho}_{\mathsf{PerO}}^{s}(\vec{x},\vec{x})$, and the first inequality follows from the triangle inequality. The bound on the norm comes from the observation that $\mathrm{Tr}_{FA\widetilde{W}}|\widetilde{\Psi}_{1}^{\mathsf{PerO}}(d;\vec{x})\rangle\langle\widetilde{\Psi}_{1}^{\mathsf{PerO}}(d;\vec{x})| + |\widetilde{\Psi}_{2}^{\mathsf{PerO}}(d;\vec{x})\rangle\langle\widetilde{\Psi}_{2}^{\mathsf{PerO}}(d;\vec{x})|$ can be interpreted as a partial trace over registers $FA\widetilde{W}B$ of the density matrix of the state $|\widetilde{\Psi}_{1}^{\mathsf{PerO}}(d;\vec{x})\rangle|1\rangle_{B} + \mathrm{Tr}_{FA\widetilde{W}}|\widetilde{\Psi}_{2}^{\mathsf{PerO}}(d;\vec{x})\rangle|2\rangle_{B}$. Both matrices are well-formed so their norm is bounded by 1.

To bound $\sum_{d=1}^{2q-2} \sum_{t=1}^{t_{\max}(d)} \left(\prod_{s=t}^{t+d-1} \frac{s}{N-s} \right)$ we first bound $\frac{s}{N-s} \leq \frac{t+d-1}{N-2q}$. Next we integrate the product by t:

$$\sum_{t=1}^{t_{\max}(d)} \left(\prod_{s=t}^{t+d-1} \frac{s}{N-s} \right) \leq \sum_{t=1}^{q} \left(\frac{t+d-1}{N-2q} \right)^d \leq \int_1^q dt \left(\frac{t+d-1}{N-2q} \right)^d \tag{43}$$

$$= \frac{\left((d+q-1)^{d+1} - d^{d+1} \right) \left(\frac{1}{N-2q} \right)^d}{d+1} \leq \frac{3q}{2} \left(\frac{3q}{N-2q} \right)^d. \tag{44}$$

Next we sum over d to get the final bound:

$$\sum_{d=1}^{2q-2} \frac{3q}{2} \left(\frac{3q}{N-2q} \right)^d = \frac{9q^2 - 9^q q^{2q} \left(\frac{1}{N-2q} \right)^{2q-2}}{2(N-5q)} \leq \frac{5q^2}{N-5q}. \tag{45}$$

This concludes the proof. □

The second important lemma provides a bound on the distance of the good state and the punctured oracle state conditioned on $\neg$Find and $\mathbb{P}\left[\text{Find} \right]$.

Lemma 13. *The good state and the compressed permutation states are close in trace distance:*

$$\frac{1}{2} \left\| \rho_{\text{Good}}^{AD} - \Phi_{\text{CPerO}}^{AD} \right\|_1 \leq \frac{1}{2} \cdot \frac{2q^{5/2}}{(N-q)}, \tag{46}$$

moreover the probability of Find *in this case is*

$$\mathbb{P}\left[\text{Find} \right] \leq \frac{2q(1+q)}{N} + \frac{3q^2(1+q)^2}{N\sqrt{N-q}} + \frac{3q^3(1+q)^2(1+q^2)}{2N(N-q)}, \tag{47}$$

that for $q \in O(\sqrt[4]{N})$ is just $\frac{7q(1+q)}{N}$.

Proof. We observe that the effective queries and the database register are essentially the same in ρ_{CPerO} and $\rho_{\text{CPhO}\setminus R_{\text{coll}}}$ and then use Claim 21 and Lemma 13 from [Cza+19]. The proof of Lemma 13 from [Cza+19] does not depend on the particular form of queries or the adversary, hence all results translate to our case. □

3.3 Punctured Compressed Permutation Oracle

In the proof of indifferentiability we puncture the compressed permutation oracle. So not only we puncture on regular collisions, but also on collisions in a part of the output (we call them *inner*-collisions). The inner-collisions relation is defined as:

$$R_{\text{inner}} := \left\{ D \in \bigcup_{s=1}^{q} \left(\{0,1\}^{r+c} \right)^{2s} : \exists i, j \neq i, y_i, y_j \in D^Y : \hat{y}_i = \hat{y}_j \vee \hat{y}_i = 0^c \right\}, \tag{48}$$

where $D = ((x_1, y_1), \ldots, (x_s, y_s))$, and by $\hat{y}$ we denote the last c bits of $y \in \{0,1\}^{r+c}$. To use the O2H lemma we just need to bound probability of Find for the combined relation (regular and inner collisions). In the proof of Lemma 13 in [Cza+19], the authors state the following bound on $\mathbb{P}\left[\text{Find} \right]$, that can be applied to the adversary interacting with $\text{CPerO}_{\{0,1\}^{r+c}}$ punctured on R_{inner}:

$$\mathbb{P}\left[\text{Find} : A[\text{CPerO}_{\{0,1\}^{r+c}} \setminus R_{\text{inner}}] \right]$$

$$\leq \sum_{i=1}^{q} \left(5 \sum_{j=1}^{i-1} \max_{s \leq j-1} \left\{ \frac{b(s+1)^{3/2}}{\sqrt{N(N-b(q))}} \right\} \right.$$

$$\left. + \max_{s \leq i-1} \left\{ \sqrt{\frac{b^3(s)}{N^2(N-b(q))}} + 2\sqrt{\frac{b(s+1)}{N}} \right\} \right)^2, \tag{49}$$

15

where $N = 2^{r+c}$.

In our case, of regular collisions combined with inner-collisions we have

$$b(s) = s - 1 + 2^r s. \tag{50}$$

So the final bound is

$$\mathbb{P}\left[\text{Find} : \text{A}[\text{CPerO}_{\{0,1\}^{r+c}} \setminus R_{\text{inner}}]\right] \leq 6\frac{q^2}{2^c} + \frac{\sqrt{30}q^4\sqrt{2^r}}{2^c\sqrt{2^c - 2q}} + \frac{50q^6 2^r}{2^c(2^c - 2q)}, \tag{51}$$

which for $q \in O(\sqrt[4]{2^{c-r}})$ simplifies to $86\frac{q^2}{2^c}$.

4 Indifferentiability of Sponges with Random Permutations

In this section we use the new technique to prove quantum indifferentiability of the sponge construction instantiated with a random permutation. We present proofs both of classical and quantum indifferentiability. Our proof of quantum security mimics the structure of the classical one. This approach is the same as in [Cza+19] and we hope will similarly lead to easier understanding of the second proof.

In our proofs we follow the (quantum) game-playing framework. Classical games make have use of lazy-sampled functions and "Bad" events. The key statement in classical game-playing proofs is the fundamental game-playing lemma (Lemma 3). It states that two games are identical until the bad event occurs. Moreover the fundamental lemma provides a bound depending on the probability of the bad event. The quantum game-playing framework gives us similar tools that are based on the compressed oracle technique and the O2H lemma.

In the quantum case, the role of bad events is played by punctured oracles, defined in Def. 2. Thanks to measurement that occur after every query we can condition on the content of quantum databases. That is a feature that allows to follow the classical proof so closely. The quantum version of the fundamental game-playing lemma is the O2H lemma. The second distinguishability bound that we use is shown in Lemma 6.

The general idea behind both indifferentiability proofs is that if there are no inner-collisions, then the outputs of the sponge construction are uniformly distributed. We condition on random permutations that do not have inner-collisions (in the q queries made by the adversary) using the (quantum) game-playing framework. All the details on sponge-specific notation used in this section are located in section 2.4.

To save space we present multiple algorithms in one. To do that we follow a convention where only the boxed algorithms perform the boxed operations. In our case there are actually more than two algorithms, so that the color of the box also matters.

In both proofs we denote the random oracle, defined in Eq. (3), by R. Moreover, we omit the second input and fix the output length to a single r-bit block.

4.1 Classical Indifferentiability of Sponges with Random Permutations

First we will present a slightly modified proof of indifferentiability from [Ber+08]. We modify the proof, so that our classical and quantum proofs are more similar. We also introduce more steps in the proof, as we want to keep a slow paste of statements so that each transition between games is clear, in that we follow the style of proof of [Cor+05]. Before we state the theorem, let us briefly go over lazy sampling of a random permutation. Specifically for permutations $f : \{0,1\}^{r+c} \to \{0,1\}^{r+c}$, sampling done in two stages, first sampling the inner part of the output and then the outer part.

By $\mathcal{D}$ we denote the set of outputs of previous queries. In the language of the sponge graph $\mathcal{D}^{-1}$ is the set of nodes with outgoing edges and $\mathcal{D}$ is the set of nodes with incoming edges. By

$\widehat{\mathcal{D}}$ we denote the set of supernodes with all nodes having an incoming edge. By $\overline{\mathcal{D}}(\hat{t})$ we denote the set of nodes in the supernode $\hat{t}$ with an incoming edge.

We need to define a procedure to lazy-sample outputs of a random permutation. The obvious solution of sampling uniformly from $\{0,1\}^{r+c} \setminus \mathcal{D}$ is not good enough as we want to sample the inner part before the outer part and retain the step-by-step structure of our proof. Instead we are going to first sample uniformly from $\{0,1\}^{r+c} \setminus \mathcal{D}$ but then discard the outer state. The value of the inner state $\hat{t}$ is then effectively sampled from $\{0,1\}^c$ with weights $\frac{|\{0,1\}^r \setminus \overline{\mathcal{D}}(\hat{t})|}{|\{0,1\}^{r+c} \setminus \mathcal{D}|}$. We call this distribution $\mathfrak{C}$. At this point we will be introducing bad events concerning the inner part of the sampled state. To sample the outer state we just sample uniformly from $\{0,1\}^r \setminus \overline{\mathcal{D}}(\hat{t})$. We denote this distribution by $\mathfrak{A}(\hat{t})$.

In the case of the inverse of the random permutation we use a similar distribution but in the above definitions we take $\mathcal{D}^{-1}$—i.e. the set of nodes with outgoing edges—in both $\mathfrak{C}$ and $\mathfrak{A}$. We denote those distributions by $\mathfrak{C}^{-1}$ and $\mathfrak{A}^{-1}(\hat{t})$ respectively.

For the proof of indifferentiability we also need an upper bound on the probability of finding a collision in the inner part of outputs of a uniformly random function $f : \{0,1\}^{r+c} \to \{0,1\}^{r+c}$. Considering how SPONGE is defined we want a bound on finding collisions and zero-preimages. We define the bound as a function of the number of queries q to f:

$$b_{\mathrm{coll}}(q) := \frac{q(q+1)}{2^{c+1}}. \tag{52}$$

The bound can be found with standard techniques, a detailed derivation is presented in section 5.1 of [Cza+19].

The bound on $\mathbb{P}[\mathrm{Bad}]$ in our proof for random permutations is the same as in the case of random functions [Cza+19]. We achieve a slightly weaker bound than in [Ber+08] because in our gradual argument we first sample from the set of inner states excluding only "full" supernodes $\widehat{\mathcal{D}}$.

With these additional comments in mind we can proceed with the classical indifferentiability result.

Theorem 14 (SPONGE with permutations, classical indifferentiability)**.** SPONGE$_f$[PAD, r, c] *calling a random permutation f is (q, ε)-indifferentiable from a random oracle (defined in Eq. (3)) for* classical *adversaries for any $q < |2^c|$ and $\varepsilon = 7\frac{q(q+1)}{2^{c+1}}$.*

Proof. In Algorithm 2 we present the indifferentiability simulators.

Game 1 We start with the real world where the distinguisher A has access to a random permutation $f : \{0,1\}^{r+c} \to \{0,1\}^{r+c}$ and the construction SPONGE$_f$ using this function. The formal definition of the first game is

$$\mathbf{Game\ 1} := \left(1 \leftarrow \mathrm{A}[\mathrm{SPONGE}_f, (f, f^{-1})]\right). \tag{53}$$

Game 2 In the second game we introduce the simulator S$_2$—defined in Algorithm 2—that lazy-samples the random permutation f. The definition of the second game is

$$\mathbf{Game\ 2} := \left(1 \leftarrow \mathrm{A}[\mathrm{SPONGE}_{\mathsf{S}_2}, (\mathsf{S}_2, \mathsf{S}_2^{-1})]\right), \tag{54}$$

where by S^{-1} we denote the backward interface of S. Because the simulator S$_2$ perfectly models a random permutation and we use the same function for the private interface we have

$$|\mathbb{P}[\mathbf{Game\ 2}] - \mathbb{P}[\mathbf{Game\ 1}]| = 0. \tag{55}$$

Algorithm 2: Classical S_2, $\boxed{\mathsf{S}_3}$, $\boxed{\mathsf{S}_4}$, permutations

---

   State   : Current sponge graph G

   Interface: f, forward queries
   Input   : $s \in \{0,1\}^{r+c}$
   Output : $f(s)$

1   **if** s has no outgoing edge **then**   `// Fresh query`
2   **if** $\hat{s} \in \mathcal{R} \wedge \mathcal{R} \cup \mathcal{U} \neq \{0,1\}^c$ **then**   `// ŝ-rooted, no saturation`
3 $\hat{t} \leftarrow \mathfrak{C}$, $\boxed{\textbf{if } \hat{t} \in \mathcal{R} \cup \mathcal{U}\text{, set } \mathrm{Bad}_1 = 1}$
4   Construct a path to s: $p := \mathsf{SpPath}(s, G)$
5   **if** $\exists x : p = \textsc{pad}(x)$ **then**
6 $\bar{t} \leftarrow \mathfrak{A}(\hat{t})$
7 $\boxed{\bar{t} := \mathsf{R}(x)}$
8   **else**
9 $\bar{t} \leftarrow \mathfrak{A}(\hat{t})$
10 $t := (\bar{t}, \hat{t})$
11   **else**
12 $t \xleftarrow{\$} (\{0,1\}^{r+c}) \setminus \mathcal{D}$
13   Add an edge (s,t) to $\mathcal{E}$.
14   Set t to the vertex at the end of the edge starting at s
15   Output t

   Interface: f^{-1}, backward queries
   Input   : $s \in \{0,1\}^{r+c}$
   Output : $f^{-1}(s)$

16   Construct the sponge graph G
17   **if** s has no incoming edge **then**   `// Fresh query`
18 $\hat{t} \leftarrow \mathfrak{C}^{-1}$, $\boxed{\textbf{if } \hat{t} \in \mathcal{R}\text{, set } \mathrm{Bad}_2 = 1}$
19 $\bar{t} \leftarrow \mathfrak{A}^{-1}(\hat{t})$
20 $t := (\bar{t}, \hat{t})$
21   Add an edge (t,s) to $\mathcal{E}$.
22   Set t to the vertex at the beginning of the edge ending at s
23   Output t

Game 3 In the next step we modify S_2 to S_3. The game is then

$$\textbf{Game 3} := \left(1 \leftarrow A[\text{SPONGE}_{S_3}, (S_3, S_3^{-1})] \right). \tag{56}$$

We made a single change in S_3 compared to S_2, we introduce the "bad" event $\text{Bad} := \text{Bad}_1 \vee \text{Bad}_2$ that marks the difference between algorithms. We use this event as the bad event in Lemma 3. With such a change of the simulators we can use Lemma 3 to bound the difference of probabilities:

$$|\mathbb{P}[\textbf{Game 3}] - \mathbb{P}[\textbf{Game 2}]| \leq \mathbb{P}[\text{Bad} = 1]. \tag{57}$$

It is now quite easy to bound $\mathbb{P}[\text{Bad} = 1]$ as this is the probability of finding a collision or a preimage of the root in the set $\{0,1\}^c$ having made q random samples. Then we have that

$$\mathbb{P}[\text{Bad} = 1] \leq b_{\text{coll}}(q), \tag{58}$$

where the inequality comes from the fact that not all queries are made to rooted nodes.

Game 4 In this step we introduce the random oracle R but only to generate the outer part of the output of f. The game is defined as

$$\textbf{Game 4} := \left(1 \leftarrow A[\text{SPONGE}_{S_4}, (S_4^R, S_4^{-1})] \right). \tag{59}$$

Now we need to observe that if $\text{Bad} = 0$ the outputs are identically distributed.

Claim 15. *Given that* $\text{Bad} = 0$, **Game 4** *and* **Game 3** *are identical:*

$$|\mathbb{P}[\textbf{Game 4} \mid \text{Bad} = 0] - \mathbb{P}[\textbf{Game 3} \mid \text{Bad} = 0]| = 0. \tag{60}$$

Proof. If $\text{Bad} = 0$, the inner-part is distributed in the same way in both games, so the only difference in distributions can come from the outer part. Given our discussion in section 2.4, though, if there are no inner-collisions, then the outer part is distributed uniformly at random. This reasoning is also presented in Lemma 1 and Lemma 2 of [Ber+07]. $\qquad \square$

The two games are identical-until-bad, this implies that the probability of setting Bad to one in both games is the same $\mathbb{P}[\text{Bad} = 1 : \textbf{Game 3}] = \mathbb{P}[\text{Bad} = 1 : \textbf{Game 4}]$. Together with the above claim we can derive the advantage:

$$|\mathbb{P}[\textbf{Game 4}] - \mathbb{P}[\textbf{Game 3}]| \stackrel{\text{Claim 15}}{=} \bigg| \mathbb{P}[\textbf{Game 4} \mid \text{Bad} = 0]$$

$$\cdot \underbrace{(\mathbb{P}[\text{Bad} = 1 : \textbf{Game 3}] - \mathbb{P}[\text{Bad} = 1 : \textbf{Game 4}])}_{=0}$$

$$+ \underbrace{\mathbb{P}[\textbf{Game 3} \mid \text{Bad} = 1]}_{\leq 1}\mathbb{P}[\text{Bad} = 1] + \underbrace{\mathbb{P}[\textbf{Game 4} \mid \text{Bad} = 1]}_{\leq 1}\mathbb{P}[\text{Bad} = 1] \bigg| \tag{61}$$

$$\leq 2\mathbb{P}[\text{Bad} = 1] \leq 2b_{\text{coll}}(q). \tag{62}$$

Game 5 In this stage of the proof we change the private interface to contain the actual random oracle. The simulator is the same and the game is

$$\textbf{Game 5} := \left(1 \leftarrow A[R, (S_4^R, S_4^{-1})] \right). \tag{63}$$

Conditioned on $\text{Bad} = 0$, the outputs of the simulator in Games 4 and 5 are the same and consistent with R. To calculate the adversary's advantage in distinguishing between the two

games we can follow the proof of Lemma 6. We change $\mathsf{R} \setminus R_1$ to **Game 5**, $\mathsf{G} \setminus R_2$ to **Game 4**, and event Find to Bad = 1. As the derivation of Lemma 6 uses no quantum mechanical arguments and the assumption holds—the games are identical conditioned on Bad = 0—the bound holds:

$$|\mathbb{P}[\textbf{Game 5}] - \mathbb{P}[\textbf{Game 4}]| \leq 4\mathbb{P}[\text{Bad}_1 = 1] \leq 4b_{\text{coll}}(q). \tag{64}$$

Collecting and adding all the differences yields the claimed $\varepsilon = 7b_{\text{coll}}(q)$. $\qquad\square$

4.2 Quantum Indifferentiability of Sponges with Random Permutations

In the proof of quantum security we will be using $\mathsf{CPerO}_{\{0,1\}^{r+c}}$ and puncture it on parts of the database, concretely on the inner parts of outputs. To maintain the structure of the proof similar to the classical case we still write the simulators to include separate procedures for sampling the inner and the outer part of the output. In fact though we sample both parts at once and puncture only on a part of the database. This change is only a semantic one, and the core of reasoning is the same. We use the fact that we can use the O2H lemma between two punctured oracles.

We denote the part of $\mathsf{CPerO}_{\{0,1\}^{r+c}}$ responsible for sampling the inner part by $\widehat{\mathsf{CPerO}}_{\{0,1\}^{r+c}}$ and the part of $\mathsf{CPerO}_{\{0,1\}^{r+c}}$ responsible for sampling the outer part by $\overline{\mathsf{CPerO}}_{\{0,1\}^{r+c}}$. The correct understanding of this notation is that in fact there is just the single $\mathsf{CPerO}_{\{0,1\}^{r+c}}$ applied and is punctured on all relations at once. In case of any statements dividing the calls to the inner and outer parts, we just move all the operations between the calls in the pseudocode to before of the first oracle.

Another important detail concerns the type of oracle we consider. For the sake of presentation we chose the basis used for the adversary's queries to be the standard basis. This is not how we defined CPerO, but the change does not influence the applicability of the compressed permutation oracle. We just need to keep in mind that the oracle used in this section is actually $\mathsf{HT}_{r+c}^{AY}\mathsf{CPerO}_{\{0,1\}^{r+c}}\mathsf{HT}_{r+c}^{AY}$, where $\mathsf{CPerO}_{\{0,1\}^{r+c}}$ is the oracle defined in Definition 8 and $\mathsf{HT}_n|y\rangle := \sum_{\eta \in \{0,1\}^n} (-1)^{y \cdot \eta}|\eta\rangle$ is the Hadamard transform.

In **Game 4** of the classical proof, we replace the sampled outer part with the output of the random oracle. Note that before the classical simulator terminates she saves the full sate in the sponge graph. There are two issues with doing the same in the quantum simulator. First of all, we cannot save information that we have not provided the adversary with. This would cause to increase the distinguishing advantage and not necessarily by a negligible amount. Second of all, we sample a random permutation, so without saving all the outputs we could cause collisions we could not track.

To solve both these problems we do not save the oracle outputs in D. However, at the beginning of each run of the simulator we reprepare the sponge graph and the missing values in D. To prepare the sponge graph we query R on all necessary inputs to $\hat{f}$, i.e. on the inputs that are consistent with a path from the root to a rooted node. This is done gradually by iterating over the length of the paths. We begin with the length-0 paths, i.e. with all inputs in the database $\widehat{D}$ where the inner part is the all zero string. If the outer part of such an input (which is not changed by the application of SpPath) is equal to a padding of an input, that input is queried to determine the outer part of the output of f, creating an edge in the sponge graph. We can continue with length-1 paths. For each entry of the database $\widehat{D}$, check whether the input register is equal to a node in the current partial sponge graph. If so, the entry corresponds to a rooted node. Using the entry and the edge connecting its input to the root, a possible padded input to Sponge is created using SpPath. If it is a valid padding, R is queried to determine the outer part of the output of f, etc.

We denote by U_G the unitary that acting on $|0\rangle$ constructs G including edges consistent with queries held by the quantum compressed database from register D. Similarly we define $\mathsf{U}_{\mathcal{R} \cup \mathcal{U}}$

to temporarily create a description of the set of supernodes that are rooted or have an outgoing edge.

It is important to note that the "IF" statements in the quantum simulator are in fact quantum controlled operations. To correctly perform the controlled operation we postpone the measurements occurring in punctured oracles to after uncomputing of G and $\mathcal{R} \cup \mathcal{U}$.

We define the function corresponding to the probability of Find, it follows from Eq. (51):

$$b_{\text{coll}}^Q(q) := 86 \frac{q^2}{2^c}. \tag{65}$$

Below we state and prove the main result of this paper, quantum indifferentiability of sponges with random permutations. Inspecting in more detail Eq. 51, we see that for $q \in O\left(2^{c/3}\right)$ the bound approaches a constant. This number of queries also trivializes the final bound in the theorem below. Considering the complexity of the inner-collision-finding algorithm from [Cza+18], which is also of the order $O\left(2^{c/3}\right)$, we can state that our bound is tight, at least in the presence of an external random oracle.

Theorem 16 (SPONGE with permutations, quantum indifferentiability). SPONGE$_f$[PAD, r, c] *calling a random permutation f is (q, ε)-indifferentiable from a random oracle (defined in Eq. (3)) for* quantum *adversaries for any $q \in O(\sqrt[4]{2^{c-r}})$ (which is trivially $q < |2^c|$) and $\varepsilon = \frac{13q^{5/2}}{2^{r+c}-q} + 688\frac{q^2}{2^c} + \sqrt{86(q+1)\frac{q^2}{2^c}}$.*

Proof. In Algorithm 3 we describe the simulators we use in this proof.

Game 1 We start with the real world where the distinguisher A has quantum access to a random permutation $f : \{0,1\}^{r+c} \to \{0,1\}^{r+c}$ and the construction SPONGE$_f$ using this function. The formal definition of the first game is

$$\textbf{Game 1} := \left(1 \leftarrow \text{A}[\text{SPONGE}_f, (f, f^{-1})]\right). \tag{66}$$

Game 2 In the second game we introduce the simulator S$_2$—defined in Algorithm 3—that lazy-samples the random permutation f. The definition of the second game is

$$\textbf{Game 2} := \left(1 \leftarrow \text{A}[\text{SPONGE}_{S_2}, (S_2, S_2^{-1})]\right), \tag{67}$$

where by S^{-1} we denote the backward interface of S. The simulator S$_2$ models the quantum random permutation with errors given in Theorem 9:

$$|\mathbb{P}[\textbf{Game 2}] - \mathbb{P}[\textbf{Game 1}]| \leq \frac{13q^{5/2}}{2^{r+c} - q}. \tag{68}$$

Game 3 In the next step we modify S$_2$ to S$_3$. The game is then

$$\textbf{Game 3} := \left(1 \leftarrow \text{A}[\text{SPONGE}_{S_3}, (S_3, S_3^{-1})]\right). \tag{69}$$

We made a single change in S$_3$ compared to S$_2$, we introduced puncturing on the relation $\mathcal{R} \cup \mathcal{U}$ in the forward direction and $\mathcal{R}$ in the backward direction. With such a change of the simulators we can use Lemma 4 to bound the difference of probabilities:

$$|\mathbb{P}[\textbf{Game 3}] - \mathbb{P}[\textbf{Game 2}]| \leq \sqrt{(q+1)\mathbb{P}[\text{Find} : \text{A}[\text{SPONGE}_{S_3}, S_3]]}, \tag{70}$$

where Find denotes the success of the measurement in the punctured oracle. We can bound $\mathbb{P}[\text{Find} : \text{A}[\text{SPONGE}_{S_3}, S_3]]$ using Eq. (51):

$$\mathbb{P}[\text{Find} : \text{A}[\text{SPONGE}_{S_3}, S_3]] \leq b_{\text{coll}}^Q(q). \tag{71}$$

Algorithm 3: Quantum $\boxed{\mathsf{S}_2}$, $\boxed{\mathsf{S}_3}$, $\boxed{\mathsf{S}_4}$, permutations

---

State : Quantum compressed database register D

Interface: f, forward queries
Input : $|s, v\rangle_{XY} \in \mathcal{H}_{\{0,1\}^{r+c}}^{\otimes 2}$
Output : $|s, v \oplus f(s)\rangle_{XY}$

1   Apply $\mathsf{U}_{\mathcal{R}\cup\mathcal{U}}\mathsf{U}_G$ to register D and two fresh registers
2   **if** $s \in \mathcal{R} \wedge \mathcal{R} \cup \mathcal{U} \neq \{0,1\}^c$ **then**   // s-rooted, no saturation
3   Apply $\boxed{\widehat{\mathsf{CPerO}}_{\{0,1\}^{r+c}}^{XYDG}}$, $\boxed{(\widehat{\mathsf{CPerO}}_{\{0,1\}^{r+c}} \setminus (\mathcal{R} \cup \mathcal{U}))^{XYDG}}$, result: $\hat{t}$
4   Construct a path to s: $p := \mathsf{SpPath}(s, G)$
5   **if** $\exists x : p = \textsc{pad}(x)$ **then**
6    $\boxed{\text{Apply } \mathsf{CStO}_{\{0,1\}^{r+c}}^{XYDG}}$, result: $\bar{t}$
7   Write x in a fresh register X_R, $\boxed{\text{apply } \mathsf{R}^{XX_RYDG}}$, uncompute x from X_R, result:
   $\bar{t}$ // Random oracle
8   **else**
9   Apply $\overline{\mathsf{CPerO}}_{\{0,1\}^{r+c}}^{XYDG}$, result: $\bar{t}$
10   $t := (\bar{t}, \hat{t})$, the value of the output stored in D
11 **else**
12   Apply $\mathsf{CPerO}_{\{0,1\}^{r+c}}^{XYDG}$, result t
13 Uncompute G and $\mathcal{R} \cup \mathcal{U}$
14 Output $|s, v + t\rangle_{XY}$

Interface: f^{-1}, backward queries
Input : $|s, v\rangle \in \mathcal{H}_{\{0,1\}^{r+c}}^{\otimes 2}$
Output : $|s, v + f^{-1}(s)\tau\rangle$

15 Apply $\mathsf{U}_{\mathcal{R}\cup\mathcal{U}} \circ \mathsf{U}_G$ to registers D and two fresh registers
16 Apply $\boxed{\widehat{\mathsf{CPerO}}_{\{0,1\}^{r+c}}^{-1\,XYDG}}$, $\boxed{(\widehat{\mathsf{CPerO}}_{\{0,1\}^{r+c}} \setminus \mathcal{R})^{-1\,XYDG}}$, result $\hat{t}$
17 Apply $\overline{\mathsf{CPerO}}_{\{0,1\}^{r+c}}^{-1\,XYDG}$, result: $\bar{t}$
18 $t := (\bar{t}, \hat{t})$, the value of the output saved to D
19 Uncompute G and $\mathcal{R} \cup \mathcal{U}$
20 Output $|s, v + t\rangle_{XY}$

Game 4 In this step we introduce the random oracle R but only to generate the outer part of the output of f. The game is defined as

$$\textbf{Game 4} := \left(1 \leftarrow A[\text{SPONGE}_{S_4}, (S_4^R, S_4^{-1})]\right). \tag{72}$$

Thanks to the classical argument, Claim 15 we have that S_4 and S_3 are identical until bad, as in Definition 5. Then we can use Lemma 6 to bound the advantage of the adversary

$$|\mathbb{P}[\textbf{Game 4}] - \mathbb{P}[\textbf{Game 3}]| \leq 4\mathbb{P}[\text{Find} : A[\text{SPONGE}_{S_3}, (S_3, S_3^{-1})]] \leq 4b_{\text{coll}}^Q(q) \tag{73}$$

Game 5 In this stage of the proof we change the private interface to contain the actual random oracle. The simulator is the same and the game is

$$\textbf{Game 5} := \left(1 \leftarrow A[R, (S_4^R, S_4^{-1})]\right). \tag{74}$$

The advantage is

$$|\mathbb{P}[\textbf{Game 5}] - \mathbb{P}[\textbf{Game 4}]| \leq 4\mathbb{P}[\text{Find} : A[\text{SPONGE}_{S_4}, (S_4^R, S_4^{-1})]] \leq 4b_{\text{coll}}^Q(q), \tag{75}$$

conditioned on ¬Find, outputs of the private interface are the same, then the games are identical until bad and we can use Lemma 6 to bound the advantage of the adversary.

Saturation certainly does not occur for $q < 2^c$ as the database in every branch of the superposition increases by at most one in every query. Collecting the differences between games yields the claimed $\varepsilon = \frac{13q^{5/2}}{2^{r+c}-q} + 8b_{\text{coll}}^Q(q) + \sqrt{(q+1)b_{\text{coll}}^Q(q)}$. □

Acknowledgments

References

[AHU19]   Andris Ambainis, Mike Hamburg, and Dominique Unruh. "Quantum Security Proofs Using Semi-classical Oracles". In: *Advances in Cryptology - CRYPTO 2019*. 2019, pp. 269–295. DOI: 10.1007/978-3-030-26951-7_10.

[BR93]   Mihir Bellare and Phillip Rogaway. "Random oracles are practical: A paradigm for designing efficient protocols". In: *Proceedings of the 1st ACM conference on Computer and communications security*. ACM. 1993, pp. 62–73. DOI: 10.1145/168588.168596.

[BR06]   Mihir Bellare and Phillip Rogaway. "The Security of Triple Encryption and a Framework for Code-Based Game-Playing Proofs". In: *Advances in Cryptology - EUROCRYPT 2006*. Springer Berlin Heidelberg, 2006, pp. 409–426. DOI: 10.1007/11761679_25.

[BBD09]   D.J. Bernstein, J. Buchmann, and E. Dahmen. *Post-Quantum Cryptography*. Springer Berlin Heidelberg, 2009.

[Ber+07] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. "Sponge functions". In: *ECRYPT hash workshop*. Vol. 2007. 9. https://keccak.team/files/SpongeFunctions.pdf. Citeseer. 2007.

[Ber+08]   Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. "On the In-differentiability of the Sponge Construction". In: *Advances in Cryptology – EURO-CRYPT 2008*. Springer Berlin Heidelberg, 2008, pp. 181–197. DOI: 10.1007/978-3-540-78967-3_11.

[Chu+20]   Kai-Min Chung, Serge Fehr, Yu-Hsuan Huang, and Tai-Ning Liao. "On the Compressed-Oracle Technique, and Post-Quantum Security of Proofs of Sequential Work". In: (2020). arXiv: 2010.11658.

[Cor+05]   Jean-Sébastien Coron, Yevgeniy Dodis, Cécile Malinaud, and Prashant Puniya. "Merkle-Damgård Revisited: How to Construct a Hash Function". In: *Advances in Cryptology – CRYPTO 2005*. Springer Berlin Heidelberg, 2005, pp. 430–448. DOI: 10.1007/11535218_26.

[Cza+18]   Jan Czajkowski, Leon Groot Bruinderink, Andreas Hülsing, Christian Schaffner, and Dominique Unruh. "Post-quantum Security of the Sponge Construction". In: *Post-Quantum Cryptography*. Springer International Publishing, 2018, pp. 185–204. DOI: 10.1007/978-3-319-79063-3_9.

[CHS19]   Jan Czajkowski, Andreas Hülsing, and Christian Schaffner. "Quantum Indistinguishability of Random Sponges". In: *Advances in Cryptology - CRYPTO 2019*. 2019, pp. 296–325. DOI: 10.1007/978-3-030-26951-7_11.

[Cza+19]   Jan Czajkowski, Christian Majenz, Christian Schaffner, and Sebastian Zur. "Quantum Lazy Sampling and Game-Playing Proofs for Quantum Indifferentiability". In: (2019). arXiv: 1904.11477.

[HI19]   Akinori Hosoyamada and Tetsu Iwata. "4-Round Luby-Rackoff Construction is a qPRP". In: *Advances in Cryptology - ASIACRYPT 2019*. 2019, pp. 145–174. DOI: 10.1007/978-3-030-34578-5_6.

[LZ19]   Qipeng Liu and Mark Zhandry. "On Finding Quantum Multi-collisions". In: *Advances in Cryptology - EUROCRYPT 2019*. 2019, pp. 189–218. DOI: 10.1007/978-3-030-17659-4_7. URL: https://doi.org/10.1007/978-3-030-17659-4_7.

[MRH04]   Ueli Maurer, Renato Renner, and Clemens Holenstein. "Indifferentiability, Impossibility Results on Reductions, and Applications to the Random Oracle Methodology". In: *Theory of Cryptography*. Springer Berlin Heidelberg, 2004, pp. 21–39. DOI: 10.1007/978-3-540-24638-1_2.

[NC10]   Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. 10th anniversary. Cambridge: Cambridge University Press, 2010. ISBN: 978-1107002173.

[NIS14]   NIST. *SHA-3 Standard: Permutation-Based Hash and Extendable-Output Functions*. Draft FIPS 202. 2014. URL: http://csrc.nist.gov/publications/drafts/fips-202/fips_202_draft.pdf.

[NIS15]   NIST. *Secure Hash Standard (SHS)*. Draft FIPS 180-4. 2015. DOI: 10.6028/NIST.FIPS.180-4.

[NIS17]   NIST. *Post-Quantum Cryptography Standardization process*. 2017. URL: https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization.

[Unr14]   Dominique Unruh. "Revocable Quantum Timed-Release Encryption". In: *Advances in Cryptology – EUROCRYPT 2014*. Springer Berlin Heidelberg, 2014, pp. 129–146. DOI: 10.1007/978-3-642-55220-5_8.

[Unr21]   Dominique Unruh. "Compressed Permutation Oracles (And the Collision-Resistance of Sponge/SHA3)". Cryptology ePrint Archive, Report 2021/062. 2021. URL: https://eprint.iacr.org/2021/062.

[Wol11]   Ronald de Wolf. "Quantum computing: Lecture notes". In: *University of Amsterdam* (2011). URL: https://arxiv.org/abs/1907.09415.

[Zha19]   Mark Zhandry. "How to Record Quantum Queries, and Applications to Quantum Indifferentiability". In: *Advances in Cryptology – CRYPTO 2019*. Springer International Publishing, 2019, pp. 239–268. DOI: 10.1007/978-3-030-26951-7_9.

A Single-query Distinguisher

One interesting, in our opinion, observation concerning the permutation oracle is that a single query conveys a different amount of information depending on the number of previous queries. Note that the first query not only informs us that $f(x) = y$ but also that the output on all other $x' \neq x$ in not y. The more queries we make, the more "negative" information we have. The importance of the above observation is confirmed with the indistinguishability bound for a single-query distinguisher.

Theorem 17. *Given quantum access to a function sampled either according to the uniform distribution $\mathfrak{U}$ over functions from $\{f : [N] \to [N]\}$ or the uniform distribution over permutations $\mathfrak{P}$, we have that for all quantum distinguishers making a single quantum query to the oracle, the distinguishability advantage is at most $\frac{1}{N-1}$:*

$$\left| \mathop{\mathbb{P}}_{f \leftarrow \mathfrak{U}} \left[1 \leftarrow A^f \right] - \mathop{\mathbb{P}}_{f \leftarrow \mathfrak{P}} \left[1 \leftarrow A^f \right] \right| \leq \frac{1}{N-1}. \tag{76}$$

Moreover there is a distinguisher B that has advantage $\frac{1}{2(N-1)}$.

Proof. Lets us inspect the indistinguishability bound for the uniform distribution over functions $\mathfrak{U}$ and the uniform distribution over permutations $\mathfrak{P}$:

$$\left| \mathop{\mathbb{P}}_{f \leftarrow \mathfrak{U}} \left[1 \leftarrow A^f \right] - \mathop{\mathbb{P}}_{f \leftarrow \mathfrak{P}} \left[1 \leftarrow A^f \right] \right| = \left| \text{Tr} \left(Q_1 \rho_{\mathfrak{U}}^A \right) - \text{Tr} \left(Q_1 \rho_{\mathfrak{P}}^A \right) \right| \tag{77}$$

$$\leq \frac{1}{2} \left\| \rho_{\mathfrak{U}}^A - \rho_{\mathfrak{P}}^A \right\|_1 \tag{78}$$

where Q_1 is the measurement corresponding to A outputting 1. The inequality above comes from the fact that trace distance gives the optimal measurement for distinguishing two states [NC10]. Register A holds all of the adversary's state, this includes her work register, query register, and any other auxiliary registers she decides to use.

We model the oracle access with the full phase oracle. The adversary's initial state is:

$$|\psi_0\rangle_A := \sum_x \sum_{\eta_x, w} \alpha_{x, \eta_x, w} |x, \eta_x, w\rangle_{AXYW}, \tag{79}$$

the states $\rho_{\mathfrak{U}}$ and $\rho_{\mathfrak{U}}$ are defined as the partial trace of the joint adversary-oracle state after a single query to $\text{PhO}_{\mathfrak{U}}$ or $\text{PhO}_{\mathfrak{P}}$, respectively. The full definitions of the reduced states are:

$$\rho_{\mathfrak{U}} = \sum_{x, x'} \sum_{\eta_x, \eta'_{x'}, w, w'} \alpha \bar{\alpha}' \sum_{y, y_{x'}} \frac{1}{N^2} \omega_N^{\eta_x y_x} \bar{\omega}_N^{\eta'_{x'} y_{x'}} |x, \eta_x, w\rangle \langle x', \eta'_{x'}, w'| \tag{80}$$

$$\rho_{\mathfrak{P}} = \sum_{x, x'} \sum_{\eta_x, \eta'_{x'}, w, w'} \alpha \bar{\alpha}' \sum_{y, y_{x'} \neq y_x} \frac{1}{N(N-1)} \omega_N^{\eta_x y_x} \bar{\omega}_N^{\eta'_{x'} y_{x'}} |x, \eta_x, w\rangle \langle x', \eta'_{x'}, w'|, \tag{81}$$

where we write α and α' for $\alpha_{x, \eta_x, w}$ and $\alpha_{x', \eta'_{x'}, w'}$ respectively.

First of all we see that terms with $x' = x$ are equal, in this case in both states the sum over y_x and $y_{x'}$ simplify to just a single sum over $y_x \in [N]$. The off-diagonal (so for $x' \neq x$) terms are a bot more complicated to calculate. Let us notice that $\sum_{y, y_{x'} \neq y_x} = \sum_{y, y_{x'}} - \sum_{y, y_{x'} = y_x}$, this

splits the difference $\rho_{\mathfrak{U}} - \rho_{\mathfrak{P}}$ into two parts:

$$\rho_{\mathfrak{U}} - \rho_{\mathfrak{P}} = \left(\frac{1}{N^2} - \frac{1}{N(N-1)} \right)$$

$$\sum_{x,x'\neq x} \sum_{\eta_x,\eta'_{x'},w,w'} \alpha\bar{\alpha}' \sum_{y,y_{x'}} \omega_N^{\eta_x y_x} \bar{\omega}_N^{\eta'_{x'} y_{x'}} |x,\eta_x,w\rangle\langle x',\eta'_{x'},w'|$$

$$+ \sum_{x,x'\neq x} \sum_{\eta_x,\eta'_{x'},w,w'} \alpha\bar{\alpha}' \underbrace{\sum_{y,y_{x'}=y_x} \frac{1}{N(N-1)} \omega_N^{\eta_x y_x} \bar{\omega}_N^{\eta'_{x'} y_{x'}}}_{=\frac{1}{N-1}\delta_{\eta_x,\eta'_{x'}}} |x,\eta_x,w\rangle\langle x',\eta'_{x'},w'| \tag{82}$$

$$= -\frac{1}{N-1}\rho_{\mathfrak{U}} + \frac{1}{N-1}\mathsf{D}_{A^Y}(|\psi_0\rangle\langle\psi_0|), \tag{83}$$

where $\mathsf{D}_A(\rho) := \sum_a |a\rangle_A\langle a|\rho|a\rangle_A\langle a|$ denotes the map that outputs the diagonal entries in register A.

Finally we have

$$\frac{1}{2}\left\| \rho_{\mathfrak{U}}^A - \rho_{\mathfrak{P}}^A \right\|_1 = \frac{1}{2}\left\| -\frac{1}{N-1}\rho_{\mathfrak{U}} + \frac{1}{N-1}\mathsf{D}_{A^Y}(|\psi_0\rangle\langle\psi_0|) \right\|_1 \tag{84}$$

$$\leq \frac{1}{2}\frac{1}{N-1}\|\rho_{\mathfrak{U}}\|_1 + \frac{1}{2}\frac{1}{N-1}\||\psi_0\rangle\langle\psi_0|\|_1 = \frac{1}{N-1}, \tag{85}$$

where we use the triangle inequality and the fact that any CPTP (Completely Positive Trace-Preserving) map, like D, can only decrease the norm of a state.

Now we present the distinguisher B. We define B such that she performs the optimal measurement, and hence has advantage $\|\rho_{\mathfrak{U}} - \rho_{\mathfrak{P}}\|_1$, as stated in Eq. (78). Now let us consider the initial state of B, it does not have a work register and is the following state:

$$|\psi_0^{\mathsf{B}}\rangle_A := \sum_{x\in\{x_1,x_2\}} \frac{1}{\sqrt{2}}|x,\eta\rangle_{A^{XY}}, \tag{86}$$

where $\eta \neq 0$ and x_1 and x_2 are any distinct inputs. We know that the diagonal (in x) terms are equal, no matter if B interacts with $\mathfrak{U}$ or $\mathfrak{P}$, and the only part of the state that influences the trace distance is the off-diagonal. We know that in the state after interaction with a random function, Eq. (80), in the branch of superposition where $x \neq x'$ the sum $\sum_{y,y_{x'}} \frac{1}{N^2}\omega_N^{\eta_x y_x}\bar{\omega}_N^{\eta'_{x'} y_{x'}} = \delta_{\eta_x,0}\delta_{\eta'_{x'},0}$. We observe that, as we have set $\eta \neq 0$, the off-diagonal part of $\rho_{\mathfrak{U}}$ (the case with $x \neq x'$) is 0.

The above discussion leads to the following equality:

$$\left| \underset{f\leftarrow\mathfrak{U}}{\mathbb{P}}\left[1 \leftarrow \mathsf{B}^f\right] - \underset{f\leftarrow\mathfrak{P}}{\mathbb{P}}\left[1 \leftarrow \mathsf{B}^f\right] \right| = \frac{1}{2}\left\| \rho_{\mathfrak{U}}^A - \rho_{\mathfrak{P}}^A \right\|_1 \tag{87}$$

$$= \frac{1}{2}\frac{1}{N-1}\left\| \sum_{x\in\{x_1,x_2\}} \sum_{x'\in\{x_1,x_2\},x'\neq x} \frac{1}{2}|x,\eta\rangle\langle x',\eta| \right\|_1 = \frac{1}{2(N-1)}, \tag{88}$$

where the last equality comes from the definition of the trace norm: $\|\mathsf{A}\|_1 = \sum_\lambda |\lambda|$, where λ are the eigenvalues of the matrix A. $\qquad\square$

Symbol Index