

How to Record Quantum Queries, and Applications to Quantum Indifferentiability

Mark Zhandry

Princeton University & NTT Research

Abstract. The quantum random oracle model (QROM) has become the standard model in which to prove the post-quantum security of random-oracle-based constructions. Unfortunately, none of the known proof techniques allow the reduction to record information about the adversary’s queries, a crucial feature of many classical ROM proofs, including all proofs of indifferentiability for hash function domain extension.

In this work, we give a new QROM proof technique that overcomes this “recording barrier”. We do so by giving a new “compressed oracle” which allows for efficient on-the-fly simulation of random oracles, roughly analogous to the usual classical simulation. We then use this new technique to give the first proof of quantum indifferentiability for the Merkle-Damgård domain extender for hash functions. We also give a proof of security for the Fujisaki-Okamoto transformation; previous proofs required modifying the scheme to include an additional hash term. Given the threat posed by quantum computers and the push toward quantum-resistant cryptosystems, our work represents an important tool for efficient post-quantum cryptosystems.

1 Introduction

The random oracle model [BR93] has proven to be a powerful tool for heuristically proving the security of schemes that otherwise lacked a security proof. In the random oracle model (ROM), a hash function H is modeled as a truly random function that can only be evaluated by querying an oracle for H . A scheme is secure in the ROM if it can be proven secure in this setting. Of course, random oracles cannot be efficiently realized; in practice, the random oracle is replaced with a concrete efficient hash function. The hope is that the ROM proof will indicate security in the real world, provided there are no structural weaknesses in the concrete hash function.

Meanwhile, given the looming threat of quantum computers [IBM17], there has been considerable interest in analyzing schemes for so called “post-quantum” security [NIS17, Son14, ATTU16, CBH⁺17, YAJ⁺17, CDG⁺17, CDG⁺15]. Many of the proposed schemes are random oracle schemes; Boneh et al. [BDF⁺11] argue that the right way of modeling the random oracle in the quantum setting is to use the quantum random oracle model, or QROM. Such a model allows a quantum attacker to query the random oracle on a quantum superposition of inputs. The idea is that a real-world quantum attacker, who knows the code for the concrete

hash function, can evaluate the hash function in superposition in order to perform tasks such as Grover search [Gro96] or collision finding [BHT98]. In order to accurately capture such real-world attacks, it is crucial to model the random oracle to allow for such superposition queries. The quantum random oracle model has been used in a variety of subsequent works to prove the post-quantum security of cryptosystems [BDF⁺11, Zha12b, Zha15, TU16, Eat17].

The Recording Barrier. Unfortunately, proving security in the quantum random oracle model can be extremely difficult. Indeed, in the classical random oracle model, one can copy down the adversary’s queries as a means to learning what points the adversary is interested in. Many classical security proofs crucially use this information in order to construct a new adversary which solves some hard underlying problem, reaching a contradiction. In the quantum setting, such copying is impossible by no-cloning. One can try to record some information about the query, but this amounts to a measurement of the adversary’s query state which can be detected by the adversary. A mischievous adversary may refuse to continue if it detects such a measurement, rendering the adversary useless for solving the underlying problem. Because of the difficulty in reading an adversary’s query, it also becomes hard to adaptively program the random oracle, another common classical proof technique.

This difficulty has led authors to develop new quantum-sound proof techniques to replace classical techniques, such as Zhandry’s small-range distributions [Zha12a] or Targhi and Unruh’s extraction technique [TU16]. These proof techniques choose the oracle from a careful distribution that allows for proofs to go through. However, every such proof technique always chooses a classical oracle at the beginning of the experiment, and leave the oracle essentially unchanged through the entire execution. The inability to change the oracle seems inherent, since if the proof gives the adversary different oracles during different queries, this is potentially easily detectable (even by classical adversaries)¹

Constraining the oracles to be fixed functions seems to limit what can be proved using such non-recording techniques. For example, Dagdelen, Fischlin, and Gagliardoni [DFG13] show that such natural proof techniques are likely incapable of proving the security of Fiat-Shamir². This leads to a natural question: *Is it possible to record information about an adversary’s quantum query without the adversary detecting*

Enter Indifferentiability. The random oracle model (quantum or otherwise) assumes the adversary treats the hash function as a monolithic object. Unfortunately, hash functions in practice are usually built from smaller building blocks, called compression functions. If one is not careful, hash functions built in this way

¹ The one exception we are aware of is Unruh’s adaptive programming [Unr15]. This proof does change the oracle adaptively, but only inputs for which adversary’s queries have only negligible “weight”. Thus, the change is not detectable. The following discussion also applies to Unruh’s technique.

² We note that if the underlying building blocks are strengthened, Fiat-Shamir was proven secure by Unruh [Unr16]

are vulnerable to attacks such as length-extension attacks. Coron et al. [CDMP05] show that a hash function built from a compression function can be as good as a monolithic oracle in many settings if it satisfies a notion of *indifferentiability*, due to Maurer, Renner, and Holenstein [MRH04]. Roughly, in indifferentiability, an adversary A has oracle access to both h and H , and the adversary is trying to distinguish two possible worlds. In the “real world”, h is a random function, and H is built from h according to the hash function construction. In the “ideal world”, H is a random function, and h is simulated so as to be consistent with H . A hash function is indifferentiable from a random oracle if no efficient adversary can distinguish the two worlds.

Coron et al.’s proof of indifferentiability for Merkle-Damgård requires the simulator to remember the queries that the adversary has made. This is actually inherent for any domain extender, by a simple counting argument discussed below. In the quantum setting, such recording presents a serious issue, as recording a query is equivalent (from the adversary’s point of view) to measuring the query. As any measurement will disturb the quantum system, such measurement may be detectable to the adversary. Note that in the case where A is interacting with a truly random h , there is no measurement happening. Therefore, if such a measurement can be detected, the adversary can distinguish the two cases, breaking indifferentiability.

Example. To illustrate what might go wrong, we will use the simple example from Coron et al. [CDMP05]. Here, we will actually assume access to two independent compression functions $h_0, h_1 : \{0, 1\}^{2n} \rightarrow \{0, 1\}^n$. We will define $H : \{0, 1\}^{3n} \rightarrow \{0, 1\}^n$ as $H(x, y) = h_1(h_0(x), y)$, where $x \in \{0, 1\}^{2n}, y \in \{0, 1\}^n$.

To argue that H is indifferentiable from a random oracle, Coron et al. use the following simulator S , which has access to H , and tries to implement the oracles h_0, h_1 . S works as follows:

- S keeps databases D_0, D_1 , which will contain tuples (x, y) . D_b containing (x, y) means that S has set $h_b(x) = y$.
- h_0 is implemented on the fly: every query on x looks up $(x, y) \in D_0$, and returns y if it is found; if no such pair is found, a random y is chosen and returned, and (x, y) is added to D_0 .
- By default, h_1 is answered randomly on the fly as in h_0 . However, it needs to make sure that $h_1(h_0(x), y)$ always evaluates to $H(x, y)$, else it is trivial to distinguish the two worlds. Therefore, on a query (z, y) , h_1 will check if there is a pair (x, z) in D_0 for some x . If so, it will reasonably guess that the adversary is trying to evaluate $H(x, y)$, and respond by making a query to $H(x, y)$. Otherwise it will resort to the default simulation.

Note that by defining the simulator in this way, if the adversary ever tries to evaluate H on (x, z) by first making a query x to h_0 to get y , and then making a query (y, z) to h_1 , the simulator will correctly set the output of h_1 to $H(x, z)$, so that the adversary will get a result that is consistent with H . However, note that it is crucial that S wrote down the queries made to h_0 , or else it will not know which point to query H when simulating h_1 .

Now consider a quantum adversary. A quantum query to, say, h_0 will be the following operation:

$$\sum_{x \in \{0,1\}^{2n}, u \in \{0,1\}^n} \alpha_{x,u} |x, u\rangle \mapsto \sum_{x \in \{0,1\}^{2n}, u \in \{0,1\}^n} \alpha_{x,u} |x, u \oplus h_0(x)\rangle$$

Now, imagine our simulator trying to answer queries to h_0 in superposition. For simplicity, suppose this is the first query to h_0 , so D_0 is empty. The natural approach is to just have S store its database D_0 in superposition, performing a map that may look like $|x, u\rangle \mapsto |x, u \oplus y\rangle \otimes |x, y\rangle$, where y is chosen randomly, and everything to the right of the \otimes is the simulators state.

But now consider the following query by an adversary. It sets up the uniform superposition $\sum_{x,u} |x, u\rangle$ and queries. In the case where h_0 is a classical function, then this state becomes

$$\sum_{x,u} |x, u \oplus h_0(x)\rangle = \sum_{x,u} |x, u\rangle$$

Namely, the state is unaffected by making the query. In contrast, the simulated query would result in

$$\sum_{x,u} |x, u \oplus y\rangle \otimes |x, y\rangle$$

Here, the adversary's state is now entangled with the simulator's. It is straightforward to detect this entanglement by applying the Quantum Fourier Transform (QFT) to the adversary's x registers, and then measuring the result. In the case where the adversary is interacting with a random h_0 , the QFT will result in a 0. In the simulated case, the QFT will result in a random string. These two cases are therefore easily distinguishable.

To remedy this issue, prior works in the quantum regime have abandoned on-the-fly simulation, instead opting for stateless simulation. Here, the simulator commits to a function to implement the oracle in the very beginning, and then sticks with this implementation throughout the entire experiment. Moreover, the simulator never records any information about the adversary's query, lest the adversary detect the entanglement with the simulator. This will certainly fix the issue above, and by carefully choosing the right implementations prior works have shown how to translate many classical results into the quantum setting.

However, for indistinguishability, choosing a single fixed function for h_0 introduces new problems. Now when the adversary makes a query to h_1 , the simulator needs to decide if the query represents an attempt at evaluating H , and if so, it must program the output of h_1 accordingly. However, without knowing what inputs the adversary has queried to h_0 , it seems impossible for the simulator to determine which point the adversary is interested in. For example, if the adversary queries h_1 on (y, z) , there will be roughly 2^n possible x that gave rise to this y (since h_0 is compressing). Therefore, the simulator must choose from one of 2^n inputs of the form (x, z) on which to query H .

To make matters even more complicated, an adversary can submit the uniform superposition $\sum_x |x, 0\rangle$, resulting in the state $\sum_x |x, h_0(x)\rangle$, which causes it to

“learn” $y = h_0(x)$. At this point, the simulator should be ready to respond to an h_1 query on (y, z) by using x , meaning the simulator *must* be entangled with x . Then, at some later time, the adversary can query again on the state $\sum_x |x, h_0(x)\rangle$, resulting in the original state $\sum_x |x, 0\rangle$ again. The adversary can test that it received the correct state using the quantum Fourier transform. Therefore, after this later query, the simulator must be un-entangled with x . Even more complex strategies are possible, where the adversary can compute and un-compute h_0 in stages, so as to try to hide what it is doing from any potential simulator.

These issues are much more general than just the simple domain extender above. Indeed, even classically domain extension with a *stateless* simulator is *impossible*, by the following simple argument. Suppose there is a hash function $H : \{0, 1\}^M \rightarrow \{0, 1\}^N$ built from a compression function $h : \{0, 1\}^m \rightarrow \{0, 1\}^n$ as $H = C^h$ for an oracle circuit C . Let $L = M + \log_2 N, \ell = m + \log_2 n$. Then L, ℓ represent the logarithm of the size of the truth tables for H, h . Since we are domain extending, we are interested in the case where $L \gg \ell$. Suppose even $L \geq \ell + 0.001$.

Suppose toward contradiction that h can be simulated statelessly, which we will represent as Sim^H (since the function can make H queries). Then h has a truth table of size 2^ℓ . In the real world, H agrees with C^h on all inputs; therefore in order for indifferentiability to hold, in the simulated world a uniformly random H must agree with $C^h = C^{\text{Sim}^H}$ on an overwhelming fraction of inputs. But this is clearly impossible, as it would allow us to compress the random truth table of H : simply output the truth table for Sim^H , along with the ϵ fraction of input/output pairs where H and C^{Sim^H} disagree. The total length of this compressed truth table is $2^\ell + (\epsilon 2^M)(MN) = 2^\ell + \epsilon N 2^L$. As ϵ is negligible (and therefore much smaller than $1/N$) the compressed truth table will be smaller than 2^L , the size of the truth table for H . But since H is a random function its truth table cannot be compressed, reaching a contradiction.

Therefore, any simulator for indifferentiability, regardless of the scheme, *must* inherently store information about the adversary. But the existing QROM techniques are utterly incapable of such recording. We therefore ask: *Is indifferentiable domain extension even possible?*

1.1 This Work

In this work, perhaps surprisingly, we answer the question above in the affirmative. Namely, we give a new *compressed oracle technique*, which allows for recording the adversary’s queries in a way that the adversary can never detect. The intuition is surprisingly simple: an adversary interacting with a random oracle can be thought of as being entangled with a uniform superposition of oracles. As entanglement is symmetric, if the adversary ever has any information about the oracle, the *oracle must also have information about the adversary*. Therefore a simulator can always record *some* information about the adversary, if done carefully.

We then use the technique to prove the indifferentiability of the Merkle-Damgård construction. We believe our new technique will be of independent

interest; for example our technique can be used to prove the security of the Fujisaki-Okamoto transformation [FO99], and also gives very short proofs of several quantum query lower bounds.

The Compressed Oracle Technique. In order to prove indifferentiability, we devise a new way of analyzing quantum query algorithms

Consider an adversary interacting with an oracle $h : \{0, 1\}^m \rightarrow \{0, 1\}^n$. It is well established that the usual quantum oracle mapping $|x, y\rangle \mapsto |x, y \oplus h(x)\rangle$ is equivalent to the “phase” oracle, which maps $|x, u\rangle \mapsto (-1)^{u \cdot h(x)}|x, u\rangle$ (we discuss this equivalence in Section 3). For simplicity, in this introduction we will focus on the phase oracle, which is without loss of generality.

Next, we note that the oracle h being chosen at random is equivalent (from the adversary’s point of view) to h being in uniform superposition $\sum_h |h\rangle$. Indeed, the superposition can be reduced to a random h by measuring, and measuring the h registers (which is outside of A ’s view) is undetectable to A . To put another way, the superposition over h is a *purification* of the adversary’s mixed state.

Therefore, we will imagine the h oracle as containing $\sum_h |h\rangle$. When A makes a query on $\sum_{x,u} \alpha_{x,u} |x, u\rangle$, the joint system of the adversary and oracle are

$$\sum_{x,u} \alpha_{x,u} |x, u\rangle \otimes \sum_h |h\rangle$$

The query introduces a phase term $(-1)^{u \cdot h(x)}$, so the joint system becomes

$$\sum_{x,u} \alpha_{x,u} |x, u\rangle \otimes \sum_h |h\rangle (-1)^{u \cdot h(x)}$$

We normally think of the phase as being returned to the adversary, but the phase really affects the entire system, so it is equivalent to think of the phase as being added to the oracle’s state.

Now, we will think of h as a vector of length $2^m \times n$ by simply writing down h ’s truth table. We will think of each x, u pair as a point function $P_{x,u}$ which outputs u on x and 0 elsewhere. Using our encoding of functions as vectors, we can write $u \cdot h(x)$ as $P_{x,u} \cdot h$. We can therefore write the post-query state as

$$\sum_{x,u} \alpha_{x,u} |x, u\rangle \otimes \sum_h |h\rangle (-1)^{h \cdot P_{x,u}}$$

In general, the state after making q queries can be written as

$$\sum_{x_1, \dots, x_q, u_1, \dots, u_q} \alpha_{x_1, \dots, x_q, u_1, \dots, u_q} |\psi_{x_1, \dots, x_q, u_1, \dots, u_q}\rangle \otimes \sum_h |h\rangle (-1)^{h \cdot (P_{x_1, u_1} + \dots + P_{x_q, u_q})}$$

Next, notice that by applying the Quantum Fourier transform to h , the h registers will now contain $(P_{x_1, u_1} + \dots + P_{x_q, u_q}) \bmod 2$. Working in the Fourier domain, we see that each query simply adds $P_{x,u}$ (modulo 2) to the result. In the Fourier domain, the initial state is 0.

Therefore, from A ’s point of view, it is indistinguishable whether the oracle for h is a random oracle, or it is implemented as follows:

- The oracle keeps as state a vector $D \in \{0,1\}^{n \times 2^m}$, initially set to 0.
- On any oracle query, the oracle performs the map $|x, u\rangle \otimes |D\rangle \mapsto |x, u\rangle \otimes |D \oplus P(x, u)\rangle$

Thus, with this remarkably simple change in perspective, the oracle can actually be implemented by recording and updating phase information about the queries being made.

We can now take this a couple steps further. Notice that after q queries, D is non-zero on at most q inputs (since it is the sum of q point functions). Therefore, we can store the database in an extremely compact form, namely the list of (x, y) pairs where $y = D(x)$ and $y \neq 0$. Notice that this allows us to efficiently simulate a random oracle, without an a priori bound on the number of queries. Previously, simulating an unbounded number of queries efficiently required computational assumptions, and simulation was only computationally secure. In contrast, simulating random oracles exactly required $2q$ -wise independent functions [Zha12b] and hence required knowing q up front. We therefore believe this simulation will have independent applications for the efficient simulation of quantum oracles. We will call this the compressed Fourier oracle.

We can then take our compressed Fourier oracle, and convert it back into a primal-domain oracle. Namely, for each (x, y) pair, we perform the QFT on the y registers. The result is a superposition of databases of (x, w) pairs, where w roughly represents $h(x)$. For any pair not in the database, $h(x)$ is implicitly a uniform superposition of inputs independent of the adversary’s view. We call this the compressed standard oracle. It intuitively represents what the adversary knows about the function h : if (x, y) is in the database then the adversary “knows” $h(x) = y$, and otherwise, the adversary “knows” nothing about $h(x)$. In Section 3, we show how to directly obtain the compressed standard oracle.

Applying Compressed Oracles to Indifferentiability. The compressed standard oracle offers a simple way to keep track of the queries the adversary has made. In particular, it tracks exactly the kind of information needed in the classical indifferentiability proof above, namely whether or not a particular value has been queried by the adversary, and what the value of the oracle at that point is. We use this to give a quantum indifferentiability proof for Merkle-Damgård construction using prefix-free encodings [CDMP05].

To illustrate our ideas, consider our simple example above with h_0, h_1 and H . Our simulator will simulate h_0 as in the compressed standard oracle, keeping a (superposition over) lists D_0 of (x, y) pairs. Next, our simulator must handle h_1 queries. When given a phase query $|y, z\rangle$, the simulator does the following. If first looks for a pair (x, y') in D_0 with $y' = y$. If one is found, it reasonably guesses that the adversary is interested in computing $H(x, z)$, and so it makes a query on (x, z) to H . Otherwise, it is reasonable to guess that the adversary is not trying to compute H on any input, since the adversary does not “know” any inputs to h_0 that would result in a query to h_1 on (y, z) .

While the above appears to work, we need to make sure the simulator does not disturb the compressed oracle. Unfortunately, some disturbance is necessary.

Indeed, determining the value of $h_0(x)$ is a measurement in the primal domain. On the other hand, the update procedure for the compressed oracles needs to decide whether or not x belongs in the database, and this corresponds to a measurement in the *Fourier* domain (since in the Fourier domain, $h_0(x)$ must be non-zero). These two measurements do not commute, so by the uncertainty principle it is impossible to perform both measurements perfectly.

Nonetheless, we show that the errors are small. Intuitively, we observe that the simulator does not actually need to know the entire value of $h_0(x)$, just whether or not it is equal to y . We call such information a “test”. Similarly, the compressed oracle implementation just needs to know whether or not $h_0(x)$ is equal to 0, but in the Fourier domain.

Now, these primal and Fourier tests still do not commute. Fortunately, they “almost” commute, which we formalize in the full version [Zha18]. The intuition is that, if a primal test of the form “is $h_0(x) = y$ ” has a non-negligible chance of succeeding, $h_0(x)$ must be very “far” from the uniform superposition. This is because a uniform superposition puts an exponentially small weight on every outcome. Recall that the uniform superposition maps to $h_0(x) = 0$ in the Fourier domain. Thus by being “far” from uniform, the Fourier domain test has a negligibly-small chance of succeeding. Therefore, one of the two tests is always “almost” determined, meaning the measurement negligibly affects the state. This means that, no matter what initial state is, the two tests “almost” commute.

Thus, the simulator can perform these tests without perturbing the state significantly. This shows that h_0 queries are correctly simulated; we also need to show that h_1 queries are correctly simulated and consistent with H . The intuition above suggests that h_1 should be consistent with H , and indeed in Section 5 we show this using a careful sequence of hybrids. Then in the full version [Zha18], we use the same ideas to prove the indifferentiability of Merkle-Damgård.

The Power of Forgetting. Surprisingly, our simulator ends up strongly resembling the classical simulator. It is natural to ask, therefore, how the simulator gets around the difficulties outlined above.

First, notice that if we translate the query $\sum_{x,u} |x,u\rangle$ in our example to a phase query, it becomes $\sum_x |x,0\rangle$. This query has no effect on the oracle’s state. This means the oracle remains un-entangled with the adversary, as desired.

Second, a query $\sum_x |x,0\rangle$ becomes $\sum_{x,u} |x,u\rangle$ for a phase query. Consider applying the query to the compressed Fourier oracle. The joint quantum system of the adversary and simulator becomes

$$\sum_{x,u \neq 0} |x,u\rangle |\{(x,u)\}\rangle + \sum_x |x,0\rangle |\{\}\rangle$$

A similar expression holds for the compressed standard oracle. Note that the simulator can clearly tell (whp) that the adversary has queried on x . Later, when the adversary queries on the same state a second time, (x,u) will get mapped to $(x,0)$, and will hence be removed from the database. Thus, after this later query, the database contains no information about x . Hence, the adversary is un-entangled with x , and so its tests will output the correct value.

Ultimately then, the key difference between our simulator and the natural quantum analog of the classical simulator is that our simulator must be ready to *forget* some of the oracle points it simulated previously. By implementing h_0 as a compressed oracle, it will forget *exactly* when it needs to so that the adversary can never detect that it is interacting with a simulated oracle.

Other results We expect our compressed oracle technique will have applications beyond indistinguishability. Here, we list two additional sets of results we are able to obtain using our technique:

Post-quantum security of Fujisaki-Okamoto. The Fujisaki-Okamoto transformation [FO99] transforms a weak public key encryption scheme into a public key encryption scheme that is secure against *chosen ciphertext attacks*, in the random oracle model. Unfortunately, the classical proof does not work in quantum random oracle model, owing to similar issues with indistinguishability proofs. Namely, in one step of the proof, the reduction looks at the queries made by the adversary in order to decrypt chosen ciphertext queries. This is crucial to allow the reduction to simulate the view of the adversary without requiring the secret decryption key. But in the quantum setting, it is no longer straightforward to read the adversary's queries without disrupting its state.

Targhi and Unruh [TU16] previously modified the transformation by including an additional random oracle hash in the ciphertext. In the proof, the hash function is set to be injective, and the reduction can invert the hash in order to decrypt.

In the full version [Zha18], we show how to adapt our compressed oracle technique to prove the security of the original transform without the extra hash. In addition, we show security against even *quantum* chosen ciphertext queries, thus proving security in the stronger model of Boneh and Zhandry [BZ13]. We note that recently, Jiang et al. [JZC⁺18] proved the security of the FO transformation when used as a key encapsulation mechanism. Their proof is tight, whereas ours is somewhat loose. On the other hand, we note that their proof does not apply if FO is used directly as an encryption scheme, and does not apply in the case of quantum chosen ciphertext queries.

Simple Quantum Query Complexity Lower Bounds. We also show that our compressed oracles can be used to give very simple and optimal quantum query complexity lower bounds for problems for *random functions*, such as pre-image search, collision finding, and more generally k -SUM.

Our proof strategy is roughly as follows. First, since intuitively the adversary has no knowledge of values of h outside of D , except with very small probability any successful algorithm will output points in D . Therefore it suffices to bound the number of queries required to get D to contain a pre-image/collision/ k -sum.

For pre-image search, we re-prove the optimal lower bound of $\Omega(2^{n/2})$ queries of [BBBV97], but for random functions; note that pre-image search for random functions and worst-case functions is equivalent using simple reductions. The proof appears superficially similar to [BBBV97]: we show that each query can

increase the “amplitude” on “good” databases by a small $O(2^{-n/2})$ amount. After q queries, this amplitude becomes $O(q/2^{n/2})$, which we then square to get the probability of a “good” database. The proof is only slightly over a page once the compressed oracle formalism has been given.

We then re-prove the optimal collision lower bound of $\Omega(2^{n/3})$ queries for random functions, matching the worst case bound [AS04] and the more recent average case bound [Zha15]. Remarkably, our proof involves only a few lines of modification to the pre-image lower bound. We show that the amplitude on “good” databases increases by $O(\sqrt{q} \times 2^{n/2})$ for each query, where the extra \sqrt{q} intuitively comes from the fact that the database has size at most q , giving q opportunities for a collision every time a new entry is added to the database³.

In contrast to our very simple extension, the prior collision bounds involved very different techniques and were much more complicated. Also note that prior works could not prove directly that finding collisions were hard. Instead, they show that distinguishing a function with many collisions from an injective function was hard. This then only works directly for expanding functions, which are of little interest to cryptographers. Zhandry [Zha15] shows for random functions a reduction from expanding functions to compressing functions, giving the desired lower bound for compressing functions. Our proof, in contrast, works directly with functions of arbitrary domain and range. These features suggests that our proof technique is fundamentally different than those of prior works.

By generalizing our collision bound slightly, we can obtain an $\Omega(2^{n/(k+1)})$ lower bound for finding a set of distinct points x_1, \dots, x_k such that $\sum_i H(x_i) = 0$. This bound is tight as long as $n \leq km$ by adapting the collision-finding algorithm of [BHT98] to this problem. Again, our proof is obtained by modifying just a few lines of the pre-image search proof.

1.2 Related Works

Ristenpart, Shacham, and Shrimpton [RSS11] shows that indifferentiability is insufficient for replacing a concrete hash function with a random oracle in the setting of multi-stage games. Nonetheless, Mittelbach [Mit14] shows that indifferentiability can still be useful in these settings. Exploring the quantum analogs of these results is an interesting direction for future research.

Acknowledgements

This work is supported in part by NSF and DARPA. Opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of NSF or DARPA.

³ and the square root comes from the fact that the norm of the sum of q unit vectors of disjoint support is \sqrt{q}

2 Preliminaries

Distinguishing quantum states. The density matrix captures all statistical information about a mixed state. That is, if two states have the same density matrix, then they are perfectly indistinguishable.

For density matrices ρ, ρ' that are not identical, we define the trace distance as $T(\rho, \rho') = \frac{1}{2} \sum_i |\lambda_i|$, where λ_i are the eigenvalues of $\rho - \rho'$. The trace distance captures the maximum distinguishing advantage amongst all possible measurements of the state.

We will need the following Theorem of Bennett et al. (which we have slightly improved, see full version [Zha18] for the improved proof):

Lemma 1 ([BBV97]). *Let $|\phi\rangle$ and $|\psi\rangle$ be quantum states with Euclidean distance ϵ . Then $T(|\phi\rangle\langle\phi|, |\psi\rangle\langle\psi|) = \epsilon\sqrt{1 - \epsilon^2/4} \leq \epsilon$.*

We will also need the following relaxation of commuting operations:

Definition 1. *Let U_0, U_1 be unitaries over the same quantum system. We say that U_0, U_1 ϵ -almost commute if, for any initial state ρ , the images of ρ under $U_0 U_1$ and $U_1 U_0$ are at most ϵ -far in trace distance.*

3 Oracle Variations

Here, we describe several oracle variations. The oracles will all be equivalent; the only difference is that the oracle registers and/or the query registers are encoded in different ways between queries. We start with the usual quantum random oracle, which comes in two flavors that we call the *standard oracle* and *phase oracle*. Then we will give our *compressed standard and phase oracles*.

Standard Oracle. Here, the oracle $H : \{0, 1\}^m \rightarrow \{0, 1\}^n$ is represented as its truth table: a vector of size 2^m where each component is an n -bit string.

The oracle takes as input a state consisting of three sets of registers: m -qubit x registers representing inputs to the function, n -qubit y registers for writing the response, and $n2^m$ -qubit H registers containing the truth table of the actual function. The x, y registers come from the adversary, and the H registers are the oracle's state, which is hidden from the adversary accept by making queries. On basis states $|x, y\rangle \otimes |H\rangle$, the oracle performs the map $|x, y\rangle \otimes |H\rangle \mapsto |x, y \oplus H(x)\rangle \otimes |H\rangle$.

For initialization, the oracle H will be initialized to the uniform superposition over all H : $\frac{1}{\sqrt{2^m \times 2^n}} \sum_H |H\rangle$. We will call this oracle StO .

The only difference between StO and the usual quantum random oracle model is that, in the usual model, H starts out as a uniformly chosen random function rather than a superposition (that is, the H registers are the completely mixed state). We will call the oracle with this different initialization StO' .

Lemma 2. *StO and StO' are perfectly indistinguishable. That is, for any adversary A making oracle queries, let $A^{\text{StO}}()$ and $A^{\text{StO}'}()$ denote the algorithm interfacing with StO and StO' , respectively. Then $\Pr[A^{\text{StO}}() = 1] = \Pr[A^{\text{StO}'}() = 1]$*

Proof. This can be seen by tracing out the oracle registers. The mixed state of the adversary in both cases will be identical. \square

Thus, our initialization is equivalent to H being a uniformly random oracle.

Phase Oracle. We will also consider the well-known phase model of oracle queries. This model technically offers a different interface to the adversary, but can be mapped to the original oracle by simple Hadamard operations.

The oracle takes as input a state consisting of three sets of registers: x registers representing inputs to the function, z phase registers, and H registers containing the truth table of the actual function. On basis states $|x, y\rangle \otimes |H\rangle$, it performs the map $|x, z\rangle \otimes |H\rangle \mapsto (-1)^{y \cdot H(x)} |x, z\rangle \otimes |H\rangle$.

For initialization, H is the uniform superposition as before. We will call this oracle PhO. Analogous to the above, this is equivalent to the case where H is uniformly random. The following Lemma is implicit in much of the literature on quantum-accessible oracles:

Lemma 3. *For any adversary A making queries to StO, let B be the adversary that is identical to A , except it performs the Hadamard transformation $H^{\otimes n}$ to the response registers before and after each query. Then $\Pr[A^{\text{StO}}() = 1] = \Pr[B^{\text{PhO}}() = 1]$*

Compressed Standard Oracles. We now define our compressed standard oracles. The intuition for our compressed standard oracle is the following. Let $|\tau\rangle$ be the uniform superposition. In the standard (uncompressed) oracle, suppose for each of the 2^m output registers, we perform the computation mapping $|\tau\rangle \mapsto |\tau\rangle|1\rangle$ and $|\phi\rangle \mapsto |\phi\rangle|0\rangle$ for any $|\phi\rangle$ orthogonal to $|\tau\rangle$. In other words, this computation tests whether or not the state of the output registers is 0 in the Fourier basis. We will write the output of the computation in some auxiliary space. Now the state of the oracle is a superposition over truth tables, and a superposition over vectors in $\{0, 1\}^{2^m}$ containing the output of the tests. A straightforward exercise (and a consequence of our analysis below) shows that if we perform these tests after q queries, all vectors in the test vector superposition have at most q positions containing a 0. The reason is, roughly, if we do the tests before any queries the vector will be identically 1 since we had a uniform superposition (which is 0 in the Fourier basis). Then, each query affects only one position of the superposition, increasing the number of 0's by at most 1.

Also notice that anywhere the vector contains a 1, the corresponding truth table component contains exactly the uniform superposition $|\tau\rangle$. Anywhere the vector contains a 0, the corresponding truth table component contains a state that is guaranteed to be orthogonal to $|\tau\rangle$.

What we can do then is compress this overall state. We will simply write down all the positions where the test vector contained a 0, and keep track of the truth table component for that position. Everywhere else we can simply ignore since we know what the truth table contains. The result is a (superposition over) database consisting of at most q input/output pairs.

In more detail, a database D will be a collection of (x, y) pairs, where $(x, y) \in D$ means the function has been specified to have value y on input x . We will write $D(x) = y$ in this case. If, for an input x there is no pair $(x, y) \in D$, then we will write $D(x) = \perp$, indicating that the function has not been specified. We will maintain that a database D only contains at most one pair for a given x .

Concretely, if we have an upper bound t on the number of specified points, a database D will be represented an element of the set S^t , where $S = (\{0, 1\}^m \cup \{\perp\}) \times \{0, 1\}^n$. Each value in S is an (x, y) pair; if $x \neq \perp$ the pair means $D(x) = y$, and $x = \perp$ means the pair is unused. For $x_1 < x_2 < \dots < x_\ell$ and y_1, \dots, y_ℓ , the database representing that input x_i has been set to y_i for $i \in [\ell]$, with all other points unspecified, will be represented as:

$$((x_1, y_1), (x_2, y_2), \dots, (x_\ell, y_\ell), (\perp, 0^n), \dots, (\perp, 0^n))$$

where the number of $(\perp, 0^n)$ pairs is equal to $t - \ell$.

After query q , the state of the oracle will be a superposition of databases in this form, using the upper bound $t = q$. So initially the state is empty. We will maintain several invariants:

- For any database in the support of the superposition, for any (x, y) pair where $x = \perp$, we have that $y = 0^n$. All $(\perp, 0^n)$ pairs are at the end of the list.
- For any database in the support of the superposition, if (x, y) occurs before (x', y') , it must be that $x < x'$.
- For any of the ℓ positions that have been specified, the y registers are in a state that is orthogonal to the uniform superposition $|\tau\rangle$ (indicating that in the Fourier domain, the registers do *not* contain 0).

We also need to describe several procedures on databases. Let $|D|$ be the number of pairs $(x, y) \in D$ for $x \neq \perp$. For a database D with $|D| < t$ and $D(x) = \perp$, write $D \cup (x, y)$ to be the new database obtained by adding the pair (x, y) to D , inserting in the appropriate spot to maintain the ordering of the x values. Since $|D|$ was originally less than t , there will be at least one $(\perp, 0^n)$ pair, which is deleted. Therefore, the overall number of pairs (including \perp s) in D and $D \cup \{(x, y)\}$ are the same.

Before describing how to process a query, we need to describe a local decompression procedure StdDecomp_x which acts on databases. This is a unitary operation. It suffices to describe its action on a set of orthonormal states. Let t be the current upper bound on the number of set points.

- For D such that $D(x) = \perp$ and $|D| < t$,

$$\text{StdDecomp}_x |D\rangle = \frac{1}{\sqrt{2^n}} \sum_y |D \cup (x, y)\rangle$$

That is, StdDecomp_x inserts into D the pair $(x, |\tau\rangle)$. This corresponds to decompressing the value of the database at position x

- For D such that $D(x) = \perp$ and $|D| = t$, $\text{StdDecomp}_x|D\rangle = |D\rangle$. This means, if there is no room to expand for decompression, StdDecomp_x does nothing. Note that these states are illegal and StdDecomp_x will never be applied to such states.
- For a D' such that $D'(x) = \perp$ and $|D'| < t$,

$$\begin{aligned}\text{StdDecomp}_x \left(\sum_y (-1)^{z \cdot y} |D' \cup (x, y)\rangle \right) &= \sum_y (-1)^{z \cdot y} |D' \cup (x, y)\rangle \text{ for } z \neq 0 \\ \text{StdDecomp}_x \left(\frac{1}{\sqrt{2^n}} \sum_y |D' \cup (x, y)\rangle \right) &= |D'\rangle\end{aligned}$$

In other words, if D already is specified on x , and moreover if the corresponding y registers are in a state orthogonal to $|\tau\rangle$ (meaning they do not contain 0 in the Fourier domain), then there is no need to decompress and StdDecomp_x is the identity. On the other hand, if D is specified at x and the corresponding y registers are in the state $|\tau\rangle$, StdDecomp_x will remove x and the y register superposition from D .

Note that the left-hand sides of last two cases form an orthonormal basis for the span of $|D\rangle$ such that $D(x) \neq \perp$. The left-hand sides of the first two cases form an orthonormal basis for the remaining D . Thus, StdDecomp_x is defined on an orthonormal basis, which by linearity defines it on all states. The right-hand sides are the same basis states just in a different order. As such, this operation maps orthogonal states to orthogonal states, and is therefore unitary. Note that StdDecomp_x is actually an involution, as applying it twice results in the identity. Let StdDecomp be the related unitary operating on a quantum system over x, y, D states, defined by it's action on the computational basis states as:

$$|x, y\rangle \otimes |D\rangle = |x, y\rangle \otimes \text{StdDecomp}_x|D\rangle$$

In other words, in superposition it applies StdDecomp_x to $|D\rangle$, where x is taken from the x registers.

For some additional notation, we will take $y \oplus \perp = y$ and $y \cdot \perp = 0$. Let Increase be the procedure which initializes a new register $|(\perp, 0^n)\rangle$ and appends it to the end. In other words, $\text{Increase}|x, y\rangle \otimes |D\rangle = |x, y\rangle \otimes |D\rangle |(\perp, 0^n)\rangle$, where $|D\rangle |(\perp, 0^n)\rangle$ is interpreted as a database computing the same partial function as D , but with the upper bound on number of points increased by 1.

Let CStO' , CPhsO' be unitaries defined on the computational basis states as

$$\begin{aligned}\text{CStO}'|x, y\rangle \otimes |D\rangle &= |x, y \oplus D(x)\rangle \otimes |D\rangle \\ \text{CPhsO}'|x, y\rangle \otimes |D\rangle &= (-1)^{y \cdot D(x)} |x, y\rangle \otimes |D\rangle\end{aligned}$$

Finally, we describe the CStO and CPhsO oracles:

$$\begin{aligned}\text{CStO} &= \text{StdDecomp} \circ \text{CStO}' \circ \text{StdDecomp} \circ \text{Increase} \\ \text{CPhsO} &= \text{StdDecomp} \circ \text{CPhsO}' \circ \text{StdDecomp} \circ \text{Increase}\end{aligned}$$

In other words, increase the bound on the number of specified points, then uncompress at x (which is ensured to have enough space since we increased the bound), apply the query (which is ensured to be specified since we decompressed), and then re-compress.

Lemma 4. *CStO and StO are perfectly indistinguishable. CPhsO and PhO are perfectly indistinguishable. That is, for any adversary A , we have $\Pr[A^{\text{CStO}}() = 1] = \Pr[A^{\text{StO}}() = 1]$, and for any adversary B , we have $\Pr[B^{\text{CPhsO}}() = 1] = \Pr[B^{\text{PhO}}() = 1]$.*

Proof. We prove the case for CStO and StO, the other case being almost identical. We prove security through a sequence of hybrids.

Hybrid 0. In this case, the adversary interacts with StO. That is, the oracle's database is initialized to the uniform superposition over all H , and each query performs the unitary mapping $|x, y\rangle \otimes |H\rangle \mapsto |x, y \oplus H(x)\rangle \otimes |H\rangle$.

Hybrid 1. In this hybrid, we use a slightly different way of representing the function H . Instead of writing H as a truth table, we represent it as a complete database $D = ((0, H(0)), (1, H(1)), \dots, (2^m - 1, H(2^m - 1)))$. Here, the upper bound on the number of determined points is exactly 2^m . The oracle's state starts out as

$$\frac{1}{\sqrt{2^{n2^m}}} \sum_H |((0, H(0)), (1, H(1)), \dots, (2^m - 1, H(2^m - 1)))\rangle$$

The update procedure for each query is simply CStO', meaning that each query maps $|x, y\rangle \otimes |((0, H(0)), (1, H(1)), \dots, (2^m - 1, H(2^m - 1)))\rangle$ to $|x, y \oplus H(x)\rangle \otimes |((0, H(0)), (1, H(1)), \dots, (2^m - 1, H(2^m - 1)))\rangle$.

Hybrid 1 is identical to **Hybrid 0**, except that we have inserted the input points $1, \dots, 2^m - 1$ into the oracle's state, which has no effect on the adversary.

Hybrid 2. Next, introduce a global decompression procedure $\text{StdDecomp}'$, which applies StdDecomp_x for all x in the domain, one at a time from 0 up to $2^m - 1$.

We observe that when the upper bound on determined points is 2^m , then StdDecomp_x commutes with $\text{StdDecomp}_{x'}$ for any x, x' . This readily follows from the fact that when the upper bound is $t = 2^m$, $D(x) = \perp$ implies $|D| < t$.

In **Hybrid 2**, the oracle starts out as the empty database with upper bound 2^m . Then, each query is implemented as $\text{StdDecomp}' \circ \text{CStO}' \circ \text{StdDecomp}'$.

Notice that $\text{StdDecomp}'$ only affects the oracle's registers and therefore commutes with the any computation on the adversary's side. Also notice that between each two queries, $\text{StdDecomp}'$ is applied twice and that it is an involution. Therefore the two applications cancel out. At the beginning, $\text{StdDecomp}'$ is applied to an empty database, which maps it to the uniform superposition

$$\frac{1}{\sqrt{2^{n2^m}}} \sum_H |((0, H(0)), (1, H(1)), \dots, (2^m - 1, H(2^m - 1)))\rangle$$

before the first application of CStO'. Therefore, this hybrid is perfectly indistinguishable from **Hybrid 1**.

Hybrid 3. This hybrid applies $\text{StdDecomp} \circ \text{CStO}' \circ \text{StdDecomp}$ for each query.

To prove indistinguishability from **Hybrid 2**, consider a database D with upper bound 2^m but where $|D| = \ell$ for some $\ell \leq 2^m$. Notice that for any D' in the support of $\text{StdDecomp}_{x'}|D\rangle$, $D'(x) = D(x)$ for all $x \neq x'$. This means

$$\begin{aligned}\text{CStO}' \circ \text{StdDecomp}_{x'}(|x, y\rangle \otimes |D\rangle) &= \text{StdDecomp}_{x'}(|x, y \oplus D(x)\rangle \otimes |D\rangle) \\ &= \text{StdDecomp}_{x'} \circ \text{CStO}'(|x, y\rangle \otimes |D\rangle)\end{aligned}$$

In other words, when the query register contains $x \neq x'$, $\text{StdDecomp}_{x'}$ and CStO' commute. Therefore,

$$\begin{aligned}\text{StdDecomp}' \circ \text{CStO}' \circ \text{StdDecomp}'(|x, y\rangle \otimes |D\rangle) &= \text{StdDecomp}_x \circ \text{CStO}' \circ \text{StdDecomp}_x(|x, y\rangle \otimes |D\rangle) \\ &= \text{StdDecomp} \circ \text{CStO}' \circ \text{StdDecomp}(|x, y\rangle \otimes |D\rangle)\end{aligned}$$

This shows that **Hybrid 2** and **Hybrid 3** are identical.

Hybrid 4. Finally, this hybrid is the compressed standard oracle: the oracle's state starts out empty, and CStO is applied for each query.

To prove equivalence, first notice that for any x, y, D , $\text{StdDecomp} \circ \text{CStO}' \circ \text{StdDecomp}(|x, y\rangle \otimes |D\rangle)$ has support on databases D' such that $|D'| \leq |D| + 1$. Indeed, all D' are defined on the same inputs except for possibly the input x .

This means that after q queries in **Hybrid 3**, the oracle's registers only have support on D containing at most q defined points; the remaining $\geq 2^m - q$ points are all $(\perp, 0^n)$. Therefore, we can discard all but the first q pairs in D , without affecting the adversary's state. The result is identical to **Hybrid 4**. \square

In the full version [Zha18], we give several more oracle variations; while not used in this work, they may be useful in other settings. These variations also provide an alternative way to arrive at the compressed standard oracles.

3.1 A Useful Lemma

Here, we provide a lemma which relates the adversary's knowledge of an oracle output to the probability that point appears in the compressed oracle database. This lemma is proved in the full version [Zha18], and follows from a straightforward (albeit delicate) analysis off the action of CStO .

Lemma 5. *Consider a quantum algorithm A making queries to a random oracle H and outputting tuples $(x_1, \dots, x_k, y_1, \dots, y_k, z)$. Let R be a collection of such tuples. Suppose with probability p , A outputs a tuple such that (1) the tuple is in R and (2) $H(x_i) = y_i$ for all i . Now consider running A with the oracle CStO , and suppose the database D is measured after A produces its output. Let p' be the probability that (1) the tuple is in R , and (2) $D(x_i) = y_i$ for all i (and in particular $D(x_i) \neq \perp$). Then $\sqrt{p} \leq \sqrt{p'} + \sqrt{k/2^n}$*

4 Quantum Query Bounds Using Compressed Oracles

In this section, we re-prove several known query complexity lower bounds, as well as provide some new bounds. All these bounds follow from simple applications of our compressed oracles.

4.1 Optimality of Grover Search

Here, we re-prove that the quadratic speed-up of Grover search is optimal. Specifically, we prove that for a random function $H : \{0, 1\}^m \rightarrow \{0, 1\}^n$, any q query algorithm has a success probability of at most $O(q^2/2^n)$ for finding a pre-image of 0^n (or any fixed value).

Theorem 1. *For any adversary making q queries to CStO or CPhsO and an arbitrary number of database read queries, if the database D is measured after the q queries, the probability it contains a pair of the form $(x, 0^n)$ is at most $O(q^2/2^n)$.*

Proof. Let $0^n \in D$ mean that D contains a pair of the form $(x, 0^n)$. The compressed oracle's database starts out empty, so the probability $0^n \in D$ is zero. We will show that the probability cannot rise too much with each query. We consider compressed phase queries, CPhsO. Compressed standard queries are handled analogously. Consider the joint state of the adversary and oracle just before the q th CPhsO query:

$$|\psi\rangle = \sum_{x,y,z,D} \alpha_{x,y,z,D} |x, y, z\rangle \otimes |D\rangle$$

Where D represents the compressed phase oracle, x, y as the query registers, and z as the adversary's private storage. Define P as the projection onto the span of basis states $|x, y, z\rangle \otimes |D\rangle$ such that $0^n \in D$. Our goal will be to relate the norms of $P|\psi\rangle$ (the magnitude before the query) to $P \cdot \text{CPhsO}|\psi\rangle$ (the magnitude after the query).

Define projections Q onto states such that (1) $0^n \notin D$ (meaning the database does not yet contain 0^n), (2) $y \neq 0$ (meaning CPhsO will affect D), and (3) $D(x) = \perp$ (meaning D has not yet been specified at x). Define projection R onto states such that $0^n \notin D$, $y \neq 0$ and $D(x) \neq \perp$; projection S onto states such that $0^n \notin D$, $y = 0$. Then $P + Q + R + S = \mathbf{I}$.

Consider $Q|\psi\rangle$. CPhsO maps basis states $|x, y, z\rangle \otimes |D\rangle$ in the support of $Q|\psi\rangle$ to $|x, y, z\rangle \otimes \frac{1}{\sqrt{2^n}} \sum_w (-1)^{y \cdot w} |D \cup (x, w)\rangle$. Since $0^n \notin D$, applying P to this state will yield $|x, y, z\rangle \otimes \frac{1}{\sqrt{2^n}} |D \cup (x, 0^n)\rangle$. Notice that the images of the different basis states are orthogonal. Therefore, $\|P \cdot \text{CPhsO} \cdot Q|\psi\rangle\| = \frac{1}{\sqrt{2^n}} \|Q|\psi\rangle\|$.

For basis vectors in the support of R , we must have $D(x) \notin \{\perp, 0^n\}$. Let D' be the database with x removed, and write $D = D' \cup (x, w)$ for $w = D(x)$. Then

some algebraic manipulations show that $\text{CPhsO}|x, y, z\rangle \otimes |D' \cup (x, w)\rangle$ is:

$$\begin{aligned} |x, y, z\rangle \otimes & \left((-1)^{y \cdot w} \left(|D' \cup (x, w)\rangle + \frac{1}{\sqrt{2^n}} |D'\rangle \right) \right. \\ & \left. + \frac{1}{2^n} \sum_{y'} (1 - (-1)^{y \cdot w} - (-1)^{y \cdot y'}) |D' \cup (x, y')\rangle \right) \end{aligned}$$

Then $P \cdot \text{CPhsO}|x, y, z\rangle \otimes |D' \cup (x, w)\rangle = \frac{-(-1)^{y \cdot w}}{2^n} |x, y, z\rangle \otimes |D' \cup (x, 0^n)\rangle$. Write $R|\psi\rangle = \sum_{x, y, z, D', w} \alpha_{x, y, z, D', w} |x, y, z\rangle \otimes |D' \cup (x, w)\rangle$. Then $\|P \cdot \text{CPhsO} \cdot R|\psi\rangle\|^2$ is equal to:

$$\frac{1}{4^n} \sum_{x, y, z, D'} \left\| \sum_w \alpha_{x, y, z, D', w} (-1)^{y \cdot w} \right\|^2 \leq \frac{1}{2^n} \sum_{x, y, z, D', w} \|\alpha_{x, y, z, D', w}\|^2 = \frac{1}{2^n} \|R|\psi\rangle\|^2$$

Finally, $\|P \cdot \text{CPhsO} \cdot P|\psi\rangle\| \leq \|P|\psi\rangle\|$ and $\text{CPhsO} \cdot S|\psi\rangle = S|\psi\rangle$. Putting it all together, we have that $\|P \cdot \text{CPhsO}|\psi\rangle\| \leq \|P|\psi\rangle\| + \frac{1}{\sqrt{2^n}} (\|Q|\psi\rangle\| + \|R|\psi\rangle\|) \leq \|P|\psi\rangle\| + \frac{1}{\sqrt{2^n}}$.

Therefore, after q queries, we have that the projection onto D containing a zero has norm at most $q/\sqrt{2^n}$. Now, the probability the database in $|\psi\rangle$ contains a 0^n is just the square of this norm, which is at most $\frac{q^2}{2^n}$. \square

The following is obtained by combining Theorem 1 with Lemma 5:

Corollary 1. *After making q quantum queries to a random oracle, the probability of finding a pre-image of 0^n is at most $O(q^2/2^n)$.*

Proof. We will assume the adversary always makes a final query on it's output x , and outputs $(x, H(x))$. This comes at the cost of at most 1 query, so it does not affect the asymptotic result. Then we can use the relation $R(x, y)$ which accepts if and only if $y = 0^n$. In the second experiment of Lemma 5, the only way for the adversary to win is to have the database contain a pre-image of 0^n . As such, Theorem 1 shows $p' = O(q^2/2^n)$. Then Lemma 5 shows that $p = O(q^2/2^n)$, which is exactly the probability the adversary outputs a pre-image of 0^n when interacting with the real random oracle.

4.2 Collision Lower Bound

Theorem 2. *For any adversary making q queries to CStO or CPhsO and an arbitrary number of database read queries, if the database D is measured after the q queries, the resulting database will contain a collision with probability at most $O(q^3/2^n)$*

Proof. The proof involves changing just a few lines of the proof of Theorem 1. We define P to project onto databases D containing a collision, and re-define Q, R, S accordingly. Write $Q|\psi\rangle = \sum_{x, y, z, D} \alpha_{x, y, z, D} |x, y, z\rangle \otimes |D\rangle$. Then

$$P \cdot \text{CPhsO} \cdot Q|\psi\rangle = \sum_{x, y, z, D} \alpha_{x, y, z, D} |x, y, z\rangle \otimes \frac{1}{\sqrt{2^n}} \sum_{w \in D} |D \cup (x, w)\rangle$$

We can write this as the $\frac{1}{\sqrt{2^n}} \sum_i |\phi_i\rangle$, where $|\phi_i\rangle$ is the partial sum which sets w to be the i th element in D (provided it exists). The $|\phi_i\rangle$ are orthogonal, and satisfy $\|\phi_i\rangle\| \leq \|Q|\psi\rangle\|$. Moreover, after q queries D has size at most q , and so there are at most q of the $|\phi_i\rangle$. Therefore, $\|P \cdot \text{CPhsO} \cdot Q|\psi\rangle\| \leq \sqrt{q/2^n} \|Q|\psi\rangle\|$.

By a similar argument, $\|P \cdot \text{CPhsO} \cdot R|\psi\rangle\| \leq \sqrt{q/2^n} \|R|\psi\rangle\|$. Putting everything together, this shows that the norm of $P|\psi\rangle$ increases by at most $\sqrt{q/2^n}$ with each query. Therefore, after q queries, the total norm is at most $\sqrt{q^3/2^n}$, giving a probability of $q^3/2^n$. \square

Corollary 2. *After making q quantum queries to a random oracle, the probability of finding a collision is at most $O(q^3/2^n)$.*

4.3 More General Settings

We can easily generalize even further. Let R be a relation on ℓ -tuples over $\{0, 1\}^n$. Say that R is *satisfied* on a database D if D contains ℓ distinct pairs (x_i, y_i) such that $R(y_1, \dots, y_\ell) = 1$. Let $k(q)$ be the maximum number of y that can be added to an unsatisfied database of size at most $q - 1$ to make it satisfied.

Theorem 3. *For any adversary making q queries to CStO or CPhsO and an arbitrary number of database read queries, if the database D is measured after the q queries, the resulting database will be satisfied with probability at most $O(q^2 k(q)/2^n)$.*

For the k -sum problem, there are at most $\binom{q}{k-1}$ incomplete tuples that can be completed by adding a new point. As such, $k(q) \leq \binom{q}{k-1} \leq q^{k-1}$. This gives:

Corollary 3. *After making q quantum queries to a random oracle, the probability of finding k distinct inputs x_i such that $\sum_i H(x_i) = 0^n$ is at most $O(q^{k+1}/2^n)$.*

5 Indifferentiability of A Simple Domain Extender

5.1 Definitions

Let $h : \{0, 1\}^m \rightarrow \{0, 1\}^n$ be a random oracle, and let $C^h : \{0, 1\}^M \rightarrow \{0, 1\}^N$ be a polynomial-sized stateless classical circuit that makes oracle queries to h .

Definition 2. *Let $H : \{0, 1\}^M \rightarrow \{0, 1\}^N$ be a random function. A stateful quantum polynomial-time simulator $\text{Sim}^H : \{0, 1\}^m \rightarrow \{0, 1\}^n$ is indifferentiable for C if, for any polynomial-time distinguisher \mathcal{D} making queries to h, H ,*

$$|\Pr[\mathcal{D}^{h, C^h}() = 1] - \Pr[\mathcal{D}^{\text{Sim}^H, H}() = 1]| < \text{negl}$$

Definition 3. *C^h is quantum indifferentiable from a random oracle if there exists an indifferentiable simulator Sim for C .*

Intuitively, in the “real” world, h is a random function and H is set to be C^h . C^h is indifferentiable if this real world is indistinguishable from an “ideal” world, where H is a random function, and h is set to be Sim^h for some efficient simulator Sim .

In order to help us prove indifferentiability of a simulator Sim , we introduce two weaker requirements. The first is *indistinguishability*, a weakened version of indifferentiability where the distinguisher is not allowed any queries to H :

Definition 4. *A simulator Sim is indistinguishable if, for any polynomial-time distinguisher \mathcal{D} making queries to h ,*

$$|\Pr[\mathcal{D}^h() = 1] - \Pr[\mathcal{D}^{\text{Sim}^H}() = 1]| < \text{negl}$$

Next, we introduce the notion of *consistency*. Here, we set h to be simulated by Sim^H , and we ask the adversary to distinguish honest evaluations of H from evaluations of C^h (where again h is still simulated by Sim^H).

Definition 5. *A simulator Sim is consistent if, for any polynomial-time distinguisher \mathcal{D} making queries to h, H , if H is simulated by Sim^H , then*

$$|\Pr[\mathcal{D}^{\text{Sim}^H, H}() = 1] - \Pr[\mathcal{D}^{\text{Sim}^H, C^{\text{Sim}^H}}() = 1]| < \text{negl}$$

Lemma 6. *Any consistent and indistinguishable simulator is indifferentiable.*

The proof of Lemma 6 is straightforward, and proved in the full version [Zha18].

Finally, it is straightforward to adapt the definitions and Lemma 6 to handle the case of many random compression functions h_1, \dots, h_ℓ . In this case, C makes queries to h_1, \dots, h_ℓ , \mathcal{D} has quantum oracle access to h_1, \dots, h_ℓ and H , while S makes quantum queries to H and simulates h_1, \dots, h_ℓ .

5.2 A Simple Domain Extender

We now consider a simple domain extender. Let $h_1 : \{0, 1\}^m \rightarrow \{0, 1\}^n, h_2 : \{0, 1\}^n \times \{0, 1\}^\ell \rightarrow \{0, 1\}^n$ be two functions. Let $C^{h_1, h_2}(x_1, x_2) = h_2(h_1(x_1), x_2)$

Theorem 4. *If h_1, h_2 are random oracles, the simple domain extender C is indifferentiable from a random oracle.*

Coron et al. [CDMP05] show that the indifferentiability of C is sufficient to prove the indifferentiability of Merkle-Damgård for a particular choice of prefix-free encoding (see paper for details). That part of the paper translates immediately to the quantum setting, so Theorem 4 then shows quantum indifferentiability for the same prefix free encoding. In the full version [Zha18], we show more generally that Merkle-Damgård is indifferentiable for *any* choice of prefix-free encoding. All the main ideas for the full proof are already contained in the proof of Theorem 4 below, just the details get a bit more complicated in the more general setting.

5.3 Our Simulator

Before describing our simulator, we need some terminology. For a database D of input/output pairs, a *collision* is two pairs $(x_1, y_1), (x_2, y_2) \in D, x_1 \neq x_2$ such that $y_1 = y_2$. For an input $(y, x_2) \in \{0, 1\}^n \times \{0, 1\}^\ell$, a *completion* in D is a pair $(x_1, y) \in D$. For such a completion, we will call $w = (x_1, x_2)$ the associated input.

We define a classical procedure `FindInput`. `FindInput` takes as input $x \in \{0, 1\}^n \times \{0, 1\}^\ell$, and a database D . It parses x as $(y, x_2) \in \{0, 1\}^n \times \{0, 1\}^\ell$. Then, it looks for a completion $(x_1, y) \in D$. If found, it will take, say, the completion with the smallest x_1 value, and output $(b = 1, w = (x_1, x_2))$. If no completion is found, it will output $(b = 0, w = 0^{m+\ell})$. Note that for the output values in D , `FindInput` only needs to apply an equality check on those values, testing if they contain y . By applying such an equality check to each output register, it can compute b and w . Looking forward, when we implement `FindInput` in superposition, this means `FindInput` only touches the output registers of D by making a computational basis test.

We are now ready to describe our simulator. `Sim` will keep a (superposition over) database D_a , which represents the simulation of the random oracle h_a that it will update according to the `CStO` update procedure. D_a is originally empty. It will also have a private random oracle h_b . For concreteness, h_b will be implemented using another instance of `CStO`, but it will be notationally convenient to treat h_b as being a uniformly random function.

On h_1 queries, `Sim` makes a query to h_a , performing the appropriate `CStO` update procedure to D_a . On h_2 queries, `Sim` performs a unitary operation with the following action on basis states:

$$|x, y\rangle \otimes |D_a\rangle \mapsto \begin{cases} |x, y \oplus h_b(x)\rangle \otimes |D_a\rangle & \text{if } \text{FindInput}(x, D_a) = (0, 0^{m+\ell}) \\ |x, y \oplus H(w)\rangle \otimes |D_a\rangle & \text{if } \text{FindInput}(x, D_a) = (1, w) \end{cases}$$

This unitary is straightforward to implement with a single query to each of h_b and H , and is detailed in the full version [Zha18].

In the next three subsections, we prove that our simulator is indifferentiable. In Section 5.4, we prove a useful commutativity lemma. Then in Sections 5.5 and 5.6, we prove the indistinguishability and consistency, respectively, of `Sim`. By Lemma 6, this proves that `Sim` is indifferentiable, proving Theorem 4.

5.4 The Almost Commutativity of `StdDecomp` and `FindInput`

Lemma 7. *Consider a quantum system over x, D, x', z . The following two unitaries $O(1/\sqrt{2^n})$ -almost commute:*

- `StdDecomp`, acting on the x, D registers.
- `FindInput`, taking as input the D, x' registers and XORing the output into z .

The intuition is that, for `StdDecomp` to have any effect, either (1) $D(x) = \perp$ or (2) $D(x)$ is in uniform superposition; `StdDecomp` will simply toggle between the two cases. Now, a uniform superposition puts a weight of $1/\sqrt{2^n}$ on each possible

y value. Since there is only a single possible y value for $D(x)$ that matches x' , it is exponentially unlikely that **FindInput** will find a match at input x in Case (2). On the other hand, it will *never* find a match at input x in Case (1). Hence, there is an exponentially small error between the action of **FindInput** on these two cases. We prove the lemma formally in the full version [Zha18].

5.5 Indistinguishability

Lemma 8. *Sim is indistinguishable. In particular, for any distinguisher \mathcal{D} making at most q queries to h_1, h_2 ,*

$$|\Pr[\mathcal{D}^{h_1, h_2}() = 1] - \Pr[\mathcal{D}^{\text{Sim}^H}() = 1]| < O(q^2 / \sqrt{2^n})$$

Proof. Recall that in the ideal world where h_1, h_2 are simulated by Sim^H , h_1 is implemented by a CStO oracle on database D_a . By applying Lemma 4, we can think of the simulator's other oracle h_b as another instance of CStO for a database D_b . Additionally, H can be simulated with yet another instance of CStO for a database E . Similarly, in the real world, h_1, h_2 will be implemented by independent instances of CStO with databases D_a, D_b . Note that, in either case, h_1 is implemented by a CStO oracle on database D_a . Therefore, the only difference between the two cases is how h_2 is implemented.

We define a classical encoding procedure **Encode** for pairs D_a, D_b of databases. Intuitively, **Encode** will scan the values $((z, x_2), y)$ in D_b , seeing if any of the (z, x_2) values correspond to a completion in D_a . If so, such a completion will have an associated input w . **Encode** will reasonably guess that such a completion corresponds to an evaluation of $H(w) = C^{h_1, h_2}(w)$. Therefore, **Encode** will remove the value $((z, x_2), y)$ in D_b , and add the pair (w, y) to a new database E , intuitively representing the oracle H . In more detail, **Encode** does the following:

- For each pair $((z, x_2), y) \in D_b$, run $\text{FindInput}((z, x_2), D_a) = (b, w)$. If $b = 1$, re-label the pair to (w, y)
- Remove all re-labeled pairs D_b (which are easily identifiable since the input will be larger) and place them in a new database E .

We define the following **Decode** procedure, which operates on triples D_a, D_b, E :

- Merge the databases D_b, E
- For each pair (w, y) that was previously in E , where $w = (x_1, x_2)$, evaluate $z = D_a(x_1)$. Re-label (w, y) to $((z, x_2), y)$. If $z = \perp$ or if the input (z, x_2) was already in the database, output \perp and abort.

Note that **Encode**, **Decode** are independent of the order elements are processed. It also follows from the descriptions above that $\text{Decode}(\text{Encode}(D_a, D_b)) = (D_a, D_b)$. Therefore, **Encode** can be implemented in superposition, giving the unitary that maps $|D_a, D_b\rangle$ to $|\text{Encode}(D_a, D_b)\rangle$. Also note that $\text{Encode}(\emptyset, \emptyset) = (\emptyset, \emptyset, \emptyset)$.

With this notation in hand, we are now ready to prove security: consider a potential distinguisher \mathcal{D} . We prove security through a sequence of hybrids.

Hybrid 0. This is the real world, where h_1, h_2 are random oracles. Let p_0 be the probability \mathcal{D} outputs 1 in this case.

Hybrid 1. This is still the real world, but we add an abort condition. Namely, after any query to h_1 , we measure if the database h_a contains a collision; if so, we immediately abort and stop the simulation. Let p_1 be the probability \mathcal{D} outputs 1 in Hybrid 1.

Lemma 9. $|p_1 - p_0| \leq O(\sqrt{q^3/2^n})$

Proof. First, suppose that before the i th query to h_1 , the superposition over h_a has support only on databases containing no collisions. Let $|\psi\rangle$ be the joint state of the adversary and simulator just after the query to h_1 . Then write $|\psi\rangle = |\psi_0\rangle + |\psi_1\rangle$ where $|\psi_0\rangle$ is the projection onto states where h_a has no collisions, and $|\psi_1\rangle$ is the projection onto states where h_a contains at least one collision. Following the proof of Theorem 2, we know that $\||\psi_1\rangle\| \leq \sqrt{i/2^n}$.

Therefore, if we let $|\psi_q\rangle$ be the joint state after the q th query in **Hybrid 0** and $|\phi_q\rangle$ the joint state in **Hybrid 2**, we would have that $\||\psi_q\rangle - |\phi_q\rangle\| \leq \sum_{i=0}^q \sqrt{i/2^n} \leq O(\sqrt{q^3/2^n})$. By Lemma 1, this means that $|p_1 - p_0| \leq O(\sqrt{q^3/2^n})$. \square

Hybrid 2. In this hybrid, there are three databases D_a, D_b, E , initialized to $|\emptyset, \emptyset, \emptyset\rangle$. Each query is answered in the following way:

- Apply `Decode` to the D_a, D_b, E registers. Measure if `Decode` gives \perp , in which case abort. Otherwise, there are now just two database registers D_a, D_b .
- Answer an h_1 (resp. h_2) query by applying the `CStO` update procedure to D_a (resp. D_b).
- Apply `Encode` to D_a, D_b .
- Apply the collision check to the database D_a .

Let p_2 be the probability \mathcal{D} outputs 1 in Hybrid 2.

Lemma 10. $p_1 = p_0$

Proof. We start with **Hybrid 1**. First, by Lemma 4, we can implement D_a, D_b in Hybrid 1 as independent instances of `CStO`. Now, between all the queries insert `Encode` followed by `Decode`. Also insert the two procedures before the first query. Now each query is preceded by a `Decode` and followed by a collision check and an `Encode`. Note that `Encode`, `Decode` do not affect the database D_a , and so commute with the collision check. Therefore, we can swap the order of the collision check and `Encode` that follow each query.

By merging the `Decode`, query, `Encode` and collision check operations together, we get exactly the update procedure of Hybrid 2. All that's left is an initial `Encode` procedure at the very beginning, which produces $|\emptyset, \emptyset, \emptyset\rangle$ as the database state, just as in Hybrid 2. \square

Hybrid 3. This hybrid is the ideal world, where h_1, h_2 queries are answered by Sim , except that we will have the abort condition if a collision in h_a is ever found. In other words, instead of decoding, applying the query, and then encoding, in Hybrid 3 we act directly on the encoded state using the algorithms specified by Sim . For h_1 queries, the difference from Hybrid 2 is just that the queries are made directly to h_a , instead of Decode , then h_a query, then Encode . For h_2 queries, the differences appear more substantial. h_2 queries, on superpositions over x, y, D_a, D_b, E , can be summarized as follows:

1. Compute the unitary mapping $|x, y, D_a, D_b, E\rangle \mapsto |x, y, D_a, D_b, E, (b, w) = \text{FindInput}(x, D_a)\rangle$
2. In superposition, apply the following conditional procedures:
 3. Conditioned on $b = 0$,
 - (a) Apply StdDecomp to uncompress D_b at x .
 - (b) Apply in superposition the map

$$|x, y, D_a, D_b, E, b, w\rangle \mapsto |x, y \oplus D_b(x), D_a, D_b, E, b, w\rangle$$

- (c) Apply StdDecomp to re-compress D_b at x .
4. Conditioned on $b = 1$,
 - (a) Apply StdDecomp to uncompress E at w .
 - (b) Apply in superposition the map

$$|x, y, D_a, D_b, E, b, w\rangle \mapsto |x, y \oplus E(w), D_a, D_b, E, b, w\rangle$$

- (c) Apply StdDecomp to re-compress E at w .
5. Uncompute (b, w) by running $\text{FindInput}(x, D_a)$ in superposition again.

Let p_3 be the probability \mathcal{D} outputs 1 in this hybrid.

Lemma 11. $|p_3 - p_2| \leq O(q^2/\sqrt{2^n})$.

Proof. We start with the very last query, and gradually change the queries one-by-one from how they were answered in Hybrid 2 to Hybrid 3.

For h_1 queries, we observe that it suffices to swap the order of Encode and CStO . Indeed, suppose we move the final Encode to come before CStO . The previous query ended with an Encode , and now the current query begins with Decode then Encode . Since $\text{Decode} \circ \text{Encode}$ is the identity, all three of these operations collapse into a single Encode , which we keep at the end of the previous query. The result is that the current query is just a direct call to CStO , as in Hybrid 3. Then it remains to show that we can swap the order of Encode and CStO . For this, notice that Encode only interacts with D_a through FindInput . As such, all steps in Encode , CStO commute except for the two StdDecomp operations in CStO and the FindInput operation in Encode for each entry in D_b (plus another FindInput operation when un-computing the scratch-space of Encode in order to implement in superposition). By Lemma 7, these $\leq 4q$ operations each $O(1/\sqrt{2^n})$ -almost commute, meaning Encode and CStO $O(q/\sqrt{2^n})$ -almost commute.

For h_2 queries, fix an x, D_a and suppose D_a contains no collisions as guaranteed. There are two cases:

- $\text{FindInput}(x, D_a) = (0, 0^{m+\ell})$. Then in Hybrid 2, decoding/encoding does not affect the labeling for an (x, z) pair in D_b . As such, Hybrid 2 will uncompress D_b at x , apply the map $|x, y, D_a, D_b, E\rangle \mapsto |x, y \oplus D_b(x), D_a, D_b, E\rangle$ and then re-compress D_b at x , for these x, D_a .
- $\text{FindInput}(x, D_a) = (1, w)$. Then in Hybrid 2, by the collision-freeness of D_a , decoding will re-label a $(w, z) \in E$ (if present) to $(x, z) \in D_b$. The effect of Hybrid 2 in this case will be to uncompress E at w , apply the map $|x, y, D_a, D_b, E\rangle \mapsto |x, y \oplus E(x), D_a, D_b, E\rangle$, and then re-compress E at w .

In either case, answering h_2 queries in Hybrid 2 and 3 act identically. Therefore, this change introduces no error.

After q h_1 or h_2 queries, the total error between Hybrid 1 and Hybrid 2 is at most $O(q^2/\sqrt{2^n})$. \square

Hybrid 4. This is the ideal world, where we remove the abort condition from Hybrid 3. Let p_4 be the probability \mathcal{D} outputs 1 in Hybrid 4. By an almost identical proof to that of Lemma 9, we have:

Lemma 12. $|p_4 - p_3| \leq O(\sqrt{q^3/2^n})$

Summing up, we have that $|p_0 - p_4| < O(q^2/\sqrt{2^n})$, proving Lemma 8. \square

5.6 Consistency

Lemma 13. *Sim is consistent. In particular, for any distinguisher \mathcal{D} making at most q quantum queries to h_1, h_2, H ,*

$$|\Pr[\mathcal{D}^{\text{Sim}^H, H}() = 1] - \Pr[\mathcal{D}^{\text{Sim}^H, C^{\text{Sim}^H}}() = 1]| < O(\sqrt{q^3/2^n})$$

In other words, h_1, h_2 are simulated as Sim^H , and the adversary cannot distinguish between H and C^{h_1, h_2} .

Proof. We first work out how H queries are answered using C^{h_1, h_2} , when we simulate h_1, h_2 using Sim^H . The input registers will be labeled with $x = (x_1, x_2)$, and the output registers labeled with y .

1. First, make an h_1 query on the x_1 registers, writing the output to some new registers initialized to $z = 0^n$. Since we are implementing h_1 using CStO , this is accomplished using the following steps:
 - Apply StdDecomp to un-compress D_a at x_1
 - Evaluate the map $|x_1, z, x_2, y\rangle \otimes |D_a\rangle \mapsto |x_1, z \oplus D_a(x_1), x_2, y\rangle \otimes |D_a\rangle$, where z is the new register that was initialized to 0.
 - Re-compress D_a at x_1 by applying StdDecomp again.
2. Next, make an h_2 query on input (z, x_2) (where z where the registers created previously) with output registers y . This has the effect of mapping to:

$$\begin{aligned} & |x_1, z, x_2, y \oplus h_b(x)\rangle \otimes |D_a\rangle \text{ if } \text{FindInput}((z, x_2), D_a) = (0, 0^{m+\ell}) \\ & |x_1, z, x_2, y \oplus H(w)\rangle \otimes |D_a\rangle \text{ if } \text{FindInput}(z, x_2), D_a) = (1, w) \end{aligned}$$

3. Finally, make another h_1 query to un-compute the value of z . This is accomplished in the following steps:
 - (a) Apply `StdDecomp` to un-compress D_a at x_1
 - (b) Evaluate the map $|x_1, z, x_2, y\rangle \otimes |D_a\rangle \mapsto |x_1, z \oplus D_a(x_1), x_2, y\rangle \otimes |D_a\rangle$.
 - (c) Re-compress D_a at x_1 by applying `StdDecomp` again.
 - (d) Then discard the z registers.

Let \mathcal{D} be a potential distinguisher. We consider the following hybrids:

Hybrid 0. In this hybrid, H queries are answered using C^{h_1, h_2} , as worked out above. Let p_0 be the probability \mathcal{D} outputs 1.

Hybrid 1. This hybrid is identical to Hybrid 0, except that Steps 1c and 3a are removed. Let p_1 be the probability \mathcal{D} outputs 1 in this hybrid.

Lemma 14. $|p_1 - p_0| < O(q/\sqrt{2^n})$.

Proof. Since Steps 1c and 3a are inverses of each other, Hybrid 1 is equivalent to moving Step 3a up to occur just after Step 1c. Note that Step 2 only interacts with D_a through two applications of `FindInput` (one for computing, one for un-computing), which in turn $O(1/\sqrt{2^n})$ -almost commutes with Step 1c. By Lemma 7, each query to H therefore creates an error $O(1/\sqrt{2^n})$, yielding a total error of $O(q/\sqrt{2^n})$. \square

Hybrid 2. This hybrid is identical to Hybrid 2, except that after each query we measure if the database D_a contains a collision. If so, we abort and stop the simulation. Let p_2 be the probability \mathcal{D} outputs 1 in this hybrid. By an almost identical proof to that of Lemma 9, we have:

Lemma 15. $|p_2 - p_1| < O(\sqrt{q^3/2^n})$

Hybrid 3. This hybrid is identical to Hybrid 2 as outlined above, except that:

- Steps 1c and 3a are removed (as in Hybrid 1 and 2)
- The operation in Step 2 is replaced with

$$|x_1, z, x_2, y\rangle \otimes |D_a\rangle \mapsto |x_1, z, x_2, y \oplus H(x_1, x_2)\rangle \otimes |D_a\rangle$$

In other words Hybrid 3 is identical to Hybrid 2, except that we change Step 2. Let p_3 be the probability \mathcal{D} outputs 1 in this hybrid.

Lemma 16. $p_3 = p_2$.

Proof. In either hybrid, since we do not apply the Steps 1c and 3a, D_a is guaranteed to contain the pair (x_1, z) , where z is the same as in Step 2. Therefore, in Hybrid 2, $\text{FindInput}((z, x_2), D_a)$ is guaranteed to find a completion. Moreover, for D_a that contain no collisions, $\text{FindInput}((z, x_2), D_a)$ will find exactly the completion (x_1, z) . In this case, $w = (x_1, x_2)$, and Hybrid 2 will make a query to H on (x_1, x_2) . The end result is that for D_a containing no collisions, Step 2 is identical in both Hybrids. Since the collision check guarantees no collisions in D_a , this shows that the two hybrids are identical. \square

Hybrid 4. In this hybrid, H queries are made directly to H , but we still have the abort condition. Let p_4 be the probability \mathcal{D} outputs 1 in this hybrid.

Lemma 17. $p_4 = p_3$

Proof. In Hybrid 3, what remains of Steps 1 and 3 are exact inverses of each other and moreover commute with the new Step 2 from Hybrid 3. Therefore, we can remove Steps 1 and 3 altogether without affecting how oracle queries are answered. The result is identical to Hybrid 4. \square

Hybrid 5. This hybrid has H queries made directly to H , but without the abort condition. Let p_5 be the probability \mathcal{D} outputs 1 in this hybrid. By an almost identical proof to that of Lemma 9, we have:

Lemma 18. $|p_5 - p_4| < O(\sqrt{q^3/2^n})$

Overall then $|p_0 - p_5| < O(\sqrt{q^3/2^n})$, finishing the proof of Lemma 13. \square

6 Fujisaki Okamoto CCA-Secure Encryption

Here, we summarize our results on the Fujisaki-Okamoto transformation [FO99]. The transformation starts with a symmetric key encryption scheme $(\text{Enc}_S, \text{Dec}_S)$ and a public key encryption scheme $(\text{Gen}_P, \text{Enc}_P, \text{Dec}_P)$. Assuming only mild security properties of these two schemes (which are much easier to obtain than strong CCA security), the conversion produces a new public key scheme $(\text{Gen}, \text{Enc}, \text{Dec})$ which is secure against chosen ciphertext attacks. Let G, H are two random oracles, where G outputs keys for Enc_S and H outputs the random coins used by Enc_P . The scheme is as follows:

- $\text{Gen} = \text{Gen}_P$.
- $\text{Enc}(\mathbf{pk}, m)$ chooses a random $\delta \in \{0, 1\}^n$, and computes $d \leftarrow \text{Enc}_S(H(\delta), m)$. Then it computes $c \leftarrow \text{Enc}_P(\mathbf{pk}, \delta; G(\delta, d))$, and outputs (c, d) .
- $\text{Dec}(\mathbf{sk}, (c, d))$ first computes $\delta' \leftarrow \text{Dec}_P(\mathbf{sk}, c)$. Then it checks that $c = \text{Enc}_P(\mathbf{pk}, \delta'; G(\delta', d))$; if not, output \perp . Finally it computes and outputs $m' \leftarrow \text{Dec}_S(H(\delta'), d)$

The main difficulty in the classical proof of security is allowing the reduction to answer decryption queries. The key idea is that, in order for the adversary to generate a valid ciphertext, it must have queried the oracles on δ . The reduction will simulate G, H on the fly by keeping track of tables of input/output pairs. When a chosen ciphertext query comes in, it will scan the tables looking for a δ that “explains” the ciphertext.

In the quantum setting, we run into a similar recording barrier as in the indistinguishability setting. Our key observation is that the output values of the G, H tables are only used for set membership tests. Just like equality tests used in our indistinguishability simulator, set membership tests in the primal and Fourier domain very nearly commute. As such, we can use our compressed oracles to mimic the classical proof following our techniques. Our reduction can even handle chosen ciphertext queries on quantum superpositions of ciphertexts. In the full version [Zha18], we prove the following theorem:

Theorem 5. *If $(\text{Enc}_S, \text{Dec}_S)$ is one-time secure and $(\text{Gen}, \text{Enc}_P, \text{Dec}_P)$ is well-spread and one-way secure, then $(\text{Gen}, \text{Enc}, \text{Dec})$ is quantum CCA secure in the quantum random oracle model.*

Bibliography

- [AS04] Scott Aaronson and Yaoyun Shi. Quantum lower bounds for the collision and the element distinctness problems. *J. ACM*, 51(4):595–605, July 2004.
- [ATTU16] Mayuresh Vivekanand Anand, Ehsan Ebrahimi Targhi, Gelo Noel Tabia, and Dominique Unruh. Post-quantum security of the CBC, CFB, OFB, CTR, and XTS modes of operation. Cryptology ePrint Archive, Report 2016/197, 2016. <http://eprint.iacr.org/2016/197>.
- [BBBV97] Charles H. Bennett, Ethan Bernstein, Gilles Brassard, and Umesh Vazirani. Strengths and weaknesses of quantum computing. *SIAM J. Comput.*, 26(5):1510–1523, October 1997.
- [BDF⁺11] Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. Random oracles in a quantum world. In Dong Hoon Lee and Xiaoyun Wang, editors, *ASIACRYPT 2011*, volume 7073 of *LNCS*, pages 41–69. Springer, Heidelberg, December 2011.
- [BHT98] Gilles Brassard, Peter Høyer, and Alain Tapp. Quantum cryptanalysis of hash and claw-free functions. In *LATIN: : Theoretical Informatics, Latin American Symposium*, pages 163–169. Springer, Heidelberg, 1998.
- [BR93] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In Dorothy E. Denning, Raymond Pyle, Ravi Ganesan, Ravi S. Sandhu, and Victoria Ashby, editors, *ACM CCS 93*, pages 62–73. ACM Press, November 1993.
- [BZ13] Dan Boneh and Mark Zhandry. Secure signatures and chosen ciphertext security in a quantum computing world. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 361–379. Springer, Heidelberg, August 2013.
- [CBH⁺17] Jan Czajkowski, Leon Groot Bruinderink, Andreas Hülsing, Christian Schaffner, and Dominique Unruh. Post-quantum security of the sponge construction. Cryptology ePrint Archive, Report 2017/771, 2017. <http://eprint.iacr.org/2017/771>.
- [CDG⁺15] Daniel Cabarcas, Denise Demirel, Florian Göpfert, Jean Lancrenon, and Thomas Wunderer. An unconditionally hiding and long-term binding post-quantum commitment scheme. Cryptology ePrint Archive, Report 2015/628, 2015. <http://eprint.iacr.org/2015/628>.
- [CDG⁺17] Melissa Chase, David Derler, Steven Goldfeder, Claudio Orlandi, Sebastian Ramacher, Christian Rechberger, Daniel Slamanig, and

Greg Zaverucha. Post-quantum zero-knowledge and signatures from symmetric-key primitives. Cryptology ePrint Archive, Report 2017/279, 2017. <http://eprint.iacr.org/2017/279>.

[CDMP05] Jean-Sébastien Coron, Yevgeniy Dodis, Cécile Malinaud, and Prashant Puniya. Merkle-Damgård revisited: How to construct a hash function. In Victor Shoup, editor, *CRYPTO 2005*, volume 3621 of *LNCS*, pages 430–448. Springer, Heidelberg, August 2005.

[DFG13] Özgür Dagdelen, Marc Fischlin, and Tommaso Gagliardoni. The Fiat-Shamir transformation in a quantum world. In Kazue Sako and Palash Sarkar, editors, *ASIACRYPT 2013, Part II*, volume 8270 of *LNCS*, pages 62–81. Springer, Heidelberg, December 2013.

[Eat17] Edward Eaton. Leighton-Micali hash-based signatures in the quantum random-oracle model. In Carlisle Adams and Jan Camenisch, editors, *SAC 2017*, volume 10719 of *LNCS*, pages 263–280. Springer, Heidelberg, August 2017.

[FO99] Eiichiro Fujisaki and Tatsuaki Okamoto. Secure integration of asymmetric and symmetric encryption schemes. In Michael J. Wiener, editor, *CRYPTO’99*, volume 1666 of *LNCS*, pages 537–554. Springer, Heidelberg, August 1999.

[Gro96] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *28th ACM STOC*, pages 212–219. ACM Press, May 1996.

[IBM17] IBM. Ibm announces advances to ibm quantum systems and ecosystem, 2017. <https://www-03.ibm.com/press/us/en/pressrelease/53374.wss>.

[JZC⁺18] Haodong Jiang, Zhenfeng Zhang, Long Chen, Hong Wang, and Zhi Ma. IND-CCA-secure key encapsulation mechanism in the quantum random oracle model, revisited. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part III*, volume 10993 of *LNCS*, pages 96–125. Springer, Heidelberg, August 2018.

[Mit14] Arno Mittelbach. Salvaging indistinguishability in a multi-stage setting. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 603–621. Springer, Heidelberg, May 2014.

[MRH04] Ueli M. Maurer, Renato Renner, and Clemens Holenstein. Indistinguishability, impossibility results on reductions, and applications to the random oracle methodology. In Moni Naor, editor, *TCC 2004*, volume 2951 of *LNCS*, pages 21–39. Springer, Heidelberg, February 2004.

[NIS17] NIST. Candidate quantum-resistant cryptographic algorithms publicly available, 2017. <https://www.nist.gov/news-events/news/2017/12/candidate-quantum-resistant-cryptographic-algorithms-publicly-available>.

[RSS11] Thomas Ristenpart, Hovav Shacham, and Thomas Shrimpton. Careful with composition: Limitations of the indistinguishability framework. In Kenneth G. Paterson, editor, *EUROCRYPT 2011*, volume 6632 of *LNCS*, pages 487–506. Springer, Heidelberg, May 2011.

- [Son14] Fang Song. A note on quantum security for post-quantum cryptography. Cryptology ePrint Archive, Report 2014/709, 2014. <http://eprint.iacr.org/2014/709>.
- [TU16] Ehsan Ebrahimi Targhi and Dominique Unruh. Post-quantum security of the Fujisaki-Okamoto and OAEP transforms. In Martin Hirt and Adam D. Smith, editors, *TCC 2016-B, Part II*, volume 9986 of *LNCS*, pages 192–216. Springer, Heidelberg, October / November 2016.
- [Unr15] Dominique Unruh. Non-interactive zero-knowledge proofs in the quantum random oracle model. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 755–784. Springer, Heidelberg, April 2015.
- [Unr16] Dominique Unruh. Collapse-binding quantum commitments without random oracles. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *ASIACRYPT 2016, Part II*, volume 10032 of *LNCS*, pages 166–195. Springer, Heidelberg, December 2016.
- [YAJ⁺17] Youngho Yoo, Reza Azarderakhsh, Amir Jalali, David Jao, and Vladimir Soukharev. A post-quantum digital signature scheme based on supersingular isogenies. Cryptology ePrint Archive, Report 2017/186, 2017. <http://eprint.iacr.org/2017/186>.
- [Zha12a] Mark Zhandry. How to construct quantum random functions. In *53rd FOCS*, pages 679–687. IEEE Computer Society Press, October 2012.
- [Zha12b] Mark Zhandry. Secure identity-based encryption in the quantum random oracle model. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 758–775. Springer, Heidelberg, August 2012.
- [Zha15] Mark Zhandry. A note on the quantum collision and set equality problems. *Quantum Information and Computation*, 15(7& 8), 2015.
- [Zha18] Mark Zhandry. How to record quantum queries, and applications to quantum indistinguishability. Cryptology ePrint Archive, Report 2018/276, 2018. <https://eprint.iacr.org/2018/276>.