

Portland State University

PDXScholar

Computer Science Faculty Publications and
Presentations

Computer Science

11-4-2022

Quantum Key-Length Extension

Joseph Jaeger

Georgia Institute of Technology

Fang Song

Portland State University

Stefano Tessaro

University of Washington

Follow this and additional works at: https://pdxscholar.library.pdx.edu/compsci_fac



Part of the [Computer Sciences Commons](#)

Let us know how access to this document benefits you.

Citation Details

Published as: Jaeger, J., Song, F., & Tessaro, S. (2021, November). Quantum key-length extension. In Theory of Cryptography Conference (pp. 209-239). Springer, Cham.

This Pre-Print is brought to you for free and open access. It has been accepted for inclusion in Computer Science Faculty Publications and Presentations by an authorized administrator of PDXScholar. Please contact us if we can make this document more accessible: pdxscholar@pdx.edu.

Quantum Key-length Extension

Joseph Jaeger¹, Fang Song², and Stefano Tessaro³

¹ Georgia Institute of Technology, Atlanta, Georgia, US
 josephjaeger@gatech.edu

² Portland State University, Portland, Oregon, US
 fang.song@pdx.edu

³ Paul G. Allen School of Computer Science & Engineering
 University of Washington, Seattle, Washington, US
 tessaro@cs.washington.edu

Abstract. Should quantum computers become available, they will reduce the effective key length of basic secret-key primitives, such as blockciphers. To address this we will either need to use blockciphers with inherently longer keys or develop key-length extension techniques to amplify the security of a blockcipher to use longer keys.

We consider the latter approach and revisit the FX and double encryption constructions. Classically, FX was proven to be a secure key-length extension technique, while double encryption fails to be more secure than single encryption due to a meet-in-the-middle attack. In this work we provide positive results, with concrete and tight bounds, for the security of *both* of these constructions against quantum attackers in ideal models.

For FX, we consider a partially-quantum model, where the attacker has quantum access to the ideal primitive, but only classical access to FX. This is a natural model and also the strongest possible, since effective quantum attacks against FX exist in the fully-quantum model when quantum access is granted to both oracles. We provide two results for FX in this model. The first establishes the security of FX against non-adaptive attackers. The second establishes security against general adaptive attackers for a variant of FX using a random oracle in place of an ideal cipher. This result relies on the techniques of Zhandry (CRYPTO '19) for lazily sampling a quantum random oracle. An extension to perfectly lazily sampling a quantum random permutation, which would help resolve the adaptive security of standard FX, is an important but challenging open question. We introduce techniques for partially-quantum proofs without relying on analyzing the classical and quantum oracles separately, which is common in existing work. This may be of broader interest.

For double encryption, we show that it amplifies strong pseudorandom permutation security in the fully-quantum model, strengthening a known result in the weaker sense of key-recovery security. This is done by adapting a technique of Tessaro and Thiruvengadam (TCC '18) to reduce the security to the difficulty of solving the list disjointness problem and then showing its hardness via a chain of reductions to the known quantum difficulty of the element distinctness problem.

1 Introduction

The looming threat of quantum computers has inspired significant efforts to design and analyze post-quantum cryptographic schemes. In the public-key setting, polynomial-time quantum algorithms for factoring and computing discrete logarithms essentially break all practically deployed primitives [28].

In the secret-key setting, Grover's quantum search algorithm [12] will reduce the effective key length of secret-key primitives by half. Thus, a primitive like the AES-128 blockcipher which may be thought to have 128 bits of security against classical computers may provide no more than 64 bits of security against a quantum computer, which would be considered significantly lacking. Even more worrisome, it was shown relatively recent that quantum computers can break several secret-key constructions completely such as the Even-Mansour blockcipher [23] and CBC-MAC [17] if we grant the attacker fully quantum access to the cryptosystem.

This would not be the first time that we find ourselves using too short of a key. A similar issue had to be addressed when the DES blockcipher was widely used and its 56 bit keylength was considered insufficient.

Following approaches considered at that time, we can either transition to using basic primitives which have longer keys (e.g. replacing AES-128 with AES-256) or design key-length extension techniques to address the loss of concrete security due to quantum computers. In this paper we analyze the latter approach. We consider two key-length extension techniques, FX [20] and double encryption, and provide provable bounds against quantum attackers in *ideal* models.

Of broader and independent interest, our study of FX focuses on a hybrid quantum model which only allows for *classical* online access to the encrypted data, whereas offline computation is quantum. This model is sometimes referred to as the “Q1 model” in the cryptanalysis literature [17,5], in contrast to the fully-quantum, so-called “Q2 model”, which allows for quantum online access. This is necessary in view of existing attacks in the Q2 model showing that FX is no more secure than the underlying cipher [24], but *also*, Q1 is arguably more realistic and less controversial than Q2. We observe that (as opposed to the plain model) ideal-model proofs in the Q1 model can be *harder* than those in the Q2 model, as we need to explicitly account for *measuring* the online queries to obtain improved bounds. In many prior ideal-model Q1 proofs, e.g. [21,22,14,4,18], this interaction is handled essentially for free because the effects of online and offline queries on an attacker’s advantage can largely be analyzed separately. Our work introduces techniques to handle the interaction between classical online queries and quantum offline ideal-model queries in Q1 proofs that cannot be analyzed separately. On the other hand, our result on double encryption considers the full Q2 model – and interestingly, restricting adversaries to the Q1 model does not improve the bound. To be self-explanatory we will often refer to the Q1 and Q2 models as the partially-quantum and fully-quantum models, respectively.

The remainder of this introduction provides a detailed overview of our results for these two constructions, and of the underlying challenges and techniques.

1.1 The FX Construction

The FX construction was originally introduced by Kilian and Rogaway [20] as a generalization of Rivest’s DESX construction. Consider a blockcipher E which uses a key $K \in \{0, 1\}^k$ to encrypt messages $M \in \{0, 1\}^n$. Then the FX construction introduces “whitening” key $K_2 \in \{0, 1\}^n$ which is xor-ed into the input and output of the blockcipher. Formally, this construction is defined by $\text{FX}[E](K \parallel K_2, M) = E_K(M \oplus K_2) \oplus K_2$. (Note that the Even-Mansour blockcipher [11] may be considered to be a special case of this construction where $k = 0$, i.e., the blockcipher is a single permutation.) This construction has negligible efficiency overhead as compared to using E directly.

Kilian and Rogaway proved this scheme secure against classical attacks in the ideal cipher model. In particular, they established that

$$\text{Adv}_{\text{FX}}^{\text{SPRP}}(\mathcal{A}) \leq pq/2^{k+n-1}.$$

Here $\text{Adv}_{\text{FX}}^{\text{SPRP}}$ measures the advantage of \mathcal{A} in breaking the strong pseudorandom permutation (SPRP) security of FX while making at most p queries to the ideal cipher and at most q queries to the FX construction. Compared to the $p/2^k$ bound achieved by E alone, this is a clear improvement so FX can be considered a successful key-length extension technique against classical attackers.

Is this construction equally effective in the face of quantum attackers? The answer is unfortunately negative. Leander and May [24], inspired by a quantum attack due to Kuwakado and Morii [23] that completely breaks Even-Mansour blockcipher, gave a quantum attack against FX, which shows that the whitening keys provide essentially no additional security over that achieved by E in isolation. Bonnetain, et al. [5] further reduced the number of *online* quantum queries in the attack. Roughly speaking, $O(n)$ quantum queries to FX construction and $O(n2^{k/2})$ local quantum computations of the blockcipher suffice to recover the secret encryption key. Note that, however, such attacks require full quantum access to both the ideal primitive and to the instance FX that is under attack, i.e. they are attacks in the fully-quantum model. The latter is rather strong and may be considered unrealistic. While we cannot prevent a quantum attacker from locally

evaluating a blockcipher in quantum superposition, honest implementations of encryption will likely continue to be classical.⁴

PARTIALLY-QUANTUM MODEL. Because of the realistic concern of the fully-quantum model and the attacks therein that void key extension in FX, we turn to the partially-quantum model in which the attacker makes quantum queries to ideal primitives, but only *classical* queries to the cryptographic constructions.

In this model there has been extensive quantum cryptanalysis on FX and related constructions [13,5]. The best existing attack [5] recovers the key of FX using roughly $2^{(k+n)/3}$ classical queries to the construction and $2^{(k+n)/3}$ quantum queries to the ideal cipher. However, to date, despite the active development in provable quantum security, we are not aware of SPRP or just PRF security analysis, which gives stronger security guarantees. Namely, it should not just be infeasible to retrieve a key, but also to merely distinguish the system from a truly random permutation (or function). We note that in the special case where the primitives are plain-model instantiations (e.g., non-random-oracle hash functions), with a bit of care many security reductions carry over to the quantum setting [25]. This is because the underlying primitives are hidden from the adversary, and hence the difficulty arising from the interaction of classical and quantum queries to two correlated oracles becomes irrelevant.

Our main contribution on FX is to prove, for the first time, indistinguishability security in the partially-quantum model, in two restricted ways. Although they do not establish the complete security, our security bounds are tight in their respective settings.⁵

NON-ADAPTIVE SECURITY. We first consider non-adaptive security where we restrict the adversary such that its classical queries to the FX construction (but not to the underlying ideal cipher) must be specified before execution has begun. We emphasize that non-adaptive security of a blockcipher suffices to prove adaptive security for many practical uses of blockciphers such as the various randomized or stateful encryption schemes (e.g. those based on counter mode or output feedback mode) in which an attacker would have no control over the inputs to the blockcipher.

In this setting the bound we prove is of the form

$$\text{Adv}_{\text{FX}}^{\text{sprp-na}}(\mathcal{A}) \leq O\left(\sqrt{p^2 q / 2^{k+n}}\right).$$

Supposing $k = n = 128$ (as with AES-128), an attacker able to make $p \approx 2^{64}$ queries to the ideal cipher could break security of E in isolation. But to attack FX, such an attacker with access to $q \approx 2^{64}$ encryptions would need to make $p \approx 2^{96}$ queries to the ideal cipher. In fact we can see from our bound that breaking the security with constant probability would require the order of $\Omega(2^{(k+n)/3})$ queries in total, matching the bound given in the attacks mentioned above [5]. Hence our bound is tight.

To prove this bound we apply a one-way to hiding (O2H) theorem of Ambainis, Hamburg, and Unruh [3], an improved version of the original one in [32]. This result provides a clean methodology for bounding the probability that an attacker can distinguish between two functions drawn from closely related distributions given quantum access. The non-adaptive setting allows us to apply this result by sampling the outputs of FX ahead of time and then considering the ideal world in which the ideal cipher is chosen independently of these outputs and the real world in which we very carefully reprogram this ideal cipher to be consistent with the outputs chosen for FX. These two ideal ciphers differ only in the $O(q)$ places where we need to reprogram.

ADAPTIVE SECURITY OF FFX. As a second approach towards understanding fully adaptive security of FX, we consider a variant construction (which we call FFX for “function FX”) that replaces the random permutation with a random function. In particular, suppose F is a function family which uses a key $K \in \{0, 1\}^k$ on input

⁴ Some may argue that maintaining purely classical states, e.g., enforcing perfect measurements is also non-trivial physically. However, we deem maintaining coherent quantum superposition significantly more challenging.

⁵ Throughout, when mentioning tightness, we mean it with respect to the resources required to achieve advantage around one. The roots in our bounds make them weaker for lower resource regimes. Removing these is an interesting future direction.

messages $M \in \{0, 1\}^n$ to produce outputs $C \in \{0, 1\}^m$. Then we define $\text{FFX}[F](K \parallel K_2, M) = F_K(M \oplus K_2)$.⁶ For this construction we prove a bound of the form

$$\text{Adv}_{\text{FFX}}^{\text{prf}}(\mathcal{A}) \leq O\left(\sqrt{pq/2^{k+n}}\right).$$

in the partially-quantum random oracle model. Note that this matches the bound we obtained for the non-adaptive security of FX. Since the same key-recovery attack [5] also applies here, it follows that our bound is tight as well. Our proof combines two techniques of analyzing a quantum random oracle, the O2H theorem above and a simulation technique by Zhandry [35]. The two techniques usually serve distinct purposes. O2H is helpful to program a random oracle, whereas Zhandry’s technique is typically convenient for (compactly) maintaining a random oracle and providing some notion of “recording” the queries. In essence, in the two function distributions of O2H for which we aim to argue indistinguishability, we apply Zhandry’s technique to simulate the functions in a compact representation. As a result, analyzing the guessing game in O2H, which implies indistinguishability, becomes intuitive and much simplified. This way of combining them could also be useful elsewhere.

To build intuition for the approach of our proof, let us first consider one way to prove the security of this construction classically. The core idea is to use lazy sampling. In the ideal world, we can independently lazily sample a random function $F : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^m$ to respond to F queries and a random function $T : \{0, 1\}^n \rightarrow \{0, 1\}^m$ to respond to FFX queries. These lazily random functions are stored in tables.

The real world can similarly be modeled by lazily sampling F and T to respond to the separate oracles. However, these oracles need to be kept consistent. So if the adversary ever queries M to FFX and $(K, M \oplus K_2)$ to F, then the game should copy values between the two tables such that the same value is returned by both oracles. (Here K and K_2 are the keys honestly sampled by the game.) Alternatively, we can think of the return value being stored only in the T table when such queries occur (rather than being copied into both tables) as long as we remember that this has happened. When represented in this manner, the two games only differ if the adversary makes such a pair of queries, where we think of the latter one as being “bad”. Thus a simple $O(pq/2^{k+n})$ bound on the probability of making such a query bounds the advantage of the adversary.

In our quantum security proof we wish to proceed analogously. First, we find a way to represent the responses to oracle queries with two (superpositions over) tables that are independent in the ideal world and dependent in the real world (the dependency occurs only for particular “bad” inputs). Then (using the O2H theorem of Ambainis, Hamburg, and Unruh) we can bound the distinguishing advantage by the probability of an attacker finding a “bad” input. In applying this theorem we will jointly think of the security game and its adversary \mathcal{A} as a combined adversary \mathcal{A}' making queries to an oracle which takes in both the input of \mathcal{A} and the tables being stored by the game – processing them appropriately.

The required representation of the oracles via two tables is a highly non-trivial step in the quantum setting. For starters, the no-cloning theorem prevents us from simply recording queries made by the adversary. This has been a recurring source of difficulty for numerous prior papers such as [33,30,9]. We make use of the recent elegant techniques of Zhandry [35] which established that, by changing the perspective (e.g., to the Fourier domain), a random function can be represented by a table which is initialized to all zeros and then xor-ed into with each oracle query made by the adversary. This makes it straightforward to represent the ideal world as two separate tables. To represent the real world similarly, we exploit the fact that the queries to FFX are classical. To check if an input to FFX is “bad” we simply check if the corresponding entry of the random oracle’s table is non-zero. To check if an input to the random oracle is “bad” we check if it overlaps with prior queries to FFX which we were able to record because they were classical. For a “bad” input we then share the storage of the two tables and this is the only case where the behavior of the real world differs from that of the ideal world. These “bad” inputs may of course be part of a superposition query and it is only for the bad components of the superposition that the games differ.

⁶ Note we have removed the external xor with K_2 . In FX this xor is necessary, but in our analysis it would not provide any benefit for FFX.

DIFFICULTY OF EXTENDING TO FX. It is possible that this proof could be extended to work for normal FX given an analogous way to lazily represent a random permutation. Unfortunately, no such representation is known.

Czajkowski, et al. [8] extended Zhandry’s lazy sampling technique to a more general class of random functions, but this does not include permutations because of the correlation between the different outputs of a random permutation. Chevalier, et al. [6] provided a framework for recording queries to quantum oracles, which enables succinctly recording queries made to an externally provided function for purposes of later responding to inverse queries. This is distinct from the lazy sampling of a permutation that we require. Rosmanis [27] introduced a new technique for analyzing random permutations in a compressed manner and applied it to the question of inverting a permutation (given only forward access to it). Additional ideas seem needed to support actions based on the oracle queries that have been performed so far. This is essential in order to extend our proof for the function variant of FX to maintain consistency for the real world in the face of “bad” queries.

Recent work of Czajkowski [7] provided an *imperfect* lazy sampling technique for permutations and used it to prove indistinguishability of SHA3. They claim that their lazy sampling strategy cannot be distinguished from a random permutation with advantage better than $O(q^2/2^n)$. Unfortunately, this bound is too weak to be useful to our FX proof. For example, if $k \geq n$ we already have $O(q^2/2^n)$ security without key-length extension. Determining if it is possible to *perfectly* lazily sample a random permutation remains an interesting future direction.

1.2 Double Encryption

The other key-extension technique we consider is double encryption. Given a blockcipher $E : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ this is defined by $DE[E](K_1 \parallel K_2, M) = E_{K_2}(E_{K_1}(M))$. This construction requires more computational overhead than FX because it requires two separate application of the blockcipher with different keys. Classically, this construction is not considered to be a successful key-length extension technique because the meet-in-the-middle attack [10,26] shows that it can be broken in essentially the same amount of time as E alone.

However, this does not rule out that double encryption is an effective key-length extension method in the quantum setting, as it is not clear that the Grover search algorithm [12] used to halve the effective keylength of blockciphers can be composed with the meet-in-the middle attack to unify their savings. The security of double encryption in the quantum setting was previously considered by Kaplan [16]. They related the key-recovery problem in double encryption to the claw-finding problem, and gave the tight quantum query bound $\Theta(N^{2/3})$ for solving key recovery (here $N = 2^k$ is the length of the lists in the claw-finding problem). This indicates that in the quantum setting double encryption is in fact useful (compare to $N^{1/2}$), although key-recovery security is fairly weak.

We strengthen their security result by proving the SPRP security, further confirming double encryption as an effective key-extension scheme against quantum attacks. This is proven in the *fully-quantum* model, and the bound we obtain matches the attack in [16] which works in the partially-quantum model. Namely restricting to the weaker partially-quantum model would not improve the bound. Our result is obtained by a reduction to list disjointness. This is a worst-case decision problem measuring how well an algorithm can distinguish between a pair of lists with zero or *exactly* one element in common, which can be viewed as a decision version of the claw-finding problem. This reduction technique was originally used by Tessaro and Thiruvengadam [31] to establish a classical time-memory trade-off for double encryption. We observe that their technique works for a quantum adversary.

We then construct a chain of reductions to show that the known quantum hardness of element distinctness [1,34] (deciding if a list of N elements are all distinct) can be used to establish the quantum hardness of solving list disjointness. Our result (ignoring log factors) implies that a highly successful attacker must make $\Omega(2^{2k/3})$ oracle queries which is more than the $\Omega(2^{k/2})$ queries needed to attack E used in isolation.

Our proof starts by observing that Zhandry’s [34] proof of the hardness of the search version of element distinctness (finding a collision in a list) in fact implies that a promise version of element distinctness (promising that there is exactly one collision) is also hard. Then a simple reduction (randomly splitting the

element distinctness list into two lists) shows the hardness of the search version of list disjointness. Next we provide a binary-search inspired algorithm showing that the decision version of list disjointness can be used to solve the search version, implying that the decision version must be hard. During our binary search we pad the lists we are considering with random elements to ensure that our lists maintain a fixed size which is necessary for our proof to go through.

The final bound we obtain for double encryption is of the form

$$\text{Adv}_{\text{DE}}^{\text{SPRP}}(\mathcal{A}) \leq O\left(\sqrt[6]{(q \cdot k \lg k)^3 / 2^{2k}}\right).$$

The sixth root arises in this bound from the final step in our chain of results analyzing list disjointness. The binary search algorithm requires its underlying decision list disjointness algorithm to have relatively high advantage. To obtain this from a given algorithm with advantage δ we need to amplify its advantage by running in on the order of $1/\delta^2$ times. The number of queries depending on the square of δ causes the root to arise in the proof.

1.3 Overview

In Section 2, we introduce preliminaries such as notation, basic cryptographic definitions, and some background on quantum computation that we will use throughout the paper. Following this, in Section 3 we consider the security of FX in the partially quantum setting. Non-adaptive SPRP security of FX is proven in Section 3.1 and adaptive PRF security of FFX is proven in Section 3.2. We conclude with Section 4 in which we prove the SPRP security of double encryption against fully quantum adaptive attacks.

2 Preliminaries

For $n, m \in \mathbb{N}$, we let $[n] = \{1, \dots, n\}$ and $[n..m] = \{n, n+1, \dots, m\}$. The set of length n bit strings is denoted $\{0, 1\}^n$. We use \parallel to denote string concatenation. We let $\text{Inj}(n, m)$ denote the set of injections $f : [n] \rightarrow [m]$.

We let $y \leftarrow^s \mathcal{A}[O_1, \dots](x_1, \dots)$ denote the (randomized) execution of algorithm \mathcal{A} with input x_1, \dots and oracle access to O_1, \dots which produces output y . For different \mathcal{A} we will specify whether it can access its oracles in quantum superposition or only classically. If \mathcal{S} is a set, then $y \leftarrow^s \mathcal{S}$ denotes randomly sampling y from \mathcal{S} .

We express security notions via pseudocode games. See Fig. 1 for some example games. In the definition of games, oracles will sometimes be specified by pseudocode with the following form.

```

Oracle O( $X_1, \dots : Z_1, \dots$ )
//Code defining  $X'_1, \dots$  and  $Z'_1, \dots$ 
Return ( $X'_1, \dots : Z'_1, \dots$ )

```

This notation indicates that X_1, \dots are variables controlled by the adversary prior to the oracle query and Z_1, \dots are variables controlled by the game itself which the adversary cannot access. At the end of the execution of the oracle, these variables are overwritten with the values indicated in the return statement. Looking ahead, we will be focusing on quantum computation so this notation will be useful to make it explicit that O can be interpreted as a unitary acting on the registers X_1, \dots and Z_1, \dots (because O will be an efficiently computable and invertible permutation over these values). If \mathbf{H} is a function stored by the game, then oracle access to \mathbf{H} represents access to the oracle that on input $(X, Y : \mathbf{H})$ returns $(X, \mathbf{H}(X) \oplus Y : \mathbf{H})$.

We define games as outputting boolean values and let $\text{Pr}[\mathbf{G}]$ denote the probability that game \mathbf{G} returns true. When not otherwise indicated, variables are implicitly initialized to store all 0's.

If \mathcal{A} is an adversary expecting access to multiple oracles we say that it is *order consistent* if the order it will alternate between queries to these different oracles is a priori fixed before execution. Note that order consistency is immediate if, e.g., \mathcal{A} is represented by a circuit where each oracle is modeled by a separate oracle gate, but is not immediate for other possible representations of an adversary.

Game $G_{F,b}^{\text{prf}}(\mathcal{A})$	$\text{Ev}(X, Y : \mathbf{H}, K, F)$
$\mathbf{H} \leftarrow_{\$} \text{Fcs}(k, n, m)$	$Y_1 \leftarrow \mathbf{F}[\mathbf{H}](K, X)$
$K \leftarrow_{\$} \{0, 1\}^{F.\text{kl}}$	$Y_0 \leftarrow F(X)$
$F \leftarrow_{\$} \text{Fcs}(0, F.\text{il}, F.\text{ol})$	Return $(X, Y_b \oplus Y : \mathbf{H}, K, F)$
$b' \leftarrow_{\$} \mathcal{A}[\text{Ev}, \mathbf{H}]$	
Return $b' = 1$	

Game $G_{E,b}^{\text{sprp}}(\mathcal{A})$	$\text{Ev}(X, Y : \mathbf{E}, K, P)$
$\mathbf{E} \leftarrow_{\$} \text{lcs}(k, n)$	$Y_1 \leftarrow \mathbf{E}[\mathbf{E}](K, X)$
$K \leftarrow_{\$} \{0, 1\}^{E.\text{kl}}$	$Y_0 \leftarrow P(X)$
$P \leftarrow_{\$} \text{lcs}(0, \mathbf{E}.\text{bl})$	Return $(X, Y_b \oplus Y : \mathbf{E}, K, P)$
$b' \leftarrow_{\$} \mathcal{A}[\text{Ev}, \text{Inv}, \mathbf{E}, \mathbf{E}^{-1}]$	$\text{Inv}(X, Y : \mathbf{E}, K, P)$
Return $b' = 1$	$Y_1 \leftarrow \mathbf{E}^{-1}[\mathbf{E}](K, X)$
	$Y_0 \leftarrow P^{-1}(X)$
	Return $(X, Y_b \oplus Y : \mathbf{E}, K, P)$

Fig. 1. Security games measuring PRF security of a family of functions \mathbf{F} and SPRP security of a blockcipher \mathbf{E} .

IDEAL MODELS. In this work we will work in ideal models – specifically, the random oracle model or the ideal cipher model. Fix $k, n, m \in \mathbb{N}$ (throughout this paper we will treat these parameters as having been fixed already). We let $\text{Fcs}(k, n, m)$ be the set of all functions $\mathbf{H} : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^m$ and $\text{lcs}(k, n) \subset \text{Fcs}(k, n, n)$ be the set of all functions $\mathbf{E} : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ such that $\mathbf{E}(K, \cdot)$ is a permutation on $\{0, 1\}^n$. When convenient, we will write $\mathbf{H}_K(x)$ in place of $\mathbf{H}(K, x)$ for $\mathbf{H} \in \text{Fcs}(k, n, m)$. Similarly, we will write $\mathbf{E}_K(x)$ for $\mathbf{E}(K, x)$ and $\mathbf{E}_K^{-1}(\cdot)$ for the inverse of $\mathbf{E}_K(\cdot)$ when $\mathbf{E} \in \text{lcs}(k, n)$. When $K = \varepsilon$ we omit the subscript to \mathbf{H} or \mathbf{E} .

In the random oracle model, honest algorithms and the adversary are given oracle access to a randomly chosen $\mathbf{H} \in \text{Fcs}(k, n, m)$. In the ideal cipher model, they are given oracle access to \mathbf{E} and \mathbf{E}^{-1} for \mathbf{E} chosen at random from $\text{lcs}(k, n)$. We refer to queries to these oracles as *primitive queries* and queries to all other oracles as *construction queries*.

FUNCTION FAMILY AND PSEUDORANDOMNESS. A function family \mathbf{F} is an efficiently computable element of $\text{Fcs}(F.\text{kl}, F.\text{il}, F.\text{ol})$. If, furthermore, $\mathbf{F} \in \text{lcs}(F.\text{kl}, F.\text{il})$ and \mathbf{F}^{-1} is efficiently computable then we say \mathbf{F} is a blockcipher and let $F.\text{bl} = F.\text{il}$.

If \mathbf{F} is a function family (constructed using oracle access to a function $\mathbf{H} \in \text{Fcs}(k, n, m)$), then its security (in the random oracle model) as a *pseudorandom function* (PRF) is measured by the game G^{prf} shown in Fig. 1. In it, the adversary \mathcal{A} attempts to distinguish between a *real world* ($b = 1$) where it is given oracle access to \mathbf{F} with a random key K and an *ideal world* ($b = 0$) where it is given access to a random function. We define the advantage function $\text{Adv}_{\mathbf{F}}^{\text{prf}}(\mathcal{A}) = \Pr[G_{\mathbf{F},1}^{\text{prf}}(\mathcal{A})] - \Pr[G_{\mathbf{F},0}^{\text{prf}}(\mathcal{A})]$.

If \mathbf{E} is a blockcipher (constructed using oracle access to a function $\mathbf{E} \in \text{lcs}(k, n)$ and its inverse), then its security (in the ideal cipher model) as a *strong pseudorandom permutation* (SPRP) is measured by the game G^{sprp} shown in Fig. 1. In it, the adversary \mathcal{A} attempts to distinguish between a *real world* ($b = 1$) where it is given oracle access to $\mathbf{E}, \mathbf{E}^{-1}$ with a random key K and an *ideal world* ($b = 0$) where it is given access to a random permutation. We define the advantage function $\text{Adv}_{\mathbf{F}}^{\text{sprp}}(\mathcal{A}) = \Pr[G_{\mathbf{F},1}^{\text{sprp}}(\mathcal{A})] - \Pr[G_{\mathbf{F},0}^{\text{sprp}}(\mathcal{A})]$.

In some examples, we will restrict attention to *non-adaptive* SPRP security. In such cases our attention is restricted to attackers whose queries to EV and INV when relevant are a priori fixed before execution. That is, \mathcal{A} is a non-adaptive attacker which makes at most q classical, non-adaptive queries to EV, INV if there exists $M_1, \dots, M_{q'}, Y_{q'+1}, \dots, Y_q \in \{0, 1\}^n$ such that \mathcal{A} only ever queries EV on M_i for $1 \leq i \leq q'$ and INV on Y_i for $q' + 1 \leq i \leq q$. Then we write $\text{Adv}_{\mathbf{F}}^{\text{sprp-na}}(\mathcal{A})$ in place of $\text{Adv}_{\mathbf{F}}^{\text{sprp}}(\mathcal{A})$.

Game $\mathsf{G}_{\mathbf{D},b}^{\text{dist}}(\mathcal{A})$	Game $\mathsf{G}_{\mathbf{D}}^{\text{guess}}(\mathcal{A})$
$(S, S', P_0, P'_0, P_1, P'_1, z) \leftarrow \mathbf{D}$	$(S, S', P_0, P'_0, P_1, P'_1, z) \leftarrow \mathbf{D}$
$b' \leftarrow \mathcal{A}[P_b, P'_b](z)$	$i \leftarrow \{1, \dots, q\}$
Return $b' = 1$	Run $\mathcal{A}[P_0, P'_0]$ until its i -th query
	Measure the input x to this query
	If the query is to P_0 then
	Return $x \in S$
	Else (the query is to P'_0)
	Return $x \in S'$

Fig. 2. Games used for O2H Theorem 1.

2.1 Quantum Background

We assume the reader has basic familiarity with quantum computation. Quantum computation proceeds by performing unitary operations on registers which each contain a fixed number of qubits. We sometimes use \circ to denote composition of unitaries. Additionally, qubits may be measured in the computational basis. We will typically use the principle of deferred measurements to without loss of generality think of such measurements as being deferred until the end of computation.

The Hadamard transform \mathcal{H} acts on a bitstring $x \in \{0, 1\}^n$ (for some $n \in \mathbb{N}$) via $\mathcal{H}|x\rangle = 1/\sqrt{2^n} \cdot \sum_{x'} (-1)^{x \cdot x'} |x'\rangle$. Here \cdot denotes inner product modulo 2 and the summation is over $x' \in \{0, 1\}^n$. The Hadamard transform is its own inverse. We sometimes use the notation $\mathcal{H}^{X_1, X_2, \dots}$ to denote the Hadamard transform applied to registers X_1, X_2, \dots .

We make use of the fact that if P is a permutation for which both P and P^{-1} can be efficiently implemented classically, then there is a comparable efficient quantumly computable unitary U_P which maps according to $U_P|x\rangle = |P(x)\rangle$ for $x \in \{0, 1\}^n$. For simplicity, we often write P in place of U_P . If $f : \{0, 1\}^m \rightarrow \{0, 1\}^m$ is a function, we define the permutation $f[\oplus](x, y) = (x, f(x) \oplus y)$.

ONE-WAY TO HIDING. We will make use of (a slight variant of) a one-way to hiding (O2H) theorem of Ambainis, Hamburg, and Unruh [3]. The theorem will consider an adversary given oracle access either to permutations (P_0, P'_0) or permutations (P_1, P'_1) . It relates the advantage of the adversary in distinguishing between these two cases to the probability that the adversary can be used to find one of those points on which P_0 differs from P_1 or P'_0 differs from P'_1 . The result considers a distribution \mathbf{D} over $(S, S', P_0, P'_0, P_1, P'_1, z)$ where S, S' are sets, P_0, P_1 are permutations on the same domain, P'_0, P'_1 are permutations on the same domain, and $z \in \{0, 1\}^*$ is some auxiliary information. Such a \mathbf{D} is *valid* if $P_0(x) = P_1(x)$ for all $x \notin S$ and $P'_0(x) = P'_1(x)$ for all $x \notin S'$. Now consider the game $\mathsf{G}_{\mathbf{D},b}^{\text{dist}}$ shown in Fig. 2. In it, an adversary \mathcal{A} is given z and tries to determine which of the oracle pairs it has access to. We define $\text{Adv}_{\mathbf{D}}^{\text{dist}}(\mathcal{A}) = \Pr[\mathsf{G}_{\mathbf{D},1}^{\text{dist}}(\mathcal{A})] - \Pr[\mathsf{G}_{\mathbf{D},0}^{\text{dist}}(\mathcal{A})]$.

The game $\mathsf{G}_{\mathbf{D}}^{\text{guess}}(\mathcal{A})$ in the same figure measures the ability of \mathcal{A} to query its oracles on inputs at which P_0 and P_1 (or P'_0 and P'_1) differ. It assumes that the adversary makes at most q oracle queries. The adversary is halted in its execution on making a random one of these queries and the input to this query is measured. If the input falls in the appropriate set S or S' , then the game returns `true`. Thus we can roughly think of this as a game in which \mathcal{A} is trying to guess a point on which the two oracles differ. We define $\text{Adv}_{\mathbf{D}}^{\text{guess}}(\mathcal{A}) = \Pr[\mathsf{G}_{\mathbf{D}}^{\text{guess}}(\mathcal{A})]$, which leads to a bound on $\text{Adv}_{\mathbf{D}}^{\text{dist}}(\mathcal{A})$.

Theorem 1 ([3], Thm.3). *Let \mathbf{D} be a valid distribution and \mathcal{A} be an adversary making at most q oracle queries. Then $\text{Adv}_{\mathbf{D}}^{\text{dist}}(\mathcal{A}) \leq 2q\sqrt{\text{Adv}_{\mathbf{D}}^{\text{guess}}(\mathcal{A})}$.*

Our statement of the theorem differs from the result as given in [3] in that we consider arbitrary permutations, rather than permutations of the form $f[\oplus]$ for some function f , and we provide the attacker

with access to two oracles rather than one.⁷ These are simply notational conveniences to match how we will be applying the theorem. The proof given in [3] suffices to establish this variant without requiring any meaningful modifications.

The most natural applications of this theorem would apply it to distributions \mathbf{D} for which the guessing advantage $\text{Adv}_{\mathbf{D}}^{\text{guess}}(\mathcal{A})$ is small for *any* efficient adversary \mathcal{A} . This will indeed be the case for our use of it in our Theorem 2. However, note that it can also be applied more broadly with a distribution \mathbf{D} where it is not necessarily difficult to guess inputs on which the oracles differ. We will do so at the end of our proof of Theorem 3. Here we will use a *deterministic* \mathbf{D} so, in particular, the sets S and S' are a priori fixed and not hard to query. The trick we will use to profitably apply the O2H result is to exploit knowledge of the particular form that \mathcal{A} will take (it will be a reduction adversary internally simulating the view of another adversary) to provide a useful bound on its guessing advantage $\text{Adv}_{\mathbf{D}}^{\text{guess}}(\mathcal{A})$.

3 The FX Construction

The FX construction (originally introduced by Kilian and Rogaway [20] as a generalization of Rivest’s DESX construction) is a keylength extension for blockciphers. In this construction, an additional key is used which is xor-ed with input and the output of the blockcipher.⁸ Formally, given a blockcipher $E \in \text{Ics}(E.\text{kl}, E.\text{bl})$, the blockcipher $\text{FX}[E]$ is defined by $\text{FX}[E](K_1 \parallel K_2, x) = E_{K_1}(x \oplus K_2) \oplus K_2$. Here $|K_1| = E.\text{kl}$ and $|K_2| = E.\text{bl}$ so $\text{FX}[E].\text{kl} = E.\text{kl} + E.\text{bl}$ and $\text{FX}[E].\text{bl} = E.\text{bl}$. Its inverse can similarly be computed as $\text{FX}[E]^{-1}(K_1 \parallel K_2, x) = E_{K_1}^{-1}(x \oplus K_2) \oplus K_2$. Let $k = E.\text{kl}$ and $n = E.\text{bl}$.

Kilian and Rogaway [19] analyzed the PRP security of FX against classical attacks, showing that $\text{Adv}_{\text{FX}}^{\text{sprp}}(\mathcal{A}) \leq 2pq/2^{k+n}$ where q is the number of E, INV queries and p is the number of E, E^{-1} queries made by \mathcal{A} (with E modeled as an ideal cipher). In [24], Leander and May showed a quantum attack against the FX construction – establishing that the added whitening keys did not provide additionally security. This attack uses a clever combination of the quantum algorithms of Grover [12] and Simon [29]. It was inspired by an attack by Kuwakado and Morii [23] showing that the Even-Mansour blockcipher [11] provides no quantum security. Thus, it seems that $\text{FX}[E]$ does not provide meaningfully more security than E against quantum attackers.

However, the attack of Leander and May requires quantum access to both the FX construction and the underlying blockcipher E . This raises the question of whether the FX is actually an effective key-length extension technique in the partially-quantum setting where the adversary performs only classical queries to the construction oracles. In this section, we approach this question from two directions. First, in Section 3.1 we apply Theorem 1 with a careful representation of the real and ideal worlds to show that FX does indeed achieve improved security against non-adaptive attacks.

Analyzing the full adaptive security of FX against classical construction queries seems beyond the capabilities of current proof techniques. Accordingly, in Section 3.2, we consider a variant of FX in which a random oracle is used in place of the ideal cipher and prove its quantum PRF security. Here we apply a new reduction technique (built on the “sparse” quantum representation of a random function introduced by Zhandry [35] and Theorem 1, the O2H theorem from Ambainis, Hamburg, and Unruh [3]) to prove that this serves as an effective key-length extension technique in our setting. It seems likely that our technique could be extended to the normal FX construction, should an appropriate sparse quantum representation of random permutations be discovered.

3.1 Security of FX Against Non-Adaptive Attacks

The following theorem bounds the security of the FX construction against non-adaptive attacks (in which the non-adaptive queries are all classical). This result is proven via a careful use of Theorem 1 in which the

⁷ Their result additionally allows the adversary to make oracle queries in parallel and bounds its advantage in terms of the “depth” of its oracle queries rather than the total number of queries. We omit this for simplicity.

⁸ Technically, the original definition of FX [20] uses distinct keys for xor-ing with the input and the output, but this would not provide any benefit in our concrete security analysis so we focus on the simplified construction.

```

Distribution D
// Step 1: Sample responses to construction queries
For  $i = 1, \dots, q'$  do
   $Y_i \leftarrow \{0, 1\}^n \setminus \{Y_1, \dots, Y_{i-1}\}$ 
   $T[M_i] \leftarrow Y_i; T^{-1}[Y_i] \leftarrow M_i$ 
For  $i = q' + 1, \dots, q$  do
  If  $T^{-1}[Y_i] \neq \perp$  then  $M_i \leftarrow T^{-1}[Y_i]$ 
  Else  $M_i \leftarrow \{0, 1\}^n \setminus \{M_1, \dots, M_{i-1}\}$ 
   $T[M_i] \leftarrow Y_i; T^{-1}[Y_i] \leftarrow M_i$ 
 $z \leftarrow (T, T^{-1})$ 
// Step 2: Sample  $f_0$  as independent ideal cipher
 $f_0 \leftarrow \text{lcs}(k, n)$ 
// Step 3: Reprogram  $f_1$  for consistency with construction queries
 $K_1 \leftarrow \{0, 1\}^k; K_2 \leftarrow \{0, 1\}^n$ 
 $\mathcal{I} \leftarrow \{M_i \oplus K_2 : 1 \leq i \leq q\}; \mathcal{O} \leftarrow \{Y_i \oplus K_2 : 1 \leq i \leq q\}$ 
 $\mathcal{I}' \leftarrow \{f_0^{-1}(K_1, y) : y \in \mathcal{O}\}; \mathcal{O}' \leftarrow \{f_0(K_1, x) : x \in \mathcal{I}\}$ 
 $S = \{(K_1, x) : x \in \mathcal{I} \cup \mathcal{I}'\}; S' = \{(K_1, y) : y \in \mathcal{O} \cup \mathcal{O}'\}$ 
For  $(K, x) \notin S$  do
   $f_1(K, x) \leftarrow f_0(K, x)$ 
For  $i = 1, \dots, q$  do
   $f_1(K_1, M_i \oplus K_2) \leftarrow Y_i \oplus K_2$ 
For  $x \in \mathcal{I}' \setminus \mathcal{I}$  do
   $f_1(K_1, x) \leftarrow \mathcal{O}' \setminus \{f_1(K_1, x) : x \in \mathcal{I} \cup \mathcal{I}', f_1(K_1, x) \neq \perp\}$ 
Return  $(S, S', f_0[\oplus], f_0^{-1}[\oplus], f_1[\oplus], f_1^{-1}[\oplus], z)$ 

```

Fig. 3. Distribution of oracles used in proof of Theorem 2.

distribution **D** is defined in terms of the non-adaptive queries that the adversary will make and defined so as to perfectly match the two worlds that \mathcal{A} is attempting to distinguish between.

Theorem 2. *Let \mathcal{A} be a quantum adversary which makes at most q classical, non-adaptive queries to EV , INV and consider $\text{FX}[\cdot]$ with the underlying blockcipher modeled by an ideal cipher drawn from $\text{lcs}(k, n)$. Then*

$$\text{Adv}_{\text{FX}}^{\text{sprp-na}}(\mathcal{A}) \leq \sqrt{8p^2q/2^{k+n}},$$

where p is the number of quantum oracle queries that \mathcal{A} makes to the ideal cipher.

Proof. We will use Theorem 1 to prove this result, so first we define a distribution **D**. Suppose that $M_1, \dots, M_{q'} \in \{0, 1\}^n$ are the distinct queries \mathcal{A} will make to EV and $Y_{q'+1}, \dots, Y_q \in \{0, 1\}^n$ are the distinct queries that \mathcal{A} will make to INV . The order in which these queries will be made does not matter. Then we define **D** as shown in Fig. 3. This distribution is valid (as required for Theorem 1) because G_1 is reprogrammed to differ from G_0 by making inputs in S map to different values in S' .

We will show that the oracles output by this distribution (described in words momentarily) can be used to perfectly simulate the views expected by \mathcal{A} . In particular, let \mathcal{A}' be an adversary (for $\mathbb{G}_{\mathbf{D}}^{\text{dist}}$) which runs \mathcal{A} , responding to $\text{EV}(M_i)$ queries with $T[M_i]$, responding to $\text{INV}(Y_i)$ queries with $T^{-1}[Y_i]$, and simulating $\mathbf{E}, \mathbf{E}^{-1}$ with its own oracles $f_b[\oplus], f_b^{-1}[\oplus]$. When \mathcal{A} halts with output b , this adversary halts with the same output. We claim that (i) $\Pr[\mathbb{G}_{\mathbf{F},1}^{\text{sprp}}(\mathcal{A})] = \Pr[\mathbb{G}_{\mathbf{D},1}^{\text{dist}}(\mathcal{A}')]$ and (ii) $\Pr[\mathbb{G}_{\mathbf{F},0}^{\text{sprp}}(\mathcal{A})] = \Pr[\mathbb{G}_{\mathbf{D},0}^{\text{dist}}(\mathcal{A}')]$. This gives $\text{Adv}_{\mathbf{F}}^{\text{sprp-na}}(\mathcal{A}) = \text{Adv}_{\mathbf{D}}^{\text{dist}}(\mathcal{A}')$.

Claim (ii) follows by noting that the view of \mathcal{A} is identical when run by $\mathbb{G}_0^{\text{sprp}}$ or by \mathcal{A}' in $\mathbb{G}_{\mathbf{D},0}^{\text{dist}}$. In $\mathbb{G}_0^{\text{sprp}}$, its construction queries are answered with the random permutation F . When it is run by \mathcal{A}' in $\mathbb{G}_{\mathbf{D},0}^{\text{dist}}$, these queries are answered with the tables T and T^{-1} which can be viewed as having just lazily sampled enough of

a random permutation to respond to the given queries. In both cases, its primitive oracle is an independently chosen ideal cipher.

Claim (i), follows by noting that the view of \mathcal{A} is identical when run by G_1^{sprp} or by \mathcal{A}' in $G_{D,1}^{\text{dist}}$. In G_1^{sprp} , construction queries are answered by the FX construction using the ideal cipher and keys K_1, K_2 . In the distribution, we first sample the responses to the construction queries and then construct the ideal cipher f_1 by picking K_1 and K_2 and setting f_1 to equal f_0 except for places where we reprogram it to be consistent with these construction queries. The construction will map M_i to Y_i for each i , which means that the condition $f_1(K_1, M_i \oplus K_2) = Y_i \oplus K_2$ should hold. These inputs and outputs are stored in the sets \mathcal{I} and \mathcal{O} . The sets \mathcal{I}' and \mathcal{O}' store the inputs mapping to \mathcal{O} and outputs mapped to \mathcal{I} by $f_0(K_1, \cdot)$, respectively. Thus while making the above condition hold, we additionally reprogram f_1 so that elements of $\mathcal{I}' \setminus \mathcal{I}$ map to (random, non-repeating) elements of $\mathcal{O}' \setminus \mathcal{O}$.

In particular, the uniformity of f_0 ensures that the map induced by $f_1(K_1, \cdot)$ between $\{0, 1\}^n \setminus (\mathcal{I} \cup \mathcal{I}')$ and $\{0, 1\}^n \setminus (\mathcal{O} \cup \mathcal{O}')$ is a random bijection. The last for loop samples a random bijection between $\mathcal{I}' \setminus \mathcal{I}$ and $\mathcal{O}' \setminus \mathcal{O}$. Because there are no biases in which values fall into these two cases among those of $\{0, 1\}^n \setminus \mathcal{I}$ and $\{0, 1\}^n \setminus \mathcal{O}$, this means the map between these two sets is a uniform bijection as desired. A more detailed probability analysis of Claim (i) is given in Appendix A.

Applying the bound on $\text{Adv}_D^{\text{dist}}(\mathcal{A}')$ from Theorem 1 gives us

$$\text{Adv}_F^{\text{sprp-na}}(\mathcal{A}) \leq 2p \sqrt{\Pr[G_D^{\text{guess}}(\mathcal{A}')]}$$

so we complete the proof by bounding this probability. Let fw denote the event that the i -th query of \mathcal{A}' in H_0^D is to f_0 and let (K, x) denote the measured value of this query so that

$$\Pr[H_D^{\text{guess}}(\mathcal{A}')] = \Pr[\text{fw}] \cdot \Pr[(K, x) \in S \mid \text{fw}] + \Pr[\neg \text{fw}] \cdot \Pr[(K, x) \in S' \mid \neg \text{fw}].$$

A union bound over the different elements of S gives

$$\begin{aligned} \Pr[(K, x) \in S \mid \text{fw}] &\leq \sum_{j=1}^q \Pr[K = K_1 \wedge x \oplus M_j = K_2 \mid \text{fw}] \\ &\quad + \sum_{j=1}^q \Pr[K = K_1 \wedge G_0(K_1, x) \oplus Y_j = K_2 \mid \text{fw}]. \end{aligned}$$

However, note that the view of \mathcal{A} in H_0^D is independent of K_1 and K_2 so we get that

$$\Pr[(K, x) \in S \mid \text{fw}] \leq 2q/2^{k+n}.$$

Applying analogous analysis to the $\neg \text{fw}$ case gives

$$\Pr[(K, x) \in S' \mid \neg \text{fw}] \leq 2q/2^{k+n}$$

and hence $\Pr[H_0^D(\mathcal{A}')] \leq 2q/2^{k+n}$. Plugging this into our earlier inequality gives the stated bound. \square

3.2 Adaptive Security of FFX

In this section we will prove the security of FFX (a variant of FX using a random oracle in place of the ideal cipher) against quantum adversaries making strictly classical queries to the construction.

Formally, given a function family $F \in \text{Fcs}(F.\text{kl}, F.\text{il}, F.\text{ol})$, we model the FFX construction by the function family $\text{FFX}[F]$ by $\text{FFX}[F](K_1 \parallel K_2, x) = F_{K_1}(x \oplus K_2)$.⁹ Here $|K_1| = F.\text{kl}$ and $|K_2| = F.\text{il}$ so $\text{FFX}[F].\text{kl} = F.\text{kl} + F.\text{il}$, $\text{FFX}[F].\text{il} = F.\text{il}$, and $\text{FFX}[F].\text{ol} = F.\text{ol}$. Let $k = F.\text{kl}$, $n = F.\text{il}$, and $m = F.\text{ol}$.

⁹ The outer xor by K_2 used in FX is omitted because it is unnecessary for our analysis.

Theorem 3. *Let \mathcal{A} be an order consistent quantum adversary which makes classical queries to Ev and consider $\text{FFX}[\cdot]$ with the underlying function family modeled by a random oracle drawn from $\text{Fcs}(k, n, m)$. Then*

$$\text{Adv}_{\text{FFX}}^{\text{prf}}(\mathcal{A}) \leq \sqrt{\frac{8(p+q)pq}{2^{k+n}}},$$

where p is the number of quantum oracle queries that \mathcal{A} makes to the random oracle and q is the number of queries it makes to Ev.

We can reasonably assume that $p > q$ so the dominant behavior of the above expression is $O\left(\sqrt{p^2q/2^{k+n}}\right)$.

The proof of this result proceeds via a sequence of hybrids which gradually transition from the real world of $\text{G}_{\text{FFX}}^{\text{prf}}$ to the ideal world. Crucial to this sequence of hybrids are the technique of Zhandry [35] which, by viewing a space under dual bases, allows one to simulate a random function using a sparse representation table and to “record” the queries to the function. For the ideal world, we can represent the random oracle and the random function underlying Ev independently using such sparse representation tables. With some careful modification, we are also able to represent the real world’s random oracle using a similar pair of sparse representation tables as if it were two separate functions. However, in this case, the tables will be slightly non-independent in that if the adversary queries Ev on an input x and the random oracle on $(K_1, x \oplus K_2)$ then the results of the latter query is stored in the Ev table, rather than the random oracle table. Beyond this minor consistency check (which we are only able to implement because the queries to Ev are classical and so can be stored by simulation), the corresponding games are identical. Having done this rewriting, we can carefully apply Theorem 1 to bound the ability of an adversary to distinguish between the two worlds by its ability to trigger this consistency check.

As mentioned in Section 2.1, our application of Theorem 1 here is somewhat atypical. Our distribution over functions \mathbf{D} will be deterministic, but we are able to still extract a meaningful bound from this by taking advantage of our knowledge of the particular behavior of the adversary we apply Theorem 1 with.

Proof. In this proof we will consider a sequence of hybrid games H_0 through H_9 . Of these games we will establish the following claims.

1. $\Pr\left[\text{G}_{\text{FFX},0}^{\text{prf}}(\mathcal{A})\right] = \Pr[H_0] = \Pr[H_1] = \Pr[H_2] = \Pr[H_3]$
2. $\Pr\left[\text{G}_{\text{FFX},1}^{\text{prf}}(\mathcal{A})\right] = \Pr[H_9] = \Pr[H_8] = \Pr[H_7] = \Pr[H_6] = \Pr[H_5] = \Pr[H_4]$
3. $\Pr[H_4] - \Pr[H_3] \leq \sqrt{8(p+q)pq/2^{\text{F.kl}+\text{F.il}}}$

Combining these claims gives the desired result.

In formally defining our hybrids we write the computation to be performed using the following quantum registers.

- W : The workspace of \mathcal{A} . The adversary’s final output is written into $W[1]$.
- K : The k -qubit register (representing the function *key/index*) \mathcal{A} uses when making oracle queries to the random oracle.
- X : The n -qubit register (representing function inputs) used when making oracle queries to the random oracle or Ev.
- Y : The m -qubit register (representing function outputs) into which the results of oracle queries are written.
- H : The $2^{k+n} \cdot m$ -qubit register which stores the function defining the random oracle (initially via its truth table).
- F : The $2^n \cdot m$ -qubit register which stores the function defining Ev.
- K_1 : The k -qubit register which stores the first key of the construction.
- K_2 : The n -qubit register which stores the second key of the construction.
- I : The $\lceil \log p \rceil$ -qubit register which tracks how many Ev queries \mathcal{A} has made.
- $\vec{X} = (\vec{X}_1, \dots, \vec{X}_p)$: The p n -qubit registers used to store the classical queries that \mathcal{A} makes to Ev.

Games H_0, H_1 $\mathbf{H} \leftarrow_s \text{Fcs}(k, n, m)$ $\mathbf{F} \leftarrow_s \text{Fcs}(0, n, m)$ $ H, F\rangle \leftarrow \mathbf{H}, \mathbf{F}\rangle$ $ H, F\rangle \leftarrow \mathcal{H} 0^{2^{k+n} \cdot m}, 0^{2^n \cdot m}\rangle$ Run $\mathcal{A}[\text{Ev}, \text{Ro}]$ Measure $W[1], H, F$ Return $W[1] = 1$	$\text{Ev}(X, Y : I, \vec{X}, F)$ $I \leftarrow I + 1 \pmod{2^{ I }}$ $\vec{X}_I \leftarrow \vec{X}_I \oplus X$ $Y \leftarrow F(X) \oplus Y$ Return $(X, Y : I, \vec{X}, F)$ $\text{Ro}(K, X, Y : H)$ $Y \leftarrow H_K(X) \oplus Y$ Return $(K, X, Y : H)$
Games H_2, H_3 $ H, F\rangle \leftarrow \mathcal{H} 0^{2^{k+n} \cdot m}, 0^{2^n \cdot m}\rangle$ $\text{Run } \mathcal{A}[\mathcal{H}^{Y,F} \circ \text{FEV} \circ \mathcal{H}^{Y,F}, \mathcal{H}^{Y,H} \circ \text{FRO} \circ \mathcal{H}^{Y,H}]$ $ H, F\rangle \leftarrow 0^{2^{k+n} \cdot m}, 0^{2^n \cdot m}\rangle$ Run $\mathcal{A}[\mathcal{H}^Y \circ \text{FEV} \circ \mathcal{H}^Y, \mathcal{H}^Y \circ \text{FRO} \circ \mathcal{H}^Y]$ $ H, F\rangle \leftarrow \mathcal{H} H, F\rangle$ Measure $W[1], H, F$ Return $W[1] = 1$	$\text{FEV}(X, Y : I, \vec{X}, F)$ $I \leftarrow I + 1 \pmod{2^{ I }}$ $\vec{X}_I \leftarrow \vec{X}_I \oplus X$ $F(X) \leftarrow F(X) \oplus Y$ Return $(X, Y : I, \vec{X}, F)$ $\text{FRO}(K, X, Y : H)$ $H_K(X) \leftarrow H_K(X) \oplus Y$ Return $(K, X, Y : H)$

Fig. 4. Hybrid games H_0 through H_3 for the proof of Theorem 3 which are equivalent to the ideal world of $G_{\text{FFX}}^{\text{prf}}$. Highlighted or boxed code is only included in the correspondingly highlighted or boxed game.

We start by changing our perspective. A quantum algorithm that makes classical queries to Ev can be modeled by thinking of a quantum algorithm that measures its X register immediately before the query. (Because the behavior of Ev is completely classical at this point, we do not need to measure the Y register as well.) Measuring the register X is indistinguishable from using a CNOT operation to copy it into a separate register (i.e. xor-ing X into the previously empty register \vec{X}_I that will never again be modified). By incorporating this CNOT operation into the behavior of our hybrid game, we treat \mathcal{A} as an attacker that makes fully quantum queries to its oracles in the hybrid game. We think of \mathcal{A} as deferring all of measurements until the end of its computation. Because all that matters is its final output $W[1]$ we can have the game measure just that register and assume that \mathcal{A} does not internally make any measurements. The principle of deferred measurement ensures that the various changes discussed here do not change the behavior of \mathcal{A} . This perspective change lets us use purely quantum analysis, rather than mixing quantum and classical.

CLAIM 1. We start by considering the hybrids H_0 through H_3 , defined in Fig. 4 which are all identical to the ideal world of $G_{\text{FFX}}^{\text{prf}}$. In these transitions we are applying the ideas of Zhandry [35] to transition to representing the random functions stored in H and F by an all zeros table which is updated whenever the adversary makes a query.

The hybrid H_0 is mostly just $G_{\text{FFX}}^{\text{prf}}$ rewritten to use the registers indicated above. So $\Pr[G_{\text{FFX},0}^{\text{prf}}(\mathcal{A})] = \Pr[H_0]$ holds.

Next consider H_1 which differs from H_0 only in the grey highlighted code which initializes H and F in the uniform superposition and then measures them at the end of execution. (Recall that the Hadamard transform applied to the all zeros state gives the uniform superposition.) Note that these register control, but are unaffected by the oracles Ev and Ro. Because they are never modified while \mathcal{A} is executing, the principle of deferred measurement tells us that this modification is undetectable by \mathcal{A} , giving $\Pr[H_0] = \Pr[H_1]$

Next consider H_2 which contains the boxed, but not the highlighted, code. This game uses the oracles FEV and FRO, the Fourier versions of Ev and Ro, which xor the Y value of \mathcal{A} 's query into the the register F or H . Note that \mathcal{A} 's access to these oracles is mitigated by $\mathcal{H}^{Y,F}$ on each query. The superscript here indicate that

<p>Games H_9, H_8</p> <p>$\mathbf{H} \leftarrow_s \text{Fcs}(k, n, m)$</p> <p>$\mathbf{K}_1 \leftarrow_s \{0, 1\}^k$</p> <p>$\mathbf{K}_2 \leftarrow_s \{0, 1\}^n$</p> <p>$H, K_1, K_2\rangle \leftarrow \mathbf{H}, \mathbf{K}_1, \mathbf{K}_2\rangle$</p> <p>$H, K_1, K_2\rangle \leftarrow \mathcal{H} 0^{2^{k+n} \cdot m}, 0^k, 0^n\rangle$</p> <p>Run $\mathcal{A}[\text{Ev}, \text{Ro}]$</p> <p>Measure $W[1], H, K_1, K_2$</p> <p>Return $W[1] = 1$</p>	<p>$\text{Ev}(X, Y : H, I, \vec{X}, K_1, K_2)$</p> <p>$I \leftarrow I + 1 \bmod 2^{ I }$</p> <p>$\vec{X}_I \leftarrow \vec{X}_I \oplus X$</p> <p>$Y \leftarrow H_{K_1}(X \oplus K_2) \oplus Y$</p> <p>Return $(X, Y : H, I, \vec{X}, K_1, K_2)$</p> <p>$\text{Ro}(K, X, Y : H)$</p> <p>$Y \leftarrow H_K(X) \oplus Y$</p> <p>Return $(K, X, Y : H)$</p>
<p>Games H_7, H_6</p> <div style="border: 1px solid black; padding: 2px; margin: 2px 0;">$H, K_1, K_2\rangle \leftarrow \mathcal{H} 0^{2^{k+n} \cdot m}, 0^k, 0^n\rangle$</div> <div style="border: 1px solid black; padding: 2px; margin: 2px 0;">$O \leftarrow \mathcal{H}^{Y, H} \circ \text{FEV} \circ \mathcal{H}^{Y, H}$</div> <div style="border: 1px solid black; padding: 2px; margin: 2px 0;">$O' \leftarrow \mathcal{H}^{Y, H} \circ \text{FRO} \circ \mathcal{H}^{Y, H}$</div> <p>Run $\mathcal{A}[O, O']$</p> <p>$K_1, K_2\rangle \leftarrow \mathcal{H} 0^k, 0^n\rangle$</p> <p>$H\rangle \leftarrow 0^{2^{k+n} \cdot m}\rangle$</p> <p>Run $\mathcal{A}[\mathcal{H}^Y \circ \text{FEV} \circ \mathcal{H}^Y, \mathcal{H}^Y \circ \text{FRO} \circ \mathcal{H}^Y]$</p> <p>$H\rangle \leftarrow \mathcal{H} H\rangle$</p> <p>Measure $W[1], H, K_1, K_2$</p> <p>Return $W[1] = 1$</p>	<p>$\text{FEV}(X, Y : H, I, \vec{X}, K_1, K_2)$</p> <p>$I \leftarrow I + 1 \bmod 2^{ I }$</p> <p>$\vec{X}_I \leftarrow \vec{X}_I \oplus X$</p> <p>$H_{K_1}(X \oplus K_2) \leftarrow H_{K_1}(X \oplus K_2) \oplus Y$</p> <p>Return $(X, Y : H, I, \vec{X}, K_1, K_2)$</p> <p>$\text{FRO}(K, X, Y : H)$</p> <p>$H_K(X) \leftarrow H_K(X) \oplus Y$</p> <p>Return $(K, X, Y : H)$</p>

Fig. 5. Hybrid games H_9 through H_6 for the proof of Theorem 3 which are equivalent to the real world of $G_{\text{FFX}}^{\text{prf}}$. Highlighted or boxed code is only included in the correspondingly highlighted or boxed game.

the Hadamard transform is being applied to the registers Y and F . We have that $\mathcal{H}^{Y, F} \circ \text{FEV} \circ \mathcal{H}^{Y, F} = \text{Ev}$ and $\mathcal{H}^{Y, H} \circ \text{FRO} \circ \mathcal{H}^{Y, H} = \text{Ro}$ both hold.¹⁰ So $\Pr[H_1] = \Pr[H_2]$ because the adversary's oracles are identical.

Next consider H_3 which contains the highlighted, but not the boxed, code. For this transition, recall that $\mathcal{H} \circ \mathcal{H}$ is the identity operator. So to transition to this game we can cancel the \mathcal{H} operator used to initialize H with the \mathcal{H}^H operator applied before \mathcal{A} 's first FRO oracle query. Similarly, we can cancel the \mathcal{H}^H operation performed after any (non-final) FRO query with the \mathcal{H}^H operation performed before the next FRO query. Finally, the \mathcal{H}^H operation that would be performed after the final FRO query is instead delayed to be performed immediately before H is measured. (We could have omitted this operation and measurement entirely because all that matters at that point is the measurement of $W[1]$.) The \mathcal{H} operators on F are similarly changed. Because \mathcal{A} does not have access to the H and F registers, we can indeed commute the \mathcal{H} operators with \mathcal{A} in this manner without changing behavior. Hence $\Pr[H_2] = \Pr[H_3]$, as desired. Note that H and F independently store tables which are initialized to all zeros and then written into by the adversary's queries.

CLAIM 2. We now consider the hybrids H_4 through H_9 (starting from H_9), which are defined in Fig. 5 and Fig. 6 using similar ideas as in the transition from H_0 through H_3 . As we will justify, these games are all equivalent to the real world of $G_{\text{FFX}}^{\text{prf}}$.

First H_9 rewrote the real world of $G_{\text{FFX}}^{\text{prf}}$ to use our specified registers and to record queries into \vec{X} in Ev. Then in H_8 , rather than sampling H , K_1 , and K_2 uniformly at the beginning of the game we put them in the uniform superposition and measure them at the end of the game. For H_7 we replace our oracles that xor

¹⁰ This follows as a consequence of the following. Let U_{\oplus} and U'_{\oplus} be the unitaries which for $y, z \in \{0, 1\}$ are defined by $U_{\oplus} |y, z\rangle = |y \oplus z, z\rangle$ and $U'_{\oplus} |y, z\rangle = |y, y \oplus z\rangle$. Then $\mathcal{H} \circ U_{\oplus} \circ \mathcal{H} = U'_{\oplus}$.

<p>Games \mathbf{H}_5, \mathbf{H}_4</p> <p>$K_1, K_2\rangle \leftarrow \mathcal{H}[0^k, 0^n]$</p> <p>$H, F\rangle \leftarrow 0^{2^{k+n} \cdot m}, 0^{2^n \cdot m}\rangle$</p> <p style="border: 1px solid black; padding: 2px;">$O \leftarrow \mathcal{H}^Y \circ \mathcal{T} \circ \text{FEV} \circ \mathcal{T} \circ \mathcal{H}^Y$</p> <p style="border: 1px solid black; padding: 2px;">$O' \leftarrow \mathcal{H}^Y \circ \mathcal{T} \circ \text{FRO} \circ \mathcal{T} \circ \mathcal{H}^Y$</p> <p style="border: 1px solid black; padding: 2px;">Run $\mathcal{A}[O, O']$</p> <p>Run $\mathcal{A}[\mathcal{H}^Y \circ \text{FEV}' \circ \mathcal{H}^Y, \mathcal{H}^Y \circ \text{FRO}' \circ \mathcal{H}^Y]$</p> <p>$H, F, I, \vec{X}, K_1, K_2\rangle \leftarrow \mathcal{T} H, F, I, \vec{X}, K_1, K_2\rangle$</p> <p>$H\rangle \leftarrow \mathcal{H} H\rangle$</p> <p>Measure $W[1], H, K_1, K_2$</p> <p>Return $W[1] = 1$</p> <p style="border: 1px solid black; padding: 2px;">$\text{FRO}'(K, X, Y : H, F, I, \vec{X}, K_1, K_2)$</p> <p>If $K = K_1$ and $X \oplus K_2 \in \{\vec{X}_1, \dots, \vec{X}_I\}$ then</p> <p style="padding-left: 20px;">// Input is bad</p> <p style="padding-left: 20px;">$F(X \oplus K_2) \leftarrow F(X \oplus K_2) \oplus Y$</p> <p>Else</p> <p style="padding-left: 20px;">$H_K(X) \leftarrow H_K(X) \oplus Y$</p> <p>Return $(K, X, Y : H, F, I, \vec{X}, K_1, K_2)$</p>	<p style="border: 1px solid black; padding: 2px;">$\text{FEV}(X, Y : H, I, \vec{X}, K_1, K_2)$</p> <p>$I \leftarrow I + 1 \pmod{2^I}$</p> <p>$\vec{X}_I \leftarrow \vec{X}_I \oplus X$</p> <p>$H_{K_1}(X \oplus K_2) \leftarrow H_{K_1}(X \oplus K_2) \oplus Y$</p> <p>Return $(X, Y : H, I, \vec{X}, K_1, K_2)$</p> <p style="border: 1px solid black; padding: 2px;">$\text{FRO}(K, X, Y : H)$</p> <p>$H_K(X) \leftarrow H_K(X) \oplus Y$</p> <p>Return $(K, X, Y : H)$</p> <p style="border: 1px solid black; padding: 2px;">$\text{FEV}'(X, Y : H, F, I, \vec{X}, K_1, K_2)$</p> <p>$I \leftarrow I + 1 \pmod{2^I}$</p> <p>$\vec{X}_I \leftarrow \vec{X}_I \oplus X$</p> <p>$\text{bool}_1 \leftarrow (X \notin \{\vec{X}_1, \dots, \vec{X}_{I-1}\})$</p> <p>$\text{bool}_2 \leftarrow (H_{K_1}(X \oplus K_2) \neq 0^m)$</p> <p>$\text{bool}_3 \leftarrow (F(X) \neq 0^m)$</p> <p>If bool_1 and $(\text{bool}_2$ or $\text{bool}_3)$ then</p> <p style="padding-left: 20px;">// Input is bad</p> <p style="padding-left: 20px;">$F'(X) \leftarrow F(X)$</p> <p style="padding-left: 20px;">$F(X) \leftarrow H_{K_1}(X \oplus K_2)$</p> <p style="padding-left: 20px;">$H_{K_1}(X \oplus K_2) \leftarrow F'(X)$</p> <p style="padding-left: 20px;">$F(X) \leftarrow F(X) \oplus Y$</p> <p>Return $(X, Y : H, I, \vec{X}, K_1, K_2)$</p>
---	--

Fig. 6. Hybrid games \mathbf{H}_5 and \mathbf{H}_4 for the proof of Theorem 3 which are equivalent to the real world of $\mathbf{G}_{\text{FFX}}^{\text{prf}}$. Highlighted or boxed code is only included in the correspondingly highlighted or boxed game. Unitary \mathcal{T} is define in the text.

into the adversary's Y register with oracles that xor into the H register using Hadamard operations, some of which we then cancel out to transition to \mathbf{H}_6 . The same arguments from Claim 1 of why these sorts of modifications do not change the behavior of the game apply here and so $\Pr[\mathbf{G}_{\text{FFX},1}^{\text{prf}}(\mathcal{A})] = \Pr[\mathbf{H}_9] = \Pr[\mathbf{H}_8] = \Pr[\mathbf{H}_7] = \Pr[\mathbf{H}_6]$.

Our next transitions are designed to make the current game's oracles identical with those of \mathbf{H}_3 , except on some "bad" inputs. In \mathbf{H}_6 we have a single all zeros table H which gets written into by queries that \mathcal{A} makes to either of its oracles, while in \mathbf{H}_3 the oracles separately wrote into either H or F . For \mathbf{H}_5 we will similarly separate the single table H into separate tables H and F . However, we cannot keep them completely independent, because if the adversary queries FEV with $X = x$ and FRO with $(K, X) = (K_1, x \oplus K_2)$ then both of these operations would be writing into the same table location in \mathbf{H}_6 . Consider the following unitary \mathcal{T} which acts on registers H and F and is controlled by the registers $I, \vec{X}, K_1,$ and K_2 . We will think of this unitary as transitioning us between a representation of H as a single table (with an all-zero F table) and a representation of it divided between H and F .

$$\begin{aligned}
 & \mathcal{T}(H, F, I, \vec{X}, K_1, K_2) \\
 & \text{For } x \in \{\vec{X}_1, \dots, \vec{X}_I\} \text{ do} \\
 & \quad F'(x) \leftarrow F(x) \\
 & \quad F(x) \leftarrow H_{K_1}(x \oplus K_2) \\
 & \quad H_{K_1}(x \oplus K_2) \leftarrow F'(x) \\
 & \text{Return } (H, F, I, \vec{X}, K_1, K_2)
 \end{aligned}$$

In words, \mathcal{T} swaps $F(x)$ and $H_{K_1}(x \oplus K_2)$ for each x that has been previous queried to FEV (as stored by \vec{X} and I). Note that \mathcal{T} is its own inverse. In \mathbf{H}_5 we (i) initialize the table F as all zeros, (ii) perform \mathcal{T} before and after each oracle query, and (iii) perform \mathcal{T} after \mathcal{A} has executed. We verify that H has the same

value in H_5 that it would have had in H_6 during each oracle query and at the end before measurement. The application of \mathcal{T} before the first oracle query does nothing (because $I = 0$) so H is all zeros for this query as required. As we've seen previously with \mathcal{H} , we can commute \mathcal{T} with the operations of \mathcal{A} because \mathcal{T} only acts on registers outside of the adversary's control. We can similarly commute \mathcal{T} with \mathcal{H}^Y . Hence the \mathcal{T} operation after every non-final oracle query can be seen to cancel with the \mathcal{T} operation before the following oracle query. The \mathcal{T} operation after the final oracle query cancels with the \mathcal{T} operation performed after \mathcal{A} halts execution. Hence, $\Pr[H_6] = \Pr[H_5]$ as claimed.

For the transition to H_4 let us dig into how our two-table representation in H_5 works so that we can incorporate the behavior of \mathcal{T} directly into the oracles. For simplicity of notation in discussion, let $\tilde{H}(x)$ denote $H_{K_1}(x \oplus K_2)$. First note that in between oracle queries the two tables representation of H_5 will satisfy the property that for each $x \in \{\vec{X}_1, \dots, \vec{X}_I\}$ we will have $\tilde{H}(x) = 0^{F.ol}$ and for all other x we will have that $F(x) = 0^{F.ol}$.¹¹ This is the case because after each query we have applied \mathcal{T} to an H which contains the same values it would have in H_6 and an F which is all zeros.

Now consider when a $\mathcal{T} \circ \text{FEV} \circ \mathcal{T}$ query is executed with some X and Y . If $X \in \{\vec{X}_1, \dots, \vec{X}_I\}$, then $F(X)$ and $\tilde{H}(X)$ are swapped, Y is xored into $\tilde{H}(X)$, then finally $F(X)$ and $\tilde{H}(X)$ are swapped back. Equivalently, we could have xored Y directly into $F(X)$ and skipped the swapping around. If $X \notin \{\vec{X}_1, \dots, \vec{X}_I\}$, then Y is xored into $\tilde{H}(X)$ before $F(X)$ and $\tilde{H}(X)$ are swapped. If $\tilde{H}(X) = 0^{F.ol}$ beforehand, then we could equivalently have xored Y directly into $F(X)$ and skipped the swapping around (because $F(X) = 0^{F.ol}$ must have held from our assumption on X). If $\tilde{H}(X) \neq 0^{F.ol}$ beforehand, then we could equivalently could have swapped $\tilde{H}(X)$ and $F(X)$ first, then xored Y into $F(X)$. The equivalent behavior we described is exactly the behavior captured by the oracle FEV' which is used in H_4 in place of $\mathcal{T} \circ \text{FEV} \circ \mathcal{T}$. It checks if $X \notin \{\vec{X}_1, \dots, \vec{X}_I\}$ (bool_1) and $\tilde{H}(X) \neq 0^m$ (bool_2), performing a swap if so. Then Y is xored into $F(X)$. A swap is also performed if $X \notin \{\vec{X}_1, \dots, \vec{X}_{I-1}\}$ and $F(X) \neq 0^m$, however this case is impossible from our earlier observation that $F(x) = 0^m$ when X is not in \vec{X} . This second case was added only to ensure that FEV' is a permutation.

Similarly, consider when a $\mathcal{T} \circ \text{FRO} \circ \mathcal{T}$ query is executed with some K, X, Y . Any swapping done by \mathcal{T} uses H_{K_1} , so when $K \neq K_1$ this just xors Y into $H_K(X)$. If $X \oplus K_2$ is not in $\{\vec{X}_1, \dots, \vec{X}_I\}$, then $H_K(X)$ is unaffected by the swapping so again this just xors Y into $H_K(X)$. When $K = K_1$ and $X \oplus K_2 \in \{\vec{X}_1, \dots, \vec{X}_I\}$, first $H_K(X)$ would have been swapped with $F(X \oplus K_2)$, then Y would be xored into $H_K(X)$, then $H_K(X)$ and $F(X \oplus K_2)$ would be swapped back. Equivalently, we could just have xored Y into $F(X \oplus K_2)$ and skipped the swapping. This behavior we described is exactly the behavior captured by the oracle FRO' which is used in H_4 in place of $\mathcal{T} \circ \text{FRO} \circ \mathcal{T}$.

We have just described that on the inputs we care about FEV' behaves identically to $\mathcal{T} \circ \text{FEV} \circ \mathcal{T}$ and FRO' behaves identically to $\mathcal{T} \circ \text{FRO} \circ \mathcal{T}$. Hence $\Pr[H_5] = \Pr[H_4]$, completing this claim.

CLAIM 3. To compare hybrids H_3 and H_4 we will note their oracles only differ on a small number of inputs (in particular those labelled as bad by comments in our code) and then apply Theorem 1 to bound the difference between them. To aid in this we have rewritten them as \tilde{H}_3 and \tilde{H}_4 in Fig. 7. For both we have removed some operations performed after \mathcal{A} halted its execution for cleanliness because these operations on registers other than $W[1]$ cannot affect the probability that it is measured to equal 1. So we have $\Pr[\tilde{H}_3] = \Pr[H_3]$ and $\Pr[\tilde{H}_4] = \Pr[H_4]$.

Let FEV_3 and FRO_3 be the permutations defining the corresponding oracle in \tilde{H}_3 . Define FEV_4 and FRO_4 analogously. These permutations differ only on the inputs we referred to as bad. So let S denote the set of bad $X, H, F, I, \vec{X}, K_1, K_2$ for FEV (i.e. those for which bool_1 and either bool_2 or bool_3 hold). Let S' denote the set of bad $K, X, H, F, I, \vec{X}, K_1, K_2$ for FRO (i.e. those for which $K = K_1$ and $X \oplus K \in \{\vec{X}_1, \dots, \vec{X}_I\}$). Let \mathbf{D} denote the distribution which always outputs $(S, S', \text{FEV}_3, \text{FRO}_3, \text{FEV}_4, \text{FRO}_4, \varepsilon)$. Clearly this is a valid distribution for Theorem 1 by our choice of S and S' .

Now we can define an adversary \mathcal{A}' for $\mathbf{G}_{\mathbf{D}, b}^{\text{dist}}$ which simulates the view of \mathcal{A} in \tilde{H}_{3+b} by locally running the code of that hybrid except for during oracle queries when it uses its f_b oracle to simulate FEV and f'_b

¹¹ More precisely, the corresponding registers hold superpositions over tables satisfying the properties we discuss.

Games \tilde{H}_3, \tilde{H}_4	$\text{FEV}(X, Y : H, F, I, \vec{X}, K_1, K_2)$
$ K_1, K_2\rangle \leftarrow \mathcal{H} 0^k, 0^n\rangle$	$I \leftarrow I + 1 \pmod{2^{l_I}}$
$ H, F\rangle \leftarrow 0^{2^{k+n} \cdot m}, 0^{2^n \cdot m}\rangle$	$\vec{X}_I \leftarrow \vec{X}_I \oplus X$
Run $\mathcal{A}[\mathcal{H}^Y \circ \text{FEV} \circ \mathcal{H}^Y, \mathcal{H}^Y \circ \text{FRO} \circ \mathcal{H}^Y]$	$\text{bool}_1 \leftarrow (X \notin \{\vec{X}_1, \dots, \vec{X}_{I-1}\})$
Measure $W[1]$	$\text{bool}_2 \leftarrow (H_{K_1}(X \oplus K_2) \neq 0^m)$
Return $W[1] = 1$	$\text{bool}_3 \leftarrow (F(X) \neq 0^m)$
$\text{FRO}(K, X, Y : H, F, I, \vec{X}, K_1, K_2)$	If bool_1 and (bool_2 or bool_3)
If $K = K_1$ and $X \oplus K_2 \in \{\vec{X}_1, \dots, \vec{X}_I\}$ then	// Input is bad
// Input is bad	$F'(X) \leftarrow F(X)$
$F(X \oplus K_2) \leftarrow F(X \oplus K_2) \oplus Y$	$F(X) \leftarrow H_{K_1}(X \oplus K_2)$
$H_K(X) \leftarrow H_K(X) \oplus Y$	$H_{K_1}(X \oplus K_2) \leftarrow F'(X)$
Else	$F(X) \leftarrow F(X) \oplus Y$
$H_K(X) \leftarrow H_K(X) \oplus Y$	Return $(X, Y : H, I, \vec{X}, K_1, K_2)$
Return $(K, X, Y : H, F, I, \vec{X}, K_1, K_2)$	

Fig. 7. Hybrid games \tilde{H}_3 and \tilde{H}_4 for the proof of Theorem 3 which are rewritten versions of H_3 and H_4 to emphasize that their oracles are identical-until-bad. Highlighted or boxed code is only included in the correspondingly highlighted or boxed game.

oracle to simulate FRO. Because the simulation of these views are perfect we have that

$$\Pr[\tilde{H}_3] - \Pr[\tilde{H}_4] = \text{Adv}_{\mathbf{D}}^{\text{dist}}(\mathcal{A}') \leq 2(p+q)\sqrt{\text{Adv}_{\mathbf{D}}^{\text{guess}}(\mathcal{A}')}$$

where the inequality follows from Theorem 1, noting that \mathcal{A}' makes $p+q$ oracle queries.

To complete the proof we bound $\text{Adv}_{\mathbf{D}}^{\text{guess}}(\mathcal{A}')$. In the following probability calculation we use x and i to denote random variables taking on the values the corresponding variables have at the end of an execution of $\mathbf{G}_{\mathbf{D}}^{\text{guess}}(\mathcal{A})$. Let \mathcal{S} denote a random variable which equals S if the measured query is to f_0 and equals S' otherwise. Then conditioning over each possible value of i gives

$$\begin{aligned} \text{Adv}_{\mathbf{D}}^{\text{guess}}(\mathcal{A}') &= \Pr[x \in \mathcal{S}] = \sum_{j=1}^{p+q} \Pr[x \in \mathcal{S} \mid i = j] \Pr[i = j] \\ &= (p+q)^{-1} \sum_{j=1}^{p+q} \Pr[x \in \mathcal{S} \mid i = j]. \end{aligned}$$

Because \mathcal{A} is order consistent we can pick disjoint sets E and R with $E \cup R = \{1, \dots, p+q\}$ such that $i \in E$ means the i -th query is to \mathcal{A}' 's FEV oracle and $i \in R$ means that the i -th query is to its FRO oracle. Note that $|E| = q$ and $|R| = p$. The view of \mathcal{A} (when run by \mathcal{A}') in $\mathbf{G}_{\mathbf{D}}^{\text{guess}}$ matches its view in \tilde{H}_3 so, in particular, it is independent of K_1 and K_2 . Hence we can think of these keys being chosen at random at the end of execution when analyzing the probability of $x \in \mathcal{S}$.

For a FRO query, the check for bad inputs is if $K_1 = K$ and $K_2 \in \{X \oplus \vec{X}_1, \dots, X \oplus \vec{X}_I\}$. The variable I is counting the number of FEV queries made so far, so $I < q$. By a union bound,

$$\Pr[x \in \mathcal{S}' \mid i = j] \leq q/2^{\text{F.kl}+\text{F.il}}$$

when $j \in R$.

For a FEV query, the check for bad inputs is if $X \notin \{\vec{X}_1, \dots, \vec{X}_{I-1}\}$ and $H_{K_1}(X \oplus K_2)$ is non-zero.¹² In \tilde{H}_3 , each query to FRO can make a single entry of H non-zero so it will never have more than p non-zero

¹² Note that the other bad inputs, for which $X \notin \{\vec{X}_1, \dots, \vec{X}_{I-1}\}$ and $F(X)$ is non-zero, will never occur.

entries. By a union bound,

$$\Pr[x \in S \mid i = j] \leq p/2^{F.kl+F.il}$$

when $j \in E$.

The proof is then completed by noting

$$\begin{aligned} \sum_{j=1}^{p+q} \Pr[x \in S \mid i = j] &= \sum_{j \in R} \Pr[x \in S' \mid i = j] + \sum_{j \in E} \Pr[x \in S \mid i = j] \\ &\leq p(q/2^{F.kl+F.il}) + q(p/2^{F.kl+F.il}) = 2pq/2^{F.kl+F.il} \end{aligned}$$

and plugging in to our earlier expression. \square

4 Double Encryption

In this section we prove the security of the double encryption key-length extension technique against fully quantum attacks. Our proof first reduces this to the ability of a quantum algorithm to solve the list disjointness problem and then extends known query lower bounds for element distinctness to list disjointness (with some modifications).

The double encryption blockcipher is constructed via two sequential application of an underlying blockcipher. Formally, given a blockcipher $E \in \text{lcs}(E.kl, E.bl)$, we define the double encryption blockcipher $\text{DE}[E]$ by $\text{DE}[E](K_1 \parallel K_2, x) = E_{K_2}(E_{K_1}(x))$. Here $|K_1| = |K_2| = E.kl$ so $\text{DE}[E].kl = 2E.kl$ and $\text{DE}[E].bl = E.bl$. Its inverse can be computed as $\text{DE}[E]^{-1}(K_1 \parallel K_2, x) = E_{K_1}^{-1}(E_{K_2}^{-1}(x))$.

Classically, the meet-in-the-middle attack [10,26] shows that this construction achieves essentially the same security a single encryption. In the quantum setting, this construction was recently considered by Kaplan [16]. They gave an attack and matching security bound for the key-recovery security of double encryption. This leaves the question of whether their security result can be extended to cover full SPRP security, which we resolve by the main theorem of this section. This theorem is proven via a reduction technique of Tessaro and Thiruvengadam [31] which they used to establish a (classical) time-memory tradeoff for the security of double encryption, by reducing its security to the list disjointness problem and conjecturing a time-memory tradeoff for that problem.

PROBLEMS AND LANGUAGES. In addition to the list disjointness problem (1LD), we will also consider two versions of the element distinctness problem (ED, 1ED). In general, a problem PB specifies a relation \mathcal{R} on set on instances \mathcal{I} (i.e. \mathcal{R} is function which maps an instance $L \in \mathcal{I}$ and witness w to a decision $\mathcal{R}(L, w) \in \{0, 1\}$). This relation induces a language $\mathcal{L} = \{L \in \mathcal{I} : \exists w, \mathcal{R}(L, w) = 1\}$. Rather than think of instances as bit strings, we will think of them as functions (to which decision and search algorithms are given oracle access). To restrict attention to functions of specific sizes we let $\mathcal{L}(D, R) = \mathcal{L} \cap \mathcal{I}(D, R)$ where $\mathcal{I}(D, R)$ denotes the restriction of \mathcal{I} to functions $L : [D] \rightarrow [R]$, where $D \leq R$. To discuss instances not in the language we let $\mathcal{L}' = \mathcal{I} \setminus \mathcal{L}$ and $\mathcal{L}'(D, R) = \mathcal{I}(D, R) \setminus \mathcal{L}$.

Problems have decision and search versions. The goal of a decision algorithm is to output 1 (representing “acceptance”) on instances in the language and 0 (representing “rejection”) otherwise. Relevant quantities are the minimum probability P^1 of accepting an instance in the language, the maximum probability P^0 of accepting an instance not in the language, and the error rater E which are formally defined by

$$\begin{aligned} P_{D,R}^1(\mathcal{A}) &= \min_{L \in \mathcal{L}(D,R)} \Pr[\mathcal{A}[L] = 1], \quad P_{D,R}^0(\mathcal{A}) = \max_{L \in \mathcal{L}'(D,R)} \Pr[\mathcal{A}[L] = 1] \\ E_{D,R}(\mathcal{A}) &= \max\{1 - P_{D,R}^1(\mathcal{A}), P_{D,R}^0(\mathcal{A})\}. \end{aligned}$$

We define the decision PB advantage of \mathcal{A} by $\text{Adv}_{D,R}^{\text{PB}}(\mathcal{A}) = P_{D,R}^1(\mathcal{A}) - P_{D,R}^0(\mathcal{A})$. In non-cryptographic contexts, instead of looking at the difference in probability that inputs that are in or out of the language are accepted, one often looks at how far these are each from 1/2. This motivates the definition $\text{Adv}_{D,R}^{\text{PB-d}}(\mathcal{A}) = \min\{2P_{D,R}^1(\mathcal{A}) - 1, 1 - 2P_{D,R}^0(\mathcal{A})\} = 1 - 2E_{D,R}(\mathcal{A})$.

The goal of a search algorithm is to output a witness for the instance. We define its advantage to be the minimum probability it succeeds, i.e., $\text{Adv}_{D,R}^{\text{PB-s}}(\mathcal{A}) = \min_{L \in \mathcal{L}(D,R)} \Pr[\mathcal{R}(L, \mathcal{A}[L]) = 1]$.

Problem	Witness	Promise
ED	$x \neq y$ s.t. $L(x) = L(y)$	-
1ED	$x \neq y$ s.t. $L(x) = L(y)$	At most one witness.
1LD	x, y s.t. $L_0(x) = L_1(y)$	At most one witness. Injective L_0, L_1 .

Fig. 8. Summary of the element distinctness and list disjointness problems we consider.

EXAMPLE PROBLEMS. The *list disjointness* problem asks how well an algorithm can distinguish between the case that is give (oracle access to) two lists which are disjoint or have one element in common. In particular, we interpret an instance L as the two functions $L_0, L_1 : \llbracket D/2 \rrbracket \rightarrow [R]$ defined by $L_b(x) = L(x + b\lfloor D/2 \rfloor)$. Let \mathcal{S}_n denote the set of L for which L_0 and L_1 are injective and which have n elements in common, i.e. for which $|\{L_0(1), \dots, L_0(\lfloor D/2 \rfloor)\} \cap \{L_1(1), \dots, L_1(\lfloor D/2 \rfloor)\}| = n$. Then 1LD is defined by the relation \mathcal{R} which on input $(L, (x, y))$ returns 1 iff $L_0(x) = L_1(y)$ and the instance set $\mathcal{I} = \mathcal{S}_0 \cup \mathcal{S}_1$ (i.e., the promise that there is at most one element in common and that the lists are individually injective). The search version of list disjointness is sometimes referred to as claw-finding.

The *element distinctness* problem asks how well an algorithm can detect whether all the elements in a list are distinct. Let \mathcal{S}'_n denote the set of L which have n collision pairs, i.e. for which $|\{\{x, y\} : x \neq y, L(x) = L(y)\}| = n$. Then ED is defined by the relation \mathcal{R} which on input $(L, (x, y))$ returns 1 iff $x \neq y$ and $L(x) = L(y)$ with the instance set $\mathcal{I} = \bigcup_{n=0}^{\infty} \mathcal{S}'_n$ consisting of all functions. We let 1ED denote restricting ED to $\mathcal{I} = \mathcal{S}'_0 \cup \mathcal{S}'_1$ (i.e., the promise that there is at most one repetition in the list).

4.1 Security result

The following theorem shows that an attacker achieving constant advantage must make $\Omega(2^{2k/3})$ oracle queries (ignoring log terms). Our bound is not tight for a lower parameter regimes, though future work may establish better bounds for list disjointness in these regimes.

Theorem 4. Consider $\text{DE}[\cdot]$ with the underlying blockcipher modeled by an ideal cipher drawn from $\text{lcs}(k, n)$. Let \mathcal{A} be a quantum adversary which makes at most q queries to the ideal cipher. Then

$$\text{Adv}_{\text{DE}}^{\text{sprp}}(\mathcal{A}) \leq 11 \sqrt[6]{(q \cdot k \lg k)^3 / 2^{2k} + 1/2^k}.$$

As mentioned earlier, our proof works by first reducing the security of double encryption against quantum queries to the security of the list disjointness problem against quantum queries. This is captured in the following theorem which we prove now.

Theorem 5. Consider $\text{DE}[\cdot]$ with the underlying blockcipher modeled by an ideal cipher drawn from $\text{lcs}(k, n)$. Let \mathcal{A} be a quantum adversary which makes at most q queries to the ideal cipher. Then for any $R \geq 2^k$ we can construct \mathcal{A}' making at most q oracle queries such that

$$\text{Adv}_{\text{DE}}^{\text{sprp}}(\mathcal{A}) \leq \text{Adv}_{2^k, R}^{\text{1LD-d}}(\mathcal{A}') + 1/2^k.$$

We state and prove a bound on $\text{Adv}^{\text{1LD-d}}$ in Section 4.2. Our proof applies the same reduction technique as Tessaro and Thiruvengadam [31], we are verifying that it works quantumly as well.

Proof. For $b \in \{0, 1\}$, let H_b be defined to be identical to $\text{G}_{\text{DE}, b}^{\text{sprp}}$ except that K_2 is chosen uniformly from $\{0, 1\}^k \setminus \{K_1\}$ rather than from $\{0, 1\}^k$. This has no effect when $b = 0$ because the keys are not used, so $\Pr[\text{G}_{\text{DE}, 0}^{\text{sprp}}(\mathcal{A})] = \Pr[\text{H}_0]$. When $b = 1$ there was only a $1/2^k$ chance that K_2 would have equalled K_1 so $\Pr[\text{G}_{\text{DE}, 1}^{\text{sprp}}(\mathcal{A})] \leq \Pr[\text{H}_1] + 1/2^k$.

Now we define a decision algorithm \mathcal{A}' for 1LD which uses its input lists to simulate a view for \mathcal{A} . When the lists are disjoint \mathcal{A}' 's view will perfectly match that of H_0 and when the lists have exactly one element in

Adversary $\mathcal{A}'[L]$	IC($K, X, Y : \rho, \pi, F$)	INV($K, X, Y : \rho, \pi, F$)
$\rho \leftarrow \text{lcs}(0, k)$	$(i, j) \leftarrow \rho(K)$	$(i, j) \leftarrow \rho(K)$
$\pi \leftarrow \text{lcs}(0, n)$	If $i = 0$ then	If $i = 0$ then
$F \leftarrow \text{lcs}(\lceil \lg R \rceil, n)$	$Y \leftarrow Y \oplus F_{L_0(j)}(X)$	$Y \leftarrow Y \oplus F_{L_0(j)}^{-1}(X)$
$b' \leftarrow \mathcal{A}[\pi, \pi^{-1}, \text{IC}, \text{INV}]$	Else	Else
Return b'	$Y \leftarrow Y \oplus \pi(F_{L_1(j)}^{-1}(X))$	$Y \leftarrow Y \oplus F_{L_1(j)}(\pi^{-1}(X))$
	Return $(K, X, Y : \rho, \pi, F)$	Return $(K, X, Y : \rho, \pi, F)$

Fig. 9. Reduction adversary used in proof of Theorem 5.

common \mathcal{A} 's view will perfectly match that of H_1 . Hence we have $\Pr[H_1] = \min_{(L, L') \in \mathcal{L}_1^{\kappa, \kappa'}} \Pr[G_{L, L'}^{\text{ld}}(\mathcal{A})]$ and $\Pr[H_0] = \max_{(L, L') \in \mathcal{L}_0^{\kappa, \kappa'}} \Pr[G_{L, L'}^{\text{ld}}(\mathcal{A})]$, so $\Pr[H_1] - \Pr[H_0] = \text{Adv}_{2^k, k'}^{\text{ld}}(\mathcal{A}')$ which gives the claimed bound.

The adversary \mathcal{A}' is defined in Fig. 9. It samples a permutation ρ on $\{0, 1\}^k$, a permutation π on $\{0, 1\}^n$, and a cipher F . The permutation ρ is used to provide a random map from the keys $K \in \{0, 1\}^k$ to the elements of the lists L_0 and L_1 . We will interpret $\rho(K)$ as a tuple (i, j) with $i \in \{0, 1\}$ and $j \in [2^k/2]$. Then K gets mapped to $L_i(j)$. Therefore either none of the keys map to the same element or a single pair of them maps to the same element.

Adversary \mathcal{A}' 's queries to EV and INV are answered using π and π^{-1} . Its queries to the ideal cipher are more complicated and are handled by the oracles IC and INV. We can verify that these oracles define permutations on bitstring inputs, so \mathcal{A}' is a well defined quantum adversary. Consider a key K and interpret $\rho(K)$ as a tuple (i, j) as described above. If $i = 0$, then ideal cipher queries for it are answered as if $\mathbf{E}_K(\cdot) = F_{L_0(j)}(\cdot)$. If $i = 1$, then ideal cipher queries for it are answered as if $\mathbf{E}_K(\cdot) = \pi(F_{L_1(j)}^{-1}(\cdot))$.¹³ If the list element K is not mapped to by any other keys, then the indexing into F ensures that $\mathbf{E}_K(\cdot)$ is independent of π and $\mathbf{E}_{K'}$ for all other K' . If K and K' map to the same list element (and $i = 0$ for K), then $\mathbf{E}_K(\cdot)$ and $\mathbf{E}'_{K'}(\cdot)$ are random permutations conditioned on $\mathbf{E}'_{K'}(\mathbf{E}_K(\cdot)) = \pi(\cdot)$ and independent of all other $\mathbf{E}_{K''}(\cdot)$.

In H_0 , the permutation of EV and INV is independent of each $\mathbf{E}_K(\cdot)$ which are themselves independent of each other. So this perfectly matches the view presented to \mathcal{A} by \mathcal{A}' when the lists are disjoint. In H_1 , each $\mathbf{E}_K(\cdot)$ is pairwise independent and the permutation of EV and INV is defined to equal $\mathbf{E}_{K_2}(\mathbf{E}_{K_1}(\cdot))$. This perfectly matches the view presented to \mathcal{A} by \mathcal{A}' when the lists have one element in common because we can think of it as just having changed the order in which the permutations $\mathbf{E}_{K_2}(\mathbf{E}_{K_1}(\cdot))$, $\mathbf{E}_{K_1}(\cdot)$, and $\mathbf{E}_{K_2}(\cdot)$ were sampled. \square

4.2 The Hardness of List Disjointness

If \mathcal{A} is an algorithm making at most q classical oracle queries, then it is not hard to prove that $\text{Adv}_{D, R}^{\text{ILD}}(\mathcal{A}) \leq q/D$. If, instead, \mathcal{A} makes at most q quantum oracle queries, the correct bound is less straightforward. In this section, we will prove the following result.

Theorem 6. *If \mathcal{A} is a quantum algorithm making at most q queries to its oracle and $D \geq 32$ is a power of 2, then*

$$\text{Adv}_{D, 3D^2}^{\text{ILD}}(\mathcal{A}) \leq 11 \sqrt[6]{(q \cdot \lg D \cdot \lg \lg D)^3 / D^2}.$$

We restrict attention to the case that D is a power of 2 only for notational simplicity in the proof. Essentially the same bound for more general D follows from the same techniques.

Ambanis's $\mathcal{O}(N^{2/3})$ query algorithm for element distinctness [2] can be used to solve list disjointness and hence shows this is tight (up to logarithmic factors) for attackers achieving constant advantage. The sixth root degrades the quality of the bound for lower parameter regimes. An interesting question we leave open is whether this could be proven without the sixth root or the logarithmic factors.

¹³ When using output of L as keys for F we are identifying the elements of $[R]$ with elements of $\{0, 1\}^{\lceil \lg R \rceil}$ in the standard manner.

PROOF SKETCH. The starting point for our reduction is that $\Omega(N^{2/3})$ lower bounds are known both the search and decision versions of ED [1,34].¹⁴ By slightly modifying Zhandry’s [34] technique for proving this, we instead get a bound on the hardness of 1ED-s. Next, a simple reduction (split the list in half at random) shows that 1LD-s is as hard as 1ED-s.

Then a “binary search” style reduction shows that 1LD-d is as hard as 1LD-s. In the reduction, the 1LD-s algorithm repeatedly splits its lists in half and uses the 1LD-d algorithm to determine which pair of lists contains the non-disjoint entries. However, we need our reduction to work by running the 1LD-d algorithm on a particular fixed size of list (the particular size we showed 1LD-d is hard for) rather than running it on numerous shrinking sizes. We achieve this by padding the lists with random elements. The choice of $R = 3D^2$ was made so that with good probability these random elements do not overlap with the actual list. This padding adds the $\lg D$ term to our bound.

Finally a generic technique allows us to relate the hardness of 1LD and 1LD-d. Given an algorithm with high 1LD advantage we can run it multiple times to get a precise estimate of how frequently it is outputting 1 and use that to determine what we want to output. This last step is the primary cause of the sixth root in our bound; it required running the 1LD algorithm on the order of $1/\delta^2$ times to get a precise enough estimate, where δ is the advantage of the 1LD algorithm. This squaring of δ in the query complexity of our 1LD-d algorithm (together with the fact that the query complexity is cubed in our 1ED-s bound) ultimately causes the sixth root.

CONSTITUENT LEMMAS. In the rest of the section we will state and prove lemmas corresponding to each of the step of the proof described above. In Section 4.3 we apply them one at a time to obtain the specific bound claimed by Theorem 6.

Lemma 1 (1ED-s is hard). *If \mathcal{A}_{1ED-s} is a quantum algorithm for 1ED-s making at most q queries to its oracle and $D \geq 32$, then $\text{Adv}_{D,3D^2}^{\text{ED-s}}(\mathcal{A}_{1ED-s}) \leq 9 \cdot (q + 2)^3 / D^2$.*

Proof. In [34], Zhandry shows that no Q -query algorithm can distinguish between a random function and a random injective function with domain $[D]$ and codomain $[R]$ with advantage better than $(\pi^2/3)Q^3/R$, as long as $D \leq R$. We could build a distinguisher \mathcal{B} that on input $L : [D] \rightarrow [R]$ runs $(x, y) \leftarrow \mathcal{A}[L]$. If (x, y) is a collision (i.e., $x \neq y$ and $L(x) = L(y)$), then \mathcal{B} outputs 1. It checks this by making two additional L queries, so $Q = q + 2$. Otherwise it outputs zero. Clearly, (x, y) cannot be a collision when L is an injection so it will always output zero. Hence, if 1coll denotes the event that L contains exactly one collision we obtain a bound of

$$\Pr[\text{1coll}] \cdot \text{Adv}_{D,3D^2}^{\text{ED-s}}(\mathcal{A}_{1ED-s}) \leq (\pi^2/3)(q + 2)^3/R. \quad (1)$$

It remains to lower bound $\Pr[\text{1coll}]$. Note there are $\binom{D}{2}$ possible pairs of inputs that could be collisions. Each pair has a $1/R$ chance of colliding. By a union bound, the probability any other input has the same output as the collision is at most $(D - 2)/R$, so there is at least a $1 - (D - 2)/R$ probability this does not occur. Given the above there are $(D - 2)$ inputs sampled from $R - 1$ possible values, so by a union bound the probability none of them collide is at least $1 - \binom{D-2}{2}/(R - 1)$. This gives

$$\Pr[\text{1coll}] \geq \frac{D(D - 1)}{2R} \cdot \left(1 - \frac{D - 2}{R}\right) \cdot \left(1 - \frac{(D - 2)(D - 3)}{2(R - 1)}\right).$$

Now setting $R = 3D^2$ and applying simple bounds (e.g. $D - 2 < D - 1 < D$ and $(D - 3)/(R - 1) < (D - 2)/R$) gives,

$$\Pr[\text{1coll}] \geq \frac{(1 - 1/D)}{6} \cdot \left(1 - \frac{1}{3D}\right) \cdot \left(1 - \frac{1}{6}\right)$$

Plugging this lower bound into equation 1, re-arranging, applying the bound $D \geq 32$, and rounding up to the nearest whole number gives the claimed bound $\text{Adv}_{D,3D^2}^{\text{ED-s}}(\mathcal{A}_{1ED-s}) \leq 9(q + 2)^3/D^3$. \square

¹⁴ In the proof we actually work with the advantage upper bounds, rather than the corresponding query lower bounds.

```

Algorithm  $\mathcal{A}_{1LD-s}[L_0, L_1]$ 
 $L' \leftarrow \text{Inj}(D, R); L'_0 \parallel L'_1 \leftarrow L$ 
 $i \leftarrow 0; L_0^0 \leftarrow L_0; L_1^0 \leftarrow L_1$ 
Repeat
   $i \leftarrow i + 1$ 
   $L_{0,l} \parallel L_{0,r} \leftarrow L_0^{i-1}$ 
   $L_{1,l} \parallel L_{1,r} \leftarrow L_1^{i-1}$ 
   $(j^*, k^*) \leftarrow (r, r)$ 
  For  $(j, k) \in \{(l, l), (l, r), (r, l)\}$  do
    //  $f \square g(x) = f(x)$  for  $x \in \text{Dom}(f)$  else  $g(x)$ 
     $b \leftarrow \mathcal{A}_{1LD-d}[L_{0,j} \square L'_0, L_{1,k} \square L'_1]$ 
    If  $b = 1$  then  $(j^*, k^*) \leftarrow (j, k)$ 
   $L_0^i \leftarrow L_{0,j^*}; L_1^i \leftarrow L_{1,k^*}$ 
Until  $|\text{Dom}(L_0^i)| = |\text{Dom}(L_1^i)| = 1$ 
Pick  $(x, y) \in \text{Dom}(L_0^i) \times \text{Dom}(L_1^i)$ 
Return  $(x, y)$ 

```

Fig. 10. Reduction algorithm \mathcal{A}_{1LD-s} for Lemma 3. For notational convenience we write the two lists as separate input, rather than combined into a single list.

Lemma 2 (1ED-s hard \Rightarrow 1LD-s hard). *Let D be even. If \mathcal{A}_{1LD-s} is a quantum algorithm for 1LD-s making at most q queries to its oracle, then there is an algorithm \mathcal{A}_{1ED-s} (described in the proof) such that $\text{Adv}_{D,R}^{1LD-s}(\mathcal{A}_{1LD-s}) \leq 2\text{Adv}_{D,R}^{1ED-s}(\mathcal{A}_{1ED-s})$. Algorithm \mathcal{A}_{1ED-s} makes at most q queries to its oracle.*

Proof. On input a list $L : [D] \rightarrow [R]$, the algorithm \mathcal{A}_{1ED-s} will pick a random permutation $\pi : [D] \rightarrow [D]$. It runs $\mathcal{A}_{1LD-s}[L \circ \pi]$ and then, on receiving output $(x, y) \in [D/2]^2$, returns $(\pi^{-1}(x), \pi^{-1}(y + D/2))$. The permutation π serves the role of splitting L into two sublists for \mathcal{A}_{1LD-s} at random. As long as the original collision of L doesn't end up being put into the same sublist (which has probability less than 1/2), \mathcal{A}_{1LD-s} will be run on a valid 1LD-s instance and \mathcal{A}_{1ED-s} will succeed whenever \mathcal{A}_{1LD-s} does. The claim follows. \square

Lemma 3 (1LD-s hard \Rightarrow 1LD-d hard). *Let D be a power of two. If \mathcal{A}_{1LD} is a quantum algorithm for 1LD making at most q queries to its oracle, then there is an 1LD-s algorithm \mathcal{A}_{1LD-s} (described in the proof) such that*

$$\text{Adv}_{D,R}^{1LD-s}(\mathcal{A}_{1LD-s}) \geq 1 - D^2/R - 1.5(\lg D - 2)(1 - \text{Adv}_{D,R}^{1LD-d}(\mathcal{A}_{1LD-d})).$$

Algorithm \mathcal{A}_{1LD-s} makes at most $3q \lg D$ queries to its oracle.

This theorem's bound is vacuous if $\text{Adv}_{D,R}^{1LD-d}(\mathcal{A}_{1LD-d})$ is too small. When applying the result we will have obtained \mathcal{A}_{1LD-d} by amplifying the advantage of another adversary to ensure it is sufficiently large.

Proof. The algorithm \mathcal{A}_{1LD-s} is given in Fig. 10. The intuition behind this algorithm is as follows. It wants to use the decision algorithm \mathcal{A}_{1LD-d} to perform a binary search to find the overlap between L_0 and L_1 . It runs for $\lg D/2 - 1$ rounds. In the i -th round, it splits the current left list L_0^{i-1} into two sublists $L_{0,l}, L_{0,r}$ and the current right lists L_1^{i-1} into two sublists $L_{1,l}, L_{1,r}$.¹⁵ If L_0^{i-1} and L_1^{i-1} have an element in common, then one of the pairs of sublists $L_{0,j}$ and $L_{1,k}$ for $j, k \in \{l, r\}$ must have an element in common. The decision algorithm \mathcal{A}_{1LD-d} is run on different pairs to determine which contains the overlap. We recurse with chosen pair, until we are left with lists that have singleton domains at which point we presume the entries therein give the overlap.

¹⁵ In code, $f \parallel g \leftarrow h$ for h with domain $\text{Dom}(h) = \{n, n+1, \dots, m\}$ defines f to be the restriction of h to domain $\text{Dom}(f) = \{n, n+1, \dots, \lfloor (n+m)/2 \rfloor\}$ and g to be the restriction of h to domain $\text{Dom}(g) = \text{Dom}(h) \setminus \text{Dom}(f)$.

Because we need to run the decision algorithm $\mathcal{A}_{\text{1LD-d}}$ on fixed size inputs we pad the sublists to be of a fixed size using lists L'_0 and L'_1 (the two halves of an injection L' that we sampled locally). As long as the image of L' does not overlap with the images of L_0 and L_1 , this padding does not introduce any additional elements repetitions between or within the list input to $\mathcal{A}_{\text{1LD-d}}$. By a union bound, overlaps occurs with probability at most D^2/R .

Conditioned on such overlaps not occurring, $\mathcal{A}_{\text{1LD-s}}$ will output the correct result if $\mathcal{A}'_{\text{1LD-d}}$ always answers correctly. To bound the probability of $\mathcal{A}'_{\text{1LD-d}}$ erring we can note that it is run $3 \cdot (\lg D - 2)$ times and, each time, has at most a $(1 - \text{Adv}_{D,R}^{\text{1LD-d}}(\mathcal{A}'_{\text{1LD-d}}))/2$ chance of error and apply a union bound.

Put together we have

$$\text{Adv}_{D,R}^{\text{1LD-s}}(\mathcal{A}_{\text{1LD-s}}) \geq 1 - D^2/R - 1.5(\lg D - 2)(1 - \text{Adv}_{D,R}^{\text{1LD-d}}(\mathcal{A}_{\text{1LD-d}})).$$

That $\mathcal{A}_{\text{1LD-s}}$ makes at most $3q(\lg D - 2)$ oracle queries is clear. \square

Lemma 4 (PB-d hard \Rightarrow PB hard). *Let PB be any problem. Suppose \mathcal{A}_{PB} is a quantum algorithm for PB making at most q queries to its oracle, with $\text{Adv}_{D,R}^{\text{PB}}(\mathcal{A}_{\text{PB}}) = \delta > 0$. Then for any $t \in \mathbb{N}$ there is an algorithm $\mathcal{A}_{\text{PB-d}}$ (described in the proof) such that $\text{Adv}_{D,R}^{\text{PB-d}}(\mathcal{A}_{\text{PB-d}}) > 1 - 2/2^t$. Algorithm $\mathcal{A}_{\text{PB-d}}$ makes $q \cdot \lceil 4.5(t+1) \ln 2/\delta^2 \rceil$ queries to its oracle.*

Proof. For compactness, let $p^1 = P_{D,R}^1(\mathcal{A}_{\text{PB}})$ denote the minimum probability that \mathcal{A}_{PB} outputs 1 on instances in the language, $p^0 = P_{D,R}^0(\mathcal{A}_{\text{PB}})$ denote the maximum probability that \mathcal{A}_{PB} outputs 1 on instances not in the language, and $\delta = \text{Adv}_{D,R}^{\text{PB}}(\mathcal{A}_{\text{PB}}) = p^1 - p^0$. We define $\mathcal{A}_{\text{PB-d}}$ to be the algorithm that, on input L , runs n independent copies of $\mathcal{A}_{\text{PB}}[L]$ (with $n = \lceil 4.5(t+1) * \ln 2/\delta^2 \rceil$) and calculates the average p of all the values output by $\mathcal{A}_{\text{PB}}[L]$. (Think of this as an estimate of the probability \mathcal{A}_{PB} outputs 1 on this input.) If $p < p^0 + \delta/2$, $\mathcal{A}_{\text{PB-d}}$ outputs 0, otherwise it outputs 1.

Let X_i denote the output of the i -th execution of $\mathcal{A}_{\text{PB-d}}$. An inequality of Hoeffding [?] bounds how far the average of independent random variables $0 \leq X_i \leq 1$ can differ from the expectation by

$$\Pr \left[\left| \sum_{i=1}^n X_i/n - \mathbb{E} \left[\sum_{i=1}^n X_i/n \right] \right| > \varepsilon \right] < 2e^{-2\varepsilon^2 n}$$

for any $\varepsilon > 0$. Let p' denote the expected value of p (which is also the expected value of X_i). If L is in the language, then $p^1 \leq p$. Otherwise $p^0 \geq p$. In either case, $\mathcal{A}_{\text{PB-d}}$ will output the correct answer if p does not differ from p' by more than $\delta/3$ (because this is strictly less than $\delta/2$). Applying the above inequality tells us that $\Pr[|p - p'| > \delta/3] < 2e^{-2\delta^2 n/9} \leq 2^{-t}$. (The value of n was chosen to make this last inequality hold.)

Hence $P_{D,R}^1(\mathcal{A}_{\text{PB-d}}) > 1 - 2^{-t}$, $P_{D,R}^0(\mathcal{A}_{\text{PB-d}}) < 2^{-t}$, and $\text{Adv}_{D,R}^{\text{PB-d}}(\mathcal{A}_{\text{PB-d}}) > 1 - 2/2^t$. The bound on $\mathcal{A}_{\text{PB-d}}$'s number of queries is clear. \square

4.3 Proof of Theorem 6

Let \mathcal{A}_{1LD} be our given list disjointness adversary which makes q oracle queries and has advantage $\delta = \text{Adv}_{D,R}^{\text{1LD}}(\mathcal{A}_{\text{1LD}}) > 0$. If we apply Lemma 4 with $t = \lg(12 \lg D)$, then we get an adversary $\mathcal{A}_{\text{1LD-d}}$ which makes $q_{\text{1LD-d}} = q \cdot \lceil 4.5(t+1) \ln 2/\delta^2 \rceil < (10q \lg \lg D)/\delta^2$ oracle queries. (We used here the assumption $D \geq 32$ to simplify constants.) This adversary has advantage

$$\text{Adv}_{D,3D^2}^{\text{1LD-d}}(\mathcal{A}_{\text{1LD-d}}) > 1 - 2/2^t = 1 - 1/(6 \lg(D)).$$

Next applying Lemma 3, gives us $\mathcal{A}_{\text{1LD-s}}$ which makes fewer than $q_{\text{1LD-s}} = 30q(\lg D)(\lg \lg D)/\delta^2$ oracle queries and has advantage

$$\begin{aligned} \text{Adv}_{D,R}^{\text{1LD-s}}(\mathcal{A}_{\text{1LD-s}}) &\geq 1 - D^2/R - 1.5(\lg D - 2)(1 - \text{Adv}_{D,R}^{\text{1LD-d}}(\mathcal{A}_{\text{1LD-d}})) \\ &> 1 - 1/3 - 1.5(\lg D - 2)(6 \lg D)^{-1} > 1 - 1/3 - 1/4 = 5/12. \end{aligned}$$

Lemmas 1 and 2 together bound this advantage from the other direction. In particular, we get that

$$5/12 < \text{Adv}_{D,R}^{\text{1LD-s}}(\mathcal{A}_{\text{1LD-s}}) \leq 18 \cdot (q_{\text{1LD-s}} + 2)^3 / D^2.$$

Using the assumption $D \geq 32$ we can bound $q_{\text{1LD-s}} + 2$ by $31q(\lg D)(\lg \lg D)/\delta^2$. Plugging this in and solving for δ gives our claimed bound of

$$\delta < 11 \sqrt[6]{(q \cdot \lg D \cdot \lg \lg D)^3 / D^2}.$$

Acknowledgements

Joseph Jaeger and Stefano Tessaro were partially supported by NSF grants CNS-1930117 (CAREER), CNS-1926324, CNS-2026774, a Sloan Research Fellowship, and a JP Morgan Faculty Award. Joseph Jaeger’s work done while at the University of Washington. Fang Song thanks Robin Kothari for helpful discussion on the element distinctness problem. Fang Song was supported by NSF grants CCF-2041841, CCF-2042414, and CCF-2054758 (CAREER).

References

1. S. Aaronson and Y. Shi. Quantum lower bounds for the collision and the element distinctness problems. *Journal of the ACM (JACM)*, 51(4):595–605, 2004.
2. A. Ambainis. Quantum walk algorithm for element distinctness. *SIAM Journal on Computing*, 37(1):210–239, 2007.
3. A. Ambainis, M. Hamburg, and D. Unruh. Quantum security proofs using semi-classical oracles. In A. Boldyreva and D. Micciancio, editors, *CRYPTO 2019, Part II*, volume 11693 of *LNCS*, pages 269–295. Springer, Heidelberg, Aug. 2019.
4. N. Bindel, M. Hamburg, K. Hövelmanns, A. Hülsing, and E. Persichetti. Tighter proofs of cca security in the quantum random oracle model. In *Theory of Cryptography Conference*, pages 61–90. Springer, 2019.
5. X. Bonnetain, A. Hosoyamada, M. Naya-Plasencia, Y. Sasaki, and A. Schrottenloher. Quantum attacks without superposition queries: the offline Simon’s algorithm. In *Advances in Cryptology – ASIACRYPT*, pages 552–583. Springer, 2019.
6. C. Chevalier, E. Ebrahimi, and Q.-H. Vu. On security notions for encryption in a quantum world. Cryptology ePrint Archive, Report 2020/237, 2020. <https://ia.cr/2020/237>.
7. J. Czajkowski. Quantum indistinguishability of sha-3. Cryptology ePrint Archive, 2021. <http://eprint.iacr.org/2021/192>.
8. J. Czajkowski, C. Majenz, C. Schaffner, and S. Zur. Quantum lazy sampling and game-playing proofs for quantum indistinguishability. Cryptology ePrint Archive, Report 2019/428, 2019. <https://eprint.iacr.org/2019/428>.
9. Ö. Dagdelen, M. Fischlin, and T. Gagliardoni. The Fiat-Shamir transformation in a quantum world. In K. Sako and P. Sarkar, editors, *ASIACRYPT 2013, Part II*, volume 8270 of *LNCS*, pages 62–81. Springer, Heidelberg, Dec. 2013.
10. W. Diffie and M. E. Hellman. Exhaustive cryptanalysis of the nbs data encryption standard. *IEEE Computer*, 10(6), 1977.
11. S. Even and Y. Mansour. A construction of a cipher from a single pseudorandom permutation. *Journal of Cryptology*, 10(3):151–162, June 1997.
12. L. K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, STOC ’96, pages 212–219, New York, NY, USA, 1996. Association for Computing Machinery.
13. A. Hosoyamada and Y. Sasaki. Cryptanalysis against symmetric-key schemes with online classical queries and offline quantum computations. In *Cryptographers’ Track at the RSA Conference*, pages 198–218. Springer, 2018.
14. K. Hövelmanns, E. Kiltz, S. Schäge, and D. Unruh. Generic authenticated key exchange in the quantum random oracle model. In *IACR International Conference on Public-Key Cryptography*, pages 389–422. Springer, 2020.
15. J. Jaeger, F. Song, and S. Tessaro. Quantum key-length extension. Cryptology ePrint Archive, 2021. <http://eprint.iacr.org/2021/579>.
16. M. Kaplan. Quantum attacks against iterated block ciphers. *arXiv preprint arXiv:1410.1434*, 2014.

17. M. Kaplan, G. Leurent, A. Leverrier, and M. Naya-Plasencia. Quantum differential and linear cryptanalysis. *IACR Trans. Symm. Cryptol.*, 2016(1):71–94, 2016. <https://tosc.iacr.org/index.php/ToSC/article/view/536>.
18. S. Katsumata, S. Yamada, and T. Yamakawa. Tighter security proofs for GPV-IBE in the quantum random oracle model. *Journal of Cryptology*, 34(1):1–46, 2021.
19. J. Kilian and P. Rogaway. How to protect DES against exhaustive key search. In N. Kobitz, editor, *CRYPTO'96*, volume 1109 of *LNCS*, pages 252–267. Springer, Heidelberg, Aug. 1996.
20. J. Kilian and P. Rogaway. How to protect DES against exhaustive key search (an analysis of DESX). *Journal of Cryptology*, 14(1):17–35, Jan. 2001.
21. E. Kiltz, V. Lyubashevsky, and C. Schaffner. A concrete treatment of Fiat-Shamir signatures in the quantum random-oracle model. In J. B. Nielsen and V. Rijmen, editors, *EUROCRYPT 2018, Part III*, volume 10822 of *LNCS*, pages 552–586. Springer, Heidelberg, Apr. / May 2018.
22. V. Kuchta, A. Sakzad, D. Stehlé, R. Steinfeld, and S.-F. Sun. Measure-rewind-measure: tighter quantum random oracle model proofs for one-way to hiding and cca security. In *Advances in Cryptology – EUROCRYPT 2020*, pages 703–728. Springer, 2020.
23. H. Kuwakado and M. Morii. Security on the quantum-type even-mansour cipher. In *2012 International Symposium on Information Theory and its Applications*, pages 312–316, 2012.
24. G. Leander and A. May. Grover meets simon - quantumly attacking the FX-construction. In T. Takagi and T. Peyrin, editors, *ASIACRYPT 2017, Part II*, volume 10625 of *LNCS*, pages 161–178. Springer, Heidelberg, Dec. 2017.
25. B. Mennink and A. Szepeieniec. XOR of PRPs in a quantum world. In T. Lange and T. Takagi, editors, *Post-Quantum Cryptography - 8th International Workshop, PQCrypto 2017*, pages 367–383. Springer, Heidelberg, 2017.
26. R. C. Merkle and M. E. Hellman. On the security of multiple encryption. *Communications of the ACM*, 24(7):465–467, 1981.
27. A. Rosmanis. Tight bounds for inverting permutations via compressed oracle arguments. *arXiv preprint arXiv:2103.08975*, 2021.
28. P. W. Shor. Algorithms for quantum computation: Discrete logarithms and factoring. In *35th FOCS*, pages 124–134. IEEE Computer Society Press, Nov. 1994.
29. D. R. Simon. On the power of quantum computation. *SIAM journal on computing*, 26(5):1474–1483, 1997.
30. E. E. Targhi and D. Unruh. Post-quantum security of the Fujisaki-Okamoto and OAEP transforms. In M. Hirt and A. D. Smith, editors, *TCC 2016-B, Part II*, volume 9986 of *LNCS*, pages 192–216. Springer, Heidelberg, Oct. / Nov. 2016.
31. S. Tessaro and A. Thiruvengadam. Provable time-memory trade-offs: Symmetric cryptography against memory-bounded adversaries. In A. Beimel and S. Dziembowski, editors, *TCC 2018, Part I*, volume 11239 of *LNCS*, pages 3–32. Springer, Heidelberg, Nov. 2018.
32. D. Unruh. Revocable quantum timed-release encryption. *J. ACM*, 62(6), Dec. 2015.
33. M. Zhandry. How to construct quantum random functions. In *53rd FOCS*, pages 679–687. IEEE Computer Society Press, Oct. 2012.
34. M. Zhandry. A note on the quantum collision and set equality problems. *Quantum Information and Computation*, 15(7& 8), 2015.
35. M. Zhandry. How to record quantum queries, and applications to quantum indistinguishability. In A. Boldyreva and D. Micciancio, editors, *CRYPTO 2019, Part II*, volume 11693 of *LNCS*, pages 239–268. Springer, Heidelberg, Aug. 2019.

A Probabilistic Analysis for Proof of Theorem 2

In this appendix we give a detailed probability analysis of Claim (i) from the proof of Theorem 2.

The view of \mathcal{A} when run by \mathcal{A}' in $\mathbf{G}_{\mathbf{D},1}^{\text{dist}}$ is determined by f_1 , f_1^{-1} , and $z = (T, T^{-1})$. These are chosen such that $T[M_i] = f_1(K_1, M_i \oplus K_2) \oplus K_2$ for $i = 1, \dots, q'$ and $T^{-1}[Y_i] = f_1^{-1}(K_1, Y_i \oplus K_2) \oplus K_2$ for $i = q' + 1, \dots, q$. Consequently to ensure this matches the view from $\mathbf{G}_{\mathbf{F},1}^{\text{sprp}}$ we need to show that \mathbf{D} chooses (f_1, K_1, K_2) uniformly from $\text{lcs}(k, n) \times \{0, 1\}^k \times \{0, 1\}^n$. It is clear that K_1 and K_2 are uniform. Furthermore, $f_1(K, \cdot) = f_0(K, \cdot)$ for $K \neq K_1$ so these are sampled correctly. Hence we can think of these values as fixed and argue that the distribution induced over $f_1(K_1, \cdot)$ is uniform over all permutations on $\{0, 1\}^n$.

Let $\mathbf{f} \in \text{lcs}(k, n)$, $N = 2^n$, and E denote the event that $f_1(K_1, \cdot) = \mathbf{f}(\cdot)$. Let E_1 denote the event that $T[M_i] = \mathbf{f}(M_i \oplus K_2) \oplus K_2$ for $i = 1, \dots, q$ after Step 1. The second for loop of Step 3 programs f_1 to satisfy $f_1(K_1, M_i \oplus K_2) = T[M_i] \oplus K_2$ so E_1 is a necessary condition for E . Note that E_1 requires Q values to have been sampled correctly in Step 1's for loops where $Q = |\{\mathbf{f}(M_i \oplus K_2) \oplus K_2 : 1 \leq i \leq q'\} \cup \{\mathbf{f}(Y_i \oplus K_2) \oplus K_2 : q' \leq i \leq q\}|$. Hence,

$$\Pr[E] = \Pr[E|E_1] \cdot \frac{(N-Q)!}{N!}.$$

Let E_2 denote the event that Step 2 samples an f_0 which is consistent with \mathbf{f} . That is to say, that f_0 is chosen such that $\mathbf{f}(\mathcal{I} \cup \mathcal{I}') = \mathcal{O} \cup \mathcal{O}'$ and $\mathbf{f}(x) = f_0(K_1, x)$ for $x \notin \mathcal{I} \cup \mathcal{I}'$. The first for loop in Step 3 programs $f_1(K_1, x) = f_0(K_1, x)$ for $x \notin \mathcal{I} \cup \mathcal{I}'$. The second and third for loop in Step 3 programs $f_1(K_1, \mathcal{I} \cup \mathcal{I}')$ to have values in $\mathcal{O} \cup \mathcal{O}'$ so E_2 is a necessary condition for E . Let $M(f_0) = |\mathcal{I} \setminus \mathcal{I}'| = |\mathcal{O} \setminus \mathcal{O}'|$ and let E_2^m denote the event that $M(f_0) = m$.

$$\Pr[E|E_1] = \sum_m \Pr[E_2^m|E_1] \cdot \Pr[E_2|E_1, E_2^m] \cdot \Pr[E|E_1, E_2^m, E_2].$$

We can think of Step 2 as lazily sampling $f_0(K_1, \cdot)$ by first sampling $f_0(K_1, x)$ for $x \in \mathcal{I}$, then sampling $f_0^{-1}(K_1, y)$ for $y \in \mathcal{O} \setminus \mathcal{O}'$, and then sampling $f_0(K_1, x)$ for $x \notin \mathcal{I} \cup \mathcal{I}'$. The event E_2 requires that the m values sampled in this second step are the elements of $\mathbf{f}^{-1}(\mathcal{O}' \setminus \mathcal{O})$ and that on the third step $\mathbf{f}(x)$ is sampled for $f_0(K_1, x)$ for each $x \notin \mathcal{I} \cup \mathcal{I}'$. Hence,

$$\Pr[E_2|E_1, E_2^m] = \frac{1}{\binom{N-Q}{m}} \cdot \frac{1}{(N-Q-m)!}.$$

For E to occur (conditioned on E_1 , E_2^m , and E_2) the third for loop in Step 3 must sample m values correctly, so $\Pr[E|E_1, E_2^m, E_2] = 1/m!$. Putting everything together we have

$$\Pr[E] = \frac{(N-Q)!}{N!} \sum_m \Pr[E_2^m|E_1] \cdot \frac{m! \cdot (N-Q-m)!}{(N-Q)!} \cdot \frac{1}{(N-Q-m)!} \cdot \frac{1}{m!} = \frac{1}{N!}.$$

So f_1 is uniformly distributed, as desired.