Grover Meets Simon – Quantumly Attacking the FX-construction

Gregor Leander[(✉)] and Alexander May

Horst Görtz Institute for IT-Security, Faculty of Mathematics,
Ruhr-University Bochum, Bochum, Germany
{gregor.leander,alex.may}@rub.de

Abstract. Using whitening keys is a well understood mean of increasing the key-length of any given cipher. Especially as it is known ever since Grover's seminal work that the effective key-length is reduced by a factor of two when considering quantum adversaries, it seems tempting to use this simple and elegant way of extending the key-length of a given cipher to increase the resistance against quantum adversaries. However, as we show in this work, using whitening keys does not increase the security in the quantum-CPA setting significantly. For this we present a quantum algorithm that breaks the construction with whitening keys in essentially the same time complexity as Grover's original algorithm breaks the underlying block cipher. Technically this result is based on the combination of the quantum algorithms of Grover and Simon for the first time in the cryptographic setting.

Keywords: Symmetric cryptography · Quantum attacks · Grover's algorithm · Simon's algorithm · FX-construction

1 Introduction

The existence of sufficiently large quantum computers has a major impact on the security of many cryptographic schemes we are using today. In particular the seminal work of Shor [24] showed that such computers would allow to factor numbers and compute discrete logs in abelian groups in polynomial time. As almost all public key schemes currently in use are build upon the assumption that those problems are intractable, the existence of quantum computers would seriously compromise the security of most of our digital communication.

This situation has triggered a whole new line of research, namely post-quantum cryptography (or quantum-resistant cryptography), that aims at developing new cryptographic primitives that would (hopefully) withstand even attackers that are equipped with quantum computers. Recently, NIST has announced a competition to eventually standardize one or several quantum-resistant public-key cryptographic algorithms [22], underlining the importance of the problem. Indeed, as NIST points out in their call for candidates, the roll out of new cryptographic schemes is a long time process and it is therefore

important to start this switch to quantum resistant cryptography well before quantum computers are actually available.[1]

In the case of symmetric cryptography, the situation seems less critical – but is also significantly less studied. For almost 20 years time, it was believed, that the only advantage an attacker would have by using a quantum computer when attacking symmetric cryptography is due to Grover's algorithm [11] for speeding up brute force search. Indeed, Grover's algorithm reduces the effective key-length of any cryptographic scheme, and thus in particular of any block-cipher, by a factor of two.

$$m \longrightarrow \boxed{E_k} \longrightarrow c$$

Given an m bit key, Grover's algorithm allows to recover the key using $\mathcal{O}(2^{m/2})$ quantum steps.

To counter that attack, it seems to be sufficient to just double the key-length of the block cipher to achieve the same security against quantum attackers. For doing so, the two main generic options are using either whitening keys or using multiple encryptions.

More recently, starting with the initial works by Kuwakado and Morii [17,18] and followed by the work by Kaplan et al. [13], it was stressed that Grover's algorithm might not be the only threat for symmetric cryptography. In particular, Kuwakado and Morii showed that the so-called Even-Mansour construction can be broken in polynomial time in the quantum CPA-setting. In this setting, the attacker is allowed to make quantum queries, that is queries to the encryption function in quantum superposition. The quantum CPA setting was first defined by Boneh, Zhandry in [5], and further intensively discussed in Kaplan et al. [13] and Anand et al. [4].

The Even-Mansour construction [10] consists of a public permutation P on n bits and two secret keys k_1 and k_2 that are used as pre- (resp. post-) whitening keys for the encryption $\mathrm{Enc}_{EM}(m)$ of some message m.

$$m \longrightarrow \overset{k_1}{\oplus} \longrightarrow \boxed{P} \longrightarrow \overset{k_2}{\oplus} \longrightarrow c$$

In a nutshell, Even and Mansour proved that, in an ideal world, where P is randomly chosen amongst all possible permutations, an attacker's advantage of distinguishing between the encryption and a random permutation is bounded by $q^2/2^n$ where q is the number of queries to P or to the encryption/decryption oracle. However, in the quantum CPA-model the scheme is completely insecure. The main idea of [18] was to consider the function

$$f(x) := \mathrm{Enc}_{EM}(x) + P(x) = P(x + k_1) + k_2 + P(x),$$

---

[1] Note that NIST states that it is "primarily concerned with attacks that use classical (rather than quantum) queries to the decryption oracle or other private-key functionality."

where $+$ is the bitwise XOR.

As this function fulfills $f(x) = f(x + k_1)$ for all x, one can use Simon's quantum algorithm [7,25], that allows to compute the unknown period k_1 of function f in linear time. Once k_1 is computed, computing k_2 is trivial even on a classical computer. It should be pointed out that Kaplan et al. [13] and Santoli, Schaffner [23] solved the technical issue of dealing with a function that does not fulfill Simon's promise, namely that $f(x) = f(y)$ iff $y \in \{x, x + k_1\}$, see Sect. 2 for more details.

The same idea was then used by Kaplan et al. [13] (and independently in [23]) to construct polynomial time quantum-CPA attacks on many modes of operations. Kaplan et al. further showed how slide attacks can profit from using a quantum computer.

The natural question that arises from the attacks on a generic cipher using Grover's algorithm and the attack on the Even-Mansour scheme using Simon's algorithm is the following: How secure is the FX construction against quantum adversaries?

This construction, proposed by Killian and Rogaway in [15,16], is an elegant way of extending the key-length of a given block cipher and is the natural combination of the Even-Mansour construction and a generic cipher. For this, we assume we are given a (secure) block cipher E, encrypting n bit messages under an m bit key k_0, and we introduce two more n bit keys k_1 and k_2 as pre- and post-whitening keys. The new block cipher is given as

$$\text{Enc}(x) = E_{k_0}(x + k_1) + k_2.$$

From an efficiency point of view, the overhead of this modification is negligible. Moreover, in an idealized model, one can prove that (using classical computers) in order to attack the FX construction scheme, the success probability of an attacker is bounded by $\frac{q^2}{2^{n+m}}$, where q is the number of queries to the encryption scheme and to the underlying block cipher.

Initially, when considering Grover's algorithm only, this scheme seems to provide significantly more resistance against quantum computers, since now $(k_0, k_1, k_2) \in \mathbb{F}_2^{m+2n}$ define the key space. Moreover, Simon's algorithm does not apply either, as the function $\text{Enc}(x) + E_k(x)$ is periodic only for the correct guess of $k = k_0$.

Our Contribution

As the main result of our paper, we show that the FX construction as described above can be broken in the quantum-CPA model in basically the same time as the scheme without whitening keys, namely in $\mathcal{O}(m + n) \cdot 2^{m/2}$ quantum steps. Thus, using whitening keys does not help at all against quantum-CPA attackers.

Technically we have to combine the quantum algorithms of Simon and Grover for this attack. Thus, in contrast to most of the other works mentioned above, we actually have to define a new quantum algorithm, rather than applying known ones to new problems. While merging Simon and Grover might seem straightforward on a high level, the main technical obstacle is that in its original form, Simon's algorithm extracts information on the secret period bit by bit while Grover's algorithm, or more generally quantum amplitude amplification [7] inherently requires all the information to be available at once. We solve this issue by running several instances of Simon's quantum circuit in parallel, which in turn comes at the price of a linear growth of the size of the quantum computer. Furthermore, we postpone the measurements in Simon's algorithm to the very end of our entire quantum algorithm using the general *deferred measurement principle* of quantum computation.

DESX, PRINCE and PRIDE. To illustrate our results on actual ciphers, we like to mention that our work implies, among others, heuristic key recovery attacks on DESX, proposed by Rivest in 1984 and formally treated in [15] as well as on PRINCE [6] and PRIDE [2]. DESX, using a 56 bit internal key and two 64 bit whitening keys, can thus be attacked in the quantum-CPA model with complexity roughly 2^{28}, while PRINCE and PRIDE, both using a 64 bit internal key and two 64 bit whitening keys, can be attacked in the quantum-CPA model with complexity roughly 2^{32}.

We like to point out that, in the analysis of the success probabilities for our attack, we actually assume that the underlying functions are random functions, which is clearly not the case for the mentioned ciphers. However, it would be very surprising if this heuristic would fail for any reasonable block cipher.

Related Work

Besides the works on Simon's algorithm already mentioned above, we like to highlight in particular the work of Kaplan [12] who shows that multiple encryption is significantly weaker in the quantum setting than in the classical setting. This work is based on quantum random walks (cf e.g. [3]).

Together with our result presented here this implies that the two most common methods for extending the key-length are far from being optimal in a quantum-CPA setting. Finally, a very recent work [1] at EUROCRYPT 2017 already explores other means of mixing the key into the state that are (potentially) more resistant against quantum attacks. It remains to be seen if our idea of combining Grover with Simon, or related algorithms, allow to attack those recent proposals as well.

Another interesting approach is to use quantum computers to improve classical attacks on block ciphers, such as linear or differential cryptanalysis. This approach has been treated for the first time in [14].

Organization of the Paper

Before we formulate and prove our main theorem (Theorem 2) in Sect. 3 in all technical details, we outline the high level idea in Sect. 2. The latter section also contains some technical lemmata that are needed in the proof of our main result. We conclude by discussing some topics for future work in Sect. 4.

2   Main Ideas of Our Quantum Algorithm

Throughout the paper, we assume that the reader is familiar with the basics of quantum algorithms, although this section is supposed to be comprehensible without deeper quantum knowledge. For a comprehensive introduction into quantum algorithms we recommend the textbooks of Mermin [20] and Lipton, Regan [19].

Recall that we are attacking the FX-construction $\mathrm{Enc}(x) = E_{k_0}(x + k_1) + k_2$. Let us look at the function

$$f(k, x) = \mathrm{Enc}(x) + E_k(x) = E_{k_0}(x + k_1) + k_2 + E_k(x).$$

For the correct key $k = k_0$, we have $f(k, x) = f(k, x + k_1)$ for all x, and thus $f(k_0, \cdot)$ has period k_1 in its second argument. However, for $k \neq k_0$ the function $f(k, \cdot)$ is not periodic with high probability.

Main Idea. We define a Grover search over $k \in \mathbb{F}_2^m$, where we test for every $f(k, \cdot)$ periodicity via Simon's algorithm. Thus, we have Grover as an outer loop with running time roughly $2^{\frac{m}{2}}$, and Simon as an inner loop with polynomial complexity.

Classically, we could define an outer loop that guesses $k = k_0$ correctly with probability $p = 2^{-m}$. This would require an expected $\frac{1}{p} = 2^m$ number of iterations until we hit the correct key. Hence, each iteration roughly increases the success probability linearly by an amount of $\frac{1}{p}$. By contrast, in a quantum setting each iteration roughly increases the *amplitude* of success by a constant. Since the probabilities are the square of their respective amplitudes, we only have to repeat approximately $\sqrt{\frac{1}{p}} = 2^{\frac{m}{2}}$ times.

This process is called *amplitude amplification* and can be seen as a natural generalization of the original Grover search [11]. The results of amplitude amplification are more accurately captured in the following theorem by Brassard et al. [8, Theorem 2].

Theorem 1 (Brassard, Hoyer, Mosca and Tapp). *Let $\mathcal{A}$ be any quantum algorithm on q qubits that uses no measurement. Let $\mathcal{B} : \mathbb{F}_2^q \rightarrow \{0, 1\}$ be a function that classifies outcomes of $\mathcal{A}$ as good or bad. Let $p > 0$ be the initial success probability that a measurement of $\mathcal{A}|\mathbf{0}\rangle$ is good. Set $k = \lceil \frac{\pi}{4\theta} \rceil$, where θ is*

defined via $\sin^2(\theta) = p$. *Moreover, define the unitary operator* $Q = -\mathcal{A}S_0\mathcal{A}^{-1}S_{\mathcal{B}}$, *where the operator* $S_{\mathcal{B}}$ *changes the sign of the good state*

$$|x\rangle \mapsto \begin{cases} -|x\rangle & \text{if } \mathcal{B}(x) = 1 \\ |x\rangle & \text{if } \mathcal{B}(x) = 0 \end{cases},$$

while S_0 *changes the sign of the amplitude only for the zero state* $|\mathbf{0}\rangle$.

Then after the computation of $Q^k\mathcal{A}|\mathbf{0}\rangle$, *a measurement yields good with probability at least* $\max\{1 - p, p\}$.

Let us describe in a high-level fashion the process of amplitude amplification behind Theorem 1. The classifier $\mathcal{B}$ partitions the Hilbert space $\mathcal{H}$ of our quantum system in a direct sum of two orthogonal subspaces, the good subspace and the bad subspace. The good one is the subspace defined by all basis states $|x\rangle$ with $\mathcal{B}(x) = 1$, the bad one is its orthogonal complement in $\mathcal{H}$.

Let $|\psi\rangle = \mathcal{A}|\mathbf{0}\rangle$ be the initial vector, and denote by $|\psi_1\rangle$, $|\psi_0\rangle$ its projection on the good and the bad subspace, respectively. Now look at the two-dimensional plane $\mathcal{H}_\psi$ spanned by $|\psi_1\rangle$, $|\psi_0\rangle$. In $\mathcal{H}_\psi$, the state $|\psi\rangle = \mathcal{A}|\mathbf{0}\rangle$ has angle θ (defined by $\sin^2(\theta) = p$) with the bad subspace. Each iteration via Q increases this angle by 2θ via the two reflections S_0 and $S_{\mathcal{B}}$. Thus, after k iterations we have angle $(2k + 1)\theta$. If this angle roughly equals $\frac{\pi}{2}$, then the resulting state is almost co-linear with the good subspace, and thus a measurement yields a good vector with high probability. This explains the choice of the number of iterations $k \approx \frac{\pi}{4\theta}$ in Theorem 1.

Let us now assume that $p = 2^{-m}$ is the probability of guessing the correct key in the FX-construction. Then $\theta = \arcsin(2^{-\frac{m}{2}}) \approx 2^{-\frac{m}{2}}$, since $\arcsin(x) \approx x$ for small x. This implies $k = \Theta(2^{-\frac{m}{2}})$, as desired. Moreover, by Theorem 1 we obtain overwhelming success probability $1 - 2^{-m}$.

Ideally, we would choose $\mathcal{A}$ as Simon's algorithm and directly apply Theorem 1 for our setting. Although Theorem 1 excludes the use of measurements, while Simon's algorithm uses measurements to extract information about the period, this slight technical problem can be easily resolved by the quantum principle of deferred measurement that postpones all measurements until the very end of the computation.

However, we still have to resolve the following problems for an application of Theorem 1.

1. **Classifier.** We need to define some classifier $\mathcal{B}$ that identifies states as good iff they correspond to the correct key $k = k_0$. However, we do not see any way of efficiently checking correctness of k_0 without the knowledge of k_1. Simon's algorithm *iteratively* computes information about the period k_1 in a bit-wise fashion, where we need a *complete* candidate k_1' for the period in order to classify states as good or bad.
2. **Simon's promise.** Simon's algorithm is originally defined for periodic functions with the promise $f(x) = f(x + k_1)$ for all x, i.e., f is a $(2 : 1)$-mapping. However, our function $f(k_0, \cdot)$ does not fulfill the promise, since some function values might have more than two preimages.

3. Success probability. Assume that we are able to define a suitable classifier $\mathcal{B}$, then we might still only be capable of lower bounding the initial success probability p (which is the case for our algorithm), instead of exactly determining it. This causes problems in properly setting the number of iterations k.

Let us briefly give an outlook how we address these problems.

Classifier. In Simon's algorithm one computes a period $k_1 \in \mathbb{F}_2^n$ bit by bit. Namely, each iteration gives a random vector u_i from the subspace $U = \{u \in \mathbb{F}_2^n \mid \langle u, k_1 \rangle = 0\}$ of all vectors orthogonal to k_1. Thus, in each iteration we obtain a linear relation $\langle u_i, k_1 \rangle = 0$. After $\mathcal{O}(n)$ iterations, we can compute the unique solution k_1.

However, there is no need to compute the u_i sequentially. In our algorithm, we compute $u_1, \ldots, u_\ell$ for some sufficiently large ℓ in parallel. We choose ℓ such that for the periodic function $f(k_0, \cdot)$ the linear span $\langle u_1, \ldots, u_\ell \rangle$ is identical to U with high probability.

Where Simon's algorithm requires $\mathcal{O}(n)$ input bits, our parallel version of Simon's algorithm $\mathcal{A}$ requires $\mathcal{O}(n^2)$ many qubits. We leave it as an open problem whether this quadratic blowup can be avoided.

Our classifier $\mathcal{B}(x)$ should now identify states $|x\rangle$ with $k = k_0$ as good. We know that $f(k_0, \cdot)$ is periodic. Thus we compute for any $f(k, \cdot)$ sufficiently many u_i's with our parallel version $\mathcal{A}$ of Simon's algorithm.

Then $\mathcal{B}$ does the classical post-processing for Simon's algorithm. Namely, we compute from the u_i's some candidate period k_1'. If $\mathcal{B}(x)$ fails in any step, we classify state $|x\rangle$ as bad.

Otherwise $\mathcal{B}$ succeeds in computing some candidate values (k, k_1') for (k_0, k_1). This allows us to check via sufficiently many plaintext/ciphertext pairs (m_i, c_i), (m_i', c_i'), whether for all i the following identity holds

$$c_i + c_i' = \text{Enc}(m) + \text{Enc}(m') = E_{k_0}(m_i + k_1) + E_{k_0}(m_i' + k_1)$$
$$\stackrel{?}{=} E_k(m_i + k_1') + E_k(m_i' + k_1').$$

Checking this identity for sufficiently many plaintext/ciphertext pairs allows us to classify all incorrect states with $(k, k_1') \neq (k_0, k_1)$ as bad.

Simon's promise. It was shown recently by Kaplan et al. [13] and Santoli, Schaffner [23] that Simon's promise can be weakened at the cost of computing more u_i. We will use yet another technique for dealing with general functions. Namely, under the mild assumption that $f(k_0, x)$ behaves as a random periodic function with period k_1, we show that any function value $f(k_0, x)$ has only two preimages with probability at least $\frac{1}{2}$. We then only argue about these proper function values $f(k_0, x)$.

This provides a simple and clean way to use only a limited number ℓ of u_i. For comparison, $\ell = 2(n + \sqrt{n})$ is sufficient for our purpose, whereas a direct application of the techniques of Kaplan et al. [13] requires $\ell > 3n$.

Success probability. We define some $\mathcal{B}$ that classifies states with $k \neq k_0$ as bad with overwhelming probability. However, for states $|x\rangle$ with $k = k_0$ our $\mathcal{B}$ classifies $|x\rangle$ as good only with a probability that is lower bounded by some constant.

Still, we choose the number k of iterations analogous to Theorem 1, basically assuming that we would classify all states with $k = k_0$ as good. This in turn implies that our choice of k might be too small to fully rotate towards the subspace of good states. Nevertheless, by adapting the analysis of Brassard et al. [8] we are still able to show that we succeed in our case with probability at least $\frac{2}{5}$.

The following two basic lemmata will be frequently used in our further analysis. The first one shows, that any $n-1$ vectors obtained from Simon's algorithm form a basis of the $n-1$-dimensional vector space U with probability at least $\frac{1}{4}$. The second one shows that this basis allows us to compute its unique orthogonal complement, and therefore the period in Simon's algorithm, in polynomial time.

Lemma 1. *Let $U \subset \mathbb{F}_2^n$ be an $(n-1)$-dimensional subspace. Suppose we obtain $\mathbf{u}_1, \ldots, \mathbf{u}_{n-1} \in U$ drawn independently at uniform from U. Then $\mathbf{u}_1, \ldots, \mathbf{u}_{n-1}$ are linearly independent and form a basis of U with probability greater than $\frac{1}{4}$.*

Proof. Let E_i, $0 \leq i < n$ be the event that the first i vectors $\mathbf{u}_1, \ldots, \mathbf{u}_i$ form an i-dimensional space. Define $\Pr[E_0] := 1$.

Let $p_1 = \Pr[E_1]$ and $p_i = \Pr[E_i \mid E_{i-1}]$ for $2 \leq i < n$. Then $p_1 = 1 - \frac{1}{2^{n-1}}$, since we only have to exclude $\mathbf{u}_1 = 0^n \in U$. Moreover for $1 < i < n$, we have

$$p_i = 1 - \frac{2^{i-1}}{2^{n-1}},$$

since $\mathbf{u}_i$ should not lie in the $(i-1)$-dimensional span $\langle \mathbf{u}_1, \ldots, \mathbf{u}_{i-1} \rangle$. We obtain

$$\Pr[E_{n-1}] = \Pr[E_{n-1} \mid E_{n-2}] \cdot \Pr[E_{n-2}] = \ldots = \prod_{i=1}^{n-1} \Pr[E_i \mid E_{i-1}]$$

$$= \prod_{i=1}^{n-1} p_i = \prod_{i=1}^{n-1} 1 - 2^{i-n} = \prod_{i=1}^{n-1} 1 - 2^{-i}.$$

Since $\Pr[E_{n-1}] \geq \lim_{n \to \infty} \prod_{i=1}^{n-1} 1 - 2^{-i} \geq 0.288$, the claim follows. $\square$

Lemma 2. *Let $\mathbf{u}_1, \ldots, \mathbf{u}_{n-1} \in \mathbb{F}_2^n$ be linearly independent. Then one can compute in time $\mathcal{O}(n^3)$ the unique vector $\mathbf{v} \in \mathbb{F}_2^n \setminus \{\mathbf{0}\}$ such that $\langle \mathbf{v}, \mathbf{u}_i \rangle = 0$ for all $i = 1, \ldots, n-1$.*

Proof. Define the matrix $U \in \mathbb{F}_2^{(n-1) \times n}$, whose rows consist of the vectors $\mathbf{u}_i$. Transform U via Gaussian elimination into

$$U' = (I_{n-1} | \bar{\mathbf{v}}) \text{ for some column vector } \bar{\mathbf{v}} \in \mathbb{F}_2^{n-1}.$$

This costs time $\mathcal{O}(n^3)$. Notice that we have $\text{span}(U) = \text{span}(U')$, since we only applied elementary row operations.

Let $\mathbf{v} = (\bar{\mathbf{v}}^{\mathbf{t}}|1) \in \mathbb{F}_2^n \setminus \{\mathbf{0}\}$. Let $\mathbf{e}_i$ denote the i-th $(n-1)$-dimensional unit vector. Then the i-th basis vector $\mathbf{u}_i' = (\mathbf{e}_i|\bar{v}_i)$ of U' satisfies

$$\langle \mathbf{v}, \mathbf{u}_i' \rangle = \langle (\bar{\mathbf{v}}^{\mathbf{t}}|1), (\mathbf{e}_i, \bar{v}_i) \rangle = \bar{v}_i + \bar{v}_i = 0 \quad \text{for } i = 1, \ldots, n-1.$$

Since $\mathbf{v}$ is orthogonal to all $\mathbf{u}_i'$, it is also orthogonal to all $\mathbf{u}_i$. $\qquad\square$

3   Combining the Algorithms of Grover and Simon

Let us now proof our main theorem, whose statement is formulated in a slightly more general fashion than in Sect. 2 to make it useful also outside the FX construction context. In the FX construction, $f_{k_0,k_1,k_2}(x)$ is $\text{Enc}(x)$, and $g(k,x)$ is $E_k(x)$.

Theorem 2. *Let $f : \mathbb{F}_2^m \times \mathbb{F}_2^{3n} \to \mathbb{F}_2^n$ with*

$$(k_0, k_1, k_2, x) \mapsto g(k_0, x + k_1) + k_2,$$

where $g : \mathbb{F}_2^m \times \mathbb{F}_2^n \to \mathbb{F}_2^n$ and $g(k, \cdot)$ is a random function $\mathbb{F}_2^n \to \mathbb{F}_2^n$ for any fixed $k \in \mathbb{F}_2^m$. Given quantum oracle access to $f_{k_0,k_1,k_2}(\cdot)$ and $g(\cdot, \cdot)$, the tuple (k_0, k_1, k_2) can be computed with success probability at least $\frac{2}{5}$ using $m + 4n(n + \sqrt{n})$ qubits and

$$2^{\frac{m}{2}} \cdot \mathcal{O}(m + n) \text{ oracle queries.}$$

Proof. Let us define the function $f' : \mathbb{F}_2^m \times \mathbb{F}_2^n \to \mathbb{F}_2^n$ with

$$(k, x) \mapsto f_{k_0,k_1,k_2}(x) + g(k, x).$$

Notice that

$$f'(k_0, x) = g(k_0, x + k_1) + k_2 + g(k_0, x).$$

Hence $f'(k_0, x) = f'(k_0, x + k_1)$, and therefore $f'(k_0, \cdot)$ is periodic with period k_1 in its second component. We use amplitude amplification to search for k_0. A generalized version of Simon's algorithm then tells us which $f'(k, \cdot)$ is periodic.

However notice that we have a non-trivial period only if $k_1 \neq 0^n$. For $k_1 = 0^n$ we obtain a constant function with $f'(k_0, \cdot) = k_2$ for all inputs x. In the case of $k \neq k_0$, by the randomness of $g(k, \cdot)$ the function $f'(k, \cdot)$ is almost balanced in each output bit. This implies that we could use a generalized version of Deutsch-Josza's algorithm [9] to decide which $f'(k, \cdot)$ is constant.

For simplicity, we will instead describe a basic Grover search for k_0. Notice that once k_0 is found, we can compute $k_2 = f_{k_0,k_1,k_2}(x) + g(k_0, x)$ for some arbitrary value of x.

Lemma 3. *Let $k_1 = 0^n$. Then one can determine k_0, k_2 with probability at least $1 - 2^{-m}$ using $2^{\frac{m}{2}} \cdot \mathcal{O}(m)$ oracle queries and $m + 1$ qubits.*

Proof. In the case $k_1 = 0^n$, we have

$$f'(k_0, x) = f_{k_0, 0^n, k_2}(x) + g(k_0, x) = g(k_0, x) + k_2 + g(k_0, x) = k_2 \text{ for any } x.$$

Thus, $f'(k_0, \cdot)$ is a constant function. Let us define a test, which checks whether $f_{k_0, 0^n, k_2}(x) + g(k, x)$ is constant to decide if $k = k_0$.

By evaluating $f_{k_0, 0^n, k_2}(\cdot)$, we compute for $1 + 2\lceil \frac{m}{n} \rceil$ random $x_i \in_R \mathbb{F}_2^n$ the function values

$$y_i = f_{k_0, 0^n, k_2}(x_i) = g(k_0, x_i) + k_2.$$

Moreover, we define the classical test $h : \mathbb{F}_2^m \to \mathbb{F}_2$ that takes as input a value $k \in \mathbb{F}_2^m$. We map k to 1 iff

$$y_i + g(k, x_i) = y_{i+1} + g(k, x_{i+1}) \text{ for all } 1 \le i \le 2\lceil \frac{m}{n} \rceil.$$

For $k = k_0$ we have $y_i + g(k, x_i) = k_2$, and therefore all identities are satisfied with probability 1. In the case $k \ne k_0$ by the randomness of $g(k, \cdot)$ any identity is fulfilled with probability 2^{-n}. Hence, all $2\lceil \frac{m}{n} \rceil$ identities are simultaneously fulfilled with probability at most 2^{-2m}.

Therefore, the probability that there is an incorrect $k \ne k_0$ that passes the test by h is at most $(2^m - 1)2^{-2m} < 2^{-m}$.

Altogether, we can use the quantum embedding of h on $m + 1$ qubits as a Grover oracle $\mathcal{B}$ from Theorem 1 that takes value 1 iff $k = k_0$ (with overwhelming probability) using $\mathcal{O}(m)$ oracle queries to $f_{k_0, k_1, k_2}(\cdot)$ and $g(\cdot, \cdot)$. Define $\mathcal{A}$ from Theorem 1 as the m-fold Hadamard transform

$$H^{\otimes m} : \mathcal{H} \to \mathcal{H} \text{ that maps } |x\rangle \mapsto \frac{1}{2^{m/2}} \sum_{k \in \mathbb{F}_2^m} (-1)^{x \cdot k} |k\rangle. \tag{1}$$

By Theorem 1, we start with $\mathcal{A}|0\rangle = H^{\otimes m}|0^m\rangle = 2^{-m/2} \sum_{k \in \mathbb{F}_2^m} |k\rangle$, the uniform superposition of all keys $k \in \mathbb{F}_2^m$. Then we have initial success probability $p = 2^{-m}$ to measure the good key $|k_0\rangle$. After $k = \lceil \frac{\pi}{4 \arcsin(\sqrt{p})} \rceil \le 2^{\frac{m}{2}}$ Grover iterations we measure $|k_0\rangle$ with probability at least $1 - 2^{-m}$.

Finally, we compute $k_2 = f_{k_0, 0^n, k_2}(x) + g(k_0, x)$ for some arbitrary value of x. □

For the remainder of the proof of Theorem 2, let us now assume $k_1 \ne 0^n$.

Let $\ell = 2(n + \sqrt{n})$. The following function h evaluates f' in parallel on ℓ arguments in the second component. Let $h : \mathbb{F}_2^m \times (\mathbb{F}_2^n)^\ell \mapsto (\mathbb{F}_2^n)^\ell$ with

$$(k, x_1, \dots x_\ell) \mapsto f'(k, x_1) || f'(k, x_2) || \dots || f'(k, x_\ell).$$

Let U_h be the universal bijective quantum embedding of h on $m + 2n\ell$ qubits, i.e. we map

$$|k, x_1, \dots x_\ell, \mathbf{0}, \dots, \mathbf{0}\rangle \mapsto |k, x_1, \dots x_\ell, h(k, x_1, \dots, x_\ell)\rangle.$$

Since we have quantum access to $f_{k_0, k_1, k_2}(\cdot)$ and $g(\cdot, \cdot)$, we can realize f' with two function queries and therefore also U_h with 2ℓ queries.

We now describe a quantum process $\mathcal{A}$, which is a parallel version of Simon's algorithm, that succeeds with initial probability $p \geq \frac{1}{5} \cdot 2^{-m}$ using $2\ell = \mathcal{O}(n)$ queries and $m + 2n\ell = m + 4n(n + \sqrt{n})$ qubits. Afterwards, we amplify $\mathcal{A}$'s success probability to at least $\frac{2}{5}$ using roughly $2^{\frac{m}{2}}$ iterations, where each iteration consumes $\mathcal{O}(m + n)$ oracle queries.

Quantum algorithm $\mathcal{A}$ on input $|0\rangle$

1. Prepare the initial $m + 2n\ell$-qubit state $|\mathbf{0}\rangle$.
2. Apply Hadamard $H^{\otimes m+n\ell}$, as defined in Eq. (1), on the first $m + n\ell$ qubits resulting in

$$\sum_{k\in\mathbb{F}_2^m,\ x_1,\ldots,x_\ell\in\mathbb{F}_2^n} |k\rangle|x_1\rangle..|x_\ell\rangle|\mathbf{0}\rangle,$$

   where we omit the amplitudes $2^{-(m+n\ell)/2}$ for ease of exposition.
3. An application of U_h yields

$$\sum_{k\in\mathbb{F}_2^m,\ x_1,\ldots,x_\ell\in\mathbb{F}_2^n} |k\rangle|x_1\rangle..|x_\ell\rangle|h(k, x_1, \ldots, x_\ell)\rangle$$

4. We now apply Hadamard on the qubits in position $m + 1 \ldots m + n\ell$ (i.e. for $|x_1\rangle..|x_\ell\rangle$), which results in

$$|\psi\rangle = \sum_{\substack{k\in\mathbb{F}_2^m,\ u_1,\ldots,u_\ell\in\mathbb{F}_2^n, \\ x_1,\ldots,x_\ell\in\mathbb{F}_2^n}} |k\rangle(-1)^{\langle u_1,x_1\rangle}|u_1\rangle\ldots(-1)^{\langle u_\ell,x_\ell\rangle}|u_{n-1}\rangle|h(k, x_1, \ldots, x_\ell)\rangle. \quad (2)$$

Assume that we would measure the last $n\ell$ qubits of state $|\psi\rangle$ from step 4. Then these qubits would collapse into

$$|h(k, x_1, \ldots, x_\ell)\rangle = |f'(k, x_1)||f'(k, x_2)||\ldots||f'(k, x_\ell)\rangle,$$

for some *fixed values* of $k, x_1, \ldots, x_\ell \in \mathbb{F}_2^n$. Assume further that $k = k_0$.

Let us look at an arbitrary n-qubit state $|z_i\rangle = (-1)^{\langle u_i,x_i\rangle}|u_i\rangle$ from $|\psi\rangle$ that is entangled with $|f'(k_0, x_i)\rangle$. Hence, $|z_i\rangle$ collapses into a superposition that is consistent with the measured $f'(k_0, x_i)$.

We call the state $|z_i\rangle$ *proper* if x_i and $x_i + k_1$ are the only preimages of $f'(k_0, x_i)$. Notice that a proper $|z_i\rangle$ collapses into the superposition

$$\left((-1)^{\langle u_i,x_i\rangle} + (-1)^{\langle u_i,x_i+k_1\rangle}\right)|u_i\rangle = (-1)^{\langle u_i,x_i\rangle}\left(1 + (-1)^{\langle u_i,k_1\rangle}\right)|u_i\rangle. \quad (3)$$

As one can see from the right-hand side of Eq. (3), the qubits $|u_i\rangle$ have a non-vanishing amplitude iff $\langle u_i, k_1\rangle = 0$. Therefore, a measurement of a proper state yields some uniformly random $u_i \in U$, where $U = \{u \in \mathbb{F}_2 \mid \langle u, k_1\rangle = 0\}$.

Notice, that in general one can have more than two preimages of $f'(k_0, x_i)$, which results in $u_i \in \mathbb{F}_2^n$ that are with a certain probability chosen from some subset of U. Although such u_i usually still provide useful information about k_1 – whenever the subset is not too small – their probability distribution is somewhat cumbersome to deal with.

For ease of exposition, we want to work with proper states $|z_i\rangle$ only. In the following lemma, we show that any $|z_i\rangle$ is proper with probability at least $\frac{1}{2}$. This in turn means that on expectation a set of vectors $S = \{u_1, \ldots, u_{2(n-1)}\}$ derived from measuring $2(n-1)$ states contains at least $n-1$ vectors $u_{i_1}, \ldots, u_{i_{n-1}}$ that are chosen independently uniformly at random from U. Notice that a priori we are not able to identify these $n-1$, since we are not able to tell which state is proper. Nevertheless, we can easily compute from S a maximal set of independent vectors. Since $u_{i_1}, \ldots, u_{i_{n-1}} \in S$, by Lemma 1 these vectors form a basis of U with probability greater than $\frac{1}{4}$.

Moreover, if we increase the cardinality of S from $2(n-1)$ to $\ell = 2(n+\sqrt{n})$, as it is done in algorithm $\mathcal{A}$, then the above does not only hold on expectation, but with constant probability.

Lemma 4. *Any state $|z_i\rangle = (-1)^{\langle u_i, x_i\rangle}|u_i\rangle$ is proper with probability at least $\frac{1}{2}$. Any set of $\ell = 2(n + \sqrt{n})$ states contains at least $n - 1$ proper states with probability greater than $\frac{4}{5}$.*

Proof. Recall that $|z_i\rangle$ is proper if $f'(k_0, x_i)$ has only two preimages. By definition of f', we have

$$f'(k_0, x_i) = g(k_0, x_i + k_1) + k_2 + g(k_0, x_i).$$

Let us denote by $S_i = i \times \{0,1\}^{n-1}$ for $i = 0,1$ the set of all n-dimensional vectors that start with bit i. Since $k_1 \neq 0^n$, assume wlog that the first bit of k_1 is 1, i.e., $k_1 \in S_1$.

Since $f'(k_0, \cdot)$ is periodic with k_1 in its second argument, the values of $f'(k_0, x)$ with $x \in S_0$ already determine the values of $f'(k_0, x + k_1)$ with $x + k_1 \in S_1$.

Let us fix some $x \in S_0$ and thus also $f'(k_0, x)$. The state $|z\rangle = (-1)^{\langle u, x\rangle}|u\rangle$ is proper if there is no other $x' \in S_0 \setminus \{x\}$ that collides with x under $f'(k_0, \cdot)$. By the randomness of $g(k_0, \cdot)$ this happens with probability

$$\Pr[|z\rangle \text{ is proper }] = 1 - \Pr[\exists x' \in S_0 \setminus \{x\} \text{ with } f'(k_0, x') = f'(k_0, x)]$$

$$\geq 1 - \frac{2^{n-1} - 1}{2^n} \geq \frac{1}{2},$$

where the first inequality follows from a union bound.

It remains to count the number of proper states within a set of $\ell = 2(n+\sqrt{n})$ states. Let X_i be an indicator variable that takes value 1 iff $|z_i\rangle$ is proper. Let $X = X_1 + \ldots X_{2(n+\sqrt{n})}$. Furthermore, define $\mu := n + \sqrt{n}$, which implies $\mathbb{E}[X] \geq \mu$. Now we apply the following Chernoff bound (see [21], Corollary 4.9 and Exercise 4.7)

$$\Pr[X \leq \mu - a] \leq e^{\frac{-2a^2}{n}} \quad \text{for all } a < \mu.$$

This implies in our case

$$\Pr[X \geq n - 1] = 1 - \Pr[X \leq n] = 1 - \Pr[X \leq \mu - \sqrt{n}] \geq 1 - e^{-2} \geq \frac{4}{5}. \qquad \square$$

Let us now go back to the quantum superposition $|\psi\rangle$ of qubits in positions $1, \ldots, m + n\ell$ (i.e. for the qubits $|k\rangle|u_1\rangle \ldots |u_{n-1}\rangle$) after applying algorithm $\mathcal{A}$, i.e. *without* measurement. Assume we had a classifier $\mathcal{B} : \mathbb{F}_2^{m+n\ell} \to \{0, 1\}$ that partitions $|\psi\rangle$ in a good subspace and a bad subspace, where the good subspace is spanned by the set of basis states $|x\rangle$ for which $\mathcal{B}(x) = 1$. We split

$$|\psi\rangle = |\psi_1\rangle + |\psi_0\rangle,$$

where $|\psi_1\rangle$ and $|\psi_0\rangle$ denotes the projection onto the good and onto the bad subspace, respectively.

Ideally we would like to define $|\psi_1\rangle$ as the sum of those basis states for which $|k\rangle = |k_0\rangle$. Unfortunately, we cannot check correctness of k_0 directly. Instead, with our classifier $\mathcal{B}$ we compute from $k, u_1, \ldots, u_\ell$ a candidate k_1' for the period k_1. This allows for an easy test $(k, k_1') \stackrel{?}{=} (k_0, k_1)$.

Classifier $\mathcal{B}$ (polynomial time computable Boolean function). Let us define the following classical Boolean function

$$\mathcal{B} : \mathbb{F}_2^{m+n\ell} \to \{0, 1\} \text{ that maps } (k, u_1, \ldots, u_\ell) \mapsto \{0, 1\}.$$

In $\mathcal{B}$, we hardwire for $\lceil \frac{3m+n\ell}{n} \rceil$ random pairs $m_i, m_i' \in_R \mathbb{F}_2^n$ with $m_i \neq m_i'$ the values

$$y_i = f_{k_0,k_1,k_2}(m_i) + f_{k_0,k_1,k_2}(m_i') = g(k_0, m_i + k_1) + g(k_0, m_i' + k_1),$$

which can be computed via $2\lceil \frac{3m+n\ell}{n} \rceil$ function evaluations of $f_{k_0,k_1,k_2}(\cdot)$. Now we check the following two steps.

1. Let $\bar{U} = \langle u_1, \ldots, u_\ell \rangle$ be the linear span of all u_i. If $\dim(\bar{U}) \neq n - 1$, output 0. Otherwise compute a basis of $\bar{U}$, and use Lemma 2 to compute the unique vector $k_1' \in \mathbb{F}_2^n \setminus \{\mathbf{0}\}$ orthogonal to $\bar{U}$.
2. Check via $2\lceil \frac{3m+n\ell}{n} \rceil$ function evaluations of $g(\cdot, \cdot)$ whether

$$y_i \stackrel{?}{=} g(k, m_i + k_1') + g(k, m_i' + k_1') \text{ for all } i = 1, \ldots, \left\lceil \frac{3m + n\ell}{n} \right\rceil.$$

If all identities hold, output GOOD. Else output BAD.

The following lemma shows that our test $\mathcal{B}$ classifies basis states with the correct $k = k_0$ as GOOD, respectively 1, with probability at least $\frac{1}{5}$. We could easily boost this into a probability close to 1 by increasing the number of qubits of $\mathcal{A}$. However, for the ease of description and for minimizing the number of qubits, we keep it this way, which eventually only slightly lowers the overall success probability of our algorithm.

More important is that $\mathcal{B}$ almost never gives false positives. Namely, whenever $\mathcal{B}$ declares a state as GOOD then indeed $k = k_0$ with overwhelming probability. This in turn implies that by our choice of parameters we safely sort out *all of the exponentially many* incorrect keys $k \neq k_0$.

Lemma 5. *If $k = k_0$ then test $\mathcal{B}$ outputs 1 with probability at least $\frac{1}{5}$.*
Vice versa, if $\mathcal{B}$ outputs 1 then $k_0 = k$ with probability at least $1 - \frac{1}{2^{2m+n\ell-4}}$.

Proof. Let us denote by GOOD the event that $\mathcal{B}$ outputs 1. We first compute

$$p_0 = \Pr[\text{GOOD} \mid k = k_0].$$

If $k = k_0$, then $f'(k, \cdot)$ is periodic with k_1 in its second argument. Moreover, we know by Lemma 4 that $u_1, \ldots, u_\ell$ contain with probability at least $\frac{4}{5}$ at least $n-1$ vectors that are uniformly at random from the subspace $U \subset \mathbb{F}_2^n$ orthogonal to k_1. From Lemma 1, these vectors form a basis of U with probability at least $\frac{1}{4}$.

In total, we pass step (1) of $\mathcal{B}$ with probability at least $\frac{4}{5} \cdot \frac{1}{4} = \frac{1}{5}$. Moreover, in the case $k = k_0$ we also have $k_1' = k_1$, i.e., $\mathcal{B}$ computes the correct k_1. Therefore, we pass all tests in step (2) of $\mathcal{B}$ with probability 1. Altogether, we obtain $p_0 \geq \frac{1}{5}$, which proves the first claim.

In order to prove the second claim, let us compute a lower bound for the probability

$$p_1 = \Pr[k = k_0 \mid \text{GOOD}] = \frac{p_0 \cdot \Pr[k = k_0]}{\Pr[\text{GOOD}]}.$$

Since $\Pr[k = k_0] = 2^{-m}$, it remains to compute

$$\begin{aligned}
\Pr[\text{GOOD}] &= \Pr[k = k_0] \cdot \Pr[\text{GOOD} \mid k = k_0] + \Pr[k \neq k_0] \cdot \Pr[\text{GOOD} \mid k \neq k_0] \\
&= 2^{-m} \cdot p_0 + (1 - 2^{-m}) \cdot \Pr[\text{GOOD} \mid k \neq k_0] \\
&\leq 2^{-m} \cdot p_0 + \Pr[\text{GOOD} \mid k \neq k_0].
\end{aligned}$$

Let us further bound the probability $\Pr[\text{GOOD} \mid k \neq k_0]$. This event means that we pass steps (1) and (2) of $\mathcal{B}$, even though we have the incorrect k. Since we need an upper bound only, we can assume that we always pass step (1) and compute some k_1'. In step (2), we check the identities

$$y_i \overset{?}{=} g(k, m_i + k_1') + g(k, m_i' + k_1') \text{ for all } i = 1, \ldots, \left\lceil \frac{3m + n\ell}{n} \right\rceil.$$

By our assumption $g(k, \cdot) : \mathbb{F}_2^n \to \mathbb{F}_2^n$ is random for any fixed $k \in \mathbb{F}_2^m$. Thus, each of these identities holds independently with probability 2^{-n}. Notice that this probability is even independent of the value of k_1' computed in step (2).

Thus, *all* the $\left\lceil \frac{3m+n\ell}{n} \right\rceil$ identities are simultaneously fulfilled with probability at most $2^{-3m-n\ell}$, which is an upper bound for $\Pr[\text{GOOD} \mid k \neq k_0]$.

This in turn implies

$$p_1 = \Pr[k = k_0 \mid \text{GOOD}] = \frac{p_0 \cdot \Pr[k = k_0]}{\Pr[\text{GOOD}]} \geq \frac{p_0 \cdot 2^{-m}}{p_0 \cdot 2^{-m}(1 + \frac{2^{-2m-n\ell}}{p_0})}$$

$$> \frac{1}{1 + 2^{3-2m-n\ell}} = \frac{1 - 2^{3-2m-n\ell}}{1 - 2^{2(3-2m-n\ell)}}$$

$$= \frac{1 - 2^{2(3-2m-n\ell)} - (2^{3-2m-n\ell} - 2^{2(3-2m-n\ell)})}{1 - 2^{2(3-2m-n\ell)}} \geq 1 - \frac{1}{2^{2m+n\ell-4}},$$

where the last inequality holds for $2m + n\ell > 3$. This concludes the proof. $\square$

By Lemma 5 our test $\mathcal{B}$ classifies a bad state $|k\rangle|u_1\rangle \ldots |u_\ell\rangle$ with $k \neq k_0$ as good with probability at most $2^{-2m-n\ell+4}$. Notice that there are at most $2^{m+n\ell}$ states in any superposition. Therefore, $\mathcal{B}$ classifies any bad state in a superposition as good with probability at most $2^{m+n\ell} \cdot 2^{-2m-n\ell+4} = 2^{-m+4}$. We will have at most $2^{\frac{m}{2}}$ iterations of $\mathcal{A}$. By a union bound, the probability that $\mathcal{B}$ classifies any bad state in a superposition in any of these iterations as good is at most $2^{\frac{m}{2}} \cdot 2^{-m+4} = 2^{-\frac{m}{2}+4}$, which is exponentially small.

This implies that $\mathcal{B}$ (almost) never yields false positives. Hence, we classify a state $|k\rangle|u_1\rangle \ldots |u_\ell\rangle$ as good iff and only if $\mathcal{B}(k, u_1, \ldots, u_\ell) = 1$. The initial success probability p of $\mathcal{A}$ in producing a good state is by Lemma 5

$$p = \Pr[|k\rangle|u_1\rangle \ldots |u_\ell\rangle \text{ is good}]$$
$$= \Pr[k = k_0] \cdot \Pr[\mathcal{B}(k, u_1, \ldots, u_\ell) = 1 \mid k = k_0] \geq \frac{1}{5} \cdot 2^{-m}. \tag{4}$$

Our Boolean function $\mathcal{B}$ defines a unitary operator $S_\mathcal{B}$ that conditionally changes the sign of the amplitudes of the good states

$$|k\rangle|u_1\rangle \ldots |u_\ell\rangle \mapsto \begin{cases} -|k\rangle|u_1\rangle \ldots |u_\ell\rangle & \text{if } \mathcal{B}(k, u_1, \ldots, u_\ell) = 1 \\ |k\rangle|u_1\rangle \ldots |u_\ell\rangle & \text{if } \mathcal{B}(k, u_1, \ldots, u_\ell) = 0 \end{cases}.$$

The complete amplification process is realized by repeatedly applying the unitary operator $Q = -\mathcal{A}S_0 \mathcal{A}^{-1}S_\mathcal{B}$ to the initial state $|\psi\rangle = \mathcal{A}|\mathbf{0}\rangle$, i.e., we compute $Q^k \mathcal{A}|\mathbf{0}\rangle$ and measure the system for some suitable number of iterations k.

Let $|\psi_1\rangle$, $|\psi_0\rangle$ be the projection of $|\psi\rangle$ onto the good and the bad subspace, respectively. Denote by $\mathcal{H}_\psi$ the 2-dimensional subspace spanned by $|\psi_1\rangle$, $|\psi_0\rangle$. Initially, in $\mathcal{H}_\psi$ the angle between $\mathcal{A}|\mathbf{0}\rangle$ and the bad subspace is θ, where

$$\sin^2(\theta) = p = \langle \psi_1 \mid \psi_1 \rangle.$$

Thus, $\theta = \arcsin(\sqrt{p}) \geq \arcsin(\sqrt{\frac{1}{5} \cdot 2^{-\frac{m}{2}}})$, where the lower bound follows from (4).

Now every Grover iteration by Q increases the angle by 2θ, i.e., to $(2k+1)\theta$ after k iterations. If this angle is roughly $\frac{\pi}{2}$, then we are almost parallel to $|\psi_1\rangle$ in $\mathcal{H}_\psi$ and measure a good state with high probabilty. Therefore, let us choose

$$k = \lceil \frac{\pi}{4 \arcsin(2^{-\frac{m}{2}})} \rceil.$$

After k iterations a final measurement produces a good state with probability $p_{\text{good}} = \sin^2((2k+1)\theta)$. Thus, we obtain

$$p_{\text{good}} \geq \sin^2 \left(\frac{\pi}{2} \cdot \frac{\arcsin\left(\sqrt{\frac{1}{5}} \cdot 2^{-\frac{m}{2}} \right)}{\arcsin(2^{-\frac{m}{2}})} \right). \tag{5}$$

Notice that $\arcsin(x) \approx x$ for small x. So for m sufficiently large, the right hand side of (5) quickly converges to $\sin^2(\frac{\pi}{2\sqrt{5}}) \approx 0.42$. Namely, for $m \geq 12$ the right hand side is already larger than $\frac{2}{5}$, which shows that we measure a good state after amplification with the claimed success probability.

This measurement reveals k_0, and an application of $\mathcal{B}$ on a good state also reveals the correct value for k_1. We can then easily compute for an arbitrary x the value

$$k_2 = f_{k_0,k_1,k_2}(x) + g(k_0, x + k_1).$$

This completes the proof of our main theorem. $\square_{\text{Theorem 2}}$

3.1   Potential Improvements

As already pointed out, we did not try to optimize all constants in Theorem 2. Many parameters are chosen in a way that allows for smooth and simple proofs. Let us comment which parameters might be subject to (small) improvements.

Memory. We use $m + 4n(n + \sqrt{n})$ many qubits. However, we use only proper states in our proof, which gives a factor 2 loss for the u_i. But states that are not proper should usually still provide enough information. Hence, roughly $m + 2n^2$ qubits should be sufficient. An open question is whether this can be lowered to $m + o(n^2)$.

Notice that one can solely consider the projection of $f'(k_0, \cdot)$ to one output bit, which is also periodic. This however still requires $m + n^2$ input qubits.

Success probability. If we have $n + \mathcal{O}(1)$ values u_i, then in the periodic case $k = k_0$ one would expect to obtain a basis of U with probability close to 1, whereas in the non-periodic case $k \neq k_0$ one would expect to obtain an n-dimensional basis with probability close to 1. This means that our classifier $\mathcal{B}$ works almost perfect, which in turn implies that our success probability is in fact close to 1.

4   Two Open Problems

Our new quantum algorithm shows that the use of whitening keys in the FX-construction does not increase security in the quantum-CPA model. This result raises at least two more natural questions to be investigated in the future.

The first and maybe most important question is the security of key-alternating ciphers against quantum-CPA attacks. Key-alternating ciphers can be seen as a multiple round generalization of the Even-Mansour construction and many popular ciphers, most importantly the AES, follow this general design principle. It would thus be of great interest to design quantum algorithms that break those ciphers, or show that this is not possible in general.

The second question is the investigation of sound techniques for extending the key length of a given cipher in a quantum setting. Of course, it is always possible to design new key-schedulings (and potentially increase the number of

rounds slightly) for larger key-sizes. The most well-known example is again the AES with its different variants of 128, 192 or 256 key bits. However, this requires an exact understanding of the internal behaviour of the cipher and it is thus of interest to investigate generic ways of increasing the key length. That is, given a cipher with an m bit key, how can we extend its key size to obtain a cipher that achieves m bit security against quantum adversaries, while tolerating only a minimal performance penalty. Initial ideas along these lines have recently been presented in [1], where the key-addition in Even-Mansour has been replaced by a different group operation.

References

1. Alagic, G., Russell, A.: Quantum-secure symmetric-key cryptography based on hidden shifts. In: Coron, J.-S., Nielsen, J.B. (eds.) EUROCRYPT 2017. LNCS, vol. 10212, pp. 65–93. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-56617-7_3
2. Albrecht, M.R., Driessen, B., Kavun, E.B., Leander, G., Paar, C., Yalçın, T.: Block ciphers – focus on the linear layer (feat. PRIDE). In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014. LNCS, vol. 8616, pp. 57–76. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-662-44371-2_4
3. Ambainis, A., Bach, E., Nayak, A., Vishwanath, A., Watrous, J.: One-dimensional quantum walks. In: Vitter, J.S., Spirakis, P.G., Yannakakis, M. (eds.) Proceedings on 33rd Annual ACM Symposium on Theory of Computing, Heraklion, Crete, Greece, 6–8 July 2001, pp. 37–49. ACM (2001)
4. Anand, M.V., Targhi, E.E., Tabia, G.N., Unruh, D.: Post-quantum security of the CBC, CFB, OFB, CTR, and XTS modes of operation. In: Takagi, T. (ed.) PQCrypto 2016. LNCS, vol. 9606, pp. 44–63. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-29360-8_4
5. Boneh, D., Zhandry, M.: Secure signatures and chosen ciphertext security in a quantum computing world. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013. LNCS, vol. 8043, pp. 361–379. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40084-1_21
6. Borghoff, J., et al.: PRINCE – a low-latency block cipher for pervasive computing applications. In: Wang, X., Sako, K. (eds.) ASIACRYPT 2012. LNCS, vol. 7658, pp. 208–225. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-34961-4_14
7. Brassard, G., Høyer, P.: An exact quantum polynomial-time algorithm for simon's problem. In: Proceedings of the Fifth Israel Symposium on Theory of Computing and Systems, ISTCS 1997, Ramat-Gan, Israel, 17–19 June 1997, pp. 12–23. IEEE Computer Society (1997)
8. Brassard, G., Hoyer, P., Mosca, M., Tapp, A.: Quantum amplitude amplification and estimation. Contemp. Math. **305**, 53–74 (2002)
9. Deutsch, D., Jozsa, R.: Rapid solution of problems by quantum computation. In: Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences, vol. 439, pp. 553–558. The Royal Society (1992)
10. Even, S., Mansour, Y.: A construction of a cipher from a single pseudorandom permutation. J. Cryptol. **10**(3), 151–162 (1997)

11. Grover, L.K.: A fast quantum mechanical algorithm for database search. In: Miller, G.L. (ed.) Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing, Philadelphia, Pennsylvania, USA, May 22–24, pp. 212–219. ACM (1996)

12. Kaplan, M.: Quantum attacks against iterated block ciphers. CoRR (2014). http://arxiv.org/abs/1410.1434

13. Kaplan, M., Leurent, G., Leverrier, A., Naya-Plasencia, M.: Breaking symmetric cryptosystems using quantum period finding. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016. LNCS, vol. 9815, pp. 207–237. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53008-5_8

14. Kaplan, M., Leurent, G., Leverrier, A., Naya-Plasencia, M.: Quantum differential and linear cryptanalysis. IACR Trans. Symmetric Cryptol. **2016**(1), 71–94 (2016)

15. Kilian, J., Rogaway, P.: How to protect DES against exhaustive key search. In: Koblitz, N. (ed.) CRYPTO 1996. LNCS, vol. 1109, pp. 252–267. Springer, Heidelberg (1996). https://doi.org/10.1007/3-540-68697-5_20

16. Kilian, J., Rogaway, P.: How to protect DES against exhaustive key search (an analysis of DESX). J. Cryptology **14**(1), 17–35 (2001)

17. Kuwakado, H., Morii, M.: Quantum distinguisher between the 3-round feistel cipher and the random permutation. In: Proceedings of the IEEE International Symposium on Information Theory, ISIT 2010, 13–18 June 2010, Austin, Texas, USA, pp. 2682–2685. IEEE (2010)

18. Kuwakado, H., Morii, M.: Security on the quantum-type even-mansour cipher. In: Proceedings of the International Symposium on Information Theory and its Applications, ISITA 2012, Honolulu, HI, USA, 28–31 October 2012, pp. 312–316. IEEE (2012)

19. Lipton, R.J., Regan, K.W.: Quantum Algorithms via Linear Algebra: A Primer. The MIT Press, Boston (2014)

20. Mermin, N.: Quantum Computer Science: An Introduction. Cambridge University Press, New York (2007)

21. Mitzenmacher, M., Upfal, E.: Probability and Computing - Randomized Algorithms and Probabilistic Analysis. Cambridge University Press, New York (2005)

22. NIST: Post quantum project. http://csrc.nist.gov/groups/ST/post-quantum-crypto/ Accessed 6 Feb 2017

23. Santoli, T., Schaffner, C.: Using simon's algorithm to attack symmetric-key cryptographic primitives. Quantum Inf. Comput. **17**(1&2), 65–78 (2017)

24. Shor, P.W.: Algorithms for quantum computation: discrete logarithms and factoring. In: 35th Annual Symposium on Foundations of Computer Science, Santa Fe, New Mexico, USA, 20–22 November 1994, pp. 124–134. IEEE Computer Society (1994)

25. Simon, D.R.: On the power of quantum computation. SIAM J. Comput. **26**(5), 1474–1483 (1997)