Using Simon's Algorithm to Attack Symmetric-Key Cryptographic Primitives

Santoli, T.; Schaffner, C.

USING SIMON'S ALGORITHM TO ATTACK SYMMETRIC-KEY CRYPTOGRAPHIC PRIMITIVES

THOMAS SANTOLI[a]

Mathematical Institute, University of Oxford;
Andrew Wiles Building, Radcliffe Observatory Quarter
Woodstock Road, Oxford, OX2 6GG

CHRISTIAN SCHAFFNER[b]
ILLC, University of Amsterdam
Centrum Wiskunde & Informatica (CWI)
QuSoft
P.O. Box 94242, 1090 GE Amsterdam, Netherlands

We present new connections between quantum information and the field of classical cryptography. In particular, we provide examples where Simon's algorithm can be used to show insecurity of commonly used cryptographic symmetric-key primitives. Specifically, these examples consist of a quantum distinguisher for the 3-round Feistel network and a forgery attack on CBC-MAC which forges a tag for a chosen-prefix message querying only other messages (of the same length). We assume that an adversary has quantum-oracle access to the respective classical primitives. Similar results have been achieved recently in independent work by Kaplan *et al.* [KLLNP16]. Our findings shed new light on the post-quantum security of cryptographic schemes and underline that classical security proofs of cryptographic constructions need to be revisited in light of quantum attackers.

Keywords: quantum cryptanalysis, Simon's algorithm, Feistel network, CBC-MAC

1 Introduction

The main goal of cryptography is secure communication. While *encryption* ensures the secrecy of a message, its integrity is guaranteed by means of *authentication*. In private-key cryptography, it is assumed that honest players Alice and Bob share a private key unknown to any attacker. The situation is symmetric between Alice and Bob, hence this scenario is also referred to as *symmetric cryptography*. Here, private-key encryption (e.g. the Advanced Encryption Standard AES in a certain encryption mode) can be used to ensure secrecy of a message, while Message-Authentication Codes (MACs) allow to verify that a message originated from the secret-key holder and that the message has not been altered in transit. Cryp-

---

[a]thsantoli1@gmail.com

[b]c.schaffner@uva.nl

tographic research has come up with satisfactory solutions for these tasks, but even if we assume that these problems are solved, the problem of pre-establishing secret keys remains.

In groundbreaking work, very elegant solutions to this key-establishing problem have been suggested by Merkle [Mer78] and Diffie and Hellman [DH76]. The key idea is to split the key up into a public and private part, where only the latter needs to be kept secret.[c] However, compared to private-key cryptography, public-key cryptoschemes often require some richer mathematical structure which makes them more vulnerable to attacks.

Post-quantum cryptography [BBD09][d] studies the security of cryptographic systems against quantum attackers. Most currently used public-key systems (such as RSA [RSA78] or elliptic-curve cryptography) are known to be insecure against attackers with large-scale quantum computers, because their security relies on the hardness of mathematical problems such as factoring large integer numbers or taking discrete logarithms. For these problems, there is an exponential gap between the quantum and classical[e] running times of the best-known algorithms: these problems can be solved in polynomial time by Shor's algorithm [Sho94] on a large enough quantum computer, whereas we do not know of any polynomial-time classical algorithms for them.

In light of the fact that quantum algorithms break most of currently used public-key cryptography, it is natural to study the impact of quantum attacks also on *private-key* (or *symmetric*) cryptosystems. However, not much is known about the security of cryptographic primitives (such as hash functions or encryption schemes) under the assumption that an attack may involve quantum queries to the primitive. So, for example, apart from querying the primitive on a single input, a quantum adversary is allowed to query it on a superposition of inputs. Such a model is easily justified in case of a hash function [BDF+11] which is given by a classical algorithm. That algorithm can then be run on a quantum superposition of possible inputs. Other settings of superposition queries have been considered as well [BZ13b, DFNS14, BZ13a, GHS16]. The common belief so far was that quantum attacks on symmetric primitives are of minor concern, as they mainly consist of employing Grover's algorithm [Gro96] to generically speed up search (sub-)problems, for instance for finding collisions in hash functions. Grover search usually allows to get at most a quadratic improvement in the number of queries required to attack the primitive, see for instance [KLLNP16] for recent cryptanalytic results in this direction. In these cases, doubling the security parameter (such as the bit-size of the output of the hash function) would be a simple solution to overcome quantum attacks.

In this article, we exhibit two examples of commonly used classical symmetric-key cryptographic primitives which have well-established classical security proofs showing that an exponential number of queries is required for an attack. However, we demonstrate that security is broken by a quantum attacker which only makes a *polynomial* number of quantum queries. This is a similar *exponential gap* between classical and quantum security as in the devastating attacks using Shor's algorithm in the public-key setting. Our results shed new light on the relation between quantum information and classical cryptography, suggesting that

---

[c] In fact, public-key cryptography was invented already in the early 70s by Ellis, Cocks and Williams at the British intelligence agency GHCQ, but these works had been classified until 1997.

[d] The term is quite well-established by now, but chosen somewhat unfortunately, because the research area is concerned with cryptography which is still secure *at the beginning* and not after the end of the era of large-scale quantum computers.

[e] Here and throughout the article, we will use the term *classical* as synonym for *non-quantum*.

many classical cryptographic security proofs have been carried out in an incomplete (namely classical) security model and need to be revisited in the light of the currently most accurate description of Nature, which is quantum mechanics.

1.1 Simon's Algorithm

Simon's algorithm was proposed in 1994 [Sim94] as first quantum algorithm exhibiting an exponential speedup in query complexity compared to any classical algorithm (in the bounded-error setting). It solves the following problem: given a function $f : \{0,1\}^n \rightarrow \{0,1\}^n$ with the promise that for some $s \in \{0,1\}^n$, it holds that $[f(x) = f(y)] \Leftrightarrow [x \oplus y \in \{0^n, s\}]$, find the n-bit string s. It can be shown [MW16] that any classical algorithm requires $\Omega(2^{n/2})$ classical queries to f in order to find s, whereas Simon's algorithm succeeds using only $O(n)$ quantum queries to f. The quantum part of Simon's algorithm consists of executing the following circuit:

where the quantum query to the classical function f is formalized in the standard way by a unitary transform

$$U_f |x\rangle |y\rangle = |x\rangle |y \oplus f(x)\rangle \tag{1}$$

where $\oplus$ denotes bitwise addition modulo 2.

The first Hadamard operation leads to the state $\sum_{x \in \{0,1\}^n} |x\rangle |0\rangle$, which after querying f becomes $\sum_{x \in \{0,1\}^n} |x\rangle |f(x)\rangle$. Measuring the second register gives a random outcome $f(x)$ and collapses the first register to $\frac{1}{\sqrt{2}}(|x\rangle + |x \oplus s\rangle)$ because of the promised structure of f. Applying another Hadamard transform leads to the state $\sum_{y \in \{0,1\}^n} (-1)^{y \cdot x}(1 + (-1)^{y \cdot s}) |y\rangle$. Note that for $y \in \{0,1\}^n$ with $y \cdot s = y_1 s_1 \oplus \cdots \oplus y_n s_n = 1$, we have destructive interference and the amplitudes of those strings y vanish. Therefore, measuring the register will result in a (random) n-bit string y with the property that $y \cdot s = 0$. Running this quantum procedure $O(n)$ of times results in a full-rank system of linear equations which can be solved efficiently classically to obtain the string s.

1.2 Our Contributions

In this article, we observe that requiring only the forward implication of Simon's promise $x = y \oplus s \Rightarrow f(x) = f(y)$ suffices for Simon's subroutine to give a j such that $j \cdot s = 0$. The difference between this weakened condition and full Simon's promise is that we cannot bound the number of runs of the subroutine which are required in order to get an independent system of n such equations, and there are cases in which such an independent system cannot be obtained at all.

Our contributions build on a technique developed by Kuwakado and Morii. They make clever use of Simon's algorithm to distinguish a three-round Feistel construction with internal permutations from a random permutation (see explanations below) [KM10], and to recover the key in the Evan-Mansour encryption scheme [KM12]. In Section 3, we slightly extend the result from [KM10]. Using our observation above, we can relax the assumption from [KM10] that the internal functions have to be permutations, replacing them with arbitrary functions. Using random functions (instead of permutations) is relevant in the application where a

Feistel network is used to convert pseudorandom functions into pseudorandom permutations. In Section 4 we present the main new result of this article, a quantum forgery attack on a commonly used scheme for message authentication called CBC-MAC (of fixed-size messages). In our attack, we assume an adversary to be able to query arbitrary quantum superpositions of messages, provided these do not contain the message for which she forges a tag in the end. In fact, our attack also applies to an extension of the basic CBC-MAC to arbitrary-length messages by prepending the length of the message.

1.3   Related work

Concurrent to this work, very similar results to ours are presented in a recent article by Kaplan, Leurent, Leverrier, and Naya-Plasencia [KLLNP16]. While there has been an initial exchange of ideas in early stages of the work, the communication stopped after a while and our respective results have been obtained independently.

One difference between our work and [KLLNP16] is that we use different ways to handle the case in which Simon's promise is not fully satisfied, that is when there is an s such that $f(x) = f(x \oplus s)$ for all inputs s, but there can be unwanted collisions $f(x) = f(y)$ where $y \notin \{x, x \oplus s\}$. These unwanted collisions may increase the number of runs of Simon's subroutine needed to obtain an independent system of equations. We observe that this does not constitute a problem when it comes to distinguishing f from a random function. Instead, the authors of [KLLNP16] give a more general analysis of Simon's algorithm, and introduce a quantity to measure the number of the aforementioned unwanted collisions, and they show how to bound the number of needed repetitions of the subroutine in terms of this quantity.

Another difference between our works is the variant of CBC-MAC which we analyze. In our attack, we focus on MACs for fixed-size messages. In particular, we manage to forge a tag for a message (with a chosen prefix) by querying the MAC with messages of the same length. In [KLLNP16], they consider a slight variant of the basic CBC-MAC which can handle arbitrary-size messages by using two separate keys. While we show how to forge a tag for a message in which each block is non-zero by querying only messages which have at least one block of zeros, what they do is showing how to forge a tag for a message which is strictly longer than all messages queried in their attack.

Moreover, the results from [KLLNP16] illustrate that Simon's algorithm can be used to break not only CBC-MAC, but a whole range of modes of operation for authentication and authenticated encryption. Furthermore, Simon's algorithm can be used to speed up *slide attacks*, a cryptanalytic technique against classical symmetric cryptoschemes.

2   Pseudorandom Functions and Permutations

In this section, we provide formal definitions and statements about classical cryptography, following the lines of [KL07].

In the following, we consider *keyed functions*. These are two-input functions $F : \{0,1\}^* \times \{0,1\}^* \to \{0,1\}^*$. Given a keyed function F, we say that it is *efficient* if there exists a polynomial-time algorithm which computes $F(k, x)$ given k and x as inputs, and we say that it is *length-preserving* if it accepts only pairs of inputs (k, x) where k and x have the same length, and if also the output $F(k, x)$ has this same length.

For each $n \in \mathbb{N}$, let $\texttt{Func}_n$ be the set of all functions mapping n-bit strings to n-bit strings.

A length-preserving keyed function F induces a distribution on $\texttt{Func}_n$, given by randomly choosing $k \leftarrow \{0,1\}^n$ and then considering the single-input function $F_k(\cdot) := F(k, \cdot)$.

Intuitively we say that F is pseudorandom if, choosing k at random in $\{0,1\}^n$, the function F_k cannot be efficiently distinguished from a function f chosen uniformly at random from the set of all functions $\texttt{Func}_n$. In order to give the formal definition of pseudorandom function, we use the notion of *oracle-distinguisher*, which consists of an algorithm $D^{\mathcal{O}(\cdot)}$ having access to an oracle function $\mathcal{O}(\cdot)$. This means that in one time step, D can query the oracle at a point x, obtaining the value of $\mathcal{O}(x)$.

Definition 1. *Let* $F : \{0,1\}^* \times \{0,1\}^* \to \{0,1\}^*$ *be an efficient length-preserving keyed function. We call F a* pseudorandom function *if for any probabilistic polynomial-time oracle-distinguisher D there exists a negligible function $\texttt{negl}(n)$ such that*

$$\left| \Pr_{k \leftarrow \{0,1\}^n}[D^{F_k(\cdot)}(1^n) = 1] - \Pr_{f \leftarrow \texttt{Func}_n}[D^{f(\cdot)}(1^n) = 1] \right| \le \texttt{negl}(n)$$

where a negligible function is one which grows slower than any inverse polynomial in n.

An efficient length-preserving function $F : \{0,1\}^* \times \{0,1\}^* \to \{0,1\}^*$ is called *keyed permutation* if for every n and for every $k \in \{0,1\}^n$ the function $F_k(\cdot) : \{0,1\}^n \to \{0,1\}^n$ is a permutation (that is, a bijection). A keyed permutation is called *efficient* if there are polynomial-time algorithms which, on inputs k and x, compute $F_k(x)$ and $F_k^{-1}(x)$ respectively.

Let $\texttt{Perm}_n$ be the set of all permutations of $\{0,1\}^n$. In the same way as a length-preserving keyed function induces a distribution on $\texttt{Func}_n$, a keyed permutation induces a distribution on $\texttt{Perm}_n$. This leads to the definition of pseudorandom permutations, analogous to pseudorandom functions:

Definition 2. *Let* $F : \{0,1\}^* \times \{0,1\}^* \to \{0,1\}^*$ *be an efficient keyed permutation. We call F a* pseudorandom permutation *if for any probabilistic polynomial-time oracle-distinguisher D there exists a negligible function $\texttt{negl}(n)$ such that*

$$\left| \Pr_{k \leftarrow \{0,1\}^n}[D^{F_k(\cdot)}(1^n) = 1] - \Pr_{\pi \leftarrow \texttt{Perm}_n}[D^{\pi(\cdot)}(1^n) = 1] \right| \le \texttt{negl}(n).$$

We can use the Feistel construction to build a pseudorandom permutation from a pseudorandom function:

Theorem 3. *[LR88] If F is a pseudorandom function, then the 3-round Feistel network (as in Figure 1) with internal round functions $F_{k_1}, F_{k_2}, F_{k_3}$ is a pseudorandom permutation mapping $2n$-bit strings to $2n$-bit strings, and using a key $k = k_1 \| k_2 \| k_3$ of length $3n$.*

The notion of pseudorandomness of classical functions can be generalized to a post-quantum setting, where we assume that adversaries are allowed to perform quantum computation. In this context, we consider *quantum oracle-distinguishers*. A quantum oracle-distinguisher is a quantum algorithm $D^{|\mathcal{O}(\cdot)\rangle}$ which is allowed to make quantum queries of the form $|x, y\rangle \mapsto |x, y \oplus \mathcal{O}(x)\rangle$. This means that D can evaluate the function $\mathcal{O}(\cdot)$ not only on classical inputs, but also on superpositions of classical inputs. This leads to the following quantum generalizations of pseudorandomness:

Definition 4. *Let* $F : \{0,1\}^* \times \{0,1\}^* \to \{0,1\}^*$ *be an efficient length-preserving keyed function. We call F a* quantum pseudorandom function *if for any polynomial-time quantum*

oracle-distinguisher D there exists a negligible function `negl`(n) *such that*

$$\left| \Pr_{k \leftarrow \{0,1\}^n}[D^{|F_k(\cdot)\rangle}(1^n) = 1] - \Pr_{f \leftarrow \mathrm{Func}_n}[D^{|f(\cdot)\rangle}(1^n) = 1] \right| \leq \mathtt{negl}(n).$$

If $F : \{0,1\}^* \times \{0,1\}^* \rightarrow \{0,1\}^*$ *is an efficient keyed permutation, we call it a* quantum pseudorandom permutation *if for any polynomial-time quantum oracle-distinguisher D there exists a negligible function* `negl`(n) *such that*

$$\left| \Pr_{k \leftarrow \{0,1\}^n}[D^{|F_k(\cdot)\rangle}(1^n) = 1] - \Pr_{\pi \leftarrow \mathrm{Perm}_n}[D^{|\pi(\cdot)\rangle}(1^n) = 1] \right| \leq \mathtt{negl}(n).$$

Quantum pseudorandomness is a strictly stronger requirement than classical pseudorandomness. For example, by the distinguisher which we showed in Section 3, the 3-round Feistel network is not a quantum pseudorandom permutation, while by Theorem 3 it is a classical pseudorandom permutation.

3 Quantum distinguisher for the 3-round Feistel network

A Feistel network is a common cryptographic way of building block ciphers. Block ciphers are the basic building blocks of encryption schemes. Ideally, a block cipher with block size n bits behaves like a fully random permutation from $\{0,1\}^n$ to $\{0,1\}^n$. Many block ciphers such as DES (Data Encryption Standard) are based on a Feistel construction.

A Feistel network operates in a series of k rounds where internal randomizing functions $P_1, \ldots, P_k$ (from n bits to n bits) are cleverly put together to result in a permutation. More specifically, if the input of the i-th round is $L_{i-1} \| R_{i-1} \in \{0,1\}^{2n}$, the output produced by the round is $L_i \| R_i \in \{0,1\}^{2n}$ defined by $L_i := R_{i-1}$ and $R_i := L_{i-1} \oplus P_i(R_{i-1})$, see Figure 1. Interestingly, this classical cryptographic construction is not alien to quantum information theorists, as it is akin to the standard way of implementing a classical function P_i in a unitary way as in Equation (1).

Thus, a k-round Feistel network takes an input $L_0 \| R_0 \in \{0,1\}^{2n}$ and produces an output $L_k \| R_k \in \{0,1\}^{2n}$ by subsequently executing the round functions. One important feature of Feistel networks is that, regardless of the internal functions $P_1, \ldots, P_k$, the resulting function $L_0 \| R_0 \mapsto L_k \| R_k$ is always a permutation. Indeed, in order to invert the function one only needs to reverse the construction without the need of computing the inverse of the internal functions $P_1, \ldots, P_k$ (in other words, the internal functions do not have to be bijective). This is also handy in practice, because it makes the decryption process very similar to the encryption process.

Classically, Theorem 3 states that a 3-round Feistel network constitutes a cryptographically secure pseudorandom permutation as long as the internal functions are pseudorandom as well. That is, a polynomial-time classical distinguisher cannot distinguish a 3-round Feistel network from a random permutation, see Section 2 for formal definitions and statements of this classical result. Quantumly, this implication does not hold, and in this section we show a polynomial-time quantum attacker which can distinguish a 3-round Feistel network (with internal random functions) from a fully random permutation. Our proof follows the lines of the results by Kuwakado and Morii [KM10]. The main difference is that they assume the internal functions of the Feistel network to be permutations, while here we can handle the more general scenario of arbitrary (random) functions.
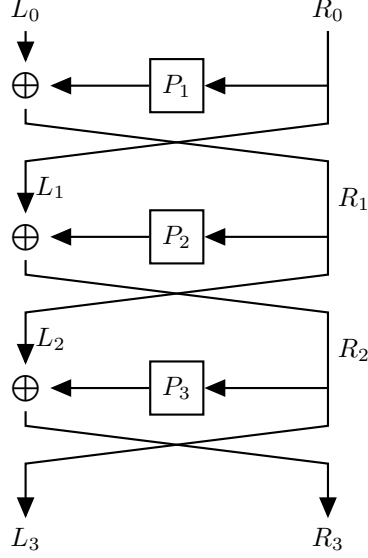
Fig. 1. A three-round Feistel network with internal functions P_1, P_2, P_3.

In this scenario, a distinguishing adversary has an oracle $V : \{0,1\}^{2n} \to \{0,1\}^{2n}$ which is either a fully random permutation of $2n$-bit strings, or a 3-round Feistel network with internal random functions P_1, P_2, P_3, that is

$$V(a\|c) = c \oplus P_2(a \oplus P_1(c)) \ \| \ a \oplus P_1(c) \oplus P_3(c \oplus P_2(a \oplus P_1(c))) \,.$$

Let $W : \{0,1\}^{2n} \to \{0,1\}^n$ be the first n bits of V: $W(a\|c) = c \oplus P_2(a \oplus P_1(c))$. Let $\alpha, \beta \in \{0,1\}^n$ be distinct fixed strings, and let us consider the function $f : \{0,1\}^{n+1} \to \{0,1\}^n$ defined as

$$f(b\|a) := \begin{cases} W(a\|\alpha) \oplus \alpha & \text{if } b = 0 \\ W(a\|\beta) \oplus \beta & \text{if } b = 1 \end{cases}$$

for all $b \in \{0,1\}$ and $a \in \{0,1\}^n$. In case V is the Feistel network, we have

$$f(b\|a) = \begin{cases} \alpha \oplus P_2(a \oplus P_1(\alpha)) \oplus \alpha & \text{if } b = 0 \\ \beta \oplus P_2(a \oplus P_1(\beta)) \oplus \beta & \text{if } b = 1 \end{cases}$$

$$= \begin{cases} P_2(a \oplus P_1(\alpha)) & \text{if } b = 0 \\ P_2(a \oplus P_1(\beta)) & \text{if } b = 1 \end{cases} .$$

For each y in the image $\text{img}(P_2)$ of P_2, denote its pre-image by $P_2^{-1}(y) = \{x^1, \ldots, x^{l_y}\}$, we have

$$f(b\|a) = y$$
$$\Leftrightarrow b\|a \in \{0\|x^i \oplus P_1(\alpha)\}_{i=1\ldots l_y} \cup \{1\|x^i \oplus P_1(\beta)\}_{i=1\ldots l_y}$$
$$= \{a^1, \ldots, a^{l_y}\} \cup \{a^1 \oplus s, \ldots, a^{l_y} \oplus s\}$$

where $s = 1 \| (P_1(\alpha) \oplus P_1(\beta))$ and $a^i = 0 \| (x^i \oplus P_1(\alpha))$ for $i = 1 \ldots l_y$. In particular, in case V is a Feistel network, for each $u \in \{0,1\}^{n+1}$, we have $f(u) = f(u \oplus s)$. On the other hand, in case V is a random permutation, no such relation is expected to hold.

Using at most two invocations of the unitary operator $U_V : |x, y\rangle \mapsto |x, y \oplus V(x)\rangle$ for computing V, we can construct the unitary operator $U_f : |u, y\rangle \mapsto |u, y \oplus f(u)\rangle$. Using this operator, one can mount the following attack, based on Simon's algorithm, to distinguish whether V is a random permutation or a 3-round Feistel network:

1. Initialize X as the empty set.

2. Run Simon's subroutine:

3. Add the result $j = (j_0, j_1, \ldots, j_n)$ to X. If X does not contain n linearly independent j's, go back to Step 2, unless X already contains $2n$ elements. In the latter case, guess V to be a 3-round Feistel network. Otherwise, solve this linear system of equations:

$$\begin{pmatrix} | & & | \\ j_1 & \cdots & j_n \\ | & & | \end{pmatrix} \begin{pmatrix} s_1 \\ \vdots \\ s_n \end{pmatrix} = \begin{pmatrix} | \\ j_0 \\ | \end{pmatrix}$$

4. Choose $u \in \{0,1\}^{n+1}$ at random, and guess V to be a 3-round Feistel network if $f(u) = f(u \oplus s)$, where $s = (1, s_1, \ldots, s_n)$. Otherwise guess V to be a random permutation.

Let us first analyze the case when V is a 3-round Feistel network. After applying U_f, we have the state

$$\frac{1}{\sqrt{2^{n+1}}} \sum_{u \in \{0,1\}^{n+1}} |u, f(u)\rangle = \frac{1}{\sqrt{2^{n+1}}} \sum_{y \in \mathrm{img}(P_2)} \sum_{i=1}^{l_y} (|a^i\rangle + |a^i \oplus s\rangle) |y\rangle \ .$$

Before measuring, we get

$$\frac{1}{2^{n+1}} \sum_{y \in \mathrm{img}(P_2)} \sum_{i=1}^{l_y} \left(\sum_{j \in \{0,1\}^{n+1}} (-1)^{a^i \cdot j} |j\rangle + \sum_{j \in \{0,1\}^{n+1}} (-1)^{(a^i \oplus s) \cdot j} |j\rangle \right) |y\rangle$$

$$= \frac{1}{2^{n+1}} \sum_{y \in \mathrm{img}(P_2)} \sum_{j \in \{0,1\}^{n+1}} \sum_{i=1}^{l_y} (-1)^{a^i \cdot j} (1 + (-1)^{s \cdot j}) |j\rangle |y\rangle$$

$$= \frac{1}{2^n} \sum_{y \in \mathrm{img}(P_2)} \left(\sum_{\substack{j \in \{0,1\}^{n+1} \\ s \cdot j = 0}} \left(\sum_{i=1}^{l_y} (-1)^{a^i \cdot j} \right) |j\rangle \right) |y\rangle \ .$$

Hence, all measurement outcomes j satisfy $s \cdot j = 0 \mod 2$, so if at some point we solve the system in Step 3 we get $(s_1, \ldots, s_n) = P_1(\alpha) \oplus P_1(\beta)$, and the condition $f(u) = f(u \oplus s)$ in Step 4 is satisfied for all $u \in \{0,1\}^{n+1}$, so the algorithm gives the correct answer. If we do not reach the step of solving the system because after collecting $2n$ vectors we still do not have n independent vectors, then we give the correct answer because we will guess V to be the 3-round Feistel network.

If V is a random permutation, then the j's will be uniform random vectors. In that case, we make an error if we do not have n independent vectors after $2n$ extractions, or if $f(u) = f(u \oplus s)$ in Step 4. Notice that the probability that after $2n$ extractions of uniform vectors, we still do not have n independent vectors is negligible in n.[f] In case we continue to solve a random system of independent equations in Step 3, we get a random vector $s_1, \ldots, s_n$ as a result. Then, at Step 4, the algorithm makes an error if the randomly chosen $u \in \{0,1\}^{n+1}$ satisfies the condition $f(u) = f(u \oplus s)$, but this also happens only with negligible probability. So we conclude that also if V is a random permutation, the algorithm gives a wrong answer by guessing it to be a 3-round Feistel network with only negligible probability.

In summary, we have constructed a quantum adversary which can efficiently distinguish a 3-round Feistel construction with internal functions from a fully random permutation. This extends the results from [KM10] to the important setting where a Feistel construction is used to turn pseudorandom functions into pseudorandom permutations.

4   Quantum forgery attack on CBC-MAC

The goal of message authentication and integrity is to allow communicating parties to verify that a received message comes from the expected sender, and that the message has not been altered in transit. Parties who share a secret key can use Message Authentication Codes (MACs) to achieve this goal: the sender feeds the key k and the message m to an algorithm $\mathsf{Mac}(k, m)$ which computes a tag t for the message, and sends the pair (m, t) to the receiver. Then the receiver, using the same key k, can check whether t is a valid tag for m. A message authentication code is considered secure if no efficient adversary who can ask arbitrary adaptive queries to the oracle $\mathsf{Mac}(k, \cdot)$ (where k is uniform and unknown to the adversary) can produce a valid tag for a message m that has not been queried before.

CBC-MAC is an efficient way of constructing a message authentication code based on a pseudorandom permutation (formally defined in Section 2). CBC-MAC is widely used in practice, and comes in many variants. Given a pseudorandom permutation $F : K \times \{0,1\}^n \to \{0,1\}^n$, the basic CBC-MAC construction is defined as

$$\mathrm{CBC}^\ell : K \times \{0,1\}^{\ell n} \to \{0,1\}^n$$
$$(k, m_1 \| m_2 \| \ldots \| m_\ell) \mapsto \mathrm{CBC}^\ell_{F_k}(m_1 \| m_2 \| \ldots \| m_\ell),$$

where

$$\mathrm{CBC}^\ell_{F_k}(m_1 \| m_2 \| \ldots \| m_\ell) = F_k(F_k(\ldots F_k(F_k(m_1) \oplus m_2) \oplus \ldots) \oplus m_\ell).$$

---

[f] This can be proven as follows: the probability that n uniform vectors of length $2n$ are linearly independent is exactly $\Pi_{i=1}^n (1 - 2^{i-1-2n}) \geq 1 - 2^{-n}$. In other words, the column rank of a random binary $2n \times n$ matrix is n except with probability at most 2^{-n}. The claim then follows directly from the equality of row and column rank, see e.g. [War05].

Fig. 2. Illustration of the computation of $t = CBC^3_{F_k}(m_1 \| m_2 \| m_3)$.

The case $\ell = 3$ is depicted in Figure 2.

Using this construction, a sender who wants to send a message m of length ℓn can compute a tag $t := \mathrm{CBC}^\ell(k, m)$ and send the tuple (m, t) to the receiver who can verify whether $t = \mathrm{CBC}^\ell(k, m)$.

One can prove that CBC^ℓ is a pseudorandom function. Therefore, as explained in [BKR00], it (classically) constitutes a secure MAC. Quantumly, even if F is a quantum-secure pseudorandom permutation, this construction does not yield a secure MAC, and not even a quantum pseudorandom function (qPRF) as defined in Section 2.

According to the definition given in [BZ13a], a signature scheme is quantum secure if there is no quantum attacker who can produce $q + 1$ valid tags by making only q queries. This notion of security is also the one used in [KLLNP16]. In this work, we use a different notion of security. Namely, we show how to forge a valid tag for a chosen-prefix message without ever querying that particular message (not even as small part of a big superposition query). We note that basic CBC-MAC is only secure for fixed-length messages, but fails in case the attacker is allowed to query and forge different-length messages. However, we would like to stress that our quantum attack does not rely on that weakness, but also breaks the extension of CBC-MAC where the length of the messages is prepended. This extension can be proven secure classically.

In the following, we describe our forgery attack on CBC-MAC for messages of fixed length ℓn, where $\ell \geq 3$. We want to forge a tag for a message which has some fixed prefix $\beta \in \{0, 1\}^{kn}$ of length kn for some $k \leq \ell - 2$. For simplicity, we restrict ourselves to the case $k = 1$. In Section 4.1, we explain how to extend the attack to prefixes of arbitrary size $k \leq \ell - 2$.

In a first step, we assume that the pseudorandom permutation F_k is replaced with a fully random permutation π. If this replacement made a difference in our forging algorithm, then this difference would allow us to use our algorithm to distinguish F_k from a fully random permutation, in contradiction with the fact that F is a qPRF.

Suppose we want to tag a message having $\alpha_1 \in \{0, 1\}^n$ as a prefix. Choose some $\alpha_0 \neq \alpha_1$. For $j = 1 \ldots \ell - 1$, let us consider the following functions:

$$g_j : \{0,1\}^{n+1} \to \{0,1\}^n$$
$$b\|x \mapsto \mathrm{CBC}_\pi^\ell(\alpha_b\|0^{(j-1)n}\|x\|0^{(\ell-j-1)n}) = \pi^{\ell-j}(\pi^j(\alpha_b) \oplus x)$$

Note that, for distinct $u, v \in \{0,1\}^{n+1}$, we have $g_j(u) = g_j(v)$ if and only if $v = u \oplus 1\|(\pi^j(\alpha_0) \oplus \pi^j(\alpha_1))$.

Using one invocation of the unitary operator $U_{\mathrm{CBC}_\pi^\ell}$ we can construct the unitary operator U_{g_j}, for $j = 1 \ldots \ell - 1$. A quantum adversary can use these operators to forge a tag on a message without querying it as follows:

1. For $j = 1 \ldots \ell - 1$, run Simon's subroutine on U_{g_j} in order to find $z^j := \pi^j(\alpha_0) \oplus \pi^j(\alpha_1)$.[g] Note that all queries made in this process are on (superpositions of) messages $m_1\|m_2\| \ldots \|m_\ell$ such that $m_i = 0^n$ for at least one $i \in \{1 \ldots \ell\}$ [h]

2. Classically query $t_0 := \mathrm{CBC}_\pi^\ell(\alpha_0\|0^{(\ell-1)n}) = \pi^\ell(\alpha_0)$ and $t_1 := \mathrm{CBC}_\pi^\ell(\alpha_1\|0^{(\ell-1)n}) = \pi^\ell(\alpha_1)$.
   Also these queries are made on messages $m_1\|m_2\| \ldots \|m_\ell$ such that $m_i = 0^n$ for at least one $i \in \{1 \ldots \ell\}$.

3. If ℓ is even, forge $(m, t) := (\alpha_1\|z^1\|z^2\| \ldots \|z^{\ell-1}, t_0)$. If ℓ is odd, forge $(m, t) := (\alpha_1\|z^1\|z^2\| \ldots \|z^{\ell-1}, t_1)$.

The message $m = \alpha_1\|z^1\|z^2\| \ldots \|z^{\ell-1}$ is never queried in the above procedure, where we query only messages $m_1\|m_2\| \ldots \|m_\ell$ such that $m_i = 0^n$ for at least one $i \in \{1 \ldots \ell\}$. In fact, as the π^i are permutations, we have $z^j = \pi^j(\alpha_0) \oplus \pi^j(\alpha_1) \neq 0^n$ for all j's.

Therefore, in order to show that the above forging attack is successful, it remains to prove that t is a valid tag for m, which means that $\mathrm{CBC}_\pi^\ell(m) = t$, where $t = t_0 = \pi^\ell(\alpha_0)$ if ℓ is even and $t = t_1 = \pi^\ell(\alpha_1)$ if ℓ is odd. This claim is proven by induction on $\ell \geq 3$:

We have two base cases, $\ell = 3$ and $\ell = 4$. For $\ell = 3$ we have:

$$\mathrm{CBC}_\pi^3(\alpha_1\|z^1\|z^2) = \pi(\pi(\pi(\alpha_1) \oplus z^1) \oplus z^2)$$
$$= \pi\left(\pi\left(\pi(\alpha_1) \oplus \pi(\alpha_0) \oplus \pi(\alpha_1)\right) \oplus \pi^2(\alpha_0) \oplus \pi^2(\alpha_1)\right)$$
$$= \pi((\pi^2(\alpha_0) \oplus \pi^2(\alpha_0) \oplus \pi^2(\alpha_1)) = \pi^3(\alpha_1) = t_1 \ .$$

For $\ell = 4$ we similarly argue that

$$\mathrm{CBC}_\pi^4(\alpha_1\|z^1\|z^2\|z^3) = \pi(\pi(\pi(\pi(\alpha_1) \oplus z^1) \oplus z^2) \oplus z^3)$$
$$= \pi\Big(\pi(\pi(\pi(\alpha_1) \oplus \pi(\alpha_0) \oplus \pi(\alpha_1)) \oplus \pi^2(\alpha_0) \oplus \pi^2(\alpha_1)) \oplus \pi^3(\alpha_0) \oplus \pi^3(\alpha_1)\Big)$$
$$= \pi(\pi(\pi^2(\alpha_0) \oplus \pi^2(\alpha_0) \oplus \pi^2(\alpha_1)) \oplus \pi^3(\alpha_0) \oplus \pi^3(\alpha_1))$$
$$= \pi(\pi^3(\alpha_1) \oplus \pi^3(\alpha_0) \oplus \pi^3(\alpha_1)) = \pi^4(\alpha_0) = t_0 \ .$$

---

[g] This step of the algorithm shows that also for the case $\ell = 2$, CBC_π^ℓ is not a quantum pseudorandom permutation: We can run Simon's algorithm with g_1 (which exists also in the case $\ell = 2$) to get a solution $1\|z^1$. Then, we pick u at random and check if $g_1(u) = g_1(u \oplus 1\|z^1)$ holds. For CBC_π^ℓ, this will always be the case, whereas for a fully random permutation, the check will only be passed with negligible probability.

[h] This would not be the case for $\ell = 2$, which is why we require $\ell \geq 3$ for the attack.

If $\ell + 1$ is even then ℓ is odd, and the induction step works as

$$\text{CBC}_\pi^{\ell+1}(\alpha_1 \| z^1 \| z^2 \| \ldots \| z^{\ell-1} \| z^\ell)$$
$$= \pi(\text{CBC}_\pi^\ell(\alpha_1 \| z^1 \| z^2 \| \ldots \| z^{\ell-1}) \oplus z^\ell)$$
$$= \pi(\pi^\ell(\alpha_1) \oplus \pi^\ell(\alpha_0) \oplus \pi^\ell(\alpha_1)) = \pi^{\ell+1}(\alpha_0) = t_0 \ .$$

Otherwise, if $\ell + 1$ is odd, then ℓ is even and one can argue similarly that

$$\text{CBC}_\pi^{\ell+1}(\alpha_1 \| z^1 \| z^2 \| \ldots \| z^{\ell-1} \| z^\ell)$$
$$= \pi(\text{CBC}_\pi^\ell(\alpha_1 \| z^1 \| z^2 \| \ldots \| z^{\ell-1}) \oplus z^\ell)$$
$$= \pi(\pi^\ell(\alpha_0) \oplus \pi^\ell(\alpha_0) \oplus \pi^\ell(\alpha_1)) = \pi^{\ell+1}(\alpha_1) = t_1 \ .$$

Thus, we can conclude by induction that the output pair (m, t) is correct.

4.1   Extension of the Attack on CBC-MAC

The algorithm presented in this section can be adapted for forging a tag for a message having arbitrary $\beta_1 \| \cdots \| \beta_k \in \{0,1\}^{kn}$ as a prefix, for some $k \leq \ell - 2$.

Choose $\alpha_1 \| \cdots \| \alpha_k$ such that

$$\overline{\alpha} := \pi(\cdots \pi(\pi(\alpha_1) \oplus \alpha_2) \cdots) \oplus \alpha_k \neq \pi(\cdots \pi(\pi(\beta_1) \oplus \beta_2) \cdots) \oplus \beta_k =: \overline{\beta} \ .$$

This inequality can be achieved by repeatedly choosing $\alpha_1 \| \cdots \| \alpha_k$ at random and querying $\text{CBC}_\pi^\ell(\alpha_1 \| \cdots \| \alpha_k \| 0^{(\ell-k)n}) = \pi^{\ell-k}\Big(\pi(\pi(\cdots \pi(\pi(\alpha_1) \oplus \alpha_2) \cdots) \oplus \alpha_k)\Big) = \pi^{\ell-k}(\pi(\overline{\alpha})) = \pi^{\ell-k+1}(\overline{\alpha})$ and $\text{CBC}_\pi^\ell(\beta_1 \| \cdots \| \beta_k \| 0^{(\ell-k)n}) = \pi^{\ell-k}\Big(\pi(\pi(\cdots \pi(\pi(\beta_1) \oplus \beta_2) \cdots) \oplus \beta_k)\Big) = \pi^{\ell-k}(\pi(\overline{\beta})) = \pi^{\ell-k+1}(\overline{\beta})$ until $\pi^{\ell-k+1}(\overline{\alpha})$ and $\pi^{\ell-k+1}(\overline{\beta})$ are different (which will already be the case after a single try except with negligible probability). Then, since $\pi^{\ell-k+1}(\overline{\alpha}) \neq \pi^{\ell-k+1}(\overline{\beta})$, we must also have $\overline{\alpha} \neq \overline{\beta}$.

We have
$$\text{CBC}_\pi^\ell(\alpha_1 \| \cdots \| \alpha_k \| 0^{(\ell-k)n}) = \text{CBC}_\pi^{\ell-k+1}(\overline{\alpha} \| 0^{(\ell-k)n})$$

and similarly
$$\text{CBC}_\pi^\ell(\beta_1 \| \cdots \| \beta_k \| 0^{(\ell-k)n}) = \text{CBC}_\pi^{\ell-k+1}(\overline{\beta} \| 0^{(\ell-k)n}) \ .$$

So, for $j = 1 \ldots \ell - k$ we define

$$g_j : \{0,1\}^{n+1} \to \{0,1\}^n$$
$$b \| x \mapsto \begin{cases} \text{CBC}_\pi^\ell(\alpha_1 \| \cdots \| \alpha_k \| 0^{(j-1)n} \| x \| 0^{(\ell-k-j)n}) & \text{if } b = 0 \\ \text{CBC}_\pi^\ell(\beta_1 \| \cdots \| \beta_k \| 0^{(j-1)n} \| x \| 0^{(\ell-k-j)n}) & \text{if } b = 1 \end{cases}$$
$$= \begin{cases} \text{CBC}_\pi^{\ell-k+1}(\overline{\alpha} \| 0^{(j-1)n} \| x \| 0^{(\ell-k-j)n}) & \text{if } b = 0 \\ \text{CBC}_\pi^{\ell-k+1}(\overline{\beta} \| 0^{(j-1)n} \| x \| 0^{(\ell-k-j)n}) & \text{if } b = 1 \end{cases}$$
$$= \begin{cases} \pi^{\ell-k+1-j}(\pi^j(\overline{\alpha}) \oplus x) & \text{if } b = 0 \\ \pi^{\ell-k+1-j}(\pi^j(\overline{\beta}) \oplus x) & \text{if } b = 1 \end{cases} \ .$$

Then we can run the algorithm from Section 4 because $\ell - k + 1 \geq 3$. Thereby, we obtain a tag for the message $m = \beta_1 \| \cdots \| \beta_k \| z^1 \| \cdots \| z^{\ell-k}$, where $z^j = \pi^j(\overline{\alpha}) \oplus \pi^j(\overline{\beta})$, and only (superpositions of) messages with at least one zero-block have been queried in the process.

5   Conclusion and Open Questions

We have shown two applications of Simon's algorithm to breaking classical symmetric cryptoschemes. As illustrated in the independent work of [KLLNP16], the same attack technique can be applied in various other cryptographic scenarios as well. These results shed new light on the post-quantum cryptographic security of classical cryptography and suggest that many classical security proofs need to be revisited in the light of quantum attackers.

Very recent work by Zhandry [Zha16] shows the existence of quantum-secure pseudo-random permutations (PRPs). However, the construction does not use Feistel networks. Therefore, it remains an intriguing open question to prove that a constant number (strictly larger than 3) of Feistel rounds suffices to obtain a quantum-secure pseudorandom permutation from a quantum secure PRF. In the case of message-authentication codes, our work suggests that one has to be very careful in using popular constructions such as CBC-MAC.

Acknowledgements

References

[BBD09]   D. J. Bernstein, J. Buchmann, and E. Dahmen, editors. *Post-Quantum Cryptography*. Springer, 2009.
DOI: `10.1007/978-3-540-88702-7`.

[BDF+11]   D. Boneh, Ö. Dagdelen, M. Fischlin, A. Lehmann, C. Schaffner, and M. Zhandry. Random oracles in a quantum world. In *Advances in Cryptology—ASIACRYPT 2011*, pages 41–69, 2011.
DOI: `10.1007/978-3-642-25385-0_3`.

[BKR00]   M. Bellare, J. Kilian, and P. Rogaway. The security of the cipher block chaining message authentication code. *Journal of Computer and System Sciences*, 61(3): 362–399, 2000.
DOI: `10.1006/jcss.1999.1694`.

[BZ13a]   D. Boneh and M. Zhandry. Quantum-secure message authentication codes. In T. Johansson and P. Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 592–608, Athens, Greece, 2013. Springer, Heidelberg, Germany.
DOI: `10.1007/978-3-642-38348-9_35`.

[BZ13b]   D. Boneh and M. Zhandry. Secure signatures and chosen ciphertext security in a quantum computing world. In *Advances in Cryptology—CRYPTO 2013*, pages 361–379, 2013.
DOI: `10.1007/978-3-642-40084-1_21`.

[DFNS14]   I. Damgård, J. Funder, J. B. Nielsen, and L. Salvail. Superposition attacks on cryptographic protocols. In *Information Theoretic Security*, pages 142–161, 2014.
DOI: `10.1007/978-3-319-04268-8_9`.

[DH76]   W. Diffie and M. Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, 22(6): 644–654, 1976.
DOI: `10.1109/tit.1976.1055638`.

[GHS16]   T. Gagliardoni, A. Hülsing, and C. Schaffner. Semantic security and indistinguishability in the quantum world. In M. Robshaw and J. Katz, editors, *CRYPTO 2016, Part III*, volume 9816 of *LNCS*, pages 60–89, Santa Barbara, CA, USA, 2016. Springer, Heidelberg, Germany.
DOI: `10.1007/978-3-662-53015-3_3`.

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

[KL07]   J. Katz and Y. Lindell. *Introduction to Modern Cryptography (Chapman & Hall/Crc Cryptography and Network Security Series)*. Chapman & Hall/CRC, 2007.

[KLLNP16]   M. Kaplan, G. Leurent, A. Leverrier, and M. Naya-Plasencia. Breaking symmetric cryptosystems using quantum period finding. In M. Robshaw and J. Katz, editors, *CRYPTO 2016, Part II*, volume 9815 of *LNCS*, pages 207–237, Santa Barbara, CA, USA, 2016. Springer, Heidelberg, Germany.
            DOI: `10.1007/978-3-662-53008-5_8`.

[KM10]   H. Kuwakado and M. Morii. Quantum distinguisher between the 3-round feistel cipher and the random permutation. In *Information Theory Proceedings (ISIT), 2010 IEEE International Symposium on*, pages 2682–2685. IEEE, 2010.

[KM12]   H. Kuwakado and M. Morii. Security on the quantum-type even-mansour cipher. In *Information Theory and its Applications (ISITA), 2012 International Symposium on*, pages 312–316. IEEE, 2012.

[LR88]   M. Luby and C. Rackoff. How to construct pseudorandom permutations from pseudorandom functions. *SIAM Journal on Computing*, 17(2), 1988.

[Mer78]   R. C. Merkle. Secure communications over insecure channels. *Communications of the ACM*, 21(4): 294–299, 1978.
            DOI: `10.1145/359460.359473`.

[MW16]   A. Montanaro and R. d. Wolf. *A Survey of Quantum Property Testing*. Number 7 in Graduate Surveys. Theory of Computing Library, 2016.
            DOI: `10.4086/toc.gs.2016.007`.

[RSA78]   R. L. Rivest, A. Shamir, and L. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Communications of the ACM*, 21(2): 120–126, 1978.
            DOI: `10.1145/359340.359342`.

[Sho94]   P. W. Shor. Algorithms for quantum computation: discrete logarithms and factoring. In *35th Annual Symposium on Foundations of Computer Science—FOCS 1994*, pages 124–134. IEEE, 1994.
            DOI: `10.1109/SFCS.1994.365700`.

[Sim94]   D. R. Simon. On the power of quantum cryptography. In *35th FOCS*, pages 116–123, Santa Fe, New Mexico, 1994. IEEE Computer Society Press.

[War05]   W. P. Wardlaw. Row rank equals column rank. *Mathematics Magazine*, 78(4): 316–318, 2005.
            Online: `http://www.jstor.org/stable/30044181`.

[Zha16]   M. Zhandry. A note on quantum-secure prps, 2016.