

Quantum-Secure Message Authentication Codes

Dan Boneh and Mark Zhandry

Stanford University
{dabo,zhandry}@cs.stanford.edu

Abstract. We construct the first Message Authentication Codes (MACs) that are existentially unforgeable against a *quantum* chosen message attack. These chosen message attacks model a quantum adversary’s ability to obtain the MAC on a superposition of messages of its choice. We begin by showing that a quantum secure PRF is sufficient for constructing a quantum secure MAC, a fact that is considerably harder to prove than its classical analogue. Next, we show that a variant of Carter-Wegman MACs can be proven to be quantum secure. Unlike the classical settings, we present an attack showing that a pair-wise independent hash family is insufficient to construct a quantum secure *one-time* MAC, but we prove that a four-wise independent family is sufficient for one-time security.

Keywords: Quantum computing, MAC, chosen message attacks, post-quantum security

1 Introduction

Message Authentication Codes (MACs) are an important building block in cryptography used to ensure data integrity. A MAC system is said to be secure if an efficient attacker capable of mounting a chosen message attack cannot produce an existential MAC forgery (see Section 2.2).

With the advent of quantum computing there is a strong interest in post-quantum cryptography, that is systems that remain secure even when the adversary has access to a quantum computer. There are two natural approaches to defining security of a MAC system against a quantum adversary. One approach is to restrict the adversary to issue classical chosen message queries, but then allow the adversary to perform quantum computations between queries. Security in this model can be achieved by basing the MAC construction on a quantum intractable problem.

The other more conservative approach to defining quantum MAC security is to model the entire security game as a quantum experiment and allow the adversary to issue *quantum* chosen message queries. That is, the adversary can submit a superposition of messages from the message space and in response receive a superposition of MAC tags on those messages. Informally, a quantum chosen message query performs the following transformation on a given superposition of messages:

$$\sum_m \psi_m |m\rangle \longrightarrow \sum_m \psi_m |m, S(k, m)\rangle$$

where $S(k, m)$ is a tag on the message m with secret key k .

To define security, let q be the number of queries that the adversary issues by the end of the game. Clearly it can now produce q classical message-tag pairs by sampling the q superpositions it received from the MAC signing oracle. We say that the MAC system is *quantum secure* if the adversary cannot produce $q + 1$ valid message-tag pairs. This captures the fact that the adversary cannot do any better than trivially sampling the responses from its MAC signing oracle and is the quantum analogue of a classical existential forgery.

1.1 Our results

In this paper we construct the first quantum secure MAC systems. We begin with a definition of quantum secure MACs and give an example of a MAC system that is secure against quantum adversaries capable of *classical* chosen message queries, but is insecure when the adversary can issue *quantum* chosen message queries. We then present a number of quantum secure MAC systems.

Quantum secure MACs. In the classical settings many MAC systems are based on the observation that a secure pseudorandom function gives rise to a secure MAC [BKR00,BCK96]. We begin by studying the same question in the quantum settings. Very recently Zhandry [Zha12b] defined the concept of a quantum secure pseudorandom function (PRF) which is a PRF that remains indistinguishable from a random function even when the adversary can issue *quantum* queries to the PRF. He showed that the classic GGM construction [GGM86] remains secure under quantum queries assuming the underlying pseudorandom generator is quantum secure.

The first question we study is whether a quantum secure PRF gives rise to a quantum secure MAC, as in the classical settings. To the MAC adversary a quantum secure PRF is indistinguishable from a random function. Therefore proving that the MAC is secure amounts to proving that with q quantum queries to a random oracle H no adversary can produce $q + 1$ input-output pairs of H with non-negligible probability. In the classical settings where the adversary can only issue *classical* queries to H this is trivial: given q evaluations of a random function, the adversary learns nothing about the value of the function at other points. Unfortunately, this argument fails under quantum queries because the response to a *single* quantum query to $H : \mathcal{X} \rightarrow \mathcal{Y}$ contains information about all of H . In fact, with a single quantum query the adversary can produce two input-output pairs of H with probability about $2/|\mathcal{Y}|$ (with classical queries the best possible is $1/|\mathcal{Y}|$). As a result, proving that q quantum queries are insufficient to produce $q + 1$ input-output pairs is quite challenging. We prove tight upper and lower bounds on this question by proving the following theorem:

Theorem 1 (informal). *Let $H : \mathcal{X} \rightarrow \mathcal{Y}$ be a random oracle. Then an adversary making at most $q < |\mathcal{X}|$ quantum queries to H will produce $q + 1$ input-output pairs of H with probability at most $(q + 1)/|\mathcal{Y}|$. Furthermore, when $q \ll |\mathcal{Y}|$ there is an algorithm that with q quantum queries to H will produce $q + 1$ input-output pairs of H with probability $1 - (1 - 1/|\mathcal{Y}|)^{q+1} \approx (q + 1)/|\mathcal{Y}|$.*

The first part of the theorem is the crucial fact needed to build quantum secure MACs and is the harder part to prove. It shows that when $|\mathcal{Y}|$ is large any algorithm has only a negligible chance in producing $q + 1$ input-output pairs of H from q quantum queries. To prove this bound we introduce a new lower-bound technique we call the *rank method* for bounding the success probability of algorithms that succeed with only small probability. Existing quantum lower bound techniques such as the polynomial method [BBC⁺01] and the adversary method [Amb00,Aar02,Amb06,ASdW09] do not give the result we need. One difficulty with existing lower bound techniques is that they generally prove asymptotic bounds on the number of queries required to solve a problem with high probability, whereas we need a bound on the success probability of an algorithm making a limited number of queries. Attempting to apply existing techniques to our problem at best only bounds the success probability away from 1 by an inverse polynomial factor, which is insufficient for our purposes. The rank method for proving quantum lower bounds overcomes these difficulties and is a general tool that can be used in other post-quantum security proofs.

The second part of Theorem 1 shows that the lower bound presented in the first part of the theorem is tight. A related algorithm was previously presented by van Dam [vD98], but only for oracles outputting one bit, namely when $\mathcal{Y} = \{0, 1\}$. For such a small range only about $|\mathcal{X}|/2$ quantum queries are needed to learn the oracle at all $|\mathcal{X}|$ points. A special case where $\mathcal{Y} = \mathcal{X} = \{0, 1\}$ and $q = 1$ was developed independently by Kerenidis and de Wolf [KdW03]. Our algorithm is a generalization of van Dam’s result to multi-bit oracles.

Quantum secure Carter-Wegman MACs. A Carter-Wegman MAC [WC81] signs a message m by computing $(r, h(m) \oplus F(k, r))$ where h is a secret hash function chosen from an xor-universal hash family, F is a secure PRF with secret key k , and r is a short random nonce. The attractive feature of Carter-Wegman MACs is that the long message m is hashed by a fast xor-universal hash h . We show that a slightly modified Carter-Wegman MAC is quantum secure assuming the underlying PRF is quantum secure in the sense of Zhandry [Zha12b].

One-time quantum secure MACs. A one-time MAC is existentially unforgeable when the adversary can make only a *single* chosen message query. Classically, one-time MACs are constructed from pair-wise independent hash functions [WC81]. These MACs are one-time secure since the value of a pair-wise independent hash at one point gives no information about its value at another point. Therefore, a single classical chosen-message query tells the adversary nothing about the MAC tag of another message.

In the quantum settings things are more complicated. Unlike the classical settings, we show that pair-wise independence does not imply existential unforgeability under a one-time quantum chosen message attack. For example, consider the simple pair-wise independent hash family $\mathcal{H} = \{h(x) = ax + b\}_{a,b \in \mathbb{F}_p}$ with domain and range \mathbb{F}_p . We show that a quantum adversary presented with an oracle for a random function $h \in \mathcal{H}$ can find both a and b with a *single* quantum query to h . Consequently, the classical one-time MAC constructed from

\mathcal{H} is completely insecure in the quantum settings. More generally we prove the following theorem:

Theorem 2 (informal). *There is a polynomial time quantum algorithm that when presented with an oracle for $h(x) = a_0 + a_1x + \dots + a_dx^d$ for random a_0, \dots, a_d in \mathbb{F}_p can recover a_0, \dots, a_d using only d quantum queries to the oracle with probability $1 - O(d/n)$.*

The $h(x) = ax + b$ attack discussed above is a special case of this theorem with $d = 1$. With classical queries finding a_0, \dots, a_d requires $d + 1$ queries, but with quantum queries the theorem shows that d queries are sufficient.

Theorem 2 is a quantum polynomial interpolation algorithm: given oracle access to the polynomial, the algorithm reconstructs its coefficients. This problem was studied previously by Kane and Kutin [KK11] who prove that $d/2$ quantum queries are insufficient to interpolate the polynomial. Interestingly, they conjecture that quantum interpolation requires $d + 1$ quantum queries as in the classical case, but Theorem 2 refutes that conjecture. Theorem 2 also applies to a quantum version of secret sharing where the shares themselves are superpositions. It shows that the classical Shamir secret sharing scheme [Sha79] is insecure if the shares are allowed to be quantum states obtained by evaluating the secret sharing polynomial on quantum superpositions. More generally, the security of secret sharing schemes in the quantum settings was analyzed by Damgård et al. [DFNS11].

As for one-time secure MACs, while pair-wise independence is insufficient for quantum one-time security, we show that four-wise independence is sufficient. That is, a four-way independent hash family gives rise to an existentially unforgeable MAC under a one-time quantum chosen message attack. It is still an open problem whether three-way independence is sufficient. More generally, we show that $(q + 1)$ -way independence is insufficient for a q -time quantum secure MAC, but $(3q + 1)$ -way independence is sufficient.

Motivation. Allowing the adversary to issue quantum chosen message queries is a natural and conservative security model and is therefore an interesting one to study. Showing that classical MAC constructions remain secure in this model gives confidence in case end-user computing devices eventually become quantum. Nevertheless, one might imagine that even in a future where computers are quantum, the last step in a MAC signing procedure is to sample the resulting quantum state so that the generated MAC is always classical. The quantum chosen message query model ensures that even if the attacker can bypass this last “classicalization” step, the MAC remains secure.

As further motivation we note that the results in this paper are the tip of a large emerging area of research with many open questions. Consider for example signature schemes. Can one design schemes that remain secure when the adversary can issue quantum chosen message queries? Similarly, can one design encryption systems that remain secure when the the adversary can issue quantum chosen ciphertext queries? More generally, for any cryptographic primitive modeled as an interactive game, one can ask how to design primitives that remain secure when the interaction between the adversary and its given oracles is quantum.

Other related work. Several recent works study the security of cryptographic primitives when the adversary can issue quantum queries [BDF⁺11,Zha12a,Zha12b]. So far these have focused on proving security of signatures, encryption, and identity-based encryption in the *quantum* random oracle model where the adversary can query the random oracle on superpositions of inputs. These works show that many, but not all, random oracle constructions remain secure in the quantum random oracle model. The quantum random oracle model has also been used to prove security of Merkle’s Puzzles in the quantum settings [BS08,BHK⁺11]. Meanwhile, Damård et al. [DFNS11] examine secret sharing and multiparty computation in a model where an adversary may corrupt a superposition of subsets of players, and build zero knowledge protocols that are secure, even when a dishonest verifier can issue challenges on superpositions.

Some progress toward identifying sufficient conditions under which classical protocols are also quantum immune has been made by Unruh [Unr10] and Hallgren et al. [HSS11]. Unruh shows that any scheme that is statistically secure in Cannetti’s universal composition (UC) framework [Can01] against classical adversaries is also statistically secure against quantum adversaries. Hallgren et al. show that for many schemes this is also true in the computational setting. These results, however, do not apply to MACs.

2 Preliminaries: Definitions and Notation

Let $[n]$ be the set $\{1, \dots, n\}$. For a prime power n , let \mathbb{F}_n be the finite field on n elements. For any positive integer n , let \mathbb{Z}_n be the ring of integers modulo n .

Functions will be denoted by capitol letters (such as F), and sets by capitol script letters (such as \mathcal{X}). We denote vectors with bold lower-case letters (such as \mathbf{v}), and the components of a vector $\mathbf{v} \in \mathcal{A}^n$ by $v_i, i \in [n]$. We denote matrices with bold capital letters (such as \mathbf{M}), and the components of a matrix $\mathbf{M} \in \mathcal{A}^{m \times n}$ by $M_{i,j}, i \in [m], j \in [n]$. Given a function $F : \mathcal{X} \rightarrow \mathcal{Y}$ and a vector $\mathbf{v} \in \mathcal{X}^n$, let $F(\mathbf{v})$ denote the vector $(F(v_1), F(v_2), \dots, F(v_k))$. Let $F([n])$ denote the vector $(F(1), F(2), \dots, F(n))$.

Given a vector space \mathcal{V} , let $\dim \mathcal{V}$ be the dimension of \mathcal{V} , or the number of vectors in any basis for \mathcal{V} . Given a set of vectors $\{\mathbf{v}_1, \dots, \mathbf{v}_k\}$, let $\text{span}\{\mathbf{v}_1, \dots, \mathbf{v}_k\}$ denote the space of all linear combinations of vectors in $\{\mathbf{v}_1, \dots, \mathbf{v}_k\}$. Given a subspace S of an inner-product space V , and a vector $\mathbf{v} \in V$, define $\text{proj}_S \mathbf{v}$ as the orthogonal projection of \mathbf{v} onto S , that is, the vector $\mathbf{w} \in S$ such that $|\mathbf{v} - \mathbf{w}|$ is minimized.

Given a matrix \mathbf{M} , we define the rank, denoted $\text{rank}(\mathbf{M})$, to be the size of the largest subset of rows (equivalently, columns) of \mathbf{M} that are linearly independent.

Given a function $F : \mathcal{X} \rightarrow \mathcal{Y}$ and a subset $\mathcal{S} \subseteq \mathcal{X}$, the restriction of F to \mathcal{S} is the function $F_{\mathcal{S}} : \mathcal{S} \rightarrow \mathcal{Y}$ where $F_{\mathcal{S}}(x) = F(x)$ for all $x \in \mathcal{S}$. A distribution D on the set of functions F from \mathcal{X} to \mathcal{Y} induces a distribution $D_{\mathcal{S}}$ on the set of functions from \mathcal{S} to \mathcal{Y} , where we sample from $D_{\mathcal{S}}$ by first sampling a function F from D , and outputting $F_{\mathcal{S}}$. We say that D is k -wise independent if, for each set \mathcal{S} of size at most k , each of the distributions $D_{\mathcal{S}}$ are truly random distributions

on functions from \mathcal{S} to \mathcal{Y} . A set \mathcal{F} of functions from \mathcal{X} to \mathcal{Y} is k -wise independent if the uniform distribution on \mathcal{F} is k -wise independent.

2.1 Quantum Computation.

The quantum system A is a complex Hilbert space \mathcal{H} with inner product $\langle \cdot | \cdot \rangle$. The state of a quantum system is given by a vector $|\psi\rangle$ of unit norm ($\langle \psi | \psi \rangle = 1$). Given quantum systems \mathcal{H}_1 and \mathcal{H}_2 , the joint quantum system is given by the tensor product $\mathcal{H}_1 \otimes \mathcal{H}_2$. Given $|\psi_1\rangle \in \mathcal{H}_1$ and $|\psi_2\rangle \in \mathcal{H}_2$, the product state is given by $|\psi_1\rangle|\psi_2\rangle \in \mathcal{H}_1 \otimes \mathcal{H}_2$. Given a quantum state $|\psi\rangle$ and an orthonormal basis $B = \{|b_0\rangle, \dots, |b_{d-1}\rangle\}$ for \mathcal{H} , a measurement of $|\psi\rangle$ in the basis B results in a value b_i with probability $|\langle b_i | \psi \rangle|^2$, and the state $|\psi\rangle$ is collapsed to the state $|b_i\rangle$. We let $b_i \leftarrow |\psi\rangle$ denote the distribution on b_i obtained by sampling $|\psi\rangle$.

A unitary transformation over a d -dimensional Hilbert space \mathcal{H} is a $d \times d$ matrix \mathbf{U} such that $\mathbf{U}\mathbf{U}^\dagger = \mathbf{I}_d$, where \mathbf{U}^\dagger represents the conjugate transpose. A quantum algorithm operates on a product space $\mathcal{H}_{in} \otimes \mathcal{H}_{out} \otimes \mathcal{H}_{work}$ and consists of n unitary transformations $\mathbf{U}_1, \dots, \mathbf{U}_n$ in this space. \mathcal{H}_{in} represents the input to the algorithm, \mathcal{H}_{out} the output, and \mathcal{H}_{work} the work space. A classical input x to the quantum algorithm is converted to the quantum state $|x, 0, 0\rangle$. Then, the unitary transformations are applied one-by-one, resulting in the final state

$$|\psi_x\rangle = \mathbf{U}_n \dots \mathbf{U}_1 |x, 0, 0\rangle .$$

The final state is measured, obtaining (a, b, c) with probability $|\langle a, b, c | \psi_x \rangle|^2$. The output of the algorithm is b .

Quantum-accessible Oracles. We will implement an oracle $O : \mathcal{X} \rightarrow \mathcal{Y}$ by a unitary transformation \mathbf{O} where

$$\mathbf{O}|x, y, z\rangle = |x, y + O(x), z\rangle$$

where $+$: $\mathcal{X} \times \mathcal{X} \rightarrow \mathcal{X}$ is some group operation on \mathcal{X} . Suppose we have a quantum algorithm that makes quantum queries to oracles O_1, \dots, O_q . Let $|\psi_0\rangle$ be the state of the algorithm before any queries, and let $\mathbf{U}_1, \dots, \mathbf{U}_q$ be the unitary transformations applied between queries. The final state of the algorithm will be

$$\mathbf{U}_q \mathbf{O}_q \dots \mathbf{U}_1 \mathbf{O}_1 |\psi_0\rangle$$

We can also have an algorithm make classical queries to O_i . In this case, the input to the oracle is measured before applying the transformation \mathbf{O}_i .

Fix an oracle $O : \mathcal{X} \rightarrow \mathcal{Y}$. Let $O^{(q)} : \mathcal{X}^q \rightarrow \mathcal{Y}^q$ be the oracle that maps \mathbf{x} into $O(\mathbf{x}) = (O(x_1), O(x_2), \dots, O(x_q))$. Observe that any quantum query to $O^{(q)}$ can be implemented using q quantum queries to O , where the unitary transformations between queries just permute the registers. We say that an algorithm that makes a single query to $O^{(q)}$ makes q non-adaptive queries to O .

2.2 Quantum secure MACs

A MAC system comprises two algorithms: a (possibly) randomized MAC signing algorithm $S(k, m)$ and a MAC verification algorithm $V(k, m, t)$. Here k denotes the secret key chosen at random from the key space, m denotes a message in the message space, and t denotes the MAC tag in the tag space on the message m . These algorithms and spaces are parameterized by a security parameter λ .

Classically, a MAC system is said to be secure if no attacker can win the following game: a random key k is chosen from the key space and the attacker is presented with a signing oracle $S(k, \cdot)$. Queries to the signing oracle are called chosen message queries. Let $\{(m_i, t_i)\}_{i=1}^q$ be the set of message-tag pairs that the attacker obtains by interacting with the signing oracle. The attacker wins the game if it can produce an existential forgery, namely a valid message-tag pair (m^*, t^*) satisfying $(m^*, t^*) \notin \{(m_i, t_i)\}_{i=1}^q$. The MAC system is said to be secure if no “efficient” adversary can win this game with non-negligible probability in λ .

Quantum chosen message queries. In the quantum settings we allow the adversary to maintain its own quantum state and issue quantum queries to the signing oracle. Let $\sum_{m,x,y} \psi_{m,x,y} |m, x, y\rangle$ be the adversary’s state just prior to issuing a signing query. The MAC signing oracle transforms this state as follows:

1. it chooses a random string r that will be used by the MAC signing algorithm,
2. it signs each “slot” in the given superposition by running $S(k, m; r)$, that is running algorithm S with randomness r . More precisely, the signing oracle performs the following transformation:

$$\sum_{m,x,y} \psi_{m,x,y} |m, x, y\rangle \longrightarrow \sum_{m,x,y} \psi_{m,x,y} |m, x \oplus S(k, m; r), y\rangle$$

When the signing algorithm is deterministic there is no need to choose an r . However, for randomized signing algorithms the same randomness is used to compute the tag for all slots in the superposition. Alternatively, we could have required fresh randomness in every slot, but this would make it harder to implement the MAC system on a quantum device. Allowing the same randomness in every slot is more conservative and frees the signer from this concern. At any rate, the two models are very close — if need be, the random string r can be used as a key for a quantum-secure PRF [Zha12b] which is used to generate a fresh pseudorandom value for every slot.

Existential forgery. After issuing q quantum chosen message queries the adversary wins the game if it can generate $q + 1$ valid classical message-tag pairs.

Definition 1. *A MAC system is existentially unforgeable under a quantum chosen message attack (EUF-qCMA) if no adversary can win the quantum MAC game with non-negligible advantage in λ .*

Zhandry [Zha12b] gives an example of a classically secure PRF that is insecure under quantum queries. This PRF gives an example MAC that is classically secure, but insecure under quantum queries. Our goal for the remainder of the paper is to construct EUF-qCMA secure MACs.

3 The Rank Method

In this section we introduce the rank method which is a general approach to proving lower bounds on quantum algorithms. The setup is as follows: we give a quantum algorithm A access to some quantity $H \in \mathcal{H}$. By access, we mean that the final state of the algorithm is some fixed function of H . In this paper, \mathcal{H} will be a set of functions, and A will be given oracle access to $H \in \mathcal{H}$ by allowing A to make q quantum oracle queries to H , for some q . For now, we will treat \mathcal{H} abstractly, and return to the specific case where \mathcal{H} is a set of functions later.

The idea behind the rank method is that, if we treat the final states of the algorithm on different H as vectors, the space spanned by these vectors will be some subspace of the overall Hilbert space. If the dimension of this subspace is small enough, the subspace (and hence all of the vectors in it) must be reasonably far from most of the vectors in the measurement basis. This allows us to bound the ability of such an algorithm to achieve some goal.

For $H \in \mathcal{H}$, let $|\psi_H\rangle$ be the final state of the quantum algorithm A , before measurement, when given access to H . Suppose the different $|\psi_H\rangle$ vectors all lie in a space of dimension d . Let $\Psi_{A,\mathcal{H}}$ be the $|\mathcal{H}| \times d$ matrix whose rows are the various vectors $|\psi_H\rangle$.

Definition 2. For a quantum algorithm A given access to some value $H \in \mathcal{H}$, we define the rank, denoted $\text{rank}(A, \mathcal{H})$, as the rank of the matrix $\Psi_{A,\mathcal{H}}$.

The rank of an algorithm A seemingly contains very little information: it gives the dimension of the subspace spanned by the $|\psi_H\rangle$ vectors, but gives no indication of the orientation of this subspace or the positions of the $|\psi_H\rangle$ vectors in the subspace. Nonetheless, we demonstrate how the success probability of an algorithm can be bounded from above knowing only the rank of $\Psi_{A,\mathcal{H}}$.

Theorem 3. Let A be a quantum algorithm that has access to some value $H \in \mathcal{H}$ drawn from some distribution D and produces some output $w \in \mathcal{W}$. Let $R : \mathcal{H} \times \mathcal{W} \rightarrow \{\text{True}, \text{False}\}$ be a binary relation. Then the probability that A outputs some w such that $R(H, w) = \text{True}$ is at most

$$\left(\max_{w \in \mathcal{W}} \Pr_{H \leftarrow D} [R(H, w)] \right) \times \text{rank}(A, \mathcal{H}) .$$

In other words, the probability that A succeeds in producing $w \in \mathcal{W}$ for which $R(H, w)$ is true is at most $\text{rank}(A, \mathcal{H})$ times the best probability of success of any algorithm that ignores H and just outputs some fixed w .

Proof. The probability that A outputs a w such that $R(H, w) = \text{True}$ is

$$\Pr_{\substack{H \leftarrow D \\ w \leftarrow |\psi_H\rangle}} [R(H, w)] = \sum_H \Pr_D[H] \sum_{w:R(H,w)} |\langle w | \psi_H \rangle|^2 = \sum_w \sum_{H:R(H,w)} \Pr_D[H] |\langle w | \psi_H \rangle|^2$$

Now, $|\langle w | \psi_H \rangle|$ is just the magnitude of the projection of $|w\rangle$ onto the space spanned by the vector $|\psi_H\rangle$, that is, $\text{proj}_{\text{span}|\psi_H\rangle}(|w\rangle)$. This is at most the

magnitude of the projection of $|w\rangle$ onto the space spanned by all of the $|\psi_{H'}\rangle$ for $H' \in \mathcal{H}$, or $\text{proj}_{\text{span}\{|\psi_{H'}\rangle\}}(|w\rangle)$. Thus,

$$\Pr_{\substack{H \leftarrow D \\ w \leftarrow |\psi_H\rangle}} [R(z, w)] \leq \sum_w \left(\sum_{H:R(H,w)} \Pr_D[H] \right) \left| \text{proj}_{\text{span}\{|\psi_{H'}\rangle\}}(|w\rangle) \right|^2$$

Now, we can perform the sum over H , which gives $\Pr_{H \leftarrow D}[R(H, w)]$. We can bound this by the maximum it attains over all w , giving us

$$\Pr_{\substack{H \leftarrow D \\ w \leftarrow |\psi_H\rangle}} [R(H, w)] \leq \left(\max_w \Pr_{H \leftarrow D}[R(H, w)] \right) \sum_w \left| \text{proj}_{\text{span}\{|\psi_{H'}\rangle\}}(|w\rangle) \right|^2$$

Now, let $|b_i\rangle$ be an orthonormal basis for $\text{span}\{|\psi_{H'}\rangle\}$. Then

$$\left| \text{proj}_{\text{span}\{|\psi_{H'}\rangle\}}(|w\rangle) \right|^2 = \sum_i |\langle b_i | w \rangle|^2$$

Summing over all w gives

$$\sum_w \left| \text{proj}_{\text{span}\{|\psi_{H'}\rangle\}}(|w\rangle) \right|^2 = \sum_w \sum_i |\langle b_i | w \rangle|^2 = \sum_i \sum_w |\langle b_i | w \rangle|^2$$

Since the w are the possible results of measurement, the vectors $|w\rangle$ form an orthonormal basis for the whole space, meaning $\sum_w |\langle b_i | w \rangle|^2 = \| |b_i\rangle \|^2 = 1$. Hence, the sum just becomes the number of $|b_i\rangle$, which is just the dimension of the space spanned by the $|\psi_{H'}\rangle$. Thus,

$$\Pr_{\substack{H \leftarrow D \\ w \leftarrow |\psi_H\rangle}} [R(H, w)] \leq \left(\max_{w \in \mathcal{W}} \Pr_{H \leftarrow D}[R(H, w)] \right) (\dim \text{span}\{|\psi_{H'}\rangle\}) .$$

But $\dim \text{span}\{|\psi_{H'}\rangle\}$ is exactly $\text{rank}(\Psi_{A, \mathcal{H}}) = \text{rank}(A, \mathcal{H})$, which finishes the proof of the theorem. \square

We now move to the specific case of oracle access. \mathcal{H} is now some set of functions from \mathcal{X} to \mathcal{Y} , and our algorithm A makes q quantum oracle queries to a function $H \in \mathcal{H}$. To use the rank method (Theorem 3) for our purposes, we need to bound the rank of such an algorithm. First, we define the following quantity:

$$C_{k,q,n} \equiv \sum_{r=0}^q \binom{k}{r} (n-1)^r .$$

Theorem 4. *Let \mathcal{X} and \mathcal{Y} be sets of size m and n and let H_0 be some function from \mathcal{X} to \mathcal{Y} . Let \mathcal{S} be a subset of \mathcal{X} of size k and let \mathcal{H} be some set of functions from \mathcal{X} to \mathcal{Y} that are equal to H_0 except possibly on points in \mathcal{S} . If A is a quantum algorithm making q queries to an oracle drawn from \mathcal{H} , then*

$$\text{rank}(A, \mathcal{H}) \leq C_{k,q,n} .$$

Proof. The proof appears in the full version [BZ13]. Here we sketch our approach. Let $|\psi_H^q\rangle$ be the final state of a quantum algorithm after q quantum oracle calls to an oracle $H \in \mathcal{H}$. We wish to bound the dimension of the space spanned by the vectors $|\psi_H^q\rangle$ for all $H \in \mathcal{H}$. We accomplish this by exhibiting a basis for this space. Our basis consists of $|\psi_{H'}^q\rangle$ vectors where H' only differs from H_0 at a maximum of q points in \mathcal{S} . A simple counting argument shows that there are exactly $C_{k,q,n}$ such H' oracles. We show in the full version that these $|\psi_{H'}^q\rangle$ vectors do indeed span the entire space. \square

3.1 An Example

Suppose our task is to, given one quantum query to an oracle $H : \mathcal{X} \rightarrow \mathcal{Y}$, produce two distinct pairs (x_0, y_0) and (x_1, y_1) such that $H(x_0) = y_0$ and $H(x_1) = y_1$. Suppose further that H is drawn from a pairwise independent set \mathcal{H} . We will now see that the rank method leads to a bound on the success probability of any quantum algorithm A .

Corollary 1. *No quantum algorithm A , making a single query to a function $H : \mathcal{X} \rightarrow \mathcal{Y}$ drawn from a pairwise independent set \mathcal{H} , can produce two distinct input/output pairs of H , except with probability at most $|\mathcal{X}|/|\mathcal{Y}|$.*

Proof. Let $m = |\mathcal{X}|$ and $n = |\mathcal{Y}|$. Since no outputs of H are fixed, we will set $\mathcal{S} = \mathcal{X}$ in Theorem 4, showing that the rank of the algorithm A is bounded by $C_{m,1,n} = 1 + m(n-1) < mn$. If an algorithm makes no queries to H , the best it can do at outputting two distinct input/output pairs is to just pick two arbitrary distinct pairs, and output those. The probability that this zero-query algorithm succeeds is at most $1/n^2$. Then Theorem 3 tells us that A succeeds with probability at most $\text{rank}(A, \mathcal{H})$ times this amount, which equates to $\frac{m}{n}$. \square

For $m > n$, this bound is trivial. However, for m smaller than n , this gives a non-trivial bound, and for m exponentially smaller than n , this bound is negligible.

4 Outputting Values of a Random Oracle

In this section, we will prove Theorem 1. We consider the following problem: given q quantum queries to a random oracle $H : \mathcal{X} \rightarrow \mathcal{Y}$, produce $k > q$ distinct pairs (x_i, y_i) such that $y_i = H(x_i)$. Let $n = |\mathcal{Y}|$ be the size of the range. Motivated by our application to quantum-accessible MACs, we are interested in the case where the range \mathcal{Y} of the oracle is large, and we want to show that to produce even one extra input/output pair ($k = q + 1$) is impossible, except with negligible probability. We are also interested in the case where the range of the oracle, though large, is far smaller than the domain. Thus, the bound we obtained in the previous section (Corollary 1) is not sufficient for our purposes, since it is only non-trivial if the range is larger than the domain.

In the classical setting, when $k \leq q$, this problem is easy, since we can just pick an arbitrary set of k different x_i values, and query the oracle on each value. For $k > q$, no adversary of even unbounded complexity can solve this problem, except with probability $1/n^{k-q}$, since for any set of k inputs, at least $k - q$ of the corresponding outputs are completely unknown to the adversary. Therefore, for large n , we have a sharp threshold: for $k \leq q$, this problem can be solved efficiently with probability 1, and for even $k = q + 1$, this problem cannot be solved, even inefficiently, except with negligible probability.

In the quantum setting, the $k \leq q$ case is the same as before, since we can still query the oracle classically. However, for $k > q$, the quantum setting is more challenging. The adversary can potentially query the random oracle on a superposition of all inputs, so he “sees” the output of the oracle on all points. Proving that it is still impossible to produce k input/output pairs is thus more complicated, and existing methods fail to prove that this problem is difficult. Therefore, it is not immediately clear that we have the same sharp threshold as before.

In Section 4.1 we use the rank method to bound the probability that any (even computationally unbounded) quantum adversary succeeds. Then in Section 4.2 we show that our bound is tight by giving an efficient algorithm for this problem that achieves the lower bound. In particular, for an oracle $H : \mathcal{X} \rightarrow \mathcal{Y}$ we consider two cases:

- Exponentially-large range \mathcal{Y} and polynomial k, q . In this case, we will see that the success probability even when $k = q + 1$ is negligible. That is, to produce even one additional input/output pair is hard. Thus, we get the same sharp threshold as in the classical case
- Constant size range \mathcal{Y} and polynomial k, q . We show that even when q is a constant fraction of k we can still produce k input/output pairs with overwhelming probability using only q quantum queries. This is in contrast to the classical case, where the success probability for $q = ck$, $c < 1$, is negligible in k .

4.1 A Tight Upper Bound

Theorem 5. *Let A be a quantum algorithm making q queries to a random oracle $H : \mathcal{X} \rightarrow \mathcal{Y}$ whose range has size n , and produces $k > q$ pairs $(x_i, y_i) \in \mathcal{X} \times \mathcal{Y}$. The probability that the x_i values are distinct and $y_i = H(x_i)$ for all $i \in [k]$ is at most $\frac{1}{n^k} C_{k,q,n}$.*

Proof. The complete proof is given in the full version [BZ13]. Here we prove the special case where k is equal to the size of the domain \mathcal{X} . In this case, any quantum algorithm that outputs k distinct input/output pairs must output *all* input/output pairs. Similar to the proof of Corollary 1, we will set $\mathcal{S} = \mathcal{X}$, and use Theorem 4 to bound the rank of A at $C_{k,q,n}$. Now, any algorithm making *zero* queries succeeds with probability at most $1/n^k$. Theorem 3 then bounds the success probability of any q query algorithm as $C_{k,q,n}/n^k$. \square

For this paper, we are interested in the case where $n = |\mathcal{Y}|$ is exponentially large, and we are only allowed a polynomial number of queries. Suppose $k = q + 1$, the easiest non-trivial case for the adversary. Then, the probability of success is

$$\frac{1}{n^{q+1}} \sum_{r=0}^q \binom{q+1}{r} (n-1)^r = 1 - \left(1 - \frac{1}{n}\right)^{q+1} \leq \frac{q+1}{n}. \quad (4.1)$$

Therefore, to produce even one extra input/output pair is impossible, except with exponentially small probability, just like in the classical case. This proves the first part of Theorem 1.

4.2 An Optimal Attack

In this section, we present a quantum algorithm for the problem of computing $H(x_i)$ for k different x_i values, given only $q < k$ queries:

Theorem 6. *Let \mathcal{X} and \mathcal{Y} be sets, and fix integers $q < k$, and k distinct values $x_1, \dots, x_k \in \mathcal{X}$. There exists a quantum algorithm A that makes q non-adaptive quantum queries to any function $H : \mathcal{X} \rightarrow \mathcal{Y}$, and produces $H(x_1), \dots, H(x_k)$ with probability $C_{k,q,n}/n^k$, where $n = |\mathcal{Y}|$.*

The proof appears in the full version [BZ13], and is similar to the algorithm of [vD98], though generalized to handle arbitrary range sizes. This algorithm has the same success probability as in Theorem 5, showing that both our attack and lower bound of Theorem 5 are optimal. This proves the second part of Theorem 1.

As we have already seen, for exponentially-large \mathcal{Y} , this attack has negligible advantage for any $k > q$. However, if $n = |\mathcal{Y}|$ is constant, we can do better. The error probability is

$$\sum_{r=q+1}^k \binom{k}{r} \left(1 - \frac{1}{n}\right)^r \left(\frac{1}{n}\right)^{k-r} = \sum_{s=0}^{k-q-1} \binom{k}{s} \left(\frac{1}{n}\right)^s \left(1 - \frac{1}{n}\right)^{k-s}.$$

This is the probability that k consecutive coin flips, where each coin is heads with probability $1/n$, yields fewer than $k - q$ heads. Using the Chernoff bound, if $q > k(1 - 1/n)$, this probability is at most

$$e^{-\frac{n}{2k}(q-k(1-1/n))^2}.$$

For a constant n , let c be any constant with $1 - 1/n < c < 1$. If we use $q = ck$ queries, the error probability is less than

$$e^{-\frac{n}{2k}(k(c+1/n-1))^2} = e^{-\frac{nk}{2}(c+1/n-1)^2},$$

which is exponentially small in k . Thus, for constant n , and any constant c with $1 - 1/n < c < 1$, using $q = ck$ quantum queries, we can determine k input/output pairs with overwhelming probability. This is in contrast to the classical case, where with any constant fraction of k queries, we can only produce k input/output pairs with negligible probability. As an example, if H outputs two bits, it is possible to produce k input/output pairs of H using only $q = 0.8k$ quantum queries. However, with $0.8k$ classical queries, we can output k input/output pairs with probability at most $4^{-0.2k} < 0.76^k$.

5 Quantum-Accessible MACs

Using Theorem 5 we can now show that a quantum secure pseudorandom function [Zha12b] gives rise to the quantum-secure MAC, namely $S(k, m) = \text{PRF}(k, m)$. We prove that this mac is secure.

Theorem 7. *If $\text{PRF} : \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$ is a quantum-secure pseudorandom function and $1/|\mathcal{Y}|$ is negligible, then $S(k, m) = \text{PRF}(k, m)$ is a EUF-qCMA-secure MAC.*

Proof. Let A be a polynomial time adversary that makes q quantum queries to $S(k, \cdot)$ and produces $q + 1$ valid input/output pairs with probability ϵ . Let Game 0 be the standard quantum MAC attack game, where A makes q quantum queries to MAC_k . By definition, A 's success probability in this game is ϵ .

Let Game 1 be the same as Game 0, except that $S(k, \cdot)$ is replaced with a truly random function $O : \mathcal{X} \rightarrow \mathcal{Y}$, and define A 's success probability as the probability that A outputs $q + 1$ input/output pairs of O . Since PRF is a quantum-secure PRF, A 's advantage in distinguishing Game 0 from Game 1 is negligible.

Now, in Game 1, A makes q quantum queries to a random oracle, and tries to produce $q + 1$ input/output pairs. However, by Theorem 5 and Eq. (4.1) we know that A 's success probability is bounded by $(q + 1)/|\mathcal{Y}|$ which is negligible. It now follows that ϵ is negligible and therefore, S is a EUF-qCMA-secure MAC. \square

5.1 Carter-Wegman MACs

In this section, we show how to modify the Carter-Wegman MAC so that it is secure in the quantum setting presented in Section 2.2. Recall that H is an XOR-universal family of hash functions from \mathcal{X} into \mathcal{Y} if for any two distinct points x and y , and any constant $c \in \mathcal{Y}$,

$$\Pr_{h \leftarrow H} [H(x) - H(y) = c] = 1/|\mathcal{Y}|$$

The Carter-Wegman construction uses a pseudorandom function family PRF with domain \mathcal{X} and range \mathcal{Y} , and an XOR-universal family of hash functions \mathcal{H} from \mathcal{M} to \mathcal{Y} . The key is a pair (k, H) , where k is a key for PRF and H is a function drawn from \mathcal{H} . To sign a message, pick a random $r \in \mathcal{X}$, and return $(r, \text{PRF}(k, r) + H(m))$.

This MAC is not, in general, secure in the quantum setting presented in Section 2.2. The reason is that the same randomness is used in all slots of a quantum chosen message query, that is the signing oracle computes:

$$\sum_m \alpha_m |m\rangle \longrightarrow \sum_m \alpha_m |m, r, \text{PRF}(k, r) + H(m)\rangle$$

where the same r is used for all classical states of the superposition. For example, suppose \mathcal{H} is the set of functions $H(x) = ax + b$ for random a and b . With even a single quantum query, the adversary will be able to obtain a and $\text{PRF}(k, r) + b$

with high probability, using the algorithm from Theorem 10 in Section 6. Knowing both of these will allow the adversary to forge any message.

We show how to modify the standard Carter-Wegman MAC to make it secure in the quantum setting.

Construction 1 (Carter-Wegman) *The Quantum Carter-Wegman MAC (QCW-MAC) is built from a pseudorandom function PRF, an XOR-universal set of functions \mathcal{H} , and a pairwise independent set of functions \mathcal{R} .*

Keys: The secret key for QCW-MAC is a pair (k, H) , where k is a key for PRF and $H : \mathcal{M} \rightarrow \mathcal{Y}$ is drawn from \mathcal{H}

Signing: To sign a message m choose a random $R \in \mathcal{R}$ and output the pair $(R(m), \text{PRF}(k, R(m)) + H(m))$ as the tag. When responding to a quantum chosen message query, the same R is used in all classical states of the superposition.

Verification: To verify that (r, s) is a valid tag for m , accept iff $\text{PRF}(k, r) + H(m) = s$.

Theorem 8. *The Quantum Carter-Wegman MAC is a EUF-qCMA secure MAC.*

The proof is given in the full version [BZ13].

6 q-time MACs

In this section, we develop quantum one-time MACs, MACs that are secure when the adversary can issue only *one* quantum chosen message query. More generally, we will study quantum q -time MACs.

Classically, any pairwise independent function is a one-time MAC. In the quantum setting, Corollary 1 shows that when the range is much larger than the domain, this still holds. However, such MACs are not useful since we want the tag to be short. We first show that when the range is not larger than the domain, pairwise independence is not enough to ensure security:

Theorem 9. *For any set \mathcal{Y} of prime-power size, and any set \mathcal{X} with $|\mathcal{X}| \geq |\mathcal{Y}|$, there exist $(q + 1)$ -wise independent functions from \mathcal{X} to \mathcal{Y} that are not q -time MACs.*

To prove this theorem, we treat \mathcal{Y} as a finite field, and assume $\mathcal{X} = \mathcal{Y}$, as our results are easy to generalize to larger domains. We use random degree q polynomials as our $(q + 1)$ -wise independent family, and show in Theorem 10 below that such polynomials can be completely recovered using only q quantum queries. It follows that the derived MAC cannot be q -time secure since once the adversary has the polynomial it can easily forge tags on new messages. The proof of the following theorem appears in the full version [BZ13]:

Theorem 10. *For any prime power n , there is an efficient quantum algorithm that makes only q quantum queries to an oracle implementing a degree- q polynomial $F : \mathbb{F}_n \rightarrow \mathbb{F}_n$, and completely determines F with probability $1 - O(qn^{-1})$.*

The theorem shows that a $(q + 1)$ -wise independence family is not necessarily a secure quantum q -time MAC since after q quantum chosen message queries the adversary extracts the entire secret key. The case $q = 1$ is particularly interesting.

6.1 Sufficient Conditions for a One-Time Mac

We show that, while pairwise independence is not enough for a one-time MAC, 4-wise independence is. We first generalize a theorem of Zhandry [Zha12a]:

Lemma 1. *Let A be any quantum algorithm that makes c classical queries and q quantum queries to an oracle H . If H is drawn from a $(c + 2q)$ -wise independent function, then the output distribution of A is identical to the case where H is truly random.*

Proof. The complete proof is given in the full version [BZ13]. If $q = 0$, then this theorem is trivial, since the c classical outputs A sees are distributed randomly. If $c = 0$, then the theorem reduces to that of Zhandry [Zha12a]. By adapting the proof of the $c = 0$ case to the general case, we get the lemma. \square

Using this lemma we show that $(3q + 1)$ -wise independence is sufficient for q -time MACs.

Theorem 11. *Any $(3q + 1)$ -wise independent family with domain \mathcal{X} and range \mathcal{Y} is a quantum q -time secure MAC provided $(q + 1)/|\mathcal{Y}|$ is negligible.*

Proof. Let D be some $(3q + 1)$ -wise independent function. Suppose we have an adversary A that makes q quantum queries to an oracle H , and attempts to produce $q + 1$ input/output pairs. Let ϵ_R be the probability of success when H is a random oracle, and let ϵ_D be the probability of success when H is drawn from D . We construct an algorithm B with access to H as follows: simulate A with oracle access to H . When A outputs $q + 1$ input/output pairs, simply make $q + 1$ queries to H to check that these are valid pairs. Output 1 if and only if all pairs are valid. Therefore, B makes q quantum queries and $c = q + 1$ classical queries to H , and outputs 1 if and only if A succeeds: if H is random, B outputs 1 with probability ϵ_R , and if H is drawn from D , B outputs 1 with probability ϵ_D . Now, since D is $(3q + 1)$ -wise independent and $3q + 1 = 2q + c$, Lemma 1 shows that the distributions of outputs when H is drawn from D is identical to that when H is random, meaning $\epsilon_D = \epsilon_R$.

Thus, when H is drawn from D , A 's succeeds with the same probability that it would if H was random. But we already know that if H is truly random, A 's success probability is less than $(q + 1)/|\mathcal{Y}|$. Therefore, when H is drawn from D , A succeeds with probability less than $(q + 1)/|\mathcal{Y}|$, which is negligible. Hence, if H is drawn from D , H is a q -time MAC. \square

7 Conclusion

We introduced the rank method as a general technique for obtaining lower bounds on quantum oracle algorithms and used this method to bound the probability that a quantum algorithm can evaluate a random oracle $\mathcal{O} : \mathcal{X} \rightarrow \mathcal{Y}$ at k points using $q < k$ queries. When the range of \mathcal{Y} is small, say $|\mathcal{Y}| = 8$, a quantum algorithm can recover k points of \mathcal{O} from only $0.9k$ queries with high probability. However, we show that when the range \mathcal{Y} is large, no algorithm can produce k input-output pairs of \mathcal{O} using only $k - 1$ queries, with non-negligible probability. We use these bounds to construct the first MACs secure against quantum chosen message attacks. We consider both PRF and Carter-Wegman constructions. For one-time MACs we showed that pair-wise independence does not ensure security, but four-way independence does.

These results suggest many directions for future work. First, can these bounds be generalized to signatures to obtain signatures secure against quantum chosen message attacks? Similarly, can we construct encryption systems secure against quantum chosen ciphertext attacks where decryption queries are superpositions of ciphertexts?

Acknowledgments

We thank Luca Trevisan and Amit Sahai for helpful conversations about this work. This work was supported by NSF, DARPA, IARPA, the Air Force Office of Scientific Research (AFO SR) under a MURI award, Samsung, and a Google Faculty Research Award. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA, IARPA, DoI/NBC, or the U.S. Government.

References

- Aar02. Scott Aaronson. Quantum lower bound for the collision problem. In *STOC*, pages 635–642, 2002.
- Amb00. Andris Ambainis. Quantum lower bounds by quantum arguments. In *STOC*, pages 636–643, 2000.
- Amb06. Andris Ambainis. Polynomial degree vs. quantum query complexity. *J. Comput. Syst. Sci.*, 72(2):220–238, 2006.
- ASdW09. Andris Ambainis, Robert Spalek, and Ronald de Wolf. A new quantum lower bound method, with applications to direct product theorems and time-space tradeoffs. *Algorithmica*, 55(3):422–461, 2009.
- BBC⁺01. Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald de Wolf. Quantum Lower Bounds by Polynomials. *Journal of the ACM (JACM)*, 48(4):778–797, July 2001.
- BCK96. Mihir Bellare, Ran Canetti, and Hugo Krawczyk. Pseudorandom functions revisited: The cascade construction and its concrete security. In *FOCS*, pages 514–523, 1996.

- BDF⁺11. Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. Random Oracles in a Quantum World. In *Advances in Cryptology — ASIACRYPT 2011*, 2011.
- BHK⁺11. Gilles Brassard, Peter Høyer, Kassem Kalach, Marc Kaplan, Sophie Laplante, and Louis Salvail. Merkle Puzzles in a Quantum World. *Advances in Cryptology - CRYPTO 2011*, pages 391–410, 2011.
- BKR00. Mihir Bellare, Joe Kilian, and Phillip Rogaway. The security of the cipher block chaining message authentication code. *J. Comput. Syst. Sci.*, 61(3), 2000.
- BS08. Gilles Brassard and Louis Salvail. Quantum Merkle Puzzles. *Second International Conference on Quantum, Nano and Micro Technologies (ICQNM 2008)*, pages 76–79, February 2008.
- BZ13. Dan Boneh and Mark Zhandry. Quantum-secure message authentication codes. In *Proceedings of EUROCRYPT*, 2013. Full version available at the Electronic Colloquium on Computational Complexity: <http://eccc.hpi-web.de/report/2012/136>.
- Can01. Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *Proc. of FOCS*. IEEE, 2001.
- DFNS11. Ivan Damgård, Jakob Funder, Jesper Buus Nielsen, and Louis Salvail. Superposition attacks on cryptographic protocols. *CoRR*, abs/1108.6313, 2011.
- GGM86. Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to Construct Random Functions. *Journal of the ACM (JACM)*, 33(4):792–807, 1986.
- HSS11. Sean Hallgren, Adam Smith, and Fang Song. Classical cryptographic protocols in a quantum world. In *Proc. of Crypto*, LNCS. Springer, 2011.
- KdW03. Iordanis Kerenidis and Ronald de Wolf. Exponential lower bound for 2-query locally decodable codes via a quantum argument. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing (STOC)*, pages 106–115, 2003.
- KK11. Daniel M. Kane and Samuel A. Kutin. Quantum interpolation of polynomials. *Quantum Information & Computation*, 11(1&2):95–103, 2011. First published in 2009.
- Sha79. Adi Shamir. How to share a secret. *Commun. ACM*, 22(11):612–613, 1979.
- Unr10. Dominique Unruh. Universally Composable Quantum Multi-Party Computation. *Advances in Cryptology — EUROCRYPT 2010*, pages 486–505, 2010.
- vD98. Wim van Dam. Quantum oracle interrogation: Getting all information for almost half the price. In *FOCS*, pages 362–367, 1998.
- WC81. Mark N. Wegman and Larry Carter. New hash functions and their use in authentication and set equality. *J. Comput. Syst. Sci.*, 22(3):265–279, 1981.
- Zha12a. Mark Zhandry. Secure Identity-Based Encryption in the Quantum Random Oracle Model. In *Advances in Cryptology — CRYPTO*, 2012. Full version available at the Cryptology ePrint Archives: <http://eprint.iacr.org/2012/076/>.
- Zha12b. Mark Zhandry. How to Construct Quantum Random Functions. In *Proceedings of FOCS*, 2012. Full version available at the Cryptology ePrint Archives: <http://eprint.iacr.org/2012/182/>.