

Improved Quantum Multicollision-Finding Algorithm

Akinori Hosoyamada¹, Yu Sasaki¹, Seiichiro Tani², and Keita Xagawa¹

¹ NTT Secure Platform Laboratories, NTT Corporation.

3-9-11, Midori-cho, Musashino-shi, Tokyo 180-8585, Japan.

² NTT Communication Science Laboratories, NTT Corporation.

3-1, Morinosato-Wakamiya, Atsugi-shi, Kanagawa 243-0198, Japan.

{hosoyamada.akinori,sasaki.yu,tani.seiichiro,xagawa.keita}@lab.ntt.co.jp

Abstract. The current paper improves the number of queries of the previous quantum multi-collision finding algorithms presented by Hosoyamada *et al.* at Asiacrypt 2017. Let an l -collision be a tuple of l distinct inputs that result in the same output of a target function. The previous algorithm finds l -collisions by recursively calling the algorithm for finding $(l - 1)$ -collisions, and it achieves the query complexity of $O(N^{(3^{l-1}-1)/(2 \cdot 3^{l-1})})$. The new algorithm removes the redundancy of the previous recursive algorithm so that different recursive calls can share a part of computations. The new algorithm achieves the query complexity of $O(N^{(2^{l-1}-1)/(2^l-1)})$. Moreover, it finds multiclaws for random functions, which are harder to find than multicollisions.

Keywords post-quantum cryptography, quantum algorithm, multiclaw, multicollision

1 Introduction

Post-quantum cryptography has recently been discussed very actively in the cryptographic community. Quantum computers would completely break many classical public-key cryptosystems. In response, NIST is now conducting a standardization to select new public-key cryptosystems that resist attacks with quantum computers. Given this background, it is now important to investigate how quantum computers can impact on other cryptographic schemes including cryptographic hash functions.

A multicollision for a function f denotes multiple inputs to f such that they are mapped to the same output value. In particular, an l -collision denotes a tuple of l distinct inputs x_1, x_2, \dots, x_l such that $f(x_1) = f(x_2) = \dots = f(x_l)$.

A multicollision is an important object in cryptography. Lower bounds on the complexity of finding a multicollision are sometimes used to derive security bounds in provable security (e.g., security bounds for the schemes based on the sponge construction [JLM14]). In a similar context, the complexity of finding a multicollision directly impacts on the best cryptanalysis against some constructions. Furthermore, multicollisions can be used as a proof-of-work for

blockchains. In digital payment schemes, a coin must be a bit-string the validity of which can be easily checked but which is hard to produce. A micro-payment scheme, MicroMint [RS96], defines coins as 4-collisions for a function. If 4-collisions can be produced quickly, a malicious user can counterfeit coins. Some recent works prove security of schemes and protocols based on the assumption that there exist functions for which it is hard to find multicollisions [BKP18,BDRV18,KNY18].

Hosoyamada *et al.* provided a survey of multicollision finding algorithms with quantum computers [HSX17]. They first showed that an l -collision can be produced with at most $O(N^{1/2})$ queries to the target function by iteratively applying the Grover search [Gro96,BBHT98] l times. They also reported that a combination of Zhandry’s algorithm with $l = 3$ [Zha15] and Belovs’ algorithm [Bel12] achieves $O(N^{10/21})$ for $l = 3$, which is faster than the simple application of Grover’s algorithm. Finally, Hosoyamada *et al.* presented their own algorithm that recursively applies the collision finding algorithm by Brassard, Høyer, and Tapp [BHT98]. Their algorithm achieves the query complexity of $O(N^{(3^{l-1}-1)/(2 \cdot 3^{l-1})})$. For $l = 3$ and $l = 4$, the complexities are $O(N^{4/9})$ and $O(N^{13/27})$, respectively, and the algorithm works as follows.

- To search for 3-collisions, it first iterates the $O(N^{1/3})$ -query quantum algorithm for finding a 2-collision $O(N^{1/9})$ times. Then, it searches for the preimage of any one of the $O(N^{1/9})$ 2-collisions by using Grover’s algorithm, which runs with $O(N^{4/9})$ queries.
- To search for 4-collisions, it iterates the $O(N^{4/9})$ -query quantum algorithm for finding a 3-collision $O(N^{1/27})$ times. Then, it searches for the preimage of any one of the $O(N^{1/27})$ 3-collisions with $O(N^{13/27})$ queries.

As demonstrated above, the recursive algorithm by Hosoyamada *et al.* [HSX17] runs $(l - 1)$ -collision algorithm multiple times, but in each invocation, the algorithm starts from scratch. This fact motivates us to consider reusing the computations when we search for multiple $(l - 1)$ -collisions.

Our Contributions. In this paper, we improve the quantum query complexity of the previous multicollision finding algorithm by removing the redundancy of the algorithm. Consider the problem of finding an l -collision of a function $f: X \rightarrow Y$, where l is a small constant and $|Y| = N$. In addition, suppose that there exists a parameter c_N such that $c_N = o(N^{\frac{1}{2^{l-1}}})$ and $|X_i| \geq |Y|/c_N$. Then, the new algorithm achieves the quantum query complexity of $O\left(c_N^{3/2} N^{\frac{2^{l-1}-1}{2^{l-1}}}\right)$. In particular, if we can take c_N as a constant, then our algorithm can find an l -collision with $O\left(N^{\frac{2^{l-1}-1}{2^{l-1}}}\right)$ queries. The complexities for small l ’s are listed in Table 1. A comparison between them can be found in Figure 1. Our algorithm finds a 2-collision, 3-collision, 4-collision, and 5-collision of SHA3-512 with $2^{170.7}$, $2^{219.4}$, $2^{238.9}$, and $2^{247.7}$ quantum queries, respectively, up to a constant factor (Table 2).

Moreover, our new algorithm finds multiclaws for random functions, which are harder to find than multicollisions: An l -claw for functions $f_i: X_i \rightarrow Y$ for $1 \leq i \leq l$ is defined as a tuple $(x_1, \dots, x_l) \in X_1 \times \dots \times X_l$ such that $f_i(x_i) = f_j(x_j)$ for all (i, j) . If there exists a parameter c_N such that $c_N = o(N^{\frac{1}{N^{2^l-1}}})$ and $|X_i| \geq |Y|/c_N$ for each i , our quantum algorithm finds an l -claw with $O\left(c_N^{3/2} N^{\frac{2^{l-1}-1}{2^l-1}}\right)$ quantum queries. In particular, if we can take c_N as a constant, then our algorithm can find an l -claw with $O\left(N^{\frac{2^{l-1}-1}{2^l-1}}\right)$ quantum queries.

Table 1. Query complexities of l -collision finding quantum algorithms. Each fraction denotes the logarithm of the number of queries to the base N . The query complexity asymptotically approaches $1/2$ as l increases.

l	2	3	4	5	6	7	8
[HSX17] : $\frac{3^{l-1}-1}{2 \cdot 3^{l-1}}$	$\frac{1}{3}$	$\frac{4}{9}$	$\frac{13}{27}$	$\frac{40}{81}$	$\frac{121}{243}$	$\frac{364}{729}$	$\frac{1093}{2187}$
Ours : $\frac{2^{l-1}-1}{2^l-1}$	$\frac{1}{3}$	$\frac{3}{7}$	$\frac{7}{15}$	$\frac{15}{31}$	$\frac{31}{63}$	$\frac{63}{127}$	$\frac{127}{255}$

l	2	3	4	5	6	7	8
[HSX17] : $\frac{3^{l-1}-1}{2 \cdot 3^{l-1}}$	0.3333..	0.4444..	0.4814..	0.4938..	0.4979..	0.4993..	0.4997..
Ours : $\frac{2^{l-1}-1}{2^l-1}$	0.3333..	0.4285..	0.4666..	0.4838..	0.4920..	0.4960..	0.4980..

Table 2. The number of queries required to find an l -collision of SHA3-512.

l	2	3	4	5
[HSX17]	$2^{170.7}$	$2^{227.6}$	$2^{246.5}$	$2^{252.8}$
Ours	$2^{170.7}$	$2^{219.4}$	$2^{238.9}$	$2^{247.7}$




Fig. 1. Quantum query complexity for finding an l -collision. “Query” denotes the logarithm of the number of queries to the base N .

Paper Outline. The remaining of this paper is organized as follows. In section 2, we describe notations, definitions and settings. In section 3, we review previous works related to the multicollision-finding problem. In section 4, we provide our new quantum algorithm and its complexity analysis. In section 5, we conclude this paper.

Concurrent Work. Very recently, Liu and Zhandry [LZ18] showed that for any constant l , $\Theta\left(N^{\frac{1}{2}(1-\frac{1}{2^l-1})}\right)$ quantum queries are both necessary and sufficient to find a l -collision with constant probability, for a random function. That is, they gave an improved upper bound and a new lower bound on the average case. The comparisons are summarized as follows:

- Liu and Zhandry consider the l -collision case that $|X| \geq l|Y|$, where X is the domain and Y is the range. We treat the case that $|X| \geq \frac{l}{c_N}|Y|$ holds for a parameter c_N which is in $o(N^{\frac{1}{2^l-1}})$. We also consider the *multiclaw* case.
- Their exponent $\frac{1}{2}(1 - \frac{1}{2^l-1})$ is the same as ours $\frac{2^l-1-1}{2^l-1}$.
- They give the upper bound $O\left(N^{\frac{1}{2}(1-\frac{1}{2^l-1})}\right)$, while we give $O\left(c_N^{3/2} N^{\frac{1}{2}(1-\frac{1}{2^l-1})}\right)$. When c_N is a constant, our bound matches their bound.
- They give a lower bound, which matches with their and our upper bound.

We finally note that our result on an improved l -collision finding algorithm for the case $|X| \geq l|Y|$ with query complexity $O\left(N^{\frac{1}{2}(1-\frac{1}{2^l-1})}\right)$ is reported in the Rump Session of Asiacrypt 2017.

2 Preliminaries

For a positive integer M , let $[M]$ denote the set $\{1, \dots, M\}$. In this paper, N denotes a positive integer. We assume that l is a small constant. We focus on reducing quantum *query* complexities for finding multicollisions and multiclaws.

We do not consider other complexity notions such as time complexity and the size of quantum circuits. Unless otherwise noted, all sets are non-empty and finite. For sets X and Y , $\text{Func}(X, Y)$ denotes the set of functions from X to Y . For each $f \in \text{Func}(X, Y)$, we denote the set $\{f(x) \mid x \in X\}$ by $\text{Im}(f)$. For a set X , let $U(X)$ denote the uniform distribution over X . For a distribution \mathcal{D} on a set X , let $x \sim \mathcal{D}$ denote that x is a random variable that takes a value drawn from X according to \mathcal{D} . When we say that an oracle of a function $f: X \rightarrow Y$ is available, we consider the situation that each elements of X and Y are encoded into suitable binary strings, and the oracle gate $O_f: |x, z\rangle \mapsto |x, z \oplus f(x)\rangle$ is available.

An l -collision for a function $f: X \rightarrow Y$ is a tuple of elements (x_1, \dots, x_l, y) in $X^l \times Y$ such that $f(x_i) = f(x_j) = y$ and $x_i \neq x_j$ for all $1 \leq i \neq j \leq l$. An l -collision is simply called a *collision* for $l = 2$, and called a *multicollision* for $l \geq 3$. Moreover, an l -claw for functions $f_i: X_i \rightarrow Y$ for $1 \leq i \leq l$ is a tuple $(x_1, \dots, x_l, y) \in X_1 \times \dots \times X_l \times Y$ such that $f_1(x_1) = \dots = f_l(x_l) = y$. An l -claw is simply called a *claw* for $l = 2$, and called a *multiclaw* for $l \geq 3$.

The problems of finding multicollisions or multiclaws are often studied in the contexts of both cryptography and quantum computation, but the problem settings of interest change depending on the contexts. In the context of quantum computation, most problems are studied in the *worst case*, and an algorithm is said to (efficiently) solve a problem only when it does (efficiently) for all functions. On the other hand, most problems in cryptography are studied in the *average case*, since randomness is one of the most crucial notions in cryptography. In particular, we say that an algorithm (efficiently) solves a problem if it does so with a high probability on average over randomly chosen functions.

This paper focuses on the settings of interest in the context of cryptography. Formally, our goal is to solve the following two problems.

Problem 1 (Multicollision-finding problem, average case). Let l be a positive integer constant, and X, Y denote non-empty finite sets. Suppose that a function $F: X \rightarrow Y$ is chosen uniformly at random and given as a quantum oracle. Then, find an l -collision for F .

Problem 2 (Multiclaw-finding problem, average case). Let l be a positive integer constant, and X_1, \dots, X_l, Y denote non-empty finite sets. Suppose that functions $f_i: X_i \rightarrow Y (1 \leq i \leq l)$ are chosen independently and uniformly at random, and given as quantum oracles. Then, find an l -claw for f_1, \dots, f_l .

Roughly speaking, Problem 1 is easier to solve than Problem 2. Suppose that $F: X \rightarrow Y$ is a function, and we want to find an l -collision for F . Let X_1, \dots, X_l be subsets of X such that $X_i \cap X_j = \emptyset$ for $i \neq j$ and $\bigcup_i X_i = X$. If (x_1, \dots, x_l, y) is an l -claw for $F|_{X_1}, \dots, F|_{X_l}$, then it is obviously an l -collision for F . In general, an algorithm for finding an l -claw can be converted into one for finding l -collisions. To be precise, the following lemma holds.

Lemma 1. *Let X, Y be non-empty finite sets, and X_1, \dots, X_l be subsets of X such that $X_i \cap X_j = \emptyset$ for $i \neq j$ and $\bigcup_i X_i = X$. If there exists a quantum*

algorithm \mathcal{A} that solves Problem 2 for the sets X_1, \dots, X_l, Y by making at most q quantum queries with probability at least p , then there exists a quantum algorithm \mathcal{B} that solves Problem 1 for the sets X, Y by making at most q quantum queries with probability at least p .

How to measure the size of a problem also changes depending on which context we are in. In the context of cryptography, the problem size is often regarded as the size of the range of functions in the problem rather than the size of the domains, since the domains of cryptographic functions such as hash functions are much larger than their ranges. Hence, we regard the range size $|Y|$ as the size of Problem 1 (and Problem 2) when we analyze the complexity of quantum algorithms.

In the context of quantum computation, there exist previous works on problems related to ours [Bel12, Amb04, Tan09, BDH⁺01] (element distinctness problem, for example), but those works usually focus on the worst case complexity and regard the domain sizes of functions as the problem size. In particular, there does not exist any previous work that studies multiclaw-finding problem for general l in the average case, to the best of authors' knowledge.

3 Previous Works

3.1 The Grover Search and Its Generalization

As a main tool for developing quantum algorithms, we use the quantum database search algorithm that was originally developed by Grover [Gro96] and later generalized by Boyer, Brassard, Høyer, and Tapp [BBHT98]. Below we introduce the generalized version.

Theorem 1. *Let X be a non-empty finite set and $f: X \rightarrow \{0, 1\}$ be a function such that $t/|X| < 17/81$, where $t = |f^{-1}(1)|$. Then, there exists a quantum algorithm BBHT that finds x such that $f(x) = 1$ with an expected number of quantum queries to f at most*

$$\frac{4|X|}{\sqrt{(|X| - t)t}} \leq \frac{9}{2} \cdot \sqrt{\frac{|X|}{t}}.$$

If $f^{-1}(1) = \emptyset$, then BBHT runs forever.

Theorem 1 implies that we can find l -collisions and l -claws for random functions with $O(\sqrt{N})$ quantum queries, if the sizes of range(s) and domain(s) of function(s) are $\Theta(N)$: Suppose that we are given functions $f_i: X_i \rightarrow Y$ for $1 \leq i \leq l$, where $|X_1|, \dots, |X_l|$, and $|Y|$ are all in $\Theta(N)$, and we want to find an l -claw for those functions. Take an element $y \in Y$ randomly, and define $F_i: X_i \rightarrow \{0, 1\}$ for each i by $F_i(x) = 1$ if and only if $f_i(x) = y$. Then, effectively, by applying BBHT to each F_i , we can find $x_i \in X_i$ such that $f_i(x_i) = y$ for each i with $O(\sqrt{N})$ quantum queries. Similarly we can find an l -collision for a random function $F: [N] \rightarrow [N]$ with $O(\sqrt{N})$ quantum queries. In particular, $O(\sqrt{N})$ is a trivial upper bound of Problem 1 and Problem 2.

3.2 The BHT Algorithm

Brassard, Høyer, and Tapp [BHT98] developed a quantum algorithm that finds 2-claws (below we call it BHT).³ BHT finds a claw for two one-to-one functions $f_1: X_1 \rightarrow Y$ and $f_2: X_2 \rightarrow Y$ as sketched in the following. For simplicity, here we assume $|X_1| = |X_2| = |Y| = N$. Under this setting, BHT finds a 2-claw with $O(N^{1/3})$ quantum queries.

Rough Sketch of BHT :

1. **Construction of a list L .** Take a subset $S \subset X_1$ of size $N^{1/3}$ arbitrarily. For each $x \in S$, compute the value $f_1(x)$ by making a query and store the pair $(x, f_1(x))$ in a list L .
2. **Extension to a claw.** Define a function $F_L: X_2 \rightarrow \{0, 1\}$ by $F_L(x') = 1$ if and only if the value $f_2(x') \in Y$ appears in the list L (i.e., there exists $x_1 \in S$ such that $f_2(x') = f_1(x_1)$). Apply BBHT to F_L and find $x_2 \in X_2$ such that $f_2(x_2)$ appears in L .
3. **Finalization.** Find $(x_1, f_1(x_1)) \in L$ such that $f_1(x_1) = f_2(x_2)$, and then output (x_1, x_2) .

Quantum query complexity. BHT finds a claw with $O(N^{1/3})$ quantum queries. In the first step, the list L is constructed by making $N^{1/3}$ queries to f_1 . In the second step, since $|F_L^{-1}(1)| = |f_2^{-1}(f_1(S))|$ is equal to $N^{1/3}$, BBHT finds x_2 with $O(\sqrt{N/N^{1/3}}) = O(N^{1/3})$ quantum queries to f_2 (note that we can evaluate F_L by making one query to f_2). The third step does not require queries. Therefore BHT finds a collision by making $O(N^{1/3})$ quantum queries in total.

Extension to a collision-finding algorithm. It is not difficult to show that BHT works for random functions. Thus, BHT can be extended to the quantum collision-finding algorithm as mentioned in section 2. Suppose we want to find a (2-)collision for a random function $F: X \rightarrow Y$. Here we assume $|X| = 2N$ and $|Y| = N$ for simplicity.

Now, choose a subset $X_1 \subset X$ of size N arbitrarily and let $X_2 := X \setminus X_1$. Then we can find a collision for F by applying the BHT algorithm introduced above to the functions $F|_{X_1}$ and $F|_{X_2}$, since a claw for them becomes a collision for F .

³ As in our case, the BHT algorithm also focus on only quantum query complexity. Although it runs in time $\tilde{O}(N^{1/3})$ on an idealized quantum computer, it requires $\tilde{O}(N^{1/3})$ qubits to store data in quantum memories. Recently Chailloux et al. [CNS17] has developed a quantum 2-collision finding algorithm that runs in time $\tilde{O}(N^{2/5})$, which is polynomially slower than the BHT algorithm but requires only $O(\log N)$ quantum memories.

3.3 The HSX Algorithm

Next, we introduce a quantum algorithm for finding multicollisions that was developed by Hosoyamada, Sasaki, and Xagawa [HSX17] (the algorithm is designed to find only multicollisions, and cannot find multiclaws). Below we call their algorithm HSX.

The main idea of HSX is to apply the strategy of BHT recursively: To find an l -collision, HSX calls itself recursively to find many $(l-1)$ -collisions, and then extend any one of those $(l-1)$ -collisions to an l -collision by applying BBHT.

Rough Sketch of HSX : In what follows, N denotes $|Y|$. Let us denote HSX(l) by the HSX algorithm for finding l -collisions. HSX(l) finds an l -collision for a random function $f: X \rightarrow Y$ with $|X| \geq l \cdot |Y|$ as follows.

Recursive call to construct a list L_{l-1} . Apply HSX($l-1$) to f $N^{1/3^{l-1}}$ times to obtain $N^{1/3^{l-1}}$ many $(l-1)$ -collisions. Store those $(l-1)$ -collisions in a list L_{l-1} .

Extension to an l -collision. Define $F_{l-1}: X \rightarrow \{0, 1\}$ by $F_{l-1}(x') = 1$ if and only if there exists an $(l-1)$ -collision $(x_1, \dots, x_{l-1}, y) \in L_{l-1}$ such that $(x_1, \dots, x_{l-1}, x', y)$ forms an l -collision for f , i.e., $f(x') = y$ and $x' \neq x_i$ for $1 \leq i \leq l-1$. Apply BBHT to F_{l-1} to find $x_l \in X$ such that $F_{l-1}(x_l) = 1$.

Finalization. Find $(x_1, \dots, x_{l-1}, y) \in L_{l-1}$ such that $F_{l-1}(x_l) = y$. Output $(x_1, \dots, x_{l-1}, x_l, y)$.

Quantum query complexity. HSX finds a l -collision with $O(N^{(3^{l-1}-1)/2 \cdot 3^{l-1}})$ quantum queries on average, which can be shown by induction as follows. For 2-collisions, HSX(2) matches the BHT algorithm. For general $l \geq 3$, suppose that HSX($l-1$) finds an $(l-1)$ -collision with $O(N^{(3^{l-2}-1)/2 \cdot 3^{l-2}})$ quantum queries on average. In its first step, HSX(l) makes $N^{1/3^{l-1}} \cdot O(N^{(3^{l-2}-1)/2 \cdot 3^{l-2}}) = O(N^{(3^{l-1}-1)/2 \cdot 3^{l-1}})$ quantum queries. Moreover, in its second step, HSX(l) makes $O(\sqrt{N/N^{(3^{l-2}-1)/2 \cdot 3^{l-2}}}) = O(N^{(3^{l-1}-1)/2 \cdot 3^{l-1}})$ quantum queries by using BBHT. The third step does not make quantum queries. Therefore it follows that HSX(l) makes $O(N^{(3^{l-1}-1)/2 \cdot 3^{l-1}})$ quantum queries in total.

4 New Quantum Algorithm Mclaw

This section gives our new quantum algorithm Mclaw that finds an l -claw for random functions $f_i: X_i \rightarrow Y$ for $1 \leq i \leq l$, where $|Y| = N$ and there exists a parameter c_N such that c_N is $o(N^{\frac{1}{2^{l-1}}})$ and $\frac{N}{c_N} \leq |X_i|$ holds for all i , with $O(c_N^{3/2} N^{(2^{l-1}-1)/(2^l-1)})$ quantum queries. Roughly speaking, this means that, an l -collision for a random function $f: X \rightarrow Y$, where $|X|$ and $|Y|$ are in $O(N)$, can be found with $O(N^{(2^{l-1}-1)/(2^l-1)})$ quantum queries, which improves the previous result [HSX17] (see section 2).

Our algorithm assumes that $|X_1|, \dots, |X_l|$ are less than or equal to $|Y|$. However, it can also be applied to the functions of interest in the context of cryptography, i.e., the functions of which domains are much larger than ranges, by restricting the domains of them to suitable subsets.

The main idea of our new algorithm is to improve HSX by getting rid of its redundancy: To find an l -collision, HSX recursively calls itself to find many $(l-1)$ -collisions. Once HSX finds an $(l-1)$ -collision $\gamma = (x_1, \dots, x_{l-1}, y)$, it stores γ in a list L_{l-1} , *discards all the data that was used to find γ* , and then start to search for another $(l-1)$ -collision γ' . It is inefficient to discard data every time an $(l-1)$ -collision is found, and our new algorithm **Mclaw** reduces the number of quantum queries by reusing those data. Moreover, HSX cannot solve multiclaw-finding problem while the BHT algorithm can also solve claw-finding problem. Our algorithm **Mclaw** can solve both of two problems.

We begin with describing intuition of our algorithm, and then give its formal description.

4.1 Intuitive Description and Complexity Analysis

We explain the idea of how to develop the BHT algorithm, how to develop a quantum algorithm to find 3-claws from BHT, and how to extend it further to the case of finding an l -claw for any l .

How to develop the BHT algorithm. Here we review how the BHT algorithm is developed. Let $f_1: X_1 \rightarrow Y$ and $f_2: X_2 \rightarrow Y$ be one-to-one functions. The goal of the BHT algorithm is to find a (2-)claw for f_1 and f_2 with $O(N^{1/3})$ quantum queries. For simplicity, below we assume that $|X_1| = |X_2| = |Y| = N$ holds. Let t_1 be a parameter that defines the size of a list of 1-claws for f_1 . It will be set as $t_1 = N^{1/3}$.

First, collect t_1 many 1-claws for f_1 and store them in a list L_1 . This first step makes t_1 queries. Second, extend 1-claws in L_1 to a 2-claw for f_1 and f_2 , by using BBHT, and output the obtained 2-claw. Since BBHT makes $O(\sqrt{N/t_1})$ queries to make a 2-claw from L_1 , this second step makes $O(\sqrt{N/t_1})$ queries (see Theorem 1). Overall, the above algorithm makes $q_2(t_1) = t_1 + \sqrt{N/t_1}$ quantum queries up to a constant factor.

The function $q_2(t_1)$ takes its minimum value $2 \cdot N^{1/3}$ when $t_1 = N^{1/3}$. The BHT algorithm is developed in this way, by setting $t_1 = N^{1/3}$.

From BHT to a 3-claw-finding algorithm. Next, we show how the above strategy to develop the BHT algorithm can be extended to develop a 3-claw-finding algorithm. Let $f_i: X_i \rightarrow Y$ be one-to-one functions for $1 \leq i \leq 3$. Our goal here is to find a 3-claw for f_1 , f_2 , and f_3 with $O(N^{3/7})$ quantum queries. For simplicity, below we assume $|X_1| = |X_2| = |X_3| = |Y| = N$. Let t_1, t_2 be parameters that define the number of 1-claws for f_1 and that of 2-claws for f_1 and f_2 , respectively. (They will be fixed later.)

First, collect t_1 many 1-claws for f_1 and store them in a list L_1 . This first step makes t_1 queries. Second, extend 1-claws in L_1 to t_2 many 2-claws for f_1 and f_2 , by using BBHT, and store them in a list L_2 . Here we do not discard the list L_1 until we construct the list L_2 of size t_2 , while the HSX algorithm does. Since BBHT makes $O(\sqrt{N/t_1})$ queries to make a 2-claw from L_1 , this second step makes $t_2 \cdot O(\sqrt{N/t_1})$ queries (see Theorem 1). Finally, extend 2-claws in L_2 to a 3-claw for f_1 , f_2 , and f_3 , by using BBHT, and output the obtained 3-claw. This final step makes $O(\sqrt{N/t_2})$ queries. Overall, the above algorithm makes $q_3(t_1, t_2) = t_1 + t_2 \cdot \sqrt{N/t_1} + \sqrt{N/t_2}$ quantum queries up to a constant factor.

The function $q_3(t_1, t_2)$ takes its minimum value $3 \cdot N^{3/7}$ when $t_1 = t_2 \cdot \sqrt{N/t_1} = \sqrt{N/t_2}$, which is equivalent to $t_1 = N^{3/7}$ and $t_2 = N^{1/7}$. We can develop a 3-claw finding algorithm with $O(N^{3/7})$ quantum queries in this way, by setting $t_1 = N^{3/7}$ and $t_2 = N^{1/7}$.

l -claw-finding algorithm for general l . Generalizing the above idea to find 3-claws, we can find l -claws for general l as follows. Let $f_i: X_i \rightarrow Y$ be one-to-one functions for $1 \leq i \leq l$. Our goal here is to find an l -claw for f_1, \dots, f_l . For simplicity, below we assume that $|X_1| = \dots = |X_l| = |Y| = N$ holds. Let t_1, \dots, t_{l-1} be parameters.

First, collect t_1 many 1-claws for f_1 and store them in a list L_1 . This first step makes t_1 queries. In the i -th step for $2 \leq i \leq l-1$, extend $(i-1)$ -claws in L_{i-1} to t_i many i -claws for f_1, \dots, f_i , by using BBHT, and store them in a list L_i . Here we do not discard the list L_{i-1} until we construct the list L_i of size t_i . Since BBHT makes $O(\sqrt{N/t_{i-1}})$ queries to make an i -claw from L_{i-1} , the i -th step makes $t_i \cdot O(\sqrt{N/t_{i-1}})$ queries. Finally, extend $(l-1)$ -claws in L_{l-1} to an l -claw for f_1, \dots, f_l , by using BBHT, and output the obtained l -claw. This final step makes $O(\sqrt{N/t_{l-1}})$ queries. Overall, this algorithm makes $q_l(t_1, \dots, t_{l-1}) = t_1 + t_2 \cdot \sqrt{N/t_1} + \dots + t_{l-1} \cdot \sqrt{N/t_{l-2}} + \sqrt{N/t_{l-1}}$ quantum queries up to a constant factor. The function $q_l(t_1, \dots, t_{l-1})$ takes its minimum value $l \cdot N^{(2^{l-1}-1)/(2^l-1)}$ when $t_1 = t_2 \cdot \sqrt{N/t_1} = \dots = t_{l-1} \cdot \sqrt{N/t_{l-2}} = \sqrt{N/t_{l-1}}$, which is equivalent to $t_i = N^{(2^{l-i}-1)/(2^l-1)}$. Hence we can find an l -claw with $O(N^{(2^{l-1}-1)/(2^l-1)})$ quantum queries, by setting $t_i = N^{(2^{l-i}-1)/(2^l-1)}$. Our new quantum algorithm Mclaw is developed based on the above strategy for random functions.

4.2 Formal Description

Here we formally describe our quantum multiclaw-finding algorithm Mclaw. A formal complexity analysis of Mclaw is given in the next subsection, and this subsection only describes how the algorithm works.

Let N be a sufficiently large integer and suppose that $|Y| = N$ holds. Below we assume that $|X_i| \leq |Y|$ holds for all i . This is a reasonable assumption since, if there is an algorithm that solves Problem 2 in the case that $|X_i| \leq |Y|$ holds for all i , then we can also solve the problem in other cases: If $|X_i| > |Y|$ holds

for some i , take a subset $S_i \subset X_i$ such that $|S_i| = |Y|$ and find an l -claw for $f_1, \dots, f_{i-1}, f_i|_{S_i}, f_{i+1}, \dots, f_l$. Then the l -claw is also an l -claw for f_1, \dots, f_l .

Here we introduce a corollary that follows from Theorem 1.

Corollary 1. *Let X, Y be non-empty finite sets, $f: X \rightarrow Y$ be a function, and $L' \subset Y$ be a non-empty subset. Then there exists a quantum algorithm **MTPS** that finds x such that $f(x) \in L'$ with an expected number of quantum queries to f at most $9\sqrt{5|X|/|f^{-1}(L')|}$.*

Let $F_{L'}: \{1, \dots, 5\} \times X \rightarrow \{0, 1\}$ be the boolean function defined by $F_{L'}(\alpha, x) = 1$ if and only if $\alpha = 1$ and $f(x) \in L'$. A quantum circuit that computes $F_{L'}$ can be implemented with two oracle calls to f . Then, run **BBHT** on $F_{L'}$. Since $|\{1, \dots, 5\} \times X| = 5|X|$ and $|F_{L'}^{-1}(1)| \leq |X| \leq 17/81 \cdot |\{1, \dots, 5\} \times X|$ always hold, we can show that the corollary follows from Theorem 1.

Our algorithm is parametrized by a positive integer k , and we denote the algorithm for the parameter k by **Mclaw $_k$** . **Mclaw $_k$** can be applied in the situation that there exists a parameter c_N such that c_N is in $o(N^{\frac{1}{2^l-1}})$ and $|X_i| \geq |Y|/c_N$ holds for each i . We impose an upper limit on the number of queries that **Mclaw $_k$** is allowed to make: We design **Mclaw $_k$** in such a way that it immediately stops and aborts if the number of queries made reaches the limit specified by the parameter $\text{Qlimit}_k = k \cdot 169lc_N^{3/2} \cdot N^{\frac{2^l-1-1}{2^l-1}}$. The upper limit Qlimit_k is necessary to prevent the algorithm from running forever, and to make the expected value of the number of queries converge. We also define the parameters controlling the sizes of the lists:

$$N_i = \begin{cases} \frac{N}{4c_N} & (i = 0), \\ N^{\frac{2^l-i-1}{2^l-1}} & (i \geq 1). \end{cases} \quad (1)$$

For ease of notation, we define L_0 and L'_0 as $L_0 = L'_0 = Y$. Then, **Mclaw $_k$** is described as in Algorithm 1.

4.3 Formal Complexity Analysis

This section gives a formal complexity analysis of **Mclaw**. The goal of this section is to show the following theorem.

Theorem 2. *Assume that there exists a parameter c_N such that c_N is in $o(N^{\frac{1}{2^l-1}})$ and $|X_i| \geq \frac{1}{c_N}|Y|$ holds for each i . If $|Y| = N$ is sufficiently large, **Mclaw $_k$** finds an l -claw with a probability at least*

$$1 - \frac{2l}{N} - l \cdot \exp\left(-\frac{1}{15} \cdot \frac{N^{\frac{1}{2^l-1}}}{c_N}\right), \quad (2)$$

by making at most

$$\text{Qlimit}_k = k \cdot 169lc_N^{3/2} \cdot N^{\frac{2^l-1-1}{2^l-1}} \quad (3)$$

quantum queries.

Algorithm 1 Mclaw_k

Input: Randomly chosen functions f_1, \dots, f_l ($f_i: X_i \rightarrow Y$ and $|X_i| \leq |Y|$).

Output: An l -multiclause for f_1, \dots, f_l or \perp .

Stop condition: If the number of queries reaches Qlimit_k, stop and output \perp .

$L_1, \dots, L_l \leftarrow \emptyset, L'_1, \dots, L'_l \leftarrow \emptyset$.

for $i = 1$ to l **do**

for $j = 1$ to $\lceil 4c_N \cdot N_i \rceil$ **do**

if $i = 1$ **then**

 Take $x_j \in X_1$ that does not appear in $L_1, y \leftarrow f_1(x_j)$. //1 query is made

else

 Find $x_j \in X_i$ whose image $y := f_i(x_j)$ is in L'_{i-1} by running MTPS on f_i and L'_{i-1} . //multiple queries are made

end if

$L_i \leftarrow L_i \cup \{(x^{(1)}, \dots, x^{(i-1)}, x_j, y)\}, L'_i \leftarrow L'_i \cup \{y\}$.

$L_{i-1} \leftarrow L_{i-1} \setminus \{(x^{(1)}, \dots, x^{(i-1)}, y)\}, L'_{i-1} \leftarrow L'_{i-1} \setminus \{y\}$.

end for

end for

Return an element $(x^{(1)}, \dots, x^{(l)}; y) \in L_l$ as an output.

This theorem shows that, for each small integer $k \geq 2$, Mclaw_k finds an l -claw with an overwhelming probability by making $O\left(c_N^{3/2} N^{\frac{2^l-1-1}{2^l-1}}\right)$ queries.

For later use, we show the following lemma.

Lemma 2. *Let X, Y be non-empty finite sets such that $|X| \leq |Y|$. Suppose that a function $f: X \rightarrow Y$ is chosen uniformly at random. Then*

$$\Pr_{f \sim U(\text{Func}(X, Y))} \left[|\text{Im}(f)| \geq \frac{|X|}{2} - \sqrt{|X| \ln |Y|/2} \right] \geq 1 - \frac{2}{|Y|} \quad (4)$$

holds.

Proof. Note that, for each $x \in X$, $f(x)$ is the random variable that takes value in Y . Moreover, $\{f(x)\}_{x \in X}$ is the set of independent random variables. Let us define a function $\Phi: Y^{\times |X|} \rightarrow \mathbb{N}$ by $\Phi(y_1, \dots, y_{|X|}) = |Y \setminus \{y_i\}_{1 \leq i \leq |X|}|$. Then Φ is 1-Lipschitz, i.e.,

$$|\Phi(y_1, \dots, y_{i-1}, y_i, y_{i+1}, \dots, y_{|X|}) - \Phi(y_1, \dots, y_{i-1}, y'_i, y_{i+1}, \dots, y_{|X|})| \leq 1 \quad (5)$$

holds for arbitrary choice of elements $y_1, \dots, y_{|X|}$, and y'_i in Y . Now we apply the following theorem to Φ .

Theorem 3 (McDiarmid's Inequality (Theorem 13.7 in [MU17])). *Let M be a positive integer, and $\Phi: Y^{\times M} \rightarrow \mathbb{N}$ be a 1-Lipschitz function. Let $\{y_i\}_{1 \leq i \leq M}$ be the set of independent random variables that take values in Y . Let μ denote the expectation value $\mathbf{E}_{y_1, \dots, y_M} [\Phi(y_1, \dots, y_M)]$. Then*

$$\Pr_{y_1, \dots, y_M} [\Phi(y_1, \dots, y_M) \geq \mu + \lambda] \leq 2e^{-2\lambda^2/M} \quad (6)$$

holds.

Apply the above theorem with $M = |X|$, $\lambda = \sqrt{|X| \ln |Y|/2}$, and $y_x = f(x)$ for each $x \in X$ (here we identify X with the set $\{1, \dots, |X|\}$). Then, since $\mathbf{E}[\Phi(y_1, \dots, y_M)] = |Y| (1 - 1/|Y|)^{|X|}$ holds, we have that

$$\Pr_{f \sim U(\text{Func}(X, Y))} \left[\Phi(y_1, \dots, y_M) \geq |Y| (1 - 1/|Y|)^{|X|} + \sqrt{|X| \ln |Y|/2} \right] \leq \frac{2}{|Y|}.$$

In addition, it follows that

$$\begin{aligned} |Y| (1 - 1/|Y|)^{|X|} &\leq |Y| e^{-|X|/|Y|} \leq |Y| \left(1 - \frac{|X|}{|Y|} + \frac{1}{2} \left(\frac{|X|}{|Y|} \right)^2 \right) \\ &= |Y| - |X| \left(1 - \frac{1}{2} \frac{|X|}{|Y|} \right) \leq |Y| - \frac{|X|}{2}, \end{aligned} \quad (7)$$

where we used the assumption that $|X| \leq |Y|$ for the last inequality. Since $\Phi(y_1, \dots, y_M) = |Y \setminus \text{Im}(f)|$ and $|\text{Im}(f)| = |Y| - |Y \setminus \text{Im}(f)|$ hold, it follows that $|Y \cap \text{Im}(f)|$ is less than or equal to

$$|Y| - \left(|Y| - \frac{|X|}{2} + \sqrt{|X| \ln |Y|/2} \right) = \frac{|X|}{2} - \sqrt{|X| \ln |Y|/2} \quad (8)$$

with a probability at least $1 - \frac{2}{|Y|}$, which completes the proof. \square

Proof (of Theorem 2). We show that Equation 2 holds. Let us define $\text{good}^{(i)}$ to be the event that

$$|\text{Im}(f_i) \cap L'_{i-1}| \geq N_{i-1} \quad (9)$$

holds just before Mclaw_k starts to construct i -multiclaws. (Intuitively, under the condition that $\text{good}^{(i)}$ occurs, the number of queries does not become too large.) We show the following claim.

Claim. For sufficiently large N ,

$$\Pr \left[\text{good}^{(i)} \right] \geq 1 - \frac{2}{N} - \exp \left(-\frac{1}{15} \cdot \frac{N_{i-1}}{c_N} \right). \quad (10)$$

holds.

Proof. In this proof we consider the situation that Mclaw_k has finished to make L_{i-1} and before starting to make i -claws. In particular, we assume that $|L_{i-1}| = |L'_{i-1}| = \lceil 4c_N N_{i-1} \rceil$.

Let $\text{pregood}^{(i)}$ be the event that $|\text{Im}(f_i)| \geq \lceil N/3c_N \rceil$ holds. Since c_N is in $o(N^{\frac{1}{2^{i-1}}})$, we have that $\frac{|X_i|}{2} - \sqrt{|X_i| \ln |Y|/2} \geq \lceil \frac{N}{3c_N} \rceil$ holds for sufficiently large N . Hence

$$\Pr \left[\text{pregood}^{(i)} \right] \geq 1 - \frac{2}{|Y|} \quad (11)$$

follows from Lemma 2.

Let us identify X and Y with the sets $\{1, \dots, |X|\}$ and $\{1, \dots, |Y|\}$, respectively. Let B_j be the j -th element in $\text{Im}(f)$. Let χ_j be the indicator variable that is defined by $\chi_j = 1$ if and only if $B_j \in L'_{i-1}$, and define a random variable χ by $\chi := \sum_j \chi_j$. Then χ follows the hypergeometric distribution under the condition that $|\text{Im}(f_i)| = \lceil N/3c_N \rceil$ holds. We use the following theorem as a fact.

Theorem 4 (Theorem 1 in [HS05]). *Let $K = K(n_1, n, m)$ denote the hypergeometric random variable describing the process of counting how many defectives are selected when n_1 items are randomly selected without replacement from n items among which there are m defective ones. Let $\lambda \geq 2$. Then*

$$\Pr[K - \mathbf{E}[K] < -\lambda] < e^{-2\alpha_{n_1, n, m}(\lambda^2 - 1)} \quad (12)$$

holds, where

$$\alpha_{n_1, m, n} = \max\left(\left(\frac{1}{n_1 + 1} + \frac{1}{n - n_1 + 1}\right), \left(\frac{1}{m + 1} + \frac{1}{n - m + 1}\right)\right). \quad (13)$$

Apply the above theorem with $n_1 = \lceil N/3c_N \rceil$, $n = N$, and $m = |L'_{i-1}| = \lceil 4c_N N_{i-1} \rceil$, for the random variable χ under the condition that $|\text{Im}(f_i)| = \lceil N/3c_N \rceil$ holds. Let **equal** denote the event that $|\text{Im}(f_i)| = \lceil N/3c_N \rceil$ holds. Then $\mathbf{E}[\chi | \text{equal}] = \frac{n_1 m}{n} \geq \frac{4}{3} N_{i-1}$ holds, and we have that

$$\begin{aligned} & \Pr\left[\chi - \mathbf{E}[\chi | \text{equal}] < -\frac{1}{4} \mathbf{E}[\chi | \text{equal}] \mid \text{equal}\right] \\ & \leq \exp\left(-2\left(\frac{1}{m+1} + \frac{1}{n-m+1}\right)\left(\frac{\mathbf{E}[\chi | \text{equal}]}{4}\right)^2 - 1\right) \\ & \leq \exp\left(-\frac{1}{15m} (\mathbf{E}[\chi | \text{equal}])^2\right) \leq \exp\left(-\frac{1}{15} \cdot \frac{N_{i-1}}{c_N}\right) \end{aligned} \quad (14)$$

for sufficiently large N . Hence

$$\Pr[\chi \geq N_{i-1} | \text{equal}] \geq 1 - \exp\left(-\frac{1}{15} \cdot \frac{N_{i-1}}{c_N}\right) \quad (15)$$

holds, which implies that

$$\begin{aligned} \Pr\left[|\text{Im}(f) \cap L'_{i-1}| \geq N_{i-1} \mid \text{pregood}^{(i)}\right] &= \Pr\left[\chi \geq N_{i-1} \mid \text{pregood}^{(i)}\right] \\ &\geq \Pr[\chi \geq N_{i-1} | \text{equal}] \\ &\geq 1 - \exp\left(-\frac{1}{15} \cdot \frac{N_{i-1}}{c_N}\right). \end{aligned} \quad (16)$$

Therefore we have that

$$\begin{aligned}
\Pr \left[\text{good}^{(i)} \right] &= \Pr \left[\text{good}^{(i)} \mid \text{pregood}^{(i)} \right] \cdot \Pr \left[\text{pregood}^{(i)} \right] \\
&= \Pr \left[|\text{Im}(f) \cap L'_{i-1}| \geq N_{i-1} \mid \text{pregood}^{(i)} \right] \cdot \Pr \left[\text{pregood}^{(i)} \right] \\
&\geq \left(1 - \frac{2}{|Y|} \right) \left(1 - \exp \left(-\frac{1}{15} \cdot \frac{N_{i-1}}{c_N} \right) \right) \\
&\geq 1 - \frac{2}{|Y|} - \exp \left(-\frac{1}{15} \cdot \frac{N_{i-1}}{c_N} \right), \tag{17}
\end{aligned}$$

which completes the proof. \square

Let good denote the event $\text{good}^{(1)} \wedge \dots \wedge \text{good}^{(l)}$. Then we can show the following claim.

Claim. For sufficiently large N , it holds that

$$\mathbf{E}[Q \mid \text{good}] \leq \frac{1}{k} \text{Qlimit}_k, \tag{18}$$

where Q is the total number of queries made by Mclaw_k .

Proof. Let us fix i and j . Let $Q_j^{(i)}$ denote the number of queries made by Mclaw_k in the j -th search to construct i -multiclaws, and $Q^{(i)}$ denote $\sum_j Q_j^{(i)}$. In the j -th search to construct i -multiclaw, we search x from X_i , where there exist at least $|L'_{i-1} \cap \text{Im}(f)| \geq N_{i-1} - j + 1$ answers under the condition that $\text{good}^{(i)}$ occurs. From Corollary 1, the expected value of the number of queries made by MTPS in the j -th search to construct i -multiclaws is upper bounded by

$$9\sqrt{5|X_i|/|f_i^{-1}(L'_{i-1})|} \leq 9\sqrt{5|X_i|/|L'_{i-1} \cap \text{Im}(f)|} \leq 21\sqrt{N/N_{i-1}} \tag{19}$$

for each j under the condition that $\text{good}^{(i)}$ occurs, for sufficiently large N (we used the condition that $N_{i-1} = \omega(c_N N_i)$ holds for the last inequality).

Hence it follows that

$$\begin{aligned}
\mathbf{E} \left[Q^{(i)} \mid \text{good}^{(i)} \right] &= \mathbf{E} \left[\sum_j Q_j^{(i)} \mid \text{good}^{(i)} \right] = \sum_j \mathbf{E} \left[Q_j^{(i)} \mid \text{good}^{(i)} \right] \\
&\leq \sum_{1 \leq j \leq \lceil 4c_N N_i \rceil} 21\sqrt{N/N_{i-1}} \leq \begin{cases} 169c_N^{3/2} N^{\frac{2^{l-1}-1}{2^{l-1}}} & (i=1) \\ 85c_N N^{\frac{2^{l-1}-1}{2^{l-1}}} & (i \geq 2) \end{cases}
\end{aligned}$$

for sufficiently large N . Hence we have that $\mathbf{E}[Q \mid \text{good}] = \sum_i \mathbf{E} \left[Q^{(i)} \mid \text{good}^{(i)} \right]$ is upper bounded by

$$169c_N^{3/2} N^{\frac{2^{l-1}-1}{2^{l-1}}} + \sum_{i=2}^l 85c_N N^{\frac{2^{l-1}-1}{2^{l-1}}} \leq 169lc_N^{3/2} \cdot N^{\frac{2^{l-1}-1}{2^{l-1}}} = \frac{1}{k} \text{Qlimit}_k,$$

which completes the proof. \square

From the above claims it follows that $\mathbf{E}[Q]$ is upper bounded by

$$\mathbf{E}[Q \mid \text{good}] + \mathbf{E}[Q \mid \neg\text{good}] \Pr[\neg\text{good}] \leq \left(\frac{1}{k} + \Pr[\neg\text{good}] \right) \cdot \text{Qlimit}_k, \quad (20)$$

and $\Pr[\neg\text{good}]$ is upper bounded by $\sum_i \Pr[\neg\text{good}^{(i)}]$, which is further upper bounded by

$$\sum_i \left(\frac{2}{N} + \exp\left(-\frac{1}{15} \cdot \frac{N_{i-1}}{c_N}\right) \right) \leq \frac{2l}{N} + l \cdot \exp\left(-\frac{1}{15} \cdot \frac{N^{\frac{1}{2^{l-1}}}}{c_N}\right). \quad (21)$$

From Markov's inequality, the probability that Q reaches Qlimit_k is at most

$$\Pr[Q \geq \text{Qlimit}_k] \leq \frac{\mathbf{E}[Q]}{\text{Qlimit}_k} \leq \frac{1}{k} + \Pr[\neg\text{good}]. \quad (22)$$

If Mclaw_k finds an l -claw, then Q does not reach Qlimit_k . Thus, from Equation 21 and Equation 22, the probability that Mclaw_k finds an l -claw is lower bounded by

$$1 - \frac{2l}{N} - l \cdot \exp\left(-\frac{1}{15} \cdot \frac{N^{\frac{1}{2^{l-1}}}}{c_N}\right), \quad (23)$$

which completes the proof. \square

5 Conclusion

This paper has developed a new quantum algorithm to find multicollisions of random functions. Our new algorithm finds an l -collision of a random function $F: [N] \rightarrow [N]$ with $O\left(N^{(2^{l-1}-1)/(2^l-1)}\right)$ quantum queries, which improves the previous upper bound $O(N^{(3^{l-1})/(2 \cdot 3^l-1)})$ by Hosoyamada *et al.* [HSX17].⁴ In fact, our algorithm can find an l -claw of random functions for $f_i: [N] \rightarrow [N]$ for $1 \leq i \leq l$ with the same complexity $O\left(N^{(2^{l-1}-1)/(2^l-1)}\right)$.

References

- Amb04. Andris Ambainis. Quantum walk algorithm for element distinctness. In *45th Symposium on Foundations of Computer Science (FOCS 2004), 17-19 October 2004, Rome, Italy, Proceedings*, pages 22–31, 2004.
- BBHT98. Michel Boyer, Gilles Brassard, Peter Høyer, and Alain Tapp. Tight bounds on quantum searching. *Fortschritte der Physik: Progress of Physics*, 46(4-5):493–505, 1998.

⁴ As we mentioned in section 1, Liu and Zhandry [LZ18] showed that this bound is tight.

- BDH⁺01. Harry Buhrman, Christoph Dürr, Mark Heiligman, Peter Høyer, Frédéric Magniez, Miklos Santha, and Ronald de Wolf. Quantum algorithms for element distinctness. In *Proceedings of the 16th Annual IEEE Conference on Computational Complexity, Chicago, Illinois, USA, June 18-21, 2001*, pages 131–137, 2001.
- BDRV18. Itay Berman, Akshay Degwekar, Ron D. Rothblum, and Prashant Nalini Vasudevan. Multi-collision resistant hash functions and their applications. In *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part II*, pages 133–161, 2018.
- Bel12. Aleksandrs Belovs. Learning-graph-based quantum algorithm for k -distinctness. In *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ, USA, October 20-23, 2012*, pages 207–216, 2012.
- BHT98. Gilles Brassard, Peter Høyer, and Alain Tapp. Quantum cryptanalysis of hash and claw-free functions. In *LATIN '98: Theoretical Informatics, Third Latin American Symposium, Campinas, Brazil, April, 20-24, 1998, Proceedings*, pages 163–169, 1998.
- BKP18. Nir Bitansky, Yael Tauman Kalai, and Omer Paneth. Multi-collision resistance: a paradigm for keyless hash functions. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 671–684, 2018.
- CNS17. André Chailloux, María Naya-Plasencia, and André Schrottenloher. An efficient quantum collision search algorithm and implications on symmetric cryptography. In *Advances in Cryptology - ASIACRYPT 2017 - 23rd International Conference on the Theory and Applications of Cryptology and Information Security, Hong Kong, China, December 3-7, 2017, Proceedings, Part II*, pages 211–240, 2017.
- Gro96. Lov K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing, Philadelphia, Pennsylvania, USA, May 22-24, 1996*, pages 212–219, 1996.
- HS05. Don Hush and Clint Scovel. Concentration of the hypergeometric distribution. *Statistics & probability letters*, 75(2):127–132, 2005.
- HSX17. Akinori Hosoyamada, Yu Sasaki, and Keita Xagawa. Quantum multicollision-finding algorithm. In *Advances in Cryptology - ASIACRYPT 2017 - 23rd International Conference on the Theory and Applications of Cryptology and Information Security, Hong Kong, China, December 3-7, 2017, Proceedings, Part II*, pages 179–210, 2017.
- JLM14. Philipp Jovanovic, Atul Luykx, and Bart Mennink. Beyond $2^{c/2}$ security in sponge-based authenticated encryption modes. In *Advances in Cryptology - ASIACRYPT 2014 - 20th International Conference on the Theory and Application of Cryptology and Information Security, Kaoshiung, Taiwan, R.O.C., December 7-11, 2014. Proceedings, Part I*, pages 85–104, 2014.
- KNY18. Ilan Komargodski, Moni Naor, and Eylon Yogev. Collision resistant hashing for paranoids: Dealing with multiple collisions. In *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part II*, pages 162–194, 2018.

- LZ18. Qipeng Liu and Mark Zhandry. On finding quantum multi-collisions. *IACR Cryptology ePrint Archive*, 2018:1096, 2018.
- MU17. Michael Mitzenmacher and Eli Upfal. *Probability and computing: Randomization and probabilistic techniques in algorithms and data analysis*. Cambridge university press, 2017.
- RS96. Ronald L. Rivest and Adi Shamir. Payword and micromint: Two simple micropayment schemes. In *Security Protocols, International Workshop, Cambridge, United Kingdom, April 10-12, 1996, Proceedings*, pages 69–87, 1996.
- Tan09. Seiichiro Tani. Claw finding algorithms using quantum walk. *Theor. Comput. Sci.*, 410(50):5285–5297, 2009.
- Zha15. Mark Zhandry. A note on the quantum collision and set equality problems. *Quantum Information & Computation*, 15(7&8):557–567, 2015.