Quantum Time-Space Tradeoff for Finding Multiple Collision Pairs*

Yassine Hamoudi and Frédéric Magniez

Université de Paris, IRIF, CNRS, F-75006 Paris, France.
{hamoudi,magniez}@irif.fr

June 7, 2021

Abstract

We study the problem of finding K collision pairs in a random function $f : [N] \to [N]$ by using a quantum computer. We prove that the number of queries to the function in the quantum random oracle model must increase significantly when the size of the available memory is limited. Namely, we demonstrate that any algorithm using S qubits of memory must perform a number T of queries that satisfies the tradeoff $T^3 S \geq \Omega(K^3 N)$. Classically, the same question has only been settled recently by Dinur [Din20, Eurocrypt'20], who showed that the Parallel Collision Search algorithm of van Oorschot and Wiener [OW99] achieves the optimal time-space tradeoff of $T^2 S = \Theta(K^2 N)$. Our result limits the extent to which quantum computing may decrease this tradeoff. Our method is based on a novel application of Zhandry's recording query technique [Zha19, Crypto'19] for proving lower bounds in the exponentially small success probability regime. As a second application, we give a simpler proof of the time-space tradeoff $T^2 S \geq \Omega(N^3)$ for sorting N numbers on a quantum computer, which was first obtained by Klauck, Špalek and de Wolf [KŠW07].

1 Introduction

The *efficiency* of a cryptographic attack is a hard-to-define concept that must express the interplay between different computational resources [Wie04, Ber05, Ber09]. Arguably, the two most used criteria are the *time* complexity, measured for instance as the number of queries to a random oracle, and the *space* complexity, which is the memory size needed to perform the attack. *Time-space tradeoffs* aim at connecting these two quantities by studying how much the time increases when the available space decreases. Devising security proofs that are sensitive to memory constraints is a challenging program. Indeed, very few tools are available to quantify the extent to which the space impacts the security level of a scheme. A recent line of work [TT18, JT19, GJT20] has made some progress for the case of *classical* attackers with bounded memory. The development of quantum computing asks the question of whether the access to quantum operations and quantum memories may lower the security levels. The answer is unclear when taking space into account. Indeed, many quantum "speed-ups" come at the cost of a dramatic increase in the space requirement [BHT98, Amb07, LZ19]. A central open question is whether a speed-up both in term of time and space complexities is achievable for such problems?

The focus of this work is to provide time-space tradeoff lower bounds for the problem of finding *multiple collision pairs* in a random function. The search for a single collision pair is one of the cornerstones of cryptanalysis. Classically, the birthday attack can be achieved by the mean

*A previous version of this paper was titled *Quantum Time-Space Tradeoffs by Recording Queries*.

of a *memoryless* (i.e. logarithmic-size memory) algorithm using Pollard's rho method [Pol75]. On the other hand, the quantum BHT algorithm [BHT98] requires fewer queries to the random function, but the product of its time and space complexities is higher than that of the classical attack! In this paper, we address the more general problem of finding *multiple* collision pairs in a random function. This task plays a central role in low-memory meet-in-the-middle attacks [OW99, Din20]. It has applications in many problems, such as double and triple encryption [OW99], subset sum [DDKS12, DEM19], k-sum [Wag02], 3-collision [JL09], etc. Recently, it has also been used to attack the post-quantum cryptography candidates NTRU [Vre16] and SIKE [ACC⁺18]. The celebrated classical Parallel Collision Search algorithm of van Oorschot and Wiener [OW99] can find as many collision pairs as desired in a time that depends on the available memory. The question of whether this algorithm achieves the optimal classical time-space tradeoff has been settled positively only recently by Chakrabarti and Chen [CC17] (for the case of 2-to-1 random functions) and Dinur [Din20] (for the case of uniformly random functions). In the quantum setting, no time-space tradeoff was known prior to our work. In other words, it could have been the case that a memoryless quantum attack outperforms the Parallel Collision Search algorithm with unlimited memory capacity.

We point out that time-space tradeoffs have been studied for a long time in the complexity community [BFK⁺81, Bea91, BFMadH⁺87, Yao94, BSSV03, Abr90, MNT93]. The few results known in the quantum circuit model are for the Sorting problem [KŠW07], Boolean Matrix-Vector and Matrix-Matrix Multiplication [KŠW07], and Evaluating Solutions to Systems of Linear Inequalities [AŠW09]. Apart from our work, all existing quantum tradeoffs are based on the hardness of Quantum Search. We use the machinery developed in our paper to give a simpler proof of the tradeoffs obtained in [KŠW07].

1.1 Our results

The *Collision Pairs Finding* problem asks to find a certain number K of disjoint collision pairs in a random function $f : [M] \to [N]$ where $M \geq N$. A *collision pair* (or simply *collision*) is a pair of values $x_1 \neq x_2$ such that $f(x_1) = f(x_2)$. Two collisions (x_1, x_2) and (x_3, x_4) are *disjoint* if $x_1, \ldots, x_4$ are all different. We measure the time T of an algorithm solving this problem as the number of query accesses to f, and the space S as the amount of memory used. We assume that the output is produced in an online fashion, meaning that a collision can be output as soon as it is discovered. The length of the output is not counted toward the space bound and the same collision may be output several times (but it contributes only once to the total count). The requirement for the collisions to be disjoint is made to simplify our proofs later on. We note that a random function $f : [N] \to [N]$ contains $(1 - 2/e)N$ disjoint collisions on average [FO89].

Classically, the single-processor Parallel Collision Search algorithm [OW99] achieves an optimal [Din20] time-space tradeoff of[1] $T^2 S = \widetilde{\Theta}(K^2 N)$ for any amount of space S between $\widetilde{\Omega}(\log N)$ and $\widetilde{O}(K)$. In the quantum setting, the BHT algorithm [BHT98] can find a single collision in time $T = \widetilde{O}(N^{1/3})$ and space $S = \widetilde{O}(N^{1/3})$. In Algorithm 2, we adapt it for finding an arbitrary number K of collisions at cost $T^2 S \leq \widetilde{O}(K^2 N)$. For the sake of simplicity in the analysis, we do not require these collisions to be disjoint. This is the same tradeoff as classically, except that the space parameter S can hold larger values up to $\widetilde{O}(K^{2/3} N^{1/3})$, hence the existence of a quantum speed-up when there is no memory constraint.

Proposition 6.7 (Restated)**.** *For any $1 \leq K \leq O(N)$ and $\widetilde{\Omega}(\log N) \leq S \leq \widetilde{O}(K^{2/3} N^{1/3})$, there exists a bounded-error quantum algorithm that can find K collisions in a random function $f : [N] \to [N]$ by making $T = \widetilde{O}(K\sqrt{N/S})$ queries and using S qubits of memory.*

The BHT algorithm achieves the optimal time complexity for finding one collision [AS04, Zha15]. Our first main result is to provide a similar lower bound for the problem of finding K

[1] The notation $\widetilde{\ }$ is used to denote the presence of hidden polynomial factors in $\log(N)$ or $1/\log(N)$.

	Classical complexity	Quantum complexity
Upper bound:	$T^2S \leq \widetilde{O}(K^2N)$ *when* $\widetilde{\Omega}(\log N) \leq S \leq \widetilde{O}(K)$ Parallel Collision Search [OW99]	$T^2S \leq \widetilde{O}(K^2N)$ *when* $\widetilde{\Omega}(\log N) \leq S \leq \widetilde{O}(K^{2/3}N^{1/3})$ Proposition 6.7
Lower bound:	$T^2S \geq \Omega(K^2N)$ [Din20]	$T^3S \geq \Omega(K^3N)$ Theorem 6.1

Table 1: Complexity to find K disjoint collisions in a random function $f : [M] \to [N]$.

disjoint collisions. We prove that the optimal time complexity must satisfy $T \geq \Omega(K^{2/3}N^{1/3})$. This bound is matched by Proposition 6.7 when $S = \Theta(K^{2/3}N^{1/3})$. More precisely, we show that the optimal success probability decreases at an exponential rate in K below this bound. This property is of crucial importance for proving our time-space tradeoff next. We note that, similarly to [Zha15], the bound is independent of the size M of the domain as long as $M \geq N$.

Theorem 4.6 (Restated)**.** *The success probability of finding K disjoint collisions in a random function $f : [M] \to [N]$ is at most $O(T^3/(K^2N))^{K/2} + 2^{-K}$ for any algorithm making T quantum queries to f and any $1 \leq K \leq N/8$.*

Our second main result is the next time-space tradeoff for the same problem of finding K collisions in a random function. We summarize the tradeoffs known for this problem in Table 1. We note that the tradeoff $T^2S \geq \Omega(K^2N)$ is always stronger than $T^3S \geq \Omega(K^3N)$ since $T \geq K$.

Theorem 6.1 (Restated)**.** *Any quantum algorithm for finding K disjoint collisions in a random function $f : [M] \to [N]$ with success probability $2/3$ must satisfy a time-space tradeoff of $T^3S \geq \Omega(K^3N)$, where $1 \leq K \leq N/8$.*

As a simple corollary, we obtain that finding almost all collisions using a memoryless algorithm (i.e. $S = O(\log N)$) must require to perform at least $T \geq \Omega(N^{4/3})$ quantum queries, whereas $T = N$ classical queries are clearly sufficient when there is no space restriction. We further show that any improvement to this lower bound would imply a breakthrough for the *Element Distinctness* problem, which consists of finding a single collision in a random function $f : [N] \to [N^2]$ (or, more generally, deciding if a function contains a collision). It is a long-standing open question to prove a time-space lower bound for this problem. Although there is some progress in the classical case [BFMadH+87, Yao94, BSSV03], no result is known in the quantum setting. We give a reduction that converts any tradeoff for finding multiple collisions into a tradeoff for Element Distinctness. We state a particular case of our reduction below.

Corollary 6.5 (Restated)**.** *Suppose that there exists $\epsilon \in (0, 1)$ such that any quantum algorithm for finding $\widetilde{\Omega}(N)$ disjoint collisions in a random function $f : [10N] \to [N]$ must satisfy a time-space tradeoff of $TS^{1/3} \geq \widetilde{\Omega}(N^{4/3+\epsilon})$. Then, any quantum algorithm for solving Element Distinctness on domain size N must satisfy a time-space tradeoff of $TS^{1/3} \geq \widetilde{\Omega}(N^{2/3+2\epsilon})$.*

We point out that $TS^{1/3} \geq \Omega(N^{2/3})$ can already be deduced from the query complexity of Element Distinctness [AS04] and $S \geq 1$. We conjecture that our current tradeoff for finding K collisions can be improved to $T^2S \geq \Omega(K^2N)$, which would imply $T^2S \geq \widetilde{\Omega}(N^2)$ for Element Distinctness (Corollary 6.6). This result would be optimal [Amb07].

Finally, we adapt the machinery developed in our paper to study the K-Search problem, which consists of finding K preimages of 1 in a random function $f : [M] \to \{0, 1\}$ where $f(x) = 1$ with probability K/N independently for each $x \in [M]$. Several variants of this problem have been considered in the literature before [KŠW07, Amb10, Špa08], where it was shown that the success probability must be exponentially small in K when the number of quantum queries is smaller than $O(\sqrt{KN})$. Our proof is the first one to consider this particular input distribution, and it is arguably simpler and more intuitive than previous work.

Theorem 5.1 (Restated)**.** *The success probability of finding $K \leq N/8$ preimages of 1 in a random function $f : [M] \to \{0, 1\}$ where $f(x) = 1$ with probability K/N for each $x \in [M]$ is at most $O(T^2/(KN))^{K/2} + 2^{-K}$ for any algorithm using T quantum queries to f.*

As an application, we reprove in Theorem 6.8 the quantum time-space tradeoff for sorting N numbers first obtained in [KŠW07].

Theorem 6.8 (Restated)**.** *Any quantum algorithm for sorting a function $f : [N] \to \{0, 1, 2\}$ with success probability 2/3 must satisfy a time-space tradeoff of $T^2S \geq \Omega(N^3)$.*

1.2 Our techniques

Recording Query Technique. We use the recording query framework of Zhandry [Zha19] to upper bound the success probability of a query-bounded algorithm in finding K collision pairs. This method intends to reproduce the classical strategy where the queries made by an algorithm (the *attacker*) are recorded and answered with on-the-fly simulation of the oracle. Zhandry brought this technique to the quantum random oracle model by showing that, for the uniform input distribution, one can record *in superposition* the queries made by a quantum algorithm. Our first technical contribution (Section 3) is to simplify the analysis of Zhandry's technique and, as a byproduct, to generalize it to any product distribution on the input. We notice that there has been other independent work on extending Zhandry's recording technique [HI19, CMSZ19, CMS19]. Our approach does not require moving to the Fourier domain (as in [CMSZ19] for instance). It is based on defining a "recording query operator" that is specific to the input distribution under consideration. This operator can replace the standard quantum query operator without changing the success probability of the algorithm, but with the effect of "recording" the quantum queries in an additional register. We detail two recording query operators corresponding to the uniform distribution (Lemma 4.2) and to the product of Bernoulli distributions (Lemma 5.3).

Finding collisions with time-bounded algorithms. Our application of the recording technique to the Collision Pairs Finding problem has two stages. We first bound the probability that the algorithm has forced the recording of many collisions after T queries. Namely, we show that the norm of the quantum state that records a new collision at the t-th query is on the order of $\sqrt{t/N}$ (Proposition 4.4). This is related to the probability that a new random value collides with one of the at most t previously recorded queries. The reason why the collisions have to be disjoint is to avoid the recording of more than one new collision in one query. By solving a simple recurrence relation, one gets that the amplitude of the basis states that have recorded at least $K/2$ collisions after T queries is at most $O(T^{3/2}/(K\sqrt{N}))^{K/2}$. We note that Liu and Zhandry [LZ19, Theorem 5] carried out a similar analysis for the multi-collision finding problem, where they obtained a similar bound of $O(T^{3/2}/\sqrt{N})^{K/2}$. The second stage of our proof relates the probability of having recorded many collisions to the actual success probability of the algorithm. If we used previous approaches (notably [Zha19, Lemma 5]), this step would degrade the upper bound on the success probability by adding a term that is polynomial in K/N. We preserve the exponentially small dependence on K by doing a more careful analysis of the relation between the recording and the standard query models (Proposition 4.5). We adopt a similar approach for analyzing the K-Search problem in Section 5.

Finding collisions with time-space bounded algorithms. We convert the above time-only bound into a time-space tradeoff by using the time-segmentation method [BFK+81, KŠW07]. Given a quantum circuit that solves the Collision Pairs Finding problem in time T and space S, we slice it into $T/(S^{2/3}N^{1/3})$ consecutive subcircuits, each of them using $S^{2/3}N^{1/3}$ queries. If no slice can output more than $\Omega(S)$ collisions with high probability then there must be at least $\Omega(K/S)$ slices in total, thus proving the desired tradeoff. Our previous lower bound implies that it is impossible to find $\Omega(S)$ collisions with probability larger than 4^{-S} in time $S^{2/3}N^{1/3}$. We must take into account that the initial memory at the beginning of each slice carries out information from previous stages. As in previous work [Aar05, KŠW07], we can "eliminate" this memory by replacing it with the completely mixed state while decreasing the success probability by a factor of 2^{-S}. Thus, if a slice outputs $\Omega(S)$ collisions then it can be used to contradict the lower bound proved before.

Element Distinctness. We connect the Collision Pairs Finding and Element Distinctness problems by showing how to transform a low-space algorithm for the latter into one for the former (Proposition 6.3). If there exists a time-$\bar{T}$ space-$\bar{S}$ algorithm for Element Distinctness on domain size $\sqrt{N}$ then we can find $\widetilde{\Omega}(N)$ collisions in a random function $f : [N] \to [N]$ by repeatedly sampling a subset $H \subset [N]$ of size $\sqrt{N}$ and using that algorithm on the function f restricted to the domain H. Among other things, we must ensure that the same collision does not occur many times and that storing the set H does not use too much memory (it turns out that 4-wise independence is sufficient for our purpose). We end up with an algorithm with time $T = O(N\bar{T})$ and space $S = O(\bar{S})$ for finding $\widetilde{\Omega}(N)$ collisions. Consequently, if the Element Distinctness problem on domain size $\sqrt{N}$ can be solved with a time-space tradeoff of $\bar{T}\bar{S}^{1/3} \leq O(N^{1/3+\epsilon})$, then there is an algorithm for finding $\widetilde{\Omega}(N)$ collisions that satisfies a time-space tradeoff of $TS^{1/3} \leq O(N^{4/3+\epsilon})$.

2 Models of computation

We first present the standard model of quantum query complexity in Section 2.1. This model is used for investigating the *time complexity* of the Collision Pairs Finding problem in Section 4, and of the K-Search problem in Section 5. Then, we describe the more general circuit model that also captures the *space complexity* in Section 2.2. It is used in Section 6 for studying time-space tradeoffs.

2.1 Query model

The (standard) model of quantum query complexity [BW02] measures the number of quantum queries an algorithm (also called an "attacker") needs to make on an input function $f : [M] \to [N]$ to find an output z satisfying some predetermined relation $\mathrm{R}(f, z)$. We present this model in more detail below.

Quantum Query Algorithm. A T-query quantum algorithm is specified by a sequence $U_0, \dots, U_T$ of unitary transformations acting on the algorithm's memory. The state $|\psi\rangle$ of the algorithm is made of three registers $\mathcal{Q}$, $\mathcal{P}$, $\mathcal{W}$ where the *query register* $\mathcal{Q}$ holds $x \in [M]$, the *phase register* $\mathcal{P}$ holds $p \in [N]$ and the *working register* $\mathcal{W}$ holds some value w. We represent a basis state in the corresponding Hilbert space as $|x, p, w\rangle_{\mathcal{QPW}}$. We may drop the subscript $\mathcal{QPW}$ when it is clear from the context. The state $|\psi_t^f\rangle$ of the algorithm after $t \leq T$ queries to some input function $f : [M] \to [N]$ is

$$|\psi_t^f\rangle = U_t \mathcal{O}_f U_{t-1} \cdots U_1 \mathcal{O}_f U_0 |0\rangle$$

where the oracle $\mathcal{O}_f$ is defined by

$$\mathcal{O}_f |x, p, w\rangle = \omega_N^{pf(x)} |x, p, w\rangle \quad \text{and} \quad \omega_N = e^{\frac{2\mathrm{i}\pi}{N}}.$$

The *output* of the algorithm is written on a substring z of the value w. The *success probability* σ_f of the quantum algorithm on f is the probability that the output value z obtained by measuring the working register of $|\psi_T^f\rangle$ in the computational basis satisfies the relation $\mathrm{R}(f,z)$. In other words, if we let Π_{succ}^f be the projector whose support consists of all basis states $|x,p,w\rangle$ such that the output substring z of w satisfies $\mathrm{R}(f,z)$, then $\sigma_f = \left\|\Pi_{\text{succ}}^f|\psi_T^f\rangle\right\|^2$.

Oracle's Register. Here, we describe the variant used in the adversary method [Amb02] and in Zhandry's work [Zha19]. It is represented as an interaction between an *algorithm* that aims at finding a correct output z, and a superposition of *oracle*'s inputs that respond to the queries from the algorithm.

The memory of the oracle is made of an *input register* $\mathcal{F}$ holding the description of a function $f : [M] \to [N]$. This register is divided into M subregisters $\mathcal{F}_1, \ldots, \mathcal{F}_M$ where $\mathcal{F}_x$ holds $f(x) \in [N]$ for each $x \in [M]$. The basis states in the corresponding Hilbert space are $|f\rangle_{\mathcal{F}} = \otimes_{x\in[M]}|f(x)\rangle_{\mathcal{F}_x}$. Given an input distribution D on the set of functions $[N]^M$, the *oracle's initial state* is the state $|\mathrm{init}\rangle_{\mathcal{F}} = \sum_{f\in[N]^M} \sqrt{\Pr[f \leftarrow D]}|f\rangle$.

The *query operator* $\mathcal{O}$ is a unitary transformation acting on the memory of the algorithm and the oracle. Its action is defined on each basis state by

$$\mathcal{O}|x,p,w\rangle|f\rangle = (\mathcal{O}_f|x,p,w\rangle)|f\rangle.$$

The joint state $|\psi_t\rangle$ of the algorithm and the oracle after t queries is equal to $|\psi_t\rangle = U_t\mathcal{O}U_{t-1}\cdots U_1\mathcal{O}U_0(|0\rangle|\mathrm{init}\rangle) = \sum_{f\in[N]^M} \sqrt{\Pr[f \leftarrow D]}|\psi_t^f\rangle|f\rangle$, where the unitaries U_i have been extended to act as the identity on $\mathcal{F}$. The *success probability* σ of a quantum algorithm on an input distribution D is the probability that the output value z and the input f obtained by measuring the working and input registers of the final state $|\psi_T\rangle$ satisfy the relation $\mathrm{R}(f,z)$. In other words, if we let Π_{succ} be the projector whose support consists of all basis states $|x,p,w\rangle|f\rangle$ such that the output substring z of w satisfies $\mathrm{R}(f,z)$, then $\sigma = \|\Pi_{\text{succ}}|\psi_T\rangle\|^2$.

2.2 Space-bounded model

Our model of space-bounded computation is identical to the one described in [KŠW07, AŠW09]. We use the quantum circuit model augmented with the oracle gates of the query model defined in the previous section. The *time complexity*, denoted by T, is the number of gates in the circuit. In practice, we lower bound it by the number of oracle gates only. The *space complexity*, denoted by S, is the number of qubits on which the circuit is operating. The result of the computation is written on some dedicated output qubits that may not be used later on, and that are *not counted toward the space bound*. In particular, the size of the output can be larger than S. Furthermore, we assume that the output qubits are updated at some predefined output gates in the circuit.

We notice that, by the deferred measurement principle, any space-bounded computation that uses T queries can be transformed into a T-query unitary algorithm as defined in Section 2.1. Thus, any lower bound on the query complexity of a problem is also a lower bound on the time complexity of that problem in the space-bounded model. This explains our use of the query model in Section 4 and Section 5.

3 Recording model

The quantum recording query model is a modification of the standard query model defined in Section 2.1 that is unnoticeable by the algorithm, but that allows us to track more easily the progress made toward solving the problem under consideration. The original recording model was formulated by Zhandry in [Zha19]. Here, we propose a simplified and more general version of this framework that only requires the initial oracle's state $|\mathrm{init}\rangle_{\mathcal{F}}$ to be a product state $\otimes_{x\in[M]}|\mathrm{init}_x\rangle_{\mathcal{F}_x}$ (instead of the uniform distribution over all basis states as in [Zha19]).

Construction. The range $[N]$ is augmented with a new symbol $\perp$. The input register $\mathcal{F}$ of the oracle can now contain $f : [M] \to [N] \cup \{\perp\}$, where $f(x) = \perp$ represents the absence of knowledge from the algorithm about the image of x. Unlike in the standard query model, the oracle's initial state is independent of the input distribution and is fixed to be $|\perp^M\rangle_{\mathcal{F}}$ (which represents the fact that the algorithm knows nothing about the input initially). We extend the query operator $\mathcal{O}$ defined in the standard query model by setting

$$\mathcal{O}|x, p, w\rangle|f\rangle = |x, p, w\rangle|f\rangle \quad \text{when } f(x) = \perp.$$

Given a product input distribution $D = D_1 \otimes \cdots \otimes D_M$ on the set $[N]^M$, the oracle's initial state in the standard query model can be decomposed as the product state $|\text{init}\rangle_{\mathcal{F}} = \otimes_{x \in [M]} |\text{init}_x\rangle_{\mathcal{F}_x}$ where $|\text{init}_x\rangle_{\mathcal{F}_x} := \sum_{y \in [N]} \sqrt{\Pr[y \leftarrow D_x]}|y\rangle_{\mathcal{F}_x}$. The "recording query operator" $\mathcal{R}$ is defined with respect to a family $(\mathcal{S}_x)_{x \in [M]}$ of unitary operators satisfying $\mathcal{S}_x|\perp\rangle_{\mathcal{F}_x} = |\text{init}_x\rangle_{\mathcal{F}_x}$ for all x as follows.

Definition 3.1. Given M unitary operators $\mathcal{S}_1, \ldots, \mathcal{S}_M$ acting on $\mathcal{F}_1, \ldots, \mathcal{F}_M$ respectively, consider the operator $\mathcal{S}$ acting on all the registers $\mathcal{QPWF}$ such that,

$$\mathcal{S} = \sum_{x \in [M]} |x\rangle\langle x|_{\mathcal{Q}} \otimes I_{\mathcal{PWF}_1 \ldots \mathcal{F}_{x-1}} \otimes \mathcal{S}_x \otimes I_{\mathcal{F}_{x+1} \ldots \mathcal{F}_M}.$$

Then, the *recording query operator* $\mathcal{R}$ with respect to $(\mathcal{S}_x)_{x \in [M]}$ is defined as $\mathcal{R} = \mathcal{S}^\dagger \mathcal{O} \mathcal{S}$.

Later in this paper, we describe two recording query operators related to the uniform distribution (Lemma 4.2) and to the product of Bernoulli distributions (Lemma 5.3).

Indistinguishability. The joint state of the algorithm and the oracle after t queries in the recording query model is defined as $|\phi_t\rangle = U_t \mathcal{R} U_{t-1} \cdots U_1 \mathcal{R} U_0 \big(|0\rangle|\perp^M\rangle\big)$. Notice that the query operator $\mathcal{R}$ can only change the value of $f(x')$ (contained in the register $\mathcal{F}_{x'}$) when it is applied to a state $|x, p, w\rangle|f\rangle$ such that $x = x'$. As a result, we have the following simple fact.

Fact 3.2. *The state $|\phi_t\rangle$ is a linear combination of basis states $|x, p, w\rangle|f\rangle$ where f contains at most t entries different from $\perp$.*

The entries of f that are different from $\perp$ represent what the oracle has learned (or "recorded") from the algorithm's queries so far. In the next theorem, we show that $|\phi_t\rangle$ is related to the state $|\psi_t\rangle$ (defined in Section 2.1) by $|\psi_t\rangle = \big(I_{\mathcal{QPW}} \otimes_{x \in [M]} \mathcal{S}_x\big)|\phi_t\rangle$. In particular, the states $|\psi_t\rangle$ and $|\phi_t\rangle$ cannot be distinguished by the algorithm since the reduced states on the algorithm's registers are identical.

Theorem 3.3. *Let $(U_0, \ldots, U_T)$ be a T-query quantum algorithm. Given M unitary operators $\mathcal{S}_1, \ldots, \mathcal{S}_M$ acting on the oracle's registers $\mathcal{F}_1, \ldots, \mathcal{F}_M$ respectively, let $\mathcal{R}$ denote the recording query operator associated with $(\mathcal{S}_x)_{x \in [M]}$, and define the initial state $|\text{init}\rangle_{\mathcal{F}} = \big(\otimes_{x \in [M]} \mathcal{S}_x\big)|\perp^M\rangle$. Then, the states*

$$\begin{cases} |\psi_t\rangle = U_t \mathcal{O} U_{t-1} \cdots U_1 \mathcal{O} U_0 \big(|0\rangle|\text{init}\rangle\big) \\ |\phi_t\rangle = U_t \mathcal{R} U_{t-1} \cdots U_1 \mathcal{R} U_0 \big(|0\rangle|\perp^M\rangle\big) \end{cases}$$

after $t \leq T$ queries in the standard and recording query models respectively satisfy

$$|\psi_t\rangle = \mathcal{T}|\phi_t\rangle \quad \textit{where} \quad \mathcal{T} = I_{\mathcal{QPW}} \bigotimes_{x \in [M]} \mathcal{S}_x.$$

Proof. We start by introducing the intermediate operator $\bar{\mathcal{R}} = \mathcal{T}^\dagger \mathcal{O} \mathcal{T}$. Observe that for any basis state $|x, p, w\rangle|f\rangle$ the operators $\bar{\mathcal{R}}$ and $\mathcal{R}$ act the same way on the registers $\mathcal{QPF}_x$ and they do not depend on the other registers. Thus, we have $\bar{\mathcal{R}} = \mathcal{R}$. We further observe that U_i and $\mathcal{T}$ commute for all i since they depend on disjoint registers. Consequently, we have that

$$\begin{aligned}
|\psi_t\rangle &= U_t \mathcal{O} U_{t-1} \mathcal{O} \cdots U_1 \mathcal{O} U_0 \cdot \mathcal{T}\big(|0\rangle|\perp^M\rangle\big) && \text{since } \mathcal{T}\big(|0\rangle|\perp^M\rangle\big) = |0\rangle|\mathrm{init}\rangle \\
&= \mathcal{T}\mathcal{T}^\dagger U_t \mathcal{O} \cdot \mathcal{T}\mathcal{T}^\dagger U_{t-1} \mathcal{O} \cdots \mathcal{T}\mathcal{T}^\dagger U_1 \mathcal{O} \cdot \mathcal{T}\mathcal{T}^\dagger U_0 \cdot \mathcal{T}\big(|0\rangle|\perp^M\rangle\big) && \text{since } \mathcal{T}\mathcal{T}^\dagger = I \\
&= \mathcal{T} U_t \mathcal{T}^\dagger \cdot \mathcal{O} \cdot \mathcal{T} U_{t-1} \mathcal{T}^\dagger \cdot \mathcal{O} \cdots \mathcal{T} U_1 \mathcal{T}^\dagger \cdot \mathcal{O} \cdot \mathcal{T} U_0 \big(|0\rangle|\perp^M\rangle\big) && \text{by commutation} \\
&= \mathcal{T} U_t \bar{\mathcal{R}} U_{t-1} \cdots U_1 \bar{\mathcal{R}} U_0 \big(|0\rangle|\perp^M\rangle\big) && \text{by definition of } \bar{\mathcal{R}} \\
&= \mathcal{T} U_t \mathcal{R} U_{t-1} \cdots U_1 \mathcal{R} U_0 \big(|0\rangle|\perp^M\rangle\big) && \text{since } \bar{\mathcal{R}} = \mathcal{R} \\
&= \mathcal{T}|\phi_t\rangle && \text{by definition of } |\phi_t\rangle.
\end{aligned}$$

□

4 Time lower bound for Collision Pairs Finding

In this section, we upper bound the success probability of finding K disjoint collisions in the query-bounded model of Section 2.1. The proof uses the recording query model of Section 3. We first describe in Section 4.1 the recording query framework associated with this problem. In Section 4.2, we study the probability that an algorithm has recorded at least $k \leq K$ collisions after $t \leq T$ queries. We prove by induction on t and k that this quantity is exponentially small in k when $t \leq O(k^{2/3}N^{1/3})$ (Proposition 4.4). Finally, in Section 4.3, we relate this progress measure to the actual success probability (Proposition 4.5), and we conclude that the latter quantity is exponentially small in K after $T \leq O(K^{2/3}N^{1/3})$ queries (Theorem 4.6).

4.1 Recording query operator

We describe a recording operator that corresponds to the uniform distribution on the set of functions $f : [M] \to [N]$. In the standard query model, the oracle's initial state is $|\mathrm{init}\rangle_{\mathcal{F}} = \otimes_{x \in [M]}\big(\frac{1}{\sqrt{N}} \sum_{y \in [N]} |y\rangle_{\mathcal{F}_x}\big)$. Consequently, in the recording query model, we choose the unitary transformations $\mathcal{S}_1, \cdots, \mathcal{S}_M$ to be defined as follows.

Definition 4.1. For any $x \in [M]$, we define the unitary $\mathcal{S}_x$ acting on the register $\mathcal{F}_x$ to be

$$\mathcal{S}_x : \begin{cases} |\perp\rangle_{\mathcal{F}_x} & \longmapsto \frac{1}{\sqrt{N}} \sum_{y \in [N]} |y\rangle_{\mathcal{F}_x} \\ \frac{1}{\sqrt{N}} \sum_{y \in [N]} |y\rangle_{\mathcal{F}_x} & \longmapsto |\perp\rangle_{\mathcal{F}_x} \\ \frac{1}{\sqrt{N}} \sum_{y \in [N]} \omega_N^{py} |y\rangle_{\mathcal{F}_x} & \longmapsto \frac{1}{\sqrt{N}} \sum_{y \in [N]} \omega_N^{py} |y\rangle_{\mathcal{F}_x} \quad \text{for } p = 1, \ldots, N-1. \end{cases}$$

These unitaries verify $\mathcal{T}|\perp^M\rangle = |\mathrm{init}\rangle$ where $\mathcal{T} = \otimes_{x \in [M]} \mathcal{S}_x$, as required by Theorem 3.3. The recording query operator is $\mathcal{R} = \mathcal{S}\mathcal{O}\mathcal{S}$ (Definition 3.1) since $\mathcal{S}^\dagger = \mathcal{S}$. The next lemma gives an explicit characterization of the action of $\mathcal{R}$ on a basis state.

Lemma 4.2. *If the recording query operator $\mathcal{R}$ associated with Definition 4.1 is applied to a basis state $|x, p, w\rangle|f\rangle$ where $p \neq 0$ then the register $|f(x)\rangle_{\mathcal{F}_x}$ is mapped to*

$$\begin{cases} \sum_{y \in [N]} \frac{\omega_N^{py}}{\sqrt{N}} |y\rangle & \text{if } f(x) = \perp \\ \frac{\omega_N^{pf(x)}}{N} |\perp\rangle + \frac{1 + \omega_N^{pf(x)}(N-2)}{N} |f(x)\rangle + \sum_{y \in [N] \setminus \{f(x)\}} \frac{1 - \omega_N^{py} - \omega_N^{pf(x)}}{N} |y\rangle & \text{otherwise} \end{cases}$$

and the other registers are unchanged. If $p = 0$ then none of the registers are changed.

Proof. By definition, the unitary $\mathcal{S}_x$ maps $|\bot\rangle_{\mathcal{F}_x} \mapsto \frac{1}{\sqrt{N}}\sum_{y\in[N]}|y\rangle$ and $|y\rangle_{\mathcal{F}_x} \mapsto \frac{1}{\sqrt{N}}|\bot\rangle + \frac{1}{\sqrt{N}}\sum_{p'\in[N]\setminus\{0\}}\omega_N^{-p'y}|\widehat{p'}\rangle$ where $y \in [N]$ and $|\widehat{p'}\rangle := \frac{1}{\sqrt{N}}\sum_{y\in[N]}\omega_N^{p'y}|y\rangle$. Thus, the action on the register $\mathcal{F}_x$ is:

- If $f(x) = \bot$ then $|f(x)\rangle_{\mathcal{F}_x} \overset{\mathcal{S}}{\mapsto} \frac{1}{\sqrt{N}}\sum\limits_{y\in[N]}|y\rangle \overset{\mathcal{O}}{\mapsto} \frac{1}{\sqrt{N}}\sum\limits_{y\in[N]}\omega_N^{py}|y\rangle \overset{\mathcal{S}}{\mapsto} \frac{1}{\sqrt{N}}\sum\limits_{y\in[N]}\omega_N^{py}|y\rangle$.
- If $f(x) \in [N]$ then $|f(x)\rangle_{\mathcal{F}_x} = \frac{1}{\sqrt{N}}\sum\limits_{p'\in[N]}\omega_N^{-p'f(x)}|\widehat{p'}\rangle \overset{\mathcal{S}}{\mapsto} \frac{1}{\sqrt{N}}|\bot\rangle + \frac{1}{\sqrt{N}}\sum\limits_{p'\in[N]\setminus\{0\}}\omega_N^{-p'f(x)}|\widehat{p'}\rangle \overset{\mathcal{O}}{\mapsto} \frac{1}{\sqrt{N}}|\bot\rangle + \frac{1}{\sqrt{N}}\sum\limits_{p'\in[N]\setminus\{0\}}\omega_N^{-p'f(x)}|\widehat{p+p'}\rangle = \frac{1}{\sqrt{N}}|\bot\rangle + \frac{\omega_N^{pf(x)}}{\sqrt{N}}\sum\limits_{p'\in[N]\setminus\{p\}}\omega_N^{-p'f(x)}|\widehat{p'}\rangle \overset{\mathcal{S}}{\mapsto} \frac{1}{N}\sum\limits_{y\in[N]}|y\rangle + \frac{\omega_N^{pf(x)}}{\sqrt{N}}|\bot\rangle + \frac{\omega_N^{pf(x)}}{\sqrt{N}}\sum\limits_{p'\in[N]\setminus\{0,p\}}\omega_N^{-p'f(x)}|\widehat{p'}\rangle = \frac{\omega_N^{pf(x)}}{N}|\bot\rangle + \frac{1+\omega_N^{pf(x)}(N-2)}{N}|f(x)\rangle + \sum\limits_{y\in[N]\setminus\{f(x)\}}\frac{1-\omega_N^{py}-\omega_N^{pf(x)}}{N}|y\rangle$.

$\square$

We note that the recording operator $\mathcal{R}$ is close to the mapping $|\bot\rangle_{\mathcal{F}_x} \mapsto \sum_{y\in[N]}\frac{\omega_N^{py}}{\sqrt{N}}|y\rangle$ and $|f(x)\rangle_{\mathcal{F}_x} \mapsto \omega_N^{pf(x)}|f(x)\rangle$ (if $f(x) \neq \bot$) up to lower-order terms of amplitude $O(1/N)$. This is analogous to a "lazy" classical oracle that would choose the value of $f(x)$ uniformly at random the first time it is queried.

4.2 Analysis of the recording progress

We define a measure of progress based on the number of disjoint collisions contained in the oracle's register of the recording query model. We first give some projectors related to this quantity.

Definition 4.3. We define the following projectors by giving the basis states on which they project:

- $\Pi_{\leq k}$, $\Pi_{=k}$ and $\Pi_{\geq k}$: all basis states $|x, p, w\rangle|f\rangle$ such that f contains respectively *at most*, *exactly* or *at least* k disjoint collisions (the entries with $\bot$ are not considered as collisions).
- $\Pi_{=k,\bot}$ and $\Pi_{=k,y}$ for $y \in [N]$: all basis states $|x, p, w\rangle|f\rangle$ such that (1) f contains *exactly* k disjoint collisions, (2) the phase multiplier p is nonzero and (3) $f(x) = \bot$ or $f(x) = y$ respectively.

We can now define the measure of progress $q_{t,k}$ for t queries and k collisions as

$$q_{t,k} = \|\Pi_{\geq k}|\phi_t\rangle\|$$

where $|\phi_t\rangle$ is the state after t queries in the recording query model. The main result of this section is the following bound on the growth of $q_{t,k}$.

Proposition 4.4. *For all t and k, we have that $q_{t,k} \leq \binom{t}{k}\left(\frac{4\sqrt{t}}{\sqrt{N}}\right)^k$.*

Proof. First, $q_{0,0} = 1$ and $q_{0,k} = 0$ for all $k \geq 1$ since the initial state is $|\phi_0\rangle = |0\rangle|\bot^M\rangle$. Then, we prove that $q_{t,k}$ satisfies the following recurrence relation

$$q_{t+1,k+1} \leq q_{t,k+1} + 4\sqrt{\frac{t}{N}}q_{t,k}. \tag{1}$$

From this result, it is trivial to conclude that $q_{t,k} \le \binom{t}{k}\left(\frac{4\sqrt{t}}{\sqrt{N}}\right)^k$. In order to prove Equation (1), we first observe that $q_{t+1,k+1} = \|\Pi_{\ge k+1} U_{t+1}\mathcal{R}|\phi_t\rangle\| = \|\Pi_{\ge k+1}\mathcal{R}|\phi_t\rangle\|$ since the unitary U_{t+1} applied by the algorithm at time $t+1$ does not modify the oracle's memory. Then, on any basis state $|x,p,w\rangle|f\rangle$, the recording query operator $\mathcal{R}$ acts as the identity on the registers $\mathcal{F}_{x'}$ for $x' \neq x$. Consequently, the basis states $|x,p,w\rangle|f\rangle$ in $|\phi_t\rangle$ that may contribute to $q_{t+1,k+1}$ must either already contain $k+1$ disjoint collisions in f, or exactly k disjoint collisions in f and $p \neq 0$. This implies that

$$q_{t+1,k+1} \le q_{t,k+1} + \|\Pi_{\ge k+1}\mathcal{R}\Pi_{=k,\perp}|\phi_t\rangle\| + \sum_{y\in[N]} \|\Pi_{\ge k+1}\mathcal{R}\Pi_{=k,y}|\phi_t\rangle\|.$$

We first bound the term $\|\Pi_{\ge k+1}\mathcal{R}\Pi_{=k,\perp}|\phi_t\rangle\|$. Consider any basis state $|x,p,w\rangle|f\rangle$ in the support of $\Pi_{=k,\perp}$ and $|\phi_t\rangle$. The function f must contain at most t entries different from $\perp$ by Fact 3.2. By Lemma 4.2, we have $\mathcal{R}|x,p,w\rangle|f\rangle = \sum_{y\in[N]} \frac{\omega_N^{py}}{\sqrt{N}}|x,p,w\rangle|y\rangle_{\mathcal{F}_x} \otimes_{x'\neq x} |f(x')\rangle_{\mathcal{F}_{x'}}$. Since there are at most t entries in f that can collide with the value contained in the register $\mathcal{F}_x$, we have $\|\Pi_{\ge k+1}\mathcal{R}|x,p,w\rangle|f\rangle\| \le \sqrt{t/N}$. Finally, since any two basis states in the support of $\Pi_{=k,\perp}$ remain orthogonal after $\Pi_{\ge k+1}\mathcal{R}$ is applied, we obtain that $\|\Pi_{\ge k+1}\mathcal{R}\Pi_{=k,\perp}|\phi_t\rangle\| \le \sqrt{t/N}\|\Pi_{=k,\perp}|\phi_t\rangle\| \le \sqrt{t/N}q_{t,k}$.

We now consider the term $\|\Pi_{\ge k+1}\mathcal{R}\Pi_{=k,y}|\phi_t\rangle\|$ for any $y \in [N]$. Again, we consider any basis state $|x,p,w\rangle|f\rangle$ in the support of $\Pi_{=k,y}$ where f has at most t entries different from $\perp$. According to Lemma 4.2, we have $\mathcal{R}|x,p,w\rangle|f\rangle = \frac{\omega_N^{pf(x)}}{N}|\perp\rangle + \frac{1+\omega_N^{pf(x)}(N-2)}{N}|f(x)\rangle + \sum_{y'\neq f(x)} \frac{1-\omega_N^{py'}-\omega_N^{pf(x)}}{N}|x,p,w\rangle|y'\rangle_{\mathcal{F}_x} \otimes_{x'\neq x}|f(x')\rangle_{\mathcal{F}_{x'}}$. As before, there are at most t terms in this sum that can be in the support of $\Pi_{\ge k+1}$. Consequently, $\|\Pi_{\ge k+1}\mathcal{R}|x,p,w\rangle|f\rangle\| \le 3\sqrt{t}/N$ and $\|\Pi_{\ge k+1}\mathcal{R}\Pi_{=k,y}|\phi_t\rangle\| \le 3\sqrt{t}/N\|\Pi_{=k,y}|\phi_t\rangle\|$.

We conclude that $q_{t+1,k+1} \le q_{t,k+1} + \sqrt{t/N}q_{t,k} + \sum_{y\in[N]} 3\sqrt{t}/N\|\Pi_{=k,y}|\phi_t\rangle\| \le q_{t,k+1} + \sqrt{t/N}q_{t,k} + 3\sqrt{t/N}\sqrt{\sum_{y\in[N]}\|\Pi_{=k,y}|\phi_t\rangle\|^2} \le q_{t,k+1} + \sqrt{t/N}q_{t,k} + 3\sqrt{t/N}q_{t,k}$, where the second step is by Cauchy-Schwarz' inequality. □

4.3 From the recording progress to the success probability

We connect the success probability $\sigma = \|\Pi_{\mathrm{succ}}|\psi_T\rangle\|^2$ in the standard query model to the final progress $q_{T,k}$ in the recording query model after T queries. We show that if the algorithm has made no significant progress for recording $k \ge K/2$ collisions then it needs to "guess" the positions of $K-k$ other collisions. Classically, the probability to find the values of $K-k$ collisions that have not been queried would be at most $(1/N^2)^{K-k}$. Here, we show similarly that if a unit state contains at most k collisions in the quantum recording model, then after mapping it to the standard query model (by applying the operator $\mathcal{T}$ of Theorem 3.3) the probability that the output register contains the correct positions of K collisions is at most $N^2(4K^2/N^2)^{K-k}$.

Proposition 4.5. *For any state $|\phi\rangle$, we have $\|\Pi_{\mathrm{succ}}\mathcal{T}\Pi_{\le k}|\phi\rangle\| \le N\left(\frac{2K}{N}\right)^{K-k}\|\Pi_{\le k}|\phi\rangle\|$.*

Proof. We assume that the output of the algorithm also contains the image of each collision pair under f. Namely, the output z is represented as a list of K triples $(x_1,x_2,y_1),\ldots,(x_{2K-1},x_{2K},y_K) \in [M]^2 \times ([N]\cup\{\perp\})$. It is correct if the input function $f:[M]\to[N]$ (in the standard query model) satisfies $f(x_{2i-1}) = f(x_{2i}) = y_i \neq \perp$ for all $1 \le i \le K$, and the values $x_1,x_2,\ldots,x_{2K}$ are all different. By definition, the support of Π_{succ} consists of all basis states $|x,p,w\rangle|f\rangle$ such that the output substring z of w satisfies these conditions.

We define a new family of projectors $\tilde{\Pi}_{a,b}$, where $0 \le a+b \le 2K$, whose supports consist of all basis states $|x,p,w\rangle|f\rangle$ satisfying the following conditions:

(A) The output substring z is made of K triples $(x_1,x_2,y_1),\ldots,(x_{2K-1},x_{2K},y_K)$ where the x_i are all different.

(B) There are exactly a indices $i \in [2K]$ such that $f(x_i) = \perp$.

(C) There are exactly b indices $i \in [2K]$ such that $f(x_i) \neq \perp$ and $f(x_i) \neq y_{\lceil i/2 \rceil}$.

For any state $|x,p,w\rangle|f\rangle$ in the support of $\tilde{\Pi}_{a,b}$, we claim that

$$\|\Pi_{\mathrm{succ}} \mathcal{T} |x,p,w\rangle|f\rangle\| \leq \left(\frac{1}{\sqrt{N}}\right)^a \left(\frac{1}{N}\right)^b. \tag{2}$$

Indeed, we have $\mathcal{T} = \otimes_{x' \in [M]} \mathcal{S}_{x'}$ and by Definition 4.1 the action of $\mathcal{S}_{x_i}$ on the register $|f(x_i)\rangle_{\mathcal{F}_{x_i}}$ is $|f(x_i)\rangle \mapsto \frac{1}{\sqrt{N}} \sum_{y \in [N]} |y\rangle$ if $f(x_i) = \perp$, and $|f(x_i)\rangle \mapsto \frac{1}{\sqrt{N}}|\perp\rangle + (1 - \frac{1}{N})|f(x_i)\rangle - \frac{1}{N} \sum_{y \in [N] \setminus f(x_i)} |y\rangle$ otherwise. The projector Π_{succ} only keeps the term $|y_{\lceil i/2 \rceil}\rangle$ in these sums, which implies Equation (2).

Let us now consider any linear combination $|\varphi\rangle = \sum_{x,p,w,f} \alpha_{x,p,w,f} |x,p,w\rangle|f\rangle$ of basis states that are in the support of $\tilde{\Pi}_{a,b}$. We claim that

$$\|\Pi_{\mathrm{succ}} \mathcal{T} |\varphi\rangle\| \leq \left(\sqrt{\frac{2K}{N}}\right)^{a+b} \||\varphi\rangle\|. \tag{3}$$

First, given two basis states $|x,p,w\rangle|f\rangle$ and $|\bar{x},\bar{p},\bar{w}\rangle|\bar{f}\rangle$ where $z = ((x_1,x_2,y_1),\ldots,(x_{2K-1}, x_{2K},y_K))$ is the output substring of w, if the tuples $\big(x,p,w,(f(x'))_{x' \notin \{x_1,\ldots,x_{2K}\}}\big)$ and $\big(\bar{x},\bar{p},\bar{w}, (\bar{f}(x'))_{x' \notin \{x_1,\ldots,x_{2K}\}}\big)$ are different then $\Pi_{\mathrm{succ}} \mathcal{T} |x,p,w\rangle|f\rangle$ must be orthogonal to $\Pi_{\mathrm{succ}} \mathcal{T} |\bar{x},\bar{p},\bar{w}\rangle|\bar{f}\rangle$. Moreover, for any $z = ((x_1,x_2,y_1),\ldots,(x_{2K-1},x_{2K},y_K))$ that satisfies condition (A), there are $\binom{2K}{a}\binom{2K-a}{b}(N-1)^b \leq (2K)^{a+b} N^b$ different ways to choose $(f(x_i))_{i \in [2K]}$ that satisfy conditions (B) and (C). Let us write $w_{\vec{x}} = \{x_1,\ldots,x_{2K}\}$ when the output substring z of w contains $x_1,\ldots,x_{2K}$. Then, by using the Cauchy-Schwarz inequality and Equation (2), we get that

$$\begin{aligned}
\|\Pi_{\mathrm{succ}} \mathcal{T} |\varphi\rangle\|^2 &= \sum_{x,p,w,(f(x'))_{x' \notin w_{\vec{x}}}} \Bigg\| \sum_{(f(x'))_{x' \in w_{\vec{x}}}} \alpha_{x,p,w,f} \Pi_{\mathrm{succ}} \mathcal{T} |x,p,w\rangle|f\rangle \Bigg\|^2 \\
&\leq \sum_{x,p,w,(f(x'))_{x' \notin w_{\vec{x}}}} \Bigg(\sum_{(f(x'))_{x' \in w_{\vec{x}}}} |\alpha_{x,p,w,f}|^2 \Bigg) \Bigg(\sum_{(f(x'))_{x' \in w_{\vec{x}}}} \|\Pi_{\mathrm{succ}} \mathcal{T} |x,p,w\rangle|f\rangle\|^2 \Bigg) \\
&\leq \||\varphi\rangle\|^2 \cdot (2K)^{a+b} N^b \left(\frac{1}{N}\right)^a \left(\frac{1}{N^2}\right)^b \\
&= \left(\frac{2K}{N}\right)^{a+b} \||\varphi\rangle\|^2,
\end{aligned}$$

which proves Equation (3). Observe now that the support of $\Pi_{\leq k}$ is contained into the union of the supports of $\tilde{\Pi}_{a,b}$ for $a + b \geq 2(K-k)$. Thus, by the triangle inequality, $\|\Pi_{\mathrm{succ}} \mathcal{T} \Pi_{\leq k} |\phi\rangle\| \leq \sum_{a+b \geq 2(K-k)} \|\Pi_{\mathrm{succ}} \mathcal{T} \tilde{\Pi}_{a,b} \Pi_{\leq k} |\phi\rangle\|$. This is at most $\sum_{a+b \geq 2(K-k)} \Big(\sqrt{\frac{2K}{N}}\Big)^{a+b} \|\tilde{\Pi}_{a,b} \Pi_{\leq k} |\phi\rangle\|$ by Equation (3). Finally, by Cauchy-Schwarz' inequality and the fact that the supports of the projectors $\tilde{\Pi}_{a,b}$ are disjoint, we have $\|\Pi_{\mathrm{succ}} \mathcal{T} \Pi_{\leq k} |\phi\rangle\| \leq \sqrt{\sum_{a+b \geq 2(K-k)} \big(\frac{2K}{N}\big)^{a+b}} \sqrt{\sum_{a,b} \|\tilde{\Pi}_{a,b} \Pi_{\leq k} |\phi\rangle\|^2} \leq N \big(\frac{2K}{N}\big)^{K-k} \|\Pi_{\leq k} |\phi\rangle\|$. ☐

We can now conclude the proof of the main result of this section.

Theorem 4.6. *The success probability of finding K disjoint collisions in a random function $f : [M] \to [N]$ is at most $O(T^3/(K^2 N))^{K/2} + 2^{-K}$ for any algorithm making T quantum queries to f and any $1 \leq K \leq N/8$.*

Proof. Let $|\psi_T\rangle$ (resp. $|\phi_T\rangle$) denote the state of the algorithm after T queries in the standard (resp. recording) query model. We recall that $|\psi_T\rangle = \mathcal{T}|\phi_T\rangle$ (Theorem 3.3). Thus, by the triangle inequality, the success probability $\sigma = \|\Pi_{\text{succ}}|\psi_T\rangle\|^2$ satisfies $\sqrt{\sigma} \leq \|\Pi_{\text{succ}}\mathcal{T}\Pi_{\geq K/2}|\phi_T\rangle\| + \|\Pi_{\text{succ}}\mathcal{T}\Pi_{\leq K/2}|\phi_T\rangle\| \leq \|\Pi_{\geq K/2}|\phi_T\rangle\| + \|\Pi_{\text{succ}}\mathcal{T}\Pi_{\leq K/2}|\phi_T\rangle\|$. Using Propositions 4.4 and 4.5, we have that $\sqrt{\sigma} \leq \binom{T}{K/2}\left(4\sqrt{T/N}\right)^{K/2} + N(2K/N)^{K/2} \leq O(T^{3/2}/(K\sqrt{N}))^{K/2} + 2^{-K/2-1}$. Finally, the upper bound on σ is derived from the standard inequality $(u+v)^2 \leq 2u^2 + 2v^2$. □

5 Time lower bound for K-Search

In this section, we illustrate the use of the recording query model to upper bound the success probability of a query-bounded algorithm on a *non-uniform* input distribution.

Theorem 5.1. *The success probability of finding $K \leq N/8$ preimages of 1 in a random function $f : [M] \to \{0,1\}$ where $f(x) = 1$ with probability K/N for each $x \in [M]$ is at most $O(T^2/(KN))^{K/2} + 2^{-K}$ for any algorithm using T quantum queries to f.*

We show that, similarly to the classical setting where a query can reveal a 1 with probability K/N, the *amplitude* of the basis states that record a new 1 increases by a factor of $\sqrt{K/N}$ after each query (Proposition 5.5). Thus, the amplitude of the basis states that have recorded at least $K/2$ ones after T queries is at most $O(T/\sqrt{KN})^{K/2}$. This implies that any algorithm with $T < O(\sqrt{KN})$ queries is likely to output at least $K/2$ ones at positions that have not been recorded. These outputs can only be correct with probability $O(K/N)^{K/2}$ (Proposition 5.6).

5.1 Recording query operator

We describe a recording operator that encodes the distribution that gives $f : [M] \to [N]$ where $f(x) = 1$ with probability K/N independently for each $x \in [M]$. In the standard query model, the oracle's initial state is $|\text{init}\rangle = \otimes_{x\in[M]}\left(\sqrt{1-K/N}|0\rangle_{\mathcal{F}_x} + \sqrt{K/N}|1\rangle_{\mathcal{F}_x}\right)$ for this distribution. Consequently, we instantiate the recording query model as follows.

Definition 5.2. For any $x \in [M]$, define the unitary $\mathcal{S}_x$ acting on the register $\mathcal{F}_x$ to be

$$\mathcal{S}_x|\bot\rangle_{\mathcal{F}_x} = |+\rangle_{\mathcal{F}_x}, \qquad \mathcal{S}_x|+\rangle_{\mathcal{F}_x} = |\bot\rangle_{\mathcal{F}_x}, \qquad \mathcal{S}_x|-\rangle_{\mathcal{F}_x} = |-\rangle_{\mathcal{F}_x}$$

where $\alpha = \sqrt{1-K/N}$, $\beta = \sqrt{K/N}$ and $|+\rangle_{\mathcal{F}_x} = \alpha|0\rangle_{\mathcal{F}_x} + \beta|1\rangle_{\mathcal{F}_x}$, $|-\rangle_{\mathcal{F}_x} = \beta|0\rangle_{\mathcal{F}_x} - \alpha|1\rangle_{\mathcal{F}_x}$.

We have $\mathcal{T}|\bot^M\rangle = |\text{init}\rangle$ when $\mathcal{T} = \otimes_{x\in[M]}\mathcal{S}_x$ as required by Theorem 3.3. The recording query operator is $\mathcal{R} = \mathcal{S}\mathcal{O}\mathcal{S}$ since $\mathcal{S}^\dagger = \mathcal{S}$, and it satisfies the next equations.

Lemma 5.3. *If the recording query operator $\mathcal{R}$ associated with Definition 5.2 is applied to a basis state $|x,p,w\rangle|f\rangle$ where $p = 1$ then the register $|f(x)\rangle_{\mathcal{F}_x}$ is mapped to*

$$\begin{cases} (1-2\beta^2)|\bot\rangle + 2\alpha\beta^2|0\rangle - 2\alpha^2\beta|1\rangle & \text{if } f(x) = \bot \\ 2\alpha\beta^2|\bot\rangle + (1-2\alpha^2\beta^2)|0\rangle + 2\alpha^3\beta|1\rangle & \text{if } f(x) = 0 \\ -2\alpha^2\beta|\bot\rangle + 2\alpha^3\beta|0\rangle + (1-2\alpha^4)|1\rangle & \text{if } f(x) = 1 \end{cases}$$

and the other registers are unchanged. If $p = 0$ then none of the registers are changed.

Proof. By definition, the unitary $\mathcal{S}_x$ maps $|\bot\rangle_{\mathcal{F}_x} \mapsto |+\rangle$, $|0\rangle_{\mathcal{F}_x} \mapsto \alpha|\bot\rangle + \beta|-\rangle$, $|1\rangle_{\mathcal{F}_x} \mapsto \beta|\bot\rangle - \alpha|-\rangle$. Thus, the action on the register $\mathcal{F}_x$ is

- If $f(x) = \bot$ then $|f(x)\rangle_{\mathcal{F}_x} \overset{\mathcal{S}}{\longmapsto} |+\rangle \overset{\mathcal{O}}{\longmapsto} \alpha|0\rangle - \beta|1\rangle \overset{\mathcal{S}}{\longmapsto} (\alpha^2-\beta^2)|\bot\rangle + 2\alpha\beta|-\rangle$.

- If $f(x) = 0$ then $|f(x)\rangle_{\mathcal{F}_x} \overset{\mathcal{S}}{\mapsto} \alpha|\perp\rangle + \beta|-\rangle \overset{\mathcal{O}}{\mapsto} \alpha|\perp\rangle + \beta(\beta|0\rangle + \alpha|1\rangle) \overset{\mathcal{S}}{\mapsto} 2\alpha\beta^2|\perp\rangle + (1 - 2\alpha^2\beta^2)|0\rangle + 2\alpha^3\beta|1\rangle$.
- If $f(x) = 1$ then $|f(x)\rangle_{\mathcal{F}_x} \overset{\mathcal{S}}{\mapsto} \beta|\perp\rangle - \alpha|-\rangle \overset{\mathcal{O}}{\mapsto} \beta|\perp\rangle - \beta(\beta|0\rangle + \alpha|1\rangle) \overset{\mathcal{S}}{\mapsto} -2\alpha^2\beta|\perp\rangle + 2\alpha^3\beta|0\rangle + (1 - 2\alpha^4)|1\rangle$.

□

If $\alpha \gg \beta$, the above lemma shows that $\mathcal{R}$ is close to the mapping $|\perp\rangle_{\mathcal{F}_x} \mapsto |\perp\rangle - 2\beta|1\rangle$, $|0\rangle_{\mathcal{F}_x} \mapsto |0\rangle + 2\beta|1\rangle$, $|1\rangle_{\mathcal{F}_x} \mapsto -|1\rangle + 2\beta(|0\rangle - |\perp\rangle)$ up to lower order terms of amplitude $O(\beta^2)$.

5.2 Analysis of the recording progress

The measure of progress is based on the number of ones contained in the oracle's register. We first give some projectors related to this quantity.

Definition 5.4. We define the following projectors by giving the basis states on which they project:

- $\Pi_{\leq k}$, $\Pi_{=k}$ and $\Pi_{\geq k}$: all basis states $|x,p,w\rangle|f\rangle$ such that f contains respectively *at most*, *exactly* or *at least* k coordinates equal to 1.
- $\Pi_{=k,\perp}$ and $\Pi_{=k,0}$: all basis states $|x,p,w\rangle|f\rangle$ such that (1) f contains *exactly* k coordinates equal to 1, (2) the phase multiplier is $p = 1$ and (3) $f(x) = \perp$ or $f(x) = 0$ respectively.

We can now define the measure of progress $q_{t,k}$ for t queries and k ones as

$$q_{t,k} = \|\Pi_{\geq k}|\phi_t\rangle\|$$

where $|\phi_t\rangle$ is the state after t queries in the recording query model. The main result of this section is the following bound on the growth of $q_{t,k}$.

Proposition 5.5. *For all t and k, we have that $q_{t,k} \leq \binom{t}{k}\left(\frac{4\sqrt{K}}{\sqrt{N}}\right)^k$.*

Proof. First, $q_{0,0} = 1$ and $q_{0,k} = 0$ for all $k \geq 1$ since the initial state is $|\phi_0\rangle = |0\rangle|\perp^M\rangle$. Then, we prove that $q_{t,k}$ satisfies the following recurrence relation

$$q_{t+1,k+1} \leq q_{t,k+1} + 4\sqrt{\frac{K}{N}}q_{t,k}. \tag{4}$$

From this result, it is trivial to conclude that $q_{t,k} \leq \binom{t}{k}\left(\frac{4\sqrt{K}}{\sqrt{N}}\right)^k$. In order to prove Equation (4), we first observe that $q_{t+1,k+1} = \|\Pi_{\geq k+1}U_{t+1}\mathcal{R}|\phi_t\rangle\| = \|\Pi_{\geq k+1}\mathcal{R}|\phi_t\rangle\|$ where U_{t+1} is the unitary applied by the algorithm at time $t+1$. Then, on a basis state $|x,p,w\rangle|f\rangle$, the recording query operator $\mathcal{R}$ acts as the identity on the registers $\mathcal{F}_{x'}$ for $x' \neq x$. Consequently, the basis states $|x,p,w\rangle|f\rangle$ in $|\phi_t\rangle$ that may contribute to $q_{t+1,k+1}$ must either already contain $k+1$ ones in f, or exactly k ones in f and $f(x) \neq 1$, $p = 1$. This implies that

$$q_{t+1,k+1} \leq q_{t,k+1} + \|\Pi_{\geq k+1}\mathcal{R}\Pi_{=k,\perp}|\phi_t\rangle\| + \|\Pi_{\geq k+1}\mathcal{R}\Pi_{=k,0}|\phi_t\rangle\|.$$

We first bound the term $\|\Pi_{\geq k+1}\mathcal{R}\Pi_{=k,\perp}|\phi_t\rangle\|$. Consider any state $|x,p,w\rangle|f\rangle$ in the support of $\Pi_{=k,\perp}$. By Lemma 5.3, we have $\Pi_{\geq k+1}\mathcal{R}|x,p,w\rangle|f\rangle = -2\alpha^2\beta|x,p,w\rangle|1\rangle_{\mathcal{F}_x} \otimes_{x' \neq x} |f(x')\rangle_{\mathcal{F}_{x'}}$. Since any two basis states in the support of $\Pi_{=k,\perp}$ remain orthogonal after $\Pi_{\geq k+1}\mathcal{R}$ is applied, we obtain that $\|\Pi_{\geq k+1}\mathcal{R}\Pi_{=k,\perp}|\phi_t\rangle\| = 2\alpha^2\beta\|\Pi_{=k,\perp}|\phi_t\rangle\| \leq 2\sqrt{K/N}(1 - K/N)q_{t,k}$.

Similarly, for $|x,p,w\rangle|f\rangle$ in the support of $\Pi_{=k,0}$ we have $\|\Pi_{\geq k+1}\mathcal{R}|x,p,w\rangle|f\rangle\| = 2\alpha^3\beta$ by Lemma 5.3. Consequently, $\|\Pi_{\geq k+1}\mathcal{R}\Pi_{=k,0}|\phi_t\rangle\| = 2\alpha^3\beta\|\Pi_{=k,0}|\phi_t\rangle\| \leq 2\sqrt{K/N}(1 - K/N)^{3/2}q_{t,k}$. We can now conclude the proof,

$$q_{t+1,k+1} \leq q_{t,k+1} + 2\sqrt{\frac{K}{N}}\left(1 - \frac{K}{N}\right)q_{t,k} + 2\sqrt{\frac{K}{N}}\left(1 - \frac{K}{N}\right)^{3/2}q_{t,k} \leq q_{t,k+1} + 4\sqrt{\frac{K}{N}}q_{t,k}.$$

□

5.3 From the recording progress to the success probability

We connect the success probability $\sigma = \|\Pi_{\mathrm{succ}}|\psi_T\rangle\|^2$ in the standard query model to the final progress $q_{T,k}$ in the recording query model after T queries. We show that if the algorithm has made no significant progress for $k \geq K/2$ then it needs to "guess" that $f(x) = 1$ for about $K - k$ positions where the $\mathcal{F}_x$ register does not contain 1. Classically, the probability to find $K - k$ preimages of 1 at positions that have not been queried would be $(K/N)^{K-k}$. Here, we show similarly that if a unit state contains at most k ones in the quantum recording model, then after mapping it to the standard query model (by applying the operator $\mathcal{T}$ of Theorem 3.3) the probability that the output register contains the correct positions of K preimages of 1 is at most $3^K \left(\frac{K}{N}\right)^{K-k}$.

Proposition 5.6. *For any $|\phi\rangle$, we have $\|\Pi_{\mathrm{succ}}\mathcal{T}\Pi_{\leq k}|\phi\rangle\| \leq 3^{K/2}\left(\sqrt{\frac{K}{N}}\right)^{K-k}\|\Pi_{\leq k}|\phi\rangle\|$.*

Proof. Let $|x, p, w\rangle|f\rangle$ be any basis state in the support of $\Pi_{\leq k}$. The output value z is a substring of w made of K distinct values $x_1, \ldots, x_K \in [M]$ indicating positions where the input f is supposed to contain ones. By definition of $\Pi_{\leq k}$, we have $f(x_i) \neq 1$ for at least $K - k$ indices $i \in [K]$. For each such index i, after applying $\mathcal{T} = \otimes_{x' \in [M]}\mathcal{S}_{x'}$, the amplitude of $|1\rangle_{\mathcal{F}_{x_i}}$ is $\sqrt{\frac{K}{N}}$ (if $f(x_i) = \perp$) or $\sqrt{\frac{K}{N}\left(1 - \frac{K}{N}\right)}$ (if $f(x_i) = 0$) by Definition 5.2. Consequently,

$$\|\Pi_{\mathrm{succ}}\mathcal{T}|x, p, w\rangle|f\rangle\| \leq \left(\sqrt{\frac{K}{N}}\right)^{K-k}. \tag{5}$$

Fix any state $|\phi\rangle$ and denote $|\varphi\rangle = \Pi_{\leq k}|\phi\rangle = \sum_{x,p,w,f}\alpha_{x,p,w,f}|x, p, w\rangle|f\rangle$. Let us write $w_{\bar{x}} = \{x_1, \ldots, x_K\}$ when the output substring z of w contains $x_1, \ldots, x_K$. For any two basis states $|x, p, w\rangle|f\rangle$ and $|\bar{x}, \bar{p}, \bar{w}\rangle|\bar{f}\rangle$, if $\left(x, p, w, (f(x'))_{x' \notin w_{\bar{x}}}\right) \neq \left(\bar{x}, \bar{p}, \bar{w}, (\bar{f}(x'))_{x' \notin w_{\bar{x}}}\right)$ then $\Pi_{\mathrm{succ}}\mathcal{T}|x, p, w\rangle|f\rangle$ is orthogonal to $\Pi_{\mathrm{succ}}\mathcal{T}|\bar{x}, \bar{p}, \bar{w}\rangle|\bar{f}\rangle$. There are 3^K choices for $|x, p, w\rangle|f\rangle$ once we set the value of $(x, p, w, (f(x'))_{x' \notin w_{\bar{x}}})$ since it remains to choose $f(x') \in \{\perp, 0, 1\}$ for $x' \in w_{\bar{x}}$. By using the Cauchy–Schwarz inequality and Equation (5), we get that

$$\begin{aligned}
\|\Pi_{\mathrm{succ}}\mathcal{T}|\varphi\rangle\|^2 &= \sum_{x,p,w,(f(x'))_{x' \notin w_{\bar{x}}}} \left\| \sum_{(f(x'))_{x' \in w_{\bar{x}}}} \alpha_{x,p,w,f}\Pi_{\mathrm{succ}}\mathcal{T}|x, p, w\rangle|f\rangle \right\|^2 \\
&\leq \sum_{x,p,w,(f(x'))_{x' \notin w_{\bar{x}}}} \left(\sum_{(f(x'))_{x' \in w_{\bar{x}}}} |\alpha_{x,p,w,f}|^2\right)\left(\sum_{(f(x'))_{x' \in w_{\bar{x}}}} \|\Pi_{\mathrm{succ}}\mathcal{T}|x, p, w\rangle|f\rangle\|^2\right) \\
&\leq \||\varphi\rangle\|^2 \cdot 3^K \left(\frac{K}{N}\right)^{K-k}.
\end{aligned}$$

□

We can now conclude the proof of the main result.

Proof of Theorem 5.1. Let $|\psi_T\rangle$ (resp. $|\phi_T\rangle$) denote the state of the algorithm after T queries in the standard (resp. recording) query model. According to Theorem 3.3, we have $|\psi_T\rangle = \mathcal{T}|\phi_T\rangle$. Thus, by the triangle inequality, the success probability $\sigma = \|\Pi_{\mathrm{succ}}|\psi_T\rangle\|^2$ satisfies $\sqrt{\sigma} \leq \|\Pi_{\mathrm{succ}}\mathcal{T}\Pi_{\geq K/2}|\phi_T\rangle\| + \|\Pi_{\mathrm{succ}}\mathcal{T}\Pi_{\leq K/2}|\phi_T\rangle\| \leq \|\Pi_{\geq K/2}|\phi_T\rangle\| + \|\Pi_{\mathrm{succ}}\mathcal{T}\Pi_{\leq K/2}|\phi_T\rangle\|$. Using Propositions 5.5 and 5.6, we have that $\sqrt{\sigma} \leq \binom{T}{K/2}\left(4\sqrt{K/N}\right)^{K/2} + 3^{K/2}\left(\sqrt{K/N}\right)^{K/2} \leq O(T/\sqrt{KN})^{K/2} + 2^{-K/2-1}$. Finally, the upper bound on σ is derived from the standard inequality $(u+v)^2 \leq 2u^2 + 2v^2$. □

6 Time-space tradeoffs

6.1 Time-space tradeoff for Collision Pairs Finding

We use the time lower bound obtained in Section 4 to derive a new time-space tradeoff for the problem of finding K disjoint collisions in a random function $f : [M] \to [N]$. We recall that the output is produced in an online fashion (Section 2.2), meaning that a collision can be output as soon as it is discovered. The length of the output is not counted toward the space bound. We allow the same collision to be output several times, but it contributes only once to the total count.

Theorem 6.1. *Any quantum algorithm for finding K disjoint collisions in a random function $f : [M] \to [N]$ with success probability $2/3$ must satisfy a time-space tradeoff of $T^3 S \geq \Omega(K^3 N)$, where $1 \leq K \leq N/8$.*

Proof. Our proof relies on the time-segmentation method for large-output problems used in [BFK⁺81, KŠW07] for instance. Fix any quantum circuit $\mathcal{C}$ in the space-bounded model of Section 2.2 running in time T and using $S > \Omega(\log N)$ qubits of memory. The circuit $\mathcal{C}$ is partitioned into $L = T/T'$ consecutive sub-circuits $\mathcal{C}_1 \parallel \mathcal{C}_2 \parallel \cdots \parallel \mathcal{C}_L$ each running in time $T' = S^{2/3}N^{1/3}$, where $\mathcal{C}_j$ takes as input the output memory of $\mathcal{C}_{j-1}$ for each $j \in [L]$. Define X_j to be the random variable that counts the number of (mutually) disjoint collisions that $\mathcal{C}$ outputs between time $(j-1)T'$ and jT' (i.e. in the sub-circuit $\mathcal{C}_j$) when the input is a random function $f : [M] \to [N]$. The algorithm must satisfy $\sum_{j=1}^{L} \mathbb{E}[X_j] \geq \Omega(K)$ to be correct. We claim that the algorithm outputs at most $3S$ collisions in expectation in each segment of the computation. Assume by contradiction that $\mathbb{E}[X_j] \geq 3S$ for some j. Since X_j is bounded between 0 and N we have $\Pr[X_j > 2S] \geq S/N$. Consequently, by running $\mathcal{C}_j$ on the completely mixed state on S qubits we obtain $2S$ disjoint collisions with probability at least $S/N \cdot 2^{-S}$ in time T' (this is akin to a union-bound argument). However, by Theorem 4.6, no quantum algorithm can find more than $2S$ disjoint collisions in time $T' = S^{2/3}N^{1/3}$ with success probability larger than 4^{-S+1}. This contradiction implies that $\mathbb{E}[X_j] \leq 3S$ for all j. Consequently, there must be at least $L \geq \Omega(K/S)$ sub-circuits in order to have $\sum_{j=1}^{L} \mathbb{E}[X_j] \geq \Omega(K)$. Since each sub-circuit runs in time $S^{2/3}N^{1/3}$ the running time of $\mathcal{C}$ is $T \geq \Omega(L \cdot S^{2/3}N^{1/3}) \geq \Omega(KN^{1/3}/S^{1/3})$. □

As an illustration of the above result, we obtain that any quantum algorithm for finding $N/8$ disjoint collisions in a random function must satisfy a time-space tradeoff of $TS^{1/3} \geq \Omega(N^{4/3})$. We prove that any improvement to this lower bound would imply a breakthrough for the Element Distinctness problem.

Definition 6.2. The Element Distinctness problem ED_N on domain size N consists of finding a collision in a random function $f : [N] \to [N^2]$.

It is well-known that the query complexity of Element Distinctness is $T = \Theta(N^{2/3})$ [AS04, Amb07]. However, it is a long-standing open problem to find any quantum time-space lower bound (even classically the question is not completely settled yet [Yao94, BSSV03]). Here, we show that *any* improvement to Theorem 6.1 would imply a non-trivial time-space tradeoff for Element Distinctness. This result relies on a reduction presented in Algorithm 1 and analyzed in Proposition 6.3 (the constants c_0, c_1, c_2 are chosen in the proof).

Proposition 6.3. *Let N be a square number. If there is an algorithm solving ED_N in time T_N and space S_N then Algorithm 1 runs in time $O(NT_{\sqrt{N}})$ and space $O(S_{\sqrt{N}})$, and it finds $c_1 N$ collisions in any function $f : [N] \to [N]$ containing at least $c_0 N$ collisions.*

The proof of Proposition 6.3 is deferred to Appendix A. We now use the above reduction to transform any low-space algorithm for Element Distinctness into one for finding $\Omega(N/\log N)$

Input: a function $f : [N] \to [N]$ containing at least $c_0 N$ collisions.
Output: at least $c_1 N$ collisions in f (not necessarily disjoint).

1. Repeat $c_2 N$ times:
   (a) Sample a 4-wise independent hash function $h : [\sqrt{N}] \to [N]$ and store it in memory.
   (b) Run an algorithm for $\mathrm{ED}_{\sqrt{N}}$ on input $f \circ h : [\sqrt{N}] \to [N]$. If it finds a collision $(f \circ h(i), f \circ h(j))$ check if $h(i) \neq h(j)$ and output the collision $(h(i), h(j))$ in this case.

Algorithm 1: Finding collisions by using $\mathrm{ED}_{\sqrt{N}}$.

disjoint collisions in a random function. Observe that Algorithm 1 does not necessarily output collisions that are mutually disjoint. Nevertheless, there is a small probability that a random function $f : [M] \to [N]$ contains multi-collisions of size larger than $\log N$ when $M \approx N$ [FO89]. Thus, there is only a $\log N$ loss in the analysis.

Proposition 6.4. *Suppose that there exists a bounded-error quantum algorithm for solving Element Distinctness on domain size N that satisfies a time-space tradeoff of $T^\alpha S^\beta \leq \widetilde{O}\big(N^{2(\gamma-\alpha)}\big)$ for some constants α, β, γ. Then, there exists a bounded-error quantum algorithm for finding $\Omega(N/\log N)$ disjoint collisions in a random function $f : [10N] \to [N]$ that satisfies a time-space tradeoff of $T^\alpha S^\beta \leq \widetilde{O}(N^\gamma)$.*

Proof. We use the constants c_0, c_1, c_2 specified in the proof of Proposition 6.3. First, we note that a random function $f : [10N] \to [N]$ contains $c_0 N$ collisions and no multi-collisions of size larger than $\log(N)$ with large probability [FO89]. Consequently, any set of $c_1 N$ collisions must contain at least $c_1 N/\log N$ mutually disjoint collisions with large probability. Assume now that there exists an algorithm solving $\mathrm{ED}_{\sqrt{10N}}$ in time $T_{\sqrt{10N}}$ and space $S_{\sqrt{10N}}$ such that $\big(T_{\sqrt{10N}}\big)^\alpha S^\beta_{\sqrt{10N}} \leq \widetilde{\Omega}(N^{\gamma-\alpha})$. Then, by plugging it into Algorithm 1, one can find $c_1 N/\log N$ disjoint collisions in a random function $f : [10N] \to [N]$ in time $T = O\big(N T_{\sqrt{10N}}\big)$ and space $S = O\big(S_{\sqrt{10N}}\big)$. We derive from the above tradeoff that $T^\alpha S^\beta \leq \widetilde{O}(N^\gamma)$. $\square$

As an application of Proposition 6.4, we obtain the following result regarding the hardness of finding $\widetilde{\Omega}(N)$ collisions.

Corollary 6.5. *Suppose that there exists $\epsilon \in (0,1)$ such that any quantum algorithm for finding $\widetilde{\Omega}(N)$ disjoint collisions in a random function $f : [10N] \to [N]$ must satisfy a time-space tradeoff of $TS^{1/3} \geq \widetilde{\Omega}(N^{4/3+\epsilon})$. Then, any quantum algorithm for solving Element Distinctness on domain size N must satisfy a time-space tradeoff of $TS^{1/3} \geq \widetilde{\Omega}(N^{2/3+2\epsilon})$.*

We conjecture that the optimal tradeoff for finding K collisions is $T^2 S = \Theta(K^2 N)$, which would imply an optimal time-space tradeoff of $T^2 S \geq \widetilde{\Omega}(N^2)$ for Element Distinctness.

Conjecture 1. Any quantum algorithm for finding K disjoint collisions in a random function $f : [M] \to [N]$ with success probability $2/3$ must satisfy a time-space tradeoff of $T^2 S \geq \Omega(K^2 N)$.

Corollary 6.6. *If Conjecture 1 is true, then any quantum algorithm for solving the Element Distinctness problem with success probability $2/3$ must satisfy a time-space tradeoff of $T^2 S \geq \widetilde{\Omega}(N^2)$.*

We describe a quantum algorithm that achieves the tradeoff of $T^2 S \leq \widetilde{O}(K^2 N)$. In order to simplify the analysis, we do not require the collisions to be disjoint.

1. Repeat $\widetilde{O}(K/S)$ times:

   (a) Sample a subset $G \subset [N]$ of size S uniformly at random.

   (b) Construct a table containing all pairs $(x, f(x))$ for $x \in G$. Sort the table according to the second entry of each pair.

   (c) Define the function $g : [N] \setminus G \to \{0,1\}$ where $g(x) = 1$ iff there exists $x' \in G$ such that $f(x) = f(x')$. Run the Grover search algorithm [BBHT98] on g, by using the table computed at step 1.b, to find all pairs $(x, x') \in G \times ([N] \setminus G)$ such that $f(x) = f(x')$. Output all of these pairs.

Algorithm 2: Finding K collision pairs in $f : [N] \to [N]$ using a memory of size S.

Proposition 6.7. *For any $1 \leq K \leq O(N)$ and $\widetilde{\Omega}(\log N) \leq S \leq \widetilde{O}(K^{2/3}N^{1/3})$, there exists a bounded-error quantum algorithm that can find K collisions in a random function $f : [N] \to [N]$ by making $T = \widetilde{O}(K\sqrt{N/S})$ queries and using S qubits of memory.*

Proof. We prove that Algorithm 2 satisfies the statement of the proposition. For simplicity, we do not try to tune the hidden factors in the big O notations.

The probability that a fixed pair (x, x') satisfies $(x, x') \in G \times ([N] \setminus G)$ for at least one iteration of step 1 is $\Omega(K/S \cdot S/N \cdot (1 - S/N)) = \Omega(K/N)$. Since a random function $f : [N] \to [N]$ contains $\Omega(N)$ collisions with high probability, the algorithm encounters $\Omega(K)$ collisions in total. Thus, if the Grover search algorithm never fails we obtain the desired number of collisions.

The expected number of pre-images of 1 under g is $O(S)$. Consequently, the complexity of Grover's search at step 1.c is $O(\sqrt{SN})$. The overall query complexity is $T = \widetilde{O}(K/S \cdot \sqrt{SN}) = \widetilde{O}(K\sqrt{N/S})$, and the space complexity is $\widetilde{O}(S)$. □

6.2 Time-space tradeoff for Sorting

We use the time lower bound obtained in Section 5 to reprove the time-space tradeoff for the Sorting problem described in [KŠW07, Theorem 21]. The input to the Sorting problem is represented as a function $f : [N] \to \{0,1,2\}$ (we do not need to consider a larger range for the proof). A quantum algorithm for the Sorting problem must output in order a sequence $x_1, \ldots, x_N \in [N]$ of distinct integers such that $f(x_1) \geq f(x_2) \geq \cdots \geq f(x_N)$ with probability at least $2/3$.

Theorem 6.8. *Any quantum algorithm for sorting a function $f : [N] \to \{0,1,2\}$ with success probability $2/3$ must satisfy a time-space tradeoff of $T^2 S \geq \Omega(N^3)$.*

Proof. The proof is a modified version of [KŠW07, Theorem 21] adapted to our version of the K-Search problem. Given a circuit $\mathcal{C}$ that runs in time T and space $\Omega(\log N) \leq S \leq N/64$, we partition it into $L = T/T'$ consecutive sub-circuits $\mathcal{C}_1 \parallel \mathcal{C}_2 \parallel \cdots \parallel \mathcal{C}_L$ each running in time $T' = \sqrt{SN}/4$. Assume by contradiction that a circuit $\mathcal{C}_j$ outputs the elements of ranks $r, r+1, \ldots, r+2S-1$ for some $r \leq N/2$. We use $\mathcal{C}_j$ to solve the K-search problem for $K = 2S$ as follows. Given an input $g : [N/2] \to \{0,1\}$ to the K-search problem where $g(x) = 1$ with probability $\frac{K}{N/4}$ for each x, define the function $f : [N] \to \{0,1,2\}$ where

$$f(x) = \begin{cases} 2 & \text{if } x < r, \\ g(x - r + 1) & \text{if } r \leq x < r + N/2, \\ 0 & \text{if } x \geq r + N/2. \end{cases}$$

Note that the function g contains at least $2S$ preimages of 1 with probability at least $2S/N$. Thus, if the circuit $\mathcal{C}$ is run on the input f, then the indices output by the sub-circuit $\mathcal{C}_j$ must contain the position of $2S$ preimages of 1 with probability at least $2/3 \cdot 2S/N$. Consequently, by running $\mathcal{C}_j$ on the completely mixed state on S qubits we can find $2S$ preimages of 1 under g with probability at least $2/3 \cdot 2S/N \cdot 2^{-S}$ in time T'. However, by Theorem 5.1, any such algorithm must succeed with probability at most 4^{-S+1}. This contradiction implies that there must be at least $L \geq \Omega(N/S)$ sub-circuits in $\mathcal{C}$. Thus, the running time of $\mathcal{C}$ is $T \geq \Omega(L \cdot \sqrt{SN}) \geq \Omega(N^{3/2}/\sqrt{S})$. □

The time-space tradeoffs for the Boolean matrix-vector product [KŠW07, Theorem 23] and the Boolean matrix product [KŠW07, Theorem 25] problems can be reproved in a similar way.

Acknowledgements

The authors want to thank the anonymous referees for their valuable comments and suggestions which helped to improve this paper. This research was supported in part by the ERA-NET Cofund in Quantum Technologies project QuantAlgo and the French ANR Blanc project RDAM.

References

[Aar05] S. Aaronson. Limitations of quantum advice and one-way communication. *Theory of Computing*, 1(1):1–28, 2005. 5

[Abr90] K. Abrahamson. A time-space tradeoff for boolean matrix multiplication. In *Proceedings of the 31st Symposium on Foundations of Computer Science (FOCS)*, pages 412–419, 1990. 2

[ACC+18] G. Adj, D. Cervantes-Vázquez, J.-J. Chi-Domínguez, A. Menezes, and F. Rodríguez-Henríquez. On the cost of computing isogenies between supersingular elliptic curves. In *Proceedings of the 25th Conference on Selected Areas in Cryptography (SAC)*, pages 322–343, 2018. 2

[Amb02] A. Ambainis. Quantum lower bounds by quantum arguments. *Journal of Computer and System Sciences*, 64(4):750–767, 2002. 6

[Amb07] A. Ambainis. Quantum walk algorithm for element distinctness. *SIAM Journal on Computing*, 37(1):210–239, 2007. 1, 3, 15

[Amb10] A. Ambainis. A new quantum lower bound method, with an application to a strong direct product theorem for quantum search. *Theory of Computing*, 6(1):1–25, 2010. 3

[AS04] S. Aaronson and Y. Shi. Quantum lower bounds for the collision and the element distinctness problems. *Journal of the ACM*, 51(4):595–605, 2004. 2, 3, 15

[AŠW09] A. Ambainis, R. Špalek, and R. de Wolf. A new quantum lower bound method, with applications to direct product theorems and time-space tradeoffs. *Algorithmica*, 55(3):422–461, 2009. 2, 6

[BBHT98] M. Boyer, G. Brassard, P. Høyer, and A. Tapp. Tight bounds on quantum searching. *Fortschritte der Physik*, 46(4-5):493–505, 1998. 17

[Bea91] P. Beame. A general sequential time-space tradeoff for finding unique elements. *SIAM Journal on Computing*, 20(2):270–277, 1991. 2

[Ber05] D. J. Bernstein. Understanding brute force, 2005. ECRYPT STVL Workshop on Symmetric Key Encryption. 1

[Ber09] D. J. Bernstein. Cost analysis of hash collisions: Will quantum computers make SHARCS obsolete? In *Proceedings of the 4th Workshop on Special-purpose Hardware for Attacking Cryptograhic Systems (SHARCS)*, pages 105–116, 2009. 1

[BFK+81] A. Borodin, M. J. Fischer, D. G. Kirkpatrick, N. A. Lynch, and M. Tompa. A time-space tradeoff for sorting on non-oblivious machines. *Journal of Computer and System Sciences*, 22(3):351–364, 1981. 2, 4, 15

[BFMadH+87] A. Borodin, F. E. Fich, F. Meyer auf der Heide, E. Upfal, and A. Wigderson. A time-space tradeoff for element distinctness. *SIAM Journal on Computing*, 16(1):97–99, 1987. 2, 3

[BHT98] G. Brassard, P. Høyer, and A. Tapp. Quantum cryptanalysis of hash and claw-free functions. In *Proceedings of the 3rd Latin American Symposium on Theoretical Informatics (LATIN)*, pages 163–169, 1998. 1, 2

[BSSV03] P. Beame, M. Saks, X. Sun, and E. Vee. Time-space trade-off lower bounds for randomized computation of decision problems. *Journal of the ACM*, 50(2):154–195, 2003. 2, 3, 15

[BW02] H. Buhrman and R. de Wolf. Complexity measures and decision tree complexity: A survey. *Theoretical Computer Science*, 288(1):21–43, 2002. 5

[CC17] A. Chakrabarti and Y. Chen. Time-space tradeoffs for the memory game, 2017. `arXiv:1712.01330 [cs.CC]`. 2

[CMS19] A. Chiesa, P. Manohar, and N. Spooner. Succinct arguments in the quantum random oracle model. In *Proceedings of the 17th Conference on Theory of Cryptography (TCC)*, pages 1–29, 2019. 4

[CMSZ19] J. Czajkowski, C. Majenz, C. Schaffner, and S. Zur. Quantum lazy sampling and game-playing proofs for quantum indifferentiability, 2019. `arXiv:1904.11477v1 [quant-ph]`. 4

[DDKS12] I. Dinur, O. Dunkelman, N. Keller, and A. Shamir. Efficient dissection of composite problems, with applications to cryptanalysis, knapsacks, and combinatorial search problems. In *Proceedings of the 32th International Cryptology Conference (CRYPTO)*, pages 719–740, 2012. 2

[DEM19] C. Delaplace, A. Esser, and A. May. Improved low-memory subset sum and LPN algorithms via multiple collisions. In *Proceedings of the 17th IMA International Conference on Cryptography and Coding (IMACC)*, pages 178–199, 2019. 2

[Din20] I. Dinur. Tight time-space lower bounds for finding multiple collision pairs and their applications. In *Proceedings of the 39th International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)*, pages 405–434, 2020. 1, 2, 3

[FO89] P. Flajolet and A. M. Odlyzko. Random mapping statistics. In *Proceedings of the 7th Workshop on the Theory and Application of Cryptographic Techniques (EUROCRYPT)*, pages 329–354, 1989. 2, 16

[GJT20] A. Ghoshal, J. Jaeger, and S. Tessaro. The memory-tightness of authenticated encryption. In *Proceedings of the 40th International Cryptology Conference (CRYPTO)*, pages 127–156, 2020. 1

[HI19] A. Hosoyamada and T. Iwata. 4-round Luby-Rackoff construction is a qPRP. In *Proceedings of the 25th International Conference on the Theory and Applications of Cryptology and Information Security (ASIACRYPT)*, pages 145–174, 2019. 4

[JL09] A. Joux and S. Lucks. Improved generic algorithms for 3-collisions. In *Proceedings of the 15th International Conference on the Theory and Applications of Cryptology and Information Security (ASIACRYPT)*, pages 347–363, 2009. 2

[JT19] J. Jaeger and S. Tessaro. Tight time-memory trade-offs for symmetric encryption. In *Proceedings of the 38th International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)*, pages 467–497, 2019. 1

[KŠW07] H. Klauck, R. Špalek, and R. de Wolf. Quantum and classical strong direct product theorems and optimal time-space tradeoffs. *SIAM Journal on Computing*, 36(5):1472–1493, 2007. 1, 2, 3, 4, 5, 6, 15, 17, 18

[LZ19] Q. Liu and M. Zhandry. On finding quantum multi-collisions. In *Proceedings of the 38th International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)*, pages 189–218, 2019. 1, 4

[MNT93] Y. Mansour, N. Nisan, and P. Tiwari. The computational complexity of universal hashing. *Theoretical Computer Science*, 107(1):121–133, 1993. 2

[OW99] P. C. van Oorschot and M. J. Wiener. Parallel collision search with cryptanalytic applications. *Journal of Cryptology*, 12(1):1–28, 1999. 1, 2, 3

[Pol75] J. M. Pollard. A Monte Carlo method for factorization. *BIT Numerical Mathematics*, 15(3):331–334, 1975. 2

[Špa08] R. Špalek. The multiplicative quantum adversary. In *Proceedings of the 23rd Computational Complexity Conference (CCC)*, pages 237–248, 2008. 3

[TT18] S. Tessaro and A. Thiruvengadam. Provable time-memory trade-offs: Symmetric cryptography against memory-bounded adversaries. In *Proceedings of the 16th Conference on Theory of Cryptography (TCC)*, pages 3–32, 2018. 1

[Vre16] C. van Vredendaal. Reduced memory meet-in-the-middle attack against the NTRU private key. *LMS Journal of Computation and Mathematics*, 19(A):43–57, 2016. 2

[Wag02] D. Wagner. A generalized birthday problem. In *Proceedings of the 22nd International Cryptology Conference (CRYPTO)*, pages 288–304, 2002. 2

[Wie04] M. J. Wiener. The full cost of cryptanalytic attacks. *Journal of Cryptology*, 17(2):105–124, 2004. 1

[Yao94] A. C.-C. Yao. Near-optimal time-space tradeoff for element distinctness. *SIAM Journal on Computing*, 23(5):966–975, 1994. 2, 3, 15

[Zha15] M. Zhandry. A note on the quantum collision and set equality problems. *Quantum Information & Computation*, 15(7&8):557–567, 2015. 2, 3

[Zha19] M. Zhandry. How to record quantum queries, and applications to quantum indifferentiability. In *Proceedings of the 39th International Cryptology Conference (CRYPTO)*, pages 239–268, 2019. 1, 4, 6

A Analysis of Algorithm 1

In this section we prove the following statement about Algorithm 1.

Proposition 6.3 (Restated)**.** *Let N be a square number. If there is an algorithm solving ED_N in time T_N and space S_N then Algorithm 1 runs in time $O(NT_{\sqrt{N}})$ and space $O(S_{\sqrt{N}})$, and it finds $c_1 N$ collisions in any function $f : [N] \to [N]$ containing at least $c_0 N$ collisions.*

Proof. We choose $c_0 = 40$, $c_1 = 1/10^4$ and $c_2 = 8$. We study the probabilities of the following events to occur in a fixed round of Algorithm 1:

- **Event A:** The hash function h is collision free (i.e. $h(i) \neq h(j)$ for all $i \neq j$).
- **Event B:** None of the collisions output during the previous rounds is present in the image of h.
- **Event C:** The function $f \circ h : [\sqrt{N}] \to [N]$ contains a collision.
- **Event D:** The algorithm for $\mathrm{ED}_{\sqrt{N}}$ finds a collision at step 2.b.

Algorithm 1 succeeds if and only if the event $A \wedge B \wedge C \wedge D$ occurs during at least $c_1 N$ rounds. We now lower bound the probability of this event happening.

For **event A**, let us consider the random variable $X = \sum_{i \neq j \in [\sqrt{N}]} 1_{h(i)=h(j)}$. Using that h is pairwise independent, we have $\mathbb{E}[X] = \binom{\sqrt{N}}{2} \frac{1}{N} \leq \frac{1}{2}$. Thus, by Markov's inequality, $\Pr[A] = 1 - \Pr[X \geq 1] \geq \frac{1}{2}$.

For **event B**, let us assume that $k < c_1 N$ collisions $(x_1, x_2), \ldots, (x_{2k-1}, x_{2k})$ have been output so far. For any $i \in [k]$, the probability that both x_{2i-1} and x_{2i} belong to $\{h(1), \ldots, h(\sqrt{N})\}$ is at most $\binom{\sqrt{N}}{2} \frac{2}{N^2} \leq \frac{1}{N}$ since h is pairwise independent. By a union bound, $\Pr[B] \geq 1 - \frac{k}{N} \geq 1 - c_1$.

For **event C**, let us consider the binary random variables $Y_{i,j} = 1_{f \circ h(i) = f \circ h(j)}$ for $i \neq j \in [\sqrt{N}]$, and let $Y = \sum_{i \neq j} Y_{i,j}$ be twice the number of collisions in $f \circ h$. Note that we may have $Y_{i,j} = 1$ because $h(i) = h(j)$ (this is taken care of in event A). For each $y \in [N]$, let $N_y = |\{x : f(x) = y\}|$ denote the number of elements that are mapped to y by f. Using that h is 4-wise independent, for any $i \neq j \neq k \neq \ell$ we have,

$$\begin{cases} \Pr[Y_{i,j} = 1] = \frac{\sum_{y \in [N]} N_y^2}{N^2} \\ \Pr[Y_{i,j} = 1 \wedge Y_{i,k} = 1] = \frac{\sum_{y \in [N]} N_y^3}{N^3} \\ \Pr[Y_{i,j} = 1 \wedge Y_{k,\ell} = 1] = \Pr[Y_{i,j} = 1] \cdot \Pr[Y_{k,\ell} = 1]. \end{cases}$$

Consequently, $\mathbb{E}[Y] = \binom{\sqrt{N}}{2} \frac{\sum_{y \in [N]} N_y^2}{N^2}$ and

$$\begin{aligned} \mathrm{Var}[Y] &= \sum_{\{i,j\}} \mathrm{Var}[Y_{i,j}] + \sum_{\{i,j\} \neq \{i,k\}} \mathrm{Cov}[Y_{i,j}, Y_{i,k}] + \sum_{\{i,j\} \cap \{k,\ell\} = \varnothing} \mathrm{Cov}[Y_{i,j}, Y_{k,\ell}] \\ &\leq \sum_{\{i,j\}} \mathbb{E}[Y_{i,j}^2] + \sum_{\{i,j\} \neq \{i,k\}} \mathbb{E}[Y_{i,j} Y_{i,k}] \\ &= \binom{\sqrt{N}}{2} \frac{\sum_{y \in [N]} N_y^2}{N^2} + 3 \binom{\sqrt{N}}{3} \frac{\sum_{y \in [N]} N_y^3}{N^3} \end{aligned}$$

where we have used that $Y_{i,j}$ and $Y_{k,\ell}$ are independent when $i \neq j \neq k \neq \ell$. The term $\sum_{y\in[N]} N_y^2$ is equal to the number of pairs $(x, x') \in [N]^2$ such that $f(x) = f(x')$. Each collision in f gives two such pairs, and we must also count the pairs (x, x). Thus, $\sum_{y\in[N]} N_y^2 \geq (1 + 2c_0)N$. Moreover, $\sum_{y\in[N]} N_y^3 \leq (\sum_{y\in[N]} N_y^2)^{3/2}$. Consequently,

$$\frac{\mathrm{Var}[Y]}{\mathbb{E}[Y]^2} \leq \frac{1 + \sqrt{N}\left(\frac{\sum_{y\in[N]} N_y^2}{N^2}\right)^{1/2}}{\binom{\sqrt{N}}{2}\frac{\sum_{y\in[N]} N_y^2}{N^2}} \leq \frac{4(1 + \sqrt{1 + 2c_0})}{1 + 2c_0}.$$

Finally, according to Chebyshev's inequality, $\Pr[Y = 0] \leq \Pr[|Y - \mathbb{E}[Y]| \geq \mathbb{E}[Y]] \leq \frac{\mathrm{Var}[Y]}{\mathbb{E}[Y]^2}$. Thus, $\Pr[C] = 1 - \Pr[Y = 0] \geq 1 - \frac{4(1+\sqrt{1+2c_0})}{1+2c_0}$.

For **event D**, we have $\Pr[D \mid A \wedge B \wedge C] \geq 2/3$ assuming the bounded-error algorithm for solving $\mathrm{ED}_{\sqrt{N}}$ succeeds with probability 2/3.

The probability of the four events happening together is $\Pr[A \wedge B \wedge C \wedge D] = \Pr[D \,|\, A \wedge B \wedge C] \cdot \Pr[A \wedge B \wedge C] \geq \Pr[D|A \wedge B \wedge C] \cdot (\Pr[A] + \Pr[B] + \Pr[C] - 2) \geq \frac{2}{3} \cdot \left(\frac{1}{2} - c_1 - \frac{4(1+\sqrt{1+2c_0})}{1+2c_0}\right) \geq 1/250$. Let τ be the number of rounds after which c_1N collisions have been found (i.e. $A \wedge B \wedge C \wedge D$ has occurred c_1N times). We have $\mathbb{E}[\tau] \leq 8c_1N$, and by Markov's inequality $\Pr[\tau \geq c_2N] \leq 250c_1/c_2 \leq 1/3$. Thus, with probability at least 2/3, Algorithm 1 outputs at least c_1N collisions in f. □