Generalized Hybrid Search with Applications to Blockchains and Hash Function Security

Alexandru Cojocaru[1](✉), Juan Garay[2], and Fang Song[3]

[1] School of Informatics, University of Edinburgh, Edinburgh, UK
a.cojocaru@ed.ac.uk

[2] Department of Computer Science and Engineering, Texas A&M University, College Station, USA
garay@tamu.edu

[3] Department of Computer Science, Portland State University, Portland, USA
fang.song@pdx.edu

Abstract. In this work we first examine the hardness of solving various search problems by hybrid quantum-classical strategies, namely, by algorithms that have both quantum and classical capabilities. We then construct a hybrid quantum-classical search algorithm and analyze its success probability.

Regarding the former, for search problems that are allowed to have multiple solutions and in which the input is sampled according to arbitrary distributions, we establish their hybrid quantum-classical query complexities—i.e., given a fixed number of classical and quantum queries, determine what is the probability of solving the search task. At a technical level, our results generalize the framework for hybrid quantum-classical search algorithms recently proposed by Rosmanis [Ros22]. Namely, for an *arbitrary* distribution D on Boolean functions, the probability that an algorithm equipped with τ_c classical queries and τ_q quantum queries succeeds in finding a preimage of 1 for a function sampled from D is at most $\nu_D \cdot (2\sqrt{\tau_c} + 2\tau_q + 1)^2$, where ν_D captures the average (over D) fraction of preimages of 1.

Regarding our second contribution, we design a hybrid algorithm which first spends all of its classical queries and in the second stage runs a "modified Grover" in which the initial state depends on the target distribution D. We then show how to analyze its success probability for arbitrary target distributions and, importantly, its optimality for the uniform and the Bernoulli distribution cases.

As applications of our hardness results, we first revisit and generalize the formal security treatment of the Bitcoin protocol called the *Bitcoin backbone* [Eurocrypt 2015], to a setting where the adversary has both quantum and classical capabilities, presenting a new *hybrid honest majority* condition necessary for the protocol to properly operate. Secondly, we re-examine the generic security of hash functions [PKC 2016] against quantum-classical hybrid adversaries.

The full version of the paper can be found at [CGS23].

K.-M. Chung and Y. Sasaki (Eds.): ASIACRYPT 2024, LNCS 15492, pp. 65–93, 2025.
https://doi.org/10.1007/978-981-96-0947-5_3

1 Introduction

The query model is an elegant abstraction and is widely adopted in cryptography. A notable example is the random oracle (RO) model [BR93], where a hash function f is modeled as a random black-box function, and all parties including the adversary can evaluate it only by issuing a query x and receiving $f(x)$ in response. Numerous cryptosystems have been designed and analyzed in the RO model—e.g., [BR94,BR96,Sho01,FOPS04,FO13].

The advent of quantum computing brings about a new query model, where *superposition* queries to the hash function f in the form of $\sum_{x,y} \alpha_{x,y} |x\rangle |y\rangle \mapsto \sum_{x,y} \alpha_{x,y} |x\rangle |y \oplus f(x)\rangle$ are permitted, which equips quantum adversaries with new capabilities. Indeed, some classically secure digital signature and public-key encryption schemes are broken in the *quantum* random oracle (QRO) model, where a quantum adversary is able to make such superposition queries to f [YZ21]. As such, a significant amount of effort has been devoted to address such quantum-query adversaries (cf. [BDF+11,ES15,Unr15,HHK17, AHU19,DFMS19,CMS19,ES20,DFMS22]), often resulting in considerable efficiency overhead, such as more complex constructions or larger key sizes, in order to maintain security.

However alarming this threat is, it does not come for free, as it requires running a large-scale quantum computer coherently for an extended amount of time, while in the near-to-intermediate term the available quantum devices are likely to be computationally restricted as well as expensive [Pre18]. This reality inspires a *hybrid* query model, where the computational entity (the adversary) is granted a quota of both classical and quantum queries, resulting in a model which subsumes the classical and quantum query models as special cases. Thus, establishing a trade-off between classical and quantum queries allows giving a more accurate estimation of security and hence optimized parameter choices for cryptosystems depending on what resources are likely to be available to near-term quantum adversaries.

Recently, Rosmanis studied the basic unstructured search problem in the hybrid query model [Ros22], where given oracle function $f : X \to \{0, 1\}$, one wants to find a "marked" input, i.e., x with $f(x) = 1$. This search problem and many variants, such as multiple or randomly chosen marked inputs, are well understood when all queries are quantum [Gro96,BBBV97,Zal99,DH09, Zha19], and where Grover's quantum algorithm gives a quadratic speedup over classical algorithms, which is also proven to be optimal [BBBV97]. To reiterate, Rosmanis's work proves the hardness of searching in the domain of a function with a *unique* marked input x^* in the hybrid query model. Specifically, any quantum algorithm with τ_c classical queries and τ_q quantum queries succeeds in finding x^* with probability at most $\frac{1}{|X|} \cdot (2\sqrt{\tau_c} + 2\tau_q + 1)^2$. This hardness bound is also shown in [HLS22], by a new recording technique tailored to the hybrid query model.

1.1 Our Contributions and Technical Overview

Bounding the Hardness of Hybrid Search

In this work, we consider an arbitrary distribution D on the function family $\mathcal{F} = \{f : X \to \{0, 1\}\}$, and prove a precise upper bound on the probability of finding a preimage x with $f(x) = 1$ when $f \leftarrow D$, for any algorithm $\mathcal{A}$ spending τ_c classical and τ_q quantum queries. Specifically, we show that:

$$\Pr_{f \leftarrow D}[f(x) = 1 : x \leftarrow \mathcal{A}^f] \leq \nu_D \cdot (2\sqrt{\tau_c} + 2\tau_q + 1)^2 ,$$

where $\nu_D \overset{\text{def}}{=} \sup_{\varphi:\|\varphi\|\leq 1} \left(\mathbb{E}_{f \leftarrow D} \left\| \left(\sum_{x:f(x)=1} |x\rangle \langle x| \right) \varphi \right\|^2 \right)$ captures the *average* fraction of preimages of 1 and is solely determined by the distribution D.

Our generalized bound then allows us to derive hardness bounds for specific relevant distributions. "All" we need to do is to analyze ν_D, and this usually can be done by simple combinatorial arguments. For example, let D be the uniform distribution over functions with exactly one marked input. Then we can observe that $\nu_D = \Pr_{f \leftarrow D}[f(x) = 1] = 1/|X|$ for an arbitrary x, which reclaims the result by Rosmanis [Ros22]. The hardness of searching given a function with $w > 1$ marked items can be similarly derived.

We further demonstrate our result on another distribution D_η, where each input is marked according to a Bernoulli trial. Namely, for every $x \in X$, we set $f(x) = 1$ with probability η *independently*. By determining ν_D in this case, we derive the hardness of search when the function is drawn from D_η. This search problem under D_η, which we call *Bernoulli Search*, is particularly useful in several cryptographic applications. Firstly, we can prove generic security bounds for hash function properties, such as preimage-resistance, second-preimage resistance and their multi-target extensions, against hybrid quantum-classical adversaries. This follows by first adapting the reductions in [HRS16], where the hash properties are connected to the *Bernoulli Search* problem in the fully quantum query setting, and then plugging in our hybrid hardness bound of *Bernoulli Search*. In another application, *Bernoulli Search* was shown to dictate the security of proofs of work (PoWs) and security properties of Bitcoin-like blockchains in the RO model (with fully quantum queries) [CGK+23]. This allows us to identify a new *honest-majority* condition under which the security of the PoW-based Bitcoin blockchain holds against hybrid adversaries equipped with both classical and quantum queries.

At a technical level, the proof of our hardness bound follows the overall strategy of [Ros22]. As in the standard optimality proof of Grover's algorithm [BBBV97], one would consider running an adversary's algorithm with respect to the input function $f \leftarrow D$ or a constant-0 function. Then one argues that each query diverges the states in these two cases, which is called a *progress measure*, by a small amount. On the other hand, in order to find a marked input in f, the final states need to differ significantly. Therefore, sufficiently many queries are necessary for the cumulative progress to grow adequately.

Now, when classical queries are mixed up with quantum queries, the quantum states would collapse after each classical query, and it becomes unclear how to measure the progress. To address this, Rosmanis considers instead an intermediate oracle named *pseudo-classical.* Namely, consider a quantum query with the output register initialized to $|0\rangle$: $\sum_x \alpha_x |x\rangle |0\rangle \mapsto \sum_x \alpha_x |x\rangle |f(x)\rangle$. We can then view a classical query as the result of measuring the input register that collapses to x and receiving $f(x)$, whereas a pseudo-classical oracle measures the output register, resulting in one of two possible outcomes: $\sum_{x:f(x)=0} \alpha_x |x\rangle |0\rangle$ (denoted as the *0-outcome branch*) or $\sum_{x:f(x)=1} \alpha_x |x\rangle |1\rangle$ (denoted as the *1-outcome branch*). With this change, one instead tracks the progress between: (i) the 0-outcome branch in case of $f \leftarrow D$, and (ii) the state in case of the constant-0 function (which always stays in the 0-outcome branch). The algorithm fails if its state stays in the 0-outcome branch and is close to the state in the constant-0 case. A key ingredient in our proof is to deliberately separate the evolution of various objects on an *individual* function and which *characteristics* of the distribution D influence the evolution and in what way. This enables us to obtain a clean and concise lower bound for the generalized hybrid search problem.

Hybrid Search Algorithms: Design and Analysis. In the second part of our work we focus on constructing a hybrid search algorithm for an arbitrary distribution D and show that in several interesting cases (e.g., Bernoulli) the algorithm is optimal. Inspired by our hardness analysis, our algorithm proceeds in a two-stage fashion:

- The first stage is purely *classical.* We query the τ_c inputs that are the most likely to be assigned the value 1 under D. More precisely, for any x in the input domain, let the function $\omega(x) = \sum_f D(f) \cdot f(x)$, which can be viewed as the (unnormalized) probability that $f(x) = 1$ with f drawn from D. Let S be the set of inputs whose $\omega(x)$ values are the τ_c-highest (ties are broken arbitrarily). Then the algorithm queries all the points $x \in S$. If none of them give a solution, we move on to the second stage.
- The second stage is fully *quantum.* We run a modified Grover algorithm $\mathcal{A}$ which is tailored to the prior knowledge on the distribution D. Instead of starting from an equal superposition of all points in the search space as in the standard Grover search algorithm, we construct an initial state in which the amplitude of each point is proportional to $\omega(x)$. Then, for each of the τ_q quantum queries, two reflection operators are applied to rotate the initial state towards a target state encoding the solutions. We give a comprehensive analysis and derive a precise lower bound for the success probability of $\mathcal{A}$ on the distribution D, which amounts to $\tau_q^2 \cdot \frac{\sum_x \omega^2(x)}{\sum_x \omega(x)}$. In other words, for the algorithm in the second stage, we define an induced distribution $\tilde{D}$ by restricting and (re-normalizing) D to functions f satisfying $f(x) = 0$ for all $x \in S$. We then invoke $\mathcal{A}$ on $\tilde{D}$ in a modular way.

Note that the hybrid algorithm needs to compute the values $\omega(x)$ from the description of the target distribution D, and during the quantum procedure, the algorithm will implement a unitary dependent on the $\omega(x)$ values, hence the algorithm does not need to be time efficient.

We can show that the success probability of the hybrid algorithm is at least the average of the success probabilities of the classical stage and of the quantum stage. In some special cases, such as the Bernoulli distribution, both the classical probability (i.e., at least one success in τ_c Bernoulli trials) and the weights $w(x)$ (hence the quantum success probability) are easy to derive. We can show that the hybrid algorithm gives matching lower bounds to the hardness bounds proven in the first part of our work.

Discussion and Directions for Future Work. We believe that the hybrid query model is both of theoretical and practical importance. Since near-term quantum computers are limited and expensive, it is to the interest of a party to supplement it with massive classical computational power. This also reflects the fact that those parties who have early access to quantum computers (e.g., large tech companies and government agencies) largely coincide with those who are capable of employing classical clusters and supercomputers. Next, we discuss some future directions.

One immediate question is to study other problems in the hybrid query model. The work of [HLS22] proves the hardness of the collision problem by their generalized recording technique in the hybrid query model. It would be useful to further develop techniques and establish more query complexity results.

Our applications to hash functions and Bitcoin-like blockchains can be seen as analyzing cryptographic constructions in the QRO model against hybrid adversaries. Many block ciphers rely on a different model, known as the *ideal cipher* model. As a simple example, the Even-Mansour cipher encrypts a message m by $E_k : m \mapsto \sigma(k \oplus m) \oplus k$, where σ is a random permutation given as an oracle and k is the secret key. As it turns out, this classically secure cipher is completely broken when quantum queries are allowed to both E_k and σ [KM10]. Since the secret key k is managed by honest users, it is debatable whether superposition access to E_k is realistic, and there has been progress in re-establishing the cipher's security under a partially quantum adversary with quantum access to σ but classical access to E_k [JST21, ABKM22]. The hybrid query model we consider in this work suggests further relaxing the queries to σ to be a hybrid of classical and quantum ones, and it would be valuable to re-examine the security of such schemes in the ideal cipher model.

Querying an oracle also occurs more broadly in many other cryptographic scenarios. Security definitions often give some algorithm as an oracle to the adversary, such as an encryption oracle in the chosen-plaintext-attack (CPA) game, and a signing oracle in formalizing the unforgeability of digital signatures. There has been a considerable effort of settling appropriate definitions and constructions (e.g., quantum-accessible pseudorandom functions, encryption and signatures) when quantum adversaries are granted superposition queries to these oracles (cf. [BZ13, Zha15, AMRS20, Zha21, CEV23]). Extending such efforts to the hybrid-adversary landscape would offer fine-grained security assessments of post-quantum cryptosystems. Finally, in the context of complexity theory, the study of hybrid algorithms is further motivated by related models focusing on the interplay between classical computation and near-future quantum devices [CCHL22], and between circuit depth and quantum queries [SZ19, CM20, CCL23].

Organization of the Paper. The rest of the paper is organized as follows. The generalized search problem we are considering, which we call *Distributional Search*, is stated in Sect. 2, together with its hybrid quantum-classical hardness; two case studies: Multi-Uniform Search and Bernoulli Search; as well as our proposed hybrid search algorithm. Detailed proofs and analyses of our main results above are presented in Sect. 3—hardness in Sect. 3.1 and the quantum algorithm analysis in Sect. 3.2, respectively. Due to space constraints, the applications of Bernoulli Search, as well as some of the proofs are presented in the full version of the paper [CGS23].

2 Problem Definition(s) and Main Results

2.1 The Distributional Search Problem

The underlying problem we consider is the search for a preimage of 1 of an arbitrarily distributed black-box boolean function.

Distributional Search Problem (Dist-Search)

Let D be an arbitrary distribution supported on the function family $\mathcal{F} = \{f : X \to \{0, 1\}\}$.

Given: Black-box access to a function f drawn from distribution D.

Goal: Find x such that $f(x) = 1$ if there exists such an x.

It is not surprising that the problem's hardness is crucially influenced by the number of solutions *on average* under D; however, what is interesting about our study is that we can show a clean quantitative relation.

Let $f : X \to \{0, 1\}$ be an arbitrary function. We define the projector on the space spanned by the preimages of 1 as: $\pi_f \overset{\text{def}}{=} \sum_{x:f(x)=1} |x\rangle\langle x|$.

Denote by $\pi_f^{\perp} \overset{\text{def}}{=} \mathbb{1} - \pi_f$, and let D be a distribution on $\mathcal{F}$. We define the value that captures the *average* fraction of preimages of 1 as:

Definition 1 (ν_D)**.** *The average fraction of solutions in $\mathcal{F}$ is defined as:*

$$\nu_D \overset{\text{def}}{=} \sup_{\varphi:\|\varphi\|\leq 1} \left(\mathbb{E}_{f \leftarrow D} \|\pi_f \varphi\|^2 \right), \tag{1}$$

where $\|\varphi\|$ denotes the Euclidean norm of the quantum state φ.

Characterization of ν_D. To better understand the ν_D value, we now derive an alternative characterization. For simplicity, assume without loss of generality that the domain of our target functions is $X = [m] \overset{\text{def}}{=} \{1, ..., m\}$, for some positive integer m. We will write down the truth table to represent each $f : [m] \rightarrow \{0, 1\}$ as a bitstring $x \in \{0, 1\}^m$ and denote by x_i the i-th bit of x.

In this way, D becomes a distribution on $\{0, 1\}^m$, and we write $d_x \overset{\text{def}}{=} D(x)$ as the probability of sampling x from the distribution D. Then, from Definition 1, we can rewrite ν_D as:

$$\nu_D = \sup_{\varphi} \left(\mathbb{E}_{x \leftarrow D} \|\pi_x \varphi\|^2 \right) , \quad \text{where } \pi_x \overset{\text{def}}{=} \sum_{i : x_i = 1} |i\rangle \langle i| \ .$$

Let $\varphi := \sum_{i=1}^m \alpha_i |i\rangle$, with $\|\alpha\| \leq 1$. We have:

$$\begin{aligned} \nu_D &= \sup_{\alpha : \|\alpha\| \leq 1} \mathbb{E}_{x \leftarrow D} \sum_{i : x_i = 1} \alpha_i^2 \\ &= \sup_{\alpha : \|\alpha\| \leq 1} \sum_{i=1}^m \alpha_i^2 \cdot \sum_{x \in \{0,1\}^m} d_x \cdot x_i \\ &= \sup_{\alpha : \|\alpha\| \leq 1} \sum_{i=1}^m \alpha_i^2 \cdot \omega_i \ , \end{aligned}$$

where, for each $i \in [m]$, we define $\omega_i \overset{\text{def}}{=} \sum_{x \in \{0,1\}^m} d_x \cdot x_i$. In other words, ω_i captures the likelihood that x_i is assigned value 1 under D. Then it becomes clear that the supremum is achieved by a vector α having 0 entries except taking 1 on i^* where ω_{i^*} is maximized: $\sup_{\alpha : \|\alpha\| \leq 1} \sum_{i=1}^m \alpha_i^2 \cdot \omega_i \leq \sup_{\alpha : \|\alpha\| \leq 1} \sum_{i=1}^m \alpha_i^2 \cdot \omega_{i^*} = \omega_{i^*} \sup_{\alpha : \|\alpha\| \leq 1} \alpha_i^2 \leq \omega_{i^*}$. Therefore,

$$\nu_D = \omega_{i^*} = \max_{i \in [m]} \omega_i \ . \tag{2}$$

We also note that for any $i \in [m]$, ω_i / ω, where $\omega \overset{\text{def}}{=} \sum_i \omega_i$, can be viewed as the probability[1] that $x_i = 1$, when x is sampled according to D.

2.2 Hardness of Dist-Search

Next, we turn to establishing the following bound for the success probability of solving Dist-Search, which constitutes one of our main results:

Theorem 1 (Hardness of Dist-Search – fixed query order). *For any algorithm $\mathcal{A}$ making up to τ_c classical queries and τ_q quantum queries (with a fixed order of the queries independent of f), $\mathcal{A}$ solves the* Dist-Search *problem with probability:*

$$\mathsf{Succ}_{\mathcal{A},D} = \Pr_{f \leftarrow D}[f(x) = 1 : x \leftarrow \mathcal{A}^f] \leq \nu_D \cdot (2\sqrt{\tau_c} + 2\tau_q + 1)^2 \ .$$

[1] We remark that normalization is required as ω might not be 1; for example, in the case of the Bernoulli distribution, $\omega = m\eta$. Hence, normalization is needed so as to view w_i / w as a probability distribution.

The proof can be found in Sect. 3.1. By relying on the result by Don *et al.* [DFH22], the above hardness result can be directly extended to any general hybrid algorithm in which the order of the classical and quantum queries can be adaptive (and can depend on the underlying oracle), at the cost of only a constant factor; i.e. increasing the number of classical and quantum queries by a factor of 2:

Theorem 2 (Hardness of Dist-Search). *For any algorithm $\mathcal{A}$ making τ_c classical queries and τ_q quantum queries, $\mathcal{A}$ solves the* Dist-Search *problem with probability:*

$$\mathsf{Succ}_{\mathcal{A},D} := \Pr_{f \leftarrow D}[f(x) = 1 : x \leftarrow \mathcal{A}^f] \leq \nu_D \cdot (2\sqrt{2\tau_c} + 4\tau_q + 1)^2 \,.$$

As this bound for general adversaries is directly derived from the hardness of hybrid algorithms with a fixed query order, in the sequel we will only focus on proving Theorem 1.

2.3 Case Studies

In this section, we will apply our hardness result to two common function distributions. As a common ingredient, it will be helpful to consider the following indicator random variable:

$$\mathbb{1}_x^f \overset{\text{def}}{=} \begin{cases} 1 & \text{, if } f(x) = 1\,; \\ 0 & \text{, if } f(x) = 0\,, \end{cases}$$

for all $f \in \mathcal{F}$ and $x \in X$. Then, for a distribution D:

$$\mathbb{E}_{f \leftarrow D}(\mathbb{1}_x^f) = \Pr_{f \leftarrow D}[f(x) = 1]\,.$$

2.3.1 Multi-uniform Search

The first interesting case is a general Grover-type search. We consider a distribution D_w which is *uniform* over functions that map exactly w inputs to 1. In other words, drawing $f \leftarrow D_w$ is equivalent to sampling a subset $S \subseteq X$ with $|S| = w$ uniformly at random and set $f(x) = 1$ if and only if $x \in S$. We consider the resulting multi-uniform search problem:

Multi-Uniform Search

Given: $f \leftarrow D_w$, which maps a uniform size-w subset to 1.
Goal: Find x such that $f(x) = 1$.

Theorem 3. *For any adversary $\mathcal{A}$ making τ_c classical queries and τ_q quantum queries,*

$$\mathsf{Succ}_{\mathcal{A},D_w} \leq \frac{w}{M} \cdot (2\sqrt{2\tau_c} + 4\tau_q + 1)^2 \,,$$

where $M = |X|$ is the domain size.

Proof. We just need to show that $\nu_D = \sup_{\varphi:\|\varphi\|\leq 1} \mathbb{E}_{f\leftarrow D_w}(\|\pi_f\varphi\|^2) \leq \frac{w}{M}$ in this case. Consider an arbitrary unit vector $\varphi = \sum_x \alpha_x |x\rangle$ with $\sum_x |\alpha_x|^2 = 1$.

$$\begin{aligned}\mathbb{E}_{f\leftarrow D_w}(\|\pi_f\varphi\|^2) &= \mathbb{E}_{f\leftarrow D_w}\left(\left|\sum_x \alpha_x \mathbb{1}_x^f |x\rangle\right|^2\right)\\ &= \sum_x |\alpha_x|^2 \cdot \mathbb{E}_{f\leftarrow D_w}(\mathbb{1}_x^f)\\ &= \sum_x |\alpha_x|^2 \cdot \Pr_{f\leftarrow D_w}[f(x)=1] = \frac{w}{M}\,.\end{aligned}$$

Alternatively, using the characterization of ν_D (Eq. 2), we can derive this result, by first noticing that for the multi-uniform distribution we have:

$$d_x = \begin{cases} \frac{1}{\binom{M}{w}} & \text{, if } \mathsf{hw}(x) = w;\\ 0 & \text{, otherwise.}\end{cases} \tag{3}$$

Then, by relying on the characterization of ν_D, we can directly conclude that:

$$\begin{aligned}\nu_D &= \max_{i\in[M]} \omega_i = \max_{i\in[M]} \sum_{x\in\{0,1\}^M} d_x \cdot x_i\\ &= \frac{1}{\binom{M}{w}} \cdot \sum_{x:\mathsf{hw}(x)=w} x_i\\ &= \frac{1}{\binom{M}{w}} \cdot \binom{M-1}{w-1} = \frac{w}{M}.\end{aligned} \tag{4}$$

□

Next, we note two special scenarios. When $w = 1$, our result reproduces Rosmanis's result [Ros22], and when $\tau_c = 0$, it reproduces the fully quantum query complexity of Grover search with multiple marked items (cf. [BBBV97, Zal99]).

2.3.2 Bernoulli Search The second interesting case we consider is what we call a Bernoulli distribution D_η on $\mathcal{F}$, as specified below:

Bernoulli Search

Given: $f \leftarrow D_\eta$ drawn via the following sampling procedure.

For each $x \in X$, *independently* set:

$$f(x) = \begin{cases} 1, & \text{with probability } \eta;\\ 0, & \text{otherwise.}\end{cases}\,.$$

Goal: Find x such that $f(x) = 1$.

Theorem 4. *For any adversary $\mathcal{A}$ making up to τ_c classical queries and τ_q quantum queries,*

$$\mathsf{Succ}_{\mathcal{A},D_\eta} \leq \eta \cdot (2\sqrt{2\tau_c} + 4\tau_q + 1)^2 \,.$$

Proof. Consider an arbitrary unit vector $\varphi = \sum_x \alpha_x |x\rangle$ with $\sum_x |\alpha_x|^2 = 1$. Again, we just need to show that $\mathbb{E}_{f\leftarrow D_\eta}(\|\pi_f \varphi\|^2) \leq \eta$. Similarly as before,

$$\mathbb{E}_{f\leftarrow D_\eta}(\|\pi_f \varphi\|^2) = \sum_x |\alpha_x|^2 \cdot \Pr_{f\leftarrow D_\eta}[f(x) = 1] = \eta \,.$$

Alternatively, using the characterization of ν_D (Eq. 2), we can derive this result directly by noting that every position is marked independently with probability η. Hence $\nu_D = \max_i w_i = \eta$. □

Note that when $\tau_c = 0$, this bound reproduces the complexity of Bernoulli Search using fully quantum queries (cf. [HRS16, ARU14]).

2.4 Designing Hybrid Search Algorithms

In the remaining of this section we propose a hybrid algorithm for the Dist-Search problem, analyze its success probability and show that in several relevant cases, the algorithm is optimal, hence leading to tight query complexity in the hybrid search model.

As a first step, we next describe a quantum search algorithm that, by adapting Grover's algorithm, takes into account a given distribution D.

2.4.1 Quantum Search Algorithm on $\boldsymbol{D}$ A main distinction from standard Grover is that the amplitudes in our initial state are proportional to the weights ω_i (capturing the likelihood that x_i is a solution under D), rather than a uniform superposition.

Quantum Search Algorithm $\mathcal{A}$ for an Arbitrary Distribution D

Given: $x \in \{0,1\}^m$ drawn from D.

Goal: Find $i \in [m]$ such that $x_i = 1$ making τ_q quantum queries to x.

Initialization: $\mathcal{A}$ constructs a unitary U_D such that

$$|\phi_0\rangle \overset{\text{def}}{=} U_D |0\rangle = \frac{1}{\sqrt{\omega}} \sum_i \sqrt{\omega_i} \, |i\rangle \ .$$

Modified Grover iteration: Repeatedly apply $G := R_0 R_x$, where

$$R_0 \overset{\text{def}}{=} -(\mathbb{1} - 2 |\phi_0\rangle \langle\phi_0|) \,,$$
$$R_x \overset{\text{def}}{=} \sum_i (-1)^{x_i} |i\rangle \langle i| \ .$$

Output: Measure the state in the computational basis and output the measurement outcome i.

Note that once U_D is available, $R_0 = -U_D(\mathbb{1} - |0\rangle\langle 0|)U_D^\dagger$ can be readily implemented, and one application of R_x can be realized by *one* query to x.

For any fixed x, we let ε_x denote the probability that $\mathcal{A}$ finds a solution (i.e., some i with $x_i = 1$); thus, $\varepsilon = \mathbb{E}_{x \leftarrow D}(\varepsilon_x)$ represents the success probability of $\mathcal{A}$ averaged over the distribution D. Next, we turn to lower-bounding this success probability; the proof is deferred to Sect. 3.2.

Theorem 5. *Algorithm $\mathcal{A}$ with τ_q quantum queries finds an i with $x_i = 1$ with probability:*

$$\varepsilon \geq \tau_q^2 \cdot \frac{\sum_i \omega_i^2}{\omega} .$$

2.4.2 A Hybrid Algorithm for Distributional Search

We are now ready to describe a hybrid algorithm equipped with τ_c classical queries and τ_q quantum queries. The basic idea is as follows: Given distribution D, let $S = \{i_1, ..., i_{\tau_c}\} \subseteq [m]$ be the set of indices with the τ_c largest values of ω_i. (In case of ties, we break them arbitrarily.) Our algorithm will first issue the τ_c classical queries on S to verify whether there exists an index $i \in S$ such that $x_i = 1$; if not, it will run the quantum search algorithm $\mathcal{A}$ from before, but on the reduced search space $[m] - S$.

In order to run the quantum algorithm in a modular fashion, we define an induced distribution $\tilde{D}$ on $\{0,1\}^{m-\tau_c}$. We will denote by x_T the substring of x of size $|T|$ obtained from concatenating the bits x_i for all $i \in T$, and by $\bar{S}$ the set defined as $\bar{S} \overset{\text{def}}{=} [m] - S$.

To define $\tilde{D}$, we first define $d \overset{\text{def}}{=} \sum_{x \in \{0,1\}^m : x_S = 0} d_x$. Then for each $\mathbf{x} \in \{0,1\}^{m-\tau_c}$, we define $\tilde{d}_{\mathbf{x}} \overset{\text{def}}{=} \frac{d_x}{d}$, where x is the unique string with $x_S = 0$ and $x_{\bar{S}} = \mathbf{x}$. Note that there is a fixed mapping that matches every index $\mathbf{i} \in \bar{S}$ with an index $i \in [m]$ such that $\mathbf{x_i} = 1$ if and only if $x_i = 1$. We assume that this mapping is performed implicitly whenever necessary. Therefore, for every $\mathbf{i} \in \bar{S}$, we can write the weight under $\tilde{D}$ as:

$$\tilde{\omega}_{\mathbf{i}} = \sum_{\mathbf{x} \in \{0,1\}^{m-\tau_c}} \tilde{d}_{\mathbf{x}} \cdot \mathbf{x_i} = \frac{\sum_{x : x_S = 0} d_x \cdot x_i}{\sum_{x : x_S = 0} d_x} .$$

Our hybrid algorithm can now be described as follows.

Hybrid Search Algorithm $\mathcal{A}_h$ for an Arbitrary Distribution D

Given: $x \in \{0,1\}^m$ drawn from D.

Goal: Find $i \in [m]$ such that $x_i = 1$ by making τ_c classical queries τ_q and quantum queries to x.

Classical Stage. $\mathcal{A}$ makes classical queries for each $i \in S$, where S, defined as above, consists of the indices with the τ_c largest ω_i. If some $x_i = 1$, output i and exit; otherwise, continue.

Quantum Stage. Run the quantum algorithm $\mathcal{A}$ on induced distribution $\tilde{D}$.

The algorithm's success probability can be split into analyzing the classical and quantum stages separately, as we show below. First, we define the following binary random variables:

- $Z_c^x = 1$ if and only if $x_i = 1$ for some $i \in S$ (i.e., the classical stage succeeds);
- $Z_q^x = 1$ if and only if the quantum stage is successful.

Lemma 1. *For any distribution D, the probability that hybrid algorithm $\mathcal{A}_h$ succeeds is:*

$$\Pr[\mathsf{Hybrid\ Success}] \geq \frac{1}{2}\left(\mathbb{E}_{x\leftarrow D}(Z_c^x) + \mathbb{E}_{x\leftarrow D}(Z_q^x)\right) .$$

Proof. The algorithm fails if both classical and quantum stages fail. Hence the failure probability is

$$\mathbb{E}_{x\leftarrow D}((1-Z_c^x)(1-Z_q^x)) = 1 - \mathbb{E}_{x\leftarrow D}(Z_c^x) - \mathbb{E}_{x\leftarrow D}(Z_q^x) + \mathbb{E}_{x\leftarrow D}(Z_c^x \cdot Z_q^x) .$$

Then, by using the Cauchy-Schwartz inequality (Lemma 5), and as Z_c^x and Z_q^x are both binary variables, we have

$$\begin{aligned} \mathbb{E}_{x\leftarrow D}(Z_c^x \cdot Z_q^x) &\leq \sqrt{\mathbb{E}_{x\leftarrow D}(Z_c^x) \cdot \mathbb{E}_{x\leftarrow D}(Z_q^x)} \\ &\leq \frac{1}{2}(\mathbb{E}_{x\leftarrow D}(Z_c^x) + \mathbb{E}_{x\leftarrow D}(Z_q^x)) . \end{aligned}$$

We can then conclude that the algorithm's success probability is

$$\Pr[\mathsf{Hybrid\ Success}] = 1 - \mathbb{E}_{x\leftarrow D}((1-Z_c^x)(1-Z_q^x)) \geq \frac{1}{2}\left(\mathbb{E}_{x\leftarrow D}(Z_c^x) + \mathbb{E}_{x\leftarrow D}(Z_q^x)\right) .$$

□

Applying Theorem 5, we can immediately give an expression for the quantum success probability. Namely:

$$\mathbb{E}_{x\leftarrow D}(Z_q^x) \geq \tau_q^2 \frac{\sum_{\mathbf{i}\in\bar{S}} \tilde{\omega}_{\mathbf{i}}^2}{\sum_{\mathbf{i}\in\bar{S}} \tilde{\omega}_{\mathbf{i}}} .$$

2.4.3 Success Probability for Special Distributions

We now show that for some special cases the hybrid algorithm above is optimal. We note that in these cases, the quantum stage actually coincides with the standard Grover search, and thus the quantum success probability can be obtained by the known result. Our analysis can be viewed as an alternative approach following the general result expressed by Theorem 5.

When $x \leftarrow D$ assigns a single i with $x_i = 1$ uniformly at random, $\tilde{D}$ can be seen as the same distribution but restricting to x with $x_S = 0$. For all $\mathbf{i} \in \bar{S}$, we have $\tilde{\omega}_{\mathbf{i}} = \frac{1}{m-c}$, and hence:

$$\mathbb{E}_{x \leftarrow D}(Z_q^x) = \tau_q^2 \cdot \frac{\sum_{\mathbf{i} \in \bar{S}} \tilde{\omega}_{\mathbf{i}}^2}{\sum_{\mathbf{i} \in \bar{S}} \tilde{\omega}_{\mathbf{i}}} = \tau_q^2 \frac{1}{m-c} .$$

It is also easy to observe that $\mathbb{E}_{x \leftarrow D}(Z_c^x) = \tau_c \frac{1}{m}$.

Lemma 2 (Uniform Search Hybrid Success and Optimality). *When D is the uniform distribution, our hybrid algorithm equipped with τ_c classical queries and τ_q quantum queries succeeds with probability at least*

$$\Pr[\mathsf{Hybrid\ Success}] \geq \frac{1}{2} \left(\frac{\tau_c}{m} + \frac{\tau_q^2}{m - \tau_c} \right) .$$

Except for constant factors and lower-order terms, this matches the hardness bound shown in Theorem 3, and hence the hybrid query complexity for the uniform distribution is $\Theta\left(\frac{1}{m}(\tau_c + \tau_q^2)\right)$.

Similarly, we can obtain a tight bound for the Bernoulli distribution, by the observation that $\tilde{D}$ in this case is just another Bernoulli distribution with the same η. Hence,

$$\mathbb{E}_{x \leftarrow D}(Z_q^x) = \eta \cdot \tau_q^2 .$$

On the other hand,

$$\mathbb{E}_{x \leftarrow D}(Z_c^x) = 1 - (1 - \eta)^{\tau_c} \geq \frac{1}{2} \eta \cdot \tau_c .$$

Lemma 3 (Bernoulli Search Hybrid Success and Optimality). *When D is the Bernoulli distribution, our hybrid algorithm equipped with τ_c classical queries and τ_q quantum queries succeeds with probability at least*

$$\Pr[\mathsf{Hybrid\ Success}] \geq \frac{1}{2} \eta \left(\frac{1}{2} \tau_c + \tau_q^2 \right) .$$

Again, except for constant factors and lower-order terms, this matches the hardness bound shown in Theorem 4, and hence the hybrid query complexity for the Bernoulli distribution is $\Theta(\eta(\tau_c + \tau_q^2))$.

3 Proofs of the Main Results

3.1 Hardness of Dist-Search

In this section we will pove the main hardness result stated in Theorem 1. For convenience, we restate it again here:

Theorem 1 (Hardness of Dist-Search – fixed query order). *For any algorithm $\mathcal{A}$ making up to τ_c classical queries and τ_q quantum queries (with a fixed order of the queries independent of f), it holds that $\mathcal{A}$ solves the* Dist-Search *problem with probability:*

$$\mathsf{Succ}_{\mathcal{A},D} := \Pr_{f \leftarrow D}[f(x) = 1 : x \leftarrow \mathcal{A}^f] \leq \nu_D \cdot (2\sqrt{\tau_c} + 2\tau_q + 1)^2 .$$

3.1.1 Preliminaries and Overview

We first formally describe an oracle function for the case of quantum and pseudo-classical queries.

Definition 2 (Query Operators). *We define the following operators, describing the actions of quantum and pseudo-classical oracles for a hybrid algorithm given a boolean function f.*

- *A pseudo-classical oracle is described by*

$$P_{f,b} \overset{\text{def}}{=} \sum_{x : f(x) = b} |x\rangle \langle x| \otimes \mathbb{1} \otimes |b\rangle$$

- *A quantum oracle is described by*

$$Q_f \overset{\text{def}}{=} \sum_{x,b} |x\rangle\langle x| \otimes \mathbb{1} \otimes |b \oplus f(x)\rangle \langle b|$$

We denote $\Pi_f \overset{\text{def}}{=} \pi_f \otimes \mathbb{1}$ ($\mathbb{1}$ operates on the output and ancilla registers) and $\Pi_f^\perp \overset{\text{def}}{=} \mathbb{1} - \Pi_f$ ($\mathbb{1}$ operates on the entire system). Then on a pseudo-classical query, the two operators $P_{f,0} = \Pi_f^\perp \otimes |0\rangle$ and $P_{f,1} = \Pi_f \otimes |1\rangle$ correspond to the two possible measurement outcomes. It is more convenient to answer quantum queries by the corresponding phase oracle:

$$Q_f \overset{\text{def}}{=} \mathbb{1} - 2\Pi_f .$$

This can be seen as setting the output register of the standard oracle in $|-\rangle$, and as a result, a quantum query flips the signs of the 1-preimages.

When running a hybrid query algorithm with f, we will keep track of the (sub-normalized) pure state $\psi_f^{(t)}$, which denotes the state of the algorithm on input f after t queries in the situation where every pseudo-classical query measures 0 (we will call this the 0-branch of $\mathcal{A}^f$). Namely, consider an arbitrary algorithm with at most τ queries (τ_q quantum and τ_c pseudo-classical) specified by a sequence of unitary operators[2] $(U^{(0)}, U^{(1)}, \ldots, U^{(\tau)})$. Let $T_c = \{t : t\text{-th query is pseudo-classical}\}$ and $T_q = \{t : t\text{-th query is quantum}\}$. Then $\psi_f^{(t)}$ is defined recursively by

$$\psi_f^{(t)} \overset{\text{def}}{=} \begin{cases} U^{(t)} P_{f,0} \psi_f^{(t-1)}, & \text{if } t \in T_c\,; \\ U^{(t)} Q_f \psi_f^{(t-1)} & \text{if } t \in T_q\,. \end{cases} \tag{5}$$

From this definition, the projection of $\psi_f^{(t)}$ under $\Pi_f^{\perp}$ characterizes the event that an algorithm fails to find a 1-preimage.

Lemma 4. *For any algorithm $\mathcal{A}$, the failure probability of finding a 1-preimage of f after t queries is*

$$\delta_f^{(t)} = \Pr[f(x) \neq 1 : x \leftarrow \mathcal{A}^f] \geq \left\| \Pi_f^{\perp} \psi_f^{(t)} \right\|^2 .$$

Hence, the failure probability with respect to distribution D satisfies

$$\delta_D^{(t)} = \mathbb{E}_{f \leftarrow D} \delta_f^{(t)} \geq \mathbb{E}_{f \leftarrow D} \left\| \Pi_f^{\perp} \psi_f^{(t)} \right\|^2 .$$

Thus, our goal becomes lower-bounding $\left\| \Pi_f^{\perp} \psi_f^{(t)} \right\|$. To do this, we consider running the same algorithm, but with a null function:

$$f_\emptyset : x \mapsto 0, \forall x \in X .$$

In this case, a quantum query is equivalent to applying identity (denoted $Q_\emptyset \overset{\text{def}}{=} \mathbb{1}$), and a pseudo-classical query does not tamper the input state either, but just appends $|0\rangle$. To be precise, we define

$$P_{\emptyset,0} \overset{\text{def}}{=} \mathbb{1} \otimes |0\rangle ,$$

and at each step $t \geq 0$, the state of the algorithm denoted by $\phi^{(t)}$ can be described as:

$$\phi^{(t)} = \begin{cases} U^{(t)} P_{\emptyset,0} \phi^{(t-1)}, & \text{if } t \in T_c\,; \\ U^{(t)} \phi^{(t-1)} & \text{if } t \in T_q\,. \end{cases}$$

Without loss of generality we assume initially $\psi_f^{(0)} = \phi^{(0)} = |0\rangle$, and hence $\left\| \Pi_f^{\perp} \psi_f^{(0)} \right\| = \left\| \Pi_f^{\perp} \phi^{(0)} \right\| = 1$. In order to succeed, algorithm $\mathcal{A}^f$ needs to move $\psi_f^{(t)}$ away from the kernel of $\Pi_f^{\perp}$ or reduce its norm. This motivates defining the progress measures below.

[2] Dimensions may grow depending on the arrangement of the pseudo-classical queries.

Table 1. Summary of variables and quantities used in our Dist-Search analysis.

π_f	$\sum_{x:f(x)=1} \|x\rangle \langle x\|$
Π_f	$\pi_f \otimes \mathbb{1}$ ($\mathbb{1}$ on ancilla registers)
δ_f	$\Pr[f(x) \neq 1 : x \leftarrow \mathcal{A}^f]$ (Failure probability with fixed f)
δ_D	$\mathbb{E}_D(\delta_f)$ (Failure probability with $f \leftarrow D$)
$\phi^{(0)} = \psi^{(0)}$	Initial state
$\phi^{(t)}$	State after t-th query in $\mathcal{A}^{f_\emptyset}$
$\psi_f^{(t)}$	State on the 0-branch after t-th query in $\mathcal{A}^f$
Q_f	$\mathbb{1} - 2\Pi_f$ (quantum oracle of f)
$Q_\emptyset$	$\mathbb{1}$ (quantum oracle of $f_\emptyset$)
$P_{f,0}$	$\Pi_f^\perp \otimes \|0\rangle$ (pseudo-classical oracle of f)
$P_{f,1}$	$\Pi_f \otimes \|1\rangle$ (pseudo-classical oracle of f)
$P_{\emptyset,0}$	$\mathbb{1} \otimes \|0\rangle$ (pseudo-classical oracle of $f_\emptyset$)
$\gamma_f^{(t)}$	$\left\lVert \Pi_f \phi^{(t)} \right\rVert^2$
$\gamma^{(t)}$	$\mathbb{E}_D(\gamma_f^{(t)})$

Definition 3 (Progress Measures). *For any function f and $t \geq 0$, define*

$$A_f^{(t)} \overset{\text{def}}{=} \left| \langle \phi^{(t)}, \psi_f^{(t)} \rangle \right|^2 , \quad B_f^{(t)} \overset{\text{def}}{=} \left\lVert \psi_f^{(t)} \right\rVert^2 - \left| \langle \phi^{(t)}, \psi_f^{(t)} \rangle \right|^2 .$$

Given a distribution D on $\mathcal{F}$, define the expected progress measures by

$$A_D^{(t)} \overset{\text{def}}{=} \mathbb{E}_{f \leftarrow D} \left(A_f^{(t)} \right) , \quad B_D^{(t)} \overset{\text{def}}{=} \mathbb{E}_{f \leftarrow D} \left(B_f^{(t)} \right) .$$

Notice that:

$$A_f^{(t)} + B_f^{(t)} = \left\lVert \psi_f^{(t)} \right\rVert^2 , \quad A_f^{(0)} = 1, \quad B_f^{(0)} = 0 .$$

We will show that $A_D^{(t)} - B_D^{(t)}$ essentially lower bounds the failure probability $\delta_D^{(t)}$ (Lemma 8). Hence, an algorithm's objective would be to *reduce* $A_D^{(t)}$ and *increase* $B_D^{(t)}$. However, we can limit how much change can occur after τ queries (Proposition 1). This is by carefully analyzing the effect of each quantum or pseudo-classical query (Lemmas 10 and 11). Roughly speaking,

- A quantum query reduces $A_D^{(t)}$ by at most $4\sqrt{\nu_D \cdot B_D^{(t)}}$ and increases $B_D^{(t)}$ by the same amount (as a quantum query does not affect $\left\lVert \psi_f^{(t)} \right\rVert^2$), and
- A pseudo-classical query increases $B_D^{(t)}$ by at most ν_D, while a part $z^{(t)}$ of $B_D^{(t)}$ can also be spent to decrease $A_D^{(t)}$ by $\sqrt{\nu_D \cdot z^{(t)}}$ (Table 1).

3.1.2 Proof of Theorem 1 First off, we state the Cauchy-Schwarz inequality for random variables and derive a corollary that is useful in several places.

Lemma 5 (Cauchy-Schwarz). *For any random variables X, Y, it holds that:* $|\mathbb{E}(XY)|^2 \leq \mathbb{E}(X^2) \cdot \mathbb{E}(Y^2)$.

Corollary 1. *Let Z be a discrete random variable, and $g(Z)$ and $h(Z)$ be two non-negative functions. Then it holds that:* $\mathbb{E}_Z\left(\sqrt{g(Z) \cdot h(Z)}\right) \leq \sqrt{\mathbb{E}_Z(g(Z)) \cdot \mathbb{E}_Z(h(Z))}$.

It will be helpful to consider a two-dimensional plane in our analysis, which we now define explicitly.

Definition 4 (Useful 2-D Plane). *For $t \geq 0$, let*

$$\phi_f^{(t)} \overset{\text{def}}{=} \frac{\Pi_f \phi^{(t)}}{\|\Pi_f \phi^{(t)}\|} = \Pi_f \phi^{(t)} / \sqrt{\gamma_f^{(t)}}\,, \qquad \phi_f^{(t)\perp} \overset{\text{def}}{=} \frac{\Pi_f^\perp \phi^{(t)}}{\left\|\Pi_f^\perp \phi^{(t)}\right\|} = \Pi_f^\perp \phi^{(t)} / \sqrt{1-\gamma_f^{(t)}}$$

be the normalized *vectors resulting of projecting $\phi^{(t)}$ on the orthogonal subspaces spanned by* 1 *and* 0 *preimages of f, respectively, and let $\Phi^{(t)}$ be the 2-dimensional plane spanned by $\{\phi_f^{(t)}, \phi_f^{(t)\perp}\}$. Then $\phi^{(t)\perp}$ is identified as the normalized state perpendicular to $\phi^{(t)}$ in $\Phi^{(t)}$, i.e.,*

$$\phi^{(t)\perp} \overset{\text{def}}{=} \phi_f^{(t)} \sqrt{1-\gamma_f^{(t)}} - \phi_f^{(t)\perp} \sqrt{\gamma_f^{(t)}}\,.$$

It is useful to decompose $\psi_f^{(t)}$ with respect to $\Phi^{(t)}$:

Lemma 6 (Decomposition of $\psi_f^{(t)}$ wrt $\Phi^{(t)}$). *Let a and b be projecting $\psi_f^{(t)}$ on the plane $\Phi^{(t)}$ and then decomposing it under basis $\{\phi^{(t)}, \phi^{(t)\perp}\}$, and let c be the remaining component of $\psi_f^{(t)}$ orthogonal to $\Phi^{(t)}$, i.e., $c \perp \Phi^{(t)}$. Then $\psi_f^{(t)}$ can be expressed as $\psi_f^{(t)} = a + b + c$ with*

$$a = \phi^{(t)} \sqrt{A_f^{(t)}}\,, \qquad b = \omega \sqrt{B_f^{(t)} - \|c\|^2} \cdot \phi^{(t)\perp}\,,$$

where ω is a complex phase ($|\omega| = 1$) of the vector $\psi_f^{(t)} - \langle \psi_f^{(t)}, \phi_f^{(t)} \rangle \cdot \phi^{(t)} - c$. Thus,

$$\Pi_f^\perp \psi_f^{(t)} = \phi_f^{(t)\perp} \left(\sqrt{1-\gamma_f^{(t)}} \sqrt{A_f^{(t)}} - \sqrt{\gamma_f^{(t)}} \cdot \omega \sqrt{B_f^{(t)} - \|c\|^2} \right) + c_f^\perp\,,$$

with $c_f^\perp := \Pi_f^\perp c$.

Intuitively, for the next result, the goal is to relate the failure probability with the progress measures A and B. To do so, we will first relate the failure probability with the norm of the non-solution component. By decomposing this norm in terms of the two progress measures A and B and an orthogonal component which can be removed, we can determine a lower bound on the failure probability as a function of the two progress measure after each performed query.

Lemma 7. *For any fixed f and $t \geq 0$,*

$$\delta_f^{(t)} \geq A_f^{(t)} - \gamma_f^{(t)} - 2\sqrt{\gamma_f^{(t)} \cdot B_f^{(t)}}\,.$$

Proof. For convenience, we omit writing the superscript (t) in this proof. We first show that $\left\|\pi_f^{\perp}\psi_f\right\| \geq \sqrt{(1-\gamma_f)A_f} - \sqrt{\gamma_f B_f}$. By Lemma 6, we have that

$$\Pi_f^{\perp}\psi_f = \phi_f^{\perp}\left(\sqrt{1-\gamma_f}\sqrt{A_f} - \sqrt{\gamma_f}\cdot\omega\sqrt{B_f - \|c\|^2}\right) + c_f^{\perp}\,,$$

with $c_f^{\perp} := \pi_f^{\perp}c$. Since $c \perp \Phi$, it follows that

$$\langle \phi_f^{\perp}, c_f^{\perp}\rangle = \langle \phi_f^{\perp}, \Pi_f^{\perp}c\rangle = \langle \Pi_f^{\perp}\phi_f^{\perp},\ \mathrm{c}\rangle = \langle \phi_f^{\perp}, c\rangle = 0\,.$$

We can then obtain:

$$\left\|\Pi_f^{\perp}\psi_f\right\| = \left|\sqrt{1-\gamma_f}\cdot\sqrt{A_f} - \sqrt{\gamma_f}\cdot\omega\sqrt{B_f - \|c\|^2}\right| + \left\|c_f^{\perp}\right\|$$

Hence by choosing $c = 0, \omega = 1$, we get: $\left\|\Pi_f^{\perp}\psi_f\right\| \geq \sqrt{(1-\gamma_f)A_f} - \sqrt{\gamma_f B_f}$. Therefore we can lower bound the failure probability:

$$\begin{aligned}\delta_f &\geq \left\|\pi_f^{\perp}\psi_f\right\|^2 \geq (1-\gamma_f)A_f - 2\sqrt{(1-\gamma_f)\gamma_f B_f}\\ &\geq A_f - \gamma_f - 2\sqrt{\gamma_f B_f} \qquad (A_f, \gamma_f \leq 1)\end{aligned}$$

□

Taking the expectation over D, we can express the failure probability with respect to the distribution.

Lemma 8. *For any distribution D and $t \geq 0$,*

$$\delta_D^{(t)} \geq A^{(t)} - \gamma^{(t)} - 2\sqrt{\gamma^{(t)} \cdot B^{(t)}}\,.$$

Proof.

$$\begin{aligned}\delta_D^{(t)} &= \mathbb{E}_{f\leftarrow D}(\delta_f^{(t)})\\ &\geq \mathbb{E}_D(A_f^{(t)}) - \mathbb{E}_D(\gamma_f^{(t)}) - 2\mathbb{E}_D\left(\sqrt{\gamma_f^{(t)}\cdot B_f^{(t)}}\right) \qquad \text{(Linearity of expectation)}\\ &\geq A^{(t)} - \gamma^{(t)} - 2\sqrt{\mathbb{E}_D(\gamma_f^{(t)})\cdot\mathbb{E}_D(B_f^{(t)})} \qquad \text{(Corollary1)}\\ &= A^{(t)} - \gamma^{(t)} - 2\sqrt{\gamma^{(t)}\cdot B^{(t)}}\end{aligned}$$

□

We can also relate $\gamma^{(t)}$ to the value ν_D determined by the distribution D:

Lemma 9. *For any $t \geq 0$ and any distribution D, we have: $\gamma^{(t)} \leq \nu_D$.*

Proof.

$$\gamma^{(t)} := \mathbb{E}_{f \leftarrow D}\left(\left\|\Pi_f \phi^{(t)}\right\|^2\right) = \mathbb{E}_{f \leftarrow D}\left(\left\|(\pi_f \otimes \mathbb{1})\phi^{(t)}\right\|^2\right)$$

We write $\phi^{(t)} = \sum_i \alpha_i |u_i\rangle \otimes |v_i\rangle$ under the Schmidt decomposition, where $\alpha_i \geq 0$ such that $\sum_i \alpha_i^2 = 1$ are the Schmidt coefficients, and $\{|u_i\rangle\}$ are orthonormal states on the system of the input register and $\{|v_i\rangle\}$ are orthonormal states on the system of output and ancilla registers. Then we can rewrite $\gamma^{(t)}$ as:

$$\begin{aligned}
\gamma^{(t)} &:= \mathbb{E}_{f \leftarrow D}\left(\left\|(\pi_f \otimes \mathbb{1})\phi^{(t)}\right\|^2\right) = \mathbb{E}_{f \leftarrow D}\left(\left\|(\pi_f \otimes \mathbb{1})\left(\sum_i \alpha_i |u_i\rangle \otimes |v_i\rangle\right)\right\|^2\right) \\
&= \mathbb{E}_{f \leftarrow D}\left(\left\|\sum_i \alpha_i(\pi_f |u_i\rangle) \otimes |v_i\rangle\right\|^2\right) \\
&= \mathbb{E}_{f \leftarrow D}\left(\sum_i \alpha_i^2 \left\|(\pi_f |u_i\rangle) \otimes |v_i\rangle\right\|^2\right) \qquad (\,|v_i\rangle \text{ are orthogonal}) \\
&= \mathbb{E}_{f \leftarrow D}\left(\sum_i \alpha_i^2 \left\|\pi_f |u_i\rangle\right\|^2 \cdot \left\||v_i\rangle\right\|^2\right) \qquad (\,\|a \otimes b\| = \|a\| \cdot \|b\|\,) \\
&= \mathbb{E}_{f \leftarrow D}\left(\sum_i \alpha_i^2 \left\|\pi_f |u_i\rangle\right\|^2\right) = \sum_i \alpha_i^2 \cdot \mathbb{E}_{f \leftarrow D}\left(\left\|\pi_f |u_i\rangle\right\|^2\right) \\
&\leq \sum_i \alpha_i^2 \nu_D \qquad (\text{definition of } \nu_D) \\
&= \nu_D \sum_i \alpha_i^2 = \nu_D
\end{aligned}$$

Proposition 1 (Bounding Progress Measures). *After $\tau = \tau_c + \tau_q$ queries,*

$$A^{(\tau)} \geq 1 - 4\nu_D \cdot (\sqrt{\tau_c} + \tau_q)^2\,, \quad B^{(\tau)} \leq \nu_D \cdot (\sqrt{\tau_c} + 2\tau_q)^2\,.$$

Proving Proposition 1 is the most involved step technically speaking. We present the details separately in Sect. 3.1.3 and here we apply it to prove Theorem 1.

Proof of Theorem 1. Assuming the bounds above on the two progress measures, we obtain that:

$$\begin{aligned}
\delta^{(\tau)} &\geq 1 - 4\gamma^{(\tau)} \cdot (\sqrt{\tau_c} + \tau_q)^2 - \gamma^{(\tau)} - 2\gamma^{(t)} \cdot (\sqrt{\tau_c} + 2\tau_q) \qquad (\text{Proposition } 1) \\
&= 1 - \gamma^{(\tau)} \cdot (4(\sqrt{\tau_c} + \tau_q) + 2\sqrt{\tau_c} + 4\tau_q + 1) \\
&\geq 1 - \gamma^{(\tau)} \cdot (2(\sqrt{\tau_c} + \tau_q) + 1)^2 \qquad (\tau_c \geq 0) \\
&\geq 1 - \nu_D \cdot (2\sqrt{\tau_c} + 2\tau_q + 1)^2 \qquad (\gamma^{(\tau)} \leq \nu_D \text{ Lemma } 9)
\end{aligned}$$

Therefore,

$$\mathsf{Succ}_{\mathcal{A},D} \leq 1 - \delta^{(\tau)} \leq \nu_D \cdot (2\sqrt{\tau_c} + 2\tau_q + 1)^2 \,.$$

□

3.1.3 Bounding the Progress Measures (Proposition 1)

We repeat the proposition statement for convenience here:

Proposition 1 (Bounding the Progress Measures). *After $\tau = \tau_c + \tau_q$ queries,*

$$A^{(\tau)} \geq 1 - 4\nu_D \cdot (\sqrt{\tau_c} + \tau_q)^2 \,, \quad B^{(\tau)} \leq \nu_D \cdot (\sqrt{\tau_c} + 2\tau_q)^2 \,.$$

Firstly, we will consider a fixed function f, and bound how much each query can possibly reduce $A_f^{(t)}$ and increase $B_f^{(t)}$.

Lemma 10 (Progress Measures for a Fixed Function). *For every t the progress measures after the $t+1$-th query satisfy the following recurrent relations:*

- *If the $t+1$-th query is* pseudo-classical*, then there exists a sequence $\left(z_f^{(t)}\right)_{t\geq 0}$, satisfying $0 \leq z_f^t \leq B_f^{(t)}$, such that:*

$$\begin{aligned} A_f^{(t+1)} &\geq A_f^{(t)} - 2\gamma_f^{(t)} - 2 \cdot \sqrt{z_f^{(t)}} \cdot \sqrt{\gamma_f^{(t)}} \\ B_f^{(t+1)} &\leq B_f^{(t)} + \gamma_f^{(t)} - z_f^{(t)} \end{aligned} \tag{6}$$

- *If the $t+1$-th query is* quantum*, then:*

$$\begin{aligned} A_f^{(t+1)} &\geq A_f^{(t)} - 4\gamma_f^{(t)} - 4 \cdot \sqrt{B_f^{(t)}} \cdot \sqrt{\gamma_f^{(t)}} \\ B_f^{(t+1)} &\leq B_f^{(t)} + 4\gamma_f^{(t)} + 4 \cdot \sqrt{B_f^{(t)}} \cdot \sqrt{\gamma_f^{(t)}} \end{aligned} \tag{7}$$

Proof. The proof can be found in the full version of the paper [CGS23]. □

Lemma 11 (Progress Measures for Dist-Search). *For every t, the progress measures after the $t+1$-th query satisfy the following recurrent relations:*

- *If the $t+1$-th query is pseudo-classical, there exists $z_t \in [0, B^{(t)}]$ such that:*

$$\begin{aligned} A^{(t+1)} &\geq A^{(t)} - 2\nu_D - 2\sqrt{\nu_D} \cdot \sqrt{z_t} \\ B^{(t+1)} &\leq B^{(t)} - z_t + \nu_D \end{aligned} \tag{8}$$

- *If the $t+1$-th query is quantum, then we have:*

$$\begin{aligned} A^{(t+1)} &\geq A^{(t)} - 4 \cdot \nu_D - 4 \cdot \sqrt{\nu_D} \cdot \sqrt{B^{(t)}} \\ B^{(t+1)} &\leq B^{(t)} + 4 \cdot \nu_D + 4 \cdot \sqrt{\nu_D} \cdot \sqrt{B^{(t)}} \end{aligned} \tag{9}$$

Proof. Letting $z_t \overset{\text{def}}{=} \mathbb{E}_{f \leftarrow D}(z_f^t)$, we can observe that $z_t \in [0, B^{(t)}]$. Taking expectations over D, and applying Corollary 1 ($\mathbb{E}(\sqrt{g(Z) \cdot h(Z)}) \leq \sqrt{\mathbb{E}(g(Z)) \cdot \mathbb{E}(h(Z))}$) and Lemma 9 ($\gamma^{(t)} \leq \nu_D$), the relations for $A^{(t)}$ and $B^{(t)}$ follow. □

Next, since we intend to lower bound $A^{(\tau)}$ and upper bound $B^{(\tau)}$, we can change the inequalities to equalities and analyze instead the new sequences (a_t, b_t) defined below. It is clear that $A^{(\tau)} \geq a_\tau$ and $B^{(\tau)} \leq b_\tau$.

Definition 5 (Sequences $(a_t)_{t\geq 0}$,$(b_t)_{t\geq 0}$). *We define the following sequences based on the evolution of the progress measures A and B:*

$$a_0 \overset{\text{def}}{=} A^{(0)} = 1 \quad ; \quad b_0 \overset{\text{def}}{=} B^{(0)} = 0$$

$$a_{t+1} \overset{\text{def}}{=} \begin{cases} a_t - 2 \cdot \nu_D - 2 \cdot \sqrt{\nu_D} \cdot \sqrt{z_t}\,, & \textit{if } t+1 \in T_c \\ a_t - 4 \cdot \nu_D - 4 \cdot \sqrt{\nu_D} \cdot \sqrt{b_t}\,, & \textit{if } t+1 \in T_q \end{cases}$$

$$b_{t+1} \overset{\text{def}}{=} \begin{cases} b_t + \nu_D - z_t\,, & \textit{if } t+1 \in T_c \\ b_t + 4 \cdot \nu_D + 4 \cdot \sqrt{\nu_D} \cdot \sqrt{b_t}\,, & \textit{if } t+1 \in T_q \end{cases}$$

where $(z_t)_{t\geq 1}$ is the sequence defined in the proof of Lemma 11, which satisfies $0 \leq z_t \leq B^{(t)}$ for any t.

Lemma 12 (Bounding a_τ and b_τ).

$$a_\tau \geq 1 - 4\nu_D \cdot (\sqrt{\tau_c} + \tau_q)^2\,, \qquad b_\tau \leq \nu_D \cdot (\sqrt{\tau_c} + 2\tau_q)^2\,. \tag{10}$$

Proof. The proof consists of four steps.

(1) First we show that $b_\tau \leq \left(\sqrt{\tau_c} + 2\tau_q\right)^2 \cdot \nu_D$.

To get an upper bound for each term of this sequence, we can let $z_t = 0$ and instead consider the sequence:

$$d_{t+1} \overset{\text{def}}{=} \begin{cases} d_t + \nu_D\,, & \text{if } t+1 \in T_c \\ d_t + 4 \cdot \nu_D + 4 \cdot \sqrt{\nu_D} \cdot \sqrt{d_t}\,, & \text{if } t+1 \in T_q \end{cases}$$

As a result we have: $b_t \leq d_t$ for any $t \in [\tau]$.

Our task is to bound the last term d_τ in the sequence. Every hybrid strategy A that uses τ_c classical queries and τ_q quantum queries can be expressed by $A = [x_1, \cdots, x_\tau]$, where if $x_i = 0$ (resp. $x_i = 1$) indicates that the i-th query of A is classical (resp. quantum), and there are exactly τ_c values of 0 and τ_q values of 1. Therefore, the sequence $(d_t)_t$ parameterized by the strategy A, denoted as $(d_t^A)_t$, can be re-written as:

$$d_{t+1}^A \overset{\text{def}}{=} \begin{cases} d_t^A + \nu_D\,, & \text{if } x_{t+1} = 0 \\ d_t^A + 4 \cdot \nu_D + 4 \cdot \sqrt{\nu_D} \cdot \sqrt{d_t}\,, & \text{if } x_{t+1} = 1 \end{cases} \tag{11}$$

Our task then becomes determining the strategy A^* which achieves the maximum $d_\tau^{A^*}$. We claim that

$$A^* \stackrel{\text{def}}{=} [0, \cdots, 0, 1, \cdots, 1],$$

namely the strategy of making all classical queries upfront is optimal. This follows from a greedy argument.

Consider two arbitrary strategies $A = [x_1, \cdots, x_i, x_{i+1}, \cdots, x_\tau]$ and $B = [y_1, \cdots, y_i, y_{i+1}, \cdots, y_\tau]$ which only differ in the i and $i+1$-th queries. Namely, $x_i = 0$, $x_{i+1} = 1$ and $y_i = 1$, $y_{i+1} = 0$ and $x_j = y_j$ for $j \in \{1, \cdots, \tau\} - \{i, i+1\}$. We next show that $d_\tau^A > d_\tau^B$. As $x_1 = y_1, \cdots x_{i-1} = y_{i-1}$, this implies directly that $d_{i-1}^A = d_{i-1}^B$. Then for the i-th and $i+1$ terms of the two sequences we have:

$$\begin{aligned}
d_i^A &= d_{i-1}^A + \nu_D & ; \quad d_{i+1}^A &= d_{i-1}^A + 5\nu_D + 4\sqrt{\nu_D}\sqrt{d_{i-1}^A + \nu_D} \\
d_i^B &= d_{i-1}^B + 4\nu_D + 4\sqrt{\nu_D}\sqrt{d_{i-1}^B} & ; \quad d_{i+1}^B &= d_{i-1}^B + 5\nu_D + 4\sqrt{\nu_D}\sqrt{d_{i-1}^B}
\end{aligned}$$

Then, as $d_{i-1}^A = d_{i-1}^B$ it is clear that $d_{i+1}^A > d_{i+1}^B$. As $x_j = y_j$ for all $i+2 \leq j \leq \tau$, this also implies that $d_\tau^A > d_\tau^B$.

Denote the following swap operation on strategies. Given as input a strategy $A = [x_1, ..., x_i, x_{i+1}, \cdots, x_\tau]$ the function swap_i outputs a strategy A':

$$\mathsf{swap}_i(A) = A' \text{ where } A' = [x_1, ..., x_{i+1}, x_i, \cdots, x_\tau]$$

Our previous argument implies that for a strategy A such that $x_i = 0$ and $x_{i+1} = 1$, we have: $d_\tau^A > d_\tau^{\mathsf{swap}_i(A)}$. Therefore, we can see that any strategy $A = [x_1, ..., x_\tau]$ can be obtained from a sequence of applications of swap_i on A^*.

$$A^* \stackrel{\text{def}}{=} [0, \cdots, 0, 1, \cdots, 1] \xrightarrow{\mathsf{swap}_{i_1}} \cdots \xrightarrow{\mathsf{swap}_{i_k}} A \text{ for some indices } i_1, ..., i_k\,.$$

It hence follows that $d_\tau^{A^*} \geq d_\tau^A$, i.e., A^* is the optimal strategy.

Now, let us compute the last term of the optimal strategy, i.e.: $d_\tau^{A^*}$. We can rewrite the sequence d_t as:

$$d_{t+1}^{A^*} = \begin{cases} d_t^{A^*} + \nu_D\,, & \text{if } 0 \leq t < \tau_c \\ d_t^{A^*} + 4 \cdot \nu_D + 4 \cdot \sqrt{\nu_D} \cdot \sqrt{d_t^{A^*}} = \left(\sqrt{d_t^{A^*}} + 2\sqrt{\nu_D}\right)^2, & \text{if } \tau_c \leq t < \tau \end{cases}$$

As $d_0^{A^*} = 0$, it is clear that we have: $d_{\tau_c}^{A^*} = \tau_c \cdot \nu_D$. For $\tau_c \leq t \leq \tau$, we will prove by induction that:

$$d_t^{A^*} = (\sqrt{\tau_c} + 2(t - \tau_c))^2 \cdot \nu_D$$

For the base case $t = \tau_c$, we already showed that $d_{\tau_c}^{A^*} = \tau_c \cdot \nu_D$. For the inductive step, we have that:

$$d_{t+1}^{A^*} = \left(\sqrt{(\sqrt{\tau_c} + 2(t - \tau_c))^2 \cdot \nu_D} + 2\sqrt{\nu_D}\right)^2 = (\sqrt{\tau_c} + 2(t - \tau_c + 1)) \cdot \nu_D$$

which concludes the inductive proof. Hence, by putting things together:

$$b_\tau \leq d_\tau \leq d_\tau^{A^*} = \left(\sqrt{\tau_c} + 2\tau_q\right)^2 \cdot \nu_D \tag{12}$$

(2) Secondly, we show that $\sum_{t\in T_q} \sqrt{b_{t-1}} \leq \sqrt{\nu_D} \cdot \tau_q(\sqrt{\tau_c} + \tau_q - 1)$.

As for b_τ, to get an upper bound we let $z_t = 0$ and use the sequence $(d_t^A)_t$. From the definition of the sequence (Eq. 11), it is clear that $(d_t^A)_t$ is a strictly increasing sequence for any strategy A. This also implies that for any strategy A we have:

$$\sum_{t\in T_q} \sqrt{d_{t-1}^A} \leq \sum_{\tau_c \leq t \leq \tau} \sqrt{d_t^A}$$

In other words, $\sum_{t\in T_q} \sqrt{d_{t-1}^A}$ is maximized when the strategy performs first all τ_c classical queries and then the τ_q quantum queries. Hence, the maximum is achieved for the strategy described above by the sequence $(d_t^{A^*})_t$.

Using the previous result in Eq. 12:

$$\sum_{\tau_c \leq t \leq \tau} d_t^{A^*} = \nu_D \cdot \sum_{\tau_c \leq t \leq \tau} \left(\sqrt{\tau_c} + 2(t-\tau_c)\right)^2$$

This gives us:

$$\begin{aligned}\sum_{t\in T_q} \sqrt{b_{t-1}} \leq \sum_{\tau_c \leq t \leq \tau} \sqrt{d_t^{A^*}} &= \sqrt{\nu_D} \sum_{\tau_c \leq t \leq \tau} \sqrt{\tau_c} + 2(t-\tau_c) \\ &\leq \sqrt{\nu_D} \left(\tau_q(\sqrt{\tau_c} - 2\tau_c) + 2 \sum_{\tau_c \leq t \leq \tau} t\right) \\ &= \sqrt{\nu_D}\tau_q(\sqrt{\tau_c} + \tau_q - 1)\end{aligned}$$

(3) Thirdly, we show that $\sum_{t\in T_c} \sqrt{z_{t-1}} \leq \sqrt{\nu_D} \cdot (\tau_c + 2\sqrt{\tau_c}\tau_q)$.

By definition of the sequence z_t (Definition 5), we know that for $t \in T_c$:

$$\sum_{t\in T_c} z_{t-1} = \nu_D \cdot \tau_c + \sum_{t\in T_c} (b_{t-1} - b_t)$$

Thus it suffices to derive an upper bound on $\sum_{t\in T_c}(b_{t-1} - b_t)$. We rewrite b_τ as:

$$b_\tau = b_0 + \sum_{t=1}^{\tau}(b_t - b_{t-1}) = \sum_{b_t \geq b_{t-1}} (b_t - b_{t-1}) + \sum_{b_t < b_{t-1}} (b_t - b_{t-1})$$

As a result, we have that:

$$\sum_{t\in T_c \ \wedge \ b_t < b_{t-1}} (b_{t-1} - b_t) < \sum_{b_t < b_{t-1}} (b_{t-1} - b_t) = \sum_{b_t \geq b_{t-1}} (b_t - b_{t-1}) - b_\tau$$

In other words we also have:

$$\sum_{t\in T_c \ \wedge \ b_t < b_{t-1}} (b_{t-1} - b_t) < \sum_{t\in T_c \ \wedge \ b_t \geq b_{t-1}} (b_t - b_{t-1}) + \sum_{t\in T_q \ \wedge \ b_t \geq b_{t-1}} (b_t - b_{t-1})$$

For $t \in T_q$, from sequence definition (Definition 5), we have that $b_t > b_{t-1}$ and hence:

$$\sum_{t\in T_c \ \wedge \ b_t < b_{t-1}} (b_{t-1} - b_t) < \sum_{t\in T_c \ \wedge \ b_t \geq b_{t-1}} (b_t - b_{t-1}) + 4\tau_q \cdot \nu_D + 4\sqrt{\nu_D} \sum_{t\in T_q} \sqrt{b_{t-1}}$$

By applying step (2), we get:

$$\sum_{t\in T_c \ \wedge \ b_t < b_{t-1}} (b_{t-1} - b_t) < \sum_{t\in T_c \ \wedge \ b_t \geq b_{t-1}} (b_t - b_{t-1}) + 4\nu_D\tau_q + 4\nu_D\tau_q(\sqrt{\tau_c} + \tau_q - 1)$$

By subtracting the first sum from the right hand side we get:

$$\sum_{t\in T_c} z_{t-1} = \nu_D \cdot \tau_c + \sum_{t\in T_c} (b_{t-1} - b_t) < \nu_D \cdot \left(\tau_c + 4\tau_q^2 + 4\tau_q\sqrt{\tau_c}\right)$$

Finally, by using the Cauchy-Schwarz inequality:

$$\sum_{t\in T_c} \sqrt{z_{t-1}} \leq \sqrt{\nu_D \cdot \left(\tau_c + 4\tau_q^2 + 4\tau_q\sqrt{\tau_c}\right)} \cdot \sqrt{\tau_c} \leq \sqrt{\nu_D} \cdot (\tau_c + 2\tau_q\sqrt{\tau_c})$$

(4) In the final step, we show that $a_\tau \geq 1 - 4\nu_D(\sqrt{\tau_c} + \tau_q)^2$.
From the definition of a_t (Definition 5):

$$\begin{aligned} a_\tau &= a_0 + \sum_{t=1}^{\tau} (a_t - a_{t-1}) \\ &= 1 - \sum_{t\in T_c} \left(2\nu_D + 2\sqrt{\nu_D} \cdot \sqrt{z_{t-1}}\right) - \sum_{t\in T_q} \left(4\nu_D + 4\sqrt{\nu_D} \cdot \sqrt{b_{t-1}}\right) \\ &= 1 - 2\tau_c\nu_D - 4\tau_q\nu_D - 2\sqrt{\nu_D} \sum_{t\in T_c} \sqrt{z_{t-1}} - 4\sqrt{\nu_D} \sum_{t\in T_q} \sqrt{b_{t-1}} \end{aligned}$$

Using the bounds derived in steps (2) and (3), we get :

$$\begin{aligned} a_\tau &\geq 1 - 2\tau_c\nu_D - 4\tau_q\nu_D - 2\nu_D \cdot (\tau_c + 2\sqrt{\tau_c}\tau_q) - 4\nu_D \cdot \tau_q(\sqrt{\tau_c} + \tau_q - 1) \\ &= 1 - 4\nu_D(\sqrt{\tau_c} + \tau_q)^2 \end{aligned}$$

□

3.2 Quantum Algorithm Analysis

In this section we will prove the success probability of our proposed quantum algorithm described in Sect. 2.4.1.

Theorem 5. *Algorithm $\mathcal{A}$ with τ_q quantum queries finds an i with $x_i = 1$ with probability:*

$$\varepsilon \geq \tau_q^2 \cdot \frac{\sum_i \omega_i^2}{\omega}.$$

Proof. We adapt the geometric analysis of standard Grover's algorithm to analyze $\mathcal{A}$. First for any x, define two states below:

$$|A_x\rangle := \frac{1}{\sqrt{\alpha_x}} \sum_{i:x_i=1} \sqrt{\omega_i}\,|i\rangle \ , \quad |B_x\rangle := \frac{1}{\sqrt{\beta_x}} \sum_{i:x_i=0} \sqrt{\omega_i}\,|i\rangle \ ,$$

with normalization factors

$$\alpha_x := \sum_{i:x_i=1} \omega_i = \sum_i \omega_i x_i, \quad \text{and} \quad \beta_x := \sum_{i:x_i=0} \omega_i = \sum_i \omega_i(1-x_i)\,.$$

We will focus on the two dimensional plane spanned by $|A_x\rangle$ and $|B_x\rangle$. Observe that ϕ_0 belongs to this plane, and can be decomposed under the basis $\{|A_x\rangle, |B_x\rangle\}$:

$$|\phi_0\rangle := \sin\theta\,|A_x\rangle + \cos\theta\,|B_x\rangle \ , \text{where:}$$

$$\sin^2\theta = |\langle\phi_0|A_x\rangle|^2 = \frac{1}{\omega\cdot\alpha_x}(\sum_i \omega_i x_i)^2 = \frac{\alpha_x}{\omega}\,.$$

We then show that on the two dimensional plane, R_0 is a reflection about $|\phi_0\rangle$ and R_x is a reflection $|B_x\rangle$. We introduce a state $|\phi_0^\perp\rangle$ on the plane orthogonal to $|\phi_0\rangle$, which can be written as

$$|\phi_0^\perp\rangle = \cos\theta\,|A_x\rangle - \sin\theta\,|B_x\rangle \ .$$

Clearly $\{\phi_0, \phi_0^\perp\}$ forms another basis on the plane, under which we can express $|A_x\rangle$ and $|B_x\rangle$ as below.

$$|A_x\rangle = \sin\theta\,|\phi_0\rangle + \cos\theta\,|\phi_0^\perp\rangle \quad , \quad |B_x\rangle = \cos\theta\,|\phi_0\rangle - \sin\theta\,|\phi_0^\perp\rangle \ .$$

It then becomes easy to verify that

$$R_0\,|A_x\rangle = \sin\theta\,|\phi_0\rangle - \cos\theta\,|\phi_0^\perp\rangle \quad , \quad R_0\,|B_x\rangle = \cos\theta\,|\phi_0\rangle + \sin\theta\,|\phi_0^\perp\rangle \ .$$

Hence R_0 reflects about ϕ_0. Similarly, R_x reflects about $|B_x\rangle$ as shown below.

$$R_x\,|\phi_0\rangle = -\sin\theta\,|A_x\rangle + \cos\theta\,|B_x\rangle \quad , \quad R_0\,|\phi_0^\perp\rangle = -\sin\theta\,|\phi_0\rangle - \cos\theta\,|\phi_0^\perp\rangle \ .$$

As a consequence, $G = R_0 R_x$ composes two reflections and effectively amounts to an rotation of 2θ. Therefore, after τ_q iterations, the state becomes

$$|\phi_{\tau_q}\rangle := \sin((2\tau_q+1)\theta)\,|A_x\rangle + \cos((2\tau_q+1)\theta)\,|B_x\rangle \ .$$

This is illustrated in Fig. 1.

When measuring $|\phi_{\tau_q}\rangle$, an outcome i with $x_i = 1$ occurs with probability

$$\varepsilon_x = \sin^2((2\tau_q+1)\theta) \geq \left(\frac{2\tau_q+1}{2}\theta\right)^2 \geq \tau_q^2 \sin^2\theta = \tau_q^2\frac{\alpha_x}{\omega}\,.$$

Thus:

$$\varepsilon = \mathbb{E}_{x\leftarrow D}\varepsilon_x \geq \tau_q^2 \frac{\mathbb{E}_x \alpha_x}{\omega} = \tau_q^2 \frac{\sum_i \omega_i \sum_x d_x x_i}{\omega} = \tau_q^2 \frac{\sum_i \omega_i^2}{\omega}\,.$$

□

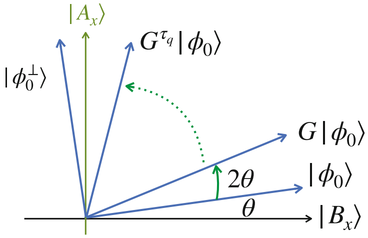

Fig. 1. Illustration of the evolution in the two-dimensional plane.

Optimality for Permutation-Invariant Distributions. Consider a special family of distributions, where ω_i are identical for all $i \in [m]$ implying that every i is mapped to 1 with equal probability. We call such a distribution D *permutation invariant*, and in this case our quantum algorithm $\mathcal{A}$ becomes identical to the standard Grover's algorithm. It also follows immediately Eq. (2) that for any $i, \omega_i = \nu_D$. Therefore we obtain that

$$\frac{\sum_i \omega_i^2}{\omega} = \frac{\sum_i \omega_i^2}{\sum_i \omega_i} = \frac{m\nu_D^2}{m\nu_D} = \nu_D \,.$$

As a result, quantum algorithm $\mathcal{A}$ succeeds with probability $\Omega(\tau_q^2 \nu_D)$ in the case of permutation-invariant distribution, which is in turn *optimal* by our hardness bound (Theorem 1). This also reproves the tight quantum query complexity for multi-uniform search and Bernoulli search. We summarize it below.

Corollary 2. *For a permutation-invariant distribution D, the quantum algorithm $\mathcal{A}$ coincides with the standard Grover's algorithm, and it succeeds with probability $\Omega(\tau_q^2 \cdot \nu_D)$ with τ_q quantum queries which is* tight.

In particular, multi-uniform search and Bernoulli search have tight quantum query complexity $\Theta(\tau_q^2 \frac{w}{m})$ and $\Theta(\tau_q^2 \eta)$ for quantum algorithms with τ_q queries.

Acknowledgements. J.G. was partially supported by NSF SaTC grants no. 2001082 and 2055694. F.S. was partially supported by NSF grant no. 1942706 (CAREER). J.G. and F.S. were also partially support by Sony by means of the Sony Research Award Program. A.C. acknowledges support from the National Science Foundation grant CCF-1813814, from the AFOSR under Award Number FA9550-20-1-0108 and the support of the Quantum Advantage Pathfinder (QAP) project, with grant reference EP/X026167/1 and the UK Engineering and Physical Sciences Research Council.

References

[ABKM22] Gorjan Alagic, Chen Bai, Jonathan Katz, and Christian Majenz. Post-quantum security of the even-mansour cipher. In *Advances in Cryptology – EUROCRYPT 2022*, pages 458–487. Springer, 2022.

[AHU19] Andris Ambainis, Mike Hamburg, and Dominique Unruh. Quantum security proofs using semi-classical oracles. In *Advances in Cryptology – CRYPTO 2019*, pages 269–295. Springer, 2019.

[AMRS20] Gorjan Alagic, Christian Majenz, Alexander Russell, and Fang Song. Quantum-secure message authentication via blind-unforgeability. In *Advances in Cryptology – EUROCRYPT 2020*. Springer, 2020.

[ARU14] Andris Ambainis, Ansis Rosmanis, and Dominique Unruh. Quantum attacks on classical proof systems: The hardness of quantum rewinding. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 474–483. IEEE, 2014.

[BBBV97] Charles H Bennett, Ethan Bernstein, Gilles Brassard, and Umesh Vazirani. Strengths and weaknesses of quantum computing. *SIAM journal on Computing*, 26(5):1510–1523, 1997.

[BDF+11] Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. Random oracles in a quantum world. In *Advances in Cryptology – ASIACRYPT 2011*, pages 41–69. Springer, 2011.

[BR93] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *Proceedings of the 1st ACM conference on Computer and Communications Security*, pages 62–73, 1993.

[BR94] Mihir Bellare and Phillip Rogaway. Optimal asymmetric encryption. In *Advances in Cryptology–EUROCRYPT 1994*, pages 92–111. Springer, 1994.

[BR96] Mihir Bellare and Phillip Rogaway. The exact security of digital signatures-how to sign with rsa and rabin. In *Advances in Cryptology–Eurocrypt 1996*, pages 399–416. Springer, 1996.

[BZ13] Dan Boneh and Mark Zhandry. Secure signatures and chosen ciphertext security in a quantum computing world. In *Advances in Cryptology – CRYPTO 2013*, pages 361–379. Springer, 2013.

[CCHL22] Sitan Chen, Jordan Cotler, Hsin-Yuan Huang, and Jerry Li. The complexity of nisq, 2022.

[CCL23] Nai-Hui Chia, Kai-Min Chung, and Ching-Yi Lai. On the need for large quantum depth. *J. ACM*, 70(1), jan 2023.

[CEV23] Céline Chevalier, Ehsan Ebrahimi, and Quoc-Huy Vu. On security notions for encryption in a quantum world. In *Progress in Cryptology – INDOCRYPT 2022*, pages 592–613. Springer, 2023.

[CGK+23] Alexandru Cojocaru, Juan Garay, Aggelos Kiayias, Fang Song, and Petros Wallden. Quantum Multi-Solution Bernoulli Search with Applications to Bitcoin's Post-Quantum Security. *Quantum*, 7:944, 2023.

[CGS23] Alexandru Cojocaru, Juan Garay, and Fang Song. Generalized hybrid search and applications. Cryptology ePrint Archive, Paper 2023/798, 2023.

[CM20] Matthew Coudron and Sanketh Menda. Computations with greater quantum depth are strictly more powerful (relative to an oracle). In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2020, page 889-901, New York, NY, USA, 2020. Association for Computing Machinery.

[CMS19] Alessandro Chiesa, Peter Manohar, and Nicholas Spooner. Succinct arguments in the quantum random oracle model. In *17th International Theory of Cryptography Conference – TCC 2019*, pages 1–29. Springer, 2019.

[DFH22] Jelle Don, Serge Fehr, and Yu-Hsuan Huang. Adaptive versus static multi-oracle algorithms, and quantum security of a split-key prf. In Eike Kiltz and Vinod Vaikuntanathan, editors, *Theory of Cryptography*, pages 33–51, Cham, 2022. Springer Nature Switzerland.

[DFMS19] Jelle Don, Serge Fehr, Christian Majenz, and Christian Schaffner. Security of the Fiat-Shamir transformation in the quantum random-oracle model. In *Advances in Cryptology – CRYPTO 2019*, pages 356–383. Springer, 2019.

[DFMS22] Jelle Don, Serge Fehr, Christian Majenz, and Christian Schaffner. Online-extractability in the quantum random-oracle model. In *Advances in Cryptology – EUROCRYPT 2022*, pages 677–706. Springer, 2022.

[DH09] Cătălin Dohotaru and Peter Høyer. Exact quantum lower bound for grover's problem. *Quantum Information & Computation*, 9(5):533–540, 2009.

[ES15] Edward Eaton and Fang Song. Making Existential-unforgeable Signatures Strongly Unforgeable in the Quantum Random-oracle Model. In *10th Conference on the Theory of Quantum Computation, Communication and Cryptography – TQC 2015*, volume 44 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 147–162. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2015.

[ES20] Edward Eaton and Fang Song. A note on the instantiability of the quantum random oracle. In *International Conference on Post-Quantum Cryptography*, pages 503–523. Springer, 2020.

[FO13] Eiichiro Fujisaki and Tatsuaki Okamoto. Secure integration of asymmetric and symmetric encryption schemes. *Journal of Cryptology*, 26(1):80–101, 2013. Preliminary version in CRYPTO 1999.

[FOPS04] Eiichiro Fujisaki, Tatsuaki Okamoto, David Pointcheval, and Jacques Stern. RSA-OAEP is secure under the rsa assumption. *Journal of Cryptology*, 17(2):81–104, 2004. Preliminary version in CRYPTO 2001.

[Gro96] Lov K Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 212–219. ACM, 1996.

[HHK17] Dennis Hofheinz, Kathrin Hövelmanns, and Eike Kiltz. A modular analysis of the fujisaki-okamoto transformation. In *15th International Theory of Cryptography Conference – TCC 2017*, pages 341–371. Springer, 2017.

[HLS22] Yassine Hamoudi, Qipeng Liu, and Makrand Sinha. Quantum-classical tradeoffs in the random oracle model. *CoRR*, abs/2211.12954, 2022.

[HRS16] Andreas Hülsing, Joost Rijneveld, and Fang Song. Mitigating multi-target attacks in hash-based signatures. In *19th IACR International Conference on Public-Key Cryptography — PKC 2016*, pages 387–416. Springer, 2016.

[JST21] Joseph Jaeger, Fang Song, and Stefano Tessaro. Quantum key-length extension. In *19th International Theory of Cryptography Conference – TCC 2021*, pages 209–239. Springer, 2021.

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

[Pre18] John Preskill. Quantum computing in the NISQ era and beyond. *Quantum*, 2:79, 2018.

[Ros22] Ansis Rosmanis. Hybrid quantum-classical search algorithms. arXiv preprint arXiv:2202.11443, 2022.

[Sho01] Victor Shoup. OAEP reconsidered. In *Advances in Cryptology–CRYPTO 2001*, pages 239–259. Springer, 2001.

[SZ19] Xiaoming Sun and Yufan Zheng. Hybrid decision trees: Longer quantum time is strictly more powerful, 2019.

[Unr15] Dominique Unruh. Non-interactive zero-knowledge proofs in the quantum random oracle model. In *Advances in Cryptology – EUROCRYPT 2015*, pages 755–784. Springer, 2015.

[YZ21] Takashi Yamakawa and Mark Zhandry. Classical vs quantum random oracles. In *Advances in Cryptology – EUROCRYPT 2021*, pages 568–597. Springer, 2021.

[Zal99] Christof Zalka. Grover's quantum searching algorithm is optimal. *Physical Review A*, 60(4):2746, 1999.

[Zha15] Mark Zhandry. Secure identity-based encryption in the quantum random oracle model. *International Journal of Quantum Information*, 13(04):1550014, 2015. Preliminary version in IACR CRYPTO 2012.

[Zha19] Mark Zhandry. How to record quantum queries, and applications to quantum indifferentiability. In *Advances in Cryptology – CRYPTO 2019*, pages 239–268. Springer, 2019.

[Zha21] Mark Zhandry. How to construct quantum random functions. *Journal of the ACM (JACM)*, 68(5):1–43, 2021. Preliminary version in FOCS 2012.