Bounds for Error Reduction with Few Quantum Queries

Sourav Chakraborty[1], Jaikumar Radhakrishnan[2,3], and Nandakumar Raghunathan[1]

[1] Department of Computer Science,
University of Chicago, Chicago,
IL 60637, USA
e-mail:{sourav,nanda}@cs.uchicago.edu
[2] Toyota Technological Institute at Chicago,
IL 60637, USA
e-mail: jaikumar@tti-c.org
[3] School of Technology and Computer Science,
Tata Institute of Fundamental Research, Mumbai 400005, India

Abstract. We consider the quantum database search problem, where we are given a function $f : [N] \to \{0,1\}$, and are required to return an $x \in [N]$ (a target address) such that $f(x) = 1$. Recently, Grover [G05] showed that there is an algorithm that after making one quantum query to the database, returns an $X \in [N]$ (a random variable) such that

$$\Pr[f(X) = 0] = \epsilon^3,$$

where $\epsilon = |f^{-1}(0)|/N$. Using the same idea, Grover derived a t-query quantum algorithm (for infinitely many t) that errs with probability only ϵ^{2t+1}. Subsequently, Tulsi, Grover and Patel [TGP05] showed, using a different algorithm, that such a reduction can be achieved for all t. This method can be placed in a more general framework, where given any algorithm that produces a target state for some database f with probability of error ϵ, one can obtain another that makes t queries to f, and errs with probability ϵ^{2t+1}. For this method to work, we do not require prior knowledge of ϵ. Note that no classical randomized algorithm can reduce the error probability to significantly below ϵ^{t+1}, even if ϵ is known. In this paper, we obtain *lower bounds* that show that the amplification achieved by these quantum algorithms is essentially optimal. We also present simple alternative algorithms that achieve the same bound as those in Grover [G05], and have some other desirable properties. We then study the best reduction in error that can be achieved by a t-query quantum algorithm, when the initial error ϵ is known to lie in an interval of the form $[\ell, u]$. We generalize our basic algorithms and lower bounds, and obtain nearly tight bounds in this setting.

1 Introduction

In this paper, we consider the problem of reducing the error in quantum search algorithms by making a small number of queries to the database. Error reduction in the form of amplitude amplification is one of the central tools in the design of efficient quantum search algorithms [G98a,G98b,BH+02]. In fact, Grover's database search algorithm [G96,G97] can be thought of as amplitude amplification applied to the trivial algorithm that queries the database at a random location and succeeds with probability at least $\frac{1}{N}$. The key feature of quantum amplitude amplification is that it can boost the success probability from a small quantity δ to a constant in $O(1/\sqrt{\delta})$ steps, whereas, in general a classical algorithm for this would require $\Omega(1/\delta)$ steps. This basic algorithm has been refined, taking into account the number of solutions and the desired final success probability $1-\epsilon$. For example, Buhrman, Cleve, de Wolf and Zalka [BC+99] obtained the following:

> **Theorem [BC+99]:** Fix $\eta \in (0,1)$, and let $N > 0$, $\epsilon \geq 2^{-N}$, and $t \leq \eta N$. Let T be the optimal number of queries a quantum computer needs to search with error $\leq \epsilon$ through an unordered list of N items containing at least t solutions. Then $\log 1/\epsilon \in \Theta(T^2/N + T\sqrt{t/N})$ (Note that the constant implicit in the Θ notation can depend on η).

Recently, Grover [G05] considered error reduction for algorithms that err with small probability. The results were subsequently refined and extended by Tulsi, Grover and Patel [TGP05]. Let us describe their results in the setting of the database search problem, where, given a database $f : [N] \rightarrow \{0,1\}$, we are asked to determine an $x \in f^{-1}(1)$. If $|f^{-1}(0)| = \epsilon N$, then choosing x uniformly at random will meet the requirements with probability at least $1-\epsilon$. This method makes no queries to the database. If one is allowed one classical query, the error can be reduced to ϵ^2 and, in general, with t classical queries one can reduce the probability of error to ϵ^{t+1}. It can be shown that no classical t-query randomized algorithm for the problem can reduce the probability of error significantly below ϵ^{t+1}, even if the value of ϵ is known in advance. Grover [G05] presented an interesting algorithm that makes one quantum query and returns an x that is in $f^{-1}(1)$ with probability $1-\epsilon^3$. Tulsi, Grover and Patel [TGP05] showed an iteration where one makes just one query to the database and performs a measurement, so that after t iterations of this operator the error is reduced to ϵ^{2t+1}. This algorithm works for all ϵ and is not based on knowing ϵ in advance. Thus this iteration can be said to exhibit a "fixed point" behavior [G05,TGP05], in that the state approaches the target state (or subspace) closer with each iteration, just as it does in randomized classical search. The iteration used in the usual Grover search algorithm [G98a,G98b] does not have this property. Note, however, that if the initial success probability is $\frac{1}{N}$, these new algorithms make $\Omega(N)$ queries to the database, whereas the original algorithm makes just $O(\sqrt{N})$ queries.

In [G05], the database is assumed to be presented by means of an oracle of the form $|x\rangle \rightarrow \exp(f(x)\pi i/3))|x\rangle$. The standard oracle for a function f used in earlier works on quantum search is $|x\rangle|b\rangle \mapsto |x\rangle|b \oplus f(x)\rangle$, where $\oplus$ is addition modulo two. It can be shown that the oracle assumed in [G05] cannot be implemented by using just one query to the standard oracle. In Tulsi, Grover and Patel [TGP05] the basic iteration uses the controlled version of the oracle, namely $|x\rangle|b\rangle|c\rangle \mapsto |x\rangle|b \oplus c \cdot f(x)\rangle|c\rangle$.

In this paper, we present a version of the algorithm that achieves the same reduction in error as in [G05], but uses the standard oracle. In fact, our basic one-query algorithm has the following natural interpretation. First, we note that for $\delta \leq \frac{3}{4}$, there is a one-query quantum algorithm $\mathcal{A}_\delta$ that makes no error if $|f^{-1}(0)| = \delta N$. Then, using a simple computation, one can show that the one-query algorithm corresponding to $\delta = \frac{1}{N}$

errs with probability less than ϵ^3 when $|f^{-1}(0)| = \epsilon N$. One can place these algorithms in a more general framework, just as later works due to Grover [G98a,G98b] and Brassard, Hoyer, Mosca and Tapp [BH+02] placed Grover's original database search algorithm [G96,G97] in the general amplitude amplification framework. The framework is as follows: Suppose there is an algorithm G that guesses a solution to a problem along with a witness, which can be checked by another algorithm T. If the guess returned by G is correct with probability $1 - \epsilon$, then there is another algorithm that uses G, G^{-1}, makes t queries to T, and guesses correctly with probability $1 - \epsilon^{2t+1}$.

These algorithms show that, in general, a t-query quantum algorithm can match the error reduction obtainable by any $2t$-query randomized algorithm. Can one do even better? The main contribution of this paper are the *lower bounds* on the error probability of t-query algorithms. We show that the amplification achieved by these algorithms is essentially optimal (see Section 2.1 for the precise statement). Our result does not follow immediately from the result of Buhrman, Cleve, de Wolf and Zalka [BC+99] cited above because of the constants implicit in the θ notation, but with a slight modification of their proofs one can derive a result similar to ours (see Section 2.1). Our lower bound result uses the polynomial method of Beals, Cleve, Buhrman, Mosca and de Wolf [BB+95] combined with an elementary analysis based on the roots of low degree polynomials, but unlike previous proofs using this method, we do not rely on any special tools for bounding the rate of growth of low degree polynomials.

2 Background, definitions and results

We first review the standard framework for quantum search. We assume that the reader is familiar with the basics of quantum circuits, especially the quantum database search algorithm of Grover [G96,G97] (see, for example, Nielsen and Chuang [NC, Chapter 6]). The database is modelled as a function $f : [N] \to S$, where $[N] \stackrel{\Delta}{=} \{0, 1, 2, \ldots, N-1\}$ is the set of addresses and S is the set of possible items to be stored in the database. For our purposes, we can take S to be $\{0, 1\}$. When thinking of bits we identify $[N]$ with $\{0, 1\}^n$. Elements of $[N]$ will be called addresses, and addresses in $f^{-1}(1)$ will be referred to as targets. In the quantum model, the database is provided to us by means of an oracle unitary transformation T_f, which acts on an $(n + 1)$-qubit space by sending the basis vector $|x\rangle|b\rangle$ to $|x\rangle|b \oplus f(x)\rangle$. For a quantum circuit $\mathcal{A}$ that makes queries to a database oracle in order to determine a target, we denote by $\mathcal{A}(f)$ the random variable (taking values in $[N]$) returned by $\mathcal{A}$ when the database oracle is T_f.

Definition 1. *Let $\mathcal{A}$ be a quantum circuit for searching databases of size N. For a database f of size N, let $\mathsf{err}_{\mathcal{A}}(f) = \Pr[\mathcal{A}(f)$ is not a target state$]$. When ϵN is an integer in $\{0, 1, 2, \ldots, N\}$, let $\mathsf{err}_{\mathcal{A}}(\epsilon) = \max_{f : |f^{-1}(0)| = \epsilon N} \mathsf{err}_{\mathcal{A}}(f)$.*

Using this notation, we can state Grover's result as follows.

Theorem 1 (Grover [G05]). *For all N, there is a one-query algorithm $\mathcal{A}$, such that for all ϵ (assuming ϵN is an integer), $\mathsf{err}_{\mathcal{A}}(\epsilon) = \epsilon^3$.*

This error reduction works in a more general setting. Let $[N]$ represent the set of possible solutions to some problem, and let $f : [N] \to \{0, 1\}$ be the function that checks that the solution is correct; as before we will assume that we are provided access to this function via the oracle T_f. Let G be a unitary transform that guesses a solution in $[N]$ that is correct with probability $1 - \epsilon$. Our goal is to devise another guessing algorithm $\mathcal{B}$ that using T_f, G and G^{-1} produces a guess that is correct with significantly better probability. Let $\mathcal{B}(T_f, G)$ be the answer returned by $\mathcal{B}$ when the checker is T_f and the guesser is G.

Theorem 2 (Grover [G05]). *There is an algorithm $\mathcal{B}$ that uses T_f once, G twice and G^{-1} once, such that $\Pr[f(\mathcal{B}(T_f, G)) = 0] = \epsilon^3$, where ϵ is the probability of error of the guessing algorithm G.*

Note that Theorem 1 follows from Theorem 2 by taking G to be the Hadamard transformation, which produces the uniform superposition on all N states when applied to the state $|0\rangle$.

Theorem 3 ([G05,?]). *For all $t \geq 0$ and all N, there is a t-query quantum database search algorithm such that, for all ϵ (ϵN is an integer), $\mathsf{err}_{\mathcal{A}}(\epsilon) = \epsilon^{2t+1}$.*

In Grover [G05], this result was obtained by recursive application of Theorem 2, and worked only for infinitely many t. Tulsi, Grover and Patel [TGP05] rederived Theorems 1 and 2 using a different one-query algorithm, which could be applied iteratively to get Theorem 3.

From now on when we consider error reduction for searching a database f and use the notation $\mathsf{err}_{\mathcal{A}}(\epsilon)$, ϵ will refer to $|f^{-1}(0)|/N$; in particular, we assume that $\epsilon N \in \{0, 1, \ldots, N - 1\}$. However, for the general framework, ϵ can be any real number in $[0, 1]$.

2.1 Our contributions

As stated earlier, in order to derive the above results, Grover [G05] and Tulsi, Grover and Patel [TGP05] assume that access to the database is available using certain special types of oracles. In the next section, we describe alternative algorithms that establish Theorem 1 while using only the standard oracle $T_f : |x\rangle|b\rangle \to |x\rangle|b \oplus f(x)\rangle$. The same idea can be used to obtain results analogous to Theorems 2. By recursively applying this algorithm we can derive a version of Theorem 3 for t of the form $\frac{3^i-1}{2}$ where i is the number of recursive applications. Our algorithms and those in Tulsi, Grover and Patel [TGP05] use similar ideas, but were obtained independently of each other.

We also consider error reduction when we are given a lower bound on the error probability ϵ, and obtain analogs of Theorems 1 and 2 in this setting.

Theorem 4 (Upper bound result).

(a) For all N and $\delta \in [0, \frac{3}{4}]$, there is a one-query algorithm A_δ such that for all $\epsilon \geq \delta$,

$$\mathsf{err}_{\mathcal{A}_\delta}(f) \leq \epsilon \left[\frac{\epsilon - \delta}{1 - \delta}\right]^2.$$

(b) *For all $\delta \in [0, \frac{3}{4}]$, there is an algorithm $\mathcal{B}_\delta$ that uses T_f once and G twice and G^{-1} once, such that*

$$\mathsf{err}_{\mathcal{B}_\delta}(T_f, G) \le \epsilon \left[\frac{\epsilon - \delta}{1 - \delta}\right]^2 .$$

The case $\epsilon = \delta$ corresponds to the fact that one can determine the target state with certainty if $|f^{-1}(0)|$ is known exactly and is at most $\frac{3N}{4}$. Furthermore, Theorems 1 and 2 can be thought of as special cases of the above Proposition corresponding to $\delta = 0$. In fact, by taking $\delta = \frac{1}{N}$ in the above proposition, we obtain the following slight improvement over Theorem 1.

Corollary 1. *For all N, there is a one-query database search algorithm $\mathcal{A}$ such that for all ϵ (where $\epsilon N \in \{0, 1, \ldots, N\}$), we have* $\mathsf{err}_{\mathcal{A}}(\epsilon) \le \epsilon \left[\frac{\epsilon - \frac{1}{N}}{1 - \frac{1}{N}}\right]^2$.

Lower bounds: The main contribution of this work is our lower bound results. We show that the reduction in error obtained in Theorem 1 and 2 are essentially optimal.

Theorem 5 (Lower bound result). *Let $0 < \ell \le u < 1$ be such that ℓN and uN are integers.*

(a) *For all one-query database search algorithms $\mathcal{A}$, for either $\epsilon = \ell$ or $\epsilon = u$,*

$$\mathsf{err}_{\mathcal{A}}(\epsilon) \ge \epsilon^3 \left(\frac{u - \ell}{u + \ell - 2\ell u}\right)^2 .$$

(b) *For all t-query database search algorithms $\mathcal{A}$, there is an $\epsilon \in [\ell, u]$ such that ϵN is an integer, and*

$$\mathsf{err}_{\mathcal{A}}(\epsilon) \ge \epsilon^{2t+1} \left(\frac{b - 1}{b + 1} - \frac{1}{N\ell(b + 1)}\right)^{2t} ,$$

where $b = (\frac{u}{\ell})^{\frac{1}{t+1}}$, and we assume that $N\ell(b - 1) > 1$.

In particular, this result shows that to achieve the same reduction in error, a quantum algorithm needs to make roughly at least half as many queries as a classical randomized algorithm. A similar result can be obtained by modifying the proof in Buhrman, Cleve, de Wolf and Zaka [BC+99]: there is a constant $c > 0$, such that for all $u > 0$, all large enough N and all t-query quantum search algorithms for databases of size N, there is an $\epsilon \in (0, u]$ (ϵN is an integer) such that $\mathsf{err}_{\mathcal{A}}(\epsilon) \ge (cu)^{2t+1}$.

3 Upper bounds: quantum algorithms

In this section, we present algorithms that justify Theorems 1, 2 and 3, but by using the standard database oracle. We then modify these algorithms to generalize and slightly improve these theorems.

3.1 Alternative algorithms using the standard oracle

We first describe an alternative algorithm $\mathcal{A}_0$ to justify Theorem 1. This simple algorithm (see Figure 1) illustrates the main idea used in all our upper bounds. We will work with n qubits corresponding to addresses in $[N]$ and one ancilla qubit. Although we do not simulate Grover's oracle directly, using the ancilla, we can reproduce the effect the complex amplitude used there by real amplitudes. As we will see, the resulting algorithm has an intuitive explanation and also some additional properties not enjoyed by the original algorithm.

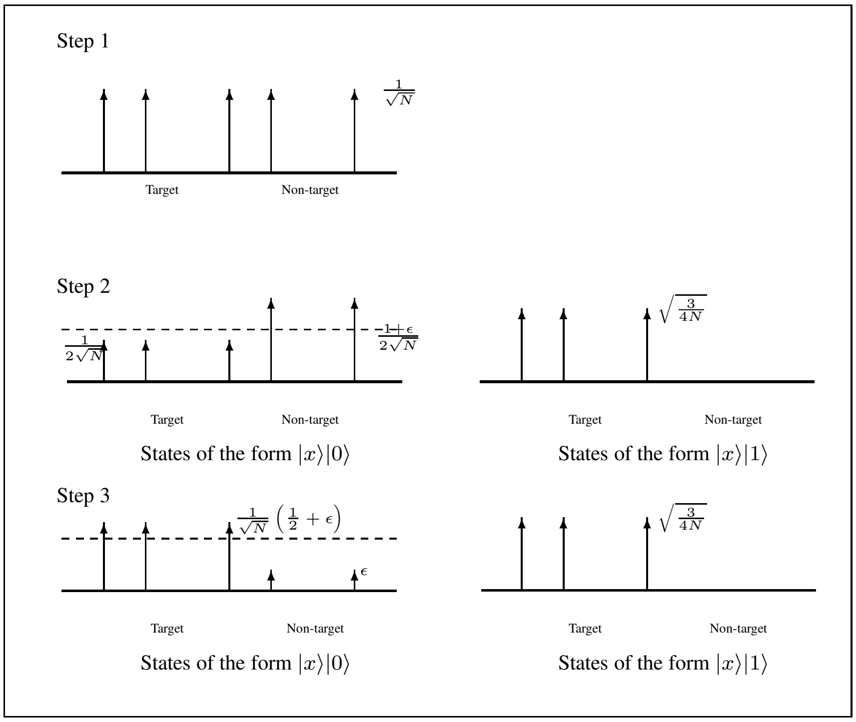

Fig. 1. The one-query algorithm

Step 1: We start with the uniform superposition on $[N]$ with the ancilla bit in the state $|0\rangle$.

Step 2: For targets x, transform $|x\rangle|0\rangle$ to $\frac{1}{2}|x\rangle|0\rangle + \sqrt{\frac{3}{4}}|x\rangle|1\rangle$. The basis states $|x\rangle|0\rangle$ for $x \in f^{-1}(0)$, are not affected by this transformation.

Step 3: Perform an inversion about the average controlled on the ancilla bit being $|0\rangle$, and then measure the address registers.

Step 1 is straightforward to implement using the n-qubit Hadamard transform H_n. For Step 2, using one-query to T_f, we implement a unitary transformation U_f, which maps $|x\rangle|0\rangle$ to $|x\rangle|0\rangle$ if $f(x)=0$ and to $|x\rangle\left(\frac{1}{2}|0\rangle+\sqrt{\frac{3}{4}}|1\rangle\right)$, if $f(x)=1$. One such transformation is $U_f=(I_n\otimes R^{-1})T_f(I_n\otimes R)$, where I_n is the n-qubit identity operator and R is the one-qubit gate for rotation by $\frac{\pi}{12}$ (that is, $|0\rangle\overset{R}{\mapsto}\cos(\frac{\pi}{12})|0\rangle+\sin(\frac{\pi}{12})|1\rangle$ and $|1\rangle\overset{R}{\mapsto}\cos(\frac{\pi}{12})|1\rangle-\sin(\frac{\pi}{12})|0\rangle$). The inversion about the average is usually implemented as $A_n=-H_n(I_n-2|\mathbf{0}\rangle\langle\mathbf{0}|)H_n$. The controlled version we need is then given by

$$A_{n,0}=A_n\otimes|0\rangle\langle 0|+I_n\otimes|1\rangle\langle 1|.$$

Let $H'=H_n\otimes I$. The final state is $|\phi_f\rangle=A_{n,0}U_fH'|\mathbf{0}\rangle|0\rangle$.

To see that the algorithm works as claimed, consider the state just before the operator $A_{n,0}$ is applied. This state is

$$\frac{1}{\sqrt{N}}\left[\sum_{x\in f^{-1}(1)}\frac{1}{2}|x\rangle|0\rangle\sum_{x\in f^{-1}(0)}|x\rangle|0\rangle\right]+\frac{1}{\sqrt{N}}\sum_{x\in f^{-1}(1)}\sqrt{\frac{3}{4}}|x\rangle|1\rangle.$$

Suppose $|f^{-1}(0)|=\epsilon N$. The "inversion about the average" is performed only on the first term, so the non-target states receive no amplitude from the second term. The average amplitude of the states in the first term is $\frac{1}{2\sqrt{N}}(1+\epsilon)$ and the amplitude of the states $|x\rangle|0\rangle$ for $x\in f^{-1}(0)$ is $\frac{1}{\sqrt{N}}$. Thus, after the inversion about the average the amplitude of $|x\rangle|0\rangle$ for $x\in f^{-1}(0)$ is $\frac{\epsilon}{\sqrt{N}}$. It follows that if we measure the address registers in the state $|\phi_f\rangle$, the probability of observing a non-target state is exactly

$$|f^{-1}(0)|\cdot\frac{\epsilon^2}{N}=\epsilon^3.$$

Remark: Note that this algorithm actually achieves more. Suppose we measure the ancilla bit in $|\phi_f\rangle$, and find a 1. Then, we are assured that we will find a target on measuring the address registers. Furthermore, the probability of the ancilla bit being 1 is exactly $\frac{3}{4}(1-\epsilon)$. One should compare this with the randomized one-query algorithm that with probability $1-\epsilon$ provides a guarantee that the solution it returns is correct. The algorithm in [G05] has no such guarantee associated with its solution. However, the algorithm obtained by Tulsi, Grover and Patel [TGP05] gives a guarantee with probability $\frac{1}{2}(1-\epsilon)$.

The general algorithm $\mathcal{B}_0$ needed to justify Theorem 2 is similar. We use G instead of H_n; that is, we let $H'=G\otimes I$, $A_n=G(2|\mathbf{0}\rangle\langle\mathbf{0}|-I_n)G^{-1}$, and, as before,

$$A_{n,0}=A_n\otimes|0\rangle\langle 0|+I_n\otimes|1\rangle\langle 1|.$$

The final state is obtained in the same way as before $|\phi_f\rangle=A_{n,0}U_fH'|\mathbf{0}\rangle|0\rangle$.

Remark: As stated, we require the controlled version of G and G^{-1} to implement $A_{n,0}$. However, we can implement G with the uncontrolled versions themselves from the following alternative expression for $A_{n,0}$:

$$A_{n,0}=(G\otimes I)[(2|\mathbf{0}\rangle\langle\mathbf{0}|-I_n)\otimes|0\rangle\langle 0|+I_n\otimes|1\rangle\langle 1|](G^{-1}\otimes I).$$

We can estimate the error probability of this algorithm using the following standard calculation. Suppose the probability of obtaining a non-target state on measuring the address registers in the state $G|\mathbf{0}\rangle$ is ϵ. Let us formally verify that the probability of obtaining a non-target state on measuring the address registers in the state $|\phi_f\rangle$ is ϵ^3. This follows using the following routine calculation. We write

$$G|\mathbf{0}\rangle = \alpha|t\rangle + \beta|t'\rangle,$$

where $|t\rangle$ is a unit vector in the "target space" spanned by $\{|x\rangle : f(x) = 1\}$, and $|t'\rangle$ is a unit vector in the orthogonal complement of the target space. By scaling $|t\rangle$ and $|t'\rangle$ by suitable phase factors, we can assume that α and β are real numbers. Furthermore $\beta^2 = \epsilon$. The state after the application of U_f is then given by

$$\left(\frac{\alpha}{2}|t\rangle + \beta|t'\rangle\right)|0\rangle + \sqrt{\frac{3}{4}}\alpha|t\rangle|1\rangle. \tag{1}$$

Now, the second term is not affected by $A_{n,0}$, so the amplitude of states in the subspace of non-target states is derived entirely from the first term, which we denote by $|u\rangle$. To analyze this contribution we write $|u\rangle$, using the basis

$$|v\rangle = \alpha|t\rangle + \beta|t'\rangle; \tag{2}$$
$$|v'\rangle = \beta|t\rangle - \alpha|t'\rangle. \tag{3}$$

That is, $|u\rangle = \left(\frac{\alpha^2}{2} + \beta^2\right)|v\rangle|0\rangle - \frac{\alpha\beta}{2}|v'\rangle|0\rangle$.

Since $A_{n,0}|v\rangle = |v\rangle$ and $A_{n,0}|v'\rangle = -|v'\rangle$, we have

$$A_{n,0}|u\rangle = \left(\frac{\alpha^2}{2} + \beta^2\right)|v\rangle|0\rangle + \frac{\alpha\beta}{2}|v'\rangle|0\rangle.$$

Returning to the basis $|t\rangle$ and $|t'\rangle$ (using (2) and (3)), we see that the amplitude associated with $|t'\rangle$ in this state is β^3. Thus, the probability that the final measurement fails to deliver a target address is exactly $\beta^6 = \epsilon^3$.

Remark: The algorithm $\mathcal{B}_0$ can be used recursively to get a t-query algorithm that achieves the bound Theorem 3. Just as in the one-query algorithm, by measuring the ancilla bits we can obtain a guarantee; this time the solution is accompanied with guarantee with probability at least $(1 - \frac{1}{t} - \frac{6\log t}{t(\log\frac{1}{\epsilon})^{\log_3 4}})$. The t-query algorithm obtained by Tulsi, Grover and Patel [TGP05] has significantly better guarantees: it certifies that its answer is correct with probability at least $1 - \epsilon^{2t}$.

3.2 Algorithms with restrictions on ϵ

As stated above, for each $\delta \in [0, 1]$, there is a one-query quantum algorithm $\mathcal{A}_\delta$ that makes no error if the $|f^{-1}(0)| = \delta N$ (or, in the general setting, if G is known to err with probability at most $\frac{3}{4}$). Let us explicitly obtain such an algorithm $\mathcal{A}_\delta$ by slightly modifying the algorithm above. The idea is to ensure that the inversion about the average performed in Step 3 reduces the amplitude of the non-target states to zero. For this, we only need to replace U_f by $U_{f,\delta}$, which maps $|x\rangle|0\rangle$ to $|x\rangle|0\rangle$ if $f(x) = 0$ and to $|x\rangle(\alpha|0\rangle + \beta|1\rangle)$, if $f(x) = 1$, where $\alpha = \dfrac{1-2\delta}{2(1-\delta)}$ and $\beta = \sqrt{1-\alpha^2}$.

Also, one can modify the implementation of U_f above, replacing $\frac{\pi}{12}$ by $\frac{\sin^{-1}(\alpha)}{2}$ (note that $\delta \le \frac{3}{4}$ implies that $|\alpha| \le 1$), and implement $U_{f,\delta}$ using just one-query to T_f.

Proposition 1. *Let* $|f^{-1}(0)| = \delta \le \frac{3}{4}$. *Then,* $\mathsf{err}_{\mathcal{A}_\delta}(f) = 0$.

An analogous modification for the general search gives us an algorithm $\mathcal{B}_\delta(T, G)$ that has no error when G produces a target state for T with probability exactly $1-\delta$. We next observe that the algorithms $\mathcal{A}_\delta$ and $\mathcal{B}_\delta$ perform well not only when the original probability is known to be δ but also if the original probability is $\epsilon \ge \delta$. This justifies Theorem 4 claimed above.

Proof of Theorem 4: We will only sketch the calculations for part (a). The average amplitude of all the states of the form $|x\rangle|0\rangle$ is $(\frac{1}{\sqrt{N}})(1-2\delta+\epsilon)/(2(1-\delta))$. From this it follows that the amplitude of a non-target state after the inversion about the average is $(\frac{1}{\sqrt{N}})(\epsilon-\delta)/(1-\delta)$. Our claim follows from this by observing that there are exactly ϵN non-target states. □

4 Lower bounds

In this section, we show that the algorithms in the previous section are essentially optimal. For the rest of this section, we fix a t-query quantum search algorithm to search a database of size N. Using the polynomial method we will show that no such algorithm can have error probability significantly less than ϵ^{t+1}, for a large range of ϵ.

The proof has two parts. First, using standard arguments we observe that $\mathsf{err}_{\mathcal{A}}(\epsilon)$ is a polynomial of degree at most $2t+1$ in ϵ.

Lemma 1. *Let $\mathcal{A}$ be a t-query quantum search algorithm for databases of size N. Then, there is a univariate polynomial $r(Z)$ with real coefficients and degree at most $2t+1$, such that for all ϵ*

$$\mathsf{err}_{\mathcal{A}}(\epsilon) \ge r(\epsilon).$$

Furthermore, $r(x) \ge 0$ for all $x \in [0,1]$.

In the second part, we analyze such low degree polynomials to obtain our lower bounds. We present this analysis first, and return to the proof of Lemma 1 after that.

4.1 Analysis of low degree polynomials

Definition 2 (Error polynomial). *We say that a univariate polynomial $r(Z)$ is an* error polynomial *if (a) $r(z) \ge 0$ for all $z \in [0,1]$, (b) $r(0) = 0$, and (c) $r(1) = 1$.*

Our goal is to show that an error polynomial of degree at most $2t+1$ cannot evaluate to significantly less than ϵ^{2t+1} for many values of ϵ. For our calculations, it will be convenient to ensure that all the roots of such a polynomial are in the interval $[0,1)$.

Lemma 2. *Let $r(Z)$ an error polynomial of degree $2t+1$ with $k < 2t+1$ roots in the interval $[0,1)$. Then, there is another error polynomial $q(Z)$ of degree at most $2t+1$ such that $q(z) \le r(z)$ for all $z \in [0,1]$, and $q(Z)$ has at least $k+1$ roots in the interval $[0,1)$.*

Proof. Let $\alpha_1, \alpha_2, \dots, \alpha_k$ be the roots of $r(x)$ in the interval $[0,1)$. Hence we can write

$$r(Z) = \prod_{i=1}^{k} (Z - \alpha_i) r'(Z),$$

where $r'(Z)$ does not have any roots in $[0,1)$. Now, by substituting $Z = 1$, we conclude that $r'(1) \ge 1$. Since $r'(Z)$ does not have any roots in $[0,1)$, it follows that $r'(z) > 0$ for all $z \in [0,1)$.

The idea now, is to subtract a suitable multiple of the polynomial $1-Z$ from $r'(Z)$ and obtain another polynomial $r''(Z)$ which has a root in $[0,1)$. Since $1-Z$ is positive in $[0,1)$, $r''(Z)$ is at most $r'(Z)$ in this interval. The polynomial $q(Z)$ will be defined by $q(Z) = \prod_{\alpha \in R}(Z-\alpha) r''(Z)$. To determine the multiple of $1-Z$ we need to subtract, consider $\lambda(c) = \min_{z\in[0,1)} r'(Z) - c(1-Z)$. Since $\lambda(c)$ is continuous, $\lambda(0) > 0$ and $\lambda(c) < 0$ for large enough c, it follows that $\lambda(c_0) = 0$ for some $c_0 > 0$. Now, let $r''(Z) = r(Z) - c_0(1-Z)$. □

By repeatedly applying Lemma 2 we obtain the following.

Lemma 3. *Let $r(Z)$ be an error polynomial of degree at most $2t+1$. Then, there is an error polynomial $q(Z)$ of degree exactly $2t+1$ such that $q(z) \le r(z)$ for all $z \in [0,1]$, and $q(Z)$ has $2t+1$ roots in the interval $[0,1)$.*

We can now present the proof of Theorem 5, our main lower bound result.

Proof of Theorem 5: Consider the case $t = 1$. By Lemma 1, it is enough to show that an error polynomial $r(Z)$ of degree at most three is bounded below as claimed. By Lemma 3, we may assume that all three roots of $r(Z)$ lie in $[0,1)$. Since $r(0) = 0$ and $r(z) \ge 0$ in $[0,1)$, we may write $r(Z) = aZ(Z-\alpha)^2$ for some $\alpha \in [0,1)$ and some positive a; since $r(1) = 1$, we conclude that $a = \frac{1}{(1-\alpha)^2}$. Thus, we need to determine the value of α so that $t(\alpha) = \max_{x\in\{\ell,u\}} \frac{r(x)}{x^3}$ is as small as possible. Consider the function $t_x(\alpha) = \frac{r(x)}{x^3} = \left(\frac{x-\alpha}{(1-\alpha)x}\right)^2$. Note that for all x, $t_x(\alpha)$ is monotonically increasing in $|x-\alpha|$. It follows that $t(\alpha)$ is minimum for some $\alpha \in [\ell, u]$. For α in this interval $t_\ell(\alpha)$ is an increasing function of α and $t_u(\alpha)$ is a decreasing function of α. So $t(\alpha)$ is minimum when $t_\ell(\alpha) = t_u(\alpha)$. It can be checked by direct computation that when $\alpha = \frac{2\ell u}{\ell+u}$,

$$t_\ell(\alpha) = t_u(\alpha) = \left(\frac{u-\ell}{u+\ell-2\ell u}\right)^2.$$

This establishes part (a) of Theorem 5.

To establish part (b), we show that an error polynomial of degree at most $2t+1$ satisfies the claim. As before, by Lemma 3, we may assume that $r(Z)$ has all its roots in $[0,1)$. Furthermore, since $r(Z) \geq 0$, we conclude that all roots in $(0,1)$ have even multiplicity. Thus we may write

$$r(Z) = \frac{Z(Z-\alpha_1)^2(Z-\alpha_2)^2\cdots(Z-\alpha_t)^2}{(1-\alpha_1)^2(1-\alpha_2)^2\cdots(1-\alpha_t)^2}.$$

Now, let $b = (\frac{u}{\ell})^{\frac{1}{t+1}}$. Consider subintervals $\{(\ell b^j, \ell b^{j+1}] : j = 0,1,\ldots,t\}$. One of these intervals say $\ell b^{j_0}, \ell b^{j_0+1}$ has no roots at all. Let ϵ be the mid point of the interval, that is, $\epsilon = (\ell b^{j_0} + \ell b^{j_0+1})/2$. Then, we have $(\epsilon - \alpha_j)^2 \geq \left(\frac{\ell b^{j_0+1} - \ell b^{j_0}}{2}\right)^2$, and since $(1-\alpha_j)^2 \leq 1$, we have $\dfrac{r(\epsilon)}{\epsilon^{2t+1}} \geq \left(\dfrac{b-1}{b+1}\right)^{2t}$.

This establishes part (b). The term $-\frac{1}{N\ell(b+1)}$ appears in the statement of Theorem 5 because we need to ensure that ϵN is an integer. □

4.2 Proof of Lemma 1

We will use the following notation. Let $p(X_1, X_2, \ldots, X_N)$ be a polynomial in N variables $X_1, X_2, \ldots, X_N$ with real coefficients. For a database $f : \{0,1\}^N \rightarrow \{0,1\}$, let

$$p(f) \stackrel{\Delta}{=} p(f(1), f(2), \ldots, f(N)).$$

Also, in the following $\mathbf{X}$ denotes the sequence of variables $X_1, X_2, \ldots, X_N$.

The key fact we need is the following.

Theorem 6 ([BB+95]). *Let $\mathcal{A}$ be a t-query quantum database search algorithm. Then, for $i = 1,2,\ldots,N$, there is a multilinear polynomial $p_i(\mathbf{X})$ of degree at most $2t$, such that for all f.*

$$\Pr[\mathcal{A}(f) = i] = p_i(f).$$

Furthermore, $p_i(\mathbf{x}) \geq 0$ for all $\mathbf{x} \in [0,1]^N$.

Lemma 4. *Let $\mathcal{A}$ be a t-query quantum database search algorithm. Then, there is a multilinear polynomial $p(\mathbf{X})$ of degree at most $2t+1$ such that for all f,*

$$\mathsf{err}_{\mathcal{A}}(f) = p_{\mathcal{A}}(f).$$

Proof. Using the polynomials $p_i(X)$ from Theorem 6, define

$$p_{\mathcal{A}}(\mathbf{X}) = \sum_{i=1}^{n} (1 - X_i) p_i(\mathbf{X}).$$

Clearly, $p(f) = \sum_{i=1}^{n} (1 - f(i)) p_i(f) = \sum_{i \in f^{-1}(0)} \Pr[\mathcal{A}(f) = i] = \mathsf{err}_{\mathcal{A}}(f)$. □

We can now prove Lemma 1. For a permutation σ of N and $f : [N] \rightarrow \{0,1\}$, let σf be the function defined by $\sigma f(i) = f(\sigma(i))$.

Note that $|f^{-1}(0)| = |(\sigma f)^{-1}(0)|$. Now,

$$\frac{1}{N!}\sum_{\sigma} p_{\mathcal{A}}(\sigma f) = \mathop{\mathrm{E}}_{\sigma}[\mathsf{err}_{\mathcal{A}}(\sigma f)] \leq \max_{\sigma} \mathsf{err}_{\mathcal{A}}(\sigma f) \leq \mathsf{err}_{\mathcal{A}}(\epsilon), \tag{4}$$

where $|f^{-1}(0)| = \epsilon N$.

Let $\sigma\mathbf{X}$ be the sequence $X_{\sigma(1)}, X_{\sigma(2)}, \ldots, X_{\sigma(N)}$, and let

$$p_{\mathcal{A}}^{\mathrm{sym}}(\mathbf{X}) = \frac{1}{N!}\sum_{\sigma} p_{\mathcal{A}}(\sigma\mathbf{X}).$$

Then, by (4), we have $p_{\mathcal{A}}^{\mathrm{sym}}(f) = \frac{1}{N!}\sum_{\sigma} p_{\mathcal{A}}(\sigma f) \leq \mathsf{err}_{\mathcal{A}}(\epsilon)$.

Now, $p_{\mathcal{A}}^{\mathrm{sym}}(\mathbf{X})$ is a symmetric multilinear polynomial in N variables of degree at most $2t+1$. For any such polynomial, there is a univariate polynomial $q(Z)$ of degree at most $2t+1$ such that if we let $\hat{p}(\mathbf{X}) = q(\sum_{i=1}^{N} X_i)/N)$, then for all f,

$$\hat{p}(f) = p_{\mathcal{A}}^{\mathrm{sym}}(f) \leq \mathsf{err}_{\mathcal{A}}(\epsilon).$$

(See Minsky and Papert [MP].) Now, $\hat{p}(f) = q((f(1) + f(2) + \ldots + f(N))/N) = q(1-\epsilon)$. To complete the proof, we take $r(Z) = q(1-Z)$. $\square$

Acknowledgements

We thank the referees for their helpful comments.

References

BB+95. R. Beals, H. Buhrman, R. Cleve, M. Mosca, R. de Wolf. Quantum lower bounds by polynomials, FOCS (1998): 352–361. quant-ph/9802049.

BC+99. H. Buhrman, R. Cleve, R. de Wolf, C. Zalka. Bounds for Small-Error and Zero-Error Quantum Algorithms, FOCS (1999): 358–368.

BH+02. G. Brassard, P. Hoyer, M. Mosca, A. Tapp: Quantum amplitude amplification and estimation. In S.J. Lomonaco and H.E. Brandt, editors, Quantum Computation and Information, AMS Contemporary mathematics Series (305), pp. 53–74, 2002. quant-ph/0005055.

G96. L.K. Grover: A fast quantum mechanical algorithm for database search. STOC (1996): 212-219. quant-ph/9605043.

G97. L.K. Grover: Quantum Mechanics helps in searching for a needle in a haystack. Physical Review Letters 79(2), pp. 325-328, July 14, 1997. quant-ph/9706033.

G98a. L.K. Grover: A framework for fast quantum mechanical algorithms, Proc. 30th ACM Symposium on Theory of Computing (STOC), 1998, 53–63.

G98b. L.K. Grover. Quantum computers can search rapidly by using almost any transformation, Phy. Rev. Letters, 80(19), 1998, 4329–4332. quant-ph/9711043.

G05. L.K. Grover. A different kind of quantum search. March 2005. quant-ph/0503205.

MP. M.L. Minsky and S.A. Papert. Perceptrons: An Introduction to Computational Geometry. MIT Press (1968).

NC. M.A. Nielsen and I.L. Chuang: Quantum Computation and Quantum Information. Cambridge University Press (2000).

TGP05. T. Tulsi, L.K. Grover, A. Patel. A new algorithm for directed quantum search. May 2005. quant-ph/0505007.