Quantum Algorithms 3/30/2023

Lecture 18: Query Lower Bounds

Instructor: Dieter van Melkebeek

We have seen a number of blackbox problems where the quantum setting offered efficiency gains over the classical setting, ranging from polynomial to exponential. As impressive as some of the gains may be, they leave the question whether there exist even better quantum algorithms. In this lecture we discuss two methods to develop quantum query lower bounds for blackbox problems, namely the adversary method and the polynomial method. Using either method we establish a lower bound of $\Omega(\sqrt{N})$ queries for quantum search, showing that the number of queries in Grover's algorithm is optimal up to a constant factor.

1 Model

Suppose we have a quantum algorithm for a classical blackbox problem that has access to the blackbox U_f for a function $f : \{0,1\}^n \to \{0,1\}^\ell$. As we are only concerned with the number of queries, we can condense the operations between successive applications of U_f into a single operation, which we can assume to be unitary by postponing intermediate measurements and introducing ancillas if need be. There can also be such an operation before the first query and after the last query. At the end, there is a measurement, we can assume without loss of generality to be of the entire system. The answer is then a subset of the bits measured. For an algorithm with q queries, this leads to the quantum circuit in Figure 1, where $V_1, \ldots, V_q, V_{\text{final}}$ represent the unitaries independent of f, $|\psi^{(0)}\rangle$ the start state, $|\psi^{(i)}\rangle$ the state right after the i-th query, and $|\psi^{\text{final}}\rangle$ the state right before the measurement at the end.

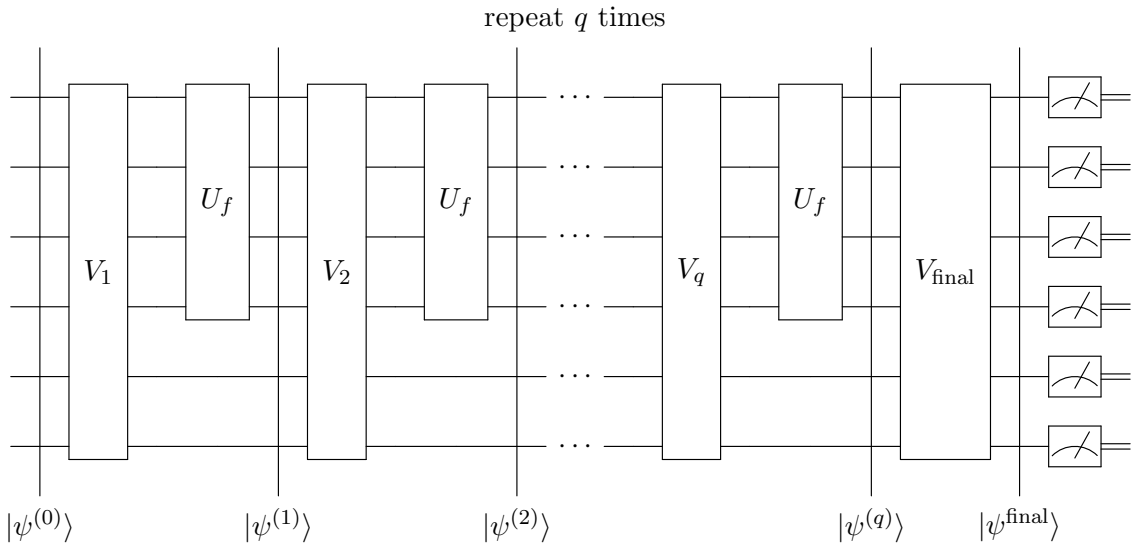

Figure 1: Model of a blackbox algorithm with q queries

The state after the i-th query is given by

$$|\psi^{(i)}\rangle = (U_f \otimes I) \cdot V_i \cdot \ldots \cdot (U_f \otimes I) \cdot V_1 \cdot |\psi^{(0)}\rangle\,,$$

where $|\psi^{(0)}\rangle = |0 \ldots 0\rangle$ without loss of generality.

2 Quantum search

The specific blackbox problem for which we will develop a lower bound is quantum search, where the underlying function f is Boolean ($\ell = 1$) and the goal is to find an $x \in \{0,1\}^n$ such that $f(x) = 1$, or else report that no such x exists. Grover's algorithm solves this problem with bounded error using $O(\sqrt{N})$ queries. We will establish a lower bound of $\Omega(\sqrt{N})$ queries for any quantum blackbox algorithm that solves the problem with bounded error, matching the upper bound given by Grover up to a constant factor.

In fact, both techniques yield the same lower bound for the weaker problem of distinguishing between the case where there is no $x \in \{0,1\}^n$ with $f(x) = 1$, and the case where there is a unique such x. The underlying decision problem can be viewed as computing the OR of the N variables $y_x \doteq f(x)$ over all $x \in \{0,1\}^n$. The lower bound applies to the promise version where there is at most one x for which $y_x = 1$.

3 Adversary method

Our first method is a quantum version of the classical adversary method. Given an algorithm that makes few queries, the method tries to adversarially construct some inputs (i.e., functions f) such that the algorithm has to behave incorrectly on at least one of them. In our case of the OR predicate, we consider the function $f = f_0$ that is identically zero and the functions $f = \delta_{x^*}$ for $x^* \in \{0,1\}^n$, where $\delta_{x^*}(x)$ is 1 at $x = x^*$ and 0 elsewhere. The function f_0 corresponds to the one and only input where the OR predicate is 0, or equivalently, there is no solution to the search problem. The functions δ_{x^*} correspond to the inputs where the OR predicate is 1 that are "closest" to f_0. Intuitively, "close" means hard to detect the difference for a query algorithm. In this case it formally means close in Hamming distance: the inputs for which there is exactly one 1.

Approach. Fix an algorithm as in Figure 1 that solves the problem with bounded error. Let $|\psi^{(i}\rangle$ denote the states in the figure on input f_0, and $|\psi^{(i)}_{x^*}\rangle$ the corresponding states on input δ_{x^*}. For any $x^* \in \{0,1\}^n$, the corresponding states start out the same ($|\psi^{(0)}\rangle = |\psi^{(0)}_{x^*}\rangle$), but in the end they need to be very different as the first one has a correct answer of 0, whereas the second one has a correct answer of 1. In fact, $|\psi^{\text{final}}\rangle$ and $|\psi^{\text{final}}_{x^*}\rangle$ need to be nearly orthogonal. The only way differences between $|\psi\rangle$ and $|\psi_{x^*}\rangle$ can arise along the way is when x^* is queried. Thus, in order for $|\psi^{\text{final}}\rangle$ and $|\psi^{\text{final}}_{x^*}\rangle$ to differ a lot, the total weight of the queries to x^* over the entire computation needs to be significant. Whereas this can happen for some x^* in a single query, for it to happen for all $x^* \in \{0,1\}^n$ requires many queries. We now work out this approach quantitatively.

Application to quantum search. Let p denote the probability distribution of the measurement at the end on input f_0, and p_{x^*} the one for input δ_{x^*}. Since the algorithm needs to answer correctly with probability at least $1 - \epsilon$ on every input, and the correct answers for f_0 and δ_{x^*} are different,

the statistical distance $d_{\text{stat}}(p, p_{x^*})$ is at least $1 - 2\epsilon$). From the lecture on quantum distances, we know that the statistical distance is at most $\| |\psi^{\text{final}}\rangle - |\psi^{\text{final}}_{x^*}\rangle \|_2$. Thus,

$$1 - 2\epsilon \leq \| |\psi^{\text{final}}\rangle - |\psi^{\text{final}}_{x^*}\rangle \|_2, \tag{1}$$

which formalizes the near-orthogonality claim.

Next we figure out by how much the difference in 2-norm can increase due to one query. Noting the $U_{f_0} = I$, we have

$$\begin{aligned} |\psi^{(i)}\rangle &= (U_{f_0} \otimes I)V_i |\psi^{(i-1)}\rangle = (I \otimes I)V_i |\psi^{(i-1)}\rangle \\ |\psi^{(i)}_{x^*}\rangle &= (U_{\delta_{x^*}} \otimes I)V_i |\psi^{(i-1)}_{x^*}\rangle , \end{aligned}$$

and therefore

$$\left\| |\psi^{(i)}\rangle - |\psi^{(i)}_{x^*}\rangle \right\| = \left\| \left((I \otimes I)V_i |\psi^{(i-1)}\rangle - (U_{\delta_{x^*}} \otimes I)V_i |\psi^{(i-1)}_{x^*}\rangle \right) \right\| .$$

This equation can be simplified by applying the triangle inequality and removing the unitary terms (which do not affect the 2-norm):

$$\begin{aligned} \left\| |\psi^{(i)}\rangle - |\psi^{(i)}_{x^*}\rangle \right\| &\leq \left\| (I \otimes I)V_i |\psi^{(i-1)}\rangle - (U_{\delta_{x^*}} \otimes I)V_i |\psi^{(i-1)}\rangle \right\| + \left\| (U_{\delta_{x^*}} \otimes I)V_i \left(|\psi^{(i-1)}\rangle - |\psi^{(i-1)}_{x^*}\rangle \right) \right\| \\ &= \underbrace{\left\| (I \otimes I)V_i |\psi^{(i-1)}\rangle - (U_{\delta_{x^*}} \otimes I)V_i |\psi^{(i-1)}\rangle \right\|}_{\Delta} + \left\| |\psi^{(i-1)}\rangle - |\psi^{(i-1)}_{x^*}\rangle \right\| . \end{aligned} \tag{2}$$

We now analyze the term Δ in Equation (2), which gives us an upper bound on the increase in 2-norm. Let

$$V_i |\psi^{(i-1)}\rangle = \sum_z \alpha_z |z\rangle ,$$

where $|z\rangle = |xbu\rangle$ such that $|xb\rangle$ represents the input into U_f and $|b\rangle$ represents the ancilla qubit in which the value of $f(x)$ is XORed. Note that

$$U_{\delta_{x^*}} \otimes I : |xbu\rangle \mapsto \begin{cases} |xbu\rangle & \text{if } x \neq x^* \\ |x\bar{b}u\rangle & \text{if } x = x^* \end{cases} .$$

We can now proceed to bound Δ:

$$\begin{aligned} \Delta^2 &= \left\| \left[(I \otimes I) - (U_{\delta_{x^*}} \otimes I) \right] \sum \alpha_z |z\rangle \right\|^2 \\ &= \left\| \left[(I \otimes I) - (U_{\delta_{x^*}} \otimes I) \right] \sum_{bu} \alpha_z |x^* bu\rangle \right\|^2 \\ &= \sum_{bu} \left| \alpha_{x^* bu} - \alpha_{x^* \bar{b} u} \right|^2 \\ &= 2 \cdot \sum_u \left| \alpha_{x^* 0u} - \alpha_{x^* 1u} \right|^2 \\ &\leq 4 \cdot \sum_u \left(|\alpha_{x^* 0u}|^2 + |\alpha_{x^* 1u}|^2 \right) \\ &= 4 \cdot \Pr[i^{\text{th}} \text{ query for } f_0 \text{ is } x^*]. \end{aligned} \tag{3}$$

We return to the term $\left\||\psi^{\text{final}}\rangle - |\psi^{\text{final}}_{x^*}\rangle\right\|$ which, by removing unitary transformations, is simply $\left\||\psi^{(q)}\rangle - |\psi^{(q)}_{x^*}\rangle\right\|$. By combining Equations (2) and (3), this gives us

$$\begin{aligned}\left\||\psi^{\text{final}}\rangle - |\psi^{\text{final}}_{x^*}\rangle\right\| &= \left\||\psi^{(q)}\rangle - |\psi^{(q)}_{x^*}\rangle\right\| \\ &\leq 2\sum_{i=1}^{q}\sqrt{\Pr[i^{\text{th}} \text{ query for } f_0 \text{ is } x^*]} + \underbrace{\left\||\psi^{(0)}\rangle - |\psi^{(0)}_{x^*}\rangle\right\|}_{=0}. \end{aligned} \tag{4}$$

which is true for every possible x^*. While the probability terms in Equation (4) might be large for particular x^*, on average they have to be small since they add up to 1. So, we sum over all x^* and apply the Cauchy-Schwarz inequality to obtain:

$$\begin{aligned}\sum_{x^*}\left\||\psi^{\text{final}}\rangle - |\psi^{\text{final}}_{x^*}\rangle\right\| &\leq 2\sum_{x^*}\sum_{i=1}^{q}\sqrt{\Pr[i^{\text{th}} \text{ query for } f_0 \text{ is } x^*]} \\ &= 2\sum_{i=1}^{q}\sum_{x^*}\sqrt{\Pr[i^{\text{th}} \text{ query for } f_0 \text{ is } x^*]} \\ &\leq 2\sum_{i=1}^{q}\left(1\cdot\sqrt{N}\right) \\ &= 2q\sqrt{N} \end{aligned} \tag{5}$$

Combining (5) with (1) summed over all $x^* \in \{0,1\}^n$, we conclude that

$$q \geq (\frac{1}{2} - \epsilon)\sqrt{N}.$$

General adversary method. As mentioned, a blackbox $f : \{0,1\}^n \to \{0,1\}$ can be viewed as a set of N Boolean variables y_x for $x \in \{0,1\}^n$, where y_x represents $f(x)$. Viewing x as a number in binary, we can also represent f as the sequence of variables $y_0, y_1, \ldots, y_{N-1}$, which corresponds to the charactertistic sequence of f. With this representation, a blackbox decision problem corresponds to a function $F : \{0,1\}^N \to \{0,1\}$. The function F is total in case of fully specified decision problems, and partial in case of promise problems. We showed that any quantum algorithm for computing $F = OR$ with bounded error has to make $\Omega(\sqrt{N})$ queries. In fact, we showed that the bound already holds for the promise version where the strings in $\{0,1\}^N$ contain at most one 1.

More generally, given a (partial) function $F : \{0,1\}^N \to \{0,1\}$, we can ask about the complexity of computing F. To establish a lower bound, we follow the same intuition – states for inputs that map to different values under F must start out identical and end up almost orthogonal. The only operations that can induce a difference between the states are the blackbox queries. If we can exhibit a collection of input pairs from $F^{-1}(0) \times F^{-1}(1)$ such that any blackbox query can only induce a small difference in the states on average over the collection, a lot of queries are needed to compute F.

In order to construct the collection, we capitalize on the fact that for a given state and two inputs x and y, a query can only induce a significant difference if the state puts a lot of weight on the positions where x and y differ. Thus, to make the difference small on average, we make sure

that any single query position $i \in [N]$ can only be different for a small fraction of the input pairs (x, y) that we consider. This approach leads to the following quantitative statement, where we view the pairs (x, y) that we consider as a bipartite graph $G = (V, E)$ where $V = X \sqcup Y$, $E \subseteq X \times Y$, $X \subseteq f^{-1}(0)$, and $Y \subset f^{-1}(1)$.

Theorem 1. *Let $F : \{0,1\}^N \to \{0,1\}$ be a partial function, $X \subseteq F^{-1}(0)$, $Y \subseteq F^{-1}(1)$, and $E \subseteq X \times Y$. Suppose that:*

- $(\forall x \in X)\ |\{y : (x, y) \in E\}| \geq d_{left}$
- $(\forall y \in Y)\ |\{x : (x, y) \in E\}| \geq d_{right}$
- $(\forall i \in [N])(\forall x \in X)\ |\{y \in Y : (x, y) \in E \text{ and} x_i \neq y_i\}| \leq c_{left}$
- $(\forall i \in [N])(\forall y \in Y)\ |\{x \in X : (x, y) \in E \text{ and} x_i \neq y_i\}| \leq c_{right}$

Then any bounded-error quantum algorithm for computing F must make $\Omega\left(\sqrt{\frac{d_{left}d_{right}}{c_{left}c_{right}}}\right)$ queries.

The set E is chosen adversarially so as to obtain as large a lower bound on the query complexity as possible. Note that to obtain a strong lower bound, the typical Hamming distance between x and y for $(x, y) \in E$ needs to be small.

We do not prove the theorem, but show how to instantiate it to obtain our lower bound for OR. We pick $X = \{0^N\}$, Y the subset of $\{0,1\}^N$ with Hamming weight 1, and $E = X \times Y$. We can set $d_{\text{left}} = N$, $d_{\text{right}} = 1$, $c_{\text{left}} = 1$, and $c_{\text{right}} = 1$. On substituting these values in the formula for lower bound in Theorem 1, we get a $\Omega(\sqrt{N})$ lower bound.

Exercise. Let F be a balanced AND of ORs. More precisely, F is given by a Boolean formula in conjunctive normal form with $\sqrt{N}$ clauses, and each clause has $\sqrt{N}$ literals. The first clause is an OR of the first $\sqrt{N}$ variables $y_0, y_1, \ldots, y_{\sqrt{N}-1}$ (called the first block), the second clause has the next $\sqrt{N}$ variables and so on. Show that the query complexity of F is $\Omega(\sqrt{N})$ and $O(\sqrt{N} \log N)$.

There exists a stronger version of Theorem 1 with weights on the edges, including a version where the weights can be negative, which gives tight results up to a constant factor for the bounded-error query complexity of any partial function F [Rei09].

4 Polynomial method

The second method turns a quantum query algorithm into a well-studied structure, namely a polynomial that approximates F, and then uses results such polynomials to obtain a lower bound on the number of queries. The method is not as universal as the quantum adversary method, but has led to interesting results, including generic relationships between quantum and classical query complexity.

Approach. The method hinges on the following lemma. We state it for the setting where the blackbox f is Boolean ($\ell = 1$).

Lemma 2 (Amplitude lemma). *Given a quantum query algorithm with q queries as in Figure 1, where $f : \{0,1\}^n \to \{0,1\}$, the amplitudes in $|\psi^{final}\rangle$ are multivariate polynomials of degree at most q in the variables $y_x \doteq f(x)$.*

Proof. We prove by induction on i that the amplitudes in $|\psi^{(i)}\rangle$ are multivariate polynomials in the variables y_x of degree at most i. The base case $i = 0$ holds as $|\psi^{(0)}\rangle$ is independent of the variables.

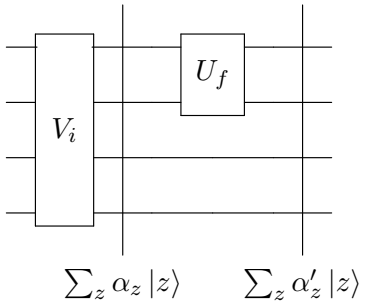

Figure 2: Inductive step of the amplitude lemma

For the induction step $i-1 \to i$, consider the circuit part in Figure 2. The state at the beginning is $|\psi^{(i-1)}\rangle$ and the one at the end is $|\psi^{(i)}\rangle$. Consider the state $V_i |\psi^{(i-1)}\rangle$, which we can decompose as $\sum_z \alpha_z |z\rangle$. By the induction hypothesis and the fact that V_i is a linear transformation, each α_z is a polynomial of degree at most $i-1$. We can similarly decompose the state $|\psi^{(i)}\rangle$ as $\sum_x \alpha'_z |z\rangle$. As $|\psi^{(i)}\rangle = (U_f \otimes I) \sum_z \alpha_z |z\rangle$ and writing $z = xbu$ as before ($|x| = n$ and $|b| = 1$), we have that

$$\alpha'_{xbu} = \begin{cases} \alpha_{xbu} & \text{if } f(x) = 0 \\ \alpha_{x\bar{b}u} & \text{if } f(x) = 1. \end{cases}$$

Note that $f(x)$ is y_x, which is one of our variables. By interpolation, we can rewrite

$$\alpha'_{xbu} = \alpha_{xbu}(1 - y_x) + \alpha_{x\bar{b}u} y_x.$$

As the α_z are polynomials of degree at most i, it follows that the α'_z are polynomials of degree at most i.

The lemma follows from the inductive claim for $i = q$ and the fact that V_{final} is a linear transformation. □

Decomposing $|\psi^{\text{final}}\rangle$ as $\sum_z \alpha_z^{\text{final}} |z\rangle$, the probability that the algorithm outputs 1 can be written as the sum of $|\alpha_z^{\text{final}}|^2 = (\alpha_z^{\text{final}})^* \alpha_z^{\text{final}}$ over all basis states $|z\rangle$ that result in output 1. For concreteness and without loss of generality, we can assume that the first qubit yields the output, in which case we sum over the basis states $|z\rangle$ where z starts with 1. By the amplitude lemma, the probability of outputting 1 is given by a multivariate real polynomial p of degree at most $2q$ in the variables y_x. We can say the following about the values $p(y)$ where $y \doteq (y_0, y_1, \ldots, y_{N-1})$.

- For every $y \in \{0,1\}^N$, $p(y) \in [0,1]$. This follows because the value $p(y)$ is a probability.

- For every $y \in \mathrm{dom}(F)$ with $F(y) = 1$, $f(y) \geq 1 - \epsilon$. This follows because the error is bounded by ϵ.
- For every $y \in \mathrm{dom}(F)$ with $F(y) = 0$, $f(y) \leq \epsilon$. This also follows from the error bound as $\Pr[\text{ output 1 on input y}] = 1 - \Pr[\text{ output 0 on input y}] \leq 1 - (1 - \epsilon) = \epsilon$.

The last two bullets mean that p approximates F to within ϵ on the domain of F.

Corollary 3. *Given a quantum query algorithm with q queries that computes a partial Boolean function F of the variables y_x where $y_x \doteq f(x)$ and $f : \{0,1\}^n \to \{0,1\}$, there exists a multivariate real polynomial p of degree at most $2q$ that has values in $[0,1]$ for all $y \in \{0,1\}^N$ and satisfies*

$$(\forall y \in dom(F))\, |p(y) - F(y)| \leq \epsilon. \tag{6}$$

Corollary 3 reduces lower bounding the number of queries of bounded-error quantum algorithms for F to lower bounding the degree of real polynomials that approximate F well on the part of the Boolean cube in the domain of F (and are in the range $[0,1]$ on the entire Boolean cube). Given the simpler mathematical structure, the latter lower bounds may be more in reach. Given the vast knowledge about polynomials and polynomial approximations, such lower bounds may already be known. In fact, the latter was the case for $F = OR$ and the promise version we consider, a result we will develop next. The method has also been used to connect quantum query complexity to classical deterministic query complexity. For any total Boolean function F, the bounded-error quantum query complexity is at least of the order of the fourth root of the deterministic query complexity [ABK+21]. If, in addition, F is symmetric (i.e., the value of F on the Boolean cube only depends on the number fo ones in the input), the fourth root can be improved to the square root. Both results are known to be tight up to lower order terms. Note that for promise problems F, the gap can be exponential.

Application to quantum search. We now show how to obtain a lower bound of $\Omega(\sqrt{N})$ for $F = OR$ on inputs with at most one 1. Consider a quantum algorithm with q queries and the multivariate real polynomial p of degree at most $2q$ from Corollary 3. Given the symmetric nature of the problem, it makes sense to consider the *symmetrization* of p:

$$p_{\mathrm{symm}}(y) \doteq \frac{1}{N!} \sum_{\pi \in S_N} p(\pi(y)),$$

where S_N denotes the set of permutations of N elements. Since the variables y_x only take values in $\{0,1\}$, each occurrence of y_x^k with $k \geq 2$ in a monomial of p can be replaced by y_x. Thus, without loss of generality, the polynomial p is multi-linear, i.e., has degree at most 1 in each variable. The same holds for its symmetrization p_{symm}. Moreover, in p_{symm} the coefficient of a multi-linear monomial only depends on its degree d. On a Boolean input y of Hamming weight $w \doteq \sum_i y_i$, there are exactly $\binom{w}{d}$ nonzero monomials of degree d, each of which evaluate to 1. It follows that $p_{\mathrm{symm}}(y) = \sum_{d=0}^{2q} c_d \binom{w}{d}$ for some reals $c_0, \cdots, c_{2q}$. Since $\binom{w}{d}$ is a univariate polynomial of degree d in w, we can write

$$p_{\mathrm{symm}}(y) = \tilde{p}\left(\sum_i y_i\right),$$

where $\tilde{p}$ is a univariate real polynomial of degree at most $2q$. By (6) and the other properties of the polynomial p in Corollary 3, we conclude:

Proposition 4. *If there exists a quantum blackbox algorithm with q queries for computing $F = OR$ on inputs from the Boolean cube of Hamming weight at most 1, then there exists a univariate real polynomial $\tilde{p}$ such that:*

$$\tilde{p}(w) \in \begin{cases} [0, \epsilon] & \text{for } w = 0 \\ [1-\epsilon, 1] & \text{for } w = 1 \\ [0, 1] & \text{for } w \in \{2, \ldots, N\}. \end{cases} \tag{7}$$

See Figure 3 for a plot of the situation.

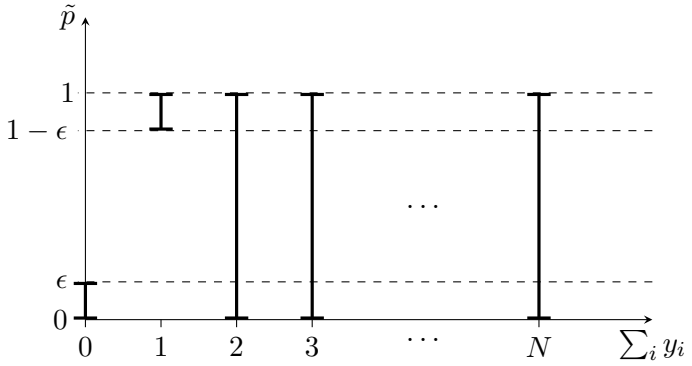

Figure 3: Approximating univariate polynomial for quantum search

Intuitively, the fact that a polynomial makes a sudden jump (from close to 0 at $w = 0$ to close to 1 at $w = 1$) and remains constrained to a small range on a nontrivial part of its domain (the range $[0, 1]$ on $\{0, 1, \ldots, N\}$) requires the polynomial to have large degree. To quantify this intuition, recall that the Chebyshev polynomialsT_d is contained in the interval $[-1, 1]$ on the interval $[-1, 1]$ and has a number of extremal properties among all such polynomials of the same degree d. In particular, the maximum derivative in absolute value, i.e., the quantity $\max_{x \in [-1,1]} |\tilde{p}(x)|$, is maximized by T_d. We will not prove this property but compute the maximum. Recall that $T_d(\cos\theta) = \cos(d\theta)$. By the chain rule, $T'_d(\cos\theta) \cdot (-\sin\theta) = -\sin(d\theta) \cdot d$, so $T'_d(\cos\theta) = d \cdot \frac{\sin(d\theta)}{\sin\theta}$, which yields a maximum absolute value of d^2 at $\theta \in \pi\mathbb{Z}$ (corresponding to $|x| = 1$). By rescaling, we have:

Lemma 5 (Markov's other inequality). *Let $\tilde{p}$ be a univariate real polynomial of degree d and $X, Y \subseteq \mathbb{R}$ intervals such that $\tilde{p}(X) \subseteq Y$. Then $|\tilde{p}'(x)| \leq \frac{|Y|}{|X|} d^2$ for all $x \in X$.*

Lemma 5 does not immediately apply to the setting of Proposition 4 depicted in Figure 3 because the lemma requires *all* points of the interval $X = [0, N]$ to have values constrained in an interval Y, whereas the proposition guarantees such a constraint at the integral points in $[0, N]$. However, in cases where $\tilde{p}$ assumes a large value at a non-integral point in $[0, N]$, $\tilde{p}$ needs to have a large derivative in absolute value somewhere between that point and the closest integral point, where the value of $\tilde{p}$ is constrained to the interval $[0, 1]$. This allows us to guarantee a large derivative in absolute value on $X[0, N]$ in both the case where $Y = \tilde{p}(X)$ is small, and the case where it it large. Here is a quantitative analysis:

- By Proposition 4, $\tilde{p}(0) \leq \epsilon$ and $\tilde{p}(1) \geq 1 - \epsilon$. By the mean value theorem, there exists $x^* \in X$ such that $\tilde{p}'(x^*) \geq 1 - 2\epsilon$.

- If $|Y| \geq 1$, there exists a point $x \in X$ such that distance of $\tilde{p}(x)$ to $[0,1]$ is $(|Y|-1)/2$. This follows because the range of $\tilde{p}(X)$ outside of $[0,1]$ has size $|Y|-1$, and at least half of it needs to be on the same side of $[0,1]$. On the other hand, at the least integral point, $\lceil x \rfloor$, $\tilde{p}$ has a value in $[0,1]$. The distance between x and $\lceil x \rfloor$ is at most $1/2$. By mean value theorem, there exists $x^* \in X$ such that $\tilde{p}'(x^*) \geq |Y|-1$.

Thus, there exists $x^* \in X$ such that $|\tilde{p}'(x^*)| \geq \max(1-2\epsilon, |Y|-1)$. By Markov's other inequality $\max(1-2\epsilon, |Y|-1) \leq \frac{|Y|}{N} d^2$, where $d = 2q$. We do not know what $|Y|$ is, but can guarantee the lower bound obtained for the case where $1-2\epsilon = |Y|-1$, or equivalently, for $|Y| = 2(1-\epsilon)$. We conclude that

$$q \geq \frac{1}{2}\sqrt{\frac{1}{2} - \frac{\epsilon}{2(1-\epsilon)}}\sqrt{N}.$$

References

[ABK+21] Scott Aaronson, Shalev Ben-David, Robin Kothari, Shravas Rao, and Avishay Tal. Degree vs. approximate degree and quantum implications of Huang's sensitivity theorem. In *53rd Annual Symposium on Theory of Computing (STOC)*, pages 1330–1342, 2021.

[Rei09] Ben Reichardt. Span programs and quantum query complexity: The general adversary bound is nearly tight for every boolean function. In *50th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 544–551, 2009.