COL7160 : Quantum Computing

Lecture 22: Quantum Query Complexity: Lower Bound Methods

Instructor: Rajendra Kumar **Scribe:** Rishabh Shakya

1 Introduction

A central question in quantum computing is: *how many times must a quantum algorithm query an input oracle to compute a Boolean function?* This lecture covers two complementary methods for proving *lower bounds* on quantum query complexity: the **polynomial method** and the **adversary method**. Both methods establish that any quantum algorithm—no matter how clever—requires a certain minimum number of oracle queries to solve a given problem with bounded error.

We work in the *black-box* or *oracle* model. The input is a string $x = (x_0, x_1, \ldots, x_{N-1}) \in \{0,1\}^N$, and the algorithm can only learn about x by querying an oracle U_x. The quantum query complexity $Q(\varphi)$ of a Boolean function $\varphi : \{0,1\}^N \to \{0,1\}$ is the minimum number of oracle queries made by any bounded-error quantum algorithm computing φ.

2 The Polynomial Method

2.1 The T-Query Quantum Algorithm

A T-query quantum algorithm on m qubits has the following structure:

$$|0\rangle^{\otimes m} \xrightarrow{U_0} \xrightarrow{U_x} \xrightarrow{U_1} \xrightarrow{U_x} \cdots \xrightarrow{U_x} \xrightarrow{U_T} \text{measure}$$

Here $U_0, U_1, \ldots, U_T$ are arbitrary unitary gates that do *not* depend on the input x, and U_x is the phase oracle for x, defined by

$$U_x |i\rangle |b\rangle \;=\; |i\rangle |b \oplus x_i\rangle ,$$

where $i \in \{0, \ldots, N-1\}$ is a basis index and $b \in \{0,1\}$. The measurement at the end outputs 1 ("good") if the measured state is a *good basis state*, and 0 otherwise. Note that a unitary transformation preserves inner products, so the structure of amplitudes across the circuit is tightly controlled.

2.2 Amplitudes Are Polynomials

Let $\alpha^t(x_1, \ldots, x_N)$ denote the amplitude of a fixed good basis state after t applications of U_x. A key structural lemma, which we state without proof, is the following.

Lemma 1 (Amplitude Polynomial Lemma)**.** *After T oracle queries, the amplitude $\alpha^T(x_1, \ldots, x_N)$ of any fixed basis state is a multivariate polynomial of degree at most T in the variables $x_1, \ldots, x_N$. Since $x_i \in \{0,1\}$, this polynomial can be taken to be* multilinear.

Proof sketch. Initially (before any queries), the amplitude of each basis state is a constant, i.e., a degree-0 polynomial. Each application of the oracle U_x maps $|i\rangle |b\rangle \mapsto |i\rangle |b \oplus x_i\rangle$, which multiplies the amplitude of certain states by x_i. This increases the polynomial degree by at most 1 per query. The fixed unitaries U_t are linear and do not depend on x, so they do not increase the degree. After T queries, the degree is therefore at most T. □

2.3 The Approximating Polynomial

For a T-query algorithm to compute φ with error at most $1/3$, the total probability of measuring a good basis state must approximate $\varphi(x)$ to within $1/3$. That is, defining

$$p(x_1, \ldots, x_N) \;:=\; \sum_{y_1, \ldots, y_m} \left|\alpha^T(x_1, \ldots, x_N)\right|^2,$$

we require

$$\left|p(x_1,\ldots,x_N) - \varphi(x_1,\ldots,x_N)\right| \;\le\; \frac{1}{3} \quad \text{for all } x \in \{0,1\}^N.$$

Since each $|\alpha^T|^2$ is a polynomial of degree at most $2T$, the function p is a real polynomial of degree at most $2T$ that $\frac{1}{3}$-approximates φ on all Boolean inputs.

Definition 2 (Approximate degree)**.** The *approximate degree* $\widetilde{\deg}(\varphi)$ of a Boolean function φ is the minimum degree of a real polynomial p such that $|p(x) - \varphi(x)| \le 1/3$ for all $x \in \{0,1\}^N$.

Theorem 3 (Polynomial Method Lower Bound [dW23])**.** *For any Boolean function φ,*

$$Q(\varphi) \;\ge\; \frac{\widetilde{\deg}(\varphi)}{2}.$$

Proof. As argued above, any T-query quantum algorithm produces a polynomial p of degree at most $2T$ that $\frac{1}{3}$-approximates φ. By the definition of approximate degree, $\widetilde{\deg}(\varphi) \le 2T$. Rearranging gives $T \ge \widetilde{\deg}(\varphi)/2$. □

To obtain a concrete lower bound, one must show that $\widetilde{\deg}(\varphi)$ is large—i.e., that no low-degree polynomial can approximate φ on Boolean inputs.

2.4 Application: The OR Function

Example 4 (OR lower bound)**.** The N-bit OR function is defined by

$$\mathrm{OR}(x_1,\ldots,x_N) \;=\; 1 - (1-x_1)(1-x_2)\cdots(1-x_N).$$

One can show (using symmetrization and Chebyshev polynomial arguments) that $\widetilde{\deg}(\mathrm{OR}) = \Theta(\sqrt{N})$. By Theorem 3, any quantum algorithm for OR requires $\Omega(\sqrt{N})$ queries. Grover's algorithm [Gro96] achieves this bound with $O(\sqrt{N})$ queries, so the bound is tight.

Remark 5. For the *exact* computation of OR (zero error), the polynomial must equal φ on all inputs, so its degree equals N. This gives a lower bound of $N/2$ queries for exact OR, consistent with the fact that classically N queries are needed. Quantumly, exact OR has query complexity $\Theta(\sqrt{N \log N})$ [dW23].

Remark 6 (Grover with unknown number of solutions). Grover's algorithm in its basic form requires knowing (or estimating) the number of marked items t. When t is unknown, one can either: (1) allow a small constant error and use amplitude estimation, or (2) use an exponential search strategy (try $T = 1, 2, 4, \ldots, 2^k$ iterations), incurring only a constant factor overhead. Both approaches still achieve $O(\sqrt{N})$ query complexity.

3 The Adversary Method

The adversary method, introduced by Ambainis [dW23], gives lower bounds by analyzing how quickly an algorithm can "distinguish" a YES input from a NO input.

3.1 Setup

Consider a Boolean function $\varphi : \{0,1\}^N \to \{0,1\}$. Let $\mathrm{YES} = \varphi^{-1}(1)$ and $\mathrm{NO} = \varphi^{-1}(0)$.
A T-query quantum algorithm processes inputs $y \in \mathrm{YES}$ and $z \in \mathrm{NO}$ starting from the same initial state $|0\rangle^{\otimes m}$. After t queries on input y, the state of the algorithm is some $|\psi_y^t\rangle$; on input z, it is $|\psi_z^t\rangle$.

- **Initially** ($t = 0$): $|\psi_y^0\rangle = |\psi_z^0\rangle = |0\rangle^{\otimes m}$, so $\langle\psi_y^0|\psi_z^0\rangle = 1$.
- **Finally** ($t = T$): the algorithm must output different answers on y and z, so $|\langle\psi_y^T|\psi_z^T\rangle| < 0.9$ (the states must be nearly orthogonal for the measurement to distinguish them reliably).

The *progress measure* is the inner product $W^t = \langle\psi_y^t|\psi_z^t\rangle$. Since unitaries preserve inner products, only the oracle queries U_x can change W^t. In particular, each query decrements $|W^t|$ by at most $O(1/\sqrt{N})$ in the worst case (depending on the problem structure). Over T queries, the total change is at most $T \cdot O(1/\sqrt{N})$. For the algorithm to succeed, the total change must be at least $|W^0| - |W^T| \approx 0.1$, forcing $T = \Omega(\sqrt{N})$ for the search problem.

3.2 The Adversary Theorem

We now state the (standard) adversary lower bound theorem.

Theorem 7 (Adversary Method [dW23])**.** *Let $\varphi : \{0,1\}^N \to \{0,1\}$ be a Boolean function. Suppose there exist sets $Y \subseteq \mathrm{YES}$ and $Z \subseteq \mathrm{NO}$, and positive integers M and M', such that*

1. *For every $y \in Y$, there are at least M strings $z \in Z$ with $\mathrm{dist}(y,z) = 1$, and*
2. *For every $z \in Z$, there are at least M' strings $y \in Y$ with $\mathrm{dist}(y,z) = 1$,*

where $\mathrm{dist}(y,z)$ denotes the Hamming distance between y and z. Then

$$Q(\varphi) \;\geq\; c\sqrt{M \cdot M'}$$

for an absolute constant $c > 0$.

Proof sketch. Define the *progress function*

$$W^t \;=\; \sum_{y \in Y,\, z \in Z} \langle \psi_y^t \mid \psi_z^t \rangle.$$

Initial value. Since $|\psi_y^0\rangle = |\psi_z^0\rangle$ for all y, z, we have $W^0 = |Y||Z|$.
Final value. For the algorithm to be correct, the states $|\psi_y^T\rangle$ and $|\psi_z^T\rangle$ must be nearly orthogonal for every (y,z) pair that differ only in the function value. One can show that $|W^T|$ is small—specifically, $|W^T| \leq \sqrt{|Y||Z|}/2$ under some normalizations.
Change per query. Each oracle query on input x applies U_x. One can bound the change $|W^{t+1} - W^t|$ using the Cauchy–Schwarz inequality and the degree-1 neighbor structure:

$$|W^{t+1} - W^t| \;\leq\; \sqrt{M \cdot M'} \cdot (\text{normalized norms}).$$

Since $W^0 - W^T \geq |Y||Z|/2$ and each query changes W by at most $O(\sqrt{M \cdot M'})$, we need at least $\Omega\Big(\sqrt{M \cdot M'}\Big)$ queries. □

Remark 8. The condition $\mathrm{dist}(y,z) = 1$ is natural: if y and z differ on only one bit i, then querying bit i is the *only* query that can distinguish them. The parameters M and M' capture how "spread out" this difficulty is: if every YES instance has many nearby NO instances and vice versa, the algorithm must do substantial work.

3.3 Application: Search / OR Lower Bound

Example 9. For the N-bit OR (or unstructured search) function, take Y to be all weight-1 strings (exactly one 1), and $Z = \{0^N\}$. Then:

- Every $y \in Y$ has exactly 1 neighbor $z = 0^N \in Z$ at Hamming distance 1, so $M = 1$.
- The unique $z = 0^N$ has N neighbors in Y (flip any one of the N bits), so $M' = N$.

By Theorem 7, $Q(\mathrm{OR}) \geq c\sqrt{1 \cdot N} = c\sqrt{N}$, i.e., $Q(\mathrm{OR}) = \Omega(\sqrt{N})$. This matches Grover's upper bound [Gro96].

References

[dW23] Ronald de Wolf. Quantum computing: Lecture notes, 2023.

[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*, STOC '96, page 212–219, New York, NY, USA, 1996. Association for Computing Machinery.