COL7160 : Quantum Computing

Lecture 21 : Amplitude Estimation and Quantum Query Lower Bounds

Instructor: Rajendra Kumar **Scribe:** Parth Gupta

1 Amplitude Estimation

Following our discussion on Grover's family of algorithms, we analyze the precision of **Amplitude Estimation**. As a recap, given a unitary A such that:

$$A|0\rangle = \sqrt{p}\,|\Psi_{good}\rangle + \sqrt{1-p}\,|\Psi_{bad}\rangle$$

where $\sqrt{p} = \sin\theta$, the Grover iterate G acts as a rotation in the 2D subspace spanned by $\{|\Psi_{good}\rangle, |\Psi_{bad}\rangle\}$.

1.1 Eigenstructure and Phase Estimation

The Grover operator $G = A(2|0\rangle\langle 0| - I)A^{-1}Z_f$ has eigenvectors:

$$|\Psi_{\pm}\rangle = \frac{1}{\sqrt{2}}(|\Psi_{good}\rangle \pm i\,|\Psi_{bad}\rangle)$$

with corresponding eigenvalues $e^{\pm i2\theta}$. When we apply Quantum Phase Estimation (QPE) using m qubits on the state $A|0\rangle$, the initial state decomposes as:

$$A|0\rangle = \alpha\,|\Psi_{+}\rangle + \beta\,|\Psi_{-}\rangle$$

where $\alpha = \frac{e^{-i(\pi/2-\theta)}}{\sqrt{2}}$ and $\beta = \frac{e^{i(\pi/2-\theta)}}{\sqrt{2}}$.

1.2 Applying Phase Estimation

We apply QPE to the operator G using the initial state $A|0\rangle$. The circuit uses m ancilla qubits. Upon measurement, we obtain a bitstring y which can be viewed as an integer $y \in \{0, \ldots, 2^m - 1\}$ or as a binary fraction $0.y_1y_2\ldots y_m$.

The probability distribution of y shows two distinct peaks. These peaks correspond to the best m-bit approximations of the phases associated with

the eigenvalues $e^{\pm i2\theta}$. Specifically, the measured value y (viewed as a fraction $\frac{y}{2^m}$) approximates:

$$\frac{y}{2^m} \approx \frac{2\theta}{2\pi} \quad \text{or} \quad \frac{y}{2^m} \approx 1 - \frac{2\theta}{2\pi}$$

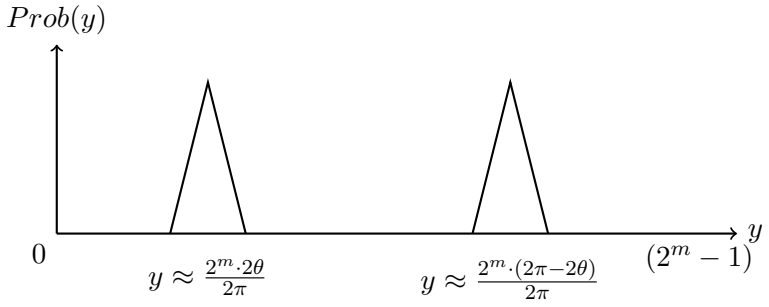

From the measured integer y, we recover our estimate θ^*:

$$\theta^* = \frac{\pi y}{2^m}$$

1.3 Error Analysis and Complexity

Applying Phase Estimation to the initial state $A|0\rangle$ effectively processes the eigenstates $|\Psi_\pm\rangle$ in parallel. Measuring the register yields an integer y such that $\frac{2\pi y}{2^m}$ approximates the phase 2θ (from $|\Psi_+\rangle$) or $2\pi - 2\theta$ (from $|\Psi_-\rangle$).

- **Precision:** With m ancilla qubits, the measurement collapses the state to the nearest integer y. This provides an estimate $\theta^* = \frac{\pi y}{2^m}$ such that $|\theta^* - \theta| \leq \frac{\pi}{2^m}$.
- **Success Probability:** This angular precision translates to an estimate of the probability $p = \sin^2\theta$ such that $|\sin^2\theta - \sin^2\theta^*| \leq C \cdot 2^{-m}$ for some constant C.
- **Query Comparison:**
  - **Classical:** To achieve an additive error ϵ in estimating p, a classical algorithm requires $O(1/\epsilon^2)$ queries via standard sampling.
  - **Quantum:** In QPE, the implementation requires controlled unitaries $G, G^2, G^4, \ldots, G^{2^{m-1}}$. The only way to create the unitary G^{2^i} is to apply G 2^i times; i.e., it requires 2^i queries to Z_f. The total number of queries is $\sum_{i=0}^{m-1} 2^i = 2^m - 1$. Since $2^m \approx 1/\epsilon$ is sufficient for the desired precision, the quantum complexity is $O(1/\epsilon)$.

2 Quantum Query Lower Bound

We now shift from designing algorithms to understanding the fundamental limits of the **black-box query model**.

2.1 The Black-Box Model

Given a function $f : \{0,1\}^n \to S$, where S is usually $\{0,1\}^m$ or $\{0,1\}$, we wish to solve a problem Φ about f using a quantum circuit that makes k calls (queries) to f.

2.2 Search vs. Decision Problems

- **Total Functions:** The function $f : \{0,1\}^n \to \{0,1\}$ is defined for *all* possible inputs in the domain. An example is Decision Grover (determining if *any* x exists such that $f(x) = 1$).
- **Partial / Promise Functions:** The function is only defined on a subset of the domain. We are "promised" that the input x satisfies a certain property (e.g., in Simon's problem, we are promised f is either 1-to-1 or 2-to-1).

We distinguish between finding a witness and deciding a property:

1. **Simon's Problem:** Find s such that $f(x) = f(y)$. The promise is that $f(x) = f(y)$ iff $y \in \{x, x \oplus s\}$.
2. **Grover's Problem:** Find x such that $f(x) = 1$.
3. **Decision Simon (Partial/Promise):** Decide if $s \neq 0^n$.
4. **Decision Grover (Total):** Decide if there exists *any* x such that $f(x) = 1$.

2.3 Query Complexity Comparison

The following table summarizes the queries required for decision versions of these problems:

Model	Decision Simon	Decision Grover
Deterministic	$\sqrt{2^n}$	2^n
Randomized	$\sqrt{2^n}$	$\frac{2}{3} \cdot 2^n$
Quantum	$O(n)$	$\sqrt{2^n}$

2.4 Complexity Bounds for Total Boolean Functions

For total Boolean functions, the gap between classical and quantum query complexities is limited to a polynomial factor.

Deterministic vs. Quantum Complexity: Based on Beals et al. (1998), the deterministic query complexity $D(f)$ is at most the sixth power of the quantum query complexity $Q(f)$:

$$D(f) \leq O(Q(f)^6) \quad \text{or equivalently} \quad Q(f) \geq \Omega\left(D(f)^{1/6}\right)$$

This theorem implies that for total functions, quantum computers cannot provide an exponential speedup; the maximum possible speedup is characterized by this sixth-power relation.

3 The Polynomial Method in Quantum Query Complexity

The **Polynomial Method** is a fundamental technique used to prove lower bounds on the number of queries required by quantum algorithms. The core principle is that the behavior of a quantum algorithm can be captured by a low-degree multilinear polynomial.

3.1 Representing Quantum Algorithms as Polynomials

A quantum algorithm making T queries to an input $x \in \{0,1\}^n$ produces a final state where the amplitudes are polynomials in $x_1, \ldots, x_n$ of degree at most T. Consequently, the probability of the algorithm accepting (outputting 1) is a polynomial of degree at most $2T$.

Let $\mathcal{A}$ be a quantum algorithm that makes T queries to an input x. Then there exists a real-valued multilinear polynomial $p(x_1, \ldots, x_n)$ of degree $\deg(p) \leq 2T$ such that for all $x \in \{0,1\}^n$:

$$p(x) = \Pr[\mathcal{A} \text{ accepts } x]$$

3.2 Lower Bounds via Approximate Degree

To compute a Boolean function f with error at most $1/3$, a quantum algorithm must correspond to a polynomial p such that $|p(x) - f(x)| \leq 1/3$ for all x. This leads to the definition of the **approximate degree**:

- **Approximate Degree** $\widetilde{\deg}(f)$: The minimum degree of a real polynomial p such that $|p(x) - f(x)| \leq 1/3$ for all $x \in \{0,1\}^n$.

The relationship between quantum query complexity $Q(f)$ and approximate degree is given by:

$$Q(f) \geq \frac{1}{2}\widetilde{\deg}(f)$$

This method has been used to prove tight lower bounds for functions such as OR, AND, and the parity function, confirming that quantum speedups for these total functions are limited to at most a polynomial factor.

References

[1] R. Beals, H. Buhrman, R. Cleve, M. Mosca, and R. de Wolf. *Quantum Lower Bounds by Polynomials*. FOCS 1998. `https://homepages.cwi.nl/~rdewolf/publ/qc/polynomials.pdf`

[2] A. Ambainis, B. Bačkurs, K. Itze, R. de Wolf, and J. Vihotss. *Separations Between Open and Closed Set Query Complexity*. arXiv:1312.0036. `https://arxiv.org/abs/1312.0036`

[3] R. de Wolf, *Quantum Computing: Lecture Notes*, 2019. `https://arxiv.org/abs/1907.09415`