Quantum Computation **(CMU 18-859BB, Fall 2015)**

Lecture 13: Lower Bounds using the Adversary Method

October 21, 2015

Lecturer: Ryan O'Donnell *Scribe: Kumail Jaffer*

1 Introduction

There are a number of known methods to show lower bounds in the quantum query model.

(0) Polynomial Method [BBC+01]

(1) Hybrid Method [BBBV97]

(2) Basic Adversary Method [Amb02]

(3) General Adversary Method [HLS07]

So far in this course we've seen (0).

In this lecture, we will develop a basic version of (2), the already basic adversary method. A generalized form of method (2), method (3), is extremely powerful. Indeed, it can be shown that every lower bound shown using this method has a corresponding upper bound (!) [Rei09]

2 The Super-Basic Adversary Method [Amb02]

To start, let's recall the quantum query model for decision problems. We are given a decision problem in the form of a function $F : \{0,1\}^n \to \{0,1\}$, and an N bit input string w, which we can access via an oracle $O_w^{\pm} : |i\rangle \mapsto (-1)^{w_i} |i\rangle$, perhaps with some promise on the form w can take. Our task is to determine $F(w)$ with high probability (say greater than $2/3$) for arbitrary w in the domain of the problem, using as few queries to $O_w^{\pm}$ as possible.

What does a circuit solving such a problem look like? A T-query circuit for such a problem, using a ancilla bits will look as follows.

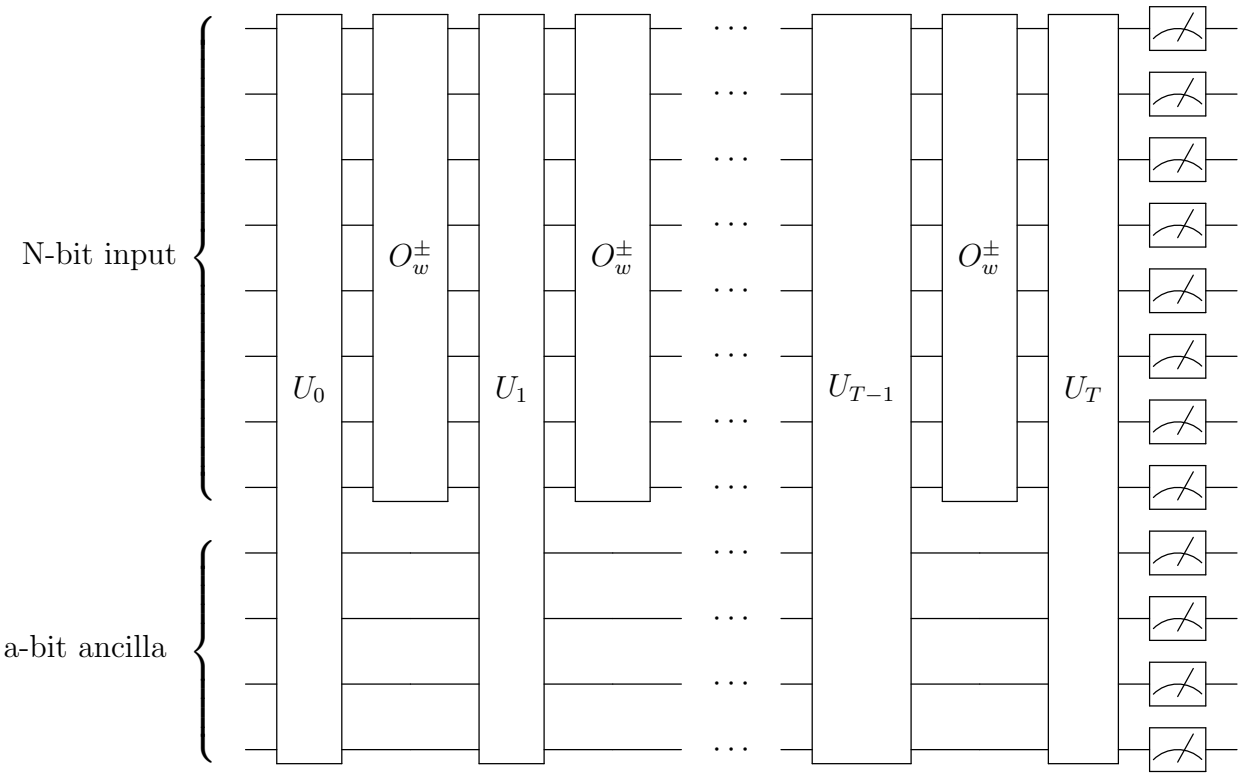

Figure 1: T-query circuit

That is, we call the oracle T times, with some unitary transformations between successive calls, and then measure at the end, and then use the measurement to determine our answer. Since all quantum operations are unitary, we can collapse any number of intermediate quantum gates into one big unitary transformation.

The super high level idea of the adversary method is to show that each call to the oracle can only give us a limited amount of progress toward our goal. More specifically, consider the following progress measure.

Definition 2.1 (Progress Measure)**.** Let $|\psi_w^t\rangle$ be the state of the system after the t'th call to the oracle, and fix some y, z such that $F(y) = 1$, $F(z) = 0$. Then the progress function is $\Phi(t) = \langle\psi_z^t|\psi_y^t\rangle$

Right off the bat we can notice some useful properties

Observation 2.2. Φ *is not affected by the non-oracle unitary transformations.*

This is because unitary transformations preserve inner product by definition. Notice that this does not apply to calls to the oracle, because y and z have different oracles. So it suffices to only examine what happens before and after calls to the oracles.

Observation 2.3. $\Phi(0) = 1$

This is because before application of the first oracle, $|\psi_y^0\rangle$ and $|\psi_z^0\rangle$ are necessarily the same.

Claim 2.4. *If a T-query circuit solves F, $\Phi(T) \leq 0.999$*

Intuitively, $|\psi_y^T\rangle$ and $|\psi_z^T\rangle$ must be sufficiently different in order for us to be able to distinguish them. If their inner product is not sufficiently small, we'll get the same result after measuring them too often to determine our response with high enough probability. More formally, we have the following.

Lemma 2.5. *Let $|\alpha\rangle = |\psi_y^T\rangle, |\beta\rangle = |\psi_z^T\rangle$. Say that C is some circuit solving F. If $|\langle\alpha|\beta\rangle| \geq 1-\epsilon$, then $|\mathbf{Pr}[C(y)=1] - \mathbf{Pr}[C(z)=1]| \leq \sqrt{2\epsilon}$.*

Proof. Assume $|\langle\alpha|\beta\rangle| \geq 1-\epsilon$. Notice that if we multiply α by a scalar with unit norm, the measurement probabilities don't change. So WLOG we can pick some scalar so that $\langle\alpha|\beta\rangle = |\langle\alpha|\beta\rangle| \geq 1-\epsilon$. Now we have

$$\|\alpha-\beta\|_2^2 = \langle\alpha-\beta|\alpha-\beta\rangle = \langle\alpha|\alpha\rangle + \langle\beta|\beta\rangle - 2\mathbf{Re}\langle\alpha|\beta\rangle$$

And we know that $\langle\alpha|\beta\rangle = \mathbf{Re}\langle\alpha|\beta\rangle \geq 1-\epsilon$, as well as that $\langle\alpha|\alpha\rangle = \langle\beta|\beta\rangle = 1$ since they are valid quantum states. So $\|\alpha-\beta\|_2^2 \leq 2\epsilon$, $\|\alpha-\beta|_2 \leq \sqrt{2\epsilon}$.

Now we think of the outcome of measurement as a probability distribution over strings, where the probability of seeing x when we are in state α is $|\alpha_x|^2$, and similarly for β. We want to show that the total variation distance is at most $\sqrt{2\epsilon}$. This is

$$\begin{aligned}
\frac{1}{2}\sum_x ||\alpha_x|^2 - |\beta_x|^2| &= \frac{1}{2}\sum_x ||\alpha_x| - |\beta_x|| \cdot ||\alpha_x| + |\beta_x|| \\
&\leq \frac{1}{2}\sum_x |\alpha_x - \beta_x| \cdot (|\alpha_x| + |\beta_x|) \\
&\leq \frac{1}{2}\sqrt{\sum_x |\alpha_x - \beta_x|^2}\sqrt{\sum_x (|\alpha_x| + |\beta_x|)^2}
\end{aligned}$$

The second and third line follow from the triangle inequality, and Cauchy-Schwarz respectively. The left factor is $\|\alpha-\beta\|_2$. Using Jensen's inequality, we know that $(a+b)^2 \leq 2a^2+2b^2$, and since $\sum_x |\alpha_x|^2 = \sum_x |\beta_x|^2 = 1$, the right factor is at most 2. So the whole thing is at most $\|\alpha-\beta\|_2 = \sqrt{2\epsilon}$ and we are done.

□

We can now prove claim 2.4

Proof. $\mathbf{Pr}[C(y)=1] \geq 2/3$ and $\mathbf{Pr}[C(z)=1] < 1/3$, so we can conclude $\sqrt{2\epsilon} \geq 1/3$, which means $\epsilon \geq 1/18$, and $\langle\alpha|\beta\rangle \leq 1-\epsilon \leq 1-1/18 \leq 0.999$. □

Now if we can show that the progress function changes by no more than Δ after each call to the oracle, then necessarily $T \leq 0.001/\Delta$, and we have a lower bound.

Of course, to come up with interesting lower bounds, we have to look at more than just one fixed y and z. Otherwise, since y and z must differ in at least one index (which is fixed since y and z are), all we have to do to differentiate between them is to make one query on the index on which they differ. It turns out that if we look instead at some specially constructed sets $Y \subseteq F^{-1}(1), Z \subseteq F^{-1}(0)$, we can come up with some very useful theorems. Here's a simple one.

Theorem 2.6 (Super-Basic Adversary Method)**.** *Let F be a decision problem, $Y \subseteq F^{-1}(1)$, $Z \subseteq F^{-1}(0)$. For every y we define $Z_y = \{z \in Z : d(y,z) = 1\}$, and for every z, $Y_z = \{y \in Y : d(y,z) = 1\}$. If $\forall y \in Y, |Z_y| \geq m$, and $\forall z \in Z, |Y_z| \geq m'$, then the quantum query complexity of F is $\Omega(\sqrt{mm'})$. Here $d(a,b)$ denotes the hamming distance between strings a and b.*

Before we get into the proof of this, let's apply it to some decision problems we understand in order to get a feel for it.

3 Illustrations

3.1 Decision Grover

In the decision version of the Grover search problem, we are given an n bit string and we are tasked with determining whether it is all 0 or not. That is, we'll return 0 if it is all zero, and 1 if not.

Let's look at $Y = \{x : x \text{ contains exactly one } 1\}$, and $Z = \{00000\ldots\}$. Notice Y has size N.

For all $y \in Y$, the all zero string is hamming distance 1 away, so $\forall y \in Y, |Z_y| = 1$. Similarly, since all the elements in Y are hamming distance 1 away from the all zero string, $\forall z \in Z, |Y_z| = N$.

So by the adversary method, we have a lower bound of $\Omega(\sqrt{N})$.

3.2 Variation on Decision Grover

In this variation of the decision Grover problem, we are again given an n bit string. We are tasked with differentiating between inputs with hamming weight $\geq k$, and those with hamming weight $\leq k-1$. We return 1 in the first case and 0 in the second.

Let's look at $Y = \{x : x \text{ has hamming weight } k\}$ and $Z = \{x : x \text{ has hamming weight } k-1\}$. It is easy to see that $\forall y \in Y, |Z_y| = k$, since for any y we can flip each of its k 1s to get something in Z). Similarly, $\forall z \in Z, |Y_z| = N - k + 1$, since we for any z we can flip each of its $N - k + 1$ 0s to get something in Y).

Using the adversary method, we have a lower bound of $\Omega(\sqrt{k(N-k+1)})$. If we make the additional assumption that $k \leq N/2$, then this is $\Omega(\sqrt{kN})$. Making this assumption is

okay, since if $k > N/2$, then we have the symmetric situation of differentiating between at least $N - k$ zeros and less than $N - k$ zeros, where $N - k \leq N/2$, and so our lower bound still applies.

3.3 Read-once CNF

In this problem, we are given an N bit string and are told to interpret it as a conjunctive normal form expression with $\sqrt{N}$ clauses and $\sqrt{N}$ variables per clause. So each bit is the value of a variable, and each set of $\sqrt{N}$ bits is to be interpreted as a clause. The task is to calculate the result.

Consider $Y = \{x : x \text{ has exactly one 1 per clause}\}$ and $Z = \{x : x \text{ has all zeros in one clause, and exactly one 1 in all the others }\}$. $\forall y \in Y, |Z_y| = \sqrt{N}$, since we can take any y and, for each of its clauses, flip the sole 1 to get something in Z. Similarly, $\forall z \in Z, |Y_z| = \sqrt{N}$, since for any z, we can flip each zero in its all zero clause to get something in Y.

Applying the adversary method, we get a lower bound of $\Omega(N)$. This lower bound in particular is very hard to show using the polynomial method, though it is nearly trivial to find with the adversary method.

4 Proof of Super-Basic Adversary Method

Proof. First of all, we need to rewrite our progress measure so that we can make use of multiple ys and zs.

Definition 4.1 (More Useful Progress Measure)**.** Let R be a binary relation defined as $R = \{(y, z) \in Y \times Z : d(y, z) = 1\}$. Then the progress measure $\Phi(t)$ is defined as $\Phi(t) = \sum_{(y,z)\in R} |\langle \psi_y^t | \psi_z^t \rangle|$.

Observation 4.2. *From our prior observations in section 2, we can immediately conclude* $\Phi(0) = |R|$, $\Phi(T) \leq 0.999|R|$, *and that, since unitary transformations don't change our progress measure, we need only concern ourselves with the effect of the oracle.*

Definition 4.3. For each $y \in Y$ let $I_y = \{i : y \text{ with } i\text{'th bit flipped} \in Z\}$.

Observation 4.4. *Notice that* $|I_y| \geq m$. *As a result,* $|R| \geq m|Y|$.

With notation out of the way, let's begin in earnest. We will show that

$$\Phi(t-1) - \Phi(t) \leq \frac{2|R|}{\sqrt{mm'}}$$

This will let us conclude that $T \in \Omega(\sqrt{mm'})$.

Fix some t, and consider some particular $(y, z) \in R$. Let $i^* \in [N]$ be the position on which y and z differ. Before the t'th oracle, a circuit solving the problem will be in some

quantum state $|\psi_y^{t\prime}\rangle$ for y and $|\psi_z^{t\prime}\rangle$ for z. That looks something like this.

$$|\psi_y^{t\prime}\rangle = \sum_i \alpha_i |i\rangle \otimes |\phi_i\rangle$$
$$|\psi_z^{t\prime}\rangle = \sum_i \beta_i |i\rangle \otimes |\chi_i\rangle$$

Where $|\phi_i\rangle, |\chi_i\rangle$ are unit vectors, $\sum_i |\alpha_i|^2 = 1$ and $\sum_i |\beta_i|^2 = 1$. Their inner product is

$$\langle\psi_y^{t\prime}|\psi_z^{t\prime}\rangle = \sum_i \alpha_i\beta_i \langle\phi_i|\chi_i\rangle$$

Now if we pass these states through the oracle to get $|\psi_y^t\rangle$ and $|\psi_z^t\rangle$, what it does is flip the sign on each α_i or β_i, whenever $y_i = 1$ or $z_i = 1$ respectively. The only index on which the sign will differ is i^*, so our new inner product is exactly

$$\langle\psi_y^t|\psi_z^t\rangle = \sum_{i\neq i^*} \alpha_i\beta_i \langle\phi_i|\chi_i\rangle - \alpha_{i^*}\beta_{i^*} \langle\phi_{i^*}|\chi_{i^*}\rangle$$

The difference between the inner products is then

$$\langle\psi_y^{t\prime}|\psi_z^{t\prime}\rangle - \langle\psi_y^t|\psi_z^t\rangle = 2\alpha_{i^*}\beta_{i^*} \langle\phi_{i^*}|\chi_{i^*}\rangle$$

The quantity we want to bound is

$$\begin{aligned}
\Phi(t-1) - \Phi(t) &= \sum_{(y,z)\in R} |\langle\psi_y^{t\prime}|\psi_z^{t\prime}\rangle| - \sum_{(y,z)\in R} |\langle\psi_y^t|\psi_z^t\rangle| \\
&\leq \sum_{(y,z)\in R} |\langle\psi_y^{t\prime}|\psi_z^{t\prime}\rangle - \langle\psi_y^t|\psi_z^t\rangle| \\
&\leq \sum_{(y,z)\in R} |2\alpha_{i^*_{(y,z)}}\beta_{i^*_{(y,z)}} \langle\phi_{i^*_{(y,z)}}|\chi_{i^*_{(y,z)}}\rangle| \\
&\leq \sum_{(y,z)\in R} 2|\alpha_{i^*_{(y,z)}}||\beta_{i^*_{(y,z)}}| \\
&\leq \sum_{(y,z)\in R} \sqrt{\frac{m}{m'}}|\alpha_{i^*_{(y,z)}}|^2 + \sum_{(y,z)\in R} \sqrt{\frac{m'}{m}}|\beta_{i^*_{(y,z)}}|^2
\end{aligned}$$

The last line follows from Jensen's inequality. Consider just the left summand. This is the

same as

$$\begin{aligned}\sqrt{\frac{m}{m'}}\sum_{(y,z)\in R}|\alpha_{i^*_{(y,z)}}|^2 &= \sqrt{\frac{m}{m'}}\sum_{y\in Y}\sum_{i\in I_y}|\alpha_i|^2\\ &\leq \sqrt{\frac{m}{m'}}|Y|\\ &\leq \sqrt{\frac{m}{m'}}\frac{|R|}{m} = \frac{|R|}{\sqrt{mm'}}\end{aligned}$$

The second line comes from the fact that $\sum_i |\alpha_i|^2 = 1$, and the third from Observation 4.4. The right summand is a symmetric situation, and so overall we get

$$\Phi(t-1) - \Phi(t) \leq \frac{2|R|}{\sqrt{mm'}}$$

as promised, concluding the proof. □

5 Generalization

This isn't quite as powerful as you might hope. In particular, it fails to give us a good lower bound when we can't find sets Y and Z which have many strings hamming distance 1 away from each other. As motiviation consider the problem of distinguishing between strings of hamming weight 0 and strings of hamming weight greater than k. In this situation, there are no $y \in F^{-1}(1), z \in F^{-1}(0)$ such that $d(y,z) = 1$.

One natural thing to do is to expand our relation so we're not restricted to just strings hamming distance 1 from each other. This line of reasoning results in the following theorem.

Theorem 5.1 (Basic Adversary Method)**.** *Let F be a decision problem, $Y \subseteq F^{-1}(1), Z \subseteq F^{-1}(0)$, and a binary relation $R \subseteq Y \times Z$. If*

1. *$\forall y \in Y$ there are m distinct $z \in Z$ such that $(y,z) \in R$*
2. *$\forall z \in Z$ there are m' distinct $y \in Y$ such that $(y,z) \in R$*
3. *$\forall y \in Y, i$ there are at most l distinct $z \in Z$ such that $y_i \neq z_i$ and $(y,z) \in R$.*
4. *$\forall z \in Z, i$ there are at most l' distinct $y \in Y$ such that $y_i \neq z_i$ and $(y,z) \in R$.*

then the quantum query complexity of F is $\Omega(\sqrt{\frac{mm'}{ll'}})$.

The proof for this is very similar to the one we just did. In later lectures, we will dive deeper into further generalizations.

References

[Amb02] A. Ambainis. Quantum Lower Bounds by Quantum Arguments. *Journal of Computer and System Sciences*, 64(4):750 – 767, 2002.

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

[BBC+01] Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald de Wolf. Quantum lower bounds by polynomials. *J. ACM*, 48(4):778–797, July 2001.

[HLS07] P. Hoyer, T. Lee, and R. Spalek. Negative Weights Make Adversaries Stronger. In *Proceedings of the 39th ACM Symposium on the Theory of Computing. ACM*, pages 526–535, November 2007.

[Rei09] B. W. Reichardt. Span Programs and Quantum Query Complexity: The General Adversary Bound is Nearly Tight for Every Boolean Function. In *Foundations of Computer Science, 2009. FOCS '09. 50th Annual IEEE Symposium on*, pages 544–551, October 2009.