6.080/6.089 GITCS	May 6-8, 2008

Lecture 22/23

Lecturer: Scott Aaronson — *Scribe: Chris Granade*

1 Quantum Mechanics

1.1 Quantum states of n qubits

If you have an object that can be in two perfectly distinguishable states $|0\rangle$ or $|1\rangle$, then it can also be in a superposition of the $|0\rangle$ and $|1\rangle$ states:

$$\alpha\,|0\rangle + \beta\,|1\rangle$$

where α and β are complex numbers such that:

$$|\alpha|^2 + |\beta|^2 = 1$$

For simplicity, let's restrict to real amplitudes only. Then, the possible states of this object–which we call a quantum bit, or qubit– lie along a circle.

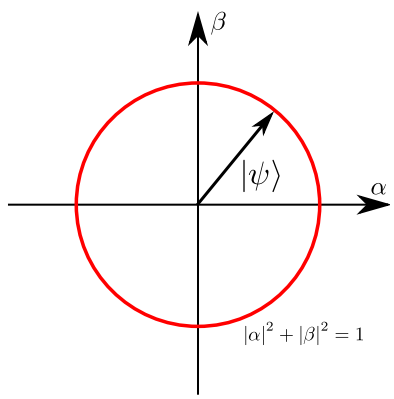

Figure 1: An arbitrary single-qubit state $|\psi\rangle$ drawn as a vector.

If you measure this object in the "standard basis," you see $|0\rangle$ with probability $|\alpha|^2$ and $|1\rangle$ with probability $|\beta|^2$. Furthermore, the object "collapses" to whichever outcome you see.

1.2 Quantum Measurements

Measurements (yielding $|x\rangle$ with probability $|\alpha_x|^2$) are *irreversible, probabilistic, and discontinuous.*

As long as you don't ask specifically what a measurement *is*—how the universe knows what constitutes a measurement and what doesn't—but just assume it as an axiom, everything is well-defined mathematically. If you do ask, you enter a no-man's land. Recently there's been an important set of ideas, known as decoherence theory, about how to explain measurement as ordinary unitary interaction, but they still don't explain where the probabilities come from.

1.3 Unitary transformations

But this is not yet interesting! The interesting part is what else we can do the qubit, besides measure it right away. It turns out that, by acting on a qubit in a suitable way–in the case of an electron, maybe shining a laser on it–we can effectively multiply the vector of amplitudes by any matrix that preserves the property that the probabilities sum to 1. By which I mean, any matrix that always maps unit vectors to other unit vectors. We call such a matrix a unitary matrix. Unitary transformations are *reversible, deterministic, and continuous.*

Examples of unitary matrices:

- The identity I.
- The NOT gate $X = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$.
- The phase-i gate $\begin{bmatrix} 1 & 0 \\ 0 & i \end{bmatrix}$.
- 45-degree counterclockwise rotation.

Physicists think of quantum states in terms of the Schrödinger equation, $\frac{d|\psi\rangle}{dt} = iH\,|\psi\rangle$ (perhaps the third most famous equation in physics after $e = mc^2$ and $F = ma$). A unitary is just the result of leaving the Schrödinger equation "on" for a while.

> **Q:** Why do we use complex numbers?
> **Scott:** The short answer is that it works! A "deeper" answer is that if we used real numbers only, it would not be possible to divide a unitary into arbitrarily small pieces. For example, the NOT gate we saw earlier can't be written as the square of a real-valued unitary matrix. We'll see in a moment that you can do this if you have complex numbers.

For each of these matrices, what does it do? Why is it unitary? How about this one?

$$\begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix}$$

Is it unitary? Given a matrix, how do you decide if it's unitary or not?

Theorem 1 *U is unitary if and only if $UU^* = I$, where U^* means you transpose the matrix and replace every entry by its complex conjugate. (A nice exercise if you've seen linear algebra.) Equivalently, $U^{-1} = U^*$. One corollary is that every unitary operation is reversible.*

As an exercise for the reader, you can apply this theorem to find which of the matricies we've already seen are unitary.

Now, let's see what happens when we take the 45-degree rotation matrix, and apply it twice to the same state.

$$\begin{aligned}
|0\rangle &\rightarrow (|0\rangle + |1\rangle)/\sqrt{2} \\
|1\rangle &\rightarrow (-|0\rangle + |1\rangle)/\sqrt{2} \\
(|0\rangle + |1\rangle)/\sqrt{2} &\rightarrow \left[\frac{|0\rangle + |1\rangle}{\sqrt{2}} + \frac{-|0\rangle + |1\rangle}{\sqrt{2}}\right]/\sqrt{2} \\
&= |1\rangle
\end{aligned}$$

This matrix acts as the "square root of NOT"! Another way to see that is by squaring the matrix.

$$\begin{bmatrix} \cos(45^\circ) & -\sin(45^\circ) \\ \sin(45^\circ) & \cos(45^\circ) \end{bmatrix}^2 |\psi\rangle = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} |\psi\rangle$$

Already, we have something that doesn't exist in the classical world.

We can also understand the action of this matrix in terms of interference of amplitudes.

2 Two Qubits

To describe two qubits, how many amplitudes do we need? Right, four – one for each possible two-bit string.

$$\alpha |00\rangle + \beta |01\rangle + \gamma |10\rangle + \delta |11\rangle$$

$$|\alpha|^2 + |\beta|^2 + |\gamma|^2 + |\delta|^2 = 1$$

If you measure both qubits, you'll get $|00\rangle$ with probability $|\alpha|^2$, $|01\rangle$ with probability $|\beta|^2$, etc. And the state will collapse to whichever 2-bit string you see.

But what happens if you measure only the first qubit, not the second? With probability $|\alpha|^2 + |\beta|^2$, you get $|0\rangle$, and the state collapses to $\frac{\alpha|00\rangle + \beta|01\rangle}{\sqrt{|\alpha|^2 + |\beta|^2}}$. With probability $|\gamma|^2 + |\delta|^2$, you get $|1\rangle$, and the state collapses to $\frac{\gamma|10\rangle + \delta|11\rangle}{\sqrt{|\gamma|^2 + |\delta|^2}}$. Any time you ask the universe a question, it makes up its mind; any time you don't it ask a question, it puts off making up its mind for as long as it can.

What happens if you apply a NOT gate to the second qubit? Answer: You get $\beta |00\rangle + \alpha |01\rangle + \delta |10\rangle + \gamma |11\rangle$. "For every possible configuration of the other qubits, what happens if I apply the gate to this qubit?" If we consider $(\alpha, \beta, \gamma, \delta)$ as a vector of four complex numbers, what does this transformation look like as a 4×4 matrix?

$$\begin{bmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

Can we always factor a two-qubit state: "here's the state of the first qubit, here's the state of the second qubit?" Sometimes we can:

- $|01\rangle = |0\rangle |1\rangle = |0\rangle \otimes |1\rangle$ (read $|0\rangle$ "tensor" $|1\rangle$).
- $|00\rangle + |01\rangle + |10\rangle + |11\rangle = \frac{1}{2} (|0\rangle + |1\rangle) (|0\rangle + |1\rangle)$.

In these cases, we say the state is **separable**. But what about $|00\rangle + |11\rangle$? This is a state that *can't* be factored. We therefore call it an **entangled** state. You might have heard about entanglement as one of the central features of quantum mechanics. Well, here it is.

Just as there are quantum states that can't be decomposed, there are also *operations* that can't be decomposed. Perhaps the simplest is the **Controlled-NOT**, which maps $|x\rangle |y\rangle$ to $|x\rangle |x \oplus y\rangle$ (i.e., flips the second bit iff the first bit is 1).

$$\begin{aligned} |00\rangle &\rightarrow |00\rangle \\ |01\rangle &\rightarrow |01\rangle \\ |10\rangle &\rightarrow |11\rangle \\ |11\rangle &\rightarrow |10\rangle \end{aligned}$$

What does this look like as a 4×4 matrix?

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

Incidentally, could we have a 2-qubit operation that mapped $|x\rangle |y\rangle$ to $|x\rangle |x \text{ AND } y\rangle$? Why not?

$$\begin{aligned} |0\rangle |0\rangle &\rightarrow |0\rangle |0\rangle \\ |0\rangle |1\rangle &\rightarrow |0\rangle |0\rangle \end{aligned}$$

This is not reversible!

2.1 Obtaining Entanglement

Before we can create a quantum computer, we need some way to entangle the qubits so they're not just a bunch of particles laying around. Perhaps the simplest such operation is the CNOT gate that we saw earlier.

So how do we use CNOT to produce entanglement? We can use a Hadamard followed by a CNOT, where the Hadamard matrix $\boxed{\text{H}}$ puts a qubit into superposition by switching between the $\{ |0\rangle, |1\rangle \}$ basis and the $\{ |+\rangle, |-\rangle \}$ basis.

$$\begin{aligned} |+\rangle &= \frac{1}{\sqrt{2}} (|0\rangle + |1\rangle) \\ |-\rangle &= \frac{1}{\sqrt{2}} (|0\rangle - |1\rangle) \\ \boxed{\text{H}} &= \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \end{aligned}$$

Applying $\boxed{\text{H}}$ to $|0\rangle$ and $|1\rangle$ results in:

$$\begin{aligned} |0\rangle &\rightarrow \frac{|0\rangle + |1\rangle}{\sqrt{2}} = |+\rangle \\ |1\rangle &\rightarrow \frac{|0\rangle - |1\rangle}{\sqrt{2}} = |-\rangle \end{aligned}$$

Already with two qubits, we're in a position to see some profound facts about quantum mechanics that took people decades to understand.

Think again about the state $|00\rangle + |11\rangle$. What happens if you measure just the first qubit? Right, with probability $1/2$ you get $|00\rangle$, with probability $1/2$ you get $|11\rangle$. Now, why might that be disturbing? Right: because the second qubit might be light-years away from the first one! For a measurement of the first qubit to *affect the second qubit* would seem to require faster-than-light communication! This is what Einstein called "spooky action at a distance."

But think about it more carefully. Can you actually use this effect to send a message faster than light? What would happen if you tried? Right, the result would be random! In fact, we're not going to prove it here, but there's something called the *no-communication theorem*, which says *nothing* you do to the first qubit only can affect the probability of any measurement outcome on the second qubit only.

But in that case, why can't we just imagine that at the moment these two qubits were created, they flipped a coin, and said, "OK, if anyone asks, we'll both be 1." Well, because in 1964, John Bell proved there are certain experiments where no explanation of that kind can possibly agree with quantum mechanics. And in the 1980s, the experiments were actually done, and they vindicated quantum mechanics and in most physicists' view, dashed Einstein's hope for a "completion" of quantum mechanics. That's on your problem set.

2.2 No-Cloning Theorem

Is it possible to duplicate a quantum state? This would be very nice, since we know we only have one chance to measure a quantum state. Here is what such a duplication would look like:

$$\alpha |0\rangle + \beta |1\rangle \rightarrow (\alpha |0\rangle + \beta |1\rangle)(\alpha |0\rangle + \beta |1\rangle) = \alpha^2 |00\rangle + \alpha\beta |01\rangle + \alpha\beta |10\rangle + \beta^2 |11\rangle$$

This operation is not possible because it is not linear. The final amplitudes α^2, β^2 and $\alpha\beta$ don't depend linearly on α and β. That's the **no-cloning theorem**, and it's really as simple as it looks.

3 n Qubits

For 60 years, *these* were the sorts of examples that drove people's intuitions about quantum mechanics: one particle, occasionally two particles. Rarely did people think abstractly about hundreds or thousands of particles all entangled with one another. But within the last 15 years, we've realized that's where things get *really* crazy. And that brings us to quantum computing. It goes without saying that I'm going to present just the theory at first. Later we can discuss where current experiments are.

How many amplitudes would we need to describe the state of 1000 qubits? Right, 2^{1000}. One for every possible string of 1000 bits:

$$\sum_{x \in \{0,1\}^{1000}} \alpha_x |x\rangle$$

Think about what this *means*. To keep track of the state of 1000 numbers, Nature, off to the side somewhere, apparently has to write down this list of 2^{1000} complex numbers. That's more numbers than there are atoms in the visible universe. Think about how much *computing power* Nature must be expending for that. What a colossal waste! The next thought: we might as well try and take advantage of it!

Q: Doesn't a single qubit already require an infinite amount of information to specify?

Scott: The answer is yes, but there is always noise and error in the real world, so we only care about approximating the amplitudes to some finite precision. In some sense, the "infinite amount of information" is just an artifact of our mathematical description of the qubit's state. By contrast, the exponent in the description of n entangled particles is not an artifact; it's real (if quantum mechanics is the right description of Nature).

3.1 Exploiting Interference

What's an immediate difficulty with taking advantage of this computational power? Well, if we simply measure n qubits, all we get is a classical n-bit string; everything else disappears. It's like the instant we look, nature tries to "hide" the fact that it's doing an exponential amount of computation.

But luckily for us, Nature doesn't always do a good job of hiding. A good example of this is the double-slit experiment: we don't measure which of the two slits the photon passed through, but rather the resulting interference pattern. In particular, we saw that the different paths taken by a quantum system can *interfere destructively* and cancel each other out.

So *that's* what we want to exploit in quantum computing. The goal is to choreograph things so that the different computational paths leading to a given wrong answer interfere destructively and cancel each out, while the different paths leading to a given right answer interfere constructively, hence the right answers are observed with high probability when we measure. You can see how this is gonna be tricky, if it's possible at all.

A key point about interference is that for two computation paths to destructively interfere with each other, they must lead to outcomes that are identical in *every respect*. To calculate the amplitude of a given outcome, you add up the amplitudes for all of the paths leading to that outcome; destructive interference is when the amplitudes cancel each other out.

3.2 Universal Set of Quantum Gates

Concretely, in a quantum computer we have n qubits, which we assume for simplicity start out all in the $|0\rangle$ state. Given these qubits, we apply a sequence of unitary transformation called "quantum gates." These gates form what's called a *quantum circuit*.

An example of such a circuit is shown below, where we apply the Hadamard to the first qubit, then do a CNOT with the second qubit acting as the control bit. Written out, the effect is $\left(\frac{|0\rangle+|1\rangle}{\sqrt{2}}\right)|0\rangle \xrightarrow{\text{CNOT}} \frac{|00\rangle+|11\rangle}{\sqrt{2}}$, the result being entangled qubits, as we discussed before. A crucial

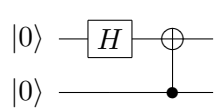

Figure 2: Entangling two qubits

point: each individual gate in a quantum circuit has to be extremely "simple", just like a classical circuit is built of AND, OR, NOT gates, the simplest imaginable building blocks. What does "simple" mean in the quantum case? Basically, that each quantum gate acts on at most (say) 2

or 3 qubits, and is the identity on all the other qubits. Why do we need to assume this? *Because physical interactions are local.*

To work with this constraint, we want a *universal set of quantum gates* that we can use to build more complex circuits, just like AND, OR, and NOT in classical computers. This universal set must contain 1-, 2-, and 3-qubit gates that can be combined to produce any unitary matrix.

We have to be careful when we say *any* unitary matrix, since there are uncountably infinitely many unitary matrices (you can rotate by any real-number angle, for instance). However, there are small sets of quantum gates that can be used to *approximate* any unitary matrix to arbitrary precision. As a technical note, the word "universal" has different meanings; for example, we usually call a set of gates universal if it can be used to approximate any unitary matrix *involving real numbers only*; this certainly suffices for quantum computation.

We've already seen the Hadamard and CNOT gates, but unfortunately these aren't sufficient to be a universal set of quantum gates. According to the Gottesman-Knill Theorem, any circuit constructed with just Hadamard and CNOT gates can be simulated efficiently with a classical computer. However, the Hadamard matrix paired with another gate called the **Toffoli gate** (also called controlled-controlled-NOT, or CCNOT) *is* sufficient to be used as a universal set of gates (for real-valued matrices).

The Toffoli gate will act similarly to the CNOT gate, except that we will control based on the first *two* qubits:

$$|x\rangle\,|y\rangle\,|z\rangle \to |x\rangle\,|y\rangle\,|z \oplus xy\rangle$$

where xy indicates the Boolean AND of x and y.

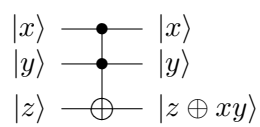

Figure 3: The Toffoli Gate diagram

Note, however, that these are not the only two gates whose combination allows for universal quantum computation. Another example of a universal pair of gates is the CNOT gate taken with the $\pi/8$ gate. We represent the $\pi/8$ gate using the following unitary:

$$T = \begin{bmatrix} \cos(\pi/8) & \sin(\pi/8) \\ -\sin(\pi/8) & \cos(\pi/8) \end{bmatrix}$$

But how many of these gates would be needed to approximate a random n-qubit unitary? Well, you remember Shannon's counting argument? What if we tried something similar in the quantum world? An n-qubit unitary has roughly $2^n \times 2^n$ degrees of freedom. On the other hand, the number of quantum circuits of size T is "merely" exponential in T. Hence, we need $T = \exp(n)$.

We, on the other hand, are only interested in the tiny subset of unitaries that can be built up out of a *polynomial* number of gates. Polynomial time is still our gold standard.

So, a quantum circuit has this polynomial number of gates, and then, at the end, *something* has to be measured. For simplicity, we assume a single qubit is measured. (Would it make a difference if there were intermediate measurements? No? Why not? Because we can simulate measurements using CNOTs.) Just like with BPP, we stipulate that if $x \in L$ (the answer is "yes"), then the

measurement outcome should be $|1\rangle$ with probability at least 2/3, while if $x \notin L$ (the answer is "no"), then the measurement outcome should be $|1\rangle$ with probability at most 1/3.

There's a final, technical requirement. We have to assume there's a classical polynomial-time algorithm to *produce* the quantum circuit in the first place. Otherwise, how do we find the circuit?

The class of all decision problems L that can be solved by such a family of quantum circuits is called BQP (Bounded-Error Quantum Polynomial Time).

4 Bounded-Error Quantum Polynomial Time (BQP)

Bounded-Error Quantum Polynomial Time (BQP) is, informally, the class of problems that can be efficiently solved by a quantum computer.

Incidentally: the idea of quantum computing occurred independently to a bunch of people in the 70s and 80s, but is usually credited to Richard Feynman and David Deutsch. BQP was defined by Bernstein and Vazirani in 1993.

4.1 Requirements for a BQP circuit

To be in BQP, a problem has to satisfy a few requirements:

Polynomial Size. How many of our building-block circuits (e.g., Hadamard and Toffoli) do we need to approximate an arbitrary n-qubit unitary? The answer is the quantum analogue to Shannon's counting argument. An n-qubit unitary has $2^n \times 2^n$ degrees of freedom, and there are doubly-exponentially many of them. On the other hand, the number of quantum circuits of size T is "merely" exponential in T. Hence, "almost all" unitaries will require an exponential number of quantum gates.

However, we are only interested in the small subset of unitaries that can be built using a *polynomial* number of gates. Polynomial time is still the gold standard.

Output. For simplicity, we assume that we measure a single qubit at the end of a quantum circuit. Just like with BPP, we stipulate that:

$$\text{Output} = \begin{cases} \text{if } x \in L: & |1\rangle \text{ with probability} \geq \frac{2}{3} \\ \text{if } x \notin L: & |1\rangle \text{ with probability} \leq \frac{1}{3} \end{cases}$$

Circuit Construction. There is a final technical requirement to constructing quantum circuits. We have to assume that there is a classical polynomial-time algorithm to *produce* the quantum circuit in the first place. Otherwise, how do we find the circuit?

4.2 BQP's Relation to Other Algorithm Families

P $\subseteq$ BQP: A quantum computer can always simulate a classical one (like using an airplane to drive down the highway). We can use the CNOT gate to simulate the NOT gate, and the Toffoli gate to simulate the AND gate.

BPP $\subseteq$ BQP: Loosely speaking, in quantum mechanics we "get randomness for free." More precisely, any time we need to make a random decision, all we need to do is apply a Hadamard

to some $|0\rangle$ qubit, putting it into an equal superposition of $|0\rangle$ and $|1\rangle$ states. Then we can CNOT that bit wherever we needed a random bit. We're not exploiting interference here; we're just using quantum mechanics as a source of random numbers.

BQP $\subseteq$ EXP: In exponential time, we can always write out a quantum state as an exponentially long vector of amplitudes, then explicitly calculate the effect of each gate in a quantum circuit.

BQP $\subseteq$ PSPACE: We can calculate the probability of each measurement outcome $|x\rangle$ by summing the amplitudes of all paths that lead to $|x\rangle$, which only takes polynomial space, as was shown by Bernstein and Vazirani. We won't give a detailed proof here.

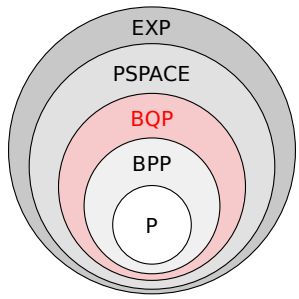

Figure 4: BQP inclusion diagram

We can draw a crucial consequence from this diagram, the first major contribution that complexity theory makes to quantum computing. Namely: in our present state of knowledge, there's little hope of proving unconditionally that quantum computers are more powerful than classical ones, since any proof of $\mathsf{P} \neq \mathsf{BQP}$ would also imply $\mathsf{P} \neq \mathsf{PSPACE}$.

5 Next Time: Quantum Algorithms

Next class we'll see some examples of quantum algorithms that actually outperform their classical counterparts:

- The Deutsch-Jozsa Algorithm
- Simon's Algorithm
- Shor's Algorithm

MIT OpenCourseWare
http://ocw.mit.edu

6.045J / 18.400J Automata, Computability, and Complexity
Spring 2011

For information about citing these materials or our Terms of Use, visit: http://ocw.mit.edu/terms.