Lecture 7, Tues Feb 7: Bloch Sphere, No-Cloning, Wiesner's Quantum Money

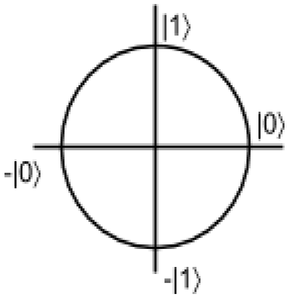

The Bloch Sphere

is a geometric representation of all possible states of a qubit. We've often drawn the state of qubits as a circle, which is already a little awkward: half of the circle is going to waste since $|0\rangle = -|0\rangle$ (both represent the same density matrix).

Instead, what if vectors that pointed in *opposite* directions were orthogonal? We get the Bloch Sphere...

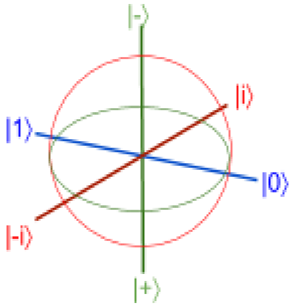

We can see that $|+\rangle$ and $|-\rangle$ should be between $|0\rangle$ and $|1\rangle$. Then we can add $|i\rangle$ and $|-i\rangle$ as a new dimension.

In this representation, points on the surface of the sphere are pure states, such that

if they're 180° apart, they're orthogonal,
and if they're 90° apart, they're conjugate.

What about mixed states?

Well we know that the maximally mixed state, $\frac{I}{2}$, can be defined as $\frac{|0\rangle+|1\rangle}{\sqrt{2}}$, $\frac{|+\rangle+|-\rangle}{\sqrt{2}}$, or $\frac{|i\rangle+|-i\rangle}{\sqrt{2}}$. The sum of any two of these vectors on the sphere is the origin.

We can in this way represent any mixed state as a point inside of the sphere.

The mixture of any states $|v\rangle$ and $|w\rangle$, represented as points on the surface of the sphere, will be a point on the line segment connecting the two.

We can show geometrically that every mixed state can be written as a mixture of only two pure states. Why? Because you can always draw a line that connects any pure state you want to some point in the sphere representing a mixed state, and then see which other pure state that the line intersects on its way out. By some vector math, the point can be described as some linear combination of the vectors representing the pure states.

Experimentalists love the Bloch sphere, because it works almost identically to how spin works with electrons and other spin-½ particles.

With these things, you can measure the particle's "spin"—a qubit attached to the particle, basically—relative to any axis of the sphere. You see if the electron is spinning clockwise or counterclockwise relative to the axis. And it behaves just like a qubit, in that the measurement collapses a more complex behavior into a binary result.

The weird part about spin-½ particles is that you *could have* asked the direction of the spin relative to any other axis. So what's really going on: what's the real spin direction? Well, the actual state is just some point on the Bloch sphere. So if the state of the electron is that it's spinning clockwise around the $(1,0,0)$ axis, we can say that it's in the $|0\rangle$ state, and if it's spinning clockwise around the $(0,1,0)$ axis, we can say that it's in the $|+\rangle$ state, and so forth. The crazy part here is how the three-dimensionality of the Bloch sphere "perfectly syncs up" with the three-dimensionality of actual physical space.

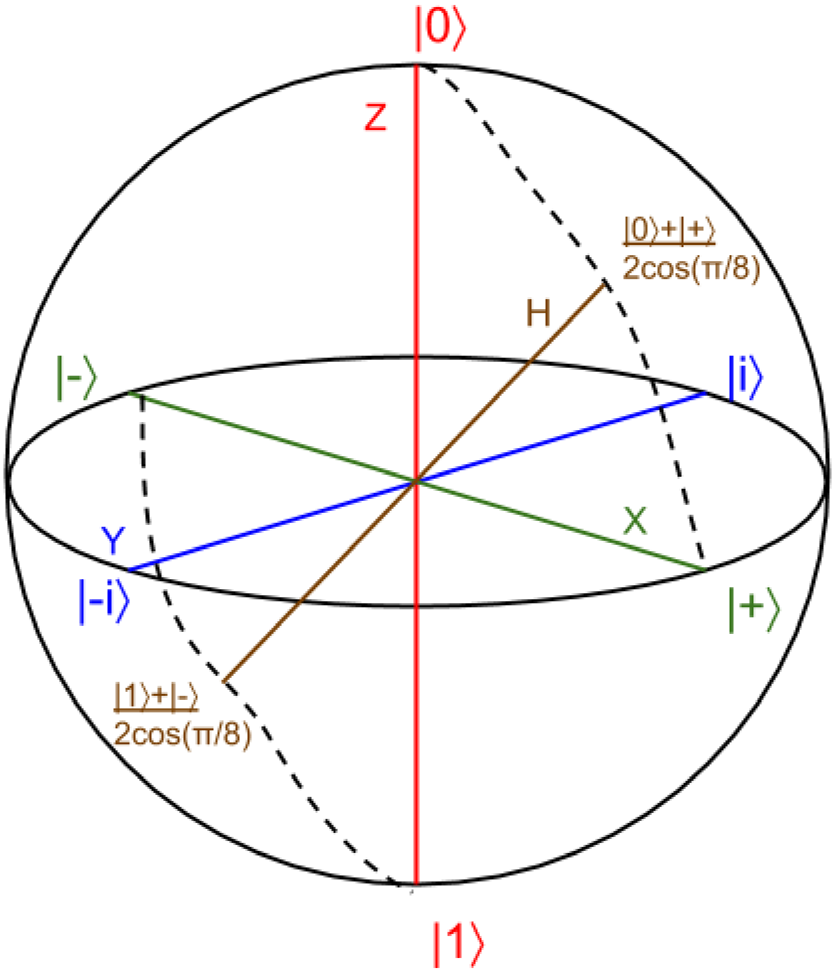

Visualizing the actions of gates on the Bloch sphere:

Applying gates X,Y,Z or H is the same as doing a half turn on their respective axis.

S corresponds to a quarter turn around Z. [in the $|+\rangle$ to $|1\rangle$ direction]

$T^2 = S$, so T corresponds to an eighth turn around Z.

$R_{\frac{\pi}{4}}$ corresponds to a quarter turn (i.e. $\frac{\pi}{4}$) on Y.

The No-Cloning Theorem

We've seen how entanglement seems to lead to "non-local effects," like for the state $\frac{|00\rangle+|11\rangle}{\sqrt{2}}$, where if Alice measures her qubit then she learns the state of Bob's. The reason that Alice isn't communicating faster than light boils down to Bob not being able to tell if his qubit's state is in the $|0\rangle$, $|1\rangle$ basis or the $|+\rangle$,$|-\rangle$ basis.

But what if Bob could make unlimited copies of his qubit? He could figure it out through repeated measurements, and so he'd be able to tell what basis Alice measured in. Faster than light communication!

Learning a classical description of a quantum state, given lots of copies of the state, is called **Quantum State Tomography**,

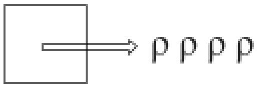

It turns out that we can prove that a procedure to reliably copy an unknown quantum state cannot exist. It's fairly easy to prove, but it's a fundamental fact about quantum mechanics.

In effect, we already saw one proof: namely, cloning would imply superluminal communication, which would violate the No-Communication Theorem that you proved in the homework! But let's see more directly why cloning is impossible.

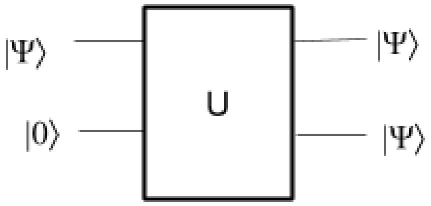

Let's try to clone a single qubit, $|\Psi\rangle = \alpha|0\rangle + \beta|1\rangle$

In our quantum circuit we want to apply some unitary transformation that takes $|\Psi\rangle$ and a $|0\rangle$ ancilla as input, and produces two copies of $|\Psi\rangle$ as output.

Algebraically, a cloner would need to do:

$$\begin{aligned}(\alpha|0\rangle + \beta|1\rangle) \otimes |0\rangle &\to (\alpha|0\rangle + \beta|1\rangle) \otimes (\alpha|0\rangle + \beta|1\rangle)\\ &= \alpha^2|00\rangle + \alpha\beta|01\rangle + \alpha\beta|10\rangle + \beta^2|11\rangle\end{aligned}$$

The cloner would need to look like:

$$\begin{bmatrix}\alpha^2\\ \alpha\beta\\ \alpha\beta\\ \beta^2\end{bmatrix} = \begin{bmatrix} & & \\ & U & \\ & & \\ & & \end{bmatrix}\begin{bmatrix}\alpha\\ \alpha\\ \alpha\\ \alpha\end{bmatrix}$$

The problem: this transformation **isn't linear** so it can't be unitary!
To clarify, a procedure that outputs some $|\Psi\rangle$ can be rerun to get $|\Psi\rangle$ repeatedly. What the No Cloning Theorem says is that if $|\Psi\rangle$ is given to you but is otherwise unknown, then you can't make a copy of it.

Another clarification:
cNOT seems like a copying gate [as it maps $|00\rangle \rightarrow |00\rangle, |10\rangle \rightarrow |11\rangle$]
So why doesn't it violate the No Cloning Theorem?

Because it only copies if the input state is $|0\rangle$ or $|1\rangle$.
Classical information CAN be copied. Just ask Richard Stallman!

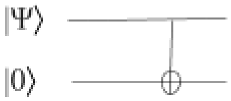

Doing cNOT on 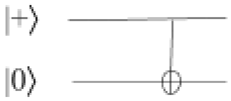 produces the Bell Pair: $\frac{|00\rangle+|11\rangle}{\sqrt{2}}$. Which sort of copies the first qubit in an entangled way, but that's different making a copy of $|+\rangle$.

Having two qubits be $\frac{I}{2}, \frac{I}{2}$ is not the same as $|+\rangle, |+\rangle$.

In general, for any orthonormal basis you can clone the basis vectors, if you know that your input state is one of them.

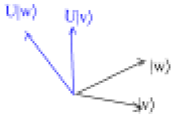

Since the No Cloning Theorem is so important, we'll present another proof of it:

A unitary transformation can be defined as a linear transformation that preserves inner product. Which is to say that the angle between $U|v\rangle$ and $U|w\rangle$ is the same as the one between $|v\rangle$ and $|w\rangle$.

Thus $\langle w|U^{\dagger}U|v\rangle = \langle w|v\rangle$.

What would a cloning map do to this inner product?

Let $|\langle v|w\rangle|^2 = c$
Then $|(\langle v| \otimes \langle v|)(|w\rangle \otimes |w\rangle)|^2 = c^2$

c only ever equals c^2 if the inner product is 0 or 1: so the transformation can only copy if v and w belong to the same orthonormal basis.

There's a fact in classical probability that provides a nice analog to the No-Cloning Theorem.

If we're given the outcome of a coin flip—from a coin that lands heads with some unknown probability —can we simulate a second, independent flip of the same coin, without having access to the coin?

You'd need $\begin{bmatrix} p^2 \\ p(1-p) \\ p(1-p) \\ (1-p)^2 \end{bmatrix} = \begin{bmatrix} & & \\ & S & \\ & & \\ & & \end{bmatrix} \begin{bmatrix} p \\ 1-p \\ 0 \\ 0 \end{bmatrix}$ to be true for some stochastic matrix S.

But once again, this transformation isn't stochastic, because it's not linear.

The No Cloning Theorem has all sorts of applications to science fiction, because you can't make arbitrary copies of a physical system (say for teleporting yourself) if any of the relevant information (say, in your brain) were encoded in quantum states that didn't belong to a known orthogonal basis.

Quantum Money

is an application of the No Cloning Theorem. In some sense it was the first idea in quantum information, and was involved in the birth of the field. The original quantum money scheme was proposed by Wiesner in 1969, though it was only published in the 80s.

Wiesner had left research by then, and had become a manual laborer.

Wiesner realized that the quantum No-Cloning Theorem--though it wasn't yet called that—could be useful to prevent counterfeiting of money. In practice, mints use special ink, watermarks, etc., but all such devices basically just lead to an arms race with the counterfeiters. So Wiesner proposed using qubits to make money that would be physically impossible to counterfeit.

The immediate problem is that a money scheme needs not only *unclonability* but also *verifiability*. How did Wiesner solve this problem?

Wiesner's Scheme

The bank prints quantum bills (we'll assume for simplicity that they're all same denomination). Each bill has:

- A classical serial number $s = \{0,1\}^n$
- A quantum state $|\Psi_{f(s)}\rangle$ (of n qubits)
  - The qubits in this state are unentangled, and each will always be in one of four states:

$$|\Psi_{00}\rangle = |0\rangle \quad |\Psi_{01}\rangle = |1\rangle \quad |\Psi_{10}\rangle = |+\rangle \quad |\Psi_{11}\rangle = |-\rangle$$

The bank maintains a giant database that stores, for each bill in circulation, the classical serial number s, as well as a string $f(s)$ that encodes what the quantum state attached to bill s is supposed to be.

$s_1, f(s_1)$
$s_2, f(s_2)$
$s_3, f(s_3)$

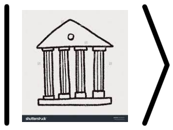

Wiesner's scheme has an important practical problem though: you need to ensure that the qubits in a bill don't lose their state (coherence). With current technology, qubits in a lab decohere in like an hour, tops. Qubits stored in a wallet would decohere much faster!

To verify a bill, you bring it back to the bank. The bank verifies the bill by looking at the serial number, and then measuring each qubit in the bill in the basis in which it was supposed to be prepared. E.g., if the qubit was supposed to be $|0\rangle$ or $|1\rangle$, then measure in the $\{|0\rangle, |1\rangle\}$ basis. For each measurement, check that you get the expected outcome.

Consider a counterfeiter who doesn't know which basis each qubit is supposed to be in, so they guess the bases uniformly at random. They only have a $\left(\frac{1}{2}\right)^n$ chance of making all n guesses correctly.

Of course one could imagine a more sophisticated counterfeiter---but it's possible to prove that, *regardless* of what the counterfeiter does, if they map a single input bill to two output bills, then the output bills will both pass verification with probability at most $\left(\frac{3}{4}\right)^n$.

Wiesner didn't actually prove the security of this scheme at the time he proposed it. Professor Aaronson asked about it on Stack Exchange a few years ago which prompted Molina, Vidick, and Watrous to write a paper that formally proved the scheme's security.