A. Quantum Wires, Circuit Elements, and Interference

At this point it is useful to begin to introduce the quantum circuit notation. The quantum circuit notation is a useful way to denote a set of actions that we apply to our quantum system. Here we'll just talk about one qubit, but in the next section we'll begin to denote more than one qubit. In the quantum circuit diagrams, time runs from left to right (unlike in physics where we often like to make time run from top to bottom). We denote a qubit by a single line, which is often called a quantum wire. A quantum wire really just represents a qubit which is not evolving, i.e. which is acted upon by I. If we initialize our qubit into a particular quantum state, then we usually write the appropriate ket on the left hand side of the appropriate quantum wire:

$$\alpha|0\rangle + \beta|1\rangle \;\text{———} \tag{12}$$

Now if we want to signify that a particular unitary evolution is to be enacted on our qubit, then we put a box with a symbol describing this unitary transform along the quantum wire. Thus for example the prescription, start in the state $|0\rangle$ and apply the transform $H = \begin{bmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \end{bmatrix}$ is denoted by

$$|0\rangle \;\text{—}\boxed{H}\text{—} \tag{13}$$

By the way, this operation, H, is called the Hadamard operation, and will be a good friend of ours for the next many lectures. Oftentimes, we will also denote the output of a quantum circuit as well. For example, the result of the above circuit is denoted by

$$|0\rangle \;\text{—}\boxed{H}\text{—}\; \tfrac{1}{\sqrt{2}}(|0\rangle + |1\rangle) \tag{14}$$

Finally we may want to denote a measurement in the computational basis on our qubit. This is denoted by a meter (or sometimes by an eyeball.) Thus, for instance we might have the circuit

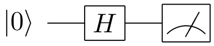

(15)

At this point, one can already introduce some neat little circuits which are very strange from a classical perspective. First consider the following circuit

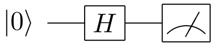

(16)

This circuit takes a system in a fixed configuration, $|0\rangle$, and the measurement after applying H, is 50% $|0\rangle$ and 50% $|1\rangle$. OK, this is not so strange: we've just produced a random bit. Certainly a classical machine can do this. Now consider applying the H twice:

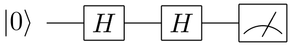

(17)

It is easy to check that $H^2 = I$, so that the state before the measurement meter is just $|0\rangle$. Thus applying H twice yields the configuration $|0\rangle$ with one hundred percent probability. Now this is a bit peculiar: applying the same physical process to our qubit once randomized it, but applying it a second time turned it back into a totally determined configuration. A bit weird. But things get a little stranger. Now suppose that after applying H you apply the Pauli Z operator (remember the Pauli's are unitary):

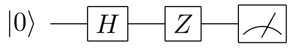

(18)

Now the state right before the measurement meter is $\frac{1}{\sqrt{2}}|0\rangle - \frac{1}{\sqrt{2}}|1\rangle$. So measuring the system results again in 50% $|0\rangle$ and 50% $|1\rangle$. But now apply a Hadamard before performing a measurement, i.e. the following circuit:

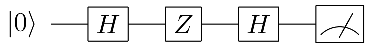

(19)

If you work through the math on this (and you should if you aren't already familiar with quantum circuits) then you will see that right before the meter the system has a description of $|1\rangle$ and so a measurement yields $|1\rangle$ with 100% probability. Now this is kind of peculiar: applying an operation which *didn't* have an observable consequence after applying the Hadamard, after applying a second Hadamard resulted in a totally different configuration. This demonstrates to use that it is just not the magnitude of the amplitudes in a quantum state that matter (i.e. if the qubit is $\alpha|0\rangle + \beta|1\rangle$, then it is not just $|\alpha|$ and $|\beta|$ which matter.) What we say is that the phases of the different configurations matter. And this is what is cool about quantum systems: because this phase, we see that the amplitudes can add in interesting ways that probabilities can't. In this particular example, we see that when we feed $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ into the Hadamard, the amplitudes to go to the $|0\rangle$ configuration constructively add up but the amplitudes to go to the $|1\rangle$ configuration destructively subtract and become zero. By changing the phase of the input to the Hadamard to $\frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$ we can reverse the pathways which are constructive and destructively adding up. This effect, where the amplitude can add in constructive or destructive manners is called interference. It is an expression, in general of the "wave nature" of quantum theory, i.e. just like waves on a pond when they cross each other can add or substract to each other the amplitudes in quantum theory can do a similar thing.

Now this last point brings us to an important and interesting question. If it is only interference which is going to make a quantum computer special then we might be in trouble. Why? Because we can imagine building a machine which uses classical waves, i.e. like the waves on a pond, to construct a computer. Why isn't a classical wave machine just like a quantum wave machine. This is a fun and interesting question that we will eventually go a long ways toward understanding in this course. Suffice it to say that you will find out that classical wave machines do not appear to be as powerful as quantum computers. But its a good question and we should keep it in the back of our minds as we continue on in this course.

II. TWO QUBITS

Having now introduced single qubits, we can start plugging onward and move up to a system with two two configuration systems, i.e. two qubits. The configurations of a two qubit system are the four configurations $|0\rangle \otimes |0\rangle$, $|0\rangle \otimes |1\rangle$, $|1\rangle \otimes |0\rangle$, and $|1\rangle \otimes |1\rangle$. Now it gets really tiresome writing these tensor products, so we often drop the $\otimes$, and express these in the slightly more compact notation $|0\rangle|0\rangle$, $|0\rangle|1\rangle$, $|1\rangle|0\rangle$, and $|1\rangle|1\rangle$. This can be confusing because the tensor product is implicit. Most often we drop the tensor product and combine the configurations together, and write $|00\rangle$, $|01\rangle$, $|10\rangle$, and $|11\rangle$. The general state of a two qubit system is given by the quantum state

$$|\psi\rangle = \alpha_{00}|00\rangle + \alpha_{01}|01\rangle + \alpha_{10}|10\rangle + \alpha_{11}|11\rangle \tag{20}$$

where $\alpha_{ij} \in \mathbb{C}$ and $|\alpha_{00}|^2 + |\alpha_{01}|^2 + |\alpha_{10}|^2 + |\alpha_{11}|^2 = 1$.

What sort of operations can we perform on two qubits? Well any 4 dimensional unitary transform. It is important, however, at this point to begin to understand the issue of locality for our quantum system. In particular, when we have two qubits, we should think that we have two separate physical systems, each with a two level system. If we perform a unitary evolution (by turning on our lasers or such) on only one of these physical sysystems, then the unitary we are implementing is of the form $U \otimes I$, where U is a two dimensional unitary matrix and I is the two dimensional identity matrix representing that we have done *nothing* to the second physical system. Similarly if we act only on the second system, then the unitary we will enact is of the form $I \otimes U$. Finally, we can act with unitaries on both qubits at the same same but with a process that does not couple to two qubits and our unitaries will be of the form $U \otimes V$. It is only when we bring the two qubits together and allow them to interact quantum mechanically that we are able to enact unitaries which cannot be expressed in the form $U \otimes V$. An example of such a unitary, and one which will be of some significance for us is the controlled-NOT operation,

$$C_X =$$

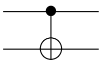

$$= \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix} \tag{21}$$

Here notice we've also begun to introduce the quantum circuit notation for two qubits. These two qubits are denoted by the two quantum wires. Since the controlled-NOT cannot be written as an action on either of these qubits along, we write it as a two qubit gate: a gate with two quantum wires input and two quantum wires output. This should be contrasted with an evolution like $U \otimes V$, which we denote in quantum circuit notation by

$$U \otimes V =$$

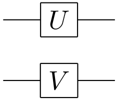

$$= \begin{bmatrix} U_{00}V_{00} & U_{00}V_{01} & U_{01}V_{00} & U_{01}V_{01} \\ U_{00}V_{10} & U_{00}V_{11} & U_{01}V_{10} & U_{01}V_{11} \\ U_{10}V_{00} & U_{10}V_{01} & U_{11}V_{00} & U_{11}V_{01} \\ U_{10}V_{10} & U_{10}V_{11} & U_{11}V_{10} & U_{11}V_{11} \end{bmatrix} \tag{22}$$

Within the class of two qubit states, an important distinction to make is between states that can be expressed as separate single qubit wave functions $|v\rangle \otimes |w\rangle$, where $|v\rangle = \alpha|0\rangle + \beta|1\rangle$ and $|w\rangle = \delta|0\rangle + \gamma|1\rangle$ and those that cannot be expressed like this. The first of these are called separable states and the latter are called entangled states. An example of an entangle state of two qubits is the state $\frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$.

How do we tell if a two qubit state is entangled? Well one way to do this is to just use the four equations that come from

$$(\alpha|0\rangle + \beta|1\rangle) \otimes (\delta|0\rangle + \gamma|1\rangle) = \alpha_{00}|00\rangle + \alpha_{01}|01\rangle + \alpha_{10}|10\rangle + \alpha_{11}|11\rangle \tag{23}$$

and see whether these expressions have a solution (notice that because we can always multiply α and β by a factor and γ and δ by the one over this factor, there will always be multiple solutions, when they exist.) Another way which is nicer is to use the Schmidt decomposition. What is the Schmidt decomposition?

CSE 599d - Quantum Computing
The No-Cloning Theorem, Classical Teleportation and Quantum Teleportation, Superdense Coding

Dave Bacon
Department of Computer Science & Engineering, University of Washington

I. THE NO-CLONING THEOREM

Cloning has been in the news a lot lately. But today we are not going to talk about *that* type of cloning, but instead of a process of duplicating quantum information. That we are not talking about cloning DNA (for our results will not be positive) is rather fortunate for the Raelians!

The no-cloning theorem is one of the earlier results in the study of quantum information. It has an interesting history, some of which is written down in a paper by Asher Peres, "How the No-Cloning Theorem Got Its Name" which is available online at `http://arxiv.org/abs/quant-ph/0205076` and is generally attributed to Wootters, Zurek, and Dieks in 1982. (When I was an undergraduate at Caltech, they had an automated system for requesting copies of articles when you were searching their publication database. I discovered, probably around my junior year, the papers of William Wootters, promptly ordered the automated system to print out *every* paper Wootters had ever written at that time and my life hasn't been the same ever since! Wootters thesis is, in my opinion, one of the most interesting resutls I've ever encountered.) So what is the no-cloning theorem?

Suppose that we have in our lab a qubit in an unknown quantum state $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$. Actually it is perhaps better to say that some external referee knows a description of this quantum state, but we in the lab don't know this quantum state. Now we can construct all sort of machines (unitaries and measurements) which act on this qubit. The question posed in the no-cloning theorem is whether it is possible to design a machine which, for all possible actual states $|\psi\rangle$, is able to take this state $|\psi\rangle$ and create two copies of this state $|\psi\rangle \otimes |\psi\rangle$, hence "cloning" the quantum state.

Let's prove a simple version of the no-cloning theorem. We want to show in this version that there is no unitary operation which can enact the evolution $|\psi\rangle \otimes |0\rangle \rightarrow |\psi\rangle \otimes |\psi\rangle$ for all possible states $|\psi\rangle$. To see this, suppose that there exists such a unitary. Then it must be able to clone $|0\rangle$ and $|1\rangle$: $U|0\rangle \otimes |0\rangle = |0\rangle \otimes |0\rangle$ and $U|1\rangle \otimes |0\rangle = |1\rangle \otimes |1\rangle$ Then if this is true, by the linearity of quantum theory, $U\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) \otimes |0\rangle = \frac{1}{\sqrt{2}}(|0\rangle \otimes |0\rangle + |1\rangle \otimes |1\rangle)$ which is not equal to $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) \otimes \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ which is what we would require if this were a cloning unitary. Thus we have a contradiction: no such unitary can exist.

We've shown the no-cloning theorem for qubits and for just unitary transformations. In fact there is a more general version of this theorem, called the no-broadcasting theorem which deals with an even more general situation. Suffice it to say, in all of these formulations a quantum machine which can *perfectly* clone all quantum states is shown to not be constructible. An interesting question which we won't talk about too much, but which has led to a lot of very nice results is what about *imperfect cloning machines*?

Finally it is interesting to ask how quantum the no-cloning theorem is. What about if we return to our classical probabilistic information processing device. Suppose we have a classical system of two configurations with the probability distribution $[p\ 1-p]^T]$. Is it possible to turn this into two copies of the probability distribution $[p\ 1-p]^T \otimes [p\ 1-p]^T$ for all possible probabilities p? At first you might think: of course we can do this: we just copy the bits! But if we take $[p\ 1-p]^T \otimes [1\ 0]$ and copy the first bit to the second, this produces the probability distribution $[p\ 0\ 0\ 1-p]^T$. This is not the same as the probability distribution $[p\ 1-p]^T \otimes [p\ 1-p]^T$. In fact one can see that if we allow stochastic evolutions, then by the same argument we used in the quantum case, there is no classical machine which can clone the probability distribution.

This latter example is one of the important points I want to make about doing quantum information science. It is important to always challenge ourselves to distinguish the difference between classical and quantum information processing machines. And it is important in these comparisons to not compare "deterministic" classical information processing machines (ones in which the probabilities are always unity) to quantum machines, but to compare probabilistic information processing devices. This doesn't mean that comparing to deterministic machines isn't important: it's just that it is often misleading. This is not to say that the quantum no-cloning theorem isn't somehow different that the classical no-cloning theorem. The quantum no-cloning theorem actually results in results which have no classical equivalent...when we are using the quantum properties of the states.

II. TELEPORTATION

If there could be teleportation, or teleportage, or whatever it is...
- Science fiction writer J. Wyndham, 1951

When a reporter asked Asher if quantum teleportation could teleport the soul as well as the body, Asher answered, characteristically, "No, not the body, just the soul."
- Obituary for Asher Peres, Physics Today, August 2005

Most of us first encounter teleportation when Captain Kirk and Spock decide to make a trip down to the surface of a planet. This form of teleportation is a method to magically transmit people or objects from one location to another. How this works is only known to television producers and Scotty. In this lecture we will talk about a different sort of teleportation. In fact we will talk about two types of teleportation, classical teleportation and quantum teleportation. Calling something teleportation makes it sound really cool and mysterious, but in this class you will find that while quantum teleportation is certainly very interesting, it is not as surprising as you might initially have expected. This is because there is a protocol in the classical world, the world of our probabilistic information processing device, which is very much in the same spirit of teleportation. We call this protocol classical teleportation and actually most of you have actually encountered this protocol before! Once we have talked about classical teleportation, I will use the classical teleportation to "derive" quantum teleportation. This derivation follows along with a very nice article written by David Mermin ("From Classical State-swapping to Quantum Teleportation", Phys. Rev. A 65, 012320 (2002); quant-ph/0105117.)

Suppose that I have in my laboratory a qubit. I don't know the state of this qubit, just that it is in some quantum state (some external referee may know what this quantum state is.) Now you live a few blocks away from me and I want to give my qubit to you such that this qubit is in the same state as it is in my lab. Well one thing I could do is that I could pack up the qubit in a bag (only theorists can do this) and take the qubit over to your house. Then you will have my qubit. But of course this is kind of a pain. For one thing, moving qubits around is not very easy to do (that bag I used is not easily constructible) and further it often happens that in moving the qubit from one location to another the qubit will lose its quantum nature (we'll learn more about this problem when we study quantum error correction.) So one thing I could ask is whether it is possible for me to pick up the phone and call you and tell you some classical information which will alow you to construct a qubit just like mine at your home. Now from the no-cloning theorem, we know that there is no way to take an arbitrary qubit $|\psi\rangle$ and clone it $|\psi\rangle \otimes |\psi\rangle$. Thus there is no way for me to get you a copy of my qubit without, in some way, disturbing my qubit. Now the goal is for me to just send you classical information. If I just perform unitary measurements on my quantum state and some ancilla, this doesn't really do anything to disturb the quantum information: it's just rotated into some new basis. Thus the only way for it to be possible for me to send the qubit using this classical channel is for me to perform measurements which involve my qubit. But now we have a bit of a paradox: such a measurement will reveal some information about the qubit (i.e. some information about the amplitudes α and β.) And this information will need to be transmitted from me to you over the classical phone line. But now if this information is enough to be able to construct the qubit, then it will be possible for multiple people who intercept our message to construct the qubit. But this would violate the no-cloning theorem. Thus if there is going to be some sort of procedure for sending my qubit from me to you with a classical phone line, the classical information transmitted over the phone line better not reveal anything about the actual amplitudes of the qubit.

A. Classical Teleportation

And this line of reasoning now leads us to classical teleportation. Is there a way to send a classical probability distribution over a bit from one party to another without revealing any information about the probability distribution? There is! It is the classical prescription used in a one-time pad for sending information securely over a classical channel. How does this work? Suppose that we have two parties, which we will now refer to as Alice and Bob in the convention invented, I think, by cryptographers, and now universally used by those in quantum computing. Alice and Bob each have one of two bits whose description is given by $[\frac{1}{2}\ 0\ 0\ \frac{1}{2}]^T$. In other words a fair coin is flipped an the value of this coin is distributed to both parties. We will often call this a shared secret random bit. Sometimes I will call it shared randomness (this will make sense when we talk about entanglement and Bell's theorem.) OK, back to our story at hand. Alice and Bob share this secret random bit. Alice has a classical system in the distribution $[p\ 1-p]^T$. What Alice does to send her distribution to Bob is that she performs an exclusive or between her bit and the shared secret random bit (recall that the exclusive or of two bits b_1 and b_2 is $b_1 + b_2$ mod 2. Now to someone who does not know the value of the shared secret random bit, the result of this result will be a bit which appears to be totally random. If this result bit is sent from Alice to Bob, then an eavesdropper, Eve, can learn nothing about the original bit $[p\ 1-p]^T$. Now when Bob gets the bit in his lab, he can use his version of the shared random bit to recover Alice's state $[p\ 1-p]^T$. He does this by performing an exclusive or between the bit he receives from Alice and his copy of the shared classical key. Why does this work? Because exclusive or is a commutative operation and doing it twice with the same value is the same as doing nothing.

Now I have made a further assumption here one which makes classical teleportation quite subtle. In particular I am assuming that when I perform the exclusive or of two bits then the value of these bits is replaced by a single value which is the output of this exclusive or. And this is what makes classical teleportation kind of subtle. Suppose that instead of using an exclusive or which takes two bits as input and produces one output bit, we perform an exclusive or which takes two bits as input and produces two bits as output. One way to do this is to use a controlled-NOT operation (which we introduced as a quantum gate, but which we can port over to the classical world.) This transform acts as $(x, y) \rightarrow (x, x \oplus y)$. We can then use this transform in a classical circuit like this

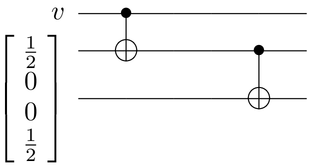
v
$\begin{bmatrix} \frac{1}{2} \\ 0 \\ 0 \\ \frac{1}{2} \end{bmatrix}$

(1)

Here the first wire is a bit which Alice has. The second and third wires are one half of the shared random secret bit. After the first controlled-NOT, the second wire is the "result" bit. At this point in the protocol it is sent from Alice to Bob. Then Bob performs his controlled-NOT and the final state which is the "teleported" system is the third bit. Now why is there a subtly here? Well because now technically the state we have after this procedure is not a state where Bob just as $[p\ 1-p]^T$, uncorrelated with Alice's system. And indeed we find that we have not violated the classical no-cloning theorem but instead Alice and Bob have correlated bits.

What should we take from this? Well that if we can use irreversible gates where information can truly (whatever that means) disappear, then we can teleport a classical probability distribution. However if we use reversible gates then, while the bit Bob has will always behave like Alice's bit, Alice will retain some correlation with this bit so we haven't exactly teleported the classical probability distribution.

B. Quantum Teleportation

Having discussed classical teleportation, we will now use it as inspiration for seeing if we can achieve a similar function in the quantum world. So again to rehash, the setup is that Alice has a qubit and she wishes to transmit this qubit to Bob using only a classical communication. Now the trick to teleporting in the classical bit from Alice to Bob was to start them out with a shared secret random bit. What will the equivalent concept be for our quantum protocol? It will be an entangled quantum state, say in analogy with the shared secret random bit, $\frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$. But we'll get to this in good time. To start our "derivation" of quantum teleportation, let's start with a simple quantum circuit.

Suppose that we want to swap two qubits. There is a unitary transform which does this and is called the SWAP gate (surprising name, eh?) This gate is sometimes denoted by

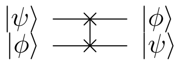

(2)

It is an easy exercise to verify that we can construct the SWAP gate using a series of controlled-NOT gates:

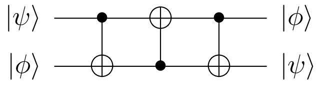

(3)

Now our goal is to SWAP our qubit from Alice to Bob (eventually using classical communication, not direct interaction like we have now) so we will think in our minds as the first quantum wire being Alice's qubit and the second being Bob's qubit. Now we don't really require Bob's qubit to be arbitrary in our attempt to do quantum teleportation. Thus we will act arbitrarily and set $|\phi\rangle = |0\rangle$. When we do this we see that one of the controlled-NOT's does not matter:

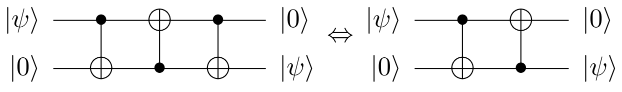

(4)

Now we are faced with a little bit of a conundrum. We know that we need to use entanglement in order to mimic the classical teleportation. But to do this in our current setup, at the very least we need another qubit. So let's just add it. Further we need to interact with this qubit in some manner. So let's do this by the following

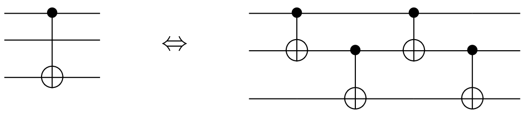

(5)

Putting this into our circuit we now have the three qubit quantum circuit

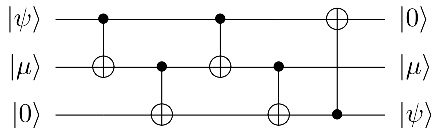

(6)

Well now we have a reasonable number of qubits, but do we have any entanglement? Recall that one way to create a secret shared random bit is to take a random bit and use it as the control in a controlled-NOT. This will produce the classical correlated state $[\frac{1}{2}\ 0\ 0\ \frac{1}{2}]^T$. A similar construction will work for creating an entangled quantum state. In particular to generate the entangled quantum state $\frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ we can do the following

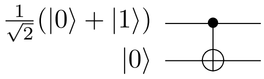

(7)

Now looking at our big circuit it seems most likely that we would like to use that second controlled-NOT to create the entanglement. But how the devil to get rid of that first controlled-NOT?

Well recall that $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ is the +1 eigenvector of X: $X|+\rangle = |+\rangle$. This in turn implies the simple identity that

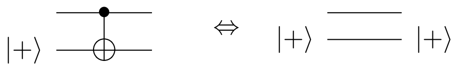

(8)

Thus if, in the big circuit we are constructing, we set $|\mu\rangle = |+\rangle$, we can eliminate the first controlled-NOT.

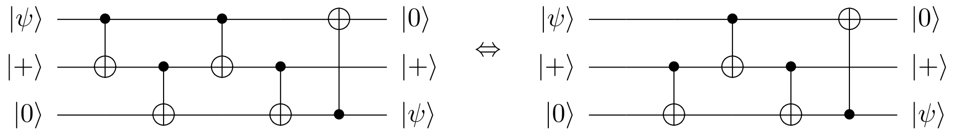

(9)

Notice how we now have, in the first controlled-NOT a way to generate entanglement in our circuit. Now we are getting somewhere! If we think about Alice having one half of this entangled quantum state and Bob having the other half, then we begin to see a circuit which resembles our classical teleportation circuit.

But one really funny thing stands out. And that is the final controlled-NOT. It is a controlled-NOT "going the wrong way," i.e. acting from Bob back to Alice. How in the world are we going to get this to turn around and act in the other direction? To see how to do this, we will need the Hadamard gate. Recall that the Hadamard matrix acts as

$$H = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \tag{10}$$

It has the properties that $H^2 = I$ and the $HXH = Z$ where X and Z are the particular Pauli matrices. Using these properties, it is easy to see that

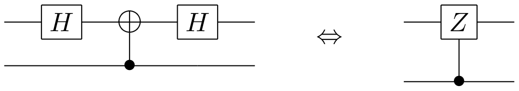

(11)

But the controlled-Z operation, is just the two by two operation

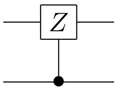

$$= \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix} = \qquad (12)$$

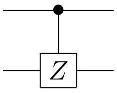

i.e. the controlled-Z, which we often call the controlled-PHASE gate is symmetric under exchange of the two qubits it acts on. Using this fact we can now being to fathom how to "flip" the last controlled-NOT around. First we insert identities I into the circuit

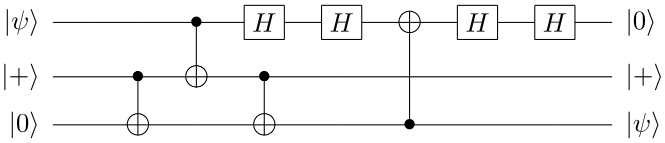

(13)

and then flip that last controlled-NOT to a controlled-PHASE

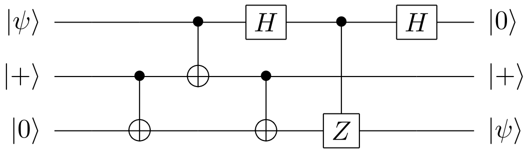

(14)

At this point it is probably just wise to act (backwards) with the last Hadamard to produce

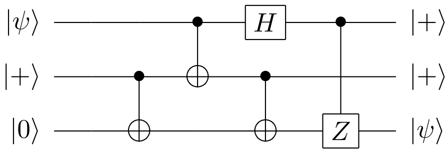

(15)

Well we are almost there. To finish things off, recall that it was our goal to teleport Alice's qubit to Bob using classical information. Thus it makes sense to measure the first two qubits in our circuit in the computational basis,

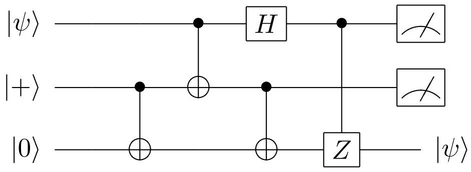

(16)

The result of this measurement will be 50% $|0\rangle$ and 50% $|1\rangle$ for both qubits, i.e these bits will be perfectly random. Next we can use the fact that measurement in the computational basis commutes with a controlled wire of a controlled quantum gate:

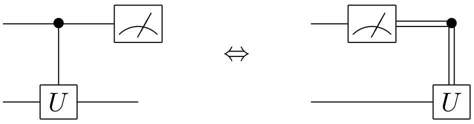

(17)

Here the double wires represent classical bits. I.e. in the second circuit we have made a measurement in the computational basis and then depending on what the output of this measurement is, used this classical bit to apply conditionally the unitary on the target qubit. Using this our circuit begins to look even more tantalizing:

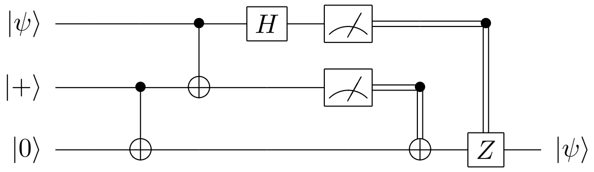

(18)

So what do we have here. The first controlled-NOT is used to create an entangled state of two qubits. One half of this entangled state can be sent to Alice and to Bob. This will function just like the shared random secret bits. Next Alice will take here qubit and her half of the entangled quantum state and perform a unitary transform followed by a measurement. What is the effect of this unitary transform followed by a measurement?

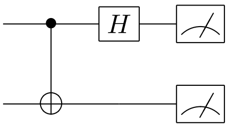

(19)

Well suppose we feed in the state $\frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ into this circuit. Then this state gets transformed to $|00\rangle$ and we will measure $|00\rangle$ with one hundred percent probability. Similarly one can check that the circuit transforms the following inputs to computational basis states

$$\begin{aligned}
|\Phi_+\rangle &= \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle) \;\rightarrow\; |00\rangle \\
|\Psi_+\rangle &= \frac{1}{\sqrt{2}}(|01\rangle + |10\rangle) \;\rightarrow\; |01\rangle \\
|\Phi_-\rangle &= \frac{1}{\sqrt{2}}(|00\rangle - |11\rangle) \;\rightarrow\; |10\rangle \\
|\Psi_-\rangle &= \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle) \;\rightarrow\; |11\rangle
\end{aligned} \tag{20}$$

These for entangled states are called the Bell states (named after John Bell who we will encounter later when we discuss entanglement.) So if any of these states are input into this circuit, then with one hundred percent probability we will get the outcome corresponding to the appropriate computational basis state. Similarly if you have a more general input which is a superposition of these states, then you will obtain, with a probability equal to the square of the overlap between the Bell basis state and the input state. This type of setup, where we perform a unitary transform followed by a measurement in the computational basis is often called "measuring in a different basis" Sometimes the inverse of this unitary is then also applied after the measurement result has been copied to some classical bits. So we now see that this part of the circuit serves to measure in the Bell basis.

Continuing on in our journey through our derived circuit, we see that having measured in the Bell basis Alice's qubit and her shared half of the entangled state, she will get one of four outcomes or two bits of classical information. Alice can then take this information and send it to Bob. Bob then, depending on which of the four possible two bit string's he received apply the appropriate Pauli operator (either I, X, Z or ZX.) Then, quite astoundingly, Bob will have Alice's qubit $|\psi\rangle$! We have "teleported" the qubit from Alice to Bob. By sharing the entangled quantum state $\frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$, Alice was able to transmit her qubit to Bob using only a classical phone line. This is rather remarkable.

There are some interesting things to think about in the quantum teleportation protocol which often throw people off when they first encounter it. The first is asking "where" Alice's qubit went immediately after Alice measures in the Bell basis. At this point in the protocol, there is no single system that has Alice's qubit. It is only after Alice tells Bob what her measurement outcome and Bob applies the appropriate transform that the qubit "reemerges." Now this will really bother you if you think about the qubit as a configuration and not as a description. At no point does the qubit really disappear: it is just that our description is no longer isolated to a single qubit at all times during the protocol.

The other thing to notice is that the measurement outcomes for the Bell basis are always perfectly random and do not depend on the qubit state $|\psi\rangle$. This is good, because this is exactly what we were aiming for. If this had not been true we could not have expected the protocol to work.

So that's it. That's quantum teleportation. Not quite as Spock-y as you might have thought, is it?

III. RULE 3 REDUX

In teleportation we've seen that performing a rotation and then making a measurement is like "measuring" in a different basis. I've defined the measurement rule as always being a measurement in the computational basis, but certainly we can make this rule a bit larger (eventually we'll expand it even further!) In particular given a basis for a quantum syste $\{|\phi_i\rangle\}$, i.e. a set of orthonormal ($\langle\phi_i|\phi_j\rangle = \delta_{i,j}$) vectors which span the Hilbert space of quantum system, we can define measurement in this basis by the following rule. If the quantum system has the quantum state

$|\psi\rangle$, then the probability of getting outcome i (which corresponds to state $|\phi_i\rangle$) is given by $Pr(i) = |\langle\phi_i|\psi\rangle|^2$. If one obtains outcome i for this measurement, then the new description of the quantum system is given by $|\phi_i\rangle$. It is easy to check that if the basis is the computational basis $\{|i\rangle\}$ then this yields are old measurement rule. Further when there is a different basis, the prescription for calculating the probability is to take the overlap of a basis element with the quantum state and take its norm squared.

It is important to emphasize at this point that exactly what measurements we can perform on our quantum system is dependent on the physics of the system. Sometimes the physics of the system will allow us to perform measurements in many different basis. But most of the time, for example in most physical implementations of qubits today, measurement in only one basis is really easy. One can effectively measure in a different basis by performing the appropriate unitary rotation, measuring in a standard basis, and then rotating back. Thus, even though measurement in any basis is not possible, we often are sloppy and talk about measuring in a different basis when it is very clear how to perform these rotations and thus effectively measure in a different basis.

IV. SUPERDENSE CODING

We have seen that sharing an entangled state between Alice and Bob is a resource that these two parties can use to teleport a qubit from Alice to Bob using two bits of classical information. What else might we use an entangled quantum state for? One other very simple application is known as superdense coding. Superdense coding allows us to use an entangled quantum state of two qubits to *double* the classical capacity of a quantum channel. Entanglement can change qubit bit into two bits!

To see how this works, we return to the entangled quantum state $|\Phi_+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$. Now suppose that Alice has one half of this entangled quantum state and Bob has the other half. If Alice applies the single qubit unitary operator X to her qubit, then the full action of this operator on both qubits is $X \otimes I$. We can then check that

$$(X \otimes I)|\Phi_+\rangle = \frac{1}{\sqrt{2}}(|10\rangle + |01\rangle) = |\Psi_+\rangle \tag{21}$$

What is neat about this is that $|\Phi_+\rangle$ and $|\Psi_+\rangle$ are orthogonal. And from our new measurement rule this means that if we can perform a measurement in a basis that contains these two states then we can distinguish between these two states. Note that this measurement will be a measurement on *both* qubits at once. How does this help us in superdense coding? Well it is easy to check that

$$\begin{aligned}(I \otimes I)|\Phi_+\rangle &= |\Phi_+\rangle \\ (X \otimes I)|\Phi_+\rangle &= |\Psi_+\rangle \\ (Z \otimes I)|\Phi_+\rangle &= |\Phi_-\rangle \\ (XZ \otimes I)|\Phi_+\rangle &= |\Psi_-\rangle\end{aligned} \tag{22}$$

Where the states indicated are the Bell basis states defined above. Recall that these states are orthogonal. Notice that we have only acted on Alice's qubit. So this is how we can use an entangled quantum state to double the classical capacity of a quantum channel from Alice to Bob. Alice and Bob start with on half of $|\Phi_+\rangle$. Alice wants to send two bits $a, b \in \{0, 1\}$ to Bob. To do this she applies X^aZ^b to her half of $|\Phi_+\rangle$. She then sends her half of the entangled quantum state to Bob. Now Bob, since he has both qubits of the entangled state can perform a measurement in the Bell basis. When he does this he will obtain one of the four outcomes corresponding to the different Bell basis states. In fact he will obtain only one of these results and these four outcomes will be cover exactly the four different bits a, b that Alice wanted to send to Bob. So by performing this measurement Bob can learn two bits of information. But Alice only sent a single qubit! This is very strange and is called superdense coding.

Now one might worry that this isn't so strange because when we send a qubit from Alice to Bob, why can't we just send an infinite amount of classical information in the α and β amplitudes of our qubit $\alpha|0\rangle + \beta|1\rangle$? Well this has certain tripped up more than its fair share of very smart people (I have seen Nobel prize winners make this mistake!) The reason, philosophically, is that the amplitudes of a qubit are more like probabilities than they are like real properties of our system. Thus you might as well ask why sending a probabilistic bit doesn't allow one to send an infinite amount of classical information in the probability p. More concretely there is a theorem, known as Holevo's theorem which says in a very precise way that a qubit can only be used to send a single bit of classical information (when there is no extra entanglement hanging around!) Now I'm using argument by authority here, which is just disgusting, but it is true and later we may get a chance to examine Holevo's beautiful theorem. But in a smaller sense we can already begin to grasp why we cannot use qubits to send more than one classical bit (without the aid of entanglement!) Suppose that I encode two bits of information into the different states $|0\rangle, |1\rangle, |+\rangle, |-\rangle$. Now if I perform a measurement on this qubit, there is no way for me to learn, with total certainty the value of these two bits? Why? Because states like $|0\rangle$ and $|+\rangle$ are not orthogonal. Thus measurements in some basis will never yield a measurement outcome with certainty (because all such basis elements that have some overlap with $|0\rangle$ must have some overlap with $|+\rangle$.) Thus we can see, intuitively at least, that we can't use a qubit to send more than a classical bit of information as long as we don't have some spare entanglement lying around. And this is what makes superdense coding even more interesting!

Acknowledgments

The diagrams in these notes were typeset using Q-circuit a latex utility written by graduate student extraordinaires Steve Flammia and Bryan Eastin.