CSE 599d - Quantum Computing
Introduction to Quantum Error Correction

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

In the last lecture we saw that open quantum systems could interact with an environment and that this coupling could turn pure states into mixed states. As we have argued before, this is a bad process, because it can lessen or destroy the interference effects which are vital to distinguishing a quantum from a classical computer. This is called the *decoherence problem*. In this lecture we will begin to see how to deal with this problem. Note, however, right off the bat, that there are other problems besides the decoherence problem. For example we still haven't address the fact that we don't have perfect control over our quantum system when attempting to perform unitary gates, preparation of states, or measurement. We will deal with all of these in good time, but for now we will focus on the problem of decoherence.

I. SIMPLE CLASSICAL ERROR CORRECTION

Well when we are stumped about what to do in the quantum world, it is often useful to look to the classical world to see if there is an equivalent problem, and if so, how that problem is dealt with. This leads us from quantum noise to classical noise.

Suppose we have the following classical situation, we have a bit which we send through a channel and with probability $1-p$ nothing happens to our bit, but with probability p the bit is flipped. This channel is called a binary symmetric channel. We can describe it by one of our doubly stochastic matrices if we feel like it

$$M = \begin{bmatrix} 1-p & p \\ p & 1-p \end{bmatrix}. \tag{1}$$

Now if we use this channel once, then the probability that we will receive the wrong bit is p. Is there a way to use this channel in such a way that we can decrease this probability? Yes, and it is rather simple. We just use the channel multiple times and use redundancy. In other words, if we want to send a 0, we use the encoding $0 \rightarrow 000$ and $1 \rightarrow 111$ and send each of these bits through the channel. Now of course there will still be errors on the channel. With probability $(1-p)^3$ no errors occur on the bits. With probability $3(1-p)^2p$ one error occurs on the bits. With probability $3(1-p)p^2$ two error occur on the bits. And with probability p^3 three errors occur on the bits. Now assume that p is small for intuitions sake (we will calculate what small means in a second.) Notice that the three probabilities we have listed above will then be in decreasing order. In particular the probability of no or one error will be greater than there being two or three errors. But if a single error occurs on our bit, we can detect this and correct it. In particular if, on the other end of the channel we decode the states by $\{000, 001, 010, 100\} \rightarrow 0$ and $\{111, 110, 101, 011\} \rightarrow 1$, then in the first two cases we will have correctly transmitted the bit, even though a single error occurred on this bit. We can thus calculate the probability that this procedure, encoding, sending the bits individually through the channel, and decoding fails. It is given by $3(1-p)p^2 + p^3 = 3p^2 - 2p^3$. Now if this is less than p we have decreased the probability of failing to transmit our bit. Indeed this occurs when $3p^2 - 2p^3 \leq p$ or when $p < \frac{1}{2}$. Thus if the probability of our bit is less than $\frac{1}{2}$, then we will have decreased our failing (from p to $3p^2 - 2p^3$.) This is great and is known as a redundancy error correcting code. The classical theory of error correcting codes is devoted to expanding on this basic observations, that redundancy can be used to protect classical information. We won't delve too deeply into classical error correction, although we will find it useful to learn some things. Because we are interested in understanding whether we can port this idea of over to quantum theory!

When we first encounter classical error correction and think about porting it over to the quantum world, there are some interesting reasons to believe that it will be impossible to make this transition. We list these here, since they are rather interesting (although claims that these were major blockades to discovering quantum error correcting are probably a bit exaggerated, at least according to the people I've talked to who were working in quantum computing when this happened.)

No Cloning We've seen that the no-cloning theorem states that there is no machine which perform $|\psi\rangle \rightarrow |\psi\rangle \otimes |\psi\rangle$. Thus a naive attempt to simply clone quantum information in the same way that we copy information in a redundancy code fails.

Measurement When we measure a quantum system, our description of the quantum system changes. Another way this is stated is that measurement disturbs the state of a quantum system. In error correction, we read out classical information in order to correctly recover our classical information. How do we perform measurements on quantum systems that don't destroy the quantum information we are trying to protect.

Quantum Noise As we noted in the last lecture, quantum noise has a continuous set of parameters to describe it. So we might think that this will cause a problem, since in classical theory we could interpret the noise as probabilities of deterministic evolutions occurring, but in quantum theory we don't have such an interpretation (at least not yet.) Of course this feels like a little bit of a red herring, since, classical noise also has continuous parameters (say the probabilities of the erring procedures) to describe it. It's just that those parameters are probabilities which makes us feel safer in the classical case, although we have learned that distinctions like this usually fade away when we examine them closer.

For the reasons we might expect that an equivalent to classical error correction in the quantum world does not exist. One of the surprising discoveries of the mid-nineties was that this is not true!

II. REVERSIBLE SIMPLE CLASSICAL ERROR CORRECTION

So where to begin. Well one place to begin is to try to understand how to perform the classical error correction we described above using classical reversible circuits. So the first part of our procedure is to encode our classical bit. Suppose that we represent our three bits by three wires. Then it is easy to check than an encoding procedure for taking a bit in the first wire to the encoded 000 and 111 configurations is

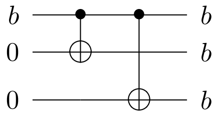

(2)

Next we send each of these bits through the bit flip channel independent of each other. We will denote this by the gate M on these bits

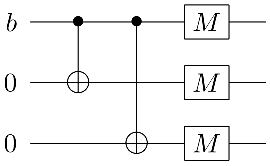

(3)

Now we need to describe the procedure for diagnosing and error and fixing this error. Consider what the two controlled-NOTs do in our encoding circuit. They take 000 → 000, 001 → 001, 010 → 010, 011 → 011,100 → 111, 101 → 110, 110 → 101, and 111 → 100. Notice that except in the case where the last two bits are 11, after this procedure, the first bit has been restored to it's proper value, given that only zero or one bit flip has occurred on our three bits. And when the last two bits are 11, then we need to flip the first bit to perform the proper correction. This implies that the decoding and fixing procedure can be done by the two controlled-NOTs, followed by a Toffoli gate. In other words

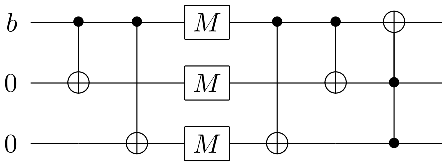

(4)

It is easy to check that if only one or no error occur where the M gates are, then the output of the first bit will be b (and in the other cases the bit will be flipped.) This is exactly the procedure we described in the previous section and we've done it using completely reversible circuit elements (except for the Ms, of course.)

III. WHEN IN ROME DO AS THE CLASSICAL CODERS WOULD DO

Now that we have a classical reversible circuit for our simple error correcting procedure, we can see what happens when we use this circuit on quantum information instead of classical information. One thing we need to do is to chose the appropriate noise channel for our code (we'll come back to more general noise eventually, don't stress about it, yet!) A natural choice is the bit flip channel we described in the last lecture which had the Kraus operator

$$A_0 = \sqrt{1-p}I \quad A_1 = \sqrt{p}X \tag{5}$$

The circuit we want to evaluate is now

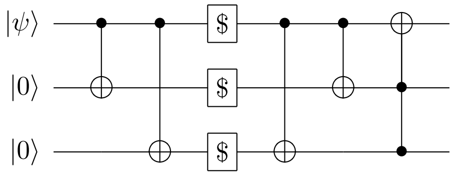

(6)

where now $|\psi\rangle$ is now an arbitrary quantum state $\alpha|0\rangle + \beta|1\rangle$ which we wish to protect. So what does this circuit do to our quantum data? Well after the first two controlled-NOTs, we can see that the state is given by

$$\alpha|000\rangle + \beta|111\rangle \tag{7}$$

Notice that we have done something like redundancy here: but we haven't copied the state, we've just "copied" it in the computational basis.

Next what happens? Well recall that we can interpret \$ as a bit flip error X happening with probability p and nothing happening with probability $1-p$. It is useful to use this interpretation to describe what will happen next. In particular it is useful to use this interpretation to say that with probability $(1-p)^3$ no error occurred on all three qubits, with probability $(1-p)^2p$ a single error occurred on the first qubit, etc. So what happens if no error occur on our system (the Kraus operator $I \otimes I \otimes I$ happens on our system.) Then we just need to run those two controlled-NOTs on our state

$$(C_X)_{13}(C_X)_{12}(\alpha|000\rangle + \beta|111\rangle) = \alpha|000\rangle + \beta|100\rangle = (\alpha|0\rangle + \beta|1\rangle)|00\rangle \tag{8}$$

The Toffoli then does nothing to this state and we see that our quantum information has survived. But this isn't too surprising since no error occurred. What about when a single bit flip error occurs? Let's say an error occurs on the second qubit. Then our state is $\alpha|010\rangle + \beta|101\rangle$. The effect of the controlled-NOTs is then

$$(C_X)_{13}(C_X)_{12}(\alpha|010\rangle + \beta|101\rangle) = \alpha|010\rangle + \beta|110\rangle = (\alpha|0\rangle + \beta|1\rangle)|10\rangle \tag{9}$$

Again, the Toffoli will do nothing to this state. And we see that our quantum information has survived its encounter with the bit flip error! One can go through the other cases of a single bit flip error. In the case where the bit flip error is on the first qubit, the Toffoli is essential in correcting the error, but in the case where it is on the third qubit, then the Toffoli does nothing. But in all three cases the quantum information is restored!

One can go through and check what happens for the cases where two or three bit flip errors occur. What ones finds out that is in these cases the resulting state in the first qubit is $\beta|0\rangle + \alpha|1\rangle$. Thus if we look at the effect of this full circuit, it will performs the evolution

$$\begin{aligned} \rho \otimes |0\rangle\langle 0| \otimes |0\rangle\langle 0| \;\rightarrow\; & (1-p)^3\rho \otimes |0\rangle\langle 0| \otimes |0\rangle\langle 0| + (1-p)^2p\rho \otimes |0\rangle\langle 0| \otimes |1\rangle\langle 1| + (1-p)^2p\rho \otimes |1\rangle\langle 1| \otimes |0\rangle\langle 0| \\ + \; & (1-p)^2p\rho \otimes |1\rangle\langle 1| \otimes |1\rangle\langle 1| + (1-p)p^2X\rho X \otimes |0\rangle\langle 0| \otimes |1\rangle\langle 1| + (1-p)p^2X\rho X \otimes |1\rangle\langle 1| \otimes |0\rangle\langle 0| \\ + \; & (1-p)p^2X\rho X \otimes |1\rangle\langle 1| \otimes |0\rangle\langle 0| + (1-p)p^2X\rho X \otimes |1\rangle\langle 1| \otimes |1\rangle\langle 1| \end{aligned} \tag{10}$$

Tracing over the second and third qubit, this ammounts to the evolution

$$\rho \rightarrow [(1-p)^3 + 3p(1-p)^2]\rho + [3p^2(1-p) + p^3]X\rho X \tag{11}$$

If we compare this with the evolution which would have occurred given no encoding,

$$\rho \rightarrow (1-p)\rho + pX\rho X \tag{12}$$

we see that if $p < \frac{1}{2}$, then our encoding acts to preserve state better than if we had not encoded the state.

Actually how do we know that we have preserved the state better? What measure should we use to deduce this and why would this be a good measure? In particular we might note that quantum error will effect different states in different ways. Thus we should probably worry about the worst case for an input state here. We'll take a little break here to discuss briefly this question.

A. Fidelity

The fidelity between two density matrices ρ and σ is defined as

$$F(\rho,\sigma)=\mathrm{Tr}\left[\sqrt{\rho^{\frac{1}{2}}\sigma\rho^{\frac{1}{2}}}\right] \tag{13}$$

We will be mostly concerned with the fidelity between a density matrix ρ and the case where σ is a pure state $\sigma=|\psi\rangle\langle\psi|$,

$$F(|\psi\rangle,\rho)=\mathrm{Tr}\left[\sqrt{\rho^{\frac{1}{2}}|\psi\rangle\langle\psi|\rho^{\frac{1}{2}}}\right]=\mathrm{Tr}\left[\sqrt{|\psi\rangle\langle\psi|\rho|\psi\rangle\langle\psi|}\right]=\sqrt{\langle\psi|\rho|\psi\rangle} \tag{14}$$

The fidelity is a measure of how close two states are. It is equal to 1 iff $\rho=\sigma$. More importantly, the fidelity can serve as an upper and lower bound on the trace distance, $D(\rho,\sigma)$,

$$1-F(\rho,\sigma)\leq D(\rho,\sigma)\leq\sqrt{1-F(\rho,\sigma)} \tag{15}$$

Well there I've gone again, talking about the trace distance without defining what it is.

The trace distance between two density matrices is

$$D(\rho,\sigma)=\frac{1}{2}\mathrm{Tr}\left[|\rho-\sigma|\right] \tag{16}$$

where $|M|=\sqrt{M^\dagger M}$. The trace distance is a beautiful metric on density matrices, because it is related directly to how different a POVM on these states can be. In particular, assume that we perform a POVM with element E_i. Then on the two state ρ and σ, this results in measurement probabilities $p_i=Pr(i|\rho)=\mathrm{Tr}[\rho E_i]$ and $q_i=Pr(i|\sigma)=\mathrm{Tr}[\sigma E_i]$. The Kolmogorov distance between the two probability distributions produced by these measurements is

$$D(p,q)=\frac{1}{2}\sum_i|p_i-q_i| \tag{17}$$

Let's show that $D(\rho,\sigma)$ is equal to the maximum over all possible POVMs of the Kolmogorov between the resulting probability distributions. Note that

$$D(p,q)=\frac{1}{2}\left|\sum_i\mathrm{Tr}[E_i(\rho-\sigma)]\right| \tag{18}$$

Now decompose $\rho-\sigma$, which is hermitian into a positive hermitian part, P and a negative hermitian part N: $\rho-\sigma=P-N$ and the support of P and N is orthogonal. Then

$$D(p,q)=\frac{1}{2}\left|\sum_i\mathrm{Tr}[E_i(P-N)]\right|\leq\frac{1}{2}\sum_i\mathrm{Tr}[E_i(P+N)]=\frac{1}{2}\sum_i\mathrm{Tr}[E_i|\rho-\sigma|]=D(\rho,\sigma) \tag{19}$$

Thus $D(p,q)\leq D(\rho,\sigma)$. But now choose a POVM with elements which project along the eigenvectors of P and N. Then

$$D(p,q)=\frac{1}{2}\left|\sum_i\mathrm{Tr}[E_i(\rho-\sigma)]\right|=\frac{1}{2}\left|\sum_i\mathrm{Tr}[E_i(P-N)]\right|=\frac{1}{2}\left|\sum_i\mathrm{Tr}[(P-N)]\right|=D(\rho,\sigma) \tag{20}$$

Thus $D(p,q)$ is always less than $D(\rho,\sigma)$ and in fact $D(p,q)$ can attain $D(\rho,\sigma)$. Thus we see that $D(\rho,\sigma)$ is equal to the maximum over all POVMS of the Kolmogorov distance of the probabilities resulting from this POVM.

Back the fidelity. We've seen that the fidelity can be used to bound the trace distance which is related to how bad our probabilities will be for the worst POVM resulting from measuring these states. Thus the fidelity is a nice quantity to use to measure how close two states are. Now one question is why not use the trace distance. One reason is that the trace distance is often hard to calculate. The other reason is that the fidelity between a density matrix and a pure state is rather easy to calculate.

B. Back To The Regularly Scheduled Program

So what happens to the fidelity of the two cases we had above, one in which no error correction was performed and one in which error correction was performed. In the first case the fidelity, assuming we start in a pure state is

$$F_1 = [\langle\psi| [(1-p)|\psi\rangle\langle\psi| + pX|\psi\rangle\langle\psi|X] |\psi\rangle]^{\frac{1}{2}} = \left[(1-p) + |\langle\psi|X|\psi\rangle|^2\right]^{\frac{1}{2}} \tag{21}$$

This is minimized (remember we want high fidelity) when the $\langle\psi|X|\psi\rangle = 0$ and is equal to

$$F_1 \geq \sqrt{1-p} \tag{22}$$

Similarly if we perform error correction we obtain

$$F_3 = \left[\langle\psi| \left[((1-p)^3 + 3p(1-p)^2)|\psi\rangle\langle\psi| + (3p^2(1-p) + p^3)X|\psi\rangle\langle\psi|X\right] |\psi\rangle\right]^{\frac{1}{2}} = \left[(1-p) + |\langle\psi|X|\psi\rangle|^2\right]^{\frac{1}{2}} \tag{23}$$

which is again bounded by

$$F_3 \geq \sqrt{(1-p)^3 + 3p(1-p)^2} \tag{24}$$

The fidelity is greater using error correction, then, when $p < \frac{1}{2}$. So our naive analysis was not so much the dunce after all.

IV. LESSONS

What lessons should we draw from our first success in quantum error correction? One observation is that instead of *copying* the quantum information we *encoded* the quantum information into a *subspace*. In particular, we have encoded into the subspace spanned by $\{|000\rangle, |111\rangle\}$, i.e. we have encoded our quantum information as $\alpha|000\rangle + \beta|111\rangle$. This is our way of getting around the no-cloning theorem.

The second problem we brought up was measurement. Somehow we have made a measurement such that we could fix our quantum data (if needed) using the Tofolli. Let's examine what happens to our subspace basis elements, $|000\rangle$ and $|111\rangle$ under the errors which we could correct. Notice that they enact the evolution

$$\begin{array}{ccc} \begin{array}{c}|000\rangle\\|111\rangle\end{array} & \overbrace{\longrightarrow}^{I\otimes I\otimes I} & \begin{array}{c}|000\rangle\\|111\rangle\end{array} \\ \begin{array}{c}|000\rangle\\|111\rangle\end{array} & \overbrace{\longrightarrow}^{X\otimes I\otimes I} & \begin{array}{c}|100\rangle\\|011\rangle\end{array} \\ \begin{array}{c}|000\rangle\\|111\rangle\end{array} & \overbrace{\longrightarrow}^{I\otimes X\otimes I} & \begin{array}{c}|010\rangle\\|101\rangle\end{array} \\ \begin{array}{c}|000\rangle\\|111\rangle\end{array} & \overbrace{\longrightarrow}^{I\otimes I\otimes X} & \begin{array}{c}|001\rangle\\|110\rangle\end{array} \end{array} \tag{25}$$

Now think about what this is doing. These error processes are mapping the subspace where we encoded the information into different *orthogonal* subspaces for each of the different errors. Further when this map is performed, the orthogonality between the basis elements is not changed (i.e. $|000\rangle$ and $|111\rangle$ are orthogonal and after the error occurs, they are mapped to states which remain orthogonal.) Now this second fact is nice, because it means that the quantum information hasn't been *distorted* in an irreversible fashion. And the first fact is nice because, if we can measure which subspace our error has taken us to, then we will be able to fix the error by applying the appropriate operation to reverse the error operation. In particular consider the operators $S_1 = Z \otimes Z \otimes I$ and $S_2 = Z \otimes I \otimes Z$. These operators square to identity and so have eigenvalues $+1$ and -1. In fact, we can see that these eigenvalues don't distinguish between states within a subspace, but do distinguish which of the four subspaces our state is in. That is to say, for example that $|000\rangle$ and $|111\rangle$ are have eignevalues $+1$ for both S_1 and S_2. Further, the subspace which occurs if a single bit flip occurs on our first qubit, $|100\rangle$ and $|011\rangle$ have eigenvalues -1 for S_1 and -1 for S_2. We can similarly calculate the other cases:

	S_1	S_2	Error
$\{\|000\rangle, \|111\rangle\}$	$+1$	$+1$	$I \otimes I \otimes I$
$\{\|100\rangle, \|011\rangle\}$	-1	-1	$X \otimes I \otimes I$
$\{\|010\rangle, \|101\rangle\}$	-1	$+1$	$I \otimes X \otimes I$
$\{\|001\rangle, \|110\rangle\}$	$+1$	-1	$I \otimes I \otimes X$

Thus we see that if we could perform a measurement which projects onto the $+1$ and -1 eigenstates of S_1 and S_2, then we could use the results of this measurement to diagnose which subspace the error has taken us to and apply the appropriate X operator to recover the original subspace. So is it possible to measure S_1 and S_2? Well we've already done it but in a destructive way in our circuit.

Consider the following circuit

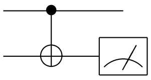

(26)

What does this circuit do? Well if the input to this circuit is $\alpha|00\rangle + \beta|11\rangle$, then the measurement outcome will be $|0\rangle$ and if the circuit is $\alpha|01\rangle + \beta|10\rangle$, then the measurement outcome is $|1\rangle$. Associating $|0\rangle$ with $+1$ and $|1\rangle$ with -1, we thus see that this is equivalent to measuring the eigenvalue of the operator $Z \otimes Z$. Notice, however, that this is a destructive measurement, i.e. it doesn't leave the subspace intact after the measurement. In the circuit we have constructed, then, the controlled-NOTs after the error have enacted things which, if we had measured the qubits, would have been the eigenvalues of S_1 and S_2. This is enough to diagnose which error as occurred. Since this also does decoding of our encoded quantum information, only in the case where the error occurred on the first qubit do we need to do anything: and this is the case where the measurement outcomes are both $|1\rangle$ and so we use the Toffoli to correct this error. This suggest that a different way to implement this error correcting circuit and that is to measure the second and third qubits. Since measurements commute through control gates turning them into classical control operations, we could thus have performed the following circuit

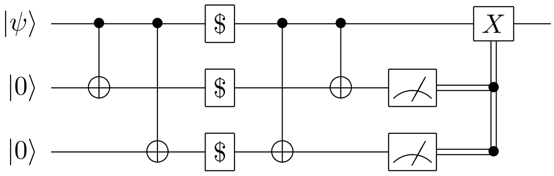

(27)

What do we learn from the above analysis? We learn that quantum error correction avoids the fact that measuring disturbs a quantum system by performing measurements which project onto subspaces. These measurements do not disturb the information encoded into the subspaces since the measurement yield degenerate values outcomes for any state in this subspace. This technique, of performing measurements which do not fully project onto a basis, is essential to being able to perform quantum error correction.

V. DEALING WITH PHASE FLIPS

Now the third problem we brought up was the fact that quantum errors for a continuous set. For now, however, lets just move on to a different error model. In particular lets consider instead of a bit flip model, a phase flip model. In this model, the Kraus operators are given by

$$A_0 = \sqrt{1-p} I \quad A_1 = \sqrt{p} Z \tag{28}$$

Now the effect of Z on a quantum state is to change the phase between $|0\rangle$ and $|1\rangle$. So how are we going to correct this error? It changes a phase, not an amplitude! Well we use the fact that phase changes are amplitude changes in a different basis. In fact, we already know the basis change to perform. The Hadamard basis. Recall that $HZH = X$. This suggest that just prior to sending our information through the quantum channel and just after receiving the quantum information we should apply Hadamard gates. This suggests the circuit

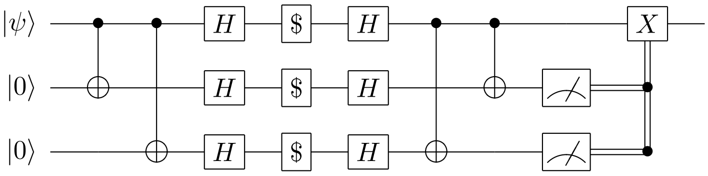

(29)

Now will this work? Certainly it will work, this is just a basis change and the fidelities will be identical to the bit flip analysis. What does this circuit do? Well instead of encoding into the subspace spanned by $|000\rangle$ and $|111\rangle$ this code encodes into a subspace spanned by $|+++\rangle$ and $|---\rangle$, where $|\pm\rangle = \frac{1}{\sqrt{2}}(|0\rangle \pm |1\rangle)$. Thus we see that by expanding our notion of encoding beyond encoding into something simple like the repeated computational basis states, we can deal with a totally different type of error, one that doesn't really have a classical analogy in the computational basis.

But we are just putting off the question of what happens when we have arbitrary errors. But notice something here. We said in a previous lecture that the bit flip error model was equivalent to the phase damping model, and in that case our Kraus operators were

$$B_0 = \begin{bmatrix} \sqrt{1-q} & 0 \\ 0 & \sqrt{1-q} \end{bmatrix} \quad B_1 = \begin{bmatrix} \sqrt{q} & 0 \\ 0 & 0 \end{bmatrix} \quad B_2 = \begin{bmatrix} 0 & 0 \\ 0 & \sqrt{q} \end{bmatrix} \tag{30}$$

When $\frac{p}{2} = q$, these were the same superoperator. We can express the above Kraus operators as

$$B_0 = \sqrt{1-q}I, \quad B_1 = \frac{\sqrt{q}}{2}(I+Z), \quad B_2 = \frac{\sqrt{q}}{2}(I-Z) \tag{31}$$

Now surely, since this is the same superoperator, our analysis of the error correcting procedure will be identical and we will be able to increase the probability of successfully correcting this model given $q \leq \frac{1}{4}$. But the B_i operators are sums of I and Z errors. Thus a code designed to correct single Z errors seems to be working on superoperators which have Kraus operator s which are sums of identity and Z errors. Why is this so? Well we have some encoded information $|\psi\rangle$. Then a Kraus operator which is the sum of terms which we can correct occurs (plus terms we can't correct.) Then we perform the measurement to distinguish which subspace we our error has taken us to. But at this point, the Kraus operators which are sum of error gets projected onto one of the error subspaces. Thus in effect the fact that the error is a sum of errors gets erased when we do this projection. We will return to this fact later, but this is an important point. While quantum errors may be a continuous set, the error correcting procedure can, in effect, digitize the errors and deal with them as if they formed a discrete set.

VI. THE SHOR CODE

So far we have dealt with two models of errors, bit flip errors and phase flip errors. We have also seen that if a code corrects an error then it will be able to correct Kraus operator errors which have a sum of these errors. Thus we expect that if we can design a quantum error correcting code which can correct a single X, Y, or Z error then this can correct an arbitrary error on a single qubit. Indeed Peter Shor designed just such a code (shortly after discovering his factoring algorithm, genius!) How does this code work? Well we've already seen that if we encode into the subspace spanned by $|000\rangle$ and $|111\rangle$, then we can correct a single bit flip error. What do single phase flip errors do to this code? Well notice that $Z \otimes I \otimes I|000\rangle = I \otimes Z \otimes I|000\rangle = I \otimes I \otimes Z|000\rangle = |000\rangle$ but $Z \otimes I \otimes I|111\rangle = I \otimes Z \otimes I|111\rangle = I \otimes I \otimes Z|111\rangle = -|111\rangle$. Thus we see that, unlike bit flip errors, single phase flip errors on this code act to *distort* the information encoded into the subspace. But also notice that these single phase flip errors act like phase flip errors on the encoded basis $|000\rangle$, $|111\rangle$. But we've seen how to deal with phase flip errors. Suppose we define the states

$$\begin{aligned} |p\rangle &= \frac{1}{\sqrt{2}}(|000\rangle + |111\rangle) \\ |m\rangle &= \frac{1}{\sqrt{2}}(|000\rangle - |111\rangle) \end{aligned} \tag{32}$$

Then a single phase flip error on these two states sends $|p\rangle$ to $|m\rangle$ and vice versa, i.e. it looks like a bit flip on these qubits. We can thus deal with these single phase flip errors by using a bit flip code. In particular define the two nine qubit states

$$\begin{aligned} |0_L\rangle &= |p\rangle \otimes |p\rangle \otimes |p\rangle = |ppp\rangle = \frac{1}{2\sqrt{2}}(|000\rangle + |111\rangle) \otimes (|000\rangle + |111\rangle) \otimes (|000\rangle + |111\rangle) \\ |1_L\rangle &= |m\rangle \otimes |m\rangle \otimes |m\rangle = |mmm\rangle = \frac{1}{2\sqrt{2}}(|000\rangle - |111\rangle) \otimes (|000\rangle - |111\rangle) \otimes (|000\rangle - |111\rangle) \end{aligned} \tag{33}$$

Suppose we encode a qubit of quantum information into these the subspace spanned by these two states. Now single phase flip errors can be fixed by diagnosing what subspace the $|ppp\rangle$ and $|mmm\rangle$ subspace has been sent to. Further single bit flip errors can be dealt with from within each $|p\rangle$ and $|m\rangle$ state: these states are encoded states in our original bit flip code. Putting this together, we see that we should be able to use this error correcting code to correct bit flips and phase flips. Actually it can do more and indeed can handle a single Y error as well

To see this lets construct the circuit for encoding into this code and then performing the decoding and correction.

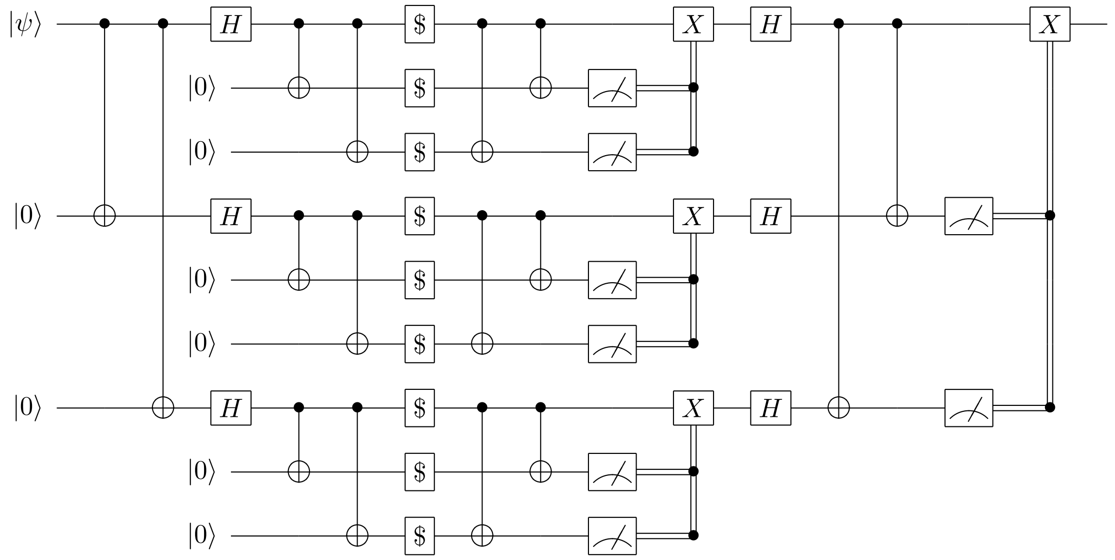

(34)

Now, first of all, isn't this beautiful! Now notice the three blocks of three in this code. Also notice how these blocks, when there is either no error or a single X error in the superoperators (of course in general our superoperators will each contain error terms, but we will simply talk about the case where there is one such error since this is, to first order, the most important error, assuming that the errors are "weak enough." Yeah, all rather vague right now, but we need to make progress without getting bogged down in the details...yet!) produces a channel on the outer three wires (just after the H and before the H) which is identity (because these are designed to correct just that error. But what about if a single Z error occurs. As we noted above, this means that on the encoded quantum information, this acts like a Z error on the $|p\rangle$, $|m\rangle$ states. Since only a Z error occurs, we are not taken out of the $|p\rangle$ and $|m\rangle$ subspace by this error, and so the effect of the error correction in an inner block will be to produce a state which has an single qubit Z error on the outer wire. Now we see how the code corrects this single Z error: this is just the phase flip error correcting code!

But what happens if, say a single $Y = iXZ$ error occurs? First notice that the global phase i doesn't really matter. So we can assume that the error is XZ. Now the Z error acting on the encoded state, acts within the encoded space as a Z error. Thus imagine performing this error, and then running the bit flip code. The bit flip code will correct the X error, but the end result will be that a Z error has occurred on the encoded information. But then the Z error will be corrected by the outer code. Thus we see that a single Y error will be corrected by this code.

So what have we in Shor's code? We have a code which can correct any single qubit error from the set $\{X, Y, Z\}$. But as we have argued above, this means that any single qubit error which is a sum of these errors will also be corrected (we will make this argument rigorous later.) So if a single arbitrary error occurs on our qubit, Shor's code will correct it. We call this a quantum error correcting code which can correct a single qubit error.

So what have we learned? We've learned that is possible to circumvent the objections we raised early: cloning, measurement, and continuous sets of errors, and to enact error correction on quantum information. What does this mean for practical implementations of a quantum computer? Well it means that more sophisticated versions of quantum error correction might be able to make a very robust computer out of many noisy components. There are many issues which we need to discuss in this theory, and the full story of how to build such a quantum computer is the subject of fault-tolerant quantum computation. A main aim in the next lectures will be to obtain an understanding of the theory of quantum error correction, building up to the theory of fault-tolerant quantum computation.

Acknowledgments

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