CS 599 P1: Introduction to Quantum Computation Boston University, Fall 2025
Instructor: Alexander Poremba

LECTURE # 20: INTRODUCTION TO QUANTUM ERROR CORRECTION

In this lecture, we give an introduction to quantum error correction, a mechanism for encoding *logical* information into *physical* information in the form of *codewords*; these will help protect fragile quantum information even in the presence of noise. Quantum error correction is a highly active research area today, and offers a pathway to *quantum fault-tolerance*.

Before we begin defining what a quantum error-correcting code is, we will first begin with a basic introduction to classical error-correcting codes.

1 Classical Error Correction

The origins of mathematical coding theory can be traced back all the way to the 1950s when mathematicians and engineers laid the foundations of *information theory*. The key problems emerged:

- Communication over a noisy channel (say, in the context of satellite communication)
- Noisy data storage (say, on a piece of noisy computing hardware)

A popular and idealized noise model is the so-called *binary symmetric channel*: Suppose that $p \in (0,1)$ is a noise parameter. Then, we can visualize the error transition probabilties as follows:

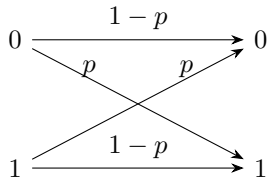

In other words, the channel performs a bit-flip with probability p and otherwise, with probability $1-p$, it leaves the information intact. How can we protect against such noise? The key idea behind an error-correcting code is to add *redundancy*, i.e., to embed the *logical information* into many more *physical bits*. A simple and yet very powerful example of an error-correcting code is the so-called *repetition code*.

Example: Repetition code

A $1 \mapsto 3$ repetition code is a natural mechanism for adding redundancy: simply repeat each logical bit many times in sequence:

$$0 \quad \mapsto \quad \bar{0} = 000 \qquad\qquad 1 \quad \mapsto \quad \bar{1} = 111.$$

Here, we use $\bar{0}$ and $\bar{1}$ to denote the encoding of the logical bits 0 and 1. For example, we can visualize the workflow of a $1 \mapsto 3$ repetition code as follows:

$$\begin{aligned} 0 &\overset{\text{encode}}{\mapsto} 000 \overset{\text{noise}}{\mapsto} 010 \overset{\text{decode}}{\mapsto} 000 \\ 1 &\overset{\text{encode}}{\mapsto} 111 \overset{\text{noise}}{\mapsto} 101 \overset{\text{decode}}{\mapsto} 111. \end{aligned}$$

To decode a noisy codeword, we simply output the most frequent bit. If the error probability p is small, we are unlikely to observe two bit-flip errors at once (since $p^2 \ll p < 1$), and hence our decoding procedure is likely to succeed.

While early computing devices (such as those based on vacuum tubes) suffered greatly from noise, today's digital computers are remarkably resilient against errors. The typical error rate for today's computers is around $p \approx 10^{-16}$. This means that error correcting codes for fault-tolerant computation are often no longer needed; occasionally, they show up in other mediums for data storage, such as compact discs (CDs).

Coding theory today is a rich area of mathematics and computer science which has had immense applications ranging from complexity theory (e.g. the hardness of approximation), to combinatorics (e.g. packing problems) and even to cryptography (e.g. as a basis of hardness for new cryptographic protocols).

Equipped with a basic understanding of error-correcting codes, we will now study error correction in the quantum realm.

2 Quantum Error Correction

Our first observation is that quantum errors are much more subtle compared to mere bit-flip errors in the classical case. Suppose we prepare a qubit in the state $|\psi\rangle$. Because quantum systems in the real world are highly susceptible to noise, such *quantum noise* could potentially corrupt our state in a unitary manner and map it to some other (and not necessarily orthogonal) state $E|\psi\rangle$. A priori E could be an arbitrary unitary which is beyond our control and also completely unknown to us.

Example: Phase error

As an illustrative example of a non-trivial single-qubit error, we consider a *phase error* of the form

$$E = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\varphi} \end{pmatrix}, \quad \text{for } \varphi \in [0, 2\pi].$$

The above phase error E has the following effect on a single-qubit state:

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle \qquad \overset{E}{\mapsto} \quad \alpha|0\rangle + e^{i\varphi}\beta|1\rangle\,.$$

This immediately suggests a number of serious challenges for quantum error correction:

Problem 1: A quantum error may come from an infinite set of possible unitary errors.

Problem 2: Detecting an error (via a quantum measurement) could destroy the logical state.

Problem 3: Adding redundancy by duplicating quantum information seems to violate no-cloning.

These challenges led people to believe that quantum error correction is fundamentally doomed to fail! Thanks to a remarkable insight, it turns out that even quantum errors can, in some sense, be discretized.

Fact: Pauli discretization

Any unitary $U \in \mathbb{C}^{2\times 2}$ (and thus any single-qubit error) can be written as a complex linear combination of the four Pauli matrices:

$$U = \alpha I + \beta X + \gamma Y + \delta Z$$

for some $\alpha, \beta, \gamma, \delta \in \mathbb{C}$ such that $|\alpha|^2 + |\beta|^2 + |\gamma|^2 + |\delta|^2 = 1$.

The four Pauli matrices can be thought of as *elementary* quantum errors:

- $I = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$ (no error)
- $X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ (bit-flip error)
- $Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}$ (bit-flip and phase flip error)
- $Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$ (phase flip error)

Because $Y = iXZ$, a Pauli Y is simply a Z error followed by an X error (up to a global sign), and hence it suffices to focus on only X and Z errors instead.

Discretizing quantum errors. The Pauli discretization observation allows us to argue that any (unknown) single-qubit error E can be always be written as

$$E = \alpha I + \beta X + \gamma Y + \delta Z,$$

for complex coefficients $\alpha, \beta, \gamma, \delta \in \mathbb{C}$ such that $|\alpha|^2 + |\beta|^2 + |\gamma|^2 + |\delta|^2 = 1$.

Applying the error E to the state $|\psi\rangle$ yields

$$\begin{aligned} E|\psi\rangle &= (\alpha I + \beta X + \gamma Y + \delta Z)|\psi\rangle \\ &= \alpha|\psi\rangle + \beta X|\psi\rangle + \gamma Y|\psi\rangle + \delta Z|\psi\rangle . \end{aligned}$$

Thus, after the error occurs, the quantum state is a *coherent superposition* of four distinct error states, corresponding to the identity error and the three Pauli errors.

At first glance, this appears problematic: quantum error correction is supposed to correct *discrete* errors, yet the post-error state is a continuous linear combination of them.

The key insight is that *quantum mechanics is linear*. Suppose a quantum error correction procedure successfully corrects each of the Pauli errors I, X, Y, Z individually. Then, by linearity, it automatically corrects any linear combination of them. More precisely, let $\mathcal{R}$ denote a recovery operation such that

$$\mathcal{R}(P\,|\psi\rangle) = |\psi\rangle \quad \text{for all } P \in \{I, X, Y, Z\}.$$

Then, this necessarily implies that

$$\begin{aligned}\mathcal{R}(E\,|\psi\rangle) &= \mathcal{R}(\alpha\,|\psi\rangle + \beta X\,|\psi\rangle + \gamma Y\,|\psi\rangle + \delta Z\,|\psi\rangle) \\ &= \alpha\mathcal{R}(|\psi\rangle) + \beta\mathcal{R}(X\,|\psi\rangle) + \gamma\mathcal{R}(Y\,|\psi\rangle) + \delta\mathcal{R}(Z\,|\psi\rangle) \\ &= (\alpha + \beta + \gamma + \delta)\,|\psi\rangle\,,\end{aligned}$$

which differs from $|\psi\rangle$ only by a global (and therefore physically irrelevant) scalar factor.

In practice, quantum error correction is implemented by performing a *syndrome measurement*. This measurement extracts information about which error occurred *without revealing information about the encoded quantum state itself*. Crucially, the syndrome measurement projects the post-error state onto one of the orthogonal error subspaces associated with the Pauli operators. As a result, the coherent superposition

$$E\,|\psi\rangle = \alpha\,|\psi\rangle + \beta X\,|\psi\rangle + \gamma Y\,|\psi\rangle + \delta Z\,|\psi\rangle$$

is *collapsed* into one of the four states

$$|\psi\rangle\,, \quad X\,|\psi\rangle\,, \quad Y\,|\psi\rangle\,, \quad Z\,|\psi\rangle\,,$$

with respective probabilities

$$|\alpha|^2, \quad |\beta|^2, \quad |\gamma|^2, \quad |\delta|^2.$$

After this measurement-induced collapse, the error is no longer continuous but has become a *classical random Pauli error*, which can be corrected by a conditional recovery operation.

3 Correcting Pauli X errors

Inspired by the classical repetition code from before, we will now attempt to construct a similar *quantum* repetition code for qubits. Once again, the trick will be to simply duplicate the logical information.

Example: Quantum repetition code (for X errors)

A $1 \mapsto 3$ quantum repetition code (for X errors) is a natural mechanism for adding redundancy in the quantum realm: simply repeat each logical basis state mulitple times:

$$|0\rangle \mapsto |\bar{0}\rangle = |000\rangle \qquad |1\rangle \mapsto |\bar{1}\rangle = |111\rangle .$$

Note that this encoding procedure does not violate no-cloning as we are merely duplicating information which is encoded in the computational basis via CNOTs. Therefore, we can use our $1 \mapsto 3$ quantum repetition code to encode an arbitrary logical qubit as follows:

$$|\psi\rangle = \alpha |0\rangle + \beta |1\rangle \qquad \mapsto \qquad |\overline{\psi}\rangle = \alpha |000\rangle + \beta |111\rangle$$

where the encoding circuit for the quantum repetition code is given by

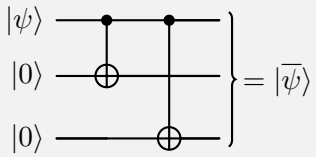

Error detection and correction. Suppose that an X error occurs on one of our three *physical* qubits, say

$$|\overline{\psi}\rangle \quad \mapsto \quad (I \otimes X \otimes I) |\overline{\psi}\rangle = \alpha |010\rangle + \beta |101\rangle .$$

How can we diagnose where the error has occurred? The naive strategy would be to simply measure $|\overline{\psi}\rangle$ in the computational basis, which would either result in 010 (with probability $|\alpha|^2$) or 101 (with probability $|\beta|^2$). In either of the two cases, we would easily be able to tell where the error has occurred by using the decoding procedure of the classical repetition code. Unfortunately, our measurement has forced the state to collapse, which means that the initial state $|\overline{\psi}\rangle$ is now lost. Can we do better?

To avoid destroying the state, we will instead perform a careful sequence of two-outcome measurements which will ensure that the state never collapses. The key idea is that it suffices to perform *parity* measurements—both on the first two, as well as the last two qubits. We will think of the parity bits as *error syndromes*. If the parities are both 0, we can conclude that no error has occurred; however, if one or both of the parities is 1, this will provide us with enough information to infer which error has occurred.

To carry out these parity measurements, we append two ancilla qubits in the state $|0\rangle$.

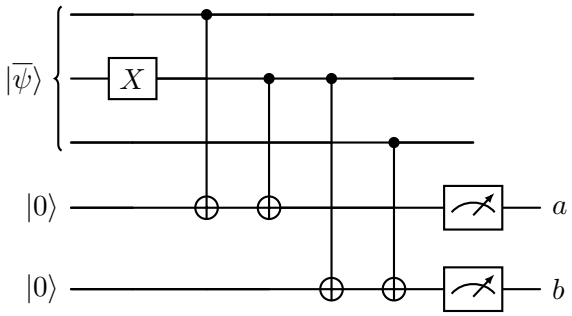

Here, the syndrome $a \in \{0, 1\}$ specifies the parity of the first two qubits, and $b \in \{0, 1\}$ specifies the parity of the last two qubits. In our example above, we would observe the outcome $a = 1$ and $b = 1$. However,

our syndrome detection mechanism also extends to an arbitrary X error on any one the three qubits. Based on the two syndrome bits, we can always infer which X error has occured:

$a \in \{0,1\}$	$b \in \{0,1\}$	**Error**
0	0	$I \otimes I \otimes I$
0	1	$I \otimes I \otimes X$
1	0	$X \otimes I \otimes I$
1	1	$I \otimes X \otimes I$

Once our syndrome measurement is complete and we successfully diagnosed the location of the Pauli X error, we can simply perform error correction by applying the appropriate X gate on the faulty qubit.

4 Correcting Pauli Z errors

We now turn our attention to correcting *phase flip* errors, i.e. Pauli Z errors. At first glance, this may appear quite different from the bit flip errors discussed previously. However, the key idea is that phase errors can be converted into bit flip errors by working in a different basis.

Action of a Z error. Recall that the Pauli Z operator acts on the computational basis as

$$Z|0\rangle = |0\rangle, \qquad Z|1\rangle = -|1\rangle.$$

Therefore, if a Z error acts on an arbitrary single-qubit state

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

then the resulting state is

$$Z|\psi\rangle = \alpha|0\rangle - \beta|1\rangle.$$

We see that a Z error does not change the computational basis states themselves, but instead flips the *relative phase* between them. This is why Pauli Z errors are often referred to as phase flip errors.

Working in the $\{|+\rangle, |-\rangle\}$ basis. To correct Z errors, it is convenient to express our states in the Hadamard basis, defined by

$$|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle), \qquad |-\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle).$$

In this basis, the action of the Z operator becomes particularly simple. Indeed, we compute

$$Z|+\rangle = \frac{1}{\sqrt{2}}(Z|0\rangle + Z|1\rangle) = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle) = |-\rangle,$$

and similarly

$$Z|-\rangle = \frac{1}{\sqrt{2}}(Z|0\rangle - Z|1\rangle) = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) = |+\rangle.$$

Thus, in the $\{|+\rangle, |-\rangle\}$ basis, a Pauli Z error acts exactly like a *bit flip*:

$$Z \equiv X \quad \text{in the } \{|+\rangle, |-\rangle\} \text{ basis.}$$

This observation allows us to reuse the quantum repetition code from before—provided we encode our logical information in the $|+\rangle$ and $|-\rangle$ states.

Example: Quantum repetition code (for Z errors)

A $1 \mapsto 3$ quantum repetition code for correcting Z errors is obtained by repeating the logical states in the Hadamard basis:

$$|+\rangle \mapsto |\bar{+}\rangle = |+++\rangle, \qquad |-\rangle \mapsto |\bar{-}\rangle = |---\rangle.$$

An arbitrary logical qubit can be written as

$$|\psi\rangle = \alpha |+\rangle + \beta |-\rangle,$$

which is encoded as

$$|\overline{\psi}\rangle = \alpha |+++\rangle + \beta |---\rangle.$$

The encoding circuit for the Z-error repetition code is obtained by first applying a Hadamard gate to each qubit, followed by the same CNOT-based repetition procedure as before:

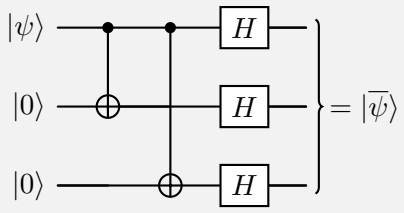

Error detection and correction. Suppose that a single Z error occurs on one of the three physical qubits, for instance on the second qubit:

$$|\overline{\psi}\rangle \quad \mapsto \quad (I \otimes Z \otimes I)\,|\overline{\psi}\rangle = \alpha |+-+\rangle + \beta |-+-\rangle.$$

Because Z errors act like bit flips in the $\{|+\rangle, |-\rangle\}$ basis, we can diagnose them using the same parity-check strategy as before—now interpreted in the Hadamard basis.

To implement this, we first apply Hadamard gates to rotate back to the computational basis, perform the usual parity checks, and then apply Hadamards again if desired.

As before, we append two ancilla qubits in the state $|0\rangle$ and measure the parity of the first two and last two qubits. Using that $ZH = HX$, we can represent the Z error as an X error in the Hadamard basis:

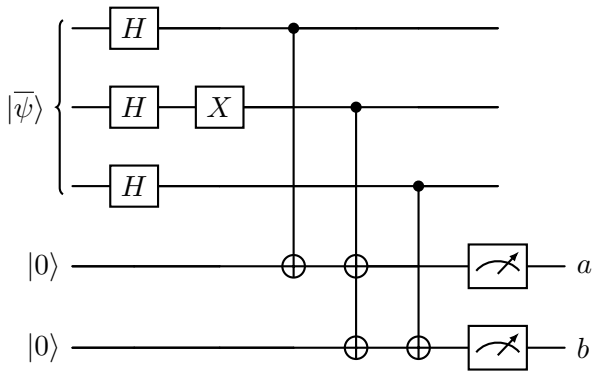

Here, the syndrome bits $a, b \in \{0, 1\}$ again encode the parity information. For a single Z error on any one of the three qubits, the syndrome outcomes uniquely identify the location of the error:

$a \in \{0, 1\}$	$b \in \{0, 1\}$	**Error**
0	0	$I \otimes I \otimes I$
0	1	$I \otimes I \otimes Z$
1	0	$Z \otimes I \otimes I$
1	1	$I \otimes Z \otimes I$

Once the syndrome measurement is complete, we can correct the error by applying a Z gate to the faulty qubit. In this way, the $1 \mapsto 3$ quantum repetition code successfully protects a logical qubit against a single Pauli Z error.

5 The Shor 9-qubit code

We have now seen how to protect a logical qubit against a single Pauli X error using a quantum repetition code in the computational basis, as well as how to protect against a single Pauli Z error by working in the Hadamard basis. We now combine these two ideas to obtain a quantum error-correcting code that can correct *both* types of errors simultaneously. This construction is known as the *Shor 9-qubit code*, and it was the first quantum error-correcting code ever discovered.

Idea of the construction. The key idea is to concatenate the two repetition codes we have already studied:

- First, we protect against Z (phase flip) errors by encoding a logical qubit using the $1 \mapsto 3$ repetition code in the $\{|+\rangle, |-\rangle\}$ basis.
- Then, we protect against X (bit flip) errors by further encoding *each* of the three resulting qubits using the $1 \mapsto 3$ repetition code in the computational basis.

In this way, a single logical qubit is encoded into $3 \times 3 = 9$ physical qubits.

Logical codewords. Recall that the Z-error repetition code maps

$$|+\rangle \mapsto |+++\rangle\,, \qquad |-\rangle \mapsto |---\rangle\,.$$

Each of the states $|+\rangle$ and $|-\rangle$ can themselves be written as

$$|\pm\rangle = \frac{1}{\sqrt{2}}(|0\rangle \pm |1\rangle).$$

Applying the X-error repetition code to each qubit then yields the Shor codewords

$$|\bar{0}\rangle = \frac{1}{\sqrt{2}}\big(|000\rangle\,|000\rangle\,|000\rangle + |111\rangle\,|111\rangle\,|111\rangle\big),$$
$$|\bar{1}\rangle = \frac{1}{\sqrt{2}}\big(|000\rangle\,|000\rangle\,|000\rangle - |111\rangle\,|111\rangle\,|111\rangle\big).$$

An arbitrary logical qubit $|\psi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle$ is therefore encoded as $|\overline{\psi}\rangle = \alpha\,|\bar{0}\rangle + \beta\,|\bar{1}\rangle\,.$

Encoding circuit. The encoding circuit for the Shor code follows directly from its construction:

1. First, encode the logical qubit using the Z-error repetition code.
2. Then, apply the X-error repetition code separately to each of the three qubits.

Schematically, we can represent both encoding and decoding in Figure 1.

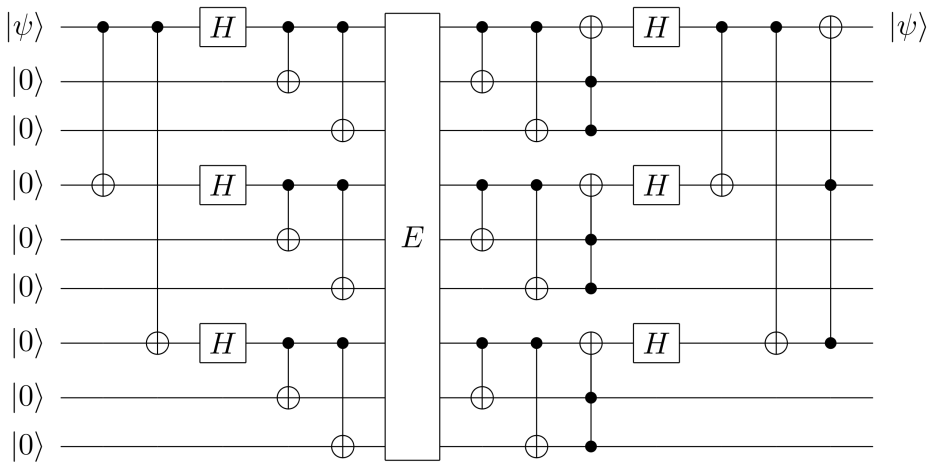

Figure 1: Shor's 9-qubit code.

Error correction. The Shor code corrects errors in two stages:

- *Bit flip correction:* Within each block of three qubits, we perform the same parity measurements as in the X-error repetition code to detect and correct a single X error.
- *Phase flip correction:* After correcting bit flip errors, we treat each block as a single logical qubit and perform the Z-error repetition code parity checks across the three blocks.

Because any single-qubit error can be written (up to a global phase) as a product of Pauli X and Z operators, correcting both types of errors suffices to protect against *arbitrary* single-qubit errors.