

Surviving as a Quantum Computer in a Classical World

Daniel Gottesman

February 21, 2024

Contents

I	Quantum Error Correcting Codes	7
1	Know Your Enemy: Quantum Errors	9
1.1	The Quantum Channel	9
1.2	Single-Qubit Example Channels	11
1.3	Multiple-Qubit Channels	17
1.4	A Peek Ahead: Errors During Computation	24
2	Redundancy Without Repetition: Basics Of Quantum Error Correction	25
2.1	Quantum Error Correction? Ridiculous!	25
2.2	The 3-Qubit Code(s)	26
2.3	The 9-Qubit Code	29
2.4	Correcting General Errors	30
2.5	The Quantum Error Correction Conditions	35
2.6	What Makes a Good QECC?	41
3	Will The Real Codeword Please Stand Still?: Stabilizer Codes	43
3.1	The 9-Qubit Code Revisited	43
3.2	Pauli Group	45
3.3	Stabilizer Codes	46
3.4	Cosets and Error Syndromes	52
3.5	Binary Symplectic Representation	56
4	Maybe I Should Have Started Here: Classical Error Correction	59
4.1	Classical Error Correction in General	59
4.2	Classical Linear Codes	61
4.3	Dual Codes	66
4.4	Non-Binary Linear Codes	67
4.5	Hamming, Gilbert-Varshamov, and Singleton Bounds, MDS Codes	70
5	Combining The Old And The New: Making Quantum Codes From Classical Codes	75
5.1	CSS Codes	75
5.2	$GF(4)$ Codes and Stabilizer Codes	80
6	Symmetries Of Symmetries: The Clifford Group	85
6.1	Definition of the Clifford Group	85
6.2	Classical Simulation of the Clifford Group	91
6.3	Generators of the Clifford Group	98
6.4	Encoding Circuits for Stabilizer Codes	102
6.5	Extending the Clifford Group to a Universal Gate Set	105

7	Tighter, Please: Upper And Lower Bounds On Quantum Codes	107
7.1	The Quantum Gilbert-Varshamov Bound	107
7.2	The Quantum Hamming Bound	109
7.3	The Quantum Singleton Bound	110
7.4	Linear Programming Bounds	112
8	Bigger Can Be Better: Qudit Codes	123
8.1	Qudit Pauli Group	123
8.2	Qudit Stabilizer Codes	129
8.3	Qudit CSS Codes	133
8.4	Qudit Clifford Group	136
9	Now, What Did I Leave Out?: Other Things You Should Know About Quantum Error Correction	139
9.1	Concatenated Codes	139
9.2	Convolutional Codes	142
9.3	Information-Theoretic Approach to QECCs	146
9.4	Approximate Quantum Error-Correcting Codes	149
II	Fault-Tolerant Quantum Computation	155
10	Everyone Makes Mistakes: Basics Of Fault Tolerance	157
10.1	The Fault-Tolerant Scenario	157
10.2	Formal Definition of Fault Tolerance	163
10.3	Statement of the Threshold Theorem for the Basic Model	169
10.4	Proofs Versus Simulations for Fault Tolerance	170
11	If Only It Were So Easy: Transversal Gates	173
11.1	What is a Transversal Gate?	173
11.2	Transversal Gates for Stabilizer Codes	176
11.3	Transversal Gates for the 7-Qubit Code	178
11.4	Transversal Gates and Measurement for CSS Codes	179
11.5	Other Topics Relating to Transversal Gates	182
12	Who Corrects The Correctors?: Fault-Tolerant Error Correction And Measurement	185
12.1	Fault Tolerant Pauli Measurement for Stabilizer Codes	185
12.2	Shor Error Correction	193
12.3	Steane Error Correction and Measurement	196
12.4	Knill Error Correction and Measurement	202
12.5	Efficiency of FTEC Protocols	205
13	Any Sufficiently Advanced Fault-Tolerant Protocol is Indistinguishable from Magic States: State Preparation And Its Applications	209
13.1	Preparation of Encoded Stabilizer States	209
13.2	Gate Teleportation	213
13.3	Compressed Gate Teleportation	216
13.4	Clifford Eigenstate Preparation by Measurement	217
13.5	Magic State Distillation	220

14 If It's Worth Doing, It's Worth Overdoing: The Threshold Theorem	227
14.1 Adversarial Errors	227
14.2 Good and Bad Extended Rectangles	229
14.3 Correctness	230
14.4 Incorrectness: Simulations With a Bad Extended Rectangle	234
14.5 Probability of Having a Bad Rectangle	237
14.6 Level Reduction	242
14.7 Concatenation and the Threshold Theorem	245
15 Now Hold On Just a Second: Assumptions Re-Examined	253
15.1 Solovay-Kitaev Theorem	253
15.2 Short-Range Gates	255
15.3 Fresh Ancilla Qubits	258
15.4 Slow Measurement or No Intermediate Measurement	262
15.5 Parallelism and Timing	265
15.6 Leakage Errors	267
15.7 Biased Errors	269
15.8 Error Independence and Non-Markovian Errors	272
16 Now, What Did I Leave Out, Part Two?: Other Things You Should Know About Fault Tolerance	287
III Miscellaneous Topics	289
17 They May Not Be Error Correction, But We Still Love Them: Other Methods Of Error Control	291
17.1 Entanglement Distillation and Quantum Repeaters	291
17.2 Subsystem Coding	291
17.3 Entanglement-Assisted Codes	291
17.4 Decoherence-Free Subspaces	291
17.5 Dynamical Decoupling	291
17.6 Error Mitigation	291
18 Wholesale Error Correction: Channel Capacity	293
19 Donuts. Is There Anything They Can't Do?: Topological Codes	295
19.1 Toric Code and Surface Codes	296
19.2 Decoding of Surface Codes	303
19.3 Fault Tolerance with Surface Codes	311
20 It Certainly Helps If You Look At Things The Right Way: Graph States	321
IV Appendices	323
A Quantum Computation	325
B Group Theory	327
C Finite Fields	329
D Linear Algebra	331

Part I

Quantum Error Correcting Codes

Chapter 1

Know Your Enemy: Quantum Errors

In this book, you will learn how to seek out and destroy errors on quantum states. Quantum errors are nasty, unforgiving things. If you don't know what you are doing, a single misstep can result in the destruction of irreplaceable quantum information. Of course, nobody's perfect, and in part II, you will learn how to handle your own fallibility. (Short answer: very carefully.) For now, we will assume you don't make any mistakes, but that doesn't mean things will be easy. The quantum errors are still out there, hungering to consume quantum information. You need to get the errors before they get you.

The key to most error hunts is preparation. There are two forms of preparation. One is dressing your quantum information properly: that is, encoding it in an appropriate quantum error-correcting code. You will learn about that in the other chapters in part I. This chapter will focus on the other aspect of preparation: studying the habits of the errors you are hunting.

1.1 The Quantum Channel

The most common way of characterizing a source of errors in a quantum system is as a quantum channel. “Quantum channel” is a general term for an opportunity for the environment to introduce errors into your quantum system. Usually, we consider a situation as depicted in figure 1.1. Alice wants to send quantum information to Bob. Alice and Bob both have perfect quantum computers, but the connection between them is possibly imperfect: the quantum channel. (We will drop the assumption that Alice and Bob have perfect quantum computers in part II.) While the quantum information is in transit, the external world (the “environment”) has a chance to interact with the system, thereby changing it in some way, introducing errors (or “noise”) relative to the state that Alice sent.

A common example of a quantum channel is an optical fiber. Single photons can pass through the optical fiber, but they may be lost or altered en route. Other possibilities are sending a photon through the air, physically moving an atom with information encoded in the state of the electron or nuclear spin, successively swapping the states of neighboring qubits arranged on a line, or even using quantum teleportation to send

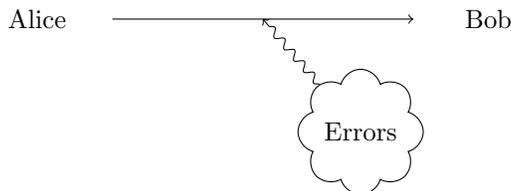


Figure 1.1: Alice, who has a perfect quantum computer, wants to send qubits to Bob, who also has a perfect quantum computer. However, the communications channel between them is *not* perfect.

a quantum state from Alice to Bob with classical communication and some sort of entangled state. This is not an exhaustive list. Indeed, any sort of communication, even a classical telephone call, can be considered as a quantum channel. If you try to send a qubit through a regular telephone connection, no amount of quantum error correction will allow you to recover the full quantum state afterwards, but that doesn't affect the telephone line's status as a quantum channel — it is simply a *very noisy* quantum channel.

Another common situation is when Bob is replaced by Alice's future self. In this case, we really want a “quantum memory”: Alice wants to prepare some quantum information, go off and do other things, and then return and manipulate her stored quantum information again. If we assume that Alice's manipulations at the beginning and the end of the process are perfect, we can consider the “memory” portion, when the qubit is stored but subject to noise, as a quantum channel.

It is worth noting that the notion of a quantum channel only applies when we can look at a single communications link in isolation. That is, to have a quantum channel, the quantum state that exits the channel should only depend on the quantum state that goes into the channel. That may seem like a tautology, but it is not. Imagine that Alice has a quantum memory, and prepares a qubit to store in the memory at time 0. At time 1, she returns and fiddles some more with the stored system, then goes away again and comes back at time 2. The storage from time 0 to time 1 is a quantum channel (assuming Alice's initial preparation of the qubit is perfect), but the storage from time 1 to time 2 might not be. The problem is that the environment might remember what happened between time 0 and 1 (and more importantly, might remember something about the *state* that was stored between time 0 and 1) and use that to influence what it does to the state during the second time interval. The error now no longer depends only on what state is stored at time 1; it also depends on what state was stored at time 0. Of course, some environments have a very short memory, and in that case, it is a very good approximation to consider the time interval 1 to 2 to be independent of the time interval 0 to 1, and with that approximation, the second time interval *is* a quantum channel. When the environment has no memory, and every time interval is independent of any other non-overlapping time interval, the environment (or the error source) is called *Markovian*. When the environment does remember over time scales long enough to matter, it is a *non-Markovian* environment.

In part I, we only consider the case depicted in figure 1.1. There the environment has only one opportunity to attack the quantum information. While that opportunity may last for an extended period of time, because Alice and Bob do not do anything with the quantum information during that time, we can lump together everything the environment does to the state into a single transformation, and consider the whole communications line as a single quantum channel. The question of whether the environment is Markovian or non-Markovian then becomes moot. If we were to generalize this picture and allow noise during Alice and Bob's processing of the qubit, then the question arises again, since the environment then gets more than one shot at the quantum information, but don't worry about that situation until part II.

Now it is time to formally define a quantum channel:

Definition 1.1. A *quantum channel* is a completely positive trace-preserving (CPTP) map.

Wasn't that easy? At least, it is if you know what a CPTP map is. If not, you should refer to appendix A, where you will learn that a CPTP map is the most general physically possible transformation for an operation where the output depends only on the input, so this is the right definition. It is frequently convenient to consider the Kraus form of a CPTP map:

$$\mathcal{E}(\rho) = \sum_k A_k \rho A_k^\dagger. \tag{1.1}$$

We can think of this channel as a collection of possible errors A_k , where error A_k occurs with probability $\text{tr}(A_k \rho A_k^\dagger)$. However, note that the probability of A_k occurring is not a single value but actually depends on the state ρ . Furthermore, remember that the decomposition into A_k operators is not unique and that $\sum_k A_k^\dagger A_k = I$.

Frequently, instead of dealing explicitly with a quantum channel, I will instead refer to the *set of possible errors*. One way to write this set is $\mathcal{E} = \{A_k\}$, but it is frequently convenient to rescale the errors, so more generally $\mathcal{E} = \{E_k\}$, where each $A_k = p_k E_k$ for some scalar p_k .

1.2 Single-Qubit Example Channels

1.2.1 Unitary Channels, Pauli Errors

Let us start by discussing some examples of channels acting on a single-qubit input. The simplest case is when there is only a single value of k in the Kraus decomposition:

$$\mathcal{E}(\rho) = A\rho A^\dagger. \quad (1.2)$$

Since $\sum_k A_k^\dagger A_k = I$, it follows that A is unitary. Typographically, I will usually represent a unitary channel the same way as a unitary matrix, e.g. $A(\rho)$ versus $A(|\psi\rangle)$, even though they formally act on different kinds of objects (density matrices versus state vectors). The one exception is that I will usually write the identity channel as \mathcal{I} , as opposed to the identity unitary I .

There are of course infinitely many unitary maps, but some are more interesting than others. One set that you will be particularly sick of by the end of this book are the Pauli matrices:

$$I = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \quad (1.3)$$

I is, of course, the identity matrix — no error. X is probably the first thing you think of when you think of an error, a classical bit flip:

$$X|0\rangle = |1\rangle \quad (1.4)$$

$$X|1\rangle = |0\rangle. \quad (1.5)$$

However, since it is a quantum operator, it also acts sensibly on superpositions:

$$X(\alpha|0\rangle + \beta|1\rangle) = \alpha|1\rangle + \beta|0\rangle. \quad (1.6)$$

Z is the most basic kind of truly quantum error, a phase flip:

$$Z(\alpha|0\rangle + \beta|1\rangle) = \alpha|0\rangle - \beta|1\rangle. \quad (1.7)$$

Y is then just a combined bit flip and phase flip error:

$$Y = iXZ \quad (1.8)$$

$$Y(\alpha|0\rangle + \beta|1\rangle) = i\alpha|1\rangle - i\beta|0\rangle. \quad (1.9)$$

Recall that an *overall phase* — one that affects all states uniformly — has no physical significance, so Y and XZ are really the same channel. In the form presented above, all the Pauli matrices are Hermitian as well as unitary, which is sometimes useful.

In other contexts, the Pauli matrices are often written as σ_x , σ_y , and σ_z or σ_1 , σ_2 , and σ_3 , but those are too much writing and less easy to read. I'll be using the Pauli matrices a *lot* in this book, so I'll use the more straightforward notation X , Y , and Z . In some of the earlier quantum error-correction literature, $Y = XZ$ instead of iXZ . There is not a huge difference, but I think this convention is somewhat nicer overall.

There are many more single qubit unitary errors, and some are even interesting. For instance, we can have phase rotation by an arbitrary angle:

$$R_\theta = \begin{pmatrix} e^{-i\theta} & 0 \\ 0 & e^{i\theta} \end{pmatrix} = e^{-i\theta} \begin{pmatrix} 1 & 0 \\ 0 & e^{2i\theta} \end{pmatrix}. \quad (1.10)$$

Again, we can ignore an overall phase, so R_θ is the same channel as $\text{diag}(1, e^{2i\theta})$. The full set of physically distinct one-qubit unitary channels is the group $\text{SU}(2)$.

In the Bloch sphere picture of the state space for a qubit, a unitary map is just a rotation of the sphere. X , Y , and Z are π rotations around the X , Y , and Z axes, as one might expect. R_θ is a rotation by angle 2θ around the Z axis. (I have defined R_θ this way in order to agree with the prevailing terminology for phase rotations, which results from talking about spin-1/2 particles.) Note that a reflection of the Bloch sphere is not a completely positive map. For instance, the transpose map is a reflection in the XZ plane, and when applied to part of an entangled state, the transpose gives something non-positive.

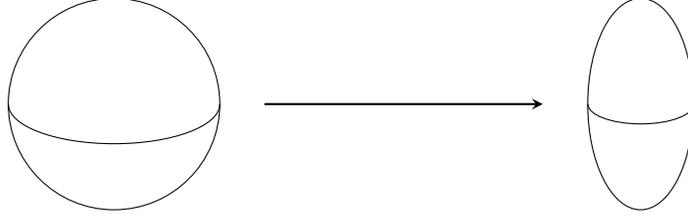


Figure 1.2: Bloch sphere transformation induced by a dephasing channel

1.2.2 Dephasing Channel

One very commonly-encountered channel is the dephasing channel.

Definition 1.2. A channel of the form

$$\mathcal{R}_p(\rho) = (1 - p)\rho + pZ\rho Z^\dagger. \quad (1.11)$$

is a *dephasing channel*. When $p = 1/2$, we have the *completely dephasing channel*. We usually restrict attention to $p \leq 1/2$ since channels with $p > 1/2$ are related to channels with $p < 1/2$ by a Z operation.

The Kraus operators of a dephasing channel in this form are $\sqrt{1-p}I$ and $\sqrt{p}Z$. Since both are proportional to unitary maps, the probabilities of these two errors occurring do not depend on the input state. Thus, this channel corresponds to no error with probability $1 - p$ and a phase flip with probability p . We can calculate how it acts on the density matrix in component form:

$$\mathcal{R}_p : \begin{pmatrix} a & b \\ c & d \end{pmatrix} \mapsto \begin{pmatrix} a & (1-2p)b \\ (1-2p)c & d \end{pmatrix}. \quad (1.12)$$

In other words, a dephasing channel shrinks the off-diagonal components of the density matrix. In the completely dephasing channel, the off-diagonal terms disappear completely.

An alternate Kraus decomposition is also edifying.

$$\mathcal{R}_p(\rho) = (1 - 2p)\rho + \frac{2p}{\pi} \int_0^\pi R_\theta \rho R_\theta^\dagger d\theta. \quad (1.13)$$

The dephasing channel is a channel for which, with probability $1 - 2p$, nothing happens to the state, and with probability $2p$, the phase between $|0\rangle$ and $|1\rangle$ is completely randomized. This seems to conflict with the decomposition into I and Z , where the state is left unchanged with probability $1 - p$, not with probability $1 - 2p$. However, when the dephasing angle θ is chosen uniformly at random, there is a reasonable chance that θ is small and the state does not change very much. Of course, the probability that the angle θ is *exactly* zero is just $1 - 2p$, but using the magic of quantum mechanics, the small- θ cases cancel out just right so that if we break the channel up into I and Z , we find the probability of error is only p .

This example illustrates a rather disturbing principle: in quantum mechanics, the notion of “probability of error” is inherently somewhat subjective. When error correction is concerned, the decomposition into I and Z is the better choice for the dephasing channel, for reasons that will become clearer in chapter 2. However, there are many channels for which none of the decompositions is particularly favored, even for the specific application of quantum error correction.

In the Bloch sphere picture, a dephasing channel shrinks the sphere into an ellipsoid, leaving the Z axis unchanged, as in figure 1.2. A completely dephasing channel shrinks the Bloch sphere down to just the segment on the Z axis.

The dephasing channel is a physically very interesting channel. A dephasing channel occurs in any system where the environment learns about the qubit in the standard basis but does not otherwise interfere with the state. Even in a more realistic system where there are more complicated interactions between the

system and the environment, there is frequently a large dephasing component to the noise. The prevalence of approximate dephasing channels is one of the reasons that the macroscopic world appears classical to us — a completely dephasing channel converts a qubit into a probabilistic classical bit.

We can make a simple model of a dephasing channel by having one environment qubit interact with the system qubit via a Hamiltonian $H = \omega Z \otimes Z$. The environment qubit starts in the pure state $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$, and the system qubit starts in state $|\psi_0\rangle = \alpha|0\rangle + \beta|1\rangle$. Then, after a time t , the state of the two qubits is

$$e^{-i\omega t Z \otimes Z} |\psi_0\rangle \otimes |+\rangle = \frac{1}{\sqrt{2}} (\alpha e^{-i\omega t} |00\rangle + \alpha e^{+i\omega t} |01\rangle + \beta e^{+i\omega t} |10\rangle + \beta e^{-i\omega t} |11\rangle). \quad (1.14)$$

(Taking $\hbar = 1$.) Tracing over the second (environment) qubit, we find that the system qubit at time t is in an equal mixture of $R_{-\omega t}|\psi_0\rangle$ and $R_{+\omega t}|\psi_0\rangle$, giving it density matrix

$$\begin{pmatrix} |\alpha|^2 & \alpha\beta^* \cos(2\omega t) \\ \alpha^*\beta \cos(2\omega t) & |\beta|^2 \end{pmatrix}. \quad (1.15)$$

In other words, at time t , we have the dephasing channel $\mathcal{R}_{(1-\cos(2\omega t))/2}$. In this simple model, the system dephases at short times, but after a longer time $t = \pi/\omega$, it returns to its starting state.

In a more realistic system, there are many different environment qubits interacting with the system, and/or there are additional qubits interacting with the environment qubits. If we take a Markovian form of our simple model, and assume that the environment's internal interaction effectively resets the environment qubit after a very short time, phase coherence instead decays steadily:

$$\begin{pmatrix} |\alpha|^2 & \alpha\beta^* e^{-t/T_2} \\ \alpha^*\beta e^{-t/T_2} & |\beta|^2 \end{pmatrix}, \quad (1.16)$$

corresponding to a dephasing channel $\mathcal{R}_{(1-e^{-t/T_2})/2}$. The time constant T_2 of the exponential decay is known as the t_2 time. (Otherwise I probably would have used a different symbol.) One way to think about this behavior is that there is a constant probability $1/T_2$ per unit time that the phase is completely randomized as an instantaneous “quantum jump”.

Another common source of dephasing is a varying energy difference between the $|0\rangle$ and $|1\rangle$ states. If $|a\rangle$ has energy E_a , after time t , $|a\rangle$ has evolved into $e^{-iE_a t}|a\rangle$. We can ignore the global phase, but there is still a relative phase difference $e^{-i(E_1-E_0)t}$ between $|0\rangle$ and $|1\rangle$. However, when E_0 and E_1 are known constants, we can generally ignore the relative phase too: if we keep track of the time t , the relative phase is known, and we can take it into account in any operation we want to perform. In some systems, there is a phase reference (such as the phase of a laser) which automatically compensates for the phase difference. However, when the energy difference varies unpredictably, we can no longer keep precise track of the relative phase, resulting in an uncompensated relative phase shift accumulating randomly over time. This results in dephasing. In some cases, the relative phase shift is not random, but is simply unknown, and more sophisticated techniques may be able to compensate.

Experimentally, the signature of dephasing is often a decay of coherent interference effects. In a Rabi oscillation experiment, the system cycles $\cos(\Omega_R t)|0\rangle + \sin(\Omega_R t)|1\rangle$. After time t , the system is measured to test if it is $|0\rangle$ or $|1\rangle$. If the system were perfect, if we were to plot the probability of 1 against the time, we would get a perfect oscillation $\sin^2(\Omega_R t)$ for all times. Instead, we get something more like figure 1.3, with oscillations decreasing in amplitude.

Let us imagine a Markovian environment, and assume that the T_2 time is much longer than the Rabi frequency Ω_R , so that we can assume the dephasing and oscillation are independent. If there is a quantum jump causing full dephasing at time t_0 , the pure state $\cos(\Omega_R t_0)|0\rangle + \sin(\Omega_R t_0)|1\rangle$ becomes the mixed state with probability $\cos^2(\Omega_R t_0)$ of $|0\rangle$ and probability $\sin^2(\Omega_R t_0)$ of $|1\rangle$. Then $|0\rangle$ and $|1\rangle$ continue with their

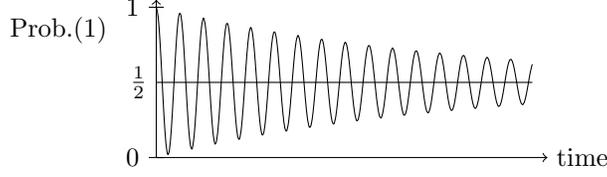


Figure 1.3: Decay of Rabi oscillations due to dephasing

own oscillations, but out of phase with each other. This mixture has density matrix

$$\begin{pmatrix} \cos^2(\Omega_R t_0) \cos^2(\Omega_R t') + \sin^2(\Omega_R t_0) \sin^2(\Omega_R t') & (\cos^2(\Omega_R t_0) - \sin^2(\Omega_R t_0)) \sin(\Omega_R t') \cos(\Omega_R t') \\ (\cos^2(\Omega_R t_0) - \sin^2(\Omega_R t_0)) \sin(\Omega_R t') \cos(\Omega_R t') & \cos^2(\Omega_R t_0) \sin^2(\Omega_R t') + \sin^2(\Omega_R t_0) \cos^2(\Omega_R t') \end{pmatrix} \quad (1.17)$$

$$= \begin{pmatrix} \sin^2(\Omega_R t_0) + \cos(2\Omega_R t_0) \cos^2(\Omega_R t') & \cos(2\Omega_R t_0) \sin(\Omega_R t') \cos(\Omega_R t') \\ \cos(2\Omega_R t_0) \sin(\Omega_R t') \cos(\Omega_R t') & \sin^2(\Omega_R t_0) + \cos(2\Omega_R t_0) \sin^2(\Omega_R t') \end{pmatrix} \quad (1.18)$$

$$= \sin^2(\Omega_R t_0) I + \cos(2\Omega_R t_0) \rho(t'), \quad (1.19)$$

where $t' = t - t_0$ and $\rho(t')$ is the density matrix of a Rabi oscillation running for time t' . We will still see Rabi oscillation, but with a reduced amplitude.

1.2.3 Depolarizing Channel and Pauli Channel

While the dephasing channel is perhaps the favorite (or least favorite, depending on your point of view) of experimentalists, it lacks a certain quantumness. It only involves one kind of non-trivial error, and indeed, in an appropriate choice of basis, it can be viewed as a purely classical noisy channel. Theorists are instead enamored of the depolarizing channel, which does experience the full range of quantum errors and is highly symmetric. It is not so common in the real world, which is less symmetric than theorists would prefer, but it is still extremely useful for exploring quantum error correction.

Definition 1.3. The *depolarizing channel* is the quantum channel

$$\mathcal{D}_p(\rho) = (1-p)\rho + \frac{p}{3}X\rho X^\dagger + \frac{p}{3}Y\rho Y^\dagger + \frac{p}{3}Z\rho Z^\dagger. \quad (1.20)$$

$\mathcal{D}_{3/4}$ is the *completely depolarizing channel*.

The depolarizing channel has an equal chance of an X , Y , or Z error, each with probability $p/3$. We can analyze what it does to the density matrix:

$$\mathcal{D}_p : \begin{pmatrix} a & b \\ c & d \end{pmatrix} \mapsto \begin{pmatrix} (1-2p/3)a + (2p/3)d & (1-4p/3)b \\ (1-4p/3)c & (1-2p/3)d + (2p/3)a \end{pmatrix}. \quad (1.21)$$

This is not particularly enlightening until you remember that $\text{tr } \rho = 1$, in which case you can see that

$$\mathcal{D}_p(\rho) = (1-4p/3)\rho + (4p/3)(I/2). \quad (1.22)$$

In other words, with probability $1-4p/3$, the qubit is left alone, but with probability $4p/3$, it is replaced with the completely mixed state. It also should now be clear why I defined $p = 3/4$ as the completely depolarizing channel: in that case, the input state is always replaced with the maximally mixed state. As with the dephasing channel, there is some ambiguity as to what the “true” error rate is for the depolarizing channel, but the two representations can be reconciled by recognizing that the completely mixed state does contain a component of the original input state.

The other thing to notice about the second representation is that it is much more symmetric than the first one. The decomposition into Paulis has a certain amount of symmetry, treating X , Y , and Z on the

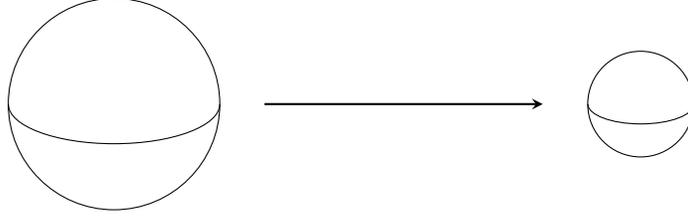


Figure 1.4: Bloch sphere transformation induced by a depolarizing channel

same footing, but the second decomposition has even more: there are no preferred bases or unitary operators at all appearing in it. This is the true beauty of the depolarizing channel — it can at once be considered as a simple mixture of the very basic Pauli errors, but also is invariant under any kind of unitary rotation. This symmetry property suggests a decomposition for the depolarizing channel akin to equation (1.13), and indeed there is one:

$$\mathcal{D}_p(\rho) = (1 - 4p/3)\rho + \int_{\text{SU}(2)} U\rho U^\dagger dU, \quad (1.23)$$

where the integral uses the Haar measure, the unitarily-invariant measure over $\text{SU}(2)$.

We can generalize the depolarizing channel by giving up its symmetry but keeping its decomposition into Paulis:

Definition 1.4. A channel of the form

$$\mathcal{E}(\rho) = p_I\rho + p_X X\rho X^\dagger + p_Y Y\rho Y^\dagger + p_Z Z\rho Z^\dagger \quad (1.24)$$

is a *Pauli channel*.

A Pauli channel has potentially different probabilities for the four Pauli matrices I , X , Y , and Z . They don't *have* to be different — dephasing channels and depolarizing channels are both examples of Pauli channels — but once you've generalized to a Pauli channel, you might as well take advantage of the opportunity to have some variety among the Paulis. Naturally, $p_I + p_X + p_Y + p_Z = 1$ so that the total probability adds up to 1.

The depolarizing channel uniformly shrinks the Bloch sphere into a smaller sphere still centered on the origin, or to a single point (the maximally mixed state) if we have the completely depolarizing channel. A more general Pauli map shrinks the Bloch sphere into an ellipsoid centered on the origin.

1.2.4 Amplitude Damping Channel

Another physically motivated channel is the amplitude damping channel. It occurs, for instance, if $|0\rangle$ and $|1\rangle$ are the ground and excited state of a two-level atom, and the excited state can decay to the ground state by spontaneously emitting a photon. The amplitude damping channel also occurs for a photonic qubit where $|0\rangle$ is no photon and $|1\rangle$ is 1 photon — damping occurs when the photon escapes or is absorbed somewhere along the way. There should be 2 Kraus operators for the amplitude damping channel, one representing the cases when there is no emission (or loss), and one representing the “no-jump” case. It is worthwhile trying to figure out yourself what the form of the Kraus operators should be before looking at the definition below.

Definition 1.5. An *amplitude damping channel* has the following form:

$$\mathcal{A}_p(\rho) = A_0\rho A_0^\dagger + A_1\rho A_1^\dagger, \quad (1.25)$$

where

$$A_0 = \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{1-p} \end{pmatrix}, \quad A_1 = \begin{pmatrix} 0 & \sqrt{p} \\ 0 & 0 \end{pmatrix}. \quad (1.26)$$

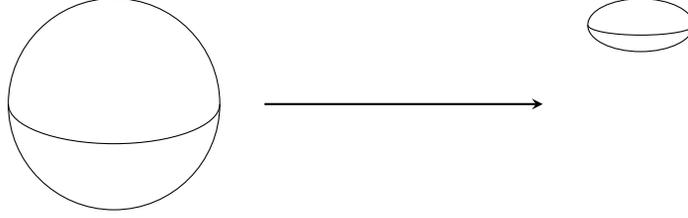


Figure 1.5: Bloch sphere transformation induced by an amplitude damping channel

Probably you were able to get the right form for A_1 (perhaps without the square root, which is a matter of convention depending on how we choose to parametrize amplitude damping channels). It represents the possibility that $|1\rangle$ can become $|0\rangle$. However, $|0\rangle$ never spontaneously gains energy in this idealized channel; even in the real world, it is fairly rare, since it requires that a stray photon of about the right energy be wandering by. Therefore the lower left corner of A_1 is 0.

A_0 might surprise you. (If not, then good work.) The most obvious guess is that since there is no decay, nothing should happen, and A_0 should be proportional to the identity. However, you won't be able to satisfy the constraint that $A_0^\dagger A_0 + A_1^\dagger A_1 = I$ if you choose A_1 as above and $A_0 \propto I$. Conceptually, the reason for this is that A_1 can only occur if the initial state was $|1\rangle$. Therefore, if A_1 *doesn't* happen, it means that the initial state was *more likely* to be $|0\rangle$, and A_0 reflects that, reducing the amplitude of $|1\rangle$ in the initial state. The same phenomenon can be found in classical probability theory, for instance in the “Monty Haul problem.”

In the Bloch sphere picture, amplitude damping results in a uniform shrinkage to a smaller sphere. It differs from the depolarizing channel in that the center of the sphere is no longer fixed. Instead, the south pole (the $|0\rangle$ state) is fixed. In the limit $p = 1$, the whole sphere shrinks down to the south pole.

The characteristic decay time for a system undergoing continuous amplitude damping is the “ T_1 time.” (After all, there had to be a T_1 to go with T_2 for the dephasing time scale.)

1.2.5 Photon Loss Channels

Amplitude damping can occur due to photon loss from a photon channel when the presence or absence of a photon represents $|0\rangle$ or $|1\rangle$. However, this is not a particularly common encoding used for photonic qubits. More often, the qubit is stored in some other state of a single photon (such as the polarization), or by being in one of two modes (a “dual-rail encoding”), or some more complicated structure possibly involving multiple photons. When we use one of these encodings, loss of a photon is no longer represented by an amplitude damping channel.

Consider, for instance, a polarization encoding, with horizontal polarization being $|0\rangle$ and vertical polarization being $|1\rangle$. Now, photon loss can affect either $|0\rangle$ or $|1\rangle$, and what results if the photon does escape is neither of those states, instead being a third state $|\text{vacuum}\rangle$. Therefore, a photon loss channel for polarization encoding takes a qubit input but its output is a 3-dimensional *qutrit*. Assuming both polarizations have equal probability p to lose a photon, the channel is then very simple. There are three Kraus operators, corresponding to no photon loss, loss of a horizontally polarized photon, and loss of a vertically polarized photon:

$$A_0 = \sqrt{1-p} \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \end{pmatrix}, A_1 = \sqrt{p} \begin{pmatrix} 0 & 0 \\ 0 & 0 \\ 1 & 0 \end{pmatrix}, \text{ and } A_2 = \sqrt{p} \begin{pmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 1 \end{pmatrix}. \quad (1.27)$$

The same channel works for photon loss from a dual-rail encoding.

When the encoding involves multiple photons, there is still a sensible description as a quantum channel, but the analysis requires a bit more quantum optics, and is beyond the scope of this book.

1.2.6 Erasure Errors

The example of the photon loss qubit channel suggests another interesting kind of error. When a photon is not lost, the state is not changed at all. When a photon *is* lost, by monitoring the photon number (but not the polarization), we can, in principle, tell that a photon has been lost. In that case, we do not know what the original state of the system was, but at least we know that something has happened to it. Measuring the photon number without destroying the polarization state is technologically difficult, so for this particular example, this is more an issue of principle than of practice. There are, however, other systems where monitoring for loss of the qubit is easier.

Definition 1.6. A *qubit erasure error* is an error A acting on a single qubit which maps all qubit states to a third state $|\perp\rangle$ orthogonal to the qubit Hilbert space.

Erasure channels may seem to be difficult channels to correct, since the information is completely lost, but that's not the case. By measuring if the state is $|\perp\rangle$ is present (without collapsing superpositions of $|0\rangle$ and $|1\rangle$) — for instance, by measuring the number of photons — we can determine if an erasure error has occurred. Because erasure channels provide some classical information about which qubits underwent errors, erasure channels are actually easier to correct than more general channels.

1.3 Multiple-Qubit Channels

It's hard to do too much with only one qubit. Even if we have only one qubit of data, we'll need more than one qubit to create a real quantum error-correcting code. Therefore, we should also think some about channels acting on multiple qubits.

1.3.1 Pauli Channels

One route to defining multiple-qubit channels is to largely forgot about the tensor product structure and just pick some completely positive map acting on the Hilbert space as a whole. Such channels can be extremely complicated, since they act on a Hilbert space of dimension 2^n when there are n qubits. Sometimes this is necessary, as it reflects the actual physics of the system, but quantum computers are usually built out of systems that naturally break up into qubits, and many-qubit interactions are rare.

Mathematically, it is usually too difficult to deal with an arbitrary n -qubit channel, but the n -qubit generalization of the Pauli channel is sufficiently well-behaved to be sometimes useful.

Definition 1.7. A channel of the form

$$\mathcal{E}(\rho) = \sum_P p_P P \rho P^\dagger, \tag{1.28}$$

where the sum is taken over P which are tensor products of I , X , Y , and Z , is a *Pauli channel*.

In a Pauli channel, the operator P occurs with probability p_P . Pauli channels are a reasonable quantum analogue of classical channels. There is a definite probability of error, and the errors that occur are large and discrete. However, since we can have a mix of bit flip and phase errors, a Pauli channel does have enough quantum features to be interesting.

1.3.2 Independent/Memoryless Channels

The simplest way to create a channel on n qubits is to just treat each qubit separately.

Definition 1.8. An *independent* channel on n qudits (each of dimension q) has the form $\otimes_{i=1}^n \mathcal{E}_i$, where each \mathcal{E}_i is a single-qudit channel.

A dimension- q qudit is a single system whose state space is a Hilbert space of dimension q , so for instance, $q = 2$ gives us a qubit, and for the moment, we will restrict attention to qubits. Often, we set all \mathcal{E}_i 's to be equal to \mathcal{E} , so that all qubits are treated equally.

As a simple example, we can consider $\mathcal{E} = \mathcal{R}_p$, the dephasing channel. Let us stick to $n = 3$ in order to be more explicit. There are a total of 8 possible Kraus operators for this channel, with the following probabilities:

Probability	Errors
$(1-p)^3$	$I \otimes I \otimes I$
$p(1-p)^2$	$Z \otimes I \otimes I, I \otimes Z \otimes I, I \otimes I \otimes Z$
$p^2(1-p)$	$Z \otimes Z \otimes I, Z \otimes I \otimes Z, I \otimes Z \otimes Z$
p^3	$Z \otimes Z \otimes Z$

The total probability of having a Z error on exactly one qubit is then $3p(1-p)^2$, and the probability of having two Z errors is $3p^2(1-p)$.

The chance of having at least one qubit with an error on it is therefore larger than the chance of having a single qubit by itself go wrong under the same dephasing channel \mathcal{R}_p . This makes sense, since there are more places for errors to occur. However, when p is small, most of the time, there will only be 0 or 1 Z error, and the other qubits will have I acting on them. This is the case we want to address through quantum error correction — errors are rare, but not negligibly so. In the case of this 3-qubit dephasing channel, we can make a good approximation by considering only the no-error or one-error possibilities, and ignoring the two- and three-qubit error cases.

1.3.3 Definition of t -Qubit Errors

More generally, we can define a t -qubit error to be one that acts on t qubits. However, there are two technical points in exactly how we want to define it:

Definition 1.9. We say that a linear operator A acting on n qubits has *weight* t if it is the tensor product of the identity on $n-t$ qubits with some matrix on the remaining t qubits. The weight of A is denoted $\text{wt } A$. We say that a weight t operator A has *support* on the t qubits on which it acts non-trivially. Generally (but not always), we cite the weight and support for the minimal set of qubits for which A acts non-trivially. A linear operator B acting on n qubits is a *t -qubit error* if it has the form

$$B = \sum_i B_i, \tag{1.29}$$

where each B_i has weight at most t , not necessarily with support on the same set of qubits for different i . An n -qubit quantum channel is a *t -qubit error channel* if it has a Kraus decomposition for which all Kraus operators are t -qubit errors. We can similarly define a *t -qubit error map* acting linearly on n -qubit density matrices but which may not be trace preserving.

Technical point number one is that even though a weight- t operator A acts non-trivially on only t qubits, it can act *arbitrarily* on those t qubits. In particular, it does not need to be a tensor product of errors on the separate qubits — it can entangle them however it likes. For instance, $\text{CNOT} \otimes I$ is a weight 2 operator acting on 3 qubits.

Technical point number two is that a t -qubit error B can be the sum of terms acting on different sets of t qubits. This naturally means that more than t qubits will be altered under the action of this error, but it turns out that they are altered in a way that is no more harmful than if we had a channel that had a possibility of altering each set of t qubits as separate Kraus operators. This point will be discussed in greater length in section 2.4.3, once we have a real quantum error-correcting code to examine.

We can also define t -qubit erasure errors. For instance, imagine we have many qubits each stored as the polarization of a photon, and each one undergoes a slight amount of photon loss.

Definition 1.10. A *t -qubit erasure error* is a weight t operator which is the tensor product of erasure errors on the qubits in its support. A *t -qubit erasure channel* is a quantum channel with a Kraus representation

where all Kraus operators are s -qubit erasure errors, with $s \leq t$. s can be different for different Kraus operators.

Note that here I am requiring that each Kraus operator has a specific set of qubits which can be erased, not a superposition of different sets of qubits. This reflects the idea that an erasure is in some sense a classical event, because the set of qubits that were erased can be measured.

1.3.4 Connection Between Independent Channel and t -Qubit Errors

An independent channel seems much more physically plausible (albeit still an approximation) than a t -qubit channel. However, it turns out to be more mathematically straightforward to design quantum error-correcting codes for t -qubit channels (or to correct a set of t -qubit errors) than for independent channels. Luckily, as suggested by the example of the 3-qubit dephasing channel, an n -qubit independent channel can be well approximated by a t -qubit channel (for some $t < n$) when the single-qubit channels making up the n -qubit channel are all close to the identity.

Theorem 1.1. *Let \mathcal{I} be the 1-qubit identity channel and $\mathcal{E} = \otimes_{i=1}^n \mathcal{E}_i$ be an n -qubit independent channel, with $\|\mathcal{E}_i - \mathcal{I}\|_\diamond < \epsilon \leq \frac{t+1}{n-t-1}$, and $\epsilon \leq 1/3$ as well. Then*

$$\|\mathcal{E} - \tilde{\mathcal{E}}\|_\diamond < 5 \binom{n}{t+1} [(4e+2)\epsilon]^{t+1} \quad (1.30)$$

for some t -qubit error map $\tilde{\mathcal{E}}$.

The significance of the theorem is that if we have a quantum error-correcting code which is designed to correct t -qubit error channels, it will automatically also correct independent channels where the single-qubit tensor factors are sufficiently close to the identity. Actually, that's not completely true: While the theorem does show that, this is not really the significance of the theorem, since there is much easier proof (which we will see in chapter 2) that a quantum error-correcting code that corrects t -qubit errors also corrects small independent error channels. Rather, this theorem provides a motivation for thinking about t -qubit error channels, which might otherwise seem rather bizarre.

When $\|\mathcal{E}_i - \mathcal{I}\|_\diamond < \epsilon$, we are getting about an ϵ chance of error per qubit, so with n qubits, we expect around ϵn errors. Therefore, we would not expect to be able to approximate \mathcal{E} well by a t -qubit error map unless $t \gtrsim \epsilon n$, which is reflected in the bound on ϵ . The other detailed constants in the theorem shouldn't be taken too seriously, as the proof could likely be tightened considerably to get better constants. But if you do insist on taking those constants seriously, you might find, for example, that when $n = 5$ and $t = 1$, you would get $\|\mathcal{E} - \tilde{\mathcal{E}}\|_\diamond \lesssim 8286\epsilon^2$, which only gives a non-trivial bound when ϵ is less than about 0.01.

Two elements of equation (1.30) that *are* significant are the combinatorial factor $\binom{n}{t+1}$, which reflects the number of $(t+1)$ -qubit subsets of the n qubits, and the exponent $t+1$ for ϵ , which tells us that the closeness of the approximation improves exponentially as we allow more qubits to have errors. In the limit of large n and any constant ratio t/n , the combination of the combinatorial factor and the exponent means that there will be a threshold value of ϵ below which we get a good approximation for all large n (and indeed, a better one as n gets larger).

Proof. We begin with a lemma on sums of error probabilities or amplitudes that will also be helpful later.

Lemma 1.2. *If $0 < t < n$, then*

- a) *For any $0 \leq \epsilon \leq 1$, $\sum_{j=t+1}^n \binom{n}{j} \epsilon^j (1-\epsilon)^{n-j} \leq \binom{n}{t+1} \epsilon^{t+1}$*
- b) *When $0 \leq \epsilon \leq \frac{t+1}{n-t-1}$, then $\sum_{j=t+1}^n \binom{n}{j} \epsilon^j \leq \binom{n}{t+1} (e\epsilon)^{t+1}$*

Proof of lemma. There are a number of ways to prove part a. One straightforward method is to interpret ϵ as a probability of some event (which is the main application we will have for this lemma). Then the sum is the probability that the event occurs at least $t+1$ times in n independent trials. We can upper bound

this probability by considering each subset of $t + 1$ trials. The total probability of having the event occur in all trials in the subset, without regard to what happens on the other $n - t - 1$ trials, is ϵ^{t+1} . There are $\binom{n}{t+1}$ subsets of size $t + 1$, so by the union bound, the total probability of having some set of $t + 1$ trials with the event is at most $\binom{n}{t+1}\epsilon^{t+1}$. Whenever the event occurs $j > t + 1$ times, we have over-counted, since we included that probability as part of all $\binom{j}{t+1}$ sets of size $t + 1$ which had the event.

For part b, note that

$$(1 - \epsilon)^{n-t-1} \geq \left(1 - \frac{t+1}{n-t-1}\right)^{n-t-1} \geq e^{-(t+1)}. \quad (1.31)$$

Then

$$\sum_{j=t+1}^n \binom{n}{j} \epsilon^j = \sum_{j=t+1}^n \binom{n}{j} \epsilon^j (1 - \epsilon)^{n-j} \frac{1}{(1 - \epsilon)^{n-j}} \quad (1.32)$$

$$\leq \sum_{j=t+1}^n \binom{n}{j} \epsilon^j (1 - \epsilon)^{n-j} \frac{1}{(1 - \epsilon)^{n-t-1}} \quad (1.33)$$

$$\leq \binom{n}{t+1} \epsilon^{t+1} e^{t+1} \quad (1.34)$$

by part a and equation (1.31). \square

As a warm-up to prove the theorem, let us consider the case when for all i , \mathcal{E}_i has a Kraus operator $(1 - \epsilon)I$. In this case, we can say that \mathcal{E}_i has probability ϵ of having an error (one of the other Kraus operators), and a probability of $1 - \epsilon$ of having no error. Part a of lemma 1.2 applies, so the probability of having at least $t + 1$ errors is at most $\binom{n}{t+1}\epsilon^{t+1}$. In this case, the t -qubit error map \mathcal{F} has all combinations of up to t “error” Kraus operators with the “good” Kraus operator $(1 - \epsilon)I$ on the other qubits. The map \mathcal{F} is completely positive, but is not trace preserving, since we have discarded the Kraus operators with more than t errors.

For the general case, first we need a better characterization of single-qubit channels \mathcal{G} which are close to the identity.

Lemma 1.3. *If \mathcal{E} is a quantum channel from \mathcal{H}_D to \mathcal{H}_D satisfying $\|\mathcal{E} - \mathcal{I}_1\|_\diamond < \epsilon \leq 1/3$, then \mathcal{E} has a Kraus representation $\mathcal{E}(\rho) = \sum_k A_k \rho A_k^\dagger$ such that $\|A_0 - I\|_\infty < \sqrt{2D}(\epsilon + \epsilon^2) + (\epsilon/2 + \epsilon^2)$ and $\sum_{k \neq 0} \|A_k\|_\infty^2 < D\epsilon(1/2 + \epsilon)$.*

Proof of lemma. Using the Choi-Jamiolkowski isomorphism, the channel \mathcal{E} corresponds to the entangled state $\Phi_\mathcal{E} = (I \otimes \mathcal{E})(|\Phi^+\rangle\langle\Phi^+|)$, with $|\Phi^+\rangle = \frac{1}{\sqrt{D}} \sum_a |aa\rangle$. (Note that I have normalized the state to make it easier to apply common identities, which is not always the convention used with the Choi-Jamiolkowski isomorphism.) Since $\|\mathcal{E} - \mathcal{I}_1\|_\diamond < \epsilon$,

$$\|\Phi_\mathcal{E} - |\Phi^+\rangle\langle\Phi^+|\|_1 < \epsilon \quad (1.35)$$

as well. Since the trace distance between these two states is small, the fidelity between them is high:

$$\langle\Phi^+|\Phi_\mathcal{E}|\Phi^+\rangle > 1 - \epsilon/2. \quad (1.36)$$

Now, $|\Phi^+\rangle\langle\Phi^+|$ has one eigenvalue $+1$ and the remaining eigenvalues 0 . Let us also write $\Phi_\mathcal{E}$ in terms of an eigenbasis,

$$\Phi_\mathcal{E} = \sum_{i=0}^{D^2-1} \lambda_i |\phi_i\rangle\langle\phi_i|. \quad (1.37)$$

$\Phi_\mathcal{E}$ is positive and has trace 1, so $\lambda_i \geq 0$ and $\sum_i \lambda_i = 1$. Now consider

$$|\langle\phi_i|\Phi^+\rangle\langle\Phi^+|\phi_j\rangle| = |\langle\phi_i|(\Phi_\mathcal{E} - |\Phi^+\rangle\langle\Phi^+|)|\phi_j\rangle| < \epsilon \quad (1.38)$$

for $i \neq j$. In addition,

$$\langle \Phi^+ | \Phi_{\mathcal{E}} | \Phi^+ \rangle = \sum_{i=0}^{D^2-1} \lambda_i |\langle \Phi^+ | \phi_i \rangle|^2 > 1 - \epsilon/2. \quad (1.39)$$

We can choose the phase of the eigenstates $|\phi_i\rangle$ to ensure that $\langle \Phi^+ | \phi_i \rangle$ is real and non-negative. Letting $a_i = \langle \Phi^+ | \phi_i \rangle$, we have $a_i \geq 0$, $\sum \lambda_i a_i^2 > 1 - \epsilon/2$, and $a_i a_j < \epsilon$ for $i \neq j$. Assume without loss of generality that a_0 is the largest of the a_i s. Then $\sum \lambda_i a_i^2 \leq \sum \lambda_i a_0^2 = a_0^2$, so

$$a_0 > \sqrt{1 - \epsilon/2} \geq 1 - \epsilon/2. \quad (1.40)$$

Thus

$$a_i < \epsilon/a_0 < \frac{\epsilon}{1 - \epsilon/2} \quad (1.41)$$

for $i \neq 0$. Then

$$\sum_{i=0}^{D^2-1} \lambda_i a_i^2 \leq \lambda_0 + \sum_{i=1}^{D^2-1} \lambda_i \frac{\epsilon}{1 - \epsilon/2} \quad (1.42)$$

$$= \lambda_0 + (1 - \lambda_0) \frac{\epsilon}{1 - \epsilon/2} \quad (1.43)$$

$$= \frac{\epsilon + (1 - 3\epsilon/2)\lambda_0}{1 - \epsilon/2}, \quad (1.44)$$

since $\sum \lambda_i = 1$. Therefore,

$$\lambda_0 \geq \frac{(1 - \epsilon/2)^2 - \epsilon}{1 - 3\epsilon/2} \geq 1 - \epsilon/2 - \epsilon^2. \quad (1.45)$$

(assuming $\epsilon \leq 1/3$), which means

$$\sum_{i \neq 0} \lambda_i = 1 - \lambda_0 \leq \epsilon/2 + \epsilon^2 \quad (1.46)$$

for $i \neq 0$.

We have now bounded all the terms we need, but we'd like a tighter bound on a_0 . We can do that by going back and plugging in the bound on $\sum \lambda_i$.

$$\| |\phi_0\rangle \langle \phi_0| - |\Phi^+\rangle \langle \Phi^+| \|_1 \leq (1 - \lambda_0) + \|\lambda_0 |\phi_0\rangle \langle \phi_0| - |\Phi^+\rangle \langle \Phi^+|\|_1 \quad (1.47)$$

$$\leq (1 - \lambda_0) + \left\| \sum_{i \neq 0} \lambda_i |\phi_i\rangle \langle \phi_i| \right\|_1 + \|\Phi_{\mathcal{E}} - |\Phi^+\rangle \langle \Phi^+|\|_1 \quad (1.48)$$

$$\leq 2\epsilon + 2\epsilon^2. \quad (1.49)$$

But the 1-norm distance between two pure states is just given by

$$\| |\phi_0\rangle \langle \phi_0| - |\Phi^+\rangle \langle \Phi^+| \|_1 = 2\sqrt{1 - |\langle \phi_0 | \Phi^+ \rangle|^2}, \quad (1.50)$$

meaning

$$\sqrt{1 - |\langle \phi_0 | \Phi^+ \rangle|^2} \leq \frac{\epsilon + \epsilon^2}{\sqrt{1 + |\langle \phi_0 | \Phi^+ \rangle|}} \leq \epsilon + \epsilon^2. \quad (1.51)$$

Thus, in the Choi-Jamiołkowski isomorphism form of the channel, the state has one large eigenvalue, for which the eigenstate is close to the maximally-entangled state, and the other eigenvalues are all small, with the eigenstates far from the maximally entangled state $|\Phi^+\rangle$. (Of course, they could be close to *other* maximally entangled states.) We can recover a set of Kraus operators for the channel by letting

$$A_k |a\rangle = \sqrt{D\lambda_k} (\langle a | \otimes I) |\phi_k\rangle \quad (1.52)$$

for basis states $|a\rangle$, extended linearly to the full Hilbert space \mathcal{H}_D . (Recall that $|\phi_i\rangle$ is a state in $\mathcal{H}_D \otimes \mathcal{H}_D$, so the right-hand side of equation (1.52) is in \mathcal{H}_D .) Then $\|A_k\|_1^2 \leq D\lambda_k$, so $\sum_{k \neq 0} \|A_k\|_\infty^2 \leq D\epsilon(1/2 + \epsilon)$, as desired.

To get the bound on A_0 , note that $A_0|\psi\rangle = \sqrt{D\lambda_0}\langle\psi^*|\phi_0\rangle$, where $|\psi^*\rangle$ is the complex conjugate of $|\psi\rangle$ in the basis $\{|a\rangle\}$. Furthermore, $|\psi\rangle = \sqrt{D}\langle\psi^*|\Phi^+\rangle$. Then

$$\|A_0 - \lambda_0 I\|_\infty = \max_{|\psi\rangle} |A_0|\psi\rangle - \lambda_0|\psi\rangle| \quad (1.53)$$

$$= \max_{|\psi\rangle} \sqrt{D\lambda_0} |\langle\psi^*|\phi_0\rangle - \langle\psi^*|\Phi^+\rangle| \quad (1.54)$$

$$= \sqrt{D\lambda_0} \left| \|\phi_0\rangle - |\Phi^+\rangle \right| \quad (1.55)$$

$$= \sqrt{D\lambda_0} \sqrt{2 - 2\operatorname{Re}\langle\Phi_+|\phi_0\rangle} \quad (1.56)$$

$$\leq \sqrt{2D}(\epsilon + \epsilon^2), \quad (1.57)$$

applying equation (1.51) in the last line, and recalling we have chosen $\langle\Phi_+|\phi_0\rangle$ to be real. Thus,

$$\|A_0 - I\|_\infty \leq \sqrt{2D}(\epsilon + \epsilon^2) + (\epsilon/2 + \epsilon^2) \quad (1.58)$$

□

For the case of qubits and applying $\epsilon \leq 1/3$, the lemma gives $\|A_0 - I\|_\infty \leq 7\epsilon/2$ and $\sum_{k \neq 0} \|A_k\|_1^2 \leq 5\epsilon/3$ for $k \neq 0$. We will round to $\|A_0 - I\|_1 < 4\epsilon$ and $\sum_{k \neq 0} \|A_k\|_1^2 \leq 2\epsilon$ for $k \neq 0$.

Given lemma 1.3, we can use a similar approach for the general case as we did for the warm-up, which was essentially classical. Channel \mathcal{E}_i has Kraus operators A_k^i , with A_0^i close to the identity and A_k^i small for $k \neq 0$. The n -qubit independent channel \mathcal{E} has Kraus operators which are all possible tensor products $\otimes_i A_{k_i}^i$. Let \mathcal{F} be the map whose Kraus operators are all tensor products $\otimes_i A_{k_i}^i$ with at most t values of i for which $k_i \neq 0$. Then

$$\|\mathcal{F} - \mathcal{E}\|_\diamond \leq \sum_{r=t+1}^n \sum_{|S|=r} \prod_{i \in S} \sum_{k_i \neq 0} \|A_{k_i}^i\|_\infty^2. \quad (1.59)$$

The sum over r represents the number of values of i for which $k_i \neq 0$, and S is the set of indices for which $k_i \neq 0$. Since $\sum_{k_i \neq 0} \|A_{k_i}^i\|_\infty^2 \leq 2\epsilon$, we get

$$\|\mathcal{F} - \mathcal{E}\|_\diamond \leq \sum_{r=t+1}^n \binom{n}{r} (2\epsilon)^r \leq \binom{n}{t+1} (2\epsilon)^{t+1} \quad (1.60)$$

by lemma 1.2.

Now \mathcal{F} is not yet a t -qubit error map because the A_0 terms include some errors. However, the A_0 terms are all near I , so we can expand them as $A_0^i = I + \delta A^i$, with $\|\delta A^i\|_\infty < 4\epsilon$. Given a Kraus operator for \mathcal{F} $A_{\{k_i\}} = \otimes_i A_{k_i}^i$ with r indices $k_i \neq 0$, expand all A_0^i s in this way and form $A'_{\{k_i\}}$ by discarding all terms in the expansion with more than $t - r$ δA^i factors. The resulting Kraus operators are composed of sums of terms with weight at most t , of which r non-identity factors come from $k_i \neq 0$ terms, and the remaining come from δA^i components of $k_i = 0$ factors.

$$\|A'_{\{k_i\}} - A_{\{k_i\}}\|_\infty \leq \left(\prod_{j|k_j \neq 0} \|A_{k_j}\|_\infty \right) \sum_{s=t+1-r}^{n-r} \sum_{|S|=s} \prod_{i \in S} \|\delta A^i\|_\infty \quad (1.61)$$

$$\leq \left(\prod_{j|k_j \neq 0} \|A_{k_j}\|_\infty \right) \sum_{s=t+1-r}^{n-r} \binom{n-r}{s} (4\epsilon)^s \quad (1.62)$$

$$\leq \binom{n-r}{t+1-r} (4\epsilon)^{t+1-r} \left(\prod_{j|k_j \neq 0} \|A_{k_j}\|_\infty \right). \quad (1.63)$$

Let ρ be an arbitrary pure state, possibly entangled between \mathcal{H}_D and a reference system. Then

$$\|A'_{\{k_i\}}\rho(A'_{\{k_i\}})^\dagger - A_{\{k_i\}}\rho A_{\{k_i\}}^\dagger\|_1 \leq \|(A'_{\{k_i\}} - A_{\{k_i\}})\rho(A'_{\{k_i\}})^\dagger\|_1 + \|A_{\{k_i\}}\rho((A'_{\{k_i\}})^\dagger - A_{\{k_i\}}^\dagger)\|_1 \quad (1.64)$$

$$\leq 2 \left(\prod_{j|k_j \neq 0} \|A_{k_j}\|_\infty \right) \|A'_{\{k_i\}} - A_{\{k_i\}}\|_\infty \quad (1.65)$$

$$\leq 2 \binom{n-r}{t+1-r} (4e\epsilon)^{t+1-r} \left(\prod_{j|k_j \neq 0} \|A_{k_j}\|_\infty^2 \right). \quad (1.66)$$

I have omitted the $\otimes I$ terms affecting the reference system in the first line; the equation is complicated enough as is. In the second line, we have used the property that $\|A|\psi\rangle\| \leq \|A\|_\infty$ and bounded $\|A_{\{k_i\}}\|_\infty$ and $\|A'_{\{k_i\}}\|_\infty$ by the norm of just the terms with $k_j \neq 0$. If we sum over all $\{k_i\}$ with the same locations for which $k_j \neq 0$, we get

$$\sum \|A'_{\{k_i\}}\rho(A'_{\{k_i\}})^\dagger - A_{\{k_i\}}\rho A_{\{k_i\}}^\dagger\|_\infty \leq 2 \binom{n-r}{t+1-r} (4e\epsilon)^{t+1-r} (2\epsilon)^r. \quad (1.67)$$

Finally, let \mathcal{G} be the linear map with Kraus operators $A'_{\{k_i\}}$. In \mathcal{G} , we have eliminated all Kraus operators with more than t errors. Then we get a bound on $\|\mathcal{G} - \mathcal{F}\|_\diamond$ by applying them to ρ , and considering the 1-norm of the resulting state, which involves summing equation (1.66) over all possible values of $\{k_i\}$ with at most t indices $k_j \neq 0$ (those with more have already been excluded from \mathcal{F}). We get

$$\|\mathcal{G} - \mathcal{F}\|_\diamond \leq 2 \sum_{r=0}^t \binom{n}{r} \binom{n-r}{t+1-r} (4e\epsilon)^{t+1-r} (2\epsilon)^r \quad (1.68)$$

$$= 2 \sum_{r=0}^t \binom{n}{t+1} \binom{t+1}{r} (4e\epsilon)^{t+1-r} (2\epsilon)^r \quad (1.69)$$

$$\leq 2 \binom{n}{t+1} [(4e+2)\epsilon]^{t+1}. \quad (1.70)$$

Combining this with equation (1.60), we get

$$\|\mathcal{G} - \mathcal{E}\|_\diamond \leq 3 \binom{n}{t+1} [(4e+2)\epsilon]^{t+1}. \quad (1.71)$$

There is one final step needed. It is possible that \mathcal{G} is no longer completely positive, since it could be that $\|A'_{\{k_i\}}\rho(A'_{\{k_i\}})^\dagger\|_1 > \|A_{\{k_i\}}\rho(A_{\{k_i\}})^\dagger\|_1$, which could result in $\text{tr } \mathcal{G}(\rho) > 1$. We should therefore scale \mathcal{G} down to $C\mathcal{G}$, for appropriately chosen constant C , to get a map that is guaranteed to be completely positive. \mathcal{F} is trace non-increasing, though, so

$$\text{tr } \mathcal{G}(\rho) \leq 1 + \|\mathcal{G} - \mathcal{F}\|_\diamond \leq 1 + 2 \binom{n}{t+1} [(4e+2)\epsilon]^{t+1} = 1/C. \quad (1.72)$$

Then

$$\|C\mathcal{G} - \mathcal{E}\|_\diamond \leq 3 \binom{n}{t+1} [(4e+2)\epsilon]^{t+1} + 1 - C \quad (1.73)$$

$$\leq 5 \binom{n}{t+1} [(4e+2)\epsilon]^{t+1}. \quad (1.74)$$

□

1.4 A Peek Ahead: Errors During Computation

In part II, we'll consider more general types of errors. In particular, quantum gates will be able to go wrong in various ways, and errors will occur multiple times during a computation. As noted above, the quantum channel picture is no longer completely general then, since the noise might be non-Markovian. Mostly we'll stick to Markovian noise, but even then matters are much more complicated. Since our control is no longer reliable, we'll have to deal with errors occurring even while we're trying to fix them. However, there are various subtler difficulties to contend with as well.

For one thing, we don't know when an error occurs, so we can't assume we do error correction immediately after each error. In particular, an error might occur right before a gate that we had intended for some other purpose. Then even if the gate itself works perfectly, it can cause the error to propagate, infecting an additional qubit with the same error. In addition, the effect of the gate can change the type of error. For instance, a Z error that occurs before a Hadamard gate will change into a X error after the gate. This phenomenon makes it much more difficult to take advantage of information we know about the errors. For instance, suppose the noise source is largely dephasing noise. We might want to use a code that is particularly good at correcting dephasing noise, but if we use Hadamard gates in our circuit, some of the Z errors that occur will have become X errors by the time we get around to correcting them, and our code won't work anywhere near as well as expected. A particularly insidious form of this phenomenon occurs when errors happen *during* the implementation of a gate, which, after all, should take a non-zero time. Even if the completed gate does not change the type of error, depending on how the gate is being implemented, the partial gate may in fact alter the structure of the noise.

We'll return to all these issues in part II, and discuss how to design fault-tolerant quantum circuits that allow reliable error correction and computation on encoded states despite the complications.

Chapter 2

Redundancy Without Repetition: Basics Of Quantum Error Correction

Now we are ready to start designing quantum error-correcting codes. A natural place to start for inspiration is to look at the theory of classical error-correcting codes, and indeed it can give us some guidance. However, we'll rapidly see that there are some major differences between classical and quantum error correction.

The simplest classical error-correcting code is the repetition code:

$$0 \mapsto 000 \tag{2.1}$$

$$1 \mapsto 111. \tag{2.2}$$

If we send this 3-bit encoding through a 1-bit error classical channel, it is clear that we will be able to correct for any error that occurs on a single bit. If all three bits are the same, we know there hasn't been an error, while if one of the three is different, for instance 010, we know that the one that's different is the one that's wrong. By enforcing a boring conformity among the bits, we can recover the original state. Recall that if we use an independent channel instead of a 1-bit error channel, then there is some chance of 2 or 3 errors, which would fool us. If there are two errors, the one bit that we think is wrong is actually the only bit that's correct, and our well-meaning attempt to fix it will actually complete the error, making all three bits wrong. Luckily, the chance of two errors occurring is only $O(p^2)$ when the probability of an error on a single bit is p (see section 1.3.2). When p is small, the chance of the encoded state ending up wrong is less than the chance that an unencoded bit can make it unchanged through the channel.

Throughout this chapter, we'll assume we have a t -qubit error channel. This is justified by theorem 1.1, which tells us that if we actually have an independent channel, it is very close to a t -qubit channel, at least when the error rate per qubit is low. We will actually reprove a version of theorem 1.1 with an easier proof and better constants specifically applicable to quantum error-correcting codes.

To make a quantum error-correcting code, we might want to imitate the classical repetition code, but that instantly runs into a few problems. We'll need to find a somewhat different way to protect our quantum information, one that adds redundancy without repeating the state.

2.1 Quantum Error Correction? Ridiculous!

Why am I so dead-set against a quantum repetition code? Perhaps we could encode

$$|\psi\rangle \mapsto |\psi\rangle|\psi\rangle|\psi\rangle. \tag{2.3}$$

If you've had much quantum information experience, you'll immediately see the problem with this encoding: it is forbidden by the No-Cloning Theorem.

1. No-cloning theorem prohibits repeating a quantum state.
2. Measuring the data while determining the error will destroy superpositions.
3. We must correct Y and Z errors in addition to X errors.
4. We must correct an infinite set of unitaries and also channels which decohere the state.

Table 2.1: Barriers to making a quantum error-correcting code.

Theorem 2.1 (No-Cloning Theorem). *There is no quantum operation which maps an arbitrary state $|\psi\rangle$ to $|\psi\rangle|\psi\rangle$.*

Proof. The problem is linearity. Suppose

$$|0\rangle \mapsto |00\rangle \tag{2.4}$$

$$|1\rangle \mapsto |11\rangle. \tag{2.5}$$

Then, because quantum operations must be linear,

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle \mapsto \alpha|00\rangle + \beta|11\rangle. \tag{2.6}$$

However, this is not equal to the cloned $|\psi\rangle$ state:

$$|\psi\rangle|\psi\rangle = \alpha^2|00\rangle + \beta^2|11\rangle + \alpha\beta(|01\rangle + |10\rangle). \tag{2.7}$$

□

Another problem has to do with how we correct errors for the repetition code. Given a 3-bit state coming out of the channel, we'd like to look at it and determine which of the three bits is different from the other two. However, when dealing with quantum states, looks really can kill. Or if not kill, at least decohere, destroying any superposition we have. And without the ability to be in a superposition, a qubit is no better than a classical bit. Somehow, to make a quantum error-correcting code, we'll need some way to correct errors without looking too closely at what we're doing.

Even if we can handle these problems, it's clear that quantum error correction will be more complicated than classical error correction. For a classical channel, basically all that can happen is a bit flip error, but quantum codes are going to need to handle phase flip errors as well, not to mention Y errors. While the 3-qubit repetition code is good at correcting bit flip errors, it doesn't do anything to help correct phase flips. Actually, it makes phase flip errors more common since there are now three qubits which could have phase flip errors instead of only one.

In addition to X , Y , and Z , we will also have to handle an infinite set of unitaries, such as the R_θ errors. Then there are more general channels, such as the dephasing channel or depolarizing channel, which turn pure states into mixed states. How can we come up with a correction operation that will turn the state from a mixed state back into a pure state?

At this point, the prospects for a quantum error-correcting code look bleak. The difficulties we've identified so far are listed in table 2.1. But don't worry — there are solutions to all of these problems. If there weren't, this book would be a lot thinner.

2.2 The 3-Qubit Code(s)

We won't try to handle all of these problems at once. Let's focus first on points 1 and 2. We'll try to make a quantum version of the three-bit repetition code. The quantum code will encode a qubit, but only protect against one kind of error, just like the classical three-bit code.

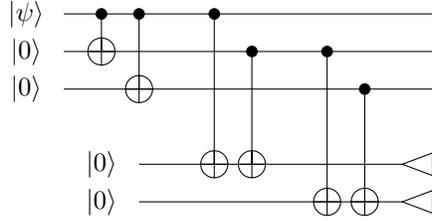


Figure 2.1: Encoding and syndrome measurement circuit for the 3-qubit bit flip correction code.

2.2.1 Correcting Bit Flips for Superposition States

Suppose we apply the classical repetition encoding to the basis states and extend to arbitrary superpositions by linearity:

$$\alpha|0\rangle + \beta|1\rangle \mapsto \alpha|000\rangle + \beta|111\rangle. \quad (2.8)$$

The first thing to notice is that this is an allowed encoding, and doesn't violate the no-cloning theorem. There is no rule against copying basis states — they act just like classical bits. The no-cloning theorem only kicks in if you try to copy superpositions as well. Rather than encoding a superposition $|\psi\rangle$ as three copies of $|\psi\rangle$ (which is impossible), we encode superpositions as *entangled* states. That takes care of difficulty 1.

Let's look at what happens to this state when there's been a bit flip error on the second qubit:

$$X_2(\alpha|000\rangle + \beta|111\rangle) = \alpha|010\rangle + \beta|101\rangle. \quad (2.9)$$

(In this book, I shall use the notation X_2 to indicate that the error X is acting on qubit 2. The same notation applies for gates. However, sometimes I shall omit the subscripts and instead write out a tensor product explicitly; in this case, that would be $I \otimes X \otimes I$.)

As with the classical repetition code, the error can be identified by noticing that the middle qubit is different from the first and third qubits. The critical point is that *this is true for both branches of the superposition*. It is therefore possible to measure the fact that the second qubit is different without measuring whether we have an encoded zero or an encoded one. If we don't measure the data that's encoded in the code, we don't destroy an encoded superposition. This is how we resolve the second barrier: we can measure the error without measuring the information we're trying to preserve.

2.2.2 Encoding Circuit and Error Syndrome Measurement

To understand this point in more detail, look at figure 2.1. The first part of the figure gives the encoding circuit for the 3-qubit code. We start with one qubit in an arbitrary unencoded state and add two additional qubits which become entangled with the first qubit. After the encoding, it is no longer particularly meaningful to ask which was the original data qubit and which are the ones we added. They are now all treated on an equal basis. The encoding of eqnrefeqn:threequbitbitflip is symmetric between all three qubits.

The second part of the circuit contains the error correction process. We wish to know whether one of the qubits is different from the other two, and if so, which one is different. We can break that down into two pieces: We ask whether the first two qubits are the same or different, and then we ask whether the second and third qubits are the same or different. In other words, we wish to know the *parity* of the first two qubits and the parity of the last two qubits. We will store the answers to these two questions on two more extra qubits. Extra qubits like these that are used for this or any other helpful purpose during a computation are known as *ancilla* qubits.

To tell whether two qubits are the same or different, the circuit in figure 2.1 uses two CNOT gates. If they are the same, either neither CNOT flips the corresponding ancilla, or both do; if they are different, exactly one of the CNOT gates flips the ancilla qubit. Thus, when we measure the ancilla qubit for one of the parity measurements, the output we get is equal to the parity: 0 for same (even parity), and 1 for

opposite (odd parity). Notice that the result does not depend at all on the encoded state, simply whether there is a bit flip error on the qubits we are measuring. That means the measurement can be done without disturbing a superposition of the encoded data.

Together, the measurement results for the two ancillas form a bit string known as the *error syndrome*. The error syndrome encapsulates all the information we have about the error. For instance, when the bit flip error is X_2 , the error syndrome is 11 because both the first pair and the second pair of qubits are different. Similarly, error syndromes 10 and 01 correspond to the errors X_1 and X_3 , respectively, and error syndrome 00 tells us that there is no error. Thus, every possible error for a one-qubit bit flip channel is accounted for. Once we know what the error is, we can simply correct it by performing another bit flip on the appropriate qubit.

Notice that the choice of error correction circuit and corresponding error syndrome is not unique. For instance, we could have measured the parity of the first and third qubits instead of measuring the parity of the second and third qubits. This would have given us the same information, but the correspondence between error syndromes and errors would have been shuffled. We could have measured the parities of all three pairs of qubits, but in that case the error syndrome (which would be 3 bits long) would be redundant: from any two error syndrome bits, we could deduce the third. It is not a coincidence that the number of syndrome bits we need is equal to the number of extra qubits we added to the data in order to perform the original encoding. This is a general property of quantum error-correcting codes, as discussed in chapter 3, although there are certain cases where we don't need all of the information encoded in the syndrome. (See section 17.2 for a description of that case.)

2.2.3 Phase Error Correction

The three-qubit code described above corrects a single bit flip error, but cannot correct any phase flip errors. It is also straightforward to make a code that works the other way: it corrects a single phase flip error, but cannot correct any bit flip errors.

The key to making this code is to notice that the Hadamard transform H switches the role of X and Z errors:

$$|0\rangle \leftrightarrow |+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) \quad (2.10)$$

$$|1\rangle \leftrightarrow |-\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle) \quad (2.11)$$

X acting on the $|0\rangle, |1\rangle$ basis is a bit flip, but in the $|+\rangle, |-\rangle$ basis, it maps

$$X|+\rangle = |+\rangle \quad (2.12)$$

$$X|-\rangle = -|-\rangle, \quad (2.13)$$

acting as a phase flip. Similarly, Z switches $|+\rangle$ and $|-\rangle$, so it acts as a bit flip in the Hadamard-rotated basis.

Therefore, we can make a three-qubit *phase* correcting code by just applying H to every qubit in the three-qubit bit flip correcting code:

$$|0\rangle \mapsto |\bar{0}\rangle = |+\rangle|+\rangle|+\rangle \quad (2.14)$$

$$|1\rangle \mapsto |\bar{1}\rangle = |-\rangle|-\rangle|-\rangle, \quad (2.15)$$

extended by linearity so that $\alpha|0\rangle + \beta|1\rangle \mapsto \alpha|\bar{0}\rangle + \beta|\bar{1}\rangle$. In this book, I shall use lines over a state or operator to indicate the encoded version of the object.

If there is a single phase flip error, for instance, Z_2 , one of the three qubits will be different than the other two in the Hadamard basis. E.g., $Z_2|\bar{0}\rangle = |+\rangle|-\rangle|+\rangle$. We can locate the error in just the same way as before, except that now we should rotate the control qubits for the CNOTs in the error correction circuit by a Hadamard transform to work in the correct basis. The modified encoding and error correction circuit is pictured in figure 2.2.

A little consideration will show you that this argument applies to Y errors as well. The X and Z error correction procedures are essentially independent, and the argument that we can identify the block of three qubits with a Z error applies equally well if there is an X error on one of the nine qubits. Thus, the nine-qubit code can correct for both an X and a Z error. In particular, it can correct for a single-qubit Y error, which corresponds to having an X and a Z on the same qubit. In the set of possible errors, we only allowed single-qubit errors, but actually the code will still work if you extend the set of possible errors to allow one X error and one Z error on *any* pair of qubits.

2.4 Correcting General Errors

The nine-qubit code addresses problem 3, but it suggests that problem 4 might be troublesome. In order to correct a single bit-flip error, we needed 3 qubits, but to correct Y and Z errors as well, we went up to nine qubits. To solve problem 4 and correct general single-qubit errors, we'll need to handle an infinite set of errors. Does this mean we will have to use infinitely many qubits? Hopefully not.

2.4.1 Continuous Phase Rotation Example

Let's examine the specific case of the continuous phase rotation R_θ in more detail. We can rewrite

$$R_\theta = \begin{pmatrix} e^{-i\theta} & 0 \\ 0 & e^{i\theta} \end{pmatrix} = \cos \theta I - i \sin \theta Z. \quad (2.20)$$

Now suppose the usual nine-qubit code has experienced an error $(R_\theta)_i$ instead of the usual X , Y , or Z error, but imagine that we do not know that and use the usual error correction circuit for the nine-qubit code. We can figure out what happens to state by again applying the linearity of quantum mechanics.

$$(R_\theta)_i |\bar{\psi}\rangle = \cos \theta I |\bar{\psi}\rangle - i \sin \theta Z_i |\bar{\psi}\rangle. \quad (2.21)$$

We know what happens to $I|\bar{\psi}\rangle$ and $Z_i|\bar{\psi}\rangle$ under the error correction circuit: First, we interact with some ancilla qubits and determine the error syndrome corresponding to these errors, then we measure the ancilla qubits. If $|I\rangle_{\text{syn}}$ and $|Z_i\rangle_{\text{syn}}$ are the error syndromes corresponding to the errors I and Z_i , when we interact with the ancillas, we get the following:

$$(R_\theta)_i |\bar{\psi}\rangle \mapsto \cos \theta I |\bar{\psi}\rangle |I\rangle_{\text{syn}} - i \sin \theta Z_i |\bar{\psi}\rangle |Z_i\rangle_{\text{syn}}. \quad (2.22)$$

Right now, we have an entangled state between the code qubits and the ancilla qubits. When we measure the ancilla register, we'll collapse the superposition, getting one of two results:

$$\text{Prob.}(\cos^2 \theta) : I |\bar{\psi}\rangle |I\rangle_{\text{syn}} \quad (2.23)$$

$$\text{Prob.}(\sin^2 \theta) : Z_i |\bar{\psi}\rangle |Z_i\rangle_{\text{syn}}. \quad (2.24)$$

Something miraculous has happened: now we have either an I error (i.e., no error at all) or a Z error, and the outcome of the error syndrome tells us which and where the error is. We can then correct it in the usual way. Somehow, by pretending that we had an I , X , Y , or Z error, we made it actually true that the error was one of those. It's a triumph of wishful thinking.

For the nine-qubit code, the same procedure works with arbitrary single-qubit unitaries, and actually works for any Kraus operator A_k . Indeed, this is not a property of just the nine-qubit code, but of quantum error-correcting codes in general: correcting I , X , Y , and Z errors is sufficient to correct general errors. At this point, we could easily prove this for the nine-qubit code, but it is worthwhile to prove the generalization. To do that, we will want to first define precisely what a quantum error-correcting code is.

2.4.2 Definition of a Quantum Error Correcting Code

Generally speaking, a quantum error-correcting code is a subspace of a larger Hilbert space. Typically, that subspace has been chosen in some complicated entangled way in order to have special properties when errors occur on states from the code space. Therefore, it is most sensible to define a quantum error-correcting code together with the set of errors it corrects. Also, we'll frequently want to consider the subspace as a Hilbert space encoding some quantum data, so we'll define the subspace as a map from a smaller Hilbert space into the big Hilbert space.

Definition 2.1. A *quantum error-correcting code* (U, \mathcal{E}) (or *QECC* for short) is a partial isometry $U : \mathcal{H}_K \rightarrow \mathcal{H}_N$ with the set of *correctable errors* \mathcal{E} (consisting of linear maps $E : \mathcal{H}_N \rightarrow \mathcal{H}_M$) with the following property: \exists quantum operation $\mathcal{D} : \mathcal{H}_M \rightarrow \mathcal{H}_K$ such that $\forall E \in \mathcal{E}, \forall |\psi\rangle \in \mathcal{H}_K$,

$$\mathcal{D}(EU|\psi\rangle\langle\psi|U^\dagger E^\dagger) = c(E, |\psi\rangle)|\psi\rangle\langle\psi|. \quad (2.25)$$

U is known as the *encoding* operation for the code (or the *encoder*), and \mathcal{D} is the *decoding* map (or *decoder*). \mathcal{D} is also known as the *recovery* operation. We say that the QECC *corrects* E if $E \in \mathcal{E}$.

Sometimes, we don't consider a specific encoding map, and refer to $\text{Image}(U)$ (more precisely known as the *code space*) as the QECC. A *codeword* is any state in the code space. Often, a QECC is mentioned without an explicit set of correctable errors, which should be determined by context. Sometimes the error set simply does not matter. When it does matter but is not specified, the set of correctable errors is often the set of all t -qubit errors for the largest possible t ; except sometimes it is instead the set of the most likely errors in the system. \mathcal{H}_K is sometimes referred to as the *logical* Hilbert space and \mathcal{H}_N is the *physical* Hilbert space. If \mathcal{E} is the set of all t -qubit errors, we say that U (or $\text{Image}(U)$) is a *t -error correcting code*, or that it *corrects t errors*.

In the above definition (and elsewhere in the book), \mathcal{H}_D is a Hilbert space of dimension D . A partial isometry is like a unitary in that it preserves inner products, but might not be invertible because it may map a Hilbert space to a Hilbert space of strictly larger dimension. Of course \mathcal{D} actually maps matrices on \mathcal{H}_M to matrices on \mathcal{H}_K . Since E can be a general linear map, not necessarily unitary, $EU|\psi\rangle$ might not be normalized properly, which is why we need the $c(E, |\psi\rangle)$ factor in equation (2.25). The definition explicitly allows $c(E, |\psi\rangle)$ to depend on $|\psi\rangle$, but it turns out that it does not (see section 2.5). It does depend on E , however.

To prove this, and in many other contexts, it is helpful to treat the decoding map as a unitary. This can be done via the Stinespring dilation, by purifying \mathcal{D} . To do so, we must add an ancilla register. Let the input ancilla to \mathcal{D} be of dimension D and let the output ancilla have dimension D' . Then $\mathcal{H}_M \otimes \mathcal{H}_D \cong \mathcal{H}_K \otimes \mathcal{H}_{D'}$. $\mathcal{H}_{D'}$ plays the role of the error syndrome. Call the purification $V : \mathcal{H}_M \otimes \mathcal{H}_D \rightarrow \mathcal{H}_K \otimes \mathcal{H}_{D'}$. Now, $U|\psi\rangle = |\bar{\psi}\rangle$, and when the QECC corrects $E \in \mathcal{E}$,

$$V(E|\bar{\psi}\rangle_N \otimes |0\rangle_D) = \sqrt{c(E, |\psi\rangle)}|\psi\rangle_K \otimes |A(E, |\psi\rangle)\rangle_{D'} \quad (2.26)$$

I have put subscripts on the kets to help you keep track of which Hilbert spaces they belong to. There might be a phase on the RHS, but it can be absorbed into the ancilla state $|A(E, |\psi\rangle)\rangle$.

Proposition 2.2. *In definition 2.1, $c(E, |\psi\rangle)$ is independent of $|\psi\rangle$, as is $|A(E, |\psi\rangle)\rangle$ when we purify the decoder.*

Proof. Consider two codewords $|\bar{\psi}\rangle$ and $|\bar{\phi}\rangle$. Imagine we purify the decoder to the unitary V , so

$$VE|\bar{\psi}\rangle \otimes |0\rangle = \sqrt{c(E, |\psi\rangle)}|\psi\rangle \otimes |A(E, |\psi\rangle)\rangle \quad (2.27)$$

$$VE|\bar{\phi}\rangle \otimes |0\rangle = \sqrt{c(E, |\phi\rangle)}|\phi\rangle \otimes |A(E, |\phi\rangle)\rangle. \quad (2.28)$$

By linearity and the definition of a QECC applied to the superposition codeword $\alpha|\bar{\psi}\rangle + \beta|\bar{\phi}\rangle$,

$$VE(\alpha|\bar{\psi}\rangle + \beta|\bar{\phi}\rangle) \otimes |0\rangle = \alpha\sqrt{c(E, |\psi\rangle)}|\psi\rangle \otimes |A(E, |\psi\rangle)\rangle + \beta\sqrt{c(E, |\phi\rangle)}|\phi\rangle \otimes |A(E, |\phi\rangle)\rangle \quad (2.29)$$

$$= \sqrt{c(E, \alpha|\psi\rangle + \beta|\phi\rangle)}(\alpha|\psi\rangle + \beta|\phi\rangle) \otimes |A(E, \alpha|\psi\rangle + \beta|\phi\rangle)\rangle. \quad (2.30)$$

These two expressions can only be equal if $|A(E, |\psi\rangle)\rangle = |A(E, |\phi\rangle)\rangle = |A(E)\rangle$ (or the discarded ancilla will be entangled with the output state) and $c(E, |\psi\rangle) = c(E, |\phi\rangle) = c(E)$ (or the decoded state for the superposition will be wrong). The conceptual point here is that in order to preserve the coherent superposition in the decoder's output, there cannot be any information about the encoded state left in the ancilla, and the amplitudes (and thus norms) of different encoded states have to match, or the Hilbert space will get distorted. \square

Note that the set of correctable errors is not a unique invariant property of an encoding map or code space, which is why I included it as part of the definition. The same code space could be used to correct different sets of errors. For instance, the 3-qubit phase correction code can correct the set $\mathcal{E} = \{I, Z_1, Z_2, Z_3\}$, but it also corrects the set $\mathcal{E} = \{I, Z_1Z_2, Z_1Z_3, Z_2Z_3\}$. In the former case, we interpret the non-zero error syndromes as caused by a single phase error, whereas in the latter case, we only consider two-qubit errors. Since Z_1 has the same error syndrome as Z_2Z_3 but they correspond to different logical states, we have to make a choice between them and cannot include both in the set of correctable errors at the same time. That is, $\mathcal{E} = \{I, Z_1, Z_2Z_3\}$ is *not* a possible set of correctable errors for the 3-qubit phase correction code.

There are many variations of the terms defined above. For instance, code space is sometimes coding space, code subspace, etc., and sometimes it is the encoded subspace. However, the encoded subspace also sometimes refers to \mathcal{H}_K , which can also be called the data or encoded data. As defined above, \mathcal{D} incorporates both the error correction procedure and the “unencoding” process of returning the data to \mathcal{H}_K , but sometimes they are considered separately.

The decoder \mathcal{D} incorporates whatever processing is necessary to correct and decode the state. For instance, it may incorporate a measurement of the error syndrome, and application of a correction operation conditional on the classical bits resulting from the syndrome measurement. The decoder \mathcal{D} may not, however, be unique. Its behavior is completely determined on the subspace spanned by $E|\bar{\psi}\rangle$, where $E \in \mathcal{E}$ and $|\bar{\psi}\rangle$ is a codeword, but outside of that subspace, \mathcal{D} can act arbitrarily, and the distinctions between different decoders can be important in a number of contexts, such as when studying the efficiency of the decoder, or its behavior on errors outside \mathcal{E} , or when considering fault tolerance.

Note that if we define a QECC just as a subspace instead of an encoding map, we haven't lost much information. In particular, the set of correctable errors is the same for all encoders:

Proposition 2.3. *Suppose we have two different encoders $U_1, U_2 : \mathcal{H}_K \rightarrow \mathcal{H}_N$ with $\text{Image}(U_1) = \text{Image}(U_2)$. Then (U_1, \mathcal{E}) is a QECC iff (U_2, \mathcal{E}) is a QECC.*

Proof. Because U_1 and U_2 are partial isometries with the same image, they can only differ by a unitary $V : \mathcal{H}_K \rightarrow \mathcal{H}_K$:

$$U_2 = U_1V. \quad (2.31)$$

If we use decoder \mathcal{D}_1 for encoder U_1 , then $V^\dagger \circ \mathcal{D}_1$ will serve as the decoding map for U_2 :

$$V^\dagger \mathcal{D}_1(EU_2|\psi\rangle) \langle \psi|U_2^\dagger E^\dagger \rangle V = V^\dagger \mathcal{D}_1(EU_1V|\psi\rangle) \langle \psi|V^\dagger U_1^\dagger E^\dagger \rangle V \quad (2.32)$$

$$= V^\dagger (c(E, V|\psi\rangle)V|\psi\rangle \langle \psi|V^\dagger) V \quad (2.33)$$

$$= c(E, |\psi\rangle)|\psi\rangle \langle \psi|. \quad (2.34)$$

\square

In the definition of a QECC, we have let the dimensions K , M , and N of the various Hilbert spaces involved be arbitrary, although $N \geq K$ or no QECC is possible. Nevertheless, there are some common restrictions. Frequently we assume $K = 2^k$ and $N = 2^n$, so there are k *logical qubits* (or *encoded qubits*) and n *physical qubits*. Sometimes we work with q -dimensional qudits instead of qubits, so $K = q^k$ and $N = q^n$. Occasionally we have even more general situations, where K and N might use different bases or not be powers at all. When there is a tensor factorization of the overall Hilbert space, but I don't want to be too specific about whether the individual factors are qubits, qudits, or maybe even different sizes from each other, I will refer to the *registers*, with each register being one tensor factor. A single error then affects a single register, whatever its size.

Most often $M = N$, so errors map the physical Hilbert space to itself. There is little lost by assuming this, since we can generally ignore any unused states in \mathcal{H}_N or \mathcal{H}_M without negative consequence. The main time when it matters is when errors occur repeatedly on the state before we perform error correction, as happens, for instance, in fault-tolerant quantum computation. In that case, you really need $M = N$, so that you can sensibly apply the same error over and over again.

One special case worth mentioning is that of erasure errors. Recall that an erasure error formally maps a qubit into a qutrit, so to treat erasure errors in the most precise way, we would take $N = 2^n$ and $M = 3^n$. However, most often we imagine that a “stop-leak” gate of some sort is performed before the decoder. The stop-leak gate will map the third $|\perp\rangle$ state of each qutrit to some state, perhaps a random one, of the corresponding qubit. Then we can consider an erasure error as mapping a qubit to a qubit, but with some additional classical information indicating that the qubit has undergone an error.

This is maybe a good place for an extremely important digression on the terminology of error correction. Specifically, I want to discuss the use of hyphens. Wait, did I say “important?” I meant “unimportant.” But I am going to discuss hyphens anyway. In English, hyphens are occasionally used in compound nouns but not that frequently. Where they *are* used is to make compound adjectives. Thus, we have a hyphen in “error-correcting code” and “fault-tolerant computation” but not in “error correction” and “fault tolerance,” unless one of the component words happens to get split between lines. (Notably, though, there is no hyphen for the “quantum” part, so “quantum error-correcting code.”) Certainly, many scientists are not native speakers of English, so they have a good excuse for making this mistake. But you no longer have an excuse.

2.4.3 Linearity of Quantum Error Correction

Now we are ready to generalize the phenomenon we observed in section 2.4.1.

Theorem 2.4. *If (U, \mathcal{E}) is a QECC, then (U, \mathcal{E}') is a QECC, where $\mathcal{E}' = \text{span } \mathcal{E}$ is the linear span of \mathcal{E} .*

In other words, if a QECC can correct E and F , then it can also correct $\alpha E + \beta F$.

Proof. Basically, the idea is to duplicate the argument used in section 2.4.1 for the general case. We haven’t defined the error syndrome for a general QECC, but we do have the decoding map to work with. We will purify it to V so that we can work with only unitary maps.

The QECC corrects $E, F \in \mathcal{E}$, so

$$V(E|\bar{\psi}\rangle_M \otimes |0\rangle_D) = c_E|\psi\rangle_K \otimes |s_E\rangle_{D'} \quad (2.35)$$

$$V(F|\bar{\psi}\rangle_M \otimes |0\rangle_D) = c_F|\psi\rangle_K \otimes |s_F\rangle_{D'}. \quad (2.36)$$

One difference from section 2.4.1 is that $|s_E\rangle$ and $|s_F\rangle$ don’t need to be orthogonal. Also note that the c ’s from definition 2.1 would be the squares of c_E and c_F .

By linearity,

$$V[(\alpha E + \beta F)|\bar{\psi}\rangle_M \otimes |0\rangle_D] = \alpha c_E|\psi\rangle_K \otimes |s_E\rangle_{D'} + \beta c_F|\psi\rangle_K \otimes |s_F\rangle_{D'} \quad (2.37)$$

$$= |\psi\rangle_K \otimes (\alpha c_E|s_E\rangle_{D'} + \beta c_F|s_F\rangle_{D'}). \quad (2.38)$$

Tracing out $\mathcal{H}_{D'}$, we find that the decoder map is of the desired form, with

$$c(\alpha E + \beta F, |\psi\rangle) = \|\alpha c_E|s_E\rangle + \beta c_F|s_F\rangle\|^2. \quad (2.39)$$

□

2.4.4 Correcting Paulis Implies Correcting All Errors

As a corollary of theorem 2.4, we can simplify the condition for a code to correct t errors:

Corollary 2.5. *If a QECC set of correctable errors includes all tensor products of I, X, Y , and Z of weight t or less, it is a t -error correcting code.*

Proof. The set of all errors with support on a given set of t qubits is the set of $2^t \times 2^t$ matrices acting on those qubits. Tensor products of I , X , Y , and Z acting on those t qubits form a basis for this set of matrices. Therefore, an arbitrary weight t error can be written as the sum (with appropriate coefficients) of weight t tensor products of I , X , Y , and Z , and an arbitrary t -qubit error can be written as a sum weight t errors. The corollary then follows from theorem 2.4. \square

As a consequence, we find that the nine-qubit code is a 1-error correcting code. Applying the definitions, we find that it encodes one logical qubit into nine physical qubits.

2.4.5 QECCs and Error Channels

We have finally answered all four objections to the existence of quantum error-correcting codes. However, it is worth thinking a little more carefully about what happens when we have a decoherent error that maps pure states to mixed states. For example, suppose that we have a dephasing channel with error probability p acting on qubit i . Then the pure codeword state $|\bar{\psi}\rangle$ becomes:

$$|\bar{\psi}\rangle \langle \bar{\psi}| \mapsto (1-p)|\bar{\psi}\rangle \langle \bar{\psi}| + pZ_i|\bar{\psi}\rangle \langle \bar{\psi}|Z_i, \quad (2.40)$$

which is a mixture of the pure states $|\bar{\psi}\rangle$ and $Z_i|\bar{\psi}\rangle$. Now, we know that the full error correction procedure corrects these two pure states:

$$|\bar{\psi}\rangle \mapsto |\bar{\psi}\rangle|I\rangle_{\text{syn}} \quad (2.41)$$

$$Z_i|\bar{\psi}\rangle \mapsto |\bar{\psi}\rangle|Z_i\rangle_{\text{syn}} \quad (2.42)$$

By the linearity of density matrices, that means that the final state after the dephasing channel followed by error correction is:

$$(1-p)|\bar{\psi}\rangle \langle \bar{\psi}| \otimes |I\rangle \langle I|_{\text{syn}} + p|\bar{\psi}\rangle \langle \bar{\psi}| \otimes |Z_i\rangle \langle Z_i|_{\text{syn}} = |\bar{\psi}\rangle \langle \bar{\psi}| \otimes [(1-p)|I\rangle \langle I|_{\text{syn}} + p|Z_i\rangle \langle Z_i|_{\text{syn}}]. \quad (2.43)$$

That is, the decoded logical qubit ends up as a pure state, but the overall state is still mixed: the randomness introduced by the error ends up in the error syndrome, or in the ancilla Hilbert space $\mathcal{H}_{D'}$ in the proof of theorem 2.4.

Recall that a t -qubit error channel is one for which there exists a Kraus decomposition where each Kraus operator is a sum of weight t linear operators. By theorem 2.4 and the argument above for the dephasing channel, any code that corrects arbitrary t -qubit linear errors will therefore also correct arbitrary t -qubit error channels.

What about independent channels?

Theorem 2.6. *Let \mathcal{I} be the 1-qubit identity channel and $\mathcal{E} = \otimes_{i=1}^n \mathcal{E}_i$ be an n -qubit independent channel, with $\|\mathcal{E}_i - \mathcal{I}\|_{\diamond} < \epsilon \leq \frac{t+1}{n-t-1}$, and let U and \mathcal{D} be the encoder and decoder for a QECC with n physical qubits that corrects t -qubit errors. Then*

$$\|\mathcal{D} \circ \mathcal{E} \circ U - \mathcal{I}\|_{\diamond} < 2 \binom{n}{t+1} (e\epsilon)^{t+1}. \quad (2.44)$$

In other words, a QECC that can correct t -qubit errors also corrects small independent error channels up to a very good approximation. A similar statement immediately follows from theorem 1.1. This should thus be viewed as a specialization of theorem 1.1 with better constant factors and a much easier proof.

Proof. The strategy is straightforward: Expand \mathcal{E} as a sum of a t -qubit error map (not normalized) and an additional small term. The t -qubit error map is corrected by the code due to linearity, and then we just have to bound the size of the additional term to prove the theorem.

Since $\|\mathcal{E}_i - \mathcal{I}\|_{\diamond} < \epsilon$, we can write $\mathcal{E}_i = \mathcal{I} + \delta\mathcal{E}_i$, with $\|\delta\mathcal{E}_i\|_{\diamond} < \epsilon$. Now,

$$\mathcal{E} = \mathcal{F} + \mathcal{G}, \quad (2.45)$$

where

$$\mathcal{F} = \sum_{r=0}^t \sum_{|S|=r} \bigotimes_{i \in S} \delta \mathcal{E}_i \quad (2.46)$$

$$\mathcal{G} = \sum_{r=t+1}^n \sum_{|S|=r} \bigotimes_{i \in S} \delta \mathcal{E}_i \quad (2.47)$$

are the sums over all tensor factors of $\leq t$ and $> t$ of the $\delta \mathcal{E}_i$ s, respectively. (The tensor with \mathcal{I} on other qubits is implicit.)

The first term \mathcal{F} is a sum of terms acting on at most t qubits. Each term in the sum is not a completely positive map. In fact, $\delta \mathcal{E}_i$ is not even positive. However, $\delta \mathcal{E}_i$ is a difference of two completely positive maps, and therefore

$$\mathcal{F}(\rho) = \sum_k \alpha_k A_k \rho A_k^\dagger, \quad (2.48)$$

where the α_k coefficients can be either positive or negative real numbers and each A_k acts on at most k qubits. By theorem 2.4 (and the linearity of density matrices, as for the dephasing channel example), the QECC corrects \mathcal{F} .

By proposition 2.2, the scaling factor $c(E, |\psi\rangle)$ in definition 2.1, the definition of a QECC, does not depend on $|\psi\rangle$, which implies that

$$\mathcal{D} \circ \mathcal{F} \circ U = c\mathcal{I} \quad (2.49)$$

for some constant c .

Meanwhile,

$$\|\mathcal{G}\|_\diamond \leq \sum_{r=t+1}^n \sum_{|S|=r} \prod_{i \in S} \|\delta \mathcal{E}_i\|_\diamond \quad (2.50)$$

$$\leq \sum_{r=t+1}^n \binom{n}{r} \epsilon^r \quad (2.51)$$

$$\leq \binom{n}{t+1} (e\epsilon)^{t+1} = \delta \quad (2.52)$$

by lemma 1.2.

Now, \mathcal{E} , \mathcal{D} , and U are CPTP maps, so $\|\mathcal{D} \circ \mathcal{E} \circ U\|_\diamond = 1$. Thus,

$$1 - \delta \leq \|\mathcal{D} \circ \mathcal{E} \circ U\|_\diamond - \|\mathcal{D} \circ \mathcal{G} \circ U\|_\diamond \leq \|\mathcal{D} \circ \mathcal{F} \circ U\|_\diamond = \|c\mathcal{I}\|_\diamond = c \leq \|\mathcal{D} \circ \mathcal{E} \circ U\|_\diamond + \|\mathcal{D} \circ \mathcal{G} \circ U\|_\diamond \leq 1 + \delta. \quad (2.53)$$

That is, $|1 - c| \leq \delta$. Therefore,

$$\|\mathcal{D} \circ \mathcal{E} \circ U - \mathcal{I}\|_\diamond = \|\mathcal{D} \circ \mathcal{F} \circ U + \mathcal{D} \circ \mathcal{G} \circ U - \mathcal{I}\|_\diamond \quad (2.54)$$

$$= \|(c-1)\mathcal{I} + \mathcal{D} \circ \mathcal{G} \circ U\|_\diamond \quad (2.55)$$

$$\leq |c-1| + \|\mathcal{G}\|_\diamond \quad (2.56)$$

$$\leq 2\delta. \quad (2.57)$$

□

2.5 The Quantum Error Correction Conditions

2.5.1 Sufficient Conditions (Non-degenerate Orthogonal Coding)

Recall the argument that we used to show that the nine-qubit code could correct phase errors. Very little about it was specific to the nine-qubit code. We said that phase errors generated subspaces which were orthogonal to the code space and to each other, and that therefore there was a measurement that distinguished

them. We can use this same argument to give a general sufficient condition to have a QECC: Suppose that $Q \subseteq \mathcal{H}_N$ is the code space, and let $E(Q)$ be the subspace formed by acting on Q with E . Then we want that for any pair of errors $E_1, E_2 \in \mathcal{E}$, the subspace $E_1(Q)$ is orthogonal to $E_2(Q)$. Then we will be able to make a measurement that identifies which subspace we are in and therefore identifies the error.

To be sure we have a QECC, we need a little bit more. After identifying the error, we need to be able to reverse it. That is only possible if $E|_Q$ is a unitary map from Q to $E(Q)$. $E|_Q$ is unitary iff $\langle \psi | E^\dagger E | \phi \rangle = \langle \psi | \phi \rangle$ for all $|\psi\rangle, |\phi\rangle \in Q$. Thus, we get the sufficient condition that (Q, \mathcal{E}) is a QECC if $\forall |\psi\rangle, |\phi\rangle \in Q, \forall E_a, E_b \in \mathcal{E}$,

$$\langle \psi | E_a^\dagger E_b | \phi \rangle = \delta_{ab} \langle \psi | \phi \rangle. \quad (2.58)$$

A code that satisfies this condition could be called a *non-degenerate orthogonal code*, although the term “orthogonal code” is not widely used.

2.5.2 The QECC Conditions

However, by looking at the nine-qubit code, we can already see that equation (2.58) is not a necessary condition. If $E_1 = Z_1$ and $E_2 = Z_2$, then equation (2.58) fails since $E_1|\bar{\psi}\rangle = E_2|\bar{\psi}\rangle$. We need to generalize slightly:

Theorem 2.7 (QECC Conditions). (Q, \mathcal{E}) is a QECC iff $\forall |\psi\rangle, |\phi\rangle \in Q, \forall E_a, E_b \in \mathcal{E}$,

$$\langle \psi | E_a^\dagger E_b | \phi \rangle = C_{ab} \langle \psi | \phi \rangle. \quad (2.59)$$

Note that C_{ab} does not depend on $|\psi\rangle$ or $|\phi\rangle$.

Proof. \Leftarrow : Recalling theorem 2.4, we might as well pick a useful spanning set for \mathcal{E} and restrict attention to that spanning set. Taking the adjoint of equation (2.59) and putting in $|\phi\rangle = |\psi\rangle$, we find that $C_{ab}^\dagger = C_{ba}^*$, i.e., the matrix C is Hermitian. Therefore C_{ab} is diagonalizable, and by choosing an appropriate spanning set $\{F_a\}$ for \mathcal{E} we can actually diagonalize C_{ab} . (Note that it is not necessarily true that $F_a \in \mathcal{E}$, but that is not particularly relevant, it just means that the original set \mathcal{E} is smaller than the maximal set of possible errors the code can correct.)

We have

$$\langle \psi | F_a^\dagger F_b | \phi \rangle = d_a \delta_{ab} \langle \psi | \phi \rangle. \quad (2.60)$$

This is almost equation (2.58), but the coefficient d_a can be different from 1. (d_a is an eigenvalue of C_{ab} , so it must be real.) However, it is still true that the different subspaces $F(Q)$ are orthogonal to each other, so we can make a measurement that determines which error F_a we have. It is not true that $F_a|_Q$ must be unitary, but if d_a is nonzero, $F_a|_Q$ can be written as a unitary followed by a uniform rescaling. The decoding then gives the original state rescaled by d_a . This is allowed by definition 2.1.

If d_a is zero, then $F_a|_Q$ cannot be inverted, since there is no state left. Formally, we are still OK, since in definition 2.1, we would just get that $c(F_a, |\psi\rangle) = 0$. Physically, this means that the error F_a *never occurs*. It has probability proportional to $\langle \psi | F_a^\dagger F_a | \psi \rangle$, which is 0.

Therefore, equation (2.59) gives sufficient conditions to have a QECC.

\Rightarrow : Recall that by proposition 2.2, $c(E, |\psi\rangle)$ and $|A(E, |\psi\rangle)\rangle$ don't depend on $|\psi\rangle$. More generally, if we purify the decoder to V , we get a unitary transformation, which preserves inner products:

$$\langle \bar{\psi} | E_a^\dagger E_b | \bar{\phi} \rangle = (\langle \bar{\psi} | \otimes \langle 0 |) E_a^\dagger V^\dagger V E_b (| \bar{\phi} \rangle \otimes | 0 \rangle) \quad (2.61)$$

$$= \sqrt{c(E_a)c(E_b)} (\langle \psi | \otimes \langle A(E_a) |) (| \phi \rangle \otimes | A(E_b) \rangle) \quad (2.62)$$

$$= \sqrt{c(E_a)c(E_b)} \langle A(E_a) | A(E_b) \rangle \langle \psi | \phi \rangle. \quad (2.63)$$

This is equation (2.59) with $C_{ab} = \sqrt{c(E_a)c(E_b)} \langle A(E_a) | A(E_b) \rangle$.

□

2.5.3 Degenerate Codes

The difference between equation (2.58) and equation (2.59) consists in the fact that C_{ab} might not be δ_{ab} . When C_{ab} has maximum rank, this difference has no deep meaning, since we have just made a bad choice of basis errors. As in the proof of theorem 2.7, we can choose a different set of errors with the same span to diagonalize C_{ab} , and we can even rescale to make $C_{ab} = \delta_{ab}$.

When C_{ab} has non-maximal rank, we cannot do this. If the error set \mathcal{E} is not linearly independent, C_{ab} cannot have maximal rank, since a linearly dependent set of errors will produce a linearly dependent set of rows in C_{ab} . The interesting case is when \mathcal{E} is linearly independent. It is possible then to still have a QECC with non-maximal rank for C_{ab} . This is known as a degenerate code.

Definition 2.2. Suppose (U, \mathcal{E}) is a QECC and \mathcal{E} is linearly independent. Then the code is *degenerate* if $\text{rank}(C_{ab}) < |\mathcal{E}|$. A code (U, \mathcal{E}) (for linearly dependent \mathcal{E}) is degenerate if (U, \mathcal{E}') is degenerate, where \mathcal{E}' is a minimal spanning set for \mathcal{E} . If a code is not degenerate, it is *non-degenerate*.

We have already seen an example of a degenerate code: the nine-qubit code. The essence of degeneracy is that different errors will produce the same result (or at least linearly dependent results) when acting on a codeword. In the nine-qubit code, any two Z errors acting on the same set of 3 qubits will produce the same result. For instance,

$$Z_1|\bar{0}\rangle = Z_2|\bar{0}\rangle = \frac{1}{2\sqrt{2}}(|000\rangle - |111\rangle)(|000\rangle + |111\rangle)(|000\rangle + |111\rangle) \quad (2.64)$$

$$Z_1|\bar{1}\rangle = Z_2|\bar{1}\rangle = \frac{1}{2\sqrt{2}}(|000\rangle + |111\rangle)(|000\rangle - |111\rangle)(|000\rangle - |111\rangle). \quad (2.65)$$

2.5.4 Distance

The most common situation when designing a QECC is to let the set of possible errors be all t -qubit errors. It is therefore worth examining theorem 2.7 more explicitly for a t -error correcting code. When E_a and E_b are both weight t errors, then $E_a^\dagger E_b$ is a weight $2t$ error.

Definition 2.3. Let $C \subseteq \mathcal{H}_N$ be a QECC. The *distance* of C is the minimum weight d of an error F such that

$$\langle \psi | F | \phi \rangle \neq c(F) \langle \psi | \phi \rangle, \quad (2.66)$$

where $|\psi\rangle$ and $|\phi\rangle$ run over all possible pairs of codewords of C .

Note that the notion of distance implies that the Hilbert space is broken up into qubits because weight is defined in terms of qubits. Naturally, we can easily generalize the distance to work with a code over qudits of dimension q by defining weight in terms of the number of qudits in the support of an operator.

We get the following corollary of theorem 2.7:

Corollary 2.8. A distance d code corrects $\lfloor (d-1)/2 \rfloor$ errors.

Inverting this formula, we find that to correct t errors, a code needs distance $2t + 1$. If the distance is even, the extra point is “wasted” for this application, but as we shall see in a moment, the extra point of distance can be helpful in alternative applications.

The three central properties of a QECC are the size of the logical subspace, the number of physical qubits, and the distance, so there is a notation encapsulating those properties.

Notation 2.4. A QECC which encodes \mathcal{H}_K into \mathcal{H}_{2^n} and has distance d is denoted as an $((n, K, d))$ code. If the physical Hilbert space is n qudits each of dimension q , it is an $((n, K, d))_q$ code. If the distance is unknown or irrelevant, it is an $((n, K))$ code or an $((n, K))_q$ code.

This notation is derived from an analogous one for classical error-correcting codes, which will be introduced in chapter 4. The classical notation uses single parentheses, and the double parenthesis indicates that we have a quantum error-correcting code.

Sometimes when we discuss a code we consider it to have a distance less than its true distance or to correct fewer errors than the maximum number it can correct. This is just a convenience indicating that we are ignoring some of the error-correcting capability of the code. However, the true distance of a QECC does give us the tightest estimate of the number of errors it can correct — that is, the converse of the corollary also holds. In order to get the definition of distance, we need F to run over all errors of weight $< d$, but the QECC conditions applied to a t -error correcting code only gives us equation (2.66) for errors of the form $E_a^\dagger E_b$, which does not include all possible weight $2t$ errors. However, equation (2.66) is linear in F , so it is sufficient to check the formula for a basis of the set of weight $2t$ errors, and the set of errors of the form $E_a^\dagger E_b$ does include such a basis.

By similar reasoning, you get the same notion of distance if you alter the definition of distance to use all d -qubit errors or just weight d tensor products of I , X , Y , and Z .

When describing a code without an explicit set of correctable errors, I said in the definition that you are supposed to determine the set of correctable errors by context. Typically, we will choose the error set to be the set of t -qubit errors, which makes the distance of a code one of its most critical properties. When the error set is given implicitly, we can define degeneracy in terms of the code's capability as a t -error correcting code.

Definition 2.5. Let $Q \subseteq \mathcal{H}_N$ be a QECC with distance $d = 2t + 1$. Then Q is *degenerate* if it is degenerate when the set of correctable errors is all t -qubit errors.

If the distance is even, $2t + 2$, and there is no more structure, it makes the most sense to define degeneracy by also considering the set of correctable errors to be t -qubit errors, but for certain families of codes (such as stabilizer codes, discussed in chapter 3), we can do better.

2.5.5 Quantum Error Detection and Erasure Errors

The distance is useful information about a QECC partially because it tells you how many errors the code can correct, but also because it encapsulates some additional error-correction-related properties of the code. In particular, the distance also tells you about the code's ability to detect errors without correcting them and about the code's ability to correct erasure errors. These applications can take advantage of even distance codes, whereas correcting general t -qubit errors only needs odd distance codes.

Definition 2.6. An encoder U (defined as for a QECC) and a set of detectable errors \mathcal{E} form a *quantum error-detecting code* if they have the following property: Let Π be the projector on the code space. Then $\Pi E |\psi\rangle = c(E, |\psi\rangle) |\psi\rangle$, for all $E \in \mathcal{E}$ and all codewords $|\psi\rangle$. The code space is defined as the image of U , as for a QECC, and we frequently refer to the code space as defining the error-detecting code instead of the encoder.

Based on the definition of an error-detecting code, the measurement $(\Pi, I - \Pi)$ will either project us back on the original state (with some probability $|c(E, |\psi\rangle)|^2$) or will identify that an error occurred. This condition can be easily rewritten in similar terms to the QECC conditions:

Theorem 2.9. (U, \mathcal{E}) is a quantum error-detecting code iff

$$\langle \psi | E | \phi \rangle = c(E) \langle \psi | \phi \rangle \tag{2.67}$$

for all codewords $|\psi\rangle$ and $|\phi\rangle$ and all $E \in \mathcal{E}$. A code with distance d detects arbitrary $(d - 1)$ -qubit errors.

Based on this theorem, we can understand the distance of the code as the minimum number of qubits on which we can act to produce an undetectable error, i.e., an error with a component taking a codeword to a *different* codeword. (An error taking a codeword to itself is considered “detectable” by the definition.)

Proof. \Leftarrow : We can write $\Pi = \sum |\psi_i\rangle \langle \psi_i|$ where the sum runs over a basis $|\psi_i\rangle$ for the code space Q . We can then calculate $\Pi E |\psi\rangle$ using equation (2.67) to see that we have a quantum error-detecting code.

\Rightarrow : Equation (2.67) follows immediately from the definition of a quantum error-detecting code if we can prove that $c(E, |\psi\rangle)$ does not depend on $|\psi\rangle$. This can be done by considering a superposition $\alpha|\psi_1\rangle + \beta|\psi_2\rangle$, as in proposition 2.2.

The definition of distance then shows that a distance d code detects $d - 1$ errors. \square

While I have presented quantum error-correcting codes and quantum error-detecting codes as different things, clearly there is a very close connection. A code able to correct t errors will also be able to detect $2t$ errors, and vice-versa. Detecting errors and correcting errors are really just two different applications for the same code, and henceforth, I won't make a distinction between a code designed to correct errors and one designed to detect errors. Both will be referred to as QECCs.

There is no unique maximal set of correctable errors for a QECC, but there is a unique maximal set of detectable errors. The quantum error-correction conditions involve a product of two errors, and there may be more than one set of errors that will run over all possibilities for the product. However, equation (2.67) is only linear in the error, so we can define a unique set of detectable errors.

Definition 2.7. Given a subspace Q , its set of *detectable errors* is

$$\{E \text{ s.t. } \langle \psi | E | \phi \rangle = c(E) \langle \psi | \phi \rangle \forall |\psi\rangle, |\phi\rangle \in Q\}. \quad (2.68)$$

If \mathcal{E}_C is the set of correctable errors and \mathcal{E}_D is the set of detectable errors for a code, then the QECC conditions can be rephrased as $\mathcal{E}_C^2 \subseteq \mathcal{E}_D$.

A code's ability to correct erasure errors can be understood just by applying the regular QECC conditions. The interesting twist in this case is that we can apply the side classical information we have about the location of the errors to correct twice as many errors:

Theorem 2.10. *A QECC with distance d can correct $d - 1$ erasure errors.*

Proof. The best way to think about an erasure-correcting code is as a set of QECCs which all have the same encoder. Each code is associated with a different error set, depending on where the erasure errors took place. Since we don't know when doing the encoding where the errors are, we have to use the same encoder in all cases. When *decoding*, however, we know where the errors are (although not what kind of errors they are), so we can choose a decoder based on the actual error set — all possible errors on the actual set of qubits erased.

Therefore we need a single code space that satisfies the QECC conditions for any error set of the form \mathcal{E}_S , which is the set of all possible errors with support on the set S of at most t qubits. But $\mathcal{E}_S^2 = \mathcal{E}_S$ since the product of two errors with support on S still has support on S . For a distance d code, the set of detectable errors includes all $(d - 1)$ -qubit errors, so when $t \leq d - 1$, all sets \mathcal{E}_S are subsets of the set of detectable errors. \square

2.5.6 Alternate Forms of the Quantum Error Correction Conditions

There are a number of variants of the QECC conditions which are useful in different contexts. You have just seen one relating a correctable set of errors with the set of detectable errors for a code. A few more appear in the following proposition:

Proposition 2.11. *The following are equivalent to the QECC conditions given in theorem 2.7:*

1. For all codewords $|\psi\rangle$, all pairs $E_a, E_b \in \mathcal{E}$,

$$\langle \psi | E_a^\dagger E_b | \psi \rangle = \text{tr}(|\psi\rangle \langle \psi | E_a^\dagger E_b) = C_{ab}. \quad (2.69)$$

2. If $\text{span}(\mathcal{E}) = \mathcal{E}$: For any pair of codewords $|\psi\rangle, |\phi\rangle$, any error $E \in \mathcal{E}$,

$$\langle \psi | E^\dagger E | \phi \rangle = C(E) \langle \psi | \phi \rangle. \quad (2.70)$$

3. If $\text{span}(\mathcal{E}) = \mathcal{E}$: For any codeword $|\psi\rangle$, any error $E \in \mathcal{E}$,

$$\text{tr}(|\psi\rangle\langle\psi|E^\dagger E) = C(E). \quad (2.71)$$

4. Let $\{|\psi_i\rangle\}$ be a basis for the code space. For any i and j and for any pair $E_a, E_b \in \mathcal{E}$,

$$\langle\psi_i|E_a^\dagger E_b|\psi_j\rangle = C_{ab}\delta_{ij}. \quad (2.72)$$

Proof. The QECC conditions from equation (2.59) immediately imply all four of these variant conditions, so we only need to show the reverse direction.

To show that the standard QECC conditions follow from the first variant above, pick arbitrary $|\phi_1\rangle$ and $|\phi_2\rangle$ in the code space and consider $|\psi\rangle = \alpha|\phi_1\rangle + \beta|\phi_2\rangle$ for different values of α and β . $|\phi_1\rangle$ and $|\phi_2\rangle$ may not be orthogonal, so we have

$$1 = |\alpha|^2 + |\beta|^2 + 2\text{Re}(\alpha^*\beta\langle\phi_1|\phi_2\rangle). \quad (2.73)$$

Now, $|\psi\rangle$ is also in the code space, and

$$C_{ab} = \langle\psi|E_a^\dagger E_b|\psi\rangle \quad (2.74)$$

$$= |\alpha|^2\langle\phi_1|E_a^\dagger E_b|\phi_1\rangle + |\beta|^2\langle\phi_2|E_a^\dagger E_b|\phi_2\rangle + \alpha^*\beta\langle\phi_1|E_a^\dagger E_b|\phi_2\rangle + \alpha\beta^*\langle\phi_2|E_a^\dagger E_b|\phi_1\rangle \quad (2.75)$$

$$= (|\alpha|^2 + |\beta|^2)C_{ab} + (\alpha^*\beta\langle\phi_1|E_a^\dagger E_b|\phi_2\rangle + \alpha\beta^*\langle\phi_2|E_a^\dagger E_b|\phi_1\rangle). \quad (2.76)$$

If we first plug in $\alpha = \beta = 1/\sqrt{2(1 + \text{Re}\langle\phi_1|\phi_2\rangle)}$, we find

$$\langle\phi_1|E_a^\dagger E_b|\phi_2\rangle + \langle\phi_2|E_a^\dagger E_b|\phi_1\rangle = C_{ab}(\langle\phi_1|\phi_2\rangle + \langle\phi_2|\phi_1\rangle). \quad (2.77)$$

Plugging in $\alpha = -i\beta = 1/\sqrt{2(1 - \text{Im}\langle\phi_1|\phi_2\rangle)}$, we get

$$\langle\phi_1|E_a^\dagger E_b|\phi_2\rangle - \langle\phi_2|E_a^\dagger E_b|\phi_1\rangle = C_{ab}(\langle\phi_1|\phi_2\rangle - \langle\phi_2|\phi_1\rangle). \quad (2.78)$$

Putting these two equations together, we find

$$\langle\phi_1|E_a^\dagger E_b|\phi_2\rangle = C_{ab}\langle\phi_1|\phi_2\rangle, \quad (2.79)$$

as desired.

The second and third variants use a similar argument for E_a and E_b . The proof of the fourth version is left as an exercise. \square

Applying a similar argument to definition of the set of detectable errors, we find that an operator E is detectable iff $\text{tr}(\rho E)$ does not depend on the codeword ρ . This has an interesting interpretation: it says that an operator is detectable if and only if measurement of that operator reveals no information about the logical state of the code. Certainly, if measuring an operator does reveal information about the logical state, it will create errors in the encoded state. What is perhaps surprising is that first, that some element of that error cannot be detected, and second, that revealing encoded information is the *only* thing that prevents an error from being detectable.

Another version of this insight appears when we apply the QECC conditions specifically to erasure errors.

Proposition 2.12. *Let \mathcal{E} be a set of erasure errors, each one corresponding to a set of erased qubits (or qudits). Then Q corrects \mathcal{E} iff ρ_S is the same for all logical states $|\psi\rangle$ whenever $S \in \mathcal{E}$. Here S is a subset of qubits that can be erased and ρ_S is the encoded state restricted to S .*

Note that the proposition just says that Q can correct for erasures on the set S iff the codeword on S has no information about the encoded state. Another interesting feature of this is that the ability to correct for erasures only refers to single sets, not pairs of sets, whereas more generally the QECC conditions refer to pairs of errors. This means that any QECC has a unique maximal set of erasure errors that it can correct, whereas, as I mentioned before, there is not a unique set of general errors that can be corrected — there is some tradeoff between correcting different kinds of errors.

Proof. As in the proof of theorem 2.10, we think of the code as a set of QECCs labelled after-the-fact by the set of erased qubits. We specialize to a single set $S \in \mathcal{E}$ and want to show that Q corrects erasures on S iff ρ_S is independent of the encoded state. Once we fix S , correcting erasures on S is equivalent to correcting arbitrary linear operators supported on S .

Let $E_a = |k\rangle\langle j|$ and $E_b = |k\rangle\langle i|$ where i, j , and k are bitstrings labelling basis vectors for the qubits just in S . We use variant 1 of the QECC conditions from proposition 2.11. Since E_a and E_b are always supported on S ,

$$\text{tr}(|\psi\rangle\langle\psi|E_a^\dagger E_b) = \langle i|\rho_S|j\rangle\langle k|k\rangle \quad (2.80)$$

$$= (\rho_S)_{ij}. \quad (2.81)$$

As we let i and j vary over all possible basis states for S , equation (2.69) says that ρ_S does not depend on which codeword $|\psi\rangle$ we have, proving the forward direction of the proposition.

The reverse direction follows easily from variant 1 or variant 3 of proposition 2.11: If ρ_S is independent of $|\psi\rangle$, then taking the trace of ρ with operators on S will also give something independent of $|\psi\rangle$. \square

2.6 What Makes a Good QECC?

Before we get any further in our discussion of quantum error-correcting codes, it is worth stopping to think about what we want out of a code. The first thing that comes to mind is that we want it to be good at correcting errors, so we want the set \mathcal{E} of correctable errors to be large. When we're dealing with a t -error channel, we want the distance of the QECC to be large.

In order to correct errors, we have to add some extra qubits. For instance, to make the 9-qubit code, we had to add 8 extra qubits to encode 1. That 9 : 1 ratio seems pretty bad, so we'd like to find codes that use less overhead. In other words, we want k , the number of encoded qubits, to be large, and n , the number of physical qubits, to be small. Certainly we need $n > k$, but we'd like the *rate* k/n to be as close to 1 as possible. It's not clear a priori whether it's better to have a code with large k or small k , so we'd like good examples of codes for any value of k .

In general, if we fix t (the number of errors corrected) and let k get larger, we can make k/n larger too. However, this is rarely a sensible thing to do. As n gets larger, there are more opportunities for errors, so the expected number of errors gets larger. A better plan is to fix t/n when we vary n . In that scenario, there is a definite limit on the rate k/n that can be achieved, and we'd like to get as close to it as possible. We'd also like to know, theoretically, what the best possible rate is so that we know what to shoot for when designing codes. A family of codes that has both k/n and t/n as constants when n gets arbitrarily large is known as a *good* family of codes, and such code families are known.

The three parameters n , k , and d encapsulate the most interesting properties of a code, but not the only interesting ones. In part II, we'll talk about fault-tolerant quantum computation, and we'd like codes that are good for doing fault-tolerant protocols. We'd like a code that has a nice succinct classical description — describing even a single state on n qubits could require specifying 2^n amplitudes, and to specify a QECC we might need to list all 2^k states in a basis for the code subspace. Preferably, the n and k values of the code will be in the correct size range for the application we have in mind — we wouldn't want to have to add lots of extra logical qubits in order to use an otherwise nice code.

We'd also like the encoding and decoding operations to be implementable with a reasonable number of quantum gates. Frequently, when there is a short description of the code, the encoder can be performed with a small circuit. For instance, this is the case for the stabilizer codes which will be introduced in chapter 3. Ideally, the encoder would run in time $O(n)$ for a code with n physical qubits, and that is harder to achieve. (An arbitrary stabilizer code takes time $\Omega(n^2/\log n)$ to encode.) We might want additional properties, for instance that the encoding circuit can be implemented easily in a two-dimensional layout.

Decoding is a much more sticky problem. Even codes with a simple description may have a very difficult decoder. For the 9-qubit code, we had an error syndrome which identified the error. The error syndrome is not well-defined for all QECCs, but even if we restrict attention to codes with an error syndrome, we are faced with the *syndrome decoding problem* of determining the actual error given the error syndrome.

The syndrome decoding problem is NP-hard for most broad classes of codes. (This is true even for classical error-correcting codes.) When we come up with new codes, we'd like ones for which the decoding problem is not nearly as hard. Again, the ideal solution would be a code for which the syndrome decoding problem can be solved in linear time. Such codes exist, but they are somewhat rare.

This is a rather long wish list of properties we'd like out of a code. It's actually possible to satisfy many of them at the same time, but we don't currently know of any QECCs that are perfect in all respects. In other words, choosing a code is a matter of trade-offs. We'll want to use one code for some purposes and a different code for other purposes, depending on which property is most desirable for our current goal.

One thing to bear in mind when looking for a new code is that codes which look different may have pretty much the same properties. For instance, if you switch the first and second qubits of a code, you probably have a different subspace than you did before, but it doesn't really deserve the name of a new code. We capture this intuition with the notion of equivalent codes:

Definition 2.8. Two QECCs C and C' on n physical qubits are *equivalent* if C' can be produced from C by performing some set of single-qubit unitaries and by permuting the qubits.

The notion of equivalence can of course be extended to codes on a q^n -dimensional Hilbert space as well. Two equivalent codes have the same basic properties.

Proposition 2.13. *If C and C' are equivalent, then C and C' have the same number of logical qubits and the same distance.*

It's not necessarily true that C and C' correct exactly the same set of errors, since the local unitaries and permutations of the qubits can scramble up the errors. However, it's always true that an equivalence will map errors into errors of the same weight. It's also true that a fast decoding procedure for C will give a fast decoding procedure for C' , so the codes don't differ in computational complexity either.

When you're looking for new codes, if you find a code that's equivalent to a known code, you probably shouldn't consider it to be new. It's not always straightforward to tell if two codes are equivalent, however. The safest thing is to find codes with new (hopefully better) parameters $((n, K, d))$ than preexisting codes.

Chapter 3

Will The Real Codeword Please Stand Still?: Stabilizer Codes

The treatment of QECCs in the last chapter was very general, but a bit unwieldy, particularly when it comes to finding new codes. We'd like a better method of discussing, manipulating, and finding QECCs, even if it comes at the cost of restricting somewhat the codes we can talk about. This chapter introduces the class of stabilizer codes, which have some nice properties and are much more tractable for most purposes than a general QECC. Stabilizer codes are built around a group of symmetries of the code that can distinguish between correct codewords and states with errors: The incorrect states will change under the action of the symmetry, while the correct ones will stay put.

3.1 The 9-Qubit Code Revisited

3.1.1 Error Syndrome Measurement for the 9-Qubit Code

The 9-qubit code is a stabilizer code, so I'll introduce the basic idea of a stabilizer code with some further examination of the nine-qubit code. Let's consider more carefully exactly what we measure when we measure the error syndrome for the nine qubit code. There are two parts to it: We measure each set of three qubits to see if one is different in the standard basis, and then we compare the three phases of the three sets of three qubits to see if one of the phases is different from the other two.

Recall that to determine if one of three qubits is different, we make two measurements. One measurement determines the parity of the first two qubits, and the other tells us the parity of the second and third qubits. Let us put that in a more quantum language: Determining the parity of two qubits is the same as measuring the eigenvalue of $Z \otimes Z$ on those qubits.

$$Z \otimes Z = \begin{pmatrix} +1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & +1 \end{pmatrix} \begin{array}{l} 00 \text{ (even)} \\ 01 \text{ (odd)} \\ 10 \text{ (odd)} \\ 11 \text{ (even)} \end{array} \quad (3.1)$$

Eigenvalue +1 corresponds to even parity (syndrome bit 0) and eigenvalue -1 corresponds to odd parity (syndrome bit 1). Thus, within each set of three, to figure out if one qubit is different in the standard basis, we should measure $Z \otimes Z \otimes I$ and $Z \otimes I \otimes Z$.

We also need to determine whether two sets of three have the same phase or opposite phase. We'd like to phrase this as measurement of an eigenvalue. It should probably involve X , since the eigenvalues of X tell us the phase in $|0\rangle \pm |1\rangle$, but we actually have an entangled state. To determine the phase of $|000\rangle \pm |111\rangle$ we should instead measure the eigenvalue of $X \otimes X \otimes X$. To determine whether two sets of three have the same phase or opposite phase, we should thus measure the eigenvalue of the tensor product of six X 's on

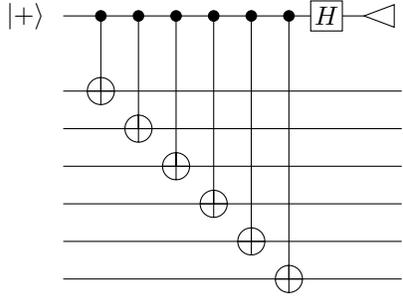


Figure 3.1: a) A non-fault-tolerant measurement procedure for $X \otimes X \otimes X \otimes X \otimes X \otimes X$.

			Z	Z															
				Z	Z														
						Z	Z												
							Z	Z											
									Z	Z									
										Z	Z								
											Z	Z							
												Z	Z						
													Z	Z					
														Z	Z				
															Z	Z			
																Z	Z		
																	Z	Z	
																		Z	Z
X	X	X		X	X	X													
			X	X	X	X	X	X	X	X									

Table 3.1: The generators of the stabilizer for the nine-qubit code. Blank spaces represent I operators.

the six relevant qubits. Remember, we want to learn whether the phases are the same or different without learning the actual value of each phase.

In case you're curious as to how one would actually measure this eigenvalue, you can look at figure 3.1. The question of how to measure the eigenvalue of a tensor product of arbitrary Paulis will be discussed in detail in section 12.1. Most of the focus of that section is on how to do the measurement fault-tolerantly, but it starts by discussing how to make a non-fault-tolerant measurement.

3.1.2 The Stabilizer for the 9-Qubit Code

The error syndrome for the 9 qubit code is 8 bits long: 2 bits for each set of three, identifying a bit flip error within that group of qubits, plus 2 additional bits telling us if and where a phase error has occurred. Putting together the list of all the operators whose eigenvalues give us those syndrome bits, we get the list in table 3.1. Each row corresponds to one operator, and the operator in column i indicates how that operator acts on qubit i . Blank spaces should be interpreted as I . Thus, the first row is the operator $Z \otimes Z \otimes I \otimes I \otimes I \otimes I \otimes I \otimes I \otimes I$. In other stabilizers, sometimes I will include the I s and sometimes omit them, whichever is most aesthetically appealing. (OK, maybe clarity plays a role too. And whimsy, definitely whimsy.)

As I mentioned in section 2.2, the choice of what qubits to measure for the error syndrome is not unique. For instance, we could have compared the parity of the first and third qubits. The parity would then be given by the eigenvalue of the operator $Z_1 Z_3$. Notice that $Z_1 Z_3$ is the product of the first two rows of table 3.1, and naturally the eigenvalue of the product will be the product of the eigenvalues. In general, we could use any product M of the operators listed in the table. The eigenvalue of M will tell us some information about the errors, but it won't be *new* information, since the eigenvalue of M can be determined from the eigenvalues of the operators used in the product.

Really, the operators we can measure to tell us about the error form a group, known as the *stabilizer* of the nine-qubit code. The eight operators listed in table 3.1 are the *generators* of the stabilizer. Choosing a

different way of defining the error syndrome of the code corresponds to choosing a different set of generators of the stabilizer.

Notice that we use tensor products of Z s to identify X (bit flip) errors, and we use tensor products of X s to identify Z (phase flip) errors. The pattern here is that the errors anticommute with the stabilizer elements used to find them, and it is this property that makes stabilizers useful for discussing quantum error-correcting codes.

3.2 Pauli Group

I warned that you would be sick of the Pauli operators by the end of this book, and this chapter will get you started. Stabilizer codes are built around the Pauli operators, as you can perhaps see from the example of the nine-qubit code. It's therefore worth examining them in more detail.

Definition 3.1. The *Pauli group* \mathbb{P}_n is composed of tensor products of I , X , Y , and Z on n qubits, with an overall phase of ± 1 or $\pm i$.

I will refer to elements of the n -qubit Pauli group as “Pauli operators,” “Pauli errors,” or just “Paulis” in the future; if I need to make a distinction between \mathbb{P}_1 and \mathbb{P}_n , I will do so explicitly.

The Pauli group, as its name suggests, is a group: It is closed under multiplication since $Y = iXZ$, and similarly, the product of any two of X , Y , and Z is equal to the third with an overall factor of $\pm i$. In addition, $X^2 = Y^2 = Z^2 = I$. Any tensor product of Paulis is its own inverse, i.e., it squares to I , and if there is an overall phase factor of $\pm i$, it instead squares to $-I$. (I use the one-qubit I and the n -qubit identity I interchangeably.)

The one-qubit Paulis anticommute with each other

$$XZ = -ZX \tag{3.2}$$

$$YZ = -ZY \tag{3.3}$$

$$XY = -YX. \tag{3.4}$$

Of course, I commutes with X , Y , and Z , and any Pauli commutes with itself. More general pairs (P, Q) of operators from \mathbb{P}_n always either commute ($PQ = QP$) or anticommute ($PQ = -QP$). The difference from the single qubit case is that for \mathbb{P}_1 , only trivial pairs commute (when one Pauli is I or both are the same Pauli), but for $n > 1$, there are non-trivial commuting pairs. For instance, $X \otimes X$ commutes with $Z \otimes Z$. I will sometimes use the notation $[P, Q] = PQ - QP$ for the commutator and $\{P, Q\} = PQ + QP$ for the anticommutator.

\mathbb{P}_n has 4^{n+1} elements since there are 4^n n -fold tensor products of I , X , Y , and Z , and 4 overall phases they could have. However, for many purposes, we can ignore the phase, giving us effectively 4^n distinct Paulis. If I wish to ignore the phase, I will refer to $\hat{\mathbb{P}}_n$. Thus, $\hat{\mathbb{P}}_n \cong \mathbb{P}_n / \{I, iI, -I, -iI\}$. The elements of $\hat{\mathbb{P}}_n$ are sets of 4 operators of the form $\{\pm P, \pm iP\}$, with P some n -qubit tensor product of I , X , Y , and Z . Each element of $\hat{\mathbb{P}}_n$ can thus be associated with an unsigned Pauli operator P , and it is more convenient by far to refer to elements of $\hat{\mathbb{P}}_n$ in terms of the associated Pauli rather than as a set of Paulis with signs.

Definition 3.2. Let $P \in \mathbb{P}_n$. Then $\hat{P} \in \hat{\mathbb{P}}_n$ is the element of $\hat{\mathbb{P}}_n$ corresponding to P . Similarly, if S is a subset of \mathbb{P}_n , then \hat{S} is the subset of $\hat{\mathbb{P}}_n$ consisting of \hat{P} for all $P \in S$.

I have introduced these conventions separating \mathbb{P}_n and $\hat{\mathbb{P}}_n$ in order to be mathematically precise. In some cases, we need to work with elements or subsets of \mathbb{P}_n , and in some cases, to get the details completely correct, it is better to work with $\hat{\mathbb{P}}_n$. However, almost always, the distinction between the two is a small technical detail. My advice is to ignore the difference between \mathbb{P}_n and $\hat{\mathbb{P}}_n$ unless you get confused.

Any Pauli operator has eigenvalues $+1$ and -1 . Each eigenspace is exactly half the Hilbert space. That is, the $+1$ and -1 eigenspaces have dimension 2^{n-1} .

One feature of the Pauli group already mentioned is that \mathbb{P}_n spans the space of all n -qubit linear operations. The weight t Paulis span the set of t -qubit errors. As discussed in corollary 2.5, this property means that to have a t -error correcting QECC, it suffices to correct all weight t Pauli operators.

This section has only discussed the most basic properties of the Pauli group. It has many more useful properties, some of which we will encounter in future sections of the book. Don't be surprised, however, if someday you are trying to solve a problem and discover a new property of the Paulis that is not mentioned in this book. In group theory, the Pauli group is an *extraspecial* group, and it truly deserves that name.

3.3 Stabilizer Codes

3.3.1 Definition and Basic Properties of the Stabilizer

Given a general quantum error-correcting code, we can define its stabilizer in more or less the same way as we did for the nine-qubit code:

Definition 3.3. Given a QECC $T \subseteq \mathcal{H}_{2^n}$, the *stabilizer* $S(T)$ is

$$S(T) = \{M \in \mathbb{P}_n \mid M|\psi\rangle = |\psi\rangle \forall |\psi\rangle \in T\}. \quad (3.5)$$

In other words, the stabilizer is composed of the Pauli operators for which all codewords are +1 eigenstates. In principle, one could define a stabilizer consisting of all unitaries which fix every codeword, but it turns out to be most useful to concentrate on the Pauli operators.

Proposition 3.1. *The stabilizer $S(T)$ of a nonempty QECC T has three basic properties:*

- a) $-I \notin S(T)$
- b) $S(T)$ is a group
- c) $S(T)$ is Abelian

The property of being Abelian is the least obvious, but it makes sense since we want a set of operators which have the codewords as simultaneous eigenstates. Normally, for this to be possible, the operators should commute; in general, they only need to commute on a subspace, but for Pauli operators this is only possible if they truly commute.

Proof.

- a) $-I$ has no +1 eigenstates, so it cannot be in the stabilizer.
- b) If $M, N \in S(T)$, it is clear that their product is also in $S(T)$: For any $|\psi\rangle \in T$,

$$MN|\psi\rangle = M|\psi\rangle = |\psi\rangle. \quad (3.6)$$

- c) By the argument for property 2, $MN|\psi\rangle = NM|\psi\rangle$ for any $|\psi\rangle \in T$, $M, N \in S(T)$. Thus, $[M, N]$ annihilates the codewords. Paulis either commute or anticommute. If M and N anticommute, then $[M, N] = 2MN$, which is again an element of \mathbb{P}_n , and therefore has no 0 eigenvalues. Therefore, the only option is that M and N commute.

□

It turns out to be most useful, when dealing with stabilizers, to work the other way around. We are given (or invent) a stabilizer with the properties we desire, and then use it deduce the code.

Definition 3.4. Let $S \subseteq \mathbb{P}_n$ be an Abelian group, with $-I \notin S$. Then define the code space corresponding to S by

$$\mathcal{T}(S) = \{|\psi\rangle \text{ s.t. } M|\psi\rangle = |\psi\rangle \forall M \in S\} \quad (3.7)$$

In general, $T \subseteq \mathcal{T}(S(T))$, but when they are actually equal, the code has special properties.

X	Z	Z	X	I
I	X	Z	Z	X
X	I	X	Z	Z
Z	X	I	X	Z

Table 3.2: The generators for the five-qubit code.

Definition 3.5. If T is a QECC with $T = \mathcal{T}(S(T))$, it is a *stabilizer code*.

Stabilizer codes are also sometimes known as *symplectic codes*, *additive codes*, or $\text{GF}(4)$ *additive codes* for reasons that will be discussed in section 3.5 and section 5.2.

If we define a code from its stabilizer, then it is always a stabilizer code. In other words, a stabilizer code is one that could be defined just by giving its stabilizer. This is a consequence of the following proposition:

Proposition 3.2. $S = S(\mathcal{T}(S))$.

I delay the proof until section 3.5 to keep you in suspense. Also, the proof will use a tool I won't introduce until then.

When dealing with stabilizer codes, I will frequently refer to the stabilizer S as the code instead of using the rather unwieldy phrase “stabilizer of the code $\mathcal{T}(S)$.” When $M \in S$, I will call M a “stabilizer element” or some variant of that phrase. However, many people seem to find that terminology unwieldy too, and just call M a “stabilizer.” I do not approve of that usage, but I can't stop you if you want to say it.

Frequently we will want to pick a particular generating set $\{M_1, \dots, M_r\}$ for the stabilizer, as we did for the nine-qubit code. In a minimal generating set no generator is a product of other generators. When we take a product of generators, the order does not matter since the generators commute. Also, note that since $-I \notin S$, all elements of S must square to I , and any power of a generator greater than 1 can be reduced. Therefore, the elements of the stabilizer are uniquely determined by taking a product

$$M_{i_1, \dots, i_r} = \prod_j M_j^{i_j}, \tag{3.8}$$

with $i_j \in \{0, 1\}$. This implies $|S| = 2^r$.

Naturally, the set of generators is not unique, but when we do have a particular generating set to work with, I will talk about its elements as “stabilizer generators” or just “generators.” When you describe a stabilizer, almost always the easiest way to do so is by listing one particular set of generators. The generators are enough to define the whole stabilizer, and it is much easier to list r generators than to list all 2^r elements of the stabilizer.

The code space of a stabilizer code is defined by requiring that all the eigenvalues be $+1$, but there is nothing magical about the $+1$ eigenstates. After all, a -1 eigenstate of M is a $+1$ eigenstate of $-M$, which is still in P_n . Given any set of generators $\{M_i\}$, we could define the code equally well by taking the -1 eigenstate, but switching M_i to $-M_i$. We *can't* arbitrarily change the eigenvalues associated with non-generators because the eigenvalues of the non-generators can be deduced from the eigenvalues of the generators. The generators are independent of each other, so we can pick their eigenvalues freely, but once we've done that, it defines the eigenvalues of all operators in the stabilizer. For more on this point, see section 3.5.

I'll conclude this subsection with a second example QECC to go with the nine-qubit code. This code has five physical qubits, and is, unsurprisingly, known as the “five-qubit code.” The stabilizer for the five-qubit code is given in table 3.2. The code is cyclic: if you permute the qubits cyclically, you'll get the same stabilizer. You can almost see this from table 3.2. Shifting by one qubit for generators one through three gives you generators two through four. However, the fifth permutation $Z \otimes Z \otimes X \otimes I \otimes X$ is not one of the generators. Nevertheless, it is still in the stabilizer, as the product of all four generators given in the table. Even though it's not a generator, it still has the same status in the code as the four operators that are listed. Indeed, we could have chosen any four of these five operators as a set of generators without changing the code.

3.3.2 Projection Operator on a Stabilizer Code Subspace

In order to convert back from the representation of a stabilizer code in terms of its stabilizer to one as a subspace of a larger physical Hilbert space, we have definition 3.4, but it is sometimes helpful to have something more concrete. In particular, we'd like a projection operator that defines the subspace.

The codewords are supposed to be +1 eigenvectors of every element of the stabilizer, although it is sufficient to check that they are +1 eigenvectors of the stabilizer generators. The projector on the +1 eigenspace of generator M_i is $(I + M_i)/2$. Therefore, the projector on the code space of S is

$$\Pi_S = \frac{1}{2^r} \prod_{i=1}^r (I + M_i), \quad (3.9)$$

when the stabilizer has r generators. Only states that are codewords — +1 eigenvectors of all generators — will avoid being annihilated by Π_S , and any codeword will be left alone by Π_S , as desired. Note that the order of the generators in the product does not matter since they all commute.

We can rewrite the projector Π_S in an interesting way by multiplying out the product. We get a sum of terms, each of which consists of a distinct product of the generators $\{M_i\}$. Based on equation (3.8), the products of the generators give us all elements of the stabilizer:

$$\Pi_S = \frac{1}{2^r} \sum_{M \in S} M. \quad (3.10)$$

We can come up with actual codewords for the code by applying Π_S to states in the physical Hilbert space and renormalizing. We might get 0, if the state we started with is orthogonal to the code space, but frequently we'll get a real codeword. For instance, for the five-qubit code, we can apply the projector to $|00000\rangle$ and $|11111\rangle$ to get two orthogonal codewords which can serve as a basis for the code space:

$$\begin{aligned} |\bar{0}\rangle &= |00000\rangle + |10010\rangle + |01001\rangle - |11011\rangle + |10100\rangle - |00110\rangle - |11101\rangle - |01111\rangle \\ &\quad + |01010\rangle - |11000\rangle - |00011\rangle - |10001\rangle - |11110\rangle - |01100\rangle - |10111\rangle + |00101\rangle \end{aligned} \quad (3.11)$$

$$\begin{aligned} |\bar{1}\rangle &= |11111\rangle + |01101\rangle + |10110\rangle - |00100\rangle + |01011\rangle - |11001\rangle - |00010\rangle - |10000\rangle \\ &\quad + |10101\rangle - |00111\rangle - |11100\rangle - |01110\rangle - |00001\rangle - |10011\rangle - |01000\rangle + |11010\rangle. \end{aligned} \quad (3.12)$$

I got these codewords by working my way through all 16 elements of the stabilizer for the five-qubit code (products of the generators given in table 3.2) and applying them to the starting states. You need to be careful of the signs involved, but otherwise the process is straightforward if tedious.

3.3.3 Distance and Size of a Stabilizer Code

When a stabilizer code is given either as a subspace or via the stabilizer, the number of physical qubits n is immediately obvious. The other two central properties (number k of logical qubits and distance d) are not so obvious, but can be deduced directly from the stabilizer.

Proposition 3.3. *If the stabilizer S on n physical qubits has $|S| = 2^r$ (i.e., S has r generators), then $\mathcal{T}(S)$ has dimension 2^{n-r} . That is, there are $k = n - r$ logical qubits.*

Intuitively the result is easy to understand. The first generator of the stabilizer divides the Hilbert space into a +1 eigenspace and a -1 eigenspace of equal size. The codewords live in the +1 eigenspace. Each additional generator divides the subspace available for codewords in half again, so the total available dimension is 2^{n-r} . Non-generators do not divide the space since their eigenvalues can be derived by looking at the eigenvalues of the generators. This argument is not a proof, since we would need to show that the additional generators divide not just the full Hilbert space in half, but also all the smaller subspaces defined by the earlier generators. The actual proof is straightforward, but less intuitive:

Proof. The dimension of $\mathcal{T}(\mathbf{S})$ is equal to the trace of $\Pi_{\mathbf{S}}$, the projection operator onto $\mathcal{T}(\mathbf{S})$. Now,

$$\text{tr } \Pi_{\mathbf{S}} = \frac{1}{2^r} \sum_{M \in \mathbf{S}} \text{tr } M. \quad (3.13)$$

However, $\text{tr } P = 0$ for all Paulis P except for the identity. Thus,

$$\text{tr } \Pi_{\mathbf{S}} = \frac{\text{tr } I}{2^r} = \frac{2^n}{2^r}. \quad (3.14)$$

□

You can check the nine-qubit code in this formula: 9 physical qubits, 8 stabilizer generators, and 1 encoded qubit. For the five-qubit code, we have four generators, so again there should be 1 encoded qubit, with the basis codewords we saw above.

There's one special case which is not, strictly speaking, a QECC, but is interesting nonetheless. Yes, it's true, there are some things which are interesting other than quantum error correction. When the stabilizer has n generators on n qubits, proposition 3.3 would tell us we have 0 encoded qubits, which is a Hilbert space of dimension 1. That is, it is a single state, up to normalization.

Definition 3.6. A *stabilizer state* is the code space of a stabilizer with n generators on n qubits.

Since $r = n$ is the maximum number of generators you can have, a stabilizer state is an extreme limit of a QECC. Some constructions of QECCs will alter the number of encoded qubits from another code, and sometimes a stabilizer state can be the starting or ending point of such a construction. Also, stabilizer states are fairly common in the theory of quantum information, even discounting states arising from quantum error correction. For instance, the GHZ state $|000\rangle + |111\rangle$ and the Bell states $|00\rangle \pm |11\rangle$, $|01\rangle \pm |10\rangle$ are stabilizer states.

To understand how to deduce the distance of a stabilizer code, let us go back to the nine-qubit code and recall how it handled errors. The generators of the stabilizer were used to give us bits of the error syndrome. In particular, some generators were able to signal the presence of certain errors while missing other errors, but by looking at the full set of generators, we were able to identify all of the single-qubit errors.

I mentioned that the property responsible for determining whether a generator M is useful for an error E is anticommutation. Let us see how this works. Suppose $M \in \mathbf{S}$ and $E \in \mathbf{P}_n$ anticommutes with M . Then for any $|\psi\rangle \in \mathcal{T}(\mathbf{S})$,

$$M(E|\psi\rangle) = -EM|\psi\rangle = -E|\psi\rangle. \quad (3.15)$$

$|\psi\rangle$ was a +1 eigenvector of M — that is the definition of the code space — but $E|\psi\rangle$ is a -1 eigenvector.

Conversely, if E commutes with M then

$$M(E|\psi\rangle) = EM|\psi\rangle = E|\psi\rangle. \quad (3.16)$$

One advantage to dealing with the Pauli group is that these are the only choices. Either M and E commute or they anticommute. If we have an error that commutes with M , M retains a +1 eigenvalue, whereas if the error anticommutes with M , M 's eigenvalue becomes -1, signaling that an error has occurred.

Definition 3.7. The *normalizer* $\mathbf{N}(\mathbf{S})$ of the stabilizer \mathbf{S} is

$$\mathbf{N}(\mathbf{S}) = \{N \in \mathbf{P}_n \mid NM = MN \ \forall M \in \mathbf{S}\}. \quad (3.17)$$

If you know some group theory, you might recognize this as the definition of the *centralizer* of \mathbf{S} (the set of things that commute with all elements of \mathbf{S}) rather than the *normalizer* (the set of things that preserve \mathbf{S} under conjugation), but because Paulis either commute or anticommute and $-I \notin \mathbf{S}$, they are the same thing for a stabilizer. I am choosing to call it the normalizer rather than the centralizer because the normalizer relates to logical operations, and this is an important function of $\mathbf{N}(\mathbf{S})$, as we shall see shortly in section 3.4.2.

Since the stabilizer is Abelian, $S \subseteq N(S)$ always. Indeed, $N(S)$ also contains $-S$ and $\pm iS$. Since we worry about eigenvectors of S , changing the sign of an operator in the stabilizer is important, but $N(S)$ is about errors, and global phase no longer matters. Therefore, we will usually want to work with $\hat{N}(S)$.

The normalizer tells us which errors can be detected by the stabilizer code. If a Pauli E is outside the normalizer $N(S)$, then $P|\psi\rangle$ has an eigenvalue -1 for some $M \in S$, and thus is detected by measuring the eigenvalues of the stabilizer elements. Note that this is true for *any* codeword $|\psi\rangle$. Also note that it is sufficient to measure the generators of the stabilizer: If N commutes with M_1 and M_2 , it also commutes with M_1M_2 . Thus, if N commutes with all generators of S , then $N \in N(S)$.

When $E \in N(S)$, then $E|\psi\rangle$ has eigenvalue $+1$ for all $M \in S$, and therefore $E|\psi\rangle \in \mathcal{T}(S)$ for codewords $|\psi\rangle$. You might think that that means it is an undetectable error, but there is actually another class of Paulis that is “detectable” by definition 2.7. If $E \in S$, then, while it’s true that $E|\psi\rangle \in \mathcal{T}(S)$ for any codeword $|\psi\rangle$, it’s also true that $E|\psi\rangle = |\psi\rangle$, so E is not actually an “error”: it acts like the identity on the code space, leaving codewords unchanged. Ignoring global phases, we can say the same if $\hat{E} \in \hat{S}$.

Putting this together, we get a characterization of the detectable errors for a stabilizer code.

Theorem 3.4. *The set of undetectable errors for a stabilizer code S is $\hat{N}(S) \setminus \hat{S}$. The distance of S is $\min\{\text{wt } E | \hat{E} \in \hat{N}(S) \setminus \hat{S}\}$.*

The slanty line is a “set minus” operation. That is, $\hat{N}(S) \setminus \hat{S}$ consists of those elements of $\hat{N}(S)$ that are not in \hat{S} .

Proof. We can prove the first statement directly from the definition of the set of detectable errors (definition 2.7). The second statement will then follow immediately from the definition of distance. We need to consider three cases for E . In cases 1 and 2, $|\psi\rangle$ and $|\phi\rangle$ are arbitrary codewords.

1. Case 1: $\hat{E} \in \hat{S}$. Then for some choice of phase, $E \in S$ and $E|\phi\rangle = |\phi\rangle$, so

$$\langle \psi | E | \phi \rangle = \langle \psi | \phi \rangle, \quad (3.18)$$

and $c(E) = 1$. \hat{E} is detectable.

2. Case 2: $\hat{E} \notin \hat{N}(S)$. Then $\exists M \in S$ such that $\{M, E\} = 0$ (for any choice of phase of E), so $M(E|\phi) = -E|\phi\rangle$, as per equation (3.15). Then

$$\langle \psi | E | \phi \rangle = \langle \psi | M E M | \phi \rangle = -\langle \psi | E | \phi \rangle = 0, \quad (3.19)$$

since $M^2 = I$ for stabilizer elements. Thus $c(E) = 0$ and E is detectable.

3. Case 3: $\hat{E} \in \hat{N}(S) \setminus \hat{S}$. Since $E \notin S$, \exists codeword $|\phi\rangle$ such that $E|\phi\rangle \neq |\phi\rangle$. Let $|\psi\rangle = E|\phi\rangle$. Since $E \in N(S)$, $|\psi\rangle$ is also a codeword. However,

$$\langle \psi | E | \phi \rangle = 1 = \frac{1}{\langle \psi | \phi \rangle} \langle \psi | \phi \rangle, \quad (3.20)$$

whereas

$$\langle \phi | E | \phi \rangle = \langle \phi | \psi \rangle = (\langle \phi | \psi \rangle) \langle \phi | \phi \rangle. \quad (3.21)$$

Comparing equations (3.20) and (3.21) to definition 2.7 tells us that \hat{E} is not detectable unless $|\langle \phi | \psi \rangle| = 1$, meaning $E|\phi\rangle = e^{i\theta}|\phi\rangle$. Furthermore, by equation (3.21), \hat{E} is undetectable unless θ is the same for all $|\phi\rangle$. But that is not possible, since that would imply $\hat{E} \in \hat{S}$.

□

The key point here is that when $E \in \hat{N}(S) \setminus \hat{S}$, then E maps some codewords to *different* codewords. That’s the essence of an undetectable error, because the code has no way of knowing if a codeword is the original encoding of the state or the result of the action of the error on a different codeword.

Moving now to correcting errors, we find

X	X	X	X
Z	Z	Z	Z

Table 3.3: The stabilizer for the four-qubit code.

Theorem 3.5. *The stabilizer code \mathcal{S} corrects a set of errors $\mathcal{E} \subseteq \mathcal{P}_n$ iff $\hat{E}^\dagger \hat{F} \notin \hat{\mathcal{N}}(\mathcal{S}) \setminus \hat{\mathcal{S}}$ for all $E, F \in \mathcal{E}$.*

The theorem follows immediately from theorem 3.4 by comparing definition 2.7 with theorem 2.7. Recall that by the linearity of QECCs, and particularly corollary 2.5, it suffices to consider Pauli errors to understand the error-correcting capabilities of the code, at least when we are interested in correcting t -qubit errors.

For stabilizer codes, we have a slightly different notation than the more general $((n, K, d))$ notation that applies to arbitrary QECCs. Since the encoded subspace has a dimension that's always a power of 2, we write the code in terms of the number of encoded qubits k rather than the dimension $K = 2^k$ of the logical Hilbert space. We also use square brackets to indicate that we are dealing with a stabilizer code.

Notation 3.8. A stabilizer code with n physical qubits, k logical qubits, and distance d is denoted as an $[[n, k, d]]$ code. If the distance is unknown or irrelevant, it is an $[[n, k]]$ code.

We know that the nine-qubit code corrects a single-qubit error, so it must have distance at least 3. In fact, there are some 3-qubit Paulis (such as $X_1 X_2 X_3$) in $\hat{\mathcal{N}}(\mathcal{S}) \setminus \hat{\mathcal{S}}$ for the code, so the distance is exactly 3. Thus, the nine-qubit code is a $[[9, 1, 3]]$ code. The five-qubit code also turns out to have distance 3, so it is a $[[5, 1, 3]]$ code. Note that these codes are also $((9, 2, 3))$ and $((5, 2, 3))$ codes, but the more specific notation for a stabilizer code is generally used for them. As it happens, the five-qubit code is the smallest distance 3 code. You'll see the proof of that, as well as other techniques for proving limits on QECCs, in chapter 7.

You might wonder about distance 2 codes, since all the QECCs we have seen so far have distance 3. Distance 2 codes tend to be much simpler, but they can still be interesting. After all, a distance 2 code can detect one error or correct one erasure. The smallest distance 2 code is a $[[4, 2, 2]]$ code given in table 3.3. It is not hard to see that any one-qubit Pauli will anticommute with one or both of the two generators, but there are some two-qubit Paulis (e.g., $X \otimes X \otimes I \otimes I$) that commute with both. Thus, the code has distance 2.

3.3.4 Degeneracy and Stabilizer Codes

When is a stabilizer code degenerate? In the proof of theorem 3.4, we implicitly calculated the matrix C_{ab} . In particular, we found that $C_{ab} = 0$ if $E^\dagger F \notin \mathcal{N}(\mathcal{S})$ and $C_{ab} = 1$ if $E^\dagger F \in \mathcal{S}$. Any set of Paulis is linearly independent, so we just need to check if the resulting C_{ab} has maximum rank.

If $E_1^\dagger E_2 \in \mathcal{S}$, then $E_1^\dagger F \in \mathcal{S} \Leftrightarrow E_2^\dagger F \in \mathcal{S}$, so in this case the rows (or columns) of C_{ab} associated with E_1 and E_2 are the same. When $E_1^\dagger E_2 \notin \mathcal{S}$, then only one (or neither) of $E_1^\dagger F$ and $E_2^\dagger F$ can be in \mathcal{S} , so in this case the rows associated with E_1 and E_2 cannot have 1s in the same column. Thus, we get the following proposition:

Proposition 3.6. *A stabilizer code \mathcal{S} is degenerate for the error set \mathcal{E} iff $\exists E_1, E_2 \in \mathcal{E}$ with $E_1^\dagger E_2 \in \mathcal{S}$.*

As with general QECCs, we can define degeneracy as a property of the code subspace without referring to specific set of errors by considering the distance. In the case of stabilizer codes, motivated by the above analysis, we can also sensibly define degeneracy when the code is used for error detection rather than error correction.

Definition 3.9. A stabilizer code \mathcal{S} with distance d is *degenerate* if $\exists M \in \mathcal{S}$, $M \neq I$, with $\text{wt } M < d$.

Note that this is slightly more general than the generic notion of a degenerate code. A stabilizer code with distance $d = 2t + 2$ can be degenerate if \mathcal{S} contains any elements of weight $2t + 1$, but degeneracy wasn't defined for general non-stabilizer QECCs of even distance.

Looking once more at the stabilizer of the nine-qubit code (table 3.1), it is immediately obvious that it is a degenerate code. There are many generators of weight 2 and the code has distance 3. However, it is not always obvious from looking at the generators whether a stabilizer code is degenerate or not. It may be that the set of generators given to you all have high weight, but there is some product of the generators that does not. The five-qubit code is non-degenerate because all of the operators in the stabilizer have weight 4, but to see that you need to do some work.

3.4 Cosets and Error Syndromes

3.4.1 The Error Syndrome and the Stabilizer

For the nine-qubit code, we figured out the stabilizer by thinking about what we wanted to measure for the error syndrome. Each generator corresponded to a bit of the error syndrome. In fact, this is completely general:

Definition 3.10. Suppose \mathbf{S} is a stabilizer code with generators M_1, \dots, M_r and $|\psi\rangle$ is an eigenvector (not necessarily with eigenvalue +1) of all $N \in \mathbf{S}$. The *error syndrome* of $|\psi\rangle$ is an r -bit string with i th bit 0 if $|\psi\rangle$ is a +1 eigenvector of M_i and i th bit 1 if $|\psi\rangle$ is a -1 eigenvector of M_i . The error syndrome $\sigma(E) : \mathbb{P}_n \rightarrow \mathbb{Z}_2^r$ is the error syndrome of $E|\phi\rangle$ for any codeword $|\phi\rangle$. When the stabilizer associated with a syndrome is in doubt, we will write $\sigma_{\mathbf{S}}(E)$.

If $P, Q \in \mathbb{P}_n$, then $c(P, Q) = 0$ if P and Q commute and $c(P, Q) = 1$ if they anticommute. ($c(P, Q) : \mathbb{P}_n \times \mathbb{P}_n \rightarrow \mathbb{Z}_2$.)

For a Pauli, bit i of the error syndrome is the same as $c(E, M_i)$. Note that the syndrome of E will be the same no matter what codeword $|\psi\rangle$ we choose to evaluate. This follows from equations (3.15) and (3.16) and means that Pauli errors map the code space, which is a subspace with +1 eigenvalue for all stabilizer elements, to a subspace that has a different set of eigenvalues, but for which each eigenvalue is still shared for all states in the subspace.

The subspaces derived this way are also error-correcting codes, but they are associated with different eigenvalues of the stabilizer. We can reinterpret them as traditional stabilizer codes (with all +1 eigenvalues) by replacing M_i with $-M_i$ whenever the i th syndrome bit is 1. Note that there are 2^r possible values of the error syndrome, and each subspace is isomorphic to the code space, which has dimension 2^k . When there are n physical qubits, $r = n - k$. Also, all the subspaces of different error syndromes are orthogonal to each other, since they have different eigenvalues. Thus, the full Hilbert space \mathcal{H}_{2^n} decomposes as a direct sum of the 2^{n-k} subspaces associated with different syndromes.

Proposition 3.7. *The $c(P, Q)$ and $\sigma(E)$ functions have the following properties. In all cases, + refers to binary addition.*

- a) $c(P_1 P_2, Q) = c(P_1, Q) + c(P_2, Q)$.
- b) $c(P, Q) = c(Q, P)$.
- c) $\sigma(EF) = \sigma(E) + \sigma(F)$.

These are straightforward to verify.

The various error syndromes are associated with different cosets of the normalizer in \mathbb{P}_n .

Proposition 3.8. *Let $E, F \in \mathbb{P}_n$, and \mathbf{S} be a stabilizer. Then F and E are in the same coset of $\mathbf{N}(\mathbf{S})$ iff E and F have the same error syndrome.*

Proof.

\Rightarrow : If E and F are in the same coset, then $F = EN$, with $N \in \mathbf{N}(\mathbf{S})$. Let $M \in \mathbf{S}$. Then $c(F, M) = c(E, M) + c(N, M)$. But $N \in \mathbf{N}(\mathbf{S})$, so $c(N, M) = 0$. Therefore the error syndromes of E and F are the same.

syndrome 00 no correction	S	\bar{X}	FS	$F\bar{X}$	syndrome 10 correct as F
	\bar{Z}	\bar{Y}	$F\bar{Z}$	$F\bar{Y}$	
syndrome 01 correct as E	ES	$E\bar{X}$	GS	$G\bar{X}$	syndrome 11 correct as G
	$E\bar{Z}$	$E\bar{Y}$	$G\bar{Z}$	$G\bar{Y}$	

Figure 3.2: Cosets of the normalizer and stabilizer in the Pauli group. The errors E , F , and G are chosen as canonical errors for the syndromes 01, 10, and 11.

\Leftarrow : The argument works the other way too: Let $N = E^\dagger F$. If the error syndromes of E and F are the same, then $c(N, M) = 0$ for all $M \in \mathbf{S}$. Therefore, $N \in \mathbf{N}(\mathbf{S})$. \square

In much the same way as the full Hilbert space can be written as a direct sum of subspaces associated to different syndromes, the whole Pauli group can be partitioned into cosets of $\mathbf{N}(\mathbf{S})$, each associated with a different syndrome. All the cosets are the same size and there are 2^r of them. As a consequence, we know the size of the normalizer:

Proposition 3.9. $|\mathbf{N}(\mathbf{S})| = 4 \cdot 2^{n+k}$ when \mathbf{S} has n physical qubits and k logical qubits.

3.4.2 Cosets inside the Normalizer

We can similarly analyze the cosets of \mathbf{S} in $\mathbf{N}(\mathbf{S})$.

Proposition 3.10. Let $N_1, N_2 \in \mathbf{N}(\mathbf{S})$. Then N_1 and N_2 are in the same coset of \mathbf{S} iff $N_1|\psi\rangle = N_2|\psi\rangle$ for all codewords $|\psi\rangle$.

Proof. $N_1|\psi\rangle = N_2|\psi\rangle$ iff $|\psi\rangle$ is a +1 eigenstate of $M = N_1^\dagger N_2$. This is true for all codewords $|\psi\rangle$ iff $M \in \mathbf{S}$. However, N_1 and N_2 are in the same coset of \mathbf{S} iff $N_2 = N_1 M$ with $M \in \mathbf{S}$. \square

Recall that when $N \in \mathbf{N}(\mathbf{S})$, then $N|\psi\rangle$ is a codeword for any input codeword $|\psi\rangle$. Thus, N is a *logical operation*: it always maps codewords to (possibly different) codewords. Since the different representatives of a coset of \mathbf{S} act the same way on codewords, they are all different realizations of the *same* logical operation. Thus, the set $\mathbf{N}(\mathbf{S})/\mathbf{S}$ is of particular interest, since it consists of the distinct logical operations performed by the normalizer. (\mathbf{S} is a normal subgroup of $\mathbf{N}(\mathbf{S})$ by definition, so we can take this quotient without any difficulties.)

Theorem 3.11. Let \mathbf{S} be a stabilizer with n physical qubits and k logical qubits. Then $\mathbf{N}(\mathbf{S})/\mathbf{S} \cong \mathbf{P}_k$.

The quotient group $\mathbf{N}(\mathbf{S})/\mathbf{S}$ can be taken to be the *logical Pauli group*, performing Paulis on the encoded qubits. The proof of theorem 3.11 will be in section 3.5.

We can also look at cosets of \mathbf{S} in the full Pauli group \mathbf{P}_n , as in figure 3.2. Each coset of $\mathbf{N}(\mathbf{S})$ breaks up into cosets of \mathbf{S} , so we can associate each coset of \mathbf{S} in \mathbf{P}_n with an error syndrome, inherited from the coset of $\mathbf{N}(\mathbf{S})$ which it sits within. If we pick a particular representative E of the coset of $\mathbf{N}(\mathbf{S})$, then we can again identify the coset of \mathbf{S} with a logical operation $\bar{P} \in \mathbf{N}(\mathbf{S})/\mathbf{S}$; the interpretation of the coset is then a combination of the logical operation \bar{P} and the error E .

When we perform decoding on a stabilizer code, we do just this. For each error syndrome s , we assign a particular error E_s with $\sigma(E_s) = s$. If our syndrome measurement gives s , we assume the error actually

	X	Z	Z	X	I
	I	X	Z	Z	X
	X	I	X	Z	Z
	Z	X	I	X	Z
\overline{X}	X	X	X	X	X
\overline{Z}	Z	Z	Z	Z	Z

Table 3.4: The generators for the five-qubit code supplemented by representatives for logical Paulis.

	X	X	X	X
	Z	Z	Z	Z
\overline{X}_1	X	X	I	I
\overline{X}_2	X	I	X	I
\overline{Z}_1	I	Z	I	Z
\overline{Z}_2	I	I	Z	Z

Table 3.5: The stabilizer for the four-qubit code supplemented by representatives for logical Paulis.

was E_s and correct that. If the error was actually some E' with the same syndrome, we hope that E_s and E' are in the same coset of S . If not, there has been a logical Pauli error, the one associated with the actual coset of S in which E' lies.

Theorem 3.12. *If S is a non-degenerate stabilizer code, the error syndromes of all errors in the correctable error set \mathcal{E} are distinct. If S is a degenerate code, E and F have the same error syndrome iff $\hat{E}^\dagger \hat{F} \in \hat{S}$.*

Proof. By proposition 3.8, two errors $E \neq F \in \mathcal{E}$ have the same error syndrome iff they are in the same coset of $N(S)$, meaning $E^\dagger F \in N(S)$. But both errors are correctable, and by theorem 3.5, $\hat{E}^\dagger \hat{F} \notin \hat{N}(S) \setminus \hat{S}$, so $E^\dagger F \in N(S)$ iff $\hat{E}^\dagger \hat{F} \in \hat{S}$. Thus, the error syndromes of two correctable errors E and F are the same iff $\hat{E}^\dagger \hat{F} \in \hat{S}$. If this ever occurs, the code is degenerate. \square

Applying this theorem is one way to see that the five-qubit code has distance 3 and is non-degenerate. There are 15 possible one-qubit errors (X , Y , Z on each of five qubits), plus the identity. There are 4 generators, so there are 16 error syndromes. If you list the syndromes of the 16 errors, you'll find that each one is unique. There are no syndromes left over, so the five-qubit code is known as a *perfect* code.

It is frequently helpful to also pick representatives of the cosets of S in the normalizer. These representatives don't have an interpretation in error correction, but they do give us concrete realizations of the logical Pauli group, which is helpful in fault-tolerant computation. For now, they are simply convenient computational tools. We don't actually need to pick representatives for every coset of S . Since $N(S)/S$ has a group structure itself, it is sufficient to pick representatives of generators of $N(S)/S$ and let the generators of other cosets be determined by multiplication. In other words, we can pick representatives for \overline{X}_i and \overline{Z}_i for $i = 1, \dots, k$. (The subscript i here means the operator acts on the i th logical qubit, not the i th physical qubit.) I have done this for the five-qubit code and the four-qubit code in tables 3.4 and 3.5. Then, for instance, you can get \overline{Y} for the five qubit code to be $i\overline{X}\overline{Z} = Y \otimes Y \otimes Y \otimes Y \otimes Y$.

You need to be careful when assigning cosets to elements of the logical Pauli group. Remember, $N(S)/S \cong P_k$, and we'd like to realize this through a choice of representatives for the generating cosets. This means that the commutation and anticommutation relationships of the logical Paulis must be realized through the coset representatives. For instance, the coset representatives for \overline{X}_i and \overline{Z}_i must anticommute, while the coset representative of \overline{X}_i must commute with the representatives for \overline{X}_j or \overline{Z}_j ($j \neq i$).

We could in principle do the same thing for $P_n/N(S)$, choosing representative errors only for the basis error syndromes and deducing the other errors by multiplication. However, in most cases, this is not a good thing to do. If our goal is to correct t errors for the maximum possible t , what we'd like to do instead is

choose the lowest weight error in each coset of $N(S)$ as its representative. If we rely on multiplication to tell us the coset representatives, we'll frequently get errors of higher weight than necessary, and then there will be some lower weight errors that won't get corrected properly.

3.4.3 Maximum Likelihood Decoding

I'd like to depart briefly from the convention of most of this chapter, for which we had some fixed set \mathcal{E} of errors that we'd like to correct, and we won't settle for anything less than correcting all of them. For the purposes of this subsection, assume we instead have some n -qubit Pauli channel, with $P \in \mathcal{P}_n$ occurring with probability p_P .

Given a stabilizer code S , we'd like to find a decoding procedure which minimizes the probability of error when the code goes through a Pauli channel. As discussed above, a decoding procedure can be understood as assigning to each error syndrome s a Pauli Q_s . When we measure s , we assume the actual error was Q_s and so perform the correction by inverting that error. If the actual error was something else, we may end up with the wrong state. In particular, if the true error P was in a different coset of S than Q_s was, we'll end up with some logical Pauli operation.

Thus, we can calculate the probability of logical error by summing over the cosets of S and $N(S)$.

Proposition 3.13. *Let C_s be the coset of $\hat{N}(S)$ with error syndrome s . $\hat{Q}_s\hat{S}$ is the coset of \hat{S} containing the canonical error \hat{Q}_s with syndrome s , so $\hat{Q}_s\hat{S} \subseteq C_s$.*

a) *The probability of error syndrome s (regardless of whether we correctly identify the error) is*

$$p_s = \sum_{\hat{P} \in C_s} p_P. \quad (3.22)$$

b) *The probability of having syndrome s and no logical error (i.e., error correction is successful) is*

$$p_{s,\text{OK}} = \sum_{\hat{P} \in \hat{Q}_s\hat{S}} p_P. \quad (3.23)$$

c) *The total probability of successfully decoding the state is*

$$p_{\text{OK}} = \sum_s \sum_{\hat{P} \in \hat{Q}_s\hat{S}} p_P. \quad (3.24)$$

There is no interaction between the different error syndromes. If we change Q_s for one value of s , the only change to p_{OK} comes through the change to $p_{s,\text{OK}}$. We can therefore treat each syndrome separately to maximize $p_{s,\text{OK}}$. It doesn't matter which element of the coset Q_sS we choose, only which coset we choose within C_s .

In order to maximize the probability of successfully decoding, for each syndrome s , we should choose the coset Q_sS which maximizes $p_{s,\text{OK}}$. As given by equation (3.23), that just means we look at each coset and add up the probability of all Paulis in the coset. Then we choose the coset with the highest value to be the representative coset for s .

The maximum likelihood decoder should be contrasted with the decoder that optimizes the distance, for which we choose the coset containing the lowest-weight Pauli. If we have an independent Pauli channel, the two procedures frequently, but not always, give the same answer: When the probability of error per qubit is small, a specific low-weight error will have higher probability than a specific high-weight error. However, a coset which contains one low-weight error and a bunch of high-weight errors may be less likely than a coset which contains a large number of medium-weight errors.

Getting the *exactly* optimal decoder, in the sense of always picking the most likely coset, is usually a computationally challenging task. Often, therefore, we are willing to settle for a good approximate decoder that gets p_{OK} to be close to the optimal value but with a much smaller computational cost.

In the Pauli group	Binary symplectic representation
$P \in \mathcal{P}_n$	$v_P = (x_P z_P) \in \mathbb{Z}_2^n \times \mathbb{Z}_2^n$
Multiplication	Addition
$c(P, Q)$	$v_P \odot v_Q$
Phase	No equivalent
Stabilizer \mathcal{S}	Weakly self-dual subspace \mathcal{S}
Normalizer $\mathcal{N}(\mathcal{S})$	Dual subspace \mathcal{S}^\perp (under \odot)
Minimal set of generators for \mathcal{S}	Basis for \mathcal{S}

Table 3.6: Equivalence between the Pauli group and its binary symplectic representation.

3.5 Binary Symplectic Representation

3.5.1 The Symplectic Representation of Paulis

One of the most useful techniques when dealing with stabilizer codes is to ignore the phases of Paulis and work with $\hat{\mathcal{P}}_n$ instead of \mathcal{P}_n . We lose a couple of things by doing this. One is the ability to define the code space, which seems like a serious loss, but isn't really — as discussed above, the other eigenspaces of the stabilizer have the same error correction properties as the code space. The other missing thing is the ability to tell whether Paulis commute or anticommute, since $\hat{\mathcal{P}}_n$ is an Abelian group. This is a much greater loss, so we'll need to find a substitute.

The structure of $\hat{\mathcal{P}}_n$ is very straightforward. There are 4^n elements, and all pairs commute under multiplication. Therefore, $\hat{\mathcal{P}}_n \cong \mathbb{Z}_2^{2n}$. The group operation for \mathbb{Z}_2^{2n} is usually written as addition. Conventionally, we choose to represent elements of $\hat{\mathcal{P}}_n$ as two n -bit binary vectors, one representing the “ X ” component and one representing the “ Z ” component.

Definition 3.11. The *binary symplectic representation* of $\hat{\mathcal{P}}_n$ is an isomorphism between $\hat{\mathcal{P}}_n$ and $\mathbb{Z}_2^n \times \mathbb{Z}_2^n$: $P \leftrightarrow v_P = (x_P|z_P)$. The i th component of x_P is 0 if P acts on qubit i as I or Z and 1 if P acts on qubit i as X or Y . The i th component of z_P is 0 if P acts on qubit i as I or X and 1 if P acts on qubit i as Y or Z . That is,

$$\begin{array}{ccc}
 I & & (0|0) \\
 X & \longleftrightarrow & (1|0) \\
 Y & & (1|1) \\
 Z & & (0|1)
 \end{array} \tag{3.25}$$

For instance, $X \otimes I \otimes Y \leftrightarrow (1\ 0\ 1|0\ 0\ 1)$. The “symplectic” refers to the substitute for commutativity/anticommutativity:

Definition 3.12. The *symplectic form* (or symplectic product) on $\mathbb{Z}_2^n \times \mathbb{Z}_2^n$ is $(x_1|z_1) \odot (x_2|z_2) = x_1 \cdot z_2 + x_2 \cdot z_1$, where the dot product is the usual scalar product in the vector space \mathbb{Z}_2^n and addition is modulo 2.

Proposition 3.14. $v_P \odot v_Q = c(P, Q)$ for $P, Q \in \mathcal{P}_n$.

Proof. This can be checked exhaustively for a single qubit. For more than one qubit, note that both \odot and $c(\cdot, \cdot)$ are equal to the parity of their single-qubit results. Therefore, since they are equal for single qubits, they are also equal for n qubits. \square

Combining the symplectic form with the binary symplectic representation recovers the primary structure of the Pauli group. We can switch back and forth between the two representations to use whichever one is the most convenient at the moment. The conversion process is summarized in table 3.6. The one thing that is lost in the transition is the phase of a Pauli, so you must be careful whenever dealing with something that depends on that.

Some of the conversions need a little more explanation. There were three conditions for \mathcal{S} to be a stabilizer: that $-I \notin \mathcal{S}$, that \mathcal{S} is a group, and that \mathcal{S} is Abelian. Since phase is lost when converting to

$$\begin{array}{c} \overline{X} \\ \overline{Z} \end{array} \left(\begin{array}{ccccc|ccccc} 1 & 0 & 0 & 1 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 1 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 1 \\ \hline 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 \end{array} \right)$$

Table 3.7: The binary symplectic representation of the stabilizer and logical Pauli operators for the five-qubit code.

the binary symplectic representation, the first condition is irrelevant. The second condition means that S becomes a linear subspace in $\mathbb{Z}_2^n \times \mathbb{Z}_2^n$, and the third means that $v \odot w = 0 \forall v, w \in S$. This analysis and the conversion of the normalizer prompts the following definition:

Definition 3.13. Let V be a linear subspace of $\mathbb{Z}_2^n \times \mathbb{Z}_2^n$. The *dual* of V (with respect to \odot) is $V^\perp = \{w \in \mathbb{Z}_2^n \times \mathbb{Z}_2^n \mid w \odot v = 0 \forall v \in V\}$. A subspace is *self-dual* if $V^\perp = V$. A subspace is *weakly self-dual* if $V \subseteq V^\perp$.

Table 3.7 gives the five-qubit code in binary symplectic representation, including the logical \overline{X} and \overline{Z} operators.

Conversion can go the other way too. For instance, it is convenient to define a notion of a set of Paulis being independent based on the standard definition for binary vector spaces:

Definition 3.14. A set of Paulis $\{P_1, \dots, P_m\}$ is *independent* iff the vectors $\{v_{P_1}, \dots, v_{P_m}\}$ are linearly independent.

A set of Paulis is independent unless one of them is a product of others; this is equivalent to saying that a set of $2n$ -bit binary vectors is independent unless one of them is a sum of others. Since the binary vector space is $2n$ -dimensional, the maximum number of independent Paulis is $2n$, which is equal to the number of generators of the Pauli group.

3.5.2 Linear Algebra Lemma

The power of the binary symplectic representation is that we can use standard results from linear algebra to prove a number of properties of the Pauli group and stabilizers. In particular, we have the following lemma:

Lemma 3.15. Let $\{P_1, \dots, P_m\}$ be an independent set of Paulis on n qubits, and let s be an m -bit vector with components s_i . Then $\exists Q \in P_n$ with $c(P_i, Q) = s_i$. In fact, there are 2^{2n-m} such Paulis.

Proof. Converting to the binary symplectic representation gives us m linearly independent vectors v_i , and we wish to find vector w such that $v_i \odot w = s_i$. Each of these conditions is a linear equation, and since the vectors v_i are independent, the system of linear equations is non-singular. There are m equations in a $2n$ -dimensional vector space, so the space of solutions has dimension $2n - m$. It is a binary vector space, so that corresponds to 2^{2n-m} solutions. \square

An important consequence of the lemma is that it tells us about the decomposition of the full physical Hilbert space into subspaces associated with the different error syndromes. I've already discussed this decomposition, but the missing piece is given by the following corollary of lemma 3.15:

Proposition 3.16. Let S be a stabilizer with generators $\{M_i\}$. Then for any error syndrome s , $\exists P \in P_n$ with $\sigma(P) = s$.

That is, not only is every Pauli outside $N(S)$ associated with an error syndrome, but every possible error syndrome is associated with some Pauli.

3.5.3 Consequences of the Lemma

Now it's time for some of the proofs I owe you. There is proposition 3.2, which claims $S = S(\mathcal{T}(S))$, and theorem 3.11, which says $N(S)/S \cong P_k$.

First, here is the proof that $N(S)/S \cong P_k$:

Proof of theorem 3.11. The most straightforward way to show this is by sequentially picking coset representatives for the cosets corresponding to \overline{X}_i and \overline{Z}_i . The Pauli group P_k is determined, like any finitely-presented group, by its generators and relations between them. In the case of the Pauli group, the relations are

$$c(X_i, X_j) = 0 \tag{3.26}$$

$$c(Z_i, Z_j) = 0 \tag{3.27}$$

$$c(X_i, Z_j) = \delta_{ij}. \tag{3.28}$$

It is sufficient to consider only the generating cosets, since we can let the representatives of a product of cosets be the product of the representatives, as discussed in section 3.4.2. Provided we can choose \overline{X}_i and \overline{Z}_i with the correct commutation relations, that gives us an injective map of P_k into $N(S)/S$. We know that $|N(S)| = 4 \cdot 2^{n+k} = |S| |P_k|$, so an injection has to be isomorphism.

Suppose we've chosen some independent set of \overline{X}_i 's and \overline{Z}_i 's with the correct commutation relations, and we wish to choose one more. The new logical Pauli must be in $N(S)$, so in particular, it must commute with all generators of the stabilizer. Second, it has a defined commutation relation with the already chosen logical Paulis. By lemma 3.15, there exists a Pauli that satisfies all of these constraints.

The only remaining thing to check is that the new logical Pauli can be chosen to be independent of the prior ones. To simplify the analysis, let us first pick all the \overline{X}_i 's, then the \overline{Z}_i 's. When we are picking \overline{Z}_j , it satisfies different commutation relations than all previously selected logical Paulis. In particular, \overline{Z}_j anticommutes with \overline{X}_j , unlike all the previous logical Paulis and all the stabilizer generators. Thus, \overline{Z}_j must be independent.

When we pick \overline{X}_j , this argument doesn't apply, since the new one will commute with all stabilizer generators and all previous logical Paulis. However, there are $n - k$ stabilizer generators and at most $k - 1$ logical Paulis already chosen, for a total of at most $n - 1$ constraints. By lemma 3.15, there are thus at least 2^{n+1} possible solutions. The group generated by the stabilizer and previous logical Paulis contains only 2^{n-1} operators, so there are possible choices for \overline{X}_j that are independent of the previous choices. □

Finally, we can prove proposition 3.2, showing that $S = S(\mathcal{T}(S))$:

Proof of proposition 3.2. If $M \in S$, then certainly $M|\psi\rangle = |\psi\rangle$ for any $|\psi\rangle \in \mathcal{T}(S)$, so $S \subseteq S(\mathcal{T}(S))$. We need to show that if $N \notin S$, then $N \notin S(\mathcal{T}(S))$. If $N \notin N(S)$ then $N|\psi\rangle$ (for any codeword $|\psi\rangle$) has a different eigenvalue for some $M \in S$, and is therefore orthogonal to the code space, which in turn means $N \notin S(\mathcal{T}(S))$.

If S has n generators, that is all we need. Otherwise, if $r < n$, we still need to show that if $N \in N(S) \setminus S$, then $\exists |\psi\rangle \in \mathcal{T}(S)$ such that $N|\psi\rangle \neq |\psi\rangle$. By lemma 3.15, we can choose a $M \in N(S) \setminus S$ such that $\{M, N\} = 0$. Consider the modified stabilizer S' formed by adding M to S as a new generator. S' has $r + 1$ generators, so its code space has dimension 2^{k-1} (by proposition 3.3). Since $k \geq 1$, there is at least one state $|\psi\rangle$ (up to normalization) in the code space of S' . $|\psi\rangle$ is left fixed by all elements of S' , and in particular by the elements of $S \subset S'$. Thus, $|\psi\rangle \in \mathcal{T}(S)$. However, N anticommutes with M , which is in the stabilizer of $|\psi\rangle$. Therefore, $N|\psi\rangle$ has eigenvalue -1 for M , and in particular is orthogonal to $|\psi\rangle$. It follows that $N \notin S(\mathcal{T}(S))$, which proves the proposition. □

Chapter 4

Maybe I Should Have Started Here: Classical Error Correction

In chapter 3, we saw the formalism of stabilizer codes, but we didn't see how to find new stabilizer codes. Indeed, finding new codes is a tricky topic. Already in the theory of classical error correction, it is quite difficult to find good new codes. Luckily, there are a few ways of taking already-discovered classical codes and turning them into quantum codes. That is not the subject of this chapter.

This chapter is instead about the theory of classical error-correcting codes. Naturally, I won't have time to go into complete detail on classical error correction, so I will focus on the particular points of the theory of classical error correction which have most relevance to QECCs. One purpose is to give you the background for chapter 5, which *is* about converting classical codes into quantum codes. The other reason is that it can give you a deeper understanding of the theory of QECCs and of stabilizer codes, since there are many parallels — and some critical differences — between the theories of classical and quantum error correction.

For those who are familiar with the theory of classical error correction, this chapter is likely to be somewhat boring, but I hope it won't be a complete waste of time. In particular, I will try to point out the parallels between classical coding theory and those aspects of quantum coding theory that we've seen so far. Some you may have already noticed yourself, but perhaps there are some parallels you missed.

4.1 Classical Error Correction in General

4.1.1 Abstract Description of a Classical Code

A general classical error-correcting code can be defined, much as we did for a QECC, as an encoding map.

Definition 4.1. A classical *error-correcting code* (e, \mathcal{E}) is a map $e : [1 \dots K] \rightarrow [1 \dots N]$ with a set of correctable errors \mathcal{E} (consisting of maps $E : [1 \dots N] \rightarrow [1 \dots M]$) with the following property: \exists map $d : [1 \dots M] \rightarrow [1 \dots K]$ such that $\forall E \in \mathcal{E}, \forall x \in [1 \dots K]$,

$$d(E(e(x))) = x. \tag{4.1}$$

The map d is the *decoder* for the code and e is the *encoder* for the code. Frequently, the code is just referred to as C , the image of the encoder e in $[1 \dots N]$.

I have written this to be completely analogous to the definition I gave for a QECC. The only real difference is that quantum states live in a Hilbert space, and maps between them should be quantum operations, or at least linear maps, whereas classically, we can consider states from any set and arbitrary functions between the sets as the encoder, decoder, and errors. But remember that classical codes are actually just a special case of quantum codes. The apparent extra constraint of linearity is really an extra freedom to include superpositions of basis states, whereas classical codes can only contain basis states.

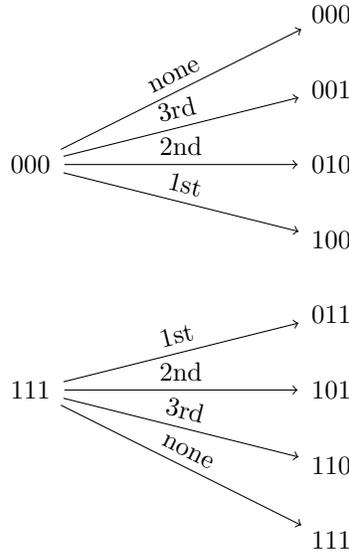


Figure 4.1: Codewords of the repetition code with errors on a single bit never get confused.

Ultimately, a classical error-correcting code is just a set of objects (frequently bit strings), and the errors diversify those objects. The goal of the decoder is to get back to the original object. We have an uncorrectable error when two different logical objects x and y get confused: $E(e(x)) = F(e(y))$. If that never happens, a decoder exists: the sets $S_x = \{E(e(x)) | E \in \mathcal{E}\}$ are all distinct, so we can define a function mapping all elements of S_x to x unambiguously.

This condition has no precise analogy for quantum codes. In some sense equation (2.58) is similar, if we interpret the correct quantum analogue of “different” as “orthogonal.” Equation (2.58) says that orthogonal quantum states and different errors must produce orthogonal states, but for QECCs, that is only a sufficient condition, not a necessary one. From that point of view, QECCs seem somewhat more generous than classical error-correcting codes, since they don’t require that it is possible to exactly determine the error. On the other hand, if you did exercise ??, you saw a condition which can be broken up into two parts. When $i \neq j$, it says that we should never confuse different quantum states under an error; that seems analogous to the condition for classical error correction. But then there is an additional non-trivial condition for a QECC when $i = j$, which makes the quantum code seem more stringent than the classical one. The extra condition can be interpreted as saying that the environment should not learn anything about the logical subspace for a QECC. Classically, of course, it is harmless if the environment simply learns the encoded data, provided it doesn’t change it. In the quantum case, learning about the data *necessarily* involves changing it.

4.1.2 Distance of a Classical Code

As with quantum codes, frequently we focus on codes where the encoded state N can be broken up into smaller pieces, $N = q^n$, and where $M = N$. We consider $[1 \dots N]$ as the direct product of n copies of $[1 \dots q]$. I will refer to each $[1 \dots q]$ factor as a *register*, a *component*, or a *coordinate*. In the common case where $q = 2$, I may just refer to the i th bit (or physical bit) of the code. Sometimes I may even forget myself and accidentally refer to the i th qubit. In the quantum case, I’ll use these same terms, but I try to avoid component since it also can refer to the component in some direction in the overall Hilbert space. In the classical case, the terminology comes from treating $[1 \dots q^n]$ as an n -dimensional vector space over the finite field $\text{GF}(q)$. We’ll get back to that approach in section 4.2.

Also, we tend to focus on cases where \mathcal{E} is composed of errors of weight t or less. The classical definitions

of *weight* and *support* are just the same as the quantum ones: the error E acts trivially on all but the coordinates in the support of E , and the weight of E is the size of the support. When we have a classical channel which causes errors independently on each coordinate with a small probability p , then the probability of having errors of weight $t + 1$ or greater is $O(p^{t+1})$, much as in the quantum case. If we can find a code that simply corrects all errors of weight t or less, then with high probability, we'll get the correct state when we decode.

Definition 4.2. The *distance* of a classical error-correcting code $C \subseteq \mathbb{Z}_N$ (with $N = q^n$) is

$$\min\{\text{wt}(E) \mid \exists x \neq y \in C, E(x) = y.\} \quad (4.2)$$

A classical code with distance d encoding K possible states in n registers of size q is an $(n, K, d)_q$ code. When $q = 2$, we just write (n, K, d) .

In other words, the distance is the minimum number of coordinates that have to be changed to get from one codeword to another. The quantum distance can be understood in the same way, but there are more subtleties in the quantum case.

As with quantum codes, the distance tells us how many errors we can correct:

Proposition 4.1. *A classical code with distance d can:*

1. correct general errors on up to $\lfloor (d - 1)/2 \rfloor$ coordinates,
2. correct erasure errors on up to $d - 1$ coordinates, or
3. detect errors on up to $d - 1$ coordinates.

4.2 Classical Linear Codes

General classical codes are unwieldy to deal with, just as are general quantum codes. The class of linear codes is very analogous to stabilizer codes — it is a much more tractable subclass of classical error-correcting codes, and many of the most interesting known codes are, in fact, linear codes.

For a linear code, we assume $N = q^n$, as above. For this section, I will further assume $q = 2$, and treat $[1 \dots N]$ as an n -dimensional vector space over \mathbb{Z}_2 . In section 4.4, I will return to the $q \neq 2$ case.

4.2.1 Generator Matrix and Parity Check Matrix

By considering $\mathbb{Z}_N = \mathbb{Z}_2^n$ as a vector space, we add an additional linear structure, and a linear code exploits that structure.

Definition 4.3. An error-correcting code $C \subseteq \mathbb{Z}_2^n$ is a *linear code* if $x, y \in C \Rightarrow x + y \in C$.

A linear code is thus a linear subspace of \mathbb{Z}_2^n . We can choose a basis x_1, \dots, x_k for the code, and all other vectors in the code will be linear combinations of the x_i 's. Since $x + x = 0 \forall x \in \mathbb{Z}_2^n$, the only question is which subset of the basis vectors are added together to make a codeword. There are thus 2^k codewords in total and k encoded bits.

Definition 4.4. The *generator matrix* G_C of a linear code C is a matrix with row i equal to x_i .

The generator matrix can be used to define the encoder of C :

Proposition 4.2. *Let $C \subseteq \mathbb{Z}_2^n$ be a linear code with k encoded bits and generator matrix G . Then the linear map $v \in \mathbb{Z}_2^k \mapsto G^T v \in \mathbb{Z}_2^n$ is the encoder for C . In other words, x is a codeword of C iff $x = G^T v$ for some $v \in \mathbb{Z}_2^k$.*

In order to check for errors, we measure the parities of bits in the codeword.

Definition 4.5. Suppose the linear code C has generator G . A *parity check matrix* H_C for C is a matrix with row i equal to $y_i \in \mathbb{Z}_2^n$, where $Gy_i = 0 \forall i$, and the set $\{y_i\}$ is a maximal linearly independent set with this property.

The parity check matrix of a code is not unique, but it is still frequently referred to as “the” parity check matrix. You can take any linear combinations of rows of the parity check matrix, and provided the new set of rows remains linearly independent, it will still function as a parity check matrix.

Theorem 4.3. *If C has k encoded bits and n physical bits, then G_C is a $k \times n$ matrix and H_C is an $(n-k) \times n$ matrix. $G_C H_C^T = 0$ and $H_C G_C^T = 0$.*

Proof. The only property which is not completely trivial is that H_C has $n - k$ rows. The constraints $G_C^T y_i = 0$ form a set of k non-singular linear equations on n bits, so the solution space has dimension $n - k$. Thus, a maximal set of $\{y_i\}$ will consist of $n - k$ of them. \square

Using the linear structure of \mathbb{Z}_2^n , we can represent a set of bit flip errors as a vector $e \in \mathbb{Z}_2^n$, with 1 for those bits which have been flipped and 0 for those bits which have not been flipped. Under this convention, $\text{wt } e$ is the weight of the error. If we start with bit string x and the error e occurs, we now have the bit string $x + e$.

The virtue of the parity check matrix is that it filters out the codewords and just tells us about the errors. Suppose $x \in C$ undergoes error e . Then, using linearity and proposition 4.2,

$$H_C(x + e) = H_C x + H_C e = H_C G_C^T v + H_C e = H_C e. \quad (4.3)$$

Definition 4.6. The *error syndrome* of an error e for linear code C is $H_C e$.

All of this probably sounds a bit familiar, since stabilizer codes and classical linear codes are closely related. The stabilizer of a stabilizer code is very analogous to the parity check matrix for a linear code, although the stabilizer also has some vague similarity to the generator matrix (in that the projector on the code space is formed from the stabilizer). In fact, linear codes are a special case of stabilizer codes, and the parity check matrix can be realized as the stabilizer:

Theorem 4.4. *Let C be a linear code with parity check matrix H , and let C correct the set of errors $\mathcal{E} \subseteq \mathbb{Z}_2^n$. Define a stabilizer code S to have stabilizer with binary symplectic representation $(0|H)$. Then S corrects the set of errors $\{(e|0) | e \in \mathcal{E}\}$ and for any $x \in \mathbb{Z}_2^n$, $x \in C$ iff $|x\rangle \in S$.*

In other words, we form a stabilizer by replacing the 1s in H with Z s, with each row of the parity check matrix becoming a generator of the stabilizer. The resulting stabilizer code corrects the same errors as C and has the same basis codewords. The one difference is that the stabilizer code encodes a quantum state, so superpositions such as $|x_1\rangle + |x_2\rangle$ are also codewords of S , even though the combination is meaningless when considering a classical code. The proof of the theorem is straightforward, and is left as an exercise.

The non-uniqueness of the parity check matrix is just the same as the non-uniqueness of the generators of a stabilizer. Replacing a row of the parity check matrix with a linear combination of rows is equivalent to replacing a generator of the stabilizer with a product of generators.

The best analogy for the rows of the generator matrix are the logical Paulis in $N(S)/S$. For a stabilizer code, the coset representatives are non-unique, which is not an issue for classical codes. It is true that the generator matrix is not completely unique — we can take a different encoder, which corresponds to a different generator matrix. However, this is a different phenomenon, and corresponds to labeling the cosets in $N(S)/S$ with different logical Paulis rather than choosing different representatives of them.

4.2.2 Distance of a Linear Code

For a linear code, the distance has a somewhat nicer form than for a general classical error-correcting code.

Proposition 4.5. *The distance of a linear code C is $\min\{\text{wt}(e) | e \in C, e \neq 0\}$.*

Proof. The general definition of distance is as $\min\{\text{wt } e \mid x \neq y \in C, x + e = y\}$ (with the error rewritten from definition 4.2 to take advantage of linearity). But when $x, y \in C$, $e = x + y$ is also in C since C is linear. Any $e \in C$ can be realized this way by simply choosing any $x \in C$ and letting $y = x + e$, which will automatically be in C , again since C is linear. $x \neq y$ is equivalent to $e \neq 0$. \square

Notation 4.7. An $(n, 2^k, d)$ linear code is denoted as an $[n, k, d]$ code.

For a linear code, the codewords themselves are the undetectable errors: adding a non-trivial codeword to another codeword results in a new codeword, and if you add a non-codeword to a codeword, you always get a non-codeword. For a set of errors to be correctable, any pair of errors must not add up to a codeword:

Theorem 4.6. A linear code C corrects the error set \mathcal{E} iff $e + f \notin C$ for all $e \neq f \in \mathcal{E}$. Equivalently, $H_C e \neq H_C f \forall e \neq f \in \mathcal{E}$ (i.e., all errors in \mathcal{E} have different error syndromes).

Proof. If all errors in \mathcal{E} have different error syndromes, then the code certainly can correct for \mathcal{E} using the following procedure: Given an erroneous codeword $x + e$, apply the parity check matrix to get $H_C e$. Since the error syndromes are unique, we can identify e and then recover $x = (x + e) + e$.

If $e + f \in C$ then $H_C(e + f) = H_C e + H_C f = 0$. Conversely, if $e + f \notin C$, then $H_C(e + f) \neq 0$. This is true because the parity check matrix is formed from a *maximal* set of vectors annihilated by the generator matrix. When $e + f \notin C$, then the set of solutions of $G_C y_i = 0$ and of $(e + f) \cdot y_i = 0$ is only $n - k - 1$ dimensional, whereas the parity check matrix has $n - k$ rows. Thus, $e + f \notin C$ iff $H_C e \neq H_C f$.

Finally, if $e + f \in C$ for some pair $e \neq f \in \mathcal{E}$, then the code cannot correct \mathcal{E} . For instance, given $x \in C$, let $y = x + (e + f)$, which will also be in C . Then $x + e = y + f$, so if this string shows up, there is no way to tell whether it should be decoded to x (with error e) or y (with error f). \square

The combination $e + f$ is reminiscent of the combination $E^\dagger F$ that appears in theorem 3.5 giving the set of correctable errors for a stabilizer code. Comparing further, you can see that for a classical code, $C \setminus \{0\}$ plays the role of $\hat{N}(S) \setminus \hat{S}$ for a stabilizer code. Indeed, if you apply the transformation theorem 4.4 to code C to get stabilizer S , and then calculate $\hat{N}(S) \setminus \hat{S}$, you find that it includes the conversion of $C \setminus \{0\}$. That's not all it contains, but it is a good exercise to work this out yourself.

Importantly, there is no good classical analogue to the concept of a degenerate quantum code. All classical codes are non-degenerate, and the presence of degenerate quantum codes complicates quantum coding theory.

4.2.3 Example: Hamming Codes

One example of a linear code is the repetition code $0 \mapsto 000, 1 \mapsto 111$. It is linear because $000 + 000 = 000$, $000 + 111 = 111$, and $111 + 111 = 000$. The generator matrix is

$$G = \begin{pmatrix} 1 & 1 & 1 \end{pmatrix}, \tag{4.4}$$

and the parity check matrix is

$$H = \begin{pmatrix} 1 & 1 & 0 \\ 1 & 0 & 1 \end{pmatrix} \tag{4.5}$$

The three-bit repetition code (a $[3, 1, 3]$ code) can of course be generalized to d -bit repetition codes, which have parameters $[d, 1, d]$, but there is another interesting generalization, starting from the observation that in the parity check matrix for the three-bit repetition code, every column is different.

Supposing we want to correct only 1-bit errors. If $\text{wt } e = 1$, with the only 1 in the i th coordinate, then for any parity check matrix H , He — the error syndrome of e — will be the i th column of H . Thus, if we pick all columns of H to be different, the error syndromes of all 1-bit errors will be different. We should not pick an all 0s column since we want the no error case ($e = 0$) to have error syndrome 0.

If we fix the number of rows of H to be r , we can choose up to $2^r - 1$ distinct columns by letting them run over all nonzero r -bit strings. This gives a Hamming code:

Definition 4.8. The $[2^r - 1, 2^r - r - 1, 3]$ Hamming code is the code whose $r \times (2^r - 1)$ parity check matrix has as columns all possible r -bit strings.

Notice that here we have defined the Hamming code by choosing a parity check matrix, much as for a stabilizer code, we choose the stabilizer to define the code. In this case, you can derive the set of actual codewords as $\{x | Hx = 0\}$. There are r linear constraints on $2^r - 1$ bits, so the linear space of solutions is $(2^r - r - 1)$ -dimensional, giving the number of encoded qubits in the definition.

As a concrete example, when $r = 3$, we get a $[7, 4, 3]$ code. It has parity check matrix

$$\begin{pmatrix} 1 & 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 1 & 1 & 0 \\ 1 & 0 & 1 & 0 & 1 & 0 & 1 \end{pmatrix}, \quad (4.6)$$

and can be taken to have generator matrix

$$\begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 1 & 1 & 0 \\ 1 & 0 & 1 & 0 & 1 & 0 & 1 \end{pmatrix}. \quad (4.7)$$

4.2.4 Example: Reed-Muller Codes

Looking at the generator matrix for the 7-bit Hamming code, it has a somewhat nice form. The form gets even nicer if you add an extra bit which is 1 for the first row and 0 for the others:

$$\begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 \end{pmatrix}. \quad (4.8)$$

Now every codeword has weight 0, 4, or 8. This is an example of a Reed-Muller code, specifically $\mathcal{R}(1, 3)$. It is an $[8, 4, 4]$ code.

Definition 4.9. The *1st order Reed-Muller code* $\mathcal{R}(1, m)$ is a linear code with $n = 2^m$ physical bits. It has a generator matrix whose rows are the all-1s vector and the vectors v_i ($i = 1, \dots, m$), where the j th coordinate of v_i is equal to the i th bit of j , when j is expanded in binary (and assuming j runs from 0 to $n - 1$).

For instance, for $n = 4$, $v_1 = (0011)$ and $v_2 = (0101)$. The generator matrix is

$$\begin{pmatrix} 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 1 \end{pmatrix} \quad (4.9)$$

It is a $[4, 3, 2]$ code. Note that in equation (4.8), the order of the bits is reversed relative to this convention.

Theorem 4.7. $\mathcal{R}(1, m)$ is a $[2^m, m + 1, 2^{m-1}]$ code.

Proof. Each of the vectors v_i has weight exactly $n/2$ ($n = 2^m$), as exactly half the numbers $0, \dots, n - 1$ will be 1 in any given bit location. We can also easily show that any sum of $s > 0$ of the v_i s will have weight exactly $n/2$: The sum will be 0 or 1 in the j th location iff the XOR of the corresponding bits of j is 0 or 1. For instance, for $v_1 + v_3$, take the XOR of the first and third bits of j . We then let j run over from $0, \dots, n - 1$. As we do this, the set of bits we are looking at can take on every possible set of values, and furthermore, each set of values appears the same number of times, corresponding to every set of values for the other bits. In particular, any particular assignment of the s bits we are interested in shows up 2^{n-s} times. For half of these assignments, the XOR will be 0, and for the other half, the XOR will be 1. Thus the weight of the sum we are looking at is exactly $n/2$.

The all-1s vector has weight n , and the all-1s vector added to any vector of weight $n/2$ is again a vector of weight $n/2$. Thus, the code has distance $n/2 = 2^{m-1}$.

This argument also shows that all the rows of the generator matrix are independent, since no linear combination gives 0. Therefore, the code has $m + 1$ encoded bits, the same as the number of rows. \square

Definition 4.10. The r th order Reed-Muller code $\mathcal{R}(r, m)$ has as rows of its generator matrix all products of up to r of the vectors v_i given in definition 4.9, where product means the bitwise product (1 in a coordinate iff all vectors in the product are 1 at that coordinate) The all-1s generator is also included. (It can be considered as the product of 0 of the v_i 's.)

For instance, $\mathcal{R}(2, 2)$ has the additional generator $v_1 v_2 = (0001)$, giving the generator matrix

$$\begin{pmatrix} 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (4.10)$$

$\mathcal{R}(2, 2)$ is a $[4, 4, 1]$ code, which actually means it contains all 4-bit vectors. $\mathcal{R}(2, 3)$ is more interesting, with generator matrix

$$\begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 \end{pmatrix} \quad (4.11)$$

$\mathcal{R}(2, 3)$ is an $[8, 7, 2]$ code.

Theorem 4.8. $\mathcal{R}(r, m)$ is a $[2^m, N(r, m), 2^{m-r}]$ code, where

$$N(r, m) = 1 + \binom{m}{1} + \binom{m}{2} + \cdots + \binom{m}{r}. \quad (4.12)$$

Proof. We must show that the code has $N(r, m)$ encoded bits and has distance 2^{m-r} .

Number of encoded bits: $\mathcal{R}(r, m)$ is spanned by a product of vectors for each subset of up to r elements of the numbers $1, \dots, m$. Counting them gives the formula for $N(r, m)$. However, we do need to show that these are linearly independent to see that this is also the number of encoded bits.

Consider the extreme case of $\mathcal{R}(m, m)$. The generator matrix for $\mathcal{R}(r, m)$ is contained in the top rows of that for $\mathcal{R}(m, m)$, so it is enough to show that the generator matrix for $\mathcal{R}(m, m)$ has full rank. $N(m, m) = 2^m = n$, so the rows are linearly independent iff they span the whole vector space.

We can imagine arbitrary vectors on 2^m bits as functions from m bits to one bit, as follows: for the j th bit of the vector, interpret j as the input of the function, with the output of the function given by the j th bit. (For instance, the vector 0110 is the function $f(00) = 0, f(01) = 1, f(10) = 1, f(11) = 0$.) In this interpretation, v_i is the function “take the i th bit of the input.”

The product of multiple vectors v_{i_1}, \dots, v_{i_s} is the function “take the AND of bits i_1, \dots, i_s of the input,” and the sum of products is the XOR of these ANDs. We can write an arbitrary function from m bits to 1 bit as the XOR of ANDs of up to all m bits, so all possible vectors are in the code $\mathcal{R}(m, m)$. Therefore all products of the v_i s are linearly independent.

Distance: It is easy to see that the product of any r vectors v_i has weight $n/2^r = 2^{m-r}$, as the product is 1 in the j th location iff the AND of the corresponding bits of j is 1, which happens for only a fraction $1/2^r$ of the possible values of j . We also have to check the distance, however, for the sums of these products, and it is not clear that taking the sum cannot cause the weight to decrease.

To show that the distance of $\mathcal{R}(r, m)$ is indeed 2^{m-r} , we perform induction on r and m . We have already calculated the distance for $\mathcal{R}(1, m)$ in theorem 4.7. As $m \geq r$, we also want as a base case to show it for a given r for the smallest possible value of m , namely $\mathcal{R}(r, r)$. In this case, there is nothing really to show, as the distance from the formula is 1, which is indeed the weight of the product of all r v_i vectors, and there is no possibility of having a shorter distance.

Claim 4.9. *Assume the distance of $\mathcal{R}(r-1, m)$ is 2^{m-r+1} and the distance of $\mathcal{R}(r, m)$ is 2^{m-r} . Then the distance of $\mathcal{R}(r, m+1)$ is 2^{m+1-r} .*

Proof of claim. We can consider any given sum of basis vectors, and break it up into a term where none of the products includes the vector v_1 and a term where all of the products include v_1 . Now, if we restrict attention to the first 2^m coordinates, the second term is uniformly 0 and the first term is a vector from $\mathcal{R}(r, m)$, which we already know has weight at least 2^{m-r} unless it is the 0 vector.

If we restrict attention to the last 2^m coordinates, v_1 is always 1, so it can be ignored, and the second term is the sum of products of at most $r-1$ vectors, and is thus a vector from $\mathcal{R}(r-1, m)$. The first term is the same on the last 2^m coordinates as it was on the first 2^m coordinates, and is again a vector from $\mathcal{R}(r, m)$. But $\mathcal{R}(r-1, m) \subseteq \mathcal{R}(r, m)$, so the sum of a term from $\mathcal{R}(r-1, m)$ and a term from $\mathcal{R}(r, m)$ is in $\mathcal{R}(r, m)$, and therefore the last 2^m coordinates have weight at least 2^{m-r} by the inductive hypothesis unless the last 2^m coordinates are all 0.

If the first term is the 0 vector, we actually have a vector from $\mathcal{R}(r-1, m)$ on the last 2^m coordinates, which therefore has weight at least 2^{m-r+1} unless it is 0. To get 0 on the last 2^m coordinates, either both terms are 0 (in which case the whole vector is 0), or the first and second terms must cancel on the last 2^m coordinates. In order to cancel the second term, the first term must actually be a vector from $\mathcal{R}(r-1, m)$ on each half of the coordinates, which again means the vector on each half of the coordinates has weight 2^{m-r+1} .

That is, we have three cases (assuming the overall vector is not 0): In case one, the first term is 0, in which case the last 2^m coordinates have weight 2^{m-r+1} . In case two, the first and last 2^m coordinates will each have weight at least 2^{m-r} . In case three, the last 2^m coordinates have weight 0, but the first 2^m coordinates have weight 2^{m-r+1} . In all of these cases, we know that the overall vector has weight at least 2^{m+1-r} , completing the induction for the distance. \square

If we have proven the formula for distance for $r-1$ and all m , then we can use induction on m . The base case is $\mathcal{R}(r, r)$ and we use the claim to prove the distance formula for this specific value of r and all m . This then allows us to use induction on r and the base case of $\mathcal{R}(1, m)$ to prove the distance formula for all values of r and m . \square

4.3 Dual Codes

4.3.1 Definition of a Dual Code

Definition 4.11. If C is a linear code, the *dual code* C^\perp is

$$C^\perp = \{y \in \mathbb{Z}_2^n \mid x \cdot y = 0 \ \forall x \in C\}. \quad (4.13)$$

The dual code switches the role of the generator and parity check matrices: The generator matrix of C^\perp is the parity check matrix of C and the parity check matrix of C^\perp is the generator matrix of C . (Note that $(C^\perp)^\perp = C$.) It follows that the dual code of an $[n, k, d]$ code is an $[n, n-k, d']$ code. The distance d' does not in general have to be related to the distance d .

We also saw a definition of a dual code when discussing the binary symplectic representation of a stabilizer code. The dot product appears in this definition and the symplectic form appeared in that one, but otherwise they are the same.

We can also define *self-dual* codes and *weakly self-dual* codes in the same way as in section 3.5. That is, a self-dual code is equal to its dual, and a weakly self-dual code is contained in its dual.

4.3.2 Dual Codes for the Examples

For the 7-bit Hamming code, we've already worked out the generator and parity check matrices. Looking at the 8 vectors in the span of the rows of the parity check matrix of the 7-bit Hamming code, we see that all the nonzero vectors in the dual code have weight 4. The dual code of the $[7, 4, 3]$ is thus a $[7, 3, 4]$ code. We've already seen that the 7-bit Hamming code is related to a Reed-Muller code. The dual is too — if take $\mathcal{R}(1, 3)$ and drop the all-1s vector, one bit is always zero. If we then discard that bit, we get the dual of the 7-bit Hamming code.

This is true in general:

Theorem 4.10. *Take $\mathcal{R}(1, m)$, remove the all-1s vector, and then puncture it: drop the bit that is always 0. The resulting $[2^m - 1, m, 2^{m-1}]$ code is the dual of the $[2^m - 1, 2^m - m - 1, 3]$ Hamming code.*

Proof. First, let us check the parameters of the punctured Reed-Muller code. Dropping the all-1s vector removes one encoded bit, leaving m logical bits. Removing codewords does not decrease the distance, and since the remaining vectors all have weight 2^{m-1} , it does not increase it in this case either. Dropping a bit that is always 0 also does not change the distance, giving the parameters $[2^m - 1, m, 2^{m-1}]$.

We are left with a generator matrix which has m rows v_1, \dots, v_m . The j th bit of v_i is the i th bit of the binary representation of j , so the j th column of the generator matrix is exactly the binary representation of v_j . This was how we constructed the parity check matrix of the $[2^m - 1, 2^m - m - 1, 3]$ code, which is the generator of the dual. \square

Naturally, this also means that the duals of the Reed-Muller codes $\mathcal{R}(1, m)$ are related to the Hamming codes. In the special case of $\mathcal{R}(1, 3)$, it was related both to the 7-bit Hamming code *and* to its dual. That is because $\mathcal{R}(1, 3)$ is a self-dual code. Most Reed-Muller codes are not self-dual, but $\mathcal{R}(1, 3)$ is not the only one that is. More importantly, the dual of every Reed-Muller code is another Reed-Muller code.

Theorem 4.11. *The dual of $\mathcal{R}(r, m)$ is $\mathcal{R}(m - r - 1, m)$ ($r \leq m - 1$).*

Proof. Let us take the dot product of two basis vectors $w \in \mathcal{R}(r, m)$ and $w' \in \mathcal{R}(r', m)$. The dot product is the parity of the pointwise product ww' . But w is the pointwise product of up to r of the v_i vectors, and w' is the pointwise product of up to r' of the v_i vectors, so ww' is the pointwise product of up to $r + r'$ of the v_i vectors. Suppose we eliminate redundant v_i s that appear twice in the product, leaving us with $s \leq r + r'$ vectors v_i that appear in at least one of the two products w and w' . We already know from theorem 4.8 that such a product has weight exactly 2^{m-s} . Therefore the dot product of w and w' (the parity of the pointwise product) is 0 unless $s = m$, which is only possible if $r + r' \geq m$. Therefore, $\mathcal{R}(m - r - 1, m)$ is orthogonal to $\mathcal{R}(r, m)$, and is contained in its dual.

Now, $\mathcal{R}(r, m)$ encodes $N(r, m) = \binom{m}{0} + \binom{m}{1} + \dots + \binom{m}{r}$ bits by theorem 4.8, so its dual encodes $2^m - N(r, m)$ bits. But $\mathcal{R}(m - r - 1, m)$ encodes

$$\binom{m}{0} + \binom{m}{1} + \dots + \binom{m}{m - r - 1} = \binom{m}{m} + \binom{m}{m - 1} + \dots + \binom{m}{r + 1} \quad (4.14)$$

bits. Therefore, $N(r, m) + N(m - r - 1, m) = 2^m$, and $\mathcal{R}(m - r - 1, m)$ is not only contained in the dual of $\mathcal{R}(r, m)$, it is the same size as the dual, and therefore equals the dual. \square

Corollary 4.12. *When $m = 2r + 1$, the Reed-Muller code $\mathcal{R}(r, m)$ is self-dual. When $m \geq 2r + 1$, $\mathcal{R}(r, m)$ is weakly self-dual.*

4.4 Non-Binary Linear Codes

While binary codes, with the number of physical states $N = 2^n$, are the most common sort of error-correcting code, non-binary codes, which have other values of N , are also used. Some non-binary codes are quite important in the theory of quantum error correction, so we'll quickly cover the concepts here. We won't get to non-binary *quantum* codes until chapter 8, but we'll use some facts about non-binary classical codes even when discussing QECCs over qubits.

4.4.1 Linear Codes Over Finite Fields

When using the binary generalization of linear codes, we usually assume that $N = q^n$, and $q = p^m$ is a prime power. We want q to be a prime power because then there is a finite field $\text{GF}(q)$ of size q , and we can consider $[1 \dots N]$ to be a vector space of dimension n over $\text{GF}(q)$. (See appendix C for an introduction to finite fields.)

Definition 4.12. An error-correcting code $C \subseteq \text{GF}(q)^n$ is a *linear code* if $x, y \in C \Rightarrow \alpha x + \beta y \in C$ for any $\alpha, \beta \in \text{GF}(q)$. C is an *additive code* if $x, y \in C \Rightarrow x + y \in C$.

For binary linear codes, we only needed to worry about adding together codewords, but for non-binary linear codes, multiplication by scalars from the field $\text{GF}(q)$ must also keep us within the code. If adding codewords gives a codeword but multiplication by scalars does not necessarily do so, then the code is merely *additive*. For bits, additive implies linear because the only scalars are 0 and 1, but for larger fields, there are more options.

We can again consider the possible errors to be vectors in $\text{GF}(q)^n$. If an error e acts on a string x , the resulting state is $x + e$, as before. Bear in mind, however, that there are $q - 1$ different errors that can affect each register.

Non-binary linear codes have generator and parity check matrices defined in the same way as for binary codes, and the distance is also defined in the same way. We use the notation $[n, k, d]_q$ for a non-binary linear code with n physical registers, each of size q . The only difference in the basic properties of the code is that we now need to be careful of the distinction between addition and subtraction, which are the same for a binary code. In particular,

Theorem 4.13. A non-binary linear code C corrects the error set \mathcal{E} iff $e - f \notin C$ for all $e \neq f \in \mathcal{E}$. Equivalently, $H_C e \neq H_C f \forall e \neq f \in \mathcal{E}$ (i.e., all errors in \mathcal{E} have different error syndromes).

We can define the dual code for a non-binary code in just the same way as for a binary code using the dot product for a vector space over $\text{GF}(q)$.

4.4.2 Example: Hamming Codes

The Hamming codes have a natural generalization to non-binary fields. Our goal in designing non-binary Hamming codes is again to choose the columns of the parity check matrix so that all single-qubit errors have different error syndromes. However, we need to be careful because there is now more than one possible single-bit error per register. If e_i is the error that is 1 in location i and 0 elsewhere, then e_i has syndrome equal to the i th column of the parity check matrix. When $\alpha \in \text{GF}(q)$, αe_i has error syndrome equal to α times the i th column of the parity check matrix. Therefore, the correct generalization of the Hamming code is not to ensure just that all columns are different; instead, we want that no column of the parity check matrix is a scalar multiple of another for any scalar in $\text{GF}(q)$.

For instance, $\text{GF}(4)$ has 4 elements 0, 1, ω , and ω^2 . Therefore, we have a $[5, 3, 3]$ Hamming code over $\text{GF}(4)$ with the following parity check matrix:

$$\begin{pmatrix} 0 & 1 & 1 & 1 & 1 \\ 1 & 0 & 1 & \omega & \omega^2 \end{pmatrix} \quad (4.15)$$

Notice that we don't include columns $(0, \omega)$ or $(0, \omega^2)$ since they are scalar multiples of $(0, 1)$, but that we include all columns of the form $(1, \alpha)$, since no two of them are scalar multiples. We don't, however, then need to include any columns (ω, α) or (ω^2, α) .

Following this logic, we find the non-binary Hamming codes:

Theorem 4.14. There exists a $[(q^r - 1)/(q - 1), (q^r - 1)/(q - 1) - r, 3]_q$ code for any prime power q and any $r > 1$. These codes are known as Hamming codes.

Proof. There are r rows in the parity check matrix. We wish to choose $(q^r - 1)/(q - 1)$ columns so that none is a scalar multiple of another. Note that

$$(q^r - 1)/(q - 1) = 1 + q + \dots + q^{r-1}. \quad (4.16)$$

We will proceed by induction on r , the number of rows in the parity check matrix. If $r = 1$, we just have the parity check matrix (1), which gives a $[1, 0, 1]_q$ code. It is a special case.

Next, we assume that we have a parity check matrix for $r - 1$. For the r -row matrix, we can choose the first column to be all 0s except for the last row, which is 1. Then any column for which the first $r - 1$ entries are 0 will be a scalar multiple of this column. The remaining columns are of the form (v, α) , where v is a column of the parity check matrix for the Hamming code for q and $r - 1$, and α is any element of $\text{GF}(q)$. When $v \neq v'$ are two different columns of the $r - 1$ Hamming code parity check matrix, then $v \neq \beta v'$, so it's certainly true that $(v, \alpha) \neq \beta(v', \alpha')$. Thus, we don't get columns which are scalar multiples that have different entries in the first $r - 1$ entries. We should also compare (v, α) with (v, α') , but since the first r entries are not all 0, the only possible scalar factor between them is 1, which implies that $\alpha = \alpha'$. Thus, this scheme gives independent columns, as desired. We can pick $(q^{r-1} - 1)/(q - 1)$ columns of the $r - 1$ Hamming code, and q entries for the last row in the column, plus we have the first column $(0, 0, \dots, 0, 1)$. The total number of columns is thus

$$1 + q[(q^{r-1} - 1)/(q - 1)] = 1 + q(1 + q + \dots + q^{r-2}) = 1 + q + \dots + q^{r-1} = (q^r - 1)/(q - 1), \quad (4.17)$$

as desired. \square

4.4.3 Example: Reed-Solomon Codes

Reed-Solomon codes are a useful class of codes which are based on polynomials over finite fields.

Definition 4.13. Let $\alpha_1, \dots, \alpha_n$ be n distinct elements of $\text{GF}(q)$, and let $k \leq n$. A *Reed-Solomon code* is

$$\{(f(\alpha_1), \dots, f(\alpha_n)) \mid f \text{ is a polynomial of degree } < k\}. \quad (4.18)$$

The encoder for this code equates the input $(\beta_0, \dots, \beta_{k-1})$ to the polynomial

$$f(x) = \beta_0 + \beta_1 x + \beta_2 x^2 + \dots + \beta_{k-1} x^{k-1}. \quad (4.19)$$

Reed-Solomon codes are widely used for practical purposes, for instance to encode DVDs to protect against minor damage and in QR codes in case some part of the code can't be read. They are useful because they have a large distance and are easy to encode and decode. But who cares about that stuff? It's all very 20th-Century. The *real* reason Reed-Solomon codes are interesting is that they are useful for making quantum codes, although we won't see how until we discuss non-binary QECCs in chapter 8.

Theorem 4.15. A *Reed-Solomon code with n points from $\text{GF}(q)$ and polynomials of degree at most $k - 1$ is an $[n, k, n - k + 1]_q$ code.*

Proof. Directly from the definition we can see that there are n physical registers and k logical registers. The distance is less obvious. The most straightforward way to see it is to consider correcting erasure errors.

Suppose that t registers have been erased. We are then left with $n - t$ registers, which are the values of the polynomial f evaluated at $n - t$ different points. We know that f has degree at most $k - 1$, and a degree $k - 1$ polynomial is determined by its value at any k points. To be more concrete, given any k points, we have k different linear equations for the coefficients $\beta_0, \dots, \beta_{k-1}$. The equations are defined by the Vandermonde matrix $V_{ij} = \alpha_i^{j-1}$. (I have assumed that the registers which are available correspond to the points $\alpha_1, \dots, \alpha_k$ for notational simplicity.) The Vandermonde matrix is non-singular — it has determinant $\prod(\alpha_i - \alpha_j)$ — so the equations have a unique solution.

In short, provided $n - t \geq k$, we can reconstruct f and decode the code. Thus, the code corrects $n - k$ erasure errors, so it has distance $n - k + 1$. The distance can't be higher than that, because if we have $n - k + 1$ erasure errors, there are only $k - 1$ points left, and there are multiple possible polynomials which go through those points since we could choose any value for a k th point and still reconstruct a valid polynomial. \square

4.5 Hamming, Gilbert-Varshamov, and Singleton Bounds, MDS Codes

To conclude the discussion of classical codes, we'll discuss some basic limits on the existence of classical codes. I won't get to the quantum analogue of these bounds until chapter 7, but the bounds help to understand the example codes that we've discussed. As you'll see, the Hamming codes and Reed-Solomon codes are optimal codes, giving a maximal k and d for minimum n .

4.5.1 Hamming Bound, Perfect Codes

One can set a simple upper bound by taking advantage of the requirement that the states must be distinguishable after errors.

Theorem 4.16 (Hamming bound). *An $(n, K, 2t + 1)_q$ code must satisfy*

$$K \left(\sum_{j=0}^t (q-1)^j \binom{n}{j} \right) \leq q^n. \quad (4.20)$$

For large n ,

$$(\log_q K)/n \leq 1 - (t/n) \log_q(q-1) - h_q(t/n), \quad (4.21)$$

where

$$h_q(x) = -x \log_q x - (1-x) \log_q(1-x). \quad (4.22)$$

Proof. Let x be a vector in the code C , and let S_x be the set of all vectors at distance at most t from x . (I.e., all strings that can be reached from x by altering up to t registers.) An error can take x to any string in S_x , so we need that $S_x \cap S_y = \emptyset$ when $x \neq y$. Otherwise, $z \in S_x \cap S_y$ can't be reliably decoded since it could have come from either x or y before the error.

Let us count the size of S_x . We break S_x into subsets $S_{x,j}$ which disagree with x on exactly $j \leq t$ registers. $S_{x,j}$ breaks down further into subsets which depend on which j registers disagree. There are $\binom{n}{j}$ possible sets of registers which disagree, and for a fixed set of registers, the strings in $S_{x,j}$ can take on any value except for the values in x . There are $q-1$ remaining choices for each of the j registers which disagree. Thus, the total size of $S_{x,j}$ is $(q-1)^j \binom{n}{j}$ and the total size of S_x is

$$\sum_{j=0}^t (q-1)^j \binom{n}{j}. \quad (4.23)$$

This is true for all x , and since $S_x \cap S_y = \emptyset$, the total number of strings in the union of all sets S_x is K times equation (4.23). This must be less than the total number of possible n -register strings, which is q^n , giving us equation (4.20).

To find the large n version of this equation, just take the logarithm base q . The h_q term comes from the log of the binomial coefficient (lemma 4.17), and only the largest value of $j = t$ contributes to the logarithm for large n . \square

The formula for $h_q(x)$ is a fairly standard information-theoretic term, but it is sufficiently useful that I'll give a proof of it:

Lemma 4.17. *For large n ,*

$$\log_q \binom{n}{j} = -j \log_q(j/n) - (n-j) \log_q(1-j/n) + o(n). \quad (4.24)$$

Proof.

$$\log \binom{n}{j} = \log(n!) - \log(j!) - \log[(n-j)!]. \quad (4.25)$$

(Assume everywhere that the base of the logarithm is q .)

Taking the logarithm of Stirling's formula and keeping only the terms at least linear in n , we have

$$\log(m!) = m \log m - m \log e. \quad (4.26)$$

Then

$$\log \binom{n}{j} = (n \log n - n \log e) - (j \log j - j \log e) - [(n-j) \log(n-j) - (n-j) \log e] \quad (4.27)$$

$$= j \log n + (n-j) \log n - j \log j - (n-j) \log(n-j) - [n-j - (n-j)] \log e \quad (4.28)$$

$$= -j \log(j/n) - (n-j) \log[(n-j)/n]. \quad (4.29)$$

□

The Hamming bound is also known as the *sphere-packing bound*. The name comes from the sets S_x , which can be viewed as “spheres” using the distance measure which counts the number of registers in which two strings differ. That metric is known as the *Hamming distance*. The spheres S_x are rather blocky for spheres, but that's what you get when you use a discrete distance.

The case when the Hamming bound is met exactly is somewhat interesting. First, it represents the best possible code you can have for a given n and $d = 2t + 1$. Second, when it is a linear code, it means that every error syndrome is used by an error of weight t or less.

Definition 4.14. A code which saturates the Hamming bound is *perfect*.

When $t = 1$, the condition for a perfect code is $K[1 + (q-1)n] = q^n$. Let us specialize to $K = q^k$. Then an $[n, k, 3]_q$ code is perfect if $(q-1)n = q^{n-k} - 1$. Letting $r = n - k$, we find that $n = (q^r - 1)/(q - 1)$, $k = n - r$. These are exactly the parameters of the Hamming codes. The Hamming codes were designed to use up every error syndrome, so it's not surprising they are perfect codes, but now you can see that these parameters are the only ones possible for distance 3 perfect codes.

4.5.2 Gilbert-Varshamov Bound

The Hamming bound tells us that codes above a certain level of efficiency cannot exist — it is an *upper bound* on the efficiency of error-correcting codes. The Gilbert-Varshamov bound is a *lower bound*. It tells us that efficient codes do exist, provided we lower our standards a bit.

Theorem 4.18 (Gilbert-Varshamov bound). *If n , K , and d satisfy*

$$K \left(\sum_{j=0}^{d-1} (q-1)^j \binom{n}{j} \right) \leq q^n, \quad (4.30)$$

then an $(n, K, d)_q$ code exists. For large n , a code exists if

$$(\log_q K)/n \leq 1 - (d/n) \log_q(q-1) - h_q(d/n), \quad (4.31)$$

with $h_q(x)$ given by equation (4.22).

The difference between the upper bound given by the Hamming bound and the lower bound given by the Gilbert-Varshamov bound is replacing t (the number of correctable errors) with $d - 1$ (the number of detectable errors), which is $2t$. The asymptotic (large n) versions of the two bounds for $q = 2$ are plotted in figure 4.2, along with the Singleton bound (section 4.5.3). Of course, there may exist codes whose efficiency is in the region between the Hamming and Gilbert-Varshamov bounds, but this is not the generic case.

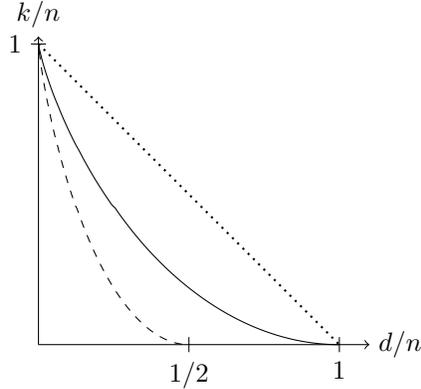


Figure 4.2: Classical Hamming bound (solid), Gilbert-Varshamov bound (dashed), and Singleton bound (dotted) for large n , $q = 2$

Proof. We will pick the codewords sequentially, thus showing the theorem by induction on K .

For the first codeword, pick any string x_1 . Then exclude the Hamming sphere of radius $d - 1$ around x_1 , that is, the set S_{x_1} of all strings that are distance $d - 1$ or less from x_1 . Choose the second codeword x_2 to be any string outside S_{x_1} . x_2 has distance at least d from x_1 , so the code $\{x_1, x_2\}$ has distance at least d .

In general, given any code $\{x_1, \dots, x_{K-1}\}$ with distance at least d , exclude the radius $d - 1$ Hamming spheres $S_{x_1}, \dots, S_{x_{K-1}}$. Provided equation (4.30) is satisfied, there is at least one string not excluded. Choose it (at random if there is more than one) as the K th codeword x_K . Then $\{x_1, \dots, x_K\}$ also has distance at least d , since x_K is no closer than a Hamming distance d to any of the previous codewords. \square

The proof can be strengthened to say not just that there exist codes with the parameters $(n, K, d)_q$, but actually that a *randomly chosen* (n, K) code has distance at least d (with high probability for large n). The Gilbert-Varshamov bound can also be improved to show the existence of codes with certain properties. For instance, when $K = 2^k$, there exists a linear code provided n , K , and d satisfy equation (4.30).

The proof of the Gilbert-Varshamov bound is sometimes referred to as *non-constructive*. Of course, in a literal sense it is constructive, since if we try all codes with the given parameters, we will eventually find one that works. In this context, the term means that it is not efficiently constructive in at least one of two ways. It can mean that picking a code according to the algorithm implicitly given by the proof produces a code whose description is exponentially large in n . In this case, that is true, since $K = \exp(O(n))$ and each codeword has to be listed separately, but if we use the version of the Gilbert-Varshamov bound applying to linear codes, there is at least an efficient description of the resulting code in terms of the generator matrix. It could also mean that the algorithm for finding the code takes exponentially long in n . That is true both for the general version and linear code version of the Gilbert-Varshamov bound, since checking that a particular code has distance d takes exponentially long.

4.5.3 Singleton Bound

The Singleton bound is another upper bound, but it has a much simpler form than the Hamming bound. In particular, the form of the Singleton bound is the same for all register sizes q .

Theorem 4.19 (Singleton bound). *If an $(n, q^k, d)_q$ code exists, then*

$$n - k \geq d - 1. \quad (4.32)$$

Proof. An $(n, q^k, d)_q$ code corrects $d - 1$ erasure errors. For instance, if we discard the last $d - 1$ registers, the remaining $n - (d - 1)$ registers can still be used to reconstruct all k logical registers. This is only possible

if $n - (d - 1) \geq k$ since otherwise two different logical strings will be represented by the same string with $n - (d - 1)$ components. \square

Codes which saturate the Singleton bound have interesting properties, and therefore get their own name.

Definition 4.15. If an $(n, q^k, d)_q$ code satisfies $n = k + d - 1$, it is a *MDS* code.

“MDS” stands for “maximum distance separable.” You can probably figure out the “maximum distance” part of the name. The “separable” part means that k coordinates of the codewords can be taken to be the logical registers when no errors are present. This is not a property that is possible for a quantum code, since it would violate the no-cloning theorem (if k qubits contain the logical state, then the remaining qubits contain no information about it, and would not be useful for error correction). I will still use the term “quantum MDS codes” for codes satisfying the analogous quantum Singleton bound.

Checking the parameters, you’ll find that the Reed-Solomon codes are MDS codes. The repetition codes $[n, 1, n]$ are too. The Hamming codes with $r \geq 3$ are not — even though they are optimal for the Hamming bound, they do not saturate the Singleton bound.

4.5.4 Properties of MDS Codes

The dual code of an $[n, k, d]$ code is an $[n, n - k]$ code, but its distance might be bad. A remarkable property of MDS codes, and one that is very useful for making quantum codes, is that their duals also have good distance.

Theorem 4.20. *The dual of an MDS code is also an MDS code. In particular, the dual of an $[n, k, n - k + 1]$ code is an $[n, n - k, k + 1]$ code.*

Proof. We wish to show that any codeword from the dual code has weight greater than k . This will show that the distance of the dual code is at least $k + 1$, but it cannot be any higher than that by the Singleton bound.

In terms of equations, we wish to show that if C is the original code and y is some string with $0 < \text{wt } y \leq k$, then $\exists x \in C$ such that $y \cdot x \neq 0$. Recalling the proof of the Singleton bound, let us erase $d - 1$ registers from an MDS code, leaving only k remaining registers, which are chosen to include the support of y . Call R the set of k registers we are considering. Since the logical state can be reconstructed from R , they must take on every possible value. That is, for any string x' with support within R , $\exists x \in C$ such that $x|_R = x'$. Let $y|_R = y'$. Since $y' \neq 0$, there certainly exists some x' such that $y' \cdot x' \neq 0$, and then $y \cdot x \neq 0$ whenever y has support on R . Since R was an arbitrary set of size k , this shows that if $y \in C^\perp$, then $\text{wt } y > k$. \square

Chapter 5

Combining The Old And The New: Making Quantum Codes From Classical Codes

Now that we've covered the basics of both classical and quantum error-correcting codes, we're ready to try combining them. By borrowing some codes from the old theory of classical error-correcting codes, we'll be able to make brand new quantum error-correcting codes, thus marrying quantum and classical error correction. But don't be blue, I'm sure there are still plenty of interesting quantum codes left to be discovered.

5.1 CSS Codes

5.1.1 CSS Construction and Parameters of a CSS Code

The first construction we'll discuss gives us a class of codes known as CSS codes after their inventors Calderbank, Shor, and Steane. The idea of CSS codes is to build on the observation that classical linear codes are a special case of stabilizer codes. In theorem 4.4, we made a stabilizer code by taking the parity check matrix for a classical linear code and replacing each 1 with a Z . The resulting stabilizer code corrects the same set of bit flip errors as the original classical code. If you instead replace the 1's in the parity check matrix with X 's, you get a stabilizer code that corrects phase flip errors. In a CSS code, we do both: some of the stabilizer generators come from one classical linear code C_1 with the parity check matrix converted to Z 's, and some generators come from a second linear code C_2 which is used to correct phase errors.

Definition 5.1. A stabilizer code is a *CSS code* if there is a choice of generators for which the stabilizer's binary symplectic representation is of the form

$$\left(\begin{array}{c|c} 0 & A \\ \hline B & 0 \end{array} \right), \quad (5.1)$$

where A is a $r_1 \times n$ matrix and B is a $r_2 \times n$ matrix for some r_1, r_2 . The generators of the form $(b|0)$ are *X generators* and the generators of the form $(0|a)$ are *Z generators*.

In other words, some generators for a CSS code are tensor products of only X and I and some are tensor products of only Z and I . This statement is of course dependent on the exact choice of generators. If you pick a different set of generators by multiplying some of the X generators with some of the Z generators, you will usually get generators which involve more than one of X , Y , and Z in the same operator. That doesn't mean the code is not a CSS code, only that you've picked a strange set of generators.

Z	Z	Z	Z	I	I	I
Z	Z	I	I	Z	Z	I
Z	I	Z	I	Z	I	Z
X	X	X	X	I	I	I
X	X	I	I	X	X	I
X	I	X	I	X	I	X
\overline{X}	X	X	X	X	X	X
\overline{Z}	Z	Z	Z	Z	Z	Z

Table 5.1: The stabilizer and logical Paulis for the 7-qubit code.

The CSS construction allows us to take two classical codes and make a quantum code. As an example, let's form a 7-qubit code from the 7-bit Hamming code discussed in section 4.2.3. The second code will also be the 7-bit Hamming code. Take the three rows in its parity check matrix and convert the 1's to Z 's, getting three generators of the stabilizer. Then take the rows a second time and convert the 1's to X 's, getting three more generators. The resulting stabilizer is shown in table 5.1.

What is the distance of the 7-qubit code? Using the Z generators of the stabilizer, we can detect and identify any single-qubit bit flip error, since those generators are derived from a classical code which can do so. The first three bits of the error syndrome tell us where a bit flip error is. Using the same logic, the last three bits tell us where any single-qubit phase error has occurred. If there is a Y error, or indeed an X error on one qubit and a Z error on another, then the error will show up in both the first three bits and the last three bits of the error syndrome, identifying it as an error combining X and Z . Thus, the code can correct any single-qubit X , Y , or Z error, and has distance 3.

How many encoded qubits do we have? Using proposition 3.3, we have $n = 7$ physical qubits and 6 generators, so there should be 1 encoded qubit. But hold on a minute. In order to have *any* encoded qubits, we can't just take an arbitrary set of Paulis and call them the generators of a stabilizer. In particular, to have a stabilizer, we need to check that the generators we've written down commute with each other. The Z generators automatically commute with each other and the X generators commute with each other, so all we need to check is that every Z generator commutes with each X generator. In the case of the 7-qubit code, they do. Therefore, the 7-qubit code is well-defined as a $[[7, 1, 3]]$ code. The 7-qubit code is also known as the *Steane code*, since Steane first proposed it.

In the case of the 7-qubit code, we derived both the X and Z generators from the same code, the $[7, 4, 3]$ Hamming code. For a more general CSS code, we don't need to do that. We can use C_1 to correct bit flip errors and C_2 to correct phase errors.

Theorem 5.1. *Let C_1 be an $[n, k_1, d_1]$ classical linear code with parity check matrix H_1 and C_2 be an $[n, k_2, d_2]$ classical linear code with parity check matrix H_2 . Suppose $C_1^\perp \subseteq C_2$. Let S be the CSS code with stabilizer*

$$\left(\begin{array}{c|c} 0 & H_1 \\ \hline H_2 & 0 \end{array} \right). \tag{5.2}$$

Then S is an $[[n, k, d]]$ quantum code with $k = k_1 + k_2 - n$ and $d \geq \min\{d_1, d_2\}$.

There are a couple of things to notice about the statement of the theorem. You might think that the condition $C_1^\perp \subseteq C_2$ implies an asymmetry between C_1 and C_2 , but that is not the case. Actually, the theorem treats them on an equal basis because $C_1^\perp \subseteq C_2 \Leftrightarrow C_2^\perp \subseteq C_1$. Also, observe that if $n > k_1 + k_2$, the theorem would predict a stabilizer code with a negative number of encoded qubits. That can't be right, and of course, it isn't. When $n > k_1 + k_2$, it is not possible that $C_1^\perp \subseteq C_2$.

While the theorem is phrased as just one way to make a CSS code, by comparing the theorem and the definition of a CSS code, you can see that it is actually the *only* way to make a CSS code. In the future, I'll refer to the codes C_1 and C_2 of a general CSS codes, indicating the classical codes that produce the Z generators and X generators, respectively. The choice of whether C_1 has the X generators or the Z

generators is an arbitrary convention, and you may find the other choice in the literature. Also, sometimes people will use the name “ C_1 ” when they mean what I am calling “ C_1^\perp .” Again, this is somewhat an arbitrary convention.

Proof. The main thing we need to check is that the X generators commute with the Z generators. Any X generator is of the form $(x|0)$, where $x \in C_2^\perp$. It is in the dual since it is derived from a row of the parity check matrix of C_2 . Any Z generator is derived from the parity check matrix of C_1 , so it has the form $(0|z)$, with $z \in C_1^\perp$. Then

$$(x|0) \odot (0|z) = x \cdot z, \quad (5.3)$$

with the usual binary inner product on the right. Thus, the stabilizer is Abelian iff $x \cdot z = 0$ for all $x \in C_2^\perp$, $z \in C_1^\perp$. Equivalently, we could say that if $z \in C_1^\perp$, then $z \in (C_2^\perp)^\perp$. Since $(C_2^\perp)^\perp = C_2$, that produces the condition $C_1^\perp \subseteq C_2$.

Now let us determine the parameters of S . H_1 has $n - k_1$ rows and H_2 has $n - k_2$ rows, so the stabilizer has $2n - k_1 - k_2$ generators. Therefore, it has $n - (2n - k_1 - k_2) = k_1 + k_2 - n$ logical qubits. The code can detect up to $d_1 - 1$ bit flip errors using the Z generators, and it can detect up to $d_2 - 1$ phase errors using the X generators, and detecting bit flip errors does not in any way interfere with detecting phase errors. Any Pauli of weight less than $\min\{d_1, d_2\}$ can be written as a product PQ , with P a tensor product of X and I with weight $< d_1$ and Q a tensor product of Z and I with weight $< d_2$. The Pauli is non-trivial if at least one of P and Q is non-trivial, in which case it can be detected by looking at the appropriate bits of the error syndrome. Thus, the distance is at least $\min\{d_1, d_2\}$. \square

From theorem 5.1, you can see why I made such a big deal about determining the duals of codes in chapter 4. The $[[7, 4, 3]]$ Hamming code contains its own dual, which is why we can use two copies of it to make the 7-qubit code. We can get other CSS codes by using the other example classical codes. For instance, let $r > m - r - 1$. Then $\mathcal{R}(r, m)^\perp = \mathcal{R}(m - r - 1, m) \subset \mathcal{R}(r, m)$, and we can make a CSS code with $C_1 = C_2 = \mathcal{R}(r, m)$. For instance, $\mathcal{R}(2, 4)$ is a $[[16, 11, 4]]$ code, and using it for C_1 and C_2 , we get a $[[16, 6, 4]]$ CSS code.

We don't have to use the same code twice. For instance, we can let $C_1 = \mathcal{R}(2, 4)$ and $C_2 = \mathcal{R}(3, 4)$, which is a $[[16, 15, 2]]$ code. $C_2^\perp = \mathcal{R}(0, 4)$, which is just the 16-bit repetition code. $\mathcal{R}(0, 4) \subset C_1$, so we can make a CSS code, getting a $[[16, 10, 2]]$ code. Of course, this particular construction is not ideal if we're interested in a distance 2 code, since by taking $C_1 = C_2 = \mathcal{R}(3, 4)$, we get a $[[16, 14, 2]]$ code, which has the same distance but more encoded qubits. Still, it might be useful if we want a code that detects any single-qubit phase error but can actually correct a bit flip error.

5.1.2 Degeneracy and CSS Codes

Theorem 5.1 says that the distance of S is greater than or equal to $\min\{d_1, d_2\}$. Why not just equal to? After all, since code C_1 has distance d_1 , there is a bit flip error of weight d_1 that cannot be detected by code C_1 and therefore the corresponding Pauli error has 0 error syndrome for S . Similarly, C_2 has distance d_2 , so there is also a Pauli with weight d_2 which is a tensor product of Z and I and has 0 error syndrome. But for a quantum code, 0 error syndrome is not enough to make an error undetectable. 0 error syndrome means that the error is in $N(S)$. An error is undetectable only if it is in $\hat{N}(S) \setminus \hat{S}$. That is, when a CSS code is degenerate, its distance could be greater than one might expect simply by examining the two classical codes that make it up.

For instance, the 9-qubit code is a degenerate CSS code. C_1 is composed of three copies of the repetition code. It has distance 3. However, C_2 has the following parity check matrix:

$$\begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 & 1 \end{pmatrix}. \quad (5.4)$$

C_2 contains strings such as $(1, 1, 0, 0, 0, 0, 0, 0, 0)$ which have weight 2. Thus, C_2 has distance 2. (It is easy to check that any single-bit error will have nonzero error syndrome.) If we just took the distance to be $\min\{d_1, d_2\}$, we would think that the nine-qubit code had distance 2. However, the phase error formed from

the short codeword above is Z_1Z_2 , which is in \mathcal{S} . Similarly for all other weight 2 codewords of C_2 . This is why the distance of the nine-qubit code can be 3.

5.1.3 Codewords of a CSS Code in Computational and Dual Basis

CSS codes also look nice when the codewords are written in the standard basis. The projector on the code space is a product of two projectors Π_1 and Π_2 . Π_i is the projector on the +1 eigenspace of the generators derived from C_i . Π_1 thus projects on the subspace spanned by codewords of C_1 . Π_2 can be written as a sum

$$\Pi_2 = \frac{1}{2^{n-k_2}} \left(\sum_{x \in C_2^\perp} P_x \right), \quad (5.5)$$

where P_x is the Pauli with binary symplectic representation $(x|0)$. The Paulis of this form (for $x \in C_2^\perp$) are the stabilizer elements formed from products of the generators derived from the parity check matrix of C_2 .

Therefore, we can find the codewords of \mathcal{S} by taking some codeword of C_1 (which is all that can pass Π_1) and applying Π_2 . In general, we get something of the form

$$|u + C_2^\perp\rangle = \sum_{v \in C_2^\perp} |u + v\rangle, \quad (5.6)$$

for $u \in C_1$. Recall that $C_2^\perp \subseteq C_1$ and C_1 is linear, so $u + v \in C_1$.

When are two such codewords $|u + C_2^\perp\rangle$ and $|u' + C_2^\perp\rangle$ equal?

$$0 = |u + C_2^\perp\rangle - |u' + C_2^\perp\rangle = \sum_{v \in C_2^\perp} |u + v\rangle - \sum_{v' \in C_2^\perp} |u' + v'\rangle. \quad (5.7)$$

This is only possible if every term in the first sum is cancelled by a term in the second sum. That is, when $v \in C_2^\perp$, $u + v = u' + v'$ for some $v' \in C_2^\perp$. Thus, $u - u' = v' - v \in C_2^\perp$ since C_2^\perp is linear. The states $|u + C_2^\perp\rangle$ only depend on cosets of C_2^\perp within C_1 , thus explaining the notation I chose to represent the basis states.

C_1 has 2^{k_1} codewords and C_2^\perp has 2^{n-k_2} codewords, so C_1/C_2^\perp has $2^{k_1+k_2-n}$ codewords, which is the same as the dimension of \mathcal{S} . The codewords $|u + C_2^\perp\rangle$ form a basis for the code space of the CSS code.

I claimed that in the CSS construction, the two classical codes used were treated equally, but in this expansion, they certainly seem unequal. The solution lies in looking in the Hadamard rotated basis. The Hadamard transform switches the role of X and Z , so you might expect that it switches the role of C_1 and C_2 . If you expected that, you are correct. Let's calculate it explicitly:

$$H^{\otimes n} |u + C_2^\perp\rangle = \sum_{v \in C_2^\perp} H^{\otimes n} |u + v\rangle \quad (5.8)$$

$$= \sum_{v \in C_2^\perp} \sum_w (-1)^{(u+v) \cdot w} |w\rangle \quad (5.9)$$

$$= \sum_w (-1)^{u \cdot w} \left(\sum_{v \in C_2^\perp} (-1)^{v \cdot w} \right) |w\rangle \quad (5.10)$$

$$= \sum_{w \in C_2} (-1)^{u \cdot w} |w\rangle \quad (5.11)$$

$$= \sum_{x \in C_2/C_1^\perp} (-1)^{u \cdot x} \sum_{y \in C_1^\perp} |x + y\rangle \quad (5.12)$$

$$= \sum_{x \in C_2/C_1^\perp} (-1)^{u \cdot x} |x + C_1^\perp\rangle \quad (5.13)$$

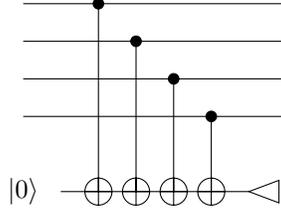


Figure 5.1: A circuit to measure the parity of 4 qubits.

To get the fourth line, observe that if $w \in C_2$, then $v \cdot w = 0$ always, but if $w \notin C_2$, then $v \cdot w = 0$ for half of the values of v and $v \cdot w = 1$ for the other half of the values of v . (It must be nonzero for some v since $w \notin C_2$. Call that v_0 . Then if $v \in C_2^\perp$, so is $v + v_0$, but exactly one of $v \cdot w = 1$ and $(v + v_0) \cdot w = 1$. Thus we can pair the elements of C_2^\perp such that within each pair one is orthogonal to w and one is not.)

In the next-to-last line, I have broken the sum over w into a sum over cosets of C_1^\perp and a sum over elements of the cosets. This is a sensible thing to do because $u \cdot w$ only depends on which coset w lies in of C_1^\perp in C_2 : If $w' = w + y$, $y \in C_1^\perp$, then $u \cdot w' = u \cdot w + u \cdot y$, but $u \in C_1$, so $u \cdot y = 0$.

In particular, if we have associated the basis codewords $|u + C_2^\perp\rangle$ with logical basis codewords, then we recognize equation (5.13) as a Hadamard transform of the logical codewords, with the new basis codewords $|x + C_1^\perp\rangle$. In the standard basis, the CSS code consists of superpositions over cosets of C_2^\perp in C_1 , while in the Hadamard-rotated basis, it consists of superpositions over cosets of C_1^\perp in C_2 .

5.1.4 Error Correction for a CSS Code

CSS codes are stabilizer codes, so any general technique for performing error correction on stabilizer codes will work for CSS codes too. However, CSS codes have additional special structure which we can take advantage of to get better error correction procedures. In section 12.3, I'll discuss a fault-tolerant error-correction procedure that is designed for CSS codes, but for now, I'll just discuss a non-fault-tolerant scheme.

For a general stabilizer code, the bits of the error syndrome come from the eigenvalues of the stabilizer generators. For a CSS code, the stabilizer generators break up into two parts, one from the code C_1 and one from the code C_2 . The first set of syndrome bits identifies bit flip errors, while the remaining syndrome bits identify phase errors.

In the standard basis, the codewords are superpositions of codewords of C_1 , so we can measure the bit flip syndrome by measuring the parity checks of C_1 . Of course, we must be careful to measure *only* the parity checks and nothing more, since we don't want to destroy any superposition of logical states. To measure a parity check, you can add a single ancilla qubit in the state $|0\rangle$ and perform CNOT gates from all the qubits involved in the parity check to the ancilla. Then measure the ancilla qubit and the outcome will be the desired parity. See figure 5.1 for an example.

The bit flip syndrome we get this way is exactly the syndrome of the classical code code C_1 if it underwent the same bit flip error. In order to decode the syndrome and learn the actual error, we can just invoke the classical procedure. If the code C_1 has an efficient syndrome decoding algorithm, then we get one for the bit flip errors of the CSS code also.

For the phase errors, we can simply rotate into the Hadamard basis. Now the codewords are superpositions of codewords of C_2 , and the phase errors have become bit flip errors. If, after the Hadamard transform, we measure the parity checks of C_2 , that is the same as measuring the phase error syndrome for the CSS code. The phase error syndrome is exactly the same as the error syndrome for C_2 if it underwent bit flip errors corresponding to the locations of the phase errors. Again, we can use the syndrome decoding algorithm for C_2 to identify the phase errors for the CSS code.

+	0	1	ω	ω^2	×	0	1	ω	ω^2
0	0	1	ω	ω^2	0	0	0	0	0
1	1	0	ω^2	ω	1	0	1	ω	ω^2
ω	ω	ω^2	0	1	ω	0	ω	ω^2	1
ω^2	ω^2	ω	1	0	ω^2	0	ω^2	1	ω

Table 5.2: The addition and multiplication tables for $\text{GF}(4)$.

5.2 $\text{GF}(4)$ Codes and Stabilizer Codes

Now we'll return to more general stabilizer codes. CSS codes have many nice properties, but they are not perfect. For instance, general stabilizer codes can be more efficient than CSS codes — the 7-qubit code is the smallest CSS code that corrects 1 error, whereas the 5-qubit code, a stabilizer code, is the best QECC of any type that corrects 1 error. While there are some non-stabilizer codes known that are slightly better than any stabilizer code, the differences are small for codes correcting Pauli channels.

In this section, I'll therefore discuss a technique for adapting some classical codes to get stabilizer codes which are not CSS codes. In order to do so, we'll have to go to non-binary codes.

5.2.1 Correspondence Between $\text{GF}(4)$ and the Pauli Group

The main “insight” of this particular technique is that the one-qubit Pauli group has 4 elements, and so does the finite field $\text{GF}(4)$. Amazing, no?

Really, we're connecting \hat{P}_n , the Pauli group sans phases, with an n -dimensional vector field over $\text{GF}(4)$. For $n = 1$, we simply associate each element of the Pauli group with the elements of $\text{GF}(4)$:

$$\begin{aligned} I &\leftrightarrow 0 & X &\leftrightarrow 1 \\ Z &\leftrightarrow \omega & Y &\leftrightarrow \omega^2 \end{aligned} \tag{5.14}$$

For n qubits, the i th tensor factor of $P \in \hat{P}_n$ becomes the i th component of a vector over $\text{GF}(4)$ using the rules in equation (5.14).

The conversion works much the same way as the conversion between \hat{P}_n and the binary symplectic representation of the Pauli group. Multiplication in \hat{P}_n becomes addition in $\text{GF}(4)$. The full correspondence is summarized in table 5.3. It's also worth recalling the addition and multiplication tables for $\text{GF}(4)$, which are given in table 5.2.

5.2.2 The $\text{GF}(4)$ Symplectic Inner Product

As with the binary symplectic representation, we need to recover somehow the notation of commutation, and again we turn to a symplectic inner product. We want some map from $\text{GF}(4) \times \text{GF}(4)$ to \mathbb{Z}_2 which gives 0 when one input is 0 or both inputs are the same, and gives 1 otherwise. It turns out that the correct formula is $\text{tr}(\bar{a}b)$. The \bar{a} and trace operations are defined as follows for $\text{GF}(4)$:

$$\begin{aligned} \bar{0} &= 0 & \bar{1} &= 1 & \text{tr } 0 &= 0 & \text{tr } 1 &= 0 \\ \bar{\omega} &= \omega^2 & \bar{\omega^2} &= \omega & \text{tr } \omega &= 1 & \text{tr } \omega^2 &= 1 \end{aligned} \tag{5.15}$$

You can check by trying all combinations that this formula works. The generalization to n -dimensional vectors is then straightforward:

$$a * b = \text{tr}(\bar{a} \cdot b), \tag{5.16}$$

where \cdot is just the usual dot product. In the classical coding literature, this inner product is sometimes known as the *trace-Hermitian inner product*.

In the Pauli group	In GF(4)
I	0
X	1
Z	ω
Y	ω^2
Multiplication	Addition
$c(P, Q)$	$a * b = \text{tr}(\bar{a} \cdot b)$
Phase	No equivalent
No equivalent	Multiplication
Unitary T (see chapter 6)	Multiplication by ω
Stabilizer S	Weakly self-dual additive code S
Normalizer $N(S)$	Dual S^\perp (under $*$)

Table 5.3: Equivalence between the Pauli group and GF(4).

5.2.3 Stabilizer Codes as GF(4) Codes

We can convert a stabilizer to sets of vectors over GF(4) just as we converted them to the binary symplectic representation. We'd like to interpret the resulting set as a classical error-correcting code over GF(4). Formally, there is no problem in doing so. However, what we get is not necessarily a linear code. Sometimes we do get a linear code, as in the example of the five-qubit code (converted to a GF(4) code in table 5.4), but it is not hard to come with stabilizer codes which don't produce linear GF(4) codes.

Because $P, Q \in S \Rightarrow PQ \in S$, the GF(4) code S which we get is additive (closed under addition). However, for it to be linear, S would also need to be closed under multiplication by ω , and that is not necessarily true. (It also needs to be closed under multiplication by ω^2 , but that follows automatically if it is closed under multiplication by ω .)

The other condition we need to satisfy is that the stabilizer is Abelian. We can express this in terms of a dual with respect to the symplectic product $*$. The dual of an additive code is additive, and the dual of the GF(4) conversion of a stabilizer is the GF(4) conversion of the normalizer. An additive code corresponds to an Abelian group if it is weakly self-dual.

We now have the main components we need to convert GF(4) codes into stabilizer codes.

Theorem 5.2. *Suppose C is an additive code over GF(4) which is weakly self-dual under $*$. Then C can be converted to a stabilizer S with parameters $[[n, k, d]]$. n is the number of physical registers in C . If C contains $K = 2^r$ codewords, then $k = n - r$. The distance d of S is the smallest weight of a member of $C^\perp \setminus C$, and in particular, d is at least equal to the distance of C^\perp .*

Note that there are a number of peculiar things involved in the conversion. In some sense C^\perp is more closely analogous to the quantum code we get, since the distance of S is at least equal to the distance of C^\perp . However, the number of encoded qubits is a formula that does not show up at all in classical coding theory (which is more interested in writing $K = 4^{r'}$ than $K = 2^r$), and of course the distance can be even larger than the distance of C^\perp because of degeneracy. And it is important to bear in mind that the dual is taken with respect to the symplectic product $*$ rather than the usual inner product.

The upshot is that if we look at the most common classical GF(4) codes and try to convert them into quantum codes, we won't get the most general stabilizer code. But who cares about that? We only want good codes, and if we can get them by looking at existing GF(4) codes, that's good enough.

5.2.4 Linear GF(4) Codes

For linear GF(4) codes, the conditions in theorem 5.2 simplify somewhat. In particular, we have the following:

Proposition 5.3. *If C is a linear GF(4) code, then its dual with respect to the inner product $\bar{x} \cdot y$ is the same as the dual with respect to $*$.*

1	ω	ω	1	0
0	1	ω	ω	1
1	0	1	ω	ω
ω	1	0	1	ω

Table 5.4: The five-qubit code converted to a GF(4) code.

Proof. We wish to show that $\bar{x} \cdot y = 0 \forall y \in C$ iff $x * y = 0 \forall y \in C$. The forward direction is trivial, so we only need to show the backwards direction.

Suppose $\text{tr}(\bar{x} \cdot y) = 0$, but $\bar{x} \cdot y = 1$. Then $\bar{x} \cdot (\omega y) = \omega$ and $x * (\omega y) = \text{tr}(\bar{x} \cdot (\omega y)) = 1$. Therefore, since C is linear, if $x * y = 0 \forall y \in C$, then $\bar{x} \cdot y = 0 \forall y \in C$. \square

This simplifies the procedure of checking for weakly self-dual codes because we can use the dual under the standard inner product and then take the conjugate rather than have to compute the dual under the unusual symplectic inner product $*$.

5.2.5 Example: Perfect Qubit Codes

Now let's look at some concrete examples of constructing stabilizer codes from GF(4) codes. It turns out that the GF(4) Hamming codes have the right properties, or rather their duals do.

Theorem 5.4. *The duals (with respect to the standard inner product) of the Hamming codes over GF(4) can be converted to stabilizer codes. The dual of the $[(4^r - 1)/3, (4^r - 1)/3 - r, 3]_4$ Hamming code becomes a $[[(4^r - 1)/3, (4^r - 1)/3 - 2r, 3]]$ qubit stabilizer code.*

Notice that the resulting stabilizer code encodes $(4^r - 1)/3 - 2r$ qubits whereas the Hamming code encodes $(4^r - 1)/3 - r$ GF(4) registers. This is a consequence of using a base of 2 instead of 4 to count registers, as mentioned in the discussion of theorem 5.2.

Proof. The Hamming codes are linear, so we need to show that a Hamming code contains its dual relative to $\bar{x} \cdot y$. Looking at the construction of the non-binary Hamming codes in the proof of theorem 4.14, we see that if we discard the last row of the parity check matrix, the first column is all 0, and then every other column is repeated four times. That means that the inner product of any two of these rows will be zero, since we end up adding the same thing four times.

Next, we should show that we also get 0 if we take the inner product of the last row with another row i . Since the first column of row i is 0, we can ignore the first column. The other columns break up into sets of four, within which row i has some value a repeated four times, while the last row runs over all the values in GF(4): 0, 1, ω , and ω^2 . Whatever the value of a ,

$$\bar{a}0 + \bar{a}1 + \bar{a}\omega + \bar{a}\omega^2 = \bar{a}(0 + 1 + \omega + \omega^2) = 0. \tag{5.17}$$

That proves that the parity check matrix of the Hamming code can be converted into a stabilizer. The number of encoded qubits follows from theorem 5.2.

The dual (with respect to $\bar{x} \cdot y$) of the dual (with respect to the standard inner product) of a Hamming code is just the conjugate of the Hamming code, produced by replacing x with \bar{x} everywhere in the code. The conjugate of a code has the same distance as the code since taking the conjugate does not change the weight, and therefore, the stabilizer codes we have derived have distance at least 3. In fact, these codes are non-degenerate, so the distance is exactly 3. \square

We get codes with parameters $[[5, 1, 3]]$, $[[21, 15, 3]]$, $[[85, 77, 3]]$, etc. The $[[5, 1, 3]]$ code produced this way (shown in table 5.5) is equivalent to the $[[5, 1, 3]]$ code we discussed before. Note that while the classical $[5, 3, 3]_4$ Hamming code has two rows in its parity check matrix, the quantum $[[5, 1, 3]]$ code has four generators of its stabilizer. The vectors $v, \omega v$ are considered linearly dependent for a code over GF(4), so we only need

					I	X	X	X	X
0	1	1	1	1	I	Z	Z	Z	Z
1	0	1	ω	ω^2	X	I	X	Z	Y
					Z	I	Z	Y	X

Table 5.5: The parity check matrix of the five-bit GF(4) Hamming code and the five-qubit code derived from it.

to include one of them in the parity check matrix, but they convert to Paulis which are independent when considered as operators on qubits, so we need to list them both in the stabilizer.

This family of quantum codes is interesting because, like the classical Hamming codes they are derived from, the codes in this family use up all of the error syndromes. The $[[\frac{4^r - 1}{3}, \frac{4^r - 1}{3} - 2r, 3]]$ code has $2r$ stabilizer generators, so 4^r error syndromes. There are $3n + 1 = 4^r$ zero and one-qubit errors, and the code is non-degenerate, so each syndrome is used exactly once.

It is unclear what the correct definition of a perfect degenerate QECC should be, but a non-degenerate quantum code is *perfect* if the number of correctable errors is exactly equal to the number of error syndromes. Thus, the QECCs derived from the GF(4) Hamming codes are perfect qubit codes.

Chapter 6

Symmetries Of Symmetries: The Clifford Group

When dealing with stabilizer codes, it is helpful to restrict attention to a set of quantum gates that is guaranteed to treat the code nicely. There exists a unitary operation that will take any subspace (for instance, a quantum error-correcting code) into any other subspace of the same dimension. Sometimes that's exactly what you want. However, if you've taken the effort to work with a code with a nice tractable description in terms of its stabilizer, you don't want your work ruined by using a poorly-thought-out unitary. In general, quantum gates will map the code space of a stabilizer code into some other code, which might not be a stabilizer code.

The Clifford group is a group of unitary gates that is specifically chosen so that it does not do this. If you start with a stabilizer code and perform a Clifford group gate, you will always have another stabilizer code. The key to this is using only unitary operations which can be thought of as permutations of the Pauli group. A Clifford group operation then just switches one stabilizer into another.

6.1 Definition of the Clifford Group

6.1.1 Motivation for the Clifford Group

Suppose we perform the unitary U on a state $|\psi\rangle$ from a stabilizer code \mathcal{S} . What happens to the stabilizer? Suppose $M \in \mathcal{S}$, so $M|\psi\rangle = |\psi\rangle$. We want to find M' for which $U|\psi\rangle$ is a +1 eigenstate. It turns out the correct choice is $M' = U M U^\dagger$:

$$M' U |\psi\rangle = U M U^\dagger U |\psi\rangle = U M |\psi\rangle = U |\psi\rangle, \tag{6.1}$$

since U is unitary. Running over all M in the stabilizer, we find that $\mathcal{S}' = \{U M U^\dagger | M \in \mathcal{S}\}$ is a set of operators for which all states $U|\psi\rangle$ are +1 eigenstates (where $|\psi\rangle$ is any element of $\mathcal{T}(\mathcal{S})$).

We'd *like* to say that \mathcal{S}' is the new stabilizer of the code, with code space $U(\mathcal{T}(\mathcal{S}))$. However, the catch is that without any additional constraint on U , \mathcal{S}' might contain many non-Paulis, and the stabilizer is supposed to be a subset of \mathcal{P}_n . In addition, the true stabilizer of the subspace $U(\mathcal{T}(\mathcal{S}))$ might contain additional Paulis that are not in \mathcal{S}' . Furthermore, $U(\mathcal{T}(\mathcal{S}))$ might not be a stabilizer code — the Paulis in $\mathcal{S}(U(\mathcal{T}(\mathcal{S})))$ might be insufficient to specify the subspace. (Recall definition 3.4.)

The Clifford group is a set of unitaries that does not have these complications. It maps stabilizers to stabilizers. Fortunately, the Clifford group contains many interesting quantum gates; unfortunately, it is not enough for a universal quantum computer. The Clifford group is sufficient for encoding stabilizer codes, and gives a good start on fault-tolerant operations, but eventually we will need to go beyond it.

6.1.2 Definition of the Clifford Group and Variants

Definition 6.1. The *Clifford group* C_n on n qubits is the normalizer of P_n in the unitary group $U(2^n)$. That is,

$$C_n = \{U \in U(2^n) \mid UPU^\dagger \in P_n \ \forall P \in P_n\}. \quad (6.2)$$

The name ‘‘Clifford group’’ is not particularly illuminating, perhaps. It is motivated by the idea that there might be a connection of some sort to Clifford algebras, but the connection is not very close. To add to the confusion, you might encounter the term ‘‘Clifford group’’ in the mathematics literature referring to a different group. In quantum information papers, Clifford group refers to definition 6.1 or one of its variants defined below. You might also encounter the terms ‘‘normalizer group,’’ ‘‘symplectic group’’ (or operations), or sometimes ‘‘stabilizer operations’’ for C_n . None of these terms is completely satisfactory, and ‘‘Clifford group’’ is the most widespread, so I will use that.

The Clifford group contains all gates of the form $e^{i\theta}I$, and if $U \in C_n$, then $e^{i\theta}U \in C_n$. As with the Paulis, global phase is frequently not significant for Clifford group elements. Indeed, it is *less* likely to matter for the Clifford group. Anticommutation is less important in the Clifford group than it is in the Pauli group, so we will usually consider not the full Clifford group, but the Clifford group with phases removed.

By definition, the Pauli group P_n is a normal subgroup of C_n . Sometimes it is better to consider the Clifford group with the Pauli subgroup modded out. The Pauli group only contains the phases $\pm 1, \pm i$, so even once the Pauli group is gone, there are still additional phases to worry about. Typically, we will want to remove those as well.

Definition 6.2.

$$\hat{C}_n = C_n / \{e^{i\theta}I\} \quad (6.3)$$

$$\check{C}_n = \hat{C}_n / \hat{P}_n. \quad (6.4)$$

I will refer to these variants as the ‘‘Clifford group,’’ sometimes without specifying which one I mean. Usually it doesn’t much matter, or can be deduced from context (or both).

6.1.3 Example Clifford Group Elements

Now let’s look at some Clifford gates. We know the Pauli group is a subgroup of C_n , so that gives us one set of examples. The Cliffords are defined to act on P_n by conjugation, and, as you’ll see shortly, the best way of characterizing an element of the Clifford group is usually by giving the action under conjugation. Suppose we have $P \in P_n \subset C_n$. What does it do to another Pauli Q under conjugation?

$$PQP^\dagger = (-1)^{c(P,Q)}QP^\dagger = (-1)^{c(P,Q)}Q. \quad (6.5)$$

Thus, $Q \mapsto \pm Q$, with the sign determined by whether P and Q commute or anticommute. The effect of conjugating by a Pauli is to rearrange the signs of other Paulis without changing their identity. This is easy to square with what we know about stabilizers from chapter 3: Applying the error P will move us from the $+1$ eigenspace of Q to the -1 eigenspace of Q iff P and Q anticommute. Moving to the -1 eigenspace is equivalent to modifying the stabilizer to contain $-Q$ instead of Q , which is the way we understand the action of Clifford group elements.

Another gate in the Clifford group is the Hadamard transform H . As we’ve discussed, the Hadamard transform switches the role of X and Z . This can be made concrete by looking at the conjugation action of H . You can work it out by multiplying together the matrices, but I’ll just tell you the answer:

$$HXH = Z \quad (6.6)$$

$$HYH = -Y \quad (6.7)$$

$$HZH = X. \quad (6.8)$$

$H^\dagger = H$, so I've skipped the adjoints in the above equations. As you can see, the action of H is indeed to switch X and Z . The effect on states is to switch Z eigenstates with X eigenstates:

$$|0\rangle \text{ (stabilizer } Z) \leftrightarrow |+\rangle = |0\rangle + |1\rangle \text{ (stabilizer } X) \quad (6.9)$$

$$|1\rangle \text{ (stabilizer } -Z) \leftrightarrow |-\rangle = |0\rangle - |1\rangle \text{ (stabilizer } -X) \quad (6.10)$$

While its action on Y is not as dramatic as its action on X and Z , H does not leave Y completely alone. Instead, it changes the sign of Y , so $+1$ eigenstates of Y get switched with -1 eigenstates of Y :

$$\frac{1}{\sqrt{2}}(|0\rangle + i|1\rangle) \leftrightarrow \frac{1}{2}[(1+i)|0\rangle + (1-i)|1\rangle] = \frac{e^{i\pi/4}}{\sqrt{2}}(|0\rangle - i|1\rangle). \quad (6.11)$$

The conjugation action doesn't tell us about the overall phase $e^{i\pi/4}$ introduced by H , but that has no physical significance anyway.

The other important thing about H 's action on Y is that I didn't need to specify it. The action of H on Y can be deduced from its action on X and Z . Conjugation is a *group homomorphism* — the conjugate of the product is the product of the conjugates:

$$UPQU^\dagger = (UPU^\dagger)(UQU^\dagger). \quad (6.12)$$

In this case,

$$HYH = H(iXZ)H = i(HXH)(HZH) = iZX = -Y. \quad (6.13)$$

Equivalently, we could deduce the action of H on X from the action on Y and Z , or the action on Z from the action on X and Y . The action on I , of course, will always be trivial ($UIU^\dagger = I$).

The Hadamard switches X and Z , but only changes the phase of Y . There are also Clifford group elements that switch other pairs of Paulis. For instance, $R_{\pi/4} \in C_n$:

$$R_{\pi/4}XR_{\pi/4}^\dagger = Y \quad (6.14)$$

$$R_{\pi/4}YR_{\pi/4}^\dagger = -X \quad (6.15)$$

$$R_{\pi/4}ZR_{\pi/4}^\dagger = Z. \quad (6.16)$$

This is not 100% analogous to H , since Z really is left alone by $R_{\pi/4}$, but Y picks up a minus sign under H . However, by multiplying with a Pauli, we can take care of that sign issue:

$$(YR_{\pi/4})X(YR_{\pi/4})^\dagger = Y \quad (6.17)$$

$$(YR_{\pi/4})Y(YR_{\pi/4})^\dagger = X \quad (6.18)$$

$$(YR_{\pi/4})Z(YR_{\pi/4})^\dagger = -Z. \quad (6.19)$$

You can deduce these equations by performing the conjugation action of $R_{\pi/4}$ and then following it with the conjugation action of Y , which switches the signs around.

You might wonder if we can completely get rid of the minus sign, rather than simply switching it to another location, by choosing an appropriate Pauli. The answer is no: any Pauli will commute with one Pauli and anticommute with two of them. Thus, it switches the sign of two of the three one-qubit Paulis X , Y , and Z . We can go from one minus sign to three minus signs, but never to zero or two. What we *can* do is change the signs of the generators of P_1 (or P_n when dealing with n qubits) in any way we like. For instance, to change $X \mapsto -X$, $Z \mapsto Z$, conjugate by Z . The change in the sign of Y is determined by the change in the signs of X and Z .

Among two-qubit Clifford group elements, the most famous gate is the CNOT gate, which has the following conjugation action on the Paulis:

$$X \otimes I \mapsto X \otimes X \quad (6.20)$$

$$Z \otimes I \mapsto Z \otimes I \quad (6.21)$$

$$I \otimes X \mapsto I \otimes X \quad (6.22)$$

$$I \otimes Z \mapsto Z \otimes Z. \quad (6.23)$$

We are dealing with the two-qubit Pauli group, and I have taken advantage of the homomorphism property of conjugation to just list the action of CNOT on a generating set of four Paulis for \mathbb{P}_2 . For any other Pauli, you can deduce the conjugation action of CNOT by multiplying as described above for the Hadamard.

6.1.4 Determining Clifford Group Element From Action on Generating Paulis

Describing the conjugation action of Clifford group elements certainly tells us a lot about the unitary we are dealing with, but you might worry that some information is being lost. Indeed, if $U' = e^{i\theta}U$, then

$$U'PU'^{\dagger} = e^{i\theta}UPU^{\dagger}e^{-i\theta} = UPU^{\dagger}, \quad (6.24)$$

so the conjugation action tells us nothing about the global phase of the unitary. Of course, the global phase doesn't have any physical significance, so this is not a great loss. But perhaps there are more important properties not captured by the conjugation action? The next theorem tells us that there are not:

Theorem 6.1. *Suppose U and V are unitaries which have the same action by conjugation on \mathbb{P}_n :*

$$UPU^{\dagger} = VPV^{\dagger} \quad (6.25)$$

for all $P \in \mathbb{P}_n$. Then $U = e^{i\theta}V$ for some θ .

The theorem does not apply only to Clifford group elements; U and V can be *any* unitary gates. When applied to the Clifford group, theorem 6.1 tells us that the conjugation action uniquely identifies elements of $\hat{\mathbb{C}}_n$. Consequently, we will be able to work with elements of $\hat{\mathbb{C}}_n$ using just the conjugation action.

Proof. Note that $V^{\dagger}U$ acts on the Pauli group trivially by conjugation ($P \mapsto P$), so it will be sufficient to consider the case of $V = I$ and show that $U = e^{i\theta}I$.

When $UPU^{\dagger} = P$ for all $P \in \mathbb{P}_n$, that means that U maps any stabilizer state to a state with the same stabilizer. For instance, the basis state $|j\rangle$ is a stabilizer state with stabilizer generated by $(-1)^{j_i}Z_i$, $i = 1, \dots, n$ (where j_i is the i th bit of j). The only states with this stabilizer are $e^{i\theta_j}|j\rangle$. We conclude that U is diagonal

$$U = \begin{pmatrix} e^{i\theta_1} & 0 & \dots & 0 \\ 0 & e^{i\theta_2} & \dots & 0 \\ \vdots & & \ddots & \\ 0 & & \dots & e^{i\theta_{2^n}} \end{pmatrix} \quad (6.26)$$

We still need to show that all θ_j are the same. Now consider stabilizer states with stabilizer generators X_1 and $(-1)^{j_{i-1}}Z_i$ for $i = 2, \dots, n$. The eigenstates of those stabilizers are of the form $e^{i\phi_j}|+\rangle|j\rangle$. Applying equation (6.26) to $|+\rangle|j\rangle$, we find

$$e^{i\theta_j} = e^{i\theta_{2^n-1+j}} = e^{i\phi_j}. \quad (6.27)$$

Similarly, looking at stabilizers with X_i in place of X_1 , we find that

$$e^{i\theta_j} = e^{i\theta_{2^n-i+j}}, \quad (6.28)$$

for any j such that the i th bit is 0. Applying all these equalities, we find that all $e^{i\theta_j}$ terms are equal, proving the theorem. \square

Which permutations of \mathbb{P}_n are allowed for a conjugation action by a Clifford group element? We know that conjugation is a group homomorphism. That means that there is no hope that we can freely choose images for any elements of \mathbb{P}_n except for a generating set. In addition, conjugation will always take $-I$ to $-I$. That means that conjugation must preserve commutation and anticommutation:

$$U(PQ)U^{\dagger} = U[(-1)^{c(P,Q)}QP]U^{\dagger} = (-1)^{c(P,Q)}(UQU^{\dagger})(UPU^{\dagger}) \quad (6.29)$$

$$= (UPU^{\dagger})(UQU^{\dagger}) = (-1)^{c(UPU^{\dagger},UQU^{\dagger})}(UQU^{\dagger})(UPU^{\dagger}). \quad (6.30)$$

Thus, $c(UPU^\dagger, UQU^\dagger) = c(P, Q)$.

The usual set of generators for \hat{P}_n is $\{X_i, Z_i\}$. To keep the right commutation relations, the images of X_i and Z_i must commute with the images of X_j and Z_j for $j \neq i$. However, the images of X_i and Z_i must anticommute with each other.

There is one additional constraint: Conjugation by U is an isomorphism, since it can be inverted by conjugating by U^\dagger . Therefore, the generators must map to an independent set of Paulis. However, this property actually follows from the commutation relations, so we don't need to impose it separately.

If these conditions are satisfied, a similar approach to the proof of theorem 6.1 gives a constructive method for finding the unitary corresponding to a given conjugation map. The same procedure applies even for conjugation by a non-Clifford unitary.

Procedure 6.1. Suppose the map $M : P_n \rightarrow U(2^n)$ is a group homomorphism, with $M(X_i) = \bar{X}_i$, $M(Z_i) = \bar{Z}_i$, such that

$$\begin{aligned} c(\bar{X}_i, \bar{X}_j) &= c(\bar{Z}_i, \bar{Z}_j) = 0, \\ c(\bar{X}_i, \bar{Z}_j) &= \delta_{ij}. \end{aligned} \tag{6.31}$$

(When \bar{X}_i and \bar{Z}_j are outside the Pauli group, the definition of $c(\cdot, \cdot)$ is the same, and $c(P, Q)$ is undefined here if P and Q neither commute or anticommute.)

The following procedure finds the matrix representation in the standard basis of a U for which conjugation by U performs M :

1. Find the state $|\psi_0\rangle$ which is a +1 eigenstate of \bar{Z}_i for all $i = 1, \dots, n$.
2. Let b be any number from 0 to $2^n - 1$, and let b_i be the i th bit of b . Let

$$\bar{X}(b) = \prod_i (\bar{X}_i)^{b_i}. \tag{6.32}$$

3. Let $|\psi_b\rangle = \bar{X}(b)|\psi_0\rangle$.
4. Let

$$U_{ab} = \langle a|\psi_b\rangle. \tag{6.33}$$

Proof of the validity of procedure 6.1. M is a group homomorphism and $Z_i = Z_i^\dagger$, so it follows that $\bar{Z}_i = \bar{Z}_i^\dagger$ as well. All of the \bar{Z}_i commute with each other, so $\Pi = \frac{1}{2^n} \prod_i (I + \bar{Z}_i)$ is a projector with trace 1. Thus, there is a unique state $|\psi_0\rangle$ (up to global phase) which is a +1 eigenstate of all \bar{Z}_i , as needed for step 1. The remaining steps of the procedure give a linear map $U|b\rangle = |\psi_b\rangle$.

I claim that $|\psi_b\rangle$ is an orthonormal basis, so U is unitary:

$$\langle \psi_b | \psi_{b'} \rangle = \langle \psi_0 | (\bar{X}(b))^\dagger \bar{X}(b') | \psi_0 \rangle \tag{6.34}$$

$$= \langle \psi_0 | \prod_i (\bar{X}_i)^{b'_i - b_i} | \psi_0 \rangle \tag{6.35}$$

using the fact that the \bar{X}_i commute with each other. As with the \bar{Z}_i s, $\bar{X}_i = \bar{X}_i^\dagger$. Then since $|\psi_0\rangle$ is a +1 eigenstate of \bar{Z}_i ,

$$\langle \psi_b | \psi_{b'} \rangle = \langle \psi_0 | \prod_i (\bar{X}_i)^{b'_i + b_i} | \psi_0 \rangle \tag{6.36}$$

$$= \langle \psi_0 | \bar{Z}_j \prod_i (\bar{X}_i)^{b_i + b'_i} | \psi_0 \rangle \tag{6.37}$$

$$= \prod_i (-1)^{(b_i + b'_i)c(\bar{Z}_j, \bar{X}_i)} \langle \psi_0 | \prod_i (\bar{X}_i)^{b_i + b'_i} \bar{Z}_j | \psi_0 \rangle \tag{6.38}$$

$$= (-1)^{b_i + b'_i} \langle \psi_0 | \prod_i (\bar{X}_i)^{b_i + b'_i} | \psi_0 \rangle \tag{6.39}$$

$$= (-1)^{b_i + b'_i} \langle \psi_b | \psi_{b'} \rangle. \tag{6.40}$$

It follows that if, for any j , $b_j \neq b'_j$, then $\langle \psi_b | \psi_{b'} \rangle = 0$ as desired. The states $|\psi_b\rangle$ we get from procedure 6.1 are normalized, so we have an orthonormal basis.

The last step is to prove that U acts by conjugation on the Pauli group according to M . Let us compute UZ_iU^\dagger and UX_iU^\dagger :

$$UZ_iU^\dagger|\psi_b\rangle = UZ_i|b\rangle \quad (6.41)$$

$$= (-1)^{b_i}|\psi_b\rangle. \quad (6.42)$$

Now,

$$\bar{Z}_i|\psi_b\rangle = \bar{Z}_i\bar{X}(b)|\psi_0\rangle \quad (6.43)$$

$$= (-1)^{b_i}\bar{X}(b)\bar{Z}_i|\psi_0\rangle \quad (6.44)$$

$$= (-1)^{b_i}|\psi_b\rangle. \quad (6.45)$$

Since the $|\psi_b\rangle$ form a basis, $UZ_iU^\dagger = \bar{Z}_i$.

Similarly,

$$UX_iU^\dagger|\psi_b\rangle = UX_i|b\rangle \quad (6.46)$$

$$= |\psi_{b'}\rangle, \quad (6.47)$$

where b' is b with the i th bit flipped ($b'_i = b_i + 1$, $b'_j = b_j$ for $j \neq i$).

$$\bar{X}_i|\psi_b\rangle = \bar{X}_i\bar{X}(b)|\psi_0\rangle \quad (6.48)$$

$$= \bar{X}(b')|\psi_0\rangle \quad (6.49)$$

$$= |\psi_{b'}\rangle, \quad (6.50)$$

so $UX_iU^\dagger = \bar{X}_i$. □

If you apply procedure 6.1 to a mapping which does not preserve commutation or anticommutation, you will still get a linear map, but it won't be unitary. Furthermore, it may not realize the conjugation action you wanted, so there is no real reason to apply procedure 6.1 unless the map you are working with satisfies equation (6.31).

6.1.5 The Clifford Group and the Symplectic Group

The conditions on the conjugation map performed by a Clifford group operation become somewhat more straightforward if we consider them in terms of the binary symplectic representation. A group homomorphism of the Paulis becomes a linear map on $2n$ -dimensional binary vectors. The requirement that the generators map to independent Paulis just says that the linear map has maximum rank. The constraint that conjugation preserves the commutation relations just means that the linear map must preserve the symplectic inner product:

$$v \odot w = (Mv) \odot (Mw) \quad (6.51)$$

$$v^T J w = v^T M^T J M w, \quad (6.52)$$

with J the $2n \times 2n$ matrix

$$J = \left(\begin{array}{c|c} 0 & I \\ \hline I & 0 \end{array} \right). \quad (6.53)$$

Running over all values of v and w , we find

$$J = M^T J M. \quad (6.54)$$

M is a member of the symplectic group $\text{Sp}(2n, \mathbb{Z}_2)$.

As usual, by going to the binary symplectic representation, we lose all information about the phases of Paulis in the stabilizer. An operation which only changes the phases of Paulis is always performed by conjugation by a Pauli:

Proposition 6.2. *If $UPU^\dagger = \pm P$ for all $P \in \mathcal{P}_n$, then $U = e^{i\theta}Q$ for $Q \in \mathcal{P}_n$ and some value of θ .*

Proof. I will show that any mapping $P \mapsto (-1)^{c_P}P$ which could possibly be performed by a unitary can also be performed by conjugation by some $Q \in \mathcal{P}_n$. Then the proposition follows from theorem 6.1.

For any unitary U , conjugation performs a group homomorphism, so

$$c_{PQ} = c_P + c_Q. \quad (6.55)$$

In particular, the mapping is completely determined by its action on X_i and Z_i for $i = 1, \dots, n$. For any c_P consistent with equation (6.55), I claim there exists $Q \in \mathcal{P}_n$ that implements it. By equation (6.5), we need to find Q such that $c(Q, X_i) = c_{X_i}$ and $c(Q, Z_i) = c_{Z_i}$. By lemma 3.15, such a Q exists, though there is only one. Specifically,

$$Q = \bigotimes_{i=1}^n Q_i \quad (6.56)$$

with

$$Q_i = \begin{cases} I & \text{if } c_{X_i} = c_{Z_i} = 0 \\ X & \text{if } c_{X_i} = 0 \text{ and } c_{Z_i} = 1 \\ Y & \text{if } c_{X_i} = 1 \text{ and } c_{Z_i} = 1 \\ Z & \text{if } c_{X_i} = 1 \text{ and } c_{Z_i} = 0 \end{cases} \quad (6.57)$$

□

It follows from proposition 6.2 that elements of $\check{\mathcal{C}}_n = \hat{\mathcal{C}}_n / \hat{\mathcal{P}}_n$ correspond uniquely to symplectic operations. Indeed, $\check{\mathcal{C}}_n$ is *exactly* the symplectic group.

Theorem 6.3. $\check{\mathcal{C}}_n \cong \text{Sp}(2n, \mathbb{Z}_2)$. *Equivalently, for every map $X_i \mapsto \bar{X}_i$, $Z_i \mapsto \bar{Z}_i$ with $\bar{X}_i, \bar{Z}_i \in \mathcal{P}_n$ and satisfying the correct commutation relations, there exists $U \in \mathcal{C}_n$ which performs that map under conjugation, and U is unique up to overall phase.*

Proof. Most of the pieces of this theorem are derived from theorem 6.1, equation (6.54), and proposition 6.2. The only remaining piece needed to complete the characterization is to show that every symplectic operation can be realized by a Clifford group operation. This is straightforward: Given any linear map from $\hat{\mathcal{P}}_n$ to $\hat{\mathcal{P}}_n$, we can lift it to a group homomorphism $M : \mathcal{P}_n \rightarrow \mathcal{P}_n$ by choosing arbitrary signs for the images of X_i and Z_i . When the original linear map is symplectic, M satisfies equation (6.31), so using procedure 6.1, we find a unitary U realizing M and thus realizing the symplectic map on the binary symplectic representation. U maps Paulis to Paulis under conjugation, so $U \in \mathcal{C}_n$. □

6.2 Classical Simulation of the Clifford Group

One of the most interesting and useful things about the Clifford group is also one of the most disappointing. If $U \in \mathcal{C}_n$, we can specify it by giving a global phase plus the images of $2n$ Paulis $X_1, Z_1, \dots, X_n, Z_n$. Each image is also an n -qubit Pauli, so requires at most $2n + 1$ bits to specify. Thus, a Clifford group element can be specified using only $(2n + 1)(2n)$ bits plus a global phase. Contrast that with a general unitary $U \in \text{U}(2^n)$, which needs 2^{2n} real parameters to specify. Clifford group elements have a much more succinct representation than general unitary gates.

Indeed, circuits made out of Clifford group gates have a much stronger property: they can be efficiently simulated on a classical computer. This is a very useful fact, since it makes working with even relatively large Clifford group circuits tractable. For instance, stabilizer codes are exactly those QECCs which can be

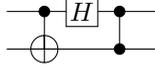


Figure 6.1: An example 2-qubit circuit made of Clifford group gates.

encoded using a circuit composed of Clifford group gates. The ability to easily describe a stabilizer code via its stabilizer is one aspect of the efficient simulatability of the Clifford group gates which form its encoder. We'll also see in part II that this property plays a big role in helping us find fault-tolerant implementations of Clifford group gates.

However, there is a price to be paid. Because Clifford group gates can classically be simulated, it means that a circuit composed only of Clifford group gates cannot access the full power of quantum computation. (Presumably; as with most statements about computational power, this one relies on some unproven complexity-theoretic assumptions, in this case the assumption that quantum computers are computationally more powerful than classical computers.) When we get to fault tolerance, that means that we'll need to venture outside the Clifford group in order to get a universal set of fault-tolerant gates.

6.2.1 Simulation of a Unitary Circuit of Clifford Group Gates

If we have a unitary circuit consisting only of Clifford group gates, the simulation procedure is quite straightforward. Note that if $U, V \in \mathcal{C}_n$, and $U : P \mapsto Q$, $V : Q \mapsto R$, then $(VU) : P \mapsto R$. That's really all there is to the simulation.

Procedure 6.2. You are given a circuit consisting of a product $U = \prod_{i=1}^m U_i$, with $U_i \in \mathcal{C}_n$. (I.e., U_1 is the first gate performed, and U_m is the last gate to be performed.) Each U_i can be specified by its action on the generators of the Pauli group, $U_i : X_j \mapsto U_i(X_j)$, $U_i : Z_j \mapsto U_i(Z_j)$. Then the action of the overall circuit U on the Pauli group can be determined as follows:

1. Initialize $\bar{X}_j = X_j$ and $\bar{Z}_j = Z_j$.
2. Starting with $i = 1$, and stepping through i up to the last gate m , repeat the following steps:
 - (a) Calculate $U_i(\bar{X}_j)$. This can be done by writing \bar{X}_j as a product of single-qubit X s and Z s and applying equation (6.12).
 - (b) Let the new value of \bar{X}_j be $U_i(\bar{X}_j)$.
 - (c) Similarly, replace \bar{Z}_j by $U_i(\bar{Z}_j)$.
3. The overall circuit U has the action $X_j \mapsto \bar{X}_j$, $Z_j \mapsto \bar{Z}_j$, using the final values of \bar{X}_j and \bar{Z}_j .

When the circuit consists of only two-qubit gates, updating an \bar{X} or \bar{Z} operator only requires updating two qubits in the decomposition, since all qubits not affected by the gate retain the same Pauli values. Therefore, each gate can be updated in a total time $O(n)$ (since we need to update $2n$ \bar{X} and \bar{Z} operators). The simulation of the full circuit thus takes a time $O(nm)$.

As an example, consider the simple circuit given in figure 6.1. Following the above procedure, we find that this circuit has the following action on Paulis:

$$\begin{aligned}
 \bar{X}_1 : X \otimes I &\rightarrow X \otimes X \rightarrow Z \otimes X \rightarrow I \otimes X \\
 \bar{Z}_1 : Z \otimes I &\rightarrow Z \otimes I \rightarrow X \otimes I \rightarrow X \otimes Z \\
 \bar{X}_2 : I \otimes X &\rightarrow I \otimes X \rightarrow I \otimes X \rightarrow Z \otimes X \\
 \bar{Z}_2 : I \otimes Z &\rightarrow Z \otimes Z \rightarrow X \otimes Z \rightarrow X \otimes I.
 \end{aligned} \tag{6.58}$$

6.2.2 Simulation of a Unitary Circuit on a Stabilizer Subspace

Another interesting case is when we have some constraints on the input state of the circuit. For instance, the circuit may involve some ancillas, or it may be intended to be performed on a state encoded in a QECC. It may even be that the full input state of the circuit is specified, and we wish to determine the exact output state of the circuit. When the input state is a stabilizer state or lies within a stabilizer code, there still exists an efficient simulation procedure.

If the initial state is completely specified as a stabilizer state, we can keep track of the behavior of the generators M_1, \dots, M_n of the stabilizer. As before, we step through gates, updating the stabilizer generators at each step just as we updated \bar{X} and \bar{Z} before. When M_j is a generator of the stabilizer before we perform gate U_i , $U_i(M_j)$ is a generator of the stabilizer after the gate, so this procedure allows us to learn the stabilizer of the output state of the circuit. There are only two small differences from the case without a constraint. First of all, we need keep track of only n generators instead of $2n$ Paulis. The other difference is that the generators of the stabilizer are not unique, so after performing any gate, if it is convenient to do so, we may choose a new set of generators of the current stabilizer. There is no requirement to do so; only do it if it is clearly going to simplify the computation.

When the input state is only partially specified by a stabilizer code, we keep track both of the stabilizer, which now has $n - k$ generators, and of $2k$ logical \bar{X} and \bar{Z} operators. The case with $k = n$ is the “no constraint” case of the previous subsection, and the case with $k = 0$ is the stabilizer state case just discussed. Again, the stabilizer generators are non-unique, so at any step, we may choose a new set of generators. The logical \bar{X} and \bar{Z} operators are also now non-unique, so at any step we may multiply them by elements of the stabilizer, which means we are choosing a different coset representative.

Procedure 6.3. You are given a circuit consisting of a product $\prod_{i=m}^1 U_i$, with $U_i \in \mathcal{C}_n$, which is to be performed on an arbitrary input state from stabilizer code \mathcal{S} . \mathcal{S} encodes k qubits ($0 \leq k \leq n$), has generators M_1, \dots, M_{n-k} , and has logical Pauli operators \bar{X}_j, \bar{Z}_j , $j = 1, \dots, k$. Each U_i can be specified by its action on the generators of the Pauli group, $U_i : X_j \mapsto U_i(X_j)$, $U_i : Z_j \mapsto U_i(Z_j)$. Then the action of the overall circuit on the stabilizer code can be determined as follows:

1. Initialize variables $N_j = M_j$ ($j = 1, \dots, n - k$), $\bar{X}'_j = \bar{X}_j$ and $\bar{Z}'_j = \bar{Z}_j$ ($j = 1, \dots, k$) to be the values given by the stabilizer code input.
2. Starting with $i = 1$, and stepping through i up to the last gate m , repeat the following steps:
 - (a) Calculate $U_i(\bar{X}'_j)$ for $j = 1, \dots, k$. This can be done by writing \bar{X}'_j as a product of single-qubit X s and Z s and applying equation (6.12).
 - (b) Let the new value of \bar{X}'_j be $U_i(\bar{X}'_j)$ for $j = 1, \dots, k$.
 - (c) Similarly, replace \bar{Z}'_j by $U_i(\bar{Z}'_j)$ for $j = 1, \dots, k$.
 - (d) Replace N_j by $U_i(N_j)$ for $j = 1, \dots, n - k$.
 - (e) The current stabilizer \mathbb{T} is generated by $\langle N_1, \dots, N_{n-k} \rangle$
 - (f) If desired, choose a new set of generators for \mathbb{T} .
 - (g) If desired, rewrite $\bar{X}'_j = N\bar{X}'_j$ or $\bar{Z}'_j = N\bar{Z}'_j$ with $N \in \mathbb{T}$.
3. The output state lies in a stabilizer code with generators equal to the final values of N_j , $j = 1, \dots, n - k$. The encoded state has undergone the Clifford group operation given by the transformation $\bar{X}_j \mapsto \bar{X}'_j$, $\bar{Z}_j \mapsto \bar{Z}'_j$.

An important special case is when some qubits are completely specified (e.g., to be $|0\rangle$), and the remaining qubits are completely unconstrained. In that case, the stabilizer is the stabilizer of the ancilla qubits, and the initial logical operators \bar{X} and \bar{Z} are the Paulis on the unconstrained input qubits.

As an example of the expanded procedure, let us consider the circuit in figure 6.1 again, but this time with the second qubit initially in the state $|0\rangle$. We now begin with $M_1 = I \otimes Z$, $\overline{X}_1 = X \otimes I$, and $\overline{Z}_1 = Z \otimes I$. Using the same analysis as before, we find that the final value of the stabilizer and logical operators are:

$$N_1 = X \otimes I \quad (6.59)$$

$$\overline{X}'_1 = I \otimes X \quad (6.60)$$

$$\overline{Z}'_1 = X \otimes Z. \quad (6.61)$$

We can choose a new coset representative for \overline{Z}'_1 : $I \otimes Z = (X \otimes I)(X \otimes Z)$. Then we can see that the overall transformation performed is to move the input qubit from the first qubit to the second qubit. In the output state, the first qubit is fixed to be $|0\rangle + |1\rangle$.

6.2.3 Measurement of Paulis

The final step is to add measurements to our simulation. I will phrase this in the most general way, where we measure the eigenvalue of an arbitrary Pauli operator on n qubits, but it would be equivalent to just consider measurement of single qubits in the standard basis. This is because measurement of any Pauli can be implemented via single-qubit measurements plus Clifford group operations. For simplicity, we restrict attention to Paulis with overall phase ± 1 , so the eigenvalues are ± 1 .

The analysis of measurements is somewhat more complicated than the analysis of unitary gates. For one thing, there are three separate cases to consider. The first case is when the Pauli P we wish to measure is in the current stabilizer \mathbb{T} up to a phase: $\pm P \in \mathbb{T}$. In that case, the measurement outcome is just ± 1 ($+1$ if $P \in \mathbb{T}$ and -1 if $-P \in \mathbb{T}$), and the state does not change, since the state being measured is an eigenstate of P .

The second case is when $P \in \mathbb{N}(\mathbb{T})$. Now, measurement of P constitutes a measurement of some of the encoded data. We must rewrite P as a product of the logical Paulis \overline{X} and \overline{Z} , which determines that P corresponds to some logical Pauli \overline{Q} . The measurement output distribution of P is the same as the measurement output distribution of Q on the original input qubit. We must also update the encoded state for the effects of the measurement. Finally, the measurement has outcome ± 1 , so we know that the post-measurement state is a $+1$ eigenstate of $\pm P$. In other words, $\pm P$ has joined the stabilizer.

The third case, when $P \notin \mathbb{N}(\mathbb{T})$, is the most complicated. Since $P \notin \mathbb{N}(\mathbb{T})$, $\exists M \in \mathbb{T}$ s.t. $\{P, M\} = 0$. Suppose $|\psi\rangle \in \mathcal{T}(\mathbb{T})$, so $M|\psi\rangle = |\psi\rangle$. The expectation value of a measurement of P on $|\psi\rangle$ is

$$\langle \psi | P | \psi \rangle = \langle \psi | P M | \psi \rangle \quad (6.62)$$

$$= \langle \psi | M (-P) | \psi \rangle \quad (6.63)$$

$$= -\langle \psi | P | \psi \rangle \quad (6.64)$$

$$= 0. \quad (6.65)$$

(since $M^2 = I$, meaning $M = M^\dagger$). Thus, the probability of a $+1$ outcome and a -1 outcome for the measurement are both $1/2$.

We can also calculate the residual state. If the measurement has outcome $(-1)^b$, the state afterwards is $|\psi'\rangle = (1/\sqrt{2})(I + (-1)^b P)|\psi\rangle$ (taking into account normalization). If $N \in \mathbb{T}$ commutes with P , then $|\psi'\rangle$ is still a $+1$ eigenstate of N :

$$N|\psi'\rangle = \frac{1}{\sqrt{2}}N(I + (-1)^b P)|\psi\rangle \quad (6.66)$$

$$= \frac{1}{\sqrt{2}}(I + (-1)^b P)N|\psi\rangle \quad (6.67)$$

$$= \frac{1}{\sqrt{2}}(I + (-1)^b P)|\psi\rangle = |\psi'\rangle. \quad (6.68)$$

$|\psi'\rangle$ is also a $+1$ eigenstate of $(-1)^b P$. That means that it is *not* an eigenstate of any M with $\{M, P\} = 0$, even if $M \in \mathbb{T}$. Define a stabilizer \mathbb{T}' generated by $(-1)^b P$ plus all $N \in \mathbb{T}$ s.t. $[N, P] = 0$. The stabilizer of the remaining state after the measurement certainly contains \mathbb{T}' . As we will see in a moment, that is all it contains.

How big is \mathbb{T}' ? To answer this, we need to know how many elements of \mathbb{T} commute with P .

Proposition 6.4. *Let \mathbb{S} be a stabilizer, and let $P \notin \mathbb{N}(\mathbb{S})$ be a Pauli that does not commute with every element of \mathbb{S} . Then exactly half of the elements of \mathbb{S} commute with P . Furthermore, for any coset $\bar{Q} \in \mathbb{N}(\mathbb{S})/\mathbb{S}$, exactly half the elements of \bar{Q} commute with P .*

Proof. If $P \notin \mathbb{N}(\mathbb{S})$, let $M \in \mathbb{S}$ be an element of the stabilizer that anticommutes with P , $\{M, P\} = 0$. Then we can pair all elements of \mathbb{S} :

$$N \leftrightarrow MN. \quad (6.69)$$

Note that $MN^2 = M$, so each element of \mathbb{S} is a member of only one pair. The reason for doing this pairing is that N commutes with P iff MN anticommutes with P :

$$c(MN, P) = c(M, P) + c(N, P) = 1 + c(N, P). \quad (6.70)$$

Thus, the number of elements of \mathbb{S} that commute with P is equal to the number of elements that anticommute with P .

Similarly, we can pair elements of the cosets representing logical Paulis, $N \leftrightarrow MN$, using the same $M \in \mathbb{S}$. Both N and MN are in the same coset \bar{Q} , and again, exactly one of them commutes with P and one anticommutes with P . Thus, half of the coset \bar{Q} commutes with P . \square

Corollary 6.5. $|\mathbb{T}'| = |\mathbb{T}|$.

The subspace of post-measurement states for a given measurement outcome can be no larger than the space of possible pre-measurement states, but it could potentially be smaller if multiple states collapse in the measurement to the same final state. However, as another corollary of proposition 6.4, we find that this cannot happen. If the code space of \mathbb{T} contains k logical qubits, one possible basis for the code space is the set of 2^k codewords which are eigenstates of $\bar{Z}_1, \dots, \bar{Z}_k$. Each of these codewords is a stabilizer state with stabilizer generated by $\mathbb{T}, \pm\bar{Z}_1, \dots, \pm\bar{Z}_k$. By proposition 6.4 and the analysis preceding it, the stabilizer of the post-measurement state is generated by \mathbb{T}' and by $\pm\bar{Z}_1, \dots, \pm\bar{Z}_k$, with the same eigenvalues. However, we must make certain that each coset is represented by an element that commutes with P ; such an element always exists by proposition 6.4. Since each of the 2^k basis states gets mapped to a distinct stabilizer state, the code space after the measurement also has dimension 2^k . This implies that \mathbb{T}' is the actual stabilizer post-measurement, as claimed.

Furthermore, this logic also tells us that the \mathbb{T} -coset of the logical \bar{Z}_i operator gets replaced by the \mathbb{T}' -coset which contains an element of the original \bar{Z}_i that commutes with P . The same reasoning tells us how to find the new versions of the other logical Paulis. In short, we now know how to update both the stabilizer and the logical Paulis after measurement of a Pauli operator.

Theorem 6.6. *Suppose our system is a codeword of the stabilizer code \mathbb{T} , with generators N_1, \dots, N_{n-k} and logical Paulis \bar{X}_i and \bar{Z}_i ($i = 1, \dots, k$), and we measure the eigenvalue of Pauli $P \notin \mathbb{N}(\mathbb{T})$. Then, conditioned on outcome $(-1)^b$, the stabilizer and logical Paulis of the system after the measurement can be determined as follows:*

1. Find generator $M \in \mathbb{T}$ such that $\{M, P\} = 0$. If necessary, reorder the generators so that $M = N_1$.
2. For each generator N_i , $i > 1$, if $[N_i, P] = 0$, let $N'_i = N_i$. If $\{N_i, P\} = 0$, let $N'_i = N_i M$.
3. Let $N'_1 = (-1)^b P$.
4. The new stabilizer \mathbb{T}' is generated by N'_1, \dots, N'_{n-k} .

5. Pick a representative \bar{Z}_i for each of the logical Z operators for T . If $[\bar{Z}_i, P] = 0$, let $\bar{Z}'_i = \bar{Z}_i$; otherwise let $\bar{Z}'_i = \bar{Z}_i M$. \bar{Z}'_i is a representative for the new logical Z_i operator.
6. Similarly, pick a representative \bar{X}_i for each of the logical X operators for T . If $[\bar{X}_i, P] = 0$, let $\bar{X}'_i = \bar{X}_i$; otherwise let $\bar{X}'_i = \bar{X}_i M$. \bar{X}'_i is a representative for the new logical X_i operator.

One interesting twist on this system is that we can essentially control the measurement outcome. In particular, suppose we measure $P \notin \mathsf{N}(\mathsf{T})$ and get outcome -1 . In theorem 6.6, we identify an element M from the old stabilizer T which anticommutes with P . However, M commutes with generators 2 through $n - k$ of T' , since they are also elements of T . M also commutes with the standard coset representatives for \bar{X}'_i and \bar{Z}'_i . Thus, if we perform M on the post-measurement state, we transform the stabilizer (according to procedure 6.3) only by changing $N'_1 = -P$ to $+P$. The logical operators are unchanged. This is exactly the same state as we would have gotten (according to theorem 6.6) if the original measurement outcome had been $+1$.

Putting everything together, we get the following procedure for simulating Clifford group gates and Pauli measurements on either a fixed initial stabilizer state or a partially specified state with some free (logical) qubits:

Procedure 6.4. You are given a circuit consisting of a m -element sequence of unitary Clifford group gates and Pauli measurements. At step number i , the circuit calls for either Clifford group gate $U_i \in \mathsf{C}_n$ (specified by its action on Paulis $U_i : Q \mapsto U_i(Q)$) or measurement of the eigenvalue of $P_i \in \mathsf{P}_n$ (assuming P_i has eigenvalues ± 1 , not $\pm i$). The initial state of the circuit is given by stabilizer S , which encodes k qubits ($0 \leq k \leq n$), has generators M_1, \dots, M_{n-k} , and has logical Pauli operators $\bar{X}_j, \bar{Z}_j, j = 1, \dots, k$.

1. Initialize variables $k' = k, N_j = M_j$ ($j = 1, \dots, n - k$), $\bar{X}'_j = \bar{X}_j$ and $\bar{Z}'_j = \bar{Z}_j$ ($j = 1, \dots, k$) to be the values given by the stabilizer code input. Let the variable state $|\bar{\psi}\rangle$ be the initial encoded state of the system, corresponding to the logical state $|\psi\rangle$.
2. Perform the following procedure for each step i in order, for $i = 1, \dots, m$:
 - (a) If at step i , the circuit calls for Clifford gate U_i :
 - i. Calculate $U_i(\bar{X}'_j)$ for $j = 1, \dots, k'$. This can be done by writing \bar{X}'_j as a product of single-qubit X s and Z s and applying equation (6.12).
 - ii. Let the new value of \bar{X}'_j be $U_i(\bar{X}'_j)$ for $j = 1, \dots, k'$.
 - iii. Similarly, replace \bar{Z}'_j by $U_i(\bar{Z}'_j)$ for $j = 1, \dots, k'$.
 - iv. Replace N_j by $U_i(N_j)$ for $j = 1, \dots, n - k'$.
 - v. The current stabilizer T is generated by $\langle N_1, \dots, N_{n-k'} \rangle$.
 - vi. If desired, choose a new set of generators for T .
 - vii. If desired, rewrite $\bar{X}'_j = N \bar{X}'_j$ or $\bar{Z}'_j = N \bar{Z}'_j$ with $N \in \mathsf{T}$.
 - (b) If at step i , the circuit calls for measurement of P_i , determine whether $\pm P_i$ is in T , if $P_i \in \mathsf{N}(\mathsf{T}) \setminus \mathsf{T}$, or if P_i is not in $\mathsf{N}(\mathsf{T})$, as follows:
 - i. Determine if P_i commutes with all generators N_j of the current stabilizer T . If not, then $P_i \notin \mathsf{N}(\mathsf{T})$.
 - ii. If P_i does commute with all generators, determine if $\pm P_i \in \mathsf{T}$. This can be done by writing P_i and the generators N_j in their binary symplectic representations and seeing if v_{P_i} is in the linear span of the v_{N_j} .
 - iii. If P_i commutes with all generators but is not in T , then $P_i \in \mathsf{N}(\mathsf{T}) \setminus \mathsf{T}$.
 - (c) If at step i , the circuit calls for measurement of P_i with $\pm P_i \in \mathsf{T}$:
 - i. Evaluate whether $+P_i \in \mathsf{T}$ or $-P_i \in \mathsf{T}$. This can be done using linear algebra and the binary symplectic representations of P_i and N_j to write $P_i = \pm \prod_{j=1}^{n-k'} N_j^{s_j}$.

- ii. If $+P_i \in \mathbb{T}$, return measurement result $+1$. If $-P_i \in \mathbb{T}$, return measurement result -1 .
- (d) If at step i , the circuit calls for measurement of P_i with $P_i \in \mathbb{N}(\mathbb{T}) \setminus \mathbb{T}$:
 - i. Write $P_i = \prod_{j=1}^{n-k'} N_j^{s_j} \prod_{j'=1}^k \overline{X}_j^{t_j} \overline{Z}_j^{u_j}$, with $s_j, t_j, u_j \in \{0, 1\}$. This can be done using linear algebra and the binary symplectic representations of the Paulis.
 - ii. Let $Q = \prod_{j'=1}^k X_j^{t_j} Z_j^{u_j}$.
 - iii. Perform a measurement of Q on the current logical state $|\psi\rangle$. Return the measurement result $(-1)^b$, and update $|\overline{\psi}\rangle$ accordingly. Note that this step may not be efficient, depending on the current state $|\psi\rangle$.
 - iv. Reduce k' by 1, and add $(-1)^b P_i$ to the stabilizer as a new generator $N_{n-k'+1}$. Update \mathbb{T} accordingly, and if desired, choose a new set of generators for \mathbb{T} .
 - v. Update the logical operators \overline{X}'_j and \overline{Z}'_j to relate to the new $|\overline{\psi}\rangle$. Again, this step may not be efficient.
- (e) If at step i , the circuit calls for measurement of P_i with $P_i \notin \mathbb{N}(\mathbb{T})$:
 - i. Choose a uniformly random bit b and return measurement result $(-1)^b$.
 - ii. Find generator $M \in \mathbb{T}$ such that $\{M, P_i\} = 0$. If necessary, reorder the generators so that $M = N_1$.
 - iii. For each generator N_j , $j > 1$, if $[N_j, P_i] = 0$, leave N_j unchanged. If $\{N_j, P_0\} = 0$, replace N_j by $N_j M$.
 - iv. Replace N_1 by $(-1)^b P_i$.
 - v. Update the stabilizer \mathbb{T} with the revised generators. If desired, choose new generators for the revised \mathbb{T} .
 - vi. Pick a representative \overline{Z}'_j for each of the logical Z operators for \mathbb{T} . If $[\overline{Z}'_j, P_i] = 0$, leave \overline{Z}'_j unchanged; otherwise replace \overline{Z}'_j by $\overline{Z}'_j M$. The updated \overline{Z}'_j is a representative for the new logical Z_j operator; choose a new representative using the updated \mathbb{T} if desired.
 - vii. Similarly, pick a representative \overline{X}'_j for each of the logical X operators for \mathbb{T} . If $[\overline{X}'_j, P_i] = 0$, leave \overline{X}'_j unchanged; otherwise replace \overline{X}'_j by $\overline{X}'_j M$. The updated \overline{X}'_j is a representative for the new logical X_j operator; choose a new representative using the updated \mathbb{T} if desired.

3. The output state lies in a stabilizer code with generators equal to the final values of N_j , $j = 1, \dots, n-k'$. The encoded state has had $k - k'$ logical qubits measured, and has undergone the transformation $\overline{X}_j \mapsto \overline{X}'_j$, $\overline{Z}_j \mapsto \overline{Z}'_j$ on the remaining qubits.

Generally, for this sort of simulation, we only consider circuits where measurement of $P_i \in \mathbb{N}(\mathbb{T}) \setminus \mathbb{T}$ does not occur, so that the simulation is an efficient one. For instance, if the circuit contains no measurements, this is not an issue, or equally if it contains no logical qubits (so $\mathbb{N}(\mathbb{T}) = \mathbb{T}$). In part II, we will see circuits that have measurements and some logical qubits, but they have been specially designed to avoid the troublesome case.

Theorem 6.7 (Gottesman-Knill). *Given a circuit starting from an initial stabilizer state, followed by a sequence of Clifford group operations and Pauli measurements, which may depend on classical computations performed on previous measurement results, there is an efficient classical simulation of the circuit.*

If you follow procedure 6.4, you will find a time of $O(n^3m)$ is needed to do the simulation, taking into account all the linear algebra manipulations needed to process measurement. However, by keeping track of slightly more information, this can be reduced to $O(n^2m)$.

To illustrate the measurement procedure, let us consider a variation of the example circuit from figure 6.1 which has some measurements in it. The revised circuit is given in figure 6.2. The first few steps are as before. After the Hadamard, we have the following stabilizer and logical Paulis:

$$\begin{aligned}
N_1 : & I \otimes Z \rightarrow X \otimes Z \\
\overline{X} : & X \otimes I \rightarrow Z \otimes X \\
\overline{Z} : & Z \otimes I \rightarrow X \otimes I.
\end{aligned} \tag{6.71}$$

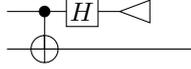


Figure 6.2: One-qubit teleportation: an example of a 2-qubit circuit made of Clifford group gates and measurement.

Then we measure $P = Z \otimes I$. In this case, there is only one element of the stabilizer, and it anticommutes with P . Suppose we get measurement outcome $+1$. By theorem 6.6, replace N_1 with P . \bar{X} commutes with P , so its representative need not change. However, we can choose a new coset representative to take advantage of the new stabilizer: $(Z \otimes X)(Z \otimes I) = I \otimes X$. \bar{Z} anticommutes with P , so we *must* choose a new coset representative

$$\bar{Z}N_1 = (X \otimes I)(X \otimes Z) = I \otimes Z. \quad (6.72)$$

After the measurement, the stabilizer and logical Paulis are thus

$$N_1 : Z \otimes I \quad (6.73)$$

$$\bar{X} : I \otimes X \quad (6.74)$$

$$\bar{Z} : I \otimes Z \quad (6.75)$$

The input qubit has been moved to the second qubit for the output, and the output value of the first qubit is $|0\rangle$.

In the case where the measurement result is -1 , the stabilizer is $-P$ instead of P . The new representative for \bar{X} also acquires this minus sign, so the final state is as follows:

$$N_1 : -Z \otimes I \quad (6.76)$$

$$\bar{X} : -I \otimes X \quad (6.77)$$

$$\bar{Z} : I \otimes Z \quad (6.78)$$

The output state of the first qubit is $|1\rangle$, and the output state of the second qubit is the input data qubit, but with a phase flip (Z) performed on it.

6.3 Generators of the Clifford Group

In section 6.1.3, we saw three common gates which are also elements of the Clifford group: H , $R_{\pi/4}$ and CNOT. In a sense, they are the *only* elements of the Clifford group, because they generate the Clifford group. That is:

Theorem 6.8. *Any gate in the n -qubit Clifford group C_n can be written as a product of $e^{i\theta}I$, H_i , $R_{\pi/4,i}$, and $\text{CNOT}_{i,j}$, with $i, j = 1, \dots, n$ and $\theta \in [0, 2\pi)$.*

$\text{CNOT}_{i,j}$ means CNOT with qubit i as the control and qubit j as the target.

Proof. First, note that the Pauli group is inside the group generated by H and $R_{\pi/4}$: $Z = (R_{\pi/4})^2$, and $X = HZH$. Second, global phases are covered by the gates $e^{i\theta}I$. To finish the proof, we thus need only to prove that the symplectic representations of H , $R_{\pi/4}$ and CNOT generate \check{C}_n .

It will be helpful to also use two additional gates: SWAP and C – Z. Both are in the Clifford group, and can easily be performed as a product of H , $R_{\pi/4}$, and CNOT, so we can use them freely without explicitly adding them to the generating set. In particular, SWAP is the product of 3 CNOT gates with alternating directions, and C – Z = $(I \otimes H)\text{CNOT}(I \otimes H)$.

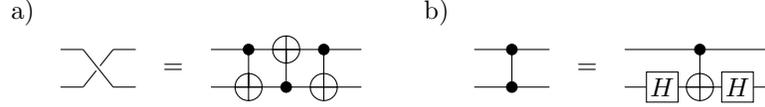


Figure 6.3: a) Clifford group circuit for the SWAP gate. b) Clifford group circuit for the C – Z gate.

We can take advantage of theorem 6.3. The symplectic representation of H is

$$\left(\begin{array}{c|c} 0 & 1 \\ \hline 1 & 0 \end{array} \right), \quad (6.79)$$

the symplectic representation of $R_{\pi/4}$ is

$$\left(\begin{array}{c|c} 1 & 0 \\ \hline 1 & 1 \end{array} \right), \quad (6.80)$$

and the symplectic representation of CNOT is

$$\left(\begin{array}{cc|cc} 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ \hline 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 \end{array} \right). \quad (6.81)$$

Additional qubits will modify these matrices by adding a direct sum with the identity matrix. For instance, when we have 3 qubits, $\text{CNOT}_{2,3}$ is

$$\left(\begin{array}{ccc|ccc} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 \\ \hline 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{array} \right). \quad (6.82)$$

In general, if we have a symplectic matrix

$$U = \left(\begin{array}{c|c} A & B \\ \hline C & D \end{array} \right), \quad (6.83)$$

representing the Clifford group gate U , then left multiplication by another symplectic matrix representing the gate V gives us the matrix for VU whereas right multiplication gives us UV . The effects of left and right multiplication by H , $R_{\pi/4}$, CNOT, C – Z, and SWAP are summarized in table 6.1. As you can see, left multiplication gives us a set of row operations and right multiplication gives us a set of column operations. We can't exactly do standard Gaussian elimination, but we can do something very similar. In particular, given an arbitrary symplectic matrix U , we will find a sequence of H , $R_{\pi/4}$, and CNOT to multiply by on the left and right to transform U into the identity:

$$\left(\prod_{k=1}^{g_L} V_{L,k} \right) U \left(\prod_{k=1}^{g_R} V_{R,k} \right) = I, \quad (6.84)$$

with $V_{L,k}$ and $V_{R,k}$ drawn from $\{H_i, R_{\pi/4,i}, \text{CNOT}_{i,j}\}$. Then

$$U = \left(\prod_{k=g_L}^1 V_{L,k}^\dagger \right) \left(\prod_{k=g_R}^1 V_{R,k}^\dagger \right), \quad (6.85)$$

Gate	Left multiplication	Right multiplication
H_i	Switches the i th rows of $(A B)$ and $(C D)$	Switches the i th columns of $\begin{pmatrix} A \\ C \end{pmatrix}$ and $\begin{pmatrix} B \\ D \end{pmatrix}$
$R_{\pi/4,i}$	Adds the i th row of $(A B)$ to the i th row of $(C D)$	Adds the i th column of $\begin{pmatrix} B \\ D \end{pmatrix}$ to the i th column of $\begin{pmatrix} A \\ C \end{pmatrix}$
$\text{CNOT}_{i,j}$	Adds the i th row of $(A B)$ to the j th row of $(A B)$, and adds the j th row of $(C D)$ to the i th row of $(C D)$	Adds the j th column of $\begin{pmatrix} A \\ C \end{pmatrix}$ to the i th column of $\begin{pmatrix} A \\ C \end{pmatrix}$, and adds the i th column of $\begin{pmatrix} B \\ D \end{pmatrix}$ to the j th column of $\begin{pmatrix} B \\ D \end{pmatrix}$
$\text{C} - \text{Z}_{i,j}$	Adds the i th row of $(A B)$ to the j th row of $(C D)$, and adds the j th row of $(A B)$ to the i th row of $(C D)$	Adds the j th column of $\begin{pmatrix} B \\ D \end{pmatrix}$ to the i th column of $\begin{pmatrix} A \\ C \end{pmatrix}$, and adds the i th column of $\begin{pmatrix} B \\ D \end{pmatrix}$ to the j th column of $\begin{pmatrix} A \\ C \end{pmatrix}$
$\text{SWAP}_{i,j}$	Swaps the i th rows of $A, B, C,$ and D with the j th rows of $A, B, C,$ and D	Swaps the i th columns of $A, B, C,$ and D with the j th columns of $A, B, C,$ and D

Table 6.1: The effects of left and right multiplication by H , $R_{\pi/4}$, CNOT, $\text{C} - \text{Z}$, and SWAP on a symplectic matrix.

proving the theorem.

In equation (6.83), let a_i , b_i , c_i , and d_i be the i th columns of A , B , C , and D , respectively, and a_{ij} , b_{ij} , c_{ij} , and d_{ij} be the i, j th entries of A , B , C , and D . Because U is symplectic, we can apply equation (6.54) to get the following properties:

1. U has full rank.
2. $C^T A + A^T C = 0$; i.e., $(a_i|c_i) \odot (a_j|c_j) = 0$.
3. $C^T B + A^T D = I$; i.e., $(a_i|c_i) \odot (b_j|d_j) = \delta_{ij}$.
4. $D^T A + B^T C = I$, which is the same as the previous property.
5. $D^T B + B^T D = 0$; i.e., $(b_i|d_i) \odot (b_j|d_j) = 0$.

These properties all remain true if we multiply U on the left or right by other symplectic matrices.

To find a sequence of the form equation (6.84), we can use the following steps:

Procedure 6.5.

1. Since U has maximum rank, somewhere in the first column of A or C is a 1. Using left multiplication by H and SWAP, we can put this 1 in the upper left corner.
2. Do column reduction on the first column of A : Use left multiplication by CNOT to add the 1 in the upper left corner of A to any other row of A with a 1 in the first column.
3. Do row reduction on the first row of A : Use right multiplication by CNOT to add the 1 in the upper left corner of A to any other column of A with a 1 in the first row.
4. Using left multiplication by $R_{\pi/4}$ and/or $C - Z$, we can similarly make the first column of C to be all 0s.
5. At this point, we find that the first row of C has also become all 0s. This must be the case since

$$0 = (a_1|c_1) \odot (a_j|c_j) = a_1 \cdot c_j + c_1 \cdot a_j = c_{1j} + 0. \quad (6.86)$$

6. Repeat steps 1 to 5 on the second row and column of A and C to make them all 0 except for a_{22} , which is 1. Continue with all other rows and columns in sequence, until we have $A = I$ and $C = 0$.
7. At this point it follows that $D = I$:

$$\delta_{ij} = (a_i|c_i) \odot (b_j|d_j) = a_i \cdot d_j + c_i \cdot b_j = d_{ij}. \quad (6.87)$$

Also, B is symmetric:

$$0 = D^T B + B^T D = B + B^T. \quad (6.88)$$

8. Right multiply by H for every qubit, switching A and B , and switching C and D , so now $B = C = I$, $D = 0$, and A is symmetric.
9. Right multiply by $R_{\pi/4}$ as needed to eliminate the diagonal of A . Right multiply by $C - Z$ to eliminate all other elements of A , leaving $A = 0$. Since $D = 0$, these operations do not change C .
10. Right multiply by H for every qubit, switching the left and right halves back again. We are left with $A = D = I$, $B = C = 0$.

□

It is worth counting just how many gates from the generating set are needed in this procedure. To eliminate all the 1s in a single row and column of A and C takes $O(n)$ gates. Doing this for all n rows and columns thus requires $O(n^2)$ gates. Steps 8 and 10 only require $O(n)$ gates, but step 8 could require one gate for each element of A , which is again $O(n^2)$. Thus, the overall number of gates used in the procedure is $O(n^2)$. It turns out that this can be slightly improved, allowing an arbitrary element of the Clifford group to be written as a product of $O(n^2/\log n)$ generators.

6.4 Encoding Circuits for Stabilizer Codes

Any stabilizer code can be encoded with a Clifford group gate. Therefore, techniques useful for decomposing a Clifford group unitary into a detailed quantum circuit are also useful for deriving encoding circuits for stabilizer codes. The main difference between encoding a stabilizer code and finding a circuit for a Clifford group operation is that there is more freedom in choosing the encoder for a stabilizer code. When you are given a full Clifford group operation, it tells you exactly what unitary you must perform (up to global phase, perhaps), whereas for a stabilizer code, you have an additional freedom to perform unitaries which leave the code space unchanged.

The procedures discussed in this section are boring but necessary. However, it is sometimes possible to come up with clever codes with more exciting encoding circuits. The most general stabilizer code uses $O(n^2)$ gates in its encoding circuit, which is not too bad, but for a big code, we'd really like an encoding circuit with only $O(n)$ gates. Some such codes exist. It's not possible to do better than that for any sensible code, since with $o(n)$ gates, you can't even touch every qubit in the code with a gate. Some qubits will end up either unprotected or unused, maybe even unloved.

The true bottleneck, however, is the *decoding* circuit, and in particular, the syndrome decoding problem. Recall that the error syndromes correspond to cosets of $N(S)$ in P_n , and that we assign a coset representative Q_s to each error syndrome s to serve as the "correct" way to correct a state with syndrome s . However, in general, it is quite difficult (NP-hard) to compute Q_s as a function of s . Efficient codes for which syndrome decoding can be performed efficiently are precious things, and one of the main points of coding theory is to find them. Even better is an efficient code for which encoding and syndrome decoding can both be done in time $O(n)$. In general, syndrome decoding, efficient or not, will use gates outside of the Clifford group, but is most often done on a classical computer since it is basically a classical problem.

6.4.1 Encoding a Stabilizer Code with Unspecified Logical Paulis

The most straightforward case is when the stabilizer is given, but not the logical Paulis. In this case, there is an additional freedom to perform unitaries within the code space, provided we do not mix codewords with non-codewords. In other words, we can use any set of logical Paulis that is convenient for us.

If the original state, pre-encoding, is $|0\rangle^{\otimes(n-k)} \otimes |\psi\rangle$ (where $|\psi\rangle$ is a k -qubit state), its initial stabilizer is generated by Z_1, Z_2, \dots, Z_{n-k} . The final stabilizer S , post-encoding, also has $n - k$ generators M_1, M_2, \dots, M_{n-k} . The choice of generators is not unique, of course, so we may pick any set of generators for S that we like. Then we must simply find some Clifford group circuit that maps $Z_i \mapsto M_i$ for $i = 1, \dots, n - k$.

We can do this using a simplified version of procedure 6.5. There are two ways we can simplify: First, we don't really care what happens to Z_i for $i > n - k$ or to X_i . Thus, we need only concern ourselves with the first $n - k$ columns of A and C . Second, the choice of generators of the stabilizer is not unique. Therefore, we may permute columns and add columns to each other within A and C without using any actual gates. We are left with the following procedure:

Procedure 6.6. Generate a list of gates using the steps given below.

1. Write the generators M_1, \dots, M_{n-k} of S as binary symplectic vectors. Produce the first $n - k$ columns of a symplectic matrix as follows: The i th column of A is $a_i = x_{M_i}$. The i th column of C is $c_i = z_{M_i}$.
2. $(a_1|c_1)$ is a non-zero vector, so there is a 1 somewhere in the first column. Using left multiplication by H and SWAP, we can put this 1 in the upper left corner.
3. Do column reduction on the first column of A : Use left multiplication by CNOT to add the 1 in the upper left corner of A to any other row of A with a 1 in the first column.
4. Using left multiplication by $R_{\pi/4}$ and/or $C - Z$, we can similarly make the first column of C to be all 0s.

5. Do row reduction on the first row of A : Replace generators of the stabilizer to add the 1 in the upper left corner of A to any other column of A with a 1 in the first row. (It does not matter much if this step is performed before or after the previous step.)

6. At this point, we find that the first row of C has also become all 0s. This must be the case since

$$0 = (a_1|c_1) \odot (a_j|c_j) = a_1 \cdot c_j + c_1 \cdot a_j = c_{1j} + 0. \quad (6.89)$$

7. Repeat the previous steps on the second row and column of A and C to make them all 0 except for a_{22} , which is 1. Continue with all other rows and columns in sequence up to column number $n - k$.

8. Perform H on qubits 1 through $n - k$ to switch A and C .

Take the inverse of this product of gates. The resulting Clifford group circuit will encode in some coset of S ; that is, the stabilizer will be almost correct, except that we may have some error syndrome different from the correct trivial syndrome. (Also note that the choice of stabilizer generators may be different from that which was originally given to us.) Calculate the action of the encoding circuit on Z_i for $i = 1, \dots, n - k$. If for any i , the resulting generator M_i has the wrong sign, add X_i to the beginning of the encoding circuit.

As an example, let us find an encoding circuit for the five-qubit code as given in table 3.2. We begin (step 1) with the matrix

$$\begin{array}{cccc} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \\ \hline 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 1 \end{array}. \quad (6.90)$$

Step 2 is unnecessary, as there is already a 1 for a_{11} . We can perform steps 3 and 4 using $\text{CNOT}_{1,4} \cdot C - Z_{1,2} \cdot C - Z_{1,3}$. Be careful: you must perform the correct transformations on the full rows, not just the first column. Then, in step 5, we can clear out the first row of A by adding the first column to the third column; this corresponds to replacing the third generator M_3 with M_1M_3 . We now have the following matrix:

$$\begin{array}{cccc} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 0 \\ \hline 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 1 \end{array}. \quad (6.91)$$

We can clear out the second column using $\text{CNOT}_{2,5} \cdot C - Z_{2,3} \cdot C - Z_{2,4}$, then reduce the second row of A by replacing M_4 with M_2M_4 . The third column requires slightly more care. First column reduce the third column of A using $\text{CNOT}_{3,4}$. Then use $C - Z_{3,4} \cdot C - Z_{3,5}$ to eliminate the third column of C . The caution is necessary because $\text{CNOT}_{3,4}$ and $C - Z_{3,4}$ do not commute; in this case, however, the difference is only Z , which will not affect the binary symplectic representation. Then we finish with the fourth column, using

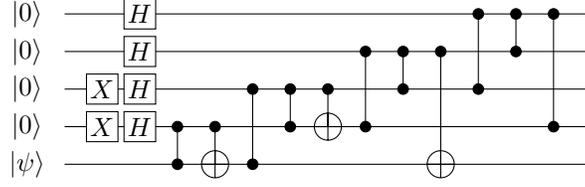


Figure 6.4: Encoding circuit for the five-qubit code derived in the text.

$\text{CNOT}_{4,5}$ followed by $C - Z_{4,5}$. We then conclude by performing H on qubits 1 through 4. These gates give us the following sequence of matrices:

$$\begin{array}{cccc}
 1 & 0 & 0 & 0 \\
 0 & 1 & 0 & 0 \\
 0 & 0 & 1 & 0 \\
 0 & 0 & 1 & 1 \\
 0 & 0 & 0 & 1 \\
 0 & 0 & 0 & 0 \\
 0 & 0 & 1 & 1 \\
 0 & 0 & 1 & 1 \\
 0 & 0 & 1 & 1
 \end{array}
 \rightarrow
 \begin{array}{cccc}
 1 & 0 & 0 & 0 \\
 0 & 1 & 0 & 0 \\
 0 & 0 & 1 & 0 \\
 0 & 0 & 0 & 1 \\
 0 & 0 & 0 & 1 \\
 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 1 \\
 0 & 0 & 0 & 1
 \end{array}
 \rightarrow
 \begin{array}{cccc}
 1 & 0 & 0 & 0 \\
 0 & 1 & 0 & 0 \\
 0 & 0 & 1 & 0 \\
 0 & 0 & 0 & 1 \\
 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0
 \end{array}
 \rightarrow
 \begin{array}{cccc}
 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 \\
 1 & 0 & 0 & 0 \\
 0 & 1 & 0 & 0 \\
 0 & 0 & 1 & 0 \\
 0 & 0 & 0 & 1
 \end{array}
 \quad (6.92)$$

Taking the inverses of these gates in reverse order, we get the following sequence of gates for the encoding circuit:

$$\begin{aligned}
 & \text{CNOT}_{1,4} \cdot C - Z_{1,2} \cdot C - Z_{1,3} \cdot \text{CNOT}_{2,5} \cdot C - Z_{2,3} \cdot C - Z_{2,4} \cdot \text{CNOT}_{3,4} \\
 & \cdot C - Z_{3,4} \cdot C - Z_{3,5} \cdot \text{CNOT}_{4,5} \cdot C - Z_{4,5} \cdot H_1 \cdot H_2 \cdot H_3 \cdot H_4
 \end{aligned}
 \quad (6.93)$$

Now calculate the effect of these gates on Z_1, \dots, Z_4 :

$$Z_1 \rightarrow X \otimes Z \otimes Z \otimes X \otimes I \quad (6.94)$$

$$Z_2 \rightarrow I \otimes X \otimes Z \otimes Z \otimes X \quad (6.95)$$

$$Z_3 \rightarrow -I \otimes Z \otimes Y \otimes Y \otimes Z \quad (6.96)$$

$$Z_4 \rightarrow -Z \otimes I \otimes Z \otimes Y \otimes Y. \quad (6.97)$$

Since

$$I \otimes Z \otimes Y \otimes Y \otimes Z = +(X \otimes Z \otimes Z \otimes X \otimes I)(X \otimes I \otimes X \otimes Z \otimes Z) \quad (6.98)$$

$$Z \otimes I \otimes Z \otimes Y \otimes Y = +(I \otimes X \otimes Z \otimes Z \otimes X)(Z \otimes X \otimes I \otimes X \otimes Z), \quad (6.99)$$

the signs of Z_3 and Z_4 need to be corrected, so we start the encoding circuit with $X_3 X_4$. The final encoding circuit is given in figure 6.4.

6.4.2 Encoding a Stabilizer Code with Specified Logical Paulis

If we are given a specific choice of logical Paulis, encoding becomes slightly more complicated. However, the same sort of techniques will work. There are two straightforward options.

Option 1 is to just use procedure 6.5, with fewer simplifications. The desired Clifford group operation maps $Z_i \rightarrow M_i$ for $i = 1, \dots, n-k$, $Z_{n-k+j} \rightarrow \bar{Z}_j$ for $j = 1, \dots, k$, and $X_{n-k+j} \rightarrow \bar{X}_j$ for $j = 1, \dots, k$. Again, it is more convenient to perform the Hadamard on every qubit so that A and C of the desired symplectic

matrix are full $n \times n$ matrices. Column operations which add the first $n - k$ columns of A and C to something can be done just by changing generators of the stabilizer or representatives of the logical Pauli cosets, but column operations involving the last k columns in either the right or left must be done with real gates.

Option 2 is to take advantage of the simplified procedure procedure 6.6 to find an encoding circuit for a stabilizer without specified logical Paulis and see what actual logical Paulis \overline{P}' it produces. Determine the logical Clifford group operation that maps \overline{P}' to the desired logical Pauli \overline{P} . Use procedure 6.5 on k qubits to find a circuit performing that Clifford group operation. Perform this circuit on qubits $n - k + 1, \dots, n$, followed by the encoding circuit given by procedure 6.6. The resulting circuit will then encode the code with the correct logical Paulis.

For instance, in the previous subsection, we found an encoding circuit for the five-qubit code. Let us determine what logical Paulis it gives:

$$X_5 \rightarrow Z \otimes I \otimes I \otimes Z \otimes X \quad (6.100)$$

$$Z_5 \rightarrow Z \otimes Z \otimes Z \otimes Z \otimes Z. \quad (6.101)$$

In this encoding, \overline{Z} is correct. However,

$$\overline{X} = Z \otimes I \otimes I \otimes Z \otimes X = -(X \otimes I \otimes X \otimes Z \otimes Z)(Z \otimes X \otimes I \otimes X \otimes Z)(X \otimes X \otimes X \otimes X \otimes X). \quad (6.102)$$

Therefore, we need the Clifford $X \rightarrow -X$, $Z \rightarrow Z$. This is just the gate Z , so we can correct the circuit of figure 6.4 by beginning with a Z_5 gate.

6.5 Extending the Clifford Group to a Universal Gate Set

Since the Clifford group can be simulated classically, to do any really interesting quantum algorithms, we will need some gates outside the Clifford group. How many gates and which ones suffice? It turns out *any* single gate will do it.

Theorem 6.9. *Let $G \subseteq \text{SU}(2^n)$ be a group of quantum gates such that $\hat{\mathcal{C}}_n \subset G$ but $G \neq \hat{\mathcal{C}}_n$. Then G is dense in $\text{SU}(2^n)$.*

That is, for any $U \in \text{SU}(2^n)$ and any $\epsilon > 0$, $\exists V \in G$ such that $\|U - V\| < \epsilon$. Thus, we can approximate all quantum gates to any desired degree of accuracy using G , even if G is generated by just the Clifford group, plus any additional gate.

People frequently pick a somewhat nice gate to use as the extra gate. Two popular choices are the Toffoli gate Tof (also known as the controlled-controlled-NOT gate) $|a\rangle|b\rangle|c\rangle \mapsto |a\rangle|b\rangle|c \oplus ab\rangle$ and the $\pi/8$ phase rotation $R_{\pi/8}$. These gates are useful extra gates because they have comparatively straightforward fault-tolerant implementations using common QECCs, as you'll see in chapter 13.

The proof of theorem 6.9 is complicated and involves some more advanced concepts. I hope to include an accessible version of it in a later draft of this book.

Chapter 7

Tighter, Please: Upper And Lower Bounds On Quantum Codes

As I've told you, the three most important properties of a quantum code are the number n of logical qubits (or qudits), the dimension K of the encoded subspace, and the distance d of the code. Together, the parameters $((n, K, d))$ tell us about the trade-off between the rate at which we can send quantum information and the tolerance we gain against errors. Ideally, we'd like to know for any given set $((n, K, d))$ whether a QECC exists with those parameters.

Unfortunately, we can't answer that question. It seems to be extremely hard. However, we can set some bounds, which set limits on where we can hope to find interesting QECCs. On the one hand, there are lower bounds, saying that codes definitely exist with certain parameters. They can be *constructive* (specifying a particular code) or *non-constructive* (proving that codes with the parameters exist without giving an efficient method of specifying one). Then we have the upper bounds. There are many different methods used for proving upper bounds, but they are uniformly destructive: they tell us there are no codes with certain parameters.

The most attractive bounds are the tightest, with the upper bounds hugging as closely as possible to the lower bounds. Occasionally we can actually make the upper and lower bounds touch, but more often there is some space between. Unfortunately, in many instances, our bounds are rather baggy and shapeless, leaving a lot of room, which may or may not contain actual QECCs, between the upper and lower bounds.

7.1 The Quantum Gilbert-Varshamov Bound

We'll start with a lower bound, the quantum Gilbert-Varshamov bound. As you might guess from the name, it is a quantum analogue of the classical Gilbert-Varshamov lower bound discussed in section 4.5. The main difference from the classical Gilbert-Varshamov bound is that there are more quantum errors that occur on a qubit than there are classical errors that can occur on a bit.

Theorem 7.1 (Quantum Gilbert-Varshamov bound). *Suppose*

$$2^k \left(\sum_{j=0}^{d-1} 3^j \binom{n}{j} \right) \leq 2^n. \quad (7.1)$$

Then there exists a $((n, 2^k, d))$ QECC. For large n , a code exists if

$$k/n \leq 1 - (d/n) \log 3 - h(d/n), \quad (7.2)$$

with $h(x) = h_2(x)$ given by equation (4.22).

Actually, the theorem will show that a stabilizer code with these parameters $[[n, k, d]]$ exists. It even can be made a little bit tighter than the statement of the theorem indicates, but it is usually phrased this way to be the closest analogue of the classical Gilbert-Varshamov bound. The proof is non-constructive. In this case, as with the classical Gilbert-Varshamov bound, that means that, while the proof specifies a procedure to find a code with these parameters, the procedure is exponentially long as a function of n , so it's not useful in practice. Since we get a stabilizer code, the code we produce can be described efficiently, but that is small consolation if we can't get all the way through the procedure to find it.

One can also prove a version of the quantum Gilbert-Varshamov bound for qudits, at least for prime power dimensions. The proof is analogous, using the idea of a qudit stabilizer code, which I'll discuss in chapter 8.

Proof. Imagine making a long list of all $[[n, k]]$ stabilizer codes. We are going to run through all possible errors of weight up to $d - 1$. For each error E , we can check which codes can detect that error and which cannot, and cross off the list any code that can't detect E . Once we finish running through all the errors, we know that any code that remains must have distance at least d . We'll prove that there must be at least one code to survive, giving us the desired $[[n, k, d]]$ code.

For any given error E , we need to count how many codes cannot detect it. For any two errors $E, F \in \mathbb{P}_n \setminus \{I\}$, there exists some Clifford group element U with $U(E) = F$. Furthermore, if stabilizer code S is a code that doesn't detect E , then $U(S)$ is a stabilizer code that doesn't detect F . That is, U will permute the list of stabilizer codes and their sets of undetectable errors. Thus, the number of codes that cannot detect E must be the same as the number of codes that cannot detect F . That is, all non-identity errors in the Pauli group are undetectable for the same number C of $[[n, k]]$ stabilizer codes.

Suppose there are N stabilizer codes with parameters $[[n, k]]$. (We can evaluate N exactly if we like, but it is not necessary for this argument.) The code S detects the errors outside $\hat{N}(S) \setminus \hat{S}$, which contains $2^{n+k} - 2^{n-k}$ elements. If we consider a bigger list of pairs (S, E) for any pair for which $E \in \hat{N}(S) \setminus \hat{S}$, we can count the number of elements on the list two ways: as the number of errors times the number of codes that cannot detect each error, or as the number of codes times the number of errors that each code cannot detect. There are $4^n - 1$ non-identity errors, and I never appears in $\hat{N}(S) \setminus \hat{S}$, so we have that

$$(4^n - 1)C = N(2^{n+k} - 2^{n-k}). \quad (7.3)$$

Now let us go back to crossing codes off our list of stabilizer codes. We work our way through the set of errors of weight up to $d - 1$. For any non-identity error E with weight less than d , there are C codes which do not detect it, so we can cross those codes off the list. Some of them may have already been eliminated because they fail to detect a previous error, but certainly we cannot eliminate more than C new codes from the list of $[[n, k]]$ QECCs. There are a total of $B - 1$ non-identity errors of weight less than d , with

$$B = \sum_{j=0}^{d-1} 3^j \binom{n}{j} \quad (7.4)$$

(not including the identity since all codes detect it). By the time we finish going through all errors of weight less than d , we have therefore crossed off at most $(B - 1)C$ codes. If $(B - 1)C < N$, then at least one code remains.

Plugging in equation (7.3) to eliminate C , the condition we get is

$$\frac{2^{n+k} - 2^{n-k}}{4^n - 1} (B - 1)N < N, \quad (7.5)$$

or

$$2^{n-k}(B - 1) < \frac{4^n - 1}{4^k - 1}. \quad (7.6)$$

This is the tightest version of the quantum Gilbert-Varshamov bound, but notice that

$$\frac{4^n - 1}{4^k - 1} > \frac{4^n}{4^k} = 4^{n-k}, \quad (7.7)$$

so if

$$2^{n-k}B \leq 4^{n-k}, \quad (7.8)$$

then equation (7.6) is satisfied too. Equation (7.8) is just what we wanted to prove. \square

7.2 The Quantum Hamming Bound

We'll now move on to an upper bound of sorts. It is an analogue of the classical Hamming bound, so is called the quantum Hamming bound.

Theorem 7.2 (Quantum Hamming bound). *If a non-degenerate $((n, K, 2t + 1))_q$ code exists, then*

$$K \left(\sum_{j=0}^t (q^2 - 1)^j \binom{n}{j} \right) \leq q^n. \quad (7.9)$$

The big catch in the statement of the quantum Hamming bound is that while it provides an upper bound, the upper bound *only holds for non-degenerate codes*. Here is where we start to see very concretely how the existence of degenerate codes makes the quantum case more complicated than the classical case. Insisting that a code be degenerate is potentially a big limitation, yet it is easier to prove bounds on the existence of non-degenerate codes. Indeed, while we don't know how to prove that the quantum Hamming bound applies to degenerate codes as well as non-degenerate codes, we also don't know any $((n, K, 2t + 1))$ codes that violate it.

Proof. The proof is a straightforward exercise in counting dimensions, and is very analogous to the proof of the classical Hamming bound. Because the code is non-degenerate, we know that in the QECC conditions, the matrix

$$C_{ab} = \langle \psi | E_a^\dagger E_b | \psi \rangle \quad (7.10)$$

(for any codeword $|\psi\rangle$) has maximum rank when E_a^\dagger and E_b run over any basis for the space of $\leq t$ -qubit errors. In particular, this means that the states $E_a|\psi\rangle$ are all linearly independent. Furthermore, if $|\psi\rangle$ and $|\phi\rangle$ are two orthogonal codewords, then the QECC conditions tell us that

$$\langle \psi | E_a^\dagger E_b | \phi \rangle = 0, \quad (7.11)$$

so $E_a|\psi\rangle$ and $E_b|\phi\rangle$ are orthogonal for any a, b .

Thus, if we let $\{|\psi_i\rangle\}$ be a basis for the code space, the set of states $\{E_a|\psi_i\rangle\}$ (over all a, i) are linearly independent. The codespace has dimension K and there are $q^2 - 1$ one-qudit non-identity errors, so there are a total of

$$B = \sum_{j=0}^t (q^2 - 1)^j \binom{n}{j} \quad (7.12)$$

linearly independent errors of weight up to t . Thus, we have a set of KB linearly independent vectors, which we must fit into a Hilbert space of n qudits. Therefore,

$$KB \leq q^n \quad (7.13)$$

\square

Looking at this proof, you'll see that the quantum Hamming bound does not make essential use of the qudit structure of the space, or the fact that we are dealing with t -qudit errors. In general, if we have a non-degenerate QECC with a K -dimensional code space living in an N -dimensional physical Hilbert space, and the code can correct \mathcal{E} , a set of linearly independent errors, then $K|\mathcal{E}| \leq N$.

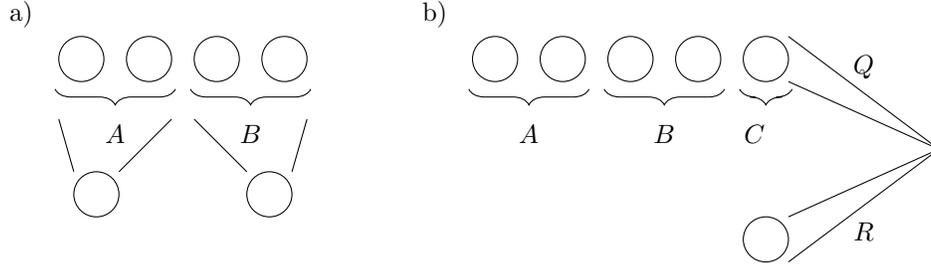


Figure 7.1: a) Dividing n registers into one set A of size $d - 1$ and one set B of size $n - (d - 1)$ for the proof of lemma 7.4. If $n - (d - 1) \leq d - 1$, we can decode each set separately to clone the encoded qubit. b) System Q is maximally entangled with system R . Q is then encoded into n registers, which are split into three sets for the proof of the quantum Singleton bound: sets A and B of size $d - 1$ and set C of size $n - 2(d - 1)$.

7.3 The Quantum Singleton Bound

The quantum Singleton bound has a similar form to the classical Singleton bound, but is significantly harder to prove. The quantum Singleton bound is also known as the Knill-Laflamme bound, after the first people to state it. The quantum Singleton bound is a real upper bound, applying to non-degenerate codes as well as degenerate ones, and to both stabilizer and non-stabilizer codes over qubits or qudits.

Theorem 7.3 (Quantum Singleton bound). *If an $((n, q^k, d))_q$ QECC exists, then*

$$n - k \geq 2(d - 1). \tag{7.14}$$

As you can see by comparing with theorem 4.19, the quantum bound is more restrictive on the distance by a factor of 2. The reason for this factor of 2 turns out to be the No-Cloning Theorem! You may recall that our very first objection in section 2.1 to the possibility of a quantum error-correcting code was that classical codes seemed to use repetition of information as a key component. We circumvented that difficulty by using entanglement to spread out quantum information without repeating it. The quantum Singleton bound can be viewed as a way of codifying just how well information can be spread around without accidentally repeating anything. For instance, the bound immediately tells us that there can be no complete quantum analogue of the classical $(3, 2, 3)$ repetition code, since a $((3, 2, 3))$ QECC would violate the quantum Singleton bound.

Proof. We'll start by studying the $k = 1$ case, which is an immediate consequence of the No-Cloning Theorem:

Lemma 7.4. *If an $((n, q, d))_q$ QECC exists, then*

$$n - 1 \geq 2(d - 1). \tag{7.15}$$

Proof of lemma. If a code has distance d , then it can correct $d - 1$ erasure errors. We can imagine dividing the n registers of the code into two sets, a set A with $d - 1$ registers and a set B with $n - (d - 1)$ registers, as pictured in figure 7.1a. Set B has experienced $d - 1$ erasures, so using just those registers, it is possible to reconstruct the original encoded state $|\psi\rangle$.

Now suppose that $n - 1 < 2(d - 1)$. Then $n - (d - 1) \leq d - 1$, and therefore set A has also experienced at most $d - 1$ erasure errors. We would then be able to reconstruct a second copy of $|\psi\rangle$ using just the registers in set A . We could then use this code to clone the state $|\psi\rangle$ via the following procedure: Encode $|\psi\rangle$ using the QECC, split up the registers into the sets A and B , and then reconstruct $|\psi\rangle$ independently for each set. We know this isn't possible, leading to a contradiction and proving the lemma. □

Now we can move on to the general case. When $k > 1$, we split the n registers into three sets, as pictured in figure 7.1b. Sets A and B will each have $d - 1$ registers, and set C will have $n - 2(d - 1)$ registers. If an

$((n, q^k, d))_q$ code exists for $k \geq 1$, then a $((n, q, d))_q$ code also exists, as we can just ignore the extra logical codewords. We therefore know from the lemma that $n > 2(d - 1)$, and set C is non-empty.

Now imagine taking a $2k$ -qudit maximally entangled state between two registers R and Q, each with dimension q^k , and encode Q in the QECC. Then divide up the registers of the QECC into the sets A, B, and C, giving us a total of 4 sets (R, A, B, C). Globally, we now have a pure state, so $S(RABC) = 0$. If we split our sets into two groups, the entropy of the two groups is equal, e.g.

$$S(RA) = S(BC) \tag{7.16}$$

$$S(RB) = S(AC). \tag{7.17}$$

We also know that the code has distance d , so we can detect any erasure error restricted to just set A or to just set B (but not necessarily errors involving both sets). Applying the alternate QECC conditions of section 2.5.6, we see that this means that any operator we measure on set A will give us the same result for all possible logical codewords. Thus, the density matrix ρ_A of set A is the same for all logical codewords, and we find that the density matrix for R and A combined is of the form

$$\sum_{i=0}^{q^k-1} |i\rangle \langle i|_R \otimes \rho_A. \tag{7.18}$$

R and A are in a tensor product state, so $S(RA) = S(R) + S(A)$. Similarly, $S(RB) = S(R) + S(B)$. We also note that $S(R) = k \log q$.

Now, $S(RA) = S(BC)$, and by subadditivity of the entropy, $S(BC) \leq S(B) + S(C)$. Therefore,

$$k \log q + S(A) = S(RA) = S(BC) \leq S(B) + S(C), \tag{7.19}$$

or

$$k \log q \leq S(C) + [S(B) - S(A)]. \tag{7.20}$$

Applying subadditivity to $S(AC)$, we also find that

$$k \log q \leq S(C) + [S(A) - S(B)]. \tag{7.21}$$

Adding together these last two equations, we find that

$$k \log q \leq S(C). \tag{7.22}$$

However, $S(C) \leq \log \dim(C) = [n - 2(d - 1)] \log q$. Therefore, we find

$$k \leq n - 2(d - 1). \tag{7.23}$$

□

As an immediate application of the quantum Singleton bound, we find that the five-qubit code is optimal. Not only can there be no $((3, 2, 3))$ code, but there can also be no $((4, 2, 3))$ QECC. The five-qubit code exactly saturates the quantum Singleton bound; it is a *quantum MDS code*. We can also saturate the quantum Singleton bound with $d = 2$. There are $[[n, n - 2, 2]]$ stabilizer codes for any even n (but not for odd n). As we go to more encoded qubits or to greater distances, it is no longer possible to exactly saturate the quantum Singleton bound with qubit codes. A similar phenomenon occurs classically, and as in the classical case, we get more MDS codes as we let the dimension of the registers increase. In section 8.3, we'll see a large family of such codes.

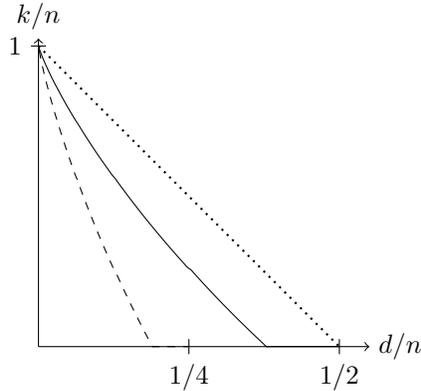


Figure 7.2: Quantum Hamming bound (solid), quantum Gilbert-Varshamov bound (dashed), and quantum Singleton bound (dotted) for large n , $q = 2$

7.4 Linear Programming Bounds

The quantum Singleton bound does provide us with some definite limits on the existence of QECCs since, unlike the quantum Hamming bound, it applies to degenerate codes too. Still, it doesn't tell us much. Figure 7.2 shows the lower bound provided by the quantum Gilbert-Varshamov bound along with the upper bound given by the quantum Singleton bound, both plotted for codes using qubits. I've also shown the quantum Hamming bound so you can see that even if we were able to prove it applied to degenerate codes, there would still be a lot of room where we don't know about the existence of codes. The quantum Singleton bound is much worse. There's no way we can be satisfied with that.

In order to do better, we'll need more powerful techniques. The main ones are known as *linear programming bounds*. A linear programming bound provides a set of linear inequalities that a QECC with a given set of parameters $((n, K, d))_q$ must satisfy. If we can prove that the set of inequalities has no solutions for a given n , K , and d , then we know that a QECC with those parameters cannot exist. If there *is* a solution, then a QECC may or may not exist, but the linear programming solution gives us a number of constraints about the structure of any possible code with those parameters, so provides a good hint as to what to look at to find a code.

In this section, I'll discuss the main linear programming bounds, but only for codes over qubits. These bounds can be generalized to apply to qudits, and it's possible to come up with more linear programming bounds, which could narrow down the set of possible codes even further.

7.4.1 Weight Enumerator and Dual Weight Enumerator

The primary tools used in coming up with linear programming bounds are known as *weight enumerators*. There are many different kinds of weight enumerators, but we'll start with the simplest, which gets no additional adjectives in its name. Weight enumerators are polynomials that encode information about the properties of a QECC.

To motivate the definition of weight enumerator, let's first consider the case of a stabilizer code. The stabilizer S consists of Pauli operators. Some of those Pauli operators are large, with high weight, and some are small, with low weight. Indeed, S always contains I , which has weight 0. We'll count up the number of elements of each weight; let A_j be the number of Paulis $M \in S$ with weight j . The A_j s are integers which tell us about the structure of the stabilizer, although they don't tell us everything. How can we get similar information about more general QECCs?

Definition 7.1. Let Π be the projector onto the code subspace of some $((n, K))$ QECC Q . Let

$$A_j = \frac{1}{K^2} \sum_{P \in \hat{\mathcal{P}}_n | \text{wt}(P)=j} |\text{tr}(P\Pi)|^2. \quad (7.24)$$

The *weight enumerator* of Q is the degree n polynomial

$$A(x) = \sum_{j=0}^n A_j x^j. \quad (7.25)$$

The motivation for defining the weight enumerator as a polynomial instead of just working with the individual coefficients will become clearer in the next subsection, where I'll discuss the quantum MacWilliams identity, which deals with the polynomial as a whole.

Proposition 7.5. *When Q is a stabilizer code with stabilizer \mathcal{S} ,*

$$A_j = |\{M \in \mathcal{S} | \text{wt } M = j\}|. \quad (7.26)$$

Proof. For a stabilizer code, $K = 2^k$, and we have an explicit description of the projector onto the code:

$$\Pi = \frac{1}{2^{n-k}} \sum_{M \in \mathcal{S}} M. \quad (7.27)$$

Now, $\text{tr } Q = 0$ if $Q \in \hat{\mathcal{P}}_n \setminus \{I\}$ and $\text{tr } I = 2^n$. Thus,

$$|\text{tr}(PM)| = \begin{cases} 2^n & \text{if } P = M \\ 0 & \text{if } P \neq M. \end{cases} \quad (7.28)$$

We will consider Paulis in $\hat{\mathcal{P}}_n$, so depending on what representatives we pick, it could actually be that $\text{tr}(PM) = -2^n$ or $\pm i2^n$. However, since we immediately take the absolute value, all of these give the same result.

Applying the general definition of A_j , we find

$$A_j = \frac{1}{2^{2k}} \sum_{P \in \hat{\mathcal{P}}_n | \text{wt}(P)=j} \left| \frac{1}{2^{n-k}} \sum_{M \in \mathcal{S}} \text{tr}(PM) \right|^2 \quad (7.29)$$

$$= \frac{1}{2^{2k}} \sum_{P \in \hat{\mathcal{P}}_n | \text{wt}(P)=j} \left| \frac{1}{2^{n-k}} \sum_{M \in \mathcal{S}} 2^n \delta_{P,M} \right|^2 \quad (7.30)$$

$$= \frac{1}{2^{2k}} \sum_{P \in \hat{\mathcal{P}}_n | \text{wt}(P)=j} 2^{2k} \chi_{\mathcal{S}}(P) \quad (7.31)$$

$$= |\{M \in \mathcal{S} | \text{wt } M = j\}|. \quad (7.32)$$

In the third line, $\chi_{\mathcal{S}}(P)$ is the indicator function, which is 1 when $P \in \mathcal{S}$ and 0 when $P \notin \mathcal{S}$. □

Similarly, we can define a dual weight enumerator. For a stabilizer code, it is built out of B_j s which give the weight distribution of the *normalizer* of \mathcal{S} . In general, we have the definition

Definition 7.2. Let Π be the projector onto the code subspace of some $((n, K))$ QECC Q . Let

$$B_j = \frac{1}{K} \sum_{P \in \hat{\mathcal{P}}_n | \text{wt}(P)=j} \text{tr}(P\Pi P^\dagger \Pi). \quad (7.33)$$

The *dual weight enumerator* of Q is the degree n polynomial

$$B(x) = \sum_{j=0}^n B_j x^j. \quad (7.34)$$

Again, the general definition reduces to the description given for stabilizer codes:

Proposition 7.6. *When Q is a stabilizer code with stabilizer S ,*

$$B_j = |\{N \in \mathbf{N}(S) \mid \text{wt } N = j\}|. \quad (7.35)$$

Proof. Again, we have

$$\Pi = \frac{1}{2^{n-k}} \sum_{M \in S} M, \quad (7.36)$$

and we see that

$$P\Pi P^\dagger = \frac{1}{2^{n-k}} \sum_{M \in S} (-1)^{c(P,M)} M. \quad (7.37)$$

Note that when $P \in \mathbf{N}(S)$, then $P\Pi P^\dagger = \Pi$. However, when $P \notin \mathbf{N}(S)$, then $P\Pi P^\dagger$ will be the projector on the subspace given by a different error syndrome of S . That subspace is orthogonal to the code space, so we have that $P\Pi P^\dagger \Pi = 0$ when $P \notin \mathbf{N}(S)$. This is also straightforward to prove through direct computation:

$$P\Pi P^\dagger \Pi = \frac{1}{2^{2(n-k)}} \sum_{M, N \in S} (-1)^{c(P,M)} M N \quad (7.38)$$

$$= \frac{1}{2^{2(n-k)}} \sum_{N' \in S} \left(\sum_{M \in S} (-1)^{c(P,M)} \right) N' \quad (7.39)$$

$$= \frac{1}{2^{2(n-k)}} \sum_{N' \in S} 2^{n-k} \chi_{\mathbf{N}(S)}(P) N' \quad (7.40)$$

$$= \chi_{\mathbf{N}(S)}(P) \Pi, \quad (7.41)$$

where $\chi_{\mathbf{N}(S)}(P)$ is the indicator function for $\mathbf{N}(S)$ and line three follows because if $P \notin \mathbf{N}(S)$, then P anticommutes with exactly half of S .

Thus, applying the general definition of B_j , we find

$$B_j = \frac{1}{2^k} \sum_{P \in \hat{\mathcal{P}}_n \mid \text{wt}(P)=j} \text{tr}(\Pi) \chi_{\mathbf{N}(S)}(P) \quad (7.42)$$

$$= |\{N \in \mathbf{N}(S) \mid \text{wt } N = j\}|. \quad (7.43)$$

□

The coefficients A_j and B_j of the weight enumerator and dual weight enumerator are always integers for a stabilizer code, but can be non-integer for general QECCs.

For a stabilizer code, the dual weight enumerator tells us about the structure of the normalizer, which in turn tells us something about which errors we can detect. If B_j is non-zero, it means there are some Paulis of weight j in the normalizer, which worries us if j is small. Actually, though, we are interested in the number of elements of $\hat{\mathbf{N}}(S) \setminus \hat{S}$, which is given by $B_j - A_j$. If *that* is non-zero, then we actually have some errors of weight j that are undetectable, giving us a bound on the distance of the code. While this argument does not hold for general QECCs, the intuition and result about A_j and B_j does carry over:

Theorem 7.7. *Let Q be a QECC with weight enumerator $A(x) = \sum A_j x^j$ and dual weight enumerator $B(x) = \sum B_j x^j$. Then*

a) $A_0 = B_0 = 1$

b) $B_j \geq A_j \geq 0$

c) Q has distance d iff $A_j = B_j$ for $j < d$.

Proof.

a) There is only one Pauli of weight 0: the identity. Thus,

$$A_0 = \frac{1}{K^2} |\text{tr}(\Pi)|^2 = 1 \quad (7.44)$$

$$B_0 = \frac{1}{K} \text{tr}(\Pi^2) = \frac{1}{K} \text{tr}(\Pi) = 1. \quad (7.45)$$

b) From the definition of A_j , we can immediately see that $A_j \geq 0$.

Let $\{|a\rangle\}$ be a basis for the code space of Q . Then

$$\Pi = \sum_a |a\rangle \langle a|, \quad (7.46)$$

so

$$A_j = \frac{1}{K^2} \sum_{P | \text{wt}(P)=j} \left| \sum_a \langle a|P|a\rangle \right|^2, \quad (7.47)$$

$$B_j = \frac{1}{K} \sum_{P | \text{wt}(P)=j} \sum_{a,b} |\langle a|P|b\rangle|^2. \quad (7.48)$$

Our next step is to apply the Cauchy-Schwarz inequality, which says that for two complex vectors \vec{x} and \vec{y} ,

$$|\vec{x} \cdot \vec{y}|^2 \leq |\vec{x}|^2 |\vec{y}|^2. \quad (7.49)$$

Let \vec{x} be the K^2 -dimensional complex vector with entries $\langle a|P|b\rangle$ (running over pairs (a, b)), and let \vec{y} be the K^2 -dimensional vector with entries equal to $(1/K)\delta_{ab}$ (again running over pairs (a, b)). Then we have

$$|\vec{x} \cdot \vec{y}|^2 = \left| \frac{1}{K} \sum_{a,b} \langle a|P|b\rangle \delta_{ab} \right|^2 \quad (7.50)$$

$$= \frac{1}{K^2} \left| \sum_a \langle a|P|a\rangle \right|^2 \quad (7.51)$$

$$\leq |\vec{x}|^2 |\vec{y}|^2 \quad (7.52)$$

$$= \left(\sum_{a,b} |\langle a|P|b\rangle|^2 \right) \left(\sum_{ab} \frac{1}{K^2} \delta_{ab} \right) \quad (7.53)$$

$$= \frac{1}{K} \sum_{a,b} |\langle a|P|b\rangle|^2. \quad (7.54)$$

Plugging into equations (7.47) and (7.48), we find that $A_j \leq B_j$.

c) From the definition of distance, equation (2.66), if the code has distance d , then for $\text{wt}(E) < d$,

$$\langle \psi|E|\phi\rangle = c(E)\langle \psi|\phi\rangle \quad (7.55)$$

for any codewords $|\psi\rangle$ and $|\phi\rangle$. Applying equations (7.47) and (7.48) for $j < d$, we get

$$A_j = \frac{1}{K^2} \sum_{E | \text{wt } E=j} K^2 |c(E)|^2 \quad (7.56)$$

$$B_j = \frac{1}{K} \sum_{E | \text{wt } E=j} K |c(E)|^2, \quad (7.57)$$

and $A_j = B_j$.

For the converse, looking at the proof of part b, the equality condition for the Cauchy-Schwarz inequality is that \vec{x} is proportional to \vec{y} . Thus, $A_j = B_j$ implies that

$$\langle a|E|b\rangle = C(E)\delta_{ab} \quad (7.58)$$

whenever $\text{wt } E = j$. This is one of the alternate forms of the QECC conditions. □

The classical theory of weight enumerators is similar, except that for classical codes, $A_j = B_j = 0$ for $j < d$. The difference is essentially a manifestation of the phenomenon of degeneracy. Let's think about the special case of stabilizer codes. In a non-degenerate stabilizer code \mathbf{S} , there are no elements of $\mathbf{N}(\mathbf{S})$ with weight less than d . Therefore $B_j = 0$ for $j < d$. In a degenerate stabilizer code, $\mathbf{N}(\mathbf{S})$ can have elements with weight less than d , but those elements must also be in \mathbf{S} . Thus, $B_j = A_j > 0$ for some $j < d$.

This intuition almost carries over to general QECCs, but it turns out to be a slightly different concept.

Definition 7.3. An $((n, K, d))$ QECC with weight enumerators $A(x)$ and $B(x)$ is *pure* if $A_j = B_j = 0$ for $j < d$. A code that is not pure is *impure*.

The argument above shows that for stabilizer codes, pure is the same as non-degenerate. However, for more general codes, the property of being pure is more stringent than the property of being non-degenerate. In other words, a code can be impure without being degenerate, but if it's degenerate, it is definitely impure.

The possibility of impure codes complicates the application of the linear programming bounds to quantum codes. Some results which are known for classical codes have not yet been adapted for quantum codes because of the additional difficulty created by dealing with impure codes.

7.4.2 Quantum MacWilliams Identity

Theorem 7.7 gives us one set of inequalities constraining the coefficients A_j and B_j . However, these inequalities by themselves are not yet enough to rule out any codes, since for any set of parameters $((n, K, d))$, we can easily come up with a set of A_j and B_j that satisfy theorem 7.7. In order to get actual upper bounds on codes this way, we'll need a stronger relation between $A(x)$ and $B(x)$.

In the case of a stabilizer code, the weight enumerator $A(x)$ encapsulates properties of the stabilizer \mathbf{S} and the dual weight enumerator $B(x)$ captures properties of $\mathbf{N}(\mathbf{S})$. Since $\mathbf{N}(\mathbf{S})$ is completely determined by \mathbf{S} , it's perhaps not too surprising that there is some relationship between $A(x)$ and $B(x)$ as well. What is remarkable is that even though the full description of \mathbf{S} contains a lot more information than is contained in $A(x)$, the weight enumerator by itself is enough to completely determine $B(x)$. Moreover, this relationship also holds for non-stabilizer codes.

Theorem 7.8 (Quantum MacWilliams identity). *Suppose we have an $((n, K))$ QECC with weight enumerator $A(x)$ and dual weight enumerator $B(x)$. Then*

$$B(x) = \frac{K}{2^n} (1 + 3x)^n A\left(\frac{1-x}{1+3x}\right) \quad (7.59)$$

The quantum MacWilliams identity above is phrased in terms of codes over qubits, but like the other bounds we've discussed, it can also be generalized to codes over qudits.

Proof. Let's eliminate the projector Π that appears in the definitions of A_j and B_j in favor of Paulis. We can do this by using the fact that $\hat{\mathcal{P}}_n$ gives a basis for the space of $2^n \times 2^n$ matrices, so

$$\Pi = \sum_{Q \in \hat{\mathcal{P}}_n} c_Q Q. \quad (7.60)$$

The coefficients $c_Q = \text{tr}(Q^\dagger \Pi)/2^n$.

Then

$$A_j = \frac{1}{K^2} \sum_{P \in \hat{\mathcal{P}}_n | \text{wt}(P)=j} \left| \sum_{Q \in \hat{\mathcal{P}}_n} c_Q \text{tr}(PQ) \right|^2 \quad (7.61)$$

$$= \frac{4^n}{K^2} \sum_{P \in \hat{\mathcal{P}}_n | \text{wt}(P)=j} |c_P|^2, \quad (7.62)$$

and

$$B_j = \frac{1}{K} \sum_{P \in \hat{\mathcal{P}}_n | \text{wt}(P)=j} \sum_{Q, R \in \hat{\mathcal{P}}_n} c_Q c_R^* \text{tr}(PQP^\dagger R^\dagger) \quad (7.63)$$

$$= \frac{1}{K} \sum_{P \in \hat{\mathcal{P}}_n | \text{wt}(P)=j} \sum_{Q, R \in \hat{\mathcal{P}}_n} c_Q c_R^* (-1)^{c(P,Q)} \delta_{Q,R} 2^n \quad (7.64)$$

$$= \frac{2^n}{K} \sum_{Q \in \hat{\mathcal{P}}_n} \left[\sum_{P \in \hat{\mathcal{P}}_n | \text{wt}(P)=j} (-1)^{c(P,Q)} \right] |c_Q|^2. \quad (7.65)$$

In the first line for B_j , I've used Π^\dagger in one place instead of Π (they are the same, after all) to finesse issues about using Paulis with the overall sign modded out. We can see that A_j and B_j both involve a sum over $|c_Q|^2$, but with two differences: In A_j , we only sum over Paulis of weight j whereas for B_j we sum over all Paulis, and in B_j we have an additional sum with alternating signs which depend on the commutation relations.

Let's take some Q with weight i and try to count how many Paulis of weight j commute with it, which we'll call C_{ij} , and how many anticommute, A_{ij} . Then we will have

$$B_j = \frac{2^n}{K} \sum_{Q \in \hat{\mathcal{P}}_n} (C_{\text{wt}(Q),j} - A_{\text{wt}(Q),j}) |c_Q|^2. \quad (7.66)$$

First of all, we know that there are a total of

$$3^j \binom{n}{j} \quad (7.67)$$

Paulis of weight j . Also note that the answer won't depend on exactly what Q we take, only its weight i , because single-qubit Clifford group operations and permutations of the qubits can relate any two Paulis of weight i , and those operations won't change the commutation relations or weights of Paulis.

Let's break up the counting problem into counting Paulis P that overlap with Q on exactly m qubits (i.e., the intersection of their supports is m qubits), and determining how many commute (C_{ij}^m) and how many anticommute (A_{ij}^m). P has total weight j , so it acts on exactly $j - m$ qubits outside the support of Q . There are thus

$$3^{j-m} \binom{n-i}{j-m} \quad (7.68)$$

ways of choosing the action of P outside the support of Q . Within the support of Q , there are $\binom{i}{m}$ ways of choosing which qubits P acts on non-trivially.

We now wish to find how many weight m Paulis commute and anticommute with a fixed Pauli of weight m , which can assume without loss of generality to be $Z^{\otimes m}$. That is, it suffices to find C_{mm}^m and A_{mm}^m , and then we know that

$$C_{ij}^m = 3^{j-m} \binom{n-i}{j-m} \binom{i}{m} C_{mm}^m \quad (7.69)$$

$$A_{ij}^m = 3^{j-m} \binom{n-i}{j-m} \binom{i}{m} A_{mm}^m. \quad (7.70)$$

Now, $C_{11}^1 = 1$ and $A_{11}^1 = 2$. We can use induction to determine C_{mm}^m and A_{mm}^m for all m :

$$C_{mm}^m = C_{(m-1)(m-1)}^{m-1} + 2A_{(m-1)(m-1)}^{m-1} \quad (7.71)$$

$$A_{mm}^m = 2C_{(m-1)(m-1)}^{m-1} + A_{(m-1)(m-1)}^{m-1}. \quad (7.72)$$

Actually, what we're really interested in is $C_{ij}^m - A_{ij}^m$, and we see that

$$C_{mm}^m - A_{mm}^m = -C_{(m-1)(m-1)}^{m-1} + A_{(m-1)(m-1)}^{m-1} = (-1)^{m-1} (C_{11}^1 - A_{11}^1) = (-1)^m. \quad (7.73)$$

Putting all of this together, we find that

$$B_j = \frac{2^n}{K} \sum_{Q \in \hat{\mathcal{P}}_n} (C_{\text{wt}(Q),j} - A_{\text{wt}(Q),j}) |c_Q|^2 \quad (7.74)$$

$$= \frac{2^n}{K} \sum_{i=0}^n \sum_{Q \in \hat{\mathcal{P}}_n \mid \text{wt } Q=i} \left[\sum_{m=0}^{\min(i,j)} (C_{ij}^m - A_{ij}^m) \right] |c_Q|^2 \quad (7.75)$$

$$= \frac{2^n}{K} \sum_{i=0}^n \sum_{Q \in \hat{\mathcal{P}}_n \mid \text{wt } Q=i} \left[\sum_{m=0}^{\min(i,j)} (-1)^m 3^{j-m} \binom{n-i}{j-m} \binom{i}{m} \right] |c_Q|^2 \quad (7.76)$$

$$= \frac{K}{2^n} \sum_{i=0}^n \left[\sum_{m=0}^{\min(i,j)} (-1)^m 3^{j-m} \binom{n-i}{j-m} \binom{i}{m} \right] A_i. \quad (7.77)$$

$$(7.78)$$

The function

$$P_j(z; n) = \sum_{m=0}^{\min(z,j)} (-1)^m 3^{j-m} \binom{n-z}{j-m} \binom{z}{m} \quad (7.79)$$

is known as a Krawtchouk polynomial, and

$$(1 + 3x)^{n-z} (1 - x)^z = \sum_j P_j(z; n) x^j, \quad (7.80)$$

as can be verified by expanding the left-hand side. That tells us

$$B(x) = \sum_{j=0}^n B_j x^j \quad (7.81)$$

$$= \frac{K}{2^n} \sum_{i=0}^n \sum_{j=0}^n P_j(i; n) A_i x^j \quad (7.82)$$

$$= \frac{K}{2^n} \sum_{i=0}^n (1+3x)^{n-i} (1-x)^i A_i \quad (7.83)$$

$$= \frac{K}{2^n} (1+3x)^n \sum_{i=0}^n \left(\frac{1-x}{1+3x} \right)^i A_i \quad (7.84)$$

$$= \frac{K}{2^n} (1+3x)^n A \left(\frac{1-x}{1+3x} \right) \quad (7.85)$$

□

The quantum MacWilliams identity combined with theorem 7.7 puts a number of restrictions on the possible values of $((n, K, d))$ for a code, enough to rule out many more possibilities than the quantum Singleton bound. However, we can set even tighter bounds by adding another weight enumerator known as the quantum shadow enumerator.

7.4.3 Quantum Shadow Enumerator

The quantum shadow enumerator has a somewhat stranger definition than the weight enumerator and dual weight enumerator. The weight enumerator and dual weight enumerator treat all Paulis of a given weight on an equal footing. The quantum shadow enumerator does not:

Definition 7.4. Suppose Π is the projector onto the code subspace of some $((n, K))$ QECC Q . Let

$$Sh_j = \frac{1}{K} \sum_{P \in \hat{P}_n | \text{wt } P=j} \text{tr}(P \Pi P^\dagger Y^{\otimes n} \Pi^* Y^{\otimes n}). \quad (7.86)$$

The *shadow enumerator* of Q is

$$Sh(x) = \sum_{j=0}^n Sh_j x^j. \quad (7.87)$$

Π^* is the complex conjugate of the projector Π . The definition of the shadow enumerator is heavily basis-dependent, both for defining the Paulis and for defining the complex conjugate of an operator.

Theorem 7.9. Suppose we have an $((n, K))$ QECC with weight enumerator $A(x)$ and shadow enumerator $Sh(x) = \sum Sh_j x^j$. Then

- a) $Sh_j \geq 0 \forall j$
- b) $Sh(x)$ can be determined from $A(x)$:

$$Sh(x) = \frac{K}{2^n} (1+3x)^n A \left(\frac{x-1}{1+3x} \right). \quad (7.88)$$

Notice that the relation between $A(x)$ and $Sh(x)$ differs from the quantum MacWilliams identity only in that there is an $x-1$ in the numerator of the argument of A instead of $1-x$.

Proof.

a) We can write $\Pi = \sum_a |a\rangle \langle a|$, for $|a\rangle$ a basis of the code space. Note that this basis might be different from the basis in which the shadow enumerator is defined, so we let $|a^*\rangle$ be the complex conjugate of $|a\rangle$ in the defining basis. Then

$$Sh_j = \frac{1}{K} \sum_{P \in \hat{\mathcal{P}}_n | \text{wt } P=j} \sum_{a,b} \text{tr}(P|a\rangle \langle a| P^\dagger Y^{\otimes n} |b^*\rangle \langle b^*| Y^{\otimes n}) \quad (7.89)$$

$$= \frac{1}{K} \sum_{P \in \hat{\mathcal{P}}_n | \text{wt } P=j} \sum_{a,b} \langle b^*| Y^{\otimes n} P |a\rangle \langle a| P^\dagger Y^{\otimes n} |b^*\rangle \quad (7.90)$$

$$\geq 0 \quad (7.91)$$

since $Y = Y^\dagger$.

b) The proof proceeds similarly to that for the quantum MacWilliams identity. We again write

$$\Pi = \sum_{Q \in \hat{\mathcal{P}}_n} c_Q Q, \quad (7.92)$$

and find

$$Sh_j = \frac{1}{K} \sum_{P \in \hat{\mathcal{P}}_n | \text{wt}(P)=j} \sum_{Q, R \in \hat{\mathcal{P}}_n} c_Q c_R^* \text{tr}(P Q P^\dagger Y^{\otimes n} R^* Y^{\otimes n}) \quad (7.93)$$

$$= \frac{1}{K} \sum_{P \in \hat{\mathcal{P}}_n | \text{wt}(P)=j} \sum_{Q, R \in \hat{\mathcal{P}}_n} c_Q c_R^* (-1)^{c(P,Q)} \text{tr}(Q Y^{\otimes n} R^* Y^{\otimes n}). \quad (7.94)$$

Now, $X^* = X$ and $Z^* = Z$, but $Y^* = -Y$, so $R^* = \pm R$, with the sign determined by whether there are an even or an odd number of Y s in the tensor decomposition of R . Furthermore,

$$Y^{\otimes n} R^* Y^{\otimes n} = (-1)^{c(R, Y^{\otimes n})} R^*. \quad (7.95)$$

$c(R, Y^{\otimes n})$ is 1 if the number of X s plus Z s in the tensor decomposition of R is odd, and 0 if it is even. Thus,

$$\text{tr}(Q Y^{\otimes n} R^* Y^{\otimes n}) = (-1)^{\text{wt}(R)} \delta_{Q,R} 2^n. \quad (7.96)$$

The sign follows because if the number of Y s and the number of X s plus Z s in R are both even or both odd, we get an overall factor of $+1$, whereas if one of them is odd and one is even, we get a factor of -1 .

We therefore have

$$Sh_j = \frac{1}{K} \sum_{P \in \hat{\mathcal{P}}_n | \text{wt}(P)=j} \sum_{Q, R \in \hat{\mathcal{P}}_n} c_Q c_R^* (-1)^{c(P,Q) + \text{wt}(R)} \delta_{Q,R} 2^n \quad (7.97)$$

$$= \frac{2^n}{K} \sum_{Q \in \hat{\mathcal{P}}_n} \left[\sum_{P \in \hat{\mathcal{P}}_n | \text{wt}(P)=j} (-1)^{c(P,Q)} \right] (-1)^{\text{wt}(Q)} |c_Q|^2. \quad (7.98)$$

The evaluation of the sum over P is just the same as before. We again get a Krawtchouk polynomial, to find

$$Sh(x) = \frac{K}{2^n} \sum_{j=0}^n \sum_{i=0}^n (-1)^i P_j(i; n) A_i x^j \quad (7.99)$$

$$= \frac{K}{2^n} \sum_{i=0}^n (1+3x)^{n-i} (1-x)^i (-1)^i A_i \quad (7.100)$$

$$= \frac{K}{2^n} (1+3x)^n \sum_{i=0}^n \left(\frac{x-1}{1+3x} \right)^i A_i \quad (7.101)$$

$$= \frac{K}{2^n} (1+3x)^n A \left(\frac{x-1}{1+3x} \right). \quad (7.102)$$

□

7.4.4 Example: There Is No $((3, 2, 2))$ Code

OK. We have these three weight enumerators, but what do we do with them? How can we put together the constraints we have to find out if certain parameters of codes are possible or not? One approach is to make approximations and use properties of the Krawtchouk polynomials to prove non-existence of certain solutions. This approach gives another method of proving the quantum Singleton bound, and can also give tighter bounds on the existence of QECCs.

Another approach is to apply the linear programming method. The procedure is to treat the parameters A_j , B_j , and Sh_j as unknown variables, and write down all the equations we have above:

$$A_0 = 1 \tag{7.103}$$

$$B_0 = 1 \tag{7.104}$$

$$B_j = A_j \text{ (for } j < d) \tag{7.105}$$

$$B_j \geq A_j \text{ (for } j \geq d) \tag{7.106}$$

$$A_j \geq 0 \tag{7.107}$$

$$Sh_j \geq 0 \tag{7.108}$$

$$B(x) = \frac{K}{2^n} (1 + 3x)^n A \left(\frac{1 - x}{1 + 3x} \right) \tag{7.109}$$

$$Sh(x) = \frac{K}{2^n} (1 + 3x)^n A \left(\frac{x - 1}{1 + 3x} \right) \tag{7.110}$$

Of course, many of the lines actually represent groups of equations, running over values of j . Notice that all the equations, even the last two groups, are linear in A_j , B_j , and Sh_j . Some are equalities and some are inequalities. The question of whether a solution exists or not is a linear programming problem, which is a common task. There are standard computer packages for solving linear programming problems, and so the usual method is to put the equations corresponding to the desired parameters $((n, K, d))$ of a code into such a package, and see whether it says a solution is possible. If not, we know no code can exist with those parameters. If so, we learn the solution(s) of the linear programming problem, which doesn't tell us whether a code exists or not, but does at least tell us possible values for the code's weight enumerator if it does exist, and therefore something about the structure of the code.

For large parameters, this is certainly best done by a computer, but to show you how the procedure works, I'll do a small example here explicitly, showing that there is no $((3, 2, 2))$ QECC. It is easy to show that there is no $[[3, 1, 2]]$ stabilizer code, but for more general codes, we need a more powerful method, and the linear programming bounds will suffice. The code is allowed by the quantum Singleton bound, and indeed, you'll see a $((3, 3, 2))_3$ code over qutrits in chapter 8.

Let us start by writing down the equations produced by the quantum MacWilliams identity:

$$B(x) = \frac{2}{2^3} (1 + 3x)^3 A \left(\frac{1 - x}{1 + 3x} \right) \tag{7.111}$$

$$= \frac{1}{4} [(1 + 3x)^3 A_0 + (1 + 3x)^2 (1 - x) A_1 + (1 + 3x)(1 - x)^2 A_2 + (1 - x)^3 A_3] \tag{7.112}$$

Collecting the coefficients of x^j for $j = 0, \dots, 3$, we find

$$4B_0 = A_0 + A_1 + A_2 + A_3 \tag{7.113}$$

$$4B_1 = 9A_0 + 5A_1 + A_2 - 3A_3 \tag{7.114}$$

$$4B_2 = 27A_0 + 3A_1 - 5A_2 + 3A_3 \tag{7.115}$$

$$4B_3 = 27A_0 - 9A_1 + 3A_2 - A_3. \tag{7.116}$$

We also know that $A_0 = A_1 = 1$, $B_1 = A_1 \geq 0$, $B_2 \geq A_2 \geq 0$, and $B_3 \geq A_3 \geq 0$. Combining these equations, we find

$$A_1 + A_2 + A_3 = 3 \quad (7.117)$$

$$A_1 + A_2 - 3A_3 = -9 \quad (7.118)$$

$$3A_1 - 9A_2 + 3A_3 \geq -27 \quad (7.119)$$

$$-9A_1 + 3A_2 - 5A_3 \geq -27 \quad (7.120)$$

From equations (7.117) and (7.118), we find $A_3 = 3$, so $A_1 + A_2 = 0$. Since $A_1, A_2 \geq 0$, this means that $A_1 = A_2 = 0$. This solution also satisfies equations (7.119) and (7.120). Using only the quantum MacWilliams identity, we cannot rule out a $((3, 2, 2))$ code, but we know that if one exists, it must have weight enumerator $A(x) = 1 + 3x^3$ and dual weight enumerator $B(x) = 1 + 9x^2 + 6x^3$.

We now invoke the shadow enumerator. Let us write down the relation between $A(x)$ and $Sh(x)$, and plug in the above value for $A(x)$:

$$Sh(x) = \frac{1}{4} [(1 + 3x)^3 A_0 + (1 + 3x)^2 (x - 1) A_1 + (1 + 3x)(x - 1)^2 A_2 + (x - 1)^3 A_3] \quad (7.121)$$

$$= \frac{1}{4} [(1 + 3x)^3 + 3(x - 1)^3] \quad (7.122)$$

$$= \frac{1}{4} (-2 + 18x + 18x^2 + 30x^3). \quad (7.123)$$

However, we see that $Sh_0 < 0$, which is not allowed. Thus, a $((3, 2, 2))$ code is not possible.

Chapter 8

Bigger Can Be Better: Qudit Codes

Computers seem to like base 2, even when they are quantum computers. Consequently, it may seem natural to stick to bits and qubits when dealing with error correcting codes, classical or quantum. Indeed, I've gone to a lot of effort in the previous chapters to build the mathematical structure of stabilizer codes and the Clifford group, all of which depends on the basic registers of our quantum computer being qubits. The problem is that there are lots of interesting codes that don't quite fit into this structure. Instead, they fit into similar structures associated to higher-dimensional registers — qudits.

The phenomenon also occurs in the theory of classical error correction, where it's helpful to work with linear codes over arbitrary finite fields rather than just binary linear codes. That's why I introduced non-binary codes in chapter 4. Now we'll see how to do the same for quantum codes. Quantum codes with bigger registers can generally correct more errors and send quantum data with a higher rate than qubit codes.

Based on the last sentence, this may appear a clear-cut case of bigger being better, but it's not quite that straightforward. You can't compare an error on a single qubit with an error on a 32-dimensional qudit on an even basis. I wouldn't say you are comparing apples to oranges; rather, you are comparing apples to boxes of apples. A 32-dimensional qudit could be written using 5 qubits, so a single error on a 32-dimensional error could be 5 single-qubit errors, and a single-qudit gate in 32 dimensions might need many one- and two-qubit gates to achieve the same transformation. Nevertheless, some qudit codes are sufficiently interesting, in terms of efficiency or other properties, to be worthwhile even once you take the size difference of the registers into account.

8.1 Qudit Pauli Group

A *qudit* is a generic term for a q -dimensional Hilbert space, considered as a fundamental unit of the overall Hilbert space of dimension q^n in much the same way that a qubit is a fundamental unit of a 2^n -dimensional Hilbert space. A 2^n -dimensional Hilbert space has many possible factorizations into two-dimensional tensor factors, but one factorization is considered physically favored, and that factorization gives us the physical qubits making up the full Hilbert space. (The *logical* qubits encoded in a quantum error-correcting code make up part of a different factorization.) Similarly, a q^n -dimensional Hilbert space is generally assumed to have some physically favored tensor product decomposition, producing the physical qudits comprising the code.

The word “qudit” for a q -dimensional Hilbert space seems to have mostly won out in the literature, but in some of the older literature you can find a variety of other terms, such as “qupit” (which usually assumes the dimension is a prime p). You can also find, in papers young and old, occasional terms referring to a qudit for a dimension of a specific size. “Qutrit” is the most common, for $D = 3$. Beyond that, you may have to delve into the word's Latin or Greek roots to figure out the dimension. For instance, a 4-dimensional qudit might be a “ququart”, a 5-dimensional one might be a “ququint” or a “qupent” or some such, and for $D = 6$, you might even get lucky with a “qusexxt.”

The first step in building quantum error-correcting codes for qudits, by whatever name, is to come up with the correct generalization of the Pauli group. As for qubits, the qudit Pauli group will play a role both in describing the types of errors we face and as a structural element for the qudit generalization of stabilizer codes. We're going to make heavy use of the machinery of finite fields in this chapter, so if you're not already familiar with it, you may want to go through appendix C first.

8.1.1 Qudit Pauli Group for Prime Dimension

The case where the qudit dimension p is prime is the simplest, so we'll start there.

Definition 8.1. The single-qudit *Pauli group* $P_1(p)$ for prime $p > 2$ consists of elements $\{\omega^a X^b Z^c\}$, where

$$\omega = e^{2\pi i/p} \quad (8.1)$$

$$X|j\rangle = |(j+1) \bmod p\rangle \quad (8.2)$$

$$Z|j\rangle = \omega^j |j\rangle, \quad (8.3)$$

and a, b, c can be anywhere from 0 to $p-1$. The n -qudit *Pauli group* $P_n(p) = P_1(p)^{\otimes n}$; it consists of tensor products of n terms, each of the form $X^b Z^c$, with an overall phase factor of ω^a . $\hat{P}_n(p) = P_n(p)/\{\omega^a I\}$ is the qudit Pauli group without phases.

That is, ω is a p th root of unity, X is “add one mod p ”, and Z is a phase shift by a p th root of unity. We can write X and Z as matrices, of course:

$$X = \begin{pmatrix} 0 & 0 & \dots & 0 & 1 \\ 1 & 0 & \dots & 0 & 0 \\ 0 & 1 & \dots & 0 & 0 \\ \vdots & & \ddots & & \vdots \\ 0 & 0 & \dots & 1 & 0 \end{pmatrix} \quad Z = \begin{pmatrix} 1 & 0 & \dots & 0 & 0 \\ 0 & \omega & \dots & 0 & 0 \\ \vdots & & \ddots & & \vdots \\ 0 & 0 & \dots & \omega^{p-2} & 0 \\ 0 & 0 & \dots & 0 & \omega^{p-1} \end{pmatrix} \quad (8.4)$$

You can see that X and Z are still unitary but are no longer Hermitian. I'd still like to talk about measuring X or Z ; for any traditionalists reading this who prefer to measure only Hermitian operators, just imagine measuring a Hermitian operator that has the same eigenbasis as X or Z . For instance, $\sum_j j|j\rangle\langle j|$ works instead of Z , and outcome j corresponds to the ω^j eigenvalue of Z .

We can calculate the commutation relations between elements of the Pauli group. We find

$$ZX = \omega XZ. \quad (8.5)$$

It follows that

$$(X^a Z^b)(X^c Z^d) = \omega^{bc-ad}(X^c Z^d)(X^a Z^b). \quad (8.6)$$

Since ω is a p th root of unity, we might as well calculate $bc - ad$ using arithmetic modulo p . We can define a new function $c(P, Q) : P_n(p) \times P_n(p) \rightarrow \mathbb{Z}_p$ by

$$PQ = \omega^{c(P, Q)}QP. \quad (8.7)$$

We no longer are choosing between commuting and anti-commuting Paulis; now commutation is measured by an integer modulo p , which describes the phase factor that appears when we move P past Q as a power of ω .

A very important difference when p is an odd prime rather than 2 is that we only need the phases to be powers of ω ; we don't get the annoying factor of i that shows up for qubits. This is because all elements of

$P_n(p)$ have order p , regardless of overall phase:

$$(\omega^a X^b Z^c)^p = \omega^{ap} (X^b Z^c) (X^b Z^c)^{p-1} \quad (8.8)$$

$$= \omega^{bc} X^{2b} Z^{2c} (X^b Z^c)^{p-2} \quad (8.9)$$

$$= \omega^{bc+2bc} X^{3b} Z^{3c} (X^b Z^c)^{p-3} \quad (8.10)$$

$$\dots = \omega^{bc+2bc+\dots+(p-1)bc} X^p Z^p \quad (8.11)$$

$$= \omega^{bcp(p-1)/2} \quad (8.12)$$

$$= 1, \quad (8.13)$$

since p is odd. This also means that the eigenvalues of all elements of the qudit Pauli group are of the form ω^j for integer j . While the qudit Pauli group is more complicated than the qubit Pauli group in a variety of ways, in this one way, at least, it is simpler.

There are alternate representations of the qudit Pauli group, just as for the qubit Pauli group. For instance, we can define a \mathbb{Z}_p symplectic representation of the qudit Pauli group by

$$P = \bigotimes_{j=1}^n X^{a_j} Z^{b_j} \mapsto v_P = (x_P | z_P), \quad (8.14)$$

with x_P and z_P being n -component vectors over \mathbb{Z}_p . x_P has entries a_j , and z_P has entries b_j . As with qubits, the overall phase of P is lost when moving to the symplectic representation.

Suppose we have v_P as defined in equation (8.14) and v_Q , with

$$Q = \bigotimes_{j=1}^n X^{c_j} Z^{d_j}. \quad (8.15)$$

We can define a symplectic inner product between v_P and v_Q as

$$v_P \odot v_Q = \sum_j (b_j c_j - a_j d_j), \quad (8.16)$$

with arithmetic taken mod p . With this definition, we have the following result:

Proposition 8.1. *Let $P, Q \in P_n(p)$. Then $c(P, Q) = v_P \odot v_Q$.*

In other words, commutation in the qudit Pauli group corresponds to the symplectic inner product.

We can also map the qudit Pauli group to $\text{GF}(p^2)$. However, the procedure is a bit more complicated than for qubits. We should think of $\text{GF}(p^2)$ as a 2-dimensional vector space over $\text{GF}(p)$ with basis $\{1, \alpha\}$. α can be any element of $\text{GF}(p^2)$ that is not in the $\text{GF}(p)$ subfield. We can write arbitrary $\beta \in \text{GF}(p^2)$ as

$$\beta = a + b\alpha, \quad (8.17)$$

with $a, b \in \text{GF}(p)$. We then map

$$X^a Z^b \mapsto a + b\alpha. \quad (8.18)$$

So far, so good. All is the same as for qubits. We lose the overall phase just as before. If we start with an element of $\hat{P}_n(p)$, we end up with an n -component vector over $\text{GF}(p^2)$ (with each element written as a 2-component vector over $\text{GF}(p)$).

The complication comes in when we look at the generalization of the symplectic inner product. We wish to define it using multiplication within $\text{GF}(p^2)$, and we want it to output an element of $\text{GF}(p)$. The latter condition could be achieved by using the trace function (see appendix C for a discussion), but it turns out we don't want to do that in this case. For the former condition, the answer is not at all obvious, and historically took some time to find. It turns out that the inner product we want is given by the following definition:

Definition 8.2. Suppose a, b are n -dimensional vectors over $\text{GF}(p^2)$. Let

$$a * b = \frac{a \cdot b^p - b \cdot a^p}{\alpha - \alpha^p}, \quad (8.19)$$

where a^p and b^p are the n -dimensional vectors whose entries are the p th power of the entries of a and b , and \cdot represents the usual dot product between vectors.

Notice that the symplectic product we define (also called a *trace-alternating* inner product) depends explicitly on the specific element $\alpha \in \text{GF}(p^2)$ that we used to give the mapping from $\hat{\mathbb{P}}_1(p)$ to $\text{GF}(p^2)$. That wasn't the case for qubits because there were only two elements of $\text{GF}(4)$ which are outside $\text{GF}(2)$, namely ω and ω^2 , but in $\text{GF}(p^2)$ there are more choices for the pair (α, α^p) .

I don't know of a better way to motivate this formula than by just calculating, but it turns out that it does work:

Theorem 8.2. Suppose $P, Q \in \hat{\mathbb{P}}_n(p)$ correspond to vectors a, b over $\text{GF}(p^2)$ according to equation (8.18). Then

$$a * b = c(P, Q). \quad (8.20)$$

Proof. Let us calculate. We have

$$P \mapsto a = x + \alpha z \quad (8.21)$$

$$Q \mapsto b = x' + \alpha z'. \quad (8.22)$$

Now,

$$(x + \alpha z)^p = x^p + \alpha^p z^p = x + \alpha^p z, \quad (8.23)$$

since the field has characteristic p and x, z are vectors over $\text{GF}(p)$. Then

$$a \cdot b^p = (x + \alpha z) \cdot (x' + \alpha z')^p \quad (8.24)$$

$$= x \cdot x' + \alpha z \cdot x' + \alpha^p x \cdot z' + \alpha^{p+1} z \cdot z' \quad (8.25)$$

$$b \cdot a^p = x \cdot x' + \alpha z' \cdot x + \alpha^p x' \cdot z + \alpha^{p+1} z \cdot z'. \quad (8.26)$$

Thus,

$$a \cdot b^p - b \cdot a^p = \alpha(z \cdot x' - x \cdot z') + \alpha^p(x \cdot z' - z \cdot x') \quad (8.27)$$

$$= (\alpha - \alpha^p)(x|z) \odot (x'|z') \quad (8.28)$$

$$= (\alpha - \alpha^p)c(P, Q). \quad (8.29)$$

□

8.1.2 Qudit Pauli Group for Prime Power Dimensions

The next most complicated case is when each qudit has dimension $q = p^m$, with p a prime. For these dimensions, the most straightforward thing to do is to let $\mathbb{P}_n(q) = \mathbb{P}_m(p)^{\otimes n}$. In other words, we consider each q -dimensional qudit as broken up into m p -dimensional qudits.

That's fine, as far as it goes, but we'd like to put more structure on the Hilbert space. In particular, we'd like to consider each q -dimensional qudit as an element of $\text{GF}(q)$. To fully take advantage of that structure, we'd like to also involve $\text{GF}(q)$ in our understanding of $\mathbb{P}_n(q)$. We can do so as follows:

Definition 8.3. Suppose $\beta \in \text{GF}(q)$. Let X^β and Z^β be defined as

$$X^\beta |\gamma\rangle = |\gamma + \beta\rangle \quad (8.30)$$

$$Z^\beta |\gamma\rangle = \omega^{\text{tr}(\beta\gamma)} |\gamma\rangle, \quad (8.31)$$

where $\gamma \in \text{GF}(q)$, $\omega = \exp(2\pi i/p)$ is a p th root of unity and tr is the $\text{GF}(q)$ trace function which maps elements of $\text{GF}(q)$ to elements of $\text{GF}(p)$. Then $\text{P}_n(q)$ consists of elements of the form

$$\eta \omega^a \bigotimes_{j=1}^n X^{\beta_j} Z^{\gamma_j}, \quad (8.32)$$

where $a \in \text{GF}(p)$, $\beta_j, \gamma_j \in \text{GF}(q)$. For odd q , η is always 1. For even q , η can be 1 or i . As usual, $\hat{\text{P}}_n(q) = \text{P}_n(q)/\{\omega^a I\}$ (for odd p) or $\hat{\text{P}}_n(q) = \text{P}_n(q)/\{i^a I\}$ (for $p = 2$).

Note that

$$Z^\gamma X^\beta = \omega^{\text{tr}(\gamma\beta)} X^\beta Z^\gamma. \quad (8.33)$$

Now let us examine how this definition plays out given a particular way of breaking up each q -dimensional register into m p -dimensional ones. Suppose we consider $\text{GF}(q)$ as an m -dimensional vector space over $\text{GF}(p)$, so

$$\gamma = \sum_{i=0}^{m-1} a_i \alpha_i, \quad (8.34)$$

with $\gamma \in \text{GF}(q)$, $a_i \in \text{GF}(p)$, and the α_i 's elements of $\text{GF}(q)$ which are linearly independent vectors when considered over $\text{GF}(p)$. It is most common to take $\alpha_i = \alpha^i$, with α some primitive element of $\text{GF}(q)$. Note that with this choice, $\alpha_0 = 1$.

We can immediately see how equation (8.34) corresponds to breaking the qudit up into pieces: $|\gamma\rangle \leftrightarrow |a_0\rangle \otimes |a_1\rangle \otimes \cdots \otimes |a_{m-1}\rangle$. X^{α_i} then has a natural interpretation as the p -dimensional X applied to the i th tensor factor:

$$X^{\alpha_i} |\gamma\rangle = |\gamma + \alpha_i\rangle = \left| \sum_j (a_j + \delta_{ij}) \alpha_j \right\rangle. \quad (8.35)$$

If we apply $X^{\sum b_i \alpha_i}$ instead, that corresponds to performing X^{b_i} in the i th tensor factor. That is, to understand the action of X^β , we expand β in the same basis used for the standard basis decomposition, and apply the appropriate power of X on each factor.

The interpretation of Z^β is a little trickier. It hinges on the notion of a *dual basis* (sometimes called a *complementary basis*): The set $\{\alpha_i\}$ forms a basis for $\text{GF}(q)$ considered as an m -dimensional vector space over $\text{GF}(p)$, and for any basis, there exists a dual basis $\{\beta_j\}$ with the property:

$$\text{tr}(\alpha_i \beta_j) = \delta_{ij}. \quad (8.36)$$

Then

$$Z^{\beta_j} |\gamma\rangle = \omega^{\text{tr}(\beta_j \sum_i a_i \alpha_i)} \bigotimes_{i=0}^{m-1} |a_i\rangle = \omega^{a_j} \bigotimes_{i=0}^{m-1} |a_i\rangle, \quad (8.37)$$

since trace is linear over $\text{GF}(p)$. That is, Z^{β_j} corresponds to performing the p -dimensional Z on the j th tensor factor. To understand the action of Z^β in general, we simply then expand β in the dual basis to the one used for the standard basis. The choice of the dual bases $\{\alpha_i\}$ and $\{\beta_j\}$ thus specifies an isomorphism $\text{P}_n(q) \cong \text{P}_{mn}(p)$.

For some fields (and in particular for $q = 2^m$, which is the most interesting case), it is possible to simplify the decomposition by choosing the basis $\{\alpha_i\}$ to be *self-dual*, i.e., $\beta_i = \alpha_i$. In that case, X^{α_i} and Z^{α_i} simply represent the Pauli matrices acting on the i th tensor factor of the q -dimensional register. Not all finite fields have self-dual bases, unfortunately, so for some values of q , we have to pick different decompositions for the exponents of X and Z . Alternatively, we could abandon definition 8.3, but the advantages of that notation greatly outweigh the inconvenience of having a slightly more complicated decomposition into p -dimensional qudits.

For either odd or even characteristic, when we drop phases from the Pauli group, we get vectors on a symplectic space:

$$P = \bigotimes_{j=1}^n X^{\eta_j} Z^{\gamma_j} \mapsto v_P = (x_P | z_P), \quad (8.38)$$

with x_P an n -dimensional vector over $\text{GF}(q)$ with entries η_j and z_P an n -dimensional vector with entries γ_j . The symplectic inner product between v_P and v_Q (with $Q = \bigotimes X^{\eta'_j} Z^{\gamma'_j}$) is

$$v_P \odot v_Q = \sum_j \text{tr}(\gamma_j \eta'_j - \eta_j \gamma'_j). \quad (8.39)$$

Multiplication now is the $\text{GF}(q)$ multiplication rule, and we take the trace to end up with an element of $\text{GF}(p)$. Once more proposition 8.1 applies.

You probably can guess the next step: We wish to map the q -dimensional Pauli group into $\text{GF}(q^2)$. The procedure is similar to that for $\text{GF}(p)$; we pick an element $\alpha \in \text{GF}(q^2) \setminus \text{GF}(q)$, and map

$$(x|z) \mapsto x + \alpha z. \quad (8.40)$$

The correct symplectic inner product in this case is

$$a * b = \text{tr}_{q/p} \left(\frac{a \cdot b^q - b \cdot a^q}{\alpha - \alpha^q} \right). \quad (8.41)$$

There are two differences from equation (8.19): We use the exponent q instead of p , and we use the trace function to give us an element of $\text{GF}(p)$ for the answer. Note that we are using the trace of $\text{GF}(q)$ over $\text{GF}(p)$, *not* the trace of $\text{GF}(q^2)$. This is because the term in parentheses already gives an element of $\text{GF}(q)$.

Theorem 8.3. *Suppose $P, Q \in \hat{\mathcal{P}}_n(q)$ correspond to vectors a, b over $\text{GF}(q^2)$ according to equation (8.40). Then*

$$a * b = c(P, Q). \quad (8.42)$$

The proof is essentially the same as theorem 8.2. We just need to add one final step where we take the trace to get the $\text{GF}(q)$ symplectic product.

8.1.3 Qudit Pauli Group for Other Dimensions

If the dimension q of the register is not a prime power, there are two sensible ways to generalize the above approaches to define a Pauli group. One idea is to break q up into its prime factorization $q = \prod p_i^{m_i}$, and treat each prime power factor as a separate sub-register of size $p_i^{m_i}$ with its own Pauli group.

The second approach is to directly generalize the Pauli group used for prime dimensions, by letting

$$\omega = e^{2\pi i/q} \quad (8.43)$$

$$X|j\rangle = |(j+1) \bmod q\rangle \quad (8.44)$$

$$Z|j\rangle = \omega^j |j\rangle. \quad (8.45)$$

As usual, the n -qubit Pauli group consists of products of the form $\omega^a \bigotimes X^b Z^c$ for odd q . For even q , we must include a possible overall factor of i as well. This version of the Pauli group is also known as the *Heisenberg-Weyl group*.

It is also possible to use the Heisenberg-Weyl group in place of the usual Pauli group for prime power dimensions. These groups are different: For instance, in the $q = 9$ Heisenberg-Weyl group, X has order 9, whereas all elements of the usual $q = 9$ Pauli group have order 3. There are some applications where this is a sensible thing to do, but the cost of using the Heisenberg-Weyl group is that we lose the field structure we normally have in prime power dimensions.

Because the mathematical structure of the Heisenberg-Weyl group is more complicated than the Pauli groups for prime dimension or prime-power dimension, the standard techniques of coding theory don't work as well. The basic structure of a stabilizer code still exists, but codes based on the Heisenberg-Weyl group lack some of the usual properties of stabilizer codes.

8.1.4 Nice Error Bases

Indeed, we can generalize even further and still have a stabilizer code structure. The most important features are that the elements of our generalized Pauli group form a group, are independent, and span the set of possible errors. This is codified by the definition of a nice error basis:

Definition 8.4. In a q -dimensional Hilbert space, let $\mathcal{E} = \{E_1, \dots, E_{q^2}\}$ be a set of unitary operators satisfying $E_1 = I$ and $\text{tr } E_i^\dagger E_j = q\delta_{ij}$. The set \mathcal{E} is a *nice error basis* if $E_i E_j = \omega_{ij} E_{f(i,j)}$ for all i, j and some phases ω_{ij} .

For instance, for a single qubit, the usual Pauli operators $\{E_1 = I, E_2 = X, E_3 = Y, E_4 = Z\}$ form a nice error basis, with $\omega_{ij} \in \{\pm 1, \pm i\}$. To get a generalized Pauli group from a nice error basis, we take elements of the form ωE_i , ω is a product of ω_{ij} s.

Because there are exactly q^2 independent elements in a nice error basis, they form a basis for the space of $q \times q$ matrices. In this sense, they can act like the Pauli matrices — we can take any error on the q -dimensional register and expand it as a sum of elements from the nice error basis. It is thus sufficient for a QECC to correct all errors in a nice error basis to correct arbitrary errors on the register. This justifies the “error basis” part of the name.

Proposition 8.4. *The indices i of the errors in a nice error basis form a group under multiplication given by the binary operation $f(i, j)$.*

Proof. By the definition of a nice error basis, the set of indices is closed under the group operation. Associativity follows from the associativity of operator multiplication:

$$E_i E_j E_k = \omega_{ij} E_{f(i,j)} E_k = \omega_{ij} \omega_{f(i,j)k} E_{f(f(i,j),k)} \quad (8.46)$$

$$= \omega_{jk} E_i E_{f(j,k)} = \omega_{jk} \omega_{if(j,k)} E_{f(i,f(j,k))}. \quad (8.47)$$

Since the E_i s are orthogonal, it follows that $f(f(i, j), k) = f(i, f(j, k))$, the statement of associativity.

The group identity is $i = 1$ since $E_1 = I$:

$$\omega_{1i} E_{f(1,i)} = I E_i = E_i. \quad (8.48)$$

To establish the existence of inverses, first note that if $f(i, j) = 1$, then

$$E_i E_{j'} = \omega_{ij'} \omega_{ij'}^{-1} E_i E_j, \quad (8.49)$$

implying $E_{j'} = \omega E_j$ for some phase ω . Since the E_i s are orthogonal, it follows that $j = j'$. Therefore, the function $j \mapsto f(i, j)$ is one-to-one, and since the domain and range both have size q^2 , it must be onto as well. In particular, for all i , there must exist i^{-1} such that $f(i, i^{-1}) = 1$, and i^{-1} is the inverse of i . \square

Definition 8.5. The group of indices $\{i\}$ under group operation $f(i, j)$ is called the *index group* of the nice error basis.

The index group can also be obtained by taking the Pauli group generated by the nice error group and modding out overall phase. For the usual qubit Pauli group, the index group is therefore isomorphic to $\mathbb{Z}_2 \times \mathbb{Z}_2$. For the Heisenberg-Weyl group in dimension q , the index group is $\mathbb{Z}_q \times \mathbb{Z}_q$, and for the prime power Pauli group, the index group is $\text{GF}(q) \times \text{GF}(q)$ under addition. All of these are fairly straightforward Abelian groups, but for large q , there exist error groups with more exotic index groups, including non-Abelian ones.

8.2 Qudit Stabilizer Codes

Now we are ready to talk about actual QECCs using qudits as registers. After going to all that effort to find qudit analogs of the Pauli group, you might guess that I'm now going to define a qudit analog of stabilizer codes. If you guessed that, you'd be correct.

8.2.1 Definition and Properties of Qudit Stabilizer Codes

The basic definition of a stabilizer and a stabilizer code is the same as for qubits. When q is a non-prime-power or when we are using a nice error group more exotic than the standard $P_n(q)$, it is still possible to define a stabilizer code, but there are some complications to the theory. I'll just stick to the simpler cases.

The definition of a stabilizer is the same as for qubits:

Definition 8.6. Let $q = p^m$ be a prime power, $P = P_n(q)$, and let Q be a subspace of the Hilbert space $\mathcal{H}_q^{\otimes n}$ (i.e., consisting of n qudits). The $\text{GF}(q)$ stabilizer of Q is the set

$$S(Q) = \{M \in P \mid \psi \text{ is an eigenvector of } M \text{ with eigenvalue } +1 \forall \psi \in Q\}. \quad (8.50)$$

Let S be a subgroup of P . We say S is a $\text{GF}(q)$ stabilizer, or just stabilizer when the $\text{GF}(q)$ is clear from context, if it is Abelian and if $e^{i\phi}I \notin S$ for any phase $\phi \neq 0$. The *code space* of a $\text{GF}(q)$ stabilizer S is the subspace

$$\mathcal{T}(S) = \{|\psi\rangle \text{ s.t. } M|\psi\rangle = |\psi\rangle \forall M \in S\}. \quad (8.51)$$

The code Q is a $\text{GF}(q)$ stabilizer code iff $Q = \mathcal{T}(S(Q))$. The *normalizer* $N(S)$ of the stabilizer S is

$$N(S) = \{N \in P \mid NM = MN \forall M \in S\}. \quad (8.52)$$

The only real differences from the qubit definition are the use of $P_n(q)$ instead of P_n and forbidding $e^{i\phi}$ for all non-zero phases ϕ , which is needed because there are more phases in $P_n(q)$ than just $\pm i, \pm 1$.

Note that a $\text{GF}(q)$ stabilizer automatically gives a $\text{GF}(p)$ stabilizer when $q = p^m$ when we reinterpret $P_n(q)$ as $P_{mn}(p)$ as discussed in section 8.1.2. Indeed, using the same isomorphism, we can interpret a $\text{GF}(p)$ stabilizer on mn qudits as a $\text{GF}(p^m)$ stabilizer on n qudits. The difference between them is the definition of weight (and thus distance of a code), which counts the number of non-trivial p -dimensional qudits in an operator for a $\text{GF}(p)$ stabilizer but the number of non-trivial q -dimensional qudits for a $\text{GF}(q)$ stabilizer. Thus, a Pauli which has weight t for $\text{GF}(q)$ might have weight up to mt for $\text{GF}(p)$, but there are also Paulis which have the same weight for $\text{GF}(p)$ and $\text{GF}(q)$.

Usually, although this is not required by definition 8.6, we deal with $\text{GF}(q)$ stabilizer codes that have an additional property:

Definition 8.7. Let $P \in P_n(q)$ have $\text{GF}(q)$ symplectic representation $(x_P|z_P)$. Let S be a $\text{GF}(q)$ stabilizer, with symplectic representation \hat{S} . Then we say S is a *true* $\text{GF}(q)$ stabilizer if $(x_P|z_P) \in \hat{S}$ implies $(\gamma x_P|\gamma z_P) \in \hat{S}$ as well for all $\gamma \in \text{GF}(q)$.

That is, for a true $\text{GF}(q)$ stabilizer code, the symplectic representation of the stabilizer is a $\text{GF}(q)$ -linear space. Note that if q is prime, any $\text{GF}(q)$ stabilizer code is a true $\text{GF}(q)$ stabilizer code since $P \in S$ implies $P^i \in S$ as well, and

$$(x_{P^i}|z_{P^i}) = i(x_P|z_P). \quad (8.53)$$

However, when $q = p^m$ with $m > 1$, then there are some elements of $\text{GF}(q)$ which are not numbers, so it is possible to have stabilizer codes which are not true $\text{GF}(q)$ stabilizer codes.

$\text{GF}(q)$ stabilizer codes have the same properties we are familiar with from qubit stabilizer codes.

Proposition 8.5. *If Q is a non-trivial subspace of the Hilbert space, then $S(Q)$ is a $\text{GF}(q)$ stabilizer (not necessarily a true $\text{GF}(q)$ stabilizer). If S is a $\text{GF}(q)$ stabilizer, then $S(\mathcal{T}(S)) = S$.*

The proof is almost identical to the qubit case. For prime qudit dimension p , we also get analogues of proposition 3.3 and theorem 3.4.

Theorem 8.6. *Let p be prime and let S be a $\text{GF}(p)$ stabilizer for the code $\mathcal{T}(S)$, which has n physical qudits. If $|\mathcal{S}| = p^r$ (i.e., S has r generators), then $\dim \mathcal{T}(S) = p^{n-r}$, so $\mathcal{T}(S)$ encodes $k = n - r$ physical qudits. The set of undetectable errors for S is $\hat{N}(S) \setminus \hat{S}$. The distance of S is $\min\{\text{wt } E \mid E \in \hat{N}(S) \setminus \hat{S}\}$.*

The proofs are closely analogous to the qubit case, so I omit the details. The main notable difference is that $\frac{1}{2}(I + M)$ is *not* the projection operator on the +1 eigenspace of M . Instead, the projector on the +1 eigenspace of M is $\frac{1}{p} \sum_{j=0}^{p-1} M^j$. The projector on the codespace can be written as

$$\Pi_S = \frac{1}{p^r} \sum_{M \in S} M. \quad (8.54)$$

(Note that the normalization is p^{-r} instead of 2^{-r} .) The other difference is that in the case where the error $E \notin \mathbf{N}(S)$, there is a phase ω^a instead of -1 , but that does not really alter the proof.

For prime power dimensions, there is a complication due to the fact that all Paulis have order p rather than q . It is most helpful to think of a $\text{GF}(q)$ stabilizer of n qudits as a $\text{GF}(p)$ stabilizer on mn qudits, in which case we can apply theorem 8.6. Considered as a $\text{GF}(p)$ stabilizer, a stabilizer with r generators has p^r elements and encodes a p^{mn-r} -dimensional Hilbert space, the straightforward analog of the qubit result. If you insist on thinking of it as a $\text{GF}(q)$ stabilizer, the same stabilizer still has r generators and p^r elements and encodes a Hilbert space of dimension $p^{mn-r} = q^{n-r/m}$ (it is, after all, the same code). In particular, a $\text{GF}(q)$ stabilizer does not need to encode an integer number of q -dimensional qudits. However, a *true* $\text{GF}(q)$ stabilizer code always has a number of generators which is a multiple of m , so does encode an integer number of q -dimensional qudits.

The other subtlety in prime power dimensions is the definition of distance, and indeed, this is the only property for which it really makes a difference whether we think of the code as a $\text{GF}(p)$ code or a $\text{GF}(q)$ code. It is still the case, of course, that the set of undetectable errors is $\hat{\mathbf{N}}(S) \setminus \hat{S}$. The distance is again the minimum weight of a Pauli in $\hat{\mathbf{N}}(S) \setminus \hat{S}$. However, when we think of it as a $\text{GF}(q)$ code, we should use weight as defined by the decomposition into q -dimensional qudits (i.e., the number of q -dimensional registers with non-trivial Paulis). If we instead want to think of the code as a $\text{GF}(p)$ code, the weight would be the number of non-trivial Paulis in the decomposition into p -dimensional qudits. The $\text{GF}(p)$ weight of a Pauli could be equal to the $\text{GF}(q)$ weight, but it could also be as high as m times the $\text{GF}(q)$ weight. Also note that the lowest-weight operator in $\hat{\mathbf{N}}(S) \setminus \hat{S}$ using the $\text{GF}(q)$ weight might even be a different operator than the lowest-weight operator using the $\text{GF}(p)$ weight.

Notation 8.8. A stabilizer code with n physical qudits of dimension q , k logical qudits (also of dimension q), and distance d is denoted as an $[[n, k, d]]_q$ code.

Compare the notation $((n, K, d))_q$ for a qudit code that is not necessarily a stabilizer code and $[[n, k, d]]$ for a qubit stabilizer code. This way of listing the properties of a qudit stabilizer code is an obvious hybrid.

8.2.2 Examples: Distance 2 Code, 5-Qudit Code

Let us start to explore the world of qudit stabilizer codes by looking at qudit versions of some specific qubit codes. First, we can derive distance 2 codes. Again, we will choose one generator to detect Z errors on any qudit and one generator to detect X errors on any qudit. For prime qudit dimension p , this leads to a stabilizer of the form

$$M_1 = \bigotimes_{i=1}^n X^{a_i} \quad (8.55)$$

$$M_2 = \bigotimes_{i=1}^n Z^{b_i}. \quad (8.56)$$

We need the generators to commute, so we must pick the exponents so that $\sum a_i b_i = 0$ using $\text{GF}(p)$ arithmetic. For instance, when n is even, it suffices if $a_i = 1$ (for all i) and $b_i = 1$ for half of the values of i and $b_i = -1$ for the other half. Note that it does *not* work to take all powers equal to 1 unless $n = 0 \pmod{p}$. It has distance 2 if all a_i and b_i are non-zero because

$$c(X_i^c Z_i^d, M_1) = da_i \quad (8.57)$$

$$c(X_i^c Z_i^d, M_2) = -cb_i. \quad (8.58)$$

$$\begin{aligned}
M_1 &= X & Z & Z^{-1} & X^{-1} & I \\
M_2 &= I & X & Z & Z^{-1} & X^{-1} \\
M_3 &= X^{-1} & I & X & Z & Z^{-1} \\
M_4 &= Z^{-1} & X^{-1} & I & X & Z
\end{aligned}$$

Table 8.1: The generators for the five-qudit code.

On the other hand, $X_1^{b_2} \otimes X_2^{-b_1}$ commutes with both generators, so we can see that the code is only distance 2. Notice also that, whereas for qubits, the smallest distance 2 code has 4 physical qubits, for larger qudits, it is possible to do it with just 3 qudits. For instance, for $p = 3$ (qutrits), there is the straightforward $[[3, 1, 2]]_3$ code with generators $X \otimes X \otimes X$ and $Z \otimes Z \otimes Z$.

For prime power qudits, the above construction does not give a distance 2 code. For instance,

$$c(Z_i^\delta, M_1) = \text{tr } \delta a_i, \quad (8.59)$$

and for any a_i , there exists a γ for which $\text{tr } \gamma a_i \neq 0$. The problem here is that X on a single dimension- q qudit does not detect Z^δ for all δ , basically because a single qudit is really m separate p -dimensional qudits. No matter what power X^α we choose, we will have the same problem, since there will be some decomposition (not necessarily the standard one) which will lead X^α to act only on a single p -dimensional tensor factor of the full q -dimensional qudit.

Luckily, we can easily fix this by making the minimal modification needed to get a true $\text{GF}(q)$ stabilizer code, repeating the same operators on all the tensor factors of a qudit. Let \mathbf{S} be the smallest stabilizer containing

$$M_1(\gamma) = \bigotimes_{i=1}^n X^{\gamma \alpha_i} \quad (8.60)$$

$$M_2(\gamma) = \bigotimes_{i=1}^n Z^{\gamma \beta_i}, \quad (8.61)$$

for all γ and any particular choice of (α_i, β_i) such that $\sum \alpha_i \beta_i = 0$ in $\text{GF}(q)$. Note that $M_1(1)$ and $M_2(1)$ commute if $\sum \text{tr}(\alpha_i \beta_i) = 0$, but that is not sufficient to make sure that all $M_1(\gamma)$ commute with all $M_2(\gamma)$. With these extra elements added to the stabilizer, the code is now distance 2. For instance, consider again the example Z_i^δ :

$$c(Z_i^\delta, M_1(\gamma)) = \text{tr}(\delta \gamma \alpha_i). \quad (8.62)$$

While this will certainly be 0 for some specific γ , $\text{tr}(\delta \gamma \alpha_i)$ can only be 0 for all γ if $\delta \alpha_i = 0$ in $\text{GF}(q)$.

The operators $M_1(\gamma)$ are not independent for all γ . Since the number of generators for \mathbf{S} is best determined by thinking of it as a $\text{GF}(p)$ code, we should represent each γ as an m -dimensional vector over $\text{GF}(p)$. (It is m -dimensional since $q = p^m$.) It's not hard to see that $M_1(\gamma)M_1(\eta) = M_1(\gamma + \eta)$ and $[M_1(\gamma)]^a = M_1(a\gamma)$ for $a \in \text{GF}(p)$. A set $\{M_1(\gamma)\}$ is thus independent for a subset of possible γ 's if the corresponding $\text{GF}(p)$ vectors are linearly independent, which means that \mathbf{S} has a total of m generators of the form $M_1(\gamma)$. Similarly, it has m generators of the form $M_2(\gamma)$. When we go back to thinking of it as a $\text{GF}(q)$ code, we get an $[[n, n-2, 2]]_q$ code. As a $\text{GF}(p)$ code, this would be a $[[mn, m(n-2), 2]]_p$ code; this is a case where the distance is the same over $\text{GF}(p)$ and $\text{GF}(q)$.

For the next example, I will look at a 5-qudit code which is the qudit generalization of the 5-qubit code. Again beginning with prime dimension, consider the stabilizer given in table 8.1. You can check directly that it is Abelian. It is a bit harder (though still not that hard) to see that it has distance 3, but trust me, it does. Thus, this is a $[[5, 1, 3]]_p$ code. We could also have used the same stabilizer as for the qubit version of the code (table 3.2), but this version has the minor advantage that it is cyclic, just like the 5-qubit code, whereas table 3.2 would give us a non-cyclic $[[5, 1, 3]]_p$ code. While the 5-qubit code is unique up to a tensor product of single-qubit unitary rotations, there are multiple inequivalent 5-qudit codes for qudit dimension

$p \geq 3$. Another difference is that the 5-qubit code is perfect (number of errors equals number of syndromes), whereas the 5-qudit codes are not (more syndromes than single-qudit errors).

Moving to prime power qudits, this 5-qudit code has the same problem as the distance 2 codes — it does not detect or correct errors on the additional tensor factors of a qudit. We can solve it in the same way. Add to the stabilizer all operators of the form $M_i(\gamma)$, with each Pauli raised to the power $\pm\gamma$ instead of ± 1 . For instance, $M_1(\gamma) = X^\gamma \otimes Z^\gamma \otimes Z^{-\gamma} \otimes X^{-\gamma} \otimes I$. The result is a true $\text{GF}(q)$ stabilizer code with parameters $[[5, 1, 3]]_q$.

8.3 Qudit CSS Codes

We can use some of the same methods to make qudit codes that we used for qubit codes. As discussed in section 8.1, we can map the q -dimensional qudit Pauli group (for either prime or prime power dimension) to vectors over $\text{GF}(q^2)$. This lets us interpret qudit stabilizer codes as $\text{GF}(q^2)$ additive codes using the symplectic inner product (8.19) or (8.41) to determine commutation. Note that $\text{GF}(q^2)$ linear codes which are weakly self-dual under the symplectic inner product are equivalent to true $\text{GF}(q)$ stabilizer codes with an additional symmetry, just as linear $\text{GF}(4)$ codes give qubit stabilizer codes with an extra symmetry (see exercise ??).

Alternatively, by writing Paulis in the $\text{GF}(q)$ symplectic form (an n -qudit Pauli written as a $2n$ -component vector over $\text{GF}(q)$), we can use the CSS construction. You'll notice that the example distance 2 codes in the last section had some generators that were all Z 's and some generators that were all X 's, whereas the 5-qudit code had a mix of X and Z in each generator. This is because those distance 2 codes are qudit CSS codes, whereas the 5-qudit code is not.

8.3.1 The CSS Construction for Qudits

For prime dimension, the CSS construction is basically the same as for qubits. Generate a stabilizer code of the form

$$\left(\begin{array}{c|c} 0 & H_1 \\ \hline H_2 & 0 \end{array} \right), \quad (8.63)$$

from the parity check matrices of two classical linear codes C_1 and C_2 . In order to define a stabilizer code, the stabilizer must commute, so use the symplectic inner product (8.16) to test that. Again we find the condition that, if $x \in C_2^\perp$ and $z \in C_1^\perp$, $x \cdot z = 0$, so in order to get a valid stabilizer code, it must be the case that $C_1^\perp \subseteq C_2$. The distance and number of encoded qudits are given by the same formulas as for qubits. The only difference from the qubit case is that the arithmetic to determine commutation is mod p instead of mod 2.

For prime power qudits, we'd like to do the same construction, but we must be a bit more careful. Given two classical linear $\text{GF}(q)$ codes C_1 and C_2 , let's define the stabilizer S to be the smallest group containing all $\bigotimes_j Z^{\gamma_j}$ and $\bigotimes_j X^{\eta_j}$, where γ runs over elements of C_1^\perp and η runs over elements of C_2^\perp . Note that it is *not* sufficient to look at the group generated in this way for basis vectors γ and η of C_1^\perp and C_2^\perp . This is because if $\gamma \in C_1^\perp$, then $\text{GF}(q)$ linearity implies that $\xi\gamma \in C_1^\perp$ as well for any $\xi \in \text{GF}(q)$, but $\bigotimes_j Z^{\gamma_j} \in S$ is not sufficient to imply that $\bigotimes_j Z^{\xi\gamma_j} \in S$. This contrasts with the $\text{GF}(p)$ case, where

$$\bigotimes_j Z^{ab_j} = \left(\bigotimes_j Z^{b_j} \right)^a, \quad (8.64)$$

and since S must be closed under multiplication, it must also be closed under integer exponentiation. Exponentiation by $\xi \in \text{GF}(q)$ doesn't make any sense, and this is responsible for the difference between $\text{GF}(p)$ and $\text{GF}(q)$ stabilizer codes (including CSS codes).

Theorem 8.7. *Let C_1 and C_2 be two classical linear codes over $\text{GF}(q)$ with parameters $[n, k_1, d_1]_q$ and $[n, k_2, d_2]_q$ and satisfying $C_1^\perp \subseteq C_2$. Then there exists a true $\text{GF}(q)$ stabilizer code with stabilizer given as above with parameters $[[n, k_1 + k_2 - n, d]]_q$, $d \geq \min(d_1, d_2)$.*

Proof. The first consideration is whether the stabilizer given above is well-defined. We need to check that $M = \bigotimes_j Z^{\gamma_j}$ and $N = \bigotimes_j X^{\eta_j}$ commute when $\gamma \in C_1^\perp$ and $\eta \in C_2^\perp$. By equation (8.39),

$$c(M, N) = \sum_j \text{tr } \gamma_j \eta_j = \text{tr } \gamma \cdot \eta \quad (8.65)$$

(using the $\text{GF}(q)$ dot product). Since $C_1^\perp \subseteq C_2$, $\gamma \in C_2$, so $\gamma \cdot \eta = 0$, and $c(M, N) = 0$ as desired.

Therefore, we have a well-defined $\text{GF}(q)$ stabilizer code. The elements of the stabilizer have symplectic representations of the form $(\eta|\gamma)$ for $\eta \in C_2^\perp$ and $\gamma \in C_1^\perp$. Since C_1 and C_2 are $\text{GF}(q)$ linear, so are C_1^\perp and C_2^\perp . Thus, the code is a true $\text{GF}(q)$ stabilizer code.

The next question is to determine how many logical qudits there are in this code. Each basis vector of C_1^\perp or C_2^\perp gives us m independent elements of \mathbf{S} ($q = p^m$ as usual) for the reasons noted above (exponentiation by $\xi \in \text{GF}(q)$ does not make sense). Thus, there are $m(n - k_1)$ generators of \mathbf{S} derived from C_1 and $m(n - k_2)$ generators derived from C_2 . Thought of as a $\text{GF}(p)$ code, the number of logical qudits is thus $mn - m(n - k_1) - m(n - k_2) = m(k_1 + k_2 - n)$. Thought of as a $\text{GF}(q)$ code again, we have $k_1 + k_2 - n$ logical qudits.

Finally, the distance can be determined just as for a usual binary CSS code. \square

It might appear at first sight that it is possible to have $\text{GF}(q)$ CSS codes that don't satisfy the condition $C_1^\perp \subseteq C_2$ since we actually only need $\text{tr } \gamma \cdot \eta = 0$ for all γ and η . However, because C_1 and C_2 are *linear* $\text{GF}(q)$ codes, $\text{tr } \gamma \cdot \eta = 0$ for all γ and η iff $\gamma \cdot \eta = 0$ is.

As with binary CSS codes, the basis codewords have a straightforward form when written out in the standard basis:

$$|\gamma + C_2^\perp\rangle = \sum_{\eta \in C_2^\perp} |\gamma + \eta\rangle, \quad (8.66)$$

for $\gamma \in C_1$. This works for both prime dimension and prime power dimension. Indeed, we could just take it as the definition of a CSS code in any dimension.

8.3.2 Polynomial Codes

One particularly interesting family of qudit CSS codes is the family of *polynomial codes*, which are CSS codes derived from classical Reed-Solomon codes or their variants. Let us pick the code C_1 as a Reed-Solomon code by choosing n distinct points $(\alpha_1, \dots, \alpha_n)$ from $\text{GF}(q) \setminus \{0\}$ ($n < q$) and a number $k_1 \leq n$. We'll use polynomials with degree up to $k_1 - 1$. We also pick a second number $0 \leq k_2 \leq k_1$ which will determine the size of C_2 .

Definition 8.9. The basis codewords of the polynomial code over $\text{GF}(q)$ with (n, k_1, k_2) are

$$|\overline{\beta_0, \dots, \beta_{k_2-1}}\rangle = \sum_{(\beta_{k_2}, \dots, \beta_{k_1-1}) \in \text{GF}(q)} \bigotimes_{i=1}^n |\beta_0 + \beta_1 \alpha_i + \beta_2 \alpha_i^2 + \dots + \beta_{k_1-1} \alpha_i^{k_1-1}\rangle. \quad (8.67)$$

The polynomial code is the span of these basis codewords.

That is, the basis codewords are indexed by k_2 values, the lowest-order coefficients of the polynomials being used. We then take the superposition over polynomials for all possible coefficients of x^{k_2} and higher, up to the maximum degree x^{k_1-1} . A particularly interesting special case is when $k_2 = 1$ so there is just one encoded qudit. Note that if $k_2 = 0$, the polynomial code can still be defined, but it is only a single state.

Clearly a polynomial code is a qudit CSS code, since the basis codewords have the correct form. However, instead of choosing both C_1 and C_2 to be Reed-Solomon codes (which wouldn't work, since they are not precisely dual to each other), we have chosen C_2^\perp to be a modified Reed-Solomon code. In particular, C_2^\perp is the code made up of vectors $(f(\alpha_1), \dots, f(\alpha_n))$, where $f(x)$ runs over degree $k_1 - 1$ or lower polynomials for which the lowest nonzero term is the x^{k_2} power or higher. With this choice, it is manifestly true that $C_2^\perp \subseteq C_1$ and that the code encodes k_2 qudits. What is not clear is the distance of the resulting polynomial code.

Theorem 8.8. *The GF(q) polynomial code with parameters (n, k_1, k_2) is a non-degenerate true $[[n, k_2, d]]_q$ stabilizer code with $d = \min(n - k_1 + 1, k_1 - k_2 + 1)$.*

Proof. The code C_1 is used to correct bit flip errors. By theorem 4.15, the distance of C_1 is $n - k_1 + 1$. To prove the formula for distance, we thus need to show that C_2 has distance $k_1 - k_2 + 1$ and that the code is non-degenerate. The fact that it is a true GF(q) code follows from theorem 8.7.

The dual code C_2^\perp has basis vectors $(\alpha_1^j, \dots, \alpha_n^j)$ for $j = (k_2, \dots, k_1 - 1)$. These form the rows of the parity check matrix H_2 . A vector orthogonal to all rows of H_2 (i.e., a vector in C_2) corresponds to a linear dependence among the columns of H_2 , so the distance of C_2 is the minimum number of columns of H_2 that are linearly dependent. There are $k_1 - k_2$ rows, so certainly no more than $k_1 - k_2$ columns can be linearly independent.

Let us look at the matrix formed by any set of $k = k_1 - k_2$ columns, say the first k . The matrix entries are $V_{ij} = \alpha_i^{k_2+j-1}$. This is not a Vandermonde matrix, but is clearly closely related. The columns of the matrix are linearly independent iff the rows are, and we can think of a linear combination of the rows as a polynomial with degree $k_1 - 1$ and all coefficients below degree k_2 being 0; that is, an element of C_2^\perp . The i th entry of the linear combination is the polynomial evaluated at α_i . A linear dependence of the rows is thus a polynomial that evaluates to 0 on all the points $\alpha_1, \dots, \alpha_k$. Because the lowest degree term of the polynomial is x^{k_2} , the polynomial also has a 0 at $x = 0$ with multiplicity k_2 . That gives a total of $k_2 + k = k_1$ zeros for the polynomial, which is too many for a degree $k_1 - 1$ polynomial; thus, for this to be true, the polynomial must be uniformly 0 everywhere (not just on $\alpha_1, \dots, \alpha_k$). In other words, there is no linear dependence of the rows and the matrix V_{ij} is non-singular. Thus, any k columns are independent. Since any $k + 1$ columns are linearly dependent, the code C_2 has distance $k + 1$.

The only remaining thing to show is that the code is non-degenerate, from which it follows that it has distance exactly $\min(n - k_1 + 1, k_1 - k_2 + 1)$ and not a greater distance. We wish to show that there exists some non-trivial logical Pauli that has this weight. I will show that there is a logical \bar{X} with weight exactly $n - k_1 + 1$ and a logical \bar{Z} with weight exactly $k_1 - k_2 + 1$.

For a CSS code (over either qubits or qudits), the logical \bar{X} operators can be taken to be products of physical X s, $\bar{X} = \bigotimes_i X^{\eta_i}$. Moreover, the vectors (η_i) must be in C_1 to commute with the Z stabilizer generators. Since the distance of C_1 is exactly $n - k_1 + 1$, there is a vector of this weight. That is, there is a non-trivial polynomial f of degree $k_1 - 1$ or less such that $\eta_i = f(\alpha_i)$ is 0 for exactly $k_1 - 1$ values of α_i . Thus $\text{wt } \bar{X} = n - k_1 + 1$. But since f has degree $k_1 - 1$, if it has more than $k_1 - 1$ zeros, then it will be uniformly zero. In particular, $f(0) \neq 0$. Now,

$$\bar{X}|\bar{0}\rangle = \sum_{g \in C_2^\perp} \bigotimes_{i=1}^n |(f + g)(\alpha_i)\rangle. \quad (8.68)$$

But the x^0 coefficient of g is 0 (since $g \in C_2^\perp$) and the x^0 coefficient of $f + g$ is *not* zero, so $\bar{X}|\bar{0}\rangle \neq |\bar{0}\rangle$ and $\bar{X} \notin \mathcal{S}$.

Similarly, the logical \bar{Z} operators are products of physical Z s, $\bar{Z} = \bigotimes_i Z^{\gamma_i}$. The vectors (γ_i) must be in C_2 , which means there is such a vector with weight exactly $k_1 - k_2 + 1$. This gives us \bar{Z} with $\text{wt } \bar{Z} = k_1 - k_2 + 1$. Now,

$$\bar{Z}|\bar{f}\rangle = \bar{Z} \sum_{g \in C_2^\perp} \bigotimes_{i=1}^{k+1} |(f + g)(\alpha_i)\rangle \quad (8.69)$$

$$= \sum_{g \in C_2^\perp} \omega^{\sum_i \text{tr } \gamma_i (f+g)(\alpha_i)} \bigotimes_i |(f + g)(\alpha_i)\rangle \quad (8.70)$$

$$= \sum_{g \in C_2^\perp} \omega^{\text{tr } \sum_i \gamma_i f(\alpha_i) + \text{tr } \sum_i \gamma_i g(\alpha_i)} \bigotimes_i |(f + g)(\alpha_i)\rangle. \quad (8.71)$$

Since $(\gamma_i) \in C_2$ and $g \in C_2^\perp$, $\sum_i \gamma_i g(\alpha_i) = 0$.

But I claim there exists some $f \in C_1$ such that $\text{tr} \sum_i \gamma_i f(\alpha_i) \neq 0$. We can assume without loss of generality that γ_i is non-zero only for $i = 1, \dots, k_1 - k_2 + 1$. Consider polynomials $f_j(x) = x^j$ for $j = 0, \dots, k_1 - k_2$. The matrix $V_{ij} = \alpha_i^j = g_j(\alpha_i)$ is a Vandermonde matrix, so the vectors $(f_j(\alpha_1), \dots, f_j(\alpha_{k_1 - k_2 + 1}))$ are linearly independent. They are vectors in a $(k_1 - k_2 + 1)$ -dimensional vector space, so there is no non-zero vector that is orthogonal to all of them. In particular, the vector (γ_i) must have non-trivial overlap with at least one of the vectors $(f_j(\alpha_i))$: $\sum_i \gamma_i f_j(\alpha_i) = \xi \neq 0$. It is possible that $\text{tr} \xi = 0$, but if we instead use the polynomial $f'(x) = \eta f_j(x)$, then $\sum_i \gamma_i f_j(\alpha_i) = \xi \eta$. η is arbitrary, so we just need to pick some η such that $\text{tr}(\xi \eta) \neq 0$ to prove the claim.

Consequently, $\bar{Z}|\bar{0}\rangle = |\bar{0}\rangle$ but $\bar{Z}|\bar{f}'\rangle \neq |\bar{f}'\rangle$. Thus, \bar{Z} , which has weight $k_1 - k_2 + 1$, is a non-trivial logical operation. This proves the polynomial code is non-degenerate and thus the distance is exactly $\min(n - k_1 + 1, k_1 - k_2 + 1)$. \square

If we choose the distances of the bit flip code and the phase code to be equal, then $n - k_1 = k_1 - k_2$, or $2k_1 = n - k_2$. In this case, $2d = n - k_1 + 1 + k_1 - k_2 + 1 = n - k_2 + 2$. Since k_2 is the number of encoded qudits, a polynomial code with these parameters saturates the quantum Singleton bound and is a quantum MDS code. For instance, we can have a $[[5, 1, 3]]_q$ polynomial code for $q > 5$. (There is also a variant with $q = 5$.) Note that this code is not equivalent to the $[[5, 1, 3]]_q$ stabilizer in table 8.1. This code is a CSS code, unlike the previous one, but it does not work for $q = 2, 3$, or 4 , whereas the code of table 8.1 works for any prime or prime power dimension.

8.4 Qudit Clifford Group

Although I've spent a number of pages discussing qudit stabilizer codes, ultimately they are not that different than qubit stabilizer codes. They're not identical, to be sure, which is why I've gone through them in detail, but the differences are small. When we get to the qudit Clifford group, however, larger differences start to appear. The same general principles hold as for the qubit case, but there is a significant divergence at the level of details. I will focus on the case of prime dimension; for prime powers, simply consider the $q = p^m$ -dimensional qudit as m p -dimensional qudits.

We start the same way as for qubits:

Definition 8.10. Let p be a prime. The *qudit Clifford group* is

$$C_n(p) = \{U \in \mathbf{U}(p^n) | UPU^\dagger \in \mathbf{P}_n(p) \ \forall P \in \mathbf{P}_n(p)\}. \quad (8.72)$$

As with qubits, $\mathbf{P}_n(p)$ is a normal subgroup of $C_n(p)$. We can also define

$$\hat{C}_n(p) = C_n(p) / \{e^{i\phi} I\} \quad (8.73)$$

$$\check{C}_n(p) = \hat{C}_n(p) / \hat{\mathbf{P}}_n(p). \quad (8.74)$$

It is still the case that $\check{C}_n(p) \cong \mathbf{Sp}(2n, \mathbb{Z}_p)$ and that there exists a unitary in $C_n(p)$ that performs any transformation of the Paulis that preserves commutation relations. It is also possible to efficiently simulate the behavior of the qudit Clifford group, including Pauli measurements, with a classical computer using essentially procedure 6.4. The main difference in the simulation is when measuring P_i with $P_i \notin \mathbf{N}(\mathbf{T})$. The measurement outcome is a uniform random value b from 0 to $p - 1$ rather than 0 or 1 . We still find a generator M of the stabilizer to replace with $\omega^b P_i$, but we choose M that has a factor of ω when commuting past P_i rather than anticommuting with it. By taking the appropriate power r of M , we can ensure that NM^r commutes with P for stabilizer generators or logical operators N ; we replace the original N with the appropriate NM^r .

The biggest difference with the qubit Pauli groups comes when comparing the actual elements of the Clifford group. Some of the elements are very closely analogous. For instance, the discrete Fourier transform over \mathbb{Z}_p is the generalization of the Hadamard transform (which is the discrete Fourier transform over \mathbb{Z}_2):

$$\mathcal{F}|a\rangle = \frac{1}{\sqrt{p}} \sum_b \omega^{ab} |b\rangle. \quad (8.75)$$

Working out the conjugation action on the Paulis, we find that Fourier does not precisely swap X and Z , but does so adding an inverse:

$$X \mapsto Z \tag{8.76}$$

$$Z \mapsto X^{-1}. \tag{8.77}$$

There is also a two-qubit SUM gate which is the direct analogue of the qubit CNOT gate:

$$\text{SUM}|a\rangle|b\rangle = |a\rangle|a + b \bmod p\rangle \tag{8.78}$$

$$X \otimes I \mapsto X \otimes X \tag{8.79}$$

$$Z \otimes I \mapsto Z \otimes I \tag{8.80}$$

$$I \otimes X \mapsto I \otimes X \tag{8.81}$$

$$I \otimes Z \mapsto Z^{-1} \otimes Z. \tag{8.82}$$

Again there is an extra inverse in one of the images in the conjugation action. The inverse powers are needed to ensure that the Clifford group gate preserves commutation relations among the Paulis. For instance, $Z \otimes Z$ does not commute with $X \otimes X$, but $Z^{-1} \otimes Z$ does.

From here, though, we start to see a bigger divergence. There are operators in $C_n(p)$ which don't have any qubit analog, such as scalar multiplication by $c \in \mathbb{Z}_p \setminus \{0\}$:

$$S_c|a\rangle = |ca\rangle \tag{8.83}$$

$$X \mapsto X^c \tag{8.84}$$

$$Z \mapsto Z^{c^{-1}}. \tag{8.85}$$

The power of Z is c^{-1} , the multiplicative inverse of c in \mathbb{Z}_p . For instance, for $p = 7$, $2^{-1} = 4$ since $2 \cdot 4 = 8 \equiv 1 \pmod{7}$. There is also no precise qudit analog of $R_{\pi/4}$. The closest is a quadratic phase gate which has a similar action on Paulis:

$$B|a\rangle = \omega^{a(a-1)/2}|a\rangle \tag{8.86}$$

$$X \mapsto XZ \tag{8.87}$$

$$Z \mapsto Z. \tag{8.88}$$

Theorem 8.9. *The gates \mathcal{F} , B , and SUM along with global phase $e^{i\theta}I$ generate $C_n(p)$.*

Proof. First, we can show that \mathcal{F} , B , SUM, S_c (for all $c \neq 0$), $P_n(p)$ and $e^{i\theta}I$ generate $C_n(p)$ using essentially procedure 6.5. In the algorithm, we can replace H with \mathcal{F} , CNOT with SUM, and $R_{\pi/4}$ with B . SWAP can be realized as

$$\text{SWAP}_{i,j} = S_{-1_i} \text{SUM}_{i,j} \text{SUM}_{j,i}^{-1} \text{SUM}_{i,j}. \tag{8.89}$$

The analog of $C - Z$ is $(I \otimes \mathcal{F})\text{SUM}(I \otimes \mathcal{F}^{-1})$. The symplectic matrices of these gates are quite similar to the qubit versions, but there are some additional minus signs that appear (when the power is an inverse). In the procedure, we will have to use some of them multiple times in order to cancel out terms. For instance, in step 2, we wish to eliminate entries in the first column of A . To do so, use left multiplication by SUM $p - a_{1j}$ to eliminate entry a_{1j} . We also add in S_c , which has the effect of multiplying rows or columns by c or c^{-1} . This is useful in step 1: we can use SWAP to move a non-zero entry into the upper left corner, but we then need S_c to make that entry 1.

The next step is to show that we don't need S_c . I claim that gates of this form can be generated using only \mathcal{F} , B , and $P_1(p)$. Let $Q = \mathcal{F}B\mathcal{F}^{-1}$. Then

$$Q : X \mapsto X \tag{8.90}$$

$$Z \mapsto ZX^{-1}. \tag{8.91}$$

The action of $B^r Q^s B^m Q^n$ is then

$$X \mapsto \omega^a X^{1-ms} Z^{m+r-rms} \quad (8.92)$$

$$Z \mapsto \omega^b X^{-n-s+mns} Z^{1-(m+r)n+(mn-1)rs} \quad (8.93)$$

The powers a and b don't matter because we can get rid of the powers of ω using an appropriate element of $P_1(p)$. Let $r = -c^{-1}$, $s = 1 - c$, $m = 1$, and $n = 1 - c^{-1}$, so $rs = rms = n = m + r$. Then

$$X \mapsto X^c \quad (8.94)$$

$$Z \mapsto Z^{c^{-1}}. \quad (8.95)$$

We also don't need the Paulis. $\mathcal{F}^2 = S_{-1}$, i.e. $\mathcal{F}^2|a\rangle = |-a\rangle$. Then

$$\mathcal{F}^2 B \mathcal{F}^2 B^{-1} |a\rangle = \mathcal{F}^2 B \omega^{-a(a-1)/2} |-a\rangle = \omega^{-a(-a-1)/2 - a(a-1)/2} |a\rangle = \omega^a |a\rangle. \quad (8.96)$$

Thus, $Z = \mathcal{F}^2 B \mathcal{F}^2 B^{-1}$, and $X = \mathcal{F}^{-1} Z \mathcal{F}$, which then can be used to generate the full Pauli group. \square