Energetics and Error Rates of Self-Correcting Quantum Memories

John Schulman

Quantum codes allow for the robust storage of quantum information despite interaction with the environment. In a quantum code, a large number of physical qubits is used to encode a smaller number of qubits. We may impose interactions between the physical qubits that make up the code so that the code space is the degenerate ground state of the resulting Hamiltonian. Then if the system is coupled to a low-temperature environment, thermodynamic relaxation will suppress errors that take the system out of the code space.

We discuss what interactions this system could have if it is based on a stabilizer code or a stabilizer subsystem code. Then we provide bounds that relate the error rate to the energy spectrum of the Hamiltonian. We consider codes on lattices and ask whether the error rate can be arbitrarily reduced by scaling up the size of the system. This reduction of error rate with scaling indeed occurs if both the size of the stabilizer group, and the number of syndromes in which local errors can't damage the encoded data, scale linearly with the system size.

Then we address the question of whether such a self-correcting system, which becomes more robust as its size is increased, is possible in two or three dimensions. We argue that it is impossible in two dimensions. We present a code in three dimensions, analyze some of its properties, and argue that it might be a self-correcting system.

I. INTRODUCTION

Quantum information is inherently more difficult to store and manipulate than classical information because of decoherence [1] [2]. Decoherence is the process of entanglement with the environment, and it is thermodynamically irreversible, since we can't access all degrees of freedom of the environment. Quantum codes are able to protect the encoded information from decoherence (and other types of errors) by hiding the information in non-local observables [3] [4]. These codes may consist of a set of physical qubits encoding a smaller number of logical qubits. The density matrix of a small subset of the physical qubits is independent of the data encoded in the logical qubits, therefore errors acting on a small number of qubits can't damage the data.

A recovery process is needed to purify the state. One approach is active error recovery. In this type of scheme, some operators are measured to diagnose what type of error occurred—the syndrome—and then an operator is applied to return the system to its error-free state. The active recovery process itself introduces noise into the system. However, when the error rate is below a certain threshold we can reliably store and manipulate the encoded quantum information.

A fundamentally different approach to error recovery is a thermodynamic one [5] [7] [10]. Here the qubits that make up the code are coupled by physical interactions described by a Hamiltonian. The encoded logical operators are a non-commuting group of symmetries of the system, so every energy level is degenerate. Errors have an energy cost, and only errors acting on a large number of qubits can damage the encoded data. If the system is coupled to an environment with low temperature, the equilibrium distribution of states will favor states with energies near the ground state energy; these states ought to be related the code space by a small number of local errors so that the encoded data is unaffected.

Section II reviews stabilizer error correction and its generalization, stabilizer subsystem error correction. Section III briefly discusses the idea of self-correcting codes. Section IV gives some conditions that ensure robust storage of quantum information. In particular, we want to know if we can reduce the rate of decoherence by increasing the system size. Section V addresses the problem of creating a physical system whose self-correcting properties get better with increasing system size.

II. STABILIZER SUBSYSTEM ERROR CORRECTION

A. Stabilizer codes

The code space $\mathcal{C}$, the subspace that the state is in when no errors have occurred, is a particular mutual eigenspace of a set of commuting Pauli operators called stabilizers. We can decompose the Hilbert space as

$$\mathcal{H} = \bigoplus_{s_1 s_2 \cdots s_t} \mathcal{L}_{s_1 s_2 \cdots s_t} \tag{1}$$

where the stabilizer S_i has eigenvalue s_i in $\mathcal{L}_{s_1 s_2 \cdots s_t}$. We measure the stabilizer operators this projects the state into a subspace $\mathcal{L}_{s_1 s_2 \cdots s_t}$. This is all the information we can gain about the state without damaging the encoded data. We associate each subspace $\mathcal{L}_{s_1 s_2 \cdots s_t}$ other than the code space $\mathcal{C}$ with a particular Pauli error: the smallest-weight error E that maps $\mathcal{C} \to \mathcal{L}_{s_1 s_2 \cdots s_t}$. (A Pauli operator is the product of single-qubit Pauli operators I, X, Y, Z acting on the physical qubits of the code; its weight is the number of non-identity operators in this product.) When we measure the stabilizers and the resulting eigenvalues are $s_1, s_2, \cdots s_t$, called the syndrome. We can return the state to subspace $\mathcal{C}$ by applying the corresponding operator E. We can use the distance d as a measure of the codes robustness: d is the smallest weight of an error that maps $\mathcal{C} \to \mathcal{C}$. It follows that no two errors with weight $< d/2$ have the same syndrome, so we can correct these errors.

B. Stabilizer subsystem error correction

In stabilizer subsystem error correction, we allow "gauge" degrees of freedom that are independent of the encoded data [12] [13]. Specifically, there is a group $\mathbf{G}$ of *gauge operators*: Pauli operators that commute with both the logical operators $\mathbf{L}$ and stabilizers $\mathbf{S}$ but do not all commute with each other. Here the Hilbert space decomposes as

$$\mathcal{H} = \bigoplus_{s_1 s_2 \cdots s_t} \mathcal{L}_{s_1 s_2 \cdots s_t} \otimes \mathcal{G}_{s_1 s_2 \cdots s_t} \tag{2}$$

In a given subspace $\mathcal{L}_{s_1 s_2 \cdots s_t} \otimes \mathcal{G}_{s_1 s_2 \cdots s_t}$ the logical operators $\mathbf{L}$ act only on $\mathcal{L}$ and the gauge operators $\mathbf{G}$ act only on $\mathcal{G}$. The code space $\mathcal{C}$ is in one of these subspaces $\mathcal{C} \otimes \mathcal{G}_0$, which we call the null-syndrome subspace. (For a stabilizer code, this $\mathcal{G}$ is trivial and the null-syndrome subspace is simply the code space $\mathcal{C}$. By measuring the syndrome, we project the state into one of these subspaces. In this case, we can have multiple errors $E_1, E_2, \cdots, E_m$ associated with each subspace, as long as $E_i E_j \in \mathcal{G}$. Here, the distance of the code as defined above is not a good measure of its strength since the gauge degrees of freedom are uncontrolled; a better measure of the strength of the code is the minimal weight in $\mathbf{LG} \backslash \mathbf{S}$.

In the appendix, it is shown how to obtain a subsystem stabilizer code from a stabilizer code by sequentially reducing the size of the stabilizer group.

III. WHAT IS A SELF-CORRECTING QUANTUM MEMORY?

As briefly described in the introduction, one approach to reducing decoherence of encoded quantum information is to choose an appropriate Hamiltonian that surpresses excitations that take the system outside of the code space. The Hamiltonian should be a sum of terms that commute with the logical operators; this way, the logical operators are symmetries of the system so the encoded data does not evolve over time. One way to accomplish this is to let the Hamiltonian be a sum of operators from the stabilizer group, i.e.

$$H = \sum_{S \in \mathbf{S}' \subset \mathbf{S}} -S \tag{3}$$

Kitaev's toric code [5] and its generalizations given in [10] are well suited to this approach. In the toric code, the stabilizers act on four sites, which are edges in a lattice; these stabilizers are in correspondence with the vertices and plaquettes of the lattice. If a system is engineered so that these four-site interactions result, the system will have a gap for excitations out of the degenerate ground state, thus errors that damage the data are suppressed.

The sites of our physical qubits have some spatial arrangement, such as a two- or three-dimensional grid. Physical considerations restrict us to two-body interactions between nearby sites.

Subsystem error-correction may be very useful for creating a physically realizable self-correcting system. When these codes are used along with an active error recovery scheme, one can determine the syndrome by measuring a set of small-weight gauge operators, rather than measuring the stabilizer operators. The value of each stabilizer operator is then determined as the product of several gauge operators. For the approach of physical self-correction, we can take our Hamiltonian to be a sum of operators in the gauge group,

$$H = \sum_{g \in \mathbf{G}' \subset \mathbf{G}} -g \tag{4}$$

thus the Hamiltonian commutes with the logical operators. We expect that for $g \in \mathbf{G}'$, the expectation in the ground state $< g >_0 \approx 1$, so it is plausible that for a stabilizer operator S which is a product of some gauge operators, the ground state will have $S = 1$, thus the the ground state will be the null-syndrome state. This must be shown mathematically, though, and determining the spectrum of these Hamiltonians is an open problem.

IV. BETTER SELF-CORRECTION WITH SCALING

Suppose a family of subsystem codes are defined on lattices of linear size L (so the number of sites $\propto L^D$, where D is the dimensionality.) The toric code [5], Bacon's code [7], and the compass-model code in section V are examples of families of codes on lattices. Suppose, moreover, that we have defined a Hamiltonian on the lattice for each L. We can ask the following question of such a family: if the noise level is low, can we arbitrarily reduce the rate of decoherence of the encoded data by increasing the system size? We will make this question more precise, and then write down a condition on the code and the Hamiltonian, which will ensure that the encoded data is preserved.

First of all, let us define what is meant by the encoded data in a system that has been acted on by errors. This discussion will refer to an active recovery process; it is just used for mathematical purposes–presumably in the systems under discussion, thermodynamics eliminates the need for active error recovery. Note that in a quantum code, a error that acts at a single site will typically anti-commute with some logical operator, thus changing its value. However, the encoded data is not damaged by such an error because we can perform a recovery procedure; we measure the stabilizer operators and remove this error by applying the smallest Pauli operator that causes the syndrome measured. This ideal recovery operator can be written as $\mathcal{R} : \rho \to \sum_i E_i P_i \rho P_i E_i$. Here the syndromes are labeled by the index i and P_i and E_i are the projectors and minimal-weight Pauli operators that correspond to syndrome i, respectively. The encoded data in a state with density matrix ρ is then defined as $P_{\mathcal{C}}(\mathcal{R}(\rho))$, where $P_{\mathcal{C}}$ is the projection $\mathcal{C} \otimes \mathcal{G}_0 \to \mathcal{C}$, where $\mathcal{C} \otimes \mathcal{G}_0$ is the null-syndrome subspace, and C and $\mathcal{G}_0$ are the components acted on by the logical and gauge operators, respectively.

If an error E is applied to the system, $\rho \to E\rho E$ and the syndrome changes. When E is a single qubit error, only certain transitions between syndromes are possible. Suppose that initially the system is in subspace i which has minimal weight error E_i, and an error E sends it to subspace j with minimal weight error E_j. Usually $E_j = EE_i$ for the single qubit operator E, but it is also possible that $E_j = LEE_i$, where L is an encoded logical operator. In this case the single qubit error followed by the recovery operator has applied a nontrivial operation to the encoded data. The pair (i, E) is called a *bad transition*, since applying the single-qubit operator E to a state with syndrome i changes the encoded data, which is defined using recovery operator $\mathcal{R}$.

Define $\mathcal{D}$ (for dangerous) to be the direct sum of all syndrome subspaces that have a bad transition. The direct sum of all syndromes that have no bad transitions is $\mathcal{D}^\perp$.

Any error that sends the code space to $\mathcal{D}$ has weight $\geq \frac{d-1}{2}$

To see this, suppose that (i, E) is a bad transition, and the new syndrome after applying E is j. Let E_i and E_j be the minimal weight errors with syndromes i and j, respectively. Then $E_j = LEE_i$, so $E_j E_i E = L$. weight$(L) \geq d$, weight$(E_j) \leq \frac{d-1}{2}$, and weight$(E) = 1$, thus weight$(E_i) \geq \frac{d-1}{2}$.

If our noise source applies errors with rate Λ, the rate $\bar{\Lambda}$ that errors are applied to the encoded data satisfies

$$\bar{\Lambda} < \Lambda \Pr(\mathcal{D}) \tag{5}$$

where $\Pr(\mathcal{D})$ is the probability that the system is in $\mathcal{D}$, assuming that the off-diagonal terms in ρ between $\mathcal{D}$ and $\mathcal{D}^\perp$ are small and can be ignored.

Let $\Delta_{\mathcal{D}}$ be the minimal energy in $\mathcal{D}$ minus the ground state energy. As before, let $\mathbf{S}$ be the stabilizer group. The number of syndrome subspaces is then $2^{\dim \mathbf{S}}$. Then we have the following bound on the error rate:

$$\bar{\Lambda}/\Lambda < e^{-\beta \Delta_{\mathcal{D}}} \cdot 2^{\dim \mathbf{S}} \tag{6}$$

Therefore if $\Delta_{\mathcal{D}} \propto L$ and $\dim \mathbf{S} \propto L$, where L is the linear dimension of the lattice, there is a temperature β for which the error rate $\bar{\Lambda} \to 0$ as $L \to \infty$. In particular, $\bar{\Lambda} \propto e^{-cL}$ for some constant c.

Note that in Bacon's subsystem code and the compass-model code given in section V, we do indeed have $\dim \mathbf{S} \propto L$. However, this is not the case in the surface codes in [5, 10], there $\dim \mathbf{S} \propto L^D$, where D is the dimension of the lattice. Note that in active error correction, a code can correct approximatly $\dim \mathbf{S}$ single-qubit errors, since this is how much the entropy is reduced through syndrome measurement and then recovery. Thus we expect that a self-correcting system can accomodate approximately $\dim \mathbf{S}$ errors; if the equilibrium distribution of states contains more errors, then decoherence is rapid. The above condition, $\Delta_{\mathcal{D}} \propto L$, ensures that the equilibrium number of errors is $\propto L$, thus if our code has $\dim \mathbf{S} \propto L$, we can reduce decoherence by scaling up the system size.

V. IMPLEMENTING SELF-CORRECTION IN PHYSICAL SYSTEMS

The previous section motivates us to look for a quantum code defined on a two- or three-dimensional lattice, along with a Hamiltonian made up of local interactions, such that high-weight errors have a high energy cost and are

thus suppressed. We will refer to a such a system, in which the rate of decoherence can be arbitrarily reduced by increasing the system size, as a self-correcting quantum memory. There is a self correcting quantum memory in four dimensions–it is given in [10]. It is a four-dimensional analog of a surface code. The logical operators are planes, so the excitations are like sheets, and they have an energy cost proportional to their boundary.

A. Difficulty of self-correction in two dimensions

The toric code and other two-dimensional surface codes are not self-correcting quantum memories. As a matter of fact, any code with a line-like logical operator is not self-correcting. Imagine applying a string of single-qubit operators that forms a piece of the logical operator. The energy cost of this error will only be due to terms in the Hamiltonian at the ends of the string, since the middle of the string is indistinguishable from the logical operator. The the energy penalty of such an error is thus bounded above by a constant that is independent of the size of the system. Thus increasing the system size will not reduce the chance that a local excitation will turn into a logical error.

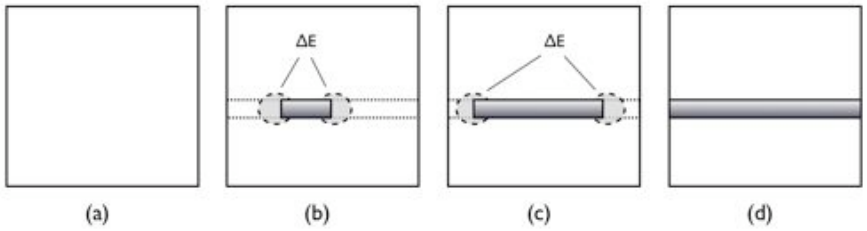

Figure 1: Energy penalty is the same for a small piece of the logical operator (b) and a large piece (c). Thus the probability that a small error will lead to a logical error (d) is not significantly decreased by increasing the system size.

The fact that a self-correcting quantum memory in two dimensions may be impossible is reminiscent of the Mermin-Wagner theorem [9]. This theorem states that under most conditions, a continuous symmetry can't be broken in two spacial dimensions. The operators $e^{i\bar{X}dt}$ are continuous symmetries of the system. If we build a very large version of our self-correcting code and cool it to below the critical temperature, this symmetry is effectively broken. If the system is prepared in a certain logical state, it will stay in this logical state.

B. 3D Compass model as a candidate self-correcting code

We will use the code below to demonstrate the method of creating a subsystem error-correcting code out of a Hamiltonian. This code is very similar to the code Bacon uses in [7]. This code is on an $L \times L \times L$ lattice of spins on a three-dimensional cube or torus. Its Hamiltonian is given by

$$H = -\sum_{p \in C} X_p X_{p+\hat{x}} + Y_p Y_{p+\hat{y}} + Z_p Z_{p+\hat{z}}. \tag{7}$$

Here, p is a vertex in the lattice, and $\hat{x}$ is the unit vector in the x direction. This Hamiltonian is known to some condensed-matter physicists as the orbital compass model.

The group of logical operators $\mathbf{L}$, is generated by planes of X perpendicular to $\hat{x}$, planes of Y perpendicular to $\hat{y}$, and planes of Z perpendicular to $\hat{z}$. This subgroup has dimension $3L$ (when considered as a vector space.) The stabilizer group $\mathbf{S}$ is generated by pairs of adjacent planes, as well as an operator containing a plane of X times a plane of Y times a plane of Z, normal to $\hat{x}, \hat{y}, \hat{z}$ respectively. This subgroup has dimension $3(L-1)+1 = 3L-2$. The quotient group $\mathbf{L}/\mathbf{S}$ has dimension $3L-(3L-2) = 2$, thus it is isomorphic to the one-qubit Pauli group, and has operators $\{\bar{I}, \bar{X}, \bar{Y}, \bar{Z}\}$ (encoded X, Y, Z), which can be implemented as single planes of X, Y, and Z with appropriate orientations.

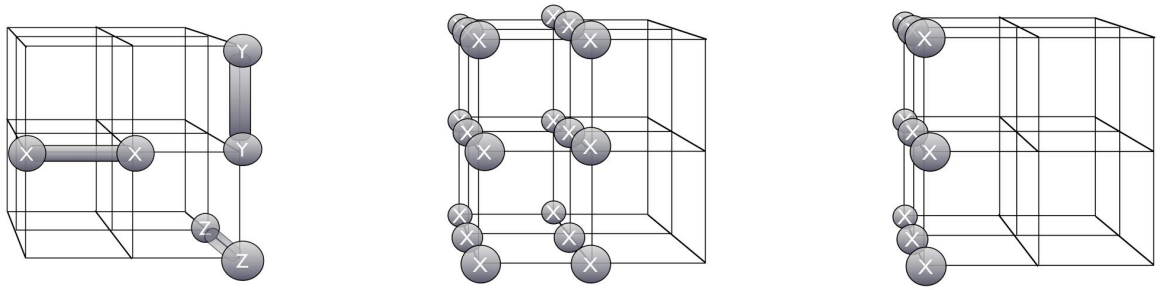

(a) Gauge operators of XX, YY, and ZZ type. The Hamiltonian is a sum of such terms.

(b) A stabilizer operator

(c) A logical operator

Figure 2: Representatitve operators corresponding to compass model code on $3 \times 3 \times 3$ lattice

This code's logical operators are gauge-equivalent to smaller Pauli operators. It turns out that the code has distance L. First we will show that the distance is at least L. Consider the logical operator $\bar{Z}$. We claim that any Pauli operator gauge-equivalent to $\bar{Z}$ has a non-identity operator in every xz plane and every xy plane. To see this, consider a given xz plane. Let χ be the product of the operators at every site of this plane. $\chi = Z$ for a logical operator $\bar{Z}$ and χ is I, X, or Y for any gauge operator. χ is multiplicative, so we have $\chi(\bar{Z}g) = \chi(Z)\chi(g) \neq I$. There are L xz planes, so $\bar{Z}$ has weight at least L.

Next we will show that the distance is precisely L, i.e. logical operator $\bar{Z}$ is gauge-equivalent to a certain Pauli operator of weight L. First note that a square made up of four adjacent Z operators lying in an xy plane is in the gauge group (shown in figure 3.)

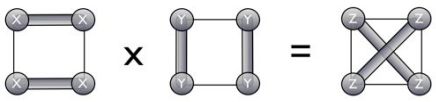

Figure 3: Product of four Z operators in an xy plane

We can take a product of several of these operators with $\bar{Z}$ to get an operator of weight L, which is shown in figure 4 for $L = 5$.

Any Pauli error E applied to the ground state has an energy cost $\langle 0|EHE|0\rangle = \langle 0|H|0\rangle - \langle 0|A(E)|0\rangle$, where $A(E)$ is the sum of the terms (e.g. $-X_p X_{p+\hat{x}}$) in the Hamiltonian that anticommute with E. Let α denote the number of terms in A, i.e. the number of terms in H that anticommute with E. Clearly α is at least the number of xy planes containing a non-trivial Pauli operator. We also have $\langle 0|EHE|0\rangle - \langle 0|H|0\rangle \geq \alpha \min \eta$, where η is the minimum expectation of XX, YY, and ZZ couplings in the Hamiltonian (e.g. $< X_p X_{p+\hat{x}} >$.) However, $E|0\rangle$ is not an eigenstate, so bounds on the energy cost of Pauli operators don't tell us about the spectrum of H; these arguments merely show that it is plausible that this system has self-correcting properties.

VI. CONCLUSION

This paper has identified how decoherence works in a self-correcting code: logical errors are caused by bad transitions between syndrome subspaces. In good codes, these bad transitions only occur between syndromes caused by a high number of errors. Thus in a self-correcting code, the energetics of the system must be such that these high-error subspaces are exponentially surpressed and rarely occupied at low temperatures. Thus we can bound the rate of decoherence if there is an energy gap between the code space and the syndrome subspaces that can undergo these bad transitions.

Then we investigated the possibility of building a self-correcting quantum memory based on local interactions in two or three dimensions. We argue that it is likely impossible in two dimensions. We present a subsystem code and corresponding Hamiltonian in three dimensions, and we argue that it is plausible that this system has self-correcting properties.

Self-correcting quantum codes might have applications in an architecture for quantum computing, in an adiabatic or circuit-based scheme, if there is a good local way for creating interactions between two encoded qubits. Their study

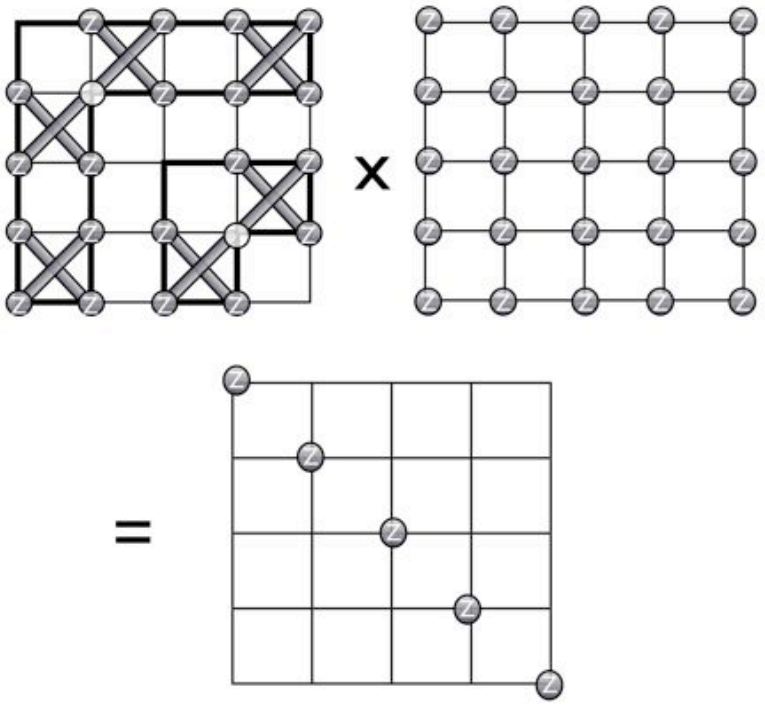

Figure 4: A weight-L operator that is gauge-equivalent to $\bar{Z}$

also hints at certain theoretical questions related to symmetry breaking in various numbers of spatial dimensions. A physical system can store a bit in two dimensions by the breaking of a discrete symmetry, one can store an analog variable in three dimensions by the breaking of a continuous symmetry (e.g. in the Heisenberg model); how many dimensions do we need to store a bit of quantum information through symmetry breaking?

Appendix: Gauge-reduction of stabilizer codes

Bacon found that it is possible to modify the Shor code to get a code with two-qubit gauge operators [7]. The stabilizers are products of these two-qubit gauge operators. The disadvantage of this procedure is that the new code doesn't correct as many errors with distance $> d/2$. However, it has major advantages for fault tolerance, both for active error correction and thermodynamic error correction. In the case of active error correction, this eliminates the need for large circuits or complicated ancilla preparations to measure the large stabilizer operators. Instead, we can measure all of the two-qubit gauge operators to determine the syndrome. Aliferis and Cross [14] rigorously proved a threshold for universal quantum computation of 1.94×10^{-4} errors per gate, an order of magnitude better than proven thresholds for the 7-qubit Steane code. In the case of thermodynamic error correction, this can make the stabilizer operators a product of local gauge operators. A Hamiltonian made from the sum of these gauge operators may give a self-correcting code.

Here we give a necessary and sufficient condition for a stabilizer code to be modified this way, which will henceforth be called a gauge reduction. A gauge reduction from stabilizer code C to C' is defined so the following conditions hold:

1. The distance of C' is the distance of C.
2. The stabilizer group $\mathcal{S}'$ of C' is a proper subgroup of the stabilizer group $\mathcal{S}$ of C.

$\mathcal{S}$ is isomorphic to the vector space F_2^s, where s is the number of stabilizer generators, and $\mathcal{S}'$ is a subspace. Then $\mathcal{S}'^{\perp}$ is also a subspace. We can choose a set of basis vectors $S_1, S_2, \cdots S_{s-r}$ that spans $\mathcal{S}'$, and $S_{s-r+1}, S_{s-r+2}, \cdots, S_s \in \mathcal{S}'^{\perp}$ so that together these bases together span $\mathcal{S}$. So $\mathcal{S}'$ can be obtained by starting with a set of generators for $\mathcal{S}$ and removing r of them.

$\mathcal{S}'$ is a subgroup of $\mathcal{S}$ if and only if we can extend a set of generators of $\mathcal{S}'$ to a set of generators of $\mathcal{S}$. In other words, $\mathcal{S}'$ can be obtained by removing some generators from $\mathcal{S}$. Under what conditions can we remove a stabilizer generator S_1? The remaining stabilizer operators must be able to detect all errors that damage the data, i.e. all errors that anticommute with a logical operator.

The following condition allows us to remove stabilizer generator S_1:

> For every error g with weight $< d$ such that g commutes with every stabilizer except S_1, g commutes with the logical operators.

This is necessary, since if it did not hold, then the new code would not detect the error g that damaged the encoded data (by anticommuting with a logical operator.) This is sufficient since the new code can detect every error with weight $< d$ that the old code can, thus it has weight d. The new code will have S_1 and g in its gauge group.

We can state the condition for gauge reduction in a different way as follows:

> There exists a syndrome such that all errors of weight less than d that cause this syndrome commute with the logical operators.

This condition is necessary for the reason given above. To see why this is sufficient, suppose there is a syndrome Θ satisfying this condition in which the stabilizers $S_1, S_2, \cdots, S_m$ have value -1. Then consider stabilizers $S_1, S_1S_2, S_1S_3, \cdots, S_1S_m$, which generate the same stabilizer group as $S_1, S_2, \cdots, S_m$. Now every error causing syndrome Θ anticommutes with $S1$ but commutes with the other stabilizer generators. Thus the previous condition is satisfied and we can remove stabilizer S_1.

Let us illustrate how a gauge reduction works with Shor's nine-qubit code. The original code is on the left of Table I. $X\mathbb{1}\mathbb{1}X\mathbb{1}\mathbb{1}\mathbb{1}\mathbb{1}\mathbb{1}\mathbb{1}$ anticommutes with S_1 and S_2 but commutes with all other stabilizer generators. Clearly no other operator with weight < 3 causes this syndrome. Now we follow the construction in the proof above: S_1, S_1S_2 generates same subgroup as S_1, S_2, and we can remove S_1. The new code is shown on the right. By iterating this procedure, we can eventually obtain Bacon's $[[9, 1, 4, 3]]$ code.

S_1	Z	Z	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$
S_2	$\mathbb{1}$	Z	Z	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$
S_3	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	Z	Z	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$
S_4	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	Z	Z	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$
S_5	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	Z	Z	$\mathbb{1}$
S_6	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	Z	Z
S_7	X	X	X	X	X	X	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$
S_8	X	X	X	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	X	X	X
$\overline{Z}$	Z	Z	Z	Z	Z	Z	Z	Z	Z
$\overline{X}$	X	X	X	X	X	X	X	X	X

S_1	Z	Z	$\mathbb{1}$	Z	Z	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$
S_2	$\mathbb{1}$	Z	Z	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$
S_3	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	Z	Z	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$
S_4	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	Z	Z	$\mathbb{1}$
S_5	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	Z	Z
S_6	X	X	X	X	X	X	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$
S_7	X	X	X	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	X	X	X
$\overline{Z}$	Z	Z	Z	Z	Z	Z	Z	Z	Z
$\overline{X}$	X	X	X	X	X	X	X	X	X
g_1^z	Z	Z	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$
g_1^x	X	$\mathbb{1}$	$\mathbb{1}$	X	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$	$\mathbb{1}$

Table I: Left: Stabilizer generators and logical operators for Shor's $[[9, 1, 0, 3]]$ code. Right: $[[9, 1, 1, 3]]$ code that results from gauge reduction.

Acknowledgments

The author received a lot of helpful guidance and input from his mentors John Preskill and David Poulin. Caltech's IQI provided a stimulating environment. This work was funded by an Axline SURF fellowship.

[1] J. Preskill, Physics 219 Lecture Notes, (2004), http://www.theory.caltech.edu/ preskill/ph229/
[2] M. Nielsen and I. Chuang, "Quantum Computation and Quantum Information," (Cambridge University Press, Cambridge, U.K., 2000)
[3] P. W. Shor, "Scheme for reducing decoherence in quantum computer memory," Phys. Rev. A **52**, 2493 (1995).
[4] A. Steane, "Error-correcting codes in quantum theory," Phys. Rev. Lett. **77**, 793 (1996).
[5] S. B. Bravyi and A. Yu. Kitaev, "Quantum codes on a lattice with boundary," quant-ph/9810052 (1998).

[6] D. Aharonov and M. Ben-Or. Fault-tolerant quantum computation with constant error rate, 1999, arXiv:quant-ph/9906129.
[7] D. Bacon, "Operator quantum error correcting subsystems for self-correcting quantum memories," (2005), Phys. Rev. A 73 012340 (2006), quant-ph/0506023.
[8] J. Kempe, A. Kitaev, and O. Regev. The complexity of the local hamiltonian problem. SIAM Journal of Computing, 35(5):1070–1097, 2006.
[9] N. D. Mermin and H. Wagner, Phys. Rev. Lett. 17 (1966) 1133;
S. Coleman, Commun. Math. Phys. 31 (1973) 259.
[10] E. Dennis, A. Yu. Kitaev, A. Landahl, J. Preskill, "Topological quantum memory," (2001), J. Math. Phys. 43 4452-4505, quant-ph/0110143
[11] D. Gottesman, "Stabilizer codes and quantum error correction," Caltech Ph.D. thesis (1997), quant-ph/9705052.
[12] D. Kribs, R. Laflamme, D. Poulin, M. Lesosky, "Operator Quantum Error Correction" (2006), quant-ph/0504189
[13] D. Poulin, "Stabilizer formalism for operator error correction," (2006), Phys. Rev. Lett. 95 230504, quant-ph/0508131.
[14] P. Aliferis, A. Cross. See http://www.its.caltech.edu/ panagiot/waterloo_july2006.pdf by P. Aliferis.