

The Heisenberg Representation of Quantum Computers

Daniel Gottesman*
T-6 Group
Los Alamos National Laboratory
Los Alamos, NM 87545

June 24, 1998

Abstract

Since Shor's discovery of an algorithm to factor numbers on a quantum computer in polynomial time, quantum computation has become a subject of immense interest. Unfortunately, one of the key features of quantum computers — the difficulty of describing them on classical computers — also makes it difficult to describe and understand precisely what can be done with them. A formalism describing the evolution of operators rather than states has proven extremely fruitful in understanding an important class of quantum operations. States used in error correction and certain communication protocols can be described by their stabilizer, a group of tensor products of Pauli matrices. Even this simple group structure is sufficient to allow a rich range of quantum effects, although it falls short of the full power of quantum computation.

1 Introduction

Computers are physical objects. While that seems obvious, it has some profound consequences. The familiar desktop and laptop computers and their big cousins, the mainframes, all work in essentially the same way. Their components are individual registers called *bits*, which can interact through various discrete operations, called *gates*. The underlying hardware varies somewhat from system to system, but currently, all bits are made up of fairly large numbers of atoms. Consequently, they behave like macroscopic *classical* objects. Modern computers are digital, so their bits take the discrete values of 0 and 1 (thus the name "bit").

*gottesma@t6-serv.lanl.gov

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

However, suppose instead that we had a computer whose bits were made up of a small number of atoms, perhaps even just one atom per bit. Then the bits in the computer would start to behave like quantum objects instead of classical objects. Instead of the individual values 0 and 1, a bit might instead be in some superposition of 0 and 1. A bit that can be in such a superposition is called a *quantum bit* or *qubit*. If a classical computer has N bits, it has a total of 2^N possible states. In contrast, a *quantum computer* with N qubits can be in any superposition of those 2^N classical basis states, resulting in an arbitrary state in a 2^N -dimensional Hilbert space. To even describe such a state on a classical computer would require $2^N - 1$ complex numbers.

Classical algorithms are frequently classified as *polynomial* or *exponential* (which actually just means more than polynomial). A polynomial algorithm to solve a problem of size N takes CN^k steps (plus lower-order terms), for some constants C and k . Such an algorithm is considered to be efficient for dealing with large problems, since the computation cost does not increase *too* much as N increases (although if C and/or k are large, the algorithm may not be useful for realistic problems). An exponential time algorithm, on the other hand, means it will be very difficult to solve large problems. Unfortunately, many important problems have no known polynomial algorithm, forcing us to resort to approximations or restricting us to small instances of the problem.

As long as we are using fundamentally classical bits, the algorithms we use will be roughly equivalent. Perhaps some system uses a factor of N more or less steps than another, or perhaps C is larger or smaller, but an exponential algorithm on one classical computer will translate to an exponential algorithm on any other classical computer. However, if we change the underlying laws of physics and work with quantum bits instead of classical bits, the situation changes. Even if there is no polynomial classical algorithm to solve a problem, there may be a polynomial quantum algorithm.

Perhaps the most dramatic known example of this is Shor's factoring algorithm [1]. Factoring an n -bit number can be very difficult, particularly if it is the product of two primes of about the same size. Despite much work, no polynomial classical algorithm is known. In fact, many classical cryptographic protocols, such as RSA, depend on the fact that factoring is hard to do.

It turns out that factoring a number N is essentially equivalent (up to some straightforward additional computation) to finding the order r of a random number x modulo N . On a quantum computer, we can do this by starting with a superposition $\sum |a\rangle$ and calculating $x^a \bmod N$ in a second register, producing the state $\sum |a\rangle |x^a\rangle$. This state is periodic in a , with period r . Therefore, by performing the discrete Fourier transform and measuring the result, we can extract r and thus factor N . The result is a polynomial quantum algorithm where all known classical algorithms are exponential.

Unfortunately, because quantum computers deal with arbitrary states in such a large Hilbert space, it can be very difficult to construct and analyze even relatively small networks of quantum gates. It turns out, however, that a

restricted class of networks can be comparatively easily described if we follow the evolution, not of the state of the quantum computer, but instead of a set of operators that could act on the computer [2]. In some ways, this is analogous to the standard Heisenberg representation of quantum mechanics, where the operators evolve in time, as opposed to the Schrödinger picture, where the states evolve. Similar, but less powerful techniques, have been used in the NMR community for years under the name “the product operator formalism.” For a strengthening of the product operator formalism with applications to quantum computing, see [3].

2 Basics of the Heisenberg Representation

Suppose we have a quantum computer in the state $|\psi\rangle$, and we apply the operator U . Then

$$UN|\psi\rangle = UNU^\dagger U|\psi\rangle, \quad (1)$$

so after the operation, the operator UNU^\dagger acts on states in just the way the operator N did before the operation. Therefore, applying U to the computer transforms an arbitrary operation N by

$$N \rightarrow UNU^\dagger. \quad (2)$$

By following the evolution of a sufficiently large set of N 's, we will be able to completely reconstruct the evolution of the state vector. The evolution (2) is linear, so it will be sufficient to follow a set that spans the set of $n \times n$ matrices \mathcal{M}_n .

For this set, we choose the *Pauli group* P , which consists of $4 \cdot 4^n$ elements. The Pauli group contains the 4^n n -qubit tensor products of the identity I and the Pauli matrices σ_x , σ_y , and σ_z :

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \quad (3)$$

Note that

$$\sigma_x^2 = \sigma_y^2 = \sigma_z^2 = I \quad (4)$$

and

$$\sigma_y = i\sigma_x\sigma_z. \quad (5)$$

Therefore, to complete the group, we must allow each of the 4^n tensor products to have an overall phase of ± 1 or $\pm i$. Below, I will write X , Y , and Z instead of σ_x , σ_y , and σ_z .¹ The operators X_i , Y_i , and Z_i are X , Y , and Z acting on the i th qubit in the computer.

¹In earlier publications, Y was instead equal to $-i\sigma_y$.

Furthermore, (2) is a multiplicative group homomorphism:

$$MN \rightarrow UMNU^\dagger = (UMU^\dagger)(UNU^\dagger). \quad (6)$$

Therefore, we can deduce the behavior of any element of the Pauli group by just following the evolution of a generating set. A convenient generating set is just $\{X_1, \dots, X_n, Z_1, \dots, Z_n\}$. To completely specify a general operator, we need only describe the evolution of $2n$ single-qubit operators.

3 The Clifford Group

Now, in general, the transformation (2) could take a Pauli matrix N to any of a rather large class of unitary operators. If we consider arbitrary quantum gates, the description of the transformation of our generating set will rapidly become unmanageably large. Instead, we will consider a restricted class of gates — the gates which transform elements of the Pauli group into other elements of the Pauli group.

The set of operators which leave the group \mathcal{P} fixed under conjugation form a group $N(\mathcal{P})$, the *normalizer* of \mathcal{P} in $U(2^n)$. $N(\mathcal{P})$ is also called the *Clifford group* \mathcal{C} for its relationship to the usual Clifford groups and Clifford algebras. While \mathcal{C} is considerably smaller than the full unitary group $U(2^n)$, it still contains a number of operators of particular interest.

For instance, the Clifford group contains the single-qubit Hadamard transform R :

$$R|j\rangle = |0\rangle + (-1)^j|1\rangle, \quad R = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}. \quad (7)$$

Another element is the phase gate P :

$$P|j\rangle = i^j|j\rangle, \quad P = \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix}. \quad (8)$$

It also contains the controlled-NOT (CNOT) gate, also known as the XOR:

$$\text{CNOT}|j\rangle|k\rangle = |j\rangle|j+k \bmod 2\rangle. \quad (9)$$

The first qubit is called the *control* qubit and the second is the *target* qubit.

In fact, R , P , and CNOT, applied to arbitrary qubits or pairs of qubits, generate the full Clifford group \mathcal{C} . The transformations these three gates induce on the Pauli matrices are given in table 1. The table also gives the symbols used to represent the Clifford group operations in gate networks, as well as the symbol for a measurement.

If we apply R or P to qubit number j , I will write $R(j)$ or $P(j)$ to represent the gate. If we apply a CNOT with qubit j as control and k as target, I will write $\text{CNOT}(j \rightarrow k)$.


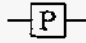
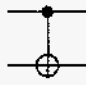
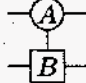
R	$X \rightarrow Z$ $Z \rightarrow X$	
P	$X \rightarrow Y$ $Z \rightarrow Z$	
CNOT	$X \otimes I \rightarrow X \otimes X$ $I \otimes X \rightarrow I \otimes X$ $Z \otimes I \rightarrow Z \otimes I$ $I \otimes Z \rightarrow Z \otimes Z$	
Measurement of A, plus operation B performed if measurement result is -1		

Table 1: Generators of the Clifford group, and their induced transformations on the Pauli matrices.

4 Applying the Heisenberg Representation

To see how to apply the Heisenberg representation, consider the following example:

Example 1 *Alice's quantum computer is working too well. Instead of performing single controlled-NOT gates, it does three at a time (figure 1a). What is it actually doing?*

To figure this out, we will follow the evolution of four operators: \bar{X}_1 , \bar{X}_2 , \bar{Z}_1 , and \bar{Z}_2 . The bars above the operators indicate that these operators represent the *logical* operators X_1 , X_2 , Z_1 , and Z_2 from the beginning of the computation. This terminology will make it easier to follow their evolution through the complete circuit.

Consider first \bar{X}_1 . It begins the computation as $X_1 = X \otimes I$. After the first CNOT (A), it becomes $X \otimes X$, as per table 1. We can rewrite this as

$$(X \otimes I)(I \otimes X), \quad (10)$$

so step B (CNOT(2 \rightarrow 1)) maps \bar{X}_1 to

$$(X \otimes I)(X \otimes X) = I \otimes X. \quad (11)$$

Then after step C, we still have $I \otimes X$, so the full circuit maps

$$\bar{X}_1 = X \otimes I \rightarrow I \otimes X. \quad (12)$$

We can perform a similar computation for \bar{X}_2 , \bar{Z}_1 , and \bar{Z}_2 . The complete calculation is summarized in figure 1b. The network of figure 1a exchanges $X_1 \leftrightarrow X_2$ and $Z_1 \leftrightarrow Z_2$. Therefore, it swaps the first and second qubits.




Figure 1: Alice's quantum computer: a) network, b) analysis.

Example 2 *The lever on Bob's quantum computer is stuck in the "forward" position, so it can only perform controlled-NOTs from qubit 1 to qubit 2. He can still perform single-qubit operations normally. How can he perform a CNOT from qubit 2 to qubit 1?*

To solve this problem, the key is to notice the symmetric behavior of the CNOT gate. If we switch X and Z and qubits 1 and 2, we get back the original transformation. Therefore, the circuit of figure 2a acts as a CNOT(2 → 1). An analysis of the circuit is given in figure 2b. We recognize the final result as CNOT(2 → 1), and so conclude that the given circuit does produce the desired gate.

Example 3 *Bob is attempting to read a paper written in ancient Hittite. All he can make sense of is a gate network, given in figure 3a. What does this network do?*

An analysis is given in figure 3b. Notice the two minus signs that appear. For instance, in step C, for \bar{X}_2 , we must apply R to Y . But

$$R(Y) = R(iXZ) = iR(X)R(Z) = iZX = -Y. \quad (13)$$

All in all, after the complete network, we get the transformation given for step D.

To convert back to the ket formalism, we can follow basis states. For instance, the initial state $|00\rangle$ starts as an eigenvector of $Z \otimes I$ and $I \otimes Z$, with both eigenvalues $+1$. Therefore, after the network, it will still be the $+1$ eigenvector of both $\bar{Z}_1 = -Y \otimes Y$ and $\bar{Z}_2 = Z \otimes X$. Thus, we deduce that this

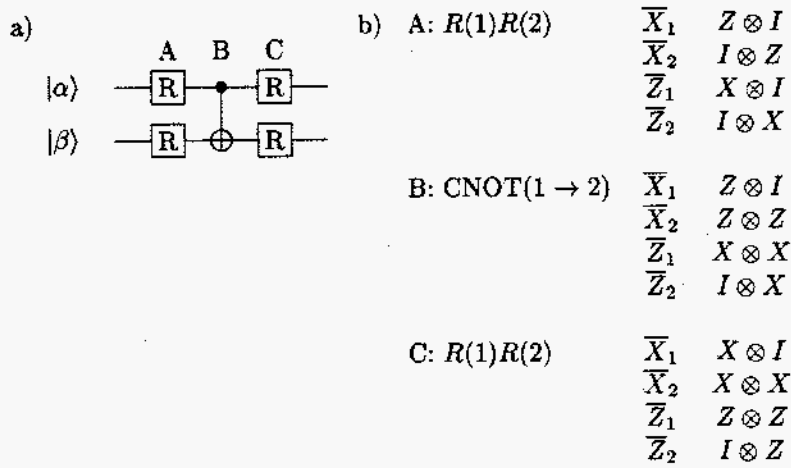


Figure 2: Bob's CNOT(2 \rightarrow 1): a) network, b) analysis.

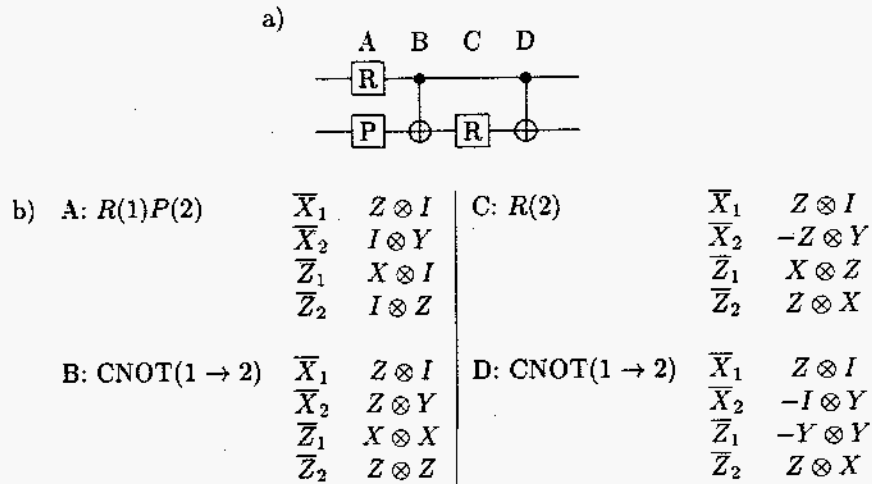


Figure 3: Ancient Hittite gate network: a) network, b) analysis




Figure 4: Alice's improved quantum computer: a) network, b) analysis.

network maps

$$|00\rangle \rightarrow \frac{1}{2} (|00\rangle + |01\rangle - |10\rangle + |11\rangle). \quad (14)$$

In addition, $|01\rangle = (I \otimes X) |00\rangle$, so $|01\rangle$ will map to \bar{X}_2 applied to the image of $|00\rangle$. Since $\bar{X}_2 \rightarrow -I \otimes Y$,

$$|01\rangle \rightarrow \frac{i}{2} (-|01\rangle + |00\rangle + |11\rangle + |10\rangle). \quad (15)$$

We can similarly find the images of $|10\rangle$ and $|11\rangle$. We can put it all together into the following unitary matrix:

$$U = \frac{1}{2} \begin{pmatrix} 1 & i & 1 & i \\ 1 & -i & 1 & -i \\ -1 & i & 1 & -i \\ 1 & i & -1 & -i \end{pmatrix}. \quad (16)$$

This does not correspond to any well-known operation, so just what the Hittites used it for must remain a mystery.

5 Stabilizers

Example 4 *By dint of no little hard work, Alice has partially fixed her quantum computer. Now it only does 2 CNOTs at a time. Unfortunately, she can only get this improvement if she puts in a $|0\rangle$ as the second input qubit (see figure 4a). What does it do now?*

The input state for this network will always be a +1 eigenvector of $I \otimes Z$. Therefore, the state of the system will always be a +1 eigenvector of \bar{Z}_2 . Since the eigenvalue of \bar{Z}_2 is fixed, there is no point in following the evolution of \bar{X}_2 as well, since it would act to move the state to the irrelevant -1 eigenvector of \bar{Z}_2 . The full analysis of this circuit is given in figure 4b.

The format of this analysis is slightly different. The unlabelled operators are those for which the state of the computer is a +1 eigenvector — in this case,

\bar{Z}_2 . In addition, since among the labelled operators, we are only following \bar{X}_1 and \bar{Z}_1 , they are instead labelled simply \bar{X} and \bar{Z} .

The final result may not be instantly recognizable. We can see that the state is in a +1 eigenvector of $Z \otimes I$, so the first qubit is $|0\rangle$. Since $Z \otimes I$ acts as the identity on all possible final states of the computer, it also follows that $\bar{Z} = Z \otimes Z$ is equivalent to $(Z \otimes I)(Z \otimes Z) = I \otimes Z$. Therefore,

$$\bar{X} \rightarrow I \otimes X, \quad (17)$$

$$\bar{Z} \rightarrow I \otimes Z. \quad (18)$$

In other words, the original first qubit has migrated to the second. Alice's computer still performs a swap of the first and second qubits.

In the previous example, only a single qubit had a fixed input value, but more generally, multiple qubits may be fixed or constrained. As above, it will be helpful to consider the set of operators in \mathcal{P} for which the input states are +1 eigenvectors. The set of such operators is closed under multiplication, and therefore forms a group, known as the *stabilizer* S [4, 5].

$$S = \{M \in \mathcal{P} \text{ such that } M|\psi\rangle = |\psi\rangle \text{ for all allowed inputs } |\psi\rangle\} \quad (19)$$

In the previous example, the stabilizer only contained two elements: $I \otimes Z$ and the identity $I \otimes I$.

The stabilizer is always an Abelian group: If $M, N \in S$, then

$$MN|\psi\rangle = |\psi\rangle \quad (20)$$

$$NM|\psi\rangle = |\psi\rangle \quad (21)$$

$$[M, N]|\psi\rangle = 0. \quad (22)$$

Any two elements of \mathcal{P} either commute or anticommute, so it follows that M and N commute. In addition, all elements of \mathcal{P} square to ± 1 (operators with an overall phase of ± 1 square to +1, operators with an overall phase of $\pm i$ square to -1). The operators that square to -1 have imaginary eigenvalues, while the operators that square to +1 have eigenvalues ± 1 . Therefore, any stabilizer is composed of operators in \mathcal{P} which have an overall phase of ± 1 , and which square to +1. It follows that S is isomorphic to $(\mathbb{Z}_2)^k$ for some k .

In fact, k is exactly the number of input qubits that are fixed. The requirement that a single qubit is fixed provides a single operator for the stabilizer. The k operators together then generate S — the 2^k elements of S are products of the generators.

The Pauli group operators that remain interesting are the operators that act on the other $n - k$ qubits (if there are n total qubits). There are a total of 2^{n-k} such independent operators. When the stabilizer puts more complicated constraints on the state, the \bar{X} and \bar{Z} operators are any 2^{n-k} operators that commute with S .




Figure 5: Making a Bell state: a) network, b) analysis.

Example 5 *Alice's quantum computer is finally fixed. She wishes to demonstrate her devotion to Bob by sending him half of a Bell state, along with a card reading "forever entangled." She performs the network of figure 5a and sends the card and second qubit to Bob.*

The analysis of the network appears in figure 5b. In this case, we only need to consider elements of the stabilizer. Again, the generators of the stabilizer are unlabelled. Their order does not matter — we could switch the generators, and they still generate the same group. In fact, we could exchange one or more of the generators with other elements of S without affecting anything, as long as we still have a set of n independent operators to act as generators. The transformation of the rest of the stabilizer is completely determined by the transformation of the generators.

The final state in this case has stabilizer $\{I \otimes I, X \otimes X, -Y \otimes Y, Z \otimes Z\}$. There is a single state which has eigenvalue +1 for all four of these operators. It is given (up to normalization) by

$$|\psi\rangle = \left(\sum_{M \in S} M \right) |00\rangle. \quad (23)$$

This is the correct state because acting on it with $N \in S$ just permutes the terms in the sum, giving the same state.² In this case, the state is

$$\frac{1}{\sqrt{2}} (|00\rangle + |11\rangle), \quad (24)$$

which is indeed a Bell state.

A state which can be completely described by specifying the stabilizer is called a *stabilizer state*. There are many states in the Hilbert space which are not stabilizer states (in fact, there are only a finite number of stabilizer states of a given size), but many of the most interesting states are stabilizer states.

²If $-Z \otimes Z$ had been in the stabilizer instead of $Z \otimes Z$, $|\psi\rangle$ would have been null. To produce the correct state, we would have had to start with some state other than $|00\rangle$ on the right.

Example 6 Alice decides she would rather send the singlet state, but doesn't want to go through the bother of another CNOT. How can she produce the singlet from what she currently has?

The singlet state is

$$\frac{1}{\sqrt{2}} (|01\rangle - |10\rangle). \quad (25)$$

It has the stabilizer $\{I \otimes I, -X \otimes X, -Y \otimes Y, -Z \otimes Z\}$ (generated by $-X \otimes X$ and $-Z \otimes Z$). This is clearly very similar to her current stabilizer (generated by $X \otimes X$ and $Z \otimes Z$). All she needs to do is apply an operator that maps $X \otimes X \rightarrow -X \otimes X$ and $Z \otimes Z \rightarrow -Z \otimes Z$.

One possible operator is $Y \otimes I$.

$$Y(X) = YXY^\dagger = -YY^\dagger X = -X, \quad (26)$$

$$Y(Z) = YZY^\dagger = -YY^\dagger Z = -Z. \quad (27)$$

Another operator could be $I \otimes Y$. $X \otimes Z$ would also work:

$$(X \otimes Z)(X \otimes X)(X \otimes Z) = -X \otimes X, \quad (28)$$

$$(X \otimes Z)(Z \otimes Z)(X \otimes Z) = -Z \otimes Z. \quad (29)$$

In fact, any operator $E \in \mathcal{P}$ that anticommutes with both $X \otimes X$ and $Z \otimes Z$ would work: if $\{E, M\} = 0$, then

$$E(M) = EME^\dagger = -EE^\dagger M = -M. \quad (30)$$

6 Quantum Error-Correcting Codes

Example 7 Alice and Bob sometimes bring quantum data home with them. Their two-year-old child Alice, Jr. is very curious. Even if the data is stored on a high shelf, Alice, Jr. sometimes gets to it. When she does, she puts a single qubit in her mouth, randomizing it, then puts it back. After that, she loses interest and goes off to create trouble elsewhere. How can Alice and Bob tell if Alice, Jr. has disturbed their state? They do not want to destroy it if the state is OK, but they are willing to discard it if it has been changed.

Alice and Bob cannot simply measure their quantum state, since that would collapse any superposition it might be in. Instead, the solution is to group the qubits in sets of two. To each pair, they add two qubits in a known state, and apply some operations resulting in the stabilizer S generated by

$$\begin{aligned} X \otimes X \otimes X \otimes X \\ Z \otimes Z \otimes Z \otimes Z. \end{aligned} \quad (31)$$




Figure 6: A single qubit error-detecting code

For instance, they could use the circuit in figure 6.

The stabilizer (31) has the advantage that any single-qubit operator anti-commutes with at least one of the generators. Therefore, if we apply any error consisting of a single-qubit operator $E \in \mathcal{P}$, at least one of the generators will change sign, as per (30). We can therefore detect that E has occurred by measuring the generators of S . If the eigenvalues are $+1$, no error has occurred (or two or more errors have occurred).

Of course, randomizing a qubit would generally not correspond to performing an operator in \mathcal{P} . However, \mathcal{P} does span the set of matrices, so the most general error will be the sum of operators from \mathcal{P} . Such an error would put the state in a superposition of eigenstates of S , but measuring the generators of S collapses the state into just one eigenstate. This has the effect of collapsing the error into some operator from \mathcal{P} .

The states described by (31) form a *quantum error-detecting code* with the ability to detect one error [7, 8]. A code that detects any error affecting less than d qubits is said to have distance d . The four-qubit code thus has distance 2. Because the code is completely determined by its stabilizer, it is said to be a *stabilizer code*.

Example 8 *Alice and Bob wish to bring some particularly important quantum data home. They worked very hard to produce it, and do not want to have to discard it if Alice, Jr. ruins a single qubit. What can they do?*

Now, Alice and Bob have to go beyond an error-detecting code to a *quantum error-correcting code*. However, the principles are the same. In order to correct a single error, it is sufficient for Alice and Bob to be able to figure out exactly what the error was. Since it is from \mathcal{P} (or has been collapsed to an error from \mathcal{P}), they can then just apply the inverse to restore the original state.

Distinguishing the state $E|\psi\rangle$ from $F|\psi\rangle$ is the same problem as distinguishing the state $EF|\psi\rangle$ ³ from the state $|\psi\rangle$. Therefore, a code that detects the error EF will distinguish the errors E and F . In other words, to correct any

³Or more precisely, $E^\dagger F|\psi\rangle$, which is the same for Hermitian E .

X	Z	Z	X	I
I	X	Z	Z	X
X	I	X	Z	Z
Z	X	I	X	Z
\bar{X}	X	X	X	X
\bar{Z}	Z	Z	Z	Z

Table 2: The five-qubit code

single-qubit error, the code must detect all two-qubit errors. More generally, to correct t arbitrary errors, the code should have distance $2t + 1$.

The stabilizer of the smallest error-correcting code to fix one error is given in table 2 [9, 10]. It encodes a single qubit in five qubits, and has distance three. We say it is a $[[5, 1, 3]]$ quantum code (or a $[[5, 1, 3]]$ code).

The \bar{X} and \bar{Z} operators are, as before, the logical operators corresponding to operations on the original data qubits. By analyzing their behavior, we could, for instance, understand how to perform operations on the data without first decoding the state (*fault-tolerant* operations [2, 6]).

I said above that a stabilizer code has distance d when for every Pauli group operator E with weight less than d , there is an element M of the stabilizer that anticommutes with E . In fact, the requirement is slightly weaker. It is possible for $E|\psi\rangle = |\psi\rangle$ for all codewords $|\psi\rangle$ for some particular E . In this case, there is no way the code will be able to detect that E has occurred — but there is no need to do so, since the state produced is exactly the correct state. If $E|\psi\rangle = |\psi\rangle \forall |\psi\rangle$, then $E \in S$. Therefore, the complete condition for a stabilizer code to have distance d is that for all $E \in \mathcal{P}$ of weight less than d , either $\exists M \in S$ with $\{M, E\} = 0$, or $E \in S$. If the second condition is ever needed, the code is said to be *degenerate*. If the second condition is not used, as for the four- and five-qubit codes above, the code is *nondegenerate*.

Quantum error-correcting codes solve a serious difficulty in the design of quantum computers. Simply having objects that behave in a more-or-less quantum fashion is insufficient to produce a quantum computer. If a qubit interacts with its environment, it will have a tendency to act as if it is in one or the other basis state (for some particular basis, which depends on the interaction), just as if it had been measured. A quantum computer with strong interactions with the environment will thus tend to act as if it is in one of the 2^N basis states — in other words, it will act just like a classical computer. Using quantum error correction and fault-tolerant operations, we can greatly simplify our task: instead of having to build a quantum computer with essentially no uncontrolled interaction with the environment, we need only build one where the interactions are small enough to give us time to do error-correction.⁴

⁴Only. Can you tell I'm a theorist?

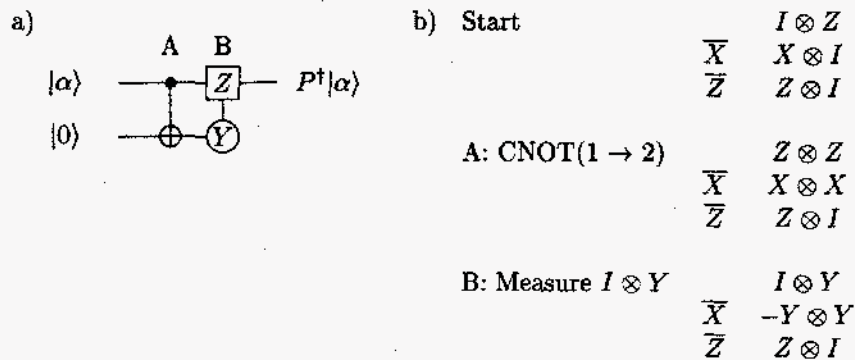


Figure 7: Creating the P gate: a) network, b) analysis.

7 Measurements

Example 9 Bob has bumped his head and cannot remember how to perform the P gate. Luckily, he does remember how to perform the CNOT, the Pauli group, and how to measure operators in the Pauli group. How can he make a P gate?

The solution is given in figure 7a. A full analysis appears in figure 7b. If the initial state $|\alpha\rangle = a|0\rangle + b|1\rangle$, then the state after completing step A is

$$a|00\rangle + b|11\rangle = \frac{1}{2}(a|0\rangle - ib|1\rangle)(|0\rangle + i|1\rangle) + \frac{1}{2}(a|0\rangle + ib|1\rangle)(|0\rangle - i|1\rangle). \quad (32)$$

Therefore, measuring $I \otimes Y$ will yield either the state

$$(a|0\rangle - ib|1\rangle)(|0\rangle + i|1\rangle), \quad (33)$$

or the state

$$(a|0\rangle + ib|1\rangle)(|0\rangle - i|1\rangle). \quad (34)$$

The first has stabilizer $I \otimes Y$ and $\bar{X} = -Y \otimes I$, while the second has stabilizer $-I \otimes Y$ and $\bar{X} = +Y \otimes I$. In both cases, $\bar{Z} = Z \otimes I$. The two possible states after the measurement are related by the action of the operator $Z \otimes Z$. Note that $Z \otimes Z$ is the stabilizer after step A.

Therefore, step B consists of measuring $I \otimes Y$ and performing $Z \otimes Z$ if the result is -1 . That produces the final state given in the table. Since $I \otimes Y$ is in the stabilizer, $Y \otimes Y$ is equivalent to $Y \otimes I$ on the states of interest. We can recognize the final operation as a P^\dagger gate. We can get a regular P gate either by performing this 3 times, or more easily by applying Z once to the output qubit.

The fact that we were able to convert the $+1$ and -1 measurement results into the same state was no fluke. Generally, suppose we have a state which

is partially (or completely) described by a stabilizer S . Suppose we wish to measure an operator $A \in \mathcal{P}$ which anticommutes with some element $M \in S$. The measurement performs one of the operators

$$P_{\pm} = \frac{1}{2}(I \pm A), \quad (35)$$

depending on whether the measurement result is ± 1 . But

$$MP_-M^\dagger = \frac{1}{2}M(I - A)M^\dagger = \frac{1}{2}(I + A)MM^\dagger = P_+. \quad (36)$$

Therefore, by applying M whenever the measurement result is -1 , we can ensure that we always get the state $P_+|\psi\rangle$.

We can go further and describe the evolution of the state using the Heisenberg formalism. The state $P_+|\psi\rangle$ will always be in a $+1$ eigenstate of A , so A is a member of the new stabilizer S' . M , on the other hand, will not be in S' . The other generators will be in S' if they commute with A , since commuting observables can have simultaneous eigenvectors. On the other hand, if $N \in S$ does not commute with A , $N \notin S'$; but MN does commute with A , so $MN \in S'$.


We can also follow the evolution of the \bar{X} and \bar{Z} operators. If they commute with A , they are unaffected by the measurement of A . Our original presentation of \bar{X}_j or \bar{Z}_j might not commute with A , but recall that operators acting on the state are equivalent up to multiplication by elements of the stabilizer. Therefore, even if \bar{X}_j does not commute with A , $M\bar{X}_j$ does, and acts the same way. Therefore, under the measurement, $\bar{X}_j \rightarrow M\bar{X}_j$.

We can sum up the rules for evolving the operators after a measurement of A (and correction if the result is -1) as follows:

1. Identify $M \in S$ satisfying $\{M, A\} = 0$.
2. Remove M from the stabilizer
3. Add A to the stabilizer
4. For each N , where N runs over the other generators of S and the \bar{X} and \bar{Z} operators, leave N alone if $[N, A] = 0$, and replace N with MN if $\{N, A\} = 0$.

Example 10 (Quantum Teleportation) *Alice needs to quickly send a qubit to Bob, but Quantum Parcel Services (QPS) is on strike. Luckily, Alice shares an EPR pair with Bob, and the regular classical phone lines are still open. How can she get him the qubit?*

The solution to this problem is quantum teleportation [11]. It is straightforward to analyze in the Heisenberg representation. I will assume Alice and Bob start with the Bell state $|00\rangle + |11\rangle$. The network is given in figure 8a. Alice and



b) Start

	I	X	X
	I	Z	Z
\bar{X}	X	I	I
\bar{Z}	Z	I	I

A: CNOT($1 \rightarrow 2$)

	I	X	X
	Z	Z	Z
\bar{X}	X	X	I
\bar{Z}	Z	I	I

B: Measure $X \otimes I \otimes I$

		X	X
\bar{X}		X	I
\bar{Z}		Z	Z

C: Measure $Z \otimes I$

\bar{X}			X
\bar{Z}			Z

Figure 8: Quantum teleportation: a) network, b) analysis.

Bob begin with one EPR pair. Alice and Bob perform some local operations, and in the process Alice sends Bob two classical bits. The net result is that Alice's qubit is sent to Bob. The Heisenberg representation analysis is given in figure 8b. The two classical bits sent from Alice to Bob are the two measurement results, which are required for Bob to perform the correct operations to turn the state into the $+1$ eigenvector of the two measured operators.

Since in this network, once a qubit is measured, it is completely determined and never reused, we drop measured qubits from the analysis. A more complicated network might include partial entangled measurements, in which case the measured qubits would still have some interesting remaining degrees of freedom and should be retained in the description. Note that when dropping qubits, the size of the stabilizer shrinks, but we keep the same number of \bar{X} and \bar{Z} operators, since we only drop qubits which are completely constrained by the stabilizer.

Example 11 (Remote XOR) *Alice forgot to perform a CNOT from another qubit of hers to the one she sent to Bob. Unfortunately, they only have one more shared EPR pair (Bob has been gambling at the entanglement casino and losing heavily). Can she still perform the CNOT?*


The network that allows Alice to perform a remote XOR to Bob's qubit is given in figure 9a. Alice and Bob start with one EPR pair. They each perform a CNOT and a measurement, then send each other their measurement results (for a total of one classical bit each direction). The result is a CNOT from Alice's qubit to Bob's. This operation has also been studied by [12]. The analysis appears in figure 9b.

8 Discussion and Summary

The methods comprising the Heisenberg representation of quantum computation allow a much more rapid analysis of many networks than can be done using the usual methods of multiplying together the full $2^n \times 2^n$ unitary matrices describing the evolution of an n qubit system, or by directly following the evolution of state vectors in such a system. For most networks, the usual methods require keeping track of an exponential number of matrix elements or coefficients.

In contrast, for a network composed only of gates from the Clifford group and measurements of Pauli group operators, the Heisenberg representation provides an excellent method of describing the system. In such networks, it is only necessary to follow the evolution of at most $2n$ \bar{X} and \bar{Z} operators, and each one is a member of the Pauli group, so can be described by $2n + 1$ bits. Thus, the time and space required to analyze such a network on a classical computer are polynomials in n , instead of exponentials.

In fact, we can go beyond just Clifford group elements and Pauli group measurements, and also add Clifford group operations conditioned on classical



b) Start

	<i>I</i>	<i>X</i>	<i>X</i>	<i>I</i>
	<i>I</i>	<i>Z</i>	<i>Z</i>	<i>I</i>
\overline{X}_A	<i>X</i>	<i>I</i>	<i>I</i>	<i>I</i>
\overline{Z}_A	<i>Z</i>	<i>I</i>	<i>I</i>	<i>I</i>
\overline{X}_B	<i>I</i>	<i>I</i>	<i>I</i>	<i>X</i>
\overline{Z}_B	<i>I</i>	<i>I</i>	<i>I</i>	<i>Z</i>

A: CNOT(1 → 2)CNOT(3 → 4)

	<i>I</i>	<i>X</i>	<i>X</i>	<i>X</i>
	<i>Z</i>	<i>Z</i>	<i>Z</i>	<i>I</i>
\overline{X}_A	<i>X</i>	<i>X</i>	<i>I</i>	<i>I</i>
\overline{Z}_A	<i>Z</i>	<i>I</i>	<i>I</i>	<i>I</i>
\overline{X}_B	<i>I</i>	<i>I</i>	<i>I</i>	<i>X</i>
\overline{Z}_B	<i>I</i>	<i>I</i>	<i>Z</i>	<i>Z</i>

B: Measure $I \otimes Z \otimes I \otimes I$

	<i>Z</i>	<i>Z</i>	<i>I</i>
\overline{X}_A	<i>X</i>	<i>X</i>	<i>X</i>
\overline{Z}_A	<i>Z</i>	<i>I</i>	<i>I</i>
\overline{X}_B	<i>I</i>	<i>I</i>	<i>X</i>
\overline{Z}_B	<i>I</i>	<i>Z</i>	<i>Z</i>

C: Measure $I \otimes X \otimes I$

\overline{X}_A	<i>X</i>	<i>X</i>
\overline{Z}_A	<i>Z</i>	<i>I</i>
\overline{X}_B	<i>I</i>	<i>X</i>
\overline{Z}_B	<i>Z</i>	<i>Z</i>

Figure 9: The remote XOR: a) network, b) analysis.

bits (which might, of course, be the results of measurements performed earlier in the computation). This collection of observations forms the proof of the following theorem [13]:

Theorem 1 (Knill's theorem) *Any quantum computer performing only: a) Clifford group gates, b) measurements of Pauli group operators, and c) Clifford group operations conditioned on classical bits, which may be the results of earlier measurements, can be perfectly simulated in polynomial time on a probabilistic classical computer.*

Of course, Clifford group operations and Pauli group measurements do not provide a universal set of quantum gates. Another gate is needed, which could be a quantum version of the classical Toffoli gate ($|a\rangle|b\rangle|c\rangle \rightarrow |a\rangle|b\rangle|c + ab\rangle$), a $\pi/8$ rotation of the Bloch sphere, or the square root of the controlled-NOT gate, among other possibilities. The Clifford group plus an appropriate extra gate generate a set of unitary operators dense in $U(2^n)$ and therefore form a universal set of gates.

Knill's theorem implies that quantum computation is only more powerful than classical computation when it uses gates outside the Clifford group. However, networks using only Clifford group gates also have a number of important applications in the area of quantum communications. Quantum error-correcting codes are an important example — stabilizer codes use only Clifford group gates for encoding and decoding, yet are extremely useful for overcoming the effects of errors and decoherence. Other important communication problems, such as quantum teleportation, also use only Clifford group gates and measurements. Finally, networks consisting of Clifford group gates and measurements may provide useful subroutines to more complex quantum computations; an efficient method of analyzing and searching for such subroutines could prove very useful.

References

- [1] Peter W. Shor, "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer," *SIAM J. Computing* **26**, 1484 (1997); quant-ph/9508027.
- [2] Daniel Gottesman, "A Theory of fault-tolerant quantum computation," *Phys. Rev. A* **57**, 127 (1998); quant-ph/9702029.
- [3] Shyamal S. Somaroo, David G. Cory, and Timothy F. Havel, "Expressing the operations of quantum computing in multiparticle geometric algebra," *Phys. Lett. A* **240**, 1 (1998); quant-ph/9801002.
- [4] Daniel Gottesman, "Class of quantum error-correcting codes saturating the quantum Hamming bound," *Phys. Rev. A* **54**, 1862 (1996); quant-ph/9604038.

- [5] A. R. Calderbank, E. M. Rains, P. W. Shor, and N. J. A. Sloane, "Quantum error correction and orthogonal geometry," *Phys. Rev. Lett.* **78**, 405 (1997); quant-ph/9605005.
- [6] P. Shor, "Fault-tolerant quantum computation," *Proceedings of the 37th Symposium on the Foundations of Computer Science*, IEEE Computer Society Press, Los Alamitos, CA, 56 (1996); quant-ph/9605011.
- [7] L. Vaidman, L. Goldenberg, and S. Weisner, "Error prevention scheme with four particles," *Phys. Rev. A* **54**, 1745 (1996); quant-ph/9603031.
- [8] M. Grassl, T. Beth, and T. Pellizzari, "Codes for the quantum erasure channel," *Phys. Rev. A* **56**, 33 (1997); quant-ph/9610042.
- [9] R. Laflamme, C. Miquel, J. P. Paz, and W. Zurek, "Perfect quantum error correction code," *Phys. Rev. Lett.* **77**, 198 (1996); quant-ph/9602019.
- [10] C. Bennett, D. DiVincenzo, J. Smolin, and W. Wootters, "Mixed state entanglement and quantum error correction," *Phys. Rev. A* **54**, 3824 (1996); quant-ph/9604024.
- [11] C. H. Bennett, G. Brassard, C. Crepeau, R. Jozsa, A. Peres, and W. K. Wootters, "Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels," *Phys. Rev. Lett.* **70**, 1895 (1993).
- [12] Wim van Dam, private communication.
- [13] E. Knill, private communication.