Introduction to quantum error correction

A. M. Steane

Phil. Trans. R. Soc. Lond. A 1998 **356**, 1739-1758
doi: 10.1098/rsta.1998.0246

The Royal Society

Introduction to quantum error correction

By A. M. Steane

Department of Atomic and Laser Physics, Clarendon Laboratory,
Parks Road, Oxford OX1 3NP, UK

An introduction to quantum error correction (QEC) is given, and some recent developments are described. QEC consists of two parts: the physics of error processes and their reversal, and the construction of quantum error-correcting codes. Errors are caused both by imperfect quantum operations, and by coupling between the quantum system and its environment. Any such process can be analysed into a sum of 'error operators', which are tensor products of Pauli spin operators. These are the analogues of classical error vectors. A quantum error-correcting code is a set of orthogonal states, 'quantum codewords', which behave in a certain useful way under the action of the most likely error operators. A computer or channel which only uses such states can be corrected by measurements which determine the error while yielding no information about which codeword or superposition of codewords is involved. Powerful codes can be found using a construction based on classical error-correcting codes. An analysis which allows even the corrective operations themselves to be imperfect leads to powerful and counter-intuitive results: the quantum coherence of a long quantum computation can be preserved even though every qubit in the computer relaxes spontaneously many times before the computation is complete.

Keywords: error-correcting code; syndrome; noise; CSS code; parity check; error operator

1. Introduction

The concept of error correction is a central component of classical information theory, and similarly quantum error correction (QEC) is one of the foundation stones of quantum information theory. The basic principles of QEC were only recently discovered, dating from 1995, but progress has been rapid since then, so that now the subject is well established and covers a large number of concepts and mathematical techniques. This paper will introduce the basic concepts of QEC mostly at a level suitable for readers unfamiliar with the subject, with some recent developments mentioned at the end.

The existence of quantum error-correcting codes was discovered independently by Shor (1995) and Steane (1996*a*). Shor proved that nine qubits could be used to protect a single qubit against general errors, while Steane described a general code construction whose simplest example does the same job using seven qubits. However, it was not clear from these early works whether the methods could be generalized. QEC was established as a general and powerful technique in the subsequent papers of Calderbank & Shor (1996) and Steane (1996*b*) in which general code constructions, existence proofs and correction methods were given. Thereafter, Knill & Laflamme (1997) provided a more general theoretical framework, describing requirements for quantum error-correcting codes (§ 3), and measures of the fidelity

Phil. Trans. R. Soc. Lond. A (1998) **356**, 1739–1758
Printed in Great Britain

of corrected states. They, and independently Ekert & Macchiavello (1996), derived the quantum Hamming bound (see § 5 *a*). The close relationship between entanglement purification (Deutsch *et al*. 1996; Bennett *et al*. 1996*a*) and quantum error correction was discussed by Bennett *et al*. (1996*b*). The theory of the construction of quantum error-correcting codes was greatly advanced by Gottesman (1996*a*) and Calderbank *et al*. (1997); their work permitted many new codes to be discovered (Gottesman 1996*a*, *b*; Calderbank *et al*. 1996, 1997; Steane 1996*c*, *d*). Codes whose efficiency exceeds (by a small amount) the possibilities of such general methods have been discovered (Shor & Smolin 1996; Rains *et al*. 1997). The subject reached a degree of maturity with the discovery by Shor & Laflamme (1996) of a beautiful quantum analogue to the classical MacWilliams identities, which provide constraints on codes and fidelity measures. An important issue which remains unresolved is the definition of the capacity of a quantum channel, though much useful work has been done (see, for example, Calderbank & Shor 1996; Steane 1996*b*; Knill & Laflamme 1997; Bennett *et al*. 1996*b*, 1997; Lloyd 1997; Schumacher & Nielsen 1996; Barnum *et al*. 1996; van Enk *et al*. 1997, and the other papers in this volume).

The results of quantum error correction are powerful and highly counter-intuitive. For example, they permit a quantum computer to produce the correct output, involving a large and complicated quantum interference between its constituent qubits, even though *every* qubit in the computer relaxes spontaneously *many times* during the computation!

2. The problem of stability in quantum computation

It is well known that quantum computers are highly sensitive to noise and imprecision. By *noise* we mean extraneous coupling between the computer and all other systems, normally referred to as the 'environment', and by *imprecision* we mean inaccuracy in the operations (quantum gates) which are deliberately applied to the computer in order to make it compute. The latter case can be regarded as noise affecting all the qubits involved in the gate, followed by a perfect implementation of the gate. This instability problem has been studied in detail by various authors (Unruh 1995; Palma *et al*. 1996; Chuang *et al*. 1995; Plenio & Knight 1996, 1997), but the essential cause of it is easy to discern.

Quantum computation relies on the manipulation of quantum states in such a way that coherence is preserved. The word 'coherence' has many meanings in physics, but in this context it means essentially that if a computer can be in states such as $|\Psi_1\rangle$ and $|\Psi_2\rangle$, then in the course of a computation it might adopt the state

$$|A\rangle = \frac{1}{\sqrt{2}}(|\Psi_1\rangle + \mathrm{e}^{\mathrm{i}\phi}|\Psi_2\rangle). \tag{2.1}$$

To preserve coherence, we require not only that $|\Psi_1\rangle$ and $|\Psi_2\rangle$ are uncorrupted by noise and imprecision, but also that the value of the phase ϕ in the superposition is well defined. However, the phase becomes undefined either if an imprecise operation produces rotation of the state through an unknown angle of order π, or if the coupling to the environment produces an entangled state such as $(|\Psi_1\rangle|e_1\rangle + \mathrm{e}^{\mathrm{i}\phi}|\Psi_2\rangle|e_2\rangle)/\sqrt{2}$, where $|e_i\rangle$ are states of the environment and $\langle e_1 \mid e_2\rangle = 0$.

Suppose there are K qubits in the computer, where we take the qubits to be physically separate systems such as atoms. Quantum computation is only more efficient than classical computation if states of the form $|A\rangle$ feature predominantly, such that $O(K)$ of the qubits are involved in the interference (e.g. $|\Psi_1\rangle = |000\ldots0\rangle$, $|\Psi_2\rangle = |111\ldots1\rangle$). Therefore, at any stage of the computation the computer fails if any one of the K qubits 'decoheres', that is, becomes entangled with the environment or randomly precesses. If the probability for a decohering process is p for any qubit, during the time it takes to perform one computational step, and there are S steps in the whole computation, then the probability that the quantum algorithm succeeds is of order $(1-p)^{KS} \simeq 1 - KSp$. For a successful computation, we therefore require $p \leqslant O(1/KS)$†.

To get a feel for the size of a quantum computer which might be needed to do powerful and versatile quantum computing, we take as an example the values of K and S required for Shor's factorization algorithm (Shor 1994) working to factorize a 130-digit number (which is currently the limit for classical computers). We need approximately $L = 400$ qubits to store the number to be factorized, and a further $4L = 1600$ as workspace (Beckman *et al.* 1996), giving $K \simeq 2000$. The algorithm requires $S \simeq 300L^3 \simeq 2 \times 10^{10}$; therefore, for success, the probability of decoherence per qubit per computational step must be $p < 2 \times 10^{-14}$.

To estimate the rate of decoherence in real systems is a difficult task, but it is safe to say that it is extremely hard (perhaps impossible) to find any system which could form a quantum computer in which the probability of an unwanted process such as spontaneous emission is less than *ca.* 10^{-5} or possibly 10^{-6} per computational step. Hence quantum systems are seven or more orders of magnitude too noisy to permit large-scale quantum computation. This is not to say that, for example, 10^4 operations on 100 qubits might not permit very valuable quantum computations (this is an open question), but it rules out most quantum algorithms.

The above conclusion leads us to ask why classical computers are as stable as they are. A single bit error in a classical computation would often be sufficient to cause failure, yet these devices are phenomenally stable, routinely performing numbers of operations of the order of Avogadro's number without a single bit error. The stability of classical computers comes not from a precise construction and isolation, but from the fact that they are insensitive to noise. This property can be identified ultimately as a property of the basic components: the electronic switches. Just as a mechanical switch is highly unlikely to spontaneously toggle, its electronic counterpart the flip-flop remains stable in one of two configurations, until toggled by an electrical impulse which is large compared to the noise in the circuit. This stability derives from a judicious combination of amplification and dissipation: in the flip-flop, transistors perform the role of amplifiers, in the mechanical switch a spring serves this role. The amplifier provides a restoring force to return the switch to one of its two stable configurations, but dissipation is also needed since in its absence the switch would oscillate and eventually toggle.

Stability in complex processes, whether computers or other man-made devices, or biological systems, is always brought about through amplification and dissipation. However, an unknown quantum state cannot be amplified ('no-cloning' theorem, Wootters & Zurek 1982; Dieks 1982; Glauber 1986; Steane 1998) and dissipation is by definition irreversible and hence incompatible with unitary quantum evolution. Therefore these methods cannot be used in quantum computers. For this reason it was thought that it would be impossible to actively correct a quantum computer while

† It is only necessary for the probability of success to be greater than $\frac{1}{2}$ for the computer to be useful, since then repetition of the whole computation can be used to generate a reliable result efficiently.

it was running, and the problem of instability of quantum computation appeared insurmountable.

3. Principles of quantum error correction

The basic ingredients of the theory of quantum error correction are the quantum states we wish to consider, and the noise processes we wish to correct. The theory can be presented in several equivalent ways, depending on whether one chooses to work with density matrices, state vectors or operators. We adopt the following mathematical approach to describe the coupling between a quantum system and its environment:

$$|e\rangle|\phi\rangle \rightarrow \sum_{\mathrm{s}} |e_{\mathrm{s}}\rangle M_{\mathrm{s}}|\phi\rangle, \tag{3.1}$$

where $|e\rangle$ ($|\phi\rangle$) is the initial state of the environment (the system), the final environment states $|e_{\mathrm{s}}\rangle$ are not necessarily orthogonal or normalized, and the operators M_{s} acting on the system are unitary. Tensor product symbols $\otimes$ are not written in order to keep the equations uncluttered. So far there is no loss of generality. Since the environment cannot be controlled, the noise process (3.1) is irreversible.

To perform quantum error correction, we now couple the system q to another system called an ancilla a, prepared in some known state $|0\rangle_a$. The unitary interaction A between q and a is carefully arranged to have the following property:

$$A(|0\rangle_a M_{\mathrm{s}}|\phi\rangle) = |s\rangle_a M_{\mathrm{s}}|\phi\rangle, \tag{3.2}$$

where the states $|s\rangle_a$ of the ancilla are orthonormal (but see below), for all M_{s} appearing in (3.1), where this is possible, or else for the dominant terms in (3.1). This interaction is termed *syndrome extraction*; the syndrome s will give us information about the noise. It is very significant that after syndrome extraction the state of the ancilla $|s\rangle_a$ depends on the noise, but not on the quantum state $|\phi\rangle$ to be corrected. This is a highly unusual property; it will only be possible for some states and for some noise terms. The set of error operators M_{s} for which syndrome extraction works will be called the set $\mathcal{S}$ of correctable errors. A major part of the theory of QEC is to identify the best syndrome extraction operators.

Applying A to the noisy state on the right-hand side of (3.1), we obtain

$$A \sum_{\mathrm{s}} |e_{\mathrm{s}}\rangle|0\rangle_a M_{\mathrm{s}}|\phi\rangle = \sum_{\mathrm{s}} |e_{\mathrm{s}}\rangle|s\rangle_a M_{\mathrm{s}}|\phi\rangle. \tag{3.3}$$

Now measure the ancilla in the basis $\{|s\rangle_a\}$. This projects the ancilla onto one particular state $|s\rangle_a$, and yields as a measurement outcome the value of s. Hence the whole system q plus a plus the environment is projected onto $|e_{\mathrm{s}}\rangle|s\rangle_a M_{\mathrm{s}}|\phi\rangle$, and furthermore we know the value of s. However, s is in one-to-one correspondence with M_{s}, so we can deduce M_{s} from s. Armed with this knowledge, we apply M_{s}^{-1} to q by means of a sequence of quantum gates, thus producing the final state

$$|e_{\mathrm{s}}\rangle|s\rangle_a|\phi\rangle. \tag{3.4}$$

The state of q has now been corrected, the state of the environment is immaterial, and we can reprepare a in $|0\rangle_a$ for further use. Actually, the requirement on s is slightly less restrictive than was just stated: rather than a one-to-one correspondence between

s and $M_{\rm s}$, if a set of error operators $\{M_j\}$ produces the same syndrome s, then it is sufficient that there exists a corrective operation $\tilde{M}_{\rm s}$ in one-to-one correspondence with s, satisfying $\langle \phi' | \tilde{M}_{\rm s} M_j | \phi \rangle = \langle \phi' \mid \phi \rangle$ for all states $|\phi\rangle, |\phi'\rangle$ under consideration, for all errors $\{M_j\}$ of syndrome s (cf. equation (3.7) below).

It is not strictly necessary to measure the ancilla, since after syndrome extraction one could arrange a further unitary interaction C between q and a, defined by $C(|s\rangle_a |\phi\rangle) = |s\rangle_a \tilde{M}_{\rm s} |\phi\rangle$. The final state would then be $|\phi\rangle \sum_{\rm s} |e_{\rm s}\rangle |s\rangle_a$. The entanglement between q and the environment is thus transferred to an entanglement between a and the environment. However, measurement of the ancilla has practical advantages compared to the use of CA.

The complete unitary operation

$$\mathcal{R}(|\alpha\rangle M_{\rm s} |\phi\rangle) = |\alpha_{\rm s}\rangle |\phi\rangle \tag{3.5}$$

is called *recovery*, where $|\alpha\rangle$ and $|\alpha_{\rm s}\rangle$ are states of any relevent systems other than q, including the environment, an ancilla and any measuring apparatuses involved.

The central problem of quantum error correction is to identify sets of states $|\phi\rangle$, and the syndrome extraction operator A, which permit recovery for the dominant noise terms $M_{\rm s}$ in the system under consideration. Note that it is sufficient to find an orthonormal set of recoverable states in order to have a recoverable Hilbert space. This is a subspace of the total Hilbert space available to the system. Error correction can be understood as a projection of the system onto the recoverable Hilbert space, or onto an orthogonal Hilbert space, followed by rotation back to the recoverable Hilbert space in the latter case.

The recoverable Hilbert space is spanned by an orthonormal set of quantum states $\{|u\rangle\}$ called *quantum codewords*. It is not hard to prove that the syndrome extraction and recovery operations described above are possible if and only if the codewords have the following properties, for all $\{M_{\rm s} \in \mathcal{S}\}$ (Knill & Laflamme 1997; Bennett *et al*. 1996*b*):

$$\langle u | M_i^\dagger M_j | v \rangle = 0, \tag{3.6}$$

$$\langle u | M_i^\dagger M_j | u \rangle = \langle v | M_i^\dagger M_j | v \rangle, \tag{3.7}$$

for $\langle u \mid v \rangle = 0$.

To illustrate the above ideas, we will now describe a simple example using a system q of three qubits, and an ancilla a of two qubits. For illustrative purposes, we restrict the noise to the simple case that the only operators $M_{\rm s}$ appearing in (3.1) are the identity, and Pauli σ_x operators acting on a single qubit in q. In this case there are two orthonormal recoverable states, namely $|000\rangle$ and $|111\rangle$, so we have a two-dimensional recoverable Hilbert space, which is sufficient to store one qubit of quantum information. A general recoverable state of q is $|\phi\rangle = a|000\rangle + b|111\rangle$ where a and b are general coefficients. The noise process entangles q with its environment, producing the noisy state

$$|e_0\rangle(a|000\rangle + b|111\rangle) + |e_1\rangle(a|001\rangle + b|110\rangle) \\ + |e_2\rangle(a|010\rangle + b|101\rangle) + |e_3\rangle(a|100\rangle + b|011\rangle).$$

The syndrome extraction operator A consists of four controlled-NOT or XOR gates, in which q acts as control and a as target (see figure 1). After this syndrome extrac-

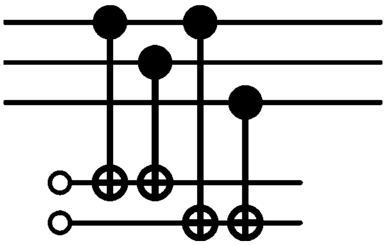

Figure 1. Example syndrome extraction for the three-bit quantum code $\{|000\rangle, |111\rangle\}$. The upper three qubits are the system q, the lower two are the ancilla a.

tion, the total state of environment, system and ancilla is

$$|e_0\rangle|00\rangle_a(a|000\rangle + b|111\rangle) + |e_1\rangle|01\rangle_a(a|001\rangle + b|110\rangle) \\ + |e_2\rangle|10\rangle_a(a|010\rangle + b|101\rangle) + |e_3\rangle|11\rangle_a(a|100\rangle + b|011\rangle).$$

We now measure a. If the measurement result is 00, then by chance the system has been projected back onto the noise-free state and nothing further is necessary. If the result is 01, 10, or 11, then we apply the σ_x operator to, respectively, the first, second, or third qubit in q. After this, the state has been returned to its noise-free form $|\phi\rangle$, irrespective of the values of a and b, and irrespective of which syndrome was obtained. Alternatively, the corrective operation can be applied without measurement by means of Toffoli (controlled-controlled-NOT) gates in which a acts as control and q as target. This three-bit code is analysed more thoroughly in Steane (1996*b*).

4. Classical error correction

To tackle the central problem of QEC, that of identifying recoverable states (quantum codewords) and the syndrome extraction operator, we first need to understand classical error correction. This is a large subject; a full introduction may be found in many readily available textbooks (Hamming 1986; MacWilliams & Sloane 1977; Jones 1979; Hill 1986). Only the most simple ideas will be given here, but they will be sufficient to open the way to powerful QEC methods.

Classical communication can be considered without loss of generality to consist in the transmission of binary words from A to B. A classical *error-correcting code* is a set of words $C = \{u\}$. We will restrict attention to the case that all the words have the same length (number of bits) n. Errors consist of bit flips. To treat errors and their correction, it is convenient to regard each binary word as an n-component vector in an n-dimensional vector space. This is a binary vector space since the components of the vectors can only be 0 or 1. Vector addition $w = u + v$ is carried out by adding the components modulo 2, so is equivalent to the EXCLUSIVE-OR operation carried out bitwise between the words. In this treatment, a general error process can be written $u \to u' = u + e$ where e is an n-component error vector which specifies which bits of u are flipped by the noise. The goal of error correction is to recover u from u'.

The weight $\mathrm{wt}(u)$ of a vector is defined to be the number of non-zero coordinates, e.g. $\mathrm{wt}(11001000) = 3$. The *distance* $d(u, v)$ between two vectors is the number of places where they differ, that is, where one has a 1 and the other a 0, e.g. $d(11001000, 01001100) = 2$. One may show that $d(u, v) = \mathrm{wt}(u + v)$.

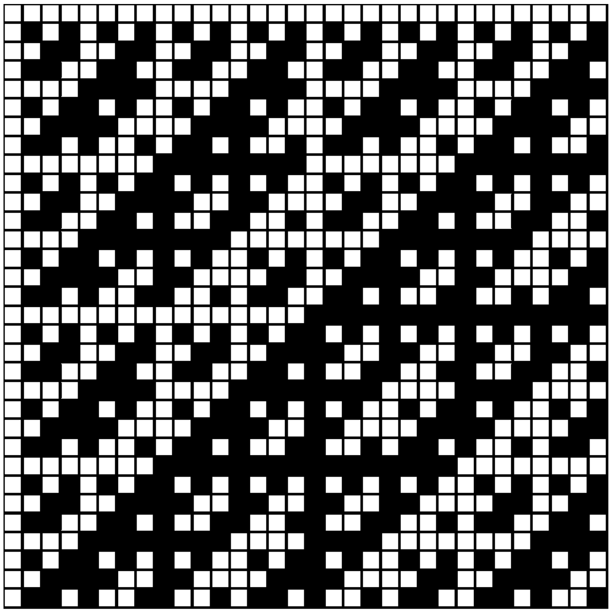

Figure 2. Example classical error-correcting code, the $[32, 5, 16]$ Reed–Muller code. The zeros and ones of the 32 codewords are represented by blank and filled squares.

A given error-correcting code C allows correction of a certain set $\mathcal{S} \equiv \{e_i\}$ of error vectors. The correctable errors are those which satisfy

$$u + e_i \neq v + e_j, \quad \forall u, v \in C \quad (u \neq v). \tag{4.1}$$

The case of no error, $e = 0$, is included in $\mathcal{S}$, so that error-free messages are 'correctable'. Equation (4.1) is the classical analogue of the conditions (3.6), (3.7) on quantum codewords. To achieve error correction, we use the fact that each message u is corrupted to $u' = u + e$ for some $e \in \mathcal{S}$. However, the receiver can deduce u unambiguously from $u + e$ since by condition (4.1), no other message v could have been corrupted to $u + e$, as long as the channel only generates correctable error vectors. In practice, a noisy channel causes both correctable and uncorrectable errors, and the problem is to match the code to the channel, so that the errors *most likely* to be generated by the channel are those the code can correct.

An example classical error-correcting code is shown in figure 2. This is the $[32, 5, 16]$ Reed–Muller code. The notation $[n, k, d]$ indicates that the length of the codewords is n bits, there are 2^k of them, and each word differs from all the others in at least d places. The parameter d is called the 'minimum distance' of the code. This code was used to send video signals back from the Mariner space probe, as follows. The intensity level for each pixel in the video signal is represented by a five-bit binary number m. This five-bit message is 'encoded' into 32 bits by choosing the mth codeword, u_m, and sending that. The receiver receives $u' = u_m + e$, where e represents the errors. If the errors are most likely to affect few qubits, then the best policy of the receiver is to assume that the codeword originally sent was whichever one most closely resembles u', that is, the codeword u with the smallest $d(u, u') = \mathrm{wt}(u + u_m + e)$. Since all the u_m are at least distance d from one another, the closest vector to u' will be u_m, whenever the error e affects fewer than $\frac{1}{2}d$ bits, which is the most likely case. Thus the receiver correctly obtains u_m, from which the message m can be deduced.

So far we have described error correction by the most obvious method, that of examining the received word u' and deducing which sent word was most likely to have produced it. This is permissible for classical information, which can be examined at will, but a more subtle method is necessary for quantum information. Fortunately, such a method is already available within classical coding theory. This is the concept of *parity checking*.

All the codewords in figure 2 have even weight, i.e. an even number of '1's. This property is called even parity. A single bit error, or an odd number of errors, will change the parity to odd. Therefore it is sufficient to evaluate the parity to be able to know whether an even or odd number of errors has taken place. Furthermore, all the codewords have even weight in the first 16 places, and in the last 16. Thus further information on errors is available by calculating the parity of these subsets of the bits. For this code there are altogether $32-5 = 27$ linearly independent parity checks satisfied by all the members of the code.

Parity checking is handled mathematically as follows. A parity check is simply the inner product (u, v) between two binary vectors, evaluated by multiplying corresponding components and summing the results (modulo 2). The positions of the 1s in one vector determine which bits of the other vector enter into the parity check. If the result is zero (one) the parity of the checked bits is even (odd). If u and v are row vectors then the inner product is uv^{T}, where 'T' denotes the transpose operation (the inner product may also be written $u \cdot v$).

An error-correcting code is termed *linear* when it has the property of closure: the sum of any two vectors in the code is also in the code, $u + v \in C\ \forall u, v \in C$. Such an $[n, k, d]$ code is completely defined by any set of k linearly independent members of the code, which together form the $k \times n$ *generator matrix* G. All the other members of the code can be generated by adding rows of G. If $h \cdot u = 0$ and $h \cdot v = 0$, then $h \cdot (u + v) = h \cdot u + h \cdot v = 0$. Therefore, if all the rows of G satisfy any given parity check h, then so do all members of the code. It can be shown that for a linear code there are $n - k$ linearly independent parity checks. These together define the $(n - k) \times n$ *parity check matrix* H. By definition,

$$HG^{\mathrm{T}} = \mathbf{0}, \tag{4.2}$$

where G^{T} is the transpose of G, and $\mathbf{0}$ is the $(n - k) \times k$ zero matrix.

We can now state the feature of classical error correction which is essential to allow a generalization to quantum error correction: the receiver can correct the received word u' without obtaining any information on which word was sent. The method is for the receiver not to examine $u' = u_m + e$ directly, but to evaluate all the parity checks:

$$\begin{aligned} H(u_m + e)^{\mathrm{T}} &= Hu_m^{\mathrm{T}} + He^{\mathrm{T}} \\ &= He^{\mathrm{T}}. \end{aligned} \tag{4.3}$$

The $n-k$ component vector He^{T} is called the error syndrome; it has two important properties: it depends only on the error e, not the codeword sent u_m, and e can be deduced from He^{T} for all correctable errors $e \in \mathcal{S}$. The former property is shown by (4.3), the latter is easily proved by using (4.1) and counting vectors with the same syndrome. The error syndrome therefore permits error correction without learning anything about which codeword was sent.

Let us conclude this section with another example of a linear binary code which will be important in what follows. It is a $[7,4,3]$ code discovered by Hamming (1950). The generator matrix is

$$G = \begin{pmatrix} 1010101 \\ 0110011 \\ 0001111 \\ 1110000 \end{pmatrix}, \tag{4.4}$$

so the 16 members of the code are

$$\begin{array}{llll} 0000000 & 1010101 & 0110011 & 1100110 \\ 0001111 & 1011010 & 0111100 & 1101001 \\ 1110000 & 0100101 & 1000011 & 0010110 \\ 1111111 & 0101010 & 1001100 & 0011001. \end{array} \tag{4.5}$$

These have been written in the following order: first the zero vector, then the first row of G. Next add the second row of G to the two vectors so far obtained, then add the third row to the four vectors previously obtained, and so on. Since the minimum distance is 3, this code permits correction of any error affecting at most 1 bit. The parity check matrix is derived from (4.2):

$$H = \begin{pmatrix} 1010101 \\ 0110011 \\ 0001111 \end{pmatrix}. \tag{4.6}$$

5. Digitizing quantum noise

So far we have identified syndrome extraction, equation (3.2), as an essential component of QEC, and we have seen in (4.3) its classical counterpart. This hints at how the classical and quantum cases are to be brought together. Before we do this, we need a new and, as it turns out, advantageous method to treat noise in quantum systems. Any interaction between a single qubit and its environment can be written in the form

$$|e\rangle(a|0\rangle + b|1\rangle) \rightarrow a(c_{00}|e_{00}\rangle|0\rangle + c_{01}|e_{01}\rangle|1\rangle) + b(c_{10}|e_{10}\rangle|1\rangle + c_{11}|e_{11}\rangle|0\rangle), \tag{5.1}$$

where $|e_{\dots}\rangle$ denotes states of the environment (not necessarily orthogonal) and $c_{\dots}$ are coefficients depending on the noise. A significant insight is that this general interaction can be written

$$|e\rangle|\phi\rangle \rightarrow (|e_I\rangle I + |e_X\rangle X + |e_Y\rangle Y + |e_Z\rangle Z)|\phi\rangle, \tag{5.2}$$

where $|\phi\rangle = a|0\rangle + b|1\rangle$ is the initial state of the qubit, and $|e_I\rangle = c_{00}|e_{00}\rangle + c_{10}|e_{10}\rangle$, $|e_X\rangle = c_{01}|e_{01}\rangle + c_{11}|e_{11}\rangle$, and so on. The operators I, X, Y, Z are, in the basis $\{|0\rangle, |1\rangle\}$,

$$I = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad Y = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}, \quad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \tag{5.3}$$

I is the identity, X is the Pauli σ_x operator, Z is the Pauli σ_z operator, and $Y \equiv XZ = -\mathrm{i}\sigma_y$. The general problem of error correction of qubits is thus reduced (Steane

1996*b*) to the problem of correcting 'bit-flip' errors (X), or 'phase' errors (Z), or both (Y). For example, phase decoherence $|e\rangle(a|0\rangle + b|1\rangle) \rightarrow a|0\rangle|e_1\rangle + b|1\rangle|e_2\rangle$ is given by $\{|e_X\rangle = |e_Y\rangle = 0,\ |e_I\rangle, |e_Z\rangle \neq 0,\ \langle e_I \mid e_Z\rangle = 0\}$. Spontaneous emission,

$$|e\rangle(a|0\rangle + b|1\rangle) \rightarrow a|e_1\rangle|0\rangle + b(|e_2\rangle|1\rangle + |e_3\rangle|0\rangle), \quad \langle e_2 \mid e_3\rangle = 0, \tag{5.4}$$

is given by $\{|e_Y\rangle = -|e_X\rangle,\ \langle e_X \mid e_I\rangle = -\langle e_X \mid e_Z\rangle,\ \langle e_I \mid e_Z\rangle = \langle e_X \mid e_X\rangle\}$. If the spontaneous emission rate is Γ, then $\langle e_X \mid e_X\rangle = \langle e_Y \mid e_Y\rangle = \langle e_Z \mid e_Z\rangle = \frac{1}{4}(1 - \mathrm{e}^{-\Gamma t})$, and $\langle e_I \mid e_I\rangle = \frac{1}{4}(1 + 3\mathrm{e}^{-\Gamma t})$. Thus a continuous error process is 'digitized' into bit and phase errors. The continuous nature of the process enters in the norms of the environment states; these are associated with the probability of obtaining each type of error.

Combining equations (3.1) and (5.2), a general noise process for a set of many qubits can be written in the form (3.1) such that each error operator M_s is a tensor product of single-qubit operators taken from the set $\{I, X, Y, Z\}$. An example error operator for a system of five qubits is

$$M = X_1 I_2 I_3 Z_4 Y_5, \tag{5.5}$$

where the subscripts indicate which qubit the $\{I, X, Y, Z\}$ operator acts on. The following notation will also be useful:

$$M = \boldsymbol{X}_x \boldsymbol{Z}_z, \tag{5.6}$$

where x and z are binary vectors which indicate where the X and Z operators appear in M. For example,

$$X_1 I_2 I_3 Z_4 Y_5 = \boldsymbol{X}_{10001} \boldsymbol{Z}_{00011}. \tag{5.7}$$

The weight of a classical error vector is the number of non-zero elements, which indicates how many bits are flipped by such an error. Similarly, we define the weight of a quantum error operator of the general form illustrated in equation (5.5) to be the number of elements in the product which are not equal to I. When noise acts on different qubits independently, the most important terms in equation (3.1) are those where the error M_s has the smallest weight. Therefore the quantum codes suitable for correcting independent noise are those which correct all errors of weight up to a maximum set as high as possible. It will be our task to find such codes.

(*a*) *Quantum Hamming bound*

A t-error-correcting quantum code is defined to be a code for which all errors of weight less than or equal to t are correctable. Since there are three possible single-qubit errors, the number of error operators of weight t acting on n qubits is $3^t C(n, t)$. Therefore a t-error-correcting code must be able to correct $\sum_{i=0}^{t} 3^i C(n, i)$ error operators. An important class of codes is those in which condition (3.7) is satisfied through both sides being equal to zero. Such codes are sometimes called non-degenerate, or orthogonal, since every erroneous (or error-free) codeword $M_j|u\rangle$ is orthogonal to every other erroneous version of itself $M_i|u\rangle$ as well as to every erroneous version of other codewords $M_i|v\rangle$. All these orthogonal states can only fit into the 2^n-dimensional Hilbert space of n qubits if

$$m \sum_{i=0}^{t} 3^i C(n, i) \leqslant 2^n. \tag{5.8}$$

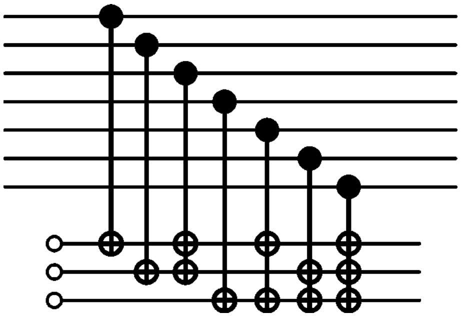

Figure 3. Syndrome extraction operation for $[\![7,1,3]\!]$ CSS code. A control with several NOTs represents several controlled-NOT operations with the same control qubit. Each of the three qubits of the ancilla begin in $|0\rangle$ and finish in $|0\rangle$ or $|1\rangle$ accordingly as the relevant parity check in H (equation (4.6)) is or is not satisfied. A further ancilla of three more qubits is required to complete the syndrome, using a similar network together with R operations (see equation (7.3).

This bound is known as the quantum Hamming bound (Knill & Laflamme 1997; Ekert & Macchiavello 1996; Bennett *et al.* 1996*b*; Gottesman 1996*a*). For $m = 2^k$ and large n, t it becomes

$$\frac{k}{n} \leqslant 1 - \frac{t}{n}\log_2 3 - H\left(\frac{t}{n}\right), \tag{5.9}$$

where $H(x) \equiv -x\log_2(x) - (1-x)\log_2(1-x)$. The rate k/n falls to zero at $t/n \simeq 0.18929$.

A code with $m = 2$ codewords represents a correctable Hilbert space of one qubit (it ‘encodes’ a single qubit). Putting $m = 2$ and $t = 1$ in the quantum Hamming bound, we have $1 + 3n \leqslant 2^{n-1}$ which is saturated by $n = 5$, and indeed a five-qubit code exists; it was first discovered independently by Laflamme *et al.* (1996) and Bennett *et al.* (1996*b*).

6. Quantum use of classical codes

The next stage in the argument is to use the classical syndrome extraction in a quantum system. For this we will need to evaluate a parity check by means of unitary operators. Suppose q is in a product state $|u\rangle$ where u is a binary vector, and we wish to evolve a single-qubit ancilla a from the state $|0\rangle_a \rightarrow |h\cdot u\rangle_a$, where h is a parity check vector. This is accomplished by operating a sequence of XOR gates, with qubits in q as controls, using the single qubit a as target. There are $\mathrm{wt}(h)$ gates, and the control qubits in q are those specified by the 1s in h (see figures 1 and 3). We will write such an operation

$$\mathrm{XOR}^{(h)}_{q,a}\,|0\rangle_a|u\rangle = |h\cdot u\rangle_a|u\rangle. \tag{6.1}$$

Suppose a quantum system is subject only to noise which causes ‘bit-flip’ errors, that is the error operators in (3.1) are $M_{\mathrm{s}} = \boldsymbol{X}_{\mathrm{s}}$. This is an artificial type of noise, but is closely related to a type commonly encountered, as will be shown shortly. In this case the quantum codewords are simply product states $|u\rangle$ in which u is a binary vector in a classical code, $u \in C$. A general state $|\psi\rangle$ in the recoverable Hilbert space can be written

$$|\psi\rangle = \sum_{u\in C} a_u |u\rangle. \tag{6.2}$$

The effect of noise (3.1) is

$$\begin{aligned}|e\rangle \sum_{u\in C} a_u|u\rangle &\rightarrow \sum_{\mathrm{s}} |e_{\mathrm{s}}\rangle \sum_{u\in C} a_u \boldsymbol{X}_{\mathrm{s}} |u\rangle \\ &= \sum_{\mathrm{s}} |e_{\mathrm{s}}\rangle \sum_{u\in C} a_u |u+s\rangle. \end{aligned} \tag{6.3}$$

Note that since $\boldsymbol{X}_{\mathrm{s}}|u\rangle = |u+s\rangle$, the quantum codeword conditions (3.6), (3.7) are satisfied for exactly the set $\{X_{\mathrm{s}}\} : s \in \mathcal{S}$ where $\mathcal{S}$ is the set of classical errors correctable by C (equation (4.1)).

The quantum syndrome extraction operation is obtained from the parity check matrix H of the classical code. We use an $n-k$ qubit ancilla, and evaluate the $n-k$ parity checks of the classical code, using the method described previously (6.1) and illustrated in figure 3. This has the effect

$$A|0\rangle_a|u+s\rangle = \mathrm{XOR}^{(H)}_{q,a}\, |0\rangle_a |u+s\rangle = |Hs^{\mathrm{T}}\rangle_a |u+s\rangle, \tag{6.4}$$

where we used (4.3). It is here that we rely on the fact that the classical syndrome is independent of the codeword, since then A has the correct effect in equation (6.3) no matter what state $|\psi\rangle$ (in the recoverable Hilbert space) we are concerned with. Recovery is completed by measuring a, deducing s from Hs^{T}, and applying $\boldsymbol{X}_{\mathrm{s}}^{-1} = \boldsymbol{X}_{\mathrm{s}}$ to q.

The noise process just considered is quite unusual in practice. However, it is closely related to a common type of noise, namely phase decoherence. Phase decoherence takes the form (3.1) with error operators consisting of tensor products of I and Z, so $M_{\mathrm{s}} = \boldsymbol{Z}_{\mathrm{s}}$. The quantum error-correcting code for this case is now simple to find because $Z = RXR$, where R is the Hadamard or basis change operation

$$R = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}. \tag{6.5}$$

The letter H is often used for this operator, but R is adopted here to avoid confusion with the parity check matrix. We will use $\boldsymbol{R} = R_1 R_2 \ldots R_n$ to denote Hadamard rotation of all the n qubits in q.

The quantum codewords in the case of phase noise are $|c_u\rangle = \boldsymbol{R}|u \in C\rangle$ where C is a classical error-correcting code. An ancilla is prepared in the state $|0\rangle_a$, and the syndrome extraction operation is $\boldsymbol{R}\,\mathrm{XOR}^{(H)}_{q,a}\,\boldsymbol{R}$ where H is the parity check matrix of C. This can be understood as exactly the same code and extraction as the previous example, only now all operations are carried out in the basis $\{R|0\rangle, R|1\rangle\}$ instead of $\{|0\rangle, |1\rangle\}$.

The simplest example of a phase-error-correcting code is a single-error-correcting code using three qubits. The two quantum codewords are

$$\begin{aligned}\boldsymbol{R}|000\rangle &= |000\rangle + |001\rangle + |010\rangle + |011\rangle + |100\rangle + |101\rangle + |110\rangle + |111\rangle, \\ \boldsymbol{R}|111\rangle &= |000\rangle - |001\rangle - |010\rangle + |011\rangle - |100\rangle + |101\rangle + |110\rangle - |111\rangle,\end{aligned}$$

where the normalization factor $1/\sqrt{8}$ has been ommitted. An equivalent code is one using codewords $\boldsymbol{R}(|000\rangle \pm |111\rangle)$ (Steane 1996*b*).

7. Correcting general quantum noise

To correct bit errors $\boldsymbol{X}_x$, we have seen that classical coding and correction methods can be translated fairly directly into the quantum context, and to correct phase errors $\boldsymbol{Z}_z$ the same ideas are used combined with a basis rotation $\boldsymbol{R}$. However, so far a combined error $\boldsymbol{X}_z\boldsymbol{Z}_z$ cannot be corrected, even if it only affects a single qubit.

The code construction and correction method discovered by Calderbank & Shor (1996) and Steane (1996*a*, *b*) works by separately correcting the X and Z errors contained in a general error operator $M_{\mathrm{s}} = \boldsymbol{X}_x\boldsymbol{Z}_z$. The essential element in the code construction is the 'dual code theorem' (Steane 1996*a*):

$$\boldsymbol{R}\sum_{i\in C}|i\rangle = \sum_{i\in C^{\perp}}|i\rangle. \tag{7.1}$$

The normalization has been omitted from this expression in order to make apparent the significant features; normalization factors will be dropped hereafter since they do not affect the argument. The content of (7.1) is that if we form a state by superposing all the members of a linear classical code, then the Hadamard transformed state is a superposition of all the members of the dual code, which is just another linear classical error-correcting code. We can correct both X and Z errors by using states like those in (7.1), as long as both C and $C^{\perp}$ have good error correction abilities.

The dual $C^{\perp}$ is defined to be the set of all vectors v satisfying $v\cdot u = 0\ \forall u \in C$. From this it can be deduced that the generator of C is the check matrix of $C^{\perp}$, and vice versa. Hence if $C = [n, k, d]$, then $C^{\perp} = [n, n-k, d^{\perp}]$. The relationship between d and $d^{\perp}$ is non-trivial, but in general if $k > \frac{1}{2}n$ then $d < d^{\perp}$.

The quantum codewords of a Calderbank, Shor and Steane (CSS) code are obtained from a pair of classical codes where one contains the dual of the other, $C_2^{\perp} \subset C_1$. The first codeword is simply the superposition of all the members of $C_2^{\perp}$ (right-hand side of (7.1)). The other codewords are equal to other subsets of C_1 obtained by displacing $C_2^{\perp}$ by linearly independent members of C_1:

$$|c_u\rangle = \sum_{i\in C_2^{\perp}}|i+u\rangle, \tag{7.2}$$

where $u \notin C_2^{\perp}$, $u \in C_1$ and $C_2^{\perp} \subset C_1$. A particularly interesting case is when $C_2 = C_1$, that is, the code $C_1 \equiv C$ contains its own dual. Since $C = [n, k, d]$ and $C^{\perp} = [n, n-k, d^{\perp}]$, the number of linearly independent u which generate new states is $k-(n-k) = 2k-n$. We can construct 2^{2k-n} orthonormal quantum codewords, which represent a Hilbert space large enough to store $2k-n$ qubits. We will show that the resulting code can correct all errors of weight less than $\frac{1}{2}d$. The parameters of the quantum code are thus $[\![n, 2k-n, d]\!]$.

To correct a CSS code obtained from a classical code $C = [n, k, d]$, $C^{\perp} \in C$ we introduce two ancillae $a(x)$ and $a(z)$, each of $n-k$ qubits, prepared in the state $|0\rangle$. The syndrome extraction operation is

$$(\boldsymbol{R}\,\mathrm{XOR}^{(H)}_{q,a(z)}\,\boldsymbol{R})\,\mathrm{XOR}^{(H)}_{q,a(x)}, \tag{7.3}$$

where H is the check matrix of C. The proof that this works correctly for all errors $M_{\mathrm{s}} = \boldsymbol{X}_x\boldsymbol{Z}_z$ of weight less than $\frac{1}{2}d$ is left as an exercise for the reader. It is straightforward through use of the relations

$$\boldsymbol{X}_x\boldsymbol{Z}_z = (-1)^{x\cdot z}\boldsymbol{Z}_z\boldsymbol{X}_z, \tag{7.4}$$

$$\mathrm{XOR}^{(H)}_{q,a}\,\boldsymbol{Z}_z = \boldsymbol{Z}_z\,\mathrm{XOR}^{(H)}_{q,a}, \tag{7.5}$$

$$\boldsymbol{R}\boldsymbol{X}_{\mathrm{s}} = \boldsymbol{Z}_{\mathrm{s}}\boldsymbol{R}, \tag{7.6}$$

$$\boldsymbol{R}\sum_{i\in C^{\perp}}|i+u\rangle = \sum_{i\in C}(-1)^{i\cdot u}|i\rangle, \tag{7.7}$$

where the latter follows from (7.1) and (7.6).

The simplest CSS code is obtained from the $[7,4,3]$ Hamming code given at the end of § 4. This is single-error-correcting and contains its dual, and therefore leads to a single-error-correcting quantum code of parameters $[\![7,1,3]\!]$. The two codewords are (cf. (4.5))

$$|c_0\rangle = |0000000\rangle + |1010101\rangle + |0110011\rangle + |1100110\rangle + |0001111\rangle + |1011010\rangle + |0111100\rangle + |1101001\rangle, \tag{7.8}$$

$$|c_1\rangle = \boldsymbol{X}_{1111111}|c_0\rangle. \tag{7.9}$$

8. Generalization: stabilizer codes

Quantum codewords are multiple-qubit entangled states with many symmetry properties. For example, the codewords of a CSS code are eigenvectors of all operators of the form $\boldsymbol{X}_a\boldsymbol{Z}_b$ for $a, b \in C^{\perp}$, with eigenvalue 1.

Proof. If $|i\rangle$ is a product state, then $\boldsymbol{Z}_b|i\rangle = (-1)^{b\cdot i}|i\rangle$. If $b \in C^{\perp}$, then $b\cdot i = 0$ $\forall i \in C$, therefore, $\boldsymbol{Z}_b|i\rangle = |i\rangle$. A CSS quantum codeword $|c_u\rangle$ is a superposition of states $|i \in C\rangle$, therefore $\boldsymbol{Z}_b|c_u\rangle = |c_u\rangle$. A similar argument applies for $\boldsymbol{X}_a$, using $\boldsymbol{X}_a = \boldsymbol{R}\boldsymbol{Z}_a\boldsymbol{R}$. ■

The argument just given points to a way to generalize the theory of construction of quantum error-correcting codes. The central ideas were discovered independently by Gottesman (1996*a*) and Calderbank *et al.* (1997); the latter authors developed the most general treatment which is elaborated in Calderbank *et al.* (1996) (see also Steane 1996*d*). A QEC code can be defined to be an orthonormal set of states which are simultaneous eigenstates, with eigenvalue 1, of all the operators in a group $\mathcal{H}$ called the stabilizer of the code. For 'additive' codes, the stabilizer has $n-k$ linearly independent members, all of the form $\boldsymbol{X}_x\boldsymbol{Z}_z$; these error operators generate the stabilizer group. For a CSS code they are $\{\boldsymbol{X}_h\boldsymbol{Z}_0, \boldsymbol{X}_0\boldsymbol{Z}_h\}$ where h is a row of the parity check matrix of the associated classical code. For orthogonal (non-degenerate) codes, the set of correctable errors $\mathcal{S}$ is any set with the property that $M_{\mathrm{s}}^{\dagger}M_t$ anticommutes with a member of $\mathcal{H}$, for all members $M_{\mathrm{s}}, M_t \in \mathcal{S}$.

Proof. where $H \in \mathcal{H}$ is the anticommuting operator (Gottesman 1996*a*),

$$\langle u|M_{\mathrm{s}}^{\dagger}M_t|v\rangle = \langle u|M_{\mathrm{s}}^{\dagger}M_tH|v\rangle = -\langle u|HM_{\mathrm{s}}^{\dagger}M_t|v\rangle = -\langle u|M_{\mathrm{s}}^{\dagger}M_t|v\rangle, \tag{8.1}$$

therefore $\langle u|M_{\mathrm{s}}^{\dagger}M_t|v\rangle = 0$, so the code conditions (3.6), (3.7) are satisfied. ■

The simplest useful code which is not a CSS code is the $[\![5,1,3]\!]$ code, the stabilizer group of which is generated by $\{\boldsymbol{X}_{11000}\boldsymbol{Z}_{00101}, \boldsymbol{X}_{01100}\boldsymbol{Z}_{10010}, \boldsymbol{X}_{00110}\boldsymbol{Z}_{01001}, \boldsymbol{X}_{00011}\boldsymbol{Z}_{10100}\}$. Since this code permits correction of all error operators of weight one or zero, it is a 'single-error-correcting code', and the minimum distance is 3 (Bennett *et al*. 1996*b*; Laflamme *et al*. 1996; Calderbank *et al*. 1997).

9. The power of error correction

We have now developed quite general methods to find recoverable quantum states and the associated syndrome extraction operation. Although the CSS construction is not the most general, it is powerful because it proves the existence of, and allows the construction of, efficient codes, in the following sense. A standard result of classical coding theory is the Gilbert–Varshamov bound (MacWilliams & Sloane 1977), which states that there exists a linear $[n, k, d]$ classical code provided

$$2^k \sum_{i=0}^{d-2} \binom{n-1}{i} < 2^n. \tag{9.1}$$

It can be shown further that there exists a linear $[n, k, d]$ code containing its own dual provided both it and the dual satisfy the bound. This means that for large n, k, d, the rate k/n of a CSS code can exceed a finite bound as $n \rightarrow \infty$ with fixed d/n:

$$k/n > 1 - 2H(d/n). \tag{9.2}$$

Suppose the qubits of a quantum computer or communication channel are subject to independent noise. In this case we can define a probability p that after a fixed time any qubit will suffer an error (X, Y or Z). The precise meaning is that p is the probability that when the syndrome is measured, the computer's state will be projected onto a Hilbert space which differs from the recoverable Hilbert space by a rotation X, Y or Z applied to any chosen qubit. In the case of spontaneous emission, $p = 3(1-\mathrm{e}^{-\Gamma t})$. For large n, the number of qubits which are found to be erroneous will be close to the mean np. The probability that more than $(1+\epsilon)np$ qubits are erroneous tends to zero as $n \rightarrow \infty$, for any $\epsilon > 0$, by the law of large numbers. Therefore error correction is almost certain to succeed, for large n, as long as $d > 2np(1+\epsilon)$. Since k/n can exceed the Gilbert–Varshomov bound, we can achieve the required distance without needing k/n to become smaller and smaller as n increases. This argument is closely related to Shannon's main theorem in classical information theory.

To estimate the success of QEC with finite resources, we need to calculate the effect of those error operators in (3.1) which are *not* correctable, for a given QEC code. One must always design the code to fit the noise, so that the uncorrectable errors are the least likely. We have concentrated on the case of t-error correcting codes, designed to cope with independent noise, in order to illustrate the methods. A useful measure of the degree to which a state has been successfully recovered is the fidelity f of the recovered state compared with the state before the noise occurred. Since QEC causes all the correctable error operators to have the same effect on q as the identity, the fidelity f is of order $f = 1 - P$, where P is the probability of an uncorrectable error. For a t-error-correcting code, and independent noise, the most likely uncorrectable errors are those which affect $t+1$ qubits, so $P \simeq C(n, t+1)p^{t+1}(1-p)^{n-t-1}$ where $p \ll 1$ is the single-qubit error probability defined in the previous paragraph, and

$C(n,x)$ is a binomial coefficient, $C(n,t+1) \simeq n^{t+1}/(t+1)!$ for $n \gg t$. Thus for small p and large n,

$$f \simeq 1 - (np)^{t+1}/(t+1)!. \tag{9.3}$$

Although t-error-correcting codes are not designed primarily to cope with systematic errors (noise affecting many qubits together), they nevertheless still perform well under some systematic effects. For example, the noise operators in the code stabilizer, which can have large weight, do not change the codewords at all. In practice, quantum computers and communication channels will be subject to a combination of systematic and uncorrelated noise. To deal with this situation different QEC codes can be combined, borrowing methods from classical coding theory, such as code concatenation (Knill & Laflamme 1996).

10. Discussion and recent developments

Quantum error correction is an information-theoretic result. That is to say, QEC does not allow us to protect all the degrees of freedom of a quantum system from noise. Rather, it allows us to protect the quantum information stored in a quantum system. Also, QEC is not a simple experimental tool, since it uses quantum information processing (syndrome extraction and recovery) in order to combat noise. However, even taking these considerations into account, we have a powerful result. This is illustrated by two apparently paradoxical truths.

Suppose a third party supplies me with a single qubit in some arbitrary quantum state $a|0\rangle + b|1\rangle$ unknown to me. I then cause this qubit to undergo a specific controlled interaction with four others initially prepared in $|0\rangle$. The following is now possible: I can invite you to choose any one of the five qubits, and change its state in any manner you choose, including measuring it or otherwise entangling it with other systems. You do not need to tell me which qubit you changed. I can nevertheless extract from the five qubits *exactly* the single-qubit state $a|0\rangle + b|1\rangle$ I started with. The large perturbation you introduced had no effect at all on the stored quantum information! This is achieved by means of the $[\![5,1,3]\!]$ single-error-correcting quantum code.

Suppose a quantum computer is constructed from qubits which all undergo spontaneous decay with lifetime τ (e.g. the $|1\rangle$ state decays to the $|0\rangle$ state). This decay could be electromagnetic, or radioactive, or any such process which it is impossible to influence unless prodigious experimental efforts are made, which we assume are not attempted. We take these unstable qubits to be the only type available, they are used both for the computer and for any ancillae required for QEC. Such a computer could nevertheless perform a quantum computation (with success probability of order 1), relying on sustained interference effects involving, say, 1000 qubits, in which the quantum coherence was preserved during a time of order $10^4\tau$. Since we cannot influence the decay process, this means that quantum coherence is preserved even though every qubit involved has spontaneously decayed (and been re-excited) about 10 000 times!

Whereas the first paradoxical result follows directly from the coding theory discussed in this paper, the second is less obvious and requires the further concepts of *fault-tolerant* computing. The main ideas were introduced by Shor (1996), and subsequently generalized by several authors (see Preskill 1998, and references therein).

The fault-tolerant methods permit syndromes to be extracted reliably even though every quantum gate and every qubit is noisy. The example figures quoted are based on the fault-tolerant method described in Steane (1997*a*), using the $[\![55, 1, 11]\!]$ code. This is a $t = 5$-error-correcting code, using $n = 55$ qubits to store each logical qubit required for the quantum computation.

The main points of the analysis in Steane (1997*a*) are as follows. The computer q is repeatedly corrected. Each complete recovery operation (syndrome extraction followed by correction) takes a time $T_{\mathcal{R}} \ll \tau$, so the bit-error probability per qubit per recovery time is $p_1 = 2(1-\exp(-T_{\mathcal{R}}/\tau)) \simeq 2T_{\mathcal{R}}/\tau$, where the factor two instead of three is because on average two-thirds of the errors need correcting by each ancilla a_x, a_z. In the method of Steane (1997*a*), on the order of $2n(t+1) = 660$ quantum gates are required to extract a syndrome for each 55-bit block in the computer. These are carried out one after another (not several in parallel), so $T_{\mathcal{R}} = 660T_{\rm g}$, where $T_{\rm g}$ is the time required for a single gate. Error propagation in the syndrome extraction networks couples errors from the ancilla into the block being corrected (and vice versa). It is found that this doubles the accumulation of errors in q, so the error probability per qubit in q per recovery time is $p = 2640T_{\rm g}/\tau$. We consider the case that $T_{\rm g} = 1.5 \times 10^{-7}\tau$, giving $p \simeq 4 \times 10^{-4}$. The fidelity of the state of any logical qubit just after the jth recovery, compared with the state just after the $j-1$th recovery, is $f_1 \simeq (1 - 1.6 \times 10^{-13})^2$, using equation (9.3), where the power of two accounts for separate extraction of X and Z syndromes. We correct all the 1000 logical qubits in parallel, so the fidelity of the state of the whole computer after the $j-1$th recovery, compared with the jth, is $f \simeq f_1^{1000}$. Since this is independent of j, the fidelity compared with the computer's initial state is f^j, which falls to 0.9 at $j \simeq 3.5 \times 10^8$. The total computation time is then $jT_{\mathcal{R}} \simeq 3 \times 10^4\tau$, in comfortable agreement with the claim.

A similar analysis can also include the effect of additional noise associated with each quantum gate. If the gate noise is of the order of n times the noise per resting qubit per gate time, it is found that the conclusions are not greatly changed. This represents a gate failure probability of order 10^{-5}, which is experimentally achievable. It is noteworthy that the block CSS codes perform better than concatenated codes even for such long computations (the example given would permit factorization of 130-digit numbers using Shor's algorithm).

A useful physical intuition into QEC is to think of it as a type of cooling, only instead of cooling towards a particular state in Hilbert space, one cools towards a particular subspace of Hilbert space. Cooling is achieved by the flow of entropy from a hotter body to a colder one, and in QEC entropy flows from the system q to the ancilla. The method relies on the low-entropy preparation of the ancilla state (Aharonov & Ben-Or 1996; Steane 1997*b*).

The author is supported by the Royal Society and by St Edmund Hall, Oxford.

References

Aharonov, D. & Ben-Or, M. 1996 Polynomial simulations of decohered quantum computers. *37th Annual Symp. on Foundations of Computer Science*, pp. 46–55.

Barnum, H., Fuchs, C. A., Jozsa, R. & Schumacher, B. 1996 A general fidelity limit for quantum channels. *Phys. Rev.* A **54**, 4707.

Beckman, D., Chari, A., Devabhaktuni, S. & Preskill, J. 1996 Efficient networks for quantum factoring. *Phys. Rev.* A **54**, 1034.

Bennett, C. H., Brassard, G., Popescu, S., Schumacher, B., Smolin, J. A. & Wootters, W. K. 1996*a* Purification of noisy entanglement and faithful teleportation via noisy channels. *Phys. Rev. Lett.* **76**, 722–725.

Bennett, C. H., DiVincenzo, D. P., Smolin, J. A. & Wootters, W. K. 1996*b* Mixed state entanglement and quantum error correction. *Phys. Rev.* A **54**, 3825.

Bennett, C. H., DiVincenzo, D. P. & Smolin, J. A. 1997 Capacities of quantum erasure channels. *Phys. Rev. Lett.* **78**, 3217–3220.

Calderbank, A. R. & Shor, P. W. 1996 Good quantum error-correcting codes exist. *Phys. Rev.* A **54**, 1098–1105.

Calderbank, A. R., Rains, E. M., Shor, P. W. & Sloane, N. J. A. 1996 Quantum error correction via codes over $GF(4)$. quant-ph/9608006.

Calderbank, A. R., Rains, E. M., Shor, P. W. & Sloane, N. J. A. 1997 Quantum error correction and orthogonal geometry. *Phys. Rev. Lett.* **78**, 405.

Chuang, I. L., Laflamme, R., Shor, P. W. & Zurek, W. H. 1995 Quantum computers, factoring and decoherence. *Science* **270**, 1633.

Deutsch, D., Ekert, A., Jozsa, R., Macchiavello, C., Popescu, S. & Sanpera, A. 1996 Quantum privacy amplification and the security of quantum cryptography over noisy channels. *Phys. Rev. Lett.* **77**, 2818.

Dieks, D. 1982 Communication by electron-paramagnetic resonance devices. *Phys. Lett.* A **92**, 271.

Ekert, A. & Macchiavello, C. 1996 Quantum error correction for communication. *Phys. Rev. Lett.* **77**, 2585–2588.

Glauber, R. J. 1986 In *Frontiers in quantum optics* (ed. E. R. Pike & S. Sarker). Bristol: Adam Hilger.

Gottesman, D. 1996*a* Class of quantum error-correcting codes saturating the quantum Hamming bound. *Phys. Rev.* A **54**, 1862–1868.

Gottesman, D. 1996*b* Pasting quantum codes. quant-ph/9607027.

Hamming, R. W. 1950 Error detecting and error correcting codes. *Bell Syst. Tech. Jl* **29**, 147.

Hamming, R. W. 1986 *Coding and information theory*, 2nd edn. Englewood Cliffs, NJ: Prentice-Hall.

Hill, R. 1986 *A first course in coding theory.* Oxford: Clarendon.

Jones, D. S. 1979 *Elementary information theory.* Oxford: Clarendon.

Knill, E. & Laflamme, R. 1996 Concatenated quantum codes. quant-ph/9608012.

Knill, E. & Laflamme, R. 1997 A theory of quantum error-correcting codes. *Phys. Rev.* A **55**, 900.

Laflamme, R., Miquel, C., Paz, J. P. & Zurek, W. H. 1996 Perfect quantum error correcting code. *Phys. Rev. Lett.* **77**, 198.

Lloyd, S. 1997 The capacity of a noisy quantum channel. *Phys. Rev.* A **55**, 1613.

MacWilliams, F. J. & Slaone, N. J. A. 1977 *The theory of error correcting codes.* Amsterdam: Elsevier.

Palma, G. M., Suominen, K.-A. & Ekert, A. K. 1996 Quantum computers and dissipation. *Proc. R. Soc. Lond.* A **452**, 567–584.

Plenio, M. B. & Knight, P. L. 1996 Realistic lower bounds for the factorization time of large numbers on a quantum computer. *Phys. Rev.* A **53**, 2986.

Plenio, M. B. & Knight, P. L. 1997 Decoherence limits to quantum computation using trapped ions. *Proc. R. Soc. Lond.* A **453**, 2017.

Preskill, J. 1998 Reliable quantum computers. *Proc. R. Soc. Lond.* A **454**, 385.

Rains, E. M., Hardin, R. H., Shor, P. W. & Sloane, N. J. A. 1997 Nonadditive quantum code. *Phys. Rev. Lett.* **79**, 953.

Schumacher, B. W. & Nielsen, M. A. 1996 Quantum data processing and error correction. *Phys. Rev.* A **54**, 2629.

Shor, P. W. 1994 Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. In *Proc. 35th Annual Symp. on Foundations of Computer Science, Santa Fe.* Los Alimatos, CA: IEEE Computer Society Press.

Shor, P. W. 1995 Scheme for reducing decoherence in quantum computer memory. *Phys. Rev.* A **52**, R2493–R2496.

Shor, P. W. 1996 Fault tolerant quantum computation. In *Proc. 37th Symp. on Foundations of Computer Science*, pp. 56–65. Los Alimatos, CA: IEEE Computer Science Press.

Shor, P. W. & Laflamme, R. 1996 Quantum analog of the MacWilliams identities in classical coding theory. *Phys. Rev. Lett.* **78**, 1600.

Shor, P. W. & Smolin, J. A. 1996 Quantum error correcting codes need not completely reveal the error syndrome. quant-ph/9604006.

Steane, A. M. 1996*a* Multiple particle interference and quantum error correction. *Proc. R. Soc. Lond.* A **452**, 2551–2577.

Steane, A. M. 1996*b* Error correcting codes in quantum theory. *Phys. Rev. Lett.* **77**, 793.

Steane, A. M. 1996*c* Simple quantum error-correcting codes. *Phys. Rev.* A **54**, 4741–4751.

Steane, A. M. 1996*d* Quantum Reed–Muller codes. *IEEE Trans. Inf. Theory* (quant-ph/9608026). (In the press.)

Steane, A. M. 1997*a* Active stabilization, quantum computation and quantum state synthesis. *Phys. Rev. Lett.* **78**, 2252.

Steane, A. M. 1997*b* Space, time and noise requirements for reliable quantum computing. *Prog. Phys.* (In the press.)

Steane, A. M. 1998 Quantum computing. *Rep. Prog. Phys.* **61**, 117.

Unruh, W. G. 1995 Maintaining coherence in quantum computers. *Phys. Rev.* A **51**, 992–997.

van Enk, S. J., Cirac, J. I. & Zoller, P. 1997 Ideal quantum communication over noisy channels: a quantum optical implementation. *Phys. Rev. Lett.* **78**, 4293.

Wootters, W. K. & Zurek, W. H. 1982 A single quantum cannot be cloned. *Nature* **299**, 802.

Discussion

P. Marcer (*BCS Cybernetic Machine Group, Keynsham, UK*). Is noise an enhancer of coherence, as is the case under the appropriate conditions in, for example, stochastic resonance or superresolution imaging? In superresolution in quantum holography, for example, increasing noise or inhomogeneity can produce a new quantum holographic superposition which increases the resolution of the holographic information previously available. That is, under the right conditions, i.e. phase conjugation, it is self-correcting in the presence of noise.

Th. Beth (*University of Karlsruhe, Germany*). Such methods of 'orthogonal imaging' to reduce the signal-to-noise ratio (Harwit & Sloane 1979) do not apply in the quantum regime. These methods essentially use the L_2-norm of probabilities after measurement, but are not possible in the unitary evolution of a quantum computer (see Beth *et al*. 1998, ch. 8).

The correction of systematic, 'all over' phase errors is done classically by filters. In quantum computation, however, this requires the transform of phases into kets. A method for doing this has recently been proposed by Kitaev (1995).

The technique of the Hadamard transform duality theorem was originally invented by Gleason and MacWilliams. It is Dr Steane, however, who saw the ingenious duality of bit-flips (σ_x) and phase-flips (σ_z) in quantum error-correcting code.

A. M. Steane. Professor Beth has answered Professor Marcer's comment. I will add that, whereas I am not familiar with the method described in holography, I suspect that the information storage going on here is essentially classical in nature. Although holography can store complete information about light fields, this is only possible for sufficiently strong fields which can in effect be 'cloned', i.e. many 'copies' of the quantum state of the light are available to form the hologram. Having said that, quantum error correction typically makes use of imperfect measurements on the relevant system, and I would not rule out the possibility that there exist situations in which introducing noise into a measuring device may improve its usefulness.

Picking up on the comment by Professor Penrose concerning systematic phase error, it is noteworthy that although codes are not necessarily designed with this type of error in mind, they nevertheless can perform very well under such noise (see the comment at the end of § 9).

Finally, Professor Beth is quite right in stressing that the connection between dual codes and the Hadamard transform is a well-known part of classical error-correction theory. The important new insight was, I think, to form a quantum state by superposing all the members of a linear code, and use the Hadamard transform to make intuitively clear that such a state will transform in a useful way under phase-flips as well as bit-flips.

Additional references

Beth, Th., Jungnickel, D. & Lenz, H. 1998 Design theory. In *Encyclopaedia of Mathematics*, 2nd edn. Cambridge University Press.

Harwit, M. & Sloane, N. J. A. 1979 *Hadamard transform optics.* London: Academic.

Kitaev, A. 1995 Quantum measurements and the Abelian stabilizer problem. quant-ph/9511026.