Stabilisation of Quantum Computations by Symmetrisation

Adriano Barenco[1], André Berthiaume[2], David Deutsch[1], Artur Ekert[1], Richard Jozsa[3] and Chiara Macchiavello[1]

[1]*Clarendon Laboratory, University of Oxford, Parks Road, Oxford OX1 3PU, U.K.*

[2]*Centrum voor Wiskunde en Informatica, Amsterdam, The Netherlands.*

[3]*School of Mathematics and Statistics, University of Plymouth, Plymouth, Devon PL4 8AA, U.K.*

Abstract

We propose a method for the stabilisation of quantum computations (including quantum state storage). The method is based on the operation of projection into $\mathcal{SYM}$, the symmetric subspace of the full state space of R redundant copies of the computer. We describe an efficient algorithm and quantum network effecting $\mathcal{SYM}$–projection and discuss the stabilising effect of the proposed method in the context of unitary errors generated by hardware imprecision, and nonunitary errors arising from external environmental interaction. Finally, limitations of the method are discussed.

§1 Introduction

Any realistic model of computation must conform to certain requirements imposed not by the mathematical properties of the model but by the laws of physics. Computations which require an exponentially increasing precision or exponential amount of time, space, energy or any other physical resource are normally regarded as unrealistic and intractable.

Any actual computational process is subject to unavoidable hardware imprecision and spurious interaction with the environment, whose nature is dictated by the laws of physics. These effects introduce errors and destabilise the progress of the desired computation. It is therefore essential to have some method of stabilising the computation against these effects.

For classical computation there is a simple and highly effective method of stabilisation. Each computational variable is represented redundantly using many more physical degrees of freedom than are logically required, and a majority vote (or average) of all the copies is taken followed by resetting all the copies to the majority answer. This process is applied periodically during the course of the computation. If we use R copies and the probability of producing the correct answer is $\frac{1}{2}+\eta$ then it can be shown ([1] page 258) that the probability E that the majority vote is wrong, is less than $\exp(-\eta^2 R/6)$. This is an extremely resource–efficient stabilisation in that the probability of error decreases *exponentially* with the degree of redundancy R. Indeed, suppose that a polynomial–time algorithm runs for M steps, each of which is correct with probability $\frac{1}{2}+\eta$ and majority voting is used after each step. The probability that the final answer will be correct is greater than $(1-E)^M$. Thus any desired success probability $1-\delta$ may be achieved using a degree of redundancy $R = O(\log(M/\delta))$, which is only logarithmic in the input size.

The majority vote method above cannot be applied in the case of quantum computation because quantum algorithms depend essentially on the maintenance of coherent superpositions of different computational states at each step. The laws of quantum mechanics forbid the identification of an unknown quantum state [17] [18] and forbid even the cloning of an unknown state [19]. Thus the majority voting method is inapplicable as we can neither determine the majority state nor reset the remaining copies to that state. In this paper we propose an alternative entirely quantum mechanical method of stabilisation utilising redundancy but which has no classical analogue. We discuss its applicability and limitations. The method was first proposed by Deutsch [2] and a brief outline of its underlying principles was given in [3].

An alternative approach to the stabilisation of classical computation involves the use of error correcting codes [20]. A quantum mechanical generalisation of this approach was recently introduced by Shor [9] and subsequently developed in [13] [10] [12] [11]. These methods are unrelated to those proposed in this paper and provide an interesting supplementary method of stabilising quantum computation.

The process of simply repeating a whole computation a sufficient number of times may serve to stabilise it in certain circumstances. Suppose that we have a quantum algorithm which succeeds with probability $1-\epsilon$ (where ϵ may increase with input size) and suppose that we know when the computation has been successful. For example the computation may produce a candidate factor of an input integer which can then be efficiently checked by trial division. For any input size L the success probability can be amplified to any prescribed level $1-\delta$ by repeating the algorithm a sufficiently large number, R, of times since the probability of at least one success in R repetitions is $1-\epsilon^R \rightarrow 1$ as $R \rightarrow \infty$. Suppose now that the success probability $1-\epsilon$ decreases with input size L as 1/poly(L). Then we can maintain any prescribed success probability by allowing R to increase as a suitable polynomial function of L. Thus if the original algorithm was efficient (i.e. polynomial time) then its R–fold repetition is still efficient i.e. the algorithm has been stabilised in an efficient manner. (Note that Shor's quantum factoring algorithm [7] [8] is of this type with success probability decreasing as $1/L$ with input size.) However if the success probability falls *exponentially* with input size L then we must use $R \sim \exp(\mathrm{L})$ to maintain any constant level of success probability. This implies an exponential increase of physical resources for stabilisation and hence this method is inefficient in this case. Unfortunately just such an exponential decay of success probability appears to be a generic feature of any physical implementation of computation, as described below.

In quantum theory, the issue of preventing information from leaking into the environment from a system ("the computer") is generally known as the decoherence problem [4, 5, 6]. According to the analysis of [6], decoherence generally causes an *exponential* decrease in success probability with input size L. Decoherence is a universal phenomenon and is expected to affect – to some extent – any physical implementation of quantum computation whatever. Thus without some *efficient* form of stabilisation, quantum algorithms which are polynomially efficient in the error-free case (like Shor's factoring algorithm) cannot be considered polynomially efficient in practice.

Consider any efficient computation which gives the correct result at each step with probability $1-\epsilon$ where ϵ is constant. This would typically be the case if each step of a computation consisted of the application of an elementary gate operation having a standard tolerance of error. Then after N steps the probability of success is (at least) $(1-\epsilon)^N \sim \exp(-\epsilon N)$ which again decreases exponentially with N. Suppose that we have a stabilisation scheme utilising redundancy R which reduces the error in each step only by a factor $1/R$ *i.e.* $\epsilon \to \epsilon/R$ (rather than the *exponential* decrease given by classical majority voting). After N steps the probability of success is now $\exp(-\epsilon N/R)$. This can be kept at any prescribed level $1-\delta$ by taking $R = \epsilon N/(-\log(1-\delta))$ which is *polynomial* in N and hence in the input size L. Thus an exponentially growing error (such as results from decoherence) in a polynomial–time computation can be efficiently stabilised by a method which reduces the error per step only as $1/R$ with the degree of redundancy. Our proposed method below will have this property.

It will be useful in the following to keep in mind the simplest possible example of the stabilisation problem where the computer consists of one qubit (*i.e.* one two–level system) and is performing no computation. In fact this simple model captures the essential features of the stabilisation problem for general quantum computations. The problem of stabilisation concerns the time evolution of an "accuracy" observable which has only two eigenvalues. As we shall see our analysis of error correction depends only on such simple observables and is independent of the substance of the computation. Thus we are addressing the problem of stabilising the *storage* of an (unknown) quantum state of one qubit against environmental interaction and (suitably random) imprecision in the construction of the storage device.

§2 The Symmetric Subspace

Our proposed stabilisation method will exploit redundancy but in contrast to the classical majority voting method, it will be based on essentially quantum mechanical properties through use of the symmetric subspace of the full state space of R copies of a physical system. Consider R copies of a quantum system each with state space $\mathcal{H}$. Denote the full state space $\mathcal{H} \otimes \mathcal{H} \otimes \ldots \mathcal{H}$ by $\mathcal{H}^R$.

Remark 1. Here we require that the R copies be distinguishable e.g. being located in separate regions of space so that the position coordinate provides an extra "external" degree of freedom for distinguishability. The state space $\mathcal{H}$ can be thought of as representing the "internal" degrees of freedom of each system. In our application these are the computational degrees of freedom of each replica of the computer. In our notation we suppress explicit mention of the distinguishing degree of freedom which is implicitly given by the written order of component states in a tensor product state (c.f. *remark 2* below).$\Box$

The symmetric subspace $\mathcal{SYM}$ of $\mathcal{H}^R$ may be characterised by either of the two following equivalent definitions.

Definition 1. $\mathcal{SYM}$ is the smallest subspace of $\mathcal{H}^R$ containing all states of the form $|\psi\rangle\,|\psi\rangle \ldots |\psi\rangle$ for all $|\psi\rangle \in \mathcal{H}$. $\Box$

Definition 2. $\mathcal{SYM}$ is the subspace of all states in $\mathcal{H}^R$ which are symmetric (*i.e.* unchanged) under the interchange of states for any pair of positions in the tensor product.(Here we interchange only the internal degrees of freedom leaving the external degrees fixed.) $\Box$

Remark 2. To clarify the notion of symmetrisation in *definition 2* note that, for example, $|\phi\rangle\,|\psi\rangle + |\psi\rangle\,|\phi\rangle \in \mathcal{H}^2$ is in $\mathcal{SYM}$. If we were to show the external degrees of freedom then this state would be written $|\phi; x_1\rangle\,|\psi; x_2\rangle + |\psi; x_1\rangle\,|\phi; x_2\rangle$. Consequently the notion of symmetrisation in *definition 2* is different from bosonic symmetrisation which requires symmetrisation of *all* degrees of freedom. For a pair of bosons the previous state would be $|\phi; x_1\rangle\,|\psi; x_2\rangle + |\psi; x_2\rangle\,|\phi; x_1\rangle$. $\Box$

Definition 1 has the following interpretation. Suppose that we have R copies of a quantum computer. If there were no errors then at each time the joint state would have the form

$$|\psi\rangle\,|\psi\rangle \ldots |\psi\rangle \in \mathcal{H}^R \tag{1}$$

In the presence of errors the states will evolve differently resulting in a joint state of the form $|\psi_1\rangle\,|\psi_2\rangle \ldots |\psi_R\rangle$ or more generally a mixture of superpositions of such states. In quantum mechanics any test ("yes/no" question) that we can apply to a physical system must correspond to a *subspace* of the total state space. States of the form (1) for all $|\psi\rangle \in \mathcal{H}$ do not, by themselves, form a subspace of $\mathcal{H}^R$. According to *definition 1*, $\mathcal{SYM}$ is the smallest subspace containing all possible error–free states.

It thus corresponds to the "most probing" test we can legitimately apply, which will be passed by all error–free states. Recall that we cannot generally identify the actual quantum state during the course of the computation or indeed gain any information about it without causing some irreparable disturbance [18]. The characterisation given in *definition 2* is especially useful in treating mathematical properties of $\mathcal{SYM}$ as below.

The equivalence of the two definitions may be proved by viewing the i^{th} component space in the tensor product $\mathcal{H}^R$ as the space of complex polynomials of degree $\leq n-1$ in the variable x_i (where n is the dimension of $\mathcal{H}$). Then *Definition 1* defines the subspace of all polynomials $p(x_1, \ldots, x_R)$ of degree $\leq n-1$ in each variable, which arise as sums of products of functions of the form

$$f_\alpha(x_1, x_2, \ldots, x_R) = \prod_{i=1}^{R}(x_i - \alpha) \tag{2}$$

for any α (where we have used the fundamental theorem of algebra). On the other hand, *definition 2* defines the space of all symmetric polynomials (of degree $\leq n-1$ in each variable). The equivalence of these subspaces then follows easily from basic properties of the standard elementary symmetric functions [14] which are defined as the coefficients of the powers of α in the expansion of (2).

The equivalence of the two definitions may also be understood via the following illustrative example which gives further insight into the structure of $\mathcal{SYM}$.

Example 1. Suppose that $\mathcal{H}$ is two–dimensional (*i.e.* a qubit) with computational basis states $|0\rangle$ and $|1\rangle$. Consider triple redundancy $R = 3$ and the symmetric subspace $\mathcal{SYM} \subset \mathcal{H}^3$. Let us tentatively denote the symmetric subspaces of *definitions 1* and *2* by $\mathcal{SYM}_{def1}$ and $\mathcal{SYM}_{def2}$ respectively. We wish to show that these coincide. Note first that $\mathcal{SYM}_{def1}$ is the span of all states of the form $|\psi\rangle\,|\psi\rangle\,|\psi\rangle$, which are clearly symmetrical in the sense of *definition 2*. Hence $\mathcal{SYM}_{def1} \subseteq \mathcal{SYM}_{def2}$. For the reverse inclusion consider a general state in $\mathcal{H}^3$:

$$\begin{aligned} |\alpha\rangle \;=\; & a_0\,|0\rangle\,|0\rangle\,|0\rangle \\ & +a_1\,|1\rangle\,|0\rangle\,|0\rangle + a_2\,|0\rangle\,|1\rangle\,|0\rangle + a_3\,|0\rangle\,|0\rangle\,|1\rangle \\ & +a_4\,|1\rangle\,|1\rangle\,|0\rangle + a_5\,|1\rangle\,|0\rangle\,|1\rangle + a_6\,|0\rangle\,|1\rangle\,|1\rangle \\ & +a_7\,|1\rangle\,|1\rangle\,|1\rangle \end{aligned} \tag{3}$$

Interchange of states for any given pair of positions (in the sense of *definition 2*) preserves the number of $|0\rangle$'s and $|1\rangle$'s in each term so that $|\alpha\rangle$ will be in $\mathcal{SYM}_{def2}$ if and only if $a_1 = a_2 = a_3$ and $a_4 = a_5 = a_6$. Indeed we see that $\mathcal{SYM}_{def2}$ is four dimensional with orthonormal basis states (labelled by the number of $|1\rangle$'s):

$$|e_0\rangle \;=\; |0\rangle\,|0\rangle\,|0\rangle$$

$$
\begin{aligned}
|e_1\rangle &= (|1\rangle|0\rangle|0\rangle + |0\rangle|1\rangle|0\rangle + |0\rangle|0\rangle|1\rangle)/\sqrt{3} \\
|e_2\rangle &= (|1\rangle|1\rangle|0\rangle + |1\rangle|0\rangle|1\rangle + |0\rangle|1\rangle|1\rangle)/\sqrt{3} \\
|e_3\rangle &= |1\rangle|1\rangle|1\rangle
\end{aligned}
\tag{4}
$$

(The four normalising factors 1, $\sqrt{3}$, $\sqrt{3}$ and 1 are square roots of the binomial coefficients ${}^3C_0, {}^3C_1, {}^3C_2, {}^3C_3$.) Now for any $|\psi_1\rangle$ of the form $a|0\rangle + |1\rangle$ we get directly that

$$|\psi_1\rangle|\psi_1\rangle|\psi_1\rangle = a^3|e_0\rangle + a^2\sqrt{3}|e_1\rangle + a\sqrt{3}|e_2\rangle + |e_3\rangle$$

Repeating this for four different values of the parameter a we get:

$$\begin{pmatrix} a^3 & a^2 & a & 1 \\ b^3 & b^2 & b & 1 \\ c^3 & c^2 & c & 1 \\ d^3 & d^2 & d & 1 \end{pmatrix} \begin{pmatrix} |e_0\rangle \\ \sqrt{3}|e_1\rangle \\ \sqrt{3}|e_2\rangle \\ |e_3\rangle \end{pmatrix} = \begin{pmatrix} |\psi_1\rangle|\psi_1\rangle|\psi_1\rangle \\ |\psi_2\rangle|\psi_2\rangle|\psi_2\rangle \\ |\psi_3\rangle|\psi_3\rangle|\psi_3\rangle \\ |\psi_4\rangle|\psi_4\rangle|\psi_4\rangle \end{pmatrix}$$

Choosing a, b, c and d so that the coefficient matrix is invertible, we see that the basis states (4) are all in $\mathcal{SYM}_{def1}$ so that $\mathcal{SYM}_{def2} \subseteq \mathcal{SYM}_{def1}$. Hence these subspaces coincide. □

From the above considerations (c.f. especially (4)) we readily see that the dimension of $\mathcal{SYM}$ for R qubits is $R+1$ so that $\mathcal{SYM}$ is an exponentially small subspace of $\mathcal{H}^R$ (of dimension 2^R). This is also true in the general case. Suppose that $\mathcal{H}$ has dimension d with orthonormal basis $|0\rangle, |1\rangle \dots |d-1\rangle$. Then $\mathcal{SYM}$ has an orthogonal basis labelled by all possible ways of making R choices from the d basis states with repetitions possible and the ordering of choices being irrelevant (c.f. (3) and (4)). The solution of this combinatorial problem gives

$$\text{Dimension of } \mathcal{SYM} = {}^{R+d-1}C_{d-1} = \frac{1}{(d-1)!}R^{d-1} + O(R^{d-2}) \tag{5}$$

which is a polynomial in R (for fixed d). Hence $\mathcal{SYM}$ is again exponentially small inside $\mathcal{H}^R$ of dimension d^R.

§3 Projection into $\mathcal{SYM}$

Our proposed method of stabilisation consists of frequently repeated projection of the joint state of R computers into the symmetric subspace $\mathcal{SYM}$. According to the interpretation of $\mathcal{SYM}$ above, the error free component of any state always lies in $\mathcal{SYM}$ so that upon successful projection this component will be unchanged and part of the error will have been removed. Note however that the projected state is generally not error–free since, for example, $\mathcal{SYM}$ contains many states which are not of the simple product form $|\psi\rangle\,|\psi\rangle\ldots|\psi\rangle$. Nevertheless the error probability will be suppressed by a factor of $1/R$ as discussed in subsequent sections. Thus the method is not one of error correction but rather of stabilisation. By choosing R sufficiently large and the rate of symmetric projection sufficiently high, the residual error at the end of a computation can, in principle, be controlled to lie within any desired small tolerance.

The operation of projection into $\mathcal{SYM}$ is a computation in itself. For our stabilisation method to be efficient it is essential that this operation be executable efficiently *i.e.* in a number of steps that increases at most polynomially with L and R where $L = \log_2 d$ is the number of qubits required to hold the state of each computer entering into the symmetrisation and R is the degree of redundancy. (Note also that clearly R can be at most a polynomial function of L in any efficient scheme.) Only then will a nominally efficient computation remain efficient after stabilisation.

We next describe an algorithm for projecting into $\mathcal{SYM}$ and show that it is efficient in the above sense. Consider first a product state $|\Psi\rangle = |a_1\rangle\,|a_2\rangle\ldots|a_R\rangle \in \mathcal{H}^R$. To project $|\Psi\rangle$ into $\mathcal{SYM}$ we carry out the following steps:

Step 1: Introduce an ancilla in a standard state $|0\rangle$ with a state space $\mathcal{A}$ of at least $R!$ dimensions.

Step 2: Make an equal amplitude superposition of the ancilla

$$\mathcal{U} : |0\rangle \rightarrow \frac{1}{\sqrt{R!}} \sum_{i=0}^{R!-1} |i\rangle$$

Step 3: Carry out the following computation: if the ancilla state is $|i\rangle$ then perform the i^{th} permutation σ_i of the component states of $|a_1\rangle\,|a_2\rangle\ldots|a_R\rangle$

$$|a_1\rangle\,|a_2\rangle\ldots|a_R\rangle\,|i\rangle \rightarrow \left|a_{\sigma_i(1)}\right\rangle \left|a_{\sigma_i(2)}\right\rangle \ldots \left|a_{\sigma_i(R)}\right\rangle |i\rangle$$

This results in the entangled state

$$\sum_i \left|a_{\sigma_i(1)}\right\rangle \left|a_{\sigma_i(2)}\right\rangle \ldots \left|a_{\sigma_i(R)}\right\rangle |i\rangle \in \mathcal{H}^R \otimes \mathcal{A}$$

Step 4: Apply the reverse computation $\mathcal{U}^{-1}$ of step 2 to the ancilla. The resulting state may be written

$$|\Upsilon\rangle = \sum_i |\xi_i\rangle\,|i\rangle \in \mathcal{H}^R \otimes \mathcal{A}$$

Since $\mathcal{U}$ transforms $|0\rangle$ to each $|i\rangle$ with equal amplitude it follows that $\mathcal{U}^{-1}$ transforms each $|i\rangle$ back to $|0\rangle$ with equal amplitude. Hence the coefficient of ancilla state $|0\rangle$ in $|\Upsilon\rangle$ is the required symmetrised state *i.e.* an equal amplitude superposition of all permutations of the R factor states of $|a_1\rangle\,|a_2\rangle\ldots|a_R\rangle$.

Step 5: Measure the ancilla in its natural basis. If the outcome is "0" then $|\Psi\rangle$ has been successfully projected into $\mathcal{SYM}$. If the outcome is not "0" then the symmetrisation has failed. (The issue of the probability of successful symmetrisation is discussed in a later section.)

Finally note that by linearity of the process, it will symmetrise a general state in $\mathcal{H}^R$ (not just the product states considered above). If the input state is already symmetric then we get it back unchanged with certainty at the end.

Consider now the computational effort involved in the above steps. Let $d = \dim \mathcal{H}$ and write $L = \log_2 d$. Step 1 requires no computational effort. The ancilla requires $\log_2(R!) = O(R \log R)$ qubits. Step 2 may be achieved by applying the discrete Fourier transform [7] [8] to the ancilla. This requires $O((R \log R)^2)$ steps. For step 3 we note that a general permutation can be effected with $O(R \log R)$ swaps. Swapping states of L qubits requires $O(L)$ operations so overall step 3 requires $O(LR \log R)$ steps. Restoring the ancilla in step 4 requires the same number of operations as step 2. In step 5 we examine separately each of the $O(R \log R)$ qubits occupied by the ancilla, requiring $O(R \log R)$ steps. Overall we require $O(LR \log R + (R \log R)^2)$ steps which is less than $O(LR^2 + R^4)$. Hence the process is efficient.

§4 A Quantum Network for $\mathcal{SYM}$ Projection

We now describe how the operation of $\mathcal{SYM}$ projection can be implemented by a network of simple quantum gates. Consider first the following inductive definition of the general permutation of $k+1$ elements $a_1, \ldots, a_k, a_{k+1}$[23]. Starting from the general permutation $a_{\sigma(1)}, \ldots, a_{\sigma(k)}$ of the k elements $a_1, \ldots, a_k$ we adjoin a_{k+1} giving $a_{\sigma(1)}, \ldots, a_{\sigma(k)}, a_{k+1}$ and then perform separately the $k+1$ operations: identity, swap $a_{\sigma(1)}$ with a_{k+1}, swap $a_{\sigma(2)}$ with a_{k+1}, ... swap $a_{\sigma(k)}$ with a_{k+1}. This generates all possible permutations of $k+1$ elements. In terms of state symmetrisation, if we have already symmetrised $|\psi_1\rangle \otimes \ldots \otimes |\psi_k\rangle$ (*i.e.* we have an equal superposition of all permutations of the states) then we can symmetrise $k+1$ states $|\psi_1\rangle \otimes \ldots \otimes |\psi_k\rangle \otimes |\psi_{k+1}\rangle$ by applying only the operation of state swapping (in suitable superposition). Thus the operation of symmetrisation of R states can be built up from $|\psi_1\rangle$ by first symmetrising $|\psi_1\rangle$ and $|\psi_2\rangle$, then successively including $|\psi_3\rangle$ up to $|\psi_R\rangle$ always using only state swappings in suitably controlled superpositions.

The basic ingredient in this process is the "controlled swap gate" or Fredkin gate, acting on three input qubits. If the first ("control") qubit is $|0\rangle$ (respectively $|1\rangle$) then the other two ("target") qubits are unaffected (respectively swapped). We describe this diagramatically in Fig. 1. The operation of state swapping itself (*i.e.* $|\psi_1\rangle \otimes |\psi_2\rangle \mapsto |\psi_2\rangle \otimes |\psi_1\rangle$) can be implemented using three controlled–NOT gates as described in [15].

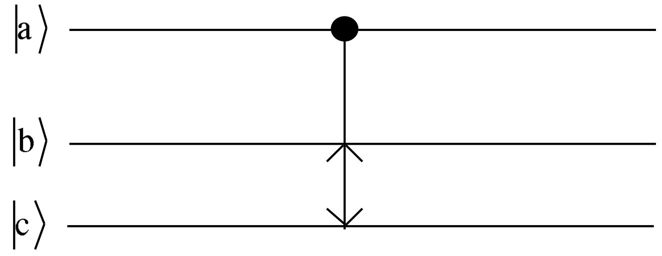

Figure 1: Schematic representation of a Fredkin gate. A Fredkin gate exchanges the state of the second and third qubit if and only if the first qubit is in state $|a\rangle = |1\rangle$.

To symmetrise $k+1$ qubits given that the first k are already symmetrised we introduce k control qubits initially in state $|0\rangle|0\rangle\ldots|0\rangle$ and apply a suitable unitary transformation, denoted U_k to generate the superposition

$$\frac{1}{\sqrt{k+1}}\left(|00\ldots0\rangle + |10\ldots0\rangle + |01\ldots0\rangle + \cdots + |00\ldots1\rangle\right). \tag{6}$$

The unitary transformation U_k can be readily obtained by a quantum network consisting of a one bit gate performing the transformation

$$\frac{1}{\sqrt{k+1}}\begin{pmatrix} 1 & -\sqrt{k} \\ \sqrt{k} & 1 \end{pmatrix} \tag{7}$$

on the first qubit and a sequence of $k-1$ two bit gates $T_{j,j+1}$ for $j=1,\ldots k-1$ acting on the j^{th} and $j+1^{th}$ qubits. In the basis $\{|0\rangle, |1\rangle\}$, $T_{j,j+1}$ is given by:

$$T_{j,j+1} = \frac{1}{\sqrt{k-j+1}}\begin{pmatrix} \sqrt{k-j+1} & 0 & 0 & 0 \\ 0 & 1 & \sqrt{k-j} & 0 \\ 0 & -\sqrt{k-j} & 1 & 0 \\ 0 & 0 & 0 & \sqrt{k-j+1} \end{pmatrix} \tag{8}$$

Having thus initialised the k control qubits, we then apply k Fredkin gates: the j^{th} Fredkin gate (for $j=1,\ldots k$) uses the j^{th} control qubit to control the swapping of the j^{th} and $k+1^{th}$ target qubits. This leads to an entangled state of the k control qubits and the $k+1$ target qubits but after applying U_k^{-1} to the control qubits, the coefficient of $|0\rangle|0\rangle\ldots|0\rangle$ will be the required symmetrisation of the $k+1$ qubits (c.f. step 4 of §3). Finally a measurement of the control qubits will effect the projection into $\mathcal{SYM}$ (c.f. step 5 of §3).

Thus to symmetrise R qubits we cascade the above construction with $k = 1, 2\ldots$ up to $k = R - 1$ requiring a total number $1 + 2 + \cdots + (R - 1) = R(R - 1)/2$ of control qubits. The size of the overall network is clearly quadratic in R. For example, for the symmetrisation of $R = 4$ qubits we obtain the network shown in Fig. 2.

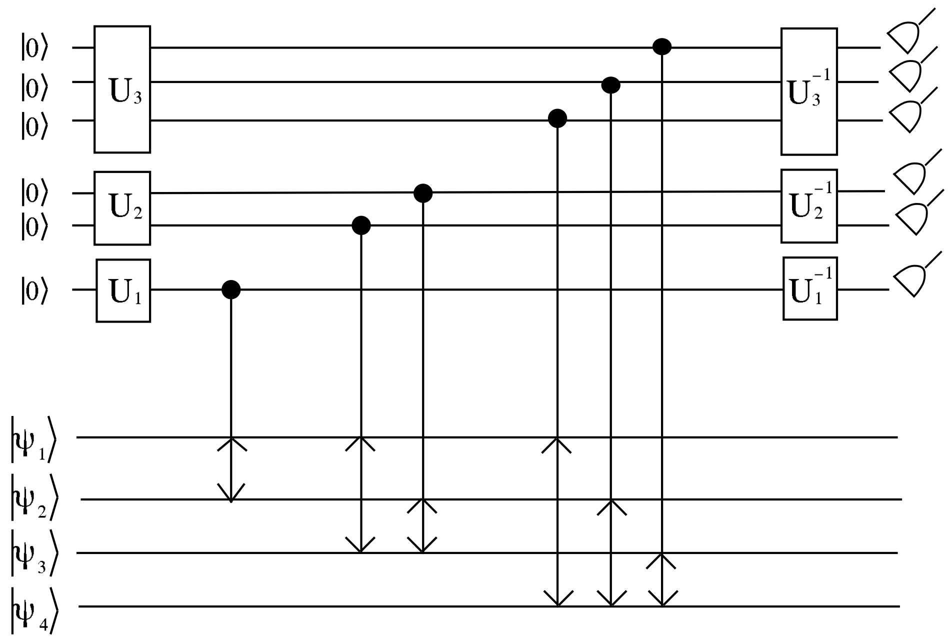

Figure 2: Quantum network for symmetrising $R = 4$ qubits. Six auxiliary qubits initially in state $|0\rangle$ are needed. The auxiliary qubits are put into an entangled state and used to control the state swapping of the four computer qubits. The operations are then undone and the auxiliary qubits measured. If every auxiliary qubit is found in state $|0\rangle$ the symmetrisation has been successful.

§5 Stabilisation Against Unitary Errors

So far we have given an efficient algorithm for projection into the symmetric subspace and provided an intuitive reason why it would be expected to reduce the error while preserving the correct computation. We now turn to a quantitative study of the effect of $\mathcal{SYM}$–projection as a basis for stabilisation in the presence of various modes of error production. It is convenient to separate the discussion into two parts considering the case where the joint state of the computers remains in a pure state, in this section, and the case of decoherence due to external environmental interaction in the following section.

Consider the simple model of R qubits initially in state (the "correct" state) $|0\rangle |0\rangle \dots |0\rangle$ with computation being the identity *i.e.* we are considering the state storage problem with R–fold redundancy. Suppose that the R storage devices are subject to independent hardware errors which cause the j^{th} state to drift as $e^{iH_j t} |0\rangle$. Here the Hamiltonians H_j are random and independent. Since the devices were intended for state storage we assume that the rate of drift is suitably bounded. This is expressed by requiring that all eigenvalues of the H_j's are suitably small:

$$|\text{eigenvalues of } H_j| \leq \epsilon \qquad j = 1, \dots, R \tag{9}$$

for some (small) constant ϵ. The stabilisation process consists of projecting the joint state of the R copies into the symmetric subspace at short time intervals δt. For simplicity we will assume that the projection can be performed essentially instantaneously. Furthermore we assume that (unlike the computation being stabilised) the projection process itself is error free. These and other assumptions will be discussed in §7 below. Under these assumptions we can readily compare the growth of errors with and without the stabilisation process.

In the basis $\{|0\rangle, |1\rangle\}$ write

$$H_j = \begin{pmatrix} a_j & c_j^* \\ c_j & b_j \end{pmatrix} \tag{10}$$

so that

$$|c_j| \leq |\lambda_1| + |\lambda_2| \leq 2\epsilon \tag{11}$$

(where λ_1, λ_2 are the eigenvalues of H_j.) We will assume that δt is small and retain only the lowest order terms in δt. After time δt the state will be

$$|\Psi(\delta t)\rangle = \bigotimes_{k=1}^{R} \{(1 + i a_k \delta t) |0\rangle + i c_k \delta t |1\rangle\} \tag{12}$$

Thus without symmetrisation the probability that the i^{th} qubit shows an error is

$$|c_i|^2 \delta t^2 \approx 4\epsilon^2 \delta t^2 \tag{13}$$

If we expand out the product in (12) we obtain 2^R terms corresponding to the exponentially large dimension of the full space of R qubits. However the amplitudes of terms involving k errors (*i.e.* products of $|0\rangle$'s and $|1\rangle$'s involving k $|1\rangle$'s) will have size $O(\delta t^k)$ and only the R terms involving one error will have size $O(\delta t)$. Thus the erroneous state (12) does not occupy these exponentially many dimensions of $\mathcal{H}^R$ equally. We noted previously that $\mathcal{SYM}$ is exponentially small inside $\mathcal{H}^R$ but the preceeding observation indicates that $\mathcal{SYM}$–projection will not generally remove *exponentially* much of the error since only R of these exponentially many dimensions are entered (to lowest order) by the erroneous evolution. We now calculate the stabilising effect of the $\mathcal{SYM}$–projection.

Consider the basis of $\mathcal{SYM}$ given by the $R+1$ orthonormal states (c.f. (4) for the case $R=3$):

$$|e_k\rangle = \frac{1}{\sqrt{^R C_k}} \sum_{\text{all "}k\text{–error" } \sigma\text{'s}} |\sigma\rangle \qquad k = 0, \ldots, R \tag{14}$$

Here the sum is over all $^R C_k$ possible strings of 0's and 1's of length R containing exactly k 1's and $R-k$ 0's. Under $\mathcal{SYM}$–projection the lowest order error terms (single–error terms) of (12) will project only onto $|e_1\rangle$. For the term with an error in the k^{th} place we get

$$ic_k\delta t\,|0\rangle \ldots |0\rangle\,|1\rangle\,|0\rangle \ldots |0\rangle \qquad \xrightarrow{\mathcal{SYM}\text{ proj}} \qquad \frac{ic_k\delta t}{\sqrt{^R C_1}}|e_1\rangle$$

Thus the normalised projected state has the form

$$\left\{1 + i\delta t \sum_{j=1}^{R} a_j\right\} |0\rangle \ldots |0\rangle + \alpha_1 |e_1\rangle + O(\delta t^2)$$

where

$$\alpha_1 = \frac{i\delta t}{\sqrt{R}} \sum_{j=1}^{R} c_j \tag{15}$$

To estimate the size of α_1, using (11) we write $c_j \approx 2\epsilon e^{i\theta_j}$ where θ_j are random phases. The expectation value of α_1 is then clearly zero but from

$$\text{Expectation value of } |\sum_{j=1}^{R} e^{i\theta_j}| = \sqrt{R} \tag{16}$$

we get [1]

$$\text{Expectation value of } |\alpha_1|^2 = 4\epsilon^2\delta t^2 \tag{17}$$

Thus, somewhat surprisingly, this probability of a single symmetrised error does not decrease with R. However it is associated with R copies and to see its residual effect on any *one* copy we use the following fact:

[1] Remark [21] : If instead of qubit systems we consider computers with dimensions large compared to the degree of redundancy R, then we would expect the individual random errors to be mutually orthogonal, so when the state is symmetrised their sum does not exhibit the cancelling effects which are present for qubits and lead to (16) and (17). However in that case, (17) and subsequent conclusions still hold because (16) may be replaced by Pythagoras' theorem *i.e.* that the sum of R orthonormal vectors has length $\sqrt{R}$.

Proposition 1 Consider the state

$$|\Xi\rangle = \sum_{k=0}^{R} \alpha_k |e_k\rangle \in \mathcal{SYM} \subseteq \mathcal{H}^R$$

where $|e_k\rangle$ are as given in (14). If one qubit is measured in the basis $\{|0\rangle, |1\rangle\}$ the probability of obtaining $|1\rangle$ is $\frac{1}{R}\sum_{k=0}^{R} |\alpha_k|^2 k$. $\Box$

Proof: Since the state is symmetric the probability of obtaining the result $|1\rangle$ for the i^{th} qubit is the same as this probability for the first qubit. Now $|e_k\rangle$ in (14) consists of ${}^RC_k$ orthogonal terms of which ${}^{R-1}C_{k-1}$ have $|1\rangle$ in the first place. Hence the term $\alpha_k |e_k\rangle$ in $|\Xi\rangle$ contributes probability

$${}^{R-1}C_{k-1} \left| \frac{\alpha_k}{\sqrt{{}^RC_k}} \right|^2 = |\alpha_k|^2 k/R$$

of obtaining outcome $|1\rangle$.$\Box$

Applying this result to (15) and using (17) we see that after successful symmetrisation the probability of error (to lowest order in δt) is $4\epsilon^2\delta t^2/R$ *i.e.* the error is suppressed by a factor of $1/R$ compared to the case (13) of no symmetrisation and in each step the amplitude of correct computation is correspondingly enhanced.

The above result is conditional on the success of the symmetrisation *i.e.* that the state projects to $\mathcal{SYM}$ rather than $\mathcal{SYM}^\perp$. If the projections are done frequently enough then the cumulative probability that they *all* succeed can be made as close as desired to unity. This is a consequence of the so–called "quantum watch-dog effect". Consider a normalised joint state $|\Xi\rangle$ of R copies initially in $\mathcal{SYM}$. Its initial probability of successful projection is 1 which is a *maximum*. Thus as the state evolves by some unitary transformation into the ambient space $\mathcal{H}^R$ the probability of successful projection will begin to change only to second order in time. If we project n times per unit time interval *i.e.* $\delta t = 1/n$ then the cumulative probability that all projections in one unit time interval succeed, is

$$(1 - k\delta t^2)^n = (1 - \frac{k}{n^2})^n \to 1 \text{ as } n \to \infty$$

Here k is a constant depending on the rate of rotation of the state out of $\mathcal{SYM}$. For redundancy degree R and the model of random unitary errors considered above we find that k grows linearly with R (as can be seen by directly calculating the length of the $\mathcal{SYM}$-projection of (12) to $O(\delta t^2)$ terms). Thus to achieve a cumulative probability of successful projection of $1 - \zeta$ in a unit time interval we would require a rate of symmetric projection which increases linearly with $-R/\log(1-\zeta)$.

The above conclusions — for a model of random independent unitary errors — will also apply to computations which are not the identity. Formally we may view the computation in a moving basis relative to which the correct computation is the identity and the previous arguments are unchanged i.e. none of the arguments depend on the actual identity of the computational basis states.

§6 Stabilisation Against Environmental Interaction

We now consider the problem of state storage with R–fold redundancy, in the presence of decoherence *i.e.* interaction with an external environment. In general each qubit will become entangled with an environment and the state of the qubit alone will no longer be describable by a pure state. It will be represented by a density matrix [16] resulting from forming a partial trace over the environment, of the joint (pure) state of the total qubit–environment system.

Consider R copies of the qubit initially all prepared in pure state $\rho_0 = |0\rangle\langle 0|$. We will assume that they interact with independent environments (this assumption is valid if the coherence length of the reservoir is less than the spatial separation between the copies [6]) so that after some short period of time δt the state of the R copies will have undergone an evolution

$$\rho^{(R)}(0) = \rho_0 \otimes \ldots \otimes \rho_0 \quad \longrightarrow \quad \rho^{(R)}(\delta t) = \rho_1 \otimes \ldots \otimes \rho_R \tag{18}$$

where $\rho_i = \rho_0 + \sigma_i$ for some Hermitian traceless σ_i and the superscript R denotes the number of states involved. We will retain only terms of first order in the perturbations σ_i so that the overall state at time δt is

$$\begin{aligned} \rho^{(R)} \; = \rho_0 \otimes \ldots \otimes \rho_0 \; & +\sigma_1 \otimes \rho_0 \otimes \ldots \otimes \rho_0 \\ & +\rho_0 \otimes \sigma_2 \otimes \ldots \otimes \rho_0 \\ & +\rho_0 \otimes \rho_0 \otimes \ldots \otimes \sigma_R \\ & +O(\sigma_i \sigma_j) \end{aligned} \tag{19}$$

and we wish to compute the projection of the state (19) into the symmetric subspace $\mathcal{SYM}$. Then we construct the state of the i^{th} qubit by partial trace over all qubits except the i^{th} and finally compare the resulting state with $\rho_0 + \sigma_i$ and see that its purity has been suitably enhanced, bringing it closer to ρ_0.

The mathematical formalism for symmetrisation of mixed states has some curious features which we digress to clarify before treating (19) itself. Consider a state $\rho_1 \otimes \rho_2$ of two qubits. The state $\frac{1}{2}(\rho_1 \otimes \rho_2 + \rho_2 \otimes \rho_1)$ is *not* a symmetric state and in fact $\rho \otimes \rho$ is not symmetric (*i.e.* it is not a density matrix supported on the subspace $\mathcal{SYM}$) unless ρ is pure! To see this consider ρ written in its diagonalising basis of orthonormal eigenstates:

$$\rho = \lambda_1 |\lambda_1\rangle\langle\lambda_1| + \lambda_2 |\lambda_2\rangle\langle\lambda_2| \tag{20}$$

Thus we can represent ρ as a mixture of its two eigenstates, and $\rho \otimes \rho$ as a mixture of the four orthonormal states $|\lambda_i\rangle \otimes |\lambda_j\rangle$ with a priori probabilities $p_{ij} = \lambda_i \lambda_j$. This latter mixture involves *nonsymmetric* states (like $|\lambda_1\rangle \otimes |\lambda_2\rangle$) so $\rho \otimes \rho$ is not symmetric.

One way of constructing the projection of $\rho \otimes \rho$ into $\mathcal{SYM}$ is to project each state of the above mixture into $\mathcal{SYM}$. Let $|\mu_{ij}\rangle$ denote the $\mathcal{SYM}$–projection of $|\lambda_i\rangle \otimes |\lambda_j\rangle$ and $|\hat{\mu_{ij}}\rangle$ denote the corresponding normalised state. The probability of successful projection is $q_{ij} = \langle \mu_{ij} | \mu_{ij} \rangle$. Then the $\mathcal{SYM}$–projection of $\rho \otimes \rho$ is the state corresponding to the mixture $|\hat{\mu_{ij}}\rangle$ with a priori probabilities $p_{ij} q_{ij} / (\sum p_{ij} q_{ij})$ which are the conditional probabilities of occurrence of states $|\hat{\mu_{ij}}\rangle$ given that the $\mathcal{SYM}$–projection was succcessful.

More formally we may introduce the (Hermitian) permutation operators $P_{12} =$ "identity" and $P_{21} =$ "swap" acting on pure states of two qubits and define the symmetrisation operator:

$$S = \frac{1}{2}(P_{12} + P_{21}) \tag{21}$$

The $\mathcal{SYM}$–projection of a pure state $|\Psi_{12}\rangle$ of two qubits is just $S\,|\Psi_{12}\rangle$, which is then renormalised to unity. It follows that the induced map on *mixed* states of two qubits (including renormalisation) is:

$$\rho_1 \otimes \rho_2 \longrightarrow \frac{S(\rho_1 \otimes \rho_2) S^\dagger}{\mathrm{Tr}\, S(\rho_1 \otimes \rho_2) S^\dagger} \tag{22}$$

The state of either qubit separately is obtained by partial trace over the other qubit.

As an example consider the symmetric projection of $\rho \otimes \rho$ followed by renormalisation and partial trace (over either qubit) to obtain the final state $\tilde{\rho}$ of one qubit, given that the $\mathcal{SYM}$–projection was successful. A direct calculation based on (22) yields:

$$\rho \mapsto \tilde{\rho} = \frac{\rho + \rho^2}{\mathrm{Tr}\,(\rho + \rho^2)} \tag{23}$$

For any mixed state ξ of a qubit the expression $\mathrm{Tr}\,\xi^2$ provides a measure of the purity of the state, ranging from $\frac{1}{4}$ for the completely mixed state $I/2$ (where I is the unit operator) to 1 for any pure state. From (23) we get

$$\mathrm{Tr}\,\tilde{\rho}^2 > \mathrm{Tr}\,\rho^2$$

so that $\tilde{\rho}$ is *purer* than ρ. This example illustrates a generic fact (c.f. below), that successful projection of a mixed state into $\mathcal{SYM}$ tends to enhance the purity of the individual systems. Indeed, consider further the state $\otimes^R \rho$ consisting of R independent copies of ρ. The symmetrisation operator is

$$S = \frac{1}{R!} \sum_{\alpha=1}^{R!} P_\alpha \tag{24}$$

where the sum ranges over all $R!$ permutations of the R indices. If we project $\otimes^R \rho$ into $\mathcal{SYM}$ and renormalise (as in (22)) and calculate the partial trace over all but one of the qubits, we obtain a reduced state $\tilde{\rho_R}$ which asymptotically tends to a *pure* state as $R \to \infty$. This limiting pure state is the eigenstate of ρ belonging to its largest eigenvalue.

Let us now return to the consideration of (19) and its $\mathcal{SYM}$–projection. The application of the symmetrisation operator (24) to each of the R terms of $\rho_0 \otimes \ldots \otimes \sigma_i \otimes \ldots \otimes \rho_0$ of Eq. (19) generates $R!^2$ terms of the form

$$\frac{1}{R!^2} P_\alpha \rho_0 \otimes \ldots \otimes \sigma_i \otimes \ldots \otimes \rho_0 P_\beta, \tag{25}$$

where P_α and P_β are permutations operators on the state space $\mathcal{H}^R$ of R qubits as above. To calculate the reduced density operator of the first qubit we take the partial trace over the $R-1$ remaining qubits. Note that the reduced states of all qubits individually are equal since the total overall state is symmetric. (To systematise the calculation of the partial traces we have found it very convenient to use the diagrammatic notation for tensor operations introduced by Penrose in [22].) For each σ_i we find that the $R!^2$ terms in (25) then reduce to the following cases:

(i) $(R-1)!^2$ terms each equal to $\sigma_i/R!^2$ corresponding to all permutations P_α and P_β which place σ_i in the first position in (25). In this case the partial trace contracts out all the ρ_0 terms leaving a coefficient of $1/R!^2$ (as the trace of any power of ρ_0 is 1).

(ii) $(R-1)!^2(R-1)R$ terms of the forms $\rho_0\sigma_i\rho_0$, $\rho_0\sigma_i$, $\sigma_i\rho_0$, or $\rho_0 \mathrm{Tr}(\sigma_i\rho_0)$, each one divided by $R!^2$. These correspond to all pairs of permutations which result in σ_i contracted onto ρ_0 in all possible ways in the partial traces.

(iii) $(R-1)!^2(R-1)$ terms which result in σ_i being contracted onto itself in the partial traces. These terms are all zero since $\mathrm{Tr}\,\sigma_i = 0$.

Note that each term in (ii) has trace given by $\mathrm{Tr}\,\sigma_i\rho/R!^2$ and each term in (i) has zero trace. Thus the resulting density operator, before normalisation, has a trace given by

$$1 + (R-1)\mathrm{Tr}(\rho_0\tilde{\sigma}), \tag{26}$$

where we have introduced $\tilde{\sigma} = \frac{1}{R}\sum_{i=1}^{R} \sigma_i$. After normalisation – which is calculated by dividing the sum of all terms in (i) and (ii) for all $i = 1, \ldots R$ by (26) – the resulting symmetrised density operator $\tilde{\rho}$ can be written

$$\begin{aligned} \tilde{\rho} &= [1-(R-1)\mathrm{Tr}(\rho_0\tilde{\sigma})]\rho_0 + \frac{1}{R}\tilde{\sigma} \\ &+ (R-1)[A\rho_0\tilde{\sigma}\rho_0 + B(\rho_0\tilde{\sigma}+\tilde{\sigma}\rho_0) + C\rho_0\mathrm{Tr}(\tilde{\sigma}\rho_0)] + O(\sigma_i\sigma_j) \end{aligned} \tag{27}$$

where A, B and C depend on R and $A + 2B + C = 1$.

If a general mixed state ξ of a qubit is measured in the basis $\{|0\rangle, |1\rangle\}$ then the probability that the outcome is 0 is given by $\langle 0|\,\xi\,|0\rangle = \mathrm{Tr}\,\rho_0\xi$. This provides the success probability in our present model. Thus the average success probability *before* symmetrisation of the perturbed qubits is:

$$\frac{1}{R}\sum_i \mathrm{Tr}\,\rho_0(\rho_0+\sigma_i) = 1 + \mathrm{Tr}\,\rho_0\tilde{\sigma} \tag{28}$$

(Note that consequently $\mathrm{Tr}\,\rho_0\tilde{\sigma}$ is necessarily negative). *After* symmetrisation, using (27) we see that

$$\mathrm{Tr}\,\rho_0\tilde{\rho} = 1 + \frac{1}{R}\mathrm{Tr}\,\rho_0\tilde{\sigma} \tag{29}$$

Hence the probability of error has again been reduced by a factor of R – exactly as found in the previous section.

We can calculate the average purity of the R copies before symmetrisation by calculating the average trace of the squared states:

$$\frac{1}{R}\sum_{i=1}^{R}\mathrm{Tr}((\rho_0+\sigma_i)^2) = 1 + 2\mathrm{Tr}(\rho_0\tilde{\sigma}). \tag{30}$$

After symmetrisation each qubit has purity

$$\mathrm{Tr}(\tilde{\rho}^2) = 1 + 2\frac{1}{R}\mathrm{Tr}(\rho_0\tilde{\sigma}). \tag{31}$$

Since $\mathrm{Tr}\,\tilde{\rho}^2$ is closer to 1 than (30), the resulting symmetrised system $\tilde{\rho}$ is left in a purer state. Indeed it follows from (29) that $\tilde{\rho}$ approaches the unperturbed state ρ_0 as R tends to infinity.

§7 Limitations

Error correction is itself a quantum computation. The above analysis has ignored the inevitable build up of errors in the computer performing that computation. Indeed for the symmetrisation of R qubits the projection algorithm requires an ancilla of at least $R!$ dimensions *i.e.* $O(R \log R)$ qubits (in fact $O(R^2)$ in our explicit network). Thus the correcting apparatus is slightly larger than the total system being corrected so the error correction ought to be subject to a similar level of error as is present in the original system. In a situation where the redundancy degree R is small compared to the number L of qubits per computer, the correcting apparatus (still of $O(R^2)$ qubits) will be small compared to the size RL of the system being corrected. However as seen in §5, the stabilisation of a linear computation on input size L requires redundancy degree $R \sim L$ so that the correcting apparatus and the computer are again of comparable size. This means that each error correcting step introduces errors of a similar, or even greater, probability than those that it is correcting. This does not however render it ineffective. Consider the following illustrative example. A certain clock is accurate to one second per day. Each day at noon it is reset using a standard time signal, the resetting operation being accurate only to one minute i.e. sixty times worse than the error being corrected. Nevertheless after ten years the corrected clock will still be in error by at most one minute. If left uncorrected the error could be almost an hour.

The main factor limiting the efficiency of our proposed method will be the frequency with which the error correcting operations can be physically performed. As noted at the end of §5, to achieve a cumulative probability $1 - \delta$ of repeated successful projection in a unit time interval, the rate of symmetric projection must increase linearly with the degree of redundancy R. Also as noted in §1, the stabilisation of a computer with input size L, running for L steps, requires R to increase linearly with L. Hence we need the overall rate of symmetric projection to increase linearly with L even for a linear time computation. Thus, beyond a certain input size, each symmetrisation will have to be performed in a a time shorter than that needed to perform the elementary quantum gate operations. Since increasing the rate of computation by a factor k presumably requires resources exponential in k, our method would necessarily require exponential resources for sufficiently large L. This property is shared by all quantum error correction schemes that have been proposed to date. Hence quantum algorithms (such as Shor's factoring algorithm) which are polynomially efficient in the absence of errors, would not be efficient if physically implemented. We wish to stress that the traditional notion of efficiency (based on the distinction between polynomial and exponential growth) is an asymptotic notion referring to computations on unboundedly large inputs. This may not be appropriate in assessing the feasibility of particular computations in practice. For example, if a quantum computer could factorise a 1000 digit integer in a reasonable time it may still exceed the abilities of any classical computer for the foreseeable future albeit that the factorisation of 2000 digit integers might be infeasible on *any* computer.

Conclusion

If the technology to implement the scheme we have described were available, it would provide a method of stabilising general coherent computations though not (exponentially) efficiently. This is because although only polynomially many steps are required in the stabilisation computation, these need to be performed in a fixed time, a characteristic time of error growth per bit.

Acknowledgements

We wish to thank Dorit Aharonov, Ethan Bernstein, Asher Peres and Umesh Vazirani for developmental discussions, and Rolf Landauer for critical appraisal.
A.E. and R.J. are supported by the Royal Society, London. C.M. is supported by the European HCM Programme. Adr.B. acknowledges support of the Berrow's Fund at Lincoln College, Oxford.
Part of this work was carried out during the Quantum Computation Workshops, conducted with the support of ELSAG-Bailey, Genova and the Institute for Scientific Interchange, Torino. We are grateful for the opportunity of collaboration provided.

References

[1] Papadimitriou, C. (1994) *Computational Complexity*, Addison–Wesley Publishing Company.

[2] Deutsch, D. (1993) talk presented at the Rank Prize Funds Mini–Symposium on Quantum Communication and Cryptography, Broadway, England.

[3] Berthiaume, A., Deutsch, D. and Jozsa, R. (1994) in *Proceedings of Workshop on Physics and Computation — PhysComp94* (IEEE Computer Society Press, Dallas, Texas).

[4] Zurek, W.H. (1991) *Physics Today* **44** (10), 36.

[5] Landauer, R. (1995) in *Proceedings of the Drexel–4 Symposium on Quantum Nonintegrability — Quantum–Classical Correspondence*, ed. D.H. Feng and B–L. Hu (International Press)

[6] Palma, G. M., Suominen, K.–A. and Ekert, A. (1996) *Proc. Roy. Soc. Lond.* **A 452**, 567.

[7] Shor, P. (1994) in *Proceedings of 35th Annual Symposium on the Foundations of Computer Science*, ed. S. Goldwasser (IEEE Computer Society Press, Los Alamitos, CA) p 116.

[8] Ekert, A. and Jozsa, R. (1996) *Rev. Mod. Phys.* July 1996.

[9] Shor, P. (1995) Phys. Rev. A **52**, R2493.

[10] Steane, A. *Multiple particle interference and quantum error correction*, preprint quant-ph/9601029, to be published in *Proc. Roy. Soc. (Lond).*

[11] Laflamme, R., Miquel, C., Paz, J.P. and Zurek, W.H. *Perfect Quantum Error Correction Codes*, preprint quant–ph/9602019.

[12] Ekert, A. and Macchiavello, C. *Quantum Error Correction for Communication*, preprint quant–ph/9602022.

[13] Calderbank, A. and Shor, P. (1996) *Good Quantum Error Correcting Codes Exist*, preprint quant-ph/9512032

[14] Lang, S. (1993) *Algebra*, 3rd edition, Addison Wesley.

[15] Barenco, A., Deutsch, D., Ekert, A. and Jozsa, R. (1995) *Phys. Rev. Lett.* **74**, 4083.

[16] Peres, A (1993) *Quantum theory: concepts and methods*, Kluwer Academic Publishers.

[17] Peres, A. (1988) *Phys. Lett. A*, **128**, 19.

[18] Fuchs, C. A. and Peres, A. (1996) *Quantum State Disturbance versus Information Gain: Uncertainty Relations for Quantum Information*, *Phys. Rev. A*, to appear.

[19] Wootters, W. K. and Zurek, W. (1982) *Nature* **299**, 802-803.

[20] MacWilliams, F. J. and Sloane, N. J. (1977) *The Theory of Error Correcting Codes*, Amsterdam:North Holland.

[21] Dorit Aharonov, private communication.

[22] Penrose, R. and Rindler, W. (1984) *Spinors and Spacetime*, Volume 1, Appendix. Cambridge University Press.

[23] Knuth, A. (1973) *The Art of Computer Programming Volume 2*, Addison Wesley.