Upper bounds on fault tolerance thresholds of noisy Clifford-based quantum computers

M B Plenio[1,3,4] and S Virmani[2,3,5]

[1] Institut für Theoretische Physik, Albert-Einstein-Allee 11, Universität Ulm,
D-89069 Ulm, Germany
[2] Department of Physics SUPA, University of Strathclyde, John Anderson
Building, 107 Rottenrow, Glasgow G4 0NG, UK
[3] Institute for Mathematical Sciences, Imperial College London,
53 Exhibition Road, London SW7 2PG, UK
[4] QOLS, Blackett Laboratory, Imperial College London, Prince Consort Road,
London SW7 2BW, UK
E-mail: shashank.virmani@strath.ac.uk

Abstract. We consider the possibility of adding noise to a quantum circuit to make it efficiently simulatable classically. In previous works, this approach has been used to derive upper bounds to fault tolerance thresholds—usually by identifying a privileged resource, such as an entangling gate or a non-Clifford operation, and then deriving the noise levels required to make it 'unprivileged'. In this work, we consider extensions of this approach where noise is added to Clifford gates too and then 'commuted' around until it concentrates on attacking the non-Clifford resource. While commuting noise around is not always straightforward, we find that easy instances can be identified in popular fault tolerance proposals, thereby enabling sharper upper bounds to be derived in these cases. For instance we find that if we take Knill's (2005 *Nature* **434** 39) fault tolerance proposal together with the ability to prepare any possible state in the XY plane of the Bloch sphere, then not more than 3.69% error-per-gate noise is sufficient to make it classical, and 13.71% of Knill's γ noise model is sufficient. These bounds have been derived without noise being added to the decoding parts of the circuits. Introducing such noise in a toy example suggests that the present approach can be optimized further to yield tighter bounds.

Contents

1. Introduction

An important open problem in the field of quantum computation is to understand the effects of noise on computational power. In particular, for various noise models and various sets of universal quantum resources, we would like to determine the so-called *fault tolerance threshold* [2], the level of noise that we can tolerate in our basic physical components before we lose the power to perform quantum computation. Most previous work on this question has focused on the problem of constructing *lower bounds* to such thresholds (see e.g. [2]). This is usually achieved by constructing explicit fault-tolerant procedures, and then *estimating* the level of noise that these schemes can tolerate. On the other hand the construction of *upper bounds* has received comparatively less attention, but has recently been the subject of an increasing number of investigations.

There are essentially two methods that have been used to derive upper bounds, which we may loosely describe as 'classical' and 'quantum' approaches. In classical approaches one tries to determine the noise levels required such that the noisy quantum computer can be efficiently simulated on a classical computer [3]–[5], [8]–[10]. In quantum approaches one tries to show that above a certain level of noise the output is bounded away from ideal in some way [6, 7], [10]–[12]. A number of investigations along these lines have resulted in upper bound estimates relevant to various architectures and noise models under varying degrees of assumption and rigour. The two approaches could ultimately be of independent interest, as it is possible that for realistic noise models an 'intermediate' form of computation could exist that is not as powerful as a quantum computer, but is more powerful than the classical.

In this work, we follow a path related to the 'magic state' question and generalizations [4, 5, 9], [18]–[21] and consider classical threshold upper bounds for a specific class of fault-tolerant quantum computational schemes—those built around a core set of *Clifford*

resources [13] such as the CNOT, Pauli rotations, Hadamard and preparation/measurement in Pauli operator eigenbases. Any device consisting of such Clifford operations can be efficiently simulated classically, and so to perform quantum computation an additional resource is needed. Common choices include phase gates of the form

$$U(\theta) := |0\rangle\langle 0| + \exp i\theta |1\rangle\langle 1|, \tag{1}$$

or a supply of single-qubit states such as

$$|\theta\rangle := \frac{1}{\sqrt{2}}(|0\rangle + e^{i\theta}|1\rangle) \tag{2}$$

that are not eigenstates of Pauli operators.

It is significant for our present discussion that devices built from Clifford operations alone can be efficiently simulated classically. This is the content of the Gottesman–Knill theorem [13, 14]. The Gottesman–Knill theorem is interesting as in many other respects Clifford operations exhibit a great deal of non-classicality—for instance they can be used to demonstrate the GHZ paradox and may generate the long-range entanglement that is considered to be a precondition for efficient quantum computation [15]. Furthermore, the entanglement that Clifford operations generate is sufficient for quantum computation—to build a full quantum computer you only need to add the ability to perform single-qubit operations to the Clifford set.

The Gottesman—Knill theorem is important for this work as it can be used [4, 5] to construct classical threshold bounds using the following simple approach. Consider a Clifford-based architecture that is made universal by the addition of a non-Clifford resource, which we refer to as R. If one could determine a noise level sufficient to turn R into an operation that lies in the convex hull of Clifford operations, then this noise level provides an *upper* bound to the fault tolerance threshold for those architectures, as such a noisy device can be efficiently simulated classically[6]. Although one would ideally like to derive threshold bounds that are valid for a wider variety of possible architectures than those based upon Clifford operations, focussing on such architectures is nevertheless quite relevant, as most *lower* bounds have been derived for Clifford-based schemes due to their significance for quantum error correction [14].

Although a number of interesting threshold bounds can be derived with this approach, the noise models considered in [4, 5] are in fact relatively weak because the Clifford operations are taken to be entirely noise free. This naturally leads to the question of whether the bounds can be improved by considering the (often more realistic) situation in which the Clifford operations are also subject to noise[7]. In this work we argue that this is indeed the case. By a straightforward modification—allowing the Clifford operations to be noisy, and then 'commuting' the noise onto the non-Clifford operations—we argue that the previous bounds of [4, 5] can be improved much further for some important families of fault-tolerant quantum computation.

'Commuting' noise around a circuit is not always straightforward, as when noise is moved through entangling gates it can lead to unmanageable error correlations between different qubits.

---

[6] In principle, this approach can also be applied to more general architectures where all resources are non-Clifford—one can ask for the noise required to turn all the basic components into probabilistic applications of Clifford operations.

[7] From an experimental point of view, this is usually a natural assumption, as in many implementations of quantum gates there is little reason to suggest that a Clifford gate should suffer far less noise than a non-Clifford gate. However, in some physical situations the Clifford operations can be much more strongly protected, and such situations were a motivation expressed by the authors of [9].

Table 1. A summary of upper bounds presented in this work. 'Injection' refers to teleportation state-injection. 'NC' means non-Clifford.

NC resources	Method	Noise model	Upper bound (%)
Phase gates/states	Any	Independent N_p^Z	7.96
Phase gates/states	Any	EPG	10.41
All gates	Any	Independent depolarizing	26.05
Phase gates	Injection	Knill	9.59
Phase states	Injection	Knill	13.71
Phase gates	Injection	EPG	3.01
Phase states	Injection	EPG	3.69
All gates	Injection	Knill	15.19
All states	Injection	Knill	21.78
All gates	Injection	EPG	5.03
All states	Injection	EPG	6.31

Hence, in this paper, we analyse only comparatively easy instances where these correlations can be made to disappear or may be accounted for easily. Fortunately, however, it turns out that the teleportation circuit is one such 'easy' instance, and so the approach works well for the various schemes (including the high-threshold proposal of Knill [1]), which use teleportation as a primitive for 'injecting' non-Clifford states into the computation to make it universal. In such cases the problem often reduces to simple geometrical considerations of the Bloch sphere. We expect that the results obtained here may be improved by developing more sophisticated techniques to tackle the correlations arising in more general situations.

A summary of the upper bounds obtained in this paper is given in table 1. In some cases they are not too far from conjectured lower bounds for some proposed schemes.

2. Basic definitions and noise models

There are a variety of noise types and models that we will consider in this work. In this section, we define some of these noise models, and establish some notation. We will only consider computational models using qubits. We will also be only considering stochastic noise models, where at each stage the ideal operation fails with some probability—we do not consider 'coherent' errors here.

- *Clifford operations*. Consider the 'Clifford group' of unitaries, defined for any number of qubits, which is generated by CNOTs, Hadamards, and the 'S' gate $S := |0\rangle\langle0| + i|1\rangle\langle1|$. Augment this group with measurements in the Pauli X, Y, Z eigenbases, and the ability to prepare any Pauli eigenstate. Any physical operation that can be generated by probabilistic application of these resources (where the probabilities can be efficiently sampled classically) will be called a *Clifford operation*.

- *Phase gates*. Phase gates, denoted $U(\theta)$, are defined as

$$U(\theta) := |0\rangle\langle0| + \exp(i\theta)|1\rangle\langle1|. \tag{3}$$

We use the same symbol U to refer to the unitary whether we consider it as a transformation on the set of pure state vectors (i.e. as $U|\psi\rangle$), or by conjugation on the set of density

matrices (i.e. as $U\rho U^{\dagger}$). Most of the bounds derived for phases gates are also valid for phase rotations in the X, Y bases too, provided that the error models are changed in the appropriate way.

- *Phase states.* Phase states, denoted $|\theta\rangle$, are defined as

$$|\theta\rangle := \frac{1}{\sqrt{2}}\left(|0\rangle + \exp(i\theta)|1\rangle\right). \tag{4}$$

  They lie in the plane of the Bloch sphere containing the Pauli X, Y eigenstates. Most of the bounds derived for phase states are also valid for pure states in the Y, Z and X, Z planes, provided that the error models are changed in the appropriate way. In fact the states used explicitly in [1, 23] are sometimes from these other planes.

- *Probabilistic application of a transformation Q.* Let Q be a quantum operation acting on a quantum system. Q could be a unitary or a completely positive (CP) map. We define N_t^Q to be the quantum operation 'with probability $(1-t)$ apply the identity operation to the system and with probability t apply the Q operation', i.e.

$$N_t^Q(\rho) = (1-t)\rho + tQ(\rho). \tag{5}$$

  Usually we will consider Q to be a Pauli operation, or some form of depolarizing operation.

- *Opposite noise.* Consider a known, given, qubit state σ. Define $\psi_\sigma^\perp$ to be the *pure* state in the polar opposite direction to σ in the Bloch sphere. We define O_t^σ to be the single-qubit operation 'with probability $(1-t)$ apply the identity operation to the system and with probability t throw it away and replace it with the pure state in the opposite direction to σ in the Bloch sphere', i.e.

$$O_t^\sigma(\rho) = (1-t)\rho + t\psi_\sigma^\perp. \tag{6}$$

  This operation will be used to construct error-per-gate (EPG, see below) noise upper bounds. Note that it is a valid quantum operation as σ is known and fixed, and does not depend upon ρ.

- *Knill's noise model.* The noise model defined in the paper by Knill [1] is parametrized by a noise strength γ in the following way: a preparation of a qubit state ψ prepares the orthogonal pure state with probability $4\gamma/15$, a measurement of the Pauli Z operator is preceded by a Pauli X rotation with probability $4\gamma/15$, a measurement of the Pauli X operator is preceded by a Pauli Z rotation with probability $4\gamma/15$, a single-qubit unitary is followed by one of the Pauli X, Y, Z rotations each with probability $4\gamma/15$, a CNOT is followed by one of the 15 non-Identity Pauli products ($I \otimes X$, $Y \otimes I$, $Z \otimes X$, etc) each with probability $\gamma/15$.

- *Independent depolarizing noise, with strength t.* After every nontrivial unitary operation, and before each nontrivial measurement, apply the noisy operation N_t^D, where D is the totally depolarizing operation:

$$D(\rho) = \frac{1}{4}\left(\rho + X\rho X + Y\rho Y + Z\rho Z\right) \tag{7}$$

- *Simultaneous depolarizing noise, with strength t.* After each nontrivial gate acting on k qubits, each possible Pauli product (including the Identity) is applied to those qubits with probability $t/(4^n)$, with the remaining probability being taken up by the Identity transformation. In particular this means that after a CNOT, for example, both qubits are 'jointly depolarized' rather than independently on each wire.

- *EPG. Probabilistic error-per-gate noise, with strength t.* Each operation (including memory steps required while waiting for measurement outcomes elsewhere) fails with probability t, but the precise manner of failure is not specified other than being probabilistic—so we are free to choose the operation that occurs in the case of failure. In particular the output of a k-qubit gate can undergo some correlated, joint, error. However, correlations in the noise affecting non-interacting qubits are not allowed in EPG, and we do not allow the noise to be conditioned on classical measurement outcomes that have occurred during post-selection elsewhere.

Although the EPG model may seem a little adversarial, it is in fact weaker than many of the noise models considered in more rigorous lower bound estimates, such as [16, 17]. In the local stochastic noise model [16], for example, a form of weak correlation is admitted, and this is where it differs from EPG. However, upper bounds for EPG models will still of course be valid upper bounds for all stronger noise models.

3. Dephasing of general phase resource circuits

To illustrate the basic approach we begin by analysing a particularly simple scenario where the non-Clifford resource and noise model are picked such that the noise can be commuted quite straightforwardly from the Clifford gates to the non-Clifford operations.

Suppose that we apply dephasing noise to Clifford-based architectures that are made universal by the addition of phase gates. In particular let us assume the following noise model: each qubit undergoes a dephasing operation of the form N_p^Z (see equation (5)) after every nontrivial gate, after every qubit state initialization, and before every measurement. Because dephasing commutes with phase gates the analysis turns out to be especially simple. This not only leads to very simple calculations but also allows us to consider the most general Clifford-based architectures with phase resources, thereby leading to the smallest upper bounds that can be proven using this approach for such architectures. We present the argument for phase gates, but with minor modification it can be shown to hold for any combination of phase gates or phase states.

In [4] it was shown that if Clifford operations are augmented with the ability to do single-qubit phase gates and if each non-Clifford gate $U(\theta)$ in the circuit is immediately followed by N_p^Z, then roughly $p \sim 14.6\%$ of noise is sufficient to make the circuit classically tractable (this is regardless of which phase gates are used). However, let us now consider how this result may be applied to more realistic noise models where the Clifford operations are also subject to noise.

On any given wire, the neighbourhood of each non-Clifford gate can always be taken to have the form,

$$CU^r C, \tag{8}$$

where C denotes a Clifford operation (which may involve more than one qubit, and could be a state preparation), and U^r denotes r applications of some non-Clifford phase-gate U (it does not matter if all these r phase gates are the same or not). The real, noisy version of this wire will be (we consider time increasing from left to right)

$$CN_p^Z (UN_p^Z)^r CN_p^Z. \tag{9}$$

Now, as our noise operation N_p^Z commutes with the Us, which also commute with each other, this equation reduces to

$$CU' (N_p^Z)^{r+1} CN_p^Z, \tag{10}$$

where we now rewrite $U' = U^r$, which is also a phase gate. This is a circuit where the leftmost Clifford operation is perfect, but we now have $r + 1$ lots of noise attacking the U' non-Clifford gate. Note that it does not actually matter if all the U gates were all the same or not; we can still commute the noise and multiply the Us to give a U' that is still a phase gate. It also does not affect the argument if the C operations are in fact qubit preparations or measurements.

With this reordering of the gates and noise we hence see that *any* circuit of Clifford operations and phase gates subject to N_p^Z noise may be reduced to an equivalent circuit where every second operation is a (usually ideal) Clifford operation, and the non-Clifford operations are subject to increased noise. If $r = 0$ then the wire is already a Clifford operation, and we do not need to attack it. The worst case from an attacker's perspective is hence when $r + 1 = 2$, i.e. there is only one U between every pair of successive Clifford operations on the same qubit, in which case we may only attack each phase gate with only *two* lots of noise. In fact this is also the most general case—it is not possible to argue that there are more than two lots of noise attacking each non-Clifford operation, as in principle every second operation on every qubit wire could be a non-Clifford operation.

We hence may now ask how high p must be for *two* lots of noise to take an arbitrary phase rotation U' into the set of Clifford operations—this is in contrast to the calculations of [4] that used only *one* lot of noise. Solving this problem is a trivial extension of the methods used in [4]. One must solve

$$(1 - 2p)^2 = (1 - 2q) \tag{11}$$

for p, where q is the minimal noise level required for a single application of the noise. The value for q found in [4] was

$$q = \frac{1}{2} \left(1 - \frac{1}{\sqrt{2}} \right) \sim 14.64\%. \tag{12}$$

To make the paper more self-contained, we explain the origin of this value in figure 1. Solving for p hence gives

$$p = \frac{1}{2} \left(1 - \sqrt{\frac{1}{\sqrt{2}}} \right) = 7.96\%. \tag{13}$$

So this result may be summarized as follows: for an architecture based upon Clifford operations, phase gates, and noise where each qubit undergoes N_p^Z independently after every nontrivial gate acting upon it (or before a measurement), then a noise level of 7.96% is sufficient to make the circuit tractable classically.

Some authors prefer to use a different parametrization of the dephasing noise model, where instead of applying N_p^Z, one applies the totally dephasing operation $\rho \to 1/2(\rho + Z\rho Z)$ with probability $\tilde{p}$. The two cases are trivially related, and so the above bound can be re-expressed for this noise model as

$$\tilde{p} = 1 - \sqrt{\frac{1}{\sqrt{2}}} = 15.91\%. \tag{14}$$

Figure 1. This figure recaps the argument in [4] leading to equation (12). The set of phase gates can be isometrically represented either as the unit circle in the complex plane or as the circumference of the XY plane of the Bloch sphere via the mappings $U(\theta) \leftrightarrow \exp(i\theta) \leftrightarrow |0\rangle + \exp(i\theta)|1\rangle$. Similarly, convex mixtures of phase gates map to the interior of the circle—for instance, the total dephasing operation is represented by the origin of the complex plane or the maximally mixed state. In such representations the Clifford unitaries $I, S := U(\pi/2), Z, S^{-1}$ map to the four points shown in the figure, and the Clifford region is the region inside the square—all operations outside the square can be shown to be non-Clifford because they take the stabilizer qubit state $|0\rangle + |1\rangle$ outside the 'Pauli octahedron' [9] of the Bloch sphere. Under an unwanted Z rotation $U(\pi/4)$ is taken to $U(5\pi/4)$, and $q \sim 14.64\%$ is given by the smallest amount of $U(5\pi/4)$ that must be mixed with $U(\pi/4)$ to take it into the square. One can see intuitively from the figure that under Z noise there is no phase gate more robust than $U(\pi/4)$. Most of the calculations later in the paper follow from similar arguments related to the geometry of the Bloch sphere.

These improvements on the bounds of [4] have been possible because noise from a neighbouring Clifford operation can be shifted on to the non-Clifford resource. It is easy to see that this attack could lead to improvements for any Clifford-based architecture provided that the noise can be shifted around the circuit easily. In the example considered above, this 'shifting' was possible because the N_p^Z noise commutes with any $U(\theta)$. Of course in general the noise models will not commute with the non-Clifford resource so straightforwardly, but as we shall see later, there are important examples of fault-tolerance schemes in the literature for which this approach can be applied quite effectively.

4. Error-per-gate (EPG) bounds for Clifford operation and phase gate or phase state schemes

The bounds of the previous section apply to a dephasing noise model that is independent of different physical qubit wires. However, a similar approach can be used to derive a bound valid for an error-per-gate model, where each nontrivial operation is subject to probabilistic noise, but different physical qubits undergoing the same multiqubit gate (e.g. a CNOT) can

Figure 2. A possible scenario in a fault-tolerant circuit.

have correlations in noise at that location in the circuit. The error-per-gate model is weaker than the 'local stochastic' type models that have been adopted in some rigorous estimates of lower bounds, so upper bounds valid for EPG are also valid upper bounds for a local stochastic model.

In this case the analysis must be changed slightly. A circuit such as that in figure 2 may occur, where the same Clifford operation, in this case the CNOT, touches two non-Clifford operations. In such cases the EPG model allows correlated noise across the two qubits (see noise model definitions), and this correlated noise may not factorize in a way that independently attacks each of the non-Clifford gates next to the CNOT. However, if we parametrize the EPG noise strength by p, then in this case at locations a, b we are allowed to choose to apply the noisy operation $N_t^Z \otimes N_t^Z$ where t is chosen so that the overall probability of any error at locations a, b (of figure 2) together, $t^2 + 2t(1-t)$, is equal to p, i.e.

$$p = t^2 + 2t(1-t) \Rightarrow t = 1 - \sqrt{1-p}. \tag{15}$$

Now the noise affecting each non-Clifford operation in circuit 2 is effectively independent. As in this section we assume the non-Clifford operations to be a phase rotation gate (a similar argument follows through for preparation of a phase state, with precisely the same numerical values), our goal is to work out when

$$(1 - 2p)(1 - 2(1 - \sqrt{1-p})) = 1 - 2q = \frac{1}{\sqrt{2}} \tag{16}$$

which has the solution for p of

$$p = 10.41\%. \tag{17}$$

So, under an EPG model, any circuit consisting of noisy Clifford operations and phase states and phase gates can have a fault tolerance threshold not higher than 10.41%.

5. Independent depolarizing thresholds for general non-Clifford unitaries

Depolarizing noise commutes with all single-qubit unitaries; hence, provided that the non-Clifford operations are single-qubit unitaries, the above reasoning applies straightforwardly. In this section, we now allow U to be any non-Clifford single-qubit unitary (although it turns out that the most robust gate will be the $\pi/8$ gate anyway). Let us ask what p is required such that (where, as before, time increases from left to right, I denotes the identity operation and

U represents the quantum operation acting by conjugation on density matrices, not just the unitary matrix itself)

$$U(N_p^D)^2 = U((1-p)I + pD)^2$$
$$= (1-p)^2 U + (1 - (1-p)^2)D$$

is a Clifford operation (the second line of this equation follows from the fact that $D^2 = DU = UD = D$ as D is the depolarizing operation). We may now apply the bound obtained in [5] (a related result has been derived independently by [8]) to find that the minimal such p satisfies

$$(1-p)^2 = 1 - \frac{6 - 2\sqrt{2}}{7} = 100\% - 45.31\% \tag{18}$$

hence

$$p = 1 - \sqrt{1 - \frac{6 - 2\sqrt{2}}{7}} = 26.05\%. \tag{19}$$

This upper bound applies regardless of the single-qubit non-Clifford gates available and the fault tolerance methods used, provided that the basic gates are Clifford + single-qubit unitaries. In [5] it is shown that the $\pi/8$ gate (i.e. $U(\pi/4)$) is the most robust gate to their depolarizing noise, so in this setting it is also the most robust.

Reichardt [22] has independently previously performed a related analysis for *simultaneous* depolarizing noise. The calculations are more difficult because in that case the noise does not commute straightforwardly. However, it is expected that an analysis of the simultaneous case would lead to upper bounds of a similar magnitude [22].

6. Specializing to teleportation state-injection schemes

If general bounds are required for architectures involving Clifford and non-Clifford operations, then it is generally not possible to shift more than two lots of noise on to each non-Clifford operation. This is because, in principle, every second operation in a given qubit wire could be a non-Clifford operation.

However, in many important proposals for fault-tolerant quantum computation the non-Clifford operations are always surrounded by specific configurations of Clifford operations. This is because non-Clifford operations are often introduced into the computation using very specific ancilla preparation constructions. For example, one method that we will consider here is the technique of 'state-injection'—see figure 3. This method involves firstly creating a physical non-Clifford qubit, either by direct access to a source of such qubits or by applying a gate such as the $\pi/8$ gate to a suitable Pauli eigenstate. This qubit is then 'teleported' into an error correcting code, by first decoding one half of an encoded Bell pair to the physical level (post-selecting on no errors), and then performing ordinary teleportation using CNOTs and X, Z measurements. Because the teleportation circuit immediately around the non-Clifford resource has around 5–7 nontrivial error locations (depending upon the precise model), one can shift up to 7 lots of noise on to the non-Clifford resource, and obtain much lower upper bounds. By allowing noise at these locations, and then remodelling this as effective noise acting only upon the non-Clifford resource, it is possible to strengthen the bounds derived in the previous sections.

Figure 3. These circuits show two possible state injection settings. The boxes denote gates, the triangles are either preparation (of the +1 eigenstate) or measurement in the X, Z bases, and the collection of wires at the bottom left of each circuit represents the postselected decoding process. The letters 'nc' denote 'non-Clifford'. The top circuit is for scenarios involving single-qubit preparation as the non-Clifford resource, and the second is for single-qubit gates. Each diagram has locations numbered from 1 to 6, with the lower figure having an extra location 7. In the text we use these numbers to define noise models for both the upper and lower circuits in this figure. For state injection protocols these circuits represent the only way in which non-Clifford operations enter the computation.

To illustrate the approach, let us consider what happens when in the top circuit of figure 3; we apply various types of Pauli errors at locations 1–6. It turns out that the teleportation circuit has the special property that Pauli errors acting anywhere in the circuit either can be discarded or can be considered as Pauli errors acting only at location 1. The following arguments explain why this is the case, by considering how Pauli errors at locations 2–6 can be discarded or reconsidered as errors effectively acting at location 1.

1. Application of Z at location 6, and no errors elsewhere. Because location 6 is immediately followed by a Z measurement, this case is essentially equivalent to no error as it is 'absorbed' by the measurement.

2. Application of X at location 6, and no errors elsewhere. An X on the target wire commutes with the CNOT, and so the X can in fact be commuted through to location 4.

3. Application of Pauli error at location 4, and no errors elsewhere. Now because the CNOT and subsequent X, Z measurements implement a Bell measurement, and because each projector B onto a Bell state satisfies $(P \otimes I)B(P \otimes I) = (I \otimes P)B(I \otimes P)$ for any Pauli operation P, a Pauli P at location 4 is equivalent to a P acting at location 1. Putting this together with the previous point shows that an X at location 6 is equivalent to an X at location 1.

4. Application of Y at location 6, and no errors elsewhere. A *Y* error as a *quantum operation* (not as a matrix) is equivalent to a *Z* error and an *X* error. The *Z* can be absorbed by the *Z* measurement on the lower wire, leaving an *X* error that can be moved to location 1 by the previous points. Hence a *Y* at location 6 is equivalent to an *X* at location 1.

5. Any Pauli errors at locations 4, 5 and 6. Because Pauli matrices either commute or anticommute, when viewed as *quantum operations* Pauli errors actually commute with each other—consider for example the identity $XZ\rho ZX = ZX\rho XZ$. Hence the previous rules can be applied to any set of Pauli errors acting at locations 5 and 6—effectively re-expressing them as Pauli errors acting with various probabilities at location 1.

6. Any Pauli errors at locations 2 and 3. Any *X* errors can be absorbed by the *X* measurement on the top wire, and any *Z* errors can be commuted through the CNOT to location 1.

We consider two variants of the state injection schemes. In *state* resource variants, the non-Clifford resource is a pure qubit state as in the top circuit of figure 3; in *gate* resource variant, the non-Clifford resource is a single-qubit unitary, as in the bottom circuit of figure 3.

The above rules can be applied to any configuration of Pauli noise at locations 2–6 to shift it all to location 1, where it can attack the non-Clifford resource. As the non-Clifford resource at location 1 is effectively a state in the Bloch sphere, we can solve the relatively simple problem of how much of this noise forces the state to enter the Clifford 'octahedron' (cf [9, 18]) formed from the convex hull of Pauli eigenstates. In the case of upper bounds for an EPG model, we are also free to try to pick the most adversarial noise we can within the EPG constraints. Consider a pure Bloch vector

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} \tag{20}$$

in the positive octant of the Bloch sphere (i.e. $x, y, z \geqslant 0$). A *Z* error flips the sign of x, y, and an *X* error flips the sign of y, z. Our goal is to find the minimal noise, for a given noise model, which takes the input Bloch vector to an output one

$$\text{Noise}: \begin{pmatrix} x \\ y \\ z \end{pmatrix} \rightarrow \begin{pmatrix} x' \\ y' \\ z' \end{pmatrix}, \tag{21}$$

satisfying $x' + y' + z' = 1$, which is the equation of the face of the octahedron in the positive octant (see e.g. figure 4).

7. Bounds for the Knill noise model and teleportation state-injection, for phase state and phase gate resources

To be perfectly clear, it is important to specify precisely how we apply Knill's noise model to the circuits in figure 3. In the top circuit of figure 3 we apply: a *Z* at locations 1 and 3 with probability $4\gamma/15$; an *X* at location 6 with probability $4\gamma/15$; and at locations 2 and 5 considered together we apply each of the 15 non-identity pairs of operators chosen from *I*, *X*, *Y* and *Z*, each pair with probability $\gamma/15$. Note that we have kept location 4 error free. The noise at location 4 will be determined by the decoding circuit that feeds it, and so in order to obtain general bounds independent of the codes used, we are adopting a noise model that is *strictly*

Figure 4. A two-dimensional projection of the Bloch sphere. The goal of the calculations in this paper is to determine the noise strength required for an input state to enter the octahedron formed from the convex hull of the Pauli eigenstates.

weaker than Knill's. Later we will discuss the effect that noise in the decoding circuits might have.

In the bottom circuit of figure 3, on the other hand, we apply the noise model: at 3 and 7 we apply a Z with probability $4\gamma/15$; at 6, an X with probability $4\gamma/15$; at 1, we apply an X, Y and Z, each with probability $4\gamma/15$; and at locations 2 and 5 considered together, we apply each of the 15 non-identity pairs of operators chosen from I, X, Y and Z, each pair with probability $\gamma/15$. Again we keep location 4 error free.

Using this noise model leads to the following effective transformation of the input Bloch vector:

$$\begin{pmatrix} x' \\ y' \\ z' \end{pmatrix} = \left(1 - \frac{16\gamma}{15}\right)^n \left(1 - \frac{8\gamma}{15}\right) \begin{pmatrix} (1 - 8\gamma/15)x \\ (1 - 8\gamma/15)^2 y \\ z \end{pmatrix},$$

where $n = 1$ for the top circuit of figure 3, and $n = 2$ for the lower circuit of figure 3. Our goal is hence to determine, for a given input resource, the minimal γ such that the output Bloch vector lies on the face of the octahedron. This involves finding the least positive root of a polynomial, which can be done rapidly using standard mathematical software.

For example, if we assume that the ideal state entering location 1 in both circuits is $|\pi/4\rangle$, either because that is the state prepared or because the non-Clifford unitary is the $\pi/8$ gate, then we find the solution

$$(1 - 16\gamma/15)(1 - 8\gamma/15)^2(1 + (1 - 8\gamma/15)) = \sqrt{2}$$
$$\Rightarrow \gamma \sim 13.6861\% \tag{22}$$

for the top circuit of figure 3, and

$$(1 - 16\gamma/15)^2(1 - 8\gamma/15)^2(1 + (1 - 8\gamma/15)) = \sqrt{2}$$
$$\gamma \quad 9.5858\% \tag{23}$$

for the lower circuit of figure 3. Similar equations can be derived and solved for any possible phase gate and phase state resource. It is not difficult to solve for the minimal γ that leads to an output vector on the face of the octahedron. Numerically scanning through these solutions suggests that the $|\pi/4\rangle$ and $U(\pi/4)$ resources are not actually the most robust non-Clifford phase state or phase gate resources in this setting, although they are very close. If we allow *all* phase state and phase gate resources, respectively, the upper bounds become

$$\gamma \sim 13.71\% \tag{24}$$

for phase states, and

$$\gamma \sim 9.59\% \tag{25}$$

for phase gates. Later we present the values obtained if any single-qubit state or gate is permitted as the non-Clifford resource.

8. Bounds for an error-per-gate (EPG) noise model and for phase gate or phase state teleportation state injection

To obtain upper bounds for an EPG model, we select noise that appears to be as detrimental as possible, yet sufficiently simple to analyse. It is quite likely that the noise we pick at each location is not the most detrimental (in the sense of pulling the teleportation circuit into a Clifford operation) within the EPG constraint, but this would require further analysis.

We choose the following noise for the top circuit of figure 3. Assume that the input non-Clifford state is $|\theta\rangle$. Apply $O_p^{|\theta\rangle\langle\theta|}$ at location 1; N_p^Z at locations 3 and 4; N_p^X at location 6; and replace the CNOT with $O_p^{|\theta\rangle\langle\theta|}$ (see equation (6)) on the top wire followed by a CNOT [8]. This gives the following transformation on the input Bloch vector (which has $z = 0$ as it is a phase state):

$$\begin{pmatrix} x \\ y \end{pmatrix} \to (1-2p)^2((1-p)(1-2p)-p)\begin{pmatrix} x \\ (1-2p)y \end{pmatrix}.$$

For an input state $|\theta = \pi/4\rangle$, the minimal p such that the output Bloch vector lies on the face of the octahedron is

$$p = 3.68124\%. \tag{26}$$

Numerics again suggest that $|\theta = \pi/4\rangle$ is not the most robust state in this setting, and in fact

$$p = 3.69\% \tag{27}$$

is a noise level sufficient to take all possible phase states into the octahedron.

For the lower circuit of figure 3 we choose the following noise. Assume that the input state is $|+\rangle$, the $+1$ eigenstate of X, and that U is the non-Clifford phase gate. Apply $O_p^{|+\rangle\langle+|}$ at location 7; $O_p^{U|+\rangle\langle+|U^\dagger}$ at location 1; N_p^Z at locations 3 and 4; N_p^X at location 6; and replace the CNOT with $O_p^{U|+\rangle\langle+|U^\dagger}$ on the top wire followed by a CNOT. This gives the following transformation on the output Bloch vector (which has $z = 0$ as it is a phase state):

$$(1-2p)^2((1-p)(1-4p+2p^2)-p)\begin{pmatrix} x \\ (1-2p)y \end{pmatrix}.$$

---

[8] In fact, for the situations considered here, the use of opposite noise only seems to make the bounds slightly tighter than using appropriately chosen Pauli errors, but we include it as it does appear to be slightly more detrimental, at least when the other errors are chosen as they are.

For an input state $U(\theta = \pi/4)$ the minimal p such that the output Bloch vector lies on the face of the octahedron is

$$p = 3.00339\%. \tag{28}$$

Numerics suggest that $U(\theta = \pi/4)$ is not the most robust phase gate in this setting, and in fact

$$p = 3.01\% \tag{29}$$

is a noise level sufficient to take all possible phase states into the octahedron.

9. Bounds for teleportation injection with general, non-Clifford states and gates

It is not difficult to perform the analysis of the previous two sections using general single-qubit non-Clifford gate and state resources, and then to numerically calculate the most robust of these resources. For the Knill noise model we find that for general unitary gates

$$\gamma = 15.19\% \tag{30}$$

is sufficient to turn the teleportation state-injection into a Clifford circuit, whereas for general states

$$\gamma = 21.78\% \tag{31}$$

is sufficient. The most robust non-Clifford state entering the circuit appears to be close to, but not exactly the same as, the so-called $|T\rangle$ state (where the $|T\rangle$ state is defined [9] as the pure state equivalent to the Bloch-sphere qubit density matrix $1/2(I + 1/\sqrt{3}(X + Y + Z)))$.

For an EPG model we use the same noise as for the phase gates/states in the previous section, except that at location 4 instead of applying N_p^Z we apply N_p^Y. Numerical analysis of the resulting equations reveals that

$$p = 5.03\% \tag{32}$$

is sufficient to turn the teleportation state-injection into a Clifford circuit, whereas for general states

$$p = 6.31\% \tag{33}$$

is sufficient.

10. Potential effects of the decoding circuits

So far, we have considered solely the injection part of the circuit as depicted in figure 3, and have not attempted to understand the effects of noise in the decoding circuit. Indeed, within the framework of Knill's noise model we have not even allowed noise at location 4. The precise form of the noise generated by the decoding circuit is difficult to determine due to the complex structure of the encoding and decoding networks, which may include many steps of concatenation. This brings with it two problems. Firstly, the entangling gates in the decoder typically generate correlations in noise, which mean that the problem becomes a multiqubit problem, rather than a simple geometrical problem on one or two qubits. Secondly, the postselection steps have an effect on the noise profile, and this is difficult to calculate.

Figure 5. This circuit shows the state-injection circuit of figure 3 (lines 1 and 2) but now we explicitly include the action of part of a decoding circuit with the Knill error model. The input of the circuit is the state $(|\psi_0\rangle|0\rangle + |\psi_1\rangle|1\rangle)/\sqrt{2}$ where the state $|\psi_i\rangle$ will be decoded by an error-free circuit to $|i\rangle$ in the second qubit. Errors from previous stages of encoding and manipulation are modelled by depolarizing errors. The error-free state-injection would then lead to the preparation of the state $|\psi\rangle$ in the output qubit at the bottom. Determining the minimal error rate that ensures that the output is a stabilizer state provides an upper bound on the fault tolerance threshold.

However, we believe that a careful analysis of the decoding circuits will lead to improvements in the bounds that we have presented. In this section, we present some very rough indications of the level of improvement that might be expected, although a more careful analysis is much more difficult and is code dependent. The estimates that we obtain here are not included in table 1 as they are not rigorously justified. We hope that by developing better techniques, perhaps inspired by analysis of lower bounds, future work may be able to analyse the decoding/encoding circuits more thoroughly than we do here.

To get a feeling for the effect of the errors that are introduced by the decoding circuit that immediately precedes position 4, we have carried out an analysis of the circuit in figure 5, which corresponds to one level of decoding prior to the injection circuit for a particular code. The input of the top arm of the teleporter is a phase state proportional to $|0\rangle + \exp(i\pi/4)|1\rangle$. This circuit is a simplified version of figure 9 of [23]. The simplification consists of neglecting essentially all errors that may occur in the preparation and encoding, as well as the decoding of the lower half of that circuit of figure 9 of [23]. We also model the error on the incoming qubits in figure 5 via depolarizing errors (it is in this assumption that most of the complexity is buried).

The computation of the error threshold in this approach is possible analytically because of the small circuit size. With the help of a computational mathematical package we find the threshold to be the sole real root (the Mathematica worksheet used to compute this may be

obtained by contacting the authors) of the polynomial

$$p(e_c) = -\frac{1}{2} - \frac{f_1(e_c)}{f_2(e_c)}$$

with

$$f_1(e_c) = (16e_c - 15)^5 (225 - 180e_c + 32e_c^2)^2 (16875 - 40500e_c + 46800e_c^2 - 23040e_c^3 + 4096e_c^4)$$

$$f_2(e_c) = 1125\sqrt{2}(576650390625 - 3536789062500e_c + 11768793750000e_c^2$$
$$- 24002325000000e_c^3 + 32367600000000e_c^4 - 29499033600000e_c^5$$
$$+ 18141511680000e_c^6 - 7375159296000e_c^7 + 1887436800000e_c^8$$
$$- 273804165120e_c^9 + 17179869184e_c^{10})$$

and we find

$$e_c^{\text{threshold}} = 9.29\%. \tag{34}$$

One can try to understand how critical is the choice of depolarizing errors entering the decoding steps. To estimate this we also analysed what happens if the depolarizing errors in the central four wires of figure 5 are removed. In this case the bound becomes

$$e_c^{\text{threshold}} = 10.06\%. \tag{35}$$

Both these bounds are quite close to the bound that would be obtained by instead modelling the noise from the decoding circuit as a depolarizing error (of strength γ using Knill's noise model) at location 4 (for which the bound would be 9.59%). All these numbers should be compared to 13.69%, which is the number obtained for precisely the same phase state input with no noise whatsoever outside locations 1–3 and 5–7 of figure 3.

A more rigorous analysis of the decoding circuits requires much more effort and may be a subject for future work. However, these rough calculations give some indication of the improvements that might be expected in the upper bounds that we have obtained so far.

11. Discussion, caveats and conclusion

We have discussed 'attacks' on fault-tolerant quantum circuits involving Clifford operations and extra resources, with the intention of adding the smallest amount of noise possible to make the circuits efficiently simulatable classically. The approach is simple—to shift noise from neighbouring Clifford gates on to the non-Clifford resources. The approach works best for situations where this 'shifting' process can be done easily, as happens for teleportation-based state injection schemes. There are many other fault-tolerance proposals in the literature (e.g. methods built around cluster states [24]) that involve a few non-Clifford resources surrounded by many Clifford operations—and so in such cases our approach could also provide improved upper bounds, depending upon the precise manner in which the 'non-Cliffordness' is 'injected' into the rest of the circuit.

We are certain that the upper bounds obtained in this work can be optimized further, particularly through consideration of the decoding circuits. The rigorous part of the analysis

performed in this work does not consider the decoding circuit present in state-injection schemes, except for a noise step at location 4 in the error-per-gate noise model. The analysis of the teleportation part of the state-injection is straightforward because the teleportation circuit simplifies correlated Pauli errors, reducing the problem to a single-qubit one. A more sophisticated attack that also applies noise to the decoding steps should lead to tighter bounds, especially for Knill's noise model.

It is also important to note that teleportation is not the only method of state injection. If the non-Clifford resource is either a measurement or a unitary, a non-Clifford measurement may be implemented directly on the decoded wire[9]. This method involves far fewer error locations outside the decoder, and hence is less susceptible to our method of attack. Hence the bounds derived in the lower part of table 1 are not generally applicable to all stabilizer schemes, but are intended more as an approach that can be fairly effective for specific fault tolerance proposals.

A summary of the upper bounds that we have obtained is given in table 1. Although the bounds that we have derived are lower than previous rigorously established upper bounds, our bounds are not typically as general, as most previous results make fewer assumptions about the architecture (and often involve incomparable noise models). For comparable fault-tolerance schemes, the estimates conjectured in [7] are lower than the values that we have obtained, and future work will be needed to see whether further optimization of our approach will go as low. However, the approach presented here is relatively simple, makes no assumptions (other than assuming that quantum computation is not classically tractable), and has a slightly different aim as it deals with 'classical' bounds rather than 'quantum'.

Acknowledgments

References

[1] Knill E 2005 *Nature* **434** 39
[2] Steane A 2003 *Phys. Rev.* A **68** 042322
   Aharonov D and Ben-Or M 1997 *Proc. 29th ACM STOC* p 176 (arXiv:quant-ph/9906129)
   Kitaev A 1997 *Russ. Math. Surveys* **52** 1191
   Gottesman D 1998 *Phys. Rev.* A **57** 127
   Reichardt B 2006 *PhD Thesis* arXiv:quant-ph/0612004
   Aliferis P 2007 *PhD Thesis* arXiv:quant-ph/0703230
[3] Harrow A and Nielsen M 2003 *Phys. Rev.* A **68** 012308 (arXiv:quant-ph/0301108)
[4] Virmani S, Huelga S F and Plenio M B 2005 *Phys. Rev.* A **71** 042328
[5] Buhrman H, Cleve R, Laurent M, Linden N, Schrijver A and Unger F 2006 *Proc. 47 Symp. FOCS*
[6] Kempe J, Regev O, Unger F and de Wolf R 2008 *Proc. 35th ICALP* arXiv:0802.1464

[9] See e.g. [1]. In fact, the ideal teleportation part of the injection circuit can be interpreted as a four outcome POVM (positive operator valued measure) acting upon location 4 in the circuits of figure 3, where each POVM element $\{M_1, .., M_4\}$ is rank-1. Doing an ordinary measurement containing the projector proportional to the ideal outcome, hence, has the same effect (and is also more efficient).

[7] Fern J 2008 arXiv:0801.2608

[8] Bravyi S 2005 private communication

[9] Bravyi S and Kitaev A 2005 *Phys. Rev.* A **71** 022316

[10] Aharonov D, Ben-Or M, Impagliazzo R and Nisan N 1996 arXiv:quant-ph/9611028

[11] Razborov A A 2004 *Quantum Inf. Comput.* **4** 222–8 (arXiv:quant-ph/0310136)

[12] Kay A 2008 *Phys. Rev.* A **77** 052319

[13] Nielsen M A and Chuang I L 2000 *Quantum Information and Computation* (Cambridge: Cambridge University Press)

[14] Gottesman D 1997 *PhD Thesis* arXiv:quant-ph/9705052

[15] Jozsa R and Linden N 2003 *Proc. R. Soc.* A **459** 2011

[16] Aliferis P, Gottesman D and Preskill J 2008 *Quantum Inf. Comput.* **8** 181–244 (arXiv:quant-ph/0703264)

[17] Reichardt B W 2006 *FOCS 2006* (arXiv:quant-ph/0608018)

[18] Reichardt B W 2005 *Quantum Inf. Proc.* **4** 251

[19] Reichardt B W 2006 arXiv:quant-ph/0608085v2

[20] Campbell E and Browne D 2009 arXiv:0908.0836v2

[21] van Dam W and Howard M 2009 *Phys. Rev. Lett.* **103** 170504

[22] Reichardt B 2008 private communication

[23] Knill E 2004 arXiv:quant-ph/0402171v1

[24] Raussendorf R and Briegel H J 2001 *Phys. Rev. Lett.* **86** 5188