Upper Bounds on the Noise Threshold for Fault-tolerant Quantum Computing

Julia Kempe[1], Oded Regev[1], Falk Unger[2], and Ronald de Wolf[2]

[1] Department of Computer Science, Tel-Aviv University, Tel-Aviv, Israel.
[2] CWI, Amsterdam, The Netherlands.

Abstract. We prove new upper bounds on the tolerable level of noise in a quantum circuit. Our circuits consist of unitary k-qubit gates each of whose input wires is subject to depolarizing noise of strength p, and arbitrary one-qubit gates that are essentially noise-free. We assume the output of the circuit is the result of measuring some designated qubit in the final state. Our main result is that for $p > 1 - \Theta(1/\sqrt{k})$, the output of any such circuit of large enough depth is essentially independent of its input, thereby making the circuit useless. For the important special case of $k = 2$, our bound is $p > 35.7\%$. Moreover, if the only gate on more than one qubit is the CNOT, then our bound becomes 29.3%. These bounds on p are notably better than previous bounds, yet incomparable because of the somewhat different circuit model that we are using. Our main technique is a Pauli basis decomposition, which we believe should lead to further progress in deriving such bounds.

1 Introduction

The field of quantum computing faces two main tasks: to build a large-scale quantum computer, and to figure out what it can do once it exists. In general the first task is best left to (experimental) physicists and engineers, but there is one crucial aspect where theorists play an important role, and that is in analyzing the level of noise that a quantum computer can tolerate before breaking down.

The physical systems in which qubits may be implemented are typically tiny and fragile (electrons, photons and the like). This raises the following paradox: On the one hand we want to isolate these systems from their environment as much as possible, in order to avoid the noise caused by unwanted interaction with the environment—so-called "decoherence". But on the other hand we need to manipulate these qubits very precisely in order to carry out computational operations. A certain level of noise and errors from the environment is therefore unavoidable in any implementation, and in order to be able to compute one would have to use techniques of error correction and fault tolerance.

Unfortunately, the techniques that are used in classical error correction and fault tolerance do not work directly in the quantum case. Moreover, extending these techniques to the quantum world seems at first sight to be nearly impossible due to the continuum of possible quantum states and error patterns. Indeed,

when the first important quantum algorithms were discovered [1–4], many dismissed the whole model of quantum computing as a pipe dream, because it was expected that decoherence would quickly destroy the necessary quantum properties of superposition and entanglement.

It thus came as a great surprise when, in the mid-1990s, *quantum error correcting codes* were developed by Shor and Steane [5–7], and these ideas later led to the development of schemes for *fault-tolerant quantum computing* [8–12]. Such schemes take any quantum algorithm designed for an ideal noiseless quantum computer, and turn it into an implementation that is robust against noise, as long as the amount of noise is below a certain threshold, known as the *fault-tolerance threshold*. The overhead introduced by the fault-tolerant schemes is typically modest (a polylogarithmic factor in the running time of the algorithm).

The existence of fault-tolerant schemes turns the problem of building a quantum computer into a hard engineering problem: if we just manage to store our qubits and operate upon them with a level of noise below the fault-tolerance threshold, then we can perform arbitrarily long quantum computations. The actual *value* of the fault-tolerance threshold is far from determined, but will have a crucial influence on the future of the area—the more noise a quantum computer can tolerate in theory, the more likely it is to be realized in practice.[3]

The first fault-tolerant schemes were only able to tolerate noise on the order of 10^{-6}, which is way below the level of accuracy that experimentalists can hope to achieve in the foreseeable future. These initial schemes have been substantially improved in the past decade. In particular, Knill has recently developed various schemes which, according to numerical calculations, seem to be able to tolerate more than 1% noise [13, 14]. If we insist on provable constructions, the best known threshold is on the order of 0.1% [15–18].

Constructions of fault-tolerant schemes provide a *lower bound* on the fault-tolerance threshold. A very interesting question, which is the topic of the current paper, is whether one can prove *upper bounds* on the fault-tolerance threshold. Such bounds give an indication on how far away we are from finding optimal fault-tolerant schemes. They can also give hints as to how one should go about constructing improved fault-tolerant schemes. Such upper bounds are statements of the form "any quantum computation performed with noise level higher than p is essentially useless", where "essentially useless" is some strong indication that interesting quantum computations are impossible in such a model. For instance, Buhrman et al. [19] quantify this by giving a classical simulation of such noisy quantum computation, and Razborov [20] shows that if the computation is too long, the output of the circuit is essentially independent of its input.

The best known upper bounds on the threshold are 50% by Razborov [20] and 45.3% by Buhrman et al. [19]. (These bounds are incomparable because they work in different models; see the end of this section for more details.) As one can see, there are still about two orders of magnitude between our best upper and lower bounds on the fault-tolerance threshold. This leaves experimentalists

[3] The "fault-tolerance threshold" is actually not a universal constant, but rather depends on the details of the circuit model (allowed set of gates, type of noise, etc.).

in the dark as to the level of accuracy they should try to achieve. In this paper, we somewhat reduce this gap. So far, much more work has been spent on lower bounds than on upper bounds. Our approach will be the less-trodden road from above, hoping to bring new techniques to bear on this problem.

Our model. In order to state our results, we need to describe our circuit model. We consider parallel circuits, composed of n *wires* and T *levels* of gates (see Figure 1). We sometimes use the term *time* to refer to one of the $T+1$ "vertical cuts" between the levels. For convenience, we assume that the number of qubits n does not change during the computation. Each level is described by a partition of the qubits, as well as a gate assigned to each set in the partition. Notice that at each level, all qubits must go through some gate (possibly the identity). For each gate the number of input qubits equals the number of output qubits.

We assume the circuit is composed of k-qubit gates that are probabilistic mixtures of unitary operations, as well as arbitrary (i.e., all completely-positive trace-preserving) one-qubit gates. In particular, it is possible to do intermediate one-qubit measurements. We assume the output of the circuit is the outcome of a measurement of a designated output qubit in the computational basis. Finally, we assume that the circuit is subject to noise as follows. Recall that p-depolarizing noise on a certain qubit replaces that qubit by the completely mixed state with probability p, and does not alter the qubit otherwise. Formally, this is described by the superoperator $\mathcal{E}$ acting on a qubit ρ as $\mathcal{E}(\rho) = (1-p)\rho + pI/2$. We assume that each one-qubit gate is followed by at least ε_1-depolarizing noise on its output qubit, where $\varepsilon_1 > 0$ is an arbitrarily small constant. Thus one-qubit gates can be essentially noise-free. We also assume that each k-qubit gate is preceded by at least ε_k-depolarizing noise on each of its input qubits, where $\varepsilon_k > 1 - \sqrt{2^{1/k} - 1} = 1 - \Theta(1/\sqrt{k})$.

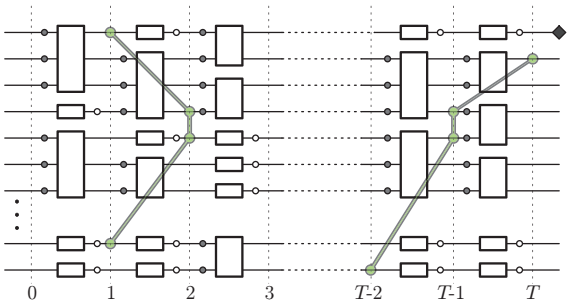

Fig. 1. Parallel circuit with $k = 3$ and T levels. Dark circles are ε_k-depolarizing noise, light circles are ε_1-depolarizing noise. We marked two consistent 4-qubit sets (defined in Section 3). The first has distance 1, the second $T-2$. The upper right qubit is the output.

Our results. In Section 3 we prove our main result:

Theorem 1. *Fix any T-level quantum circuit as above. Then for any two states ρ and τ, the probabilities of obtaining measurement outcome 1 at the output qubit starting from ρ and starting from τ, respectively, differ by at most $2^{-\Omega(T)}$.*

In other words, for any $\eta > 0$, the probability of measuring 1 at the output qubit of a circuit running for $T = O(\log(1/\eta))$ levels is independent of the input (up

to $\pm\eta$). This makes the output essentially independent of the starting state, and renders long computations "essentially useless".

Of special interest from an experimental point of view is the case $k = 2$, for which our bound becomes about 35.7%. Furthermore, for the case in which the only allowed two-qubit gate is the controlled-NOT (CNOT) gate, we can improve our bound further to about 29.3%, as we show in the full version of this paper [21]. This case is interesting both theoretically and experimentally. Note also that the CNOT gate together with all one-qubit gates forms a universal set [22]. The same noise-bound applies if we also allow controlled-Y and controlled-Z gates.

Significance of results. First, it is known that fault-tolerant quantum computation is impossible (for any positive noise level) without a source of fresh qubits. Our model takes care of this by allowing arbitrary one-qubit gates—in particular, this includes gates that take any input, and output a fixed one-qubit state, for instance $|0\rangle$. This justifies our assumption that the number of qubits in the circuit remains the same throughout the computation: all qubits can be present from the start, since we can reset them to whatever we want whenever needed.

Second, our assumption that all k-qubit gates are mixtures of unitaries does slightly restrict generality. Not every completely-positive trace-preserving map can be written as a mixture of unitaries.[4] However, we believe that it is still a reasonable assumption. For instance, to the best of our knowledge, all known fault-tolerant constructions can be implemented using such gates (in addition to arbitrary one-qubit gates). Moreover, all known quantum algorithms obtain their speed-up over classical algorithms by using only unitary gates.

Third, we only analyze depolarizing noise acting independently on each qubit. Depolarizing noise is the "most symmetric" type of one-qubit noise and therefore a natural choice for our analysis. Also, it is a relatively weak type of noise: it is not adversarial and does not have correlations between the errors occurring on different qubits. However, since we are proving an *upper bound* on the fault-tolerance threshold, this weakness is actually a good thing, making our result stronger. In principle one can extend our results to various other one-qubit noise models, using an analysis similar to the one developed in Lemma 2. However, not all noise models can actually yield a result like ours. For instance, if we have Toffoli gates with only phaseflip errors, then we can do perfect classical computation. Statements like Theorem 1 are just false in that case.

A more severe restriction is the assumption that the output consists of one qubit. However, we believe that in many instances this is still a reasonable assumption, for instance when the circuit is solving a decision problem. Moreover, our results can easily be extended to the case where the output is obtained by a measurement on a small number of qubits, instead of only one.

By allowing essentially noise-free one-qubit gates, our model addresses the fact that gates on more than one qubit are generally much harder to implement than one-qubit gates. It should also be noted that the exact value of the constant

[4] One can implement an arbitrary gate by a unitary gate on the original qubits and additional ancilla qubits in a fixed pure state. However, this increases the arity of the gate, and the ancilla qubits will be affected by the noise before the unitary.

ε_1 is inessential and can be chosen arbitrarily small, as this just affects the constant in the $\Omega(\cdot)$ of Theorem 1. In fact, $\varepsilon_1 > 0$ is only necessary because otherwise it would be possible to let $\rho := |0\rangle\langle 0| \otimes \rho'$ and $\tau := |1\rangle\langle 1| \otimes \tau'$, do nothing for T levels (i.e., apply noise-free identity gates on all wires) and then measure the first qubit. The resulting difference between output probabilities is 1. Instead of assuming $\varepsilon_1 > 0$ noise, we could alternatively deal with this issue by requiring that every path from the input to the output qubit goes through enough k-qubit gates. Our proof can easily be adapted to this case.

Since our theorem applies to arbitrary starting states, it applies to the case that the initial state is encoded in a good quantum error-correcting code, or is some sort of "magic state" [23, 24]. Also in these case, the computation becomes essentially independent of the input after sufficiently many levels.

Finally, it is interesting to note that our bound on the threshold behaves like $1 - \Theta(1/\sqrt{k})$. This matches what is known for classical circuits [25, 26], and therefore probably represents the correct asymptotic behavior. Previous bounds only achieved an asymptotic behavior of $1 - \Theta(1/k)$ [20].

Techniques. We believe that a main part of our contribution is introducing a new technique for obtaining upper bounds on the fault-tolerance threshold. Namely, we use a Pauli basis decomposition in order to track the state of the computation. A finer analysis of the Pauli coefficients might improve the bounds we achieve here, and possibly obtain bounds for other computational models.

Related work. The work most closely related to ours is that of Razborov [20]. There, he proves an upper bound of $\varepsilon_k = 1-1/k$ on the fault-tolerance threshold. On one hand, his result is stronger than ours as it allows arbitrary k-qubit gates and not just mixtures of unitaries. Razborov also has a second result, namely the trace distance between the two states obtained by applying the circuit to starting states ρ and τ, respectively, is upper bounded by $n2^{-\Omega(T)}$. Hence even the results of arbitrary n-qubit measurement on the full final state become essentially independent of the initial state after $T = O(\log n)$ levels. On the other hand, the value of our bound is better for all values of k, and we also allow essentially noise-free one-qubit gates. Hence the two results are incomparable. Razborov's proof is based on tracking how the trace distance evolves during the computation. Our proof is similar in flavor, but instead of working with the trace distance, we work with the Frobenius distance (since it can easily be expressed in terms of the Pauli decomposition).

Buhrman et al. [19] show that classical circuits can efficiently simulate any quantum circuit that consists of perfect, noise-free *stabilizer operations* (meaning Clifford gates (Hadamard, phase gate, CNOT), preparations of states in the computational basis, and measurements in the computational basis) and arbitrary one-qubit unitary gates that are followed by 45.3% depolarizing noise. Hence such circuits are not significantly more powerful than classical circuits.[5] This

[5] The 45.3%-bound of [19] is in fact *tight* if one additionally allows perfect classical control (i.e., the ability to condition future gates on earlier classical measurement outcomes): circuits with perfect stabilizer operations and arbitrary one-qubits gates suffering from less than 45.3% noise, can simulate perfect quantum circuits. See [27]

result is incomparable to ours: the noise models and the set of allowed gates are different (and we feel ours is more realistic). In particular, in our case noise hits the qubits going into the k-qubit gates but barely affects the one-qubit gates, while in their case the noise only hits the non-Clifford one-qubit unitaries.

Another related result is by Virmani et al. [28]. Instead of depolarizing noise, they consider "dephasing noise". This models phase-errors: rather than replacing a qubit by the completely mixed state with some probability p, dephasing noise applies the Z-gate with probability $p/2$. Virmani et al. [28] show, among other results, that we can efficiently classically simulate any quantum circuit consisting of perfect stabilizer operations, and one-qubit unitary gates that are diagonal in the computational basis and are followed by more than 29.3% dephasing noise. Their result is incomparable to ours for essentially the same reasons as why the Buhrman et al. result is incomparable: a different noise model and a different statement about the resulting power of their noisy quantum circuits.

Finally, it is known that it is impossible to transmit quantum information through a p-depolarizing channel for $p > 1/3$ [29]. As Razborov [20] noticed, this seems to suggest that quantum computation is impossible with depolarizing noise of strength greater than $1/3$, but there is no proof that this is the case.

2 Preliminaries

Let $\mathcal{P} = \{I, X, Y, Z\}$ be the set of one-qubit Pauli matrices,

$$I = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \; X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \; Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \; Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

and let $\mathcal{P}_* = \{X, Y, Z\}$. We use $\mathcal{P}^n$ to denote the set of all tensor products of n one-qubit Pauli matrices. For a Pauli matrix $S \in \mathcal{P}^n$ we define its *support*, denoted $\mathrm{supp}(S)$, to be the qubits on which S is not identity. We sometimes use superscripts to indicate the qubits on which certain operators act. Thus $I^{\mathcal{A}}$ denotes the identity operator applied to the qubits in set $\mathcal{A}$.

The set of all $2^n \times 2^n$ Hermitian matrices forms a 4^n-dimensional real vector space. On this space we consider the Hilbert-Schmidt inner product, given by $\langle A, B\rangle := \mathrm{Tr}(A^\dagger B) = \mathrm{Tr}(AB)$. Note that for any $S, S' \in \mathcal{P}^n$, $\mathrm{Tr}(SS') = 2^n$ if $S = S'$ and 0 otherwise, and hence $\mathcal{P}^n$ is an orthogonal basis of this space. It follows that we can uniquely express any Hermitian matrix δ in this basis as

$$\delta = \frac{1}{2^n} \sum_{S \in \mathcal{P}^n} \widehat{\delta}(S) S$$

where $\widehat{\delta}(S) := \mathrm{Tr}(\delta S)$ are the (real) coefficients.

We now state some observations. By the orthogonality of $\mathcal{P}^n$, for any δ,

$$\mathrm{Tr}(\delta^2) = \frac{1}{2^n} \sum_{S \in \mathcal{P}^n} \widehat{\delta}(S)^2.$$

and [19, Section 5]. These assumptions are not very realistic: in particular, assuming perfect, noise-free CNOTs is a far cry from experimental practice.

Observation 2 (Unitary preserves sum of squares) *For any unitary matrix U and any Hermitian matrix δ, if we denote $\delta' = U\delta U^\dagger$, then*

$$\sum_{S\in\mathcal{P}^n} \widehat{\delta'}(S)^2 = 2^n \mathrm{Tr}(\delta'^2) = 2^n \mathrm{Tr}(U\delta U^\dagger U\delta U^\dagger) = 2^n \mathrm{Tr}(\delta^2) = \sum_{S\in\mathcal{P}^n} \widehat{\delta}(S)^2.$$

This also shows that conjugating by a unitary matrix, when viewed as a linear operation on the vector of Pauli coefficients, is an orthogonal transformation.

Observation 3 (Tracing out qubits) *Let δ be some Hermitian matrix on a set of qubits W. For $V \subseteq W$, let $\delta_V = \mathrm{Tr}_{W\setminus V}(\delta)$. Then,*

$$\widehat{\delta}(SI^{W\setminus V}) = \mathrm{Tr}(\delta \cdot SI^{W\setminus V}) = \mathrm{Tr}(\delta_V \cdot S) = \widehat{\delta_V}(S).$$

Observation 4 (Noise in the Pauli basis) *Applying a p-depolarizing noise $\mathcal{E}$ to the j-th qubit of Hermitian matrix δ changes the coefficients as follows:*

$$\widehat{\mathcal{E}(\delta)}(S) = \begin{cases} \widehat{\delta}(S) & \text{if } S_j = I \\ (1-p)\widehat{\delta}(S) & \text{if } S_j \neq I \end{cases}$$

In other words, $\mathcal{E}$ "shrinks" by a factor $1-p$ all coefficients that have support on the j-th coordinate.

Observation 5 *Let ρ and τ be two one-qubit states and let $\delta = \rho - \tau$. Consider the two probability distributions obtained by performing a measurement in the computational basis on ρ and τ, respectively. Then the variation distance between these two distributions is $\frac{1}{2}|\widehat{\delta}(Z)|$.*

Proof: Since there are only two possible outcomes for the measurements, the variation distance between the two distributions is exactly the difference in the probabilities of obtaining the outcome 0, which (using $\mathrm{Tr}(\delta) = 0$) is given by

$$|\mathrm{Tr}((\rho-\tau)\cdot|0\rangle\langle 0|)| = \left|\mathrm{Tr}\left(\delta \cdot \frac{I+Z}{2}\right)\right| = \frac{1}{2}|\mathrm{Tr}(\delta\cdot Z)| = \frac{1}{2}|\widehat{\delta}(Z)|. \qquad \blacksquare$$

Our final observation follows immediately from the convexity of the function x^2.

Observation 6 (Convexity) *Let p_i be any probability distribution, and δ_i a set of Hermitian matrices. Let $\delta = \sum_i p_i\delta_i$. Then* $\sum_{S\in\mathcal{P}^n} \widehat{\delta}(S)^2 \leq \sum_i p_i \sum_{S\in\mathcal{P}^n} \widehat{\delta_i}(S)^2$.

3 Proof of Theorem 1

In this section we prove Theorem 1. The idea is the following. Fix two arbitrary initial states ρ and τ. Our goal is to show that after applying the noisy circuit, the state of the output qubit is nearly the same with both starting states. Equivalently, we can define $\delta = \rho - \tau$ and show that after applying the noisy circuit to δ, the "state" of the output qubit is essentially 0 (the noisy circuit is a

linear operation, and hence there is no problem in applying it to δ, which is the difference of two density matrices). In order to show this, we will examine how the coefficients of δ in the Pauli basis develop through the circuit. Initially we might have many large coefficients. Our goal is to show that the coefficients of the output qubit are essentially 0. This is established by analyzing the balance between two opposing forces: noise, which shrinks coefficients by a constant factor (as in Observation 4), and gates, which can increase coefficients. As we saw in Observation 2, unitary gates preserve the sum of squares of coefficients. They can, however, "concentrate" several small coefficients into one large coefficient. One-qubit operations need not preserve the sum of squares (a good example is the gate that resets a qubit to the $|0\rangle$ state), but we can still deal with them by using a known characterization of one-qubit gates. This allows us to bound the amount by which one-qubit gates can increase the Pauli coefficients, and (roughly) shows that the gate that resets a qubit to $|0\rangle$ is "as bad as it gets".

We introduce some terminology. From now on we use the term *qubit* to mean a wire at a specific time, so there are $(T+1)n$ qubits (although during the proof we will also consider qubits that are located between a gate and its associated noise). We say that a set of qubits V is *consistent* if we can meaningfully talk about a "state of the qubits of V" (see Figure 1). More formally, we define a consistent set as follows. The set of all qubits at time 0 and all its subsets are consistent. If V is some consistent set of qubits, which contains all input qubits IN of some gate (possibly a one-qubit identity gate), then also $(V \setminus IN) \cup OUT$ and all its subsets are consistent, where OUT denotes the gate's output qubits. Note that here we think of the noise as being part of the gate. For a consistent set V and a state (or more generally, a Hermitian matrix) ρ, we denote the state of V when the circuit is applied with the initial state ρ, by ρ_V. In other words, ρ_V is the state one obtains by applying some initial part of the circuit to ρ, and then tracing out from the resulting state all qubits that are not in V.

If v is a qubit, we use $\mathrm{dist}(v)$ to denote its distance from the input, i.e., the level of the gate just preceding it. The qubits of the starting state have $\mathrm{dist}(v) = 0$. For a nonempty set V of qubits we define $\mathrm{dist}(V) = \min\{\mathrm{dist}(v) \mid v \in V\}$, and extend it to the empty set by $\mathrm{dist}(\emptyset) = \infty$. Note that $\mathrm{dist}(V)$ does not increase if we add qubits to V. In the rest of this section we prove the following lemma, showing that a certain invariant holds for all consistent sets V.

Lemma 1. *For all $\varepsilon_1 > 0$ and $\varepsilon_k > 1 - \sqrt{2^{1/k} - 1}$ there is a $\theta < 1$ such that the following holds. For any T-level circuit in our model, initial states ρ and τ, $\delta = \rho - \tau$, and any consistent V, we have $\mathrm{Tr}(\delta_V^2) \leq 2 \cdot \theta^{\mathrm{dist}(V)}$. Equivalently:*

$$\sum_{S \in \mathcal{P}^V} \widehat{\delta_V}(S)^2 \leq 2 \cdot 2^{|V|} \cdot \theta^{\mathrm{dist}(V)}. \tag{1}$$

If we consider the consistent set V containing the output qubit at time T, then we get that $\widehat{\delta_V}(Z)^2 \leq 4\theta^T$. By Observation 5, this implies Theorem 1.

3.1 Proof of Lemma 1

The proof of the invariant is by induction on the sets V. At the base are all sets V contained entirely within time 0. All other sets are handled in the induction step. To justify the inductive proof, we need an ordering on the consistent sets V such that for each V, the proof for V uses the inductive hypothesis only on sets V' that appear before V. As will become apparent from the proof, if we denote by $\mathrm{latest}(V)$ the maximum time at which V contains a qubit, then each V' for which we use the induction hypothesis has strictly less qubits than V at time $\mathrm{latest}(V)$. Therefore, we can order the sets V first in increasing order of $\mathrm{latest}(V)$ and then in increasing order of the number of qubits at time $\mathrm{latest}(V)$.

Base case. Here V is fully contained within time 0. If $V = \emptyset$ then both sides of the invariant are zero, so from now on assume V is nonempty. In this case $\mathrm{dist}(V) = 0$. The matrix δ_V is the difference of two density matrices ρ_V and τ_V. Hence $\mathrm{Tr}(\delta_V^2) = \mathrm{Tr}(\rho_V^2) + \mathrm{Tr}(\tau_V^2) - 2\mathrm{Tr}(\rho_V \tau_V) \leq 2$, and the invariant is satisfied.

Induction step. Let V'' be any consistent set containing at least one qubit at time greater than zero. Our goal in this section is to prove the invariant for V''. Consider any of the qubits of V'' located at time $\mathrm{latest}(V'')$ and let G be the gate that has this qubit as one of its output qubits. We now consider two cases, depending on whether G is a k-qubit gate or a one-qubit gate.

<u>*Case 1*</u>*: G is a k-qubit gate.* Here G is a probabilistic mixture of k-qubit unitaries. First, by Observation 6 it suffices to prove the invariant for k-qubit unitaries. So assume G is a k-qubit unitary acting on the qubits $\mathcal{A} = \{A_1, \ldots, A_k\}$. Let $\mathcal{A}' = \{A'_1, \ldots, A'_k\}$ be the qubits after the ε_k-noise but before the gate G and $\mathcal{A}'' = \{A''_1, \ldots, A''_k\}$ the qubits after G (see Figure 2). By our choice of G, $\mathcal{A}'' \cap V'' \neq \emptyset$. Define $V' = (V'' \setminus \mathcal{A}'') \cup \mathcal{A}'$ and $V = (V'' \setminus \mathcal{A}'') \cup \mathcal{A}$. Note that V and its subsets are consistent sets with strictly fewer qubits than V'' at time $\mathrm{latest}(V'')$, hence we can apply the induction hypothesis to them. Our goal is to prove the invariant Eq. (1) for V''. First, by Observation 3,

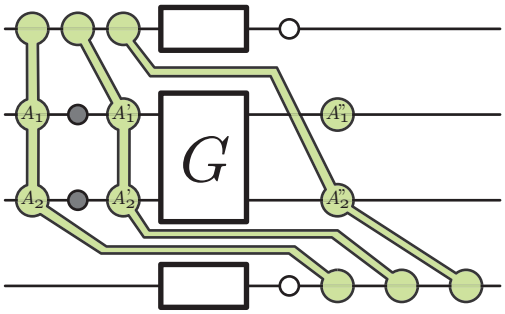

Fig. 2. An example showing the sets V, V', and V'' for a two-qubit gate G.

$$\sum_{S \in \mathcal{P}^{V''}} \widehat{\delta_{V''}}(S)^2 \leq \sum_{S \in \mathcal{P}^{V'' \cup \mathcal{A}''}} \widehat{\delta_{V'' \cup \mathcal{A}''}}(S)^2. \tag{2}$$

Because G (which maps $\delta_{V'}$ to $\delta_{V'' \cup \mathcal{A}''}$) is unitary, it preserves the sum of squares of $\widehat{\delta}$-coefficients (see Observation 2), so the right hand side of (2) is equal to

$$\sum_{S \in \mathcal{P}^{V'}} \widehat{\delta_{V'}}(S)^2 = \sum_{S \in \mathcal{P}^{V' \setminus \mathcal{A}'}} \sum_{R \in \mathcal{P}^{\mathcal{A}'}} \widehat{\delta_{V'}}(RS)^2.$$

Since the only difference between δ_V and $\delta_{V'}$ is noise on the qubits $A_1, \ldots, A_k$, using Observation 4 and denoting $\mu = 1 - \varepsilon_k$, we get that the above is at most

$$\begin{aligned}
&\sum_{S \in \mathcal{P}^{V \setminus \mathcal{A}}} \sum_{R \in \mathcal{P}^{\mathcal{A}}} \mu^{2|\mathrm{supp}(R)|} \widehat{\delta_V}(RS)^2 \\
&= \sum_{S \in \mathcal{P}^{V \setminus \mathcal{A}}} \sum_{a \subseteq \mathcal{A}} \mu^{2|a|}(1-\mu^2)^{k-|a|} \sum_{R \in \mathcal{P}^a \otimes I^{\mathcal{A} \setminus a}} \widehat{\delta_V}(RS)^2,
\end{aligned}$$

where the equality is because for any fixed S and any $R \in \mathcal{P}^{\mathcal{A}}$, the term $\widehat{\delta_V}(RS)^2$, which appears with coefficient $\mu^{2|\mathrm{supp}(R)|}$ on the left, appears with the same coefficient $\sum_{a \supseteq \mathrm{supp}(R)} \mu^{2|a|}(1-\mu^2)^{k-|a|} = \mu^{2|\mathrm{supp}(R)|}$ on the right. By rearranging and using Observation 3 we get that the above is equal to

$$\begin{aligned}
&\sum_{a \subseteq \mathcal{A}} \mu^{2|a|}(1-\mu^2)^{k-|a|} \sum_{S \in \mathcal{P}^{(V \setminus \mathcal{A}) \cup a}} \widehat{\delta_{(V \setminus \mathcal{A}) \cup a}}(S)^2 \\
&\leq \sum_{a \subseteq \mathcal{A}} \mu^{2|a|}(1-\mu^2)^{k-|a|} 2 \cdot 2^{|(V \setminus \mathcal{A}) \cup a|} \cdot \theta^{\mathrm{dist}((V \setminus \mathcal{A}) \cup a)}
\end{aligned}$$

where we used the inductive hypothesis. Note that $\mathrm{dist}((V \setminus \mathcal{A}) \cup a) \geq \mathrm{dist}(V)$, so the above is

$$\begin{aligned}
&\leq 2 \cdot 2^{|V \setminus \mathcal{A}|} \cdot \theta^{\mathrm{dist}(V)} \sum_{a \subseteq \mathcal{A}} 2^{|a|} \mu^{2|a|}(1-\mu^2)^{k-|a|} \\
&= 2 \cdot 2^{|V \setminus \mathcal{A}|} \cdot \theta^{\mathrm{dist}(V)}((1-\mu^2) + 2\mu^2)^k = \ 2 \cdot 2^{|V \setminus \mathcal{A}|} \cdot \theta^{\mathrm{dist}(V)}(1+\mu^2)^k. \qquad (3)
\end{aligned}$$

Note that $|V \setminus \mathcal{A}| \leq |V''| - 1$ and $\mathrm{dist}(V'') - 1 \leq \mathrm{dist}(V)$, so the right hand side is bounded by

$$\leq 2 \cdot 2^{|V''|-1} \cdot \theta^{\mathrm{dist}(V'')-1}(1+\mu^2)^k.$$

Since $\varepsilon_k > 1 - \sqrt{2^{1/k} - 1}$, we have that $(1+\mu^2)^k \leq 2\theta$ if θ is close enough to 1, so we can finally bound the last expression to prove the invariant for V''

$$\leq 2 \cdot 2^{|V''|} \cdot \theta^{\mathrm{dist}(V'')}.$$

Case 2: G is a one-qubit gate. Before proving the invariant, we need to prove the following property of completely-positive trace-preserving (CPTP) maps on one qubit. The proof appears in the full version of this paper [21].

Lemma 2. *For any CPTP map G on one qubit there exists a $\beta \in [0,1]$ such that the following holds. For any Hermitian matrix δ, if we let δ' denote the result of applying G to δ, then we have*

$$\widehat{\delta'}(X)^2 + \widehat{\delta'}(Y)^2 + \widehat{\delta'}(Z)^2 \leq (1-\beta) \cdot \widehat{\delta}(I)^2 + \beta \cdot (\widehat{\delta}(X)^2 + \widehat{\delta}(Y)^2 + \widehat{\delta}(Z)^2).$$

Let A be the qubit G is acting on, and recall that our goal is to prove the invariant for the set V''. Denote by A' the qubit of G after the gate but before the ε_1 noise, and by A'' the qubit after the noise. As before, by our choice of G, we

have $A'' \in V''$. Let $\mathcal{A} = \{A\}$, $\mathcal{A}' = \{A'\}$, $\mathcal{A}'' = \{A''\}$. Define $V' = (V'' \setminus \mathcal{A}'') \cup \mathcal{A}'$ and $V = (V'' \setminus \mathcal{A}'') \cup \mathcal{A}$ and notice that $|V| = |V'| = |V''|$. By using Lemma 2, we obtain a $\beta \in [0,1]$ such that

$$\sum_{S \in \mathcal{P}^{V''}} \widehat{\delta_{V''}}(S)^2 \leq \sum_{S \in \mathcal{P}^{V' \setminus \mathcal{A}'}} \left(\widehat{\delta_{V'}}(IS)^2 + (1-\varepsilon_1)^2 \sum_{R \in \mathcal{P}_*^{\mathcal{A}'}} \widehat{\delta_{V'}}(RS)^2 \right)$$
$$\leq \sum_{S \in \mathcal{P}^{V \setminus \mathcal{A}}} \left((1+(1-\varepsilon_1)^2(1-2\beta))\widehat{\delta_V}(IS)^2 + (1-\varepsilon_1)^2\beta \sum_{R \in \mathcal{P}^{\mathcal{A}}} \widehat{\delta_V}(RS)^2 \right).$$

By applying the induction hypothesis to both $V \setminus \mathcal{A}$ and V, we can upper bound the above by

$$(1 + (1-\varepsilon_1)^2(1-2\beta)) \cdot 2 \cdot 2^{|V|-1} \cdot \theta^{\mathrm{dist}(V \setminus \mathcal{A})} + (1-\varepsilon_1)^2 \beta \cdot 2 \cdot 2^{|V|} \cdot \theta^{\mathrm{dist}(V)}$$
$$\leq \frac{1 + (1-\varepsilon_1)^2}{2\theta} \cdot 2 \cdot 2^{|V''|} \cdot \theta^{\mathrm{dist}(V'')}$$

where we used that $|V| = |V''|$, and $\mathrm{dist}(V'') - 1 \leq \mathrm{dist}(V) \leq \mathrm{dist}(V \setminus \mathcal{A})$. Hence the invariant remains valid if we choose $\theta < 1$ such that $1 + (1-\varepsilon_1)^2 \leq 2\theta$.

Acknowledgment We thank Mary Beth Ruskai for a pointer to [30] and for sharing her insights on one-qubit operations; Peter Shor for a discussion on entanglement-breaking channels which is related to the discussion of [29] at the end of Section 1; and an anonymous ICALP referee for helpful comments.

All authors acknowledge support by the European Commission under the Integrated Project Qubit Applications (QAP) funded by the IST directorate as Contract Number 015848. JK is supported by an Alon Fellowship and by the Israeli Science Foundation, OR by the Binational Science Foundation and by the Israel Science Foundation, and RdW is partially supported by a Veni grant from the Netherlands Organization for Scientific Research (NWO).

References

1. Bernstein, E., Vazirani, U.: Quantum complexity theory. SIAM Journal on Computing **26**(5) (1997) 1411–1473 Earlier version in STOC'93.
2. Simon, D.: On the power of quantum computation. SIAM Journal on Computing **26**(5) (1997) 1474–1483 Earlier version in FOCS'94.
3. Shor, P.W.: Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. SIAM Journal on Computing **26**(5) (1997) 1484–1509 Earlier version in FOCS'94.
4. Grover, L.K.: A fast quantum mechanical algorithm for database search. In: Proceedings of 28th ACM STOC. (1996) 212–219
5. Shor, P.W.: Scheme for reducing decoherence in quantum memory. Physical Review A **52** (1995) 2493
6. Shor, P.W.: Fault-tolerant quantum computation. In: 37th FOCS (1996) 56–65
7. Steane, A.: Multiple particle interference and quantum error correction. In: Proceedings of the Royal Society of London. Volume A452. (1996) 2551–2577

8. Knill, M., Laflamme, R., Zurek, W.: Accuracy threshold for quantum computation. quant-ph/9610011 (15 Oct 1996)
9. Knill, E., Laflamme, R., Zurek, W.H.: Resilient quantum computation. Science **279**(5349) (1998) 342–345
10. Aharonov, D., Ben-Or, M.: Fault tolerant quantum computation with constant error. In: Proceedings of 29th ACM STOC. (1997) 176–188
11. Kitaev, A.Y.: Quantum computations: Algorithms and error correction. Russian Mathematical Surveys **52**(6) (1997) 1191–1249
12. Gottesman, D.: Stabilizer Codes and Quantum Error Correction. PhD thesis, Caltech (1997) quant-ph/9702052.
13. Knill, M.: Quantum computing with realistically noisy devices. Nature **434** (2005) 39–44
14. Knill, M.: Fault-tolerant postselected quantum computation: Threshold analysis. quant-ph/0404104 (19 Apr 2004)
15. Aliferis, P., Gottesman, D., Preskill, J.: Accuracy threshold for postselected quantum computation. Quantum Information and Computation **8**(3) (2008) 181–244
16. Aliferis, P.: Threshold lower bounds for Knill's Fibonacci scheme. quant-ph/0709.3603 (22 Sep 2007)
17. Aliferis, P.: Level Reduction and the Quantum Threshold Theorem. PhD thesis, Caltech (2007) quant-ph/0703264.
18. Reichardt, B.: Error-Detection-Based Quantum Fault Tolerance Against Discrete Pauli Noise. PhD thesis, UC Berkeley (2006) quant-ph/0612004.
19. Buhrman, H., Cleve, R., Laurent, M., Linden, N., Schrijver, A., Unger, F.: New limits on fault-tolerant quantum computation. In: 47th FOCS. (2006) 411–419
20. Razborov, A.: An upper bound on the threshold quantum decoherence rate. Quantum Information and Computation **4**(3) (2004) 222–228
21. Kempe, J., Regev, O., Unger, F., de Wolf, R.: Upper bounds on the noise threshold for fault-tolerant quantum computing (2008) quant-ph/0802.1464.
22. Barenco, A., Bennett, C., Cleve, R., DiVincenzo, D., Margolus, N., Shor, P., Sleator, T., Smolin, J., Weinfurter, H.: Elementary gates for quantum computation. Physical Review A **52** (1995) 3457–3467
23. Bravyi, S., Kitaev, A.: Universal quantum computation with ideal Clifford gates and noisy ancillas. Physical Review A **71**(022316) (2005)
24. Reichardt, B.: Quantum universality from Magic States Distillation applied to CSS codes. Quantum Information Processing **4** (2005) 251–264
25. Evans, W.S., Schulman, L.J.: Signal propagation and noisy circuits. IEEE Trans. Inform. Theory **45**(7) (1999) 2367–2373
26. Evans, W.S., Schulman, L.J.: On the maximum tolerable noise of k-input gates for reliable computation by formulas. IEEE Trans. Inform. Theory **49**(11) (2003) 3094–3098
27. Reichardt, B.: Quantum universality by distilling certain one- and two-qubit states with stabilizer operations. quant-ph/0608085 (2006)
28. Virmani, S., Huelga, S., Plenio, M.: Classical simulability, entanglement breaking, and quantum computation thresholds. Physical Review A **71**(042328) (2005)
29. Bruss, D., DiVincenzo, D., Ekert, A., Fuchs, C., Macchiavello, C., Smolin, J.: Optimal universal and state-dependent quantum cloning. Physical Review A **43** (1998) 2368–2378
30. Ruskai, M.B., Szarek, S., Werner, E.: An analysis of completely-positive trace-preserving maps on $\mathcal{M}_2$. Linear Algebra and its Applications **347** (2002) 159–187