FAULT-TOLERANT QUANTUM COMPUTATION WITH CONSTANT ERROR RATE *

DORIT AHARONOV† AND MICHAEL BEN-OR‡

Key Words: Quantum computation, Noise and Decoherence, Density matrices, Concatenated quantum error correcting codes, Polynomial codes, Universal quantum gates

Abstract. Shor has showed how to perform fault tolerant quantum computation when the probability for an error in a qubit or a gate, η, decays with the size of the computation polylogarithmically, an assumption which is physically unreasonable. This paper improves this result and shows that quantum computation can be made robust against errors and inaccuracies, when the error rate, η, is smaller than a constant threshold, η_c. The cost is polylogarithmic in space and time. The result holds for a very general noise model, which includes probabilistic errors, decoherence, amplitude damping, depolarization, and systematic inaccuracies in the gates. Moreover, we allow exponentially decaying correlations between the errors both in space and in time. Fault tolerant computation can be performed with any universal set of gates. The result also holds for quantum particles with $p > 2$ states, namely qupits, and is also generalized to one dimensional quantum computers with only nearest neighbor interactions. No measurements, or classical operations, are required during the quantum computation.

We use Calderbank-Shor Steane (CSS) quantum error correcting codes, generalized to qupits, and introduce a new class of CSS codes over F_p, called polynomial codes. It is shown how to apply a universal set of gates fault tolerantly on states encoded by general CSS codes, based on modifications of Shor's procedures, and on states encoded by polynomial codes, where the procedures for polynomial codes have a simple and systematic structure based on the algebraic properties of the code. Geometrical and group theoretical arguments are used to prove the universality of the two sets of gates. Our key theorem asserts that applying computation on encoded states recursively achieves fault tolerance against constant error rate. The generalization to general noise models is done using the framework of quantum circuits with density matrices. Finally, we calculate the threshold to be $\eta_c \simeq 10^{-6}$, in the best case.

The paper contains new and significantly simpler proofs for most of the known results which we use. For example, we give a simple proof that it suffices to correct bit and phase flips, we significantly simplify Calderbank and Shor's original proof of the correctness of CSS codes. We also give a simple proof of the fact that two-qubit gates are universal. The paper thus provides a self contained and complete proof for universal fault tolerant quantum computation.

1. Outline. Quantum computation has recently gained a lot of attention, due to oracle results[65, 10, 11] and quantum algorithms[25, 65, 36], in particular Shor's factorization algorithm [62]. These results indicate a possibility that quantum computers are exponentially more powerful than classical computers. It is yet unclear whether and how quantum computers will be physically realizable[19, 22, 30, 37, 50, 51, 55, 59, 69] but as any physical system, they *in principle* will be subjected to noise, such as decoherence [76, 71], and inaccuracies. Thus, the question of correcting noise cannot be separated from the complexity questions. Without error corrections, the effect of noise will accumulate and ruin the entire computation[72, 17, 54, 52, 7], and hence the computation must be protected. For classical circuits, von Neumann has shown already in 1956

*A preliminary version of this paper, under the name "Fault-Tolerant Quantum Computation With Constant Error", was published in *Proceedings of the 29th Annual ACM Symposium on Theory of Computing (STOC)* 1997

†School of Mathematics, Institute for Advanced Studies, Princeton, New Jersey

‡Department of Computer Science, The Hebrew University, Jerusalem, Israel

that the computation can be made robust to noise[53]. However, the similar question for quantum systems is much more complicated. Even the simpler question of protecting quantum information is harder than the classical analogue because one must also protect the quantum correlations between the quantum bits (qubits). It was argued by scientists that due to the continuity of quantum states, and to the fact that quantum states cannot be cloned[74], it will be impossible to protect quantum information from noise[72, 47]. Despite these pessimistic beliefs, Shor discovered a scheme to reduce the effect of decoherence[63]. Immediately after that, Calderbank and Shor[13] and Steane[68] showed that good quantum error correcting codes exist, a result which was followed by many explicit examples of quantum codes (e.g.[46, 68]). A theory for quantum error correcting codes was developed[43], and a group theoretical framework for almost all codes was found [14, 15, 31]. The existence of quantum error correcting codes, however, does not imply the existence of noise resistant quantum computation, since due to computation the faults propagate. One must be able to compute without allowing the errors to propagate too much, while error corrections should be made in the presence of noise. Recently Shor[64] showed how to use quantum codes in order to perform fault tolerant quantum computation in the presence of probabilistic errors, when the *error rate*, or the fault probability each time step, per qubit or gate, is polylogarithmically small. This assumption is physically unrealistic. In this paper we improve this result and show how to perform fault tolerant quantum computation in the presence of constant error rate, as long as the error rate η, is smaller than some constant threshold, η_0. The result holds for a very general noise model, which includes, besides probabilistic errors, also decoherence, amplitude and phase damping, depolarization, and systematic inaccuracies in the gates. Fault tolerance can be achieved using any universal set of gates. The result also holds when working with quantum particles of more than two states, instead of qubits. Our scheme can be generalized to work also in the case in which the quantum computer is a one dimensional array of qubits or qupits, with only nearest neighbor interactions. No measurements, or classical operations, are required during the quantum computation. The cost is polylogarithmic in the depth and size of the quantum circuit. The main assumption on the noise is that of locality, i.e. the noise process in different gates and qubits is independent in time and space. This assumption can be slightly relaxed, by allowing exponentially decaying correlations in both space and time. Such assumptions are made also in the classical scenario, and are likely to hold in physical realizations of quantum computers. Thus, this paper settles the question of quantum computation in the presence of local noise.

Let us first describe the computational model with which we work. The standard model[23, 24, 75] of quantum circuits with unitary gates, allows only unitary operations on qubits. However, noisy quantum systems are not isolated from other quantum systems, usually referred to as the environment. Their interactions with the environment are unitary, but when restricting the system to the quantum computer alone the operation on the system is no longer unitary, and the state of the quantum circuit is no longer a vector in the Hilbert space. It is possible to stay in the unitary model, keeping track of the state of the environment. However, this environment is not part of the computer, and it is assumed that we have no control or information on its state. We find it more elegant to follow the framework of the physicists, and to work within the model of quantum circuits with mixed states defined by Aharonov, Kitaev and Nisan[4]. In this model, the state of the set of qubits is always defined: It is a probability distribution over pure states, i.e. a *mixed state*, or a density matrix, and not merely a pure state as in the standard model. The quantum gates in this model are not necessarily unitary: any physically allowed operator on qubits is a quantum gate. In general, it is easiest to think of a general physical operator as a unitary operation on the system of qubits and any number of extra qubits, followed by an operator which discards the extra qubits. In particular, one can describe in this model an operation which adds a blank qubit to the system, or discards a qubit. The model of quantum circuits with mixed states is equivalent in computational power to the standard model of quantum circuits[4], but is more appropriate to work with when dealing with errors.

The noise process can be described naturally in this model as follows. Since noise is a dynamical process which depends on time, the circuit will be leveled, i.e. gates will be applied at discrete time steps. Between the time steps, we add the noise process. The simplest model for noise is the probabilistic process: Each time step, each qubit or gate undergoes a fault (i.e. an arbitrary quantum operation) with independent probability η, and η is referred to as the *error rate*. The probabilistic noise process can be generalized to a more realistic model of noise, which is the following: Each qubit or (qubits participating in the same gate), each time step, undergoes a physical operator which is at most η far from the identity, in some metric on operators. This model includes, apart from probabilistic errors, also decoherence, amplitude and phase damping and systematic inaccuracies in the gates. Two important assumptions were made in this definition: Independence between different faults in space, i.e. locality, and independence in time, which is called the Markovian assumption. It turns out that for our purposes, we can release these two restrictions slightly, to allow exponentially decaying correlations in both time and space, and all the results of this paper will still hold. However, we will turn to the generalized noise model in a very late stage of this paper. Meanwhile, it is simpler to keep in mind the independent probabilistic noise model.

Before we explain how to make quantum circuits fault tolerant, let us discuss how to protect quantum information against noise using quantum error correcting codes. As in classical linear block codes, a quantum error correcting code encodes the state of each qubit on a *block* of, say, m qubits. The encoded state is sometimes called the *logical* state. The code is said to correct d errors if the logical state is recoverable given that not more than d errors occurred in each block. The difference from classical codes is that quantum superpositions should be recoverable as well, so one should be able to protect quantum coherence, and not only basis states. Calderbank and Shor[13] and Steane [68] were the first to construct such quantum codes, which are now called CSS codes. The basic ingredients are two observations. First, the most general fault, or quantum operation, on a qubit can be described as a linear combination of four simple operations: The identity, i.e. no error at all, a bit flip ($|0\rangle \leftrightarrow |1\rangle$), a phase flip ($|0\rangle \longmapsto |0\rangle, |1\rangle \longmapsto -|1\rangle$) or both a bit flip and a phase flip. We give here a simple proof of the fact, first proven by Bennett *at. al.*[8], that it suffices to correct only these four basic errors. In order to correct bit flips, one can use analogues of classical error correcting codes. To correct phase flips, another observation comes in handy: A phase flip is actually a bit flip in the Fourier transformed basis, and hence, one can correct bit flips in the original basis, and then correct bit flips in the Fourier transformed basis, which translates to correcting phase flips in the original basis. In this paper we give an alternative proof for CSS codes, which generalized these ideas to quantum particles with $p > 2$ states, and is significantly simpler than the original proof of Calderbank and Shor[13].

We define a new class of quantum codes, which are called *polynomial codes.* The idea is based on a theorem by Schumacher and Nielsen[61] which asserts that a quantum state can be corrected only if no information about the state has leaked to the environment through the noise process. More precisely, if the reduced density matrix on any t qubits does not depend on the logical qubit, then there exists a unitary operation which recovers the original state even if the environment interacted with t qubits, i.e. t errors occurred. This is reminiscent of the situation in classical secret sharing schemes: We should divide the "secret", i.e. the logical qubit, among many parties (i.e. physical qupits) such that no t parties share any information about the secret. This analogy between secret sharing and quantum codes suggests that secret sharing schemes might prove useful in quantum coding theory. Here, we adopt the scheme suggested by Ben-Or, Goldwasser and Wigderson[9], who suggested to use random polynomials, evaluated at different points in a field of p elements, as a way to divide a secret among a few parties. A random polynomial of degree d is chosen, and then each party gets the evaluation of the polynomial at a different point in the field F_p. The secret is the value of the polynomial at 0. To adopt this scheme to the quantum setting, we simply replace the random polynomial by a superposition of all polynomials, to get a quantum code. It turns out that we get a special case of the CSS codes. Polynomial codes are useful from many aspects. First, they have a very nice algebraic structure, which allows to manipulate them easily. This allows us to apply fault tolerant operations on states encoded by polynomial codes in a simple and systematic way, as we will see later. Polynomial codes might also have nice applications to other areas in quantum information theory, as was demonstrated recently by Gottesman *et al*[35] who used polynomial codes for quantum secret sharing. We will use both polynomial codes and general CSS codes in our fault tolerant scheme.

In order to protect a quantum computation against faults, one can try to compute not on the states themselves, but on states encoded by quantum error correcting codes. Each gate in the original computation will be replaced by a "procedure", which applies the encoded gate on the encoded state. Naturally, in order to prevent accumulation of errors, error corrections should be applied frequently. Unfortunately, computing on encoded quantum states does not automatically provide protection against faults, even if error corrections are applied after every procedure. The problem lies in the fact that during the computation of the procedure, faults can *propagate* to "correct" qubits. This can happen if a gate operates on a damaged qubit and some "correct" qubits- in general, this can cause all the qubits that participate in the gate to be damaged. The procedures should thus be designed carefully, in such a way that a fault during the operation of the procedure can only effect a small number of qubits in each block. We refer to such procedures as fault tolerant procedures. It turns out that many gates can be applied *bitwise*, meaning that applying the gate on each qubit separately, achieves the desired gate on the encoded state. Unfortunately, not all gates can be applied fault tolerantly in such a simple way, and sometimes a very complicated procedure is needed. We are actually looking for a pair consisting of a quantum code which can correct d errors, and a corresponding *universal* set of gates, such that their procedures, with respect to the code, allow one fault to spread to at most d qubits. Since the error corrections, encoding, and decoding procedures are also subjected to faults, we need them to be fault tolerant too.

In [64], Shor introduced a universal set of gates, which we denote by $\mathcal{G}_1$, and showed how to apply the gates fault tolerantly on encoded states. It turns out that almost all of the gates in $\mathcal{G}_1$ can be applied bitwise. The complicated procedures are the fault tolerant error correction and encoding, and the Toffoli gate. Shor used noiseless measurements and classical computation extensively in these procedures. Here, we modify Shor's procedures so that they can be performed within the framework of our model, i.e. all operations are quantum operations, subjected to noise. Basically, we follow Shor's constructions, with some additional tricks. We remark here that in this paper we make an effort to show that fault tolerance can be achieved without using classical operations or measurements. This fact is desirable because of two reasons. One is purely theoretical: One would like to know that measurements and classical operations are not essential, and that the quantum model is complete in the sense that it can be made fault tolerant within itself. The other reason is practical: In some suggestions for physical realizations of quantum computers, such as the NMR computer, it might be very hard to incorporate measurements and classical computations during the quantum computation. Thus, we require that all procedures are done using only quantum computation.

For polynomial codes, we introduce a new universal set of gates which we denote by $\mathcal{G}_2$, and show how to perform these gates fault tolerantly on states encoded with polynomial codes. Here is where the advantage of polynomial codes come into play. Unlike the construction of the fault tolerant Toffoli gate for CSS codes, (and other fault tolerant constructions) which is a long and complicated sequence of tricks, the fault tolerant procedures for polynomial codes all have a common relatively simple structure, which uses the algebraic properties of the polynomials. To perform a gates in $\mathcal{G}_2$, we apply two steps. First the gate is applied pit-wise. This always achieves the correct operation of the gate, but sometimes the final state is encoded by polynomial codes with degree which is twice the original degree d. The second step is therefore a degree reduction, which can be applied using interpolation techniques, as the quantum analogue of the degree reductions used in [9]. Thus, the fault tolerant procedures for polynomial codes have a simple systematic structure.

Next, we prove the universality of the two sets of gates which we use, $\mathcal{G}_1$ for the CSS codes and $\mathcal{G}_2$ for the polynomial codes. Universality of a set of gates means that the subgroup generated by the set of gates is dense in the group of unitary operations on n qubits, $U(2^n)$, (for qupits the group is $U(p^n)$). Kitaev[39] and Solovay[66] showed, by a beautiful proof which uses Lie algebra and Lie groups, that universality of a set $\mathcal{G}$ guarantees that any quantum computation can be approximated by a circuit which uses $\mathcal{G}$ with only polylogarithmic cost. Our proof of universality of $\mathcal{G}_1$ is a simple reduction to a set of gates, which was shown to be universal by Kitaev[39]. The proof that $\mathcal{G}_2$ is universal is more involved. The idea is first to generate, using gates from $\mathcal{G}_2$, matrices with eigenvalues which are not integer roots of unity, a fact which is proved using field theoretical arguments. This enables us to show that certain $U(2)$ subgroups are contained in the subgroup generated by $\mathcal{G}_2$, where these $U(2)$ operate on non orthogonal subspaces, which together span the entire space. To complete the proof, we prove a geometrical lemma which asserts that if A and B are non orthogonal then $U(A) \cup U(B)$ generate the whole group $U(A \oplus B)$. Thus we have that both sets of gates which are used in our scheme are universal.

We proceed to explain how a quantum code accompanied with fault tolerant procedures for a universal set of gates can be used to enhance the reliability of a quantum circuit. We start with our original quantum circuit, M_0. The circuit M_1 which will compute on encoded states is defined as a simulation of M_0. A qubit in M_0 transforms to a block of m qubits in M_1, and each gate in M_0 transforms to a fault tolerant procedure in M_1 applied on the corresponding blocks. Note that procedures might take more than one time step to perform. Thus, a time step in M_0 is mapped to a time interval in M_1, which is called a *working period*. In order to prevent accumulation of errors, at the end of each working period an error correction is applied on each block. The working period now consists of two stages: a computation stage and a correction stage. The idea is therefore to apply alternately computation stages and correction stages, hoping that during the computation stage the damage that accumulated is still small enough so that the corrections are still able to correct it.

We now want to show that the reliability of the simulating circuit is larger than that of the original circuit. In the original circuit, the occurrence of one fault may cause the failure of the computation. In contrast, the simulating circuit can tolerate a number of errors, say k, in each procedure, since they are immediately corrected in the following error correcting stage. The effective noise rate of M_1 is thus the probability for more than k faults in a procedure. This effective error rate will be smaller than the actual error rate η, if η is smaller than a certain threshold, η_c, which depends on the size of the procedures, and on the number of errors which the error correcting code can correct. If $\eta < \eta_c$, M_1 will be more reliable than M_0. For fault tolerance, we need the effective error rate to be polynomially small, to ensure that with constant probability no effective error occurred. In [64] Shor uses quantum CSS codes which encode one qubit on polylog(n) to achieve fault tolerance with the error rate η being polylogarithmically small.

The main goal of this paper is to improve this result and to show that quantum computation can be robust even in the presence of constant error rate. Such an improvement from a constant error rate η to polynomially small effective noise rate is impossible to achieve in the above method. Instead, reliability in the presence of constant error rate is achieved by applying the above scheme recursively. The idea is simple: as long as the error rate is below the threshold, the effective noise rate of the simulating circuit can be decreased by simulating it again, and so on for several levels, say r, to get the final circuit M_r. It will suffice that each level of simulation improves only slightly the effective noise rate, since the improvement is exponential in the number of levels. Such a small improvement can be achieved when using a code of constant block size.

The final circuit, M_r, has an hierarchical structure. Each qubit in the original circuit transforms to a block of qubits in the next level, and they in their turn transform to a block of blocks in the second simulation and so on. A gate in the original circuit transforms to a procedure in the next level, which transforms to a larger procedure containing smaller procedures in the next level and so on. The final circuit computes in all the levels: The largest procedures, computing on the largest (highest level) blocks, correspond to operations on qubits in the original circuit. The smaller procedures, operating on smaller blocks, correspond to computation in lower levels. Note, that each level simulates the error corrections in the previous level, and adds error corrections in the current level. The final circuit, thus, includes error corrections of all the levels, where during the computation of error corrections of larger blocks smaller blocks of lower levels are being corrected. The lower the level, the more often error corrections of this level are applied, which is in correspondence with the fact that smaller blocks are more likely to be quickly damaged.

The analysis of the scheme turns out to be quite complicated. To do this, we generalizes the works of Tsirelson [38] and Gác's [28] to the quantum case. First, we distinguish between two notions: The error in the state of the qubits, i.e. the set of qubits which have errors, and the actual faults which took place during the computation, namely the fault path, which is a list of points in time and place where faults had occurred (in a specific run of the computation.) We then define the notion of sparse errors and sparse fault paths. Naturally, due to the hierarchical structure of this scheme, these notions are defined recursively. A block in the lowest level is said to be close to it's correct state if it does not have too many errors, and a higher level block is "close" to it's correct state if it does not contain too many blocks of the previous level which are far from their correct state. If all the blocks of the highest level are close to being correct, we say that the set of errors is *sparse*. Which set of faults does not cause the state to be too far from correct in the above metric? The answer is recursive too: A computation of the lowest level procedure is said to be undamaged if not too many faults occurred in it. Computation of higher level procedures are not damaged if they do not contain too many lower level procedures which are damaged. A fault path will be called *sparse*, if the computations of all the highest level procedures are not damaged. The proof of the threshold result is done by showing that the probability for "bad" faults, i.e. the probability for the set of faults not to be sparse decays exponentially with the number of levels, r. Thus, bad faults are negligible, if the error rate is below the threshold. This is the easy part. The more complicated task is to show that the "good" faults are indeed good: i.e. if the set of faults is sparse enough, then the set of errors is also kept sparse until the end of the computation. This part is done using induction on the number of levels r, and is quite involved.

So far, we have dealt only with probabilistic faults. Actually, the circuit generated by the above recursive scheme is robust also against general local noise. We prove this by writing the noise operator on each qubit as the identity plus a small error term. Expanding the error terms in powers of η, we get a sum of terms, each corresponding to a different fault path. The threshold result for general noise is again proved by dividing the faults to good and bad parts. First, we show that the bad part is negligible i.e. the norm of the sum of all terms which correspond to non-sparse fault paths is small. The proof that the good faults are indeed good, i.e. that the error in the terms corresponding to sparse fault paths is sparse, is based on the proof for probabilistic errors, together with some linearity considerations.

So far, what we have shown is that fault tolerant computation can be achieved using one of the two sets of gates, $\mathcal{G}_1, \mathcal{G}_2$. In fact, fault tolerant computation can be performed using any universal set of gates. This last generalization is shown by first designing a circuit which uses gates from $\mathcal{G}_1$ or $\mathcal{G}_2$, and then approximate each gate in the final circuit, up to constant accuracy, by a constant number of gates from the desired set of gates which we want to work with. This completes the proof of the threshold result in full generality.

The threshold result is shown to hold also for one dimensional quantum circuits. These are systems of n qubits arranged in a line, where gates are applied only on nearest neighbor qubits, and all qubits are present from the beginning. In this model, one cannot insert a new qubit between two existing qubits, in order to preserve the geometry of the arrangements. However, the non unitary operation which initializes a qubit in the middle of the computation is allowed. To apply fault tolerant procedures in this case, we use swap gates in order to bring far away qubits together. We show that these swap gates do not cause the error to propagate too much.

Finally, it is left to estimate the exact threshold in all the above cases. We give a general simple formula to calculate the threshold in the error rate, given the fault tolerant procedures. This formula depends only on two parameters of these procedures: the size of the largest procedure used in our scheme, and the number of errors which the code corrects. The size of the procedures depend strongly on many variables, such as the geometry of the system, the exact universal set of gates which is used, and whether noiseless classical computation can be used during the computation. We give here an estimation of the threshold in one case, in which no geometry is imposed on the system, error free classical operations and measurements can be performed during the computation, and the codes which are used are polynomial codes. The threshold which is achieved in this case is $\simeq 10^{-6}$. We did not attempt to optimize the threshold, and probably many improvements can be made which can save a lot of gates in our suggested procedures.

Some of the assumptions we made in this work cannot be released. First, the choice to work with circuits which allow parallelism is necessary, because sequential quantum computation, such as quantum Turing machines, cannot be noise resistant[3], since error corrections cannot be applied fast enough to prevent accumulation of errors. It is also crucial that we allow qubits to be input at different times to the circuit unlike (Yao[75]), because this provides a way to discard entropy which is accumulating in the circuit. Without this assumption, i.e. under the restriction that all qubits are initialized at time 0, it is impossible to compute fault tolerantly without an exponential blowup in the size of the circuit[5]. An interesting question is whether other assumptions, such as the locality assumption, can be relaxed. A partial positive answer in this direction was given by Chuang *et. al.*[48]. In connection with the assumptions made for Fault tolerance, see also [1, 70].

A substantial part of the results in this paper were published in a proceeding version in [2]. The threshold result was independently described also by Knill, Laflamme and Zurek[42], and by Kitaev[40, 39]. This was done using different quantum correcting codes, and required measurements in the middle of the quantum computation. In this work we provide a complete proof to all the details of the fault tolerant quantum computation codes construction and in particular we rigorously show that the recursive scheme works. Kitaev has shown another, very elegant way to achieve fault tolerance in a different model of quantum computation, which works with anyons[41]. After the results were first published, it was shown by Gottesman[32] that fault tolerant procedures for the universal set of gates introduced by Shor can be constructed for any stabilizer code, using basically the same technique as presented by Shor for the CSS codes, a result which was generalized to qupits[33]. Fault tolerant quantum computation in low dimensions is also discussed by Gottesman[34]. Non binary codes were defined independently also by Chuang[18] and Knill[44]. For a survey of this area see Preskill[58].

Organization of paper: In section 2 we define the model of noisy quantum circuits. Section 3 is devoted to CSS and polynomial codes. In section 4 we describe how to apply fault tolerant operations on states encoded by CSS codes. In section 5 we describe fault tolerant procedures for polynomial codes. Section 6 proves the universality of the two sets of gates used in the previous two sections. In section 7 we present recursive simulations and the threshold result for probabilistic errors. In section 8 we generalize the result to general local noise, and in section 9 we generalize to circuits which use any universal set of gates. In section 10 we explain how to deal with exponentially decaying correlations in the noise. Section 11 proves the threshold result for one dimensional quantum circuit, and in 12 we conclude with possible implications to physics, and some open questions.

2. The Model of Noisy Quantum Circuits. In this section we recall the definitions of quantum circuits with mixed states as defined by Aharonov, Kitaev and Nisan[4], and define noisy quantum circuits. The model is based on basic quantum physics, and good references are Cohen-Tanudji[21], Sakurai [60], and Peres[56].

2.1. Pure States. A quantum physical system in a *pure state* is described by a unit vector in a Hilbert space, i.e a vector space with an inner product. In the *Dirac* notation a pure state is denoted by $|\alpha\rangle$. The physical system which corresponds to a quantum circuit consists of n quantum two-state particles, and the Hilbert space of such a system is $\mathcal{H}_2^n = \mathcal{C}^{\{0,1\}^n}$ i.e. a 2^n dimensional complex vector space. $\mathcal{H}_2^n$ is viewed as a tensor product of n Hilbert spaces of one two-state particle: $\mathcal{H}_2^n = \mathcal{H}_2 \otimes \mathcal{H}_2 \otimes \ldots \otimes \mathcal{H}_2$. The k'th copy of $\mathcal{H}_2$ will be referred to as the k'th **qubit**. We choose a special basis for $\mathcal{H}_2^n$, which is called the computational basis. It consists of the 2^n orthogonal states: $|i\rangle, 0 \leq i < 2^n$, where i is in binary representation. $|i\rangle$ can be seen as a tensor product of states in $\mathcal{H}_2$: $|i\rangle = |i_1\rangle|i_2\rangle....|i_n\rangle = |i_1 i_2..i_n\rangle$, where each i_j gets 0 or 1. Such a state, $|i\rangle$, corresponds to the j'th particle being in the state $|i_j\rangle$. A pure state $|\alpha\rangle \in \mathcal{H}_2^n$ is generally a *superposition* of the basis states: $|\alpha\rangle = \sum_{i=1}^{2^n} c_i|i\rangle$, with $\sum_{i=1}^{2^n} |c_i|^2 = 1$. A vector in $\mathcal{H}_2^n$, $v_\alpha = (c_1, c_2, ..., c_{2^n})$, written in the computational basis representation, with $\sum_{i=1}^{2^n} |c_i|^2 = 1$, corresponds to the pure state: $|\alpha\rangle = \sum_{i=1}^{2^n} c_i|i\rangle$. $v_\alpha^\dagger$, the transposed-complex conjugate of v_α, is denoted $\langle\alpha|$. The inner product between $|\alpha\rangle$ and $|\beta\rangle$ is denoted $\langle\alpha|\beta\rangle = (v_\alpha^\dagger, v_\beta)$. The matrix $v_\alpha^\dagger v_\beta$ is denoted as $|\alpha\rangle\langle\beta|$.

2.2. Mixed States. In general, a quantum system is not in a pure state. This may be attributed to the fact that we have only partial knowledge about the system, or that the system is not isolated from the rest of the universe. We say that the system is in a *mixed state*, and assign with the system a probability distribution, or *mixture* of pure states, denoted by $\{\alpha\} = \{p_k, |\alpha_k\rangle\}$. This means that the system is with probability p_k in the pure state $|\alpha_k\rangle$. This description is not unique, as different mixtures might represent the same physical system. As an alternative description, physicists use the notion of **density matrices**, which is an equivalent description but has many advantages. A density matrix ρ on $\mathcal{H}_2^n$ is an hermitian (i.e. $\rho = \rho^\dagger$) semi positive definite matrix of dimension $2^n \times 2^n$ with trace $\mathrm{Tr}(\rho) = 1$. A pure state $|\alpha\rangle = \sum_i c_i|i\rangle$ is represented by the density matrix: $\rho_{|\alpha\rangle} = |\alpha\rangle\langle\alpha|$, i.e. $\rho_{|\alpha\rangle}(i,j) = c_i c_j^*$. (By definition, $\rho(i,j) = \langle i|\rho|j\rangle$). A mixture $\{\alpha\} = \{p_l, |\alpha_l\rangle\}$, is associated with the density matrix $\rho_{\{\alpha\}} = \sum_l p_l|\alpha_l\rangle\langle\alpha_l|$. This association is not one-to-one, but it is **onto** the density matrices, because any density matrix describes the mixture of it's eigenvectors, with the probabilities being the corresponding eigenvalues. Note that diagonal density matrices correspond to probability distributions over classical states. Density matrices are linear operators on their Hilbert spaces.

2.3. Quantum Operators and Quantum Gates. Transformations of density matrices are linear operators on operators (sometimes called *super-operators*). $\mathbf{L}(\mathcal{N})$ denotes the set of all linear operators on $\mathcal{N}$, a finite dimensional Hilbert space. The most general quantum operator is any operator $T : \mathbf{L}(\mathcal{N}) \rightarrow \mathbf{L}(\mathcal{M})$ which sends density matrices to density matrices. A super-operator is called *positive* if it sends positive semi-definite Hermitian matrices to positive semi-definite Hermitian matrices, so T must be positive. Super-operators can be extended to operate on larger spaces by taking tensor product with the identity operator: $T : \mathbf{L}(\mathcal{N}) \rightarrow \mathbf{L}(\mathcal{M})$ will be extended to $T \otimes I : \mathbf{L}(\mathcal{N} \otimes \mathcal{R}) \rightarrow \mathbf{L}(\mathcal{M} \otimes \mathcal{R})$, and it is also required that the extended operator remains positive, and such T is called a *completely positive map*. Clearly, T should also be trace-preserving. It turns out that any super operator which satisfies these conditions indeed take density matrices to density matrices[45]. Thus quantum super-operators are defined as follows:

DEFINITION 1. *A* **physical super-operator***, T, from k to l qubits is a trace preserving, completely positive, linear map from density matrices on k qubits to density matrices on l qubits.*

Linear operations on mixed states preserve the probabilistic interpretation of the mixture, because $T \circ \rho = T \circ (\sum_l p_l |\alpha_l\rangle\langle\alpha_l|) = \sum_l p_l T \circ (|\alpha_l\rangle\langle\alpha_l|)$.

A very important example of a quantum super-operator is the super-operator corresponding to the standard unitary transformation on a Hilbert space $\mathcal{N}$, $|\alpha\rangle \mapsto U|\alpha\rangle$, which sends a quantum state $\rho = |\alpha\rangle\langle\alpha|$ to the state $U\rho U^\dagger$. Another important super-operator is discarding a set of qubits. A density matrix of n qubits can be reduced to a subset, A, of $m < n$ qubits. We say that the rest of the system, represented by the Hilbert space $\mathcal{F} = \mathbf{C}^{2^{n-m}}$, is *traced out*, and denote the new matrix by $\rho|_A = \mathrm{Tr}_{\mathcal{F}}\,\rho$, where $\mathrm{Tr}_{\mathcal{F}} : \mathbf{L}(\mathcal{N} \otimes \mathcal{F}) \to \mathbf{L}(\mathcal{N})$ is defined by: $\rho|_A(i,j) = \sum_{k=1}^{2^{n-m}} \rho(ik, jk)$, and k runs over a basis for $\mathcal{F}$. In words, it means averaging over $\mathcal{F}$. If the state of the qubits which are traced out is in tensor product with the state of the other qubits, then discarding the qubits means simply erasing their state. However, if the state of the discarded qubits is not in tensor product with the rest, the reduced density matrix is always a mixed state. In this paper, a discarding qubit gate will always be applied on qubits which are in tensor product with the rest of the system. Any quantum operation which does not operate on $\mathcal{F}$ commutes with this tracing out. One more useful quantum super-operator is the one which adds a blank qubit to the system, $V_0 : |\xi\rangle \mapsto |\xi\rangle \otimes |0\rangle : \mathbf{C}^{2^n} \to \mathbf{C}^{2^{n+1}}$. This is described by the super-operator $T : \rho \mapsto \rho \otimes |0\rangle\langle 0|$.

A lemma by Choi[16], and Hellwig and Kraus[45] asserts that any physically allowed super-operator from k to l qubits can be described as a combination of the above three operations: add $k + l$ blank qubits, apply a unitary transformation on the $2k + l$ qubits and then trace out $2k$ qubits. One can think of this as the definition of a quantum super-operator. We define a quantum gate to be the most general quantum operator.

DEFINITION 2. *A quantum gate of order (k, l) is a super-operator from k to l qubits.*

In this paper we will only use three types of quantum gates: discard or add qubits, and of course unitary gates. However, quantum noise will be allowed to be a general physical operator.

2.4. Measurements. A quantum system can be **measured**, or observed. Let us consider a set of positive semi-definite Hermitian operators $\{P_m\}$, such that $\sum_m P_m = I$. The measurement is a process which yields a probabilistic classical output. For a given density matrix ρ, the output is m with probability $\Pr(m) = \mathrm{Tr}(P_m \rho)$. Any measurement thus defines a probability distribution over the possible outputs. In this paper we will only need a basic measurement of r qubits. In this case, P_m (with $1 \leq i \leq 2^r$) are the projection on the subspace S_m, which is the subspace spanned by basic vectors on which the measured qubits have the values corresponding to the string m: $S_m = \mathrm{span}\{\,|m, j\rangle,\ j = 1, \ldots, 2^{n-r}\}$. This process corresponds to measuring the value of r qubits, in the basic basis, where here, for simplicity, we considered measuring the first r qubits.

To describe a measurement as a super-operator, we replace the classical result of a measurement, m, by the density matrix $|m\rangle\langle m|$ in an appropriate Hilbert space $\mathcal{M}$. The state of the quantum system after a projection measurement is also defined; it is equal to $\Pr(m)^{-1}P_m\rho P_m$. (It has the same meaning as a conditional probability distribution). Thus, the projection measurement can be described by a super-operator T which maps quantum states on the space $\mathcal{N}$ to quantum states on the space $\mathcal{N}\otimes\mathcal{M}$, the result being diagonal with respect to the second variable:

$$T\rho \ = \ \sum_m (P_m\rho P_m)\otimes\Big(|m\rangle\langle m|\Big) \tag{2.1}$$

2.5. The Trace Metric on Density Operators. A useful metric on density matrices, the *trace* metric, was defined in [4]. It is induced from the *trace norm* for general Hermitian operators. The trace norm on Hermitian matrices is the sum of the absolute values of the eigenvalues. $\|\rho\| = \sum_i |\lambda_i|$ It was shown in [4] that the distance between two matrices in this metric $\|\rho_1-\rho_2\|$ is exactly the measurable distance between them: It equals to the maximal total variation distance between all possible outcome distributions of the same measurement done on the two matrices.

2.6. Quantum Circuits with Mixed States. We now define a quantum circuit:

DEFINITION 3. *Let $\mathcal{G}$ be a family of quantum gates.* **A Quantum circuit that uses gates from** $\mathcal{G}$ *is a directed acyclic graph. Each node v in the graph is labeled by a gate $\mathfrak{g}_v\in\mathcal{G}$ of order (k_v,l_v). The in-degree and out-degree of v are equal k_v and l_v, respectively.*

Here is a schematic example of such a circuit.

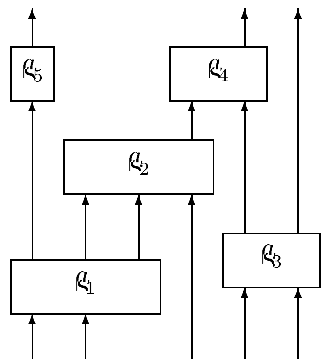

The circuit operates on a density matrix as follows:

DEFINITION 4. **final density matrix***: Let Q be a quantum circuit. Choose a topological sort for Q: $\mathfrak{g}_t...\mathfrak{g}_1$. Where $\mathfrak{g}_j$ are the gates used in the circuit. The final density matrix for an initial density matrix ρ is $Q\circ\rho = \mathfrak{g}_t\circ...\circ\mathfrak{g}_2\circ\mathfrak{g}_1\circ\rho$.*

$Q\circ\rho$ is well defined and does not depend on the topological sort of the circuit [4]. At the end of the computation, the r output qubits are measured, and the classical outcome, which is a string of r bits, is the output of the circuit. For an input string i, the probability for an output string j is $\langle j|\Big(Q\circ|i\rangle\langle i|\Big)|j\rangle$.

DEFINITION 5. **Computed function:** *Let Q be a quantum circuit, with n inputs and r outputs. Q is said to compute a function f: $\{0,1\}^n\longmapsto\{0,1\}^r$, with accuracy ϵ if for any input $i\in\{0,1\}^n$, the probability for the output to be $f(i)$ is greater than $1-\epsilon$.*

2.7. Quantum Circuits with probabilistic Errors. . To implement noise in the quantum circuits, we divide the quantum circuit to time steps, and the gates are applied at integer time steps. We permit that qubits are input and output at different time steps. Faults occur in between time steps, in the *locations* of the circuit.

DEFINITION 6. *A set $(q_1, q_2, ..q_l, t)$ is a location in the quantum circuit Q if the qubits $q_1, ...q_l$ participated in the same gate in Q, in time step t. and no other qubit participated in that gate. For this matter, the absence of a gate is viewed as the identity gate on one qubit, so if a qubit q did not participate in any gate at time t, then (q, t) is a location in Q as well.*

Each location in the circuit exhibits a fault with with independent probability η.

DEFINITION 7. *The list of times and places where faults had occurred, (in a specific run of the computation) is called a fault path.*

Each fault path, $\mathcal{F}$, is assigned a probability $Pr(\mathcal{F})$, which is a function of the number of faults in it. A fault at time t at a certain location means that an (arbitrary) general quantum operation is applied on the faulty qubits after time t. Given a certain fault path, we still have to determine the quantum operations that were applied on the faulty qubits. Denote by $\mathcal{E}(\mathcal{F})$ the choice of faults for the fault path, say $\mathcal{E}_t, ...\mathcal{E}_2, \mathcal{E}_1$. The final density matrix for this given choice $\mathcal{E}(\mathcal{F})$ is easily defined by applying the faults in between the gates: $Q(\mathcal{F}) \circ \rho = \mathcal{E}_t \circ Q_t \circ ...\mathcal{E}_2 \circ Q_2 \circ \mathcal{E}_1 \circ Q_1 \circ \rho$. The final density matrix, $Q \circ \rho$ of the circuit is defined as a weighted average over the final density matrices for each fault path:

$$Q \circ \rho = \sum_{\mathcal{F}} Pr(\mathcal{F}) Q(\mathcal{F}) \circ \rho. \tag{2.2}$$

This final density matrix depends on the exact choice of quantum operations applied on the faulty gates. However, a noisy quantum computer is said to compute a function f with confidence ϵ in the presence of error rate η, if the circuit outputs the function f with accuracy ϵ regardless of the exact specification of the quantum operations which occurred during the faults.

2.8. Quantum Circuits with General Local Noise. In most of the paper we will use only the probabilistic noise model described above. However, our result will be generalized in section 8 to a much more general noise model. In the general noise model, we replace the probabilistic faults which occur with probability η in each location by applying in each location in the quantum circuit an arbitrary general quantum operator, on which the only restriction is that it within η distance to the identity, in a certain norm of super-operators. Thus, in between time steps, a *noise operator* operates on all the qubits. For the time step t the noise operator is of the form

$$\mathcal{E}(t) = \mathcal{E}_{A_{1,t}}(t) \otimes \mathcal{E}_{A_{2,t}}(t) \otimes \cdots \otimes \mathcal{E}_{A_{l_t,t}}(t). \tag{2.3}$$

$A_{i,t}$ runs over all possible locations at time t, and for each one of them,

$$\|\mathcal{E}_{A_{i,t}}(t) - I\| \leq \eta \tag{2.4}$$

The exact definition of the norm on super-operators is not used in this paper. We merely need the following properties: