

Resilient Quantum Computation

Emanuel Knill, Raymond Laflamme, Wojciech H. Zurek

Practical realization of quantum computers will require overcoming decoherence and operational errors, which lead to problems that are more severe than in classical computation. It is shown that arbitrarily accurate quantum computation is possible provided that the error per operation is below a threshold value.

The discovery that quantum computers can be much more powerful than their classical counterparts (1–4) and recent advances in quantum device technology (5, 6) have brought the field of quantum computation into the limelight. However, until recently, the hope of taming quantum systems has been overshadowed by the fragility of quantum information. This fragility comes from two seemingly contradictory requirements. The system must be well insulated from the environment, whereas its components need to interact strongly to perform the computation. As a result of the interaction, the computer necessarily becomes increasingly entangled with the outside world, and quantum information is gradually lost through decoherence (7). Moreover, quantum computation requires the application of precise unitary operations to the system, and these operations cannot be implemented exactly. For these reasons, some have concluded that the potential power of quantum computers cannot be harnessed in practice (8, 9).

The first indication that such early assessments of the practicality of quantum computation might be overly pessimistic was the discovery of quantum error-correcting codes by Shor (10) and Steane (11). These codes imply that it is possible to overcome memory errors provided that the operations required for encoding, decoding, and error correction are implemented without errors (12–16). This requirement of error-free error-correction operations was relaxed by Shor (17) with the use of fault tolerance in the case of stochastic errors (errors occurring independently and incoherently) and was essentially eliminated in the case of quantum channels in (18).

Two obstacles to quantum computation remained. First, as the number of elementary quantum operations grow, Shor's fault-tolerant implementation still requires asymptotically zero errors per operation. Second, many error types expected

in real devices cannot be represented with stochastic errors. In particular, unitary overrotation of operations (19) and small but nonnegligible interactions between nearby bits give rise to such errors. Our purpose is to show how these obstacles can be surmounted.

The first obstacle is overcome by using concatenated codes that involve re-encoding already encoded bits. This process reduces the effective error rate at each level, with the final accuracy being dependent on how many levels of the hierarchy are used. To overcome the second obstacle, we show by explicit construction that an error threshold exists such that if each gate in a physical implementation of a quantum network has error less than this threshold, it is possible to perform any quantum computation with arbitrary accuracy. Therefore, noise, if it is below a certain level, is not an obstacle to unlimited resilient quantum computation.

We begin by emphasizing the various assumptions about errors and then describe four fundamental elements of fault tolerance: quantum error-correcting codes, fault-tolerant error-correction methods, encoded operations, and concatenation. We combine these elements to implement any computation resiliently. Finally, the resilient networks are analyzed to obtain rigorous thresholds for the quasi-independent error model.

Quantum networks and operations. Any quantum algorithm can be described by means of a quantum network, that is, a space-time diagram of the operations that are to be applied to each bit. A quantum bit is the prototypical two-state system spanned by $|0\rangle$ and $|1\rangle$ (the classical states). We use the set of operations consisting of the Pauli matrices σ_x (bit flip) and σ_z (sign flip), the 90° rotations around the x and y axes $[(I - i\sigma_x)/\sqrt{2}]$ and $[(I - i\sigma_y)/\sqrt{2}]$, where I is the identity matrix and $i = \sqrt{-1}$, and the controlled not (c-not) (20). These generate the normalizer group. To complete the set of operations, we add preparation of $|0\rangle$ and $|\pi/8\rangle = \cos(\pi/8)|0\rangle + \sin(\pi/8)|1\rangle$ and measurement of $|0\rangle$ and $|1\rangle$ (21). Classical com-

putation can be used to process measurement outcomes and control future quantum gates.

Assumptions and error models. To describe a noisy quantum network, we introduce the notion of error locations, which are sites where errors can occur. The behavior of a quantum network can be represented as a sum of networks representing all combinations of possible errors. We call this an error expansion of the network. The final state of the computation can be evaluated by summing the states associated with each part of the error expansion. Operational error locations are placed after each gate (including state preparation but not measurements). Memory errors exist on each bit between operations. We give an analysis for operational errors; the generalization to other errors is straightforward.

Under imperfect evolution of a network, we can distinguish the part of the error expansion that has the desired behavior and the erroneous part. The overall amplitude of failure is the maximum amplitude of the incorrect part over all possible initial states. This quantity is closely related to the standard notions of fidelity (22) but is easier to analyze in the present context. For stochastic models, error probability is given by the square of the error amplitude. Thus, noise limits are generally much more stringent when considering nonstochastic models.

Assumptions on error behavior are conveniently expressed as constraints on allowed error expansions. There are three classes of constraints that we consider. The first class involves the types of error operators that can occur at a given location. We assume that no bits are lost during the computation. The second class concerns restrictions on coherences between operators in an error expansion. Both stochastic and nonstochastic errors are considered. The last class constrains correlations between errors at different locations. We consider two error models.

1) Independent stochastic errors. This is the simplest model. It assumes that errors are distributed independently and randomly at each error location. The probability associated with the model is the probability of having an error at an error location.

2) Quasi-independent (monotonic) errors. This model relaxes the assumptions of the previous model by allowing some coherence and correlation between errors at different locations. It requires that the maximum amplitude of those parts of the error expansion having errors at a given k location, or any subset of those parts, is bounded by a constant times p^k , where p is

E. Knill is at CIC-3, Mail Stop B265, and R. Laflamme and W. H. Zurek are at T-6, Mail Stop B288, Los Alamos National Laboratory, Los Alamos, NM 87545, USA.

the amplitude parameter of the model. One can show that the probability (or amplitude) of failure of the whole computation in these error models is bounded by a constant times np , where n is the number of error-locations in the network.

How realistic are these error models? Because physics is described by local interactions, it is reasonable to assume that error events on bits are caused by independent environments except when they are intentionally modified by an interaction implementing a multibit operation. This is the independent-error model in physical terms. It covers not only independent stochastic errors but also coherence of small errors at different locations. For example, a modification of the internal energies of each bit by a weak external field is allowed, provided the deviations are small enough. The quasi-independent error model is more general. In addition to other error types, it can also support weak pairwise interactions between adjacent bits.

Elements of fault tolerance. The first element required for resilience is quantum error correction. In classical computation, error correction is usually accomplished by redundantly encoding information. The simplest method involves copying the information at least three times and using majority voting to recover it after errors have occurred. This method cannot be straightforwardly applied to quantum computers for three reasons. First, it is not possible to clone unknown quantum states (23). Second, in order to take a majority vote, it is naively thought that we must first learn the encoded information by measurement; the act of measuring would destroy any quantum coherence of the state. Finally, only one type of error needs to be considered for classical binary information, the bit flip, whereas quantum states can be modified by a continuum of possible errors.

Shor (10) and Steane (11, 24) discovered how these objections could be overcome. To avoid copying information to introduce redundancy, it is possible to exploit entangled states supported

by additional bits. To avoid collapse of the quantum information in the process of correcting errors, it is possible to make a partial measurement that extracts only error information (the syndrome) and leaves the encoded state untouched. To deal with the continuum of possible errors, it is sufficient to recognize that every error can be represented as a linear combination of the standard errors (11) (no error, bit flip, sign flip, or both). Together with the observation that linear combinations of correctable errors are also correctable, this process allows discretization of the error possibilities. The general theory of error correction is discussed in (25).

Error-correcting codes work under the assumptions that encoding and error correction are error free. This situation can be a good approximation in the case of quantum memory. However, we cannot expect that the operations involved in recovering the encoded state are exact. A method to implement these operations with fault tolerance is given in (17, 26). For the one-error-correcting codes to be used in this work, the method has the property that the error correction succeeds provided one error at most occurred in the network. In the worst case, such an error is equivalent to introducing one error either before or after a successful error-correction step.

Quantum error-correcting methods protect information in memory. To maintain the protection while performing computations requires operating directly on the encoded state. How this can be done without introducing uncorrectable errors is discussed in (17, 27–29). To ensure that errors in encoded operations are corrected, each such operation is preceded by two attempts at fault tolerantly correcting errors in each of the incoming encoded bits. Encoded state preparation is similar but requires multiple preparation attempts to ensure an improved chance of success (30).

The final ingredient required to implement resilience is concatenation. The combination of quantum error correction, fault-tolerant error correction, and encoded operations can be viewed as a technique for exploiting bits on which we can operate with error p per gate to define abstract (encoded) bits with a smaller error per gate. The effective error probability is reduced from p to at most cp^2 , where c is a constant to be determined below. Concatenation involves applying this combination of techniques hierarchically. The abstract bits defined at one level are used for encoding bits at the next level. Figure 1 shows a network for concatenat-

ing a 3-bit code twice. The effective error on the encoded state after h levels of concatenation is at most $c^{2h-1}p^{2h}$. Therefore, by increasing the number of levels in the hierarchy, we can reduce the probability of error to any desired value.

Analysis. There is a simple method for estimating the thresholds associated with hierarchical encodings (Box 1). The threshold can be obtained by exploiting “single-error elimination” networks. They have the property that if an error occurs at a location with no other errors present in the “region of influence,” then it will not affect the encoded state. A naïve estimate for the probability of error at the next level of a hierarchy is given by the number of pairs of errors that can occur within any region of influence times the square of the probability of error at this level. We use the 7-bit code described in (11, 16). When implemented with fault-tolerant error correction, the number of error locations in a region of influence is bounded by ~ 1204 , which is the product of the number of bits involved in an encoded 2-bit gate (2×7) and the number of operations affecting a bit before an error at a given location can be eliminated (~ 86). This yields a probability of error at the next level of less than $10^6 p^2$, which is less than p when $p < 10^{-6}$.

There are two problems with the simple estimate of the previous paragraph. The first is that it is necessary to define what it means for an encoded gate to fail. The second is that in order to apply the analysis recursively at each level of the hierarchy, the induced errors must satisfy the same error model with the desired error probability. For the independent stochastic error models, this goal cannot be achieved.

To define encoded gate failure, we consider the components of the error expansion. Each component associates a specific error with each error location. These errors can be moved to other locations by means of error propagation without changing the overall behavior of the network. For example, σ_x before a 90° rotation around the y axis is equivalent to σ_z after the rotation. Encoded gate errors result from the propagation of errors from bits at lower levels. In order to perform this propagation correctly while preserving the error model, it is necessary to associate to each encoded gate a specific region of the network.

The algorithm for associating regions to each encoded gate and determining gate failure proceeds in several steps. First, each encoded gate is associated with the encoding network and the two preceding error-correction attempts for each of




Fig. 1. Concatenation of the 3-bit code. If the error rate is p , the encoding will give a rate of $c^{2h-1}p^{2h}$ for the h th level of the hierarchy.

the bits involved. If at least two errors occurred in these networks, the encoded gate has failed. (Because the failure may not actually propagate to the next level, the probability of failure is overestimated.) Next, we reconsider each gate not yet failed and reallocate their regions. For each pair of successive error-correction attempts belonging to one of these gates, if the first one has an error, its network is allocated to the previous encoded gate. Any encoded gate with at least two errors in the newly allocated region has failed.

These allocations of regions ensure that (i) a failed gate has at least two errors and (ii) the first of the error-correction attempts of a successful gate is error-free. Therefore, any single, preexisting error in the bits is corrected, and there is at most one error in the gate's region, which does not affect the encoded information. As a result, the encoded gate after error propa-

gation is error-free.

To determine the probability of error at the next level in the quasi-independent error models, it now suffices to observe that for a given k encoded gates to fail, two different errors must have occurred for each gate in its allocated region. Because these regions are bounded, this probability can be estimated as $(cp^2)^k$, where c is the number of pairs of locations among those that can be allocated to a given gate. If the number of such locations is m , then $c \leq m^2/2$. In our case, the gate with the largest value of m is preparation of the $|\pi/8\rangle$ state, with $m \leq 792$, which gives $c \leq 313894$ and a threshold of 3×10^{-6} . The overhead, in the number of extra operations, required for implementing this version of resilient quantum computation is polylogarithmic in the number of steps of the original algorithm and the inverse of the maximum tolerable probability of failure.

Discussion. We have demonstrated that quantum computation can be performed arbitrarily accurately provided that the noise per operation is sufficiently small and satisfies suitable independence assumptions. Because the overheads are asymptotically well behaved, the threshold results demonstrate that quantum computation is possible in the presence of physically reasonable sources of noise.

Threshold results have been obtained independently by Kitaev (31) and Aharonov and Ben-Or (32). They analyzed independent stochastic error models and obtained completeness of operations by adopting Shor's implementation of the Toffoli gate. Kitaev used a different method for fault-tolerant extraction of the syndrome. His method is less efficient and consequently yielded substantially worse thresholds. Aharonov and Ben-Or provided an analysis that does not require accurate classical computation for syndrome calculations, and they estimated a threshold of around 10^{-6} for the independent stochastic-error model. There may be some cases where this extension is needed, for example when performing ensemble quantum computation such as that envisioned for nuclear magnetic resonance (33, 34). Our techniques for generalizing arguments to quasi-independent error models can in principle be used to extend their analysis.

Not only do the threshold theorems show that quantum computation is possible in principle, but they demonstrate that the apparent distance limitations of quantum cryptography can be overcome. It suffices to be able to transmit the bits over some reasonable distance before error-correction operations must be applied to avoid loss of encoded information.


The actual values of the thresholds we have obtained are rigorous but overly pessimistic in several ways. First, the error models used are the most adversarial satisfying independence assumptions. In practice, we need not worry about such adversarial error behavior. The actual error types at the physical level are likely to be much more constrained than assumed by our error-blind analysis. That known error behavior can be exploited to reduce error has been demonstrated in a specific example (35). In addition, simulations suggest that for the depolarizing channel, thresholds are substantially better than suggested by our calculations (28, 36). Second, we have made no attempt to optimize the implementation of fault tolerance. Suggestions for optimization can be found in (28, 30, 36). Nevertheless, the results suggest that nonstochastic errors such as those studied in (19) must be

Estimate of the threshold: An example. To illustrate our method for estimating the threshold, we assume that only phase errors are relevant and a 3-bit code is sufficient. We consider the hierarchical encoding of a c-not gate (i). To determine the reduction of error rate in the fault tolerant implementation (ii), we calculate the total number of error locations. This implementation includes two error-correction steps and an encoded version of the c-not.

The error correction step (iii) involves two independent determinations of the error to ensure that it was measured correctly. Cat states $|00\rangle + |11\rangle$ (iv) are prepared for fault-tolerant measurement of the error. The error is fixed by applying a sign flip σ_z to the erroneous bit.

The yellow diamonds correspond to error locations. Diamonds connected by a vertical line are considered one location. The triangle surrounding $|0\rangle$ denotes state preparation; the one surrounding M , a measurement. The dotted lines indicate operations conditional on the results of measurements. H is a sign flip followed by a 90° rotation around the y axis.

We can now count the number of error locations: there are $c = 9$ in the cat state preparation. Two of these states are needed for each error determination. Each error determination contains $e = 4$ error locations and 1 in the error-correction step. Four error-correction attempts are needed, and $c_n = 3$ error locations are present in the encoded c-not gate. This gives $c_n + 4[2(e + 2c) + 1] = 183$ error locations. Thus, an upper bound for pairs of error locations is $183 \cdot 182/2 = 16,653$, which gives a threshold of $1/16,653 \sim 6 \times 10^{-5}$.



controlled more carefully than stochastic noise.

Whether resilient quantum computation can be implemented in practice remains to be seen. However, the results obtained here show that, in principle, noise of a level below the error threshold is not an obstacle for quantum computation.

REFERENCES

1. D. Deutsch and R. Jozsa, *Proc. R. Soc. London Ser. A* **439**, 553 (1985).
2. E. Bernstein and U. Vazirani, in *Proceedings of the 25th ACM Symposium on the Theory of Computation* (ACM Press, New York, 1993), pp. 11–20.
3. D. R. Simon, in *Proceedings of the 35th Annual Symposium on Foundations of Computer Science* (IEEE Press, Los Alamitos, CA, 1994), pp. 116–123.
4. P. W. Shor, in *ibid.*, pp. 124–134.
5. J. I. Cirac and P. Zoller, *Phys. Rev. Lett.* **74**, 4091 (1995).
6. C. Monroe *et al.*, *ibid.* **75**, 4714 (1995).
7. W. H. Zurek, *Phys. Today* **44**, 36 (October 1991).
8. R. Landauer, *Philos. Trans. R. Soc. London* **353**, 367 (1995).
9. W. G. Unruh, *Phys. Rev. A* **51**, 992 (1995).
10. P. W. Shor, *ibid.* **52**, 2493 (1995).
11. A. Steane, *Proc. R. Soc. London Ser. A* **452**, 2551 (1996).
12. A. R. Calderbank and P. W. Shor, *Phys. Rev. A* **54**, 1098 (1996).
13. R. Laflamme, C. Miquel, J. P. Paz, W. H. Zurek, *Phys. Rev. Lett.* **77**, 198 (1996).
14. C. H. Bennett, D. P. DiVincenzo, J. A. Smolin, W. K. Wootters, *Phys. Rev. A* **54**, 3824 (1996).
15. D. Gottesman, *ibid.*, p. 1862.
16. A. R. Calderbank, E. M. Rains, P. W. Shor, N. J. A. Sloane, *ibid.* **78**, 405 (1997).
17. P. W. Shor, in *Proceedings of the Symposium on the Foundations of Computer Science* (IEEE Press, Los Alamitos, CA, 1996), pp. 56–65.
18. E. Knill and R. Laflamme, quant-ph/9608012, Los Alamos e-Print Archive at xxx.lanl.gov (1996).
19. C. Miquel, J. P. Paz, W. H. Zurek, *Phys. Rev. Lett.* **78**, 3971 (1997).
20. A. Barenco *et al.*, *Phys. Rev. A* **52**, 3457 (1995).
21. E. Knill, R. Laflamme, W. H. Zurek, quant-ph/9702058, Los Alamos e-Print Archive at xxx.lanl.gov (1997).
22. R. Jozsa, *J. Mod. Opt.* **41**, 2315 (1995).
23. W. K. Wootters and W. H. Zurek, *Nature* **229**, 802 (1982).
24. A. Steane, *Phys. Rev. Lett.* **77**, 793 (1996).
25. E. Knill and R. Laflamme, *Phys. Rev. A* **55**, 900 (1996).
26. D. P. DiVincenzo and P. W. Shor, *Phys. Rev. Lett.* **77**, 3260 (1996).
27. W. H. Zurek and R. Laflamme, *ibid.*, p. 4683.
28. J. Preskill, *Proc. R. Soc. London Ser. A*, in press.
29. D. Gottesman, quant-ph/9702029, Los Alamos e-Print Archive at xxx.lanl.gov (1997).
30. E. Knill, R. Laflamme, W. Zurek, quant-ph/9610011 and quant-ph/9705031, Los Alamos e-Print Archive at xxx.lanl.gov (1996, 1997).
31. A. Y. Kitaev, in *Quantum Communication, Computing, and Measurement*, O. Hirota *et al.*, Eds. (Plenum, New York, 1997).
32. D. Aharonov and M. Ben-Or, quant-ph/9611025, Los Alamos e-Print Archive at xxx.lanl.gov (1996).
33. D. G. Cory, A. F. Fahmy, T. F. Havel, in *Proceedings of the 4th Workshop on Physics and Computation* (New England Complex Systems Institute, Boston, MA, 1996).
34. N. A. Gershenfeld and I. L. Chuang, *Science* **275**, 350 (1997).
35. J. I. Cirac, T. Pellizzari, P. Zoller, *ibid.* **273**, 1207 (1996).
36. C. Zalka, quant-ph/9612028, Los Alamos e-Print Archive at xxx.lanl.gov (1996).
37. We thank the National Security Agency for support.

15 May 1997; accepted 21 November 1997

Absence of a Spin Gap in the Superconducting Ladder Compound $\text{Sr}_2\text{Ca}_{12}\text{Cu}_{24}\text{O}_{41}$

H. Mayaffre, P. Auban-Senzier, M. Nardone, D. Jérôme,*
D. Poilblanc, C. Bourbonnais, U. Ammerahl,
G. Dhalenne, A. Revcolevschi

Transport and copper-63 nuclear magnetic resonance measurements of the Knight shift and relaxation time T_1 performed on the two-leg spin ladders of $\text{Sr}_2\text{Ca}_{12}\text{Cu}_{24}\text{O}_{41}$ single crystals as a function of pressure show a collapse of the gap in ladder spin excitations when superconductivity is stabilized at 31 kilobars. This result suggests that the superconducting phase in these materials may be connected to this transition and the collapse of the spin gap, and support the prediction made with exact diagonalization techniques in two-leg isotropic $t - J$ ladder models of a transition between a low-doping spin gap phase and a gapless regime.

The existence of superconductivity in two families of materials where this property was not expected at first sight [the low-

dimensional organic conductors and the high-transition temperature (T_c) cuprates] has been a major achievement of condensed matter research of the last two decades. The mechanism of superconductivity for both classes of compounds is still under intense debate, but there is already a consensus about their low-dimensional electronic structure that may be the clue governing superconducting (SC) pairing correlations. The recent finding of new SC copper oxide structures (1) exhibiting one-dimensional (1D) features with both isolated CuO_2 chains and Cu_2O_3 ladders—that is, pairs of CuO_2 chains linked by oxygen atoms between the coppers—has profoundly revived the interest for superconductivity in cuprates and 1D

materials.

We discuss the ladder compound $\text{Sr}_2\text{Ca}_{12}\text{Cu}_{24}\text{O}_{41}$, which derives from the parent compound $\text{Sr}_{14}\text{Cu}_{24}\text{O}_{41}$ through Ca substitution. The structure of $\text{Sr}_{14}\text{Cu}_{24}\text{O}_{41}$ displays CuO_2 chains and Cu_2O_3 two-leg ladders parallel to the c axis of the structure (2); other insulating 1D materials like SrCu_2O_3 contain only Cu_2O_3 ladders and no chains. In contrast, the undoped parent compound for the high- T_c cuprates exhibits a 2D CuO_2 layer structure. In both systems, all copper sites belonging to the ladders or to the planes are occupied by a spin $1/2$ Cu^{2+} ion. However, although long-range antiferromagnetism is stabilized at low T in the 2D spin system, the properties of the spin ladder materials can be drastically different. In two-leg ladder systems, dominant antiferromagnetic (AF) coupling J between the copper spins on the same rung leads to the formation of a spin singlet on each rung. Consequently, the ground state of the whole ladder is a singlet spin state, and a finite energy is needed to excite a rung spin singlet to a spin triplet state. A spin gap situation is obtained with a characteristic exponential drop of the spin susceptibility upon cooling down.

The existence of a spin gap in a spin-ladder structure has been first proposed theoretically (3) and found experimentally in several even-leg ladder copper oxide systems [SrCu_2O_3 (4, 5), $\text{LaCuO}_{2.5}$ (6)] or organic materials (7). The spin gap is expected to be quite robust to various perturbations. For example, it is predicted to be stable up to arbitrary small magnetic coupling along the rungs of the ladders (8) or in the presence of a small interladder coupling (9).

H. Mayaffre, P. Auban-Senzier, M. Nardone, D. Jérôme, Laboratoire de Physique des Solides URA 002 (associé au CNRS), Université Paris-Sud, 91405 Orsay, France.

D. Poilblanc, Laboratoire de Physique Quantique UMR 5626 (associé au CNRS), Université Paul Sabatier, 31062 Toulouse, France.

C. Bourbonnais, Centre de Recherche en Physique du Solide, Département de Physique, Université de Sherbrooke, Sherbrooke, Québec, Canada J1K2R1.

U. Ammerahl, Laboratoire de Chimie des Solides, URA 446 (associé au CNRS), Université Paris-Sud, 91405 Orsay, France, and II Physikalisches Institut, Universität zu Köln, Zùlpicher Strasse 77 D-50937 Köln, Germany.

G. Dhalenne and A. Revcolevschi, Laboratoire de Chimie des Solides, URA 446 (associé au CNRS), Université Paris-Sud, 91405 Orsay, France.

*To whom correspondence should be addressed.