Robust quantum computation

vs.

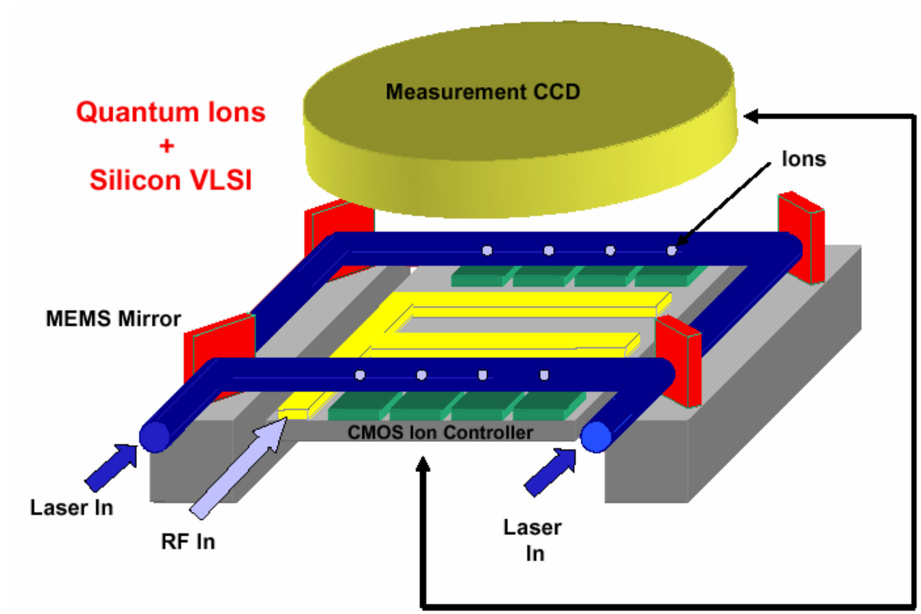

Institute for Quantum Information

Sponsored by the NSF

John Preskill, Caltech

27 February 2007
http://www.iqi.caltech.edu/

Topological quantum computer:

The Dream

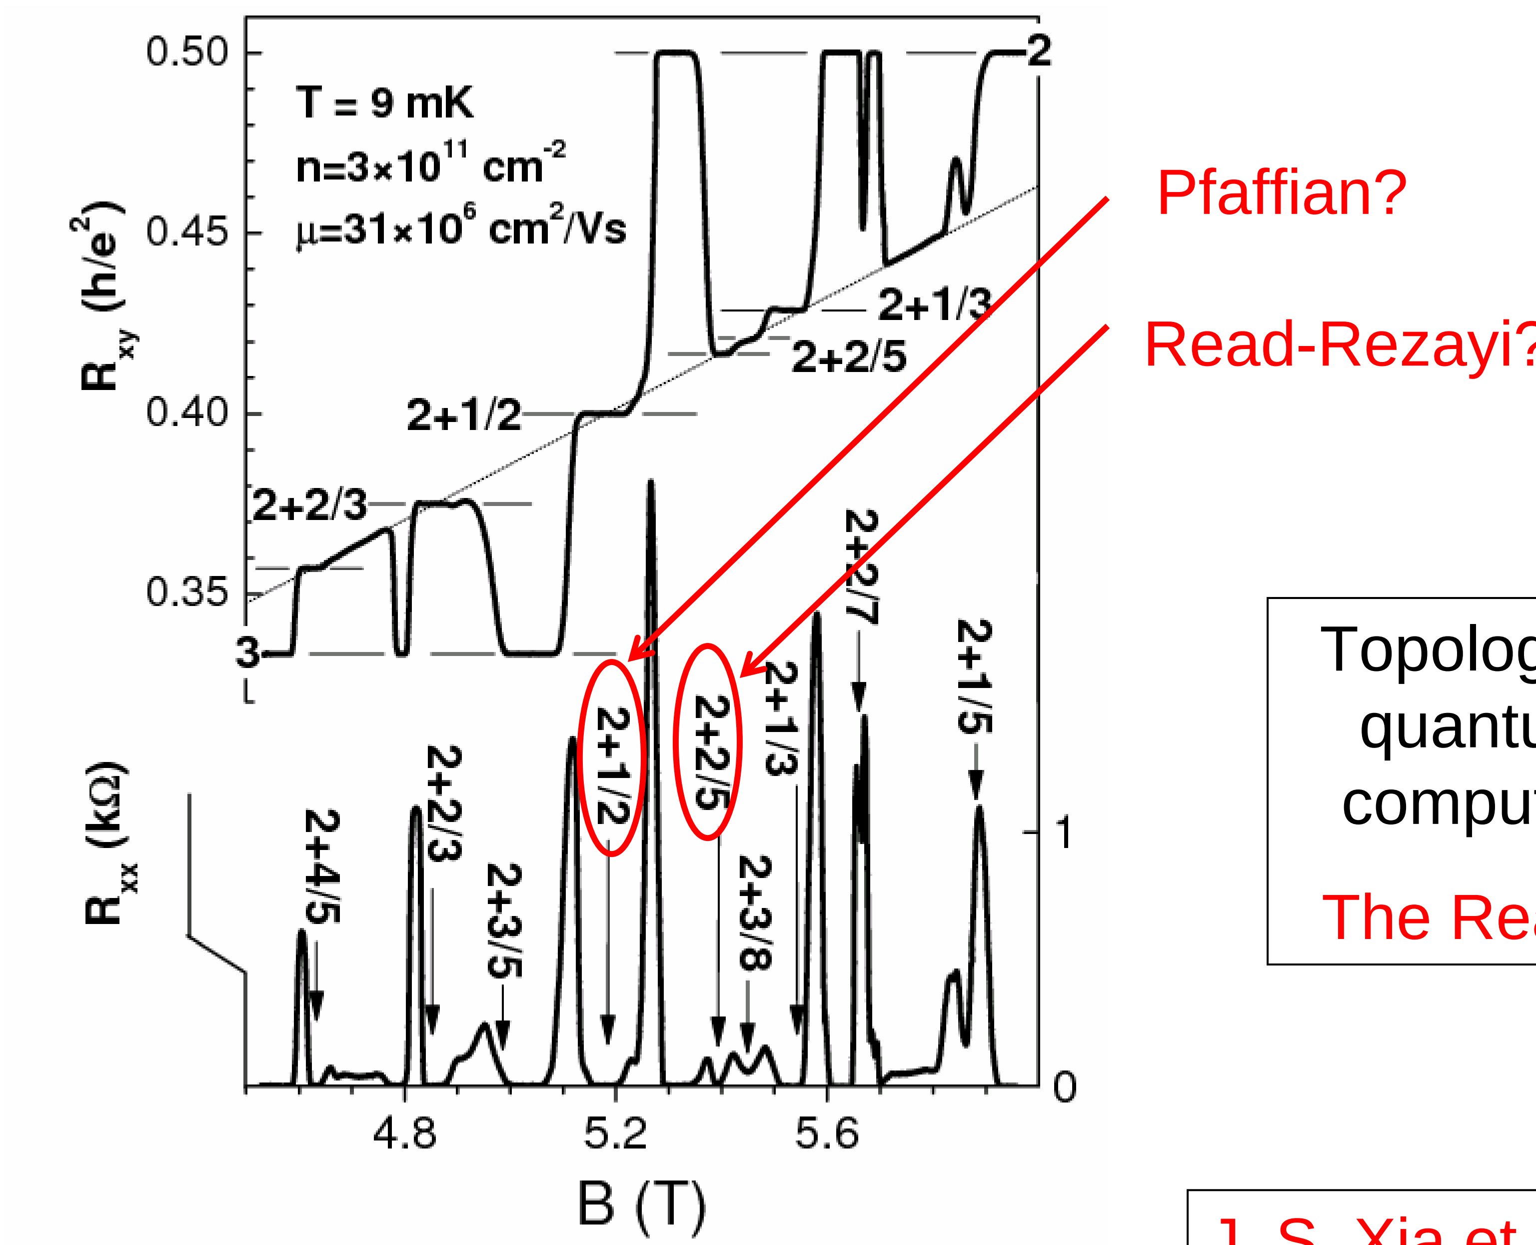

Topological quantum computer:

The Reality

J. S. Xia et al. (2004)

Scalable ion trap quantum computer: The Dream

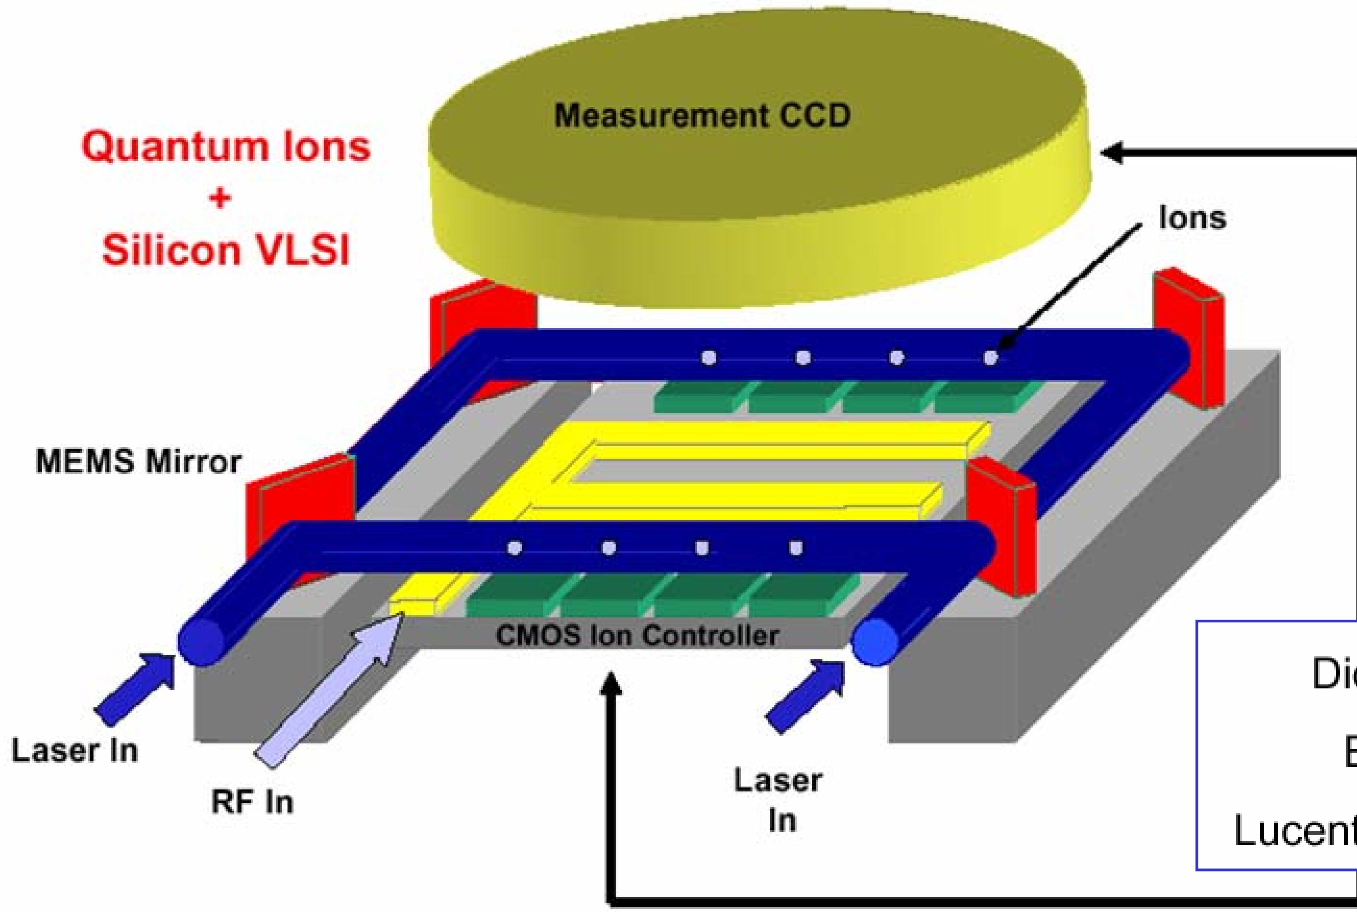

Dick Slusher

Bell Labs

Lucent Technologies

System Compatibility of Quantum & Classical: Spatial Pitch, Clock Speed
Operating Temperature, Power Dissipation

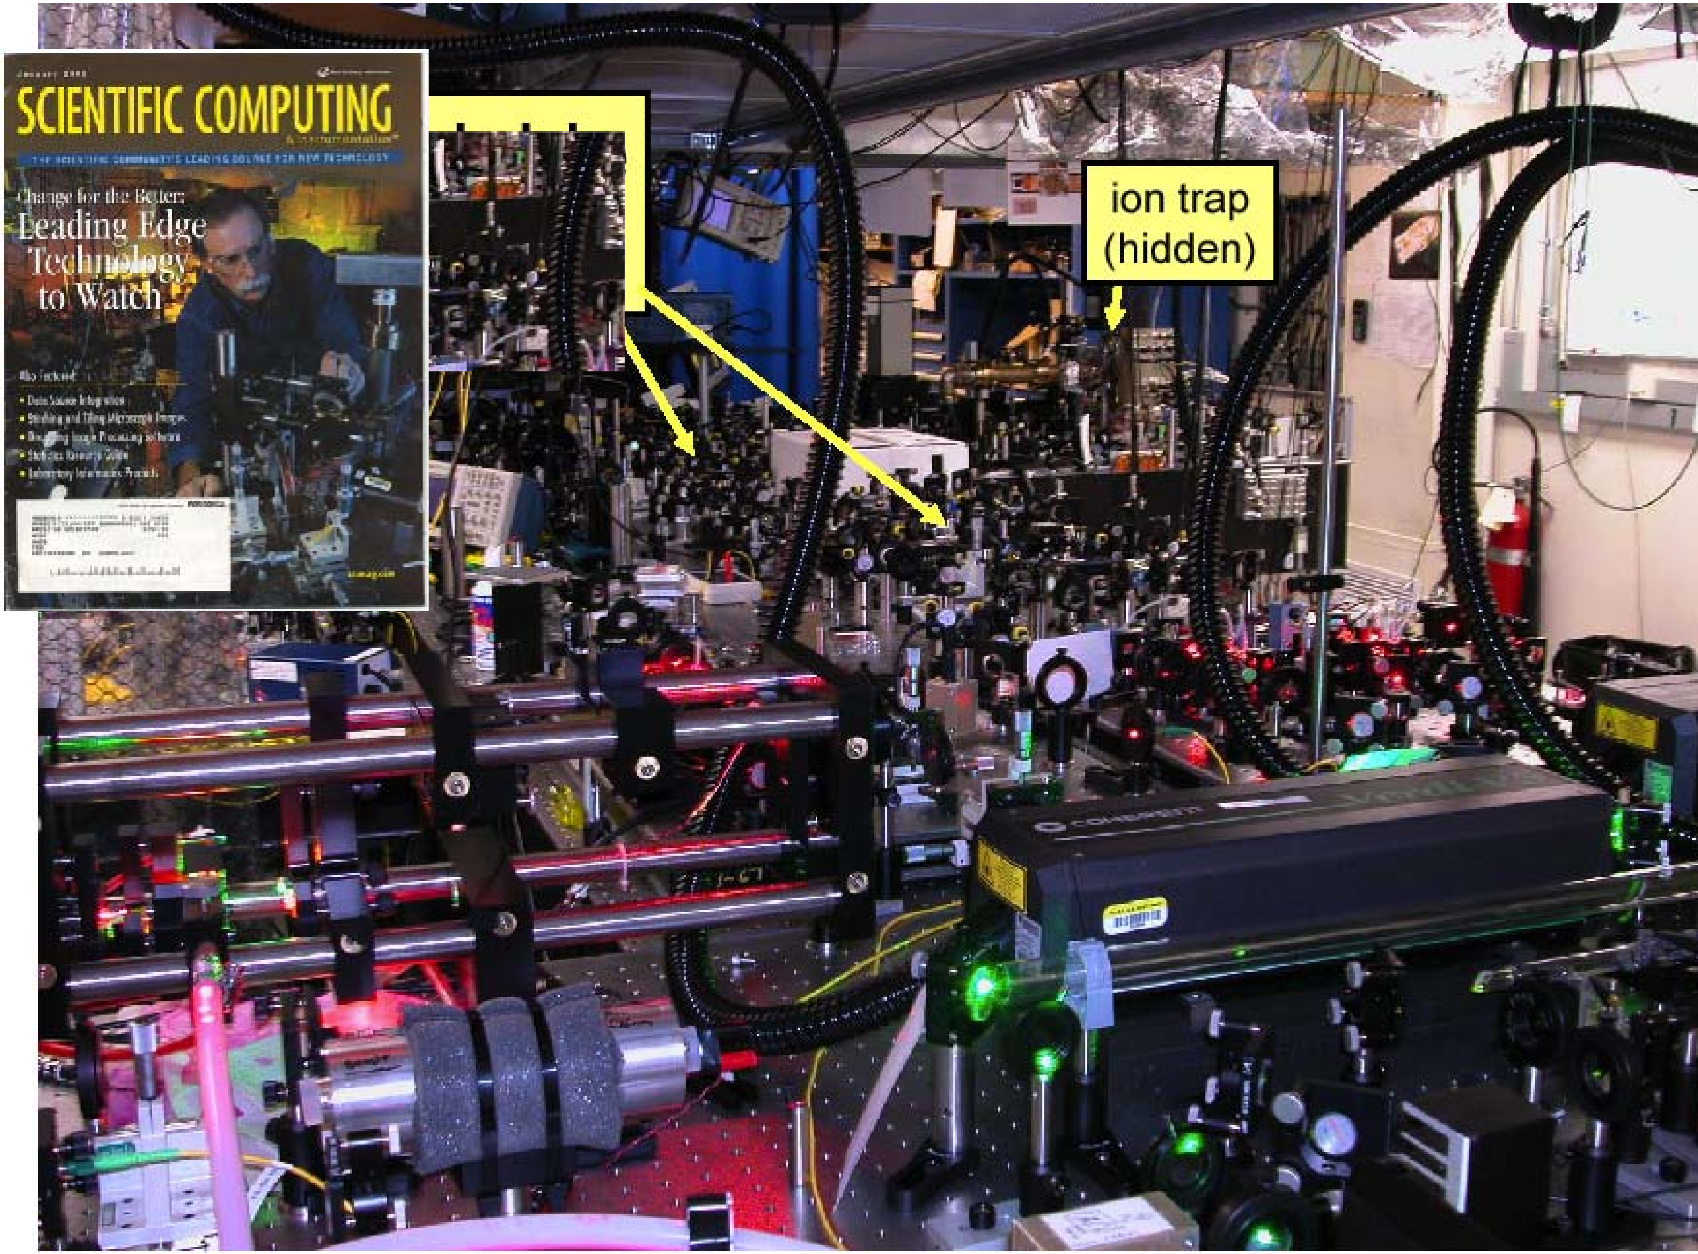

Ion trap quantum computer: The Reality

GaAs Ion Trap

D. Stick, W. Hensinger, S. Olmschenk, M. Madsen (Michigan)
K. Schwab (Laboratory for Physical Sciences)

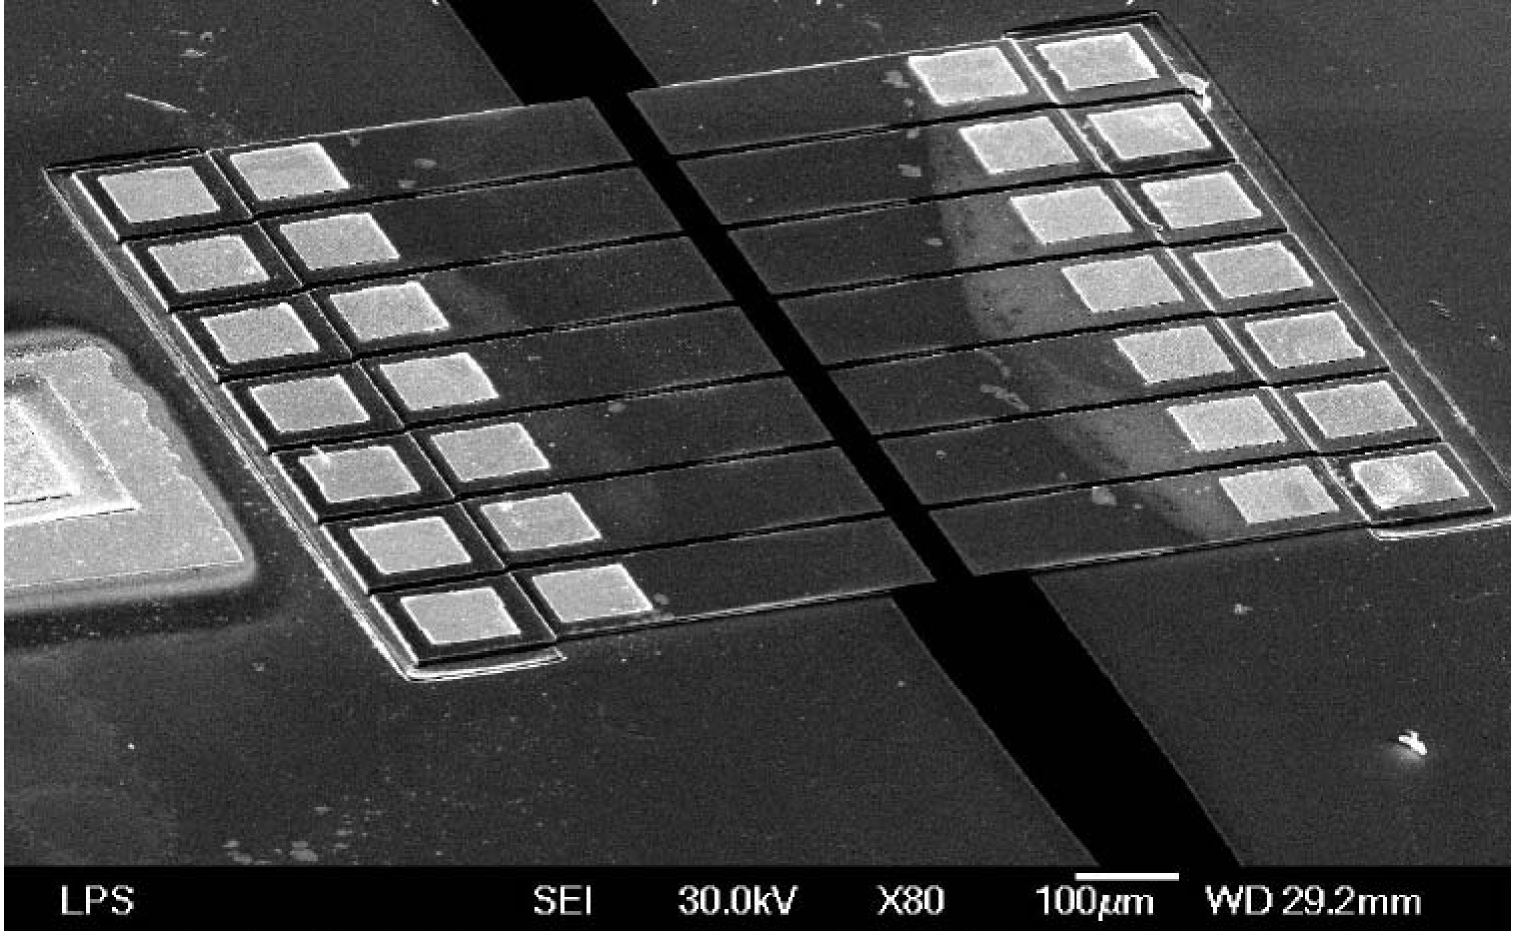

Quantum hardware:

- The future does not necessarily belong to the ion trappers: for example, electron spins in quantum dots, superconducting qubits, ultracold neutral atoms are all making impressive progress.

- But ion traps have a head start, and some serious effort has been devoted to conceiving scalable architectures.

- "Ion trap chips look well placed to create useful computers before other methods." – Andrew Steane

- "There is progress, but it's still very slow." – Chris Monroe

- "I'd say almost any prediction about what a quantum computer will look like will, with high probability, be wrong. Ion trappers are encouraged because we can at least see a straightforward path to making a large processor, but the technical problems are extremely challenging. It might be fair to say that ion traps are currently in the lead; however, a good analogy might be that we're leading a marathon race, but only one meter from the start line."

  – Dave Wineland

Quantum fault tolerance: Topological vs. “Brute force”

- Error correction and fault tolerance will be essential in the operation of large-scale quantum computers, both to prevent decoherence and to control the accumulating effects of small errors in unitary quantum gates.

- Topological quantum computing is the elegant approach, in which the “hardware” is intrinsically robust due to principles of local quantum physics (if operated at a temperature well below the mass gap). We hope it will work, but a physical realization of a topological quantum computer may be hard to achieve.

- There is also a “standard” approach to fault-tolerant quantum computing, which uses clever circuit design to overcome the deficiencies of quantum hardware. It, too, works in principle, if the hardware is not too noisy.

- Either approach (or perhaps a combination of the two) might eventually lead to quantum computers capable of solving hard problems. Which path we eventually follow will depend on which turns out to be more feasible technologically, and we don’t know that yet.

Robust quantum computation

1. Quantum error-correcting codes
2. Fault-tolerant quantum computing
3. Quantum accuracy threshold theorem
4. New developments:
   a) subsystem codes
   b) local gates
   c) slow measurements
   d) postselected simulation
5. Questions:
   a) high-frequency noise
   b) asymmetric noise
   c) 3D topological order

Quantum computer: the standard model

(1) Hilbert space of n qubits: $\mathfrak{H} = \mathbb{C}^{2^n}$

(2) prepare initial state: $|0\rangle^{\otimes n} = |000\ldots0\rangle$

(3) execute circuit built from set of universal quantum gates: $\{U_1, U_2, U_3, \ldots U_{n_G}\}$

(4) measure in basis $\{|0\rangle, |1\rangle\}$

The model can be simulated by a classical computer with access to a random number generator. But there is an exponential slowdown, since the simulation involves matrices of exponential size… Thus we believe that quantum model is intrinsically more powerful than the corresponding classical model.

Our goal is to simulate accurately the ideal quantum circuit model using the imperfect noisy gates that can be executed by an actual device (assuming the noise is not too strong).

Errors

The most general type of error acting on n qubits can be expressed as a unitary transformation acting on the qubits and their environment:

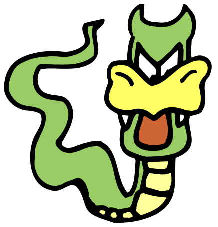

$$|\psi\rangle$$

Errors

The most general type of error acting on n qubits can be expressed as a unitary transformation acting on the qubits and their environment:

$$U : |\psi\rangle \otimes |0\rangle_E \to \sum_a E_a |\psi\rangle \otimes |a\rangle_E$$

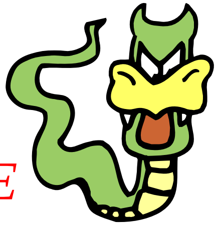

Errors

The most general type of error acting on n qubits can be expressed as a unitary transformation acting on the qubits and their environment:

$$U : |\psi\rangle \otimes |0\rangle_E \rightarrow \sum_a E_a |\psi\rangle \otimes |a\rangle_E$$

The states $|a\rangle_E$ of the environment are neither normalized nor mutually orthogonal. The operators $\{E_a\}$ are a basis for operators acting on n qubits, conveniently chosen to be "Pauli operators": $\{I, X, Y, Z\}^{\otimes n}$,

where

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

The errors could be "unitary errors" if $|a\rangle_E = C_a |0\rangle_E$ or decoherence errors if the states of the environment are mutually orthogonal.

Errors

$$U: |\psi\rangle \otimes |0\rangle_E \rightarrow \sum_a E_a |\psi\rangle \otimes |a\rangle_E$$

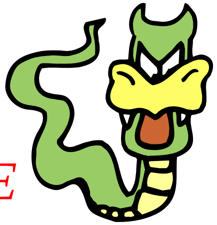

Our objective is to recover the (unknown) state $|\psi\rangle$ of the quantum computer. We can't expect to succeed for arbitrary errors, but we might succeed if the errors are of a restricted type. In fact, since the interactions with the environment are *local*, it is reasonable to expect that the errors are not too strongly correlated.

Define the "weight" w of a Pauli operator to be the number of qubits on which it acts nontrivially; that is $X,Y,$ or Z is applied to w of the qubits, and I is applied to n-w qubits. If errors are rare and weakly correlated, then Pauli operators E_a with large weight have small amplitude $\| |a\rangle_E \|$.

Error recovery

We would like to devise a recovery procedure that acts on the data and an *ancilla*:

$$V : E_a |\psi\rangle \otimes |0\rangle_A \to |\psi\rangle \otimes |a\rangle_A$$

which works for any $E_a \in \{\text{Pauli operators of weight} \le t\}$.

Then we say that we can "correct t errors" in the block of n qubits. Information about the error that occurred gets transferred to the ancilla and can be discarded:

$$|\psi\rangle \otimes |0\rangle_E \otimes |0\rangle_A \xrightarrow{\text{error}} \sum_a E_a |\psi\rangle \otimes |a\rangle_E \otimes |0\rangle_A$$

$$\xrightarrow{\text{recover}} \sum_a |\psi\rangle \otimes |a\rangle_A \otimes |a\rangle_E = |\psi\rangle \otimes |\varphi\rangle_{EA}$$

Error recovery

$$|\psi\rangle\otimes|0\rangle_E\otimes|0\rangle_A \xrightarrow{\text{error}} \sum_a E_a|\psi\rangle\otimes|a\rangle_E\otimes|0\rangle_A$$

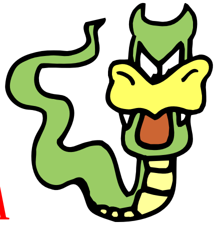

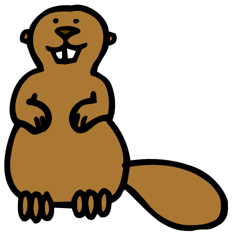

$$\xrightarrow{\text{recover}} \sum_a |\psi\rangle\otimes|a\rangle_A\otimes|a\rangle_E = |\psi\rangle\otimes|\varphi\rangle_{EA}$$

Entropy

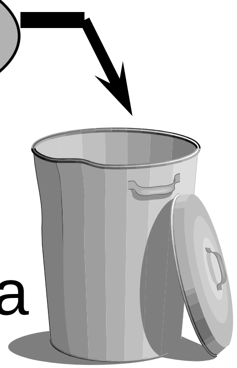

Errors entangle the data with the environment, producing *decoherence*. Recovery transforms entanglement of the data with the environment into entanglement of the ancilla with the environment, "purifying" the data. Decoherence is thus reversed. Entropy introduced in the data is transferred to the ancilla and can be discarded --- we "refrigerate" the data at the expense of "heating" the ancilla. If we wish to erase the ancilla (cool it to $T\approx 0$, so that we can use it again) we need to pay a power bill.

Quantum error-correcting code

We won't be able to correct all errors of weight up to t for arbitrary states $|\psi\rangle \in \mathfrak{H}_{n \text{ qubits}}$. But perhaps we can succeed for states contained in a *code subspace* of the full Hilbert space,

$$\mathfrak{H}_{\text{code}} \in \mathfrak{H}_{n \text{ qubits}}.$$

If the code subspace has dimension 2^k, then we say that k encoded qubits are embedded in the block of n qubits.

How can such a code be constructed? It will *suffice* if

$$\left\{ E_a \mathfrak{H}_{\text{code}}, \quad E_a \in \{\text{Pauli operators of weight} \leq t\} \right\}$$

are mutually orthogonal.

If so, then it is possible in principle to perform an (incomplete) orthogonal measurement that determines the error $\boldsymbol{E_a}$ (without revealing any information about the encoded state). We recover by applying the unitary transformation $\boldsymbol{E_a^{-1}}$.

5-qubit code

Suppose we would like to encode $k=1$ protected qubits in a block of n qubits, and be able to correct all weight-1 errors. How large must n be?

There are two mutually orthogonal "codewords" $|\bar{0}\rangle, |\bar{1}\rangle$ that span the code subspace. Furthermore all $E_a|\bar{0}\rangle, E_b|\bar{1}\rangle$ should be mutually orthogonal.

If n=5, then there are 3×5+1=16 Pauli operators of weight ≤ 1, and the Hilbert space of 5 qubits has dimension 2^5=32. Therefore, for n=5, there is just barely enough room: 16×2 ≤ 2^5=32 .

To see that the code really exists, we can construct it explicitly.

5-qubit code

The code is the simultaneous eigenspace with eigenvalue 1 of 4 commuting *check operators (stabilizer generators):*

$$M_1 = X\ Z\ Z\ X\ I = +1$$
$$M_2 = I\ X\ Z\ Z\ X = +1$$
$$M_3 = X\ I\ X\ Z\ Z = +1$$
$$M_4 = Z\ X\ I\ X\ Z = +1$$

All of these stabilizer generators square to I; they are mutually commuting because there are two collisions between X and Z.

The other three generators are obtained from the first by cyclic permutations. (Note that $M_5 = Z\ Z\ X\ I\ X = M_1 M_2 M_3 M_4$ is not independent.) Therefore, the code is cyclic (cyclic permutations of the qubits preserve the code space).

Claim: no Pauli operator E of weight 1 or 2 commutes with all of the check operators. Weight 1: each column contains an X and a Z. Weight 2: Because the code is cyclic, it suffices to consider $??I\ I\ I$ and $?I\ ?I\ I$ ….

5-qubit code

-- $k=1$ protected qubit

-- corrects $t=1$ error

The code is the simultaneous eigenspace with eigenvalue 1 of 4 commuting *check operators:*

$M_1 = X\,Z\,Z\,X\,I = +1$

$M_2 = I\,X\,Z\,Z\,X = +1$

$M_3 = X\,I\,X\,Z\,Z = +1$

$M_4 = Z\,X\,I\,X\,Z = +1$

By these operators, we can distinguish all possible weight-one errors. Each "syndrome" points to a unique Pauli operator of weight 0 or 1.

	M_1	M_2	M_3	M_4
X_1	+	+	+	-
Y_1	-	+	-	-
Z_1	-	+	-	+
X_2	-	+	+	+
Y_2	-	-	+	-
Z_2	+	-	+	-
X_3	-	-	+	+
Y_3	-	-	-	+
Z_3	+	+	-	+
X_4	+	-	-	+
Y_4	-	-	-	-
Z_4	-	+	+	-
X_5	+	+	-	-
Y_5	+	-	-	-
Z_5	+	-	+	+
I	+	+	+	+

5-qubit code

$M_1 = X\,Z\,Z\,X\,I$

$M_2 = I\,X\,Z\,Z\,X$

$M_3 = X\,I\,X\,Z\,Z$

$M_4 = Z\,X\,I\,X\,Z$

How do we measure the stabilizer generators without destroying the encoded state?

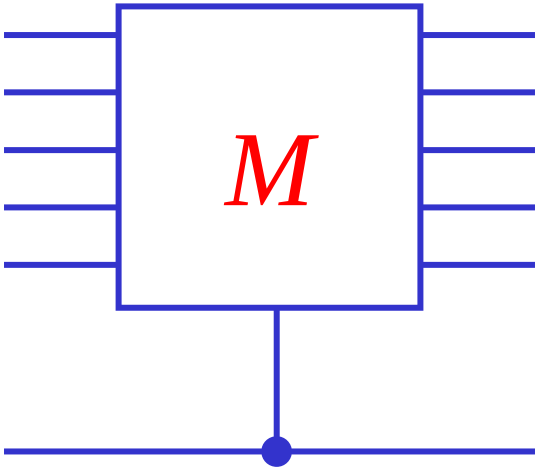

Apply M conditioned on value of an ancilla qubit.

$X=1$ Eigenstate: $|0\rangle_A + |1\rangle_A$ → $|0\rangle_A + M\,|1\rangle_A$

Measure X

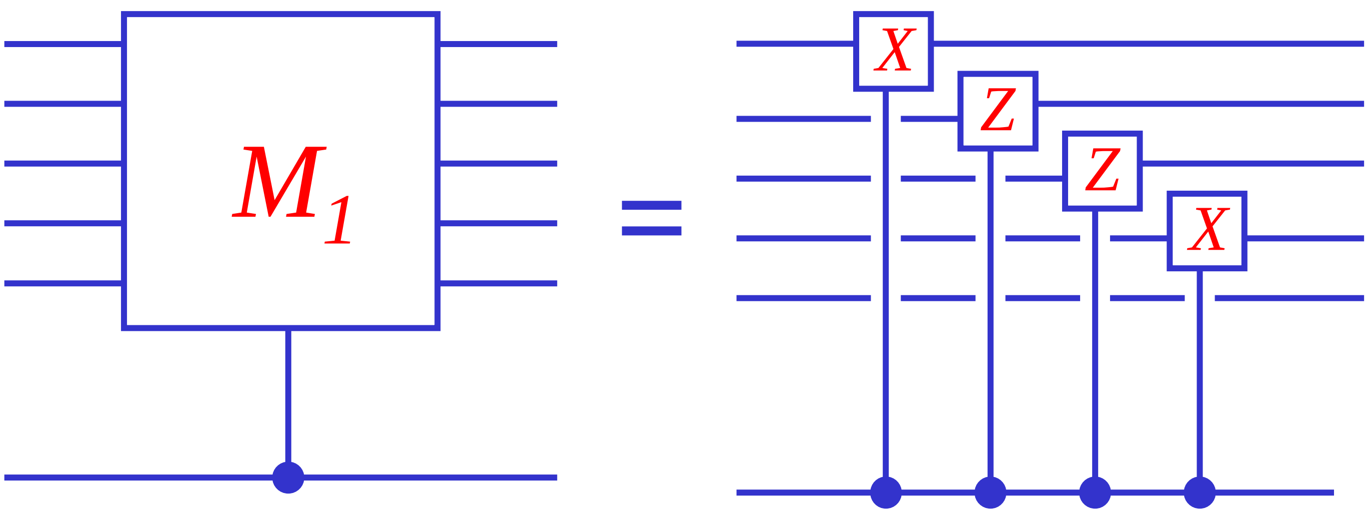

Fault tolerance

- The measured error syndrome (*i.e.*, the eigenvalues of the check operators) might be inaccurate.

- Errors might propagate during syndrome measurement.

- We need to implement a universal set of quantum gates that act on encoded quantum states, without unacceptable error propagation.

- We need codes that can correct many errors in the code block.

Fault-tolerant error correction

Fault: a location in a circuit where a gate or storage error occurs.
Error: a qubit in a block that deviates from the ideal state.

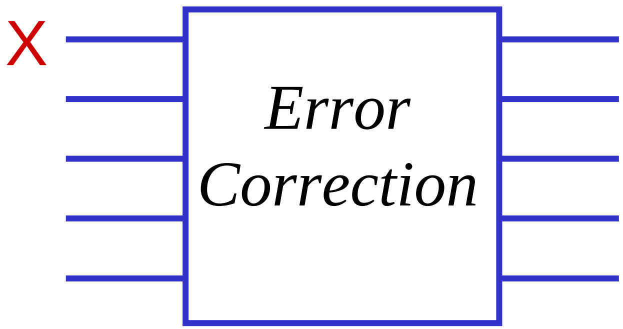

If input has at most one error, and circuit has no faults, output has no errors.

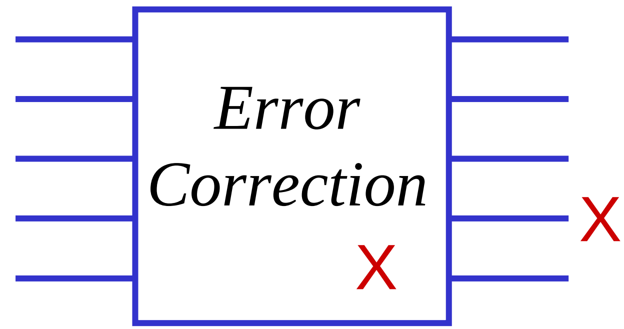

If input has no errors, and circuit has at most one fault, output has at most one error.

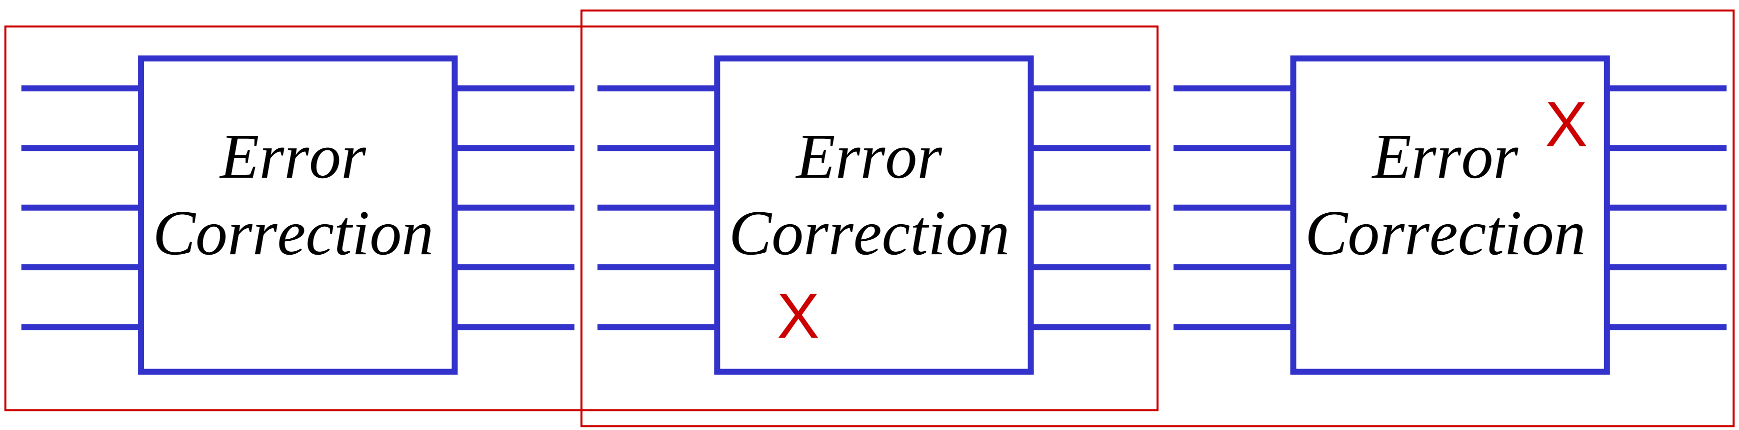

A quantum memory fails only if two faults occur in some "extended rectangle."

Fault-tolerant quantum gates

Fault: a location in a circuit where a gate or storage error occurs.
Error: a qubit in a block that deviates from the ideal state.

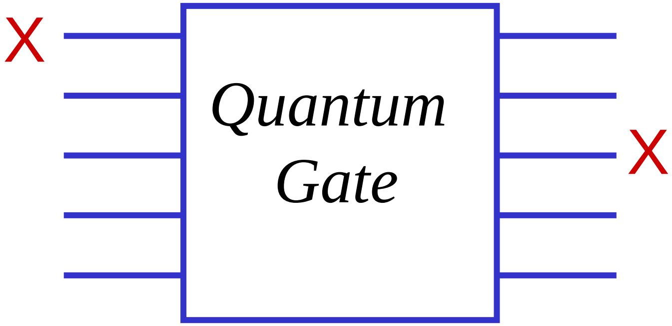

If input has at most one error, and circuit has no faults, output has at most one error in each block.

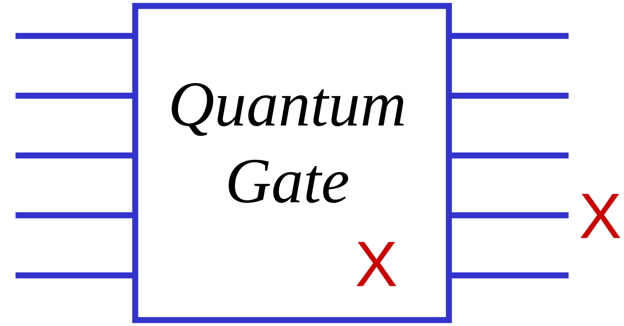

If input has no errors, and circuit has at most one fault, output has at most one error in each block.

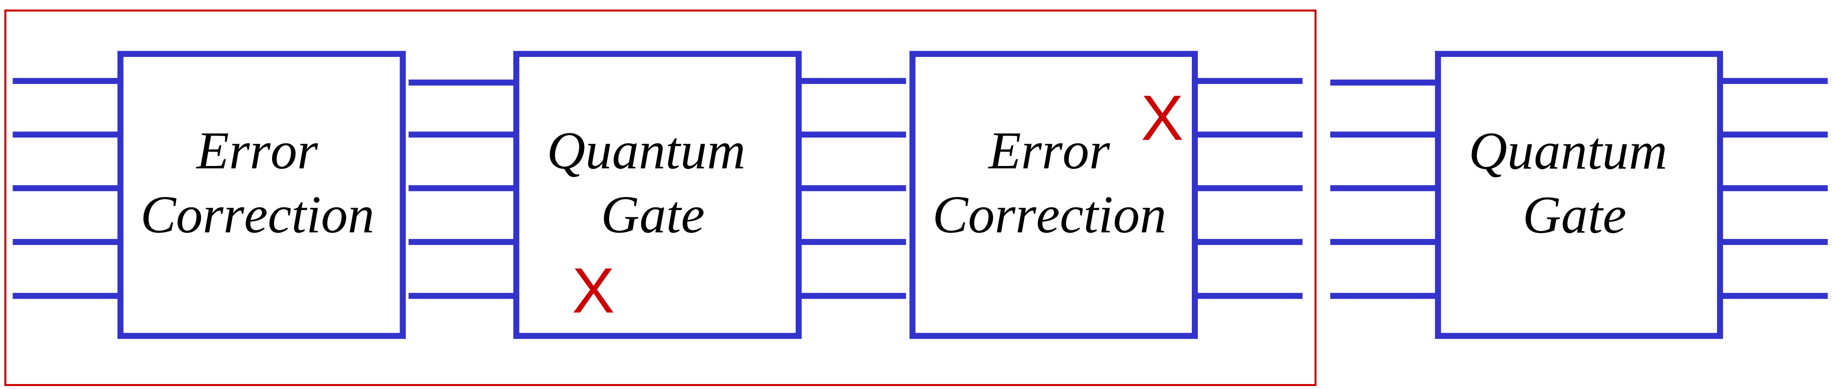

Each gate is preceded by an error correction step. The circuit simulation fails only if two faults occur in some “extended rectangle.”

Fault-tolerant quantum gates

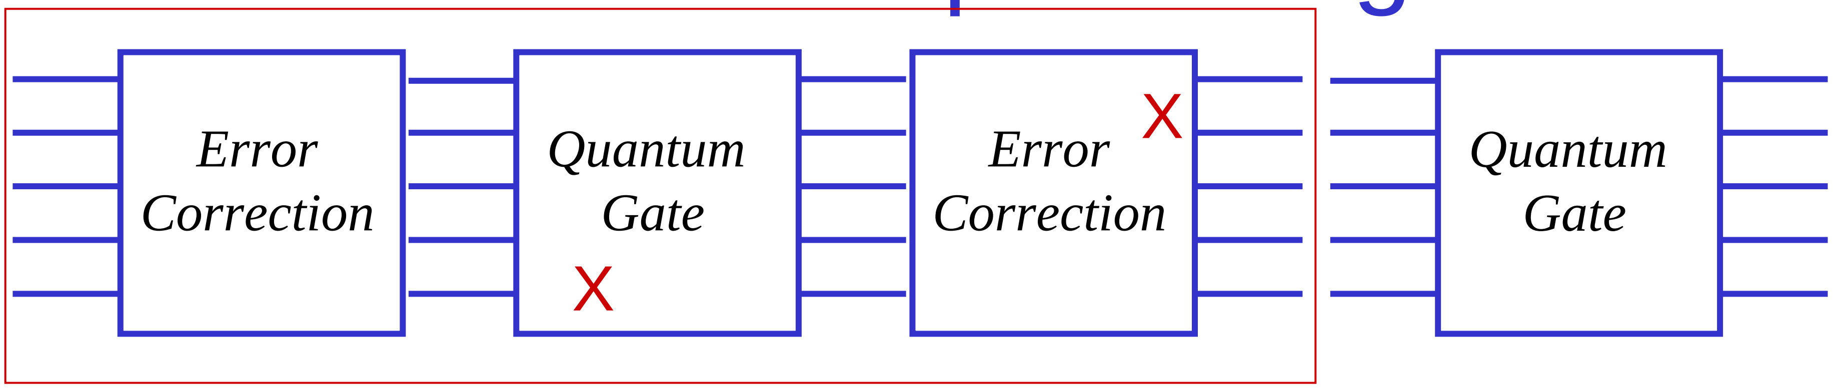

Each gate is followed by an error correction step. The circuit simulation fails only if two faults occur in some "extended rectangle."

If we simulate an ideal circuit with L quantum gates, and faults occur independently with probability ε at each circuit location, then the probability of failure is

$$P_{\text{fail}} \le L A_{\max} \varepsilon^2$$

where $A_{\max}$ is an upper bound on the number of pairs of circuit locations in each extended rectangle. Therefore, by using a quantum code that corrects one error and fault-tolerant quantum gates, we can improve the circuit size that can be simulated reliably to $L=O(\varepsilon^{-2})$, compared to $L=O(\varepsilon^{-1})$ for an unprotected quantum circuit.

Recursive simulation

In a fault-tolerant simulation, each (level-0) ideal gate is replaced by a *1-Rectangle*: a (level-1) gate gadget followed by (level-1) error correction on each output block. In a level-k simulation, this replacement is repeated k times --- the ideal gate is replaced by a k-*Rectangle*.

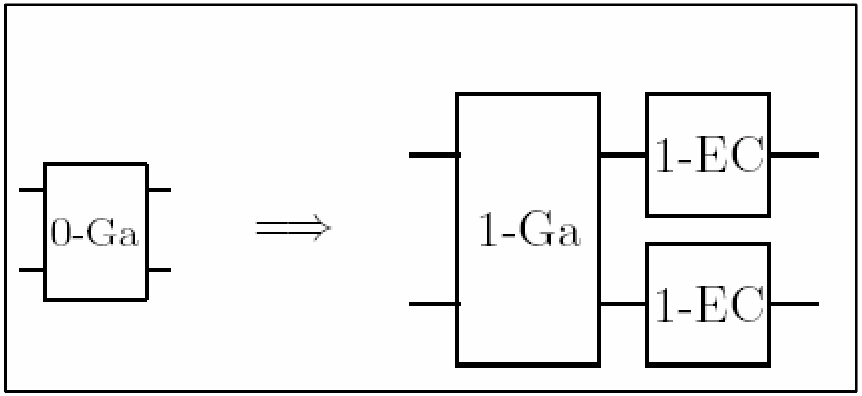

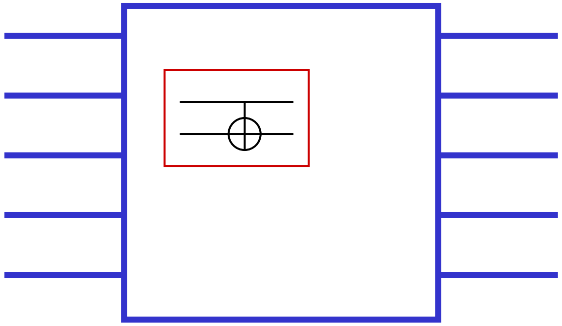

A *1-rectangle* is built from quantum gates.

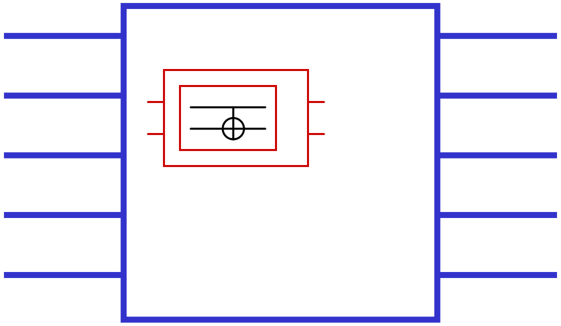

A *2-rectangle* is built from 1-rectangles.

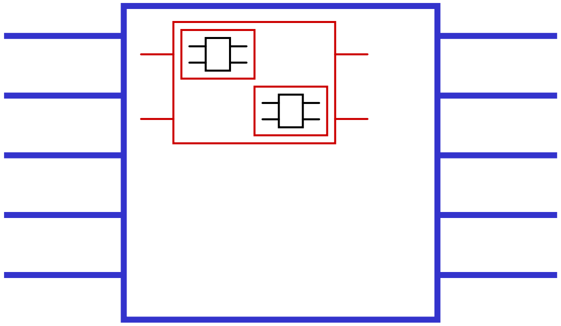

A *3-rectangle* is built from 2-rectangles.

(1) The computation is accurate if the faults in a level-k simulation are *sparse.*
(2) A non-sparse distribution of faults is *very* unlikely if the noise is *weak*.

There is *threshold of accuracy*. If the fault rate is below the threshold, then an arbitrarily long quantum computation can be executed with good reliability.

Level Reduction: “coarse-grained” computation

Simulated gate is *correct* if:

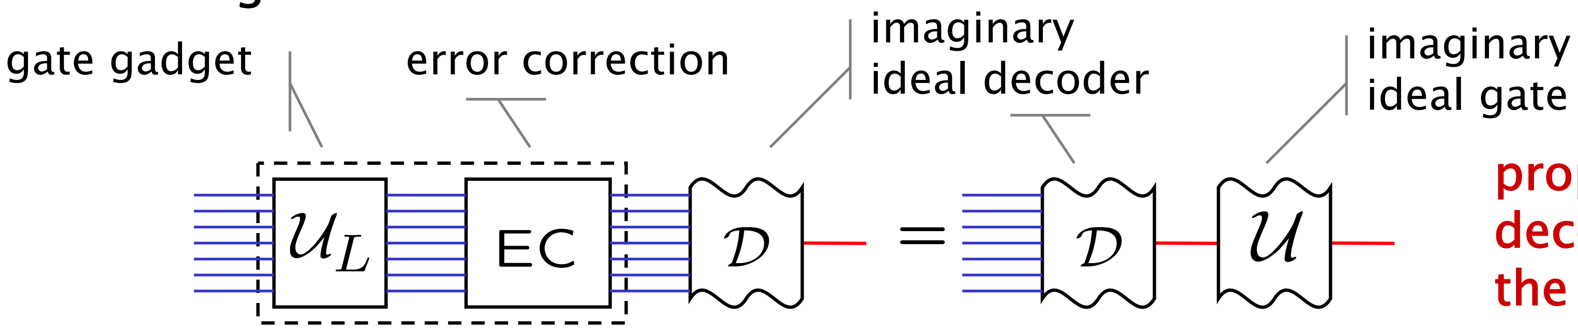

propagate decoders to the left

Simulated measurements and preparations are *correct* if:

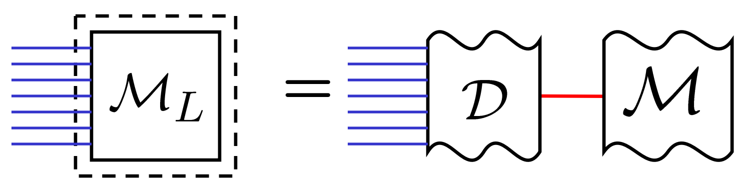

create decoders

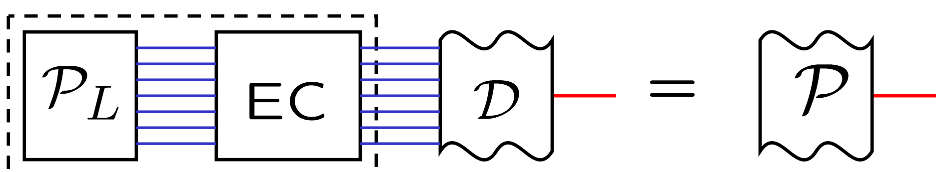

annihilate decoders

Decoders sweeping from right to left transform a level-1 computation to an equivalent level-0 computation. Each “good” level-1 extended rectangle (with no more than one fault) becomes an ideal level-0 gate, and each “bad” level-1 extended rectangle (with two or more faults) becomes a faulty level-0 gate. *If* our noise model is stable under level reduction, the coarse-graining can be repeated many times.

Local Stochastic Noise

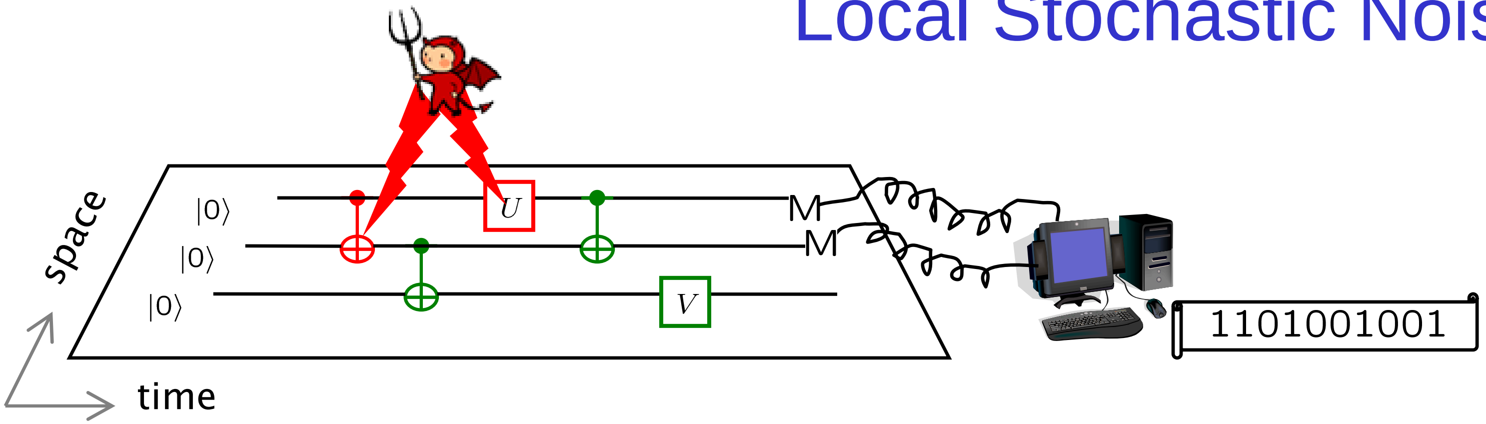

Noisy Circuit = Σ "Fault Paths"

For *local stochastic noise* with strength ε, the sum of the probabilities of all fault paths such that r specified gates are faulty is at most ε^r.

(For each fault path, the operations at the faulty locations are chosen by the adversary.)

After one level reduction step, the circuit is still subject to local stochastic noise with a "renormalized" strength:

$$\varepsilon^{(1)} \leq \varepsilon^2 / \varepsilon_0 = \varepsilon_0 (\varepsilon / \varepsilon_0)^2$$

The constant ε_0 is estimated by counting the number of "malignant" pairs of fault locations that can cause a 1-rectangle to be incorrect. If level reduction is repeated k times, the renormalized strength becomes:

$$\varepsilon^{(k)} < \varepsilon_0 (\varepsilon / \varepsilon_0)^{2^k}$$

Accuracy Threshold

Quantum Accuracy Threshold Theorem: Consider a quantum computer subject to local stochastic noise with strength ε. There exists a constant $\varepsilon_0 > 0$ such that for a fixed $\varepsilon < \varepsilon_0$ and fixed $\delta > 0$, any circuit of size L can be simulated by a circuit of size L^* with accuracy greater than 1-δ, where, for some constant c,

$$L^* = O\left[L (\log L)^c \right]$$

Aharonov, Ben-Or (1996)
Kitaev (1996)

The numerical value of the *accuracy threshold* ε_0 is of practical interest!

$$\varepsilon_0 > 2.73 \times 10^{-5}$$

Aliferis, Gottesman, Preskill (2005).

assuming:

parallelism, fresh ancillas (necessary assumptions)

nonlocal gates, fast measurements, fast and accurate classical processing, no leakage (convenient assumptions).

Four noteworthy developments

1) Improved thresholds with subsystem codes – Aliferis, Cross (2006)
2) Threshold for local gates in 2D – Svore, DiVincenzo, Terhal (2006)
3) Threshold when measurements are slow – DiVincenzo, Aliferis (2006)
4) Threshold for postselected computation – Reichardt (2006), Aliferis, Gottesman, Preskill (2007)

Three questions

1) Threshold in terms of noise power spectrum?
2) Threshold for asymmetric noise?
3) Self-correcting quantum memory (finite-temperature topological order)?

Subsystem codes

Hilbert space decomposes as:

$$\mathcal{H} = \bigoplus_s \mathcal{H}_{\mathcal{L}}^{(s)} \otimes \mathcal{H}_{\mathcal{T}}^{(s)}$$

syndrome → s

"logical" subsystem → $\mathcal{H}_{\mathcal{L}}^{(s)}$

"gauge" subsystem → $\mathcal{H}_{\mathcal{T}}^{(s)}$

- A subsystem code becomes a standard stabilizer code when the gauge subsystem is trivial (e.g., if we "fix the gauge").
- But there is no need to fix the gauge, as errors acting on gauge qubits do not damage the protected information.
- Maintaining the gauge freedom reduces the number of check operators.
- Syndrome information can be extracted by measuring the gauge qubits, and for some codes the gauge-qubit operators have lower weight than the stabilizer generators, so it is easier to measure the gauge operators fault tolerantly.

3 × 3 "Bacon-Shor code"

Shor (1995)
Bacon (2005)

check operators:

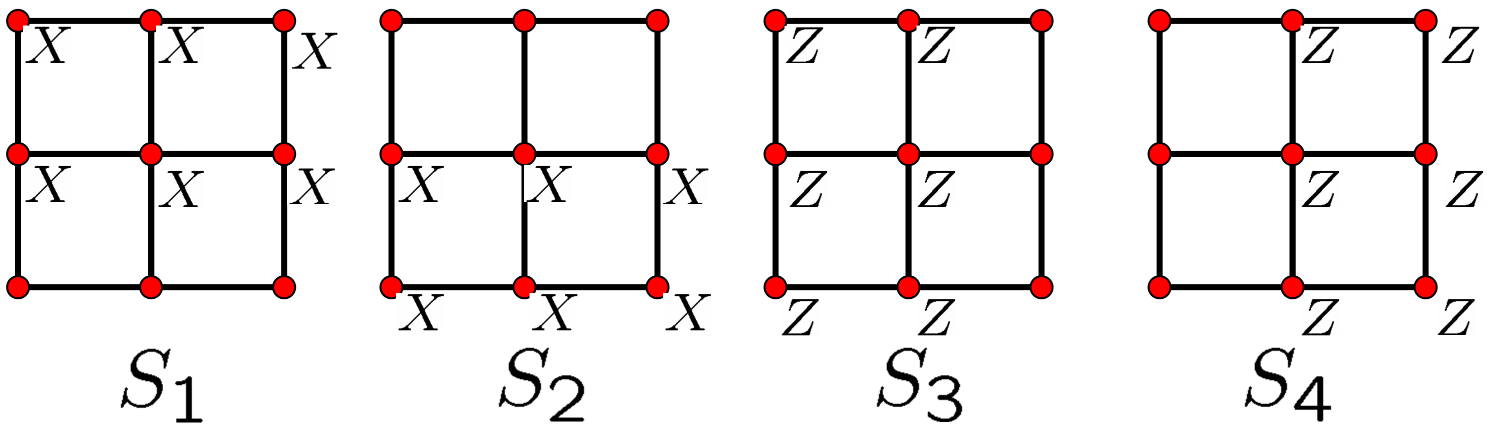

logical operators:

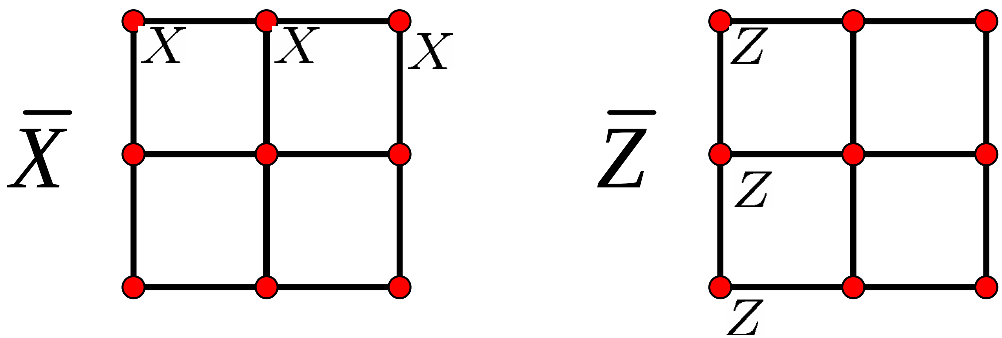

corrects one error

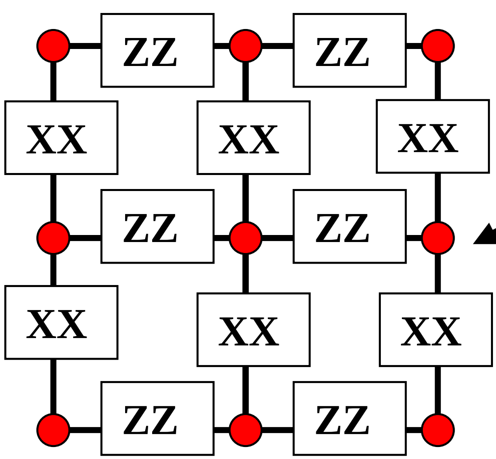

These weight-two *gauge* Pauli operators commute with the logical operations, and measuring them determines the check operators in the stabilizer. Because only weight-two operators are measured, error correction is efficient and easily made fault tolerant.

3 × 3 "Bacon-Shor code"

Shor (1995)
Bacon (2005)

check operators:

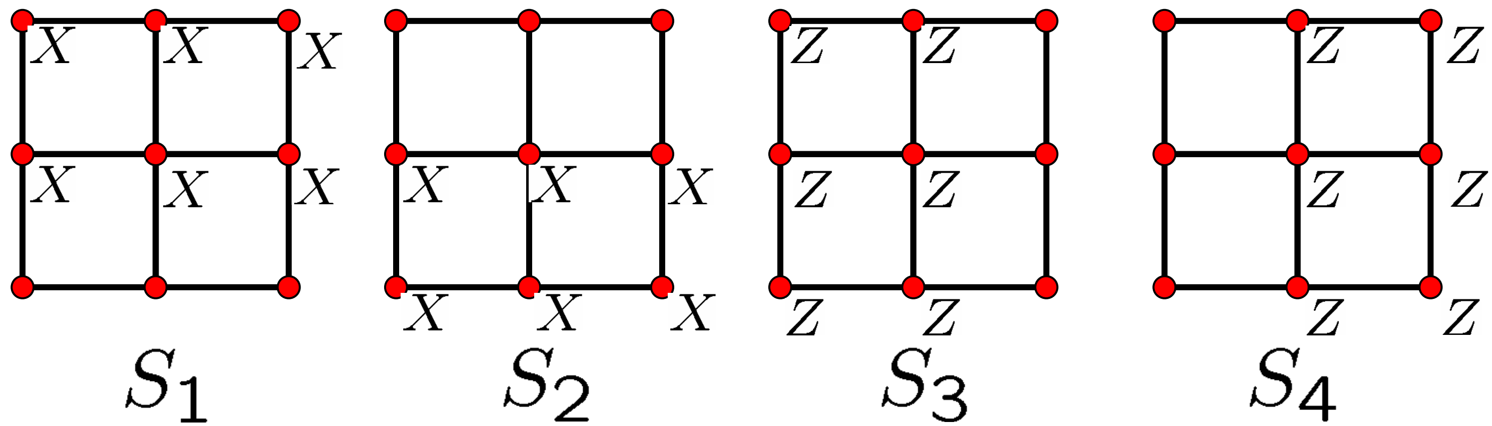

logical operators:

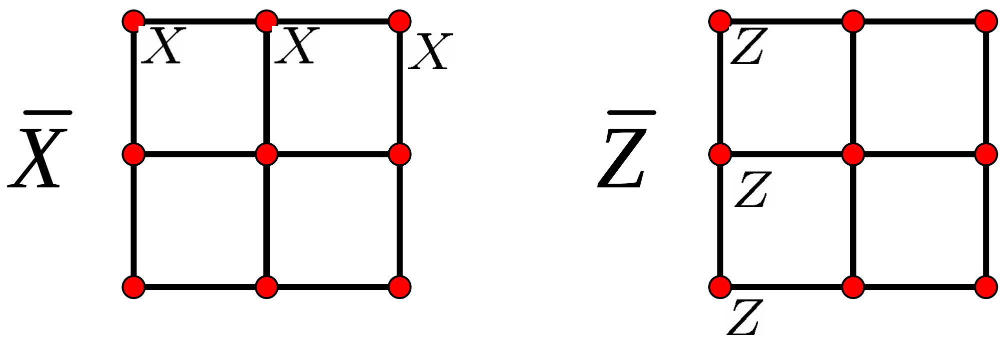

corrects one error

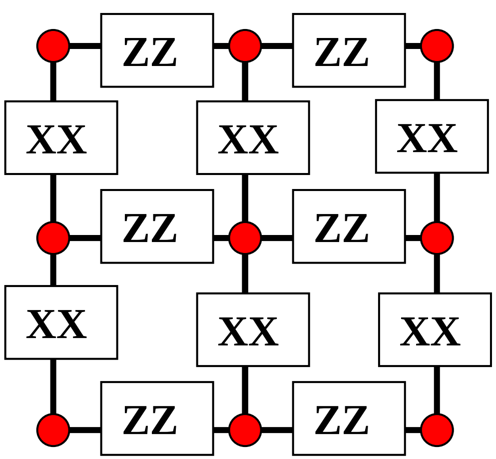

The optimal threshold estimate is found using the 5 X 5 Bacon-Shor code (which corrects two errors):

$$\varepsilon_0 \geq 1.9 \times 10^{-4}$$

Aliferis, Cross (2006)

Logical Qubit in a 2D Lattice

Svore, DiVincenzo, Terhal (2006)

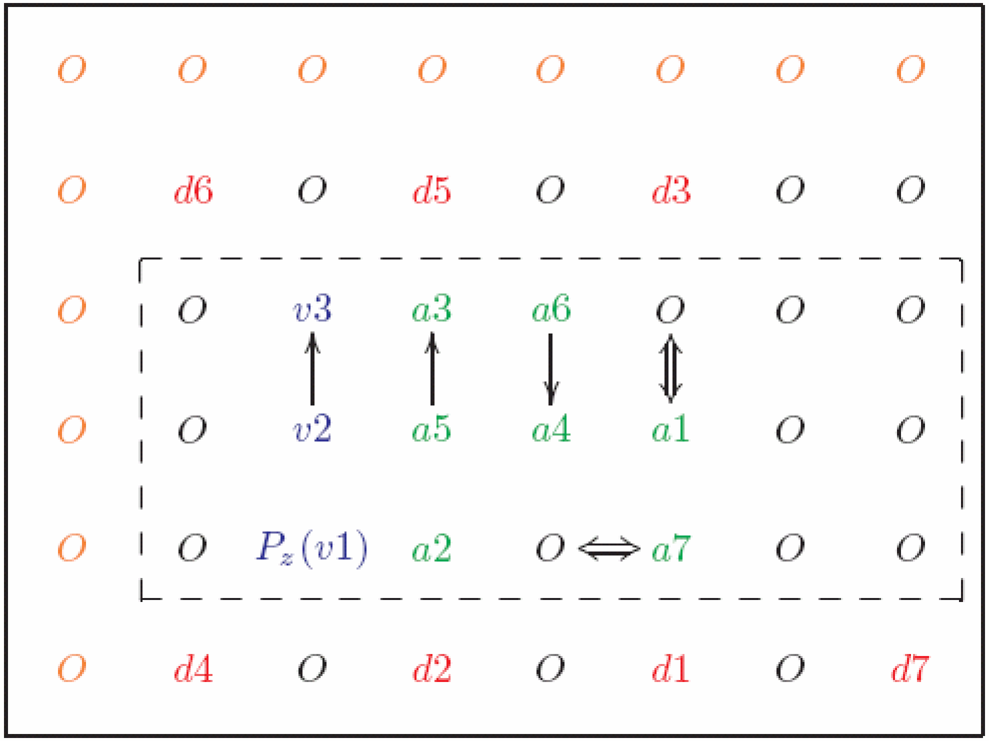

One logical qubit is stored in a 6 X 8 lattice cell. The cell contains a 7-qubit code block (*d*) , a 7-qubit ancilla block (*a*), and 3 qubits (*v*) that are used for ancilla verification; the rest (*O*) are "dummy qubits" that serve as communication channels. We include a (noisy) swap in our gate set at each level of the recursive hierarchy. There are 61 time steps in the CNOT extended rectangle.

Logical Qubit in a 2D Lattice

Svore, DiVincenzo, Terhal (2006)

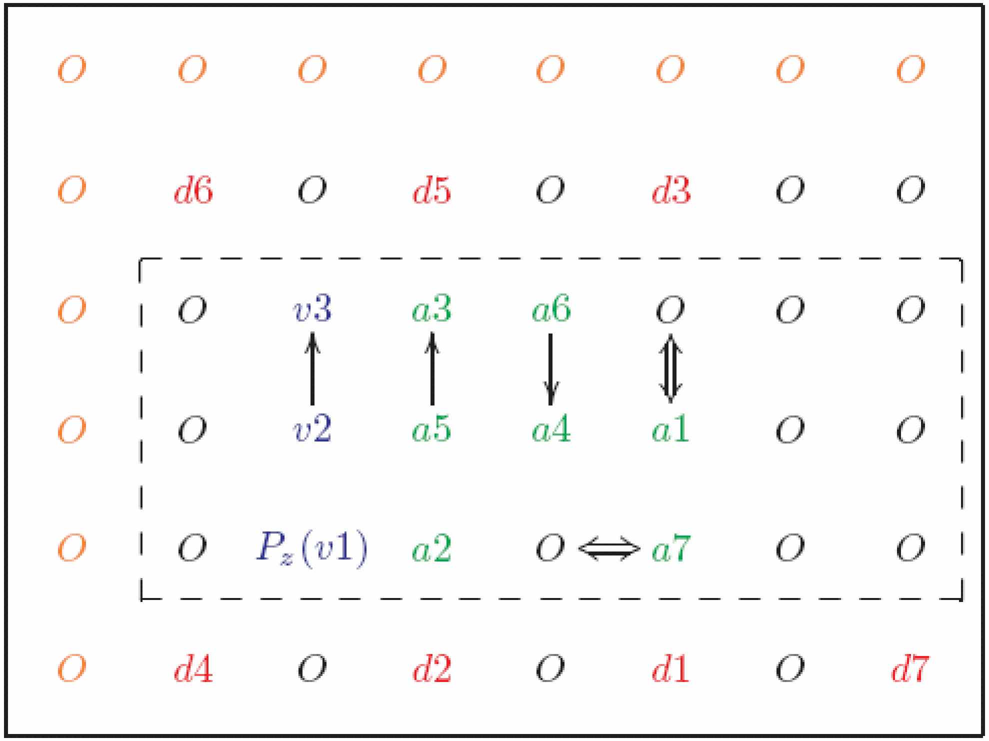

When the gates are required to be local rather than nonlocal, the threshold worsens from 3.61×10^{-5} to 1.85×10^{-5}, assuming that the storage fault rate is 1/10 the fault rate for gates, measurements, and preparations. For a two-rung "qubit ladder", Stephens, Fowler, Hollenberg (2007) find 1.97×10^{-6}. Further improvements may be achieved with subsystem codes.

Fault tolerance with slow measurements

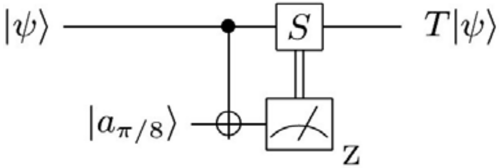

In some systems (e.g., spins in quantum dots) measurements take much longer than gates. Yet fast measurements are desirable because:

1) Measurements extract the error syndrome (the measurements can be done "coherently" but the threshold suffers).
2) Measurements verify ancillas used for error correction.
3) Measurements allow "teleportation" of gates that are needed to complete a universal fault-tolerant gate set.

But …

1) Don't wait for the syndrome, or apply recovery operations. The syndrome, once known, can be propagated through subsequent gates by an efficient classical computation.
2) Decode the ancilla, measure it eventually, and infer encoded errors that propagated through the circuit.
3) Teleport only at high levels in the recursive hierarchy, where encoded gates take as long as measurements.

The threshold is little effected even if measurements take ~ 1000 times longer than gates. -- DiVincenzo and Aliferis (2006)

Accuracy threshold using error-*detecting* codes

Using Bacon-Shor codes, we obtain a lower bound on the accuracy threshold (for *adversarial independent stochastic noise, nonlocal gates*)

$\varepsilon_0 > 1.9 \times 10^{-4}$

We can improve the threshold further if we can simulate gates with $\varepsilon_{\text{eff}} < \varepsilon_0$ using gates with $\varepsilon > \varepsilon_0$.

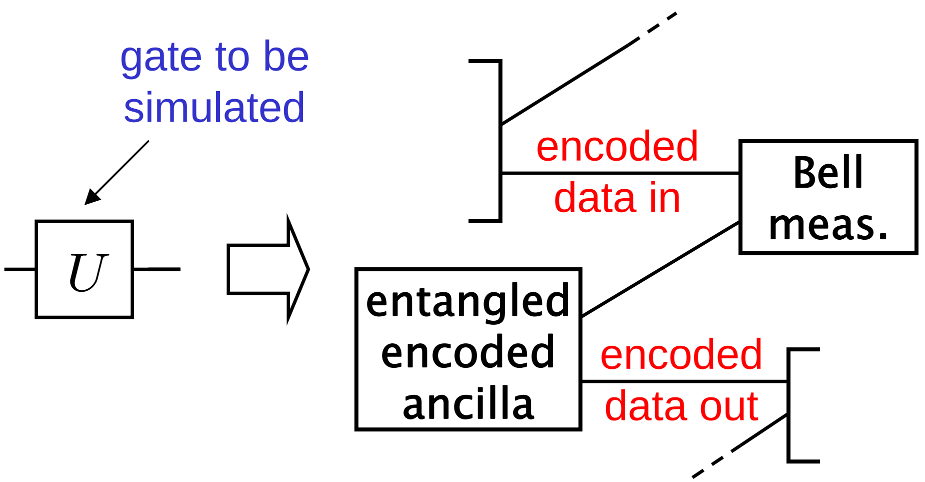

Knill's idea (2004):

Prepare suitable ancillas offline and teleport gates. Encoded error rate $\varepsilon_{\text{eff}} < \varepsilon_0$ can be achieved if the errors in the ancilla are *nearly independent* and have error rate below e.g. 5%.

Protect the ancilla-preparation circuit using a (recursive) error-*detecting* code and accept the ancilla only if no errors are detected. Errors occurring during decoding are independent.

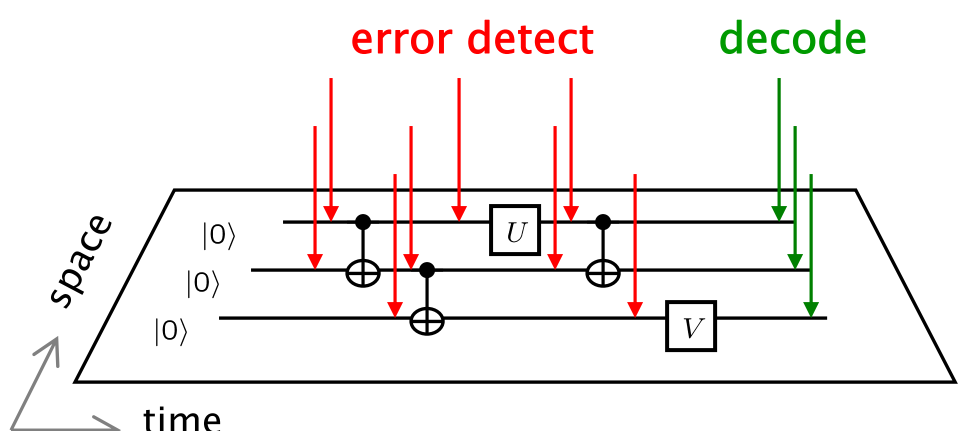

Fault-tolerant gadgets

If we use a (distance-three) quantum error-*correcting* code:

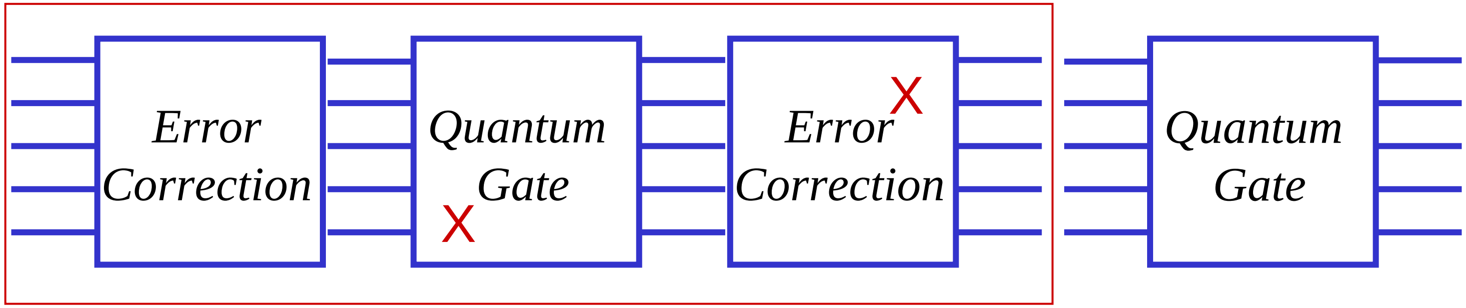

Each gate is followed by an error correction step. The circuit simulation fails only if two faults occur in some "extended rectangle." If faults occur with probability ε, then the gadget fails with prob $O(\varepsilon^2)$.

If we use a (distance-two) quantum error-*detecting* code:

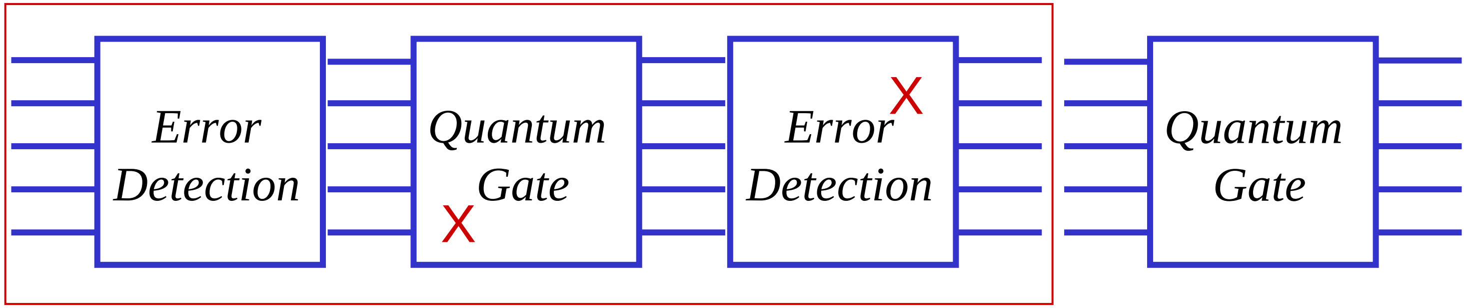

Each gate is followed by an error *detection* step, and the computation is aborted if an error is detected. The circuit simulation fails only if two faults occur in some "extended rectangle." If faults occur with probability ε, then the gadget fails with prob $O(\varepsilon^2)$.

Threshold for postselected quantum computation

We can boost the reliability by building a hierarchy of gadgets within gadgets --- the fault-tolerant circuit simulates the ideal circuit if the faults are *sparse*.

However … to assess the reliability of the postselected circuit, we must estimate the probability that it fails conditioned on *global* acceptance --- i.e., acceptance by every error detection in the entire circuit.

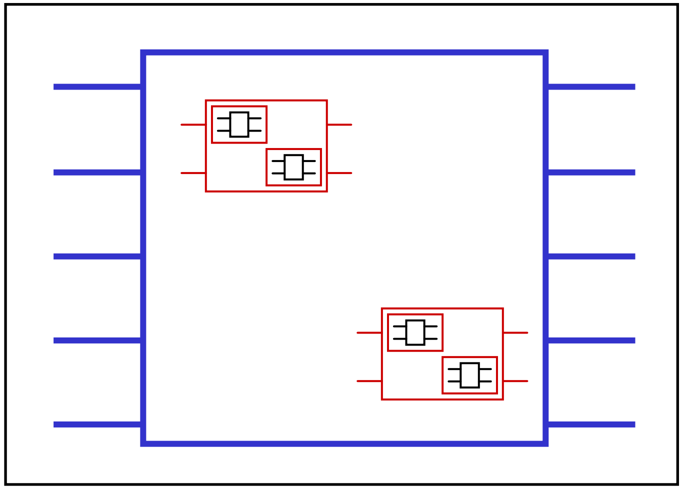

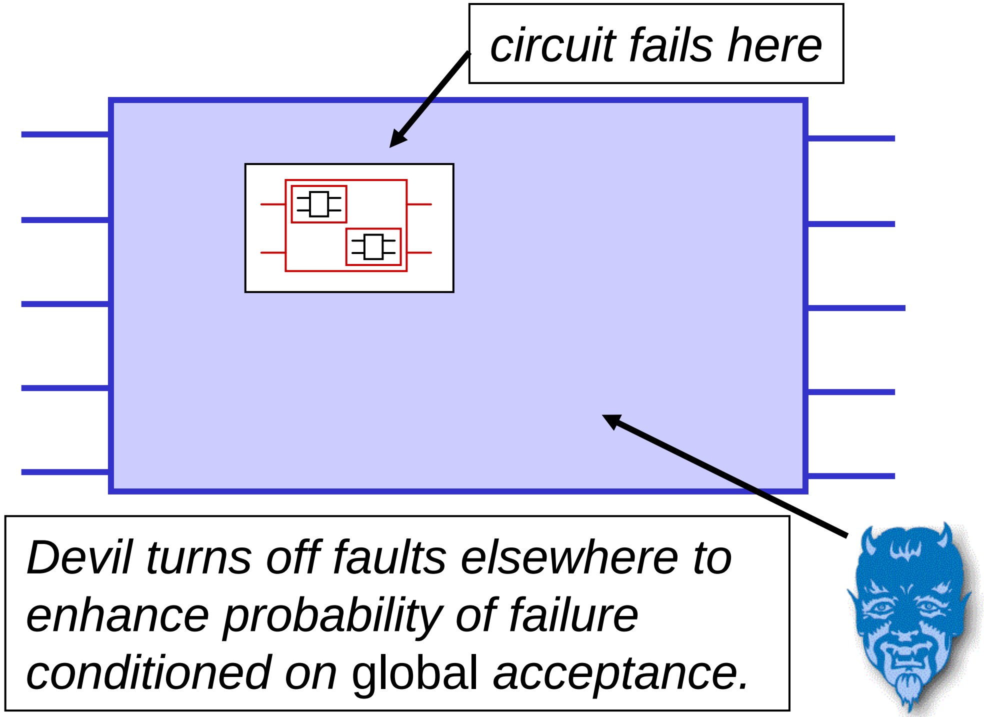

To obtain a threshold theorem for postselected computation, we must disallow correlations in the noise that could be tolerated if error correction were used instead. Otherwise, the devil could enhance greatly the conditional probability of failure in one part of the circuit by *turning off* faults elsewhere.

Threshold for postselected quantum computation

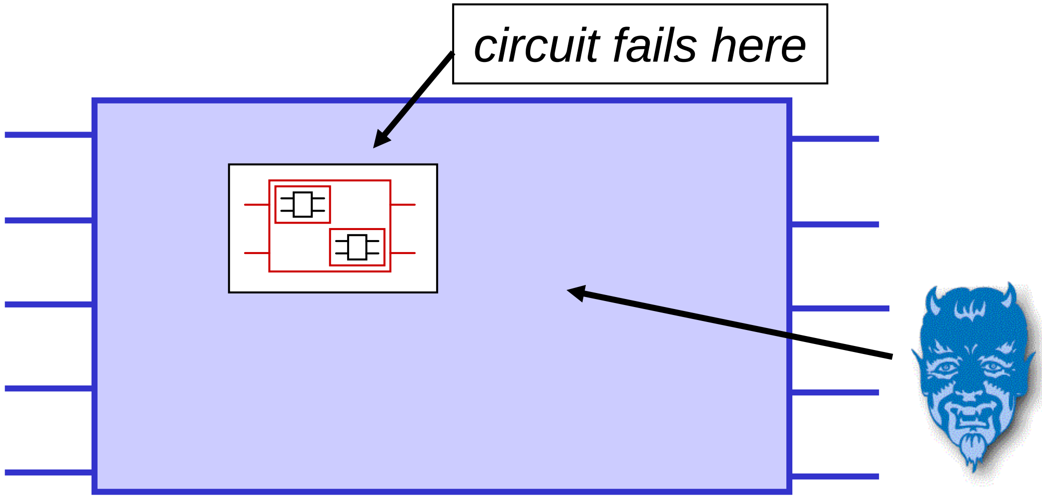

Devil turns off faults elsewhere to enhance probability of failure conditioned on global *acceptance.*

We need a noise model that

a) Limits the adversary's global control.
b) Is stable under level reduction.

Local stochastic noise has (b) but not (a). Independent noise has (a) but not (b).

in between is *locally correlated* stochastic noise:

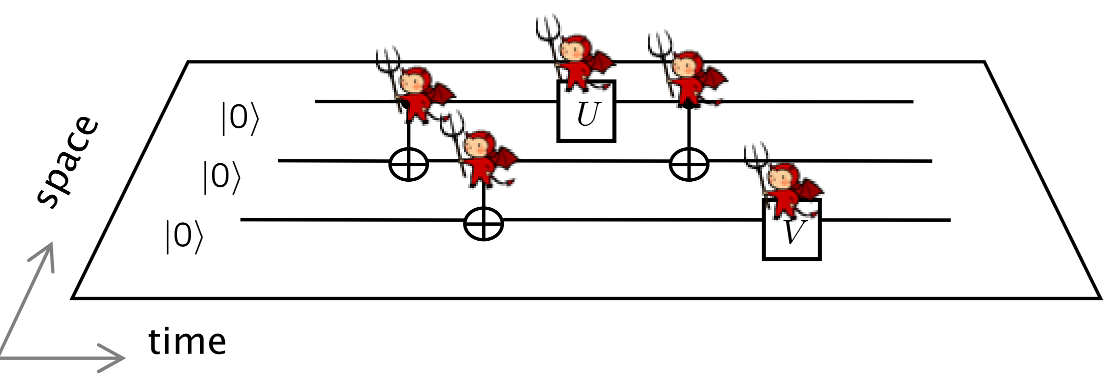

- different adversaries control each noisy operation,
- adversaries can communicate only "locally,"
- messages are erased by ideal gates.

Threshold for postselected quantum computation

The bad gadgets in the postselected circuit form connected clusters, surrounded by error detections with no faults. Thus the clusters (which typically contain just one or a small number of bad gadgets) are isolated from one another, enabling us to relate the probability of failure of a gadget conditioned on local acceptance (within the cluster) to its probability of failure conditioned on global acceptance. This means that error detection and (global) postselection improves reliability, and we can show by an inductive step that the probability of failure in a recursive simulation gets arbitrarily small if the noise is sufficiently weak..

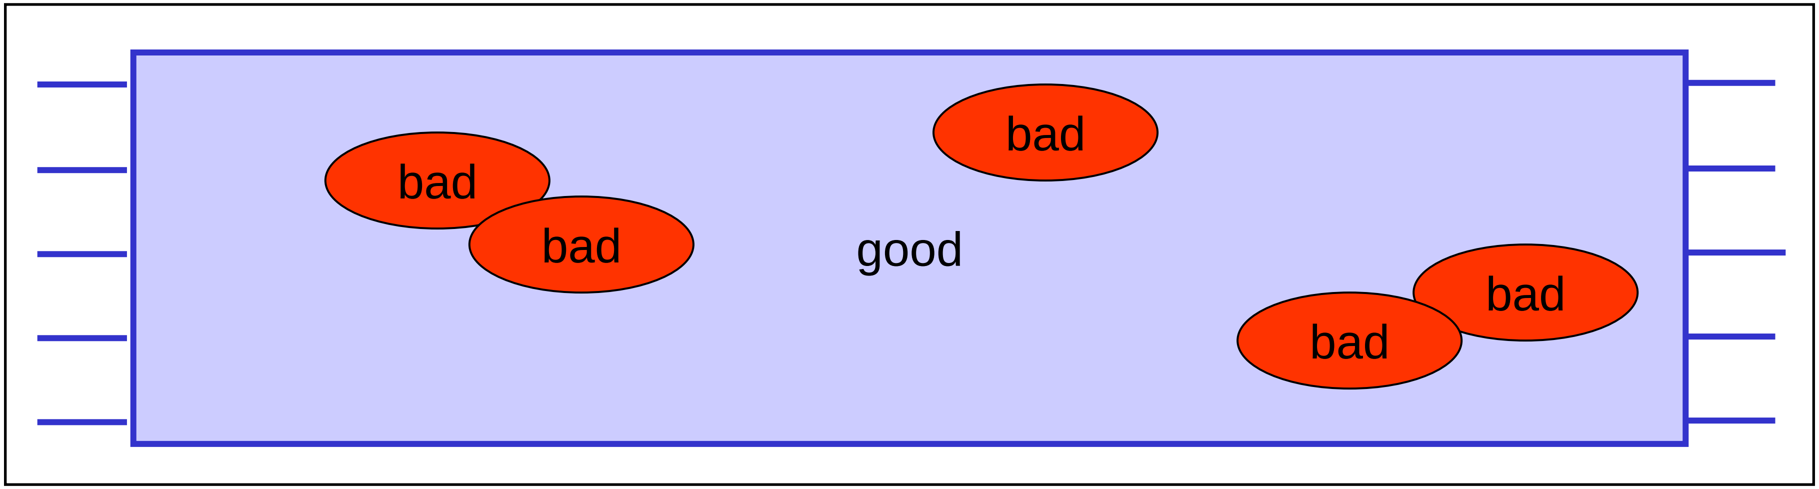

Counting the ways for error-detecting gadgets to fail, we find $\varepsilon_{0,\mathrm{ED}} > 1.04 \times 10^{-3}$ (Aliferis-Gottesman-Preskill 2007, Reichardt 2006). This is the best rigorously established lower bound on the accuracy threshold so far, but still a factor of 30 below Knill's estimate based on simulations. Note that the overhead cost of postselected simulation may be prohibitive for ε close to $\varepsilon_{0,\mathrm{ED}}$ (but acceptable for $\varepsilon > \varepsilon_0$).

Local non-Markovian noise

Terhal, Burkard (2004)
Aliferis, Gottesman, Preskill (2005)
Aharonov, Kitaev, Preskill (2005)

From a physics perspective, it is natural to formulate the noise model in terms of a Hamiltonian that couples the system to the environment.

Non-Markovian noise with a *nonlocal* bath.

$$H = H_{System} + H_{Bath} + H_{System-Bath}$$

where

$$H_{System-Bath} = \sum_{\substack{\text{terms } a \text{ acting} \\ \text{locally on the system}}} H^{(a)}_{System-Bath}$$

Then

$$U_{SB} = \sum \text{“}Fault\ Paths\text{”}$$

For *local noise* with strength ε, the norm of the sum of all fault paths such that r specified gates are faulty is at most ε^r.

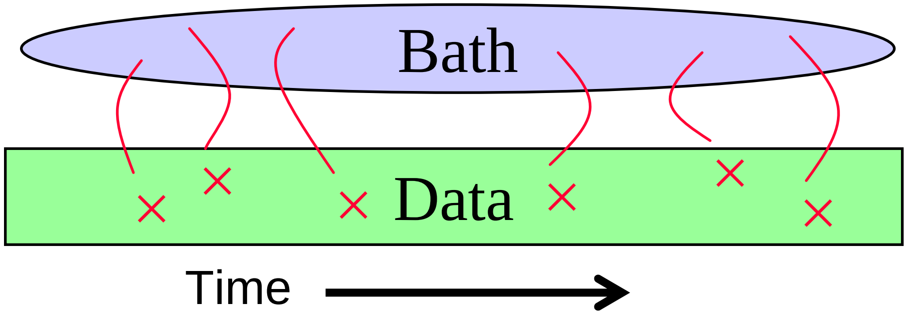

$$\varepsilon = \max \left\| H^{(a)}_{System-Bath} \right\| t_0$$

max: over all times and locations

t_0: time to execute a gate

Fault-tolerant recursive simulation

Non-Markovian noise with a *nonlocal* bath.

$$H = H_{System} + H_{Bath} + H_{System-Bath}$$

Quantum error correction works as long as the coupling of the system to the bath is *local* (only a few system qubits are jointly coupled to the bath) and *weak* (sum of terms, each with a small norm). Arbitrary (nonlocal) couplings among the bath degrees of freedom are allowed.

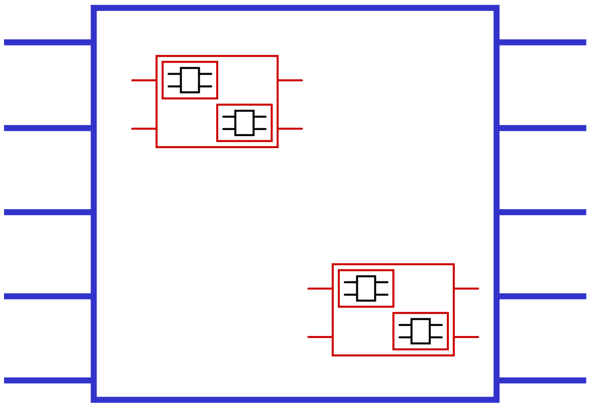

A hierarchy of “gadgets within gadgets” is reliable if the faults are *sparse*.

We find a rigorous upper bound on the norm of the sum of all “bad” diagrams (such that the faults are *not* sparsely distributed in spacetime). Actually, this works even for interactions among the system qubits that decay algebraically with distance…

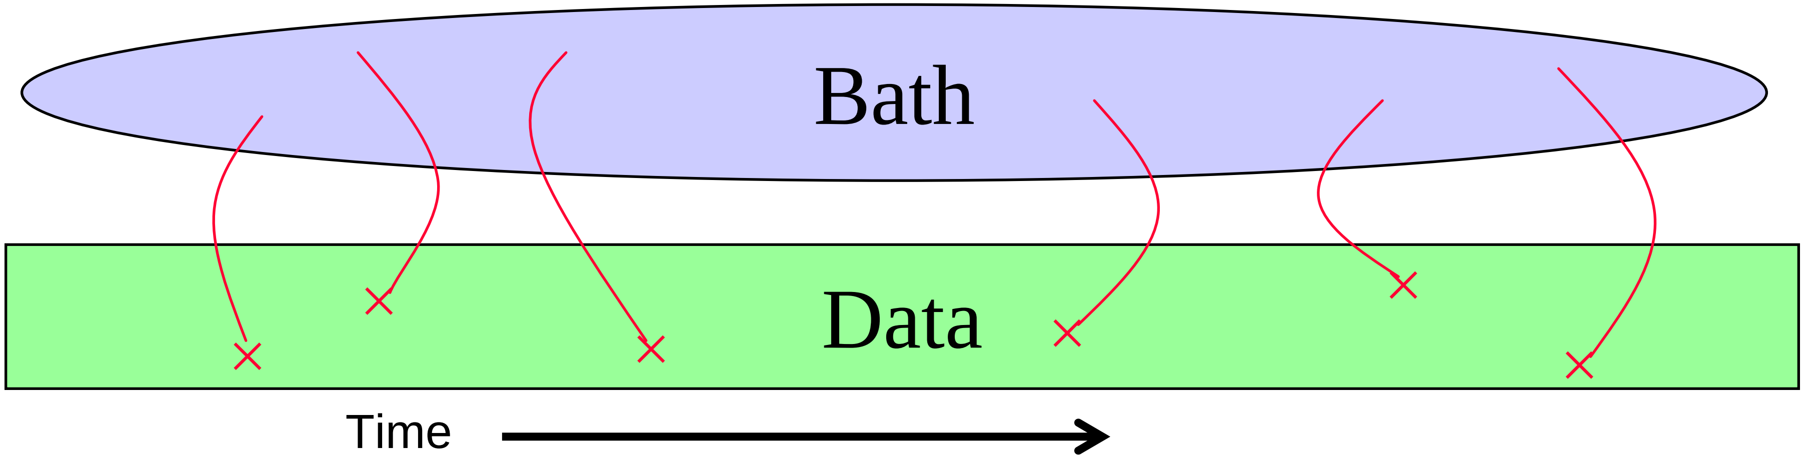

Local non-Markovian noise

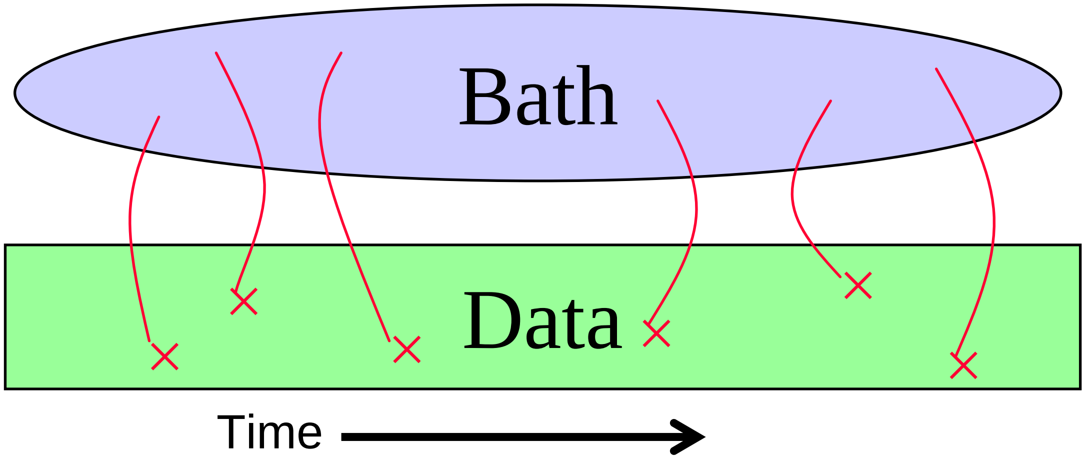

$$\left\| H^{(a)}_{System-Bath} \right\| t_0 < \varepsilon_0$$

However, expressing the threshold condition in terms of the norm of the system-bath coupling has disadvantages.

E.g., this noise strength is not directly measurable in experiments, and furthermore in the case of a bath of harmonic oscillators, the norm is *infinite*.

$$H_B + H_{SB} = \sum_k \frac{1}{2}\omega_k a_k^\dagger a_k + \frac{1}{2}\sigma_z \left(\sum_k g_k a_k + g_k^* a_k^\dagger \right) \qquad \sum_k |g_k|^2 \approx \int_0^\infty d\omega J(\omega)$$

It would be more natural, and more broadly applicable, if we could express the threshold condition in terms of the *correlation functions* of the bath.

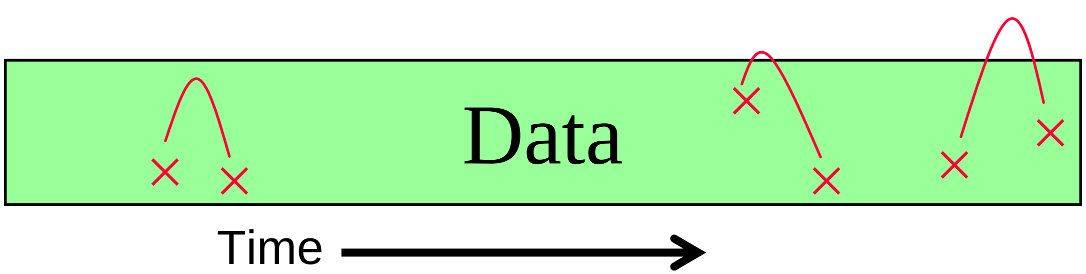

Local non-Markovian noise

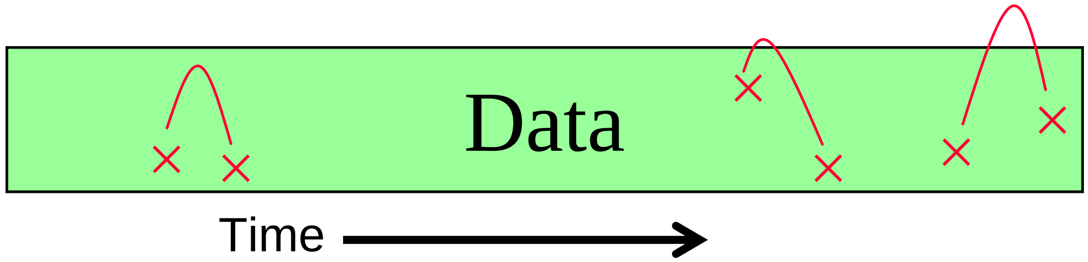

$$\left\| H^{(a)}_{System-Bath} \right\| t_0 < \varepsilon_0$$

The norm condition constrains the very-high-frequency fluctuations of the bath (the time-correlators at very short times). Intuitively, fluctuations with a time scale much shorter than the time it takes to execute a quantum gate should average out.

The threshold condition should be formulated in terms of an effective description of the noise, with high frequencies integrated out. If expressed in terms of e.g. the power spectrum of the noise, this criterion could be more directly applied to real systems (and to e.g. the spin-boson model, where the norm condition is not useful).

(Low-frequency noise, on the other hand, can be addressed with other methods, such as spin echoes, composite pulses, “decoherence-free subsystems” …)

Example: 1D Ising model (repetition code)

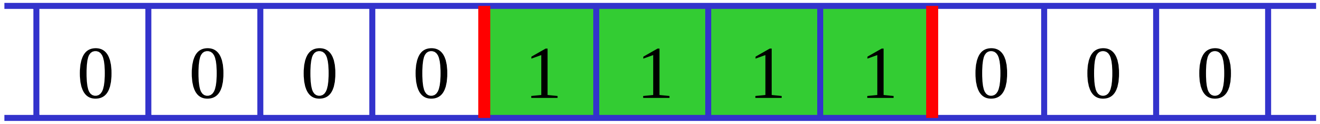

When a connected (one-dimensional) droplet of flipped spins arises due to a thermal fluctuation, only the (zero-dimensional) boundary of the droplet contributes to the energy; thus the energy cost is independent of the size of the droplet.

Therefore, thermal fluctuations disorder the spins at any nonzero termperature. A one-dimensional ferromagnet is not a robust (classical) memory.

2D Ising model (repetition code)

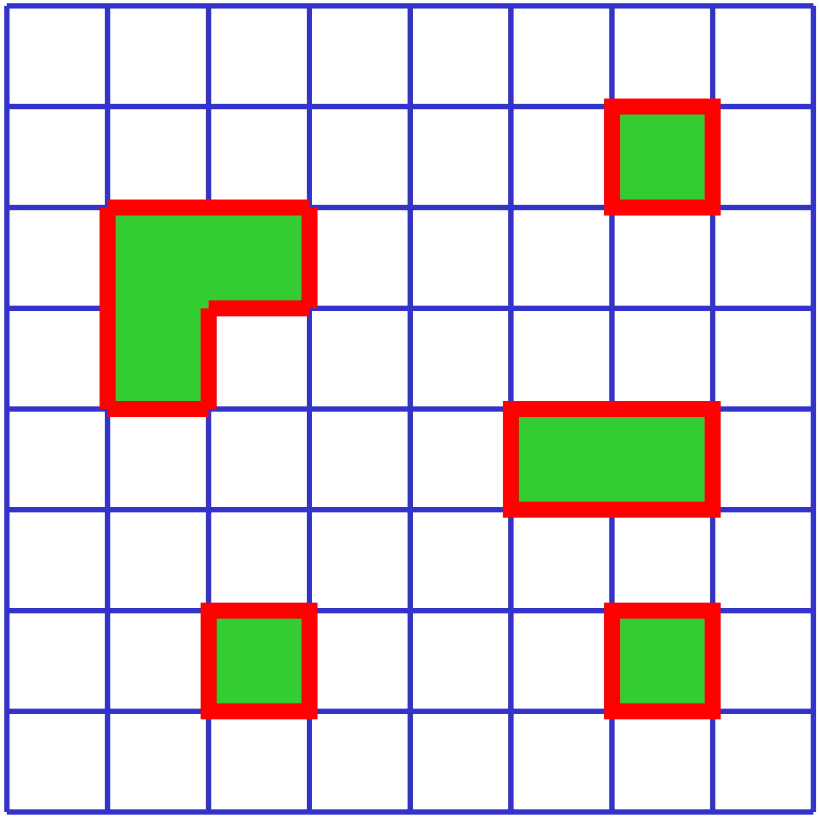

This memory is a repetition code, but with redundant (hence robust) parity checks.

Again, droplets of flipped spins arise as thermal fluctuations. But now the energy cost of a (two-dimensional) droplet is proportional to the length of its (one-dimensional) boundary.

Therefore, droplets with linear size L are suppressed at sufficiently low nonzero temperature T by the Boltzman factor $\exp(-L / T)$, and are rare.

The probability of a memory error becomes exponentially small when the block size is large. (Actual storage media, which are robust at room temperature, rely on this physical principle.)

Toric Code

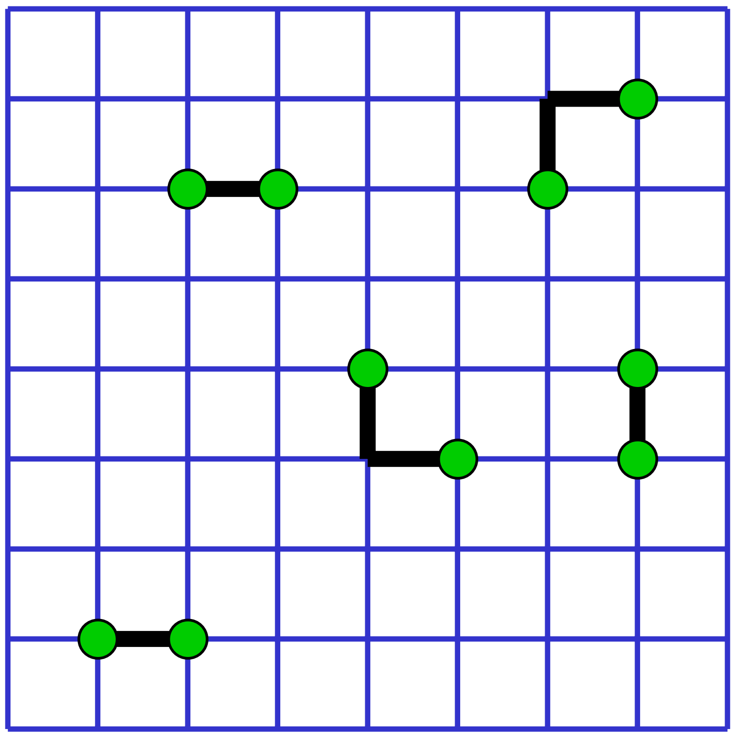

When a connected (one-dimensional) chain of flipped qubits arises due to a thermal fluctuation, only the (zero-dimensional) boundary of the chain contributes to the energy; thus the energy cost is independent of the length of the chain.

Therefore, thermal fluctuations disorder the system at any nonzero termperature. A two-dimensional topological medium is not a robust quantum memory.

On the other hand, if we continuously observe the defect gas, then at low temperature the chain segments are typically short, and the defect positions are strongly correlated. Therefore, it is easy to guess how to "pair" the defects and to infer when encoded errors occur. The probability of an encoded error is $e^{-O(L)}$ on an L X L torus if in each round, the probability of a qubit error and of a syndrome measurement error are both below 3%. Dennis, Kitaev, Landahl, Preskill (2001).

Toric code in *four* dimensions

-- Qubits are on *plaquettes* (2-cells):

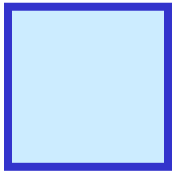

-- 6-qubit X and Z check operators at *edges* and *cubes* (dual links):

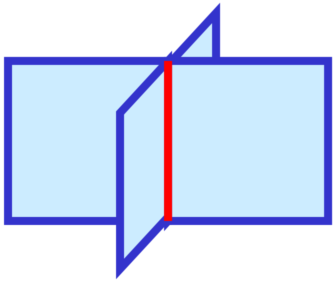

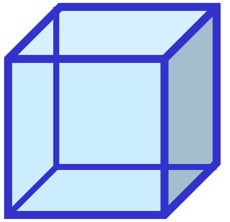

-- Logical operations: homologically nontrivial 2-surfaces of lattice and dual lattice:

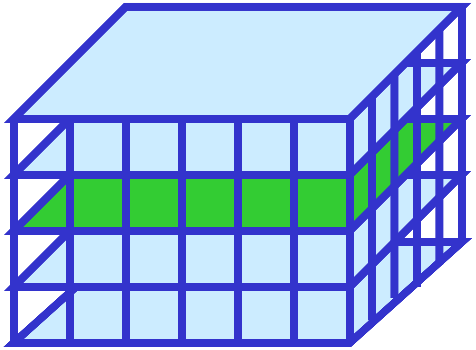

-- Defects are closed *loops of string,* (or dual loops) which bound droplets of flipped qubits:

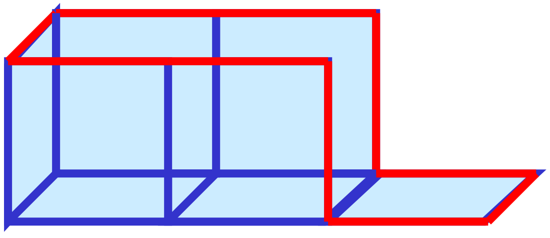

Topological order at finite temperature

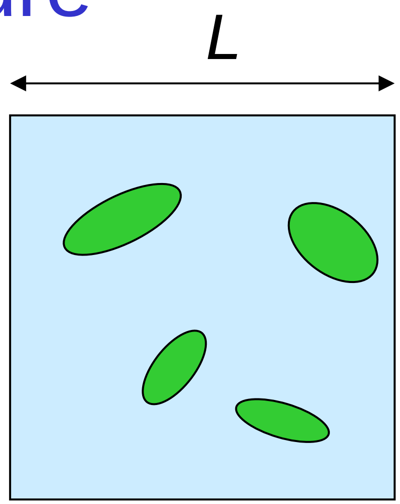

In the 4D toric code, the energy cost of a 2D droplet of flipped qubits is proportional to the length of its 1D boundary.

To cause encoded errors, Droplets of linear size *L*, which could cause encoded errors, are suppressed at sufficiently low nonzero temperature *T* by the Boltzman factor exp(-*L* / *T*), and are rare.

Question: Is “finite-temperature topological order” possible in 3D?

In the 3D toric code, we can choose to have *point* defects at the boundary of 1D bit-flip error chains and *string* defects at the boundary of 2D phase-error droplets, or the other way around.

Absence of an obvious exactly solvable model (corresponding to an RG fixed point) makes one suspect that robust 3D topological memory is not possible.

But what about the 3D compass model. (Note that a gap may not be necessary --- Cf. Bacon 2005.)

Four noteworthy developments

1) Improved thresholds with subsystem codes – Aliferis, Cross (2006)
2) Threshold for local gates in 2D – Svore, DiVincenzo, Terhal (2006)
3) Threshold when measurements are slow – DiVincenzo, Aliferis (2006)
4) Threshold for postselected computation – Reichardt (2006), Aliferis, Gottesman, Preskill (2007)

Three questions

1) Threshold in terms of noise power spectrum?
2) Threshold for asymmetric noise?
3) Self-correcting quantum memory (finite-temperature topological order)?

Robust quantum computation

vs.

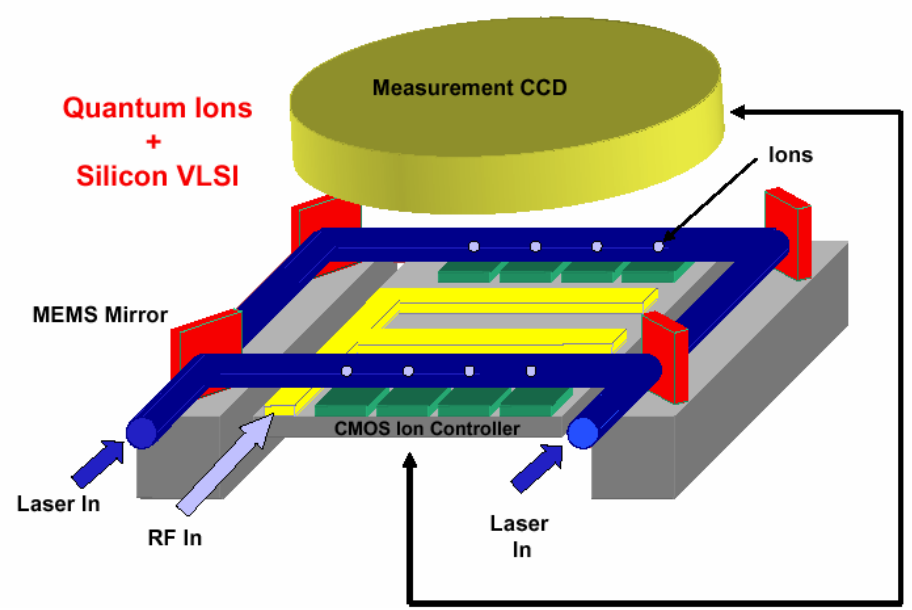

Institute for Quantum Information

Sponsored by the NSF

John Preskill, Caltech
27 February 2007
http://www.iqi.caltech.edu/