Fault-tolerant quantum computation versus realistic noise

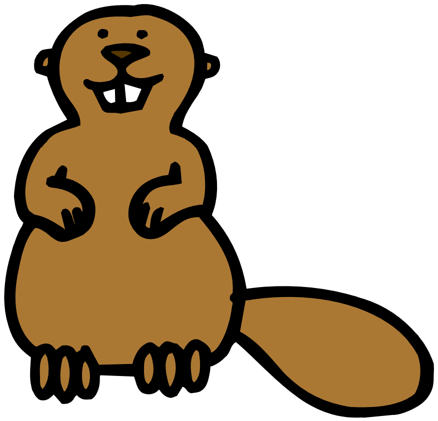

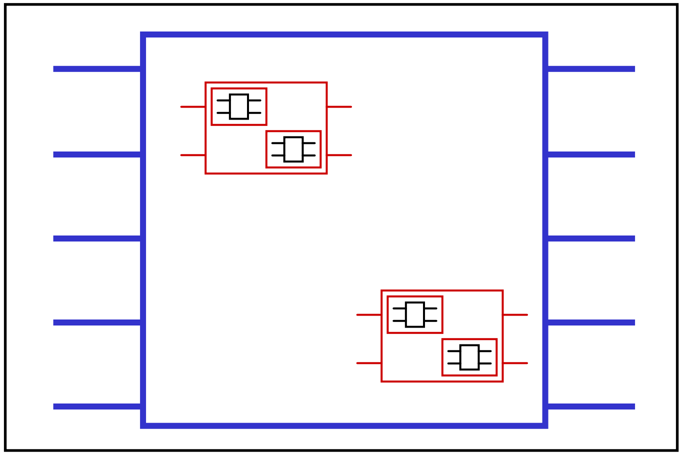

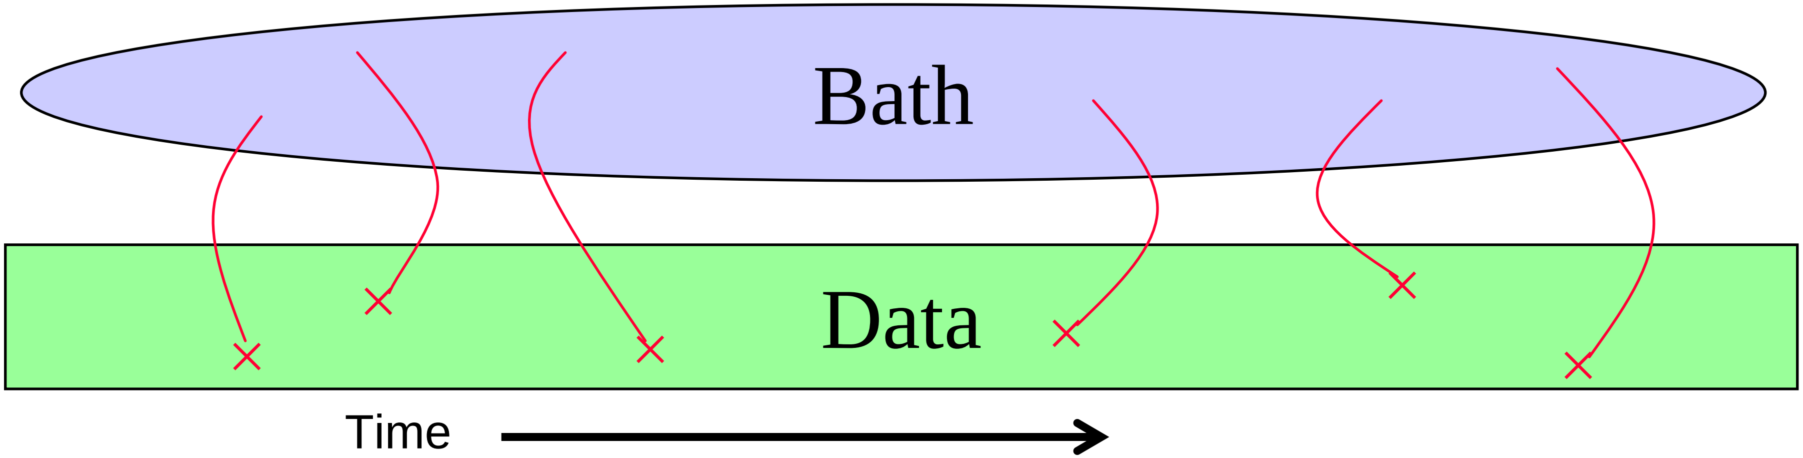

Institute for Quantum Information
Sponsored by the NSF

John Preskill, Caltech
24 March 2008

Quantum error correction and fault tolerance

Can large-scale quantum computers really be built and operated? Surely there are daunting technical challenges to be overcome. But are there obstacles *in principle* that might prevent us from ever attacking hard computational problems with quantum computers?

What comes to mind in particular is the problem of *errors.* Quantum computers will be far more susceptible to error than conventional digital computers. A particular challenge is to prevent *decoherence* due to interactions of the computer with the environment. Even aside from decoherence, the unitary quantum gates will not be perfect, and small imperfections will accumulate over time...

Our confidence that large-scale quantum computations will someday be possible has been bolstered by developments in the theory of quantum error correction and fault tolerance.

Quantum error correction and fault tolerance

1. Modeling errors and error correction
2. Fault-tolerant quantum computation
3. Quantum accuracy threshold theorem
4. Biased noise (Aliferis-Preskill 2007)
5. Non-Markovian (Gaussian) noise.

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.

The goal of fault-tolerant quantum computing 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:

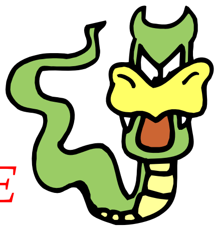

$$U : |\psi\rangle \otimes |0\rangle_E \to \sum_a E_a |\psi\rangle \otimes |a\rangle_E$$

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 strong decoherence errors if the states of the environment are mutually orthogonal.

Errors

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

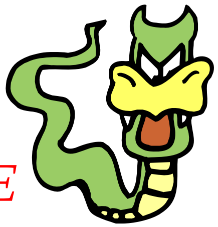

The objective of quantum error correction is to recover the (unknown) state $|\psi\rangle$ of the quantum computer. We can't expect to succeed for arbitrary errors, but we might succeed if the errors are of a restricted type. In fact, since the interactions with the environment are *local*, it is reasonable 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 weakly correlated (and rare), 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*:

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

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

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

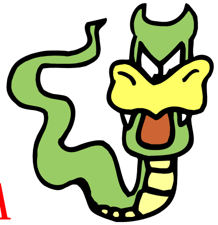

$$|\psi\rangle \otimes |0\rangle_E \otimes |0\rangle_A \xrightarrow{\text{error}} \sum_a E_a |\psi\rangle \otimes |a\rangle_E \otimes |0\rangle_A$$

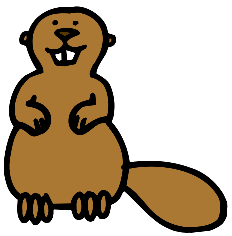

$$\xrightarrow{\text{recover}} \sum_a |\psi\rangle \otimes |a\rangle_A \otimes |a\rangle_E = |\psi\rangle \otimes |\varphi\rangle_{EA}$$

Entropy

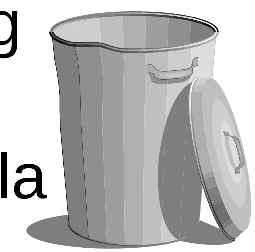

Errors entangle the data with the environment, producing *decoherence*. Recovery transforms entanglement of the data with the environment into entanglement of the ancilla with the environment, "purifying" the data. Decoherence is thus reversed. Entropy introduced in the data is transferred to the ancilla 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

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

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 $\boldsymbol{E_a^{-1}}$.

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.

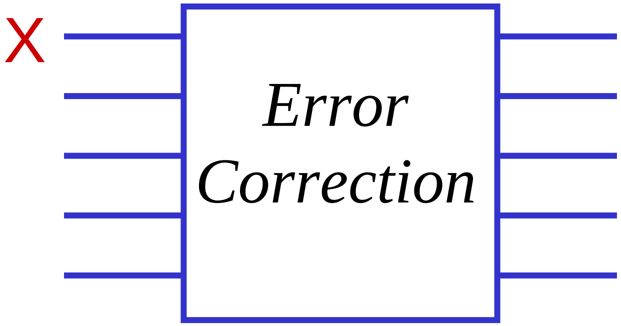

If input has at most one error, and circuit has no faults, output has no errors.

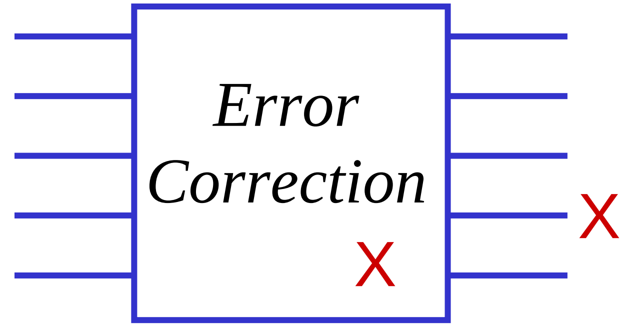

If input has no errors, and circuit has at most one fault, output has at most one error.

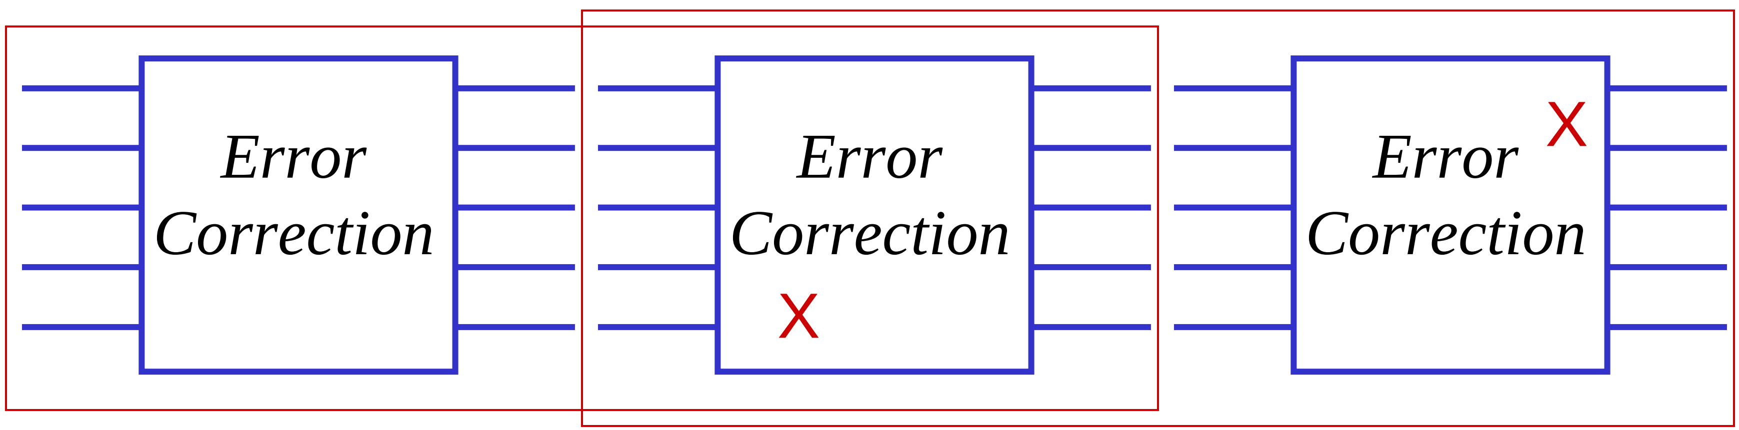

A quantum memory fails only if two faults occur in some “extended rectangle.”

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.

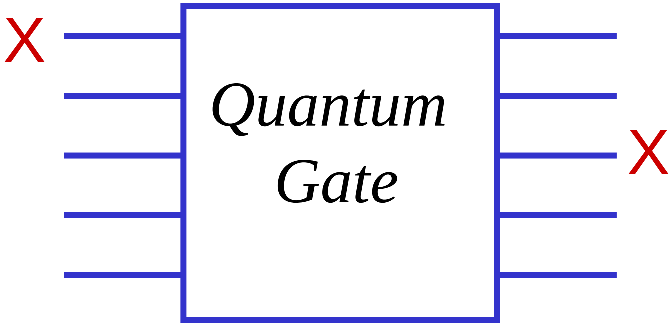

If input has at most one error, and circuit has no faults, output has at most one error in each block.

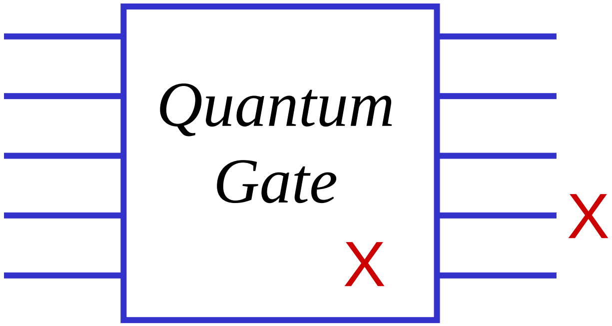

If input has no errors, and circuit has at most one fault, output has at most one error in each block.

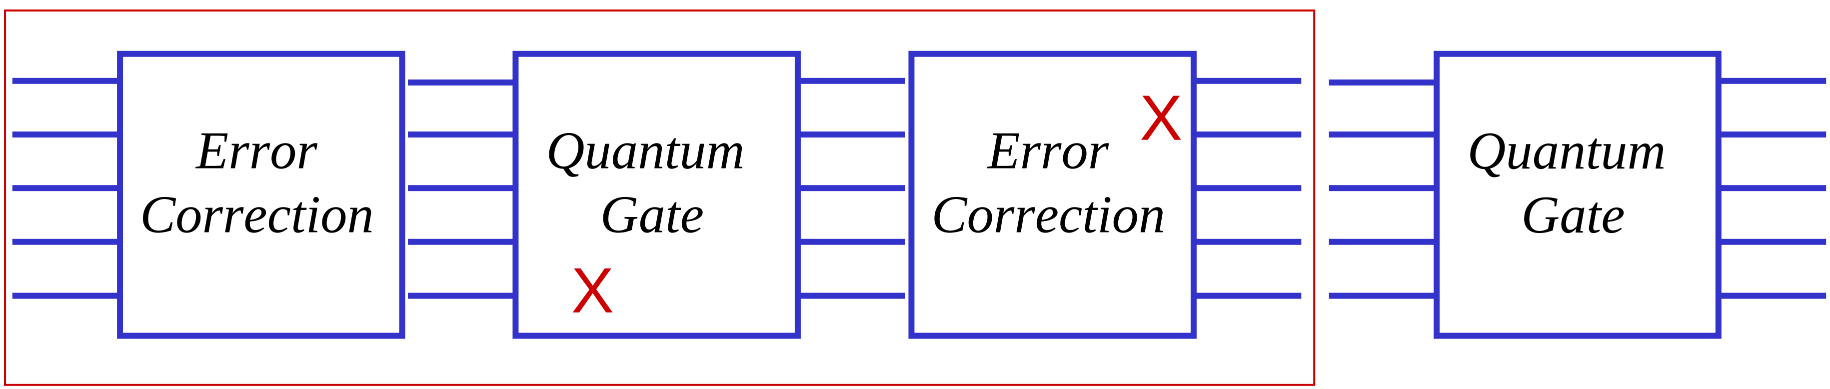

Each gate is preceded by an error correction step. The circuit simulation fails only if two faults occur in some "extended rectangle."

Fault-tolerant quantum gates

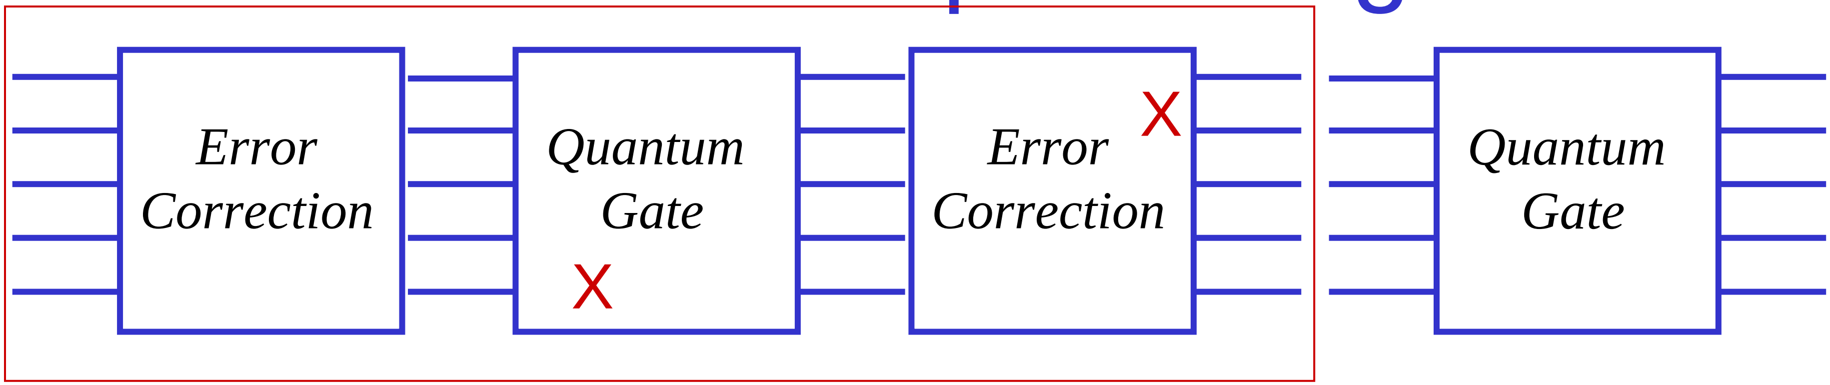

Each gate is followed by an error correction step. The circuit simulation fails only if two faults occur in some "extended rectangle."

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

$$P_{\mathrm{fail}} \leq L A_{\mathrm{max}} \varepsilon^2$$

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(ε^{-2}), compared to L=O(ε^{-1}) for an unprotected quantum circuit.

Fault-tolerant quantum gates

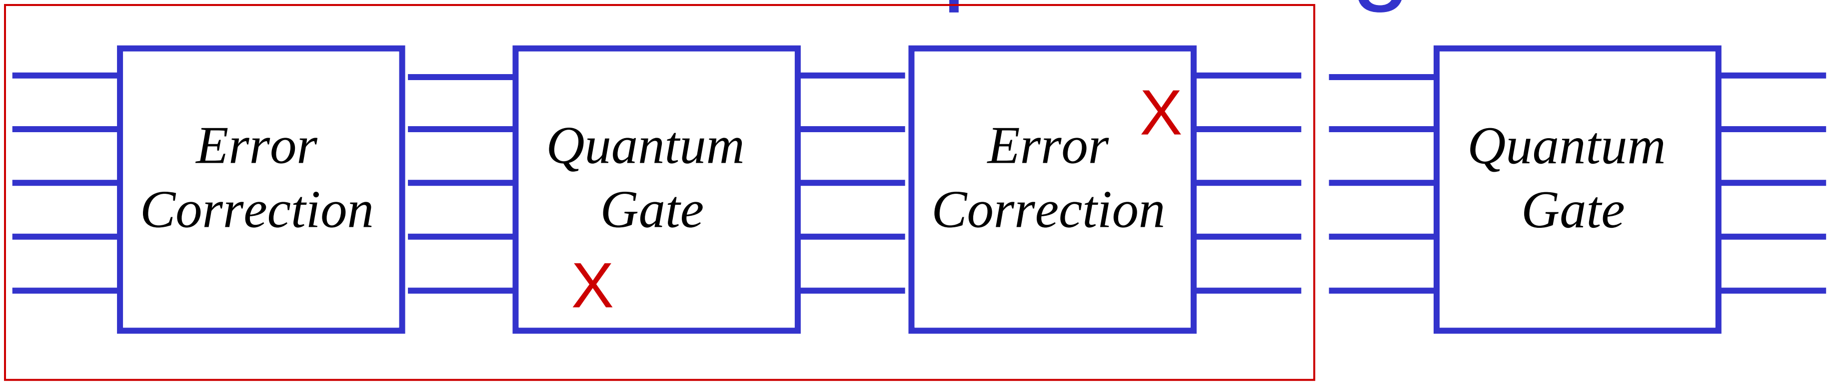

Each gate is followed by an error correction step. The circuit simulation fails only if two faults occur in some "extended rectangle."

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

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

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(ε^{-2}), compared to L=O(ε^{-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*.

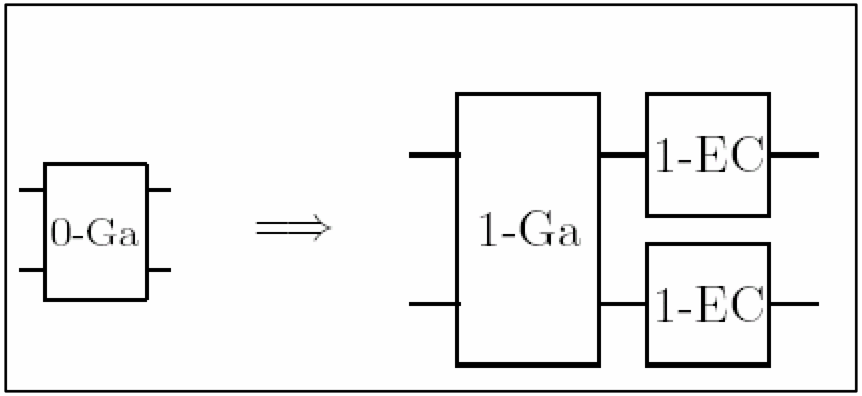

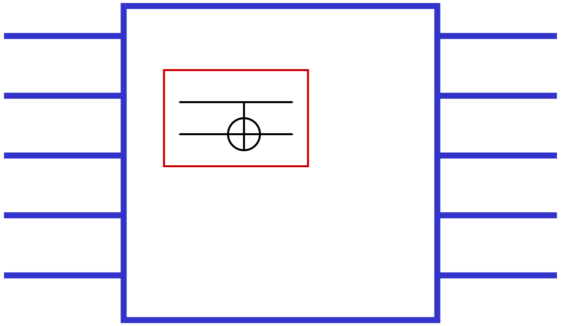

A *1-rectangle* is built from quantum gates.

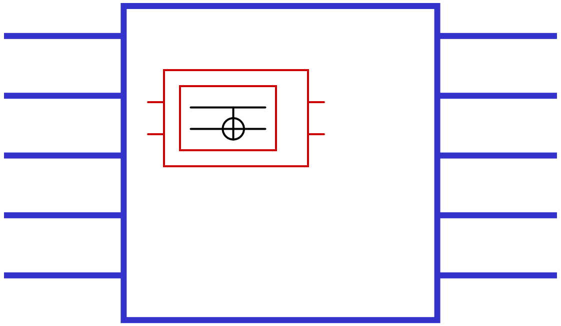

A *2-rectangle* is built from 1-rectangles.

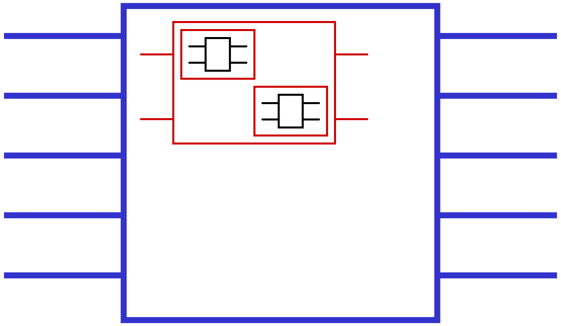

A *3-rectangle* is built from 2-rectangles.

(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:

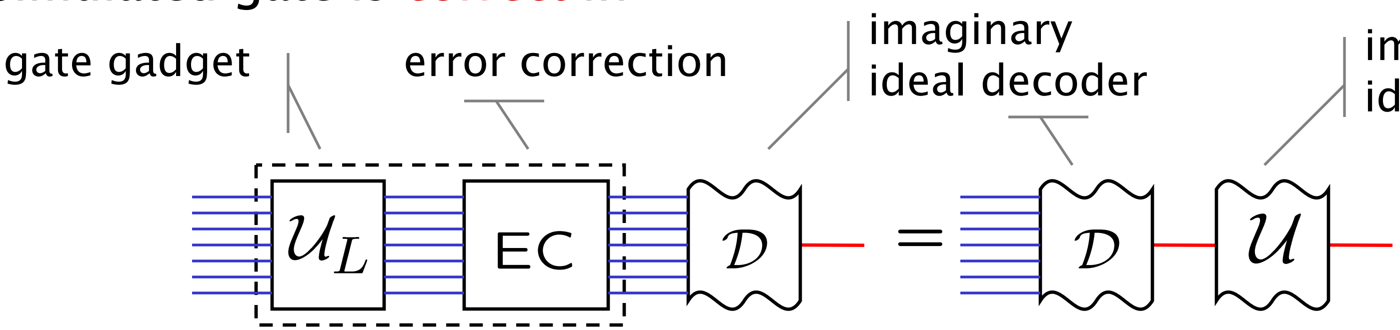

propagate decoders to the left

Simulated measurements and preparations are *correct* if:

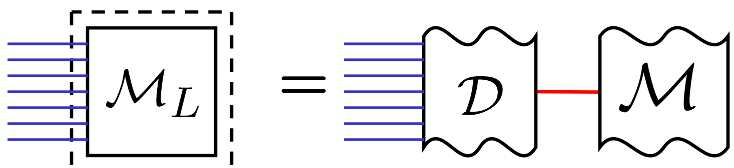

create decoders

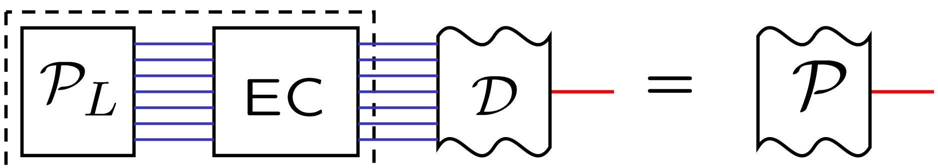

annihilate decoders

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

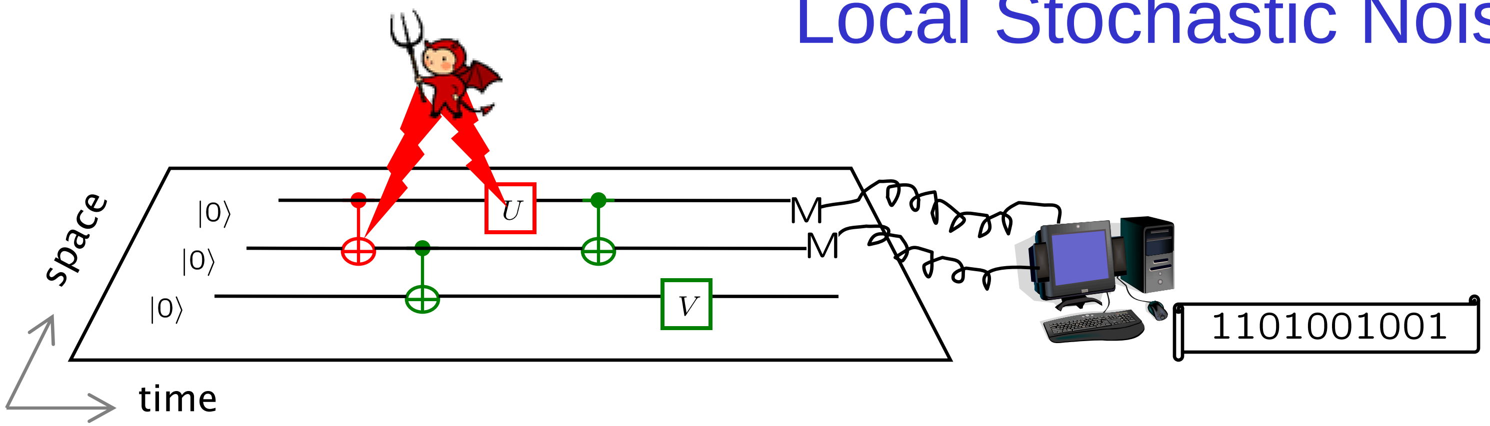

Noisy Circuit = Σ "*Fault Paths*"

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,

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

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

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

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

Aliferis, Gottesman, Preskill (2005)

Reichardt (2005)

assuming:

parallelism, fresh ancillas (necessary assumptions)

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

Some noteworthy recent developments

1) Threshold for local gates in 2D – Svore, DiVincenzo, Terhal (2006)
2) Threshold when measurements are slow – DiVincenzo, Aliferis (2006)
3) Improved thresholds with subsystem codes – Aliferis, Cross (2006)
4) Threshold for postselected computation – Knill (2004), Reichardt (2006), Aliferis, Gottesman, Preskill (2007)
5) Improved threshold via flagging and message passing – Knill (2004), Aliferis (2007)
6) Topological protection with cluster states – Raussendorf, Harrington, Goyal (2005, 2007)

Rigorous threshold estimate for local stochastic noise:

$$\boxed{\varepsilon_0 > 1.0 \times 10^{-3}}$$

Two issues

1) The local stochastic noise model describes generic noise with no special structure. Can we improve the threshold estimate by exploiting the structure of the noise in actual devices (such as a bias in favor of dephasing errors over bit flip errors)? (Aliferis-Preskill arXiv:0710.1301 --- I'll skip this to save time.)

2) The local stochastic noise model is handy for analysis and has some quasi-realistic features, but is still rather artificial; as usually formulated it is not founded on a physical (e.g. Hamiltonian) description of the origin of the noise. Can we prove threshold theorems for noise models that are better motivated physically, and how is the numerical value of the threshold affected by *coherence* and *memory* in the interaction with the environment?

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

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

For *local (coherent) 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.

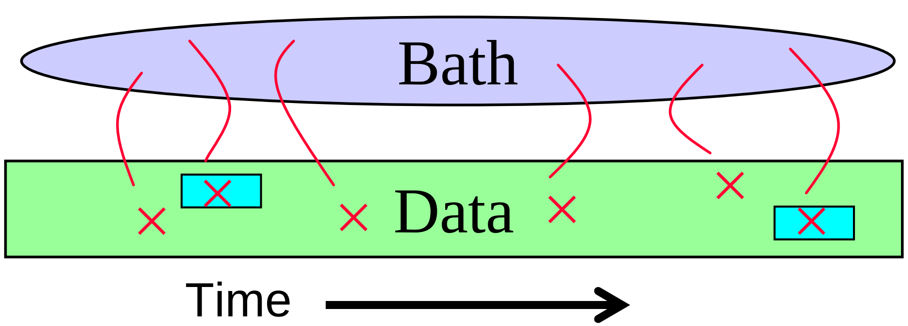

$$\boxed{\varepsilon = \max \left\| H^{(a)}_{System-Bath} \right\| t_0}$$

over all times and locations

time to execute a gate

Local non-Markovian noise

Non-Markovian noise with a *nonlocal* bath.

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

We can 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). Fault-tolerant quantum computation is effective if the noise strength ε is small enough, e.g., $\varepsilon < 10^{-4}$.

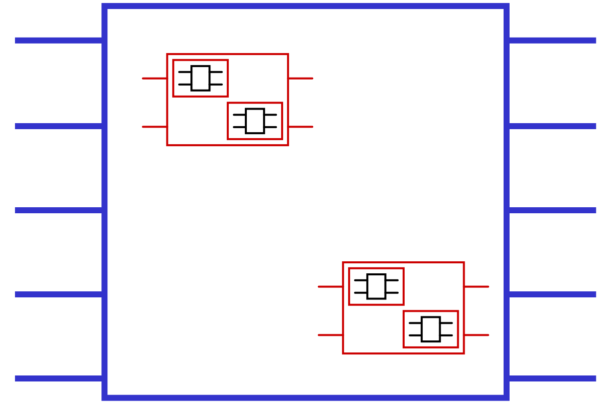

A hierarchy of "gadgets within gadgets" is reliable if the faults are *sparse*.

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.

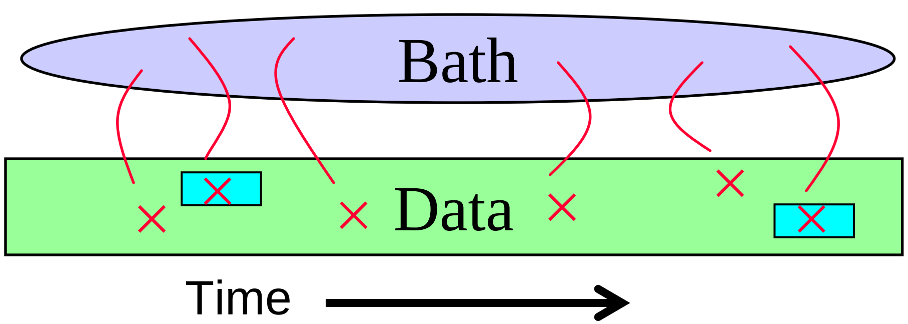

$$\varepsilon = \max \left\| H^{(a)}_{System-Bath} \right\| t_0$$

over all times and locations

time to execute a gate

Local non-Markovian noise

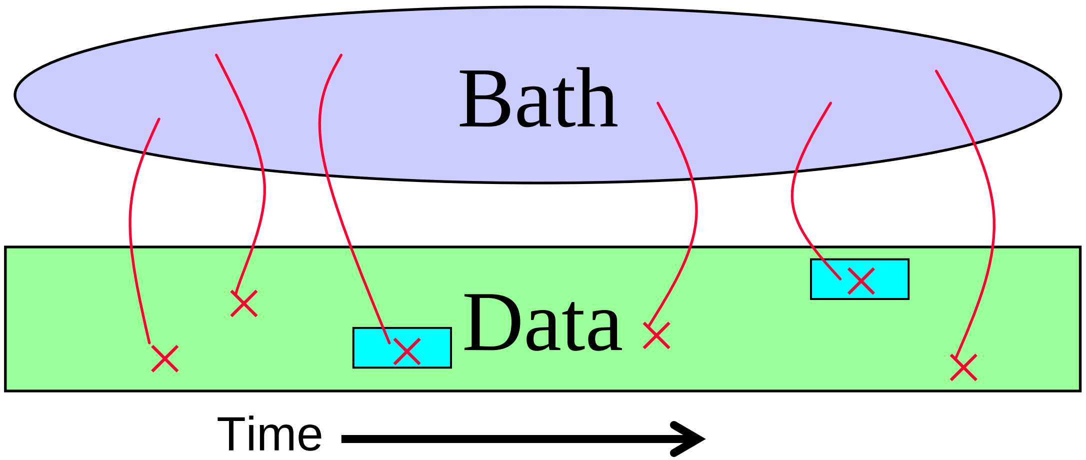

$$\left\| H^{(a)}_{System-Bath} \right\| t_0 < \varepsilon_0 \approx 10^{-4}$$

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

1) Interference: This condition (which applies even if there is no coupling to the bath at all, and the perturbation describes imperfect control of the qubits) seems discouraging because it requires an *amplitude* rather than a *probability* (square of an amplitude) to be small. (We pessimistically allow the bad fault paths to interfere constructively.) Under a plausible “randomization” hypothesis this estimate could be improved, but it is not so obvious what further assumptions we should make about the noise model to justify a rigorous argument that incorporates “randomization”.

2) Memory: The norm of the system-bath Hamiltonian is not directly measurable in experiments, and in fact for some noise models (e.g. coupling to a bath of harmonic oscillators) the norm is infinite. 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

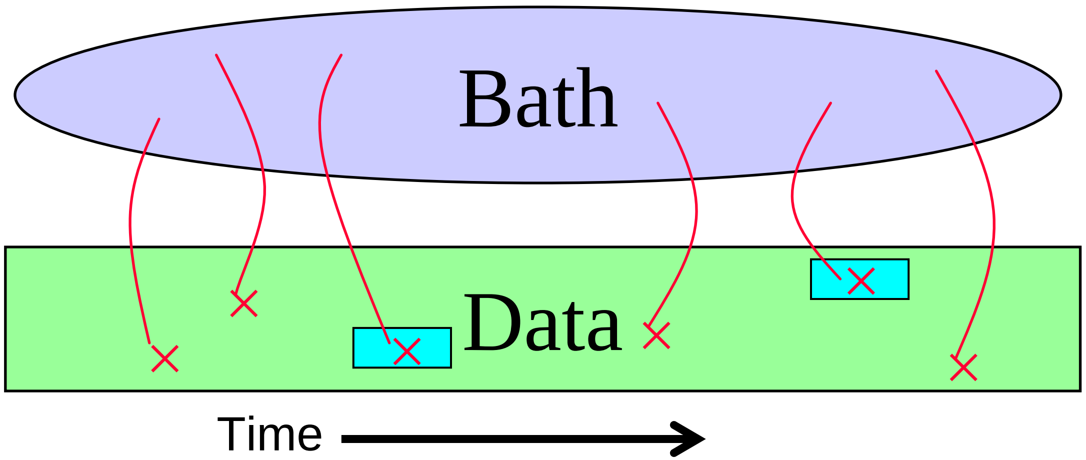

$$\left\| H^{(a)}_{System-Bath} \right\| t_0 < \varepsilon_0 \approx 10^{-4}$$

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

The derivation of the norm condition has the advantage that it does not require any assumption about the bath Hamiltonian, or about the state of the bath. (However, it does require that we model qubit preparation as an ideal preparation followed by interaction with the bath, and that we model qubit measurement as interaction with the bath followed by ideal measurement.)

But the norm condition has the disadvantage that it severely 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. But we should be cautious: perhaps during a long computation an initially benign state of the environment is driven to a new state that inflicts worse damage on the system than naively expected (“Alicki’s nightmare”).

Local non-Markovian noise

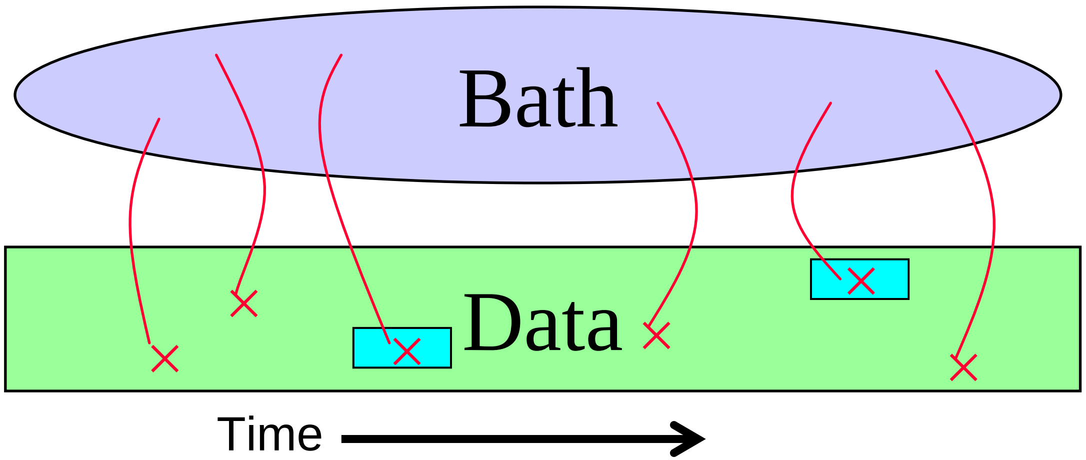

$$\boxed{\left\| H^{(a)}_{System-Bath} \right\| t_0 < \varepsilon_0 \approx 10^{-4}}$$

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

If we are willing to make further assumptions about the noise model, we *can* formulate a threshold condition in terms of the power spectrum of the bath fluctuations, which places less stringent constraints on the high frequency noise than the operator norm condition.

We will consider the case where each qubit couples to a thermal bath of harmonic oscillators. Our task is to estimate the the norm squared of the bad part of the system-bath wave function:

$$\left\| \, |\psi_{SB}^{bad}\rangle \right\|^2 = \langle \psi_{SB}^{0} | U_{SB}^{bad\,\dagger} U_{SB}^{bad} | \psi_{SB}^{0} \rangle \le \varepsilon^{2r}$$

At least one insertion of perturbation at each of r marked locations

Gaussian noise model

In the *Gaussian noise model*, each system qubit couples to a bath of harmonic oscillators:

$$H = H_S + H_B + H_{SB} \qquad H_B = \sum_k \frac{1}{2}\omega_k a_k^\dagger a_k$$

(uncoupled oscillators)

$$H_{SB}(t) = \sum_x \sum_\alpha \lambda_\alpha(x,t)\phi_\alpha(x)\sigma_\alpha(x)$$

$$\phi_\alpha(x) = \sum_k g_{k,\alpha}(x)a_k + g_{k,\alpha}(x)^* a_k^\dagger$$

(x labels qubit position, ϕ is a Hermitian bath operator, σ is Pauli operator acting on the system qubit, λ is a coupling constant.)

$$\phi_\alpha(x,t) \equiv e^{iH_B t}\phi_\alpha(x)e^{-iH_B t} = \sum_k g_{k,\alpha}(x)a_k e^{-i\omega_k t} + g_{k,\alpha}(x)^* a_k^\dagger e^{i\omega_k t}$$

(“interaction picture” field)

In the bath’s “vacuum” state, annihilated by each a_k,

$$_B\langle 0 | \phi_\alpha(x_1,t_1)\phi_\beta(x_2,t_2) | 0\rangle_B \equiv \int_0^\infty d\omega\, J_{\alpha\beta}(x_1,x_2,\omega)e^{-i\omega(t_1-t_2)}$$

$$\sum_k g_{k,\alpha}(x_1)g_{k,\beta}(x_2)^* \approx \int d\omega\, J_{\alpha\beta}(x_1,x_2,\omega)$$

(noise power spectrum, where sum over modes has been approximated by a frequency integral.)

Gaussian noise model

We say that the noise is *Gaussian* because the fluctuations of the bath obey Gaussian statistics: all correlation functions are determined by the two-point correlators. For a shorthand, denote $\langle \phi_\alpha(x_1,t_1)\phi_\beta(x_2,t_2)\rangle_B \equiv \Delta(1,2)$

Then

$$\langle \phi(1)\phi(2)\phi(3)\phi(4)\rangle_B \equiv \Delta(1,2)\Delta(3,4)+\Delta(1,3)\Delta(2,4)+\Delta(1,4)\Delta(2,3)$$

(a sum of "contractions"). Applies not just to vacuum expectation value, but also to expectation value in a thermal state of the bath ("Wick's theorem").

$$\langle \phi(1)\phi(2)\phi(3)\phi(4)\rangle_B =$$

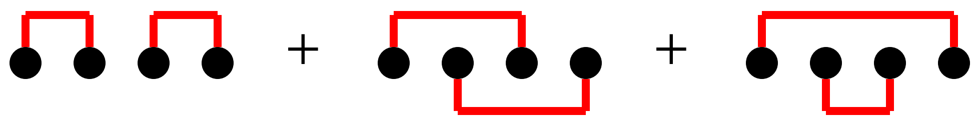

Similarly, the 2n-point correlation function can be expressed as a product of two-point correlators, summed over all possible pairwise contractions.

$$\langle \phi(1)\phi(2)\cdots\phi(2n)\rangle_B = \sum_{contractions} \Delta(i_1,j_1)\Delta(i_2,j_3)\cdots\Delta(i_n,j_n)$$

Gaussian noise model

Now we consider the case where $r = 1$ location(s) in the quantum circuit is "bad"; i.e., has at least one insertion of the perturbation. We are to sum all the "bad" contributions to the norm squared of the (pure) state of system and bath.

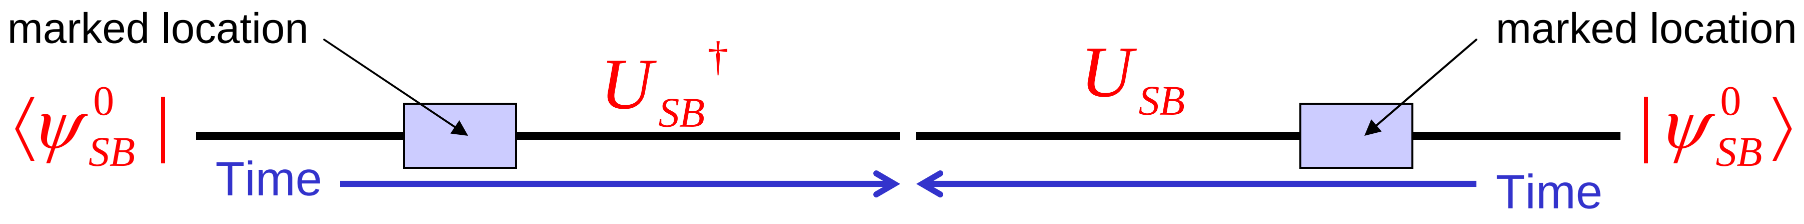

It is convenient to bend this picture into a hairpin shape ("Schwinger-Keldysh diagram")

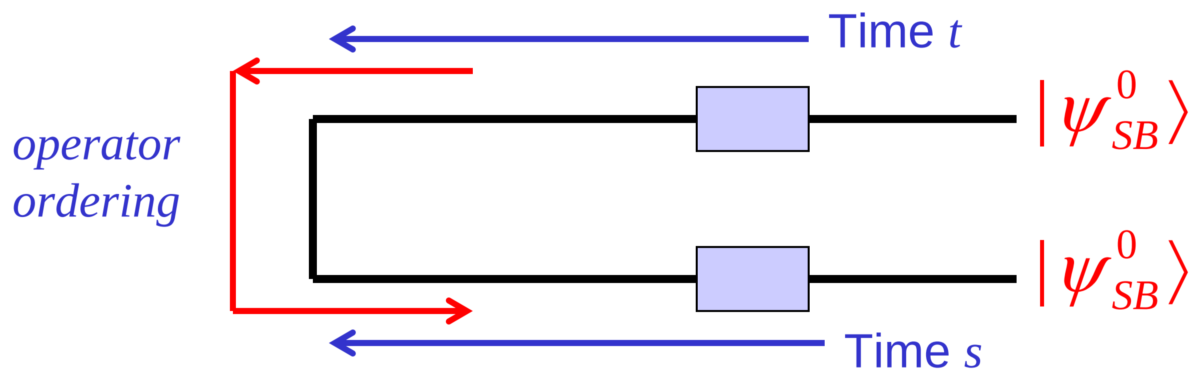

Time increases to the left on both branches, but "time-ordered" operators on the "upper branch" act "before" "anti-time-ordered" operators on the "lower branch".

Gaussian noise model

Now we consider expanding the time evolution operator U_{SB} in powers of the perturbation H_{SB}, summed to all orders. For a fixed term in this expansion, the system and the bath are uncoupled in between insertions of H_{SB}: the system evolves ideally between insertions, as determined by H_S, and the bath evolves as determined by H_B ("interaction picture").

Thus tracing out the bath generates the expectation value of a product of bath fields in the interaction picture, which can be evaluated using Wick's theorem (i.e., using the Gaussian statistics of the bath fluctuations). This is accompanied by the expectation value in the system's initial state of a product of interaction picture operators acting on the system qubits.

We are to sum up all the diagrams with at least one insertion of the perturbation inside the marked location on each branch of the Keldysh diagram.

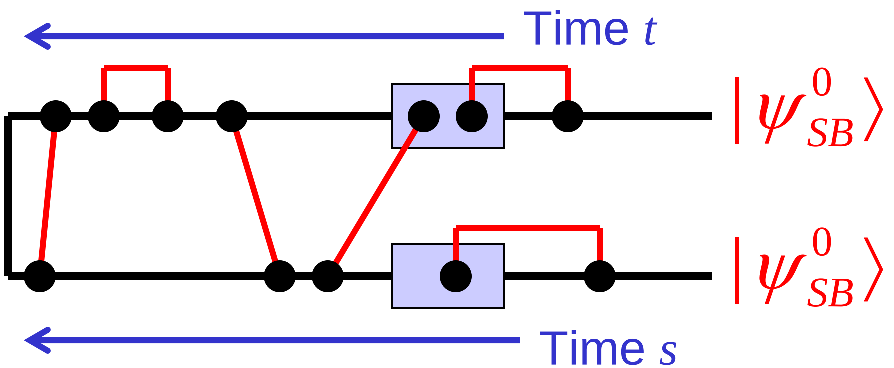

This sum is the norm squared of the bad part of the system-bath state:

$$\langle\psi_{SB}^0|U_{SB}^{bad\,\dagger}U_{SB}^{bad}|\psi_{SB}^0\rangle$$

Gaussian noise model

We can do the sum exactly only in some special cases (more about that later). But we can get a useful upper bound on the sum by this method (Cf, Terhal-Burkard, AGP, AKP)

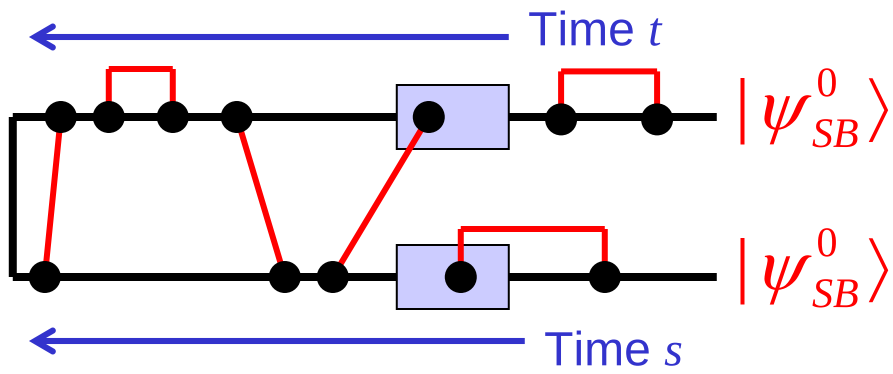

Suppose we fix the *earliest* insertion of the perturbation inside the marked location *on both branches*. These insertions might be contracted with one another; otherwise, each is contracted with another insertion somewhere else. Now we are to:

(1) "Dress" these diagrams with all possible additional insertions and contractions. But these additional insertions, in order to be "legal," must not occur in the marked location earlier than the fixed earliest insertion.

(2) Integrate over the position of the earliest insertion inside the marked location on both branches.

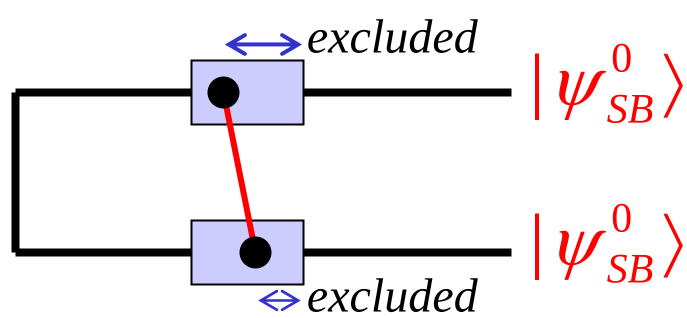

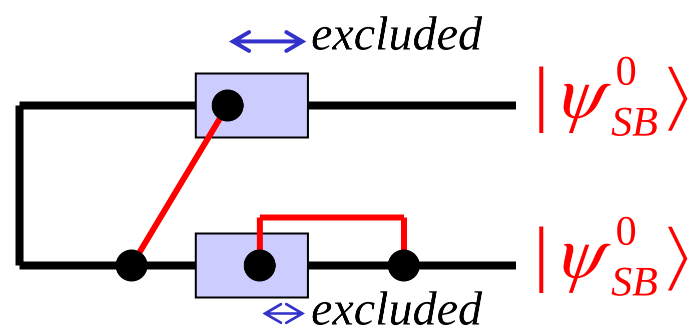

Gaussian noise model

Suppose, for example, that the earliest insertions inside the marked location on the two branches are contracted with each other.

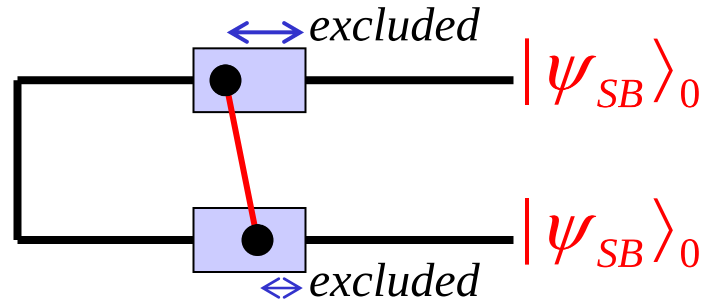

Then, the resummation of all the legal ways to dress this diagram is equivalent to evolving the state using a "hybrid" Hamiltonian, which is $H_{hybrid} = H_S + H_B$ in the marked location after the fixed first insertion, and $H_{hybrid} = H_S + H_B + H_{SB}$ everywhere else.

When we integrate over the position in the marked location of the first insertion, then, we have

$$\left| \int_{\square} ds \int_{\square} dt \, \langle \psi_{SB}^{0} | \sigma_\alpha(s) \sigma_\beta(t) | \psi_{SB}^{0} \rangle \, \lambda_\alpha(s) \lambda_\beta(t) \, {}_B\langle 0 | \phi_\alpha(s) \phi_\beta(t) | 0 \rangle_B \right|$$

$$\leq \int_{\square} ds \int_{\square} dt \, \left| \lambda_\alpha(s) \lambda_\beta(t) \, {}_B\langle 0 | \phi_\alpha(s) \phi_\beta(t) | 0 \rangle_B \right|$$

Here $\sigma_\alpha(t) = U_{SB}^{hybrid\,\dagger} \sigma_\alpha(t) \, U_{SB}^{hybrid}$ is evaluated in the "hybrid picture", and s and t are integrated over the marked location (denoted by $\square$). The sum over α and β is understood.

Gaussian noise model

By similar reasoning, using the hybrid picture, we can bound the sum of diagrams such that the earliest insertions of the perturbation inside the marked locations are not contracted with one another:

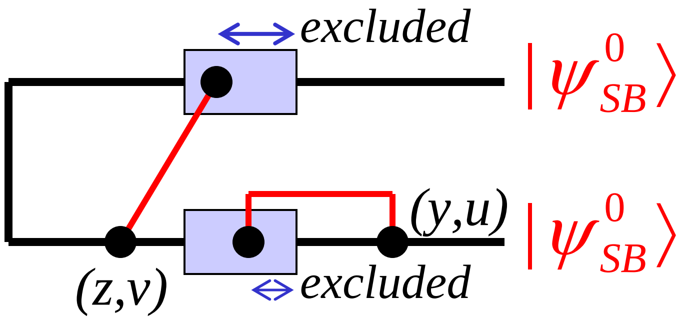

$$\leq \int_{\square} ds \int_{all} du \sum_{y} \left| \lambda_\alpha(x,s)\lambda_\gamma(y,u)\,{}_B\langle 0|\bar{T}\left[\phi_\alpha(x,s)\phi_\gamma(y,u)\right]|0\rangle_B \right|$$

$$\times \int_{\square} dt \int_{all} dv \sum_{z} \left| \lambda_\beta(x,t)\lambda_\delta(z,v)\,{}_B\langle 0|T\left[\phi_\beta(x,t)\phi_\delta(z,v)\right]|0\rangle_B \right|$$

Here (y,u) is the spacetime position of the insertion that is contracted with the first insertion inside the marked location on the lower branch, and (z,v) is the spacetime position of the insertion that is contracted with the first insertion inside the marked location on the upper branch. These can be anywhere except for the excluded region (and can be on either branch); we still have an upper bound if we integrate over all of spacetime. The T denotes time-ordering and $\bar{T}$ denotes anti-time ordering (needed to ensure the proper order of operators).

Gaussian noise model

When there are r marked locations in the circuit, we get a bound on norm squared of the bad part by summing over all ways to contract the marked locations, either with one another or with external locations (shown for $r=2$).

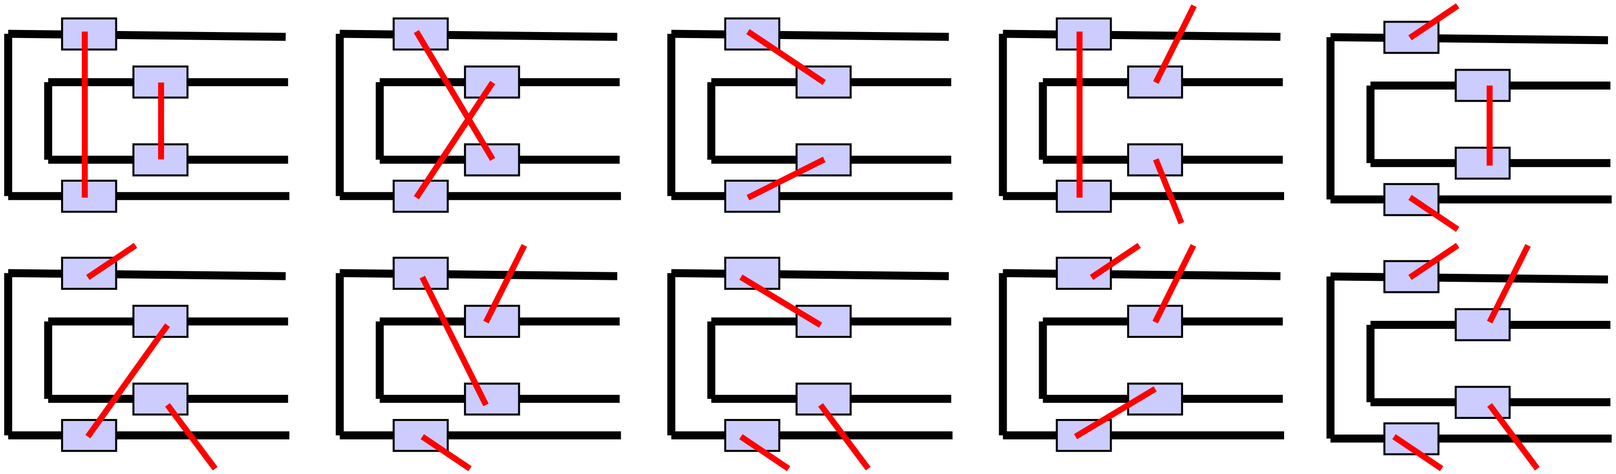

Using the same methods as in AKP05, we can bound the sum of the absolute values of all the diagrams, finding: $\langle \psi_{SB}^{0} | U_{SB}^{bad\dagger} U_{SB}^{bad} | \psi_{SB}^{0} \rangle \leq \varepsilon^{2r}$

where: $\varepsilon^2 = C \int_{\square} ds \int_{all} du \sum_{y} \left| \lambda_\alpha(x,s) \lambda_\gamma(y,u)\, {}_B\langle 0 | \tilde{T}\left[\phi_\alpha(x,s) \phi_\gamma(y,u) \right] | 0 \rangle_B \right|$

and $C = e^{2+1/e}$

In this noise model, fault-tolerant quantum computing works if ε is small enough (e.g. smaller than 10^{-4}).

Gaussian noise model

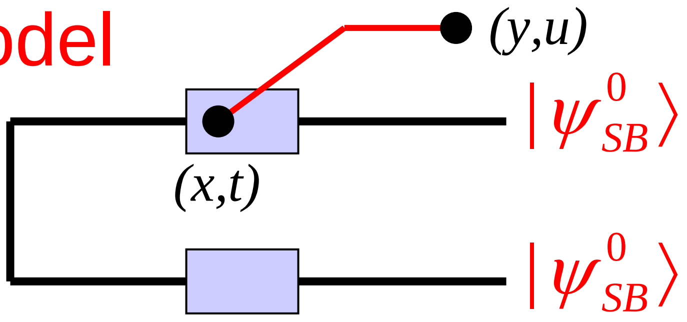

In this noise model, fault-tolerant quantum computing works if ε is small enough (e.g. smaller than 10^{-4}).

$$\varepsilon = \max\left(C\int_{\square} dt \int_{all} du \sum_{y} \left| \lambda_\alpha(x,t)\lambda_\gamma(y,u)\, {}_B\langle 0 | \tilde{T}\left[\phi_\alpha(x,t)\phi_\gamma(y,u)\right] | 0\rangle_B \right| \right)^{1/2}$$

If correlations are critical (decay like a power), then this expression converges provided

$$\int_{all} du \sum_{y} \left| \lambda_\alpha(x,t)\lambda_\gamma(y,u)\, {}_B\Delta(x,t;y,u) \right| < \infty$$

or $\displaystyle \int_{all} dt \int_{all} d^D x \frac{1}{\left(x^2 + t^{2/z}\right)^{\delta}} < \infty$ i.e. $D + z < 2\delta$

(D is the spatial dimension, δ is the scaling dimension of the bath field, and z is the dynamical critical exponent. This is the same criterion as cited by Novais et al.; however, here we have not used (at least not directly) the idea that fault paths that generate distinct syndrome histories should not be added coherently.

Gaussian noise model

$$\varepsilon = \left[C \int_{\square} dt \int_{all} du \sum_{y} \left| \lambda_\alpha(x,t) \lambda_\gamma(y,u)_B \Delta(x,t;y,u) \right| \right]^{1/2}$$

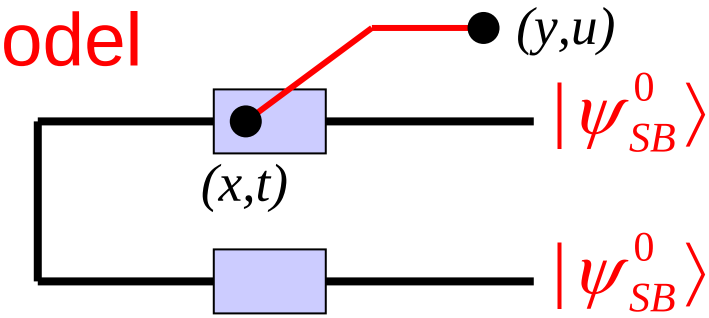

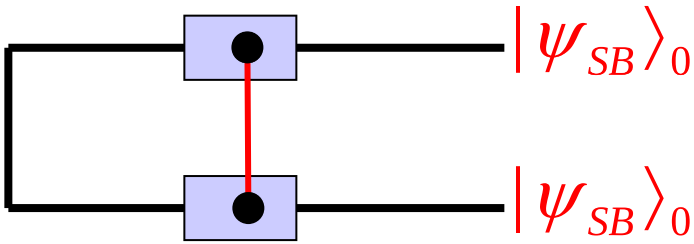

In the Markovian limit, the correlator is a delta function, with support at vanishing time difference:

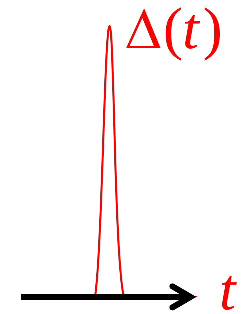

Thus: $$\varepsilon = \left(\Gamma t_0 \right)^{1/2}$$

where Γ is an error *rate*, and t_0 is the time to execute a gate. In the Markovian case, fault paths really do decohere, and errors can be assigned probabilities rather than amplitudes. But our argument is not clever enough to exploit this property, and hence our threshold condition requires the error amplitude to be small, rather than the square of the amplitude.

This result applies to "high temperature" Ohmic noise, which has a flat power spectrum up to a cutoff frequency (i.e. the inverse width of the peak). The norm condition, on the other hand, requires the *height* of the peak in the correlator to be small, a quantity that depends on the frequency cutoff.

Gaussian noise model

In the case of zero-temperature Ohmic noise,

$$\tilde{\Delta}(\omega) = 2\pi A \omega e^{-\omega\tau_c} \quad \text{and} \quad \Delta(t_1 - t_2) = \frac{-A}{\left[(t_1 - t_2) - i\tau_c\right]^2}$$

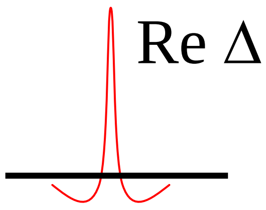

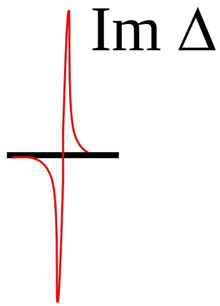

Both the real and the imaginary part of the correlator wiggle, and therefore the integral of the correlator has only a logarithmic sensitivity to the cutoff (cf. Novais et al.).

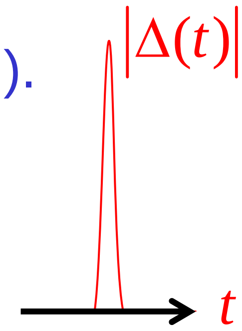

However, unfortunately when we take the absolute value of the correlator, we lose the benefit of the wiggles, and the cutoff dependence is stronger:

$$\varepsilon^2 \approx \int_0^{t_0} dt_1 \int_{-\infty}^{\infty} dt_2 \left|\Delta(t_1 - t_2)\right| = \int_0^{t_0} dt_1 \int_{-\infty}^{\infty} dt_2 \; A / \left[(t_1 - t_2)^2 + \tau_c{}^2\right] = \pi A \left(t_0 / \tau_c\right)$$

The height of the peak is τ_c^{-2} and its width is τ_c. By integrating, we improve the value of ε relative to to our original norm condition by a factor $\left(\tau_c / t_0\right)^{1/2}$. Still, rather strong sensitivity to the cutoff remains (in the zero-temperature Ohmic case).

Gaussian noise model

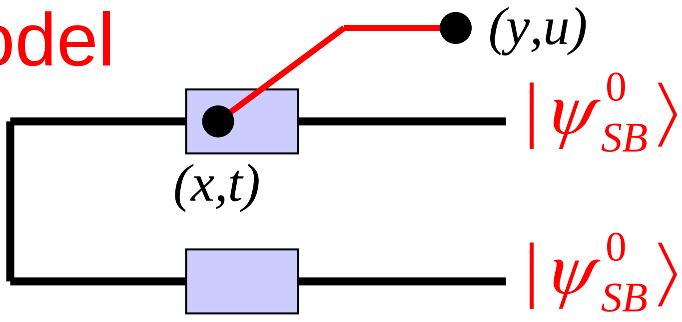

Thus in some cases (like high-temperature Ohmic noise) our new threshold condition for Gaussian noise has no artificial sensitivity to very-high-frequency fluctuations of the bath,

while in other cases (like zero-temperature Ohmic noise) sensitivity to the cutoff remains, yet is improved compared to the norm condition of Terhal-Burkard04, AGP05, AKP05;

i.e., $\varepsilon \approx \sqrt{A}\left(t_0/\tau_c\right)^{1/2}$ (new) vs. $\varepsilon \approx \sqrt{A}\left(t_0/\tau_c\right)$ (old).

Even this weaker dependence on the ratio of the working period of a gate to the cutoff time scale may be spurious. However, I have been able to prove this only for the extreme case of diagonal noise and diagonal gates (as in AP07).

For the diagonal case, the faults commute with the system-bath evolution operator and can be propagated forward to the measurements. The diagrams can be summed explicitly, and only logarithmic dependence on the cutoff is found. Even this logarithmic divergence arises because of the way preparations and measurements are modeled (e.g. an instantaneous ideal measurement preceded by interactions with the bath), and might be avoided by using a more realistic measurement model.

Toward “realistic noise”

1) We can improve the threshold estimate by exploiting the structure of the noise in actual devices. Diagonal two-qubit gates, which plausibly have highly biased noise, along with single-qubit preparations and measurements, suffice for universal fault-tolerant quantum computation.

2) We can formulate a threshold condition for non-Markovian noise in terms of the norm of the system-bath Hamiltonian, but this condition places severe constraints on very-high-frequency noise. For the special case of Gaussian non-Markovian noise, the threshold condition is less sensitive to the very-high-frequency noise. The condition can be improved further for diagonal Gaussian noise, and perhaps in other cases. Is it a mathematical technicality, or a real potential obstacle to large-scale fault tolerance (Alicki’s nightmare)?