Fault-tolerant quantum computation against realistic noise

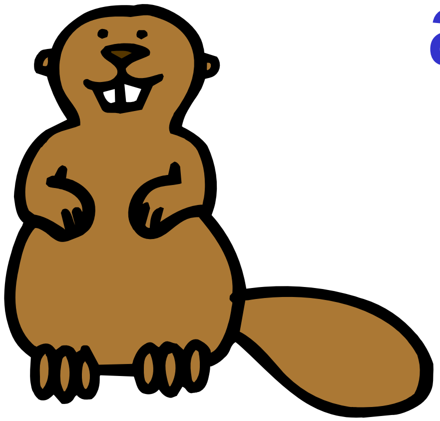

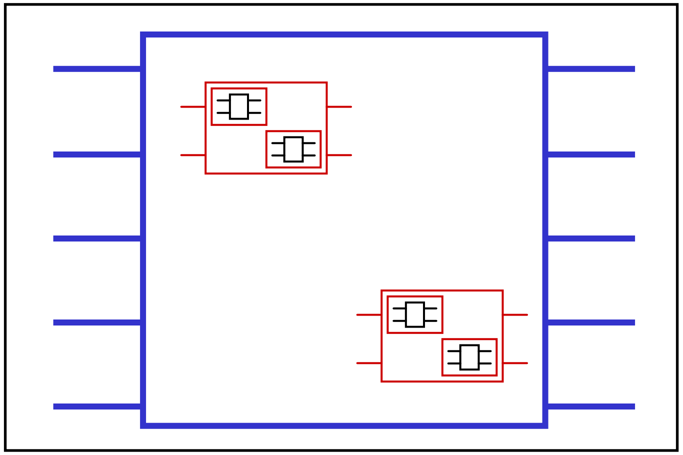

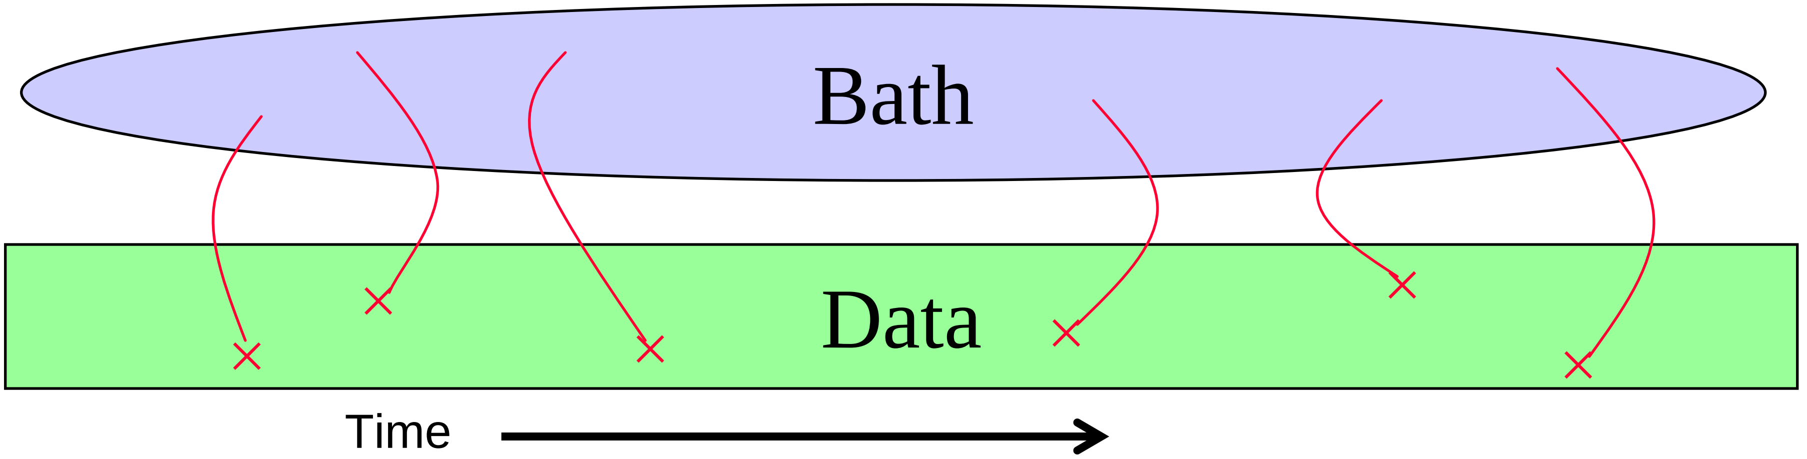

Institute for Quantum Information

Sponsored by the NSF

John Preskill, Caltech

19 December 2007

Quantum fault tolerance

- Error correction and fault tolerance will be essential in the operation of large-scale quantum computers, both to prevent decoherence and to control the cumulative effects of small errors in unitary quantum gates.
- This talk focuses on fault-tolerant processing of quantum information using quantum error-correcting codes (the foundation for our belief that scalable quantum computers are possible).
- There are a variety of other useful ideas concerning protecting quantum computers from noise … e.g., dynamical decoupling, noiseless subsystems, protection arising from (nonabelian) topological order, …
- I won't say much about these other ideas, even though they are important, they are related to my main topic, and they can be fruitfully combined with the methods I'll discuss.

Fault-tolerant quantum computation

1. Fault-tolerant quantum computing
2. Quantum accuracy threshold theorem
3. Biased noise (Aliferis-Preskill 2007)
4. Non-Markovian (Gaussian) noise.

Fault tolerance

Quantum information can be protected by a quantum error-correcting code. *But the quantum gates that we use to encode and to recover from error are themselves noisy.*

• 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.

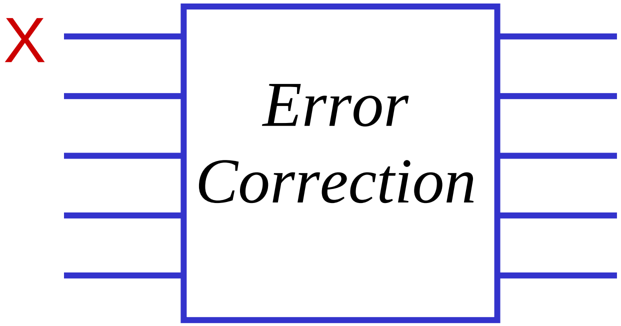

If input has at most one error, and circuit has no faults, output has no errors.

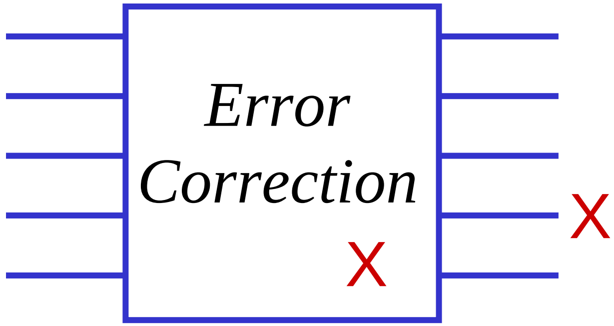

If input has no errors, and circuit has at most one fault, output has at most one error.

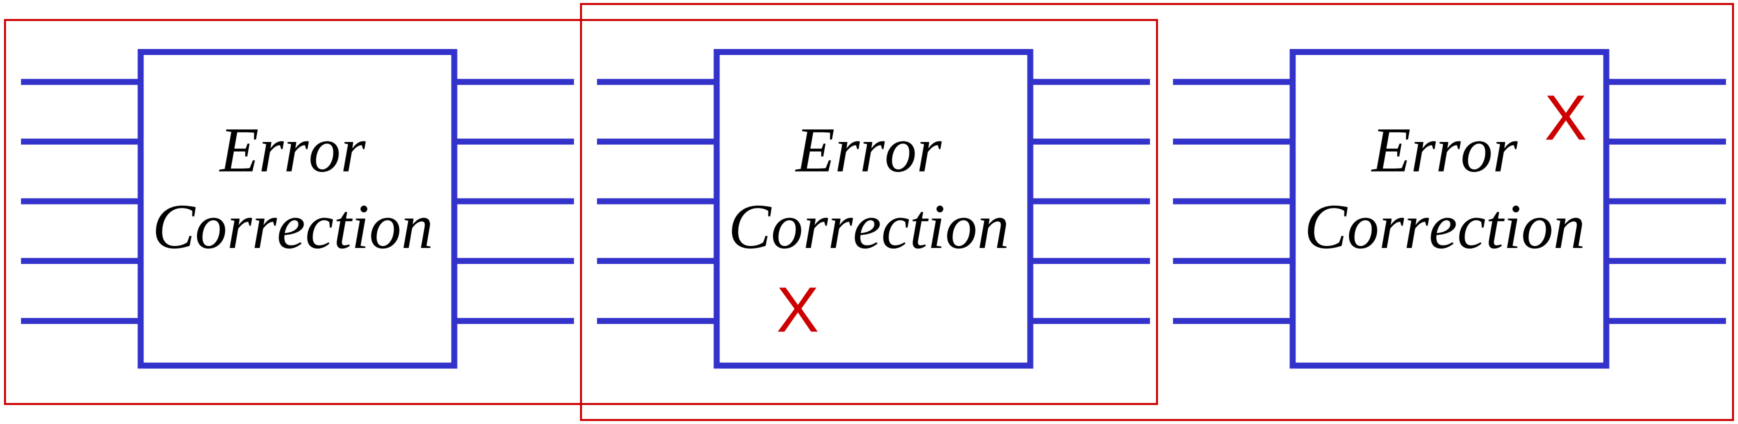

A quantum memory fails only if two faults occur in some “extended rectangle.”

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.

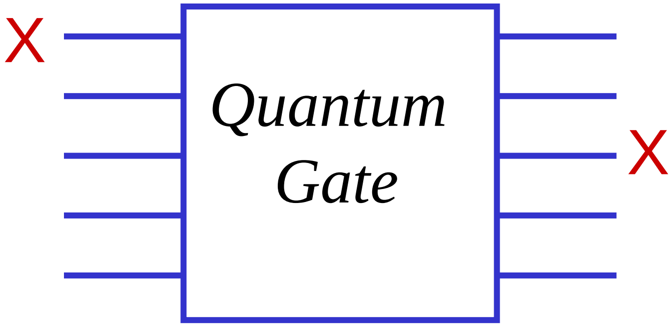

If input has at most one error, and circuit has no faults, output has at most one error in each block.

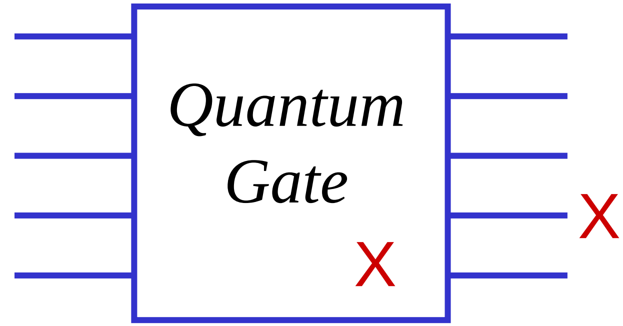

If input has no errors, and circuit has at most one fault, output has at most one error in each block.

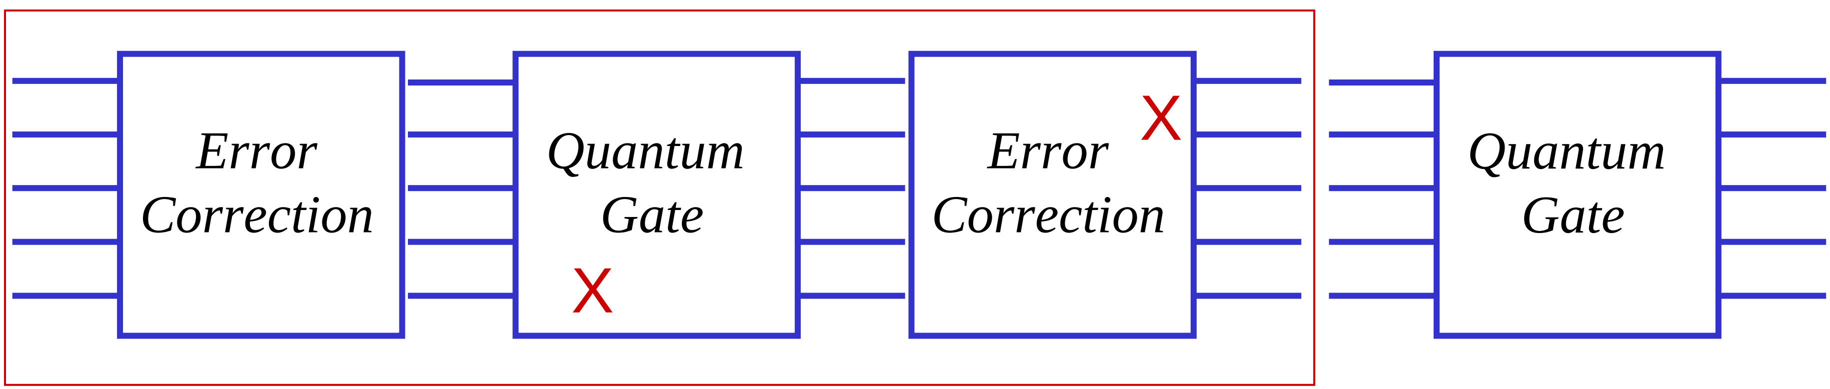

Each gate is preceded by an error correction step. The circuit simulation fails only if two faults occur in some “extended rectangle.”

Fault-tolerant quantum gates

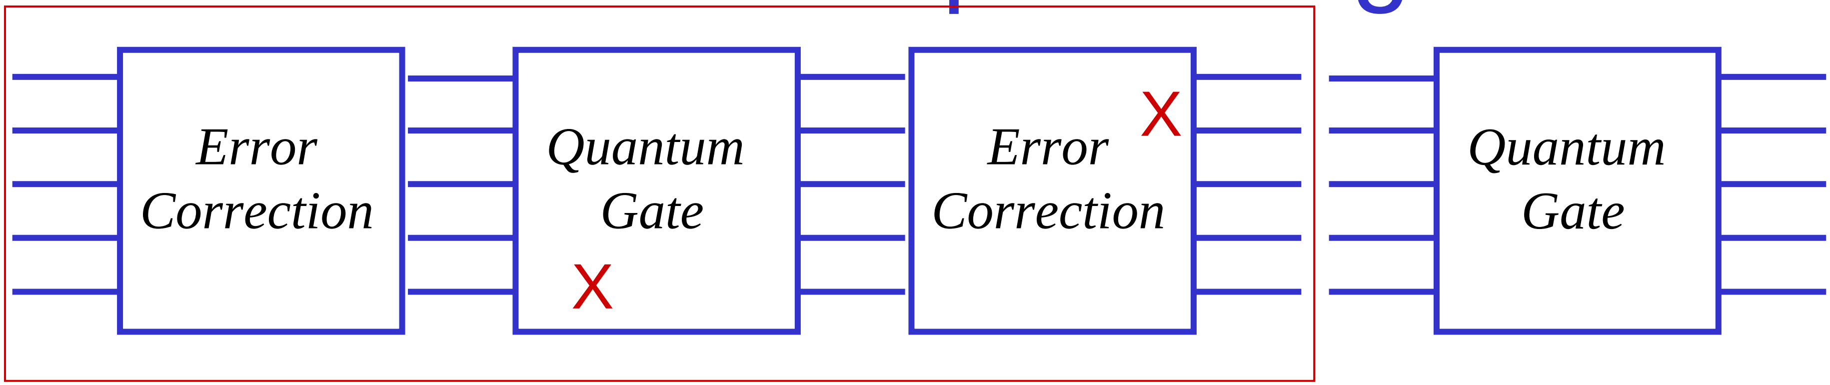

Each gate is followed by an error correction step. The circuit simulation fails only if two faults occur in some "extended rectangle."

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

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

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

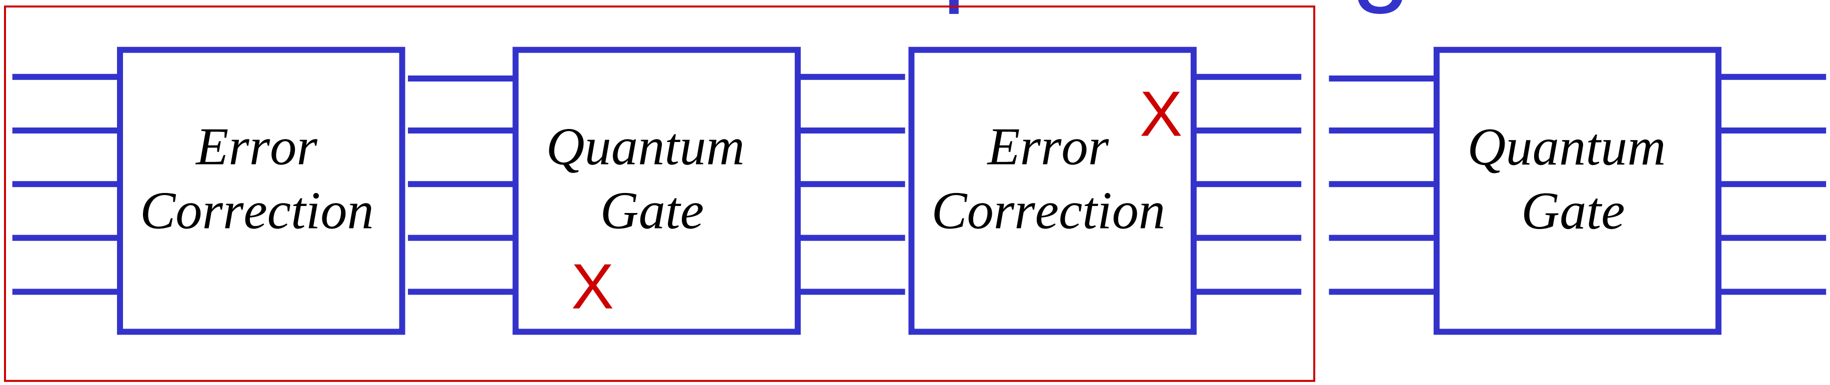

Each gate is followed by an error correction step. The circuit simulation fails only if two faults occur in some “extended rectangle.”

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

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

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*.

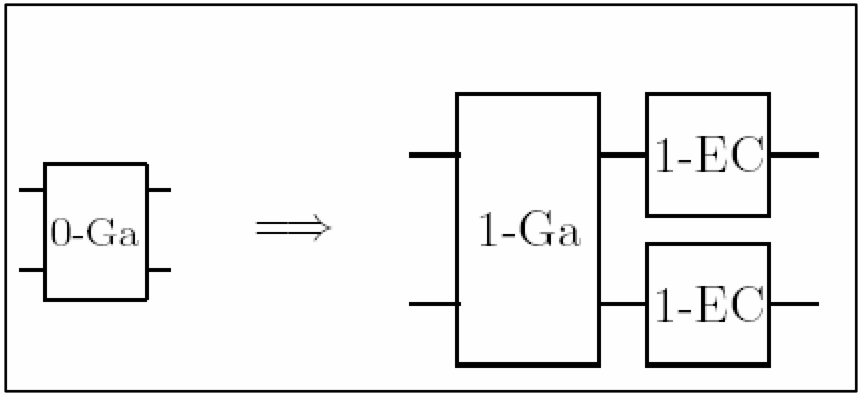

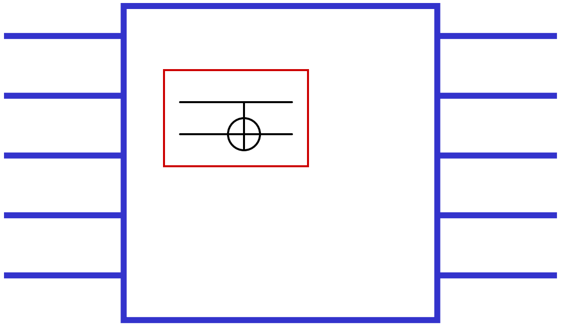

A *1-rectangle* is built from quantum gates.

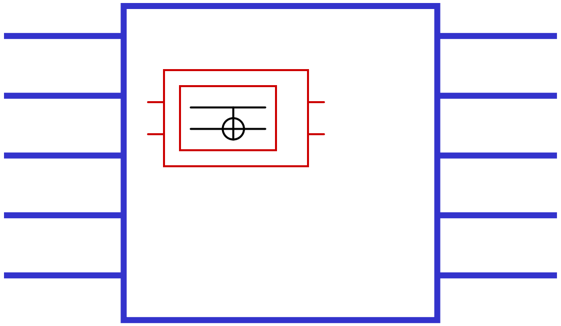

A *2-rectangle* is built from 1-rectangles.

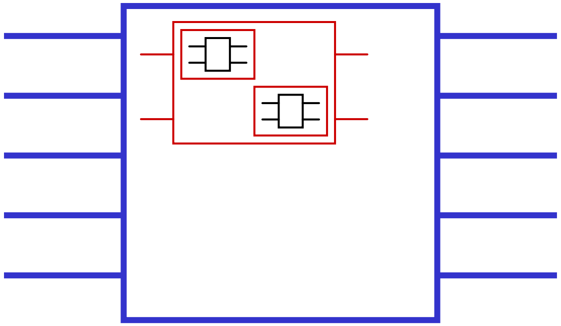

A *3-rectangle* is built from 2-rectangles.

(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:

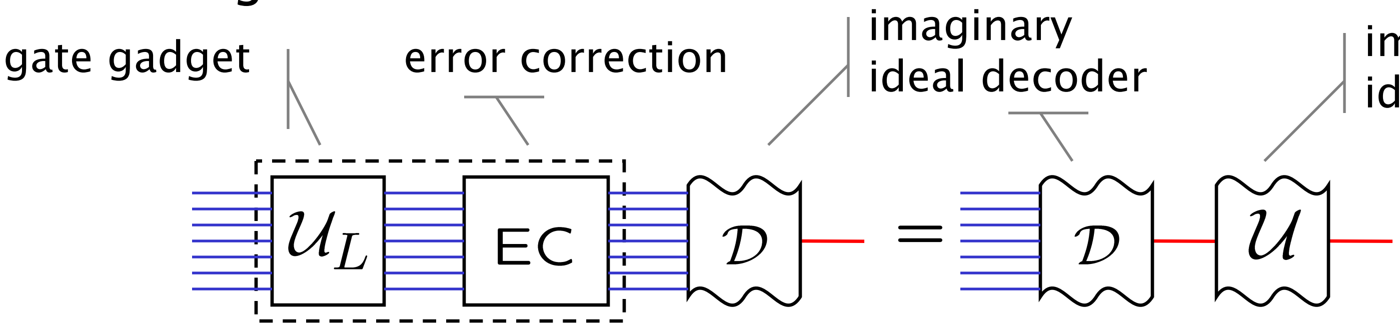

propagate decoders to the left

Simulated measurements and preparations are *correct* if:

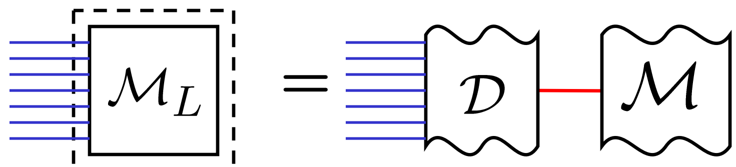

create decoders

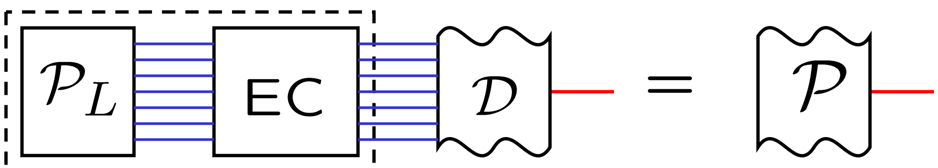

annihilate decoders

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

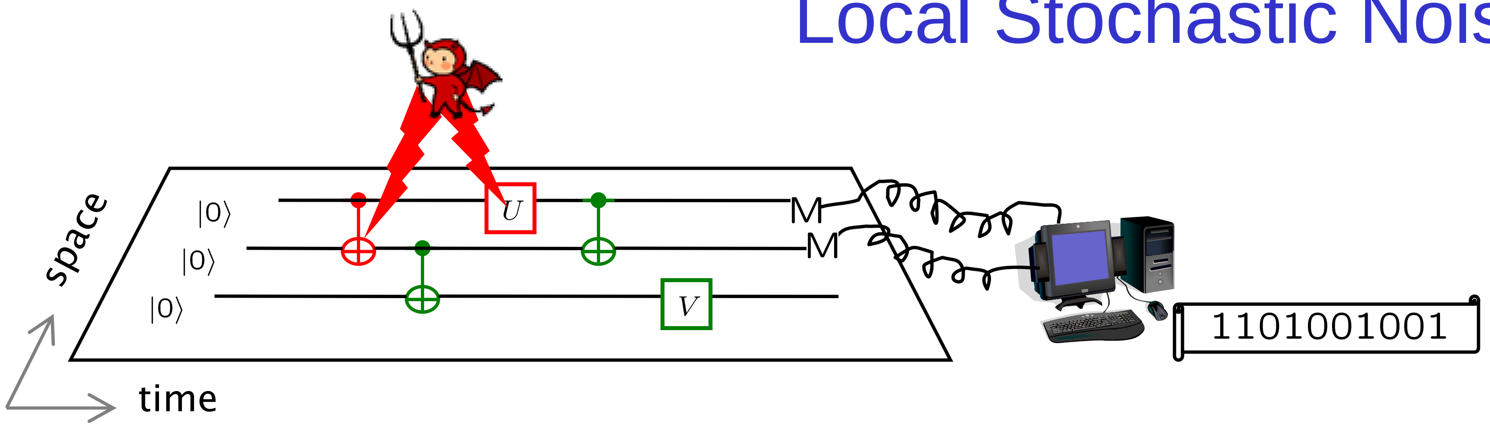

Noisy Circuit = Σ "Fault Paths"

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!

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

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:

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

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?

2) The local stochastic noise model is handy for analysis and has some quasi-realistic features, but 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?

Biased Noise

Aliferis-Preskill 2007

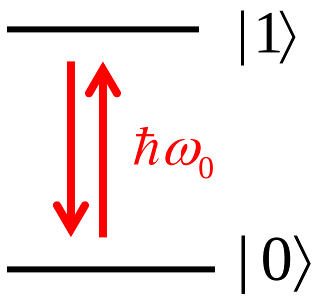

In many physical settings, Z noise (dephasing) is much stronger than X noise (relaxation). Dephasing arises from low frequency noise, while relaxation arises from noise with frequency ~ω_0 (the energy splitting between computational basis states). Typically, the higher-frequency noise has a different physical origin than the low-frequency noise, and it can be orders of magnitude weaker.

Can fault-tolerant gadgets exploit this bias? We should use a code that corrects more Z errors than X errors (for example, a phase-error correcting repetition code at the bottom of a concatenated scheme). But our universal set of gate gadgets must have the property that the (common) Z errors do not propagate to become (rare) X errors. (For example, transversal Hadamard gates should be avoided.)

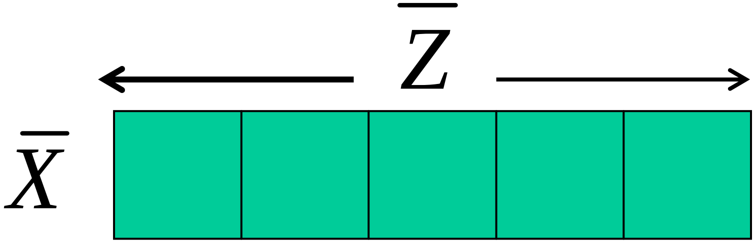

Biased Noise

Our universal set of gadgets must have the property that the (common) Z errors do not propagate to become (rare) X errors. For example, transversal Hadamard gates should be avoided. E.g., CNOT gates (which are transversal for CSS codes) do not propagate Z to X, and other gates in the universal set can be teleported. We could estimate the threshold under the assumption that CNOT gates have highly biased noise (Gourlay & Snowdon 2000, Stephens et al. 2007, Evans et al. 2007).

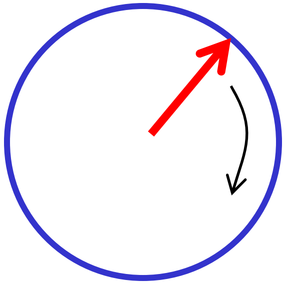

But is it reasonable to assume that the CNOT gate has highly biased noise? Gates are realized by a time-dependent control Hamiltonian that turns on and off, and a Z error occuring while a qubit is rotating on the Bloch sphere can generate an X error. Furthermore, if a qubit is rotated about the X axis, then an over-rotation or under-rotation can generate an X error.

We should distinguish between *diagonal* gates (diagonal matrices in the computational basis) and nondiagonal gates; it is reasonable to assume that the noise in diagonal gates is dominated by dephasing, while the noise in nondiagonal gates has no spectial structure. For diagonal gates, the ideal control Hamiltonian commutes with Z at all times; these gates do not propapagate Z noise to X noise.

Local Stochastic Biased Noise Model

$$\textit{Noisy Circuit} = \sum \text{"}\textit{Fault Paths}\text{"}$$

In our biased noise model, there are two different values of the noise strength: ε quantifies the rate for dephasing faults in diagonal gates and for unstructured faults in non-diagonal gates, and $\varepsilon' << \varepsilon$ quantifies the rate for unstructured faults in diagonal gates.

For *local stochastic biased noise* with strength $(\varepsilon, \varepsilon')$, the sum of the probabilities of all fault paths that have faults at r specified locations, or which s are unstructured faults at diagonal locations, is at most

$$\varepsilon^{r-s}\left(\varepsilon'\right)^{s}$$

(For each fault path, the operations at the faulty locations are chosen by the adversary.) For the unstructured faults, the Kraus operators are arbitrary, while for the structured "dephasing" faults, the Kraus operators are diagonal in the computational basis.

$\varepsilon' / \varepsilon \ll 1$ is the "bias factor".

Fault-tolerant scheme

We will build our fault-tolerant scheme using single-qubit preparations of the X eigenstate $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$, single-qubit measurements in the X basis, and the diagonal two-qubit CPHASE gate:

CPHASE = controlled-Z = diag(1, 1, 1, -1)

The n-qubit repetition code C_1 will protect against dephasing (Z errors) at the lowest level of a concatenated scheme, but it provides no protection against bit flips (X errors). Using our fundamental operations, we will construct a fault-tolerant CNOT gate and fault-tolerant encoded measurements, where the effective noise in the encoded CNOT is approximately balanced. Then we may use well-known fault-tolerant constructions and a concatenated code C_2 to realized highly accurate encoded CNOT gates protected by $C_1 \triangleright C_2$. Finally, we use injection by teleportation of the states

$$|+i\rangle = (|0\rangle + i|1\rangle)/\sqrt{2} \qquad |T\rangle = (|0\rangle + e^{i\pi/4}|1\rangle)/\sqrt{2}$$

and state distillation to complete the set of universal gadgets at the top level.

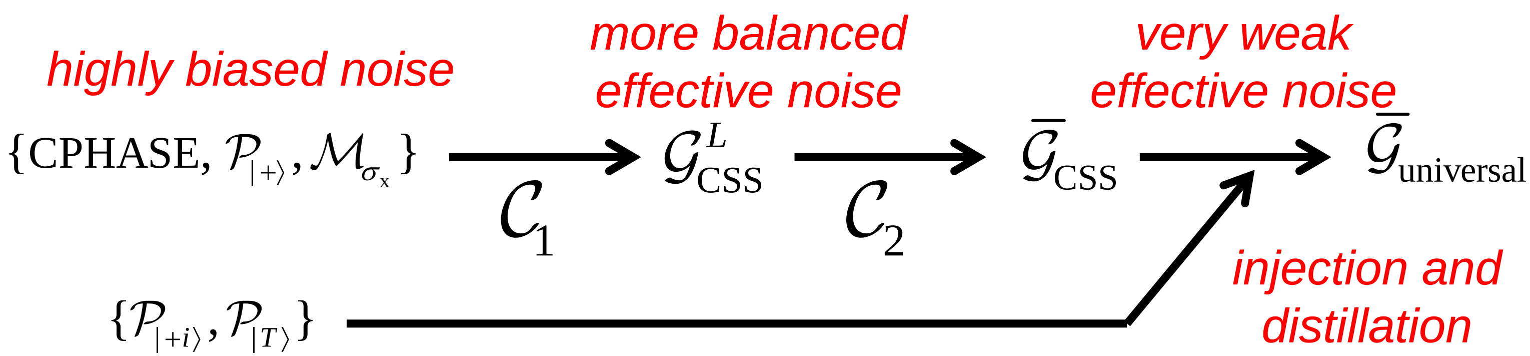

Encoded Preparations and Measurements

To be specific, consider the three-qubit repetition code (though we may actually want to use a longer code). The codewords are $|{+}{+}{+}\rangle$ and $|{-}{-}{-}\rangle$; these can be encoded and measured transversally:

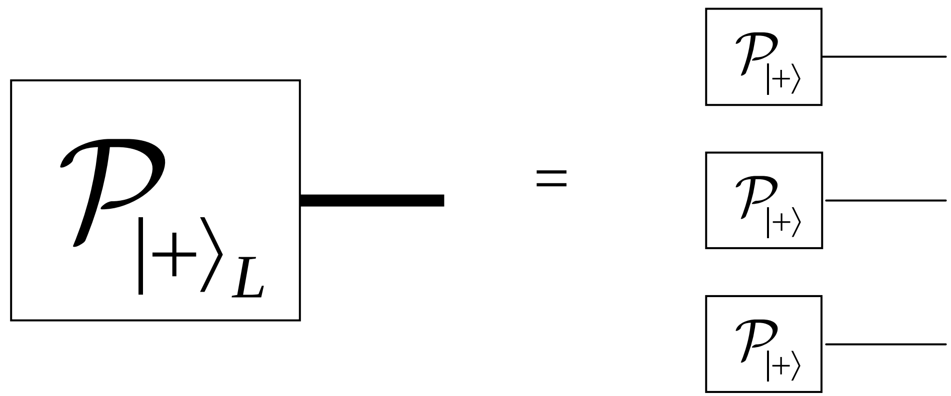

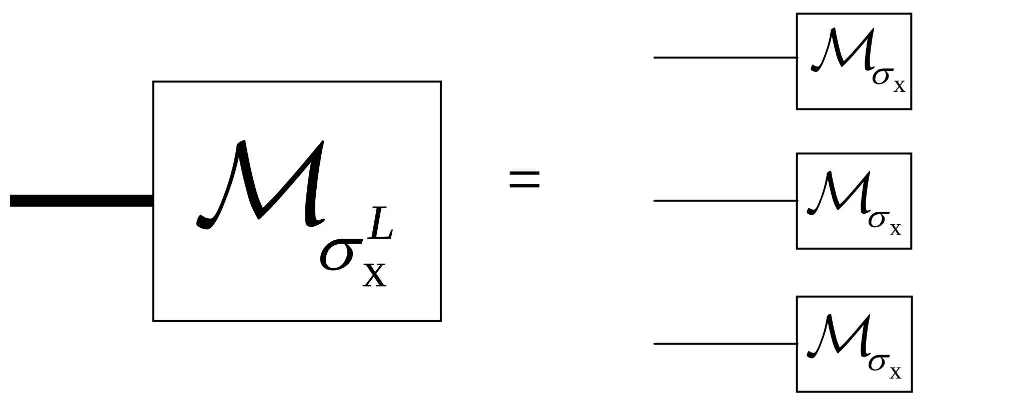

For the encoded measurement, a majority vote is performed on the outcomes. These operations are fault-tolerant, because two faults are needed to cause an encoded error.

Encoded Measurement

The encoded $Z_L = Z\,Z\,Z$ can be measured by preparing an ancilla qubit, which interacts with each qubit in the data block via three consecutive CPHASE gates, and is then measured in the X basis. However the measurement is not fault-tolerant because a single Z fault acting on the ancilla can flip the outcome. To ensure fault tolerance, the measurement is repeated three times, and the majority of the outcomes is computed. (Fault-tolerance means protection against a single Z fault in the circuit; there is no protection against X faults.)

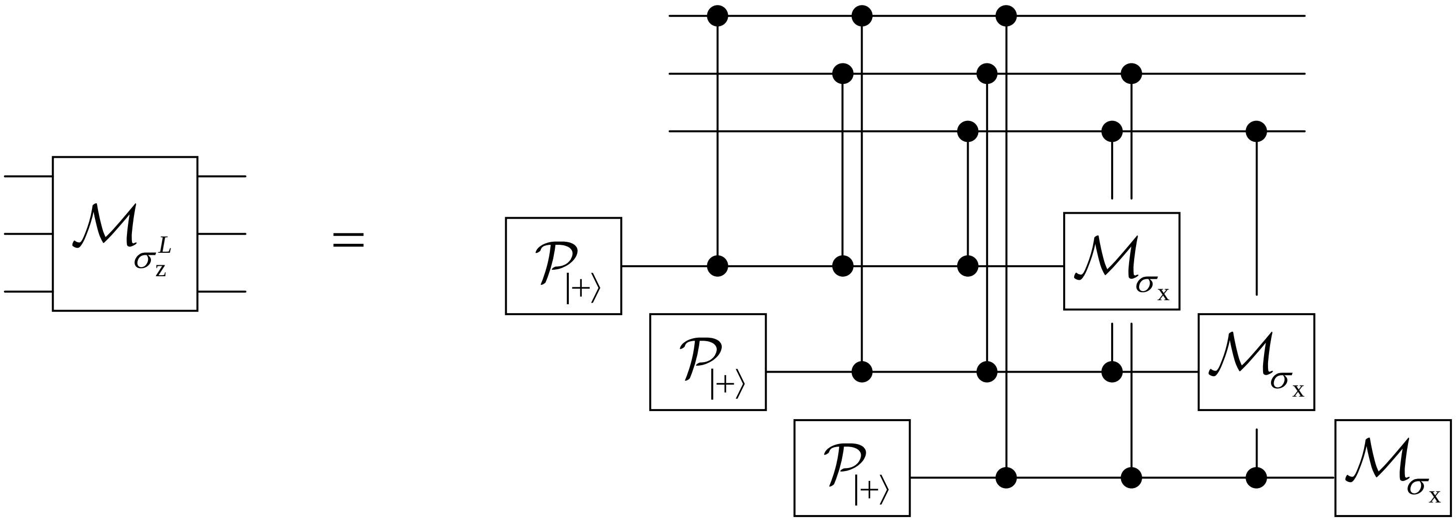

Note that the measurement repetitions can be staggered so that the data qubits are never idle in between consecutive CPHASE gates.

Error Correction by One-Bit Teleportation

Error correction can be performed by teleporting an encoded block. Because we only need to correct Z errors, the error correction gadget can be simplified to an encoded version of "one-bit teleportation":

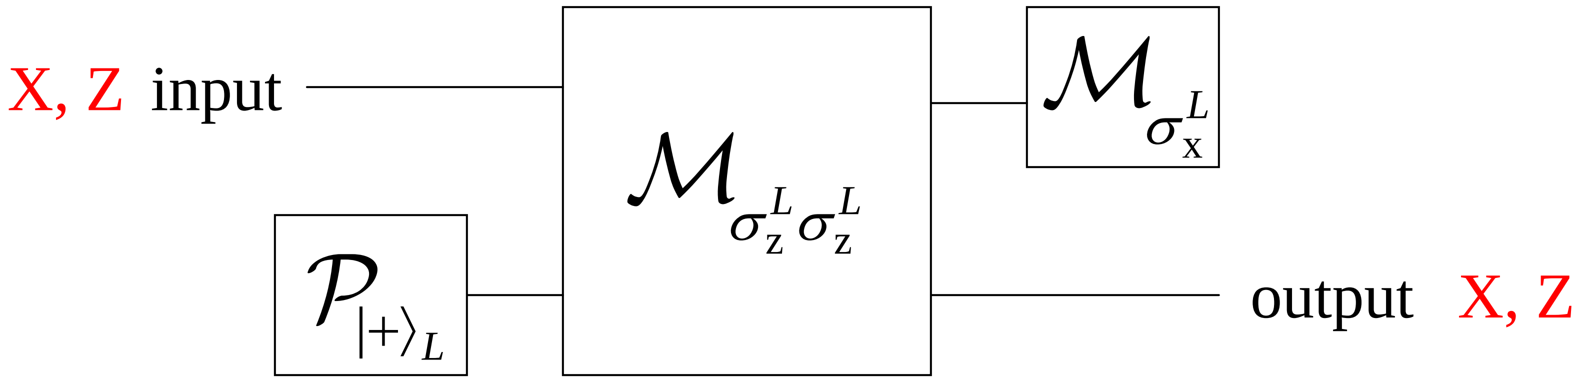

If all operations are ideal, the output matches the input, apart from a possible Z_L (if the outcome of X_L measurement is -1) and a possible X_L (if the outcome of the Z_L measurement is -1). If the n-qubit repetition code is used, then the encoded $Z_L Z_L$ measurement is performed using 2n phase gates and is repeated n times for fault tolerance: if the input block has m Z errors and the error correction gadget has s Z faults, then the outcome of the X_L measurement agrees with the ideal case for m + s $\leq$ (n-1)/2, and the outcome of the $Z_L Z_L$ measurement agrees with the ideal case for s $\leq$ (n-1)/2; furthermore the number of Z errors in the output block is no more than s.

Teleported CNOT gate

By combining one-bit teleportation on both output blocks with the transversal encoded CNOT gate, we obtain this gadget, which corrects errors and executes the gate (up to Pauli operators determined by the measurements):

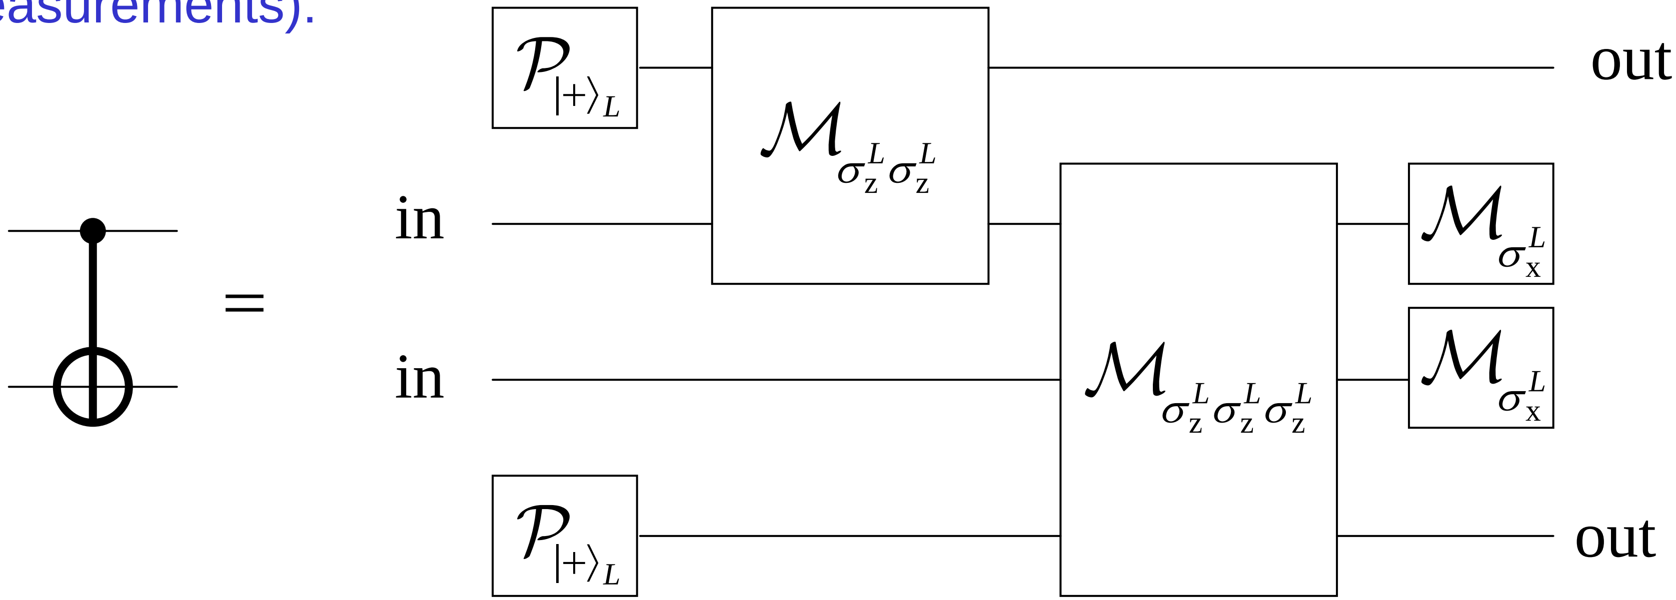

$$XI \longrightarrow XX, \quad IX \longrightarrow IX, \quad ZI \longrightarrow ZI, \quad IZ \longrightarrow ZZ$$

The encoded $Z_L Z_L$ and $Z_L Z_L Z_L$ measurements are performed repeatedly, using ancilla qubits and CPHASE gates. Thus we have succeeded in constructing a fault-tolerant encoded CNOT gate from fundamental CPHASE gates, preparations, and measurements.

Analysis

To analyze the reliability of a circuit constructed from these fault-tolerant gadgets, we propagate faults forward until they reach X measurements. Propagating the Z errors is simple because these commute with the CPHASE gates. A gadget operates correctly if all measurement outcomes agree with the ideal case (the case in which the input blocks have no errors and the gadget contains no faults).

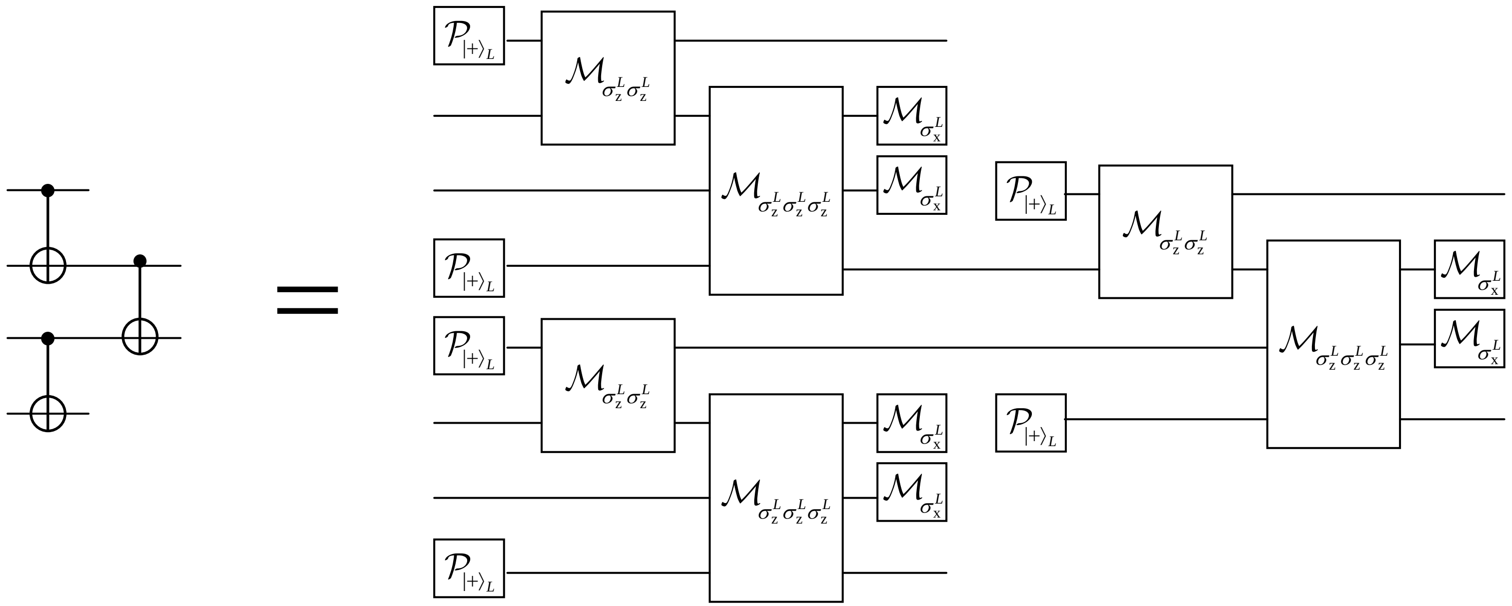

We pessimistically assume that a single unstructured fault in a CPHASE gate causes gadget failure. But for a Z faults to cause failure, there must be at least (n+1)/2 Z errors in a code block, causing measurement of X_L to fail, or else there is at least one Z fault in each of at least (n+1)/2 repeated measurements, causing meaurement of $Z_L Z_L$ or $Z_L Z_L Z_L$ to fail.

Analysis

To bound the probability of failure for the “level-1” encoded CNOT, we add together upper bounds on all the possible failure modes:

$$\varepsilon^{(1)} \leq \varepsilon\left(\mathfrak{M}^{[1]}_{\sigma_x^L}\right) + \varepsilon\left(\mathfrak{M}^{[2]}_{\sigma_x^L}\right) + \varepsilon\left(\mathfrak{M}_{\sigma_z^L\sigma_z^L}\right) + \varepsilon\left(\mathfrak{M}_{\sigma_z^L\sigma_z^L\sigma_z^L}\right) + \varepsilon^{(1)}_{\text{unstructured}}$$

Any of the four encoded measurements might fail, due to multiple Z faults, or an unstructured fault might occur in the gadget.

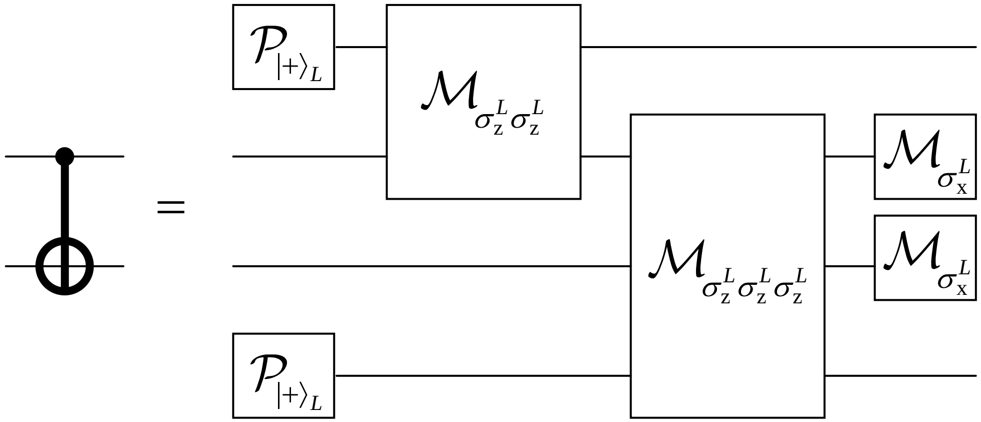

Now we may ask, for what values of ε and ε' will the local stochastic biased noise model yield an encoded CNOT with an effective fault rate below the previously established threshold value of 10^{-3}?

Analysis

Actually, one more idea helps us to improve the result further. Majority voting is used in each encoded measurement, and an encoded error is more likely if the vote is “close”. E.g., for n=5 (which turns out to be optimal), if the vote is 3-to-2, then there must be at least three faults for the majority to be wrong. But if the vote is 4-to-1, then there must be at least four faults for the majority to be wrong.

Therefore, we “flag” blocks with close votes. This information is used at higher levels of the concatenated hierarchy to improve the reliability of decoding. (An error-*detecting* code is used at the next level up, which can *correct* errors that occur at known positions in the code block.)

We then find, that if $\varepsilon'/\varepsilon < 10^{-4}$, we can simulate reliable CNOT gates provided that

$$\varepsilon < 5.1\times 10^{-3}$$

After demanding distillability of high-fidelity states that are needed to complete the universal gate set at the top level of the concatenated code, we obtain a threshold estimate for the local stochastic biased noise model:

$$\varepsilon_0 > 4.7\times 10^{-3}$$

(about a factor of 5 better than the threshold for unbiased noise).

Accuracy threshold for biased noise

We obtain a threshold estimate for the local stochastic biased noise model:

$$\varepsilon_0 > 4.7 \times 10^{-3}$$

(about a factor of 5 better than the threshold for unbiased noise).

This result illustrates how properties of the noise can be exploited to improve fault-tolerance. Perhaps just as important, it shows that fundamental CPHASE gates, along with preparations and measurements, suffice for universal fault-tolerant quantum computation. This feature is appealing from the perspective of some proposed physical implementations (Aliferis talk).

Here geometric constraints have not been considered; we have assumed long-range couplings. For a scheme with local gates, SWAP gates with biased noise are needed.

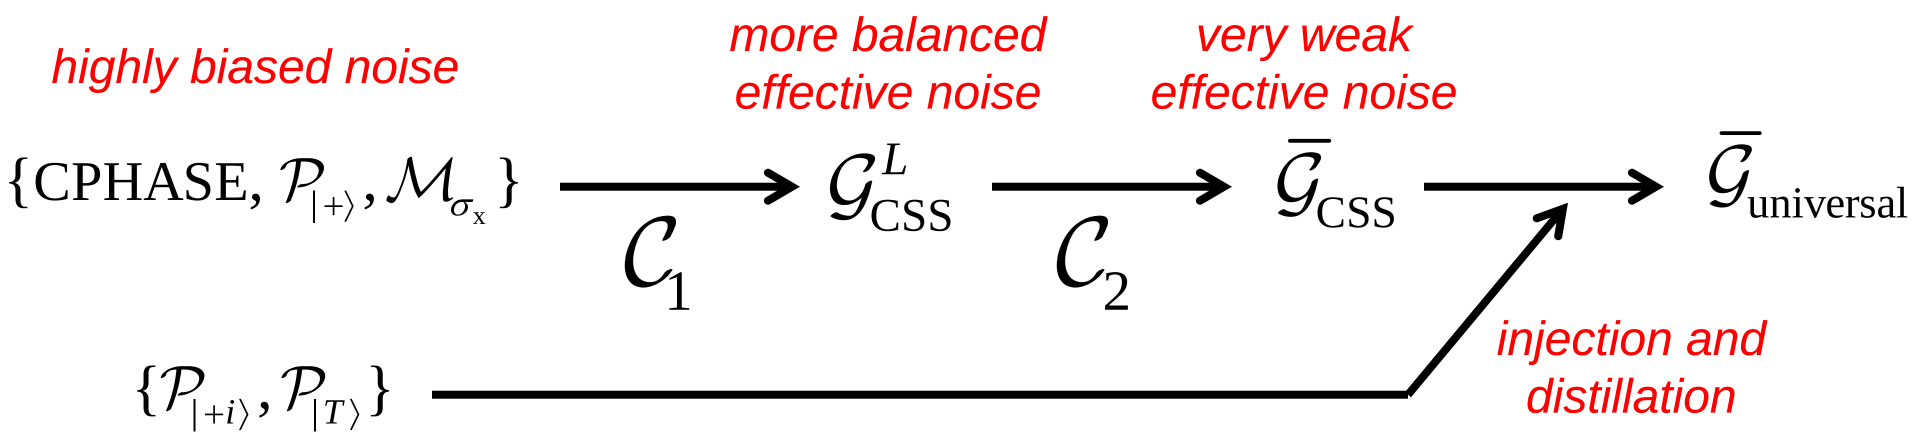

Local non-Markovian noise

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

To analyze the performance of a fault-tolerant protocol, we describe a noisy circuit using a fault path expansion. In each term in the expansion, the circuit locations (i.e. gates) in a particular set are regarded as faulty, while all other locations are assumed to be ideal.

$$\text{Noisy Circuit} = \sum \text{“Fault Paths”}$$

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.

However, distinct fault paths do not necessarily decohere, so it may be more appropriate to assign *amplitudes* rather than probabilities to the fault paths. The unitary operator that describes the joint evolution of the computer and its environment (the “bath”) has a fault path expansion:

$$|\psi\rangle_{SB} = U_{SB}\,|\psi^0_{SB}\rangle = \sum \text{“Fault Paths”}$$

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.

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.

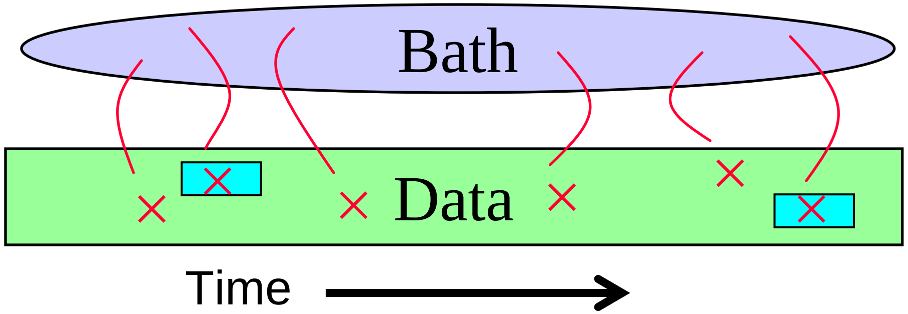

$$\varepsilon = \max \left\| H^{(a)}_{System-Bath} \right\| t_0$$

over all times and locations (max)

time to execute a gate (t_0)

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., $\varepsilon < 10^{-4}$.

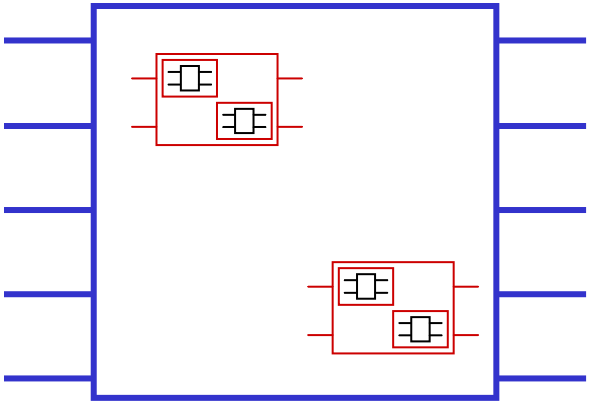

A hierarchy of “gadgets within gadgets” is reliable if the faults are *sparse*.

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.

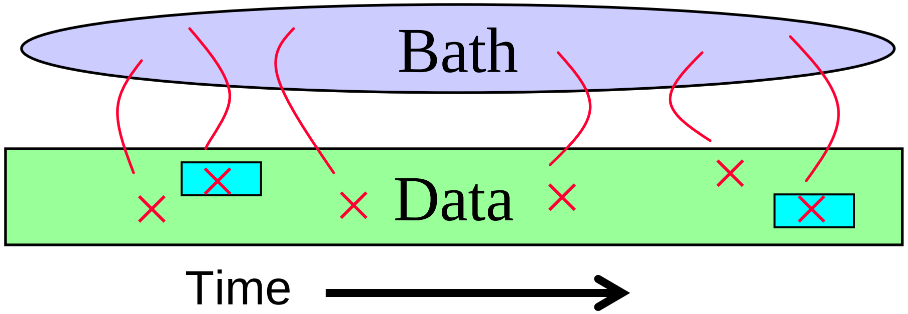

$$\varepsilon = \max \left\| H^{(a)}_{System-Bath} \right\| t_0$$

over all times and locations

time to execute a gate

Local non-Markovian noise

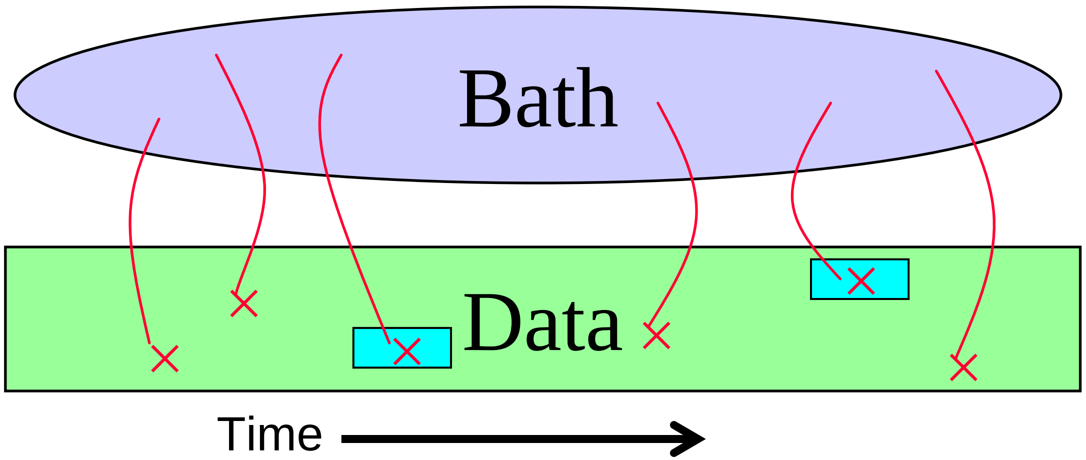

$$\left\| H^{(a)}_{System-Bath} \right\| t_0 < \varepsilon_0 \approx 10^{-4}$$

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

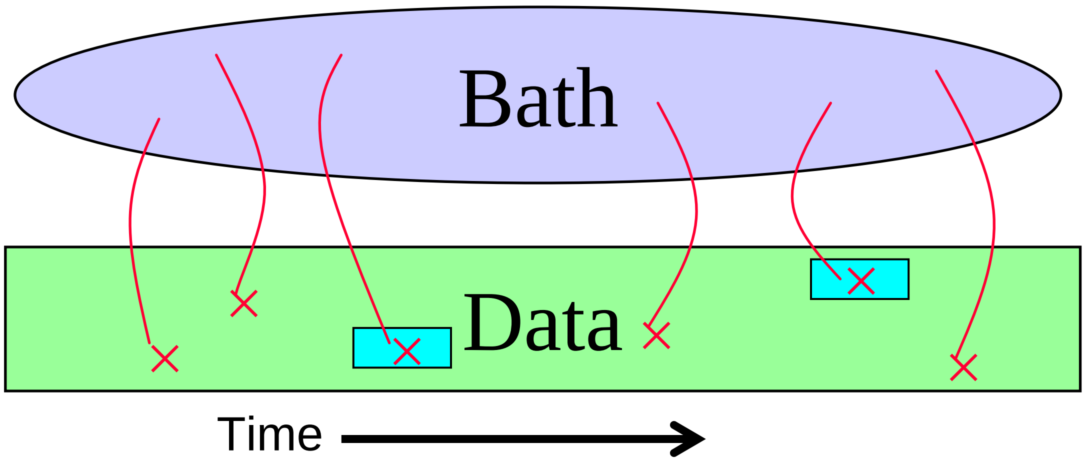

$$\left\| H^{(a)}_{System-Bath} \right\| t_0 < \varepsilon_0 \approx 10^{-4}$$

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

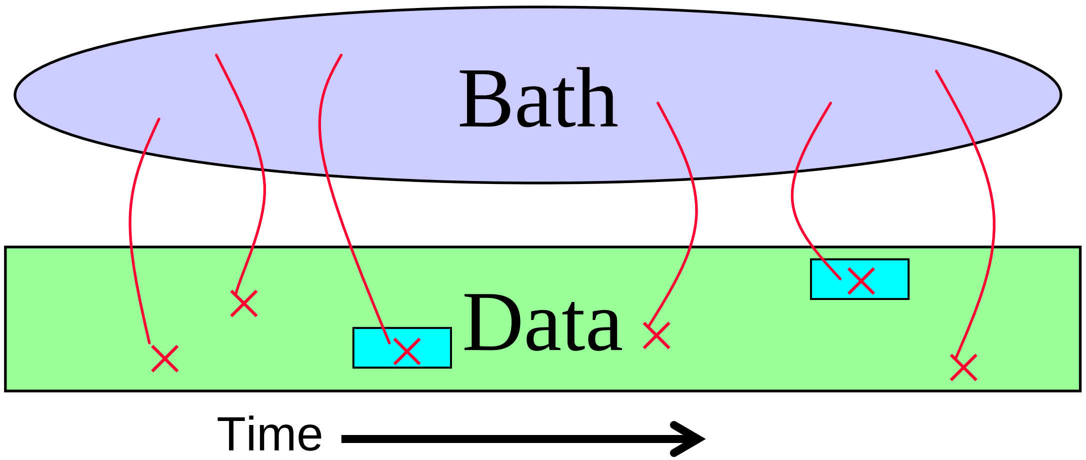

$$\left\| H^{(a)}_{System-Bath} \right\| t_0 < \varepsilon_0 \approx 10^{-4}$$

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 bath of harmonic oscillators. Our task is to estimate the the norm squared of the bad part of the system-bath wave function:

$$\left\| |\psi^{bad}_{SB}\rangle \right\|^2 = \langle \psi^{0}_{SB} | U^{bad\dagger}_{SB} U^{bad}_{SB} | \psi^{0}_{SB} \rangle \le \varepsilon^{2r}$$

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.)

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

("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”).

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

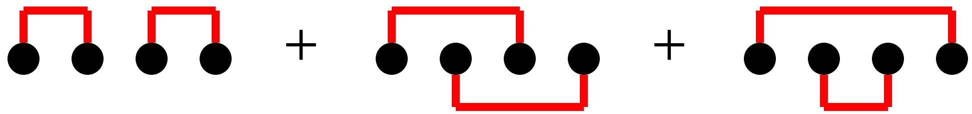

Similarly, the 2n-point correlation function can be expressed as a product of two-point correlators, summed over all possible pairwise contractions.

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

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.

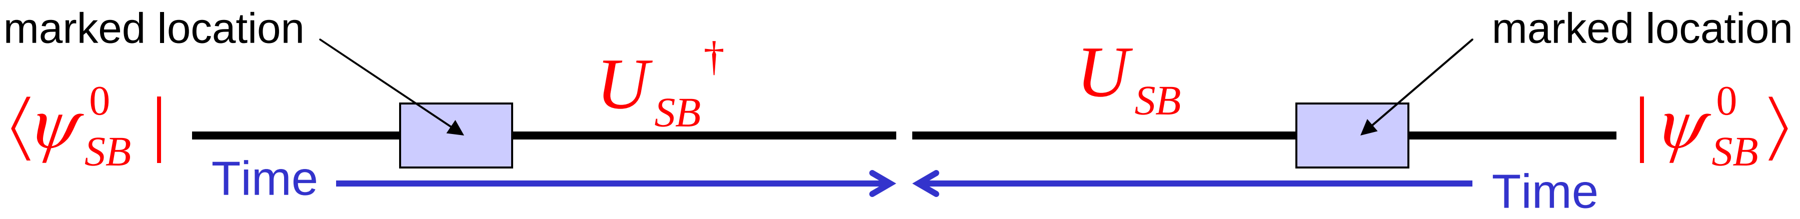

It is convenient to bend this picture into a hairpin shape ("Schwinger-Keldysh diagram")

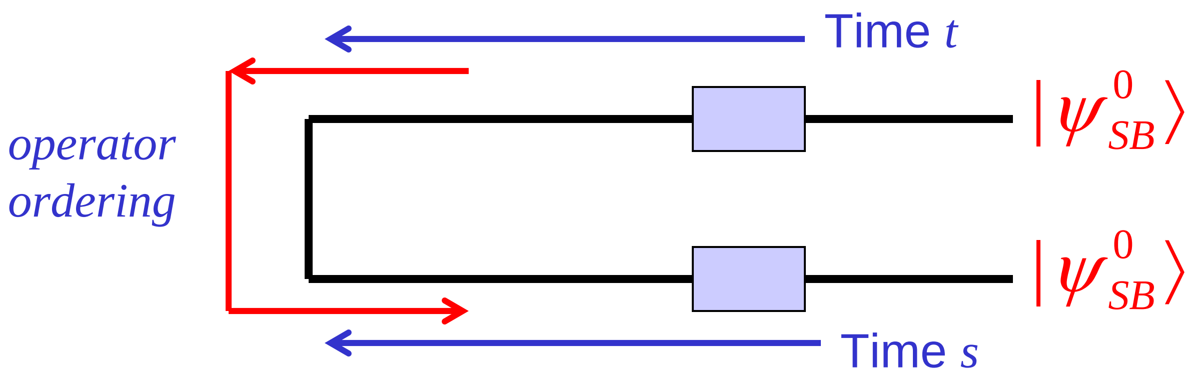

Time increases to the left on both branches, but "time-ordered" operators on the "upper branch" act "before" "anti-time-ordered" operators on the "lower branch".

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.

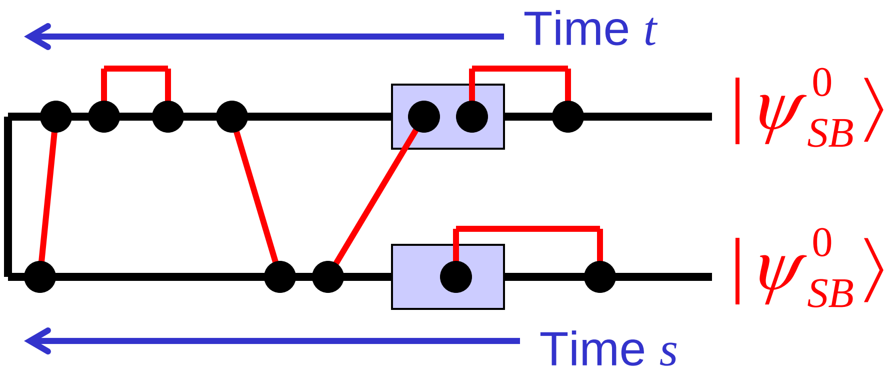

This sum is the norm squared of the bad part of the system-bath state:

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

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)

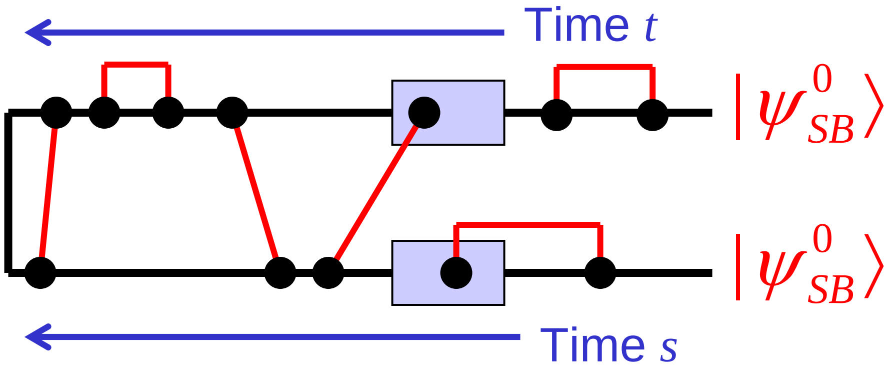

Suppose we fix the *earliest* insertion of the perturbation inside the marked location *on both branches*. These insertions might be contracted with one another; otherwise, each is contracted with another insertion somewhere else. Now we are to:

(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.

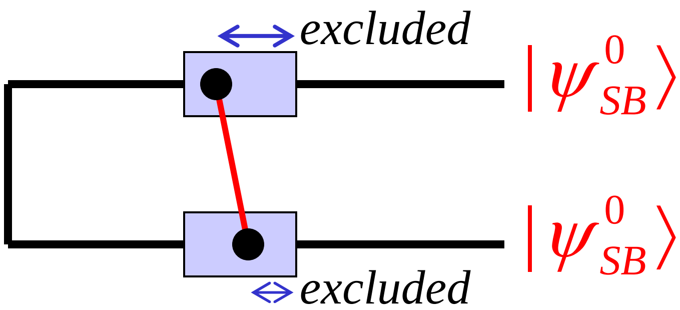

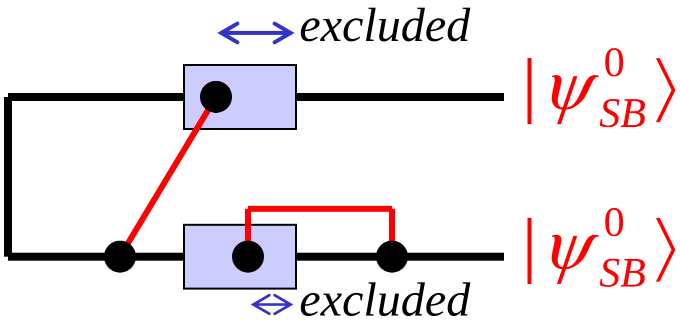

Gaussian noise model

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

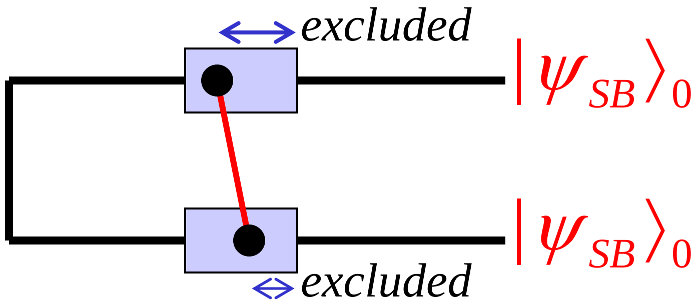

Then, the resummation of all the legal ways to dress this diagram is equivalent to evolving the state using a "hybrid" Hamiltonian, which is

$H_{hybrid} = H_S + H_B$ 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

$$\int_{\square} ds \int_{\square} dt \; {}_S\langle 0 | \sigma_\alpha(s)\sigma_\beta(t) | 0\rangle_S \; \lambda_\alpha(s)\lambda_\beta(t) \, {}_B\langle 0 | \phi_\alpha(s)\phi_\beta(t) | 0\rangle_B$$

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

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:

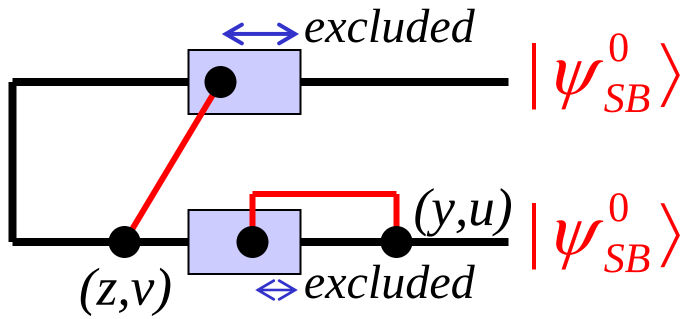

$$\le \int_{\square} ds \int_{all} du \sum_{y} \left| \lambda_\alpha(x,s)\lambda_\gamma(y,u)\, {}_B\langle 0 | \bar{T}\left[\phi_\alpha(x,s)\phi_\gamma(y,u)\right] | 0\rangle_B \right|$$

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

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 $r=2$).

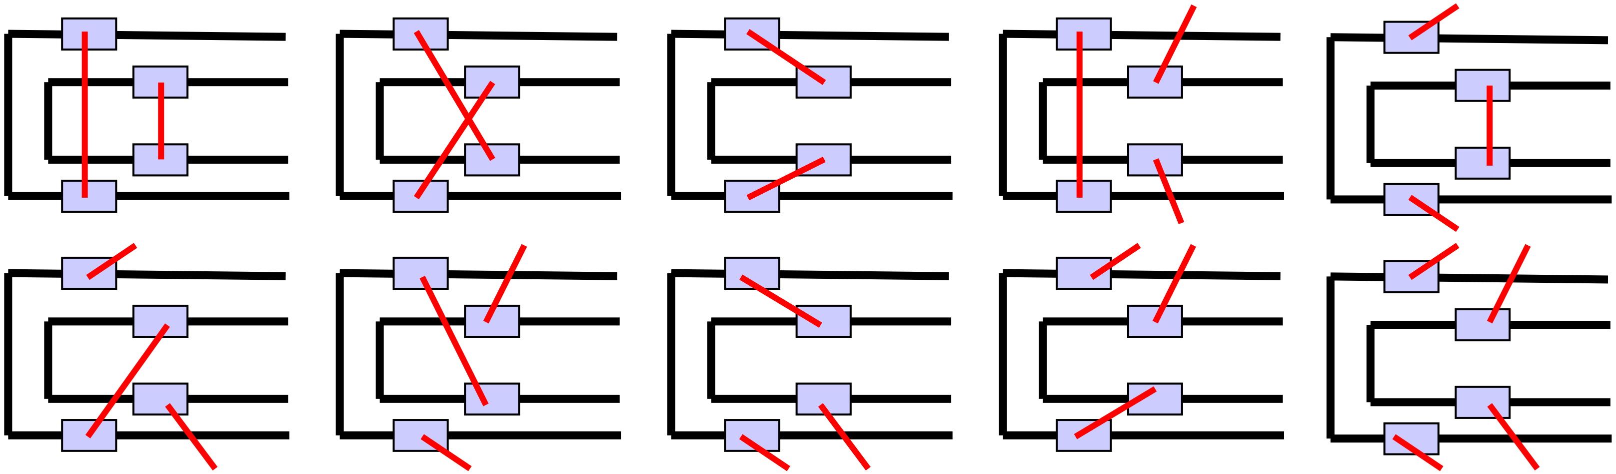

Using the same methods as in AKP05, we can bound the sum of the absolute values of all the diagrams, finding: $\langle \psi^0_{SB} | U^{bad\,\dagger}_{SB} U^{bad}_{SB} | \psi^0_{SB} \rangle \le \varepsilon^{2r}$

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

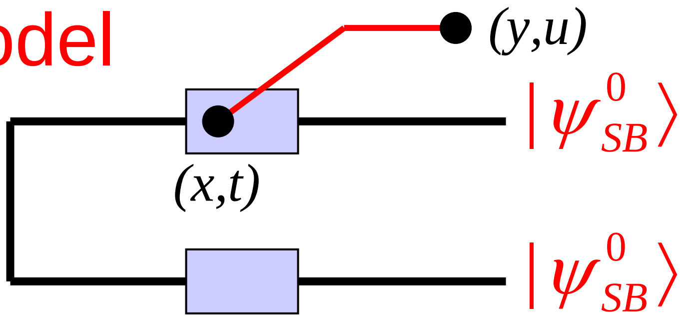

In this noise model, fault-tolerant quantum computing works if ε is small enough (e.g. smaller than 10^{-4}).

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

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 $\displaystyle \int_{all} dt \int_{all} d^D x \frac{1}{\left(x^2 + t^{2/z}\right)^{2\delta}} < \infty$ 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 ideal that fault paths that generate distinct syndrome histories should not be added coherently.

Gaussian noise model

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

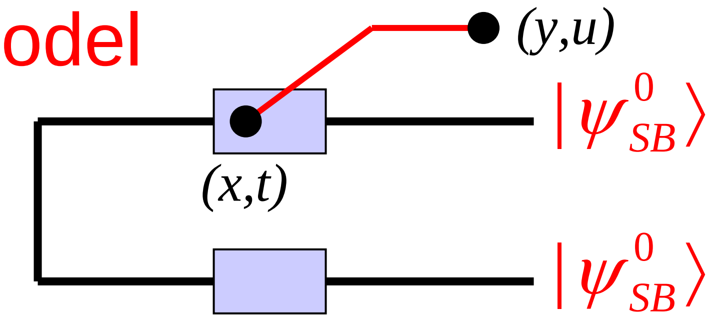

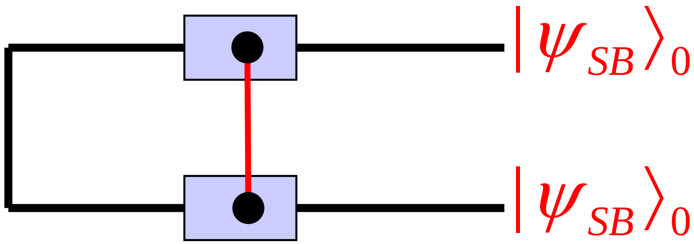

In the Markovian limit, the correlator is a delta function, with support at vanishing time difference:

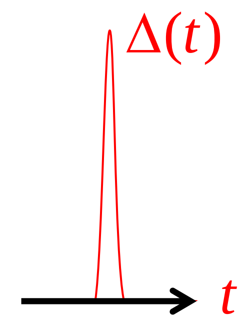

Thus: $\varepsilon = \left(\Gamma t_0\right)^{1/2}$

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,

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

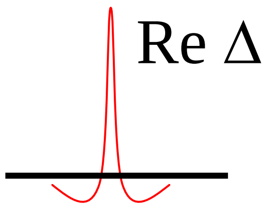

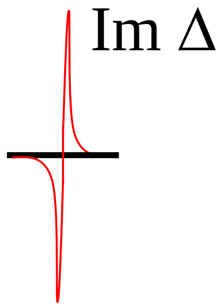

Both the real and the imaginary part of the correlator wiggle, and therefore the integral of the correlator has only a logarithmic sensitivity to the cutoff (cf. Novais et al.).

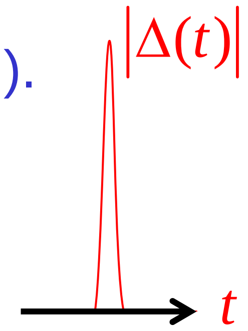

However, unfortunately when we take the absolute value of the correlator, we lose the benefit of the wiggles, and the cutoff dependence is stronger:

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

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

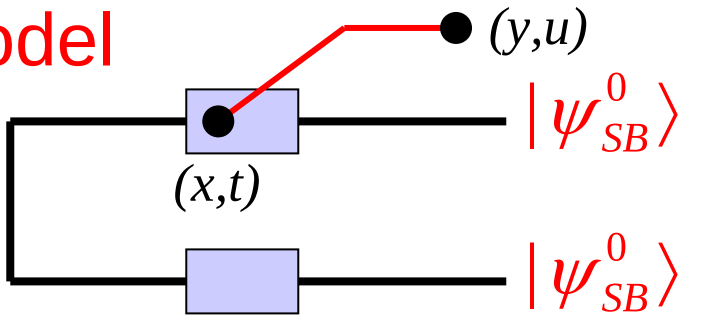

Thus in some cases (like high-temperature Ohmic noise) our new threshold condition for Gaussian noise has no artificial sensitivity to very-high-frequency fluctuations of the bath,

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 plausible 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)?