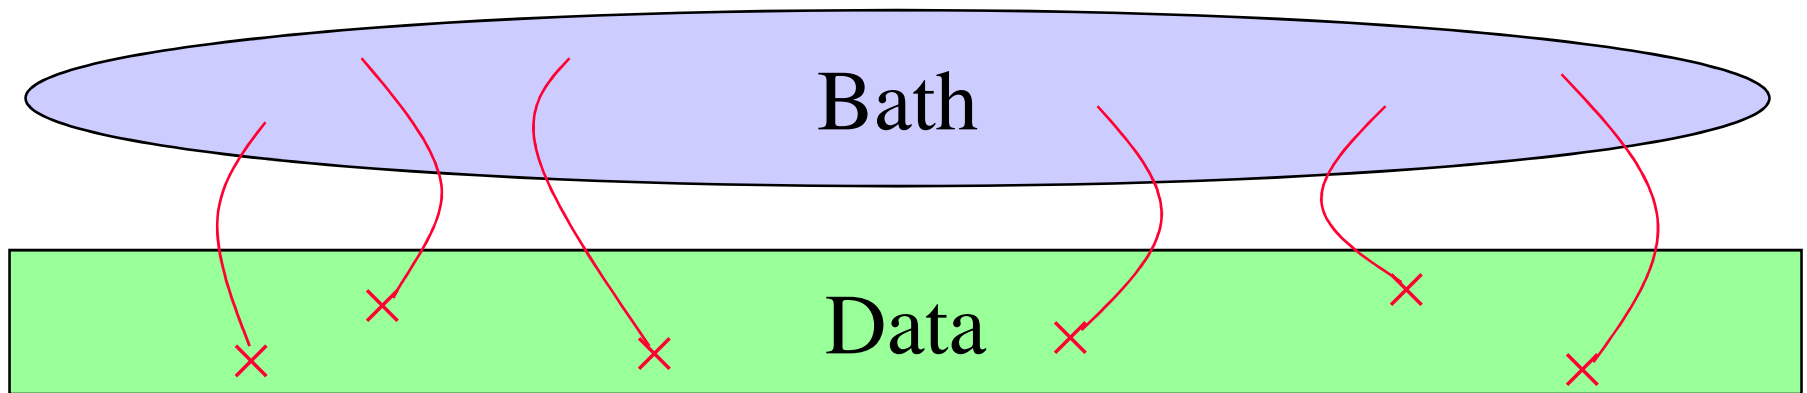
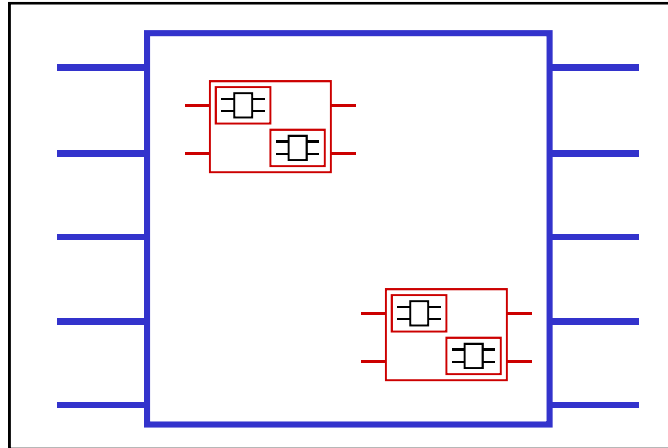
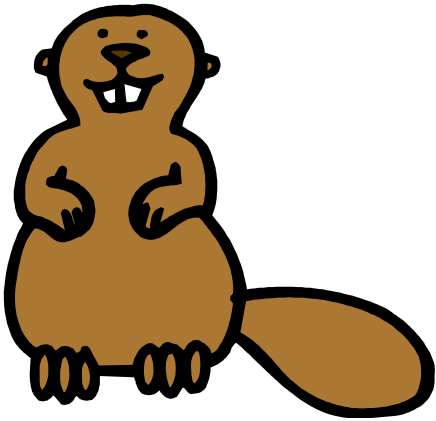


Quantum error correction and fault tolerance



Time →

Institute for Quantum
Information and Matter



John Preskill, Caltech
NSF, 29 September 2012

Truism:

the macroscopic world is classical.

the microscopic world is quantum.

Goal of QIS:

controllable quantum behavior in scalable systems

Why?

Classical systems cannot simulate quantum systems efficiently (a widely believed but unproven conjecture).

But to control quantum systems we must slay the dragon of decoherence ...

Is this merely *really, really hard*?

Or is it *ridiculously* hard?

Why quantum computing is hard

We want qubits to interact strongly with one another.

We don't want qubits to interact with the environment.

Until we measure them.

Quantum Hardware



Schoelkopf



Martinis



Yacoby

Two-level ions in a Paul trap, coupled to “phonons.”

Superconducting circuits with Josephson junctions.

Electron spin (or charge) in quantum dots.

Cold neutral atoms in optical lattices.

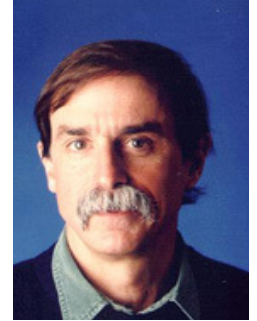
Two-level atoms in a high-finesse microcavity, strongly coupled to cavity modes of the electromagnetic field.

Linear optics with efficient single-photon sources and detectors.

Nuclear spins in semiconductors, and in liquid state NMR.

Nitrogen vacancy centers in diamond.

Anyons in fractional quantum Hall systems, quantum wires, etc.



Wineland



Blatt



Marcus

Quantum computer: the standard model

- (1) Hilbert space of n qubits: $\mathfrak{H} = \mathbb{C}^{2^n}$
- (2) prepare initial state: $|0\rangle^{\otimes n} = |000\dots 0\rangle$
- (3) execute circuit built from set of universal quantum gates: $\{U_1, U_2, U_3, \dots, U_{n_G}\}$
- (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).

Quantum fault tolerance

The goal of fault-tolerant quantum computing is to operate a large-scale (quantum) computer reliably, even though the components of the computer are noisy.

Reliability can be enhanced by encoding the computer's state in the blocks of a quantum error-correcting code. Each “logical” qubit is stored nonlocally, shared by many physical qubits, and can be protected if the noise is sufficiently weak and also sufficiently weakly correlated in space and time.

Two central questions are:

- 1) For what noise models does fault-tolerant quantum computing work effectively?
- 2) For a given noise model, what is the overhead cost of simulating an ideal quantum computation with noisy hardware?

Quantum fault tolerance

To *really* operate a large-scale quantum computer, many implementation-specific systems engineering issues will need to be addressed.

Though I am a theoretical physicist, not an engineer, I have devoted much of my research effort since the mid-1990s to quantum fault tolerance, because I believe that this subject raises questions and stimulates insights that are of broad and fundamental interest in quantum information science.

Whatever the applications turn out to be, the quest for a large-scale quantum computer is one of the grand scientific challenges of the 21st century.

Quantum error correction

Protect not just against bit flips, but also against the environment “watching the computer,” so that computational paths can interfere.

If a quantum computation works, and you ask the quantum computer later what it did, it should answer: “I forget..”

The computation is encrypted, i.e. hidden from the environment. (Not the answer, which is classical, but the path followed by the computer to reach the answer.)

And even a properly “encrypted” computation may fail, unless the gates are sufficiently accurate.

Irony: Macroscopic systems are usually highly vulnerable to decoherence, but we can protect information better by encoding it nonlocally, in a “macroscopic” memory.

Quantum error correction (and topological order)

A “logical qubit” is encoded using many “physical qubits.” We want to protect the logical qubit, with orthonormal basis states $|0\rangle$ and $|1\rangle$, from a set of possible error operators $\{ E_a \}$.

For protection against bit flips:

$$E_a |0\rangle \perp E_b |1\rangle .$$

For protection against phase errors:

$$E_a (|0\rangle + |1\rangle) \perp E_b (|0\rangle - |1\rangle) .$$

In fact, these conditions suffice to ensure the existence of a recovery map.

It follows that

$$\langle 0 | E_b^\dagger E_a | 0 \rangle = \langle 1 | E_b^\dagger E_a | 1 \rangle .$$

Compare the definition of topological order: if V is a (quasi-)local operator and $|0\rangle, |1\rangle$ are ground states of a local Hamiltonian, then

$$\langle 1 | V | 0 \rangle = 0, \text{ and } \langle 0 | V | 0 \rangle = \langle 1 | V | 1 \rangle .$$

up to corrections exponentially small in the system size. (Ground states are locally indistinguishable.)

Scalable quantum computing

Quantum Accuracy Threshold Theorem: Consider a quantum computer subject to **quasi-independent 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 $1 - \delta$, where, for some constant c ,

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

Aharonov, Ben-Or

Kitaev

Laflamme, Knill, Zurek

Aliferis, Gottesman, Preskill

Reichardt

assuming:

parallelism, fresh qubits (*necessary* assumptions)

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

“Practical” considerations:

Resource requirements, systems engineering issues

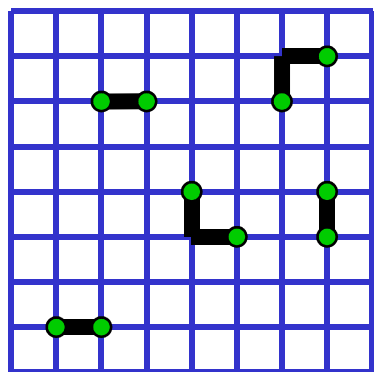
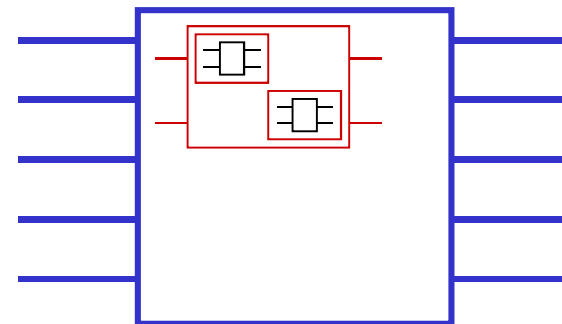
Matters of “principle”:

Conditions on the noise model, what schemes are scalable, etc.

Accuracy Threshold

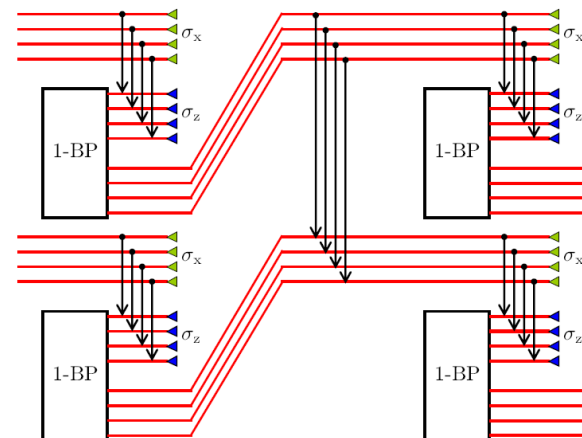
Accuracy threshold theorems have been proved for three types of fault-tolerant schemes:

Recursive: hierarchy of gadgets within gadgets, with logical error rate decreasing rapidly with level.



Topological: check operators are local on a two-dimensional surface, and detect the *boundary* points of error *chains*. Logical error rate decays exponentially with block's linear size.

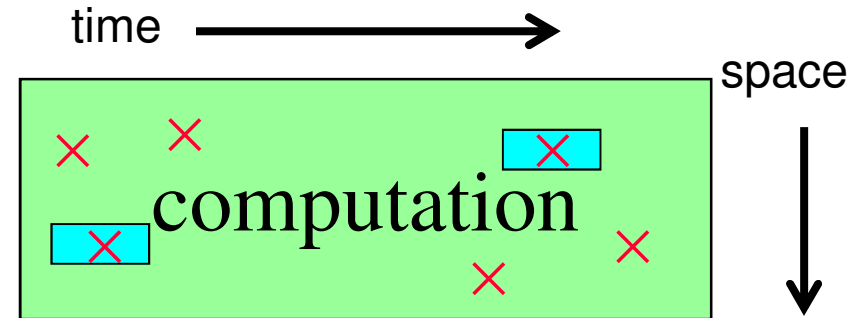
Teleported: Encoded Bell pairs are prepared recursively, but used only at the top level. The (quantum) depth blowup of the simulation is a constant factor.



Noise models

Two types of noise models are most commonly considered in rigorous estimates of the accuracy threshold.

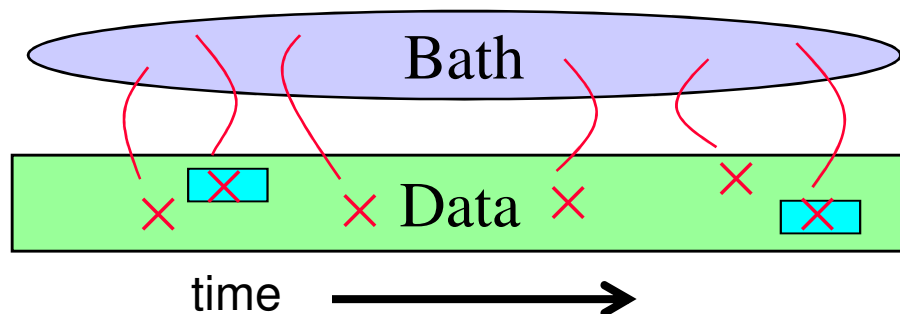
In the **local stochastic noise model**, “fault paths” are assigned *probabilities*. For any set of r gates in the circuit, the probability that all r of the gates have faults is no larger than ϵ^r .



The threshold theorem shows that fault-tolerance works for $\epsilon < \epsilon_0$. Though not fully realistic, these models provide a useful caricature of noise in actual devices, and can be compared with simulations.

In more realistic **Hamiltonian noise models**, fault paths can add *coherently*. The joint dynamics of the system and “bath” is determined by a Hamiltonian

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



that acts *locally* on the system. Fault tolerance works if the system-bath coupling responsible for the noise is sufficiently weak.

Non-Markovian noise

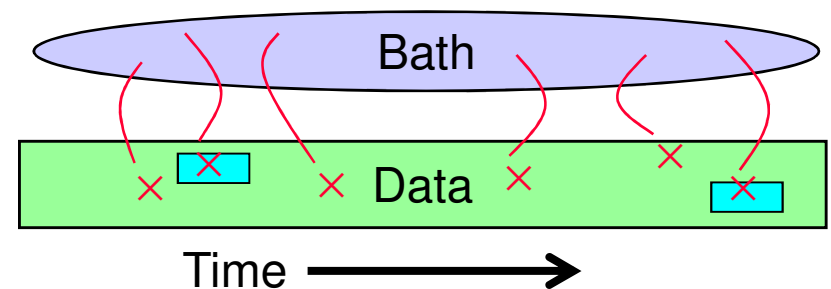
Terhal, Burkard 2005; Aliferis, Gottesman, Preskill 2006; Aharonov, Kitaev, Preskill 2006; Ng, Preskill 2009

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

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

where

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



Threshold condition can be formulated as $\varepsilon \leq \varepsilon_0 \cong 10^{-4}$, where *noise strength* ε can be defined in either of two ways:

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

over all times
and locations

gate execution time

Internal bath dynamics can
be strong and nonlocal

$$\varepsilon = \max \left(\int_{1, \text{circuit location}} \int_{2, \text{all spacetime}} |\Delta_{Bath}(1, 2)| \right)^{1/2}$$

bath correlation function

applies for a Gaussian
(harmonic oscillator) bath

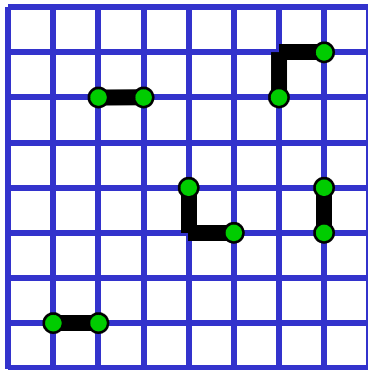
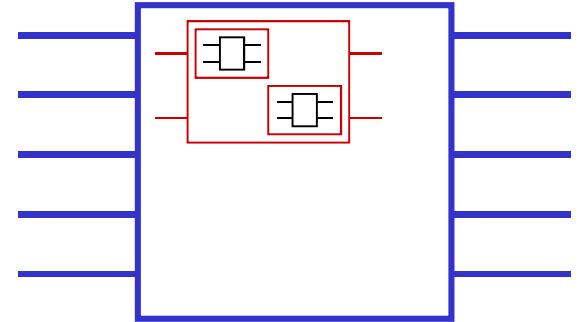
In either scenario, noise Hamiltonian is assumed to act locally on the system

Accuracy Threshold

Some threshold estimates for stochastic noise:

Recursive: $\epsilon_0 > 1.94 \times 10^{-4}$ proven for local stochastic noise using “Bacon-Shor codes.

-- Aliferis, Cross

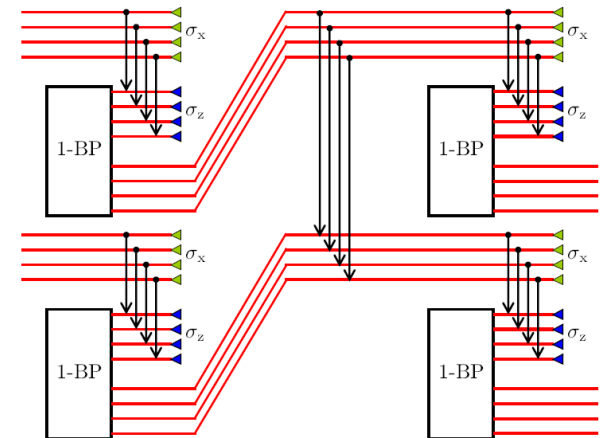


Topological: $\epsilon_0 \sim 7.5 \times 10^{-3}$ estimated for independent depolarizing noise in a *local two-dimensional* measurement-based scheme (combination of numerics and analytic argument).

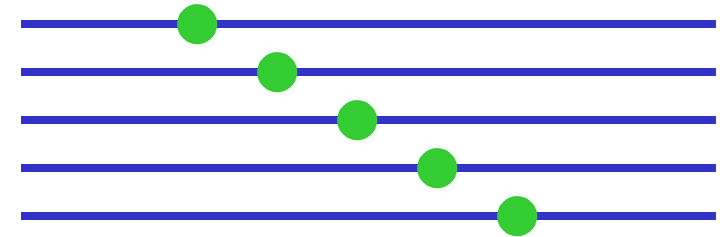
-- Raussendorf, Harrington, Goyal

Teleported: $\epsilon_0 > 6.7 \times 10^{-4}$ proven for local stochastic noise using concatenated error-detecting codes ($\epsilon_0 > 1.25 \times 10^{-3}$ for depolarizing noise); simulations indicate $\epsilon_0 \sim 1 \times 10^{-2}$ for depolarizing noise.

-- Knill; Aliferis, Preskill



Limitations on transversal (local unitary) logical gates



The logical gates close to the identity that can be executed with local unitary transformations form a (perhaps trivial) Lie algebra.

$$U = I + \mathcal{E}(A_1 + A_2 + \cdots + A_n)$$

If the code can “detect” a weight-one error, then for each i :

$$\Pi A_i \Pi \propto \Pi$$

(Π = projector onto code space)

If U preserves the code space,
then U acts trivially on code space:

$$U \Pi = \Pi U \Pi = \Pi$$

There are no transversal gates close to the identity. The group generated by transversal gates is finite and hence nonuniversal. Which logical gates can be executed transversally depends on what code we use.

Limitations on constant-depth logical gates in topological stabilizer codes

Topological stabilizer code: check operators are geometrically local Pauli operators; code distance is “macroscopic.”

What logical gates can be executed using *constant depth* circuits of geometrically local gates, which are inherently fault-tolerant (each fault affects $O(1)$ qubits).

$$2D : U(Pauli)U^{-1} \in Pauli, \quad U \in C_2 \quad (\text{Clifford})$$

$$3D : U(Pauli)U^{-1} \in C_2, \quad U \in C_3$$

etc.

Applies to only a special class of quantum codes, but to a broader class of protected gates than transversal gates.

Transversal gates are topological

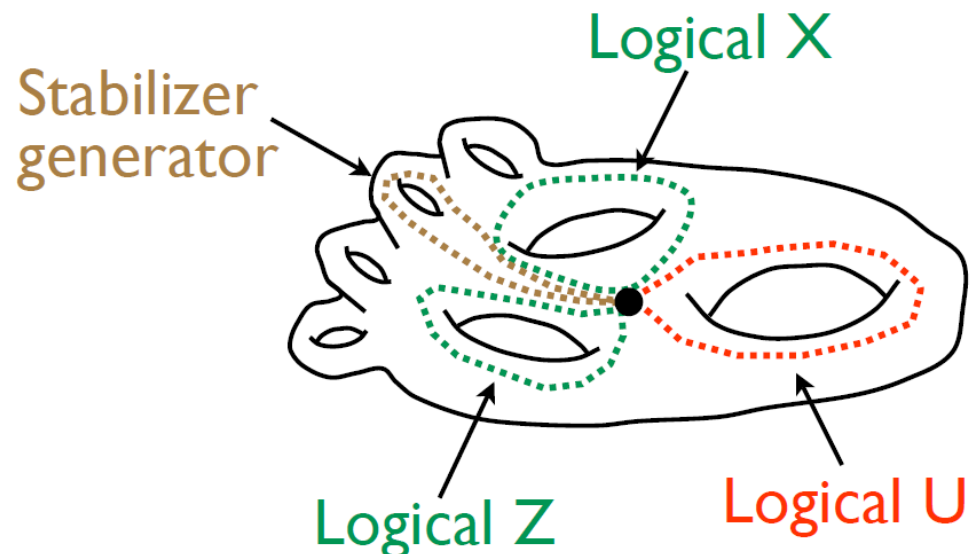
Under local unitaries, an initial K -dimensional quantum code gets mapped to a new K -dimensional code. The family of LU-equivalent codes is a manifold, the base space of a fiber bundle with a $U(K)$ structure group. (LU defines a notion of parallel transport of the code space.)

A transversal gate corresponds to a loop in the base space, where we return to the original code, but with a nontrivial “twist.”

But small loops define trivial gates (no transversal logical gates close to the identity). Hence the connection is *flat*.

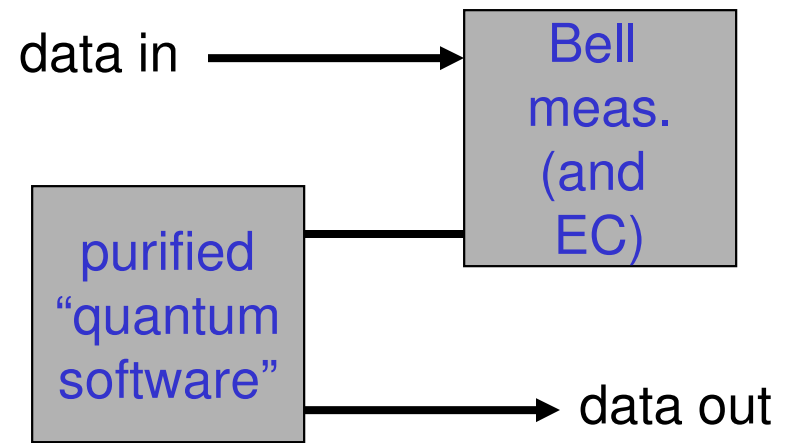
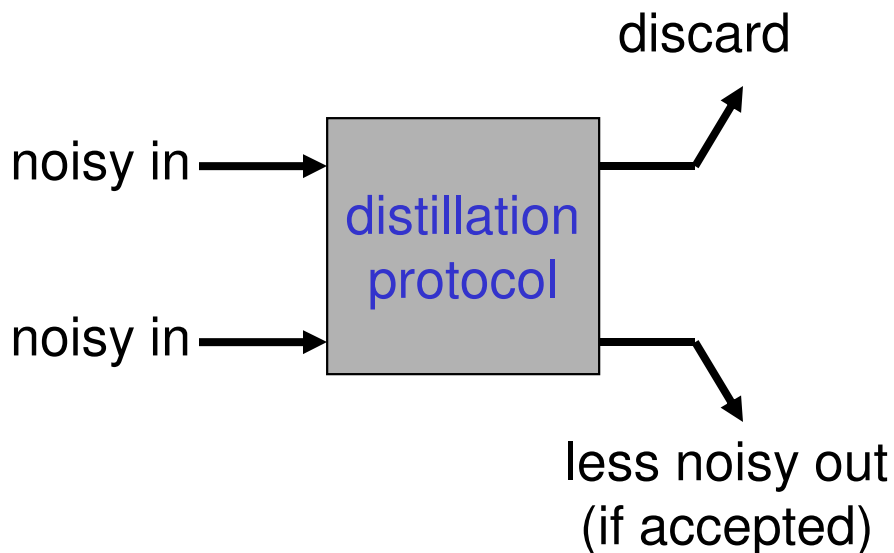
A transversal gate is executed by traversing a noncontractible path on the manifold of LU-equivalent codes.

Gottesman & Zhang 2012.



Gate teleportation and state distillation

In fault-tolerant schemes, a version of quantum teleportation is used to complete a universal set of protected quantum gates. Suitable “quantum software” is prepared and verified offline, then measurements are performed that transform the incoming data to outgoing data, with a “twist” (an encoded operation) determined by the software.



Reliable software is obtained from noisy software via a multi-round state distillation protocol. In each round (which uses CNOT gates and measurements), there are n noisy input copies of the software of which $n-k$ copies are destroyed. The remaining output k copies, if accepted, are less noisy than the input copies

Gottesman, Chuang; Knill; Bravyi, Kitaev

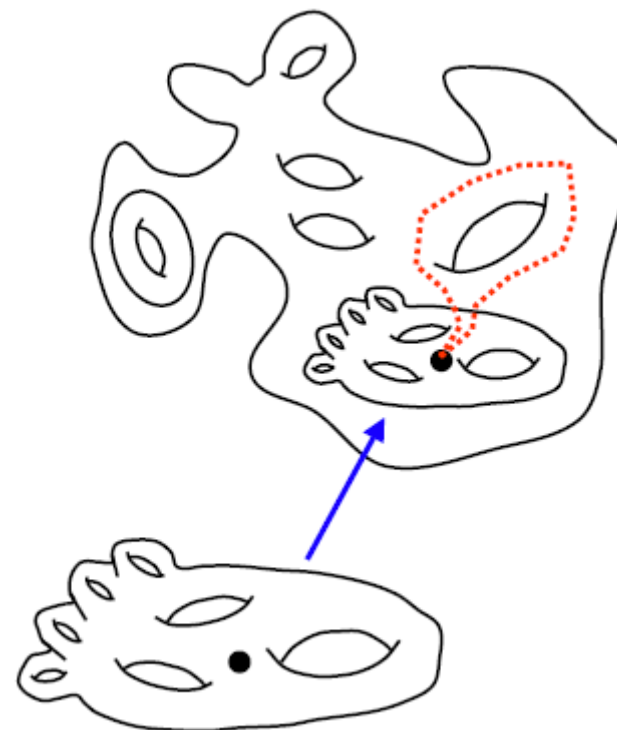
Beyond transversal gates?

For gates needed to complete a universal set, we need to go outside the manifold of LU-equivalent codes.

We embed the manifold in a larger manifold, and traverse a noncontractible path in the larger manifold.

We should be able to understand ancilla constructions including magic state distillation in this framework.

In general, our gate set includes gates that are constructed by temporarily leaving the manifold of LU-equivalent codes, and then returning to it eventually.



Gottesman & Zhang 2012.

Self-Correcting Quantum Memory?

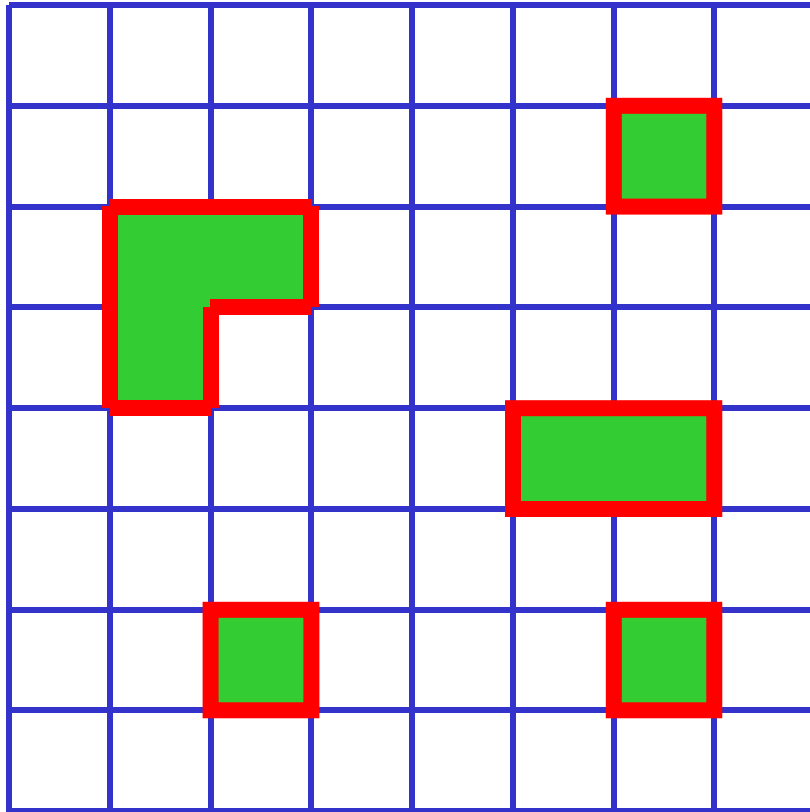
Example: 1D ferromagnet (repetition code)



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 temperature. A one-dimensional ferromagnet is not a robust (classical) memory.

2D ferromagnet (repetition code)



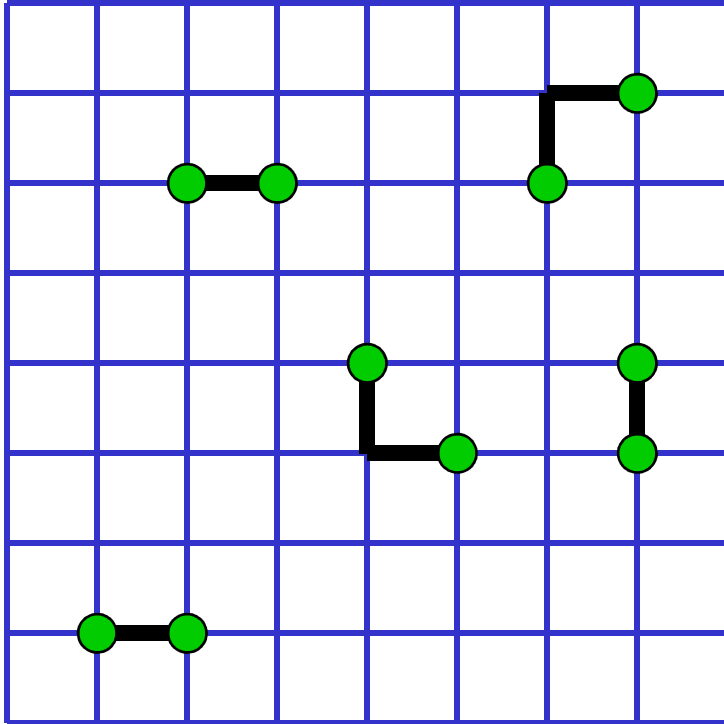
This memory is a repetition code, but with redundant (hence robust) parity checks.

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 Boltzmann factor $\exp(-L / T)$, and are rare.

The storage time for classical information becomes exponentially long when the block size is large. (Actual storage media, which are robust at room temperature, rely on this physical principle.)

Topological Code



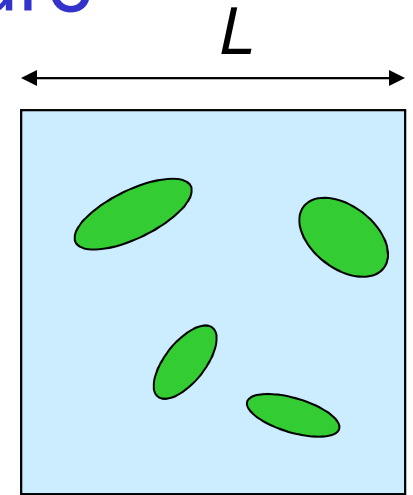
A topological medium in 2D is similar to the 1D Ising model: pairs of anyons are produced by thermal fluctuations at a rate that does not depend on the system size. These anyons can then diffuse apart without any additional energy cost. When anyons diffuse a distance comparable to the distance between pairs, logical errors arise.

Therefore, thermal fluctuations disorder the system at any nonzero temperature. A two-dimensional topological medium is not self-correcting quantum memory.

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 (Dennis et al. 2002).



Question: Is “finite-temperature topological order” possible in less than 4D?

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.

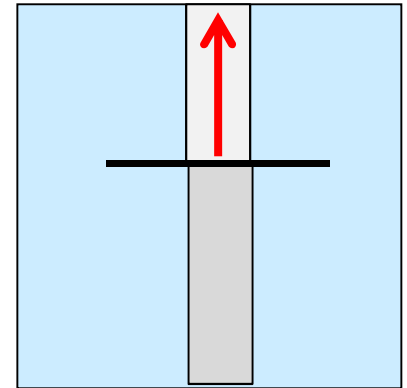
The 3D toric code is a self-correcting *classical* memory, which unlike the 2D Ising model, is stable with respect to an applied “magnetic field.”

But the 3D toric code is not a self-correcting quantum memory.

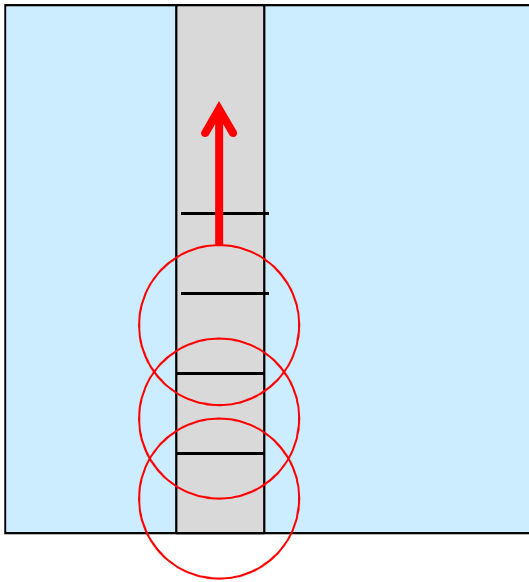
Self-correction in two dimensions?

Local commuting projector code:

The code is the simultaneous eigenspace of a set of commuting projectors. Again, there is a “string logical operator,” only “slightly entangling” across a cut through the string. (Haah-Preskill, Bravyi-Poulin-Terhal, Bravyi-Terhal).



Logical error due to particle sliding along strip?



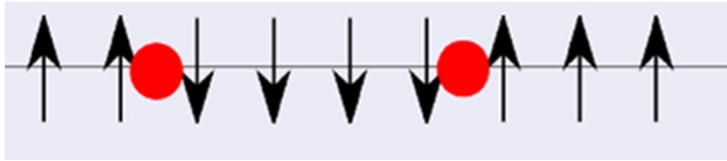
Build the string operator:

Divide the strip into constant-length segments. Each time we wish to extend the string, “twirl” the next segment and project. If we fail, twirl and try again.

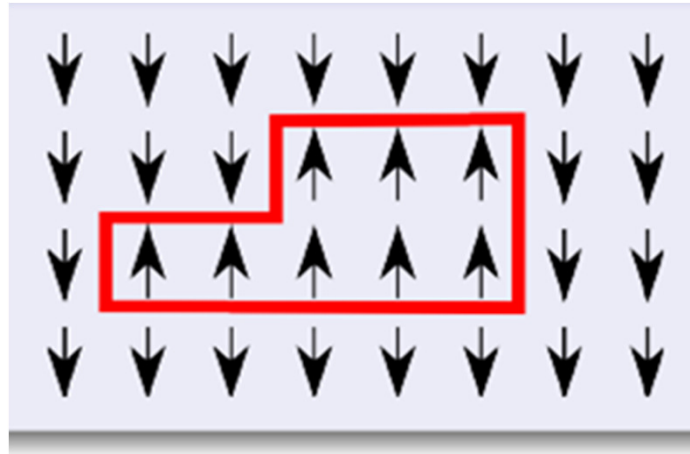
If the code satisfies the “local topological order” condition (proposed to ensure stability with respect to generical perturbations), then the projection succeeds with constant probability (no “blind alleys”).

Landon-Cardinal & Poulin 2012

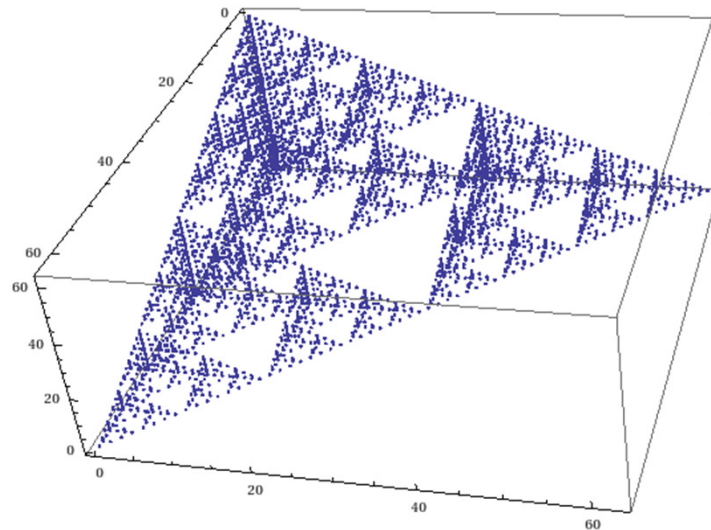
Excitations in local classical and quantum codes



Mobile pointlike excitations:
1D Ising model, 2D toric code



No pointlike excitations:
2D Ising model, 4D toric code

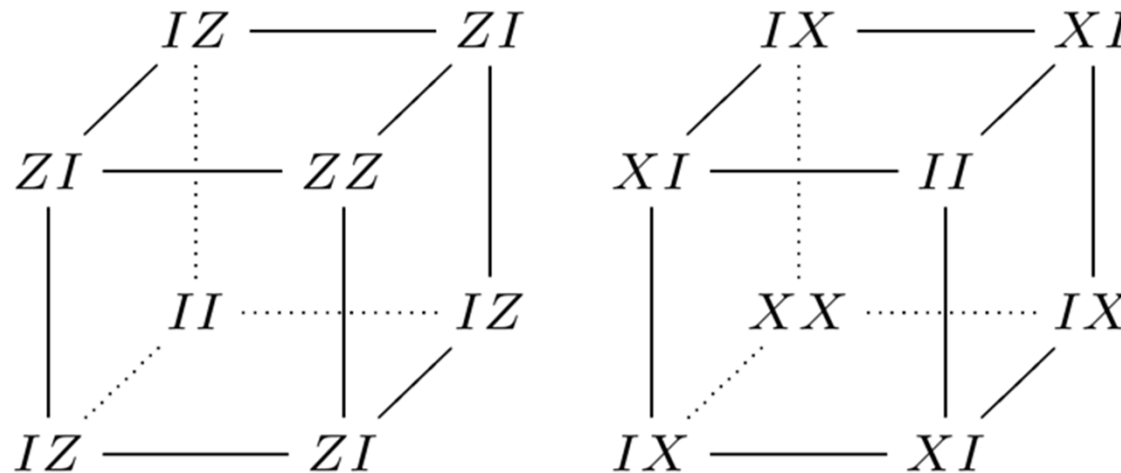


Immobile pointlike excitations:
3D Haah code (2011).

There are mobile point defects in any
“*scale-invariant*” translation-invariant 3D
stabilizer code (Yoshida 2011).

Haah's code

Haah 2011



A local stabilizer code with two qubits per site on a simple cubic lattice.

Two stabilizer generators on each cube.

No logical string operators.

Code distance grows faster than linearly with linear system size L .

The barrier height for a logical error is $O(\log L)$.

Topologically ordered: code states look the same locally.

Equilibrates slowly when cooled from high to low temperature (glass).

For weak noise, annealing corrects errors with high success probability.

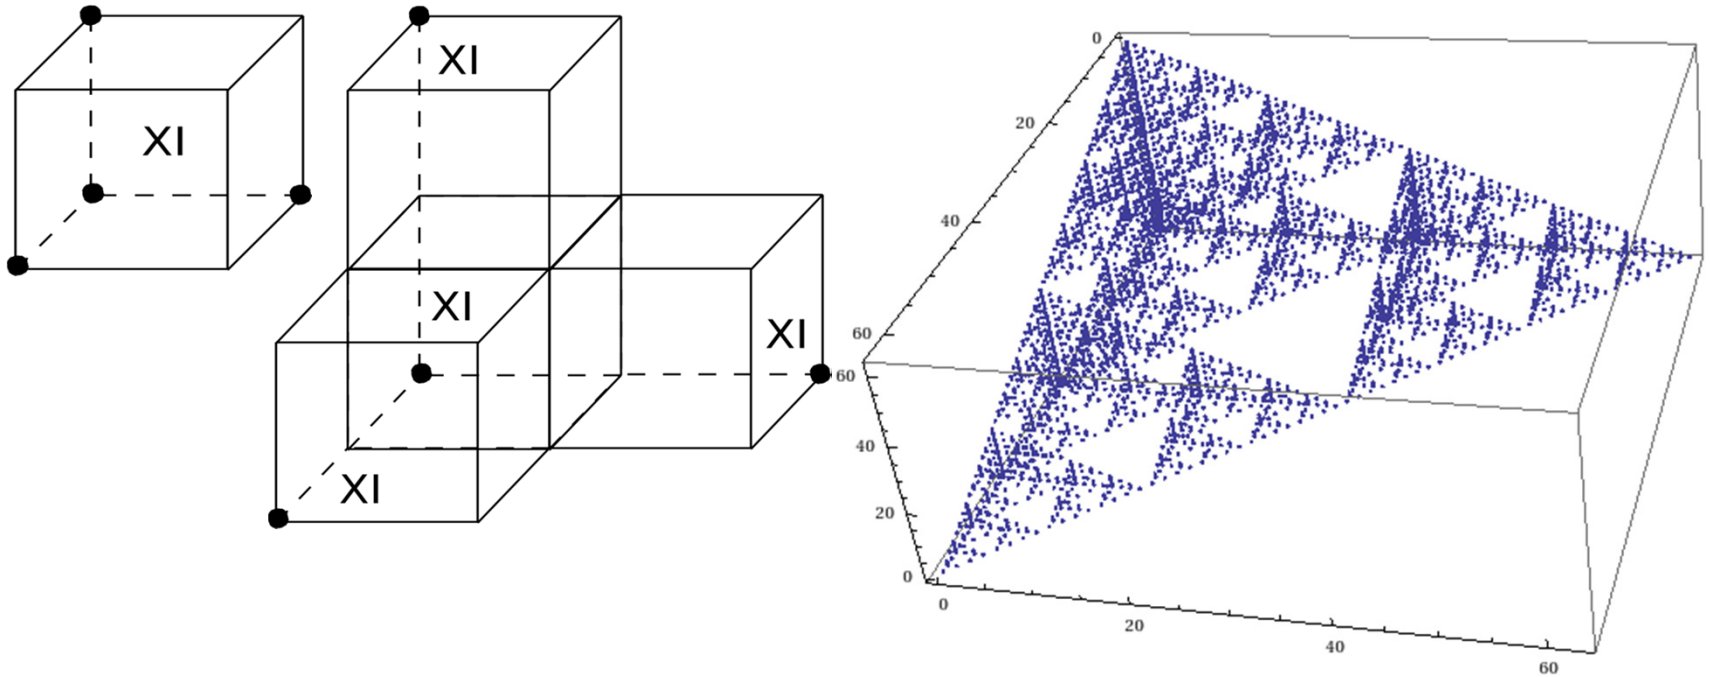
Degeneracy on $L \times L \times L$ torus

L	k	L	k
2	6	3	2
4	14	5	2
6	6	7	2
8	30	9	2
10	6	11	2
12	14	13	2
14	6	15	50
16	62	17	2

$$\frac{k+2}{4} = \deg_x \gcd [1 + (1+x)^L, 1 + (1+tx)^L, 1 + (1+t^2x)^L]_{\mathbb{F}_4}$$

where $t^2 + t + 1 = 0$.

Isolated charges

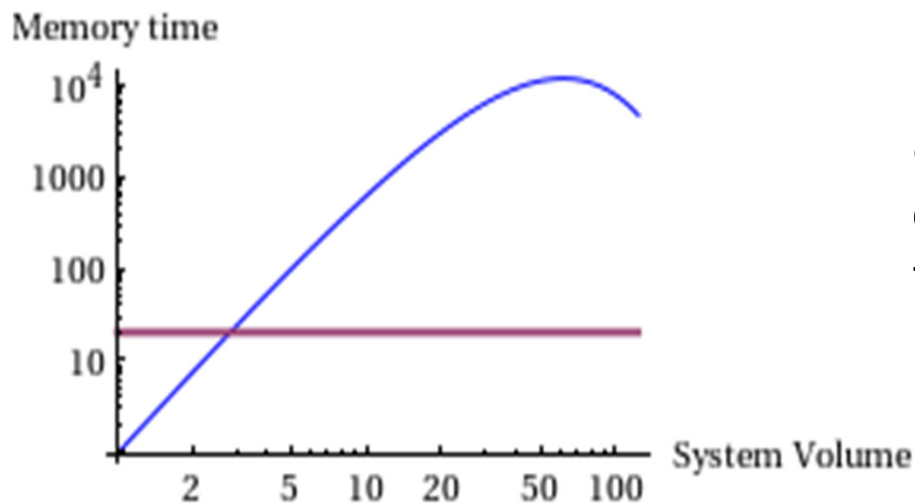


A local process starting from the “vacuum” (no excitations) and arriving at a state where a single topological defect is isolated from all others by distance at least R , must pass through a state whose “energy” is logarithmic in R .

This energy barrier impedes thermal defect diffusion, enhancing the stability of the quantum memory.

Memory time

Bravyi-Haah 2011



Because of the logarithmically increasing height of the logical energy barrier, the memory time grows like a power of volume for small system size.

$$t_{mem} \sim V^{\Omega(\beta)}$$
$$(\beta = (\text{temperature})^{-1})$$

But once the system size grows beyond an optimal size, the entropy of the defects grows exponentially with volume, overwhelming the logarithmic energy cost. Thus the memory time is a constant depending on the temperature.

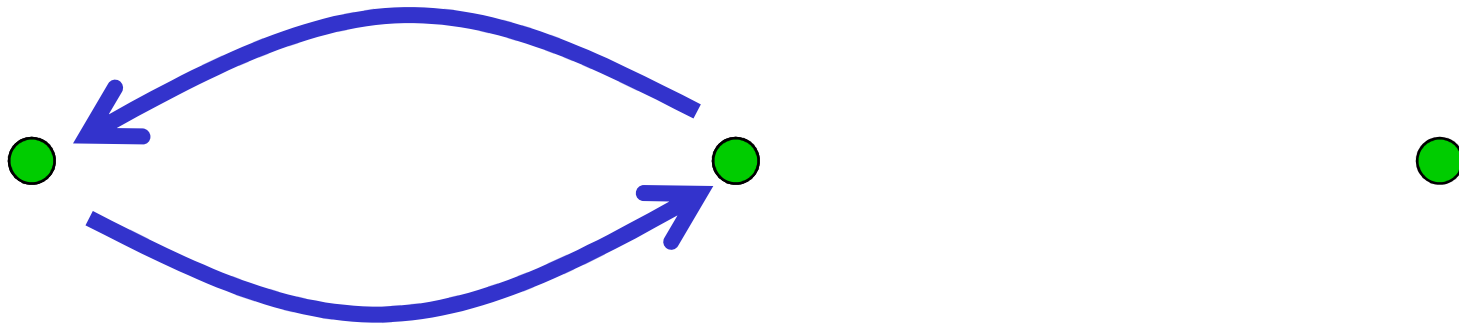
$$V^* \sim e^{\Omega(\beta)}$$

$$t_{mem}^* \sim e^{\Omega(\beta^2)}$$

Michnicki 2012: Energy barrier $O(L^{2/3})$ in 3D code w/o translation invariance.

Nonabelian anyons

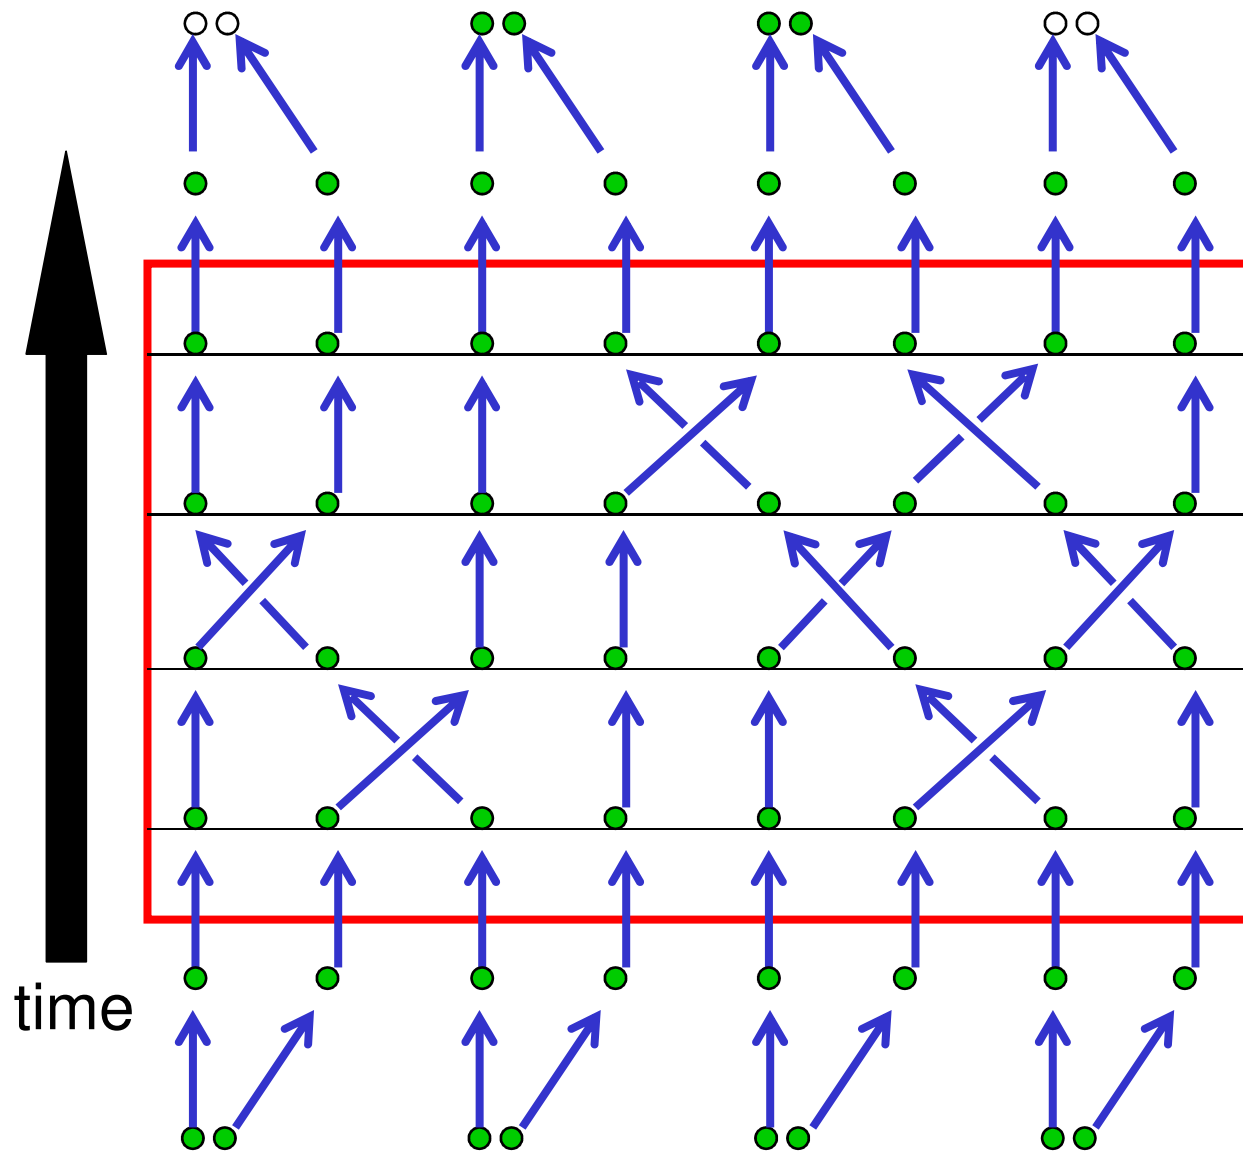
Quantum information can be stored in the collective state of exotic particles in two spatial dimensions (“anyons”).



The information can be processed by exchanging the positions of the anyons (even though the anyons never come close to one another).

Topological quantum computation

(Kitaev '97, FLW '00)



annihilate pairs?

braid

braid

braid

create pairs

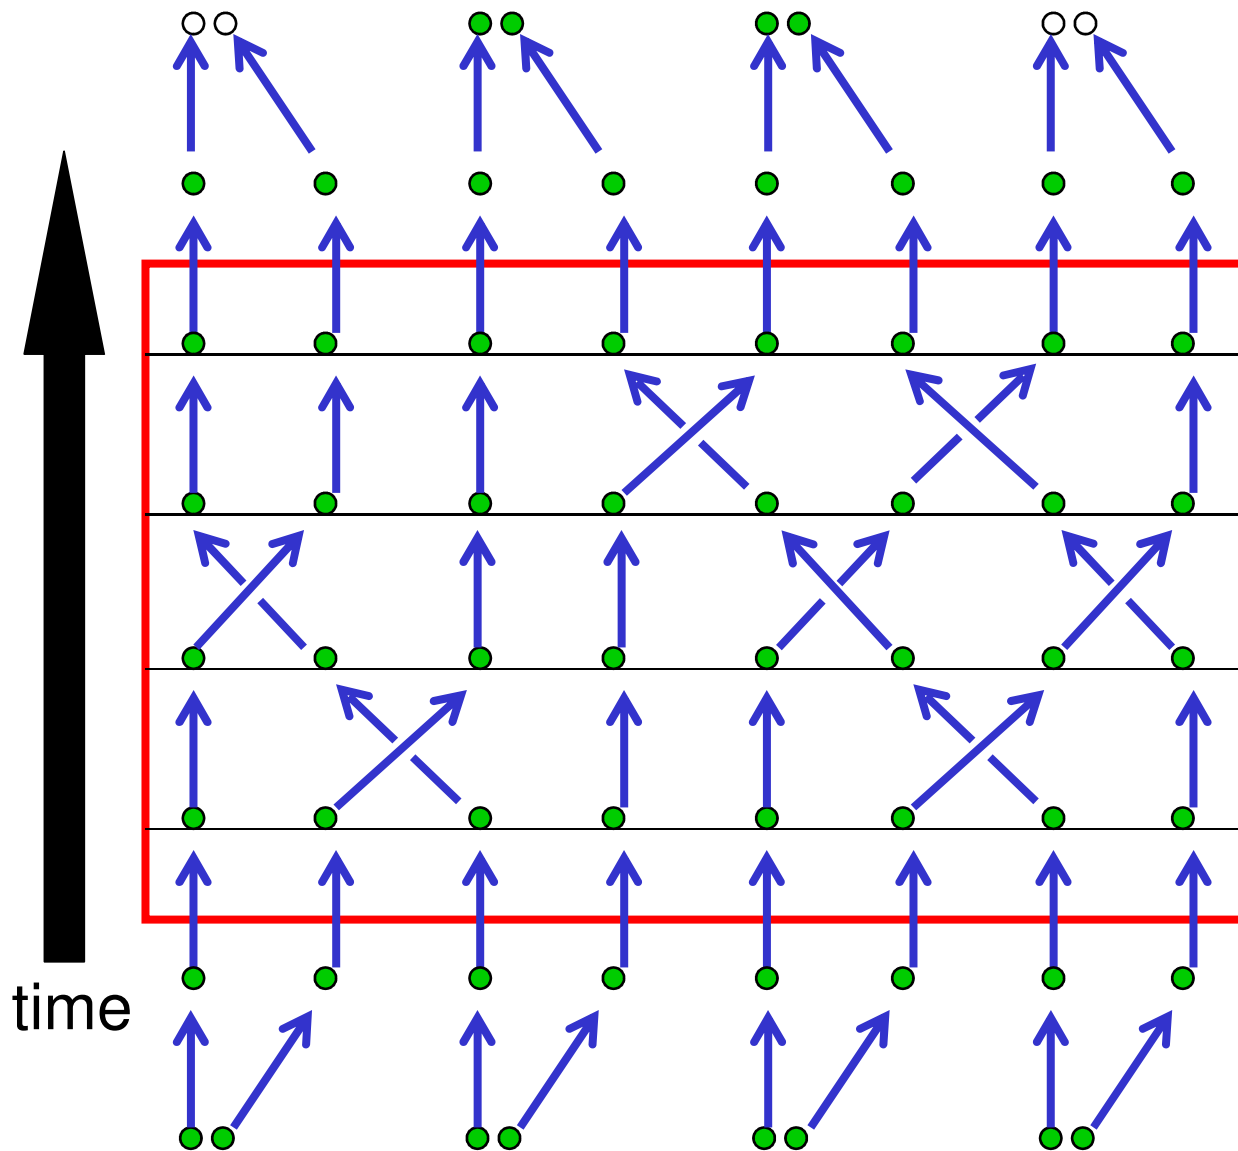


Kitaev



Freedman

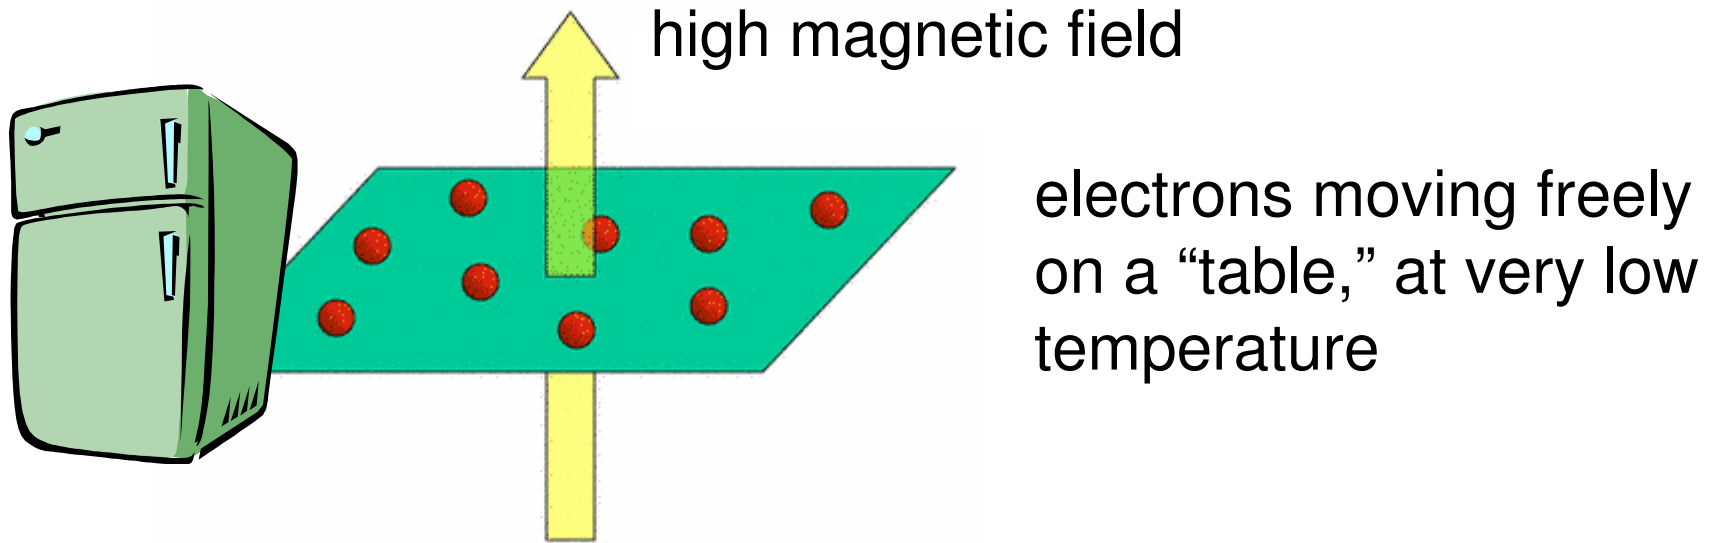
Topological quantum computation



The computation is intrinsically resistant to decoherence.

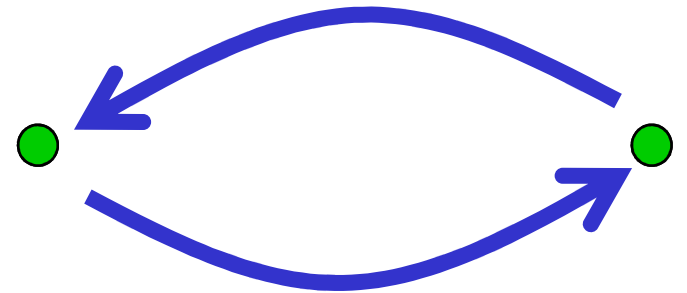
If the paths followed by the particles in spacetime execute the right braid, then the quantum computation is guaranteed to give the right answer!

Anyons: the fractional quantum Hall effect

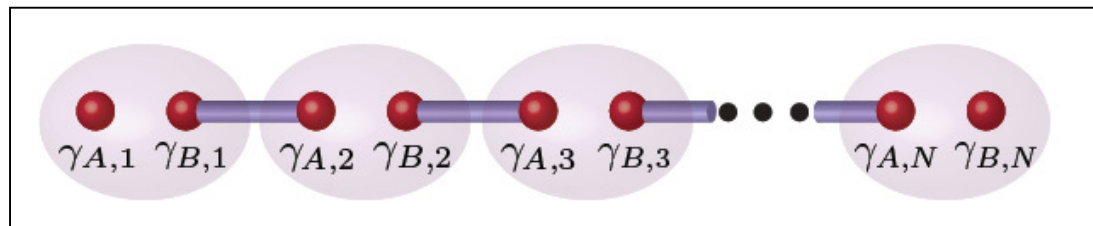
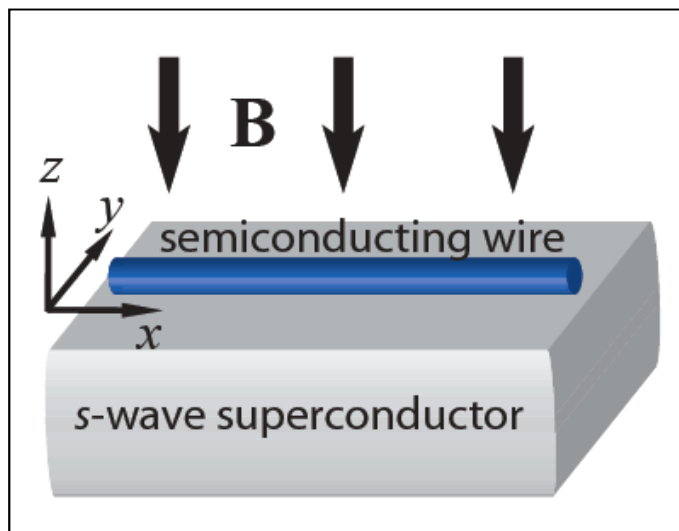


An exotic new phase of matter, with particle excitations that are profoundly different than electrons.

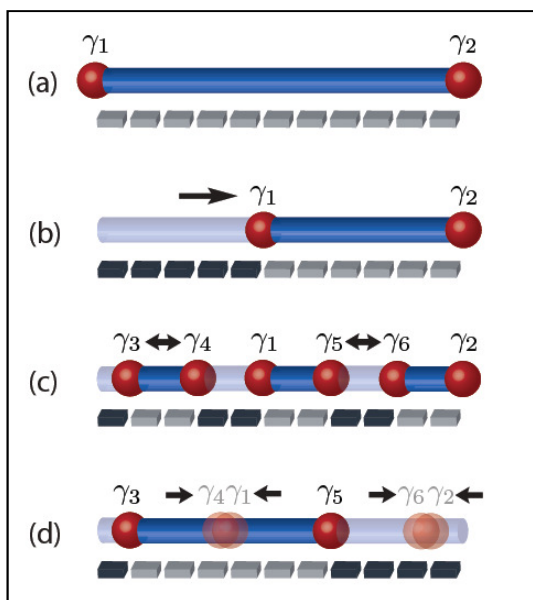
These particles are *anyons*: they have topological interactions.



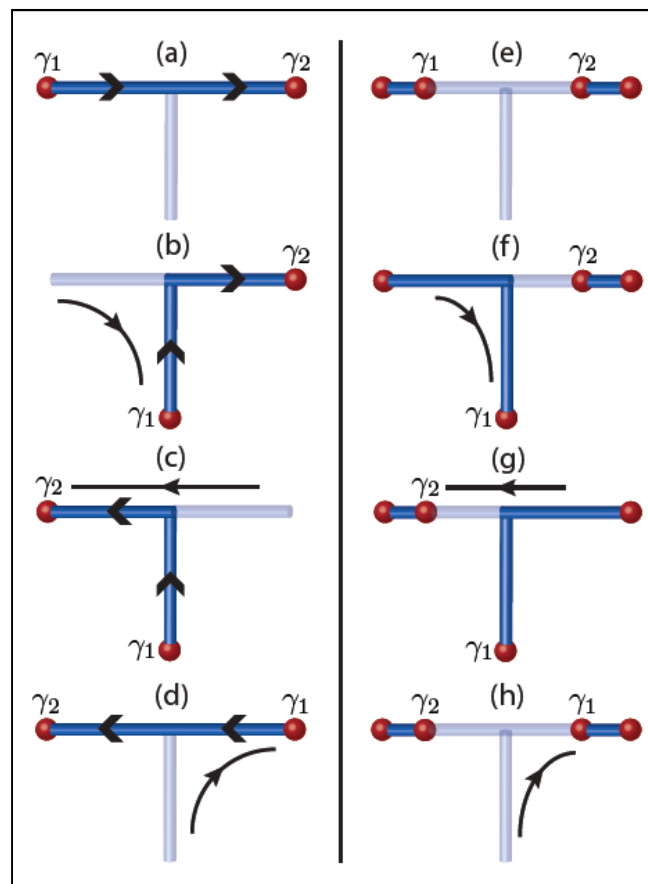
Majorana fermions at the ends of quantum wires (Kitaev 2001)



Topological superconductor:
charge can be *odd* multiple of e .



Anyons, appearing where topological
and normal superconductors meet,
move as chemical potential is adjusted.



Alicea et al.:

To exchange
particles, park
one using a
T-junction.

(The topologically
protected gates
are not universal.)

Hardware

- Robust devices (e.g. “0- Π ” superconducting qubit).
- Topological protection and processing (e.g. Majorana fermions in quantum wires).

Software

- Optimized threshold and overhead.
- Adapting fault tolerance to noise.
- Dynamical decoupling.

Systems engineering (wires, power, cooling, etc.)

Matters of principle

- Limitations on noise correlations
- Justifying error phase randomization \rightarrow error probabilities (e.g., relating error benchmarking to fault tolerance requirements).
- Self-correcting hardware (e.g., favorable scaling of storage time with system size, in fewer than four dimensions?).
- Other scalable schemes besides concatenated codes and topological codes (perhaps fault-tolerant adiabatic quantum computing?).
- Broader implications of quantum error correction in physical science.

Quantum fault tolerance

Operating a large-scale quantum computer will be a grand scientific and engineering achievement.

Judicious application of the principles of fault-tolerant quantum computing will be the key to making it happen.

Fascinating connections with statistical physics, quantum many-body theory, device physics, and decoherence make the study of quantum fault tolerance highly rewarding.

Additional Slides

Noise correlations and scalability

In general, the noise Hamiltonian may contain terms acting on m system qubits, for $m = 1, 2, 3, \dots$

$$H_{\text{System-Bath}} = \sum_i H_i^{(1)} + \sum_{\langle ij \rangle} H_{ij}^{(2)} + \sum_{\langle ijk \rangle} H_{ijk}^{(3)} + \dots$$

Quantum computing is provably scalable if $\varepsilon \leq \varepsilon_0 \cong 10^{-4}$, where

$$\varepsilon = (9.44) \times \max_m \eta_m^{1/m} \quad \text{and}$$

$$\eta_m = \max_{j_1} \sum_{j_2, j_3, \dots, j_m} \| H_{j_1 j_2 j_3 \dots j_m}^{(m)} \| t_0$$

[t_0 is the maximal duration of any quantum gate.]

over all times
and qubits

interactions fall
off with distance

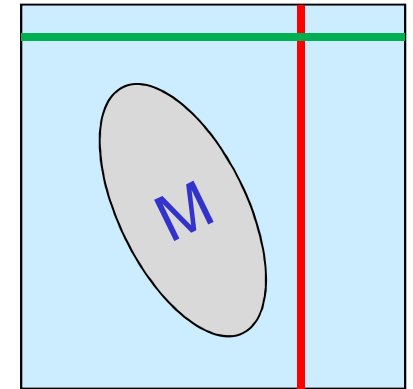
term that acts collectively on m system
qubits should be exponentially small in m .

Currently known proofs of the threshold theorem require the noise to be “quasi-local” in the sense that the m -qubit noise term in the Hamiltonian decays exponentially with m . Can experiments verify this scaling?

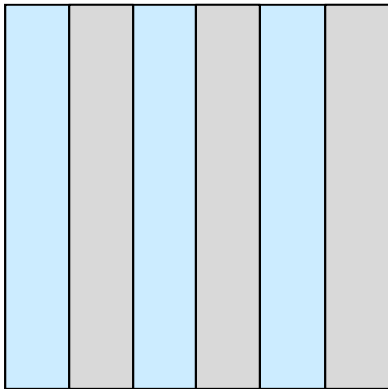
Self-correction in two dimensions?

Consider a stabilizer (or subsystem) code.

Cleaning Lemma: If a set of qubits M is correctable (supports no nontrivial logical operator), then any nontrivial logical Pauli operator can be “cleaned” by applying stabilizer generators, so it acts trivially on M .

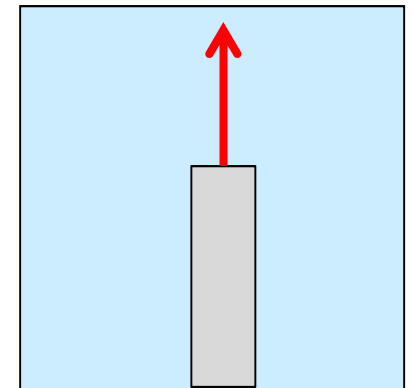


Theorem (Bravyi-Terhal 2009, Kay-Colbeck 2008): A 2D stabilizer code has a “string logical operator.”



Cover the code block with stripes or alternating color, wide enough so that no check operator acts on two stripes of the same color. If the gray stripes are correctable, they can be cleaned, so there is a nontrivial logical operator supported on blue stripes.

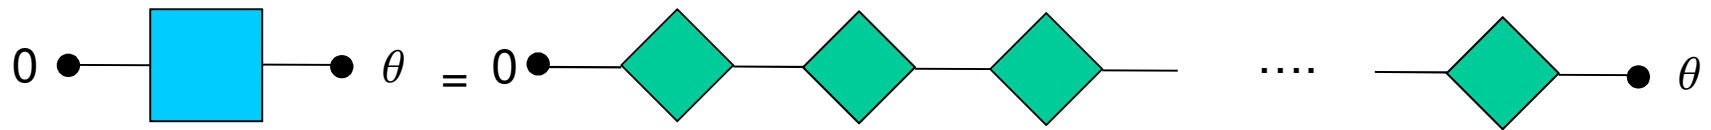
Each check operator overlaps with just one blue stripe, so the operator on each blue stripe is logical, and for at least one stripe is nontrivial.



Protected superconducting qubit

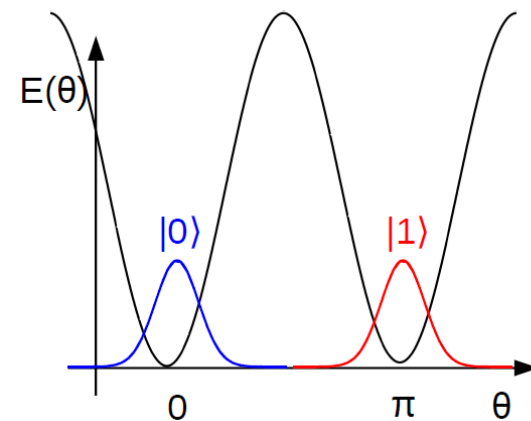
Ioffe et al.
Kitaev

One way to make a robust superconducting (0- π) qubit is to build a long chain of devices. Each individual device favors a phase change of 0 or π across its leads. The phase difference between the two ends of the chain can likewise be either 0 or π but with large local phase fluctuations along the chain.



The two basis states of the qubit are distinguished by a global property of the chain --- both look the same locally. For long chain, the breaking of the degeneracy of the two states due to a generic local perturbation occurs in a high order of perturbation theory and is strongly suppressed.

The barrier is high enough to suppress bit flips, and the stable degeneracy suppresses phase errors. Protection arises because the encoding of quantum information is highly nonlocal, and splitting of degeneracy scales exponentially with (square root of) size of the device.



$$E \approx f(2\theta) + O\left(\exp\left(-c\sqrt{\text{size}}\right)\right)$$

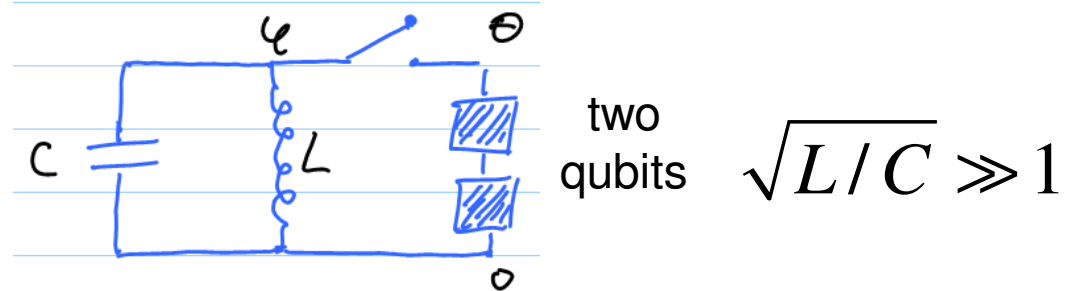
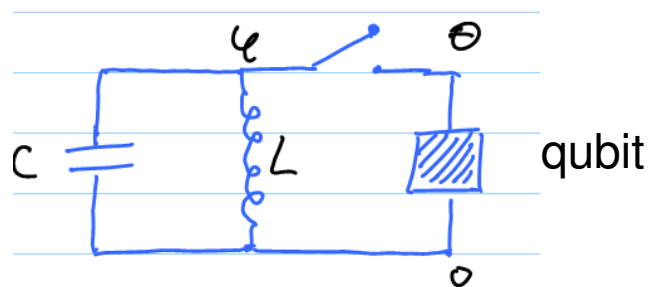
Protected superconducting qubit

Kitaev,
Brooks,
Preskill

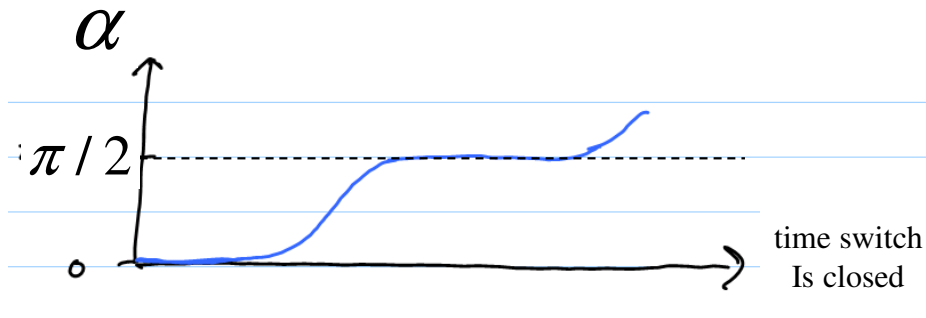
Some gates are also protected: we can execute

$$\exp\left(i\frac{\pi}{4}Z\right) \text{ and } \exp\left(i\frac{\pi}{4}Z_1 \otimes Z_2\right)$$

with exponential precision. This is achieved by coupling a qubit or a pair of qubits to a “superinductor” with large phase fluctuations:



To execute the gate, we (1) close the switch, (2) keep it closed for awhile, (3) open the switch. This procedure alters the relative phase of the two basis states of the qubit:

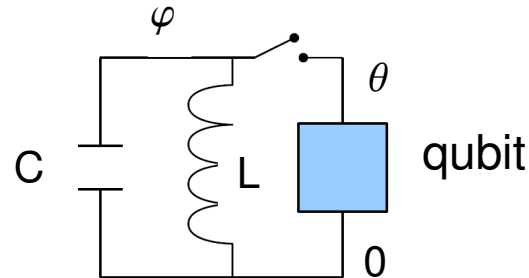
$$(a|0\rangle + b|1\rangle) \otimes |\text{init}\rangle \rightarrow (a|0\rangle + be^{-i\alpha}|1\rangle) \otimes |\text{final}\rangle$$


The relative phase induced by the gate “locks” at $\pi/2$. For $\sqrt{L/C} \approx 80$ phase error $\sim \text{few} \times 10^{-8}$ is achieved for timing error of order 1 percent. Why?

Protected phase gate

Kitaev,
Brooks,
Preskill

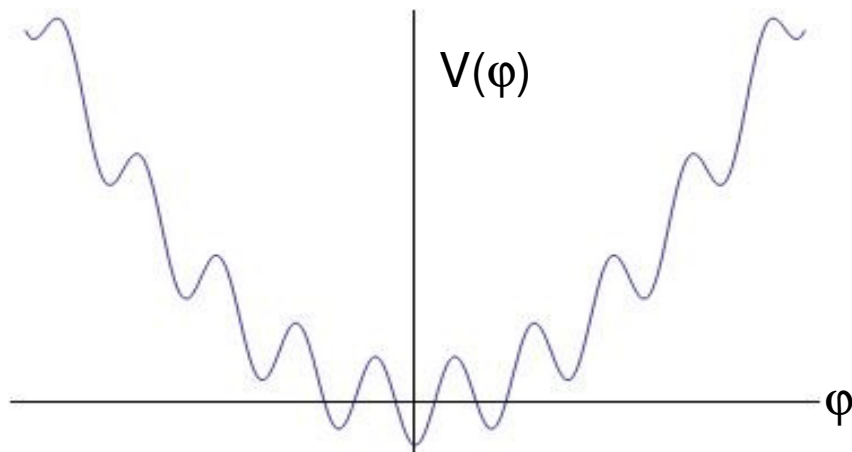
$$\exp\left(i\frac{\pi}{4}Z\right)$$



$$\sqrt{L/C} \gg \hbar / (2e)^2 \approx 1 k\Omega$$

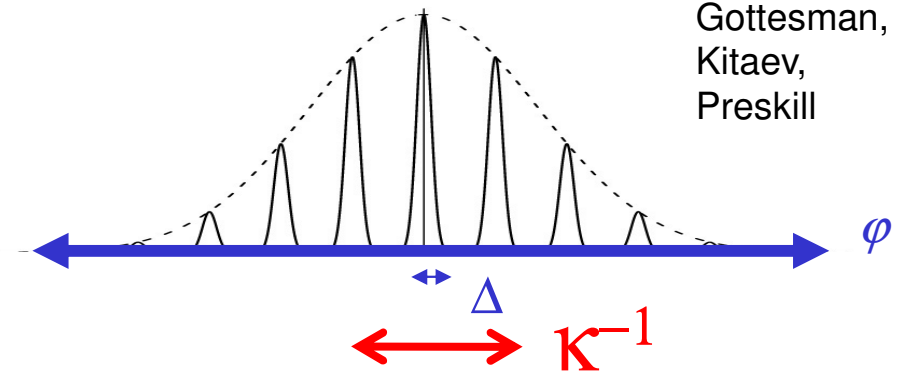
Switch is really a tunable Josephson junction:

$$H = \frac{Q^2}{2C} + \frac{\phi^2}{2L} - J(t) \cos(\phi - \theta)$$



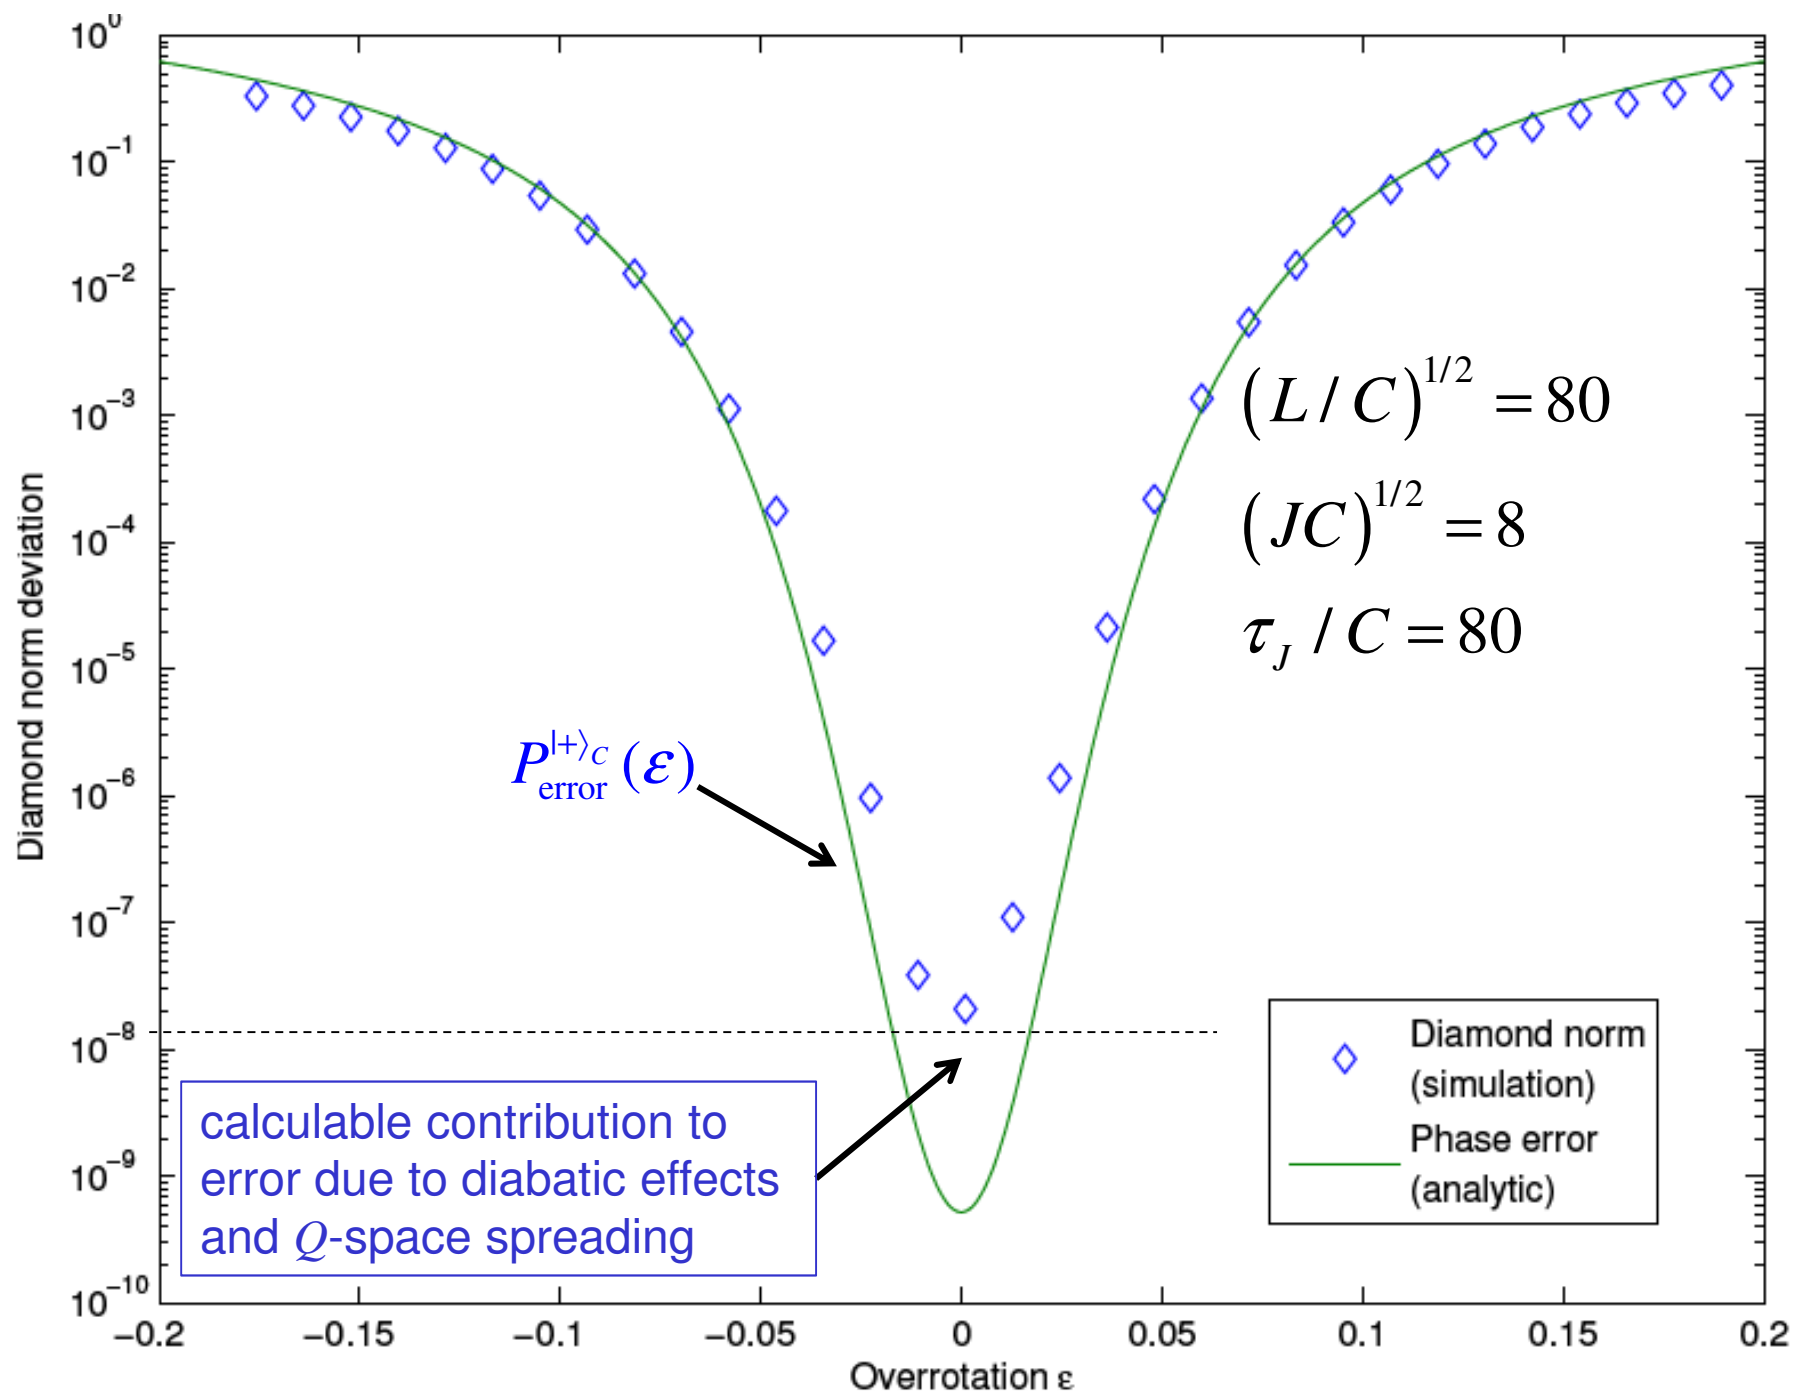
Under suitable adiabaticity conditions, closing the switch transforms a broad oscillator state (e.g. the ground state) into a grid state (approximate codeword).

Gottesman,
Kitaev,
Preskill



Peaks are at even or odd multiples of π depending on whether θ is 0 or π , i.e. on whether qubit is 0 or 1. Inner width squared is $(JC)^{-1/2}$ and outer width is $(L/C)^{1/2}$

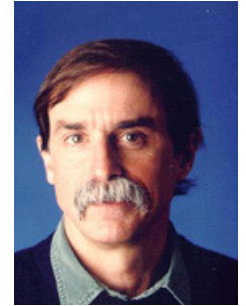
$$\omega_J^{-1} = \sqrt{C/J} \ll \text{switching time} \ll \omega^{-1} = \sqrt{LC} \gg 1$$



Some recently reported error rates

Ion trap – one-qubit gates:

$\sim 2 \times 10^{-5}$ [NIST]



Wineland

Ion trap – two-qubit gates:

$\sim 5 \times 10^{-3}$ [Innsbruck]



Blatt

Superconducting circuits – one-qubit gate

$\sim 2.5 \times 10^{-3}$ [Yale]

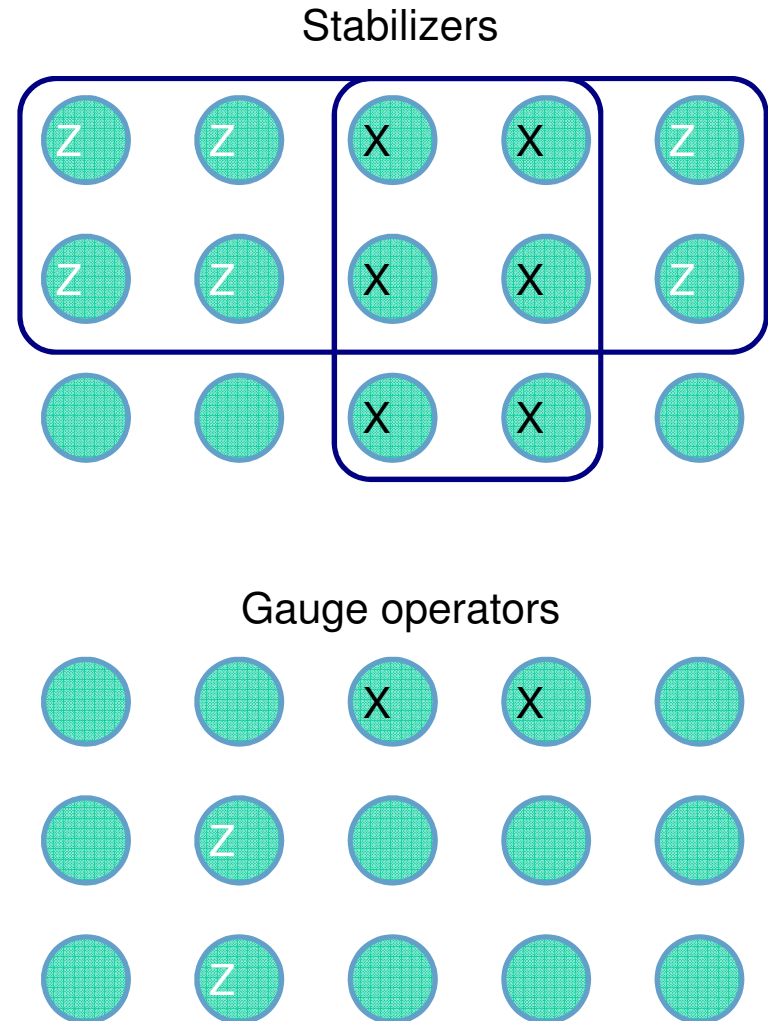
Error rates are estimated by performing “circuits” of variable size, and observing how the error in the final readout grows with circuit size.



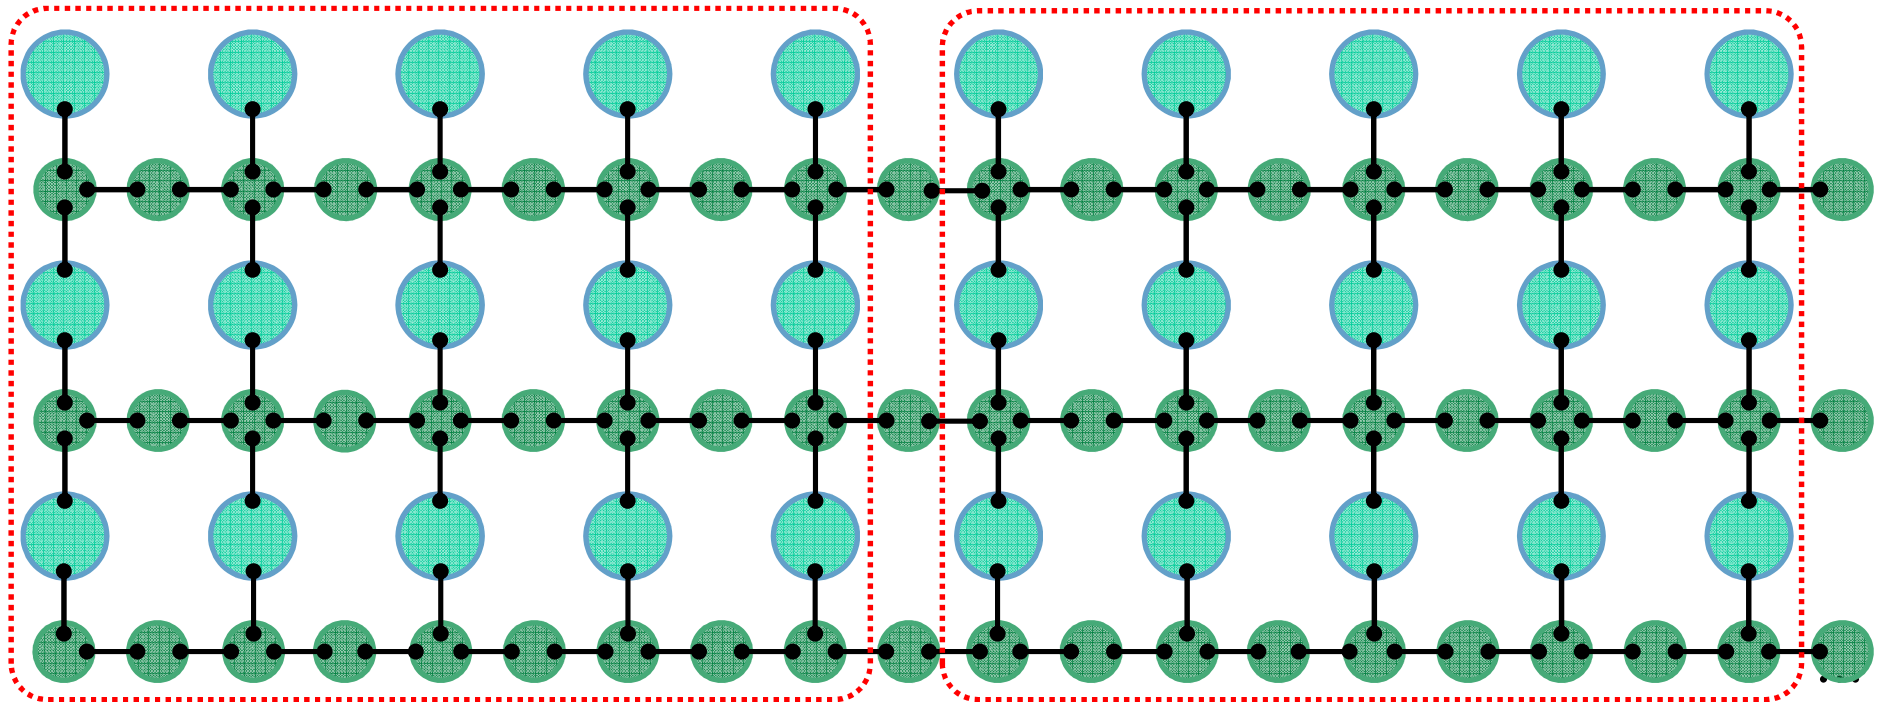
Schoelkopf

Asymmetric Bacon-Shor codes

- Goal: design circuits for fault-tolerant quantum computation that take advantage of dephasing bias in noise
- Encode within $n \times m$ block of Bacon-Shor code
- Independently tunable levels of X and Z error protection
- Gauge structure allows measuring high-weight stabilizers with weight 2 gauge operators

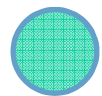


Asymmetric Bacon-Shor codes



Block 1

Block 2



Data qubit



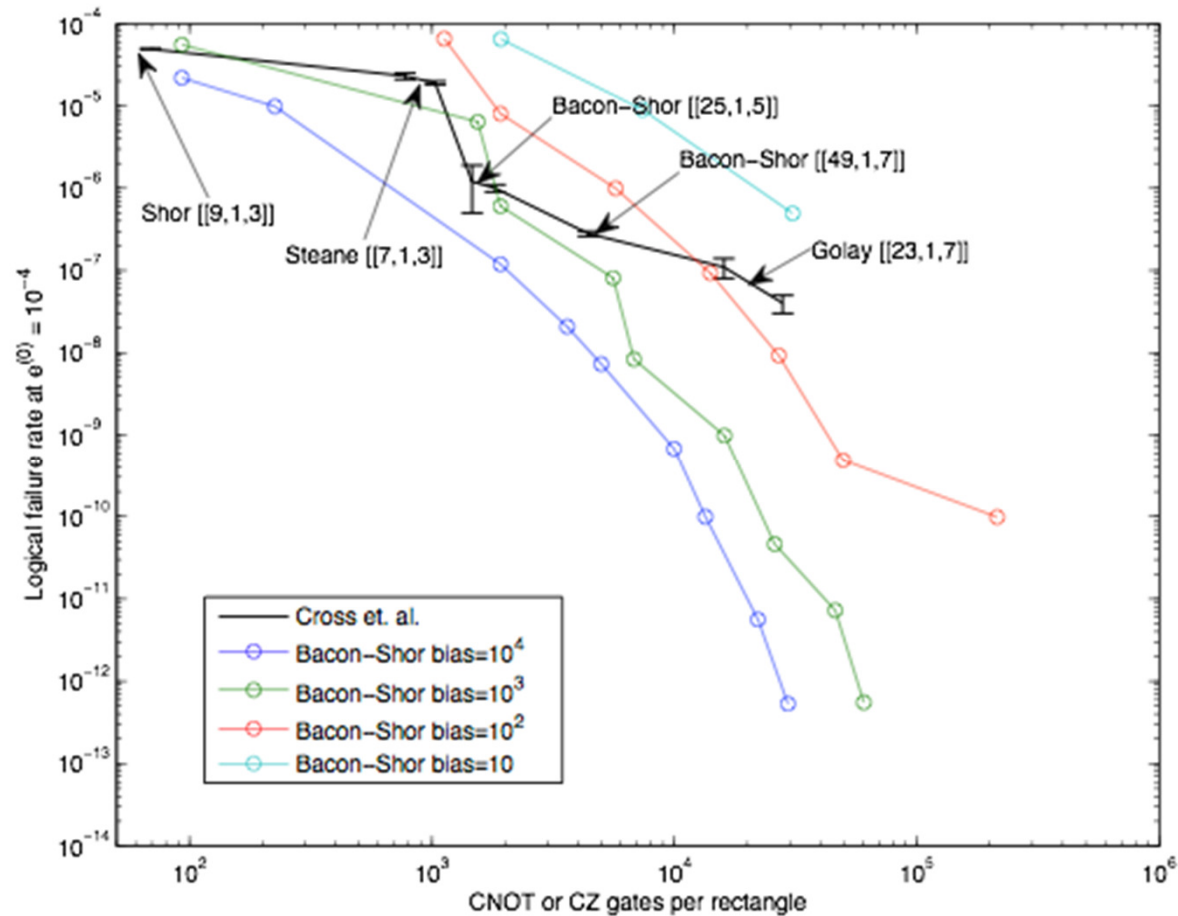
Ancilla qubit



Potential CZ location

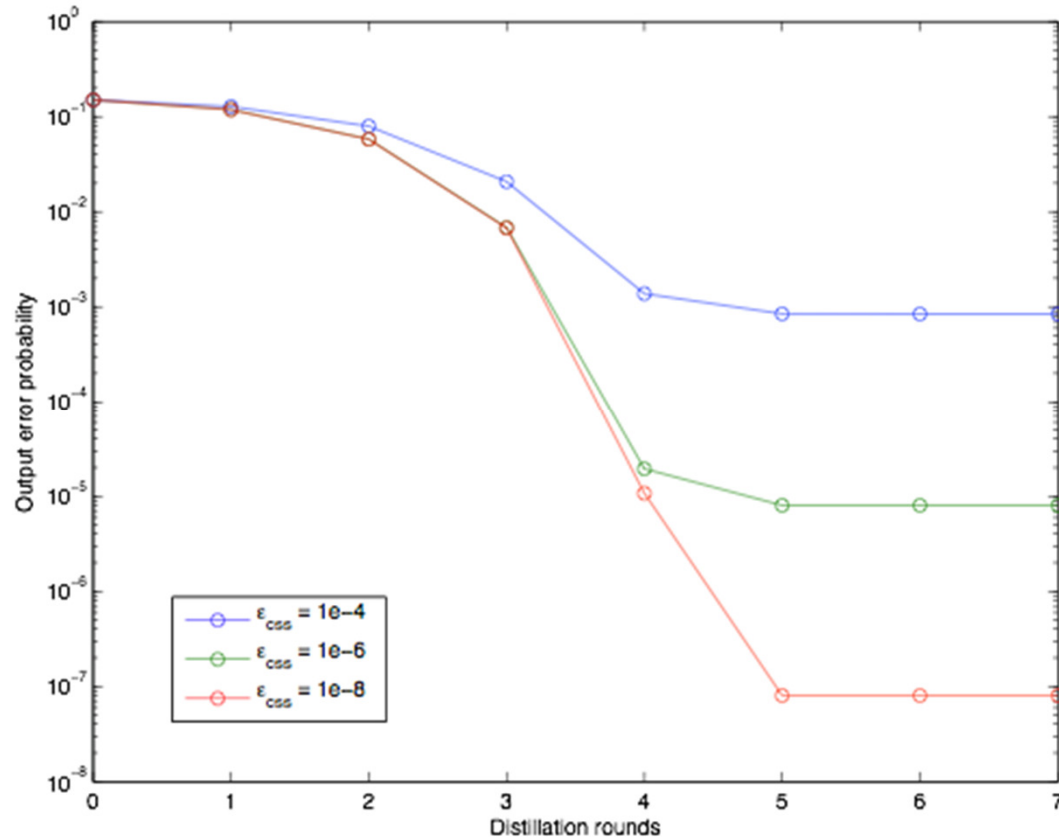
Bacon-Shor gauge structure permits a geometrically local architecture with gates only acting on nearest neighbors

Asymmetric Bacon-Shor codes



- When bias is high, can achieve a low error rate with significantly reduced overhead compared to the unbiased case

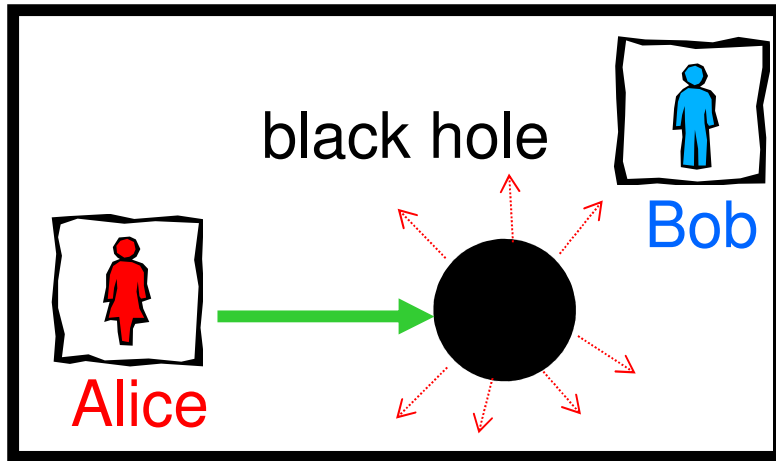
Asymmetric Bacon-Shor codes



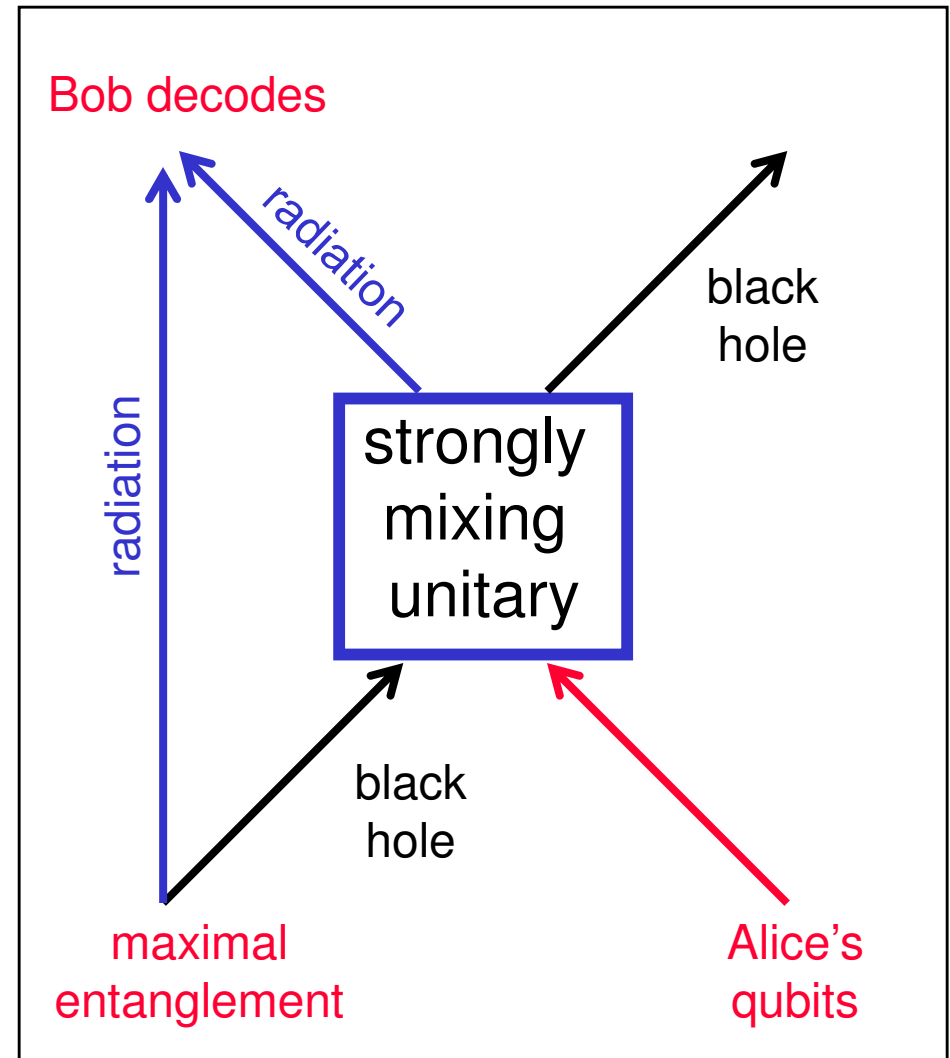
- To achieve universality, use magic state injection and distillation protocols
- Analyze distilled state output error probability while explicitly accounting for imperfect Clifford gates
- Starting with input error probability of 0.01, 2-3 rounds of distillation suffice to reach maximum amount of distillation for essentially all practical encoded Clifford error rates

How fast does information escape from a black hole?

Hayden,
Preskill



Black holes are (we believe) efficient quantum information processors. How long do we have to wait for information absorbed by a black hole to be revealed in its emitted Hawking radiation? We reconsidered this question using tools from quantum information theory.



Our (tentative) conclusion is that the retention time can be surprisingly short. The analysis uses the theory of quantum error-correcting codes and quantum circuits.