Clifford Limits and Magic States: Gottesman–Knill, Eastin–Knill and Distillation

Why fault-tolerant quantum computers need a non-Clifford resource, and how magic states supply it

Technical note · October 26, 2026 · D.F. Hakala

The gates that are easy to protect in an error-corrected computer, the Clifford gates, can be simulated efficiently on a classical computer (Gottesman–Knill), so on their own they give no quantum advantage. No error-correcting code can supply the missing non-Clifford gate transversally either (Eastin–Knill). The standard way out is to prepare special resource states, magic states, purify them by distillation using only Clifford operations, and consume one for each T gate by gate teleportation. That extra machinery dominates the hardware overhead of large fault-tolerant machines.

1. Gottesman–Knill theorem (Gottesman, 1998, crediting Knill)

Circuits built only from Clifford operations can be simulated efficiently on a classical computer, however much entanglement they create.

Statement

A quantum circuit built only from the following pieces can be simulated classically in polynomial time and memory:

- preparing qubits in computational basis states,
- **Clifford gates:** Hadamard (H), the phase gate (S = √Z) and CNOT,
- measurements in the Pauli (X, Y, Z) bases,
- classical control based on measurement results.

Why it works

Clifford gates map Pauli operators to Pauli operators. Instead of tracking 2^n amplitudes, a simulator tracks the n Pauli "stabilizer" operators that define the state. That takes about n^2 bits (a binary tableau), and each gate updates it in time polynomial in n. Measurement outcomes and their probabilities can be read off the tableau. Aaronson and Gottesman's CHP simulator does this for millions of qubits.

What it means

- Clifford circuits can produce large amounts of entanglement (GHZ states, cluster states, the code states of the surface code), yet they give no quantum speedup. **Entanglement alone is not the source of quantum advantage.**
- Going beyond classical simulation needs at least one **non-Clifford gate**, usually the T gate or a Toffoli. Clifford + T is a universal gate set. This non-Clifford ingredient is called "magic" (Section 4).
- Classical simulation cost grows roughly exponentially in the number of T gates, not the number of qubits. That is why T-count is the standard cost metric for quantum algorithms.
- In practice, it is also why error-correction research can simulate huge codes and decoders classically: those circuits are almost entirely Clifford.

2. Eastin–Knill theorem (Eastin and Knill, 2009)

No useful error-correcting code has a universal set of transversal logical gates; every code is missing something.

Statement

No quantum error-correcting code that can detect arbitrary errors on any single physical qubit (or subsystem) can implement a **universal** set of logical gates **transversally**.

"Transversal" means the logical gate is applied as a product of gates that each act on only one qubit of each code block, with no interactions within a block. Transversal gates are naturally fault-tolerant: one faulty physical gate creates at most one error per block, which the code can correct.

Why it holds (the intuition)

1. The transversal logical gates form a group.
2. If that group were universal, it would contain logical operations arbitrarily close to the identity, continuously connected to it.
3. A transversal operation infinitesimally close to the identity is a sum of tiny single-qubit (local) perturbations.
4. The code must detect any local error, so local perturbations cannot act nontrivially on the encoded information. Yet these operations do change the logical state, which is a contradiction.

So for any useful code the transversal gate set is **finite (discrete)** and therefore not universal.

Examples

Code	Transversal	Not transversal
Surface code	Paulis and CNOT between blocks; other Cliffords via lattice surgery or code deformation	T
2D color code	Full Clifford group	T
3D color code / 15-qubit Reed–Muller code [[15,1,3]]	T	H

3. How the two theorems fit together

The two theorems box in fault-tolerant design from opposite sides:

- **Gottesman–Knill:** the gates that are easy to do fault-tolerantly (Cliffords) are, by themselves, classically simulable and useless for quantum advantage.
- **Eastin–Knill:** the missing non-Clifford gate cannot be had for free through transversal operations in the same code.

The non-Clifford part therefore needs extra machinery, and in practice that machinery dominates hardware overhead. The rest of this note covers the main answer, magic states, and the alternatives.

4. Magic states

The missing gate is usually the T gate, a phase rotation by $\pi/4$ on $|1\rangle$. It is also called the $\pi/8$ gate, because up to a global phase it is a rotation of $\pm\pi/8$ about Z:

$$T = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{pmatrix} = e^{i\pi/8} \begin{pmatrix} e^{-i\pi/8} & 0 \\ 0 & e^{i\pi/8} \end{pmatrix} \quad (1)$$

Instead of performing T directly on the data, the computer prepares a special resource state and consumes it. The standard one is:

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

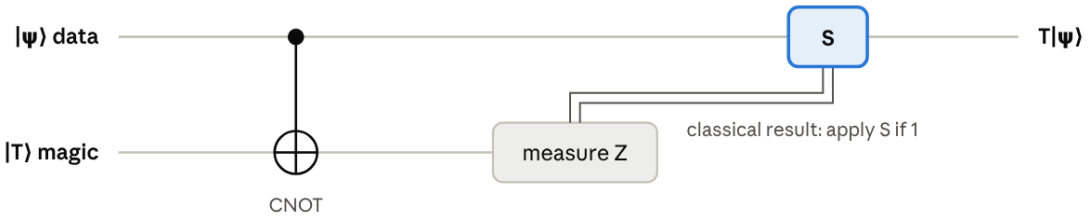
It is called "magic" because it lies outside the set of states Clifford operations can reach (the stabilizer states); that extra resource is exactly what Clifford-only circuits lack. For Toffoli-heavy algorithms, a three-qubit CCZ magic state plays the same role.

5. Using a magic state: gate teleportation

Each T gate in an algorithm consumes one magic state:

1. Apply a CNOT with the data qubit as control and the magic state |T⟩ as target.
2. Measure the magic qubit in the Z basis.
3. If the result is 1, apply a corrective S gate to the data qubit; S is a cheap Clifford gate.
4. The data qubit now has T applied to it, and the magic state has been used up.

One CNOT, one measurement and a conditional S apply T to the data



The magic state is consumed; only Clifford operations touch the data, so the step stays fault-tolerant.

Figure 1. Gate-teleportation circuit for one T gate.

Consequences:

- **Algorithm cost is counted in T gates.** The T-count largely sets an algorithm's cost.
- **Dedicated factories.** Large machines devote regions to "magic-state factories" that produce |T⟩ continuously.
- **Hardware requirements.** Mid-circuit measurement and real-time feed-forward are mandatory: the S correction is decided from a measurement result within the same circuit.

6. Routes to the non-Clifford gate

Method	How it works	Output error	Cost
Injection	Prepare the state on one physical qubit with a physical T rotation, then encode it into the logical code	About the physical error rate p (far too high to use directly)	Very low

Method	How it works	Output error	Cost
Distillation	Many noisy injected states are purified into fewer clean ones with Clifford circuits and checks (Section 7)	$35p^3$ per 15-to-1 level; repeatable	High: about 15 inputs per output per level
Cultivation (Gidney, Shutty and Jones, 2024)	The state is grown and repeatedly checked inside a single surface-code patch	Low enough for many algorithms	Much cheaper than distillation
Code switching / gauge fixing	Move logical information into a code with a transversal T, such as the [[15,1,3]] Reed–Muller or 3D color code, apply T, move back	Set by the code distance	Depends on switching overhead
Non-transversal fault-tolerant constructions	Pieceable fault tolerance, measurement-based tricks or symmetry-breaking approaches	Construction-dependent	Varies; none is transversal in the strict sense, so none contradicts Eastin–Knill

Harvard used 3D [[15,1,3]] codes in 2025 for teleportation-based arbitrary-angle rotations on neutral atoms.

7. Distillation

Distillation takes many noisy copies of a resource state, runs a circuit built from reliable operations, checks the results, and keeps only the outputs where every check passes. The result is fewer copies, each with much lower errors. It works because the checks reveal which inputs were bad, and outputs are accepted only when all checks agree. The idea began with entanglement distillation: purifying noisy Bell pairs into fewer, nearly perfect ones (Bennett et al., 1996).

Magic-state distillation (Bravyi and Kitaev, 2005) applies this to $|T\rangle$ states using only Clifford operations, which are already cheap and reliable at the logical level.

15-to-1 protocol

1. Take 15 noisy $|T\rangle$ states, each with error p.
2. Run the Clifford decoding circuit of the [[15,1,3]] Reed–Muller code, which has a transversal T.
3. Measure 14 of the qubits. If any check fails, discard the batch.
4. If every check passes, the remaining qubit is a $|T\rangle$ state with lower error. The smallest error pattern the code cannot detect involves 3 inputs, and there are 35 such patterns:

$$p_{out} \approx 35\,p^3, \qquad P_{success} \approx 1 - 15p \tag{3}$$

Repeating it L times compounds the gain, at a cost that grows as 15^L:

$$p_{out}^{(L)} \approx 35^{(3^L-1)/2}\,p^{3^L}, \qquad N_{in} = 15^L \tag{4}$$

Level L	Output error (from $p = 10^{-3}$)	Noisy inputs per clean output
0 (injected)	1×10^{-3}	1
1	3.5×10^{-8}	15
2	1.5×10^{-21}	225

That cost is why magic-state factories can take up a large share of a fault-tolerant machine.

5-to-1 protocol

Based on the [[5,1,3]] code, it suppresses error quadratically ($p \rightarrow O(p^2)$) using fewer inputs. QuEra, Harvard and MIT demonstrated it on neutral atoms in 2025, on distance-3 and distance-5 color-code logical qubits: the first magic-state distillation performed with logical qubits.

8. Relevance to neutral-atom arrays

- **Transversal Cliffords are natural.** A whole code block can be moved into pairing with another and given a global Rydberg pulse, performing a transversal CNOT in one step; recent logical-qubit demonstrations used exactly this. Eastin–Knill is the reason such machines still need magic states or code switching alongside these fast transversal Cliffords.
- **A magic-state factory zone.** A neutral-atom machine would give magic states their own zone, alongside storage, gate and readout zones, with AODs carrying finished states to where they are consumed.
- **What it relies on:** parallel transversal Clifford gates, mid-circuit measurement and fast feed-forward from the error-correction control stack.
- **Where the savings are:** cutting factory cost through cultivation, transversal-T codes or code switching is one of the main ways to reduce total qubit counts and runtime.

References

D. Gottesman, “The Heisenberg representation of quantum computers,” arXiv:quant-ph/9807006 (1998).

S. Aaronson and D. Gottesman, “Improved simulation of stabilizer circuits,” Phys. Rev. A 70, 052328 (2004).

B. Eastin and E. Knill, “Restrictions on transversal encoded quantum gate sets,” Phys. Rev. Lett. 102, 110502 (2009).

S. Bravyi and A. Kitaev, “Universal quantum computation with ideal Clifford gates and noisy ancillas,” Phys. Rev. A 71, 022316 (2005).

C. H. Bennett et al., “Purification of noisy entanglement and faithful teleportation via noisy channels,” Phys. Rev. Lett. 76, 722 (1996).

C. Gidney, N. Shutty and C. Jones, “Magic state cultivation: growing T states as cheap as CNOT gates,” arXiv (2024).

QuEra, Harvard and MIT, logical-level magic-state distillation on neutral atoms, Nature (2025); summary: Converge Digest — https://convergedigest.com/quera-harvard-mit-achieve-first-logical-level-magic-state-distillation-on-neutral-atom-quantum/

D. Bluvstein et al., “Architectural mechanisms of a universal fault-tolerant quantum computer,” arXiv:2506.20661 (2025) — https://arxiv.org/abs/2506.20661