A Gentle Mathematical Introduction to Quantum Error Correcting Codes

Jaya Sharma

Revised teaching draft October 2026
Original survey dated May 2023

Abstract

Quantum error correction protects an unknown quantum state by encoding it into a larger system and measuring information about errors without measuring the encoded amplitudes. This survey develops that idea from three-qubit repetition codes to stabilizer, Calderbank–Shor–Steane, subsystem, and bosonic codes. The central tools are the Pauli operator basis, the Knill–Laflamme condition, and binary symplectic linear algebra. Worked examples include coherent-error recovery, degeneracy in the nine-qubit Shor code, syndrome decoding in the seven-qubit Steane code, logical operators in the toric code, the Bacon–Shor gauge subsystem, and single-photon-loss recovery in a binomial encoding. Cat and Gottesman–Kitaev–Preskill codes introduce oscillator-based protection. A treatment of universal fault-tolerant computation explains transversal gates, their limitations, and magic-state injection. A final perspective distinguishes code parameters, decoding, and fault-tolerant implementation, with selected developments in surface codes, quantum low-density parity-check codes, and holographic codes. The text includes exercises with solutions and assumes undergraduate linear algebra; no previous course in quantum error correction is required.

Scope. A tutorial on subspace, stabilizer, subsystem, and bosonic encodings, with an introduction to universal fault-tolerant computation and selected research examples published through 2025. This is not an exhaustive review of the field or a claim of new mathematical results. The bosonic and gate constructions illustrate principles rather than specify complete hardware protocols.

Prerequisites. Complex vectors, matrices, eigenvalues, tensor products, and elementary probability. Binary linear codes and the needed quantum notation are introduced here.

Version note. This student revision supersedes the exposition of the original survey for review purposes. It is an author-review draft, not a peer-reviewed publication. Reference links were checked in October 2026.

Editing assistance. ChatGPT assisted with organization, exposition, and editing. The model designation GPT-6.0 is supplied by the author and has not been independently verified.

Contents

1 The problem and the basic notation 3

2 Noise and the error model 4

3 Repetition codes and syndrome measurement 5

4 The condition for quantum error correction 6

5 Stabilizer codes and binary linear algebra 8

6 The Shor code and the role of degeneracy 9

7 Classical linear codes and the CSS construction 10

8 From stabilizers to geometry 12

9 Subsystem codes and gauge freedom 14

10 Bosonic codes: using an oscillator 15

11 Universal fault-tolerant quantum computation 17

12 Interpreting modern code families and performance claims 20

13 Exercises 22

14 A guide to further study 24

A Solutions to the exercises 24

B A compact mathematical checklist 26

References 26

How to use this survey

For a first reading, study Sections 1–3, then the examples in Sections 5–7. Return to the correction theorem in Section 4 after you have computed a syndrome yourself. Section 8 gives a geometric route into modern codes. Sections 9–11 extend the framework to gauge freedom, oscillators, and universal computation. Section 12 explains how to read performance claims.

After the core sections, you should be able to compute a Pauli syndrome, test whether stabilizer generators commute, count encoded qubits, explain degeneracy, construct and decode the Steane code, and distinguish an encoding from a fault-tolerant implementation.

1 The problem and the basic notation

The purpose of error correction is to preserve quantum information even when the physical system carrying it is disturbed. Shor's nine-qubit construction established that this is possible despite the no-cloning theorem [1]. The essential idea is to distribute information across a larger Hilbert space so that suitable measurements reveal an error syndrome while leaving the logical state recoverable.

It is helpful to separate three objects from the outset: a *code* is a subspace, a *decoder* processes measurement outcomes, and a *fault-tolerant protocol* specifies how noisy physical operations implement correction and computation. Good parameters for the subspace alone do not establish the performance of the full protocol.

1.1 States and operators

A pure qubit state is a unit vector

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle, \qquad |\alpha|^2 + |\beta|^2 = 1. \tag{1.1}$$

Multiplying the entire vector by a complex phase does not change its physical state. An n-qubit system has Hilbert space $(\mathbb{C}^2)^{\otimes n}$. Its computational basis is $\{|x\rangle : x \in \mathbb{F}_2^n\}$. The dimension is 2^n, but a measurement does not reveal all 2^n amplitudes.

A density operator ρ is positive semidefinite with $\operatorname{Tr}\rho = 1$. Pure states have $\rho = |\psi\rangle\langle\psi|$; mixtures have the form $\sum_j p_j|\psi_j\rangle\langle\psi_j|$. A projective measurement with mutually orthogonal projectors P_s satisfying $\sum_s P_s = I$ produces outcome s with probability $\operatorname{Tr}(P_s\rho)$. Conditional on an outcome with nonzero probability, the state is

$$\rho_s = \frac{P_s\rho P_s}{\operatorname{Tr}(P_s\rho)}. \tag{1.2}$$

Measurement may change a state. Error correction requires carefully chosen measurements, not the absence of measurement. Standard background is available in Nielsen and Chuang [2].

Symbol	Meaning and convention
$\mathcal{C}$, P	Codespace and its orthogonal projector
$\|j_L\rangle$	Encoded logical basis state
$[n, k, d]$	Classical binary code parameters
$[[n, k, d]]$	Quantum code: n data qubits, k logical qubits, distance d
X_i, Z_i	Pauli operator on qubit i, identity elsewhere
$\bar{X}, \bar{Z}$	Logical Pauli operators on an encoded qubit
S, g_j	Stabilizer group and chosen generators
$s_j \in \mathbb{F}_2$	Syndrome bit: eigenvalue $(-1)^{s_j}$ of g_j
$(x \mid z)$	Binary Pauli representation, with X coordinates first
H_X, H_Z	Rows specify X checks and Z checks, respectively
p, p_L	Physical and logical error probabilities; their time unit must be stated

Table 1: Notation used throughout. The Hadamard gate is denoted H; a classical parity-check matrix is denoted H_{cl} when confusion is possible.

1.2 Pauli and Hadamard operators

The Pauli matrices and the Hadamard matrix are

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad Y = \begin{pmatrix} 0 & -\mathrm{i} \\ \mathrm{i} & 0 \end{pmatrix}, \quad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \quad H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}. \tag{1.3}$$

They satisfy $X^2 = Y^2 = Z^2 = H^2 = I$, $XZ = -ZX$, $Y = \mathrm{i}XZ$, $HXH = Z$, and $HZH = X$. Write $|\pm\rangle = (|0\rangle \pm |1\rangle)/\sqrt{2}$. Then $Z|+\rangle = |-\rangle$ and $Z|-\rangle = |+\rangle$: a phase flip in the computational basis is a bit flip in the Hadamard basis.

The *weight* of a tensor-product Pauli is the number of its nonidentity factors. Thus $X_2 Z_5$ has weight two, whereas Y_2 has weight one.

2 Noise and the error model

2.1 Quantum channels

A quantum channel is a completely positive trace-preserving linear map. In finite dimension it admits a Kraus representation

$$\mathcal{N}(\rho) = \sum_a E_a \rho E_a^\dagger, \qquad \sum_a E_a^\dagger E_a = I. \tag{2.1}$$

The Kraus operators already include probability amplitudes. For example, a bit-flip channel has $E_0 = \sqrt{1-p}\, I$ and $E_1 = \sqrt{p}\, X$, giving $\mathcal{N}(\rho) = (1-p)\rho + pX\rho X$. An arbitrary channel need not be a mixture of one "no-error" unitary and one "error" unitary.

Channel	Action on a single-qubit density operator
Bit flip	$(1-p)\rho + pX\rho X$
Phase flip	$(1-p)\rho + pZ\rho Z$
Bit and phase flip	$(1-p)\rho + pY\rho Y$
Depolarizing, error-probability convention	$(1-p)\rho + \frac{p}{3}(X\rho X + Y\rho Y + Z\rho Z)$

Table 2: Stochastic Pauli channels, with $0 \le p \le 1$. Another convention writes depolarizing noise as $(1-q)\rho + qI/2$; in their common parameter range, $q = 4p/3$. Numerical noise rates must specify the convention.

Example 2.1 (Amplitude damping is not a Pauli mixture)**.** For $0 \le \gamma \le 1$, amplitude damping has Kraus operators

$$E_0 = \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{1-\gamma} \end{pmatrix}, \qquad E_1 = \begin{pmatrix} 0 & \sqrt{\gamma} \\ 0 & 0 \end{pmatrix}. \tag{2.2}$$

They satisfy Eq. (2.1). In particular, $|1\rangle\langle 1|$ becomes $(1-\gamma)|1\rangle\langle 1| + \gamma|0\rangle\langle 0|$, whereas $|0\rangle\langle 0|$ is unchanged. For $\gamma > 0$, this channel does not preserve $I/2$, so it cannot be a mixture of Pauli unitaries, all of which preserve $I/2$.

2.2 Why Pauli errors are sufficient for exact correction

The matrices I, X, Y, Z form a basis of all complex 2×2 matrices. Consequently any operator supported on a specified set of t qubits is a linear combination of Paulis supported there. This

statement concerns a *linear combination of operators*, not a probabilistic mixture of Pauli faults.

Section 4 proves the relevant closure property: a recovery correcting a set of errors also corrects channels whose Kraus operators lie in their linear span. Therefore correcting all Pauli errors of weight at most t suffices for arbitrary errors supported on at most t qubits. It does not correct unrestricted noise on all n qubits. Independent physical noise generally contains higher-weight components as well; their effect produces a nonzero logical failure probability.

Assumptions behind a probability calculation

Unless stated otherwise, the elementary examples assume ideal encoding, perfect syndrome measurements, and perfect recovery. When a formula uses independent physical errors, that assumption is explicit. Circuit faults, measurement errors, leakage, and correlated noise require additional analysis.

3 Repetition codes and syndrome measurement

3.1 Encoding is not cloning

The three-qubit bit-flip code is the subspace

$$\mathcal{C}_{\text{rep}} = \text{span}\{|000\rangle, |111\rangle\}, \qquad \alpha|0\rangle + \beta|1\rangle \longmapsto \alpha|000\rangle + \beta|111\rangle. \tag{3.1}$$

Two controlled-NOT gates, from the input qubit to two ancillas initialized in $|0\rangle$, implement this encoding. Here a controlled-NOT maps $|a\rangle|b\rangle$ to $|a\rangle|a \oplus b\rangle$. The result in Eq. (3.1) is generally not $(\alpha|0\rangle + \beta|1\rangle)^{\otimes 3}$, so no unknown state has been copied.

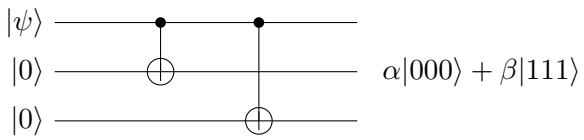


Figure 1: Encoding the bit-flip repetition code. A filled dot is the control and a circled plus is the target. Wires are read from left to right.

3.2 Parity checks reveal an error without revealing the amplitudes

Both codewords are $+1$ eigenvectors of $g_1 = Z_1Z_2$ and $g_2 = Z_2Z_3$. These commuting observables compare pairs of bits. They do not distinguish $|000\rangle$ from $|111\rangle$. Measure them and record their eigenvalues as bits s_1, s_2.

The joint syndrome projectors are

$$P_{s_1s_2} = \frac{1}{4}(I + (-1)^{s_1}Z_1Z_2)(I + (-1)^{s_2}Z_2Z_3). \tag{3.2}$$

For instance, $P_{11} = |010\rangle\langle 010| + |101\rangle\langle 101|$. If X_2 occurred, the state is $\alpha|010\rangle + \beta|101\rangle$, and outcome 11 occurs with certainty. Applying X_2 restores the encoded state. No measurement has revealed α or β.

To measure Z_1Z_2 physically, initialize an ancilla in $|0\rangle$, apply controlled-NOT gates from data qubits 1 and 2 to it, and measure the ancilla in the computational basis. On basis input $|b_1b_2\rangle$

Assumed error	g_1 eigenvalue	g_2 eigenvalue	Syndrome and recovery
I	$+1$	$+1$	00, apply I
X_1	-1	$+1$	10, apply X_1
X_2	-1	-1	11, apply X_2
X_3	$+1$	-1	01, apply X_3

Table 3: Syndrome decoding under the assumption of at most one bit flip. A syndrome is a measurement result; a stabilizer is an operator.

the ancilla stores $b_1 \oplus b_2$. This is a measurement of parity, not of the two individual data bits. The circuit describes ideal syndrome extraction; its fault tolerance is a separate question.

Example 3.1 (A coherent bit rotation)**.** Let $U = \exp(-\mathrm{i}\theta X_2) = \cos\theta\, I - \mathrm{i}\sin\theta\, X_2$. On $|\psi_L\rangle$ it produces a coherent superposition of the 00 and 11 syndrome sectors. Measurement gives these outcomes with probabilities $\cos^2\theta$ and $\sin^2\theta$. Applying the corresponding recovery from Table 3 restores $|\psi_L\rangle$ in either branch, up to a global phase. The measurement does change the corrupted state, but it does not destroy the recoverable logical amplitudes.

3.3 What the repetition code does and does not protect

Suppose each physical qubit independently undergoes X with probability p. With the decoder in Table 3, two or three flips produce a logical bit flip, so

$$p_L = \binom{3}{2} p^2(1-p) + p^3 = 3p^2 - 2p^3. \tag{3.3}$$

For example, X_1X_2 has the same syndrome as X_3. The decoder applies X_3, leaving $X_1X_2X_3$, which swaps the two codewords. For $0 < p < 1/2$, Eq. (3.3) gives $p_L < p$ under the stated ideal assumptions.

A single Z_1 changes the sign of the $|111\rangle$ component while commuting with both checks. Thus the code does not correct arbitrary single-qubit errors. Its quantum parameters are $[[3, 1, 1]]$, even though it corrects one error from the restricted bit-flip model.

Conjugating the encoding and checks by $H^{\otimes 3}$ gives the phase-flip code with codewords $|+++\rangle$ and $|---\rangle$, checks X_1X_2, X_2X_3, and recovery Z_i for a phase flip on qubit i. After recovery the state remains in this phase-flip codespace. To recover a *single unencoded qubit*, apply the inverse of the full encoding circuit; Hadamards alone do not decode it.

4 The condition for quantum error correction

Definition 4.1. A code encoding k qubits into n qubits is a 2^k-dimensional subspace $\mathcal{C} \subseteq (\mathbb{C}^2)^{\otimes n}$. A channel $\mathcal{N}$ is exactly correctable on $\mathcal{C}$ if there is a channel $\mathcal{R}$ such that $(\mathcal{R}\circ\mathcal{N})(\rho) = \rho$ for every state ρ supported on $\mathcal{C}$.

This equality also preserves entanglement with an external reference system: the composite acts as the identity on the encoded subsystem.

Theorem 4.2 (Knill–Laflamme [3])**.** *Let P project onto $\mathcal{C}$ and let E_a be Kraus operators of a channel. Exact correction is possible if and only if there is a matrix $c = (c_{ab})$ such that*

$$PE_a^\dagger E_b P = c_{ab} P \qquad \text{for all } a, b. \tag{4.1}$$

Equivalently, for any orthonormal logical basis $\{|i_L\rangle\}$, $\langle i_L|E_a^\dagger E_b|j_L\rangle = c_{ab}\delta_{ij}$.

The coefficient matrix is positive semidefinite. The condition says that the overlaps of error images contain no information that distinguishes logical basis states. In particular, for a fixed Kraus event a, the probability $\langle\psi_L|E_a^\dagger E_a|\psi_L\rangle = c_{aa}$ is independent of the logical state.

Why the condition is sufficient. Diagonalize c and make the corresponding unitary change of Kraus operators, obtaining F_a with $PF_a^\dagger F_b P = \lambda_a \delta_{ab} P$. For $\lambda_a > 0$, define $V_a = F_a P/\sqrt{\lambda_a}$. Restricted to $\mathcal{C}$, each V_a is an isometry, and different $V_a\mathcal{C}$ are orthogonal. Measure which of these error subspaces contains the state and apply the inverse isometry there. The output is the original logical state. Zero-eigenvalue terms vanish on $\mathcal{C}$, and the recovery can be completed arbitrarily on the remaining orthogonal complement to form a channel. The necessity direction is proved in [3]. □

Proposition 4.3 (Closure under linear combinations)**.** *If Eq. (4.1) holds for a set $\{E_a\}$, it holds for every set $\{F_\mu\}$ contained in its linear span.*

Proof. Write $F_\mu = \sum_a u_{\mu a} E_a$. Then

$$PF_\mu^\dagger F_\nu P = \left(\sum_{a,b} \overline{u_{\mu a}} u_{\nu b} c_{ab} \right) P.$$

The recovery obtained by separating orthogonal error subspaces works for all channels with Kraus operators in this span. □

This explains why continuous error amplitudes do not require continuously many distinct correction rules. It is the operator span and its action on the code that matter.

4.1 Detection, correction, and distance

An operator E is *detectable* by $\mathcal{C}$ when PEP is a scalar multiple of P. If the scalar is nonzero, this does not promise an alarm: an error acting as a scalar is harmless on the code. Exact correction of a family $\{E_a\}$ requires detectability of all products $E_a^\dagger E_b$.

For a nontrivial qubit code, its distance d is the smallest Pauli weight for which PEP is not a scalar multiple of P. A $[[n, k, d]]$ code detects every Pauli error of weight below d and corrects arbitrary errors on at most

$$t = \left\lfloor \frac{d-1}{2} \right\rfloor \tag{4.2}$$

unknown locations. Indeed, the product of two errors of weight at most t has weight at most $2t < d$, so the Knill–Laflamme condition holds. If the faulty locations are known, the code can instead correct erasures on any set of at most $d-1$ qubits: the relevant products remain supported on that same known set. The distinction is between unknown and known locations, not between two different values of the code distance.

5 Stabilizer codes and binary linear algebra

5.1 Commuting checks define the codespace

The n-qubit Pauli group consists of tensor products of I, X, Y, Z with phases $\{\pm 1, \pm \mathrm{i}\}$. A stabilizer group S is an abelian subgroup not containing $-I$. Its codespace is the joint $+1$ eigenspace

$$\mathcal{C}(S) = \{|\psi\rangle : g|\psi\rangle = |\psi\rangle \text{ for every } g \in S\}. \tag{5.1}$$

The exclusion of $-I$ is essential: no nonzero vector is its $+1$ eigenvector. Stabilizer elements are Hermitian with square I. If $g_1, \ldots, g_r$ are independent generators, the codespace projector is

$$P = \prod_{j=1}^{r} \frac{I + g_j}{2} = \frac{1}{|S|} \sum_{g \in S} g. \tag{5.2}$$

Only the identity Pauli has nonzero trace, so $\operatorname{Tr} P = 2^{n-r}$. Thus the code encodes $k = n - r$ logical qubits [4]. Redundant checks may also be measured, but they do not further reduce k.

For a Pauli error E, define $s_j(E)$ by $g_j E = (-1)^{s_j(E)} E g_j$. If $|\psi\rangle \in \mathcal{C}(S)$, then

$$g_j E|\psi\rangle = (-1)^{s_j(E)} E|\psi\rangle. \tag{5.3}$$

This gives the syndrome directly from commutation, without expanding state vectors of dimension 2^n.

5.2 Logical operators and degeneracy

Let $\mathrm{Cent}(S)$ denote the Pauli operators commuting with every stabilizer. Throughout distance and equivalence statements, ignore overall Pauli phases. There are three cases:

1. $E \notin \mathrm{Cent}(S)$: some check has eigenvalue -1, giving a nonzero syndrome.
2. E is proportional to an element of S: it acts as a scalar on the codespace and is physically harmless there.
3. $E \in \mathrm{Cent}(S)$ but is not proportional to any stabilizer: it acts as a nontrivial logical Pauli and has zero syndrome.

In the third case, some logical eigenstates may be unchanged up to phase, even though the operator is nontrivial on the entire codespace. For $k \geq 1$,

$$d = \min\{\mathrm{wt}(E) : E \in \mathrm{Cent}(S),\ E \not\propto S\}. \tag{5.4}$$

Two Pauli errors E, F have the same syndrome precisely when $E^\dagger F$ commutes with all checks. They have the same correctable effect when $E^\dagger F$ is proportional to a stabilizer. Otherwise they differ by a nontrivial logical operator. A code is *degenerate for an error set* if distinct errors in that set act identically on the code up to phase. Correcting an error therefore need not identify its precise physical location.

If a decoder returns a recovery Pauli R with the observed syndrome, then RE has zero syndrome. Recovery succeeds if $RE \propto S$; a zero residual syndrome alone is not enough.

5.3 The symplectic representation

For $x, z \in \mathbb{F}_2^n$, write $X^x Z^z = \bigotimes_i X_i^{x_i} Z_i^{z_i}$. Modulo phases, encode it by the row vector $(x \mid z)$. The commutation rule is

$$(X^x Z^z)(X^{x'} Z^{z'}) = (-1)^{x \cdot z' + z \cdot x'} (X^{x'} Z^{z'})(X^x Z^z). \tag{5.5}$$

All binary sums and products here are in $\mathbb{F}_2$. If $M = (A \mid B)$ has the generator vectors as rows, the checks commute exactly when

$$AB^T + BA^T = 0. \tag{5.6}$$

The syndrome of $E \leftrightarrow (x \mid z)$ is

$$s(E) = Az^T + Bx^T. \tag{5.7}$$

For a valid stabilizer, $k = n - \operatorname{rank} M$. Binary rows do not encode generator signs: one must also choose consistent Hermitian representatives whose group excludes $-I$. For independent commuting rows, assigning signs consistently defines a codespace; with dependent rows, phase consistency must be checked.

Example 5.1 (A four-qubit detection code)**.** Take $g_1 = XXXX$ and $g_2 = ZZZZ$. They commute because their supports overlap on four qubits. There are two independent checks, hence $k = 4 - 2 = 2$. One orthonormal basis is

$$\frac{|0000\rangle + |1111\rangle}{\sqrt{2}}, \quad \frac{|0110\rangle + |1001\rangle}{\sqrt{2}}, \quad \frac{|1010\rangle + |0101\rangle}{\sqrt{2}}, \quad \frac{|1100\rangle + |0011\rangle}{\sqrt{2}}.$$

Every single-qubit X, Y, or Z anticommutes with a check. However, $Z_1 Z_2$ commutes with both and is not a stabilizer, so $d = 2$. This is a $[[4, 2, 2]]$ code: it detects one arbitrary Pauli error and corrects one known erasure, but cannot correct an arbitrary error at an unknown location. For example, X_1 and X_2 have the same syndrome and differ by a nontrivial logical Pauli.

6 The Shor code and the role of degeneracy

The Shor code [1] combines the phase-flip repetition code with an inner bit-flip repetition encoding. With three blocks of three qubits, define

$$|0_L\rangle = \left(\frac{|000\rangle + |111\rangle}{\sqrt{2}}\right)^{\otimes 3}, \qquad |1_L\rangle = \left(\frac{|000\rangle - |111\rangle}{\sqrt{2}}\right)^{\otimes 3}. \tag{6.1}$$

Six checks locate bit flips within blocks, and two checks compare block phases:

$$\begin{aligned} &Z_1Z_2,\ Z_2Z_3,\ Z_4Z_5,\ Z_5Z_6,\ Z_7Z_8,\ Z_8Z_9, \\ &X_1X_2X_3X_4X_5X_6, \quad X_4X_5X_6X_7X_8X_9. \end{aligned} \tag{6.2}$$

They are independent and commuting, so $k = 9 - 8 = 1$. A single X error is identified within a block; a single Z error is identified only up to its block. A Y error has both contributions to the syndrome.

Example 6.1 (A shared syndrome is harmless when the difference is a stabilizer)**.** The errors Z_1, Z_2, Z_3 all commute with the six Z checks, anticommute with the first long X check, and commute with the second. They therefore share a syndrome. Yet $Z_1Z_2 \in S$ and $Z_1Z_3 = (Z_1Z_2)(Z_2Z_3) \in S$. If the actual error is Z_2 and recovery applies Z_1, then Z_1Z_2 acts as the identity on every encoded state. This is degeneracy, not a decoding failure.

Logical representatives consistent with Eq. (6.1) are

$$\bar{X} = Z_1 Z_4 Z_7, \qquad \bar{Z} = X_1 X_2 X_3. \tag{6.3}$$

The first swaps the two codewords; the second gives eigenvalues $+1, -1$. Both commute with all checks and anticommute with each other. Their weight is three.

To see that no smaller nontrivial logical Pauli exists, use $(x \mid z)$. Commutation with the two Z checks in each block forces the three x entries in that block to agree. A commuting Pauli of weight at most two must therefore have $x = 0$. Commutation with the long X checks forces the three block parities of z to agree. With at most two nonzero entries, all these parities are zero; the resulting Z operator is a product of within-block ZZ stabilizers. Thus $d = 3$, and the code is $[[9, 1, 3]]$.

The syndrome analysis shows that every pair of single-qubit Pauli errors is either distinguishable or equivalent modulo the stabilizer. The code therefore satisfies the correction condition for all weight-one Paulis, and Proposition 4.3 extends this to arbitrary single-qubit noise. Neither this statement nor the distance-three label implies that every error of weight two is uncorrectable: some such errors are stabilizers or otherwise happen to be recovered.

7 Classical linear codes and the CSS construction

7.1 The classical ingredients

A binary linear code $C \subseteq \mathbb{F}_2^n$ of dimension k is an $[n, k]$ code. Its Hamming distance is $d(C) = \min_{0 \neq c \in C} \mathrm{wt}(c)$, where $\mathrm{wt}(c)$ counts nonzero coordinates. The distance between words u, v is $\mathrm{wt}(u + v)$. A generator matrix has rows spanning C; a parity-check matrix H_{cl} has kernel C, with codewords viewed as column vectors when multiplying by H_{cl}.

The dual code is $C^{\perp} = \{u : u \cdot c = 0 \ \forall c \in C\}$. It has dimension $n - k$ and is the row space of any full-rank parity-check matrix. If $c \in C$ is sent and $c + e$ is received, its syndrome is $H_{\mathrm{cl}}(c + e)^T = H_{\mathrm{cl}} e^T$. Distance d allows unique correction of at most $\lfloor (d - 1)/2 \rfloor$ arbitrary bit errors.

7.2 Nested codes and quantum codewords

Let $C_2 \subseteq C_1 \subseteq \mathbb{F}_2^n$ have dimensions k_2, k_1. The CSS construction [5, 6] associates a logical basis state with each coset of C_2 in C_1:

$$|x + C_2\rangle = \frac{1}{\sqrt{|C_2|}} \sum_{y \in C_2} |x + y\rangle, \qquad x \in C_1/C_2. \tag{7.1}$$

The notation $x \in C_1/C_2$ means that one representative is selected per coset. Distinct cosets are disjoint, so the states are orthonormal. There are $2^{k_1 - k_2}$ of them; hence $k = k_1 - k_2$ logical qubits.

For $u \in C_2$, the operator X^u permutes the terms in Eq. (7.1). For $v \in C_1^{\perp}$, the phase $(-1)^{v \cdot (x+y)}$ is one. Thus the stabilizers are generated by

$$X^u \quad (u \text{ in a basis of } C_2), \qquad Z^v \quad (v \text{ in a basis of } C_1^{\perp}). \tag{7.2}$$

They commute because $u \cdot v = 0$. This calculation gives a direct bridge between coset states and stabilizer matrices.

7.3 Matrix convention and distance

Choose H_X with row space C_2 and H_Z with row space $C_1^\perp$. The full binary check matrix, in the $(x \mid z)$ convention, is the *block matrix*

$$M_{\mathrm{CSS}} = \begin{pmatrix} H_X & 0 \\ 0 & H_Z \end{pmatrix}, \qquad H_X H_Z^T = 0. \tag{7.3}$$

The matrices H_X, H_Z can have different numbers of rows. They are not concatenated horizontally as if each row were a mixed X/Z check. With redundant rows allowed,

$$k = n - \operatorname{rank} H_X - \operatorname{rank} H_Z. \tag{7.4}$$

For $k > 0$, the minimum weights of nontrivial pure logical operators are

$$\begin{aligned} d_X &= \min\{\mathrm{wt}(x) : x \in \ker H_Z \setminus \operatorname{row} H_X\} = \min_{x \in C_1 \setminus C_2} \mathrm{wt}(x), \\ d_Z &= \min\{\mathrm{wt}(z) : z \in \ker H_X \setminus \operatorname{row} H_Z\} = \min_{z \in C_2^\perp \setminus C_1^\perp} \mathrm{wt}(z), \\ d &= \min(d_X, d_Z). \end{aligned} \tag{7.5}$$

To justify the last line, a mixed commuting Pauli $X^x Z^z$ is nontrivial only if at least one component is outside its stabilizer row space. Its support contains the support of that component. Conversely, a pure logical operator of weight d_X or d_Z realizes the upper bound. In particular, $d \geq \min\{d(C_1), d(C_2^\perp)\}$, when these classical distances are defined. The exclusion of stabilizer subspaces in Eq. (7.5) is essential for the exact distance.

For $E = X^e Z^f$, the measured syndromes are

$$s_X = H_X f^T, \qquad s_Z = H_Z e^T. \tag{7.6}$$

Thus X checks reveal the phase-error component, and Z checks reveal the bit-error component. The subscript names the check, not the error it detects. There is no need to physically apply Hadamards to the data to define CSS phase decoding. If a circuit does change basis with Hadamards, the recovery operators must also be conjugated consistently.

7.4 A complete seven-qubit example

Consider the Hamming parity-check matrix

$$H_{\mathrm{Ham}} = \begin{pmatrix} 1 & 0 & 1 & 0 & 1 & 0 & 1 \\ 0 & 1 & 1 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 & 1 & 1 & 1 \end{pmatrix}. \tag{7.7}$$

Its seven columns are the distinct nonzero three-bit vectors. Its rows are independent, have even weight, and overlap pairwise on two positions, so $H_{\mathrm{Ham}} H_{\mathrm{Ham}}^T = 0$. Put $C_1 = \ker H_{\mathrm{Ham}}$ and $C_2 = \operatorname{row} H_{\mathrm{Ham}} = C_1^\perp$. Then $C_2 \subset C_1$, with $k_1 = 4, k_2 = 3$.

Taking $H_X = H_Z = H_{\mathrm{Ham}}$ gives the Steane code. Its six checks are

$$\begin{aligned} &X_1X_3X_5X_7, \quad X_2X_3X_6X_7, \quad X_4X_5X_6X_7, \\ &Z_1Z_3Z_5Z_7, \quad Z_2Z_3Z_6Z_7, \quad Z_4Z_5Z_6Z_7. \end{aligned} \tag{7.8}$$

There are $7 - 3 - 3 = 1$ logical qubits. No weight-one or weight-two nonzero vector lies in $\ker H_{\mathrm{Ham}}$: such a vector would make a column zero or two columns equal. The vector

$(1,1,1,0,0,0,0)$ is in the kernel because the first three columns sum to zero. Every nonzero vector of row H_{Ham} has weight four, so that weight-three vector is not a stabilizer vector. Therefore $d_X = d_Z = 3$ and the code is $[[7,1,3]]$.

A convenient logical basis and pair of logical operators are

$$|0_L\rangle = \frac{1}{\sqrt{8}} \sum_{y \in C_2} |y\rangle, \qquad |1_L\rangle = X^{\otimes 7}|0_L\rangle, \tag{7.9}$$
$$\bar{X} = X^{\otimes 7}, \qquad \bar{Z} = Z^{\otimes 7}.$$

The all-ones vector belongs to C_1 but not C_2. The weight-seven representatives above are convenient, but not minimal: $X_1X_2X_3$ and $Z_1Z_2Z_3$ are equivalent representatives after multiplying by the corresponding check on qubits $4,5,6,7$.

Example 7.1 (Decoding Y_5)**.** Ignoring its overall phase, $Y_5 = X_5Z_5$. Both syndrome halves equal column 5 of H_{Ham}:

$$s_X = (1,0,1)^T, \qquad s_Z = (1,0,1)^T.$$

Looking up that column identifies qubit 5 in both halves. Applying Y_5, or the corresponding X and Z corrections in either order, restores the state up to global phase. For X_5 alone, $s_X = 0$ and $s_Z = (1,0,1)^T$; for Z_5 the roles reverse.

The same separate decoder can correct some errors outside the distance guarantee. For example, X_5Z_2 has a single error in each component, so the two column lookups recover it despite its total weight two. This does not contradict distance three: a guarantee for *all* errors up to a weight is stronger than successful recovery of some heavier errors. Under correlated Pauli noise, decoding the two components independently need not be statistically optimal because a Y fault correlates them.

8 From stabilizers to geometry

8.1 The toric code

Kitaev's toric code [7] gives a geometric realization of CSS checks. Take an $L \times L$ square lattice with periodic boundaries, with $L \geq 3$, and put a data qubit on every edge. There are $n = 2L^2$ data qubits. For a vertex v and face f, define

$$A_v = \prod_{e \ni v} X_e, \qquad B_f = \prod_{e \in \partial f} Z_e. \tag{8.1}$$

A star and a plaquette overlap on zero or two edges, so they commute. Products of all stars and all plaquettes are both identity because each edge occurs twice. On the connected torus these are the only dependencies within the two check types. Hence the stabilizer rank is $2L^2 - 2$ and $k = 2$.

8.2 Syndromes are boundaries of error chains

A Z error on one edge anticommutes with the stars at its endpoints. More generally, a Z chain has violated star checks at its boundary vertices. A closed Z loop has no such endpoints and hence no syndrome. If it bounds a collection of faces, it is a product of their plaquette



A_v: four X factors B_f: four Z factors

Figure 2: Local supports of a star check and a plaquette check. Dots on edges mark the qubits involved. The panels are local lattice fragments, not separate planar codes; the toric construction identifies opposite boundaries of the full lattice.

stabilizers and is harmless. If it winds nontrivially around the torus, it is a logical operator. X chains have the analogous description on the dual lattice, where vertices correspond to original faces.

For a compact algebraic formulation, let ∂_2 map faces to their boundary edges, and ∂_1 map edges to their endpoint vertices, all over $\mathbb{F}_2$. A boundary has no boundary:

$$\partial_1 \partial_2 = 0. \tag{8.2}$$

The CSS matrices are $H_X = \partial_1$ and $H_Z = \partial_2^T$, so this identity is exactly $H_X H_Z^T = 0$. Logical Z classes form $\ker \partial_1 / \operatorname{im} \partial_2$, the first homology of the torus over $\mathbb{F}_2$. Its dimension is two, agreeing with the stabilizer count.

A shortest noncontractible loop has length L, on both primal and dual lattices. Thus the parameters are

$$[[2L^2, 2, L]]. \tag{8.3}$$

One can choose logical Z loops in the two independent winding directions, and logical X loops on the dual lattice crossing the corresponding Z loops once. A single crossing supplies the required logical anticommutation.

Example 8.1 (Why matching the syndrome is not sufficient)**.** Suppose a Z error occupies a path E between two vertices and a decoder chooses another path R with the same endpoints. The residual RE is a closed cycle, with edges added modulo two. If the cycle is a boundary, correction succeeds. If it winds around the torus, a logical error remains. The decoder's task is therefore to choose the right equivalence class of chains, not merely to remove the endpoints.

8.3 Planar surface codes

Surface-code patches replace periodic boundaries by suitable boundary checks. A standard rotated distance-d planar patch encodes one logical qubit into d^2 data qubits. A common extraction layout uses a further $d^2 - 1$ measurement ancillas, for $2d^2 - 1$ qubits before extra hardware-specific resources. This count is different from the toric code's $2L^2$ *data* qubits and two logical qubits. Boundary conventions and layouts matter [8].

Minimum-weight perfect matching is one common decoding approach for suitable surface-code noise models. Faulty syndrome measurements require information across repeated rounds: changes in check outcomes, often called *detection events*, are decoded in spacetime. Neither two-dimensional placement nor distance alone specifies the extraction circuit, decoder, or achievable logical error rate.

9 Subsystem codes and gauge freedom

An ordinary subspace code protects every degree of freedom in its codespace. A subsystem code instead separates the information we want to preserve from auxiliary degrees of freedom whose state may change. This is the operator quantum error-correction viewpoint [9, 10].

9.1 What is protected?

Write the physical Hilbert space as

$$\mathcal{H}_{\text{phys}} = (\mathcal{H}_L \otimes \mathcal{H}_G) \oplus \mathcal{H}_{\text{rest}}, \qquad \dim \mathcal{H}_L = 2^k, \quad \dim \mathcal{H}_G = 2^r. \tag{9.1}$$

Here L is the logical subsystem and G is the *gauge* subsystem. Recovery must restore the logical information, including its entanglement with a reference system, but need not restore the gauge state. These tensor factors describe an encoded decomposition; gauge qubits need not be a particular subset of the physical qubits.

Let P project onto $\mathcal{H}_L \otimes \mathcal{H}_G$. With this identification, the subsystem version of the correction condition is

$$P E_a^\dagger E_b P\Big|_{\mathcal{H}_L \otimes \mathcal{H}_G} = I_L \otimes M_{ab} \tag{9.2}$$

for operators M_{ab} on the gauge subsystem. Compare this with the scalar condition in Section 4: error overlaps may distinguish gauge states, but not logical states. For $r = 0$, Eq. (9.2) reduces to the ordinary Knill–Laflamme condition.

9.2 Gauge generators, stabilizers, and logical operators

For a Pauli subsystem code, the *gauge group* $\mathcal{G}$ may contain anticommuting operators. Its elements act trivially on L within the chosen codespace, while possibly changing G. The stabilizer S consists of signed commuting representatives from the center of $\mathcal{G}$; as usual, choose consistent eigenvalues and exclude $-I$ from S. All ranks below refer to binary Pauli vectors, so scalar phases are removed.

If the binary ranks of $\mathcal{G}$ and S are g and s, respectively, then

$$r = \frac{g - s}{2}, \qquad k = n - s - r = n - \frac{g + s}{2}. \tag{9.3}$$

The stabilizer fixes a 2^{n-s}-dimensional space, partitioned into k logical and r gauge qubits. A *bare* logical Pauli commutes with every gauge generator. Multiplying it by a gauge operator gives a *dressed* representative with the same logical action. The distance relevant to arbitrary local Pauli errors is the dressed distance

$$d = \min\{\text{wt}(E) : E \in \text{Cent}(S) \setminus \mathcal{G}\}, \tag{9.4}$$

where phases are again ignored. A low-weight gauge operator is harmless to L, even when it is not a stabilizer. We use $[[n, k, r, d]]$ for these four subsystem parameters [10].

9.3 Worked example: the nine-qubit Bacon–Shor code

Place nine qubits at positions (i, j) of a 3×3 array, with i the row and j the column. Choose horizontal XX and vertical ZZ gauge generators:

$$g^X_{i,j} = X_{i,j} X_{i,j+1} \qquad (i = 1, 2, 3;\ j = 1, 2), \tag{9.5}$$

$$g^Z_{i,j} = Z_{i,j} Z_{i+1,j} \qquad (i = 1, 2;\ j = 1, 2, 3). \tag{9.6}$$

An XX pair and a ZZ pair sharing exactly one site anticommute. Consequently, the code is not the simultaneous $+1$ eigenspace of all twelve gauge generators. Instead its four independent stabilizers are

$$S_j^X = \prod_{i=1}^{3} g_{i,j}^X \quad (j = 1, 2), \qquad S_i^Z = \prod_{j=1}^{3} g_{i,j}^Z \quad (i = 1, 2). \tag{9.7}$$

Each has weight six and commutes with every gauge generator. There are $g = 12$ independent gauge vectors and $s = 4$ independent central vectors. Equation (9.3) gives $r = 4$ and $k = 1$. This convention for the Bacon–Shor construction is equivalent to interchanging rows, columns, and Pauli types in other presentations [11].

Bare logical representatives are a column of X operators and a row of Z operators:

$$\bar{X} = \prod_{i=1}^{3} X_{i,1}, \qquad \bar{Z} = \prod_{j=1}^{3} Z_{1,j}. \tag{9.8}$$

Each commutes with every gauge generator; their single intersection makes them anticommute. Thus $d \leq 3$. To prove equality, let a Pauli E of weight at most two commute with all four stabilizers. Commutation with the two Z stabilizers forces the three row parities of the X component of E to agree. They must all be zero, since three odd rows would require at least three occupied sites. Every even X pattern in a row is generated by the horizontal XX pairs. Similarly, the Z component has even parity in every column and is generated by the vertical ZZ pairs. Therefore $E \in \mathcal{G}$, so no lighter dressed logical exists. The code has parameters $[[9, 1, 4, 3]]$.

This example reduces the weight of the operators one measures, not the weight of the stabilizers themselves. In an ideal round of mutually commuting XX gauge measurements, multiplying the appropriate outcomes yields each S_j^X eigenvalue. A separate ZZ round gives the Z stabilizer outcomes. Gauge measurements from different types can disturb one another's individual outcomes, but preserve the stabilizer sector. With noisy measurements, the extraction schedule, outcome processing, and repeated information must be analyzed explicitly.

Gauge freedom and gauge fixing

A recovery is successful when its residual action changes only the gauge subsystem. *Gauge fixing* selects additional commuting gauge observables and fixes their eigenvalues, reducing gauge freedom while retaining the protected logical information. Suitable changes of gauge can make different logical gates available; Section 11 returns to this use. The existence of gauge freedom alone does not imply single-shot correction or a universal gate set.

10 Bosonic codes: using an oscillator

A bosonic mode has the infinite-dimensional Fock basis $\{|m\rangle : m = 0, 1, 2, \ldots\}$. Encoding into a few carefully chosen superpositions of these levels can protect a logical qubit without using a tensor product of many two-level data systems. The resource accounting changes: oscillator energy, control operations, measurement ancillas, and couplings all matter. "One mode" is not the same resource as one physical qubit.

10.1 Photon loss and the choice of error set

The annihilation and number operators satisfy

$$a|m\rangle = \sqrt{m}|m-1\rangle, \qquad \hat{n} = a^\dagger a, \qquad \hat{n}|m\rangle = m|m\rangle. \tag{10.1}$$

For a pure-loss channel with transmissivity $0 \le \eta \le 1$, one Kraus representation is

$$E_\ell = \sqrt{\frac{(1-\eta)^\ell}{\ell!}}\,\eta^{\hat{n}/2}a^\ell, \qquad \ell = 0, 1, 2, \ldots. \tag{10.2}$$

Indeed, $E_\ell|m\rangle = \sqrt{\binom{m}{\ell}}(1-\eta)^{\ell/2}\eta^{(m-\ell)/2}|m-\ell\rangle$ for $m \ge \ell$. Squaring and summing these amplitudes gives one by the binomial theorem. The no-loss operator $E_0 = \eta^{\hat{n}/2}$ is not the identity. Thus correcting the selected error set $\{I, a\}$ exactly is a useful mathematical property, but does not mean exact correction of the full finite-time loss channel [12].

10.2 Worked example: a single-loss binomial code

A simple binomial encoding is

$$|0_L\rangle = \frac{|0\rangle + |4\rangle}{\sqrt{2}}, \qquad |1_L\rangle = |2\rangle. \tag{10.3}$$

Both codewords have even photon-number parity. Their single-loss images are

$$a|0_L\rangle = \sqrt{2}|3\rangle, \qquad a|1_L\rangle = \sqrt{2}|1\rangle. \tag{10.4}$$

With P the code projector, direct calculation gives

$$PIP = P, \qquad PaP = Pa^\dagger P = 0, \qquad Pa^\dagger aP = 2P. \tag{10.5}$$

These are precisely the four Knill–Laflamme products for $\{I, a\}$. The equal mean photon number, two in both codewords, prevents the loss probability from identifying the logical basis state in this error model. Measuring parity distinguishes the even code sector from the odd single-loss sector without distinguishing α from β in either. In the odd sector, the isometry $|3\rangle \mapsto |0_L\rangle$, $|1\rangle \mapsto |1_L\rangle$ recovers the logical state. This map can be extended to a physical recovery using an ancilla and a reset.

To see the limitation of the selected error set, apply the actual no-loss operator:

$$E_0|0_L\rangle = \frac{|0\rangle + \eta^2|4\rangle}{\sqrt{2}}, \qquad E_0|1_L\rangle = \eta|2\rangle. \tag{10.6}$$

The corresponding squared norms are $(1+\eta^4)/2$ and η^2. Their difference is $(1-\eta^2)^2/2$, which is nonzero for $0 \le \eta < 1$. Therefore $PE_0^\dagger E_0 P$ is not a scalar multiple of P, and this encoding does not exactly correct the full channel of Eq. (10.2). For a short interval the difference is second order in $1-\eta$; realistic recovery must still address no-loss deformation, multiple losses, and imperfections in the correction cycle. Larger binomial constructions are designed for specified sets of loss, gain, and dephasing errors [12].

10.3 Cat encodings and why parity needs care

A coherent state obeys $a|\alpha\rangle = \alpha|\alpha\rangle$. For $\alpha \ne 0$, normalized even and odd cat states are

$$|C_\alpha^\pm\rangle = N_\pm(|\alpha\rangle \pm |-\alpha\rangle), \qquad N_\pm = \left[2(1 \pm e^{-2|\alpha|^2})\right]^{-1/2}. \tag{10.7}$$

They are orthogonal because they have opposite parity. However,

$$a|C_\alpha^+\rangle = \alpha\frac{N_+}{N_-}|C_\alpha^-\rangle, \qquad a|C_\alpha^-\rangle = \alpha\frac{N_-}{N_+}|C_\alpha^+\rangle. \tag{10.8}$$

If these two states form the logical basis, a single loss stays inside the code and changes the logical state. Here PaP has nonzero off-diagonal entries, so it fails the scalar correction condition for $\{I, a\}$. Measuring parity would measure logical Z, collapsing a general encoded superposition. This contrasts with Eq. (10.3), where both logical basis states have the same parity.

"Cat code" therefore does not specify one error model or one recovery. Stabilized two-component encodings can offer strongly asymmetric protection, while suitable four-component constructions place the logical states in the same parity sector and use parity changes to monitor loss. Mirrahimi and coauthors analyze both approaches and associated protected operations [13]. An additional outer code may be needed to correct the dominant remaining logical error. Always specify the basis, stabilization mechanism, noise bias, and measurements before claiming protection.

10.4 Gottesman–Kitaev–Preskill grid codes

The Gottesman–Kitaev–Preskill (GKP) encoding protects against small displacements in phase space [14]. Set $\hbar = 1$ and use dimensionless quadratures

$$q = \frac{a + a^\dagger}{\sqrt{2}}, \qquad p = \frac{a - a^\dagger}{\mathrm{i}\sqrt{2}}, \qquad [q, p] = \mathrm{i}. \tag{10.9}$$

The ideal square-lattice code has displacement stabilizers

$$U_q = e^{-2\mathrm{i}\sqrt{\pi}p}, \qquad U_p = e^{2\mathrm{i}\sqrt{\pi}q}. \tag{10.10}$$

The Weyl relation $e^{-\mathrm{i}up}e^{\mathrm{i}vq} = e^{-\mathrm{i}uv}e^{\mathrm{i}vq}e^{-\mathrm{i}up}$ makes these commute, since $uv = 4\pi$. Formal logical codewords are position combs

$$|\mu_L\rangle \propto \sum_{s\in\mathbb{Z}} |q = (2s + \mu)\sqrt{\pi}\rangle, \qquad \mu \in \{0, 1\}. \tag{10.11}$$

They are distributions, not normalizable finite-energy states. Physical approximations replace the peaks by finite-width wave packets and impose an envelope. Finite squeezing and finite energy must be included in an implementation's noise analysis.

Logical Paulis are half-stabilizer displacements,

$$\bar{X} = e^{-\mathrm{i}\sqrt{\pi}p}, \qquad \bar{Z} = e^{\mathrm{i}\sqrt{\pi}q}. \tag{10.12}$$

Their Weyl phase is $e^{-\mathrm{i}\pi} = -1$, so they anticommute. For an error $D(u, v) = e^{-\mathrm{i}up}e^{\mathrm{i}vq}$, stabilizer measurements determine the shifts modulo $\sqrt{\pi}$. Ideal nearest-lattice recovery succeeds throughout the open cell

$$|u| < \frac{\sqrt{\pi}}{2}, \qquad |v| < \frac{\sqrt{\pi}}{2}. \tag{10.13}$$

Larger displacements may be mistaken for another lattice point and leave a logical Pauli. This cell statement assumes ideal code states and exact syndrome information. It is not a threshold for finite-energy GKP hardware or a direct guarantee for arbitrary photon loss.

11 Universal fault-tolerant quantum computation

Protecting a state while it waits is one task; manipulating it reliably throughout an algorithm is another. A universal logical gate set can approximate any unitary on the encoded qubits to arbitrary accuracy. A fault-tolerant computer must combine such gates with reliable preparation, syndrome extraction, classical processing, and readout under a specified noise model [15].

Encoding	Useful mathematical picture	Essential qualification
Binomial	Finite Fock superpositions; prescribed error products	A finite correctable error set is not the full loss channel
Cat	Coherent-state superpositions and stabilization	The basis and number of components determine which errors are detected or biased
GKP	A lattice of displacement stabilizers	Ideal combs have infinite energy; physical states give approximate protection

Table 4: Three bosonic approaches. Qubit Hamming weight and $[[n, k, d]]$ alone do not describe their noise protection or physical cost.

11.1 Fault propagation and transversal gates

Code distance describes errors on a block at an instant. Circuit faults can create several such errors. For example, conjugation by a controlled-NOT with control c and target t gives

$$X_c \mapsto X_c X_t, \quad Z_c \mapsto Z_c, \quad X_t \mapsto X_t, \quad Z_t \mapsto Z_c Z_t. \tag{11.1}$$

An ancilla interacting with many data qubits can therefore spread a fault. Fault-tolerant gadgets must limit this spread and remain reliable when the incoming data and syndrome measurements are imperfect. For a distance-$2t + 1$ code, merely saying "at most t faults" is insufficient unless the gadget analysis bounds the residual data errors those faults can produce.

A transversal gate factors into operations on disjoint sets, with each set containing at most one physical subsystem from each participating code block. A single faulty paired gate then affects at most one qubit per block, under the usual local gate-fault model. This makes transversality a useful construction principle. Preparation, error correction, and measurement still require their own analysis; locality of one logical gate does not prove a complete threshold theorem.

The Steane code provides a concrete example. Since its X and Z check matrices agree, applying H to all seven data qubits exchanges the check types and implements logical H. Pairwise CNOT between two blocks implements logical CNOT. Write the phase gate as $S_{\mathrm{ph}} = \mathrm{diag}(1, \mathrm{i})$ to distinguish it from a stabilizer group. In the logical basis of Section 7, the $|0_L\rangle$ strings have weights zero or four, while the $|1_L\rangle$ strings have weights three or seven. Hence $S_{\mathrm{ph}}^{\otimes 7}$ acts as logical $S_{\mathrm{ph}}^{\dagger}$, and $(S_{\mathrm{ph}}^{\dagger})^{\otimes 7}$ implements logical S_{ph}. Physical and logical phase conventions matter.

11.2 Why Clifford gates are not the whole answer

The Clifford group consists of unitaries that map Pauli operators to Pauli operators under conjugation. It is generated by H, S_{ph}, and CNOT. These operations alone are not a universal gate set. Adding

$$T = \mathrm{diag}(1, e^{\mathrm{i}\pi/4}) \tag{11.2}$$

gives the familiar universal Clifford+T set. For example, $TXT^{\dagger} = (X + Y)/\sqrt{2}$, so T is not Clifford.

The Eastin–Knill theorem rules out a universal set consisting entirely of transversal encoded gates on a fixed finite-dimensional code that detects arbitrary errors on any one physical subsystem [16]. It does not rule out universal fault-tolerant computation. It tells us that the full construction needs something beyond transversal gates on that one fixed code, such as resource-state preparation and measurement or a change of code. The finite-dimensional assumptions should not be silently applied to an ideal infinite-energy GKP construction.

11.3 Worked example: injecting a non-Clifford gate

Suppose a high-quality resource state

$$|A\rangle = T|+\rangle = \frac{|0\rangle + e^{i\pi/4}|1\rangle}{\sqrt{2}} \tag{11.3}$$

is available. Let the data be $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$. Apply CNOT with the data as control and $|A\rangle$ as target, then measure the target in the computational basis. The unnormalized data branches are

$$\begin{aligned} m = 0: &\quad \frac{\alpha|0\rangle + e^{i\pi/4}\beta|1\rangle}{\sqrt{2}} = \frac{T|\psi\rangle}{\sqrt{2}}, \\ m = 1: &\quad \frac{e^{i\pi/4}\alpha|0\rangle + \beta|1\rangle}{\sqrt{2}} = \frac{e^{i\pi/4}T^\dagger|\psi\rangle}{\sqrt{2}}. \end{aligned} \tag{11.4}$$

Each outcome occurs with probability $1/2$. Applying S_{ph} to the data when $m = 1$ gives $T|\psi\rangle$ up to a global phase, since $S_{\mathrm{ph}}T^\dagger = T$. Thus Clifford gates, a Pauli-basis measurement, feed-forward, and a suitable non-stabilizer resource state implement the desired non-Clifford gate.

At the logical level this algebra is unchanged, but the quality of the encoded resource matters. Injecting a noisy state can transfer its error to the data. *Magic-state distillation* uses multiple imperfect resource states and Clifford processing to produce fewer, better states under specified assumptions [17, 18]. Its acceptance probability, input-noise requirements, residual error, and space–time cost depend on the protocol. The ideal calculation above demonstrates gate injection; by itself it is not a fault-tolerance proof for raw noisy ancillas.

11.4 Other routes and the resource question

Code switching and gauge fixing can change which gates are easy to implement while preserving the logical information. Gauge color codes provide a construction in which suitable gauge changes enable a universal set [19]. There is no conflict with Eastin–Knill: the full procedure includes changing the code constraints and performing measurements, rather than applying only transversal unitaries on one fixed code. Those transitions must themselves be implemented fault tolerantly.

Threshold theorems establish the possibility of arbitrarily reliable computation below a noise bound under stated assumptions; an early result is due to Aharonov and Ben-Or [20]. They do not assign one hardware-independent threshold to every $[[n, k, d]]$ code. Nor does a memory experiment establish a universal processor. A computation-level resource estimate must include logical preparation and readout, non-Clifford operations or resource-state factories, routing, measurement latency, decoding, and the total algorithm failure budget. Bosonic encodings and subsystem codes can contribute to these constructions, but their encoding alone does not settle the universal gate problem.

From correctable errors to a usable computation

Specify the code and noise model; construct reliable syndrome circuits; identify a complete logical gate set; explain how non-Clifford resources or code changes are protected; and count the resources needed to meet the algorithm's error budget. Each step is a separate claim that needs evidence.

12 Interpreting modern code families and performance claims

This section gives selected examples through 2025 to orient further reading. It separates rigorous parameter statements, model-dependent simulations, and experiments. It should not be read as a comprehensive account of every subsequent development.

12.1 Thresholds and finite-size scaling

Section 11 explained why a complete fault-tolerant construction needs more than a code. A threshold belongs to a family of implementations and a noise model; it is not a universal constant attached solely to the symbols $[[n, k, d]]$. In particular, a memory threshold and the resources for a universal computation answer different questions.

For a specified odd-distance surface-code implementation, a useful empirical parametrization in an appropriate low-noise regime is

$$p_L(d, p) \approx A \left(\frac{p}{p_{\text{th}}} \right)^{(d+1)/2}. \tag{12.1}$$

Here p_L might mean error per correction cycle, p summarizes the assumed physical noise, and A and p_{th} depend on the model and protocol. This is an approximate fit, not a theorem for every code or for arbitrary correlations. A finite-size *pseudothreshold*, such as a crossing where a chosen encoded error rate equals its unencoded comparator, is also distinct from an asymptotic threshold.

12.2 A surface-code memory experiment

Google Quantum AI and Collaborators reported below-threshold surface-code memories in a paper published online in December 2024 and in the 2025 volume of *Nature* [21]. The reported suppression factor when increasing distance by two was $\Lambda = 2.14 \pm 0.02$, with distance-seven error per correction cycle $(1.43 \pm 0.03) \times 10^{-3}$. Its 101-qubit memory comprised 49 data qubits, 48 measurement qubits, and four additional leakage-removal qubits. These are results for the reported device, protocol, and decoder; they do not establish arbitrary-depth universal fault-tolerant computation or a universal surface-code threshold.

This example also shows why resource counting matters. The nominal rotated distance-seven patch uses $49 + 48 = 97$ data and measurement qubits. The experiment's additional resources account for the larger total. A quotation of "101 qubits" should not be used as the code's n in $[[n, k, d]]$.

12.3 Quantum low-density parity-check codes

A family of quantum LDPC codes has a choice of stabilizer generators with both bounded weight per generator and bounded number of generators incident on each qubit, uniformly as the block length grows. Both row and column sparsity matter. The toric and surface-code families are themselves quantum LDPC families; LDPC and topological are not disjoint categories.

A family is *asymptotically good* when $k = \Omega(n)$ and $d = \Omega(n)$, so both the rate k/n and relative distance d/n remain bounded away from zero. Panteleev and Kalachev established such quantum LDPC families [22], and Leverrier and Zémor gave the quantum Tanner

construction [23]. These are existence and construction results about asymptotic families, not immediate finite-device overhead guarantees.

For contrast, Eq. (8.3) gives rate $1/L^2$ and distance proportional to $\sqrt{n}$. Sparsity does not require short-range interactions on a fixed two-dimensional grid; good quantum LDPC constructions may need connectivity that a planar device cannot implement directly. The perspective by Breuckmann and Eberhardt [24] is useful background, though it predates the good-code breakthrough.

12.4 An example of a finite-size overhead comparison

Bravyi and coauthors studied bivariate bicycle codes with explicit syndrome circuits and numerical decoding [25]. Their $[[144, 12, 12]]$ example uses 144 data qubits and 144 measurement ancillas. Under the paper's circuit-level noise model at physical error probability 10^{-3}, they reported logical performance near 2×10^{-7} per cycle for the encoded block, with substantially fewer qubits than the surface-code comparison in that study. These are simulation and modelling results, with the specified nonlocal connectivity, circuits, and decoder. They are not a hardware demonstration or a guarantee for every quantum LDPC code. The logical error metric is for the block of 12 encoded qubits, so it must not silently be equated with a one-logical-qubit metric.

12.5 Single-shot correction is an additional property

In single-shot error correction, one noisy round of suitable check measurements can be sufficient to produce a recovery with controlled residual data noise. The goal is not to infer a perfectly clean syndrome or guarantee zero residual errors. Stability under repeated rounds and the permitted noise class are part of the statement.

Bombín developed this idea for suitable topological constructions [26]. Gu and coauthors proved single-shot decoding guarantees for quantum Tanner codes under stated hypotheses [27]. The property must be established for a code family, measurement scheme, noise assumptions, and decoder. It does not follow from sparse checks alone, nor from the label "higher-dimensional" or "hypergraph-product." An efficient decoder for one family is likewise not a polynomial-time solution to every quantum decoding problem.

12.6 Holographic codes as a different application

Holographic error correction concerns how information encoded on a boundary can admit reconstruction from different boundary regions. Almheiri, Dong, and Harlow formulated a connection with bulk locality in AdS/CFT [28]; Pastawski and coauthors constructed tensor-network toy models using perfect tensors [29]. These models illuminate erasure recovery, entanglement, and operator reconstruction. They are not, merely by being holographic, proposals with established circuit-level thresholds for practical quantum memories. Their purpose and performance questions differ from those of hardware-oriented surface-code and LDPC studies.

Example	Parameter statement	Main lesson or qualification
Bit-flip repetition	$[[3,1,1]]$	Corrects one X fault only under a restricted error model
Four-qubit code	$[[4,2,2]]$	Detection differs from unknown-location correction
Shor code	$[[9,1,3]]$	Degenerate errors need not have distinct syndromes
Steane code	$[[7,1,3]]$	CSS decoding reduces to two classical syndrome problems
Bacon–Shor code	$[[9,1,4,3]]$	Four gauge qubits; distance refers to dressed logical operators
Square toric family	$[[2L^2,2,L]]$	Local checks; noncontractible logical loops
Rotated planar patch	$[[d^2,1,d]]$	Parameter n counts data qubits; ancillas are additional
Good quantum LDPC families	$k=\Omega(n)$, $d=\Omega(n)$	Asymptotic parameters do not specify finite-size hardware cost

Table 5: Qubit code parameters under the conventions of this survey; Bacon–Shor uses $[[n,k,r,d]]$. Bosonic encodings are compared separately in Table 4. No row specifies a hardware-independent threshold or a universal decoding guarantee.

Questions to ask when reading a performance claim

What is the noise model? What operations are assumed perfect? Which decoder and extraction circuit are used? Is the result a theorem, simulation, extrapolation, or experiment? Does n count only data qubits? Is failure measured per cycle, per logical gate, per logical qubit, or per encoded block? Are connectivity and classical decoding costs included?

13 Exercises

The exercises use only material introduced in the text. Solutions appear in Appendix A. A useful first sequence is Exercises 1, 3, 5, 8, and 10; the remaining questions consolidate definitions and assumptions.

1. **Encoding and cloning.** Expand $(\alpha|0\rangle+\beta|1\rangle)^{\otimes 3}$ and compare it with Eq. (3.1). For $\alpha=\beta=1/\sqrt{2}$, which computational-basis outcomes can occur in each state?

2. **Kraus normalization.** Verify the completeness relation for amplitude damping. Compute $\mathcal{N}(I/2)$ and explain why $\gamma>0$ rules out a Pauli-mixture description.

3. **A repetition-code failure.** Compute the syndrome of X_1X_2 and the residual after Table 3 recovery. Derive Eq. (3.3) and show that $p_L<p$ exactly for $0<p<1/2$ within $0<p<1$.

4. **A phase error.** Evaluate $Z_2(\alpha|000\rangle+\beta|111\rangle)$. Explain why the bit-flip repetition code has quantum distance one.

5. **Symplectic commutation.** Encode $g=XYI$ and $h=ZYZ$ as binary vectors. Use Eq. (5.5) to decide whether they commute. Can both stabilize a nonzero vector with eigenvalue $+1$?

6. **Shor degeneracy.** Using the ordered checks in Eq. (6.2), compute the syndromes of Z_1 and Z_2. Show directly that applying Z_1 after Z_2 recovers every code state.

7. **Two CSS conventions.** Starting from $C_2 \subseteq C_1$, derive row $H_X = C_2$, row $H_Z = C_1^{\perp}$, and Eq. (7.4). Why do Z checks detect X errors?

8. **Steane decoding.** Compute both syndrome halves for X_5Z_2 and give a recovery. Why does success here not mean the code corrects every weight-two error?

9. **Minimal logical representatives.** Show that $X_1X_2X_3$ and $Z_1Z_2Z_3$ commute with all Steane checks, are nontrivial, and anticommute with each other.

10. **Counting the toric code.** For $L = 3$, calculate the number of data qubits, independent checks, logical qubits, and distance. Explain why the code does not have $k = 0$ even though there are $2L^2$ displayed checks.

11. **A conditional scaling calculation.** Assume Eq. (12.1) with the same A and p_{th} at distances d and $d + 2$. Derive the ratio of logical error probabilities. List two reasons that an experiment might depart from this ratio.

12. **Known erasures.** Use the Knill–Laflamme condition to explain why a distance-d code corrects any known set of $d - 1$ erased qubits but is guaranteed to correct only $\lfloor (d - 1)/2 \rfloor$ errors at unknown locations.

13. **Gauge counting and distance.** In the Bacon–Shor example, use $g = 12$ and $s = 4$ to count logical and gauge qubits. Show that $X_{1,1}X_{1,2}$ is a gauge operator but not a stabilizer. Why does its weight not contradict distance three?

14. **A bosonic error set.** Verify Eq. (10.5). Compare the two no-loss probabilities in Eq. (10.6), and explain why exact correction of $\{I, a\}$ does not imply exact correction of the full pure-loss channel.

15. **When parity measures the data.** For $\alpha|C_\gamma^+\rangle + \beta|C_\gamma^-\rangle$ with $\gamma \neq 0$, write the parity-measurement probabilities and post-measurement states. Why is this not a harmless loss syndrome for this two-component logical basis?

16. **GKP rounding.** In the ideal square-lattice code, take $u = 0.6\sqrt{\pi}$ and $v = 0$. Which representative of u in $[-\sqrt{\pi}/2, \sqrt{\pi}/2)$ is reported by the syndrome? What is left after nearest-lattice correction?

17. **Ancilla propagation.** An ancilla a is the control of CNOT gates to data qubits 1 and 2. An X_a fault occurs before both gates. Use Eq. (11.1) to find the propagated Pauli. What feature of transversal gates avoids this particular within-block spread?

18. **Gate injection.** Derive both branches in Eq. (11.4), including their probabilities. Explain why the correction is S_{ph} for outcome one and why a noisy resource state needs additional analysis.

19. **Steane phase convention.** Using the weights of the two logical cosets, verify that $S_{\text{ph}}^{\otimes 7}$ implements logical $S_{\text{ph}}^{\dagger}$. Why does this transversal Clifford construction not violate the Eastin–Knill theorem?

14 A guide to further study

For a compact accompanying introduction, read Roffe [30]. For background on states, channels, and quantum circuits, use Chapter 10 and the relevant preliminary chapters of Nielsen and Chuang [2]. Gottesman's thesis [4] develops the stabilizer formalism beyond the examples given here. These are useful learning sources; foundational results should also be attributed to their original papers.

A six-session reading group can organize the text as follows:

1. **Noise and measurement.** Sections 1–3; Exercises 1–4. Calculate a parity measurement on a superposition.
2. **Correction and stabilizers.** Sections 4–6; Exercises 5–6. Explain why a shared syndrome can be either harmless or fatal.
3. **Constructing a code.** Section 7; Exercises 7–9. Build and decode the Steane code from the displayed matrix.
4. **Geometry and gauge freedom.** Sections 8–9; Exercises 10, 12, and 13. Compare a logical loop with a dressed logical operator.
5. **Oscillator encodings.** Section 10; Exercises 14–16. Check a finite correctable error set and interpret ideal grid-code recovery.
6. **Universal computation and evidence.** Sections 11–12; Exercises 11 and 17–19. Derive gate injection, then compare a theorem, a simulation, and an experiment.

For a next project, implement the binary syndrome calculations and verify distances for the small codes. The accompanying `verify_examples.py` enumerates low-weight Paulis and checks the repetition, four-qubit, Shor, Steane, and $L = 3$ toric examples. It also checks the Bacon–Shor dressed distance, the binomial error products and no-loss mismatch, the Steane phase convention, and both injection branches. It requires only Python's standard library. These finite checks support the worked examples; they are not a numerical proof of asymptotic performance or a substitute for the general arguments.

A Solutions to the exercises

1. The tensor cube is $\alpha^3|000\rangle+\alpha^2\beta(|001\rangle+|010\rangle+|100\rangle)+\alpha\beta^2(|011\rangle+|101\rangle+|110\rangle)+\beta^3|111\rangle$. For equal amplitudes it assigns probability $1/8$ to each of the eight strings. The encoded state assigns probability $1/2$ each to 000 and 111. The encoding is a linear map of the input amplitudes, not cloning.

2. $E_0^\dagger E_0 = \mathrm{diag}(1, 1-\gamma)$ and $E_1^\dagger E_1 = \mathrm{diag}(0, \gamma)$, whose sum is I. Also $\mathcal{N}(I/2) = \mathrm{diag}((1+\gamma)/2, (1-\gamma)/2)$. For $\gamma > 0$ this differs from $I/2$, whereas every convex mixture of unitary Pauli conjugations fixes $I/2$.

3. X_1X_2 commutes with Z_1Z_2 and anticommutes with Z_2Z_3, so its syndrome is 01. Recovery gives $X_3X_1X_2 = \bar{X}$. Two flips have total probability $3p^2(1-p)$ and three flips probability p^3. Finally, $p_L - p = -p(1-p)(1-2p)$, which is negative exactly on $0 < p < 1/2$ within the stated interval.

4. The state becomes $\alpha|000\rangle - \beta|111\rangle$. Both checks commute with Z_2, but Z_2 acts nontrivially on the logical space. It is a weight-one logical Z, establishing distance one.

5. Ignoring phases, $g \leftrightarrow (110 \mid 010)$ and $h \leftrightarrow (010 \mid 111)$. Their symplectic product is $110 \cdot 111 + 010 \cdot 010 = 0 + 1 = 1$ in $\mathbb{F}_2$. They anticommute. If a nonzero vector had eigenvalue $+1$ for both, then $gh|\psi\rangle = |\psi\rangle$ and $hg|\psi\rangle = |\psi\rangle$, contradicting $gh = -hg$.

6. Both syndromes are $(0,0,0,0,0,0,1,0)^T$. Because Z_1Z_2 is a listed generator, $Z_1Z_2|\psi_L\rangle = |\psi_L\rangle$ for every encoded state. Identification of the exact position within the first block is unnecessary.

7. Translation by C_2 leaves each coset sum fixed; vectors in $C_1^{\perp}$ contribute trivial phases. Hence the check row spaces are as stated and $k = n - k_2 - (n - k_1) = k_1 - k_2$. The relation $ZX = -XZ$ makes each Z check report the parity of X errors in its support.

8. $s_Z = H_{\text{Ham}} e_5^T = (1,0,1)^T$ and $s_X = H_{\text{Ham}} e_2^T = (0,1,0)^T$. A recovery is Z_2X_5. Some other weight-two errors are ambiguous: X_1X_2 has the same syndrome as X_3, so that recovery leaves the nontrivial logical $X_1X_2X_3$. Thus there is no guarantee for all weight-two errors.

9. The vector 1110000 is in $\ker H_{\text{Ham}}$ because columns 1, 2, and 3 sum to zero. It is not in the row space, whose nonzero vectors all have weight four. This proves commutation and nontriviality for both pure operators. Their supports overlap in three places, so they anticommute.

10. There are 18 data qubits, $18 - 2 = 16$ independent checks, $k = 2$ logical qubits, and $d = 3$. The products of all stars and of all plaquettes are identity, giving two check dependencies.

11. The ratio is $p_L(d{+}2, p)/p_L(d, p) \approx p/p_{\text{th}}$, and the inverse suppression factor is approximately p_{th}/p. Correlated faults or a logical error floor may spoil the fit. Other reasons include finite-size effects, a changing prefactor or decoder, and a physical noise distribution not summarized by a single fixed parameter p.

12. Operators supported on a known erased set of size e have products supported on that same set, so $e < d$ ensures the correction condition. For two errors each of weight at most t at unknown positions, their supports may be disjoint, giving product weight up to $2t$. The uniform guarantee is therefore $2t < d$.

13. Equation (9.3) gives $r = (12-4)/2 = 4$ and $k = 9-4-4 = 1$. The operator $X_{1,1}X_{1,2}$ is one of the gauge generators; it anticommutes with $Z_{1,1}Z_{2,1}$ and therefore is not central in the gauge group. Distance excludes all gauge operators, not just stabilizers. This weight-two operator acts only on the gauge subsystem.

14. Single loss maps the even codewords to orthogonal odd states, so $PaP = 0$ and $Pa^{\dagger}P = 0$. The diagonal matrix elements of $\hat{n}$ are both two and its off-diagonal elements vanish, giving $P\hat{n}P = 2P$. For full loss, $\langle 0_L|E_0^{\dagger}E_0|0_L\rangle = (1+\eta^4)/2$ and $\langle 1_L|E_0^{\dagger}E_0|1_L\rangle = \eta^2$. The difference $(1-\eta^2)^2/2$ is positive for $0 \le \eta < 1$, so this Kraus-product condition fails.

15. The even and odd outcomes have probabilities $|\alpha|^2$ and $|\beta|^2$. Their normalized states are $|C_{\gamma}^{+}\rangle$ and $|C_{\gamma}^{-}\rangle$ up to global phase. The measurement reads the encoded basis value and removes coherence between the two components. It is a logical measurement, unlike parity measurement in the binomial example.

16. The representative is $-0.4\sqrt{\pi}$. Recovery shifts by $+0.4\sqrt{\pi}$, leaving a total shift $\sqrt{\pi}$. This is the logical $\bar{X}$, so the ideal decoder has made a logical error even though the final stabilizer syndrome is trivial.

17. After the first gate the fault is X_aX_1, and after the second it is $X_aX_1X_2$. Two data errors have resulted from one ancilla fault. A transversal pairing never lets that one participating physical subsystem couple to two data qubits in the same block during that gate.

18. After CNOT, the joint state is $[\alpha|0\rangle(|0\rangle + e^{i\pi/4}|1\rangle) + \beta|1\rangle(|1\rangle + e^{i\pi/4}|0\rangle)]/\sqrt{2}$. Selecting each target outcome gives Eq. (11.4); each branch has squared norm $1/2$. For outcome one, $S_{\mathrm{ph}}T^\dagger = T$ gives the desired gate up to global phase. A resource-state error can become a data error, so this ideal identity does not bound the noisy gate's failure rate.

19. Physical $S_{\mathrm{ph}}^{\otimes 7}$ multiplies a computational string of weight w by i^w. This is one for $w = 0, 4$ and $-i$ for $w = 3, 7$. Its logical matrix is therefore $\mathrm{diag}(1, -i) = S_{\mathrm{ph}}^\dagger$. The resulting transversal Clifford gates are not universal; an additional protected non-Clifford resource or construction is needed.

B A compact mathematical checklist

When presenting a new stabilizer example, give the check generators and their signs, verify commutation and independence, calculate k, and exhibit logical operators with the claimed commutation relations. To claim distance d, show both a nontrivial logical of weight d and the absence of lighter ones. A displayed encoding map alone does not establish distance.

For a CSS example, specify the $(x \mid z)$ ordering, the meanings of H_X, H_Z, and $H_X H_Z^T = 0$. Use ranks rather than the raw number of displayed rows. State whether a distance formula is exact or only a bound. If a decoding argument assumes at most one bit error and at most one phase error, distinguish that set from all errors of a given total Pauli weight.

For a subsystem code, distinguish the gauge group from its commuting stabilizer, count both logical and gauge qubits, and use dressed distance. For a bosonic code, specify the mode, quadrature convention, physical energy assumptions, and exact or approximate error model. A nonnormalizable ideal state or a truncated error set requires explicit interpretation. For a logical-gate claim, verify the action on the encoded basis and account for circuit faults, resource states, and feed-forward.

For a probability claim, specify the channel and its parameter convention, independence or correlation assumptions, the decoder, treatment of noisy measurements, and the unit of logical failure. For a resource claim, distinguish data qubits, ancillas, and other implementation resources. These details make calculations reproducible and comparisons meaningful.

References

[1] Peter W. Shor. Scheme for reducing decoherence in quantum computer memory. *Physical Review A*, 52:R2493–R2496, 1995. doi: 10.1103/PhysRevA.52.R2493. URL https://doi.org/10.1103/PhysRevA.52.R2493.

[2] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 10th anniversary edition, 2010. doi: 10.1017/CBO9780511976667. URL https://doi.org/10.1017/CBO9780511976667.

[3] Emanuel Knill and Raymond Laflamme. Theory of quantum error-correcting codes. *Physical Review A*, 55:900–911, 1997. doi: 10.1103/PhysRevA.55.900. URL https://doi.org/10.1103/PhysRevA.55.900.

[4] Daniel Gottesman. *Stabilizer Codes and Quantum Error Correction*. PhD thesis, California Institute of Technology, 1997. URL https://arxiv.org/abs/quant-ph/9705052. arXiv:quant-ph/9705052.

[5] A. R. Calderbank and Peter W. Shor. Good quantum error-correcting codes exist. *Physical Review A*, 54:1098–1105, 1996. doi: 10.1103/PhysRevA.54.1098. URL https://arxiv.org/abs/quant-ph/9512032.

[6] Andrew M. Steane. Multiple-particle interference and quantum error correction. *Proceedings of the Royal Society of London A*, 452:2551–2577, 1996. doi: 10.1098/rspa.1996.0136. URL https://arxiv.org/abs/quant-ph/9601029.

[7] A. Yu. Kitaev. Fault-tolerant quantum computation by anyons. *Annals of Physics*, 303:2–30, 2003. doi: 10.1016/S0003-4916(02)00018-0. URL https://arxiv.org/abs/quant-ph/9707021. Preprint first posted in 1997.

[8] Austin G. Fowler, Matteo Mariantoni, John M. Martinis, and Andrew N. Cleland. Surface codes: Towards practical large-scale quantum computation. *Physical Review A*, 86:032324, 2012. doi: 10.1103/PhysRevA.86.032324. URL https://arxiv.org/abs/1208.0928.

[9] David Kribs, Raymond Laflamme, and David Poulin. Unified and generalized approach to quantum error correction. *Physical Review Letters*, 94:180501, 2005. doi: 10.1103/PhysRevLett.94.180501. URL https://arxiv.org/abs/quant-ph/0412076.

[10] David Poulin. Stabilizer formalism for operator quantum error correction. *Physical Review Letters*, 95:230504, 2005. doi: 10.1103/PhysRevLett.95.230504. URL https://arxiv.org/abs/quant-ph/0508131.

[11] Dave Bacon. Operator quantum error-correcting subsystems for self-correcting quantum memories. *Physical Review A*, 73:012340, 2006. doi: 10.1103/PhysRevA.73.012340. URL https://arxiv.org/abs/quant-ph/0506023.

[12] Marios H. Michael, Matti Silveri, R. T. Brierley, Victor V. Albert, Juha Salmilehto, Liang Jiang, and S. M. Girvin. New class of quantum error-correcting codes for a bosonic mode. *Physical Review X*, 6:031006, 2016. doi: 10.1103/PhysRevX.6.031006. URL https://arxiv.org/abs/1602.00008.

[13] Mazyar Mirrahimi, Zaki Leghtas, Victor V. Albert, Steven Touzard, Robert J. Schoelkopf, Liang Jiang, and Michel H. Devoret. Dynamically protected cat-qubits: a new paradigm for universal quantum computation. *New Journal of Physics*, 16:045014, 2014. doi: 10.1088/1367-2630/16/4/045014. URL https://arxiv.org/abs/1312.2017.

[14] Daniel Gottesman, Alexei Kitaev, and John Preskill. Encoding a qubit in an oscillator. *Physical Review A*, 64:012310, 2001. doi: 10.1103/PhysRevA.64.012310. URL https://arxiv.org/abs/quant-ph/0008040.

[15] Daniel Gottesman. Theory of fault-tolerant quantum computation. *Physical Review A*, 57:127–137, 1998. doi: 10.1103/PhysRevA.57.127. URL https://arxiv.org/abs/quant-ph/9702029.

[16] Bryan Eastin and Emanuel Knill. Restrictions on transversal encoded quantum gate sets. *Physical Review Letters*, 102:110502, 2009. doi: 10.1103/PhysRevLett.102.110502. URL https://arxiv.org/abs/0811.4262.

[17] Sergey Bravyi and Alexei Kitaev. Universal quantum computation with ideal Clifford gates and noisy ancillas. *Physical Review A*, 71:022316, 2005. doi: 10.1103/PhysRevA.71.022316. URL https://arxiv.org/abs/quant-ph/0403025.

[18] Sergey Bravyi and Jeongwan Haah. Magic-state distillation with low overhead. *Physical Review A*, 86:052329, 2012. doi: 10.1103/PhysRevA.86.052329. URL https://arxiv.org/abs/1209.2426.

[19] Héctor Bombín. Gauge color codes: optimal transversal gates and gauge fixing in topological stabilizer codes. *New Journal of Physics*, 17:083002, 2015. URL https://arxiv.org/abs/1311.0879.

[20] Dorit Aharonov and Michael Ben-Or. Fault tolerant quantum computation with constant error, 1996. URL https://arxiv.org/abs/quant-ph/9611025. arXiv:quant-ph/9611025.

[21] Google Quantum AI and Collaborators. Quantum error correction below the surface code threshold. *Nature*, 638:920–926, 2025. doi: 10.1038/s41586-024-08449-y. URL https://doi.org/10.1038/s41586-024-08449-y. Published online 9 December 2024; a 2026 author correction concerns labels in Fig. 3a, doi:10.1038/s41586-026-10559-8.

[22] Pavel Panteleev and Gleb Kalachev. Asymptotically good quantum and locally testable classical LDPC codes. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 375–388. ACM, 2022. doi: 10.1145/3519935.3520017. URL https://arxiv.org/abs/2111.03654. Preprint arXiv:2111.03654 first posted in 2021.

[23] Anthony Leverrier and Gilles Zémor. Quantum Tanner codes. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 872–883. IEEE, 2022. doi: 10.1109/FOCS54457.2022.00117. URL https://arxiv.org/abs/2202.13641. Full version available as arXiv:2202.13641.

[24] Nikolas P. Breuckmann and Jens Niklas Eberhardt. Quantum low-density parity-check codes. *PRX Quantum*, 2:040101, 2021. doi: 10.1103/PRXQuantum.2.040101. URL https://arxiv.org/abs/2103.06309.

[25] Sergey Bravyi, Andrew W. Cross, Jay M. Gambetta, Dmitri Maslov, Patrick Rall, and Theodore J. Yoder. High-threshold and low-overhead fault-tolerant quantum memory. *Nature*, 627:778–782, 2024. doi: 10.1038/s41586-024-07107-7. URL https://doi.org/10.1038/s41586-024-07107-7.

[26] Héctor Bombín. Single-shot fault-tolerant quantum error correction. *Physical Review X*, 5:031043, 2015. doi: 10.1103/PhysRevX.5.031043. URL https://arxiv.org/abs/1404.5504.

[27] Shouzhen Gu, Eugene Tang, Libor Caha, Shin Ho Choe, Zhiyang He, and Aleksander Kubica. Single-shot decoding of good quantum LDPC codes. *Communications in Mathematical Physics*, 405:85, 2024. doi: 10.1007/s00220-024-04951-6. URL https://arxiv.org/abs/2306.12470.

[28] Ahmed Almheiri, Xi Dong, and Daniel Harlow. Bulk locality and quantum error correction in AdS/CFT. *Journal of High Energy Physics*, 2015(04):163, 2015. doi: 10.1007/JHEP04(2015)163. URL https://arxiv.org/abs/1411.7041.

[29] Fernando Pastawski, Beni Yoshida, Daniel Harlow, and John Preskill. Holographic quantum error-correcting codes: Toy models for the bulk/boundary correspondence. *Journal of High Energy Physics*, 2015(06):149, 2015. doi: 10.1007/JHEP06(2015)149. URL https://arxiv.org/abs/1503.06237.

[30] Joschka Roffe. Quantum error correction: An introductory guide. *Contemporary Physics*, 60(3):226–245, 2019. doi: 10.1080/00107514.2019.1667078. URL https://arxiv.org/abs/1907.11157.