Module 5 — Decoherence

Quantum Foundations

hbar.university

Density Matrices and Mixed States

Position in Knowledge Graph

System: Quantum Foundations **Structural Domain:** Decoherence **Node:** 5.1

The state postulate (Node 2.1) identified quantum states with rays in Hilbert space — pure states. This node generalizes the state concept to **mixed states**: statistical ensembles of pure states. The mathematical object is the **density matrix**, and it is the right state object whenever the system is open, noisy, partially traced out, or simply uncertain.

Density matrices are the most important object in this course after the Hilbert space itself. Modules 5–9 all live on density matrices.

Formal Definition

A **density operator** (or **density matrix**) on a Hilbert space $\mathcal{H}$ is an operator $\rho : \mathcal{H} \to \mathcal{H}$ satisfying:

1. **Hermiticity:** $\rho = \rho^\dagger$.
2. **Positivity:** $\langle\psi|\rho|\psi\rangle \geq 0$ for all $|\psi\rangle \in \mathcal{H}$.
3. **Unit trace:** $\mathrm{tr}(\rho) = 1$.

A density operator is **pure** if $\mathrm{tr}(\rho^2) = 1$, and **mixed** otherwise.

For a pure state $|\psi\rangle$, the corresponding density operator is the projector

$$\rho = |\psi\rangle\langle\psi|.$$

For a statistical ensemble $\{(p_k, |\psi_k\rangle)\}$ with $p_k \geq 0$ and $\sum_k p_k = 1$, the density operator is the convex combination

$$\rho = \sum_k p_k \, |\psi_k\rangle\langle\psi_k|.$$

The set of density operators is closed and convex; the extreme points are the pure states.

Born Rule for Density Matrices

For an observable A and density matrix ρ, the expectation value is

$$\langle A\rangle_\rho = \mathrm{tr}(\rho A).$$

For a projective measurement $\{P_i\}$, the probability of outcome i is

$$p(i) = \mathrm{tr}(\rho P_i),$$

and the post-measurement state, conditional on outcome i, is

$$\rho' = \frac{P_i \, \rho \, P_i}{\mathrm{tr}(\rho P_i)}.$$

Both of these reduce to the pure-state Born rule (Node 2.3) when $\rho = |\psi\rangle\langle\psi|$.

Intuition

A density matrix is the most general quantum state. It encodes:

- **Quantum uncertainty** (intrinsic to a pure state — non-zero variance of incompatible observables).
- **Classical uncertainty** (about which pure state the system is actually in).

The two are conceptually distinct but the formalism does not separate them. This is the structural reason that density matrices have multiple ensemble decompositions: many different convex combinations $\{(p_k, |\psi_k\rangle)\}$ can give the same ρ.

The textbook example: the maximally mixed qubit state

$$\rho = \frac{I}{2} = \frac{1}{2}|0\rangle\langle 0| + \frac{1}{2}|1\rangle\langle 1| = \frac{1}{2}|+\rangle\langle +| + \frac{1}{2}|-\rangle\langle -|$$

— two completely different "ensembles," same density matrix, identical predictions for every measurement.

Structural Role

Density matrices are required wherever the pure-state formalism breaks down:

- **Subsystems of entangled states** (Node 5.2): tracing out part of an entangled pure state gives a mixed reduced state.
- **Open systems** (Module 8): system-environment coupling drives pure states to mixed states.
- **Noisy quantum hardware** (relevant to VQA): real circuits do not produce pure outputs.
- **Measurement statistics**: pre-readout, the system is described by a density matrix conditional on the previous measurement record.
- **Quantum information**: channels (Module 7) act on density matrices, not on rays.

The density matrix is also the right object for **quantum statistical mechanics**: the Gibbs state $\rho = e^{-\beta H}/Z$ is a density matrix (Node 9.1), and quantum thermodynamics builds entirely on this.

Mathematical Development

Bloch ball representation (qubit)

Every 2×2 density matrix can be written

$$\rho = \frac{1}{2}(I + \vec{r} \cdot \vec{\sigma})$$

where $\vec{r} \in \mathbb{R}^3$ with $\|\vec{r}\| \leq 1$, and $\vec{\sigma} = (X, Y, Z)$ is the vector of Pauli matrices.

- $\|\vec{r}\| = 1$: pure state, on the Bloch sphere.
- $\|\vec{r}\| < 1$: mixed state, in the interior of the Bloch ball.
- $\vec{r} = 0$: maximally mixed state $I/2$, at the center.

The pure-state postulate (Node 2.1) maps states to the *surface* of the ball; the density-matrix generalization fills in the interior.

Purity

The **purity** of ρ is

$$\gamma(\rho) = \mathrm{tr}(\rho^2) \in [1/d, 1],$$

where $d = \dim \mathcal{H}$. It equals 1 for pure states and $1/d$ for the maximally mixed state.

Worked Example

Mixture of $|0\rangle$ and $|1\rangle$ versus $|+\rangle$

Consider two preparations:

1. With probability 1/2, the qubit is in $|0\rangle$; with probability 1/2, it is in $|1\rangle$. Density matrix: $\rho_1 = I/2$.
2. The qubit is definitely in $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$. Density matrix: $\rho_2 = |+\rangle\langle+|$.

These are *very* different states. Measuring X on ρ_1 gives outcome ± 1 with equal probability; measuring X on ρ_2 gives $+1$ with certainty.

The lesson: a mixed state and a pure superposition look similar in computational-basis measurements but diverge sharply when measured in other bases. The density matrix captures this correctly.

Cross-Links

- **Linear Algebra:** Hermitian, positive semidefinite operators with unit trace.
- **Statistics:** Density matrices generalize probability distributions; the diagonal of ρ in any basis is a classical distribution over that basis.
- **Quantum Information:** All channels and entropies are defined on density matrices.
- **VQA:** Cost functions $\langle H \rangle = \mathrm{tr}(\rho H)$ on noisy circuits use this trace formula.
- **Statistics → Bayesian inference:** The density matrix is to a quantum state what the posterior is to a Bayesian belief — a complete summary of "what we know."

Structural Compression

Invariant: A density operator is positive, Hermitian, and unit-trace; the set of density operators is convex with pure states as extreme points.

Mechanism: Convex hull of $\{|\psi\rangle\langle\psi| : |\psi\rangle \in \mathcal{H}, \||\psi\rangle\| = 1\}$.

Cross-System Echo: Probability simplices in classical statistics are also convex with pure (delta) measures as extreme points. The density operator is the non-commutative generalization. Both are the most general state objects in their respective theories.

Exercises

1. Show that for any density matrix ρ, $\mathrm{tr}(\rho^2) \leq 1$, with equality iff ρ is pure.
2. Verify that $\rho = \frac{1}{2}|+\rangle\langle+| + \frac{1}{2}|-\rangle\langle-| = I/2$.
3. Show that the Bloch ball parameterization $\rho = \frac{1}{2}(I + \vec{r}\cdot\vec{\sigma})$ is positive iff $\|\vec{r}\| \leq 1$.
4. Compute the purity of $\rho = p|0\rangle\langle 0| + (1-p)|+\rangle\langle+|$ as a function of p. For what p is it most mixed?
5. Construct two distinct ensembles $\{(p_k, |\psi_k\rangle)\}$ and $\{(q_l, |\phi_l\rangle)\}$ that produce the same density matrix ρ, and verify they make identical predictions for every observable.

Reduced Density Operator and Partial Trace

Position in Knowledge Graph

System: Quantum Foundations **Structural Domain:** Decoherence **Node:** 5.2

The reduced density operator is the answer to the question: *given a state ρ_{AB} on a composite system, what is the state of subsystem A alone?* The answer is non-trivial precisely because entanglement exists — tracing out B from an entangled pure state does not leave a pure state on A. The partial trace is the mathematical operation that implements this restriction, and it is the bridge between the closed-system unitarity of Module 2 and the open-system dynamics of Modules 5–8.

Formal Definition

Let ρ_{AB} be a density operator on $\mathcal{H}_A \otimes \mathcal{H}_B$. The **reduced density operator** on A, also called the marginal state on A, is defined by the **partial trace** over B:

$$\rho_A := \mathrm{Tr}_B(\rho_{AB}) := \sum_j (I_A \otimes \langle j_B|)\, \rho_{AB}\, (I_A \otimes |j_B\rangle),$$

where $\{|j_B\rangle\}$ is any orthonormal basis of $\mathcal{H}_B$. The result is independent of the chosen basis and is a valid density operator on $\mathcal{H}_A$ (positive, Hermitian, unit trace).

Equivalently, on product operators the partial trace acts as

$$\mathrm{Tr}_B(A_S \otimes B_S) = A_S \,\mathrm{Tr}(B_S),$$

and extends by linearity to arbitrary operators on $\mathcal{H}_A \otimes \mathcal{H}_B$.

The defining operational property

The reduced density operator is the *unique* operator on $\mathcal{H}_A$ such that, for every observable M_A acting on A alone,

$$\mathrm{Tr}(\rho_A M_A) = \mathrm{Tr}\big(\rho_{AB}(M_A \otimes I_B)\big).$$

Expressed differently: ρ_A reproduces the expectation values of all A-only observables computed on the full joint state. This is the sense in which ρ_A is "the state of A": it is everything an observer confined to A needs to compute predictions.

Intuition

If Alice holds subsystem A and Bob holds subsystem B, and Alice has no access to B at all, then every measurement Alice can perform is of the form $M_A \otimes I_B$. The density operator ρ_A is the state that reproduces the statistics of all such measurements — it is Alice's complete local description of reality.

The operational meaning of the partial trace is *marginalization*: discarding subsystem B and asking what remains. In classical probability, the marginal of a joint distribution is obtained by summing over the discarded variable. In quantum mechanics, the analogous operation is the partial trace, which sums over a basis of the discarded subsystem.

The crucial structural fact: **tracing out a subsystem of a pure entangled state generally gives a mixed state.** This is the structural origin of decoherence. Unitary evolution preserves purity at the level of

the full system, but "purity" can be lost in the transition to the subsystem description whenever the full state is entangled. A pure quantum state globally can look maximally mixed locally — that's entanglement.

The reverse operation, **purification**, says that every mixed state can be viewed as the marginal of a pure state on a larger Hilbert space. Mixed states are precisely the image of pure bipartite states under the partial trace.

Structural Role

Reduced density operators are the load-bearing object of open-system quantum mechanics:

- **Local observers use reduced states.** Any observer confined to a subsystem describes the world with a reduced density operator. This is the structural origin of why the world can *look* classical (mixed) even though the global wavefunction is pure.
- **Entanglement detection.** If $|\Psi\rangle_{AB}$ is pure, then $|\Psi\rangle$ is entangled iff ρ_A is mixed. Purity of the reduced state is the single cleanest test of pure-state entanglement.
- **Decoherence (Nodes 5.3–5.5).** Decoherence is formalized as the reduced dynamics of a system coupled to an environment: $\rho_S(t) = \mathrm{Tr}_E\big(U(t)(\rho_S(0) \otimes \rho_E)U^\dagger(t)\big)$. The map $\rho_S(0) \to \rho_S(t)$ is the open-system evolution.
- **Channels (Node 7.3).** Every quantum channel is a partial trace of a unitary on a larger space (Stinespring dilation). The partial trace is the structural universal building block of "generalized quantum dynamics."
- **Purification.** Every mixed state ρ_A can be written as $\mathrm{Tr}_B\,|\Psi\rangle\langle\Psi|$ for some pure $|\Psi\rangle \in \mathcal{H}_A \otimes \mathcal{H}_B$. This embedding is one of the main tools of quantum information theory: to study a mixed state, lift it to a pure state on a larger space and work there.
- **Entanglement entropy (Node 4.5).** $S(\rho_A)$ of the reduced state is the canonical pure-state entanglement measure.

Mathematical Development

Linearity and basis-independence

The partial trace is a linear map $\mathrm{Tr}_B : \mathcal{B}(\mathcal{H}_A \otimes \mathcal{H}_B) \to \mathcal{B}(\mathcal{H}_A)$ that preserves Hermiticity, positivity, and trace. Basis-independence follows from the cyclic property: for two bases $\{|j_B\rangle\}, \{|j'_B\rangle\}$ related by a unitary U_B, the sums

$$\sum_j (I \otimes \langle j_B|)\rho(I \otimes |j_B\rangle) = \sum_{j,k,k'} U^*_{jk}U_{jk'}(I \otimes \langle k'_B|)\rho(I \otimes |k_B\rangle) = \sum_k (I \otimes \langle k_B|)\rho(I \otimes |k_B\rangle)$$

agree, using $\sum_j U^*_{jk}U_{jk'} = \delta_{kk'}$. The result depends only on the subsystem, not the choice of basis.

Schmidt decomposition and the spectrum of ρ_A

For a *pure* bipartite state $|\Psi\rangle_{AB} = \sum_i \lambda_i |i_A\rangle \otimes |i_B\rangle$ in Schmidt form (Node 4.3),

$$\rho_A = \mathrm{Tr}_B\,|\Psi\rangle\langle\Psi| = \sum_i \lambda_i^2\,|i_A\rangle\langle i_A|, \qquad \rho_B = \sum_i \lambda_i^2\,|i_B\rangle\langle i_B|.$$

Both reduced states are automatically diagonal in the Schmidt basis of their respective sides, and they share the same nonzero eigenvalues (the squared Schmidt coefficients). This is why entanglement entropy $S(\rho_A) = S(\rho_B)$ for pure bipartite states — the two marginals carry the same spectral information.

Purification

Theorem. Every density operator ρ_A on $\mathcal{H}_A$ is the reduced state of a pure state $|\Psi\rangle \in \mathcal{H}_A \otimes \mathcal{H}_B$ for some auxiliary space $\mathcal{H}_B$ with $\dim \mathcal{H}_B \geq \operatorname{rank} \rho_A$.

Construction. Diagonalize $\rho_A = \sum_i p_i |i\rangle\langle i|$. Take $\mathcal{H}_B$ with orthonormal basis $\{|i\rangle_B\}$ of the same dimension as the support of ρ_A, and set

$$|\Psi\rangle := \sum_i \sqrt{p_i}\, |i\rangle_A \otimes |i\rangle_B.$$

Then $\operatorname{Tr}_B |\Psi\rangle\langle\Psi| = \sum_i p_i |i\rangle_A\langle i|_A = \rho_A$. ■

Purifications are unique up to isometry on the auxiliary space $\mathcal{H}_B$, and the minimal purification has $\dim \mathcal{H}_B = \operatorname{rank} \rho_A$.

Why the diagonal-in-B prescription works

The formula $\operatorname{Tr}_B(\rho) = \sum_j \langle j_B|\rho|j_B\rangle_B$ (with implicit identity on A) may look ad hoc. The structural reason: it is the *unique* linear map satisfying $\operatorname{Tr}_B(A \otimes B) = A \operatorname{Tr}(B)$, and this in turn is forced by the operational requirement that expectation values of A-only observables be preserved. In category-theoretic language, the partial trace is the unique natural transformation making the corresponding diagram of C^*-algebras commute.

Reduced dynamics

If the joint state evolves unitarily under $U_{AB}(t)$, the reduced state on A evolves as

$$\rho_A(t) = \operatorname{Tr}_B\big(U_{AB}(t)\,(\rho_A(0) \otimes \rho_B(0))\,U_{AB}(t)^\dagger\big),$$

where we assumed an initially factorized state. This map $\rho_A(0) \mapsto \rho_A(t)$ is a **quantum channel** — completely positive and trace-preserving — but it is generally *not unitary*. The non-unitarity is the structural content of decoherence: the system's state evolution, in isolation, has lost information.

Worked Example

One qubit of a Bell pair

Take $|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$. The density matrix of the joint state is

$$|\Phi^+\rangle\langle\Phi^+| = \tfrac{1}{2}(|00\rangle\langle 00| + |00\rangle\langle 11| + |11\rangle\langle 00| + |11\rangle\langle 11|).$$

Trace out qubit B:

$$\rho_A = \tfrac{1}{2}(|0\rangle\langle 0| \operatorname{Tr}|0\rangle\langle 0| + |0\rangle\langle 1| \operatorname{Tr}|0\rangle\langle 1| + |1\rangle\langle 0| \operatorname{Tr}|1\rangle\langle 0| + |1\rangle\langle 1| \operatorname{Tr}|1\rangle\langle 1|),$$

$$= \tfrac{1}{2}(|0\rangle\langle 0| \cdot 1 + |0\rangle\langle 1| \cdot 0 + |1\rangle\langle 0| \cdot 0 + |1\rangle\langle 1| \cdot 1) = \tfrac{1}{2}(|0\rangle\langle 0| + |1\rangle\langle 1|) = \tfrac{I}{2}.$$

The reduced state of one qubit of a Bell pair is **maximally mixed** — a completely uninformative state, indistinguishable from a fair coin flip between $|0\rangle$ and $|1\rangle$. All of the original pure state's information is stored in the *correlation* between the two qubits, not in either qubit alone.

This is the cleanest structural illustration of decoherence: a pure global state can look maximally mixed locally.

Reduced state of a partially entangled pure state

Take $|\Psi\rangle = \cos\theta|00\rangle + \sin\theta|11\rangle$. By the same calculation,

$$\rho_A = \cos^2\theta\,|0\rangle\langle 0| + \sin^2\theta\,|1\rangle\langle 1|.$$

The purity is $\operatorname{Tr}\rho_A^2 = \cos^4\theta + \sin^4\theta$, which ranges from 1 (at $\theta = 0$ or $\pi/2$, product state) to 1/2 (at $\theta = \pi/4$, Bell state). The entanglement entropy $h(\cos^2\theta)$ (binary entropy) behaves similarly.

Reduced state of a product state

For a product state $\rho_{AB} = \rho_A \otimes \rho_B$, the partial trace trivially gives back ρ_A:

$$\operatorname{Tr}_B(\rho_A \otimes \rho_B) = \rho_A \operatorname{Tr}(\rho_B) = \rho_A \cdot 1 = \rho_A.$$

No information is lost in the reduction, because ρ_B carries nothing that depends on A. Purity is preserved: if ρ_{AB} is a product of pure states, so is the reduction.

Classically correlated state

Consider $\rho_{AB} = \frac{1}{2}(|00\rangle\langle 00| + |11\rangle\langle 11|)$, a classically correlated mixed state. Tracing out B:

$$\rho_A = \tfrac{1}{2}(|0\rangle\langle 0| + |1\rangle\langle 1|) = \tfrac{I}{2}.$$

The reduced state is again $I/2$, *exactly the same* as the reduced state of $|\Phi^+\rangle$. From the marginal perspective, the Bell state and the classically correlated state are indistinguishable — the distinction lies entirely in the *correlations*, which are washed out by the partial trace. This is why entanglement cannot be detected from reduced states alone: you need access to cross-correlators like $\langle X_A X_B\rangle$.

Cross-Links

- **Module 4.3 (Schmidt decomposition):** Reduced states of bipartite pure states are diagonal in the Schmidt basis.
- **Module 4.5 (Entanglement measures):** The von Neumann entropy of ρ_A is the canonical pure-state entanglement measure.
- **Module 5.3 (System–environment coupling):** Reduced dynamics are defined by tracing out an unobserved environment.
- **Module 7.3 (Quantum channels):** Every CPTP channel arises as a partial trace of a unitary on a dilated space (Stinespring).
- **Statistics:** Marginalization of joint probability distributions. The partial trace generalizes $p(x) = \sum_y p(x,y)$ to non-commuting marginals.

Structural Compression

Invariant: ρ_A is the unique density operator on $\mathcal{H}_A$ such that $\operatorname{Tr}(\rho_A M_A) = \operatorname{Tr}(\rho_{AB}(M_A \otimes I_B))$ for every A-only observable M_A.

Mechanism: Sum over an orthonormal basis of $\mathcal{H}_B$: $\rho_A = \sum_j \langle j_B|\rho_{AB}|j_B\rangle$.

Cross-System Echo: Marginalization in classical probability integrates out the unobserved variable to give the marginal density. The partial trace is the non-commutative generalization — the operational restriction of a joint state to a subsystem that has no access to the rest.

Exercises

1. Verify directly that the partial trace of a density matrix is a density matrix: positive, Hermitian, unit-trace.
2. Show that Tr_B is independent of the basis chosen for $\mathcal{H}_B$.
3. Compute the reduced state of qubit A for each of the four Bell states. Verify that all four give $\rho_A = I/2$.
4. For $|\Psi\rangle = \cos\theta|00\rangle + \sin\theta|11\rangle$, compute the purity of ρ_A as a function of θ and plot it.
5. Given the mixed state $\rho_A = 0.7|0\rangle\langle 0| + 0.3|1\rangle\langle 1|$, construct an explicit purification in $\mathbb{C}^2 \otimes \mathbb{C}^2$ and verify that the reduced state on A matches.
6. Show that for a pure bipartite state, ρ_A is pure iff $|\Psi\rangle_{AB}$ is a product state.
7. Argue, structurally, why reduced density operators are the right objects for describing open quantum systems — what goes wrong if you try to use pure states instead?

System–Environment Coupling

Position in Knowledge Graph

System: Quantum Foundations **Structural Domain:** Decoherence **Node:** 5.3

Every real quantum system is coupled to an environment. Photons, phonons, stray electromagnetic fields, nuclear spin baths, vibrational modes — no laboratory degree of freedom is truly isolated. **Decoherence** is the name for what happens when this coupling is taken seriously: the system's reduced dynamics become non-unitary, coherences between pointer states decay, and the quantum-mechanical description at the level of the system alone acquires classical features. This node sets up the general framework — total Hamiltonian $H = H_S + H_E + H_{SE}$, initial product state, unitary joint evolution, partial trace over the environment — that will be used throughout Modules 5 and 8.

Formal Definition

Consider a **system** S with Hilbert space $\mathcal{H}_S$ and an **environment** E with Hilbert space $\mathcal{H}_E$. The total Hilbert space is $\mathcal{H} = \mathcal{H}_S \otimes \mathcal{H}_E$, and the total Hamiltonian decomposes as

$$H = H_S \otimes I_E + I_S \otimes H_E + H_{SE},$$

where H_S acts on the system alone, H_E on the environment alone, and H_{SE} is the interaction (generally non-factorizable).

Initial condition. At $t = 0$ we assume the joint state is factorized:

$$\rho_{\text{tot}}(0) = \rho_S(0) \otimes \rho_E(0).$$

This is an idealization (in realistic situations the system and environment have been coupled for some time prior to $t = 0$ and are not factorizable), but it is the starting point for the Born–Markov derivations of master equations (Node 8.2).

Joint unitary evolution. The total state evolves unitarily:

$$\rho_{\text{tot}}(t) = U(t)\rho_{\text{tot}}(0)U(t)^\dagger, \qquad U(t) = e^{-iHt/\hbar}.$$

Reduced system dynamics. The system state at time t is obtained by tracing out the environment:

$$\rho_S(t) := \text{Tr}_E\big(U(t)[\rho_S(0) \otimes \rho_E(0)]U(t)^\dagger\big).$$

The map $\Phi_t : \rho_S(0) \mapsto \rho_S(t)$ is a **quantum channel** (completely positive, trace-preserving), but is generally **not unitary**. The non-unitarity is the mathematical signature of decoherence.

Intuition

The closed-system postulate (Node 2.4) said that quantum states evolve unitarily under the Schrödinger equation. This remains true for the *total* system + environment — the joint state is always pure, always evolves unitarily. What breaks is the assumption that the system alone has unitary dynamics. Once you couple the system to degrees of freedom you cannot observe, the reduced state stops behaving like a closed-system state.

The structural picture: imagine the system and environment are two dancers. Before the music starts they are independent; the state is a product $\rho_S \otimes \rho_E$. Once the interaction turns on they become entangled — the total state is no longer a product, even though it is still pure. If you watch only one dancer (the system), their motion becomes *correlated with* the other dancer's motion, and from the perspective of the isolated observation their behavior becomes stochastic. The "randomness" is not fundamental; it is a reflection of the unobserved correlations with the environment.

Operationally, decoherence is the process by which information *leaks* from system to environment. Phases of the system wavefunction that depend on which state the system is in get copied into environmental degrees of freedom, and once they are copied, coherent interference between those states becomes impossible from the system's perspective — the environment "knows which branch" and averaging over it destroys the off-diagonal terms.

The structural slogan: *The environment is an unasked-for measurement apparatus.* Its coupling to the system amounts to a continuous weak measurement in whatever basis the coupling selects, and the reduced system dynamics are exactly the measurement-induced dephasing of Node 3.6, averaged over an unread measurement record.

Structural Role

System–environment coupling is the universal frame for open quantum dynamics:

- **Decoherence (Nodes 5.4–5.6).** The specific reduced dynamics obtained by tracing out different kinds of environments — photon baths, phonon baths, spin baths — give the canonical dephasing and relaxation channels.
- **Pointer basis (Node 5.5).** The basis in which H_{SE} is diagonal is the one the environment "measures." It becomes the pointer basis of the decoherence process.
- **Open-system dynamics (Module 8).** The Lindblad master equation is the Markovian approximation of the exact reduced dynamics derived here. The non-Markovian corrections are the structural content of Node 8.5.
- **Decoherence-free subspaces.** A subspace of $\mathcal{H}_S$ that is invariant under H_{SE} does not decohere — it is protected from the environment by symmetry. This is the structural design principle behind several error-correcting codes and quantum memory architectures.
- **Stinespring / Naimark (Nodes 3.3, 7.3).** The "unitary on a larger space + partial trace" structure is the universal form of generalized quantum dynamics. System–environment coupling is the physical instantiation.

Mathematical Development

The reduced dynamics is a channel

Theorem. Given a factorized initial state $\rho_{\text{tot}}(0) = \rho_S(0) \otimes \rho_E(0)$ with fixed $\rho_E(0)$, the map $\Phi_t : \rho_S(0) \mapsto \rho_S(t)$ is completely positive and trace-preserving.

Proof sketch. Write the initial environment state in its spectral decomposition $\rho_E(0) = \sum_\mu p_\mu |e_\mu\rangle\langle e_\mu|$. Then

$$\rho_S(t) = \sum_\mu p_\mu \sum_\nu \langle e_\nu | U(t) | e_\mu \rangle \rho_S(0) \langle e_\mu | U(t)^\dagger | e_\nu \rangle = \sum_{\mu,\nu} K_{\mu\nu}(t) \rho_S(0) K_{\mu\nu}(t)^\dagger,$$

where $K_{\mu\nu}(t) := \sqrt{p_\mu}\langle e_\nu | U(t) | e_\mu \rangle$ are operators on $\mathcal{H}_S$. The completeness relation

$$\sum_{\mu,\nu} K_{\mu\nu}(t)^\dagger K_{\mu\nu}(t) = \sum_\mu p_\mu \sum_\nu \langle e_\mu|U(t)^\dagger|e_\nu\rangle\langle e_\nu|U(t)|e_\mu\rangle = \sum_\mu p_\mu \langle e_\mu|I|e_\mu\rangle = 1$$

confirms trace preservation. The Kraus representation makes complete positivity manifest. ■

Non-uniqueness of the Kraus decomposition

Different environment-basis choices give different Kraus operators, but the same channel Φ_t. This is the measurement-theoretic content of the Naimark dilation (Node 3.3): the physical channel is intrinsic to the joint unitary and the initial environment state, but its Kraus decomposition is basis-dependent.

Failure of divisibility

For general H_{SE}, the family of channels $\{\Phi_t\}_{t\geq 0}$ is *not* a one-parameter semigroup — that is, $\Phi_{t+s} \neq \Phi_t \circ \Phi_s$ in general. This is the structural reason that **non-Markovian** dynamics exist: information can flow from environment back to system, and the reduced dynamics carry memory. Markovian approximations (Node 8.2) are precisely those that discard this memory and treat $\{\Phi_t\}$ as a semigroup.

The Born approximation

In many physical situations the environment is vastly larger than the system and its state is barely perturbed by the interaction. The **Born approximation** treats $\rho_E(t) \approx \rho_E(0)$ throughout the evolution, and writes $\rho_{\text{tot}}(t) \approx \rho_S(t) \otimes \rho_E(0)$. Combined with a short-memory (Markov) assumption, this yields the Lindblad master equation of Node 8.2. The structural content of the Born approximation is: *the environment is much larger than the system and absorbs perturbations without reacting.*

Weak coupling and linear-response form

For $H_{SE} = g\, S \otimes E$ (a single system operator S coupled to a single environment operator E with strength g), second-order perturbation theory gives the leading decoherence rate

$$\frac{d\rho_S}{dt} \sim -g^2 \int_0^\infty d\tau\, C_E(\tau)[S, [S(-\tau), \rho_S]],$$

where $C_E(\tau) = \text{Tr}(\rho_E E(\tau) E(0))$ is the environment's autocorrelation function. The decoherence rate is proportional to the environment's response function at the system's characteristic frequencies; the functional form of C_E determines whether the dynamics are Markovian (short correlations) or memory-laden (long correlations).

Worked Example

Single qubit dephasing via a bosonic bath

Take a qubit system with $H_S = \frac{\omega_0}{2} Z$ coupled to a bosonic environment (infinite set of harmonic oscillators) via

$$H_{SE} = Z \otimes \sum_k (g_k a_k^\dagger + g_k^* a_k),$$

and $H_E = \sum_k \omega_k a_k^\dagger a_k$. The total Hamiltonian is the **spin–boson model in the dephasing limit** — the coupling is diagonal in Z, so there is no energy exchange between system and environment, and the qubit's populations ($|0\rangle$ and $|1\rangle$ probabilities) are conserved. Only the phase relationship between them is affected.

Exact solution. The coupling is diagonal in Z, so the total evolution can be computed exactly. The off-diagonal matrix element of $\rho_S(t)$ evolves as

$$\rho_{01}(t) = \rho_{01}(0)\, e^{-i\omega_0 t}\, e^{-\Gamma(t)},$$

with

$$\Gamma(t) = 4\int_0^\infty d\omega\, \frac{J(\omega)}{\omega^2}(1-\cos\omega t)\coth\frac{\beta\omega}{2},$$

where $J(\omega) = \sum_k |g_k|^2\delta(\omega-\omega_k)$ is the environment's **spectral density** and β is the inverse temperature. The diagonal elements (populations) are unchanged.

- **Ohmic spectral density** $J(\omega) = \eta\omega e^{-\omega/\omega_c}$: $\Gamma(t) \sim \eta\ln(\omega_c t)$ at long times — logarithmic decoherence at zero temperature, linear in t at finite temperature.
- **Super-Ohmic** $J(\omega) \propto \omega^s$, $s > 1$: $\Gamma(t)$ saturates — the decoherence is bounded and reversible in the long-time limit. Entanglement is protected.
- **Sub-Ohmic** $J(\omega) \propto \omega^s$, $s < 1$: $\Gamma(t)$ grows faster than linearly — decoherence is stronger than Markovian.

The spectral density therefore determines whether the dynamics are Markovian, super-decoherent, or decoherence-protected; the functional form of $J(\omega)$ is a basic design parameter of quantum hardware.

The short-time linear regime

At very short times $t \ll 1/\omega_c$, regardless of the spectral density,

$$\Gamma(t) \approx \gamma\, t^2, \qquad \gamma = 2\int_0^\infty d\omega\, J(\omega)\coth\frac{\beta\omega}{2},$$

giving a quadratic-in-time decay. This is the **Zeno regime**, where frequent measurement (even weak) can freeze the dynamics and suppress decoherence. At slightly longer times $t \gg 1/\omega_c$ but still within the Markov regime, the decay becomes exponential: $\rho_{01}(t) \sim e^{-\Gamma_\infty t}$. The crossover is a structural feature of the coupling and is the reason Markovian approximations fail at short times.

The spin–boson model in the relaxing case

If instead the coupling were $H_{SE} = X \otimes \sum_k (g_k a_k^\dagger + g_k^* a_k)$ — a σ_x-type transverse coupling — the environment exchanges energy with the qubit, causing *relaxation* (amplitude damping) in addition to dephasing. The populations ρ_{00}, ρ_{11} now evolve toward thermal equilibrium, not just the off-diagonal elements. This is the structural distinction between T_2 **processes** (pure dephasing) and T_1 **processes** (relaxation), made precise in Node 5.4 and quantified in Node 5.6.

Cross-Links

- **Module 2.4 (Unitary evolution):** Closed-system dynamics are unitary; open-system dynamics are the partial trace of unitary dynamics on a larger space.
- **Module 3.6 (Back-action):** Decoherence is the back-action of the environment on the system averaged over an unread measurement.
- **Module 5.4 (Decoherence mechanisms):** Canonical channels (dephasing, amplitude damping) are the reduced dynamics in specific limits of system–environment coupling.
- **Module 8 (Open systems):** The Lindblad master equation is the Markovian approximation of the reduced dynamics derived here.

- **Module 7.3 (Stinespring dilation):** The structure "joint unitary + partial trace" is the universal form of CPTP maps; system–environment coupling is its physical realization.

Structural Compression

Invariant: The joint system-plus-environment state evolves unitarily; the reduced system state is obtained by partial trace. Reduced dynamics are in general non-unitary and non-Markovian.

Mechanism:

$$\rho_S(t) = \mathrm{Tr}_E\big(U(t)[\rho_S(0)\otimes\rho_E(0)]U(t)^\dagger\big), \qquad U(t) = e^{-iHt/\hbar}, \quad H = H_S + H_E + H_{SE}.$$

Cross-System Echo: Coarse-graining in statistical mechanics — integrating out microscopic degrees of freedom to obtain effective macroscopic dynamics — is the same structural operation. Decoherence is the quantum-mechanical analogue of the second law's thermodynamic arrow: unitary global evolution, irreversible-looking local dynamics after tracing out the unobserved bath.

Exercises

1. Show that for a trivial interaction $H_{SE} = 0$, the reduced dynamics are unitary $\Phi_t(\rho_S) = e^{-iH_St/\hbar}\rho_S e^{iH_St/\hbar}$.
2. For $H_{SE} = gZ_S\otimes\sigma_z^E$ with a single-qubit environment initially in $(|0\rangle+|1\rangle)/\sqrt{2}$, compute the reduced system dynamics exactly and show that the off-diagonal element of ρ_S oscillates in time without decaying. Why doesn't a single-qubit environment cause decoherence?
3. For an N-qubit environment in $|+\rangle^{\otimes N}$ coupled as $H_{SE} = Z_S\otimes\sum_k\sigma_z^{(k)}$, compute $\rho_{01}(t)$ and show that it decays as $\cos^N(gt)$. Estimate how fast $N\to\infty$ causes decoherence.
4. Derive the short-time quadratic decay $\Gamma(t)\approx\gamma t^2$ for the spin–boson model with Ohmic spectral density.
5. Show that for a decoherence-free subspace $\mathcal{H}_{\mathrm{DFS}}\subseteq\mathcal{H}_S$ (a subspace on which H_{SE} acts as a multiple of the identity), states in the subspace evolve unitarily under H_S alone.
6. Compute the reduced dynamics for an initially entangled system–environment state $|\Psi_{SE}(0)\rangle = (|00\rangle+|11\rangle)/\sqrt{2}$ with $H_{SE}=0$. Show that the "dynamics" look non-trivial even though the Hamiltonian is trivial. What goes wrong with the factorized-initial-state assumption?
7. Argue, structurally, why the factorized initial state $\rho_S\otimes\rho_E$ is an idealization, and what breaks down when the assumption fails.

Decoherence Mechanisms

Position in Knowledge Graph

System: Quantum Foundations **Structural Domain:** Decoherence **Node:** 5.4

Generic system–environment coupling (Node 5.3) produces messy reduced dynamics, but most physical situations are dominated by one of a few canonical channels. The two most important are **dephasing** (pure loss of phase information, no energy exchange) and **amplitude damping** (energy exchange with a zero-temperature environment, leading to relaxation to the ground state). Together they cover the dominant noise processes in almost every quantum hardware platform — superconducting qubits, trapped ions, nitrogen-vacancy centres, photonic qubits — and they are the Lindblad primitives from which more general noise models are built.

Formal Definition

Dephasing channel

The **qubit dephasing channel** with parameter $p \in [0,1]$ is the CPTP map

$$\Phi_{\text{deph}}(\rho) = (1-p)\rho + p\, Z\rho Z,$$

with Kraus operators

$$K_0 = \sqrt{1-p/2}\, I, \qquad K_1 = \sqrt{p/2}\, Z.$$

In matrix form on a qubit,

$$\rho = \begin{pmatrix} \rho_{00} & \rho_{01} \\ \rho_{10} & \rho_{11} \end{pmatrix} \quad \mapsto \quad \Phi_{\text{deph}}(\rho) = \begin{pmatrix} \rho_{00} & (1-p)\rho_{01} \\ (1-p)\rho_{10} & \rho_{11} \end{pmatrix}.$$

Populations (diagonal elements) are preserved. Coherences (off-diagonal elements) are scaled by $1-p$. In the continuous-time limit, $\rho_{01}(t) = \rho_{01}(0)\, e^{-t/T_2}$, where T_2 is the **dephasing time**.

Amplitude damping channel

The **amplitude damping channel** with parameter $\gamma \in [0,1]$ is the CPTP map with Kraus operators

$$K_0 = \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{1-\gamma} \end{pmatrix}, \qquad K_1 = \begin{pmatrix} 0 & \sqrt{\gamma} \\ 0 & 0 \end{pmatrix}.$$

Here $|1\rangle$ is the excited state and $|0\rangle$ the ground state. The map sends

$$\rho = \begin{pmatrix} \rho_{00} & \rho_{01} \\ \rho_{10} & \rho_{11} \end{pmatrix} \quad \mapsto \quad \begin{pmatrix} \rho_{00}+\gamma\rho_{11} & \sqrt{1-\gamma}\,\rho_{01} \\ \sqrt{1-\gamma}\,\rho_{10} & (1-\gamma)\rho_{11} \end{pmatrix}.$$

Excited-state population ρ_{11} decays to zero, transferring probability to ρ_{00}; coherences decay by a factor $\sqrt{1-\gamma}$. In the continuous-time limit, $\rho_{11}(t) = \rho_{11}(0)e^{-t/T_1}$, where T_1 is the **relaxation time**.

Generalized amplitude damping

At finite temperature, excited states can be *excited* back from the environment. The generalized channel sends ρ_{11} toward the thermal equilibrium population p_∞ rather than to zero, with Kraus operators parameterized by both γ and p_∞. The ground state eventually becomes a Gibbs-weighted mixture.

Intuition

Every realistic noise source on a quantum system can be schematically split into:

- **Phase noise (T_2 processes).** The environment "reads" which energy eigenstate the system is in, without exchanging energy. This looks like a weak Z-measurement: off-diagonal coherences decay, populations are preserved. The information that leaks to the environment is *which basis state* the system is in; the information that is lost from the system is the *relative phase* between basis states.
- **Amplitude noise (T_1 processes).** The environment exchanges energy with the system, pulling the population from the excited state to the ground state (or toward thermal equilibrium). This is what you get from spontaneous emission into a zero-temperature photon bath, or from any dissipative coupling.

Both happen simultaneously in real systems, and the dominant one depends on the platform: superconducting qubits are typically T_1-dominated (dielectric loss, radiative emission into the resonator), trapped-ion qubits are typically T_2-dominated (magnetic-field fluctuations, laser phase noise), photonic qubits are dominated by photon loss (a specific form of amplitude damping).

The two processes have very different structural character:

- **Dephasing** is *reversible* in principle — it is caused by the environment coupling to a diagonal observable, and if the environment's state were recorded and fed back the coherences could in principle be recovered. It is the quantum analogue of pure "measurement back-action without readout."
- **Amplitude damping** is *irreversible* — energy has actually left the system and entered the environment. Undoing it would require reversing the energy flow, which is an exact time-reversal of the environment's dynamics.

The distinction between the two is the structural content of "$T_2 \neq T_1$" — you can have phase noise without energy loss, but you cannot have energy loss without phase noise.

Structural Role

Dephasing and amplitude damping are the load-bearing noise primitives of quantum information:

- **Noise models for VQA and NISQ.** Circuit-level noise is typically modeled as single-qubit dephasing plus amplitude damping between gates, with two-qubit gate-specific error channels added on top. Benchmarking and error mitigation protocols target these two channels in particular.
- **Lindblad generators.** The continuous-time limits of these channels give the Lindblad dissipators $\mathcal{L}_Z$ (dephasing) and $\mathcal{L}_{\sigma_-}$ (relaxation) of Module 8.
- **Coherence time budgeting.** T_1 and T_2 are the headline hardware figures of merit: the number of coherent gates that can be executed before noise dominates is $\sim \min(T_1, T_2)/\tau_g$, where τ_g is the gate time.
- **$T_2 \leq 2T_1$.** A general relationship: pure dephasing is always bounded below by the dephasing induced by relaxation. Any T_1 process contributes to T_2, so the pure dephasing rate is $1/T_\phi = 1/T_2 - 1/(2T_1) \geq 0$.
- **Error correction (Module 7).** Stabilizer codes target the bit-flip, phase-flip, and combined errors that these channels produce in the Pauli decomposition.
- **Bloch ball picture.** Dephasing collapses the Bloch ball onto the z-axis; amplitude damping collapses it onto the south pole (ground state). The geometric picture is cleanest for a qubit.

Mathematical Development

Unitary dilation of the dephasing channel

The dephasing channel $\Phi_{\text{deph}}(\rho) = (1-p)\rho + pZ\rho Z$ arises from a CNOT-like coupling to an ancilla qubit:

$$U = |0\rangle\langle 0|_S \otimes I_E + |1\rangle\langle 1|_S \otimes X_E,$$

applied to a system state $\rho_S \otimes |e_0\rangle\langle e_0|$ with ancilla initial state $|e_0\rangle = \sqrt{1-p/2}|0\rangle_E + \sqrt{p/2}|1\rangle_E$. Tracing out the ancilla gives the dephasing channel. This is a direct instance of Stinespring dilation: a concrete joint unitary + partial trace realization of a noise channel.

Unitary dilation of the amplitude damping channel

The amplitude damping channel arises from a beam-splitter-like coupling between system and environment (initially in the ground state):

$$U = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \sqrt{1-\gamma} & \sqrt{\gamma} & 0 \\ 0 & -\sqrt{\gamma} & \sqrt{1-\gamma} & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix},$$

acting on $(|0\rangle, |1\rangle)_S \otimes (|0\rangle, |1\rangle)_E$ with initial environment state $|0\rangle_E$. Tracing out the environment recovers the Kraus operators above. Physically: if the system is in the excited state $|1\rangle$, it transfers a quantum of energy to the environment with probability γ, leaving both in their respective ground/excited states entangled.

Bloch ball dynamics

Under dephasing with parameter p, the Bloch vector (r_x, r_y, r_z) evolves as

$$(r_x, r_y, r_z) \mapsto ((1-2p)r_x, (1-2p)r_y, r_z).$$

The ball is compressed along the xy-plane, collapsing onto the z-axis as $p \to 1/2$. The z-axis is the *decoherence-free* direction: eigenstates of Z are fixed points.

Under amplitude damping with parameter γ,

$$r_x \mapsto \sqrt{1-\gamma}\, r_x, \quad r_y \mapsto \sqrt{1-\gamma}\, r_y, \quad r_z \mapsto (1-\gamma)r_z + \gamma.$$

The ball is compressed and shifted: xy-components decay by $\sqrt{1-\gamma}$, z-component decays toward the ground state $r_z = 1$ (the fixed point is the south pole if we orient $|0\rangle$ downward, or the north pole if upward). As $\gamma \to 1$ every state is sent to the ground state $|0\rangle$.

The $T_2 \leq 2T_1$ bound

Pure dephasing has rate $1/T_\phi$; amplitude damping has rate $1/T_1$. The total coherence decay rate is

$$\frac{1}{T_2} = \frac{1}{2T_1} + \frac{1}{T_\phi}.$$

The factor of $1/2$ in the T_1 contribution is because amplitude damping decays the coherences at *half* the rate at which it decays the populations (this follows from the Kraus operator K_0, which has $\sqrt{1-\gamma}$ on the excited entry but no decay on the ground entry). Hence $T_2 \leq 2T_1$, with equality when pure dephasing vanishes. In practice, pure dephasing is usually significant, so $T_2 < 2T_1$.

Composition of channels

The product of dephasing and amplitude damping is computed by composing Kraus operators. The result is another CPTP channel with six Kraus operators (the tensor-product structure), describing combined noise. In the continuous-time Lindblad limit, dephasing and amplitude damping combine additively at the level of dissipators:

$$\mathcal{L}(\rho) = \gamma_\phi\big(Z\rho Z - \rho\big) + \gamma_1\big(\sigma_-\rho\sigma_+ - \tfrac{1}{2}\{\sigma_+\sigma_-, \rho\}\big).$$

The two generators commute only if γ_ϕ or γ_1 vanishes; otherwise the dynamics must be solved jointly.

Worked Example

Qubit Bloch ball under dephasing

Start with a pure state on the equator of the Bloch sphere, $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$, density matrix

$$|+\rangle\langle+| = \tfrac{1}{2}(I + X) = \tfrac{1}{2}\begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}.$$

Bloch vector $\vec{r} = (1, 0, 0)$, pure (length 1).

Apply dephasing at rate $1/T_2$ for time t. The off-diagonal element shrinks by e^{-t/T_2}:

$$\rho(t) = \tfrac{1}{2}\begin{pmatrix} 1 & e^{-t/T_2} \\ e^{-t/T_2} & 1 \end{pmatrix}.$$

Bloch vector $\vec{r}(t) = (e^{-t/T_2}, 0, 0)$. At $t \to \infty$, $\rho \to I/2$, maximally mixed. Entropy grows from 0 to 1 bit; purity decreases from 1 to 1/2. The state migrates along the x-axis from the Bloch sphere to the center.

Crucially, the state $|0\rangle$ (Bloch vector $(0, 0, 1)$) is completely unchanged by dephasing — it is a pointer state of the Z-dephasing channel. Any classical mixture of $|0\rangle, |1\rangle$ is also unchanged.

Excited state under amplitude damping

Start with $\rho(0) = |1\rangle\langle1|$, the excited state. Under amplitude damping with rate $1/T_1$,

$$\rho_{11}(t) = e^{-t/T_1}, \qquad \rho_{00}(t) = 1 - e^{-t/T_1}.$$

The excited-state population decays exponentially to zero, and the lost probability goes entirely to the ground state. At $t \to \infty$, $\rho \to |0\rangle\langle0|$, the pure ground state. Purity is preserved at infinity (the fixed point is pure), even though the trajectory passes through mixed states at intermediate times.

For an initial superposition $|+\rangle$, amplitude damping drives the state toward the ground state $|0\rangle$, not the equatorial center $I/2$ — a key difference from dephasing. The asymptotic state is pure because zero-temperature amplitude damping has a pure fixed point.

Combined dynamics for superconducting transmon qubit

For a typical superconducting qubit: $T_1 \approx 100\,\mu s$, $T_2 \approx 50\,\mu s$, gate time $\tau_g \approx 10$ ns. The Bloch vector at time t starting from $|+\rangle$:

$$\vec{r}(t) = \big(e^{-t/T_2}, 0, 1 - e^{-t/(2T_1)}\big) \quad \text{(approximately)}.$$

In the long-time limit the state converges to $(0, 0, 1) = |0\rangle$, the ground state. The decoherence time T_2 sets the timescale for the xy-plane components; the relaxation time T_1 sets the timescale for the z-component's approach to the ground state. The ratio $T_2/\tau_g \approx 5000$ gives a rough upper bound on circuit depth before noise dominates.

Cross-Links

- **Module 5.3 (System–environment coupling):** Dephasing and amplitude damping are the reduced dynamics for specific choices of H_{SE} (diagonal vs. off-diagonal coupling).
- **Module 5.5 (Pointer basis):** Dephasing makes the Z-basis the pointer basis. Amplitude damping makes the ground state the unique fixed point.
- **Module 5.6 (Decoherence timescales):** T_1 and T_2 parameterize these channels in continuous-time form.
- **Module 8.2 (Lindblad equation):** The continuous-time generators $\mathcal{L}_Z$ and $\mathcal{L}_{\sigma_-}$ are the infinitesimal versions of these channels.
- **Module 7 (Error correction):** Stabilizer codes are designed to correct the Pauli decomposition of these noise channels.
- **VQA:** Circuit-level noise is typically modeled as compositions of these two primitives.

Structural Compression

Invariant: Dephasing destroys phase information and preserves energy; amplitude damping destroys excitation and exchanges energy.

Mechanism: Two Lindblad generators: $\mathcal{L}_Z(\rho) = Z\rho Z - \rho$ (dephasing) and $\mathcal{L}_{\sigma_-}(\rho) = \sigma_-\rho\sigma_+ - \frac{1}{2}\{\sigma_+\sigma_-, \rho\}$ (relaxation).

Cross-System Echo: Friction in classical mechanics dissipates kinetic energy; phase noise in classical oscillators shifts relative phases without changing amplitude. The two are quantum-mechanically the same distinction made precise: T_1 = classical friction, T_2 = classical phase noise, with the additional quantum constraint $T_2 \leq 2T_1$ that has no classical analogue.

Exercises

1. Verify that the dephasing Kraus operators $K_0 = \sqrt{1 - p/2}I, K_1 = \sqrt{p/2}Z$ satisfy $K_0^\dagger K_0 + K_1^\dagger K_1 = I$.
2. Verify that the amplitude damping Kraus operators satisfy completeness and compute their action on the Bloch vector.
3. Show that the dephasing channel is unital (preserves I), but the amplitude damping channel is not.
4. Compute the composition of dephasing with parameter p_1 and dephasing with parameter p_2. Verify it is dephasing with parameter $p_1 + p_2 - 2p_1p_2$.
5. Show that amplitude damping has a unique fixed point $|0\rangle\langle 0|$, while dephasing has a one-parameter family of fixed points (all Z-diagonal states).
6. Derive the $T_2 \leq 2T_1$ bound by computing how amplitude damping alone affects the off-diagonal elements of ρ and comparing with the population decay rate.
7. Argue, structurally, why no single decoherence channel can be both energy-conserving *and* produce a pure-state fixed point other than an energy eigenstate.

Pointer Basis and Einselection

Position in Knowledge Graph

System: Quantum Foundations **Structural Domain:** Decoherence **Node:** 5.5

Not every basis of a system's Hilbert space is equally "classical." Decoherence dynamically selects a preferred basis — the **pointer basis** — in which superpositions are rapidly destroyed and eigenstates are approximately stable. Wojciech Zurek's term for this dynamical selection is **einselection** ("environment-induced superselection"). The pointer basis is the *dynamical* answer to the question "why does the world look classical?" — it explains why we observe definite positions and momenta rather than arbitrary superpositions of them, without resolving the deeper measurement problem of Module 6.

Formal Definition

The pointer basis (schematic)

Given a system S coupled to an environment E via an interaction Hamiltonian H_{SE}, the **pointer basis** of S is the basis $\{|p_i\rangle\}$ of $\mathcal{H}_S$ in which:

1. **States are stable.** If $\rho_S(0) = |p_i\rangle\langle p_i|$, then $\rho_S(t) \approx |p_i\rangle\langle p_i|$ for all t of interest, up to small corrections from the system's own dynamics.
2. **Superpositions decohere.** If $\rho_S(0) = |\psi\rangle\langle\psi|$ with $|\psi\rangle = \sum_i c_i|p_i\rangle$, the reduced state rapidly becomes $\rho_S(t) \approx \sum_i |c_i|^2 |p_i\rangle\langle p_i|$ — diagonal in the pointer basis — on a timescale τ_D much shorter than the system's own dynamical timescales.

The commutativity criterion (idealized)

In the limit of dominant system–environment coupling (when H_{SE} overwhelms H_S), the pointer basis is the basis in which a **pointer observable** O_S satisfies

$$[O_S, H_{SE}] = 0.$$

This is the ideal case. States that are eigenstates of O_S are unaffected by the interaction and are stable under environmental coupling; states in superposition across different eigenvalues of O_S are dephased. The eigenbasis of O_S is the pointer basis, and the process of environment-induced diagonalization in this basis is **einselection**.

The predictability sieve

More generally, when H_S is non-negligible and H_{SE} does not strictly commute with any system observable, the pointer basis is defined by the **predictability sieve** (Zurek 1993): the basis that minimizes the growth of entropy of the reduced state over a fixed timescale. The pointer states are the *most classical* pure states in the sense that, of all pure states, they maintain the most classical-looking behavior under the open dynamics.

Formally: define $S(\psi, t) := -\operatorname{Tr}(\rho_\psi(t) \log \rho_\psi(t))$ where $\rho_\psi(t)$ is the reduced state at time t starting from $|\psi\rangle$. The pointer states are the pure states $|\psi\rangle$ minimizing $S(\psi, t)$ in some averaged sense. Equivalently, they are the states most robust against entanglement with the environment.

Intuition

A classical object has a definite position, a definite momentum, a definite energy — we don't observe a car being in a superposition of "parked" and "driving." Why not? The textbook answer "measurement collapses

the wavefunction" is circular: it presupposes a pre-existing classical apparatus doing the measuring, without explaining *why* that apparatus behaves classically.

Einselection gives a dynamical answer. The environment is constantly "measuring" every macroscopic system, in whatever basis the interaction Hamiltonian picks out. For an object like a car, the dominant coupling is through position (every scattered photon carries position information), so the environment incessantly dephases the car in the position basis. Superpositions of "parked here" and "parked there" decohere on astronomically short timescales — typically 10^{-23} seconds or faster for macroscopic objects. The result: macroscopic systems *dynamically* live in the pointer basis, and the pointer basis is the position basis.

For smaller systems, different couplings dominate. An atom in a thermal electromagnetic bath decoheres in the energy basis (the environment couples to its dipole moment via photon exchange). A spin qubit in a nuclear bath decoheres in the Z basis (hyperfine coupling is diagonal in spin-z). The pointer basis depends on what the environment is looking at.

The structural insight: **there is no universal pointer basis. The pointer basis is determined by the structure of the system–environment interaction, not by the system's own Hamiltonian.** Different environments pick out different pointer bases for the same system, and different bases for the same system's states in different physical setups.

Einselection does not explain *why* there are definite outcomes (that is the measurement problem, Module 6). It explains *why the outcomes we see are the ones we see* — why we see positions rather than position superpositions, why we see energies rather than energy superpositions. The selection is dynamical, driven by the structure of the coupling to the environment.

Structural Role

Einselection is the structural story behind the quantum-to-classical transition:

- **The classical world as a dynamical phenomenon.** Classical reality is not a separate layer of physics; it is the limit in which decoherence is so fast that the pointer basis is effectively exact. Module 10.4 (quantum-to-classical transition) builds on this.
- **The interpretive status of "collapse."** Einselection provides a dynamical mechanism for *apparent* collapse: reduced states become diagonal in the pointer basis, losing their off-diagonal coherences exponentially fast. It does not provide a mechanism for the *selection* of a particular outcome — that is still the measurement problem.
- **Decoherence-free subspaces (DFS).** A subspace on which H_{SE} acts trivially (as a multiple of the identity) is unaffected by decoherence. DFS engineering is a passive error-correction technique: states prepared in DFS are protected without active intervention. The structural characterization of DFS is the complementary side of the pointer basis — where einselection selects classical pointer states, DFS selects quantum-protected states.
- **Quantum Darwinism (Node 5.6 in Zurek's framework, Node 5.6 in this course uses the timescales slot instead).** Pointer states are "objective" in the sense that multiple environmental fragments all carry the same information about them, so multiple observers can agree on the system's state without disturbing it. This is Zurek's proposed mechanism for the emergence of objective classical reality.
- **Schrödinger's cat.** The pointer basis for a macroscopic cat is the {alive, dead} basis (or more precisely, the position-like basis of macroscopic configurations). Superpositions $(|\text{alive}\rangle + |\text{dead}\rangle)/\sqrt{2}$ decohere in fantastically short times — Zurek estimates $\sim 10^{-23}$ s for a realistic cat in a realistic environment. This is why we do not observe macroscopic superpositions: they exist for impossibly short times before being einselected.

Mathematical Development

Exact commutativity case

Suppose $[O_S, H_{SE}] = 0$ exactly, and H_{SE} can be written as

$$H_{SE} = \sum_i |p_i\rangle\langle p_i| \otimes B_i,$$

where $\{|p_i\rangle\}$ is the eigenbasis of O_S (and of H_{SE} on the system side), and B_i are environment operators. The joint evolution $U(t) = e^{-iHt/\hbar}$ then acts on a system basis state $|p_i\rangle \otimes |e_0\rangle$ as

$$U(t)|p_i\rangle \otimes |e_0\rangle = |p_i\rangle \otimes |e_i(t)\rangle,$$

with $|e_i(t)\rangle := e^{-iB_i t/\hbar}|e_0\rangle$ (ignoring the system's own Hamiltonian for clarity). That is, the pointer states $|p_i\rangle$ are *frozen* by the interaction — they do not evolve into superpositions of other pointer states.

For a superposition $|\psi\rangle = \sum_i c_i|p_i\rangle$, the joint state at time t is

$$|\Psi(t)\rangle = \sum_i c_i|p_i\rangle \otimes |e_i(t)\rangle.$$

Tracing out the environment:

$$\rho_S(t) = \sum_{i,j} c_i c_j^* \langle e_j(t)|e_i(t)\rangle |p_i\rangle\langle p_j|.$$

The diagonal elements $i = j$ give $\langle e_i(t)|e_i(t)\rangle = 1$, independent of time. The off-diagonal elements $i \neq j$ give $\langle e_j(t)|e_i(t)\rangle$, the overlap of two environmental branches — the **decoherence factor**. For any reasonable environment, these overlaps decay rapidly to zero as the branches become distinguishable, leaving $\rho_S(t)$ approximately diagonal in the pointer basis.

The structural conclusion: *pointer states are the states whose environmental "image" does not change under the dynamics*, so their coherences survive, while superpositions across them are destroyed by the environment learning which branch the system is in.

The predictability sieve (schematic)

When $[O_S, H_{SE}] \neq 0$ exactly but H_{SE} dominates over H_S, the pointer basis is only approximately defined. The predictability sieve works as follows:

1. For each pure state $|\psi\rangle$, compute the reduced-state entropy $S(t)_{|\psi\rangle}$ after a short time t.
2. The pointer states are those $|\psi\rangle$ for which $S(t)$ grows most slowly.

Equivalently: pointer states are states of maximum robustness — states whose purity decays most slowly under the joint dynamics.

For ordinary physical couplings the pointer basis identified by the predictability sieve agrees with the $[O_S, H_{SE}] = 0$ basis to leading order, and corrections are perturbative in H_S/H_{SE}.

Decoherence-free subspaces

A subspace $\mathcal{H}_{\mathrm{DFS}} \subseteq \mathcal{H}_S$ is a **decoherence-free subspace** with respect to $H_{SE} = \sum_k S_k \otimes E_k$ if every system operator S_k acts as a multiple of the identity on $\mathcal{H}_{\mathrm{DFS}}$:

$$S_k|\psi\rangle = s_k|\psi\rangle \qquad \forall|\psi\rangle \in \mathcal{H}_{\mathrm{DFS}},$$

for scalars s_k. States in $\mathcal{H}_{\text{DFS}}$ are unentangled with the environment under the interaction (up to a global phase), and they evolve unitarily under the effective Hamiltonian

$$H_{\text{eff}} = H_S + \sum_k s_k E_k.$$

DFS can be explicitly constructed when the interaction has a symmetry. For instance, if H_{SE} is permutation-symmetric on a set of qubits (e.g., collective dephasing), the total spin states $|j, m\rangle$ with $m = 0$ constitute a DFS. This is the structural basis of several quantum-memory protocols.

The timescale separation

For einselection to work, decoherence must be much faster than the system's own internal dynamics:

$$\tau_D \ll \tau_S,$$

where τ_D is the decoherence time (the timescale on which off-diagonal coherences of the reduced state decay) and τ_S is the system's own dynamical timescale ($\hbar/\|H_S\|$ or similar). When this inequality holds, the system has no time to evolve coherently before being decohered, and the pointer basis is well-defined. When it fails, the system's own Hamiltonian can mix pointer states and the basis loses its classical character.

For macroscopic systems the inequality is satisfied by enormous margins — decoherence times are on the order of 10^{-23} s while dynamical times are on the order of seconds — so the pointer basis is essentially exact. For microscopic systems the inequality may not hold at all, and decoherence produces subtler effects (e.g., line broadening rather than classical selection).

Worked Example

Position as the pointer basis for a massive particle

Consider a particle of mass m coupled to a thermal bath of photons (ambient radiation) or air molecules. The dominant scattering process is elastic scattering off the particle, with the scattered photons/molecules carrying information about the particle's position. The interaction is schematically

$$H_{SE} \sim V(\hat{x}_S, \hat{X}_E),$$

where $\hat{x}_S$ is the particle's position and $\hat{X}_E$ summarizes the environment's degrees of freedom. The interaction is *diagonal* in the position basis — scattering preserves the particle's position at leading order — so the pointer basis is the position eigenbasis.

The quantitative result (Joos and Zeh 1985): for a macroscopic object of cross-section σ, embedded in a gas of particles of density n and thermal wavelength λ, the decoherence rate for a spatial superposition of extent Δx is

$$\Gamma \sim n\sigma v \cdot \frac{(\Delta x)^2}{\lambda^2},$$

where v is the thermal velocity. For realistic values, $\Gamma \gtrsim 10^{23}$ per second — decoherence on Planck-like timescales. Position superpositions of macroscopic objects do not survive.

Energy as the pointer basis for an atom

Consider an atom coupled to the vacuum electromagnetic field via the dipole interaction

$$H_{SE} = -\vec{d} \cdot \vec{E},$$

where $\vec{d}$ is the atomic dipole moment and $\vec{E}$ is the field. The dipole operator is off-diagonal in the atomic energy basis (it couples different energy eigenstates), so H_{SE} does *not* commute with the atomic Hamiltonian H_S. Nevertheless, at weak coupling the predictability sieve identifies the energy basis as the pointer basis: superpositions of different energy eigenstates are the ones that get dephased, and energy eigenstates are the most robust under the dynamics (they only decay by spontaneous emission, on the T_1 timescale).

The intuition: the environment "sees" the atom through its dipole radiation, which is diagonal in transition frequencies (energy differences). So different energy eigenstates couple differently to the field and are distinguishable by the environment, while superpositions of them are dephased.

Qubit pointer basis under Z-coupling

For a qubit with $H_S = (\omega_0/2)X$ (transverse magnetic field) and $H_{SE} = Z \otimes B$ (pure dephasing coupling), the pointer basis is the Z-eigenbasis $\{|0\rangle, |1\rangle\}$ — the eigenbasis of the operator that appears in H_{SE}.

However, note that $[H_S, Z] \neq 0$: the system Hamiltonian tries to rotate Z-eigenstates around the X-axis, away from the pointer basis. The resulting dynamics is a competition between the environment einselecting the Z-basis and the system Hamiltonian trying to rotate out of it. When the dephasing rate $1/T_2$ is much larger than the Larmor frequency ω_0, the environment wins and the qubit is pinned in the Z-basis (quantum Zeno regime). When $\omega_0 \gg 1/T_2$, the system wins and rotates freely.

Schrödinger's cat

Zurek's canonical estimate: a 1 kg cat in a superposition of $\Delta x \sim 10$ cm at room temperature decoheres in

$$\tau_D \sim 10^{-23} \text{ seconds.}$$

This is 13 orders of magnitude shorter than the Planck time for the room-temperature gas bath, and staggeringly shorter than any conceivable observational timescale. In this regime, the pointer basis (macroscopic position configurations) is exact to any accuracy we could care about, and there is no physical possibility of observing a cat in a coherent superposition of alive and dead. Schrödinger's cat is a thought experiment, not a prediction.

Cross-Links

- **Module 5.3 (System–environment coupling):** Einselection is the structural consequence of the coupling Hamiltonian, computed by tracing out the environment.
- **Module 5.4 (Decoherence mechanisms):** Dephasing is einselection in the commuting observable's basis; amplitude damping is einselection + energy relaxation.
- **Module 5.6 (Decoherence timescales):** The decoherence time τ_D is what determines how fast the pointer basis takes over.
- **Module 6 (Interpretations):** Einselection provides a dynamical story for *why the pointer basis is what it is*, but does not solve the measurement problem (why there are definite outcomes at all).
- **Module 10.4 (Quantum-to-classical transition):** The pointer basis is the structural bridge from quantum to classical mechanics.
- **Systems Thinking:** Stable equilibria in dynamical systems — states that are robust under perturbations from the environment.

Structural Compression

Invariant: The pointer basis is the basis dynamically selected by the system–environment coupling. States in the pointer basis are stable; superpositions across it are exponentially suppressed.

Mechanism: In the strong-coupling limit, pointer states are eigenstates of a system observable commuting with H_{SE}. More generally (predictability sieve), they are the pure states whose reduced-state entropy grows most slowly under the joint dynamics.

Cross-System Echo: Attractors in dynamical systems and stable equilibria in control theory are the same structural idea: states selected by the dynamics as being robust under perturbation. Einselection specifies "robust under perturbation from the environment" with a quantum-specific notion of decoherence. The existence of a classical world is, from this point of view, the statement that the pointer basis is exact enough, fast enough, that its selection is indistinguishable from a fundamental law.

Exercises

1. For $H_{SE} = Z \otimes B$ with B an arbitrary environment operator, show that the pointer basis is $\{|0\rangle, |1\rangle\}$ (the Z-eigenbasis).
2. For $H_{SE} = X \otimes B$, identify the pointer basis. How does it relate to the Z-basis?
3. Show that a subspace spanned by $\{|01\rangle, |10\rangle\}$ of two qubits is a decoherence-free subspace for collective dephasing $H_{SE} = (Z_1 + Z_2) \otimes B$.
4. Compute the reduced dynamics of a qubit with $H_S = (\omega_0/2)X$ and $H_{SE} = Z \otimes B$ in the limit of strong dephasing (Zeno regime). Show that the qubit is frozen in the Z-basis.
5. Argue why the position basis is the pointer basis for a macroscopic object scattering photons, and why the energy basis would be the pointer basis for an atom emitting radiation.
6. Zurek's predictability sieve minimizes entropy growth of the reduced state. Argue why this is equivalent, for short times and weak dynamics, to the commutativity criterion $[O_S, H_{SE}] = 0$.
7. Argue, structurally, why einselection explains the apparent classicality of the macroscopic world but does not solve the measurement problem.

Decoherence Timescales

Position in Knowledge Graph

System: Quantum Foundations **Structural Domain:** Decoherence **Node:** 5.6

How fast does decoherence happen? For macroscopic objects the answer is "unobservably fast" — Schrödinger's cat decoheres in something like 10^{-23} seconds. For microscopic systems the answer is "fast enough to matter but slow enough to exploit" — superconducting qubits maintain coherence for microseconds to milliseconds, long enough to run quantum circuits. The structural scaling that explains both regimes, and everything in between, is the topic of this node. Decoherence rates depend on the coupling strength, the spectral density of the environment at the relevant frequencies, and the "size" of the superposition in whatever basis is being einselected. Understanding this scaling is the key to both the apparent classicality of the macroscopic world and the engineering challenge of building quantum computers.

Formal Definition

T_1, T_2, T_ϕ

For a qubit coupled to an environment, the decoherence times are defined operationally from the decay of density-matrix elements in the energy eigenbasis:

- **Relaxation time T_1:** The characteristic timescale of population decay. If the qubit starts in the excited state, $\rho_{11}(t) - \rho_{11}^{(eq)} = (\rho_{11}(0) - \rho_{11}^{(eq)})e^{-t/T_1}$, where $\rho_{11}^{(eq)}$ is the thermal equilibrium excited-state population.
- **Dephasing time T_2:** The characteristic timescale of off-diagonal coherence decay. $|\rho_{01}(t)| = |\rho_{01}(0)|e^{-t/T_2}$.
- **Pure dephasing time T_ϕ:** The coherence decay due to phase noise alone, not counting the contribution from relaxation. Related by

$$\frac{1}{T_2} = \frac{1}{2T_1} + \frac{1}{T_\phi}.$$

The factor of $1/2$ in the T_1 contribution reflects the fact that amplitude damping decays off-diagonal elements at half the rate it decays populations (the Kraus operator $K_0 = \mathrm{diag}(1, \sqrt{1-\gamma})$ gives $\sqrt{1-\gamma}$ on the off-diagonal but $(1-\gamma)$ on the population).

General decoherence rate from Fermi's golden rule

For a system weakly coupled to a bath through $H_{SE} = A \otimes B$ (single system operator A, single bath operator B), the decoherence rate at second order in perturbation theory is

$$\Gamma = \frac{1}{\hbar^2} \int_{-\infty}^{\infty} d\tau \, \langle B(\tau) B(0) \rangle_{\rho_E} \, e^{i\omega\tau} = \frac{1}{\hbar^2} \, S_B(\omega),$$

where $S_B(\omega)$ is the **spectral density** of the bath operator B at the transition frequency ω corresponding to the system process of interest. This is Fermi's golden rule for decoherence: the rate is the Fourier transform of the bath's autocorrelation function evaluated at the system's characteristic frequency.

For the qubit relaxation rate,

$$\frac{1}{T_1} = \frac{|g|^2}{\hbar^2} \, J(\omega_0) \left[\, n_{\mathrm{th}}(\omega_0) + 1 \, \right],$$

where g is the coupling strength, $J(\omega)$ is the environment's spectral density, ω_0 is the qubit transition frequency, and $n_{\text{th}}(\omega) = (e^{\beta\omega} - 1)^{-1}$ is the thermal occupation.

For the pure dephasing rate at low frequencies,

$$\frac{1}{T_\phi} = \frac{|g|^2}{\hbar^2} S_z(0),$$

where $S_z(0)$ is the zero-frequency spectral density of the environmental longitudinal noise. Dephasing is sensitive to *low-frequency* noise (flicker noise, $1/f$ noise) while relaxation is sensitive to noise at the qubit frequency.

Decoherence rate for macroscopic superpositions

For a massive object in a spatial superposition of separation Δx, scattering off environmental particles of thermal wavelength λ, the decoherence rate is (Joos–Zeh 1985, Gallis–Fleming 1990)

$$\Gamma \sim \Lambda \, (\Delta x)^2,$$

with Λ a scattering constant that scales linearly with the environmental density and scattering cross-section. The $(\Delta x)^2$ scaling is the structural signature: **decoherence rate grows as the square of the superposition size**. For any macroscopic Δx, this gives astronomically large rates.

Intuition

Three structural factors determine how fast a system decoheres:

1. **Coupling strength.** Stronger coupling to the environment means faster decoherence. The rate is typically quadratic in the coupling, $\Gamma \propto |g|^2$.
2. **Bath spectral density.** The environment must have modes at the relevant frequency to absorb the system's information. A bath with no modes at the system's transition frequency cannot drive relaxation, no matter how strong the coupling.
3. **Superposition "size".** For a spatially extended superposition, the decoherence rate grows with the square of the separation, because more environmental scattering events are needed to resolve a larger superposition. This is the reason Schrödinger's cat decoheres so fast: its superposition covers macroscopic distances.

The combined message: **the more macroscopic, the more coupled, and the noisier the environment, the faster decoherence.** This is a quantitative gradient, not a sharp boundary — smaller, better-isolated, colder systems decohere slower, and the line between "quantum" and "classical" is determined by where the decoherence rate crosses the threshold of observational irrelevance.

The key intuitive distinction between T_1 and T_2:

- T_1 **(relaxation)**: requires *energy exchange* with the environment. The bath must have modes at the qubit transition frequency to accept a quantum of energy. Typically limited by the density of states and temperature of the environment at ω_0.
- T_ϕ **(pure dephasing)**: requires only *phase fluctuations* of the qubit's energy levels. No energy exchange, so any low-frequency noise (temperature fluctuations, flux noise, charge noise) contributes. Usually dominated by $1/f$ noise in solid-state qubits.

In most experimental platforms, pure dephasing dominates: $T_\phi \ll 2T_1$, so $T_2 \approx T_\phi$. Exceptions are systems with good shielding from low-frequency noise (trapped ions in magnetic shielding, NV centres in isotopically purified diamond), where $T_2 \approx 2T_1$ can be approached.

Structural Role

Decoherence timescales are the load-bearing hardware parameters of every quantum technology:

- **Quantum computing depth budget.** The number of gates that can be executed coherently is roughly $N \sim T_2/\tau_g$, where τ_g is the gate time. For current superconducting qubits, $T_2 \sim 50\mu s$ and $\tau_g \sim 10$ ns, giving $N \sim 5000$. Increasing either parameter is a major engineering target.
- **Error correction threshold.** Fault-tolerant quantum computing requires physical error rates below $\sim 10^{-3}$ per gate. For gates with time τ_g, this means $T_2/\tau_g \gtrsim 1000$. Reaching this threshold is the structural gate to scalable quantum computing.
- **Quantum sensing and metrology.** The Ramsey interferometry uncertainty on a frequency measurement scales as $1/(T_2\sqrt{N})$ where N is the number of shots. Longer T_2 means better sensors.
- **The classical limit.** For macroscopic objects, $\Gamma \sim 10^{23}\,\mathrm{s}^{-1}$ or larger. This is the structural reason "classicality" is the overwhelming norm: the decoherence rate is faster than any physical timescale. Quantum coherence is the exception, requiring deliberate isolation.
- **Quantum Zeno effect.** When the environment couples strongly enough, $T_2 \ll$ any internal dynamical time, the system is pinned in the pointer basis by continuous einselection. This is the structural mechanism behind the Zeno effect (Node 3.5/3.6 continuous limit).
- **Hardware platform comparison.** Different platforms have wildly different T_1, T_2 profiles. Trapped ions: $T_1 \sim$ hours, $T_2 \sim$ seconds. Superconducting qubits: $T_1 \sim 100\mu s$, $T_2 \sim 50\mu s$. Photonic qubits: no relaxation (photons don't decay), but photon loss (amplitude damping) over fiber is the effective coherence limit. NV centres: $T_2 \sim$ ms at room temperature, seconds at cryogenic.

Mathematical Development

Derivation of Fermi's golden rule for decoherence

Consider $H = H_S + H_E + g\, A_S \otimes B_E$ and a system initially in $|\psi_i\rangle$ and environment in ρ_E (thermal state of H_E). At second order in g, the rate of transition out of $|\psi_i\rangle$ is

$$\Gamma_{i\to f} = \frac{2\pi}{\hbar}|g|^2|\langle f|A_S|i\rangle|^2\,\delta(E_f - E_i - \omega)\,S_B(\omega),$$

summed over final states $|f\rangle$. The transition rate depends on the matrix elements of the system operator, the density of final states, and the bath spectral density at the right frequency. This is the standard Fermi's golden rule for transitions into a continuum.

For decoherence specifically, the coherence between two states $|i\rangle, |j\rangle$ decays at rate

$$\frac{1}{T_{ij}} = \frac{1}{2}\left(\frac{1}{T_1^{(i)}} + \frac{1}{T_1^{(j)}}\right) + \frac{1}{T_\phi^{(ij)}},$$

where the first term is the average of the two states' relaxation rates (amplitude damping) and the second is pure dephasing between the pair. For a two-level system this reduces to the standard $1/T_2 = 1/(2T_1) + 1/T_\phi$.

Short-time vs. long-time behavior

At very short times, decoherence is *quadratic* in time rather than exponential:

$$|\rho_{01}(t)| \approx |\rho_{01}(0)|\,(1 - \tfrac{1}{2}\gamma_0 t^2 + O(t^4)),$$

with $\gamma_0 = \mathrm{Var}(H_{SE})/\hbar^2$ the variance of the coupling. This is the **Zeno regime**. As t increases past the bath correlation time τ_c, the dynamics transition to exponential decay:

$$|\rho_{01}(t)| \sim e^{-t/T_2}.$$

The transition happens at $t \sim \tau_c$, the inverse of the environment's characteristic frequency cutoff. For $t \ll \tau_c$, frequent measurement (or even the Zeno-type self-interruption) can suppress decoherence; for $t \gg \tau_c$, exponential decay is unavoidable.

Scaling with system size

For a collection of N qubits coupled independently to the same environment, the *collective* decoherence rate can scale as N^2 rather than N — **superdecoherence**. This is the structural reason that decoherence-free subspaces (which are immune to collective noise) become increasingly valuable at large N: they are the only way to resist the superdecoherence scaling.

For a spatial superposition of extent Δx involving many constituents, the decoherence rate is bounded below by the per-constituent rate times $(\Delta x/\lambda)^2$ (with λ the environmental thermal wavelength). For a cat of mass M in a superposition of macroscopic separation, the product gives $\Gamma \propto M(\Delta x)^2$, rendering coherent superpositions of classical objects unobservable.

Spectral density and noise types

Different noise spectra give different scaling of T_2 with measurement time:

- **White noise** (flat $S(\omega)$): exponential decay, well-defined T_2.
- $1/f$ **noise** (ubiquitous in solid-state systems): $|\rho_{01}(t)| \sim e^{-t^2/T_\phi^2} \cdot$ corrections — Gaussian decay, no fixed T_2 but an effective one at a given timescale.
- **Ohmic spectral density** $J(\omega) \propto \omega$ at low frequency: exponential decay at finite temperature, power-law at zero temperature.

The structural consequence: the figure of merit T_2 is a simplification that hides the true nature of the noise. Realistic systems show non-exponential decay, and the effective T_2 depends on which dynamical range is being probed. This is why dynamical decoupling sequences (CPMG, UDD) can *extend* apparent T_2 — they filter out certain frequency ranges of noise that would otherwise contribute.

Worked Example

Superconducting transmon qubit

Typical transmon parameters (as of 2025):

- Qubit frequency $\omega_0/2\pi \approx 5$ GHz
- $T_1 \approx 100\,\mu s$ (limited by dielectric loss and Purcell decay into the readout resonator)
- $T_2 \approx 50\,\mu s$ (limited by flux noise and charge noise at low frequencies)
- Single-qubit gate time $\tau_g \approx 20$ ns
- Two-qubit gate time $\tau_{2g} \approx 200$ ns

Gates per T_2: $T_2/\tau_g \approx 2500$. **Circuits of depth 1000** are possible with some error; deeper circuits require error correction.

Threshold for fault tolerance: Surface code threshold is approximately $p_{\text{th}} \sim 10^{-2}$ for circuit-level noise. This requires $T_2/\tau_g \gtrsim 100$, easily satisfied. The harder challenge is two-qubit gate fidelity, which must be $> 99\%$ and currently hovers at 99.5–99.9% in the best systems.

Macroscopic Schrödinger's cat

Following Zurek's estimate: take a 1 kg cat in a superposition of two positions 10 cm apart, embedded in a room-temperature gas bath. The relevant parameters:

- Molecular scattering cross-section $\sigma \sim 10^{-19}\,\mathrm{m}^2$
- Gas density $n \sim 10^{25}\,\mathrm{m}^{-3}$
- Thermal velocity $v \sim 500\,\mathrm{m/s}$
- Thermal de Broglie wavelength $\lambda \sim 10^{-11}\,\mathrm{m}$
- Superposition extent $\Delta x \sim 10^{-1}\,\mathrm{m}$

Decoherence rate:

$$\Gamma \sim n\sigma v\,\frac{(\Delta x)^2}{\lambda^2} \sim 10^{25} \cdot 10^{-19} \cdot 10^{2} \cdot 10^{20} \approx 10^{28}\,\mathrm{s}^{-1}.$$

This gives $\tau_D \sim 10^{-28}$ seconds. The spatial superposition is coherent for less than the Planck time. It is unobservable in principle, not just in practice.

Nuclear spin qubit

For a ^{31}P nuclear spin in silicon (Kane qubit concept), the relevant coupling is the hyperfine interaction with the bath of ^{29}Si nuclei. Isotopically purified ^{28}Si reduces the bath. Measured values (Tyryshkin et al. 2012):

- $T_1 \sim$ hours (nuclear spin relaxation is very slow — nuclear magnetic moments are tiny)
- $T_2 \sim 30$ seconds (dynamical decoupling extends this further)

These are among the longest coherence times ever measured for solid-state qubits. The structural reason is that nuclear spins couple weakly to everything: the hyperfine coupling is small, the gyromagnetic ratio is $1000\times$ smaller than for electrons, and the relevant noise frequencies are far from thermal fluctuations. This is a design-principle illustration of "minimize coupling + minimize spectral density at relevant frequencies."

Cross-Links

- **Module 5.4 (Decoherence mechanisms):** T_1 and T_2 parameterize the canonical dephasing and amplitude damping channels.
- **Module 5.5 (Pointer basis):** The decoherence time sets the rate at which the reduced state becomes diagonal in the pointer basis.
- **Module 8.2 (Lindblad equation):** The continuous-time Lindblad dynamics have T_1^{-1} and T_2^{-1} as the dissipator rates.
- **VQA:** Coherence time sets the circuit depth budget. Barren plateaus and trainability depend on whether the circuit fits within the coherence window.
- **Module 10.4 (Quantum-to-classical transition):** The structural fact that decoherence is fast explains the emergence of classicality.
- **Quantum sensing:** Metrological sensitivity scales with T_2; longer coherence means better measurements.

Structural Compression

Invariant: Decoherence rate scales with the *squared coupling strength*, the *spectral density of the environment at the relevant frequency*, and (for spatial superpositions) the *squared separation*.

Mechanism:

$$\Gamma \sim \frac{|g|^2}{\hbar^2}\,S_B(\omega) \quad \text{(Fermi's golden rule for a local coupling)},$$

with S_B the bath spectral density and ω the system's characteristic frequency.

Cross-System Echo: Damping rates in classical oscillators follow the same structural form — coupling squared times spectral density at the resonance frequency. The distinctive quantum feature is the $(\Delta x)^2$

scaling for spatial superpositions, which has no classical analogue and is why macroscopic objects are "classical" to overwhelming precision: the decoherence rate is driven catastrophically high by the macroscopic size of any would-be superposition.

Exercises

1. Verify the relation $1/T_2 = 1/(2T_1) + 1/T_\phi$ by computing how the amplitude damping channel affects the off-diagonal element of a qubit density matrix.
2. For a qubit with $T_1 = 100\,\mu s$ and pure dephasing $T_\phi = 70\,\mu s$, compute T_2.
3. Derive the $(\Delta x)^2$ scaling of the decoherence rate for a spatial superposition in a thermal bath, under the assumption that scattering events are localized on the thermal wavelength.
4. For an Ohmic bath with spectral density $J(\omega) = \eta\omega e^{-\omega/\omega_c}$, compute the pure dephasing rate at zero temperature and at finite temperature β.
5. Estimate the decoherence time of a single atom in a magneto-optical trap at room temperature, and compare to a trapped ion in an ultra-high-vacuum chamber.
6. For a transmon qubit with $T_1 = 100\,\mu s$ and gate time 20 ns, compute the gate error rate if all error is coherence-limited, and assess whether fault tolerance thresholds are met.
7. Argue, structurally, why macroscopic objects cannot exhibit coherent quantum behavior under realistic environmental conditions, and quantify the "macroscopicity" at which decoherence becomes faster than any observational timescale.