Lecture 10: Physical Implementations

From Abstract Qubits to Real Hardware

Moody T. Chu

North Carolina State University

Spring 2026

Lecture Overview

1. From Abstract to Physical: Motivation
2. The DiVincenzo Criteria
3. The Harmonic Oscillator and the Leakage Problem
4. Subspace Isolation via Invariant Commutators (Jaynes–Cummings)
5. High-Temperature Mixed-State Ensembles (NMR)
6. Pseudo-Pure State Construction

From Abstract to Physical: Why Implementation Matters

- Lectures 1–9 built a complete, mathematically rigorous theory of quantum computation: gates, circuits, algorithms, and complexity.
- **The central challenge:** every result assumed we can prepare $|0\rangle$, apply any unitary U, and measure in the computational basis. In reality, none of these steps is free.
  - Quantum states decohere due to unwanted environmental interactions.
  - Physical systems have far more than two energy levels — choosing which two to use as $|0\rangle$ and $|1\rangle$ is itself a non-trivial design problem.
  - Measurement disturbs the system; calibrating gates is a precision engineering challenge.
- This lecture asks: **what does a physical system need to satisfy in order to function as a reliable qubit?** We study two concrete approaches: cavity QED (the Jaynes–Cummings model) and nuclear magnetic resonance (NMR).

The DiVincenzo Criteria: Overview

- Any physical system proposed as a quantum computer must satisfy the five **DiVincenzo criteria** (2000), which constitute the minimal necessary conditions for scalable quantum computation.

Criterion	Requirement
I	A scalable physical system with well-characterized qubits
II	The ability to initialize qubits to a fiducial state (e.g., $\|00\cdots0\rangle$)
III	Coherence times much longer than gate operation times
IV	A universal set of quantum gates
V	A qubit-specific measurement capability

- Two additional criteria are sometimes added for distributed computing: the ability to interconvert stationary and flying qubits, and to faithfully transmit flying qubits between locations.

DiVincenzo Criterion I: Scalable Qubits

- **A scalable physical system with well-characterized qubits.**
  - **Two-level system:** we need a reliable, *isolatable* quantum two-level system. Most physical systems have more than two relevant energy levels; the key design challenge is confining dynamics to a chosen pair.
    - **Subspace embedding:** place the computational states $|0\rangle$ and $|1\rangle$ in a 2D subspace $\mathscr{H}_c$ of the full Hilbert space $\mathscr{H}$. The computational subspace must be *invariant* under all operations used in the algorithm, or leakage to other levels will corrupt the computation.
  - **Scalability:** adding more qubits should not exponentially increase the engineering complexity. Each qubit must be addressable independently.
  - **Well-characterized:** the Hamiltonian parameters (frequencies, couplings) must be known precisely enough to calibrate high-fidelity gates.

DiVincenzo Criterion II: State Initialization

- **The ability to initialize the register to a known state.**
  - Algorithms require a definite starting state (typically $|00\cdots0\rangle$).
  - In practice, initialization is achieved by cooling the system to its ground state (cryogenic dilution refrigerators for superconducting qubits), by optical pumping (trapped ions), or by pseudo-pure state preparation at high temperatures (NMR; Section 5).
- **DiVincenzo Criterion III: Long Coherence Times.**
  - Quantum information is fragile. Unwanted coupling to the environment causes decoherence — the irreversible decay of quantum superpositions into classical mixtures.
  - The coherence time T_2 (dephasing) and the relaxation time T_1 (energy decay) must both be much longer than the gate time t_{gate}:

$$T_1,\ T_2 \ \gg\ t_{\text{gate}}.$$

  - The ratio T_2/t_{gate} is the number of gates one can execute before the computation is scrambled by decoherence. Fault-tolerant thresholds require this ratio to be at least $\sim 10^4$.

DiVincenzo Criteria IV and V

- **Criterion IV: Universal Gate Set.**
  - The physical interactions available must be sufficient to approximate any unitary to arbitrary precision.
  - By the universality results of Lecture 6, it suffices to implement arbitrary single-qubit rotations and one entangling two-qubit gate (e.g., CNOT).
  - In practice, this means engineering *controllable couplings* between qubits that can be switched on and off on demand.
- **Criterion V: Qubit-Specific Measurement.**
  - We must be able to measure individual qubits in the computational basis without disturbing the rest of the register.
  - This is technically demanding: a quantum measurement must project onto $|0\rangle$ or $|1\rangle$ with high fidelity, and the result must be amplified to a macroscopic classical signal.
  - **The overarching challenge:** Criteria III and V are in direct tension — strong environmental coupling enables fast, high-fidelity measurement (Criterion V) but also accelerates decoherence (against Criterion III). Every physical implementation must navigate this trade-off.

Why Not Just Use a Harmonic Oscillator?

- The quantum harmonic oscillator is one of the most controllable and well-understood quantum systems available in the laboratory (LC circuits, mechanical resonators, optical cavities).
- Its Hamiltonian and equally-spaced energy levels:

$$\mathscr{H}_{\text{osc}} = \hbar\omega\left(a^\dagger a + \tfrac{1}{2}\right), \qquad E_n = \hbar\omega\left(n + \tfrac{1}{2}\right), \quad n = 0, 1, 2, \ldots \tag{1}$$

- It seems natural to encode $|0\rangle \leftrightarrow |n=0\rangle$ and $|1\rangle \leftrightarrow |n=1\rangle$.
- **The fatal problem:** any external drive (control field) that induces transitions between $|0\rangle$ and $|1\rangle$ will *simultaneously* drive transitions between $|1\rangle$ and $|2\rangle$, $|2\rangle$ and $|3\rangle$, and so on, with the same coupling strength $\sqrt{n+1}$.
- This is the **leakage problem**: amplitude leaks out of the computational subspace $\{|0\rangle, |1\rangle\}$ into higher levels, destroying the qubit structure.

External Control and Tridiagonal Coupling

- An external control field (e.g., a microwave drive) adds a time-dependent perturbation $V(t) = \Omega(t)(a + a^\dagger)$ to the Hamiltonian.
- The key matrix elements are:

$$\langle m | a^\dagger + a | n \rangle = \sqrt{n}\,\delta_{m,n-1} + \sqrt{n+1}\,\delta_{m,n+1}. \tag{2}$$

- Written as a matrix in the $\{|n\rangle\}$ basis, the control operator is tridiagonal:

$$a + a^\dagger = \begin{pmatrix} 0 & \sqrt{1} & 0 & 0 & \cdots \\ \sqrt{1} & 0 & \sqrt{2} & 0 & \cdots \\ 0 & \sqrt{2} & 0 & \sqrt{3} & \cdots \\ 0 & 0 & \sqrt{3} & 0 & \cdots \\ \vdots & & & & \ddots \end{pmatrix} \tag{3}$$

- **Consequence:** any pulse that drives $|0\rangle \leftrightarrow |1\rangle$ (coupling $\sqrt{1}$) also drives $|1\rangle \leftrightarrow |2\rangle$ (coupling $\sqrt{2}$) with even greater strength. There is no way to restrict the dynamics to the computational subspace.

The Impossibility of Subspace Isolation

- **The core difficulty, stated precisely:** DiVincenzo Criterion I demands that the computational subspace $\mathscr{H}_c = \text{span}\{|0\rangle, |1\rangle\}$ be *invariant* under the full time evolution operator e^{-iHt}.
- A subspace $\mathscr{H}_c$ is invariant under e^{-iHt} if and only if it is invariant under H itself, i.e., $H\mathscr{H}_c \subseteq \mathscr{H}_c$.
- For the driven harmonic oscillator, $H|1\rangle$ contains a component along $|2\rangle$ (from the $\sqrt{2}$ matrix element above). Hence $\mathscr{H}_c$ is *not* invariant, and leakage is unavoidable.
- **Two physical solutions:**
  1. Break the equal-spacing of energy levels so the $|0\rangle \leftrightarrow |1\rangle$ transition can be driven *selectively* by a frequency-tuned pulse (the idea behind transmon and superconducting qubits).
  2. Couple the oscillator to a two-level system with a strong natural frequency selectivity — the Jaynes–Cummings model, developed in the next section.

The Jaynes–Cummings Model: Physical Setup

- The Jaynes–Cummings model describes a two-level atom (the qubit) coupled to a single quantized mode of the electromagnetic field (a cavity). It is the foundational model of *cavity quantum electrodynamics* (cavity QED).
- Physical systems: superconducting qubits in microwave cavities, atoms in optical cavities, quantum dots in photonic crystal resonators.
- The full Hilbert space is $\mathscr{H} = \mathscr{H}_{\text{qubit}} \otimes \ell^2(\mathbb{Z}_{\geq 0})$, where the second factor is the infinite-dimensional Fock space of the cavity mode.
- **The key idea:** by designing the coupling between the atom and the cavity, we can arrange for the computational subspace to be *dynamically decoupled* from all higher levels — not by eliminating the coupling, but by exploiting a conserved quantity.

The Jaynes–Cummings Hamiltonian

- The Hamiltonian (in units $\hbar = 1$) is:

$$H = \frac{\omega_0}{2}(\sigma_z \otimes I) + \omega(I \otimes a^\dagger a) + g(\sigma_+ \otimes a + \sigma_- \otimes a^\dagger), \qquad (4)$$

  - ω_0: qubit transition frequency (atom energy splitting).
  - ω: cavity mode frequency.
  - g: atom–cavity coupling strength.
  - $\sigma_\pm = (\sigma_x \pm i\sigma_y)/2$: atomic raising/lowering operators.
  - $a, a^\dagger$: cavity photon annihilation/creation operators.
- The interaction term $g(\sigma_+ \otimes a + \sigma_- \otimes a^\dagger)$ describes energy exchange: the atom absorbs a photon from the cavity ($\sigma_- a^\dagger$, atom de-excites, photon created) or emits one ($\sigma_+ a$, atom excites, photon annihilated).

The Total Excitation Number Operator

- Define the **total excitation number operator**:

$$\hat{N} = \frac{1}{2}(\sigma_z + I) \otimes I + I \otimes a^\dagger a. \tag{5}$$

  $\hat{N}$ counts the total number of excitations: 1 if the atom is excited ($|1\rangle$), plus the number of photons in the cavity.

- **Claim: $[H, \hat{N}] = 0$.** The key brackets are:

$$\begin{aligned}
[\sigma_+ \otimes a,\ \sigma_z \otimes I] &= [\sigma_+, \sigma_z] \otimes a = (-2\sigma_+) \otimes a, \\
[\sigma_+ \otimes a,\ I \otimes a^\dagger a] &= I \otimes [a, a^\dagger a] = I \otimes a = \sigma_+ \otimes a, \\
[\sigma_- \otimes a^\dagger,\ \sigma_z \otimes I] &= [\sigma_-, \sigma_z] \otimes a^\dagger = (2\sigma_-) \otimes a^\dagger, \\
[\sigma_- \otimes a^\dagger,\ I \otimes a^\dagger a] &= I \otimes [a^\dagger, a^\dagger a] = -I \otimes a^\dagger.
\end{aligned}$$

  Collecting all terms from $[H, \hat{N}]$: the $\frac{\omega_0}{2}$ and ω terms commute with $\hat{N}$ trivially (same operators). The interaction terms contribute $\frac{\omega_0}{2}(-2) + \omega(+1)$ from the $\sigma_+ \otimes a$ side and $\frac{\omega_0}{2}(+2) + \omega(-1)$ from $\sigma_- \otimes a^\dagger$ — each pair cancels exactly, confirming $[H, \hat{N}] = 0$.

Block-Diagonal Structure

- Since $[H, \hat{N}] = 0$, the eigenspaces of $\hat{N}$ are *invariant* under the time evolution e^{-iHt}.
- $\hat{N}$ has eigenvalues $n = 0, 1, 2, \ldots$ with eigenspaces:
  - ◇ $n = 0$: only $|\downarrow, 0\rangle$ (atom in ground state, zero photons). **1-dimensional.**
  - ◇ $n \geq 1$: spanned by $|\downarrow, n\rangle$ and $|\uparrow, n-1\rangle$. **2-dimensional.**
- The computational subspace $\mathscr{H}_c = \text{span}\{|\downarrow, 0\rangle, |\uparrow, 0\rangle\}$... but wait: $|\uparrow, 0\rangle$ has $\hat{N} = 1$ while $|\downarrow, 0\rangle$ has $\hat{N} = 0$. They are in *different* invariant subspaces!
- **The resolution:** the computational subspace for cavity QED is taken as $\{n = 0 \text{ sector}\} = \{|\downarrow, 0\rangle\}$ and the $n = 1$ sector $\{|\uparrow, 0\rangle, |\downarrow, 1\rangle\}$, with qubit states encoded as $|0\rangle_L := |\downarrow, 0\rangle$ and $|1\rangle_L := |\uparrow, 0\rangle$. The $n = 1$ sector is a **2D invariant subspace**: H cannot couple it to the $n = 0$ or $n \geq 2$ sectors.

The Restricted Hamiltonian and Time Evolution

- Within the n-th sector (spanned by $|\uparrow, n-1\rangle$ and $|\downarrow, n\rangle$), the Hamiltonian restricts to the 2×2 block:

$$H^{(n)} = \left(n - \tfrac{1}{2}\right) \omega\, I + \frac{\Delta}{2} \sigma_z + g\sqrt{n}\, \sigma_x, \tag{6}$$

  where $\Delta = \omega_0 - \omega$ is the atom–cavity detuning.
- At **resonance** ($\Delta = 0$, $\omega_0 = \omega$), this simplifies to:

$$H^{(n)}\Big|_{\Delta=0} = \left(n - \tfrac{1}{2}\right) \omega\, I + g\sqrt{n}\, \sigma_x. \tag{7}$$

- The time evolution within each sector is exact and analytic:

$$U^{(n)}(t) = e^{-i(n-1/2)\omega t} \begin{pmatrix} \cos\left(g\sqrt{n}\, t\right) & -i \sin\left(g\sqrt{n}\, t\right) \\ -i \sin\left(g\sqrt{n}\, t\right) & \cos\left(g\sqrt{n}\, t\right) \end{pmatrix}. \tag{8}$$

- **The Jaynes–Cummings solution:** the sectors evolve independently. The qubit oscillates between $|\uparrow, n-1\rangle$ and $|\downarrow, n\rangle$ at the *vacuum Rabi frequency* $\Omega_n = 2g\sqrt{n}$ — the characteristic oscillation frequency of the atom–cavity exchange.

Nuclear Magnetic Resonance (NMR) Quantum Computing

- **NMR** exploits the nuclear spin of atoms in a molecule as qubits. In a strong static magnetic field B_0, nuclear spins split into energy eigenstates $|\uparrow\rangle$ and $|\downarrow\rangle$ (the computational basis $|0\rangle$ and $|1\rangle$).
- Gates are implemented by radio-frequency (RF) pulses that resonantly drive spin-flip transitions. Two-qubit gates exploit the natural J-coupling between nuclei on adjacent atoms.
- **The fundamental obstacle:** at room temperature, NMR samples contain $\sim 10^{23}$ molecules in a thermal mixed state, not a pure quantum state.
- The thermal density matrix at inverse temperature $\beta = 1/(k_B T)$ is:

$$\rho_{\text{th}} = \frac{e^{-\beta H}}{\mathcal{Z}}, \quad \mathcal{Z} = \mathrm{Tr}[e^{-\beta H}]. \tag{9}$$

- At room temperature $k_B T \gg \hbar\omega_{\text{NMR}}$, so $\beta H \ll 1$, and the thermal state is nearly completely mixed. Pure-state quantum computing seems impossible.

High-Temperature Approximation

- For an n-qubit NMR system, the Hamiltonian in the external field is $H = -\sum_{j=1}^{n} \frac{\hbar\omega_j}{2}\sigma_z^{(j)}$.
- In the high-temperature regime $k_B T \gg \|\hbar\omega_j\|$, expand to first order:

$$\rho_{\text{th}} = \frac{e^{-\beta H}}{\mathcal{Z}} \approx \frac{1}{2^n}\left(I + \beta\sum_{j=1}^{n}\frac{\hbar\omega_j}{2}\sigma_z^{(j)}\right). \tag{10}$$

- Define the **polarization** per spin $\varepsilon_j = \hbar\omega_j\beta/2 \ll 1$. Then:

$$\rho_{\text{th}} \approx \frac{1}{2^n}\left(I + \sum_{j=1}^{n}\varepsilon_j\,\sigma_z^{(j)}\right) = \frac{I}{2^n} + \frac{1}{2^n}\sum_{j=1}^{n}\varepsilon_j\,\sigma_z^{(j)}. \tag{11}$$

- The state is the maximally mixed state plus a tiny *deviation matrix* proportional to $\varepsilon_j \sim 10^{-5}$ at room temperature.
- **Key insight:** the identity component $I/2^n$ is invariant under any unitary $U\rho U^\dagger = I/2^n$. It contributes nothing to the *signal* (which comes from the deviation matrix alone) but dominates the *noise*.

Pseudo-Pure States: The Idea

- **Definition:** a state of the form
$$\rho_{\text{pp}} = \frac{1-\varepsilon}{2^n} I + \varepsilon \left|\psi\right\rangle \left\langle\psi\right| \tag{12}$$
is called a **pseudo-pure state** (PPS) with polarization ε and effective pure state $|\psi\rangle$.
- **Why this is useful:** any unitary U acts as:
$$U\rho_{\text{pp}}U^\dagger = \frac{1-\varepsilon}{2^n} I + \varepsilon\, U \left|\psi\right\rangle \left\langle\psi\right| U^\dagger. \tag{13}$$
The identity part is unchanged; only the ε-weighted pure component evolves. The NMR signal is proportional to $\varepsilon\langle\psi|U^\dagger \mathcal{O} U|\psi\rangle$, exactly as if we had run the algorithm on the pure state $|\psi\rangle$.
- The NMR measurement cannot distinguish ρ_{pp} from $\left|\psi\right\rangle \left\langle\psi\right|$ — the identity offset is a constant background that cancels in the final observable.
- **The challenge:** how to prepare ρ_{pp} starting from the thermal state ρ_{th}?

Pseudo-Pure State Construction (I)

- Starting from a diagonal thermal state $\rho_{\text{th}} = \text{diag}(\lambda_0, \lambda_1, \ldots, \lambda_{2^n-1})$ with $\lambda_0 \geq \lambda_1 \geq \cdots$, we wish to construct $\rho_{\text{pp}} \propto (c \,|0\cdots 0\rangle \langle 0 \cdots 0| + I)$ for some $c > 0$.
- **Temporal averaging method (for $2^n = 4$):** apply three unitary operations $\{I, P_1, P_2\}$ (cyclic permutations of the non-ground populations) and average:

$$\bar{\rho} = \frac{1}{3}\left(\rho + P_1 \rho P_1^\dagger + P_2 \rho P_2^\dagger\right). \tag{14}$$

- With $\rho_{\text{th}} = \text{diag}(a, b, c, d)$, cyclic permutations $P_1 : (b, c, d) \to (c, d, b)$ and $P_2 : (b, c, d) \to (d, b, c)$ give:

$$\bar{\rho} = \text{diag}\left(a,\ \tfrac{b+c+d}{3},\ \tfrac{b+c+d}{3},\ \tfrac{b+c+d}{3}\right). \tag{15}$$

Pseudo-Pure State Construction (II)

- Setting $s = b + c + d$, the averaged state is:

$$\bar{\rho} = \frac{s}{4} I + \left(a - \frac{s}{3}\right) |00\rangle \langle 00| \,, \tag{16}$$

  which is a pseudo-pure state with effective pure state $|00\rangle$ and polarization $a - s/3 > 0$ (since $\lambda_0 = a$ is the largest eigenvalue).
- **Verification (numerical):** with $(a, b, c, d) = (0.4, 0.2, 0.2, 0.2)$, the averaged state has diagonal $(1.2,\ 0.6,\ 0.6,\ 0.6)/3$, and the formula gives $\frac{s}{4} I + (a - s/3) |00\rangle \langle 00|$ with $s = 0.6$, $a - s/3 = 0.2$ — matching exactly.
- This construction requires only classical control (unitary permutations, averaging over repeated experiments) and works for any 2^n-dimensional diagonal ρ_{th} by choosing $n(2^n - 1)/2$ permutations.

Scalability Limitations of NMR

- **Polarization shrinks exponentially:** for n qubits, $\varepsilon \approx n\hbar\omega/(2^n k_B T) \sim n/2^n \to 0$. Maintaining SNR requires averaging exponentially many experiments, capping NMR at small n.
- **The entanglement question:** ρ_{pp} is separable (non-entangled) for any $\varepsilon < 1/(1+2^n)$. Room-temperature polarizations always satisfy this bound, so NMR quantum computers never generate true quantum entanglement and cannot provide genuine quantum speedup.
  - ◇ NMR demonstrations were valuable proof-of-concept implementations, but not a viable path to large-scale quantum advantage.
- Current platforms (superconducting circuits, trapped ions, photonics) achieve genuine entanglement at millikelvin temperatures and do not suffer from this exponential polarization penalty.

Summary of Lecture 10 (I)

- **DiVincenzo criteria:** five necessary conditions for any physical quantum computing platform — scalable qubits, state initialization, coherence time $\gg$ gate time, universal gates, and qubit-specific measurement.
- **Leakage problem:** a harmonic oscillator's equally-spaced levels make the computational subspace non-invariant under any control field; leakage to higher levels is unavoidable.
- **Jaynes–Cummings:** coupling the oscillator to a two-level atom gives a conserved excitation number $\hat{N}$ with $[H, \hat{N}] = 0$, producing invariant 2D sectors with exact analytic dynamics (vacuum Rabi oscillations at $2g\sqrt{n}$).

Summary of Lecture 10 (II)

- **NMR high-temperature approximation:** $\rho_{\text{th}} \approx I/2^n + \varepsilon \sum_j \sigma_z^{(j)}/2^n$; the tiny deviation $\varepsilon \sim 10^{-5}$ behaves as a pure state under unitary evolution and carries the entire observable signal.
- **Pseudo-pure states:** constructed by temporal averaging over cyclic permutations of non-ground populations. Polarization decays as $\varepsilon \sim n/2^n$, making NMR fundamentally unscalable beyond small qubit numbers.

Looking Ahead: Lecture 11

- This lecture grounded abstract qubits in physical reality. Lecture 11 turns to the practical consequences of imperfection.
- **Quantum Error Correction:**
  - Can quantum information be protected against decoherence and gate errors without violating the no-cloning theorem?
  - The three-qubit bit-flip code and the Shor code as the first constructive answers.
  - **Stabilizer formalism:** a unified algebraic framework for describing and constructing error-correcting codes efficiently.
- The bridge between this lecture and the next: the coherence time T_2 sets a clock; error correction is the engineering discipline that extends the effective coherence time of a logical qubit far beyond the physical coherence time of any individual qubit.