Gate Errors and Coherence Times

Variational Quantum Algorithms · Module 9 — Hardware Noise Models

Node 9.3

Position in Knowledge Graph

System: Variational Quantum Algorithms **Structural Domain:** Hardware Noise Models **Node:** 9.3

Node 9.2 treated depolarizing noise as the canonical model. This node gets **physical**: it explains where noise actually comes from in real hardware, what "gate error" and "coherence time" mean operationally, and how the abstract quantum channels of 9.1 connect to the physical parameters T_1, T_2, gate duration, and fidelity that appear on hardware spec sheets.

The goal: after this node, you can read a hardware specification (e.g., "IBM Heron r1 processor, 133 qubits, median T_1 150 µs, median T_2 80 µs, median 2Q fidelity 99.5%") and know exactly what it implies for VQA depth, shot budget, and expected performance.

This is the **bridge between theory and practice** — the point where abstract channels become concrete numbers.

The Sources Of Noise

Real quantum hardware has several distinct noise sources, each with its own physics and its own channel representation.

1. Energy relaxation (amplitude damping)

Physics: the excited state $|1\rangle$ can decay to the ground state $|0\rangle$ via spontaneous emission (superconducting), photon scattering (ions, atoms), or other dissipative mechanisms. The rate is characterized by the **energy relaxation time** T_1.

Channel. The amplitude damping channel at rate $\gamma = 1 - e^{-t/T_1}$ (where t is the elapsed time):

$$\mathcal{A}_\gamma(\rho) = E_0 \rho E_0^\dagger + E_1 \rho E_1^\dagger,$$

with

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

Key properties: - **Non-unital:** amplitude damping moves states *toward* $|0\rangle$, not toward $I/2$. There is a preferred direction. - **Breaks symmetry:** a state initialized to $|1\rangle$ decays, but $|0\rangle$ does not. Noise depends on the state. - **Cumulative:** at time t, the channel parameter is $\gamma(t) = 1 - e^{-t/T_1}$.

2. Dephasing (phase damping)

Physics: the coherence between $|0\rangle$ and $|1\rangle$ is destroyed by fluctuations in the environment (charge noise, flux noise, magnetic field variations). The rate is T_ϕ, and combines with T_1 to give the **total coherence time** T_2:

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

Channel. The phase damping channel at rate $\lambda = 1 - e^{-t/T_\phi}$:

$$\mathcal{P}_\lambda(\rho) = \begin{pmatrix} \rho_{00} & \sqrt{1-\lambda}\rho_{01} \\ \sqrt{1-\lambda}\rho_{10} & \rho_{11} \end{pmatrix}.$$

Key properties: - **Unital:** it preserves $I/2$. - **Diagonal-preserving:** computational basis populations are unchanged. - **Eigenvalue structure:** exponential decay of off-diagonal elements (coherences).

3. Coherent errors

Physics: miscalibrations in the control pulses cause systematic over- or under-rotations. Not random — deterministic but unknown.

Channel (example: over-rotation). $U_{\text{actual}} = R_Y(\theta + \delta)$ instead of $R_Y(\theta)$. No probability mixing — just a wrong unitary.

Key properties: - **Unitary:** still a valid quantum operation. - **Coherent accumulation:** many small errors can add up to a large error if they are correlated. - **Not captured by Pauli channels:** the standard depolarizing model misses coherent errors, which is a frequent source of theory-experiment discrepancy.

4. Crosstalk

Physics: a gate on qubit A affects qubit B due to unintended couplings. Purely a multi-qubit effect.

Channel. The ideal single-qubit gate U_A is replaced by a correlated two-qubit channel that depends on the state of qubit B. Hard to characterize, hard to model.

Key properties: - **Non-local:** creates correlations between distant errors. - **Context-dependent:** the effect depends on what other gates are happening in parallel. - **One of the dominant error sources** in 2026 superconducting hardware.

5. Measurement errors

Physics: the measurement operation itself is imperfect. A qubit in state $|0\rangle$ may be reported as "1" with some probability, and vice versa.

Channel. Modeled as a classical confusion matrix applied to the measurement outcome:

$$\begin{pmatrix} P(0_{\text{reported}}|0_{\text{true}}) & P(0_{\text{reported}}|1_{\text{true}}) \\ P(1_{\text{reported}}|0_{\text{true}}) & P(1_{\text{reported}}|1_{\text{true}}) \end{pmatrix}.$$

Typical values: 1-5% error rate for current superconducting hardware, $\sim 0.5\%$ for trapped ions.

Key properties: - **Classical:** can be corrected by post-processing if the confusion matrix is known. - **Asymmetric:** typically different rates for $0 \to 1$ and $1 \to 0$ errors. - **Dominates at shallow depth:** for very shallow circuits, measurement error can be the largest contribution.

The Relaxation-Time Parameters

The headline numbers for any piece of quantum hardware:

T_1 **(energy relaxation time):** the characteristic timescale for the excited state to decay. Typical values: - Superconducting (transmon): 50-300 µs. - Trapped ion: seconds. - Neutral atom (Rydberg): hundreds of µs. - Photonic: effectively infinite (no relaxation). - Silicon spin: milliseconds.

T_2 **(total dephasing time):** the characteristic timescale for coherences to decay, always $\leq 2T_1$. Typical values: - Superconducting: 20-200 µs (often $T_2 \ll 2T_1$ because of dephasing). - Trapped ion: hundreds of ms to seconds. - Neutral atom: 0.5-5 ms. - Photonic: effectively infinite. - Silicon spin: hundreds of µs (isotope-enriched).

T_2^* **(pure dephasing time without echo):** the dephasing time measured by a Ramsey experiment (no echo pulse). Usually shorter than T_2.

Gate time: the duration of a gate operation. Typical values: - Superconducting single-qubit: 20-50 ns. - Superconducting two-qubit: 100-500 ns. - Trapped ion: 10-500 µs. - Neutral atom: 100 ns-1 µs.

The gate error budget: a gate of duration τ on a qubit with coherence time T_2 has coherence-limited error $\sim \tau/T_2$. For superconducting hardware with $\tau = 300$ ns and $T_2 = 100$ µs, this is 3×10^{-3} — a ~0.3% error floor from coherence alone. Actual two-qubit errors are usually 2-3x larger due to calibration and other non-idealities.

The Circuit Time Budget

For a circuit with L gates of duration τ, the total circuit time is $L\tau$. For the circuit to be useful, this must be much less than T_2:

$$L\tau \ll T_2 \quad \Longleftrightarrow \quad L \ll T_2/\tau.$$

Examples:

- Superconducting ($T_2 = 100$ µs, $\tau = 300$ ns): $L \ll 333$ gates. Usable depth ~ 100 gates.
- Trapped ions ($T_2 = 1$ s, $\tau = 100$ µs): $L \ll 10^4$ gates. Usable depth $\sim 1000 - 3000$ gates.
- Neutral atoms ($T_2 = 2$ ms, $\tau = 500$ ns): $L \ll 4000$ gates.

Observation: the gate-count-per-coherence-time is surprisingly similar across technologies in the low hundreds. The different technologies have different absolute speeds but similar "relative" circuit budgets.

For VQAs: the usable depth $\sim 100 - 1000$ gates is consistent with the depolarizing analysis of 9.2. Real depth is more like $\sim 10 - 50$ layers of an HEA for 20-30 qubits — limited by both coherence and noise-induced plateaus.

Fidelity As An Operational Metric

Gate fidelity is the headline metric, but it can mean several things:

1. Average gate fidelity $\bar{F}$: the average over Haar-random input states of the fidelity between the ideal and actual output. The number reported on spec sheets.

2. Process fidelity F_{proc}: the overlap between the ideal and actual process, normalized. Related to average fidelity by $\bar{F} = (dF_{\text{proc}} + 1)/(d + 1)$.

3. Entanglement fidelity: for two-qubit gates, the fidelity when the input is a maximally entangled state. Sometimes more informative.

4. Benchmarking fidelity: what **randomized benchmarking** reports. This is a specific operational protocol that effectively measures the twirled depolarizing rate, which is close to but not identical with the average gate fidelity.

For most practical purposes, **the randomized benchmarking fidelity is what hardware vendors report**, and that's what you should use in your depth-fidelity calculations. It's the depolarizing-equivalent fidelity after Pauli twirling.

The Amplitude Damping Vs Depolarizing Comparison

A practically important comparison: how different is amplitude damping from depolarizing?

At low rates ($\gamma \ll 1$): very similar. Both are approximately "error with probability $\sim \gamma$," and the detailed structure doesn't matter much for shallow circuits.

At higher rates: amplitude damping is *preferentially* toward $|0\rangle$, so it biases expectation values of Z toward $+1$ (the $|0\rangle$ eigenvalue). Depolarizing biases all Pauli expectations toward zero symmetrically.

For VQAs: if your cost function has significant Z character, amplitude damping will shift the cost systematically (upward for Z with positive coefficient). The optimizer will find a *different* minimum than the ideal one. This is a real practical issue with superconducting hardware.

Quantitative comparison: at average gate fidelity 99.5%, depolarizing noise shifts $\langle Z\rangle$ to $0.995\langle Z\rangle$; amplitude damping shifts $\langle Z\rangle$ to $\langle Z\rangle + (\text{const}) \cdot (1 - \text{rate})$ — a state-dependent additive shift.

Upshot: for theoretical analysis, assume depolarizing for clean bounds. For real hardware, use amplitude damping + dephasing. For quick practical estimates, depolarizing is usually accurate enough.

The Coherence-Limited Regime

A useful conceptual division: a VQA circuit is either **coherence-limited** or **gate-error-limited**.

Coherence-limited: the total circuit time is close to T_2, so the whole state decays before the circuit finishes. Further reductions in per-gate error don't help — the state is already decohering. This is the regime for ions (where individual gate fidelities are excellent but circuits are slow).

Gate-error-limited: the circuit is fast compared to T_2, but individual gates have non-trivial error rates. The total error is dominated by summing per-gate errors. This is the regime for superconducting qubits (fast gates, shorter coherence).

For VQA design:

- In the coherence-limited regime, **reducing gate count** is the right strategy. Every gate fewer is better because it saves circuit time.
- In the gate-error-limited regime, **reducing per-gate error** is the right strategy. Every calibration improvement compounds.

Most 2026 superconducting hardware is gate-error-limited for shallow circuits and coherence-limited for deep ones. The crossover depth depends on the specific machine.

Calibration And Drift

Real hardware parameters **drift** on timescales of minutes to hours:

- **Qubit frequencies** drift by tens of kHz over hours.
- **Gate fidelities** drift by 10-20% (relative) over a day.

- T_1, T_2 can vary 2x over a week.

Hardware providers run periodic **recalibration** (typically every 1-4 hours) to keep parameters within spec. Between calibrations, the noise model is slowly changing.

For VQAs: if a VQA training run takes 12 hours, the noise model at the end of the run is different from the model at the beginning. A static noise-aware compilation won't be valid throughout. The right response is **adaptive noise characterization** — periodically re-measuring the noise parameters and updating the compilation or the optimizer accordingly.

This is one of the practical deployment challenges that Module 10 (Adaptive Control Systems) addresses.

Structural Role

This node is the **physical-hardware** node of Module 9. It connects the abstract noise channels of 9.1-9.2 to the concrete parameters (T_1, T_2, gate time, fidelity) that appear on hardware spec sheets, and gives the practitioner the information needed to translate between hardware specs and VQA performance estimates.

It is the prerequisite for nodes 9.4 (mitigation) and 9.6 (noise-adaptive ansatz design), which build on the physical understanding.

Relations to Other Concepts

- **Bloch (1946), Redfield (1957)** — foundational work on relaxation theory.
- **Randomized benchmarking (Magesan et al. 2011)** — the standard fidelity measurement.
- **Quantum process tomography** — full characterization of a gate.
- **Module 9.1 (Quantum noise channels)** — the mathematical framework.
- **Module 10 (Adaptive control)** — handling drift via adaptation.

Worked Example — Noise Budget For A 15-Qubit QAOA

Target hardware: superconducting, $T_1 = 150$ µs, $T_2 = 80$ µs, single-qubit gate time 30 ns, two-qubit gate time 300 ns, single-qubit fidelity 99.9%, two-qubit fidelity 99.5%.

Circuit: 15-qubit QAOA at depth $p = 3$. Each layer has 15 single-qubit gates + 14 two-qubit gates = 29 gates. Three layers = 87 gates, of which 42 are two-qubit.

Time budget: 87 gates times average time ~ 200 ns ≈ 17 µs. **Well within $T_2 = 80$ µs.** Coherence is not the bottleneck.

Fidelity budget: - Single-qubit error contribution: $45 \cdot 10^{-3} = 0.045$. - Two-qubit error contribution: $42 \cdot 5 \times 10^{-3} = 0.21$. - Total depolarizing error: ≈ 0.26. Circuit fidelity ≈ 0.74.

Measurement error: typically 2% per qubit, so 15 qubits of measurement error contribute another ~ 0.3 to the total error.

Net: about 40-50% of the signal survives. Shots needed to resolve signal above shot noise: for a cost that is ~ 0.5 in magnitude on a good run and 0.25 after noise, shot noise of 10^{-3} (i.e., 10^6 shots) is comfortably below the signal.

Conclusion: this QAOA is workable. For depth 6, we'd double the gate count and the signal fidelity would drop to ~ 0.3. At depth 10, ~ 0.1 — workable but getting marginal. At depth 20, signal is lost.

This is the kind of calculation every experimentalist does before running a VQA.

Visualization

Pending. A chart showing, for several hardware platforms (superconducting, trapped ion, neutral atom, silicon), the typical T_1, T_2, gate time, and gate error. Overlaid: the usable circuit depth for each platform. Caption: "Hardware comes in flavors. Each flavor has a different depth budget."

Exercises

1. A qubit has $T_1 = 100$ µs, $T_2 = 60$ µs. What is the pure dephasing time T_ϕ?
2. For a superconducting two-qubit gate of duration 400 ns and fidelity 99.2%, what fraction of the error is coherence-limited vs gate-error-limited?
3. (VQA) For a 30-qubit VQE on trapped-ion hardware ($T_2 = 1$ s, two-qubit gate time 300 µs, fidelity 99.9%), estimate the maximum circuit depth.

Research Extensions

- **Time-dependent noise models** — capturing drift between calibrations.
- **Crosstalk characterization** — systematic methods for identifying and mitigating crosstalk.
- **Coherent-error-aware compilation** — reducing coherent errors via gate-sequence design.
- **Noise tomography** — full process characterization of hardware gates.

Cross-links to Other Courses

- **Open Quantum Systems** — Lindblad master equations.
- **Quantum Measurement Theory** — measurement errors and corrections.
- **Hardware Physics (superconducting/ion/atom)** — the underlying physics of specific platforms.
- **Module 9.4 (Noise mitigation)** — the response to the errors characterized here.

Variational Quantum Algorithms · Module 9 — Hardware Noise Models · Node 9.3
Concepts: quantum-noise, decoherence, quantum-channels
Prerequisites: 9.1, 9.2
Enables: 9.4, 9.8

hbar.university — own.your.cognition