



Noise-dependent limits of dynamical decoupling: fidelity and variance dynamics in single-qubit systems

Kim E

Submitted: April 2, 2026, Revised: version 1, April 22, 2026, version 2, April 23, 2026, version 3, April 24, 2026

Accepted: April 25, 2026

Abstract

Decoherence is a primary barrier to practical quantum computation, causing quantum states to degrade through phase randomization and irreversible energy relaxation. This study presents a simulation-based evaluation of echo-based dynamical decoupling (DD) pulse sequences as a software-level mitigation strategy in single-qubit systems. Using the Qiskit Aer density matrix simulator, four strategies—baseline, Hahn Echo, Carr-Purcell-Meiboom-Gill (CPMG), and Mixed Synchronization—were evaluated across low-frequency phase noise, high-frequency phase noise, and combined $T_1 + T_2$ noise regimes. Hahn Echo restored mean fidelity to 100.00% under low-frequency noise (baseline: 57.90%), while CPMG achieved 97.62% under high-frequency noise (baseline: 2.45%). Under combined noise, all strategies showed limited improvement, with Mixed-Sync reaching 65.34% versus 60.95%. Coefficient of variation (CV) analysis across 500 Monte Carlo trials showed that CPMG reduced purity-decay variance threefold under high-frequency noise, with bootstrap confidence intervals confirming statistical significance. Time-resolved analysis showed that CPMG suppressed variance growth throughout evolution, reducing variance growth from +267% (baseline) to +89%, while baseline and Hahn Echo exhibited progressive amplification of variance. These results show that dynamical decoupling influences the temporal structure of noise, not just mean fidelity. In T_2 -dominated regimes, pulse sequences constrain variability; in mixed regimes, they partially shape it; and in T_1 -dominated regimes, both mean performance and variability converge to hardware-limited behavior. DD effectively corrects reversible phase errors but cannot overcome irreversible T_1 loss, with implications for NISQ circuit design.

Keywords

Dynamical decoupling, Decoherence mitigation, Quantum noise, Qubit fidelity, Purity decay, Variance analysis, Monte Carlo simulation, Bootstrap confidence intervals, NISQ quantum computing, Error mitigation

Eli Kim, Urbana High School, 1002 S Race St, Urbana, IL 61801, USA. eli.kd19@gmail.com

1. Introduction

Quantum computers exploit superposition and entanglement to solve problems that are computationally intractable for classical machines. A qubit can exist as a coherent combination of basis states $|0\rangle$ and $|1\rangle$, enabling algorithms such as Shor's (1) integer factorization, Grover's (2) unstructured search, and quantum chemistry simulations relevant to pharmaceutical discovery and materials science (3). These applications require sustaining quantum coherence across many gate operations — a requirement that decoherence fundamentally threatens.

Decoherence arises from unavoidable qubit-environment coupling. Two distinct mechanisms dominate. Phase decoherence (T_2 dephasing) randomizes the relative phase between $|0\rangle$ and $|1\rangle$ without changing the qubit's energy, manifesting as equatorial dispersion on the Bloch sphere. Energy relaxation (T_1 decay) drives the qubit irreversibly toward the ground state through spontaneous emission. These mechanisms obey the constraint $T_2 \leq 2T_1$ (4). The critical distinction is reversibility: T_2 errors can in principle be undone by refocusing pulses, while T_1 loss cannot. Figure 1 illustrates both processes on the Bloch sphere.




Figure 1. Bloch sphere illustrating T_1 (blue, energy relaxation toward $|1\rangle$) and T_2 (red, equatorial phase spreading). T_1 decay is irreversible; T_2 dephasing is reversible and is the target of echo-based correction.

Dynamical decoupling (DD) exploits the reversibility of phase errors by applying sequences of precisely timed π pulses during idle circuit intervals, effectively time-averaging the qubit-environment interaction and refocusing accumulated phase (5). The concept originates in nuclear magnetic resonance (NMR), where Hahn (6) first demonstrated that a single refocusing pulse could reverse static

dephasing. The Carr-Purcell-Meiboom-Gill (CPMG) sequence later extended this to time-varying noise through rapid pulse repetition (7). Khodjasteh and Lidar (8) formally established fault-tolerant DD conditions, showing that pulse sequences can suppress decoherence to arbitrary order when the interpulse interval is shorter than the environmental correlation time.

Understanding the Hahn Echo mechanism in detail is essential for interpreting all subsequent results. Figure 2 shows the six-stage Bloch sphere evolution of a qubit undergoing Hahn Echo: initialization to $|0\rangle$ (①), Hadamard rotation to the equatorial superposition (②), free-precession phase accumulation (③), π -pulse application that mirrors the Bloch vector (④), reversed free-precession that cancels the accumulated phase (⑤), and final state recovery (⑥).




Figure 2. Six-stage Bloch sphere evolution during Hahn Echo. States ①–⑥ show: $|0\rangle$ initialization, Hadamard superposition, free-precession phase accumulation ($e^{i\phi}$), π -pulse reflection, reversed precession ($e^{-i\phi}$), and full refocusing to $|0\rangle$. The π pulse mirrors the phase error so that the second free-evolution interval precisely cancels the first, restoring coherence.

The practical relevance of decoherence (9,10). In the Noisy Intermediate-Scale Quantum (NISQ) era, where full fault-tolerant error correction remains out of reach (11), DD represents one of the most accessible error mitigation strategies available. The present study investigated four DD strategies across

three noise regimes, quantifying not only mean fidelity and purity but also the variance of purity decay — a robustness metric that prior simulation work has largely overlooked.

2. Methods

2.1 Simulation framework

All simulations were implemented in Python 3.10 using Qiskit (version 0.44) and the Qiskit Aer density matrix simulator (12). Density matrix simulation was chosen because it fully represents both pure and mixed quantum states, enabling computation of purity $\text{Tr}(\rho^2)$ in addition to fidelity. For a pure state $|\psi\rangle$, $\text{Tr}(\rho^2) = 1$; for a mixed state, $\text{Tr}(\rho^2) < 1$, with the deficit quantifying irreversible coherence loss (13).

2.2 Assumptions

Four key assumptions were made explicitly. First, noise was treated as Markovian — the environment was assumed memoryless, excluding the non-Markovian temporal correlations that characterize $1/f$ flux noise in superconducting qubits (14). Second, all π pulses were treated as instantaneous and ideal; finite pulse duration and calibration errors were excluded. Third, noise parameters were held constant across trials, excluding hardware drift. Fourth, simulations were restricted to a single qubit, thus excluding multi-qubit crosstalk. These assumptions isolated the theoretical performance of each strategy in a controlled setting.

2.3 Qubit initialization and noise models

Each simulation began with the qubit initialized to $|0\rangle$, then rotated to the maximally sensitive superposition $|+\rangle = (|0\rangle + |1\rangle) / \sqrt{2}$ by a Hadamard gate. Phase noise (T_2) was modeled as a stochastic Z-rotation with magnitude drawn from a Gaussian distribution (σ -parameterized), applied during each idle time step. Two phase noise regimes were evaluated: low-frequency (quasi-static, representing flux noise in superconducting circuits) and high-frequency (rapidly fluctuating, representing charge noise and photon shot noise). Energy relaxation (T_1) was modeled using an amplitude damping channel with decay probability $\gamma = 1 - \exp(-\Delta t / T_1)$ per time step (4). Combined simulations applied both noise channels simultaneously.

2.4 Circuit strategies

Four strategies were evaluated. Figure 3 shows the Baseline circuit (no DD), which served as the reference condition. A Hadamard gate initialized the superposition, idle steps accumulated noise unimpeded, and a final Hadamard converted back for measurement.

Figure 4 shows the Hahn Echo strategy, which inserted a single π pulse (X gate) at the midpoint of the idle interval. As illustrated in Figure 2, this single pulse refocuses static and low-frequency phase errors by time-reversing phase accumulation. Hahn Echo is maximally effective when noise is quasi-static relative to the total idle time (6).

Strategy: Baseline




Figure 3. Baseline circuit: Hadamard initialization, unprotected idle interval ($\times 30$ repetitions), and final Hadamard readout. No refocusing pulses are applied; phase errors accumulate freely.

Strategy: Hahn (1-X)




Figure 4. Hahn Echo circuit: a single X-gate (π pulse) applied at the midpoint of the idle interval. The refocusing pulse reverses phase accumulated in the first half, canceling it during the second half.

Figure 5 shows the CPMG (Dense) strategy, rapidly fluctuating phase errors that a single Hahn Echo pulse cannot track. CPMG is particularly effective against broadband and high-frequency noise (7). By placing pulses at a repetition rate exceeding the dominant noise frequency, CPMG effectively time-averages

Strategy: CPMG (Dense)




Figure 5. CPMG circuit with $N = 8$ equally spaced X-gate (π) pulses within each idle block ($\times 15$ repetitions). Dense pulse placement suppresses high-frequency phase fluctuations that a single Hahn Echo pulse cannot refocus.

Strategy: Mixed-Sync




Figure 6. Mixed Synchronization circuit with geometrically spaced X pulses within each block ($\times 5$ repetitions). Non-uniform spacing targets both low- and high-frequency noise simultaneously, trading peak performance for spectral robustness.

Figure 6 shows the Mixed Synchronization Strategy, which employed non-uniformly spaced π pulses arranged in a geometric progression within each repetition block. This variable spacing was designed to simultaneously address low- and high-frequency noise components, providing broader spectral coverage than either Hahn Echo or CPMG alone (8).

2.5 Performance metrics and statistical analysis

State fidelity was defined as $F = \langle +|\rho|+ \rangle$, ranging from 0 to 1. Purity was defined as $P = \text{Tr}(\rho^2)$, ranging from 0.5 (maximally mixed) to 1 (pure). Each condition was evaluated across 500 independent Monte Carlo trials. Mean, standard deviation (Std), and standard error of the mean (SEM) were computed across all trials; 95% confidence intervals were constructed using the normal approximation. In addition, the coefficient of variation of purity decay ($CV = \sigma_p / \bar{P}$) was computed for each strategy and noise condition, providing a normalized robustness metric that captured trial-to-trial consistency independently of absolute purity level.

2.6 Bootstrap analysis of variance

To establish statistical significance of variance differences between strategies, bootstrap resampling ($n = 10,000$ iterations) was performed to construct 95% confidence intervals for purity variance. Non-overlapping confidence intervals provide strong evidence that observed variance differences are robust rather than artifacts of sampling variability. This analysis was applied to the high-frequency and combined noise regimes where variance differences were the most pronounced.

2.7 Time-resolved variance analysis

To characterize the temporal evolution of stochastic variability, purity variance and coefficient of variation were computed at each time step during the decoherence process for the high-frequency noise regime. This analysis distinguished whether DD strategies suppressed variance primarily during early-time evolution, late-time evolution, or throughout the entire process, providing mechanistic insight beyond aggregate endpoint statistics.

3. Results

3.1 Low-frequency phase noise

Under low-frequency phase noise, the unprotected baseline exhibited a mean fidelity of 57.90% (Std = 0.48%, SEM = 0.05%, 95% CI [57.80, 58.00]). The Hahn Echo sequence achieved perfect mean fidelity of 100.00% (Std = 0.00%, 95% CI [100.00, 100.00]), representing an improvement of +42.10 percentage points. Purity under Hahn Echo remained at 1.00, confirming complete phase error refocusing with no residual mixing. This result is consistent with the Hahn Echo mechanism illustrated in Figure 2: under quasi-static noise, the phase accumulated during the first idle half was precisely negated during the second, restoring the ideal state with zero variance across all 500 trials.

3.2 High-frequency phase noise

Under high-frequency phase noise, the unprotected baseline collapsed to a mean fidelity of 2.45% (Std = 0.38%, SEM = 0.04%, 95% CI [2.37, 2.53]), indicative of near complete loss of coherence. CPMG recovered mean fidelity to 97.62% (Std = 0.45%, SEM = 0.05%, 95% CI [97.53, 97.71]), an improvement of +95.17 percentage points. Hahn Echo was

ineffective in this regime, as rapidly varying phase fluctuations between its two pulses could not be refocused by a single midpoint inversion. The CV of purity was 0.18 for the baseline control, 0.13 for Hahn Echo, and 0.06 for CPMG under high-frequency phase noise, demonstrating that the denser CPMG pulse sequence not only improved mean performance but had also reduced trial-to-trial variability by a factor of three.

3.3 Combined $T_1 + T_2$ noise

Under combined noise, the baseline exhibited a mean fidelity of 60.95% (Std = 0.39%, SEM = 0.04%, 95% CI [60.87, 61.03]). The best-performing strategy, Mixed-Sync, reached 65.34% (Std = 1.43%, SEM = 0.14%, 95% CI [65.06, 65.62]), an improvement of +4.39 percentage points. Critically, purity declined substantially across all strategies under

combined noise, and CV converged to between 0.07 and 0.09 regardless of strategy. This convergence established that T_1 amplitude damping, not phase noise, dominated stochastic variability under combined conditions. The dissociation between fidelity and purity in this regime reflected a T_1 -induced ground-state bias that partially mimicked the target state projector, inflating fidelity while true coherence was irreversibly lost.

3.4 Statistical significance of variance reduction

Bootstrap confidence interval analysis (10,000 iterations) confirmed that the observed variance reductions were statistically robust. Under high-frequency phase noise, the 95% bootstrap CIs for purity variance were completely non-overlapping across strategies (Table 1).

Table 1. Variance across strategies using high-frequency phase noise.

Strategy	Variance	95% CI Lower	95% CI Upper
Baseline	0.0324	0.0298	0.0351
Hahn Echo	0.0171	0.0152	0.0191
CPMG	0.0036	0.0030	0.0043
Mixed-Sync	0.0089	0.0077	0.0102

The complete separation of confidence intervals confirmed that CPMG's threefold variance reduction was not a sampling artifact but a robust consequence of dense pulse placement. Figure 7 presents these bootstrap CIs, showing that CPMG achieved significantly tighter variance bounds than both baseline and Hahn Echo strategies.

Under combined $T_1 + T_2$ noise, bootstrap CIs for baseline and Mixed-Sync strategies overlapped substantially, confirming that variance convergence across strategies reflects a fundamental T_1 -dominated regime where pulse strategy has minimal impact on stochastic variability. This convergence is consistent with the theoretical prediction that amplitude damping introduces irreducible shot noise that cannot be suppressed by coherent pulse sequences.




Figure 7. Bootstrap confidence intervals (95%, $n = 10,000$) for purity variance under high-frequency phase noise (left) and combined $T_1 + T_2$ noise (right). Non-overlapping CIs in the high-frequency regime confirm statistically significant variance reduction by CPMG. Under combined noise, CIs overlap substantially, confirming convergence driven by irreversible T_1 amplitude damping. Error bars show $[\text{variance} - CI_{\text{lower}}, CI_{\text{upper}} - \text{variance}]$.

3.5 Temporal structure of variance suppression
Time-resolved analysis of purity variance evolution during decoherence revealed distinct temporal signatures for each strategy under high-frequency phase noise (Figure 8). The baseline strategy exhibited progressive variance amplification, with variance increasing by +267% from early time ($t = 0.17$) to late time ($t = 0.83$). Hahn Echo showed intermediate behavior with +153% variance growth. In contrast, CPMG maintained consistently low variance throughout evolution, with +89% increase—demonstrating that dense pulse placement suppressed variance accumulation at all time scales, not merely at the final measurement.

Panel D of Figure 8 shows that CPMG's early-time variance was already $4.2\times$ lower than baseline, and this advantage was maintained through late times ($3.8\times$ reduction). This temporal consistency has important implications for multi-qubit circuits: CPMG provides reliable variance suppression regardless of when intermediate measurements or gates are applied during idle intervals, whereas baseline and single-pulse strategies exhibit time-dependent vulnerability to variance suppression.

Figures 9 and 10 and Tables 2 through 4 present the complete results.




Figure 8. Time-resolved purity variance analysis under high-frequency phase noise. (A) Mean purity decay with $\pm 1\sigma$ bands showing CPMG's sustained coherence. (B) Variance evolution showing CPMG maintains low variance throughout, while baseline exhibits progressive amplification. (C) Coefficient of variation (CV) evolution demonstrating normalized variability trends. (D) Early-time ($t = 0.17$) vs. late-time ($t = 0.83$) variance comparison. Percentage labels above each late bar show variance growth from early to late time (Baseline: +267%, Hahn Echo: +153%, CPMG: +89%), confirming temporal consistency of CPMG's variance suppression.

Figure 9 shows Bloch sphere visualizations of representative final states for each noise condition, comparing unprotected and best-protected outcomes. Table 2 provides the statistical summary (mean, Std, SEM, and 95% CI) for all strategy–noise combinations. Table 3 reports absolute fidelity gains in percentage points (pp). Table 4 presents the variance and the coefficient of variation (for both purity and fidelity) across noise regimes. Figure 10 presents Monte Carlo comparisons between baseline and optimized strategies across all three noise regimes.




Figure 9. Bloch sphere visualizations of representative final qubit states for each noise scenario (rows) comparing unprotected baseline (left) and best-performing protected strategy (right). Green dots show Monte Carlo trial spread; red arrows show the mean final Bloch vector. Tighter clustering indicates lower purity-decay variance. Under combined $T_1 + T_2$ noise, both conditions show dispersed, low-purity states, consistent with irreversible amplitude damping.

Table 2. Detailed statistical results for all strategy-noise combinations: mean fidelity (%), standard deviation, standard error of the mean, and 95% confidence interval across 500 Monte Carlo trials. Hahn Echo and CPMG achieve near-zero standard deviation under their optimal noise regimes, while Mixed-Sync shows the highest standard deviation under combined noise (Std = 1.43%), reflecting sensitivity to T_1 stochasticity.

Noise Scenario	Strategy	Mean Fidelity (%)	Std (%)	SEM (%)
Low-Freq Phase	Baseline	57.90	1.45	0.14
Low-Freq Phase	Hahn (1-X)	100.00	0.00	0.00
High-Freq Phase	Baseline	2.43	0.47	0.05
High-Freq Phase	CPMG (Dense)	97.62	0.45	0.05
Mixed (T1+T2)	Baseline	60.95	1.45	0.15
Mixed (T1+T2)	Mixed-Sync	65.34	1.43	0.14

Table 3. Mean fidelity and 95% confidence intervals for baseline and best-protected strategy in each noise regime, with absolute improvement in percentage points (pp). Hahn Echo and CPMG achieve gains of +42.10 pp and +95.17 pp respectively under their optimal conditions. Under combined $T_1 + T_2$ noise, the maximum gain is only +4.39 pp, demonstrating the ceiling imposed by irreversible energy relaxation.

Noise Regime	Strategy	Mean Fidelity (%)	95% CI	Improvement (pp)
LOW-FREQUENCY PHASE	BASELINE	57.90	[57.53, 58.27]	-
LOW-FREQUENCY PHASE	HAHN ECHO	100.00	[100.00, 100.00]	+42.10
HIGH-FREQUENCY PHASE	BASELINE	2.45	[2.18, 2.72]	-
HIGH-FREQUENCY PHASE	CPMG	97.62	[97.53, 97.71]	+95.17
COMBINED $T_1 + T_2$	BASELINE	60.95	[60.67, 61.23]	-
COMBINED $T_1 + T_2$	MIXED-SYNC	65.34	[65.06, 65.62]	+4.39

Table 4. The variance as well as the coefficient of variation (both for purity and fidelity) under different noise regimes.

Noise Regime	Strategy	Mean (%)	Std (%)	Variance (%) ²	SEM (%)	CV (Fidelity)	CV (Purity)	95% CI (Mean)
LOW-FREQ PHASE	BASELINE	57.90	0.48	0.2304	0.05	0.0083	—	[57.80, 58.00]
LOW-FREQ PHASE	HAHN ECHO	100.00	0.00	0.0000	0.00	0.0000	—	[100.00, 100.00]
HIGH-FREQ PHASE	BASELINE	2.45	0.38	0.1444	0.04	0.1551	0.18	[2.37, 2.53]
HIGH-FREQ PHASE	CPMG	97.62	0.45	0.2025	0.05	0.0046	0.06	[97.53, 97.71]
COMBINED $T_1 + T_2$	BASELINE	60.95	0.39	0.1521	0.04	0.0064	~0.08	[60.87, 61.03]
COMBINED $T_1 + T_2$	MIXED-SYNC	65.34	1.43	2.0449	0.14	0.0219	~0.08	[65.06, 65.62]




Figure 10. Monte Carlo simulation results (100 trials, 1024 shots per trial) showing estimated fidelity to $|0\rangle$ for baseline (blue) and best-protected strategy (orange) across all three noise regimes, with ± 1 SEM error bars. Echo-based strategies recover near-perfect fidelity under phase-dominant noise but provide only marginal improvement under combined $T_1 + T_2$ conditions.

3.6 Summary of results

This study evaluated four dynamical decoupling strategies across three noise regimes, quantifying both mean performance and stochastic robustness through coefficient of variation (CV) analysis, bootstrap confidence interval testing, and time-resolved variance characterization. The complete empirical findings establish both the power and fundamental limits of software-level decoherence mitigation.

Under phase-dominant noise regimes, echo-based DD achieved near-complete fidelity restoration. Hahn Echo recovered mean fidelity from 57.90% to 100.00% under low-frequency noise (+42.10 pp), while CPMG recovered fidelity from 2.45% to 97.62% under high-frequency noise (+95.17 pp). These gains reflect the reversibility of phase errors and the effectiveness of time-reversal pulse sequences when matched to the noise spectrum. Under combined $T_1 + T_2$ noise, all strategies converged to modest gains, with Mixed-Sync achieving

only +4.39 pp improvement (65.34% vs. 60.95% baseline), establishing irreversible amplitude damping as the ceiling for coherent pulse correction.

Variance analysis reinforced this performance gap. CV analysis found that CPMG reduced purity-decay variance threefold under high-frequency noise ($CV = 0.06$) compared to baseline ($CV = 0.18$). Bootstrap confidence interval analysis with 10,000 resampling iterations confirmed that this reduction was statistically robust, with completely non-overlapping 95% CIs between CPMG and baseline. Under combined noise, variance CIs overlapped substantially across all strategies ($CV \approx 0.07\text{--}0.09$), confirming that T_1 amplitude damping—not pulse strategy—dominates stochastic variability when energy relaxation becomes the primary decoherence mechanism.

4. Discussion

The results of this study provide strong empirical validation of dynamical decoupling

theory across qualitatively distinct noise regimes. The near-perfect fidelity restoration by Hahn Echo under low-frequency noise and by CPMG under high-frequency noise confirm the theoretical predictions of Viola and Lloyd (5) and Souza et al. (7). The mechanism underlying these results is the time-reversal symmetry of Hamiltonian evolution: a π pulse at the midpoint of an idle interval reverses the sign of the qubit-environment coupling, causing phase accumulated in the first half to be precisely canceled during the second.

The bootstrap confidence interval analysis provided statistical evidence that CPMG's variance reduction was not a sampling artifact but a fundamental consequence of dense pulse placement. Non-overlapping 95% CIs under high-frequency noise established that the observed threefold variance reduction could be reproduced across independent experimental realizations with high confidence. This statistical rigor strengthens the practical recommendation that CPMG should be the default choice for high-frequency noise mitigation when variance consistency is critical.

The time-resolved variance analysis revealed a dimension of DD performance that aggregate endpoint metrics cannot capture; i.e. at/from what precise time during the decoherence process variance suppression occurs. CPMG's ability to maintain low variance from early through late evolution times has direct implications for fault-tolerant quantum computing. In multi-qubit circuits, intermediate syndrome measurements or conditional gates may access qubit states at arbitrary points during idle intervals. Strategies that provide time-

dependent variance suppression (such as baseline or single-pulse DD) introduce timing-dependent error correlations that complicate error correction. CPMG's temporal consistency makes error budgets more predictable and syndrome extraction more reliable.

The convergence of variance across all strategies under combined $T_1 + T_2$ noise, confirmed by overlapping bootstrap CIs, establishes a practical guideline: practitioners should benchmark T_1 / T_2 ratios before investing in complex pulse sequences. When $T_1 \leq 2T_2$ (the theoretical maximum), amplitude damping dominates and DD benefit diminishes. Current state-of-the-art superconducting qubits exhibit T_1 times of 100–500 μs (9), but circuit idle times for complex algorithms can approach these limits, making hardware-level T_1 extension the critical bottleneck for near-term quantum advantage.

Time-resolved variance analysis further showed that CPMG suppressed variance consistently throughout the decoherence process, not just at endpoint measurement. The baseline strategy exhibited a +267% variance increase from early time ($t = 0.17$) to late time ($t = 0.83$), Hahn Echo showed +153% variance increase, while CPMG maintained low variance with a +89% increase. CPMG's early-time variance was $4.2\times$ lower than baseline, and this advantage persisted through late times ($3.8\times$ reduction). This consistency demonstrates that dense pulse placement provides reliable variance suppression regardless of when intermediate measurements or conditional gates access the qubit state, a critical property for fault-tolerant multi-qubit circuits.

The results show that dynamical decoupling sequences do not merely reduce average error, but reshape the temporal and statistical geometry of noise in a way that determines whether quantum information remains operationally usable. In particular, the time-resolved variance results showed that CPMG converts a regime of progressive stochastic divergence (baseline: +267% variance growth) into one of bounded, near-stationary fluctuations (+89%), implying that the error process transitions from variance-amplifying to variance-bounded dynamics under dense pulse control. This has a nontrivial consequence: even if two strategies yield similar endpoint fidelity or comparable mean purity at certain times, the one with suppressed variance growth ensures that intermediate quantum states remain within a predictable error envelope, which is essential for composability in multi-step circuits. Conversely, the convergence of variance across all strategies in the T_1 -dominated regime indicates that amplitude damping imposes a floor not just on mean coherence but on controllability of stochasticity itself—that is, beyond a certain point, the system loses not only information but also the ability to stabilize its uncertainty, rendering pulse-level interventions fundamentally incapable of improving reliability. Taken together, the results suggest a hierarchy: (1) in T_2 -dominated regimes, pulse sequences act as stochastic regulators that enforce temporal consistency; (2) in mixed regimes, they transiently shape but cannot bound variability; and (3) in T_1 -dominated regimes, both mean performance and higher-order statistical structure collapse to hardware-limited behavior. This reframes the role of dynamical decoupling from error “correction” to noise-structure engineering, where the key

design objective is not simply maximizing fidelity, but ensuring that error remains temporally uniform, bounded, and therefore schedulable within larger quantum algorithms.

Mechanistically, the dissociation between fidelity and purity under combined noise reveals that partial ground-state bias induced by T_1 decay can inflate fidelity metrics while true quantum coherence is irreversibly lost. The convergence of both mean performance and variance across strategies in this regime confirms that software-level pulse sequences cannot overcome hardware-limited energy relaxation. These findings establish clear design principles: DD should be deployed when $T_2 < T_1$ (phase noise dominates), noise spectrum characterization should guide pulse sequence selection (Hahn for low-frequency, CPMG for high-frequency), and variance consistency should be prioritized in multi-qubit applications where error correlations compound.

5. Limitations

This study was restricted to a single-qubit system with Markovian, idealized noise models. Multi-qubit crosstalk, leakage to non-computational states, and non-Markovian noise correlations characteristic of $1/f$ flux noise in superconducting hardware was excluded (13). All π pulses were treated as instantaneous and error-free; finite pulse duration and calibration imperfections accumulate over many pulses and may reduce CPMG’s advantage in practice. Noise parameters were held constant across trials, excluding hardware drift and slow parameter fluctuations. These idealizations favored DD performance relative to realistic hardware and should be addressed in future

work using experimentally measured noise spectra and multi-qubit circuit models.

6. Conclusion

This study presented a systematic simulation-based evaluation of dynamical decoupling (DD) as a software-level decoherence mitigation strategy for single-qubit quantum systems. Hahn Echo achieved perfect fidelity restoration (+42.10 pp) under low-frequency phase noise, while CPMG achieved near-perfect recovery (+95.17 pp) under high-frequency phase noise. Bootstrap confidence interval analysis ($n = 10,000$) confirmed that CPMG's threefold variance reduction under high-frequency noise as being statistically robust, with completely non-overlapping 95% CIs providing strong evidence of reproducibility. Time-resolved variance analysis further demonstrated that CPMG suppresses variance consistently throughout evolution (+89% early-to-late increase) compared to baseline's progressive amplification (+267%), establishing temporal reliability relevant for multi-qubit circuit operations.

Critically, these results show that dynamical decoupling sequences influence the temporal and statistical structure of noise, converting variance-amplifying error dynamics into variance-bounded behavior under dense pulse control. This extends the role of DD beyond

fidelity optimization toward shaping the temporal consistency of errors. Under combined $T_1 + T_2$ noise, all strategies converged to modest gains (+4.39 pp maximum) with overlapping variance confidence intervals, establishing amplitude damping as the fundamental limit of software-level correction.

Taken together, the results define a three-tier regime hierarchy: DD is most effective in T_2 -dominated environments, where pulse sequences act as stochastic regulators that constrain variability; it retains partial efficacy in mixed regimes; and it becomes fundamentally limited in T_1 -dominated regimes, where both mean performance and variability converge to hardware-imposed behavior. These findings indicate that composability in multi-step quantum circuits depends not only on high mean fidelity but also on suppressed variance growth, making temporal consistency a key design criterion.

These results have direct implications for NISQ-era quantum circuit design: DD should be applied during idle intervals when phase noise dominates, with noise characterization incorporating both mean and variance metrics to predict benefit. Future work should extend this framework to multi-qubit systems with correlated noise, experimentally calibrated non-Markovian models, and integration with quantum error correction protocols.

7. References

1. Shor, P. W. (1994). Algorithms for quantum computation: Discrete logarithms and factoring. In *Proceedings of the 35th Annual Symposium on Foundations of Computer Science* (pp. 124–134). <https://doi.org/10.1109/SFCS.1994.365700>

2. Grover, L. K. (1996). A fast quantum mechanical algorithm for database search. In *Proceedings of the 28th Annual ACM Symposium on Theory of Computing* (pp. 212–219). <https://doi.org/10.1145/237814.237866>
3. Cao, Y., Romero, J., Olson, J. P., Degroote, M., Johnson, P. D., Kieferová, M., et al. (2019). Quantum chemistry in the age of quantum computing. *Chemical Reviews*, 119(19), 10856–10915. <https://doi.org/10.1021/acs.chemrev.8b00803>
4. Nielsen, M. A., & Chuang, I. L. (2009). *Quantum computation and quantum information*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511976667>
5. Viola, L., & Lloyd, S. (1998). Dynamical suppression of decoherence in two-state quantum systems. *Physical Review A*, 58(4), 2733–2744. <https://doi.org/10.1103/PhysRevA.58.2733>
6. Hahn, E. L. (1950). Spin echoes. *Physical Review*, 80(4), 580–594. <https://doi.org/10.1103/PhysRev.80.580>
7. Souza, A. M., Álvarez, G. A., & Suter, D. (2011). Robust dynamical decoupling for quantum computing and quantum memory. *Physical Review Letters*, 106(24), 240501. <https://doi.org/10.1103/PhysRevLett.106.240501>
8. Khodjasteh, K., & Lidar, D. A. (2005). Fault-tolerant quantum dynamical decoupling. *Physical Review Letters*, 95(18), 180501. <https://doi.org/10.1103/PhysRevLett.95.180501>
9. Krantz, P., Kjaergaard, M., Yan, F., Orlando, T. P., Gustavsson, S., & Oliver, W. D. (2019). A quantum engineer's guide to superconducting qubits. *Applied Physics Reviews*, 6(2), 021318.
10. Yan, F., Gustavsson, S., Kamal, A., Birenbaum, J., Sears, A., Hover, D., et al. (2016). The flux qubit revisited to enhance coherence and reproducibility. *Nature Communications*, 7(1), 12964. <https://doi.org/10.1038/ncomms12964>
11. Preskill, J. (2018). Quantum computing in the NISQ era and beyond. *Quantum*, 2, 79. <https://doi.org/10.22331/q-2018-08-06-79>
12. Qiskit contributors. (2023). *Qiskit: An open-source framework for quantum computing*. <https://doi.org/10.5281/zenodo.2573505>
13. Breuer, H.-P., & Petruccione, F. (2006). *The theory of open quantum systems*. Oxford University Press.

14. Paladino, E., Galperin, Y. M., Falci, G., & Altshuler, B. L. (2014). $1/f$ noise: Implications for solid-state quantum information. *Reviews of Modern Physics*, 86(2), 361–418.
<https://doi.org/10.1103/RevModPhys.86.361>