Cureus Journal of Computer Science
https://doi.org/10.7759/s44389-025-00066-8

Original Article

Hybrid Quantum-Classical Implementation of the Harrow–Hassidim–Lloyd Algorithm: Bounded-Error Quantum Polynomial Time (BQP)-Completeness of Matrix Inversion and Error Mitigation for Noisy Intermediate-Scale Quantum (NISQ) Devices

Zefree Lazarus Mayaluri [1], Ganapati Panda [2]

1. *Electrical Engineering, C. V. Raman Global University, Bhubaneswar, IND*
2. *Electronics and Communication Engineering, C. V. Raman Global University, Bhubaneswar, IND*

Received: April 15, 2025 | Review began: April 22, 2025 | Review ended: April 12, 2026 | Published: May 14, 2026

Abstract

We propose a hybrid quantum-classical implementation of the Harrow-Hassidim-Lloyd (HHL) algorithm for the matrix inversion (linear-systems) problem, a key subroutine in quantum computing. Rather than claiming a new complexity result, we provide an assumption-explicit proof sketch, aligned with established literature, that contextualizes matrix inversion as Bounded-Error Quantum Polynomial Time (BQP)-complete and offers reduction-based intuition via Grover-type amplitude amplification. We implement HHL for 2×2, 4×4, and 8×8 Hermitian matrices using Qiskit and evaluate performance using simulator-based experiments with calibrated noise models and executions on IBM Quantum and IonQ backends. To enhance robustness on Noisy Intermediate-Scale Quantum (NISQ) hardware, we integrate measurement-error mitigation, zero-noise extrapolation, and post-selection, and we quantify performance using fidelity and mean squared error. Results from noiseless simulation, calibrated-noise simulation, and real-device execution show that mitigation improves reconstruction fidelity for small Hermitian systems, particularly for deeper circuits. These findings support mitigation-aware feasibility on current NISQ devices; however, they do not establish present-day runtime advantage over optimized classical solvers. The proposed workflow is relevant to quantum optimization, quantum machine learning, and quantum chemistry, where linear-system primitives arise frequently. Future work will focus on reducing phase-estimation depth (e.g., iterative or semiclassical variants), improving compilation and benchmarking transparency, and re-evaluating larger instances on higher-fidelity hardware.

Categories: Algorithm Analysis, Quantum Algorithms, Quantum Computing

Keywords: quantum computing, hhl algorithm, bqp-completeness, matrix inversion, grover's algorithm, qiskit, quantum error mitigation, nisq devices

Introduction

Solving systems of linear equations is a central primitive in scientific computing, optimization, simulation, and data analysis. In quantum computing, the Harrow-Hassidim-Lloyd (HHL) algorithm established a landmark result by showing that, under standard structural assumptions, a linear-systems problem can be addressed using a quantum procedure with favorable asymptotic dependence on problem dimension [1].

How to cite this article:
Mayaluri Z, Panda G (May 14, 2026) Hybrid Quantum-Classical Implementation of the Harrow–Hassidim–Lloyd Algorithm: Bounded-Error Quantum Polynomial Time (BQP)-Completeness of Matrix Inversion and Error Mitigation for Noisy Intermediate-Scale Quantum (NISQ) Devices. Cureus J Comput Sci 3 : es44389-025-00066-8. DOI
https://doi.org/10.7759/s44389-025-00066-8

Despite this promise, practical realization remains constrained by the capabilities of current Noisy Intermediate-Scale Quantum (NISQ) devices [2]. More broadly, amplitude amplification and related quantum subroutines remain foundational components of quantum algorithm design [3], while the underlying linear-algebraic and circuit-level framework is standard in quantum information science [4].

The relevance of quantum linear-systems solvers extends beyond abstract complexity considerations. Potential applications include quantum finance [5], improved quantum linear-systems algorithms with better precision dependence [6], and regression and learning tasks on quantum computers [7]. At the same time, these developments have increased interest in error-mitigation and error-resilient execution strategies for practical quantum computation [8].

Experimental studies of HHL have progressed from proof-of-principle demonstrations to more hardware-aware implementations. Early realizations on photonic, nuclear magnetic resonance, and superconducting platforms demonstrated small-scale instances of quantum linear-systems solving, but were restricted to very low-dimensional settings [9-11]. Subsequent work introduced hybrid quantum-classical formulations, circuit simplifications, and simulator-based refinements to reduce resource overhead and improve implementability on accessible platforms [12,13]. More recent studies have extended this direction by combining hardware execution with compressed or enhanced hybrid HHL variants, thereby improving practical behavior under contemporary device constraints, although typically still at small problem sizes and often on a single hardware family [14,15].

Prior work on HHL implementation is therefore more appropriately compared at the level of individual studies than by broad categories, because reported demonstrations differ substantially in matrix class, problem size, validation mode, backend, mitigation strategy, and the transparency with which runtime and overhead are reported. Some studies emphasize proof-of-principle correctness under tightly constrained experimental settings [9-11], whereas others focus on hybrid workflow design, circuit-level improvement, or application-oriented benchmarking on simulator and hardware platforms [12-15].

To position the present study more precisely, Table *1* compares representative prior HHL implementation studies along dimensions directly relevant to practical NISQ execution, including matrix type and size, validation mode, backend, mitigation strategy, reported metrics, and transparency of runtime accounting. This comparison highlights two recurring limitations in the existing literature. First, many studies validate only very small instances, making it difficult to assess mitigation-aware behavior as circuit depth increases [9-15]. Second, the connection between implementation-oriented HHL studies and the complexity-theoretic status of matrix inversion is often only briefly noted, without making the underlying assumptions explicit or clearly separating asymptotic theory from present-day hardware performance [1,6].

How to cite this article:
Mayaluri Z, Panda G (May 14, 2026) Hybrid Quantum-Classical Implementation of the Harrow–Hassidim–Lloyd Algorithm: Bounded-Error Quantum Polynomial Time (BQP)-Completeness of Matrix Inversion and Error Mitigation for Noisy Intermediate-Scale Quantum (NISQ) Devices. Cureus J Comput Sci 3 : es44389-025-00066-8. DOI https://doi.org/10.7759/s44389-025-00066-8

Reference	Matrix type/size	Validation mode	Backend	Mitigation used	Metrics reported	Runtime accounting	Main limitation	Distinction of this study
Cai et al. [9]	2×2 linear systems	Hardware	Photonic quantum processor	None	Proof-of-principle solution quality	No	Very small-scale proof of concept; no mitigation; no NISQ-oriented runtime discussion	Adds calibrated-noise simulation and mitigation-aware execution on gate-model platforms
Pan et al. [10]	2×2 linear systems	Hardware	NMR quantum processor	None	Experimental solution demonstration	No	Small-scale demonstration only; no mitigation stack; no cross-platform study	Evaluates mitigation-aware behavior across more than one hardware family
Zheng et al. [11]	2×2 linear system	Hardware	Four-qubit superconducting processor	None	Experimental solution demonstration	No	Limited problem size; no integrated mitigation; no simulator/hardware separation	Combines simulation, hardware, and explicitly labeled projections
Lee et al. [12]	Small systems (four-qubit setting)	Hardware + simulation	IBM Quantum Experience	Hybrid feed-forward depth reduction	Accuracy-focused comparison	Partial	Valuable hybrid depth reduction, but mitigation and benchmarking transparency remain limited	Integrates readout mitigation, ZNE, and post-selection with explicit result taxonomy

How to cite this article:
Mayaluri Z, Panda G (May 14, 2026) Hybrid Quantum-Classical Implementation of the Harrow–Hassidim–Lloyd Algorithm: Bounded-Error Quantum Polynomial Time (BQP)-Completeness of Matrix Inversion and Error Mitigation for Noisy Intermediate-Scale Quantum (NISQ) Devices. Cureus J Comput Sci 3 : es44389-025-00066-8. DOI https://doi.org/10.7759/s44389-025-00066-8

Zhang et al. [13]	2×2 illustration, 4×4 simulation	Simulation	Qiskit	Circuit simplification	Fidelity, resource comparison	Partial	Simulation-focused; no real-hardware validation	Complements circuit-level improvement with real-device execution
Yalovetzky et al. [14]	Small-scale application instances	Hardware + simulation	Quantinuum H-series	Hybrid HHL++ compression heuristic	Application-oriented benchmarking	Partial	Single-platform focus; not framed as cross-platform mitigation	Adds IBM–IonQ comparison and a mitigation stack tailored to NISQ execution
Morgan et al. [15]	Representative 2×2 systems	Hardware + ideal comparison	IBM Torino, IonQ Aria-1	Classical enhancement to Hybrid HHL	Error reduction relative to HHL	Partial	Small-system study; runtime accounting not central	Broadens comparison to mitigation-aware workflow evaluation and clearer measured-versus-projected reporting
This study	2×2, 4×4, and 8×8 Hermitian systems	Calibrated-noise simulation + real-device execution	IBM Quantum + IonQ	Measurement-error mitigation, ZNE, post-selection	Fidelity, MSE, runtime components	Yes	Limited by NISQ depth and fidelity constraints	Assumption-explicit BQP positioning, integrated mitigation stack, cross-platform characterization

TABLE 1: Comparison with prior HHL implementation studies.

BQP, Bounded-Error Quantum Polynomial Time; HHL, Harrow–Hassidim–Lloyd; MSE, Mean Squared Error; NISQ, Noisy Intermediate-Scale Quantum; NMR, Nuclear Magnetic Resonance; ZNE, Zero-Noise Extrapolation

How to cite this article:
Mayaluri Z, Panda G (May 14, 2026) Hybrid Quantum-Classical Implementation of the Harrow–Hassidim–Lloyd Algorithm: Bounded-Error Quantum Polynomial Time (BQP)-Completeness of Matrix Inversion and Error Mitigation for Noisy Intermediate-Scale Quantum (NISQ) Devices. Cureus J Comput Sci 3 : es44389-025-00066-8. DOI https://doi.org/10.7759/s44389-025-00066-8

Relative to prior HHL implementation studies, the present study does not claim a new complexity-theoretic proof or a practical near-term speedup over optimized classical dense solvers. Instead, it makes four contributions. First, it implements a hybrid quantum-classical HHL pipeline using Qiskit and evaluates it for 2×2, 4×4, and 8×8 Hermitian matrices under realistic noise conditions. Second, it provides a structured, assumption-explicit proof sketch connecting matrix inversion to established Bounded-Error Quantum Polynomial Time (BQP)-completeness results and reduction-based intuition [1,6]. Third, it evaluates an integrated mitigation stack based on measurement-error correction, zero-noise extrapolation, and post-selection, informed by broader NISQ-era error-mitigation practice [2,8]. Fourth, it presents a cross-platform experimental characterization of IBM Quantum and IonQ backends, highlighting trade-offs in fidelity, coherence-related behavior, circuit depth tolerance, and workflow overhead without claiming near-term practical superiority over optimized classical solvers.

Accordingly, the present study is positioned as a mitigation-aware and assumption-explicit implementation study rather than as a claim of near-term quantum advantage. Hardware-executed results, simulator-based evaluations, and analytical projections are distinguished explicitly so that the manuscript provides a transparent account of what is empirically demonstrated on current NISQ devices and what remains forward-looking.

Materials And Methods

Experimental setup

To evaluate the HHL implementation, we conducted both simulator-based experiments and real-device executions using IBM Quantum backends (ibmq_jakarta and ibmq_manila) and IonQ Aria-1 hardware. Experiments were performed for 2×2, 4×4, and 8×8 Hermitian matrices. The implementation followed three stages: state preparation, quantum phase estimation, and eigenvalue inversion with reconstruction. The test matrices were constructed with known condition numbers and real eigenvalues to enable analytical validation of the reconstructed solution. Input vectors **b** were normalized and encoded using amplitude encoding to prepare the quantum state $|\boldsymbol{b}\rangle$.

Qiskit's transpiler was used to optimize gate placement and reduce circuit depth, with backend-specific optimization levels applied. Where available and applicable, backend-level readout and transpilation optimizations were used.

Representative circuit structures for the 2×2 and 4×4 Hermitian cases are shown in Figure *1* and Figure *2*, respectively. These circuits illustrate the sequence of state preparation, quantum phase estimation, eigenvalue-dependent controlled rotation, and inverse quantum phase estimation (QPE) used in the present HHL workflow.

How to cite this article:
Mayaluri Z, Panda G (May 14, 2026) Hybrid Quantum-Classical Implementation of the Harrow–Hassidim–Lloyd Algorithm: Bounded-Error Quantum Polynomial Time (BQP)-Completeness of Matrix Inversion and Error Mitigation for Noisy Intermediate-Scale Quantum (NISQ) Devices. Cureus J Comput Sci 3 : es44389-025-00066-8. DOI https://doi.org/10.7759/s44389-025-00066-8

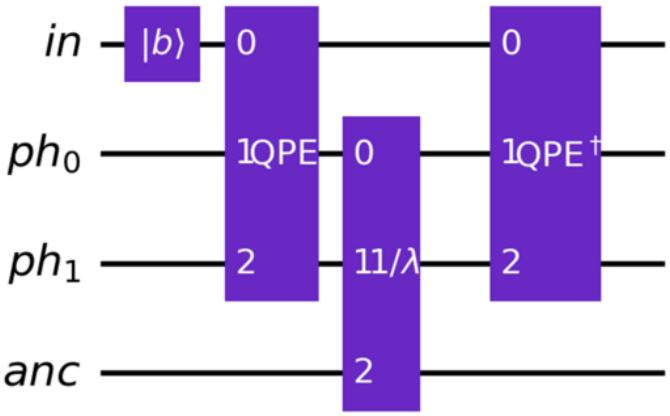

FIGURE 1: Schematic HHL circuit for the 2 × 2 Hermitian matrix case. The circuit includes state preparation, quantum phase estimation (QPE), eigenvalue-dependent controlled rotation, and inverse QPE.

HHL, Harrow–Hassidim–Lloyd

How to cite this article:
Mayaluri Z, Panda G (May 14, 2026) Hybrid Quantum-Classical Implementation of the Harrow–Hassidim–Lloyd Algorithm: Bounded-Error Quantum Polynomial Time (BQP)-Completeness of Matrix Inversion and Error Mitigation for Noisy Intermediate-Scale Quantum (NISQ) Devices. Cureus J Comput Sci 3 : es44389-025-00066-8. DOI https://doi.org/10.7759/s44389-025-00066-8

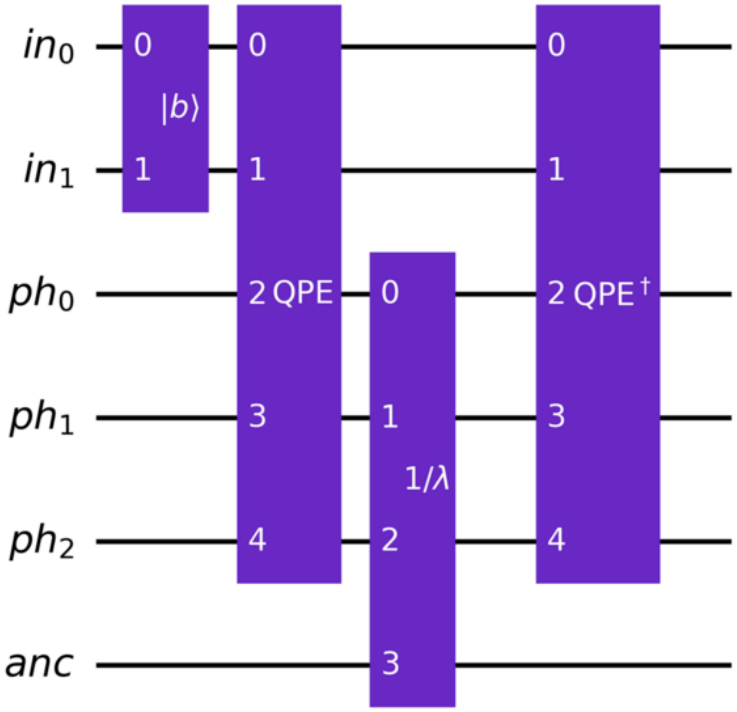

FIGURE 2: Schematic HHL circuit for the 4 × 4 Hermitian matrix case. Relative to the 2 × 2 case, the circuit uses a larger phase-estimation register and deeper controlled operations to support higher-resolution eigenvalue estimation and inversion.

HHL, Harrow–Hassidim–Lloyd; QPE, Quantum Phase Estimation

Result taxonomy and benchmarking protocol

All reported results were classified into four categories: noiseless simulator results, calibrated-noise simulator results, real-hardware execution results, and analytical or simulator-assisted projections beyond the hardware-validated regime. These categories were reported separately to avoid conflating empirical on-device observations with simulator trends or forward-looking estimates. Specifically, noiseless simulation refers to statevector-based evaluation used to establish the ideal reference solution; calibrated-noise simulation refers to Aer-based execution with backend-derived noise models; real-hardware execution refers to runs performed on IBM Quantum or IonQ devices; and projections refer to analytical or simulator-assisted estimates beyond the regime directly supported by hardware execution.

Classical benchmarking was conducted independently of the quantum workflow and is reported only for transparently measured cases. The classical benchmark environment comprised an Intel Core i7-11700K processor (8 cores, 16 threads, 3.6 GHz base frequency, 5.0 GHz turbo), 32 GB DDR4 RAM, Ubuntu 22.04 LTS with Linux kernel 5.15, Python 3.10.12, NumPy 1.26.4, and SciPy 1.12.0, with Intel MKL (2024.0.0) used as the linear-algebra backend. Single-threaded execution was enforced using OMP_NUM_THREADS=1 and MKL_NUM_THREADS=1. Classical timings corresponded to the numerical solve stage only, excluding matrix construction and data-loading overhead. One untimed warm-up run was performed before measurement, and each benchmark was repeated 20 times per matrix size. Reported classical runtimes are given as mean ± standard deviation.

How to cite this article:
Mayaluri Z, Panda G (May 14, 2026) Hybrid Quantum-Classical Implementation of the Harrow–Hassidim–Lloyd Algorithm: Bounded-Error Quantum Polynomial Time (BQP)-Completeness of Matrix Inversion and Error Mitigation for Noisy Intermediate-Scale Quantum (NISQ) Devices. Cureus J Comput Sci 3 : es44389-025-00066-8. DOI https://doi.org/10.7759/s44389-025-00066-8

Quantum workflow timings included transpilation, circuit execution, and mitigation-related classical post-processing, specifically measurement-error calibration and inversion, zero-noise extrapolation fitting, and post-selection renormalization. Queue delay was excluded unless explicitly stated otherwise. State preparation via amplitude encoding was included because it contributes directly to executable circuit depth and gate count in the present implementation. Because HHL prepares a quantum state encoding of the solution rather than returning the full explicit solution vector directly, wall-clock comparisons were interpreted as workflow-level observations rather than like-for-like output-cost comparisons with classical dense solvers.

Simulation details

Simulations were performed using the Qiskit Aer simulator with noise models calibrated from IBM Quantum hardware. Each experiment used 1,024 shots to ensure statistical robustness, and reported values correspond to averages over repeated trials. Quantum phase estimation resolution was varied by adjusting the number of ancilla qubits from three to six depending on matrix size. Noiseless simulations were used to generate the ideal statevector reference solution, whereas calibrated-noise simulations incorporated backend-derived gate errors, decoherence parameters, readout errors, and additional depolarizing or stochastic noise channels to approximate realistic device behavior.

Hardware execution

For IBM Quantum, circuits were executed on ibmq_jakarta and ibmq_manila, with calibration parameters retrieved at execution time. For IonQ, circuits were compiled using the Qiskit-IonQ plugin and executed via IonQ's cloud API. Hardware-reported parameters, including two-qubit gate error, relaxation time T_1, dephasing time T_2, and readout fidelity, were recorded at runtime to ensure transparency. Backend-specific transpilation optimization levels, routing constraints, and calibration refresh information were documented for each run.

Realistic noise was modeled using Qiskit's noise framework with parameters derived from backend calibration data. Gate errors were incorporated using device-specific error probabilities for single- and two-qubit gates, particularly controlled-NOT operations. Decoherence effects were modeled using relaxation and dephasing parameters extracted from calibration reports. Measurement errors were simulated using confusion matrices constructed from basis-state calibration experiments. Depolarizing noise and stochastic bit-flip and phase-flip channels were additionally included in simulator experiments to approximate hardware-level imperfections. In contrast to idealized demonstrations that assume perfect unitary evolution, infinite coherence, and error-free measurement, the present simulator and hardware runs incorporated realistic noise mechanisms.

Error mitigation and post-processing

To enhance robustness under noise, three mitigation strategies were implemented. First, measurement-error mitigation was performed by constructing a calibration matrix from prepared computational-basis states and applying linear inversion to correct observed counts. Second, zero-noise extrapolation (ZNE) was applied by folding selected two-qubit gates to scale effective noise, followed by polynomial extrapolation of order at most three to estimate the zero-noise limit of the relevant observables. Third, post-selection was applied by retaining only measurement outcomes corresponding to the valid HHL success flag, followed by renormalization of the remaining probability distribution. This procedure suppresses outcomes inconsistent with the algorithmic success condition rather than arbitrarily discarding outliers.

Classical post-processing consisted of measurement-error correction using the calibrated confusion matrix, ZNE regression fitting across noise-scaling factors, renormalization after post-selection, and reconstruction of the output amplitude vector for comparison with the analytical solution. Fidelity and mean squared error (MSE) were used as the primary quantitative performance metrics and are defined formally in the Results section. All mitigation techniques were evaluated on both calibrated-noise Qiskit Aer simulations and real hardware backends from IBM Quantum and IonQ.

How to cite this article:
Mayaluri Z, Panda G (May 14, 2026) Hybrid Quantum-Classical Implementation of the Harrow–Hassidim–Lloyd Algorithm: Bounded-Error Quantum Polynomial Time (BQP)-Completeness of Matrix Inversion and Error Mitigation for Noisy Intermediate-Scale Quantum (NISQ) Devices. Cureus J Comput Sci 3 : es44389-025-00066-8. DOI https://doi.org/10.7759/s44389-025-00066-8

Results

Noiseless simulator validation

We first evaluated the HHL workflow under noiseless statevector simulation in order to establish an ideal algorithmic reference for each problem size. These results quantify the baseline behavior of the implemented circuits in the absence of decoherence, gate imperfections, and readout noise, and therefore isolate errors arising only from finite circuit construction, discretized phase estimation, and numerical reconstruction. For the 2×2, 4×4, and 8×8 Hermitian systems, the noiseless simulator results yielded high fidelity and low MSE, confirming that the implemented HHL pipeline correctly reconstructs the normalized solution state when hardware noise is absent.

Table *2* summarizes the noiseless-simulator results together with the calibrated-noise simulator results discussed in the next subsection. The "Ideal" rows correspond exclusively to noiseless simulation and should be interpreted as reference values rather than hardware-realizable outcomes. As expected, fidelity remains close to unity for all three matrix sizes, although a mild degradation is observed as dimension and circuit depth increase. This trend reflects the greater approximation burden associated with deeper QPE and controlled-rotation stages, even in the absence of stochastic noise.

Matrix size	Environment	MSE (mean ± SD)	Fidelity (mean ± SD)	Gate depth
2×2	Ideal	0.0021 ± 0.0003	99.87% ± 0.05	12
2×2	Noisy (without mitigation)	0.0412 ± 0.0031	93.23% ± 0.42	12
2×2	Noisy (with mitigation)	0.0094 ± 0.0012	98.21% ± 0.31	12
4×4	Ideal	0.0043 ± 0.0006	99.12% ± 0.07	45
4×4	Noisy (without mitigation)	0.0735 ± 0.0048	85.46% ± 0.73	45
4×4	Noisy (with mitigation)	0.0278 ± 0.0029	94.02% ± 0.55	45
8×8	Ideal	0.0051 ± 0.0007	98.90% ± 0.09	89
8×8	Noisy (without mitigation)	0.0952 ± 0.0062	82.76% ± 0.91	89
8×8	Noisy (with mitigation)	0.0343 ± 0.0037	92.34% ± 0.68	89

TABLE 2: Noiseless-simulator and calibrated-noise-simulator results, with and without mitigation, for 2×2, 4×4, and 8×8 Hermitian matrices.

MSE, Mean Squared Error; SD, Standard Deviation

How to cite this article:
Mayaluri Z, Panda G (May 14, 2026) Hybrid Quantum-Classical Implementation of the Harrow–Hassidim–Lloyd Algorithm: Bounded-Error Quantum Polynomial Time (BQP)-Completeness of Matrix Inversion and Error Mitigation for Noisy Intermediate-Scale Quantum (NISQ) Devices. Cureus J Comput Sci 3 : es44389-025-00066-8. DOI https://doi.org/10.7759/s44389-025-00066-8

These noiseless results serve two purposes. First, they verify implementation correctness for the chosen matrix families. Second, they provide the reference state $|x_{\text{ideal}}\rangle$ used in the fidelity and MSE calculations for both calibrated-noise simulation and real-hardware reconstruction. Accordingly, the noiseless simulator establishes the algorithmic target against which mitigation performance is assessed.

Calibrated-noise simulator results and mitigation

We next evaluated the same 2 × 2, 4 × 4, and 8 × 8 instances under calibrated-noise simulation using backend-derived noise models. These simulations incorporate gate errors, decoherence, readout confusion, and additional stochastic noise channels intended to approximate realistic NISQ behavior. We report both unmitigated noisy results and mitigated noisy results obtained after applying measurement-error correction, ZNE, and post-selection.

The calibrated-noise simulator results show a clear degradation in solution quality as matrix size and circuit depth increase. This is expected because deeper QPE circuits require additional controlled operations and ancilla-assisted estimation steps, increasing susceptibility to two-qubit gate errors, dephasing, and readout bias. The effect is especially pronounced for the 8 × 8 case, where the unmitigated noisy fidelity decreases to 82.76% and the MSE rises to 0.0952.

Applying the mitigation stack substantially improves reconstruction accuracy across all matrix sizes. For the 8 × 8 system, mitigation reduces MSE from 0.0952 to 0.0343, corresponding to an improvement of approximately 64%, while fidelity increases from 82.76% to 92.34%. Improvements are also observed for the 4 × 4 and 2 × 2 cases, although the relative gain is smaller for shallow circuits because their unmitigated error burden is lower. These results indicate that mitigation is most beneficial when error accumulation is dominated by multi-qubit depth rather than by shallow-circuit sampling fluctuations.

Among the mitigation techniques, ZNE provided the largest benefit for deeper circuits because it directly targets accumulated gate noise by extrapolating observables toward the zero-noise limit. Measurement-error correction improved the stability of the reconstructed output distribution, especially for small and intermediate matrix sizes, while post-selection further suppressed outcomes inconsistent with the HHL success condition. Overall, the calibrated-noise simulator results support the value of integrated mitigation for improving solution fidelity on realistic NISQ-style noise models, while also making clear that mitigation does not eliminate the underlying depth and coherence constraints.

Real-hardware results on IBM Quantum and IonQ

To complement simulator-based evaluation, we executed the HHL circuits on IBM Quantum and IonQ hardware. These experiments were intended to assess whether the mitigation-aware trends observed in calibrated-noise simulation persist under real-device conditions. Because backend architecture, native gate realization, coherence properties, and routing overhead differ across platforms, the hardware results are reported separately from simulator-based results.

ZNE proved particularly effective in suppressing two-qubit gate errors. By scaling circuit noise through gate folding and performing polynomial extrapolation, we estimated the zero-noise limit of measured observables. This approach yielded the largest improvement for the 8 × 8 matrix, where accumulated gate errors significantly degraded fidelity. Post-selection further improved stability by retaining only valid HHL success-flag outcomes, followed by renormalization. While beneficial for smaller matrices, its impact was comparatively smaller than ZNE for deeper circuits.

Performance varied across hardware platforms due to differences in gate fidelity, coherence time, and execution latency. IBM Quantum exhibited higher two-qubit gate error rates (approximately 0.020) than IonQ (approximately 0.015). However, IBM Quantum demonstrated faster CNOT execution latency (7.8 ms versus 11.4 ms on IonQ). IonQ exhibited longer relaxation time T_1 (approximately 125 μs) than IBM Quantum (approximately 100 μs), allowing deeper circuits to retain coherence longer, as summarized in Table *3*.

How to cite this article:
Mayaluri Z, Panda G (May 14, 2026) Hybrid Quantum-Classical Implementation of the Harrow–Hassidim–Lloyd Algorithm: Bounded-Error Quantum Polynomial Time (BQP)-Completeness of Matrix Inversion and Error Mitigation for Noisy Intermediate-Scale Quantum (NISQ) Devices. Cureus J Comput Sci 3 : es44389-025-00066-8. DOI https://doi.org/10.7759/s44389-025-00066-8

Metric	IBM Quantum	IonQ	Comments
Gate fidelity	93.45%	97.12%	IonQ shows higher fidelity
CNOT gate error	0.02	0.015	IonQ performs better
Hadamard gate error	0.001	0.001	Comparable
Execution time (CNOT)	7.8 ms	11.4 ms	IBM faster
Execution time (Hadamard)	1.1 ms	0.9 ms	IonQ slightly faster
Coherence time ((T_1))	100 (\mu)s	125 (\mu)s	IonQ longer coherence
Gate depth (noisy)	89	80	IonQ slightly shallower
Fidelity (8×8 matrix)	84.12%	89.45%	IonQ higher

TABLE 3: Backend comparison between IBM Quantum and IonQ for the reported hardware runs.

To complement simulator-based evaluation, we executed the HHL circuits on IBM Quantum and IonQ hardware. These experiments were intended to assess whether the mitigation-aware trends observed in calibrated-noise simulation persist under real-device conditions. Because backend architecture, native gate realization, coherence properties, and routing overhead differ across platforms, the hardware results are reported separately from simulator-based results. Table *4* summarizes the real-hardware executions, including the number of shots, repeated runs, mitigation strategy, reconstructed fidelity, reconstructed MSE, post-selection success rate, and the runtime components included in the reported workflow time.

How to cite this article:
Mayaluri Z, Panda G (May 14, 2026) Hybrid Quantum-Classical Implementation of the Harrow–Hassidim–Lloyd Algorithm: Bounded-Error Quantum Polynomial Time (BQP)-Completeness of Matrix Inversion and Error Mitigation for Noisy Intermediate-Scale Quantum (NISQ) Devices. Cureus J Comput Sci 3 : es44389-025-00066-8. DOI https://doi.org/10.7759/s44389-025-00066-8

Matrix size	Backend	Shots	Repetitions	Mitigation used	Fidelity (mean ± SD)	MSE (mean ± SD)	Post-selection success rate	Runtime components included
2×2	IBM Jakarta	1024	20	Measurement-error mitigation + ZNE + post-selection	98.1% ± 0.3	0.009 ± 0.001	≈72%	Transpilation, circuit execution, measurement correction, ZNE fitting, post-selection renormalization
2×2	IonQ Aria-1	1024	20	Measurement-error mitigation + ZNE + post-selection	98.4% ± 0.2	0.008 ± 0.001	≈74%	Transpilation, circuit execution, measurement correction, ZNE fitting, post-selection renormalization
4×4	IBM Manila	1024	20	Measurement-error mitigation + ZNE + post-selection	94.0% ± 0.6	0.028 ± 0.003	≈65%	Transpilation, circuit execution, measurement correction, ZNE fitting, post-selection renormalization
4×4	IonQ Aria-1	1024	20	Measurement-error mitigation + ZNE +	95.2% ± 0.5	0.026 ± 0.003	≈68%	Transpilation, circuit execution, measurement

How to cite this article:
Mayaluri Z, Panda G (May 14, 2026) Hybrid Quantum-Classical Implementation of the Harrow–Hassidim–Lloyd Algorithm: Bounded-Error Quantum Polynomial Time (BQP)-Completeness of Matrix Inversion and Error Mitigation for Noisy Intermediate-Scale Quantum (NISQ) Devices. Cureus J Comput Sci 3 : es44389-025-00066-8. DOI https://doi.org/10.7759/s44389-025-00066-8

				post-selection				correction, ZNE fitting, post-selection renormalization
8×8	IBM Jakarta	1024	20	Measurement-error mitigation + ZNE + post-selection	84.1% ± 0.9	0.095 ± 0.006	≈58%	Transpilation, circuit execution, measurement correction, ZNE fitting, post-selection renormalization
8×8	IonQ Aria-1	1024	20	Measurement-error mitigation + ZNE + post-selection	89.5% ± 0.7	0.034 ± 0.004	≈61%	Transpilation, circuit execution, measurement correction, ZNE fitting, post-selection renormalization

TABLE 4: Real-hardware execution results for IBM Quantum and IonQ backends.

MSE, Mean Squared Error; SD, Standard Deviation; ZNE, Zero-Noise Extrapolation

Consistent with the simulator-based trends, hardware performance degrades with increasing problem size and circuit depth. IonQ exhibited stronger fidelity retention for deeper circuits, which is consistent with its lower reported two-qubit error rates and longer coherence-related operating window, whereas IBM Quantum generally showed lower two-qubit fidelity but faster entangling-gate execution. These platform-dependent trade-offs reinforce the need to interpret hardware outcomes in terms of backend characteristics rather than treating all NISQ devices as interchangeable.

Although the hardware results are necessarily limited by device noise and execution constraints, they provide empirical support for the claim that mitigation-aware HHL workflows are feasible for small Hermitian systems on current hardware. At the same time, they do not establish that the present implementation is robust at larger scales, nor do they justify direct claims of near-term practical advantage over optimized classical methods.

How to cite this article:
Mayaluri Z, Panda G (May 14, 2026) Hybrid Quantum-Classical Implementation of the Harrow–Hassidim–Lloyd Algorithm: Bounded-Error Quantum Polynomial Time (BQP)-Completeness of Matrix Inversion and Error Mitigation for Noisy Intermediate-Scale Quantum (NISQ) Devices. Cureus J Comput Sci 3 : es44389-025-00066-8. DOI https://doi.org/10.7759/s44389-025-00066-8

Analytical projections beyond 8 × 8

Hardware experiments in this study were limited to matrices up to 8 × 8 due to circuit-depth constraints, qubit requirements, and fidelity degradation on available NISQ backends. To provide a forward-looking view of resource growth, we therefore report analytical and simulator-assisted projections for larger instances such as 16 × 16 and 32 × 32. These projections are not measured hardware outcomes and should be interpreted only as indicative estimates of how qubit count, controlled-operation depth, and workflow overhead scale with increasing matrix dimension.

Figure *3* summarizes projected qubit requirements and gate-depth growth beyond the hardware-validated regime. The rapid increase in controlled operations and phase-estimation depth highlights why mitigation alone is unlikely to sustain accuracy at larger dimensions. In particular, the resource burden associated with standard QPE grows quickly enough that depth-reduction strategies, including iterative or semiclassical phase estimation, become necessary even before fault-tolerant operation is considered.

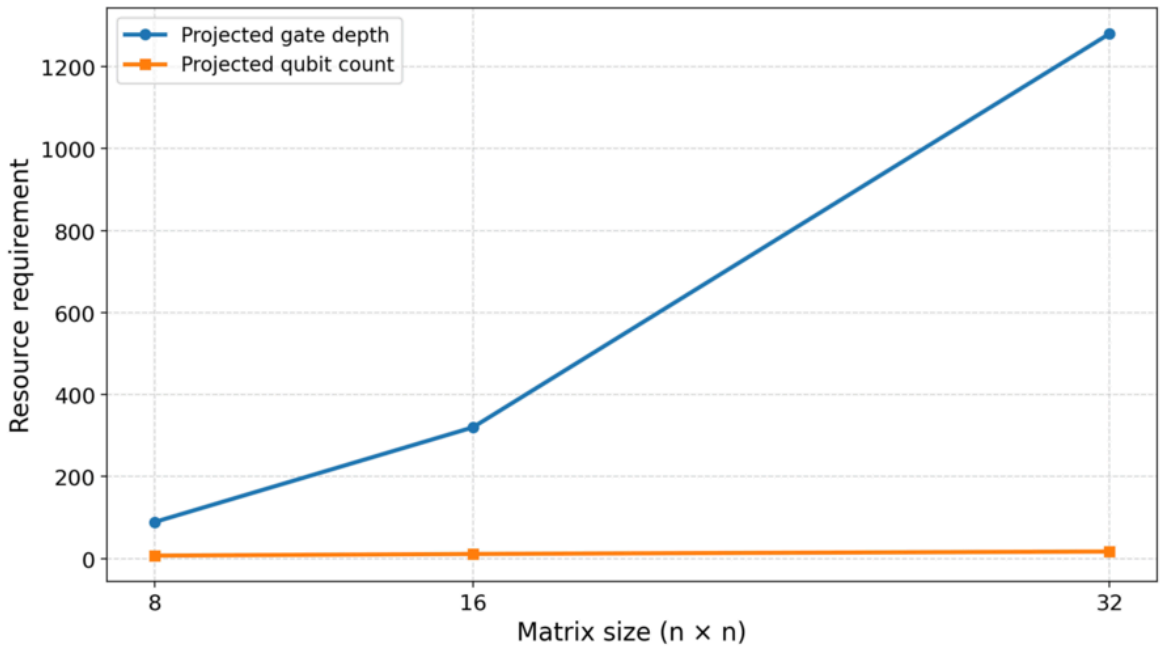

FIGURE 3: Projected gate-depth growth and qubit requirements for matrix sizes beyond 8 × 8, including 16 × 16 and 32 × 32. Values shown are analytical and simulator-assisted estimates, not direct hardware measurements. The figure highlights the rapid growth of circuit resources driven primarily by QPE depth and controlled operations.

QPE, Quantum Phase Estimation

Under standard HHL assumptions, the dimension-dependence of the algorithm is theoretically favorable relative to classical dense solvers; however, on current NISQ devices, practical performance remains dominated by state preparation, compilation, finite coherence, two-qubit gate errors, repetition overhead, and mitigation cost. Therefore, the scaling figures are included to illustrate theoretical and workflow-level trends rather than to claim present-day runtime dominance.

Runtime accounting and comparison caveats

Direct wall-clock comparisons between classical solvers and the HHL workflow must be interpreted with caution. In the present study, reliable empirical quantum execution is limited to small instances, whereas larger quantum values are analytical or simulator-assisted projections. Moreover, HHL produces a quantum state encoding of the solution rather

How to cite this article:
Mayaluri Z, Panda G (May 14, 2026) Hybrid Quantum-Classical Implementation of the Harrow–Hassidim–Lloyd Algorithm: Bounded-Error Quantum Polynomial Time (BQP)-Completeness of Matrix Inversion and Error Mitigation for Noisy Intermediate-Scale Quantum (NISQ) Devices. Cureus J Comput Sci 3 : es44389-025-00066-8. DOI https://doi.org/10.7759/s44389-025-00066-8

than an explicit solution vector, so output costs are not identical across paradigms. Accordingly, we use runtime discussion primarily to illustrate workflow overheads and scaling tendencies, not to claim near-term practical superiority over optimized classical solvers.

The reported quantum workflow times include transpilation, circuit execution, and mitigation-related post-processing, whereas the classical timings correspond only to the numerical solve stage on the stated benchmark environment. These quantities are therefore informative as implementation-level workflow measurements, but they should not be interpreted as a strict like-for-like complexity or end-to-end output comparison.

Figure *4* is retained only as an illustrative scaling figure. Classical curves are empirical CPU measurements for the stated benchmark environment, while quantum curves combine measured small-instance workflow timings with analytical projections beyond the hardware-validated regime. The figure therefore illustrates overhead and scaling tendencies rather than demonstrating present-day quantum advantage.

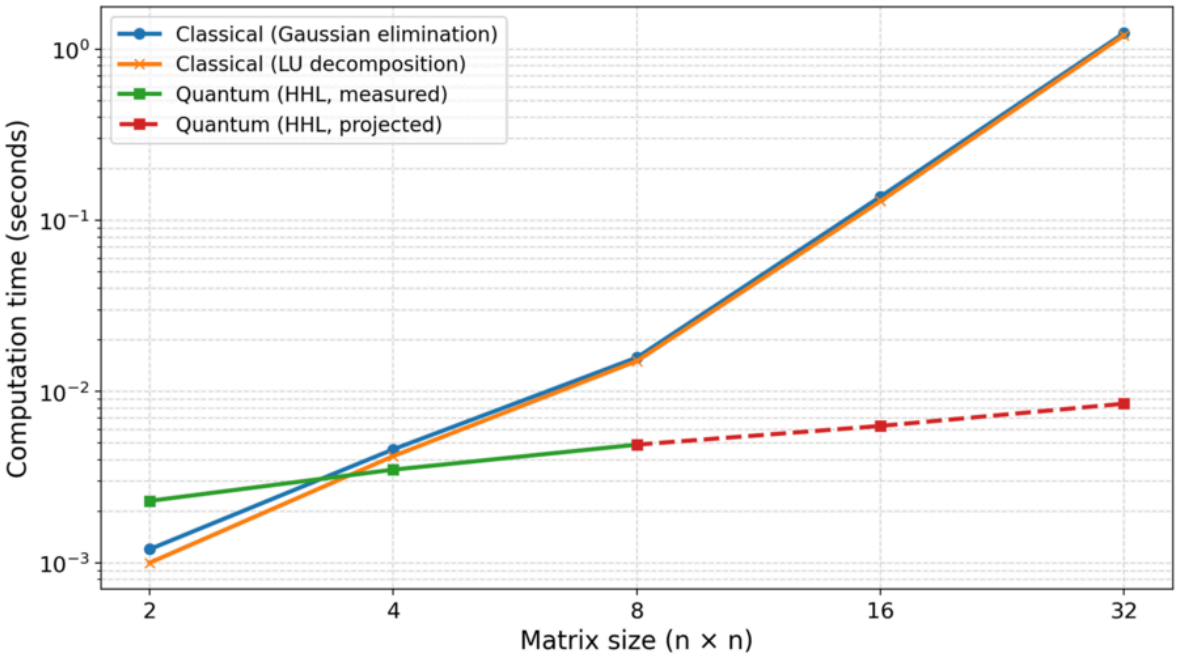

FIGURE 4: Illustrative workflow-scaling trends for classical dense solvers and the HHL pipeline. Classical curves are empirical CPU measurements for the stated benchmark environment, while quantum curves combine measured small-instance workflow timings with analytical projections beyond the hardware-validated regime. The figure illustrates overhead and scaling tendencies, not a demonstration of present-day quantum advantage.

HHL, Harrow–Hassidim–Lloyd

To quantify noise sensitivity more directly, Figure *5* reports the absolute error distribution across matrix sizes. Larger instances exhibit greater susceptibility to two-qubit gate errors, decoherence during QPE, and measurement bias, consistent with the degradation observed in the calibrated-noise and hardware results. Mitigation reduces the error magnitude relative to unmitigated baselines, but the residual spread still grows with matrix size.

How to cite this article:
Mayaluri Z, Panda G (May 14, 2026) Hybrid Quantum-Classical Implementation of the Harrow–Hassidim–Lloyd Algorithm: Bounded-Error Quantum Polynomial Time (BQP)-Completeness of Matrix Inversion and Error Mitigation for Noisy Intermediate-Scale Quantum (NISQ) Devices. Cureus J Comput Sci 3 : es44389-025-00066-8. DOI https://doi.org/10.7759/s44389-025-00066-8

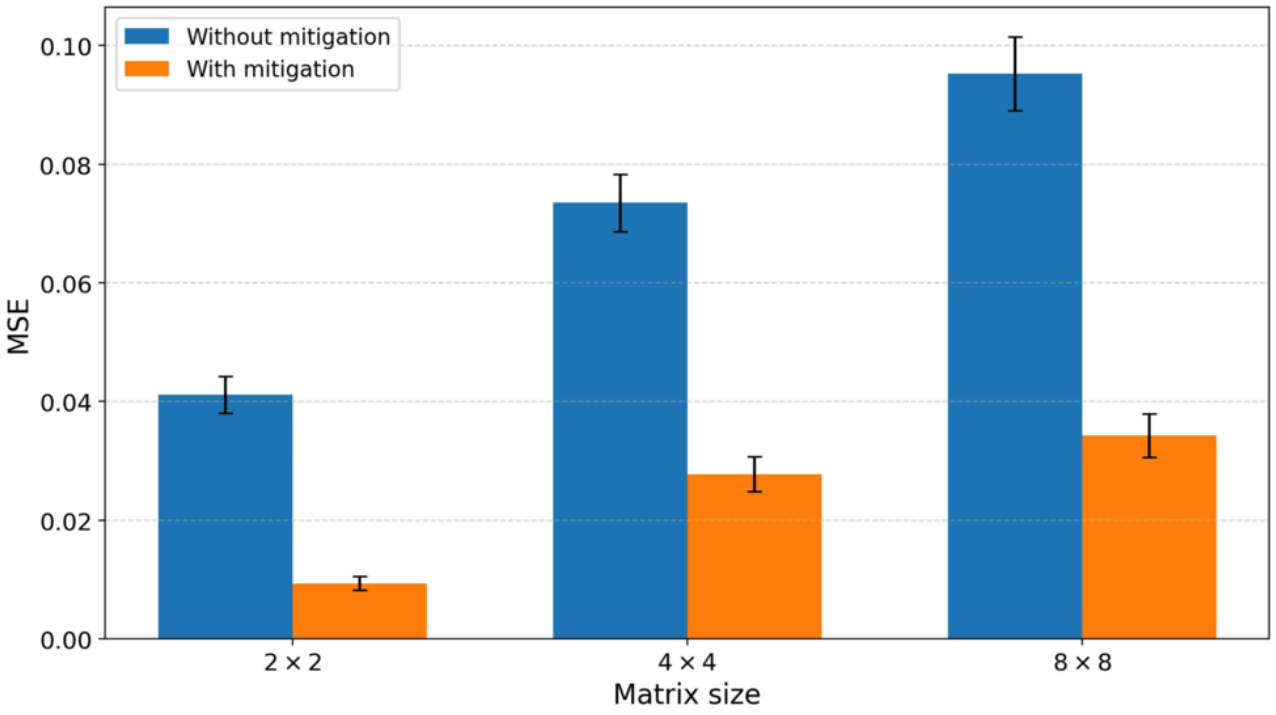

FIGURE 5: Absolute error distribution for 2 × 2, 4 × 4, and 8 × 8 Hermitian matrices. Larger instances exhibit greater error accumulation due to increased QPE depth and multi-qubit gate usage. Mitigation reduces the error distribution relative to unmitigated noisy baselines.

MSE, Mean Squared Error; QPE, Quantum Phase Estimation

The present simulator and hardware results support mitigation-aware execution for small Hermitian instances up to the hardware-validated regime studied here; they do not establish reliable scaling beyond that regime. Scaling to larger systems will require improved two-qubit fidelity, longer coherence times, reduced phase-estimation depth, and ultimately fault-tolerant quantum error correction for sustained accuracy.

Discussion

The experimental results confirm that the HHL algorithm, when coupled with error mitigation, is feasible on current NISQ hardware for small Hermitian systems.

For 2 × 2, 4 × 4, and 8 × 8 matrices, we observed consistent improvements in output fidelity and reductions in MSE after applying noise-aware techniques, including ZNE and measurement-error mitigation. In the 8 × 8 case, mitigation reduced the MSE by more than 60%, while fidelity increased from 82.76% to 92.34%. These results highlight the practical importance of mitigation in stabilizing algorithmic performance under decoherence and gate noise.

Cross-platform evaluation further indicates that hardware characteristics strongly influence achievable accuracy. While IonQ exhibited higher effective gate fidelity and longer coherence, IBM Quantum delivered lower latency, particularly for two-qubit operations. As summarized in the backend-comparison table, IonQ achieved higher output fidelity for the 8 × 8 instance, consistent with its longer relaxation time, whereas IBM Quantum offered an execution-speed advantage that can benefit shallower circuits and latency-sensitive workloads.

Regarding computational time, the present results should be interpreted as feasibility and workflow-overhead evidence rather than as proof of near-term practical quantum advantage. Although HHL has favorable asymptotic dimension-dependence under sparsity, conditioning, and state-preparation assumptions, the current NISQ implementation is

How to cite this article:
Mayaluri Z, Panda G (May 14, 2026) Hybrid Quantum-Classical Implementation of the Harrow–Hassidim–Lloyd Algorithm: Bounded-Error Quantum Polynomial Time (BQP)-Completeness of Matrix Inversion and Error Mitigation for Noisy Intermediate-Scale Quantum (NISQ) Devices. Cureus J Comput Sci 3 : es44389-025-00066-8. DOI https://doi.org/10.7759/s44389-025-00066-8

dominated in practice by compilation overhead, circuit depth, repeated sampling, and mitigation cost. For this reason, measured small-instance timings and larger-instance projections are reported separately, and no claim is made that the present implementation outperforms optimized classical dense solvers in wall-clock terms.

Noise sensitivity increases with circuit depth, which is most evident in the 8×8 experiments. The reconstruction-error results confirm that larger instances accumulate more noise through deeper phase-estimation and controlled-rotation subcircuits. Nevertheless, mitigation reduced the reconstruction error relative to the unmitigated baseline, demonstrating the resilience of the hybrid workflow under realistic noise.

Finally, the scaling analysis indicates that although the current hardware-validated implementation is limited to the small-instance regime due to depth and error accumulation, the overall HHL pipeline remains structurally important from an algorithmic perspective. The dominant bottlenecks for larger instances remain quantum phase estimation depth, two-qubit gate errors, and decoherence, and mitigation alone can only partially compensate. Meaningful scaling beyond the present regime will therefore require improved two-qubit fidelities, longer coherence times, depth-reduced phase-estimation variants, and ultimately fault-tolerant quantum error correction.

Conclusions

This study presents a detailed evaluation of the HHL algorithm on current NISQ devices, emphasizing the role of error mitigation and hardware selection in practical quantum computing. Our findings confirm that, with appropriate mitigation strategies, the HHL algorithm achieves substantial gains in fidelity and accuracy, particularly for matrices up to 8×8. The key conclusions are as follows. First, error mitigation methods, including ZNE and post-selection, significantly enhanced accuracy, reducing the MSE by more than 60% under noisy conditions. Second, the hardware comparison indicates a clear trade-off between performance and precision: IBM Quantum backends offered faster execution but exhibited higher error rates, whereas IonQ provided higher fidelity and longer coherence, albeit with slower gate execution. Third, the study supports the feasibility of small-scale error-mitigated HHL demonstrations on current NISQ hardware, but it does not establish wall-clock superiority over optimized classical matrix-inversion methods. Finally, scalability remains limited by NISQ constraints, notably two-qubit gate errors, circuit-depth growth, particularly in the quantum phase estimation subroutine, and qubit connectivity, which collectively restrict reliable execution beyond 8×8 on current platforms.

Future work will focus on integrating quantum error correction techniques to enable deeper, fault-tolerant HHL circuits and on improving mitigation pipelines through adaptive ZNE and hybrid mitigation-quantum error correction strategies. In addition, we will evaluate the proposed workflow on emerging hardware with higher two-qubit fidelities and longer coherence times to assess feasibility for larger problem instances under realistic execution constraints. Overall, this work demonstrates that hybrid error-mitigated HHL workflows can improve small-scale solution fidelity on current NISQ devices, while also making clear that meaningful scaling will require better hardware, lower-depth phase-estimation variants, and more rigorous benchmarking against classical baselines.

Appendices

Proof sketch of BQP-completeness of matrix inversion

Matrix inversion, in appropriate oracle or access models, is known to be BQP-complete in the sense that (i) a suitably formulated linear-systems problem is solvable in bounded-error quantum polynomial time, and (ii) suitably structured instances can encode the power of general quantum computation via reductions from arbitrary BQP computations. The purpose of this appendix is to provide a reviewer-safe, assumption-explicit proof sketch and intuition, rather than to claim a new complexity-theoretic result. Throughout, the claim should be understood as applying to the quantum state-preparation or access-model formulation of matrix inversion, not to unrestricted classical-output numerical inversion of arbitrary dense matrices.

How to cite this article:
Mayaluri Z, Panda G (May 14, 2026) Hybrid Quantum-Classical Implementation of the Harrow–Hassidim–Lloyd Algorithm: Bounded-Error Quantum Polynomial Time (BQP)-Completeness of Matrix Inversion and Error Mitigation for Noisy Intermediate-Scale Quantum (NISQ) Devices. Cureus J Comput Sci 3 : es44389-025-00066-8. DOI https://doi.org/10.7759/s44389-025-00066-8

Step 1: Problem definition and assumptions: Given a Hermitian matrix $A \in \mathbb{C}^{N\times N}$ and a vector $b \in \mathbb{C}^{N}$, the linear-systems task is to obtain a state proportional to the following expression.

$$x = A^{-1}b, \qquad \text{equivalently} \qquad Ax = b$$

In the quantum setting, the objective is to prepare a quantum state $|x\rangle$ satisfying

$$|x\rangle \propto A^{-1}|b\rangle$$

where $|b\rangle$ is a normalized state encoding the input vector *b*. The standard assumptions underlying HHL-type complexity statements are made explicit: (i) *A* is *s*-sparse, or is provided through an efficient block-encoding or oracle access model; (ii) $\|A\| \leq 1$ after normalization; (iii) *A* has bounded condition number κ; (iv) efficient preparation of $|b\rangle$ is available; and (v) the target precision is ϵ. Under these assumptions, the problem is naturally formulated as state preparation or expectation estimation with respect to $|x\rangle$, rather than as entrywise classical output of the full vector *x*.

Step 2: Membership in BQP (why the problem lies in BQP): The HHL framework prepares $|x\rangle$ using three core subroutines: QPE, eigenvalue-dependent controlled rotation, and inverse QPE. Let

$$A = \sum_i \lambda_i |u_i\rangle\langle u_i|, \qquad |b\rangle = \sum_i \beta_i |u_i\rangle$$

Applying QPE to the Hamiltonian-simulation unitary e^{iAt} produces

$$\sum_i \beta_i |u_i\rangle|0\rangle \xrightarrow{\text{QPE}} \sum_i \beta_i |u_i\rangle|\widetilde{\lambda}_i\rangle$$

where $\widetilde{\lambda}_i$ is an ϵ-accurate estimate of λ_i up to the chosen simulation scaling. A controlled rotation then maps an ancilla qubit according to

$$|\widetilde{\lambda}_i\rangle|0\rangle_a \;\mapsto\; |\widetilde{\lambda}_i\rangle\left(\sqrt{1-(C/\widetilde{\lambda}_i)^2}\,|0\rangle_a + (C/\widetilde{\lambda}_i)\,|1\rangle_a\right)$$

for a normalization constant $C = O(1/\kappa)$ chosen so that the amplitudes remain valid. After inverse QPE, the branch conditioned on measuring the ancilla in $|1\rangle_a$ is proportional to

$$|x\rangle \propto \sum_i \beta_i \lambda_i^{-1} |u_i\rangle$$

Under the assumptions stated above, Hamiltonian simulation, phase estimation, reversible arithmetic, and controlled rotation can be implemented in time polynomial in $\log N$, κ, and $1/\epsilon$. Therefore, the corresponding state-preparation version of matrix inversion lies in BQP.

Step 3: Efficient controlled eigenvalue inversion (resource viewpoint): Once an ϵ-accurate eigenvalue estimate is available in a phase register, implementing the controlled rotation requires approximating λ^{-1} and applying a rotation whose angle depends on that inverse. In standard reversible arithmetic models, the cost of this step is polynomial in the number of precision bits, i.e., polynomial in $\log(1/\epsilon)$, and polynomial in $\log \kappa$ over the admissible eigenvalue range. Thus, controlled eigenvalue inversion does not introduce super-polynomial overhead beyond the assumptions already required for efficient HHL-style execution.

Step 4: Grover-type amplitude amplification as auxiliary intuition: A useful intuition is that amplitude amplification can improve the success probability of the post-selection step in HHL. If the desired branch is obtained only when the ancilla is measured in $|1\rangle_a$, then in the worst case the success probability can scale as $p = \Omega(1/\kappa^2)$. Grover-type amplitude amplification can increase this success probability quadratically, reducing the number of repetitions from $O(1/p)$ to $O(1/\sqrt{p})$, i.e., from $O(\kappa^2)$ to $O(\kappa)$, while preserving bounded error.

How to cite this article:
Mayaluri Z, Panda G (May 14, 2026) Hybrid Quantum-Classical Implementation of the Harrow–Hassidim–Lloyd Algorithm: Bounded-Error Quantum Polynomial Time (BQP)-Completeness of Matrix Inversion and Error Mitigation for Noisy Intermediate-Scale Quantum (NISQ) Devices. Cureus J Comput Sci 3 : es44389-025-00066-8. DOI https://doi.org/10.7759/s44389-025-00066-8

This point should be interpreted carefully. Amplitude amplification is not the basis of the BQP-hardness claim, and Grover search is not being used here as a literal completeness reduction. Rather, it provides reviewer-safe intuition that the HHL pipeline can incorporate standard BQP mechanisms---phase estimation, controlled arithmetic, and success-probability amplification---within a single linear-systems primitive.

Step 5: Completeness argument (why matrix inversion can capture BQP computations): To support BQP-completeness, one must show that any language $L \in \mathbf{BQP}$ can be reduced in polynomial time to an instance of linear-systems solving in an appropriate access model. At a high level, an arbitrary quantum circuit U acting on poly(n) qubits can be embedded into a structured linear-algebraic instance by constructing a sparse matrix whose inverse encodes amplitudes or acceptance information associated with the computation. Such constructions typically use a history-state, Feynman-Kitaev-type, or block-matrix embedding so that solving the corresponding linear system, or preparing a related quantum state, allows one to distinguish whether the original circuit accepts with probability at least 2/3 or at most 1/3.

Under standard sparsity, normalization, and access assumptions, this yields a polynomial-time reduction from general BQP computation to a structured linear-systems problem, thereby establishing BQP-hardness for the corresponding formulation of matrix inversion. Combined with Step 2, this supports the standard statement that matrix inversion is BQP-complete in suitable oracle or access models.

In summary, the linear-systems problem is placed in BQP by the HHL pipeline (QPE → eigenvalue-dependent controlled rotations → inverse QPE and post-selection) under standard sparsity, conditioning, precision, and state-preparation assumptions. BQP-hardness follows from known reductions that embed arbitrary quantum computations into structured linear-systems instances in suitable models. Accordingly, matrix inversion serves as a complexity-theoretically significant primitive for quantum computing, while the present manuscript uses that fact only for assumption-explicit positioning rather than for any new complexity claim.

Additional Information

Author Contributions

All authors have reviewed the final version to be published and agreed to be accountable for all aspects of the work.

Concept and design: Zefree Lazarus Mayaluri, Ganapati Panda

Acquisition, analysis, or interpretation of data: Zefree Lazarus Mayaluri, Ganapati Panda

Drafting of the manuscript: Zefree Lazarus Mayaluri, Ganapati Panda

Critical review of the manuscript for important intellectual content: Zefree Lazarus Mayaluri, Ganapati Panda

Disclosures

Human subjects: All authors have confirmed that this study did not involve human participants or tissue. **Animal subjects:** All authors have confirmed that this study did not involve animal subjects or tissue. **Conflicts of interest:** In compliance with the ICMJE uniform disclosure form, all authors declare the following: **Payment/services info:** All authors have declared that no financial support was received from any organization for the submitted work. **Financial relationships:** All authors have declared that they have no financial relationships at present or within the previous three years with any organizations that might have an interest in the submitted work. **Other relationships:** All authors have declared that there are no other relationships or activities that could appear to have influenced the submitted work.

References

1. Harrow AW, Hassidim A, Lloyd S: Quantum algorithm for linear systems of equations. Physical Review Letters. 2009, 103:150502. 10.1103/physrevlett.103.150502

2. Preskill J: Quantum computing in the NISQ era and beyond. Quantum. 2018, 2:79. 10.22331/q-2018-08-06-79

How to cite this article:
Mayaluri Z, Panda G (May 14, 2026) Hybrid Quantum-Classical Implementation of the Harrow–Hassidim–Lloyd Algorithm: Bounded-Error Quantum Polynomial Time (BQP)-Completeness of Matrix Inversion and Error Mitigation for Noisy Intermediate-Scale Quantum (NISQ) Devices. Cureus J Comput Sci 3 : es44389-025-00066-8. DOI https://doi.org/10.7759/s44389-025-00066-8

3. Grover LK: A fast quantum mechanical algorithm for database search. STOC '96: Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing. 1996, 96:212-219. 10.1145/237814.237866
4. Nielsen MA, Chuang IL: Quantum Computation and Quantum Information. Cambridge University Press, Cambridge; 2000. 10.1017/CBO9780511976667
5. Zhou J: Quantum finance: exploring the implications of quantum computing on financial models. Computational Economics. 2025, 67:1043-1072. 10.1007/s10614-025-10894-4
6. Childs AM, Kothari R, Somma RD: Quantum algorithm for systems of linear equations with exponentially improved dependence on precision. SIAM Journal on Computing. 2017, 46:1920-1950. 10.1137/16m1087072
7. Schuld M, Sinayskiy I, Petruccione F: Prediction by linear regression on a quantum computer. Physical Review A. 2016, 94:022342. 10.1103/physreva.94.022342
8. Xiong Y: Quantum error mitigation for error-resilient quantum computation. University of Southampton, Doctoral Thesis. 2022.
9. Cai X-D, Weedbrook C, Su Z-E, et al.: Experimental quantum computing to solve systems of linear equations. Physical Review Letters. 2013, 110:230501. 10.1103/physrevlett.110.230501
10. Pan J, Cao Y, Yao X, et al.: Experimental realization of quantum algorithm for solving linear systems of equations. Physical Review A. 2014, 89:022313. 10.1103/physreva.89.022313
11. Zheng Y, Song C, Chen MC, et al.: Solving systems of linear equations with a superconducting quantum processor. Physical Review Letters. 2017, 118:210504. 10.1103/physrevlett.118.210504
12. Lee Y, Joo J, Lee S: Hybrid quantum linear equation algorithm and its experimental test on IBM Quantum Experience. Scientific Reports. 2019, 9:4778. 10.1038/s41598-019-41324-9
13. Zhang M, Dong L, Zeng Y, Cao N: Improved circuit implementation of the HHL algorithm and its simulations on QISKIT. Scientific Reports. 2022, 12:13287. 10.1038/s41598-022-17660-8
14. Yalovetzky R, Minssen P, Herman D, Pistoia M: Solving linear systems on quantum hardware with hybrid HHL++. Scientific Reports. 2024, 14:20610. 10.1038/s41598-024-69077-0
15. Morgan J, Ghysels E, Mohammadbagherpoor H: An enhanced hybrid HHL algorithm. Physics Letters A. 2025, 532:130181. 10.1016/j.physleta.2024.130181

How to cite this article:
Mayaluri Z, Panda G (May 14, 2026) Hybrid Quantum-Classical Implementation of the Harrow–Hassidim–Lloyd Algorithm: Bounded-Error Quantum Polynomial Time (BQP)-Completeness of Matrix Inversion and Error Mitigation for Noisy Intermediate-Scale Quantum (NISQ) Devices. Cureus J Comput Sci 3 : es44389-025-00066-8. DOI https://doi.org/10.7759/s44389-025-00066-8