

Exploratory Quantum Linear-Solver Application to an MHD Nanofluid Couette Flow Problem: HHL and VQLS Benchmarked Against a Validated Classical Baseline

S. I. Mosala^{1,2}, O. D. Makinde³, and A. Muronga¹

¹Department of Physics, Nelson Mandela University, Gqeberha, South Africa

²Corresponding author: mrmosalas.i@gmail.com

³Faculty of Military Science, Stellenbosch University, Saldanha, South Africa

Motivation and Scope

Magnetohydrodynamic (MHD) nanofluid flows are governed by coupled, nonlinear momentum and energy equations whose discretisation produces large, structured (typically tridiagonal or banded) linear systems at every time step. Quantum linear-solver algorithms, notably the Harrow–Hassidim–Lloyd (HHL) algorithm and the Variational Quantum Linear Solver (VQLS), offer a theoretically distinct route to solving such systems, but their practical behaviour on physically realistic, sparse, moderately conditioned matrices arising from thermo-fluid problems is not yet well characterised. This work reports an exploratory application of HHL and VQLS to the linear systems arising from an unsteady MHD nanofluid Couette flow problem, with results benchmarked against classical solvers and, critically, against an independently derived closed-form exact solution.

We emphasise that this is a linear-algebra-level benchmark: the quantum algorithms are applied to the *discretised* system after classical formulation of the governing partial differential equations, not to a quantum reformulation of the PDEs themselves. This scope is stated explicitly, and the exact-solution validation described below is used specifically to give the quantum results a rigorous, physically grounded point of comparison rather than a purely internal one.

Problem and Classical Baseline

The flow considered is unsteady, incompressible, electrically conducting nanofluid Couette flow between two parallel plates under partial slip and convective heat exchange, in the presence of a transverse magnetic field. In dimensionless form, the momentum equation is

$$A_1 \frac{\partial^2 W}{\partial \eta^2} - A_2 Ha^2 W + G = A_4 \frac{\partial W}{\partial \tau}, \quad (1)$$

where $W(\eta, \tau)$ is the dimensionless velocity, Ha the Hartmann number, G the pressure-gradient parameter, and A_1, A_2, A_4 are nanofluid property ratios. The energy equation is coupled through viscous dissipation and Joule heating source terms. The system is solved classically using both a Bivariate Spectral Quasilinearisation Method (BI-SQLM) and a Crank–Nicolson (CN) finite-difference scheme; the latter is used specifically to export the linear system at each time step for the quantum solvers, since its structured, tridiagonal form is directly amenable to sparse Hamiltonian simulation.

As an independent validation step, the steady-state form of the momentum equation admits a closed-form exact solution,

$$W(\eta) = C_1 \cosh(k\eta) + C_2 \sinh(k\eta) + \frac{G}{A_2 H a^2}, \quad k = H a \sqrt{A_2/A_1}, \quad (2)$$

with C_1, C_2 determined by the no-slip and Navier-slip boundary conditions. This solution was derived and symbolically verified (substituted back into the governing ODE and both boundary conditions, confirming exact satisfaction to machine precision) and used to validate BI-SQLM, CN, and (for the first time in this study) the quantum solvers directly, over the full domain $\eta \in [0, 1]$.

Quantum Implementation

Both HHL and VQLS were implemented in Qiskit (Aer simulator backend) and applied to small linear systems ($n = 4$) extracted from the CN discretisation, encoded as Hermitian matrices via symmetric preprocessing. HHL was implemented with quantum phase estimation using an 8-qubit clock register (11 qubits total, 2 for the solution register), giving a statevector relative error of 0.41% against the classical reference, but a transpiled circuit depth of 43,127 gates, roughly 86 times the coherence budget of current superconducting hardware, making near-term direct deployment infeasible. VQLS was implemented with a 2-layer hardware-efficient ansatz (2 qubits total) and a COBYLA classical optimiser minimising the standard Hadamard-test cost function; it converged to a relative error of 0.007% at a transpiled depth of only 5 gates, a $58\times$ reduction in circuit depth relative to HHL for comparable or better accuracy, and well within near-term hardware constraints.

To assess whether these results depend on the specific demonstration case, an additional, harder-conditioned system ($\kappa = 8.46$, versus $\kappa = 5.92$ for the original demonstration) was constructed by assembling a symmetric-by-construction Galerkin finite-element discretisation directly over the full domain, and solving it with the same unmodified HHL and VQLS circuits. Against the independently derived exact solution, VQLS again substantially outperformed HHL (maximum error 0.032% versus 2.84%), indicating that its advantage persists as conditioning becomes more demanding.

Summary of Results

Metric	Classical	HHL	VQLS
Relative error (vs. classical, $n = 4$ demo)	—	0.41%	0.007%
Transpiled circuit depth	—	43,127	5
Total qubits	—	11	2
NISQ hardware feasible	—	No	Yes
Max. error vs. exact solution (full domain, $\kappa = 8.46$)	0.004%	2.84%	0.032%

Table 1: Summary comparison across classical and quantum solvers.

A condition-number sweep over Hartmann number ($Ha = 0\text{--}7$) at fixed grid resolution shows $\kappa(\mathbf{A})$ varying by less than 0.5%, indicating that for this problem class, quantum resource requirements are governed primarily by spatial discretisation rather than magnetic field strength, a finding relevant to preconditioning strategy design for future, larger-scale quantum treatments of this problem class.

Conclusion and Outlook

VQLS demonstrates consistent, hardware-relevant advantage over HHL for this class of thermo-fluid linear system, both in a favourable demonstration case and under harder, more representative conditioning, with accuracy independently confirmed against a closed-form exact solution rather than only self-consistently against classical solvers. HHL retains a theoretical asymptotic advantage that is expected to dominate at larger problem sizes, but remains hardware-infeasible at current circuit depths for this problem. This exploratory linear-solver-level study establishes a validated foundation; a full quantum reformulation of the governing MHD nanofluid PDEs is identified as the primary direction for future work.