Scalable Variational Quantum Linear Solvers for CFD via an Efficient Laplacian Block Encoding and CMA-ES Optimization

Viraj Dsouza[*,1], Alex Khan[*,2], Rut Lineswala[*,3] and Abhishek Chopra[*,4]
[*] BosonQ Psi Corp., Syracuse, New York, 13202, USA.

Abstract

Solving the pressure-Poisson equation efficiently is a central computational bottleneck in incompressible Computational Fluid Dynamics (CFD), and quantum linear solvers have been identified as a potential long-term avenue for addressing this challenge. Building on our prior work [1], we present two concrete improvements to the Variational Quantum Linear Solver (VQLS) for structured CFD linear systems.

First, we integrate a new unified block-encoding framework for discrete Laplacian operators [2], a companion contribution by the authors, which reduces two-qubit gate counts by up to 3.2x and circuit depth by up to 2.5x relative to [1], while also improving post-selection success probability. This enables near-exact recovery of the classical solution at $N = 128$, compared to $N = 32$ in our prior work.

Second, we show that the Covariance Matrix Adaptation Evolution Strategy (CMA-ES) outperforms COBYLA as a VQLS optimizer at scale: both converge for $N \leq 64$, but CMA-ES continues to converge at $N = 128$ where COBYLA stalls, and at $N = 256$ CMA-ES reduces the cost function to $C_G \approx 0.59$ while COBYLA makes no progress beyond $C_G \approx 0.90$.

The methodology is validated on two CFD-relevant benchmarks: a 2D Poisson problem with sinusoidal forcing, representative of incompressible pressure-Poisson subproblems, and a steady-state heat conduction problem with a localized Gaussian source term. All results are obtained on the Qiskit Aer statevector simulator, and together establish a clearer pathway toward practical quantum-assisted CFD solvers.

Keywords: Variational quantum linear solver, Block encoding, Discrete Laplacian, Pressure-Poisson equation, CMA-ES, Barren plateau, Hybrid quantum-classical computing, Computational fluid dynamics.

1. Introduction

Computational Fluid Dynamics (CFD) relies heavily on the repeated solution of large sparse linear systems arising from the discretization of partial differential equations. In incompressible flow solvers, projection and fractional-step methods require the solution of a pressure-Poisson equation at every time step, making elliptic linear solvers a central computational component of many CFD workflows. Although modern high-fidelity methods such as Large-Eddy Simulation (LES) and Direct Numerical Simulation (DNS) have achieved remarkable predictive capability, their computational cost remains a major challenge for industrial-scale simulations and certification workflows [3,4]. The continued growth in computational demand has motivated investigation into emerging computational paradigms, and the CFD Vision 2030 studies identify quantum computing as a potential long-term research direction for future CFD capabilities [5,6].

Quantum approaches to computational fluid dynamics (QCFD) have developed along several directions, including Schrödinger-type formulations [7], Quantum Lattice Boltzmann Methods (QLBM) [8,9], and

[1]Quantum Computing Researcher, BosonQ Psi Corp.
[2]VP of Quantum R&D, BosonQ Psi Corp.
[3]Founder and CTO, BosonQ Psi Corp.
[4]Founder and CEO, BosonQ Psi Corp.

quantum linear algebra approaches for solving the sparse systems arising from PDE discretizations. Among these, direct matrix inversion methods are particularly appealing because they admit theoretically efficient quantum algorithms. The theoretical foundation for quantum linear solvers was established by the Harrow–Hassidim–Lloyd (HHL) algorithm [10], which demonstrated an exponential speedup over classical methods for sparse, well-conditioned systems under idealized assumptions on matrix sparsity, condition number, and data access. While HHL has been demonstrated experimentally on small systems [11], its practical deployment requires deep circuits and fault-tolerant quantum hardware, which current industry roadmaps project to remain several years away [12–14]. Similar limitations apply to Quantum Singular Value Transformation (QSVT) [15], which provides a unified framework for matrix inversion, Hamiltonian simulation, and related linear algebra operations [16].

For the near-term Noisy Intermediate-Scale Quantum (NISQ) era [17], the Variational Quantum Linear Solver (VQLS) [18] has emerged as the most practically accessible quantum linear solver. VQLS recasts the linear system $A\boldsymbol{x} = \boldsymbol{b}$ as a variational optimization problem: a classical optimizer iteratively updates the parameters $\boldsymbol{\theta}$ of a parametrized quantum circuit $V(\boldsymbol{\theta})$ to minimize a cost function that quantifies the mismatch between $A\,|x(\boldsymbol{\theta})\rangle$ and $|b\rangle$. The approach is hardware-friendly in that it requires only shallow, parametrized circuits and classical post-processing, and it has been explored for several PDE systems relevant to CFD, including the heat conduction equation [1, 19] and the Poisson equation [20].

A fundamental challenge shared by VQLS and quantum linear solvers more broadly is the classical readout bottleneck: the quantum computer encodes the solution as a quantum state $|x\rangle$, and extracting the full classical solution vector requires $O(N)$ measurements, which eliminates any potential exponential speedup [21]. Practical quantum advantage for CFD therefore requires either that the workflow be restructured to query low-dimensional functionals of the solution such as forces, fluxes, or global energy norms rather than the full field, or that the quantum linear solver be embedded in a larger quantum pipeline that never requires classical readout of the full state. Both directions remain active areas of research, and the present work does not resolve them. Rather, our focus is on building the algorithmic and circuit-level infrastructure- efficient block-encoding circuits, scalable optimizers, and validated benchmark results, that any future quantum CFD solver will require, and on characterizing the current performance limits of VQLS on structured CFD problems.

Despite its promise for the NISQ era, the practical scalability of VQLS for CFD applications remains limited by two major challenges. First, the standard VQLS formulation [18] requires the system matrix A to be expressed as a linear combination of unitaries (LCU), with the cost function estimated through $O(L^2)$ independent Hadamard-test circuits, where L is the number of terms in the LCU decomposition. For structured sparse matrices arising from finite-difference discretizations of the Poisson operator, L grows polynomially with the matrix dimension N, leading to prohibitive circuit-evaluation costs as the system size increases [1]. Second, VQLS training is affected by barren plateaus [22], where the cost-function gradients become exponentially suppressed with increasing qubit number, causing both gradient-based and common gradient-free optimizers to stagnate. The severity of this phenomenon is well established theoretically: for a broad class of parametrized quantum circuits, the variance of the cost-function gradient decays exponentially with the number of qubits, making it exponentially unlikely to sample a parameter configuration with a non-negligible gradient [22, 23]. In practice, this manifests as the optimizer appearing to make no progress despite continued circuit evaluations, and it is particularly acute for the deep circuits and large qubit counts encountered in CFD-scale Poisson problems.

In our previous work [1], we explored potential strategies for mitigating these limitations in structured CFD matrices using an oracle-based block-encoding construction [24, 25] that replaces the Pauli decomposition and reduces the number of quantum circuit evaluations per cost-function call from $O(L^2)$ to one. To improve trainability, a Quantum-Inspired Evolutionary Optimizer (QIEO) [26] was introduced in place of COBYLA, showing encouraging convergence behavior on 1D and 2D Poisson benchmark systems. Those results demonstrated that block encoding is a viable strategy for reducing the quantum circuit overhead in VQLS for CFD, but also identified that the block-encoding circuit itself and the choice of optimizer remain important factors limiting scalability.

The present work extends [1] along two complementary directions aimed at improving the practicality and scalability of VQLS-based CFD solvers. First, we replace the specialized tridiagonal/pentadiagonal block-encoding constructions used in [1] with a recently developed unified framework for discrete Laplacian operators [2], supporting Dirichlet, Neumann, and periodic boundary conditions across arbitrary spatial dimensions within a single modular circuit. This construction, developed by the authors as a companion contribution [2], achieves an exact $(1, m, 0)$-block encoding with $m = 3 + \lceil \log_2 D \rceil$ ancilla qubits for a D-dimensional problem, and gate complexity scaling as $O(\log N \cdot \log D)$. In the present study we integrate this construction into the VQLS workflow for the two-dimensional Poisson problem with Dirichlet boundary conditions. Second, motivated by the optimization difficulties associated with barren-plateau landscapes, we investigate the Covariance Matrix Adaptation Evolution Strategy (CMA-ES) [27, 28] as an alternative to COBYLA. CMA-ES is a population-based, derivative-free optimizer that adapts the full covariance matrix of its search distribution, enabling it to locate and follow narrow descent directions in high-dimensional flat cost landscapes. Through a comparative study across system sizes from $N = 4$ to $N = 256$, we show that CMA-ES delays the onset of barren-plateau stagnation and remains effective in regimes where COBYLA makes no progress.

The methodology is validated on two benchmark problems of direct relevance to CFD. The first is a two-dimensional Poisson problem with sinusoidal forcing $f(x, y) = \sin(\pi x)\sin(\pi y)$, representative of the pressure-Poisson subproblem that arises in incompressible projection-method solvers and widely used as a standard PDE benchmark due to its known analytical solution. The second is a steady-state heat conduction problem with a localized Gaussian source term, modelling a heat-generating component embedded in a conducting plate and representative of elliptic diffusion problems encountered in thermal management applications. Both problems are computationally tractable by classical methods and therefore provide meaningful references for assessing quantum solution accuracy.

The remainder of this paper is structured as follows. Section 2 introduces the benchmark problem used to validate the proposed approach, reviews the block encoding of the discretized 2D Laplacian for the Dirichlet case, presents the VQLS framework with the block encoding, and discusses the barren-plateau phenomenon and its implications for optimizer design. Section 3 presents the numerical results: cost-function convergence and solution accuracy for the two grid configurations, a comparative study of COBYLA and CMA-ES, and circuit resource estimates. Section 4 summarises the findings and outlines directions for future work.

2. VQLS via Block Encoding: Theory and Methodology

2.1 The 2D Pressure-Poisson Problem

In incompressible CFD, projection and fractional-step methods enforce the divergence-free condition on the velocity field by solving a pressure-Poisson equation of the form

$$\nabla^2 p = \frac{\rho}{\Delta t} \nabla \cdot \mathbf{u}^*, \tag{1}$$

where $\mathbf{u}^*$ is the intermediate (non-solenoidal) velocity field, ρ is the fluid density, and Δt is the time step. The right-hand side represents the local divergence error introduced during the velocity prediction step, and solving (1) yields the pressure correction that restores incompressibility. This elliptic subproblem is one of the dominant computational bottlenecks in many CFD workflows, and it is therefore a natural and well-studied benchmark for linear solvers.

Rather than simulating the complete Navier–Stokes system, we isolate the elliptic subproblem and consider the general two-dimensional Poisson equation on the unit square with homogeneous Dirichlet boundary conditions,

$$-\nabla^2 u(x, y) = f(x, y), \quad (x, y) \in [0, 1]^2, \qquad u\big|_{\partial\Omega} = 0, \tag{2}$$

where u is the unknown scalar field and f is a prescribed source term. Depending on the physical application, u may represent a pressure field arising from an incompressible flow projection step or a temperature field arising from steady-state heat conduction. The specific source terms used to benchmark the VQLS solver in this work are described in Section 3.

The unit square is discretized on a uniform Cartesian grid with N_x interior points along x and N_y interior points along y, with uniform mesh spacings $\Delta x = 1/(N_x + 1)$ and $\Delta y = 1/(N_y + 1)$. Dirichlet boundary conditions are imposed on all four domain boundaries, with the boundary values not necessarily restricted to zero. As a result, only the $N = N_x \times N_y$ interior grid points are treated as unknowns. Applying a standard second-order central finite-difference stencil to the interior nodes yields the linear system

$$A\,\mathbf{x} = \mathbf{b}, \tag{3}$$

where $\mathbf{x} \in \mathbb{R}^N$ contains the interior pressure unknowns, and $\mathbf{b} \in \mathbb{R}^N$ contains the discretized forcing term together with any boundary contributions arising from the finite-difference stencil. The system matrix $A \in \mathbb{R}^{N\times N}$ inherits the separable Kronecker-sum structure of the Laplacian operator,

$$A = I_{N_y} \otimes T_{N_x} + T_{N_y} \otimes I_{N_x}, \tag{4}$$

where I_{N_d} denotes the $N_d \times N_d$ identity matrix and $T_{N_d} \in \mathbb{R}^{N_d\times N_d}$ is the standard one-dimensional second-order finite-difference Laplacian,

$$T_{N_d} = \frac{1}{(\Delta x_d)^2}\begin{pmatrix} -2 & 1 & & \\ 1 & -2 & 1 & \\ & \ddots & \ddots & \ddots \\ & & 1 & -2 \end{pmatrix} \in \mathbb{R}^{N_d\times N_d}. \tag{5}$$

Here $\Delta x_d = \{\Delta x, \Delta y\}$ based on the d involved. The matrix A is sparse and symmetric. Its sparsity pattern reflects the standard five-point finite-difference stencil connecting each interior node to its four nearest neighbours. In the present implementation, the negative sign associated with the continuous Poisson operator is absorbed into the right-hand side vector.

The block-encoding framework requires $\|A\|_2 \le 1$. We therefore normalize the system matrix by its spectral norm,

$$A' = \frac{A}{\|A\|_2}, \tag{6}$$

and rescale the right-hand side accordingly, $\mathbf{b}' = \mathbf{b}/\|A\|_2$. The normalized system $A'\mathbf{x} = \mathbf{b}'$ is then passed to the VQLS circuit. The spectral norm $\|A\|_2$ can be computed classically and scales as $O(1/h^2_{\min})$, where $h_{\min} = \min(\Delta x, \Delta y)$.

We consider a range of proof-of-concept discretization sizes spanning from $N = 4$ to $N = 256$ interior unknowns, corresponding to uniform grids ranging from 2×2 to 16×16 interior points. These problem sizes are intentionally modest and are chosen to remain compatible with the qubit and circuit-depth limitations of current quantum simulators. In all cases, the classical reference solution is obtained using the Successive Over-Relaxation (SOR) iterative method, against which the VQLS solution is compared to assess solution accuracy.

2.2 Block Encoding of the 2D Dirichlet Laplacian

Block encoding provides the standard mechanism for embedding a generally non-unitary matrix into a larger unitary operator that can be realised as a quantum circuit. Formally, a unitary $U_A \in \mathbb{C}^{2^{n+m}\times 2^{n+m}}$ acting on m ancilla qubits and n system qubits is called an (α, m, ε)-block encoding of $A \in \mathbb{C}^{2^n\times 2^n}$ if

$$\left\|A - \alpha\left(\langle 0|^{\otimes m} \otimes I\right) U_A \left(|0\rangle^{\otimes m} \otimes I\right)\right\|_2 \le \varepsilon, \tag{7}$$

where $\alpha > 0$ is the sub-normalization factor and $\varepsilon \geq 0$ is the approximation error. When $\varepsilon = 0$ the encoding is said to be *exact*. In block-matrix form, U_A takes the structure

$$U_A = \begin{pmatrix} A/\alpha & * \\ * & * \end{pmatrix}, \tag{8}$$

where the asterisks denote irrelevant blocks. Post-selecting the ancilla register on the outcome $|0\rangle^{\otimes m}$ projects the system register onto a state proportional to $A|\psi\rangle$, with success probability $\|A|\psi\rangle\|^2/\alpha^2$.

The block-encoding requirement $\|A\|_2 \leq 1$ is not satisfied by the system matrix A introduced in (4), whose spectral norm grows as $O(1/h_{\min}^2)$. We therefore introduce the globally scaled 2D Dirichlet Laplacian

$$\widetilde{A} := \frac{\Lambda}{4} A, \qquad \Lambda := \left(\frac{1}{\Delta x^2} + \frac{1}{\Delta y^2}\right)^{-1}, \tag{9}$$

which satisfies $\|\widetilde{A}\|_2 \leq 1$ by construction. The Kronecker-sum structure of (4) is preserved under this scaling, giving

$$\widetilde{A} = \omega_1 \widetilde{A}^{(N_x)} \otimes I_{N_y} + \omega_2 I_{N_x} \otimes \widetilde{A}^{(N_y)}, \tag{10}$$

where $\widetilde{A}^{(N_d)} := \frac{\Delta x_d^2}{4} T_{N_d}$ is the scaled 1D component along dimension d, satisfying $\|\widetilde{A}^{(N_d)}\|_2 \leq 1$, and the dimension-dependent weights

$$\omega_1 = \frac{1/\Delta x^2}{1/\Delta x^2 + 1/\Delta y^2}, \qquad \omega_2 = 1 - \omega_1, \tag{11}$$

satisfy $\omega_1 + \omega_2 = 1$. The linear system to be solved by VQLS is accordingly $\widetilde{A}\,\mathbf{x} = \widetilde{\mathbf{b}}$, where $\widetilde{\mathbf{b}} = (\Lambda/4)\,\mathbf{b}$. For a uniform grid, $\Delta x = \Delta y$, hence we get $\omega_1 = \omega_2 = 1/2$.

We employ the unified construction of [2] to build an explicit quantum circuit that block encodes $\widetilde{A}$. The construction exploits the Kronecker-sum separability of (10) by introducing a one-qubit *selector* register $|k\rangle$ that is prepared in the superposition

$$U_{\text{prep}_k}|0\rangle = \sqrt{\omega_1}|0\rangle + \sqrt{\omega_2}|1\rangle, \tag{12}$$

implemented by a single $R_y(\theta_0)$ rotation with $\theta_0 = 2\arccos\sqrt{\omega_1}$. Conditioned on $|k\rangle = |d\rangle$ ($d \in \{1, 2\}$), a 1D Dirichlet block-encoding sub-circuit U_d is applied to the system register $|j^{(d)}\rangle$, after which $U_{\text{prep}_k}^{\dagger}$ uncomputes the selector.

Each 1D sub-circuit U_d is itself constructed from cyclic shift operators $S^-|j\rangle = |j - 1 \bmod N_d\rangle$ and $S^+|j\rangle = |j + 1 \bmod N_d\rangle$, a Hadamard–$Z$–Hadamard ladder on two ancilla qubits $|\ell_0\rangle$, $|\ell_1\rangle$, and one additional $|\text{del}\rangle$ ancilla that suppresses the unwanted wrap-around coupling at the Dirichlet boundaries $j = 0$ and $j = N_d - 1$; see [2] for the full construction and proof.

The complete circuit, shown in Figure 1, uses $m = 4$ ancilla qubits in total and constitutes an exact $(1, 4, 0)$-block encoding of $\widetilde{A}$.

Following the analytical resource estimates of [2], the T-gate count of the 2D Dirichlet block-encoding circuit scales as $O(\log N \cdot \log D)$, which for fixed dimension $D = 2$ reduces to $O(\log N)$. This poly-logarithmic scaling in the matrix dimension is the central efficiency property of the construction.

In our prior work [1], the 2D Laplacian was block encoded using the pentadiagonal oracle construction of [25], which required $m = 5$ ancilla qubits and a circuit depth that, while also poly-logarithmic in N, was not formulated within a unified framework and did not fully exploit the Kronecker-sum separability of the operator. The construction of [2] improves on this in three concrete ways that are directly relevant to the present VQLS application. **(i)** The post-selection success probability $p_{\text{succ}} = \|\widetilde{L}^{(2)}|\psi\rangle\|^2$ is optimal compared to previous explicit constructions [25]. **(ii)** The circuit-level benchmarks after transpilation

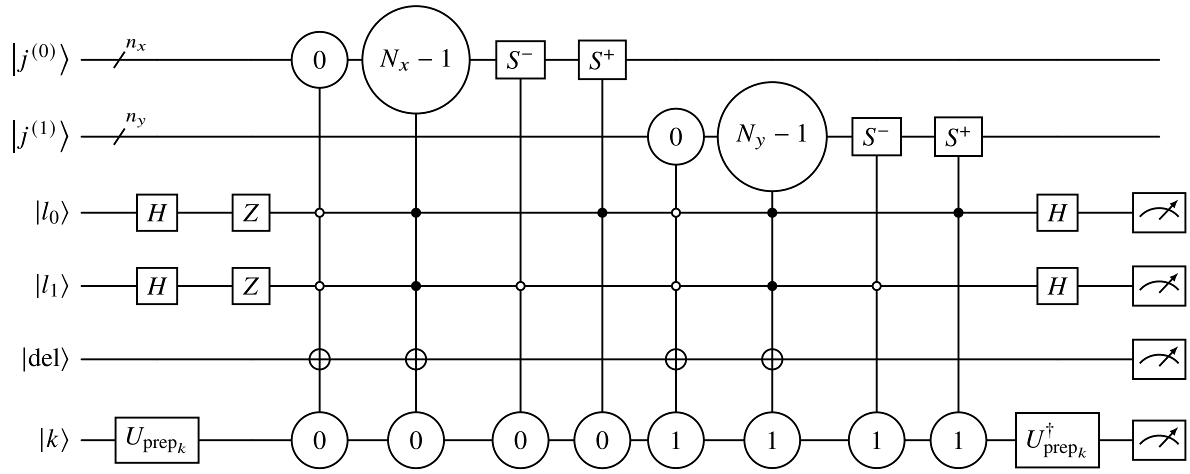

Figure 1: Block encoding circuit for the 2D Laplacian with Dirichlet boundary conditions along both spatial directions, following the construction of [2]. Post-selecting all four ancilla qubits on $|0\rangle^{\otimes 4}$ realises the action of $\widetilde{A}$ on the system register.

to an IBM hardware gate set show substantially lower two-qubit gate counts and circuit depth compared with both [25] and the general-purpose methods FABLE [29] and BITBLE [30] at equivalent matrix dimensions. **(iii)** Statevector simulations confirm that the success probability of [2] decays significantly more slowly with N than all compared methods, an advantage that becomes increasingly pronounced as the system size grows. Together, these properties make the construction of [2] the most resource-efficient explicit block encoding currently available for the 2D Dirichlet Laplacian, and hence the natural choice for the VQLS implementation in this work.

2.3 Cost Function Evaluation

Given the scaled linear system $\widetilde{A}\,\mathbf{x} = \widetilde{\mathbf{b}}$, VQLS seeks a quantum state proportional to the solution vector by preparing a parametrized ansatz state

$$|x(\boldsymbol{\theta})\rangle = V(\boldsymbol{\theta})|0\rangle^{\otimes n}, \tag{13}$$

where $V(\boldsymbol{\theta})$ is a parametrized unitary circuit acting on n system qubits and $\boldsymbol{\theta} \in \mathbb{R}^p$ is the vector of variational parameters to be optimised. The goal is to find $\boldsymbol{\theta}^*$ such that $|x(\boldsymbol{\theta}^*)\rangle \propto \widetilde{A}^{-1}|\widetilde{b}\rangle$, where $|\widetilde{b}\rangle = U|0\rangle^{\otimes n}$ is the normalized right-hand side state prepared by a unitary U. To drive the optimization towards this target, one needs a scalar cost function $C_G(\boldsymbol{\theta})$ that is (i) efficiently computable on a quantum circuit, (ii) equal to zero if and only if $\widetilde{A}|x(\boldsymbol{\theta})\rangle \propto |\widetilde{b}\rangle$, and (iii) positive otherwise.

A natural candidate is the normalized global cost function [18]

$$C_G(\boldsymbol{\theta}) = 1 - \frac{\left|\langle\widetilde{b}|\widetilde{A}|x(\boldsymbol{\theta})\rangle\right|^2}{\langle x(\boldsymbol{\theta})|\widetilde{A}^\dagger\widetilde{A}|x(\boldsymbol{\theta})\rangle}, \tag{14}$$

which measures the squared overlap between the normalized vectors $\widetilde{A}|x(\boldsymbol{\theta})\rangle/\|\widetilde{A}|x(\boldsymbol{\theta})\rangle\|$ and $|\widetilde{b}\rangle$. By the Cauchy–Schwarz inequality, $C_G(\boldsymbol{\theta}) \geq 0$ always holds. Moreover, $C_G(\boldsymbol{\theta}) = 0$ if and only if $\widetilde{A}|x(\boldsymbol{\theta})\rangle$ is proportional to $|\widetilde{b}\rangle$, i.e. if and only if $|x(\boldsymbol{\theta})\rangle \propto \widetilde{A}^{-1}|\widetilde{b}\rangle$, which is the desired solution state. Minimising $C_G(\boldsymbol{\theta})$ therefore directly drives the ansatz state towards the solution of the linear system. An equivalent and more convenient expression is obtained by recognising that the denominator is simply $\|A|x(\boldsymbol{\theta})\rangle\|^2$, so that the cost can be written as

$$C_G(\boldsymbol{\theta}) = 1 - \left| \langle \widetilde{b} | \frac{\widetilde{A}|x(\boldsymbol{\theta})\rangle}{\|\widetilde{A}|x(\boldsymbol{\theta})\rangle\|} \right|^2 = 1 - \left| \langle 0^n | \Psi_{\text{post}}(\boldsymbol{\theta})\rangle \right|^2, \tag{15}$$

where the post-selected state $|\Psi_{\text{post}}(\boldsymbol{\theta})\rangle$ is defined in (21) below. The cost is therefore the complement of the fidelity between the normalized state $\widetilde{A}|x(\boldsymbol{\theta})\rangle$ and the target $|\widetilde{b}\rangle$: it equals zero at the solution and increases monotonically as the ansatz deviates from it.

In the block-encoding formulation [1], we assume the $(\alpha, m, 0)$-block-encoding unitary $U_{\widetilde{A}}$ of Section 2.2, with $\alpha = 1$ and $m = 4$ ancilla qubits. The extended Hamiltonian operator is defined as

$$H_G^{(\text{ext})} := U_{\widetilde{A}}^{\dagger} \Big(|0^m\rangle\langle 0^m|_a \otimes \big(I_s - |\widetilde{b}\rangle\langle \widetilde{b}|_s \big) \Big) U_{\widetilde{A}}, \tag{16}$$

where the subscripts a and s refer to the ancilla and system subspaces respectively. For the input state $|X_{\text{ext}}\rangle = |0^m\rangle_a \otimes |x(\boldsymbol{\theta})\rangle_s$, the expectation value of this Hamiltonian is

$$\langle X_{\text{ext}} | H_G^{(\text{ext})} | X_{\text{ext}} \rangle = \frac{1}{\alpha^2} \langle x(\boldsymbol{\theta}) | \widetilde{A}^{\dagger} \big(I - |\widetilde{b}\rangle\langle\widetilde{b}| \big) \widetilde{A} | x(\boldsymbol{\theta})\rangle, \tag{17}$$

which corresponds to the numerator of $C_G(\boldsymbol{\theta})$ up to the factor α^{-2}. Evaluating this expectation value on a quantum circuit proceeds as follows. One prepares

$$|\Psi'(\boldsymbol{\theta})\rangle := U_{\widetilde{A}} \big(|0^m\rangle \otimes |x(\boldsymbol{\theta})\rangle \big), \qquad U|0^n\rangle = |\widetilde{b}\rangle, \tag{18}$$

and post-selects the ancilla register on $|0^m\rangle$, yielding the unnormalized system state

$$\big(\langle 0^m | \otimes U^{\dagger} \big) |\Psi'(\boldsymbol{\theta})\rangle = \frac{1}{\alpha} U^{\dagger} \widetilde{A} |x(\boldsymbol{\theta})\rangle. \tag{19}$$

The probability of obtaining the ancilla outcome $|0^m\rangle$ is

$$p(\boldsymbol{\theta}) = \frac{\langle x(\boldsymbol{\theta}) | \widetilde{A}^{\dagger} \widetilde{A} | x(\boldsymbol{\theta})\rangle}{\alpha^2}, \tag{20}$$

and the normalized post-selected state is

$$|\Psi_{\text{post}}(\boldsymbol{\theta})\rangle = \frac{U^{\dagger} \widetilde{A} |x(\boldsymbol{\theta})\rangle}{\sqrt{\langle x(\boldsymbol{\theta}) | \widetilde{A}^{\dagger} \widetilde{A} | x(\boldsymbol{\theta})\rangle}}. \tag{21}$$

Substituting into (15) and using Bayes' theorem, the fidelity term becomes

$$\left| \langle 0^n | \Psi_{\text{post}}(\boldsymbol{\theta})\rangle \right|^2 = P(\text{sys} = |0^n\rangle \mid \text{anc} = |0^m\rangle) = \frac{P\big(|0\rangle^{\otimes(n+m)}\big)}{P(|0\rangle^{\otimes m})}, \tag{22}$$

so that the global cost function is entirely determined by two measurement probabilities obtainable from a single quantum circuit, and the final expression for the cost is

$$C_G(\boldsymbol{\theta}) = 1 - \frac{P\big(|0\rangle^{\otimes(n+m)}\big)}{P(|0\rangle^{\otimes m})}. \tag{23}$$

A key advantage of the block-encoding formulation is that $C_G(\boldsymbol{\theta})$ is estimated from the measurement statistics of a *single* quantum circuit via (23), with $m = 4$ for the 2D Dirichlet case. Post-selecting the ancilla register on $|0\rangle^{\otimes 4}$ and measuring all qubits yields the numerator and denominator of Eq. (23), from which $C_G(\boldsymbol{\theta})$ is computed classically. This contrasts with the standard VQLS formulation [18], which requires $O(L^2)$ separate Hadamard-test circuits for computing the cost function. The full VQLS circuit implementing this evaluation is shown in Figure 2.

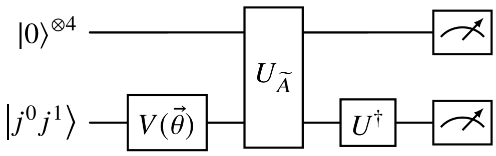

Figure 2: Full VQLS circuit for the 2D Poisson problem via block encoding.

The VQLS algorithm returns the normalized state $|x(\boldsymbol{\theta}^*)\rangle$ such that $\widetilde{A}|x(\boldsymbol{\theta}^*)\rangle \propto |\widetilde{b}\rangle$ or $\widetilde{A}|x(\boldsymbol{\theta}^*)\rangle = \|\mathbf{x}\||\widetilde{b}\rangle$, which encodes the direction of the solution but not its magnitude. To recover the exact classical solution vector $\mathbf{x}$ for (3), the unknown norm $\|\mathbf{x}\|$ must be determined. Using $A|x(\boldsymbol{\theta}^*)\rangle = \frac{\|\mathbf{x}\|}{\|\mathbf{b}\|}\mathbf{b}$, one can show that

$$\|\mathbf{x}\| = \sqrt{\langle x(\boldsymbol{\theta}^*)|A^\dagger A|x(\boldsymbol{\theta}^*)\rangle}, \tag{24}$$

Eq. (24) is estimated classically from the optimised ansatz state and the known matrix A. The classical solution vector is then reconstructed as $\mathbf{x} = \frac{\|\mathbf{b}\|}{\|\mathbf{x}\|} \cdot |x(\boldsymbol{\theta}^*)\rangle$, from which the interior solution field is assembled and compared against the SOR reference in Section 3. Note that this post-processing step is performed entirely classically and introduces no additional quantum circuit evaluations.

For the variational circuit $V(\boldsymbol{\theta})$, we utilize the hardware-efficient ansatz (HEA) [31] comprising alternating single-qubit R_y rotation gates and two-qubit controlled-Z (CZ) entangling gates. Each layer consists of $R_y(\theta_i)$ rotations applied to all n system qubits, followed by CZ gates on alternating pairs of neighbouring qubits, followed by a second round of R_y rotations, as shown in the Figure 3. For an ansatz with l layers, the total number of variational parameters is $p = 2l(n-2) + n$, which scales linearly with the number of system qubits.

In all experiments reported in Section 3, we use $l = 2$, keeping the circuit depth minimal while retaining sufficient expressivity for the problem sizes considered. The variational parameters $\boldsymbol{\theta}$ are initialized by drawing each component independently and uniformly from $[0, 2\pi)$ at the start of each optimization run. A run is considered converged when the best-so-far cost satisfies $C_G(\boldsymbol{\theta}) \le 10^{-6}$; runs that do not reach this threshold within the allocated evaluation budget are considered partially converged, and the best cost achieved is reported.

2.4 The Barren-Plateau Phenomenon

A fundamental challenge shared by all variational quantum algorithms, including VQLS, is the barren-plateau phenomenon. As the number of qubits n or the circuit depth increases, the gradient of the cost function with respect to the variational parameters becomes exponentially suppressed, making optimization increasingly difficult. Formally, a cost function $C_G(\boldsymbol{\theta})$ is said to exhibit a *probabilistic barren plateau* [22] if, when $\boldsymbol{\theta}$ is sampled from some distribution,

$$\mathrm{Var}_{\boldsymbol{\theta}}[C_G(\boldsymbol{\theta})] \quad \text{or} \quad \mathrm{Var}_{\boldsymbol{\theta}}[\partial_\mu C_G(\boldsymbol{\theta})] \in O(b^{-n}), \tag{25}$$

for some $b > 1$ and some parameter $\theta_\mu \in \boldsymbol{\theta}$. This implies that as n increases, the probability of sampling a parameter vector at which the gradient is non-negligible becomes exponentially small, a behavior attributed to the emergence of unitary 2-design characteristics in sufficiently deep or expressive parametric circuits [22, 32].

It is important to note, however, that barren plateaus describe an *average-case* property of the cost landscape rather than a uniform absence of gradients everywhere. Even in high-dimensional parameter spaces dominated by flat regions, narrow "gorges" [23], localized directions of non-negligible gradient magnitude can persist around the global minimum, as illustrated in Figure 4. The practical consequence

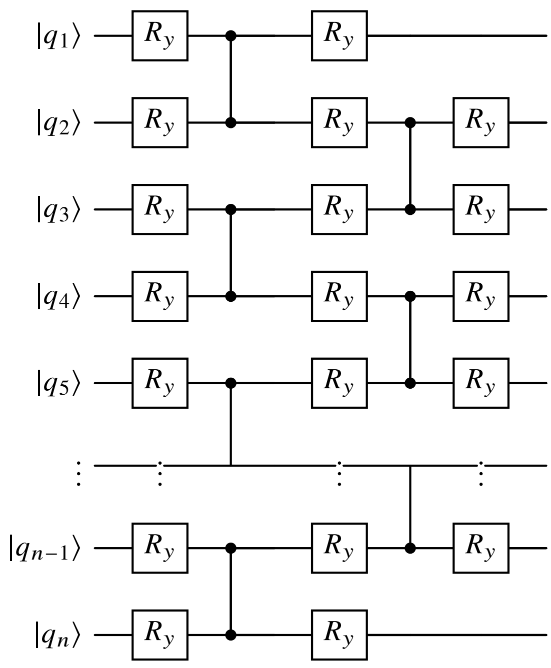

Figure 3: Hardware-Efficient ansatz with one layer for $V(\vec{\theta})$.

is that optimization success depends not only on the expected gradient magnitude but critically on the optimizer's ability to locate and exploit these rare descent pathways. Gradient-based and deterministic gradient-free methods that rely on local differential information tend to stagnate once gradients become exponentially suppressed. Population-based optimizers, by contrast, maintain global exploration of the parameter space and can occasionally discover gorge-like descent directions even when point-wise gradient estimates are negligibly small. The comparative performance of COBYLA and CMA-ES in this regard is studied empirically in Section 3.3.

3. Results

This section presents the numerical results of the proposed VQLS implementation across two benchmark problems. Section 3.1 defines the two benchmark problems and their discretization parameters. Section 3.2 assesses solution accuracy through field comparisons and a systematic error study across system sizes. Section 3.3 analyzes the effect of optimizer choice on cost-function convergence and barren-plateau onset. Section 3.4 reports transpiled circuit resource estimates and compares them against the prior implementation of [1].

3.1 Benchmark Problems

Both benchmark problems are instances of the general 2D Poisson framework introduced in Section 2.1, differing only in the choice of source term $f(x, y)$ and the imposed boundary conditions.

3.1.1 Sinusoidal source (Poisson benchmark)

We consider (2) with the prescribed smooth source term

$$f(x, y) = -\sin(\pi x)\sin(\pi y), \tag{26}$$

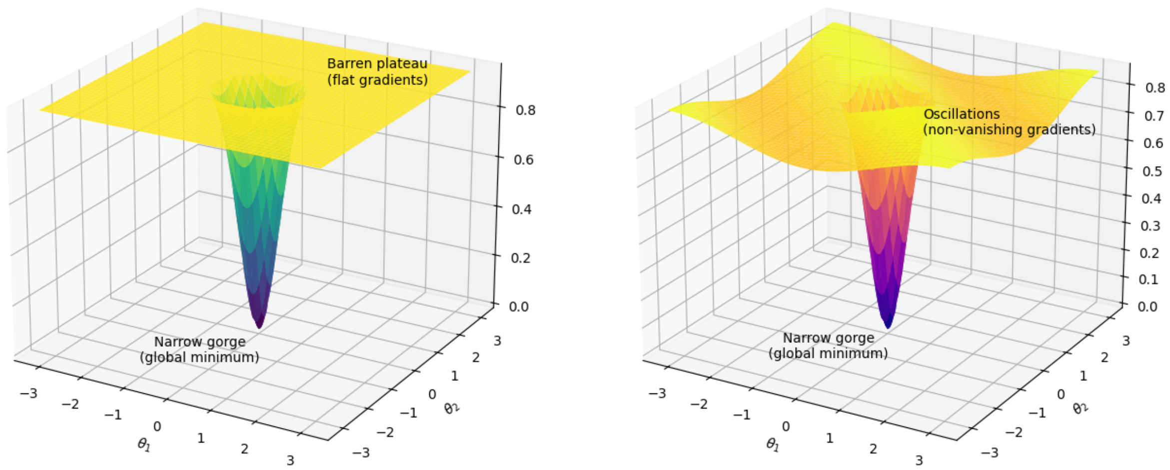

(a) (b)

Figure 4: Illustrative theoretical cost landscapes exhibiting narrow gorges in variational optimization. Figure (a) shows a global cost function $C_{\text{barren}} = 1 - \left[\cos^2(\theta_1/2)\cos^2(\theta_2/2)\right]^k$, which features a narrow global minimum surrounded by a barren plateau. Figure (b) shows a modified cost $0.9C_{\text{barren}} - 0.1\sin^2\big((\theta_1+\theta_2)/2\big)$, which retains the narrow gorge while introducing oscillatory structure that prevents gradient collapse. For parametric quantum circuits it is know that the landscape shown in Figure (b) cannot arise [23].

and homogeneous Dirichlet boundary conditions $u|_{\partial\Omega} = 0$ on all four walls. This is a standard PDE benchmark: the forcing is smooth, free of singularities, and admits a known analytical solution $u(x, y) = \sin(\pi x)\sin(\pi y)/(2\pi^2)$, providing a rigorous accuracy check.

3.1.2 Gaussian source (heat conduction benchmark)

We consider the steady-state heat conduction equation $-k\nabla^2 T = Q(x, y)$ on $[0, 1]^2$, which is a direct instance of (2) with $u \equiv T$ and source term

$$f(x, y) = Q(x, y) = Q_0 \exp\left(-\frac{(x - x_0)^2 + (y - y_0)^2}{2\sigma^2}\right), \tag{27}$$

where Q_0 is the peak heat generation rate, (x_0, y_0) is the source centre, and σ is the source width. The boundary conditions are Dirichlet on all four walls: bottom wall at T_{hot}, top wall at T_{cold}, and side walls at $T_{\text{sides}} = (T_{\text{hot}} + T_{\text{cold}})/2$, as shown in the illustrative Fig 5. This configuration models a thin conducting plate with a localized heat-generating component, representative of an electronic cooling scenario, held between a hot base and a cold lid, with the side walls maintained at the mean temperature. The Gaussian source provides a physically non-trivial, spatially localized forcing that exercises the solver's ability to resolve sharp interior gradients. The parameters used in this work are $T_{\text{hot}} = 1$, $T_{\text{cold}} = 0$, $Q_0 = 5$, $(x_0, y_0) = (0.5, 0.3)$, and $\sigma = 0.1$, placing the source in the lower half of the domain near the hot wall.

Both problems are discretized following the procedure of Section 2.1. The system matrix A is identical for both benchmarks at the same grid resolution; only the right-hand side vector $\mathbf{b}$ differs. For the sinusoidal source benchmark, detailed solution comparisons are presented for two grid configurations that together probe both rectangular and square operator geometries:

- **Rectangular grid:** $N_x = 16$, $N_y = 8$, yielding $N = 128$ interior unknowns and a 128×128 system

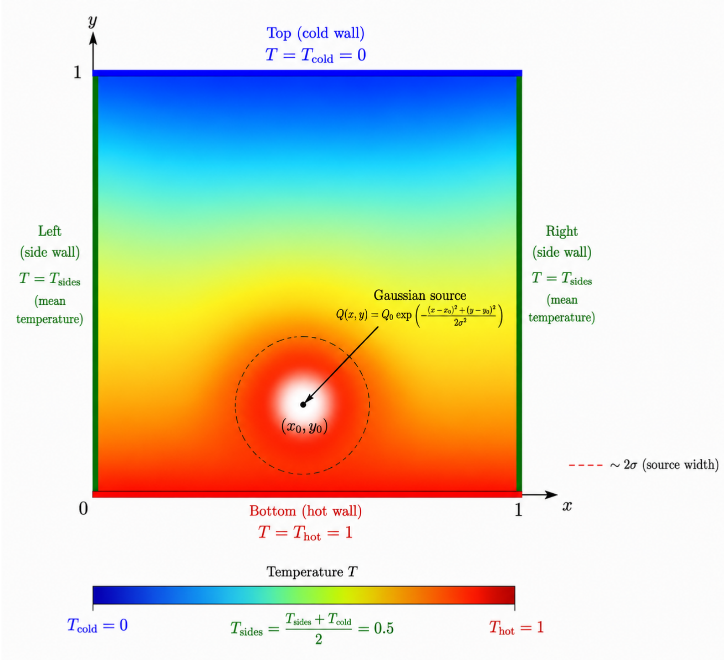

Figure 5: Schematic of the steady-state heat conduction benchmark problem. The interior color map is an illustrative representation.

matrix ($n = 7$ system qubits, 11 total qubits).

- **Square grid:** $N_x = N_y = 16$, yielding $N = 256$ interior unknowns and a 256×256 system matrix ($n = 8$ system qubits, 12 total qubits).

For the heat conduction benchmark, results are presented for the square grid ($N_x = N_y = 16$, $N = 256$) only, as this configuration is the more demanding of the two and most clearly illustrates the performance of the solver in the presence of a spatially localized source. A systematic study of accuracy across system sizes from $N = 4$ to $N = 256$ is reported in Table 1 for both benchmarks. All simulations are performed on the Qiskit Aer statevector simulator using the CMA-ES optimizer. Classical reference solutions are obtained using the SOR iterative solver.

3.2 Solution Accuracy

Solution accuracy is assessed by the maximum absolute pointwise error $\varepsilon_\infty = \max_i |u_i^{\text{VQLS}} - u_i^{\text{SOR}}|$ over all interior nodes, computed after sign alignment and unit-norm normalization of both solution vectors.

3.2.1 Sinusoidal Source Benchmark

Figures 6–7 show the classical and VQLS solution fields together with the absolute error fields for the sinusoidal source benchmark.

For the rectangular grid (Figure 6) the two solution fields are visually indistinguishable: the VQLS solver correctly reproduces the smooth bell-shaped pressure distribution that is characteristic of the sinusoidal source, peaking at the domain centre and decaying to zero at all four walls in accordance with the homogeneous Dirichlet conditions. The absolute error field confirms this visual agreement quantitatively: the maximum pointwise error is $\varepsilon_\infty \approx 10^{-5}$, with residuals arranged in a horizontal banding pattern whose peaks align with the nodal structure of the sinusoidal source along the y-direction. This near-machine-precision accuracy at $N = 128$ ($n = 7$ qubits) demonstrates that the CMA-ES optimizer has converged to the correct solution to high fidelity at this scale.

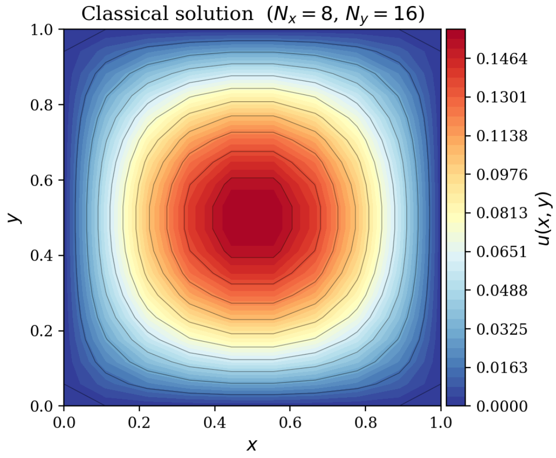

(a) Classical (SOR) solution.

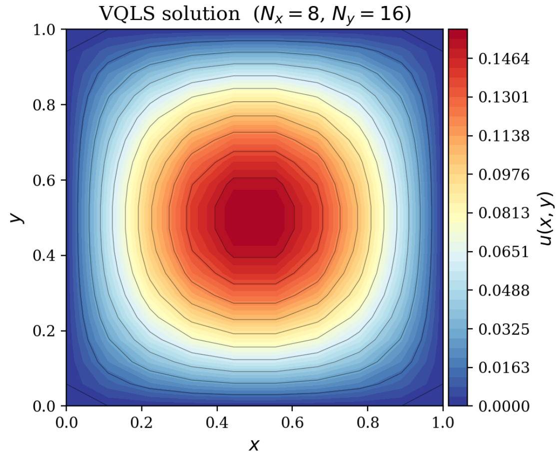

(b) Quantum (VQLS) solution.

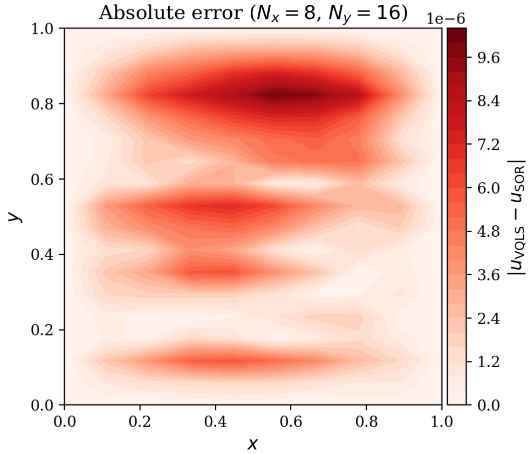

(c) Absolute error field $|u^{\mathrm{VQLS}} - u^{\mathrm{SOR}}|$.

Figure 6: Sinusoidal source benchmark, rectangular grid ($N_x = 8$, $N_y = 16$, $N = 128$, $n = 7$ system qubits, 11 total qubits), CMA-ES optimizer with 25,000 function evaluations. The VQLS solution is visually indistinguishable from the classical reference; the maximum pointwise error $\varepsilon_\infty \approx 10^{-5}$.

For the square grid (Figure 7), the VQLS solution again reproduces the qualitative features of the classical field, though at this larger problem size ($N = 256$, $n = 8$ qubits) the maximum absolute error increases to $\varepsilon_\infty \approx 0.032$. The result is consistent with the known barren-plateau behavior, where the cost landscape becomes increasingly flat and harder to traverse, and motivates the use of more advanced optimizers for larger system sizes.

3.2.2 Heat Conduction Benchmark

Figure 8 shows the classical and VQLS temperature fields and the absolute error field for the heat conduction benchmark with Gaussian source, for the square grid ($N_x = N_y = 16$, $N = 256$, $n = 8$ system qubits). The VQLS solution correctly reproduces the dominant large-scale thermal stratification: the hot bottom wall ($T_{\mathrm{hot}} = 1$) and cold top wall ($T_{\mathrm{cold}} = 0$) are both faithfully represented, and the overall direction of the vertical temperature gradient is preserved. The isotherm contours in the VQLS field are however more horizontally uniform than in the classical solution, lacking the lateral curvature near the side walls and the subtle local temperature elevation induced by the Gaussian source at $(x_0, y_0) = (0.5, 0.3)$. The absolute error field confirms this: the maximum pointwise error reaches $\varepsilon_\infty \approx 0.144$, concentrated

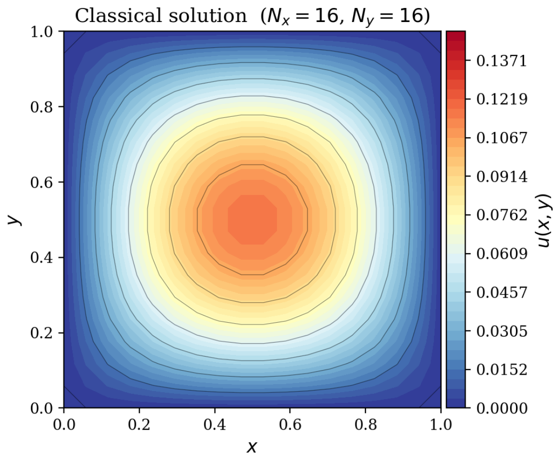

(a) Classical (SOR) solution.

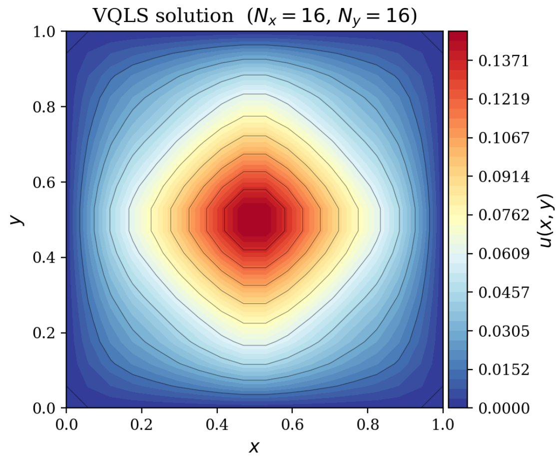

(b) Quantum (VQLS) solution.

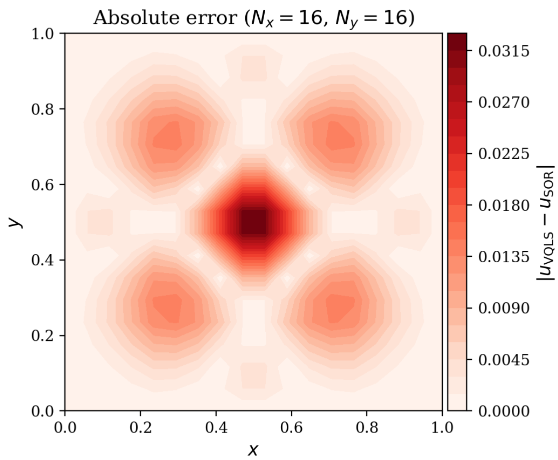

(c) Absolute error field $|u^{\mathrm{VQLS}} - u^{\mathrm{SOR}}|$.

Figure 7: Sinusoidal source benchmark, square grid ($N_x = N_y = 16$, $N = 256$, $n = 8$ system qubits, 12 total qubits), CMA-ES optimizer with 25,000 function evaluations. The large-scale solution structure is fairly reproduced; the maximum pointwise error is $\varepsilon_\infty \approx 0.032$, concentrated in a five-lobe pattern consistent with partial optimizer convergence at the onset of the barren-plateau regime.

in a broad band across the upper half of the domain ($y \gtrsim 0.5$) and in a lower arc near the bottom wall, with minimum error in the mid-domain strip around $y \approx 0.4$. This behavior is consistent with partial convergence at $n = 8$ qubits, where the onset of the barren-plateau phenomenon makes full optimization increasingly difficult within the given evaluation budget, and is qualitatively similar to what is observed for the sinusoidal source benchmark at the same system size. A systematic summary of solution accuracy across all tested system sizes for both benchmarks is provided in Table 1.

3.3 Effect of Optimizer Choice on Convergence

Optimizer choice has a material effect on the convergence behavior of VQLS as the system size grows. To isolate this effect, the comparison is carried out on the sinusoidal source benchmark; qualitatively identical behavior is observed for the heat conduction benchmark. Both optimizers are run with a maximum budget of upto $\sim$ 25,000 iterations.

Figure 9a shows the cost function evolution for $N = 64$ ($N_x = N_y = 8$, $n = 6$ qubits). At this system size, both optimizers converge to near-zero cost. COBYLA reaches this regime more rapidly, with its best-so-far

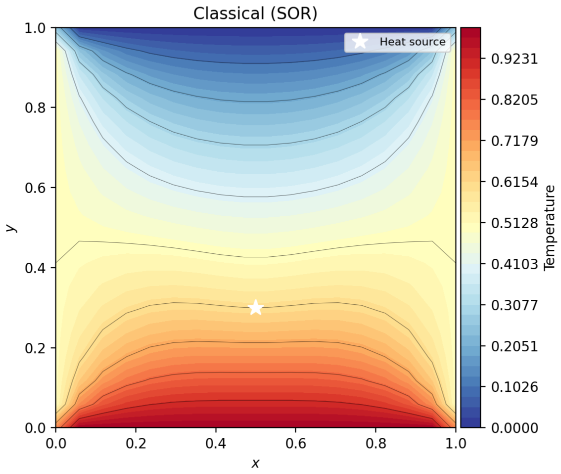

(a) Classical (SOR) solution.

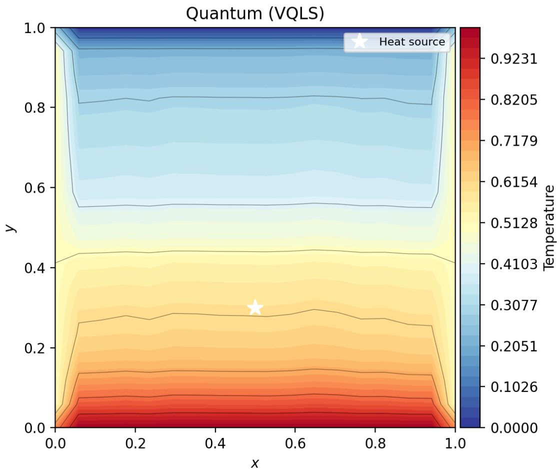

(b) Quantum (VQLS) solution.

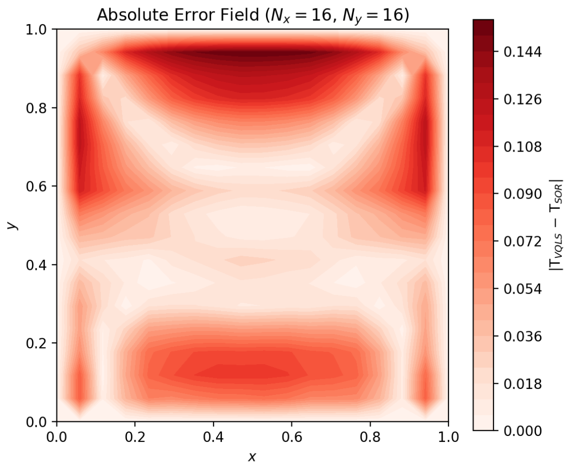

(c) Absolute error field $|T^{\mathrm{VQLS}} - T^{\mathrm{SOR}}|$.

Figure 8: Heat conduction benchmark with Gaussian source, square grid ($N_x = N_y = 16$, $N = 256$, $n = 8$ system qubits, 12 total qubits), CMA-ES optimizer with 25,000 function evaluations. The white star marks the heat source at $(x_0, y_0) = (0.5, 0.3)$. The large-scale thermal stratification is correctly reproduced; the maximum pointwise error $\varepsilon_\infty \approx 0.144$ is concentrated in the upper domain and near the bottom wall.

envelope descending to near-machine-precision values by approximately 5,000 function evaluations. CMA-ES achieves a comparable final cost but requires $\sim 9{,}000$ evaluations to do so. This is consistent with the known behavior of CMA-ES on low-dimensional problems: the population-based search incurs an overhead per generation that is unnecessary when the landscape is not yet dominated by barren-plateau effects.

The picture changes substantially at $N = 256$ ($N_x = N_y = 16$, $n = 8$ qubits), shown in Figure 9b. COBYLA's best-so-far envelope descends from $C_G \approx 1$ to $C_G \approx 0.90$ over its 25,000-evaluation budget and then stalls completely, making no further progress. This is a clear manifestation of the barren-plateau phenomenon: the cost landscape at $n = 8$ qubits is sufficiently flat that COBYLA's local linear model cannot identify productive descent directions, and the optimizer becomes trapped. CMA-ES, which begins its effective descent later (around evaluation 10,000 when its population has sufficiently explored the landscape), continues to make monotonic progress throughout its budget, ultimately reaching $C_G \approx 0.59$. While this is still far from convergence, the contrast with COBYLA is significant: CMA-ES reduces the cost by approximately 35% relative to its starting value in the region where COBYLA makes zero progress.

Table 1: VQLS solution accuracy across system sizes for both benchmark problems. $\varepsilon_\infty = \max_i |u_i^{\text{VQLS}} - u_i^{\text{SOR}}|$ is the maximum absolute pointwise error over interior nodes, computed after sign alignment; the solution norm is recovered via Eq. (24) . n denotes system qubits; total circuit width is $n + 4$.The non-monotone behaviour in the heat conduction column at intermediate system sizes is currently under investigation;

$N_x \times N_y$	N	n	ε_∞	
			Sine	Heat
2×2	4	2	7×10^{-6}	1.62×10^{-5}
2×4	8	3	2.8×10^{-5}	3.1×10^{-1}
4×4	16	4	1.08×10^{-5}	4.05×10^{-5}
4×8	32	5	1.12×10^{-5}	4.2×10^{-1}
8×8	64	6	3.6×10^{-6}	1.9×10^{-1}
8×16	128	7	9.6×10^{-6}	4.8×10^{-1}
16×16	256	8	3.15×10^{-2}	1.44×10^{-1}

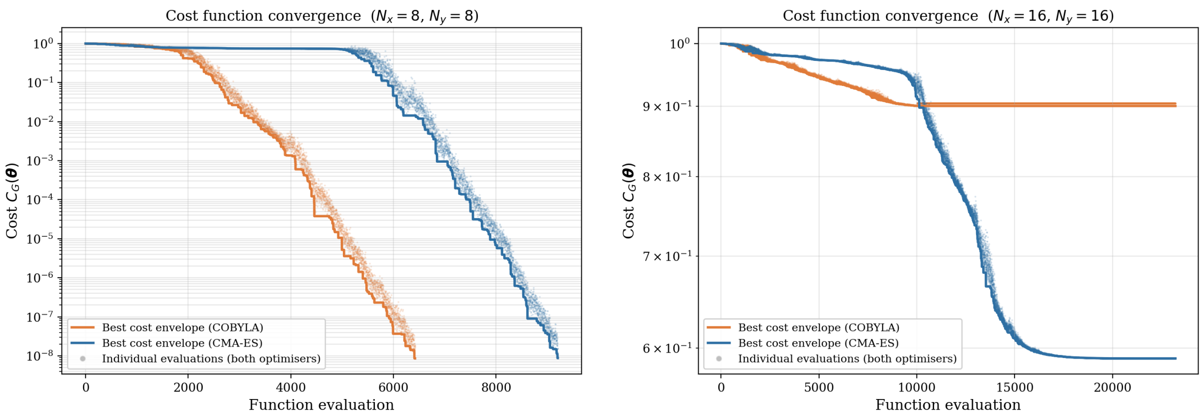

(a) $N_x = N_y = 8$, $N = 64$, $n = 6$ qubits.

(b) $N_x = N_y = 16$, $N = 256$, $n = 8$ qubits.

Figure 9: Cost function convergence for COBYLA and CMA-ES at two representative system sizes on the sinusoidal source benchmark. The solid lines show the best-so-far envelope; the faint dots show individual function evaluations.

Figure 10 summarizes the best cost achieved by each optimizer across all tested system sizes. Both optimizers achieve $C_G \approx 0$ for $N \leq 64$, confirming that the barren-plateau effect is negligible at these scales within the given evaluation budget. The divergence begins at $N = 128$: COBYLA's best cost rises to $C_G \approx 0.36$ while CMA-ES still achieves $C_G \approx 0$, indicating that CMA-ES delays the effective onset of barren-plateau stagnation by at least one system-size step relative to COBYLA. At $N = 256$, both optimizers fail to reach zero cost, but CMA-ES achieves a substantially lower best cost ($C_G \approx 0.59$ vs $C_G \approx 0.90$). These results demonstrate that while CMA-ES does not eliminate the barren-plateau problem, it mitigates its onset and continues to make progress in regimes where COBYLA stalls entirely, making it the more appropriate choice for larger CFD-relevant system sizes.

3.4 Circuit Resource Estimates

Table 2 reports the circuit resources for the VQLS implementation using the 2D Dirichlet Laplacian block encoding of [2], transpiled to the {CX, RZ, SX, X} basis set. The circuit width is $n + 4$ qubits in all cases, one fewer ancilla than the pentadiagonal oracle implementation of [1], which required $n + 5$ qubits.

Figures 11b and 11c compare the two-qubit gate count and circuit depth of the present implementation

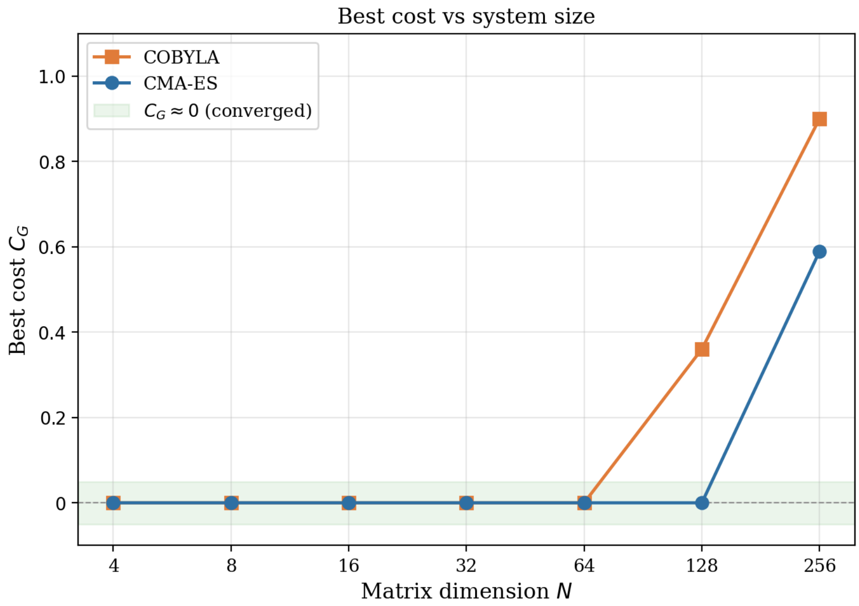

Figure 10: Best cost C_G achieved within the evaluation budget as a function of matrix dimension N, for COBYLA and CMA-ES on the sinusoidal source benchmark. A run is considered converged when the best-so-far cost satisfies $C_G(\boldsymbol{\theta}) \leq 10^{-6}$.

Table 2: Circuit resources for the 2D Dirichlet Laplacian VQLS implementation of this work, using the block encoding of [2]. Width = $n + 4$ (system qubits + 4 ancillae). Gate counts and depth are after transpilation to the {CX, RZ, SX, X} basis.

N	n	Width	Depth	1-qubit gates	2-qubit gates
4	2	6	339	292	146
8	3	7	493	456	207
16	4	8	812	910	334
32	5	9	2358	2650	910
64	6	10	7227	8565	2540
128	7	11	25909	31315	8495
256	8	12	98347	120127	31304

against the pentadiagonal oracle used in our prior work [1], across all tested system sizes. The improvement is consistent and substantial: at $N = 64$ ($n = 6$ qubits), the VQLS circuit of the present work requires 3.2x fewer two-qubit gates and produces a 2.5x shallower circuit; at $N = 256$ ($n = 8$ qubits), the reductions are 1.6x and 1.4x respectively. These gains arise from the more efficient exploitation of the Kronecker-sum separability of the 2D Laplacian in the construction of [2], which reduces both the ancilla overhead and the depth of the boundary-handling sub-circuits compared to the general pentadiagonal construction of [25].

Beyond gate counts, the new block encoding also achieves a higher post-selection success probability than [25], as demonstrated numerically in [2]. For the same input state, the probability of the ancilla register returning $|0\rangle^{\otimes m}$ decays significantly more slowly with N under the new construction. This directly translates to a practical improvement in VQLS performance: a higher success probability means more useful measurement outcomes per circuit execution, reducing the number of shots required to estimate $C_G(\boldsymbol{\theta})$ to a given statistical precision. We attribute the improved solution accuracy observed in this work relative to [1] (near-exact recovery at $N = 128$ compared to $N = 32$ in the prior work) in part to this higher success probability enabling more reliable cost-function estimation at larger system sizes.

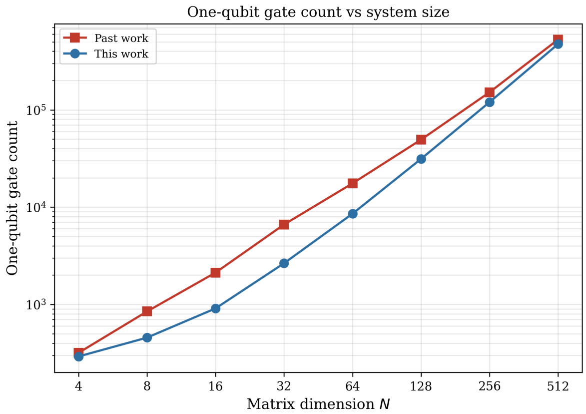

(a) One-qubit gate count.

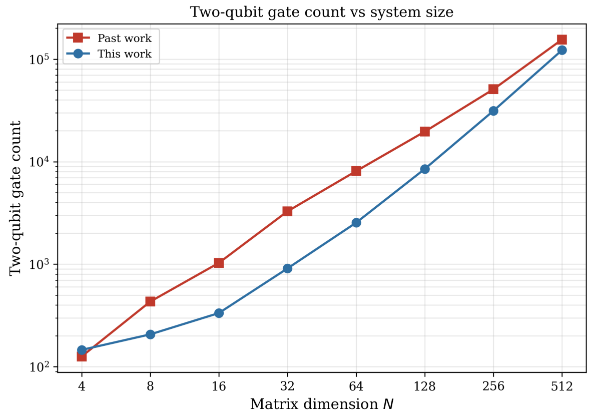

(b) Two-qubit gate count.

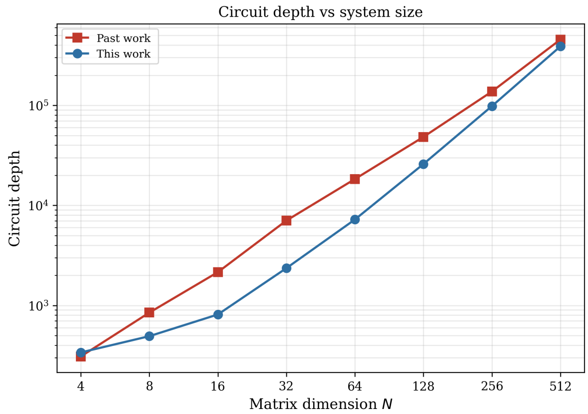

(c) Circuit depth.

Figure 11: VQLS circuit resource comparison between the block encoding of [2] used in this work (blue circles, $n+4$ qubits) and the pentadiagonal oracle of [25] used in our prior work [1] (red squares, $n+5$ qubits), across all tested system sizes. All counts are after transpilation to the {CX, RZ, SX, X} basis gate set.

4. Conclusion and Future Work

This work has demonstrated that two targeted changes to the VQLS pipeline of [1]— replacing the block-encoding circuit and replacing the optimizer, produce measurable, quantified improvements in solution accuracy and convergence behavior for CFD-relevant Poisson systems.

The results across both benchmarks tell a consistent story. For small systems ($N \leq 64$), the solver is well-behaved regardless of optimizer choice: both COBYLA and CMA-ES converge to near-zero cost, with COBYLA doing so more efficiently within the evaluation budget. The block-encoding improvement is the dominant factor at these scales, with the 3.2× reduction in two-qubit gate count and the higher post-selection success probability of [2] enabling reliable cost-function estimation that was not achievable with the pentadiagonal oracle of [1] at comparable sizes. The sinusoidal source benchmark confirms this directly: near-exact recovery ($\varepsilon_\infty \sim 10^{-5}$) is achieved at $N = 128$.

Beyond $N = 64$, the barren-plateau onset becomes the limiting factor, and here the choice of optimizer matters. COBYLA stalls at $N = 128$ and makes no progress at $N = 256$, while CMA-ES continues to reduce the cost at both sizes. At $N = 256$, the VQLS solution for both benchmarks correctly reproduces the large-scale structure of the classical reference — the bell-shaped pressure distribution for the sinusoidal source and the vertical thermal stratification for the heat conduction problem — but the residuals are $O(10^{-2})$ and $O(10^{-1})$ respectively rather than $O(10^{-5})$, reflecting the incomplete optimization rather than a failure of the block-encoding circuit. The resource comparison with [1] reinforces the VQLS circuit-level gains quantitatively. Across all tested sizes, the VQLS circuit used in this work requires fewer two-qubit gates and produces shallower circuits.

Several directions present themselves for future work. On the algorithmic side, the most pressing open question is the development of optimizers that are explicitly aware of the variational landscape geometry during training. The quantum natural gradient (QNG) framework offers a principled starting point, and adapting it to the VQLS setting in a computationally tractable way remains an active direction. Equally important is ansatz design: hardware-efficient layered ansätze are known to be susceptible to barren plateaus at depth, and physically informed or problem-structure-aware ansätze that incorporate the sparsity and Kronecker-sum structure of the Laplacian may improve trainability and reduce the parameter count required for accurate solutions.

On the problem side, the block-encoding framework of [2] natively supports mixed boundary conditions and extensions to three spatial dimensions, both directly relevant to realistic CFD geometries. A further natural extension is the integration of the VQLS pressure solver into a time-accurate incompressible Navier–Stokes solver via a projection-method loop, providing the first demonstration of quantum-assisted time-stepping for CFD. Finally, near-term demonstrations on IBM and IonQ quantum processors with appropriate noise mitigation strategies are essential for establishing the practical fidelity of the approach beyond statevector simulation, and the circuit resource estimates presented here confirm that modest system sizes are within reach of current hardware.

Looking further ahead, the block-encoding circuits used in this work are directly reusable as subroutines in Quantum Singular Value Transformation (QSVT)-based fault-tolerant quantum linear solvers [**?**], which provide provable polynomial speedups over classical methods under well-defined conditions on matrix sparsity and condition number. In the nearer term, the most promising path to practical quantum advantage is the restructuring of CFD workflows to query low-dimensional functionals of the solution rather than the full solution field, thereby circumventing the classical readout bottleneck. Identifying which quantities of engineering interest in CFD admit efficient quantum observable estimation is an important open problem that we intend to pursue in future work.

Acknowledgements

The authors thank Alexandre Boutot (University of Waterloo) for his contributions to the block-encoding framework [2]. The authors also thank Rajas K. Dalvi for his insights on the role of optimizers in delaying barren plateaus, and Kunal Garg for reviewing the manuscript.

References

[1] Viraj D'Souza, Dhamotharan S, Tanishque Kumar, Ramesh Kolluru, Ferdin Don Bosco, Abhishek Chopra, Rut Lineswala, and Onkar Sahni. *Towards Hybrid Quantum-Classical Computing for Large-Scale CFD: Tackling Scalability Challenges of the Variational Quantum Linear Solver*.

[2] Alexandre Boutot and Viraj Dsouza. Explicit block encodings of discrete laplacians with mixed boundary conditions, 2026.

[3] M. Mani and A. J. Dorgan. A perspective on the state of aerospace computational fluid dynamics technology. *Annual Review of Fluid Mechanics*, 55(1):431–457, 2023.

[4] Timothy Mauery, Juan Alonso, Andrew Cary, Vincent Lee, Robert Malecki, Dimitri Mavriplis, Gorazd Medic, John Schaefer, and Jeffrey Slotnick. A Guide for Aircraft Certification by Analysis. NASA Contractor Report, NNL16AA04B/80LARC19F0018, May 2021.

[5] Jeffrey P. Slotnick, Abdollah Khodadoust, Juan Alonso, David Darmofal, William Gropp, Elizabeth Lurie, and Dimitri J. Mavriplis. CFD Vision 2030 Study: A Path to Revolutionary Computational Aerosciences. NASA Contractor Report, NNL08AA16B, NNL12AD05T, WBS 794072.02.07.02.03, March 2014.

[6] Andrew W. Cary, John Chawner, Earl P. Duque, William Gropp, William L. Kleb, Raymond M. Kolonay, Eric Nielsen, and Brian Smith. *CFD Vision 2030 Road Map: Progress and Perspectives*.

[7] Zhaoyuan Meng and Yue Yang. Quantum computing of fluid dynamics using the hydrodynamic schrödinger equation. *Physical Review Research*, 5(3), 2023.

[8] Blaga N. Todorova and René Steijl. Quantum algorithm for the collisionless boltzmann equation. *Journal of Computational Physics*, 409:109347, 2020.

[9] Wael Itani and Sauro Succi. Analysis of carlemann linearization of lattice boltzmann, 2021.

[10] Aram W. Harrow, Avinatan Hassidim, and Seth Lloyd. Quantum algorithm for linear systems of equations. *Phys. Rev. Lett.*, 103:150502, Oct 2009.

[11] Yarui Zheng, Chao Song, Ming-Cheng Chen, Benxiang Xia, Wuxin Liu, Qiujiang Guo, Libo Zhang, Da Xu, Hui Deng, Keqiang Huang, Yulin Wu, Zhiguang Yan, Dongning Zheng, Li Lu, Jian-Wei Pan, H. Wang, Chao-Yang Lu, and Xiaobo Zhu. Solving systems of linear equations with a superconducting quantum processor. *Phys. Rev. Lett.*, 118:210504, May 2017.

[12] IBM Quantum. Ibm quantum technology roadmap, 2024.

[13] Quantinuum. Quantinuum accelerated roadmap to universal fault-tolerant quantum computing by 2030, 2024.

[14] IonQ. Ionq accelerated roadmap: Turning quantum ambition into reality, 2024.

[15] András Gilyén, Yuan Su, Guang Hao Low, and Nathan Wiebe. Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics. In *Proceedings of the 51st Annual ACM Symposium on Theory of Computing (STOC)*, pages 193–204. ACM, 2019.

[16] John M. Martyn, Zane M. Rossi, Andrew K. Tan, and Isaac L. Chuang. Grand unification of quantum algorithms. *PRX Quantum*, 2(4):040203, 2021.

[17] John Preskill. Quantum computing in the nisq era and beyond. *Quantum*, 2:79, August 2018.

[18] Carlos Bravo-Prieto, Ryan LaRose, M. Cerezo, Yigit Subasi, Lukasz Cincio, and Patrick J. Coles. Variational quantum linear solver. *Quantum*, 7:1188, November 2023.

[19] Y. Y. Liu, Z. Chen, C. Shu, S. C. Chew, B. C. Khoo, X. Zhao, and Y. D. Cui. Application of a variational hybrid quantum-classical algorithm to heat conduction equation and analysis of time complexity. *Physics of Fluids*, 34(11), November 2022.

[20] Yudong Cao, Anargyros Papageorgiou, Iasonas Petras, Joseph Traub, and Sabre Kais. Quantum algorithm and circuit design solving the poisson equation. *New Journal of Physics*, 15(1):013021, January 2013.

[21] Scott Aaronson. Read the fine print. *Nature Physics*, 11:291–293, April 2015.

[22] Jarrod R. McClean, Sergio Boixo, Vadim N. Smelyanskiy, Ryan Babbush, and Hartmut Neven. Barren plateaus in quantum neural network training landscapes. *Nature Communications*, 9(1), November 2018.

[23] Andrew Arrasmith, Zoe Holmes, M. Cerezo, and Patrick J. Coles. Equivalence of quantum barren plateaus to cost concentration and narrow gorges. *Quantum Science and Technology*, 7(4):045015, 2022.

[24] Daan Camps, Lin Lin, Roel Van Beeumen, and Chao Yang. Explicit quantum circuits for block encodings of certain sparse matrices. *SIAM Journal on Matrix Analysis and Applications*, 45(1):801–827, 2024.

[25] Christoph Sünderhauf, Earl Campbell, and Joan Camps. Block-encoding structured matrices for data input in quantum computing. *Quantum*, 8:1226, January 2024.

[26] Kandula Eswara Sai Kumar, Supreeth B S, Rajas Dalvi, Aman Mittal, Aakif Akhtar, Ferdin Don Bosco, Rut Lineswala, and Abhishek Chopra. Benchmarking of gpu-optimized quantum-inspired evolutionary optimization algorithm using functional analysis, 2024.

[27] Nikolaus Hansen and Andreas Ostermeier. Completely derandomized self-adaptation in evolution strategies. *Evolutionary Computation*, 9(2):159–195, 2001.

[28] Nikolaus Hansen. *The CMA evolution strategy: A comparing review*, pages 75–102. Studies in Fuzziness and Soft Computing. December 2006.

[29] Daan Camps and Roel Van Beeumen. Fable: Fast approximate quantum circuits for block-encodings.

In *2022 IEEE International Conference on Quantum Computing and Engineering (QCE)*, page 104–113. IEEE, 2022.

[30] Zexian Li, Xiao-Ming Zhang, Chunlin Yang, and Guofeng Zhang. Binary tree block encoding of classical matrix, 2025.

[31] Abhinav Kandala, Antonio Mezzacapo, Kristan Temme, Maika Takita, Markus Brink, Jerry M. Chow, and Jay M. Gambetta. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature*, 549(7671):242–246, 2017.

[32] M. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles. Cost function dependent barren plateaus in shallow parametrized quantum circuits. *Nature Communications*, 12(1):1791, 2021.