

INFORMS 2026 ANNUAL MEETING

HPDECO

Quantum-Accelerated Adjoint Method for PDE-Constrained Optimization

Anish Ghosh

JOINT WORK WITH

Dr. Mohammadhossein Mohammadisaihroudi

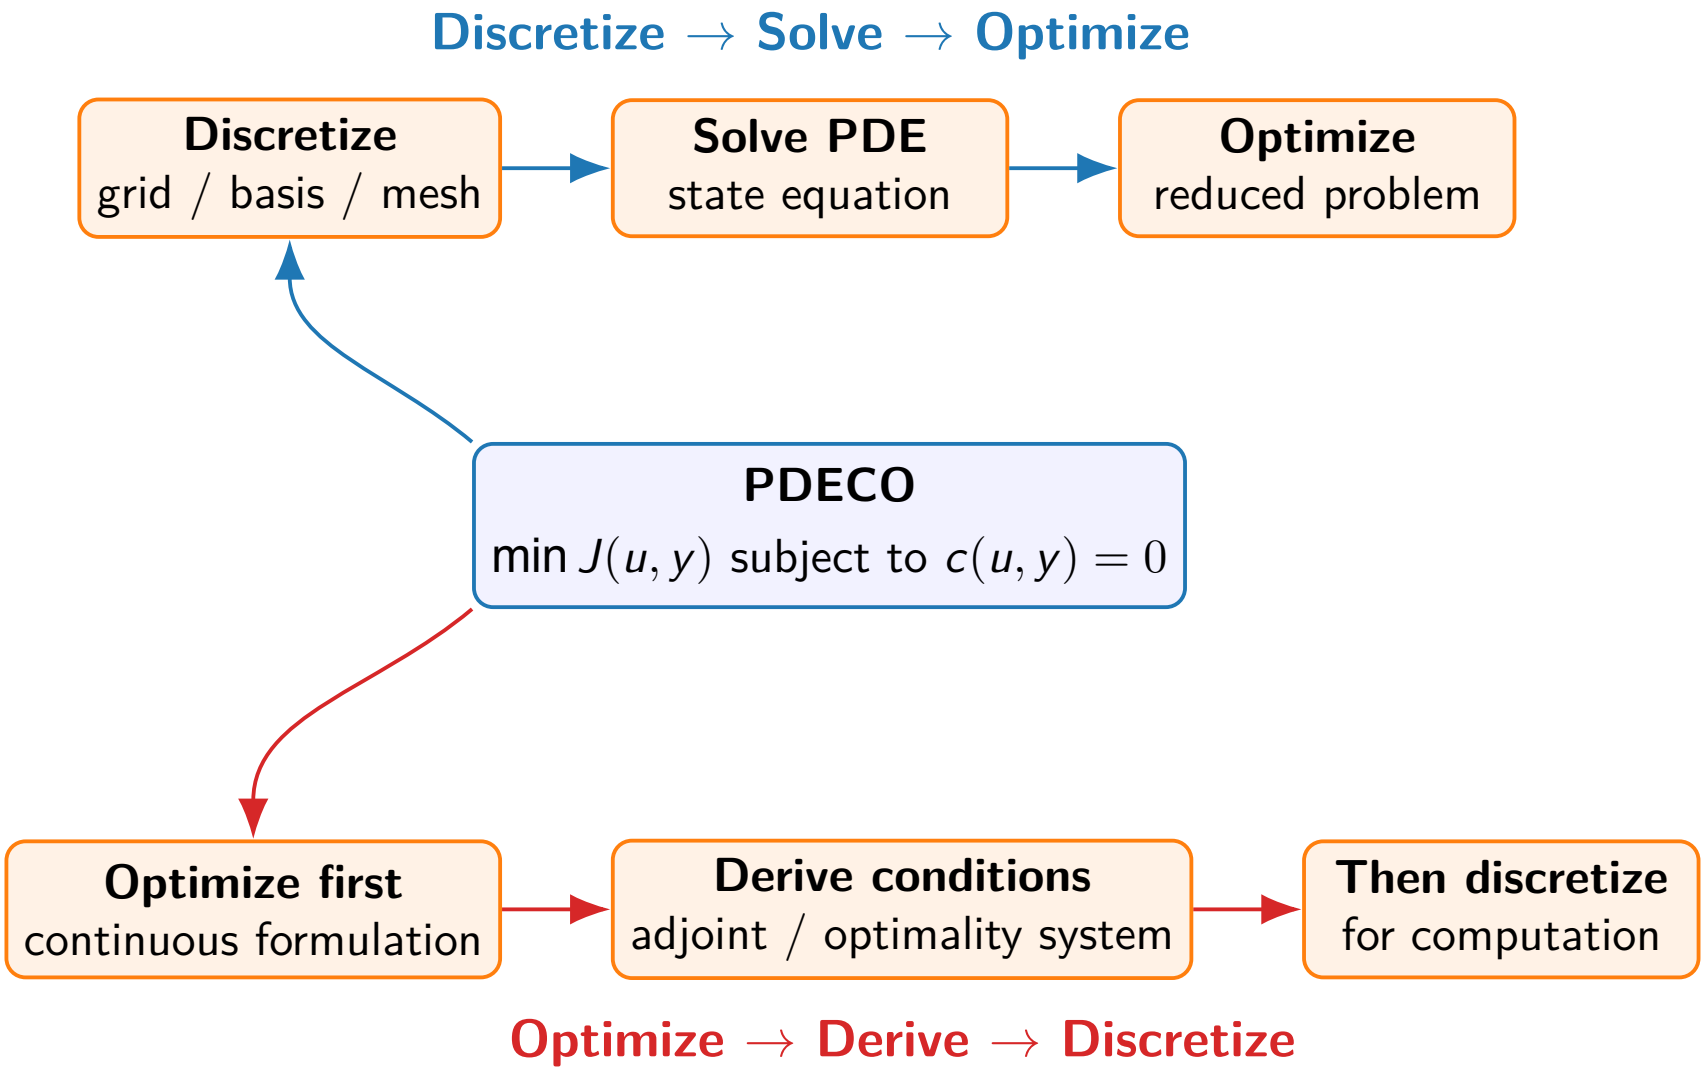
UMBC QSI Quantum Science Institute

What is PDE-Constrained Optimization?

PDE-constrained optimization (PDECO) is an optimization problem where we seek a control or design that minimizes an objective functional while satisfying a governing PDE constraint.

In many applications, the goal is to bring the state u as close as possible to a desired target u_d.

Two viewpoints for solving PDECO




Aerodynamics
(shape optimization)
Battery Cooling
(thermal management)
Pollution Control
(air dispersion)
PDE-Constrained Optimization
optimize decisions under physical laws
Wildfire Response
(spread / containment)
Chip Thermal Design
(thermal management)
PINNs / Scientific AI
(physics modeling)
$\frac{\partial u}{\partial t} = \mathcal{L}(u)$

PDE-Constrained Optimization: Formal Definition

Let $\Omega \subset \mathbb{R}^d$ be a bounded domain. Given a fixed desired state u_d, find an admissible control function $y(x)$ such that the corresponding state $u(x)$ minimizes an objective functional subject to a PDE constraint.

Formal Definition

$$\begin{aligned} \min_{u,y} \quad & J(u,y) \\ \text{subject to} \quad & c(u,y) = 0, \\ & u \in V_{\text{ad}}, \\ & y \in Y_{\text{ad}}. \end{aligned}$$

Notation

$u = u(x)$: State function,

u_d : Desired target state,

$y = y(x)$: Control function,

V_{ad} : Admissible state set,

Y_{ad} : Admissible control set.

Reduced-Space Formulation

Discretized PDECO Problem

$$\begin{aligned} \min_{u_h, x} \quad & J_h(u_h, x) \\ \text{s.t.} \quad & c_{h,i}(u_h, x) = 0, \quad i = 1, \ldots, n_u, \\ & x \in X_{\text{ad}}. \end{aligned}$$

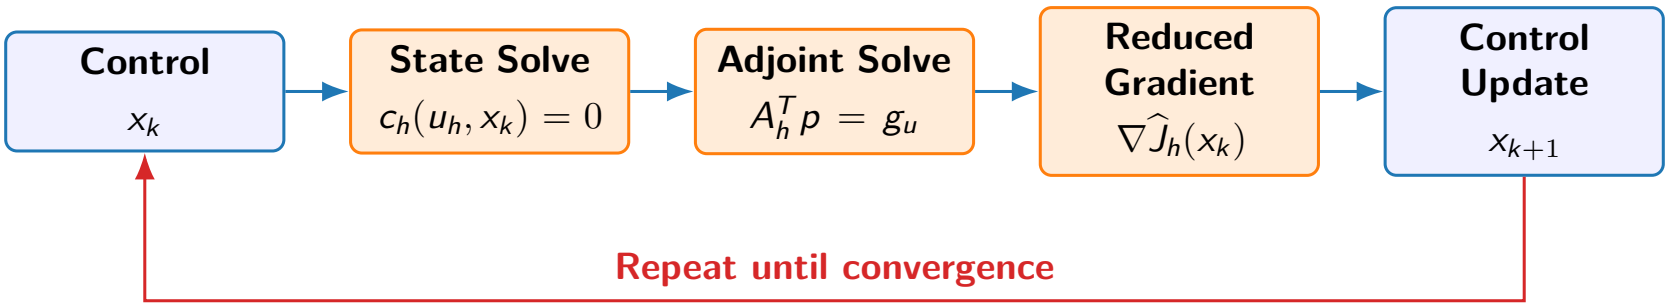
$$\Downarrow \quad u_h = u_h(x), \qquad c_h(u_h(x), x) = 0$$

Reduced Optimization Problem

$$\begin{aligned} \min_{x \in X_{\text{ad}}} \quad & \widehat{J}_h(x), \\ & \widehat{J}_h(x) := J_h(u_h(x), x). \end{aligned}$$

$u_h \in \mathbb{R}^{n_u}$, $x \in \mathbb{R}^{n_x}$, Ω_h : fixed grid of mesh size h

Computational Bottlenecks in PDECO



Main Bottlenecks: State solve, adjoint linear solve, and reduced-gradient evaluation.

Adjoint Method and Reduced Gradient

State Equation

$$c_h(u_h, x) = 0$$
$$u_h = u_h(x)$$

Adjoint Equation

$$A_h^T p = g_u$$
$$A_h = \frac{\partial c_h}{\partial u_h}, \qquad g_u = \nabla_{u_h} J_h(u_h, x)$$

Reduced Objective and Gradient

$$\widehat{J}_h(x) = J_h(u_h(x), x)$$

$$\boxed{\nabla \widehat{J}_h(x) = \nabla_x J_h - \left(\frac{\partial c_h}{\partial x}\right)^T p}$$

$$(\nabla \widehat{J}_h(x))_i = \frac{\partial J_h}{\partial x_i} - \left\langle p, \frac{\partial c_h}{\partial x_i} \right\rangle, \quad i = 1, \ldots, n_x$$

Key observation: One adjoint solve and n_x inner-product evaluations.

Quantum Linear System Algorithms (QLSAs)

Quantum Linear System Algorithms (QLSAs) are quantum algorithms designed to solve large linear systems by preparing a quantum state proportional to the solution vector.

Classical Linear Solver

$$Az = b$$

$$z = A^{-1}b$$

Returns the complete solution vector $z \in \mathbb{R}^N$.

Quantum Linear Solver

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

$$|z\rangle \approx \frac{A^{-1}|b\rangle}{\|A^{-1}|b\rangle\|}$$

Prepares a normalized quantum state, not the full classical vector.

QLSA Complexity and Motivation

- **HHL-based QLSA:** Prepares a quantum solution state with complexity

$$\widetilde{O}\left(\frac{s^2\kappa^2 \log N}{\varepsilon}\right).$$

QLSA Complexity and Motivation

- **HHL-based QLSA:** Prepares a quantum solution state with complexity

$$\widetilde{O}\left(\frac{s^2\kappa^2 \log N}{\varepsilon}\right).$$

- **Classical solvers:** Efficient multigrid methods can achieve $O(N)$ complexity under suitable assumptions.

QLSA Complexity and Motivation

- **HHL-based QLSA:** Prepares a quantum solution state with complexity

$$\widetilde{O}\left(\frac{s^2\kappa^2 \log N}{\varepsilon}\right).$$

- **Classical solvers:** Efficient multigrid methods can achieve $O(N)$ complexity under suitable assumptions.

- **HHL with full tomography:** Reconstructing the complete solution incurs at least $\Omega(N)$ cost, eliminating the exponential speedup in system dimension.

QLSA Complexity and Motivation

- **HHL-based QLSA:** Prepares a quantum solution state with complexity

$$\widetilde{O}\left(\frac{s^2\kappa^2 \log N}{\varepsilon}\right).$$

- **Classical solvers:** Efficient multigrid methods can achieve $O(N)$ complexity under suitable assumptions.

- **HHL with full tomography:** Reconstructing the complete solution incurs at least $\Omega(N)$ cost, eliminating the exponential speedup in system dimension.

- **Motivation for our PDECO algorithm:** Can we exploit QLSAs to compute the required reduced gradients **without reconstructing the full adjoint vector**?

SWAP Test

The **SWAP test** estimates the squared magnitude of the normalized inner product between two vectors $\phi, \psi \in \mathbb{C}^N$ using quantum measurements.

Quantum States

$$|\phi\rangle = \frac{\phi}{\|\phi\|}, \qquad |\psi\rangle = \frac{\psi}{\|\psi\|}$$

$$\langle\phi|\psi\rangle = \frac{\phi^\dagger \psi}{\|\phi\|\|\psi\|}$$

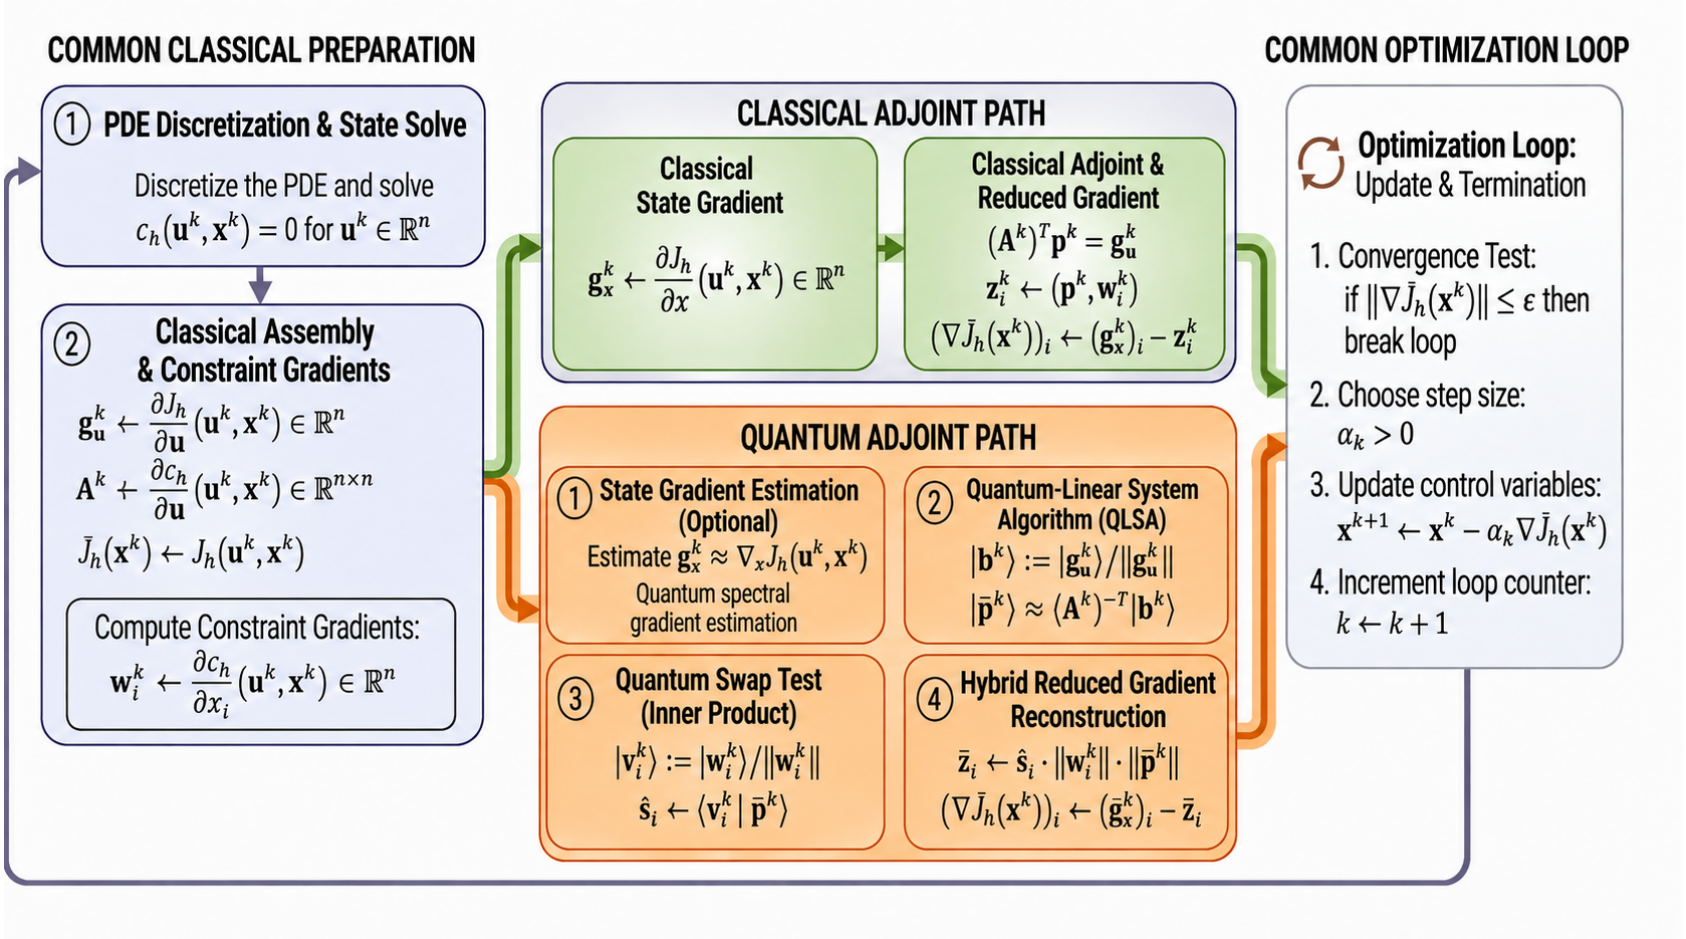
SWAP Test Measurement

$$P(0) = \frac{1 + |\langle\phi|\psi\rangle|^2}{2}$$

$$\boxed{|\langle\phi|\psi\rangle|^2 = 2P(0) - 1}$$

Key idea: Estimate the squared overlap without reconstructing either vector.

Classical vs. Quantum Paths



Classical vs. Hybrid Adjoint Complexity

Classical Adjoint Method (CG)

$$T_{\text{classical}} = O\left(s\, n_u \sqrt{\kappa} \log \frac{1}{\varepsilon} + n_x n_u\right)$$

Hybrid Quantum Adjoint Method (HHL)

$$T_{\text{quantum}} = O\left(\frac{n_x s^2 \kappa^2 \log n_u}{\varepsilon^3} + \frac{n_x}{\varepsilon^2}\right)$$

Potential quantum advantage: Polylogarithmic dependence on n_u under suitable assumptions.

Notation: n_u: state dimension, n_x: control dimension, s: sparsity, κ: condition number. $\varepsilon > 0$: target accuracy tolerance for the adjoint solve and inner-product estimation.

Numerical Experiments: Two PDECO Models

Experiment A: Nonlinear Heat

$$\frac{k}{\Delta t^2}(u_{i-1} - 2u_i + u_{i+1}) + \varepsilon(u_i)(T_\infty^4 - u_i^4) + h(T_\infty - u_i) + c_0 \sum_{j=0}^{n_x-1} x_j t_i^j = 0$$

$$J(u, x) = \frac{1}{2}\|u - u_d\|^2 + \frac{\beta}{2}\|x\|^2, \qquad u_d = 0.$$

Experiment B: Linear Elliptic Control

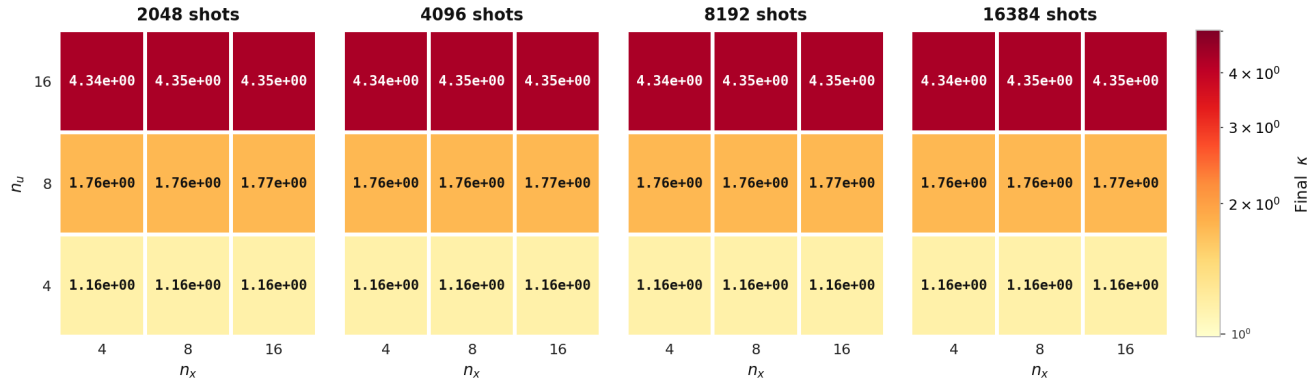
$$-\frac{d}{dt}\left(k(t)\frac{du}{dt}\right) + \mu u = Bx + f(t), \qquad u(0) = u(1) = 0.$$

$$J(u, x) = \frac{1}{2}\|u - u_d\|^2 + \frac{\alpha}{2}\|x\|^2, \qquad u_d = 2u^{\text{ref}}.$$

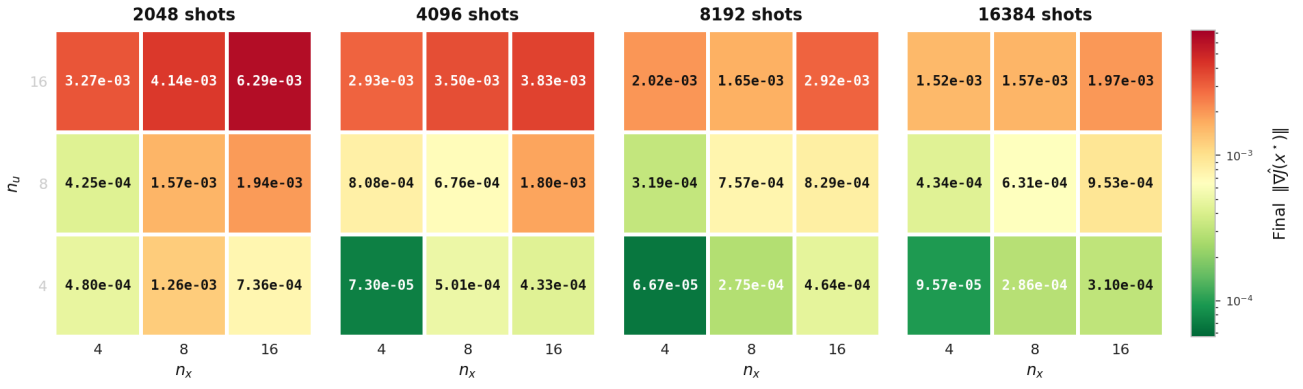
Five cases: Uniform diffusion; variable diffusion; diffusion contrast; localized control; nonzero source.

$$n_u, n_x \in \{4, 8, 16\}, \quad \text{shots} \in \{2048, 4096, 8192, 16384\}.$$

Nonlinear Heat: Effect of Shot Count and Conditioning



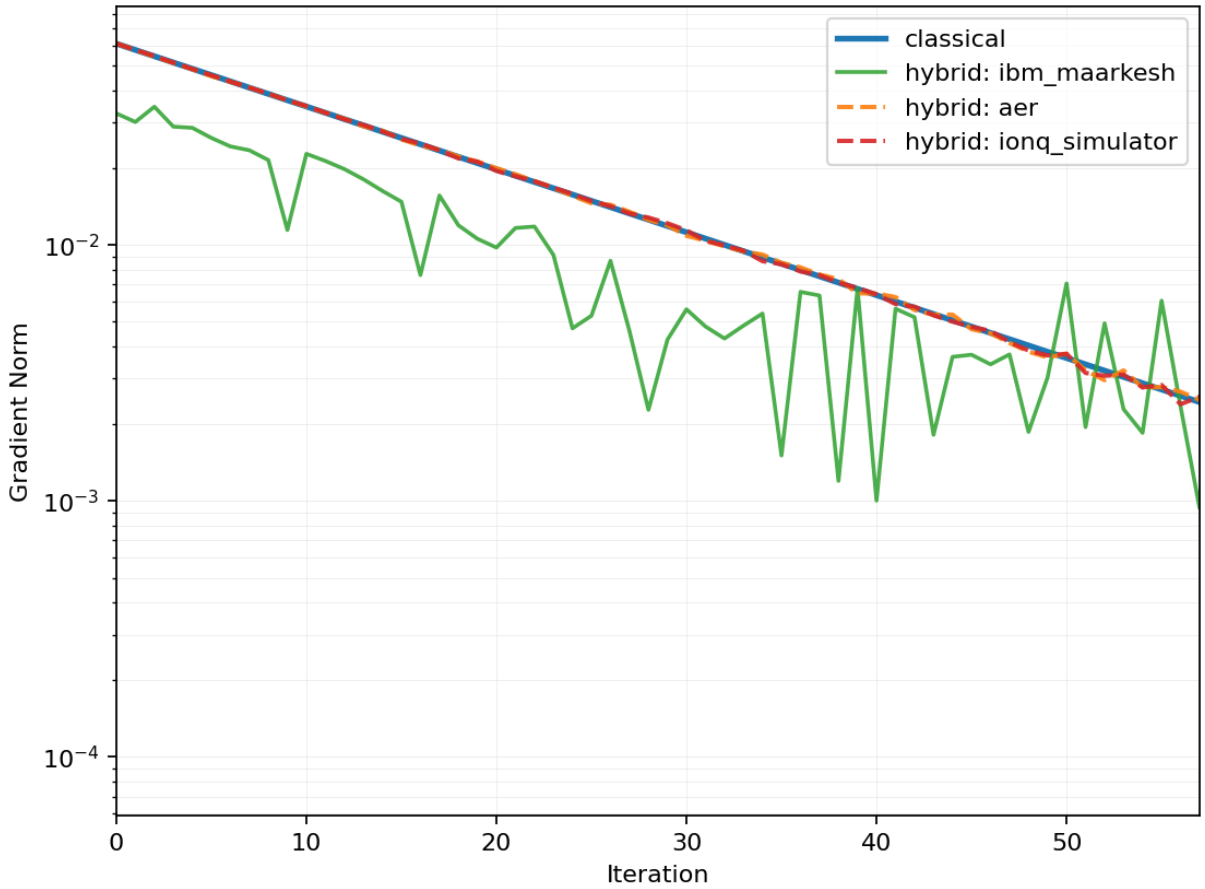
Condition number $\kappa_2(A)$



Final gradient norm $\|\nabla \hat{J}(x^\star)\|$

Observation: More shots improve gradient accuracy, while conditioning depends primarily on the state dimension.

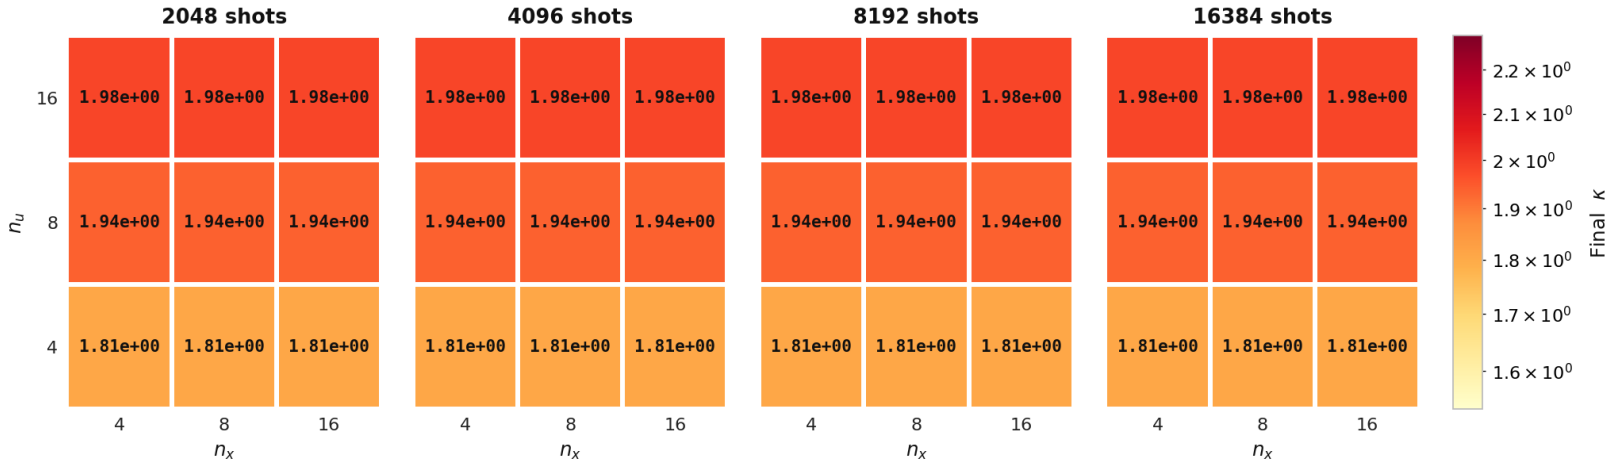
Nonlinear Heat: Results on Actual Quantum Hardware



Nonlinear heat : $n_x = 2$, $n_u = 4$; IBM Marrakesh (real quantum hardware, 2048 shots), Qiskit Aer simulator, IonQ simulator, and classical adjoint method.

Observation: Both simulators closely match classical convergence; IBM quantum hardware exhibits fluctuations but follows the same overall decreasing trend.

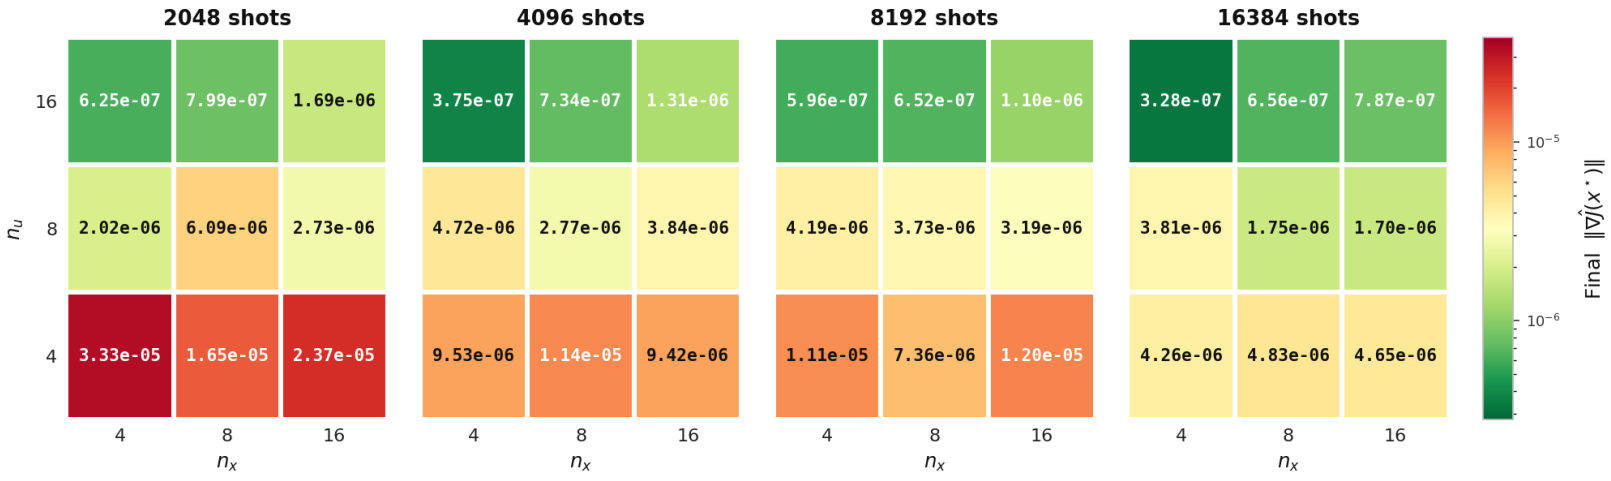
Elliptic: Condition Number



Experiment 1: $k(t) = 1,\ f(t) = 0.$

Observation: The condition number remains below 2 and is independent of the control dimension and shot count.

Elliptic: Gradient Accuracy



Experiment 1: $k(t) = 1$, $f(t) = 0$.

Observation: The final gradient norm decreases approximately as $O(n_u^{-2})$ in this experiment, primarily associated with the normalization of the discrete elliptic operator.

Conclusions and Future Directions

What We Accomplished

- Developed a hybrid quantum–classical adjoint method for PDECO.
- Used QLSA and quantum inner-product estimation to avoid full tomography.
- Established complexity and inexact-gradient convergence results.
- Validated the method using simulators and real IBM quantum hardware.

Future Directions

- Extend the framework to second-order adjoint methods.
- Investigate quantum-accelerated Hessian-vector products.
- Develop quantum preconditioning strategies to improve the efficiency of QLSA-based adjoint solves.

Thank You

Questions & Discussion

Quantum-Accelerated Adjoint Method for PDE-Constrained Optimization

Project Webpage

Project Details & Resources

Anish Ghosh

Dr.Mohammadhossein Mohammadisaihroudi

INFORMS 2026 Annual Meeting