Jun 8, 2025

Assessing the Variational Quantum Linear Solver for Fluid Dynamics on a Hybrid Quantum-HPC Stack

ACM ICS 2025

Chao Lu, Muralikrishnan Gopalakrishnan Meena,
Eduardo Antonio Coello Pérez, Amir Shehata,
Seongmin Kim, Kalyana Chakravarthi Gottiparthi, In-Saeng Suh

Oak Ridge National Laboratory

U.S. DEPARTMENT OF ENERGY

ORNL IS MANAGED BY UT-BATTELLE LLC
FOR THE US DEPARTMENT OF ENERGY

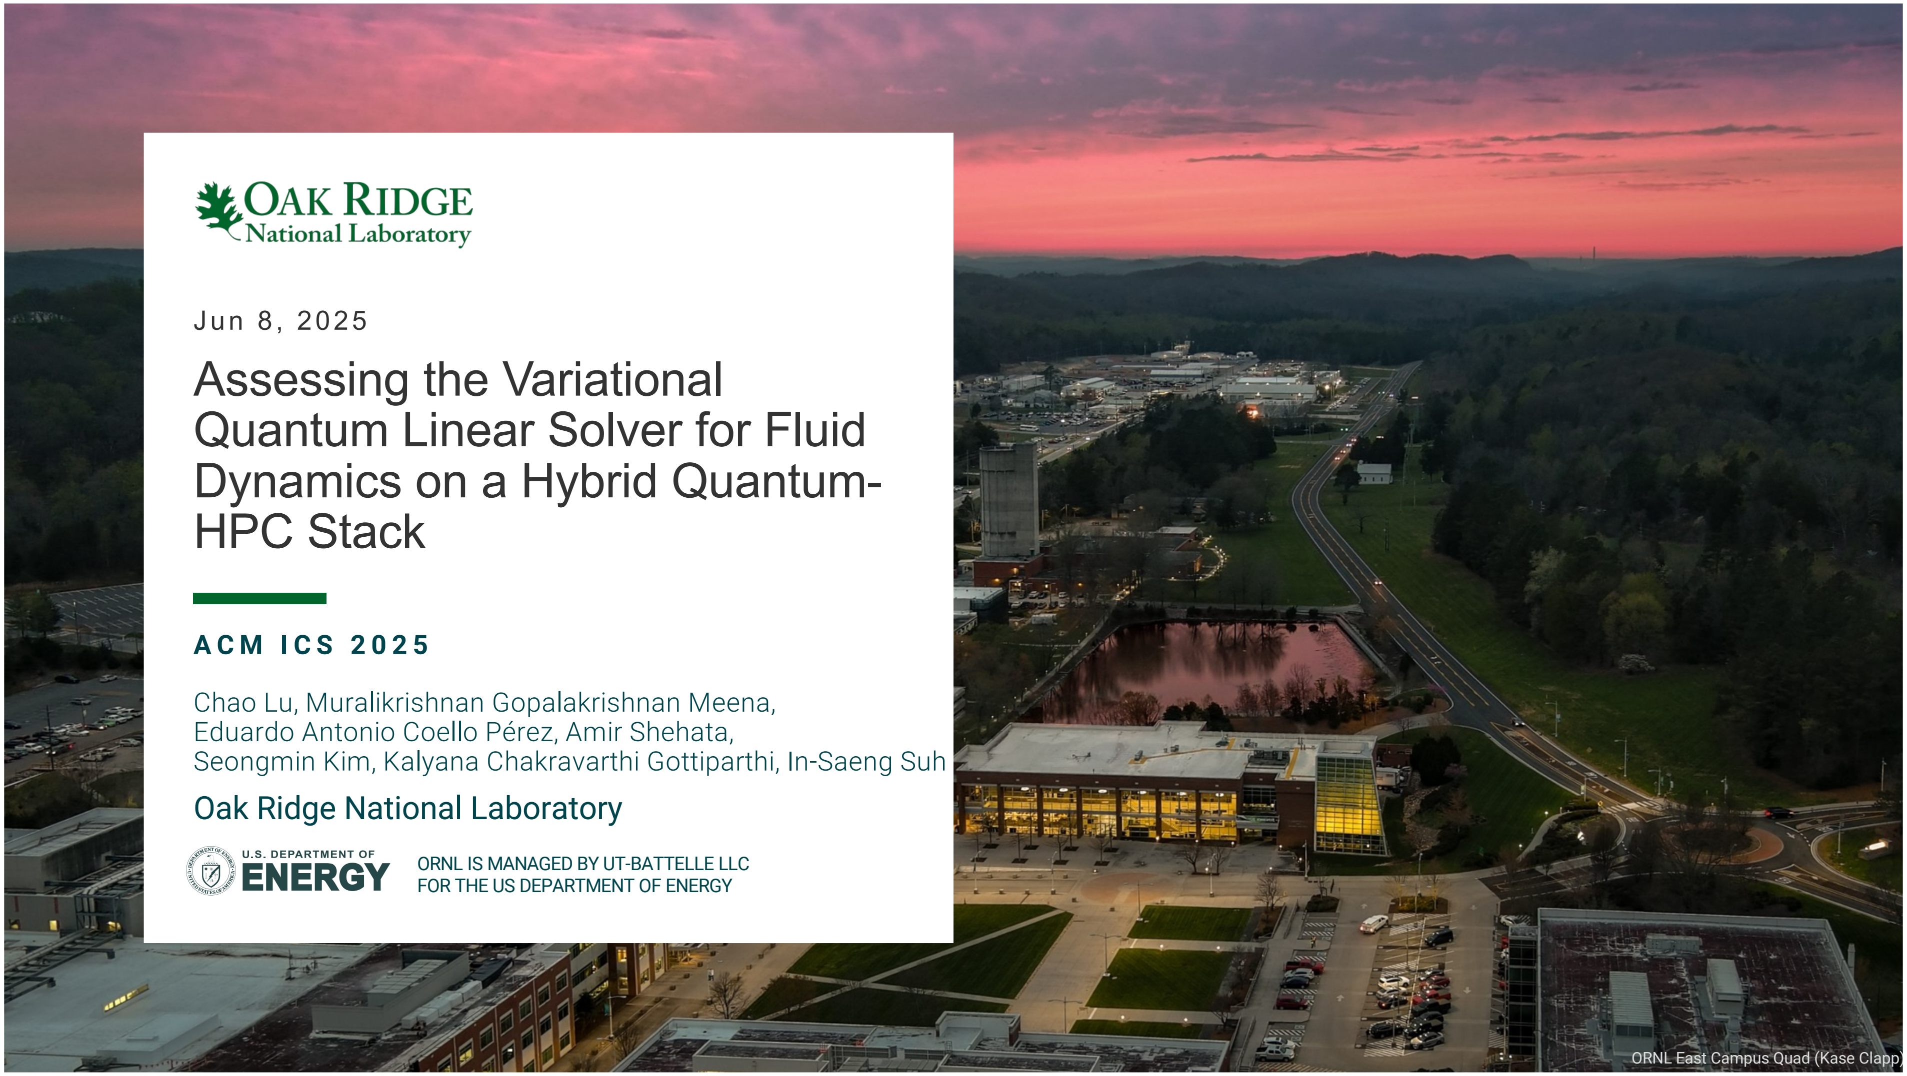

ORNL East Campus Quad (Kase Clapp)

Outline

Introduction

Challenges and Preliminary result

Conclusion and Future Works

Fluid dynamics is everywhere & its “pretty” turbulent

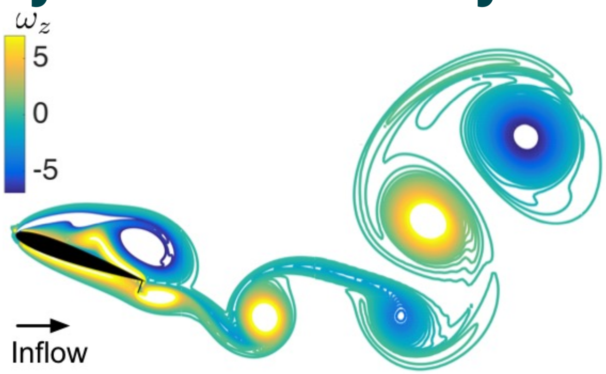

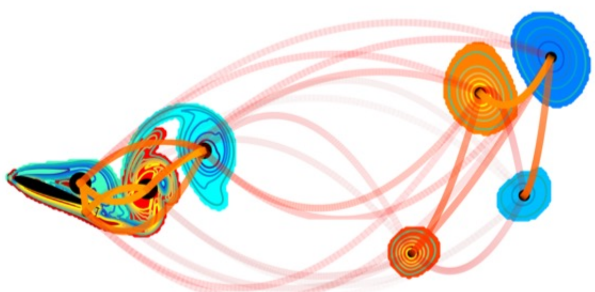

Reduce-order modeling of laminar flow over an airfoil[1]

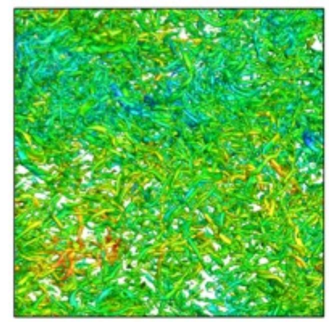

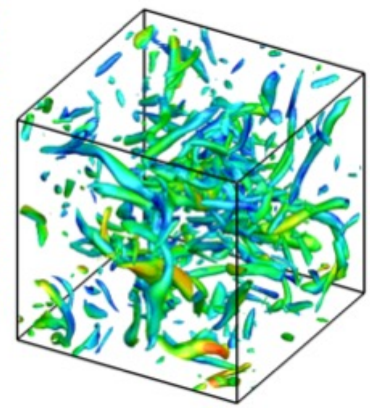

Community structures in 3D isotropic turbulence[2]

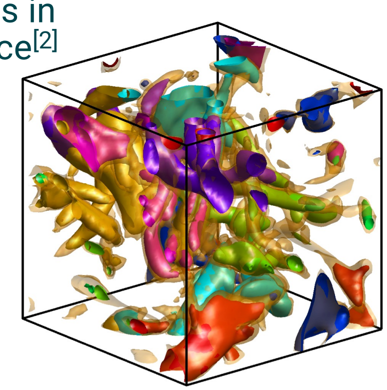

[1] Gopalakrishnan Meena, et al., *PRE*, 2018
[2] Gopalakrishnan Meena & Taira, *J. Fluid Mech.*, 2021

Grid resolution can get prohibitively expensive to simulate for practical fluid flow problems

- 3D stratified turbulence model for oceanographic flow[1]
- $22400 \times 22400 \times 5600 \approx 3 \times 10^{12} > 2^{41}$ grid points
- Turbulence related projects: 35-45% of 2023 OLCF Frontier allocation

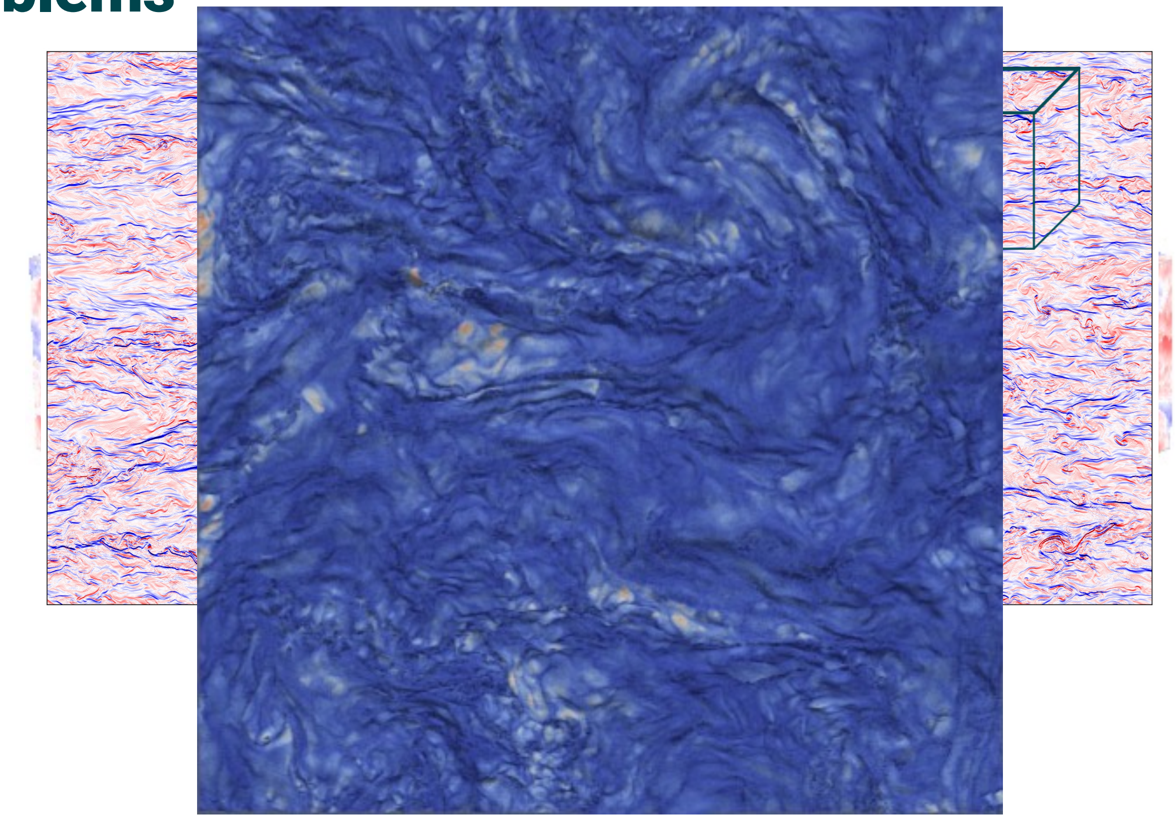

Isosurface of scalar dissipation rate for 1/25th of the domain, constructed using 127×10^6 triangles.

[1] J. J. Riley, et al., *J. Turbulence*, 2023

Quantum linear solvers have the potential to exponentially reduce cost of solving large problems

- Quantum computing applications to fluid flow problems:
  - Lattice simulations: fluid motion modeled as the motion of discrete particles
  - **Continuum simulations**: fluid motion modeled as a continuous field
  - Linear flow problems
- Linear (ideal) flow problems[1-4] : N—S equations with assumptions
- Use **Quantum Linear System Algorithms (QLSA)**
- Classical: $\boldsymbol{O}(\boldsymbol{N})$ (or higher for denser non-symmetric matrices)
- Quantum: $\boldsymbol{O}(\boldsymbol{log}(\boldsymbol{N}))$
- Disadvantage: Enlarged solution space

OAK RIDGE National Laboratory

[1] Yepez, *PRE,* 2001
[2] Xu, Daley, Givi, Somma, *AIAAJ,* 2018
[3] Bharadwaj & Sreenivasan, *PNAS*, 2023
[4] Gopalakrishnan Meena at al., *PoF*, 2024

Benchmarking LuGo for QLSA circuits

- Benchmark on QLSA to calculate tridiagnal Toeplitz linear systems
- Analysis on Frontier & Perlmutter supercomputers with a two-hour time limitation
- LuGo achieves reduction for: (1) time to generate and run circuits, and (2) circuit depth

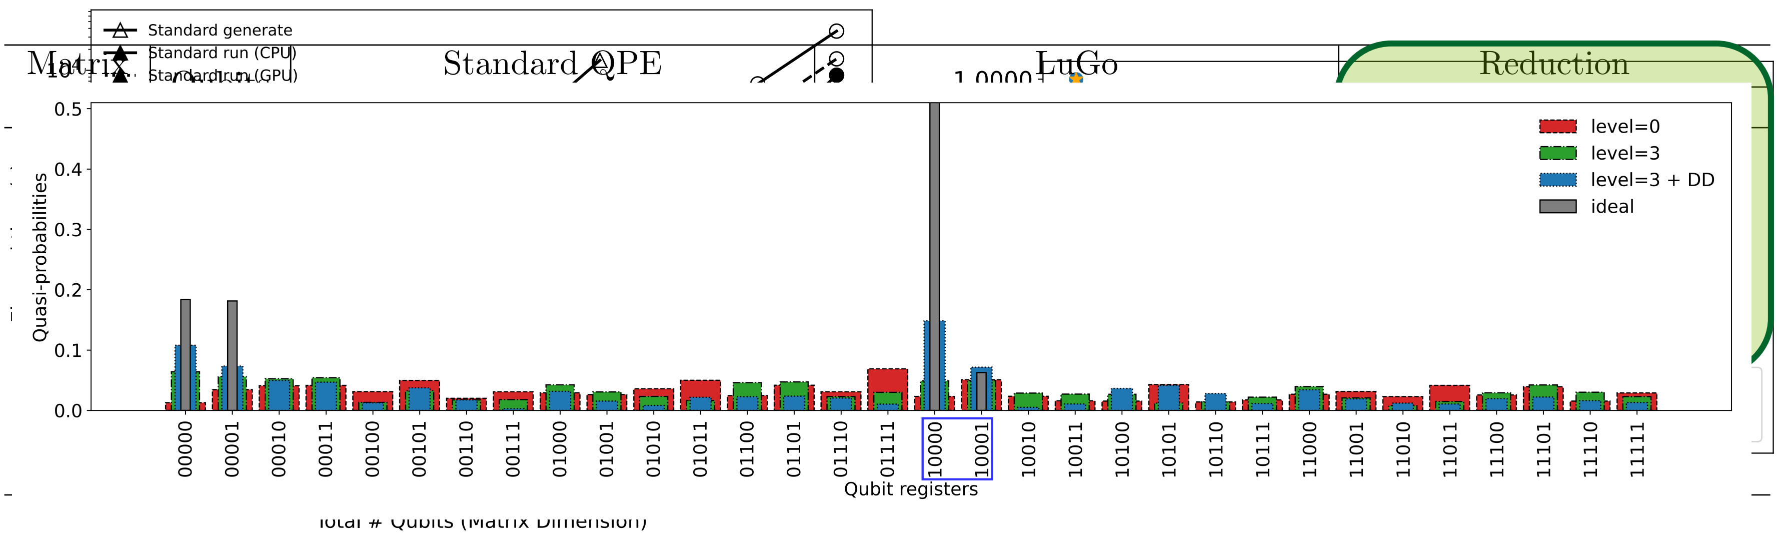

This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231.

LuGo-based HHL enables scaling to larger flow problems

Results using classical simulators

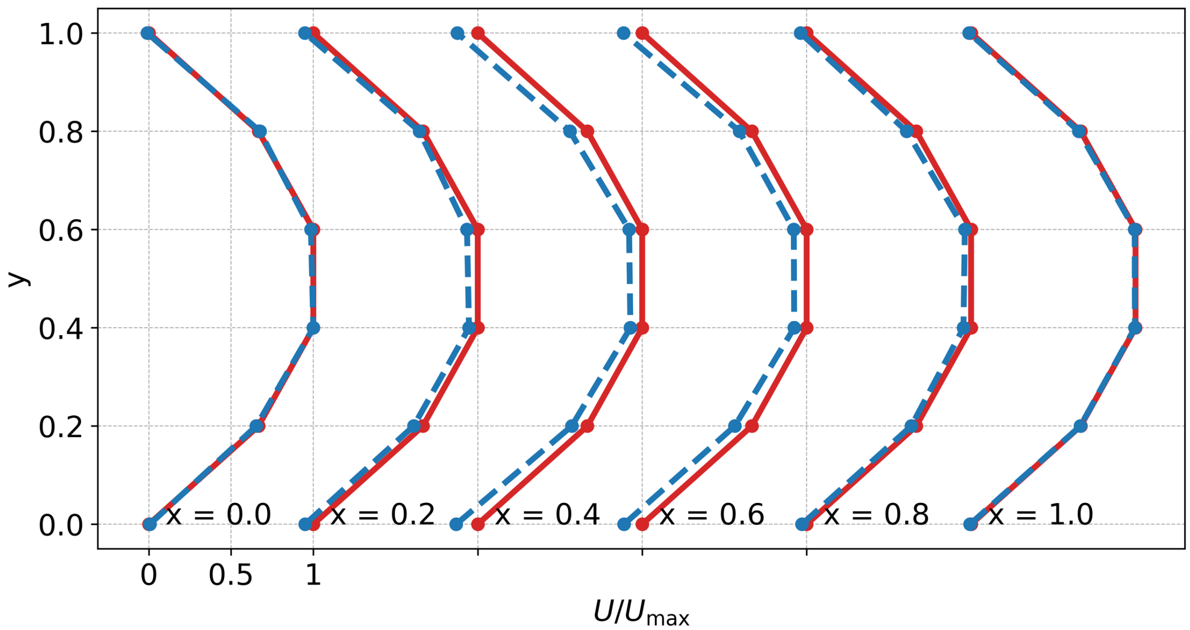

Oak Ridge National Laboratory

LuGo-based HHL solver better scales on superconducting & trapped-ion quantum hardware

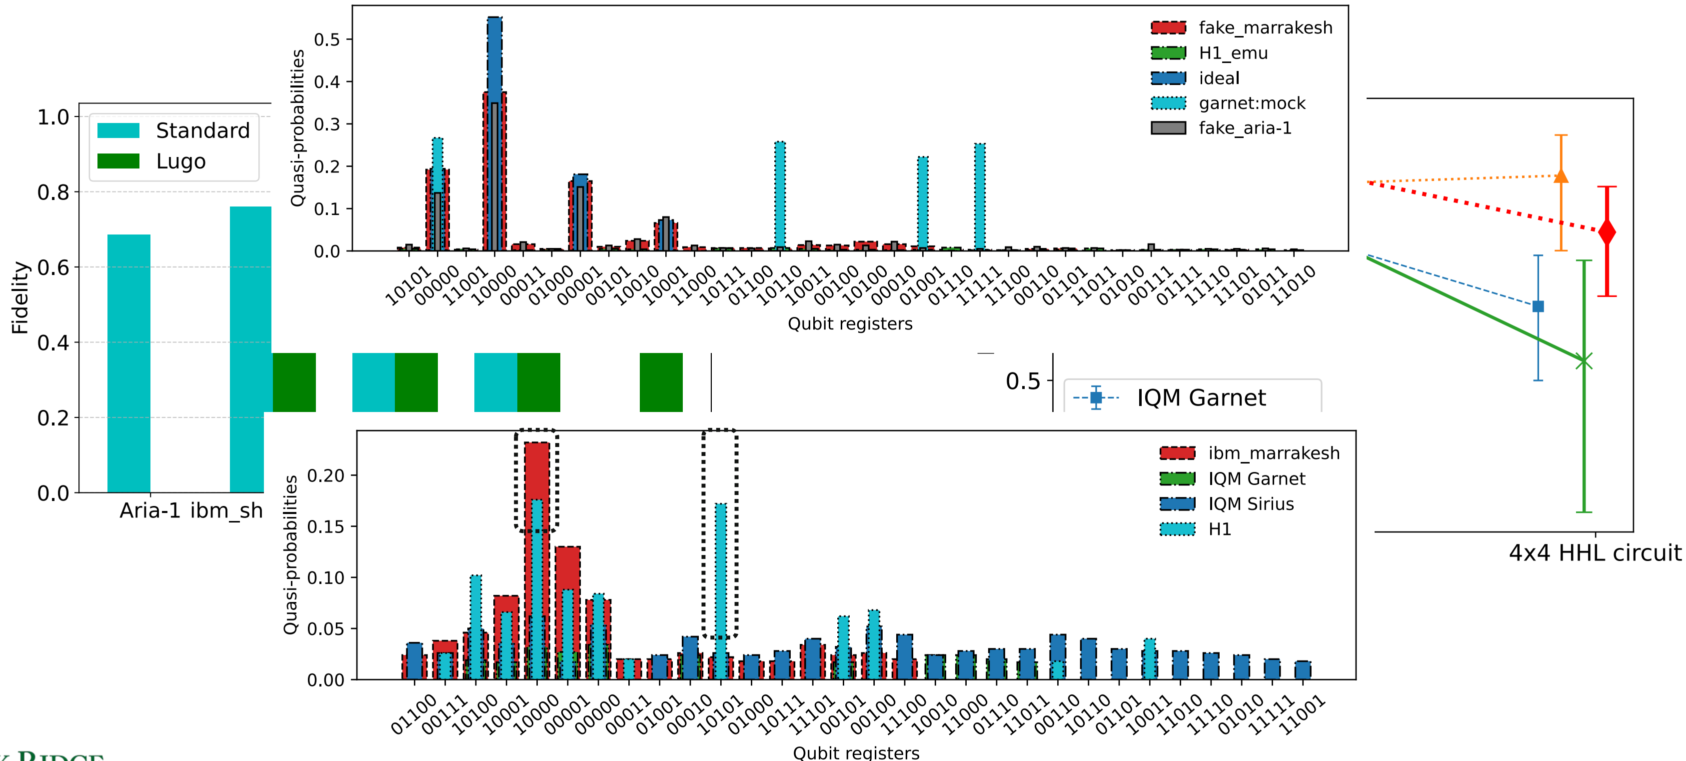

Due to the limitation of HHL algorithms, Variational Quantum Linear solver are considered as an alternative.

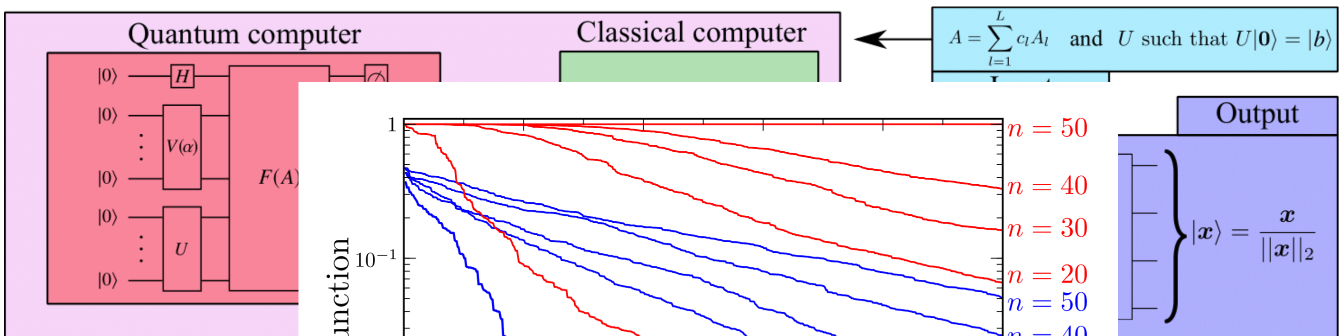

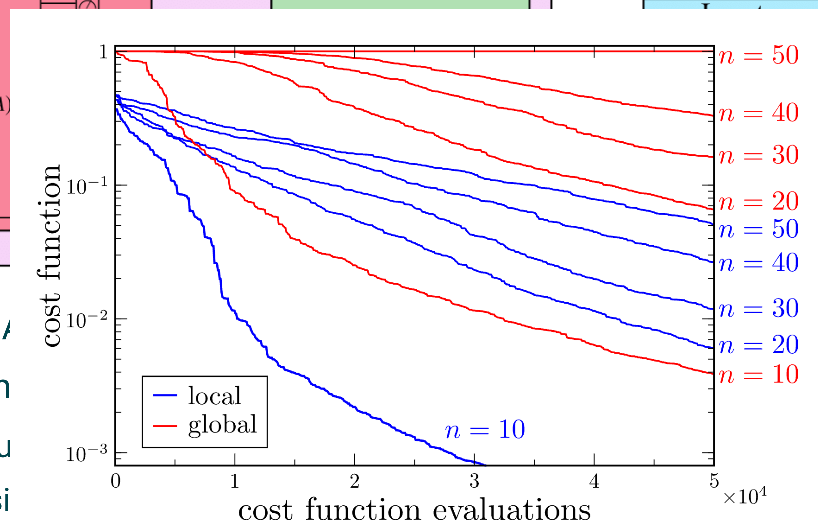

- Hybrid Quantum-Classical A
  - Practical for NISQ m
    - More robust to qu
    - Reduced circuit si
- The VQLS algorithm attempt to minimize the result of $1 - \langle b|A|x(\theta)\rangle^2$ to estimate the closeness of $Ax(\theta)$ is to the vector $|b\rangle$.
- The classical side will assign a set of weight to the quantum computer to estimate the cost.

[1] Bravo-Prieto, C. et, al. *Quantum*, *2020*.

Challenges

- VQLS for generalized matrices
- Scaling simple matrix
- Complexity involved with CFD matrix
- Ansatz complexity
- Optimizer needed
- Large # circuits
- Distributed framework
- Batched circuit submission

VQLS for generalized matrices – LCU decomposition

- To enable the generalized matrix input for VQLS, we use Linear Combination of Unitaries (LCU) decomposition.
- LCU decompose an arbitrary matrix to a combination of Pauli Matrices $\sum_j c_j U_j$
    - c_j refers to the coefficient;
    - U_j refers to the corresponding Pauli-matrices.

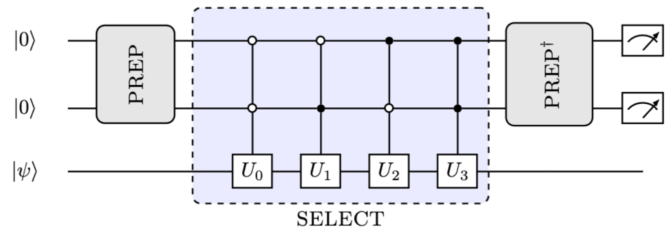

- For the VQLS algorithm, the matrix is initialized by multiple cost circuits with each circuit containing a set of Pauli-gates.
- The cost function is then obtained by adding up cost of all circuits multiply with corresponded coefficients.

VQLS for generalized matrices

- We constructed relatively simple matrices to test the scalability of the algorithm with LCU decomposition.
- For simplicity, we constructed the matrix with only three Pauli terms
    - with coefficient of 1, 0.2, 0.2 corresponds to Pauli-Z, and two Pauli-X terms.
- The matrix is now: $\begin{bmatrix} 1 & 0 & 0.4 & 0 \\ 0 & 1 & 0 & 0 \\ 0.4 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$ for 2-qubit matrix.
- The vector input is $[1,1,\dots,1]^T$.

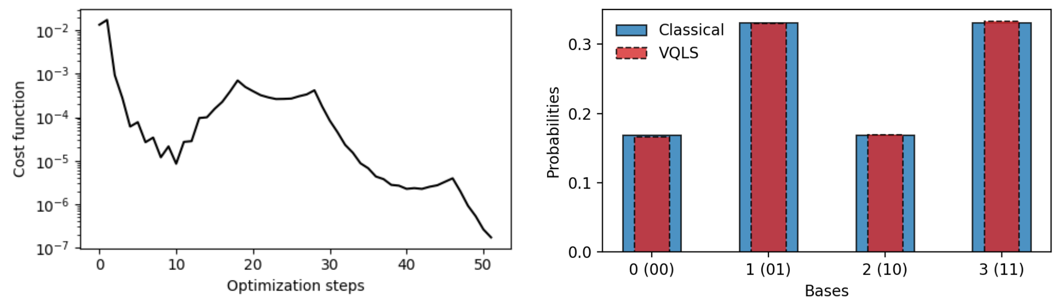

Validating the VQLS implementation on a tridiagonal Toeplitz system with matrix size 4 × 4. (Top) The convergence of the optimization process and (bottom) the comparison of probabilities of the states from the classical and VQLS methods.

Challenges

- VQLS for generalized matrices
- Scaling simple matrix
- Complexity involved with CFD matrix
- Ansatz complexity
- Optimizer needed
- Large # circuits
- Distributed framework
- Batched circuit submission

Scaling simple matrix

- For a matrix decomposition with a simple LCU decomposition method, we successfully scale the algorithm to matrix size of $2^8 \times 2^8$ with high fidelity over 99.9% on quantum simulator.

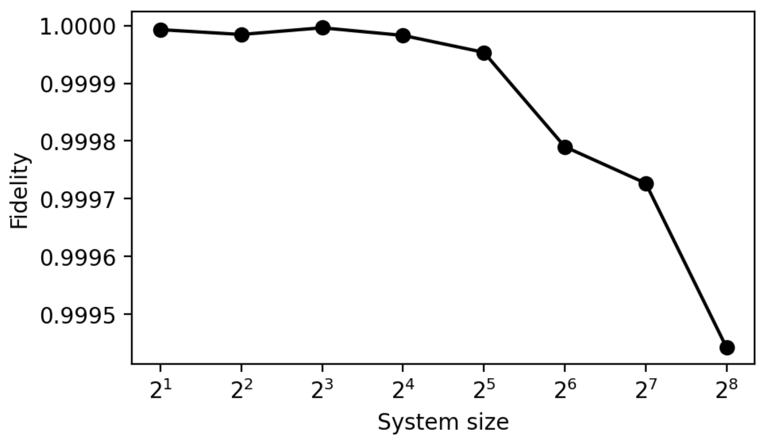

Effect of problem size on fidelity demonstrated for a tridiagonal Toeplitz system.

OAK RIDGE National Laboratory

Challenges

VQLS for generalized matrices

Scaling simple matrix

Complexity involved with CFD matrices

Ansatz complexity

Optimizer needed

Large # circuits

Distributed framework

Batched circuit submission

Complexity involved with CFD matrices

- Not with dimension of $2^N \times 2^N$ – pad to the nearest square matrix of $2^N \times 2^N$.
- Large condition numbers – VQLS can handle input matrix with ill-conditioned matrices where HHL cannot.
- An increased number of LCU terms for complex – more quantum circuits required for each cost computation.
- Post-quantum process – from quantum state to actual state.
- Complex result state $|x\rangle$ that is difficult to express – more complex quantum ansatz for expressibility.

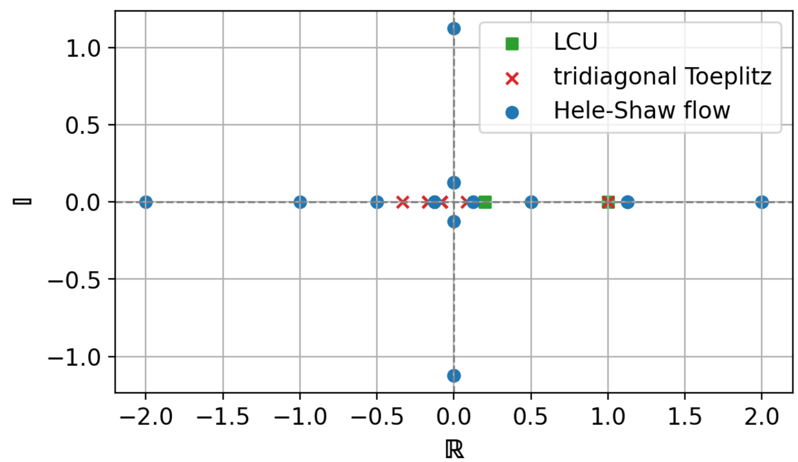

Comparing the coefficients of the LCU decomposition of the system matrices for the tridiagonal Toeplitz system and Hele−Shaw flow on the complex plane. Both problems are chosen to have a system size of 24

OAK RIDGE National Laboratory

Challenges

VQLS for generalized matrices

Scaling simple matrix

Complexity involved with CFD matrices

Ansatz complexity

Optimizer needed

Large # circuits

Distributed framework

Batched circuit submission

Ansatz complexity

- Due to the complexity of real-world problems with the possible solution $|x\rangle$, ansatz with low depth and complexity might fail to converge.
- We are considering multiple different types of quantum ansatz for efficient convergence and high fidelity.

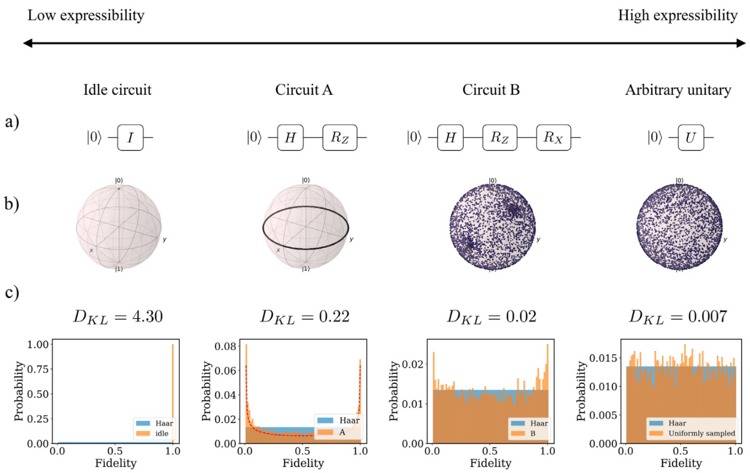

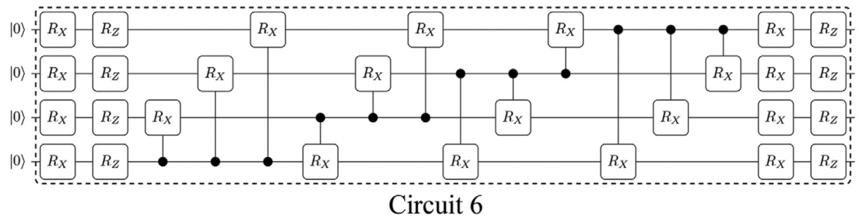

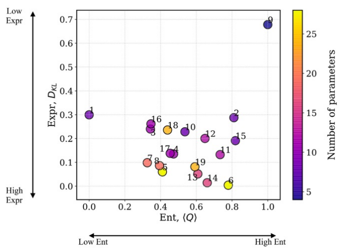

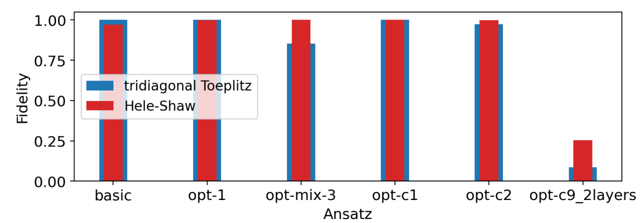

Ansatz expressibility, entanglement, and number of parameters explanations, and preliminary results obtained of the performance achieved on tridiagnal Toeplitz and Hele-Shaw experiments conducted using different quantum ansatz. Comparing effect of various ansatz on the fidelity for the tridiagonal Toeplitz system and Hele–Shaw flow. Both problems are chosen to have a system size of 8.

[1] SukinSim, et al., *Adv. QuantumTechnol. 2019*

Challenges

VQLS for generalized matrices

Scaling simple matrix

Complexity involved with CFD matrices

Ansatz complexity

Optimizer needed

Large # circuits

Distributed framework

Batched circuit submission

Optimizer needed and large number of circuits

- To minimize the cost function with less optimization steps, we used gradient and gradient-free optimizers to compare the number of steps and minimum cost function achieved.
- Currently, gradient-based optimizer ‘Adam’ and ‘NAdam’ optimizer in Pytorch outperformed gradient-free optimizer with lowest steps and minimum cost value obtained.
- Due to many steps and decomposed Pauli terms because of LCU decomposition, large number of circuits are required to perform the optimization algorithm.

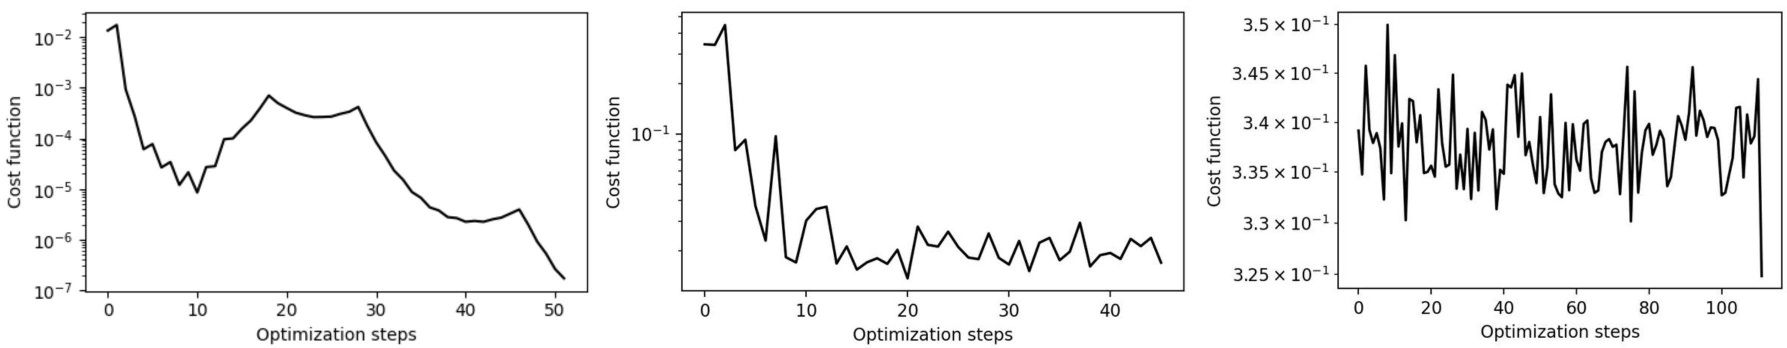

(a) ‘NAdam’ (gradient-based) (b) COBYLA (gradient-free) (c) L-BFGS-B

Learning curve by selecting different optimizer for a sample LCU decomposition matrix input using gradient and gradient-free approach.

Challenges

VQLS for generalized matrices

Scaling simple matrix

Complexity involved with CFD matrices

Ansatz complexity

Optimizer needed

Large # circuits

Distributed framework

Batched circuit submission

Distributed framework

- To cope with HPC-Quantum stack and release the pressure of large amount of quantum circuits for VQLS, we embedded QFW framework to enable Quantum-HPC

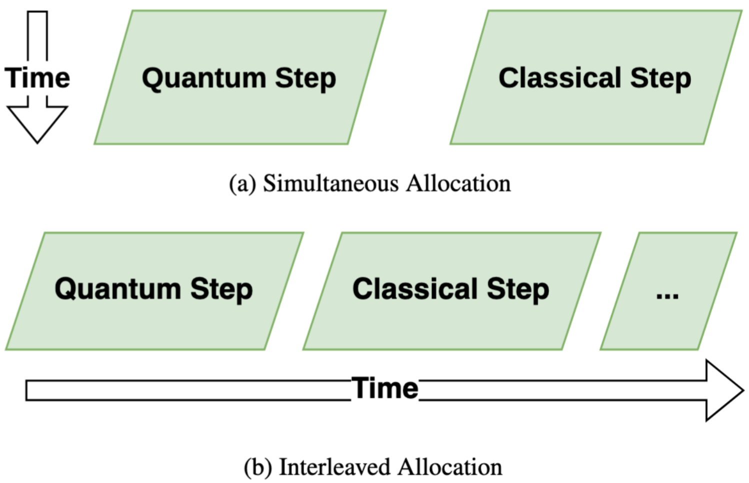

Allocation strategies to support hybrid QC/HPC applications.

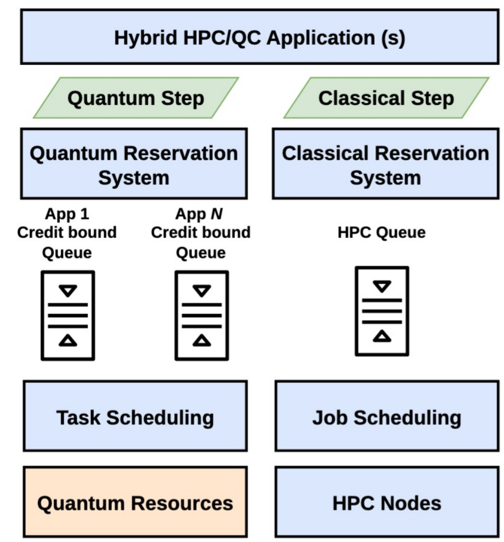

Illustration of hybrid QC/HPC application resource management.

Challenges

VQLS for generalized matrices

Scaling simple matrix

Complexity involved with CFD matrices

Ansatz complexity

Optimizer needed

Large # circuits

Distributed framework

Batched circuit submission

Batched circuit submission

- To reduce the interaction between classical and quantum computer interaction and maximize the quantum hardware qubit capacity, we enabled batched circuit input for a quantum computer/simulator.

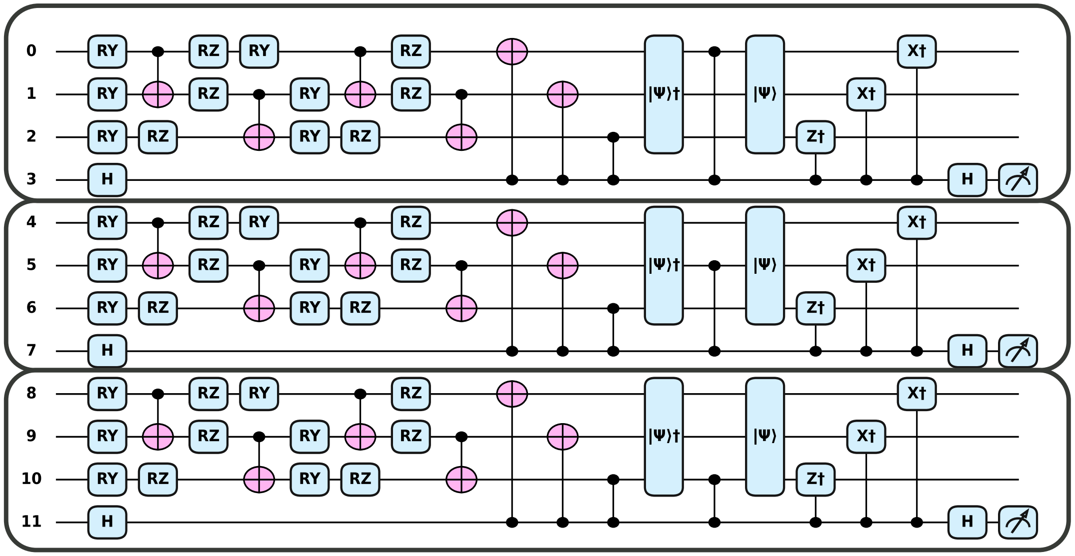

Batched VQLS circuit a batched submission of three circuits. Three cost values will be return by one execution

OAK RIDGE National Laboratory

Outline

Background and Challenges

Preliminary Results

Conclusion and Future Works

OAK RIDGE National Laboratory

Conclusion and Future Works

- In the future, we intend to
  - perform more detailed analysis of quantum ansatz with respect to expressibility, weights and entanglement capability;
  - Reduce the total number of circuits required for VQLS.

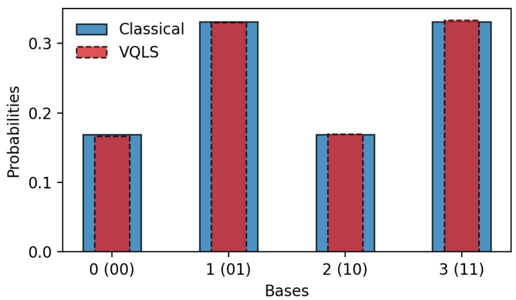

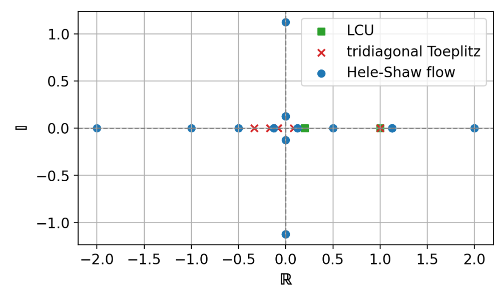

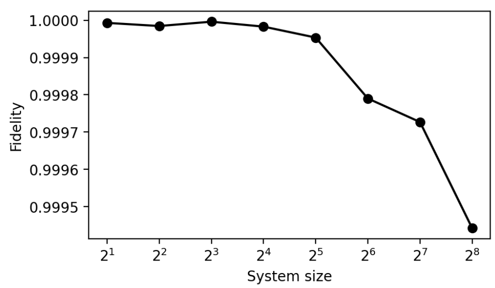

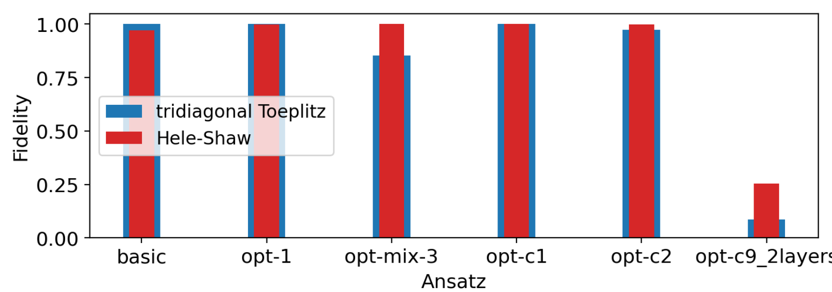

OAK RIDGE National Laboratory

Acknowledgement

This research used resources of the **Oak Ridge Leadership Computing Facility (OLCF) at the Oak Ridge National Laboratory**, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

OAK RIDGE
National Laboratory

Oak Ridge National Laboratory

U.S. DEPARTMENT OF ENERGY

ORNL IS MANAGED BY UT-BATTELLE LLC
FOR THE US DEPARTMENT OF ENERGY