Performance Analysis of Phaseless Auxiliary Field Quantum Monte Carlo

Namita Shah[1], Neil Mehta[2], Pooja Rao[3], Nickolas Joyner[2], Katie Klymko[2], Ermal Rrapaj[2], Lisa Claus[2], Charles Lively[2]

[1]University of Michigan, [2]Lawrence Berkeley National Laboratory, [3]NVIDIA Corporation

U.S. DEPARTMENT of ENERGY
Office of Science

Research Question

What are the cost–accuracy trade-offs among quantum state-preparation methods for determining the lowest ground state energy of a molecule?

Introduction

Why the ground-state energy

A molecule's behavior depends on how its electrons interact. To do this, we need to know its **ground-state energy**: the lowest energy the electrons can reach. However, this is computationally very difficult to measure exactly so we use Monte Carlo methods.

ph-AFQMC: Monte Carlo with a guide

Phaseless auxiliary-field quantum Monte Carlo (ph-AFQMC) [1] estimates the ground-state energy by simulating many *walkers* (candidate electron configurations) that evolve through random steps toward the ground state (Fig. 1d). A **trial wavefunction** provides an approximate "map" to help guide the walkers. A better "map" could lead to a more accurate energy estimate and a lower computational cost.

This experiment explores the tradeoffs between 2 different methods:

- Variational Quantum Eigensolver (VQE)
  - UCCSD, ADAPT-UCCSD, LUCJ
- Generative Quantum Eigensolver (GQE)

Methods Compared

Table 1. Trial-wavefunction preparation methods

Method	How the circuit is built
VQE (Variational Quantum Eigensolver) [2]: fixed circuit template; a classical optimizer tunes its angles to minimize energy.	
UCCSD [3]	Unitary Coupled Cluster, Singles and Doubles: all one- and two-electron excitations. Deepest circuit.
ADAPT-UCCSD [4]	Adds the excitation with the largest energy gradient, one per round.
LUCJ [5]	Local Unitary Cluster Jastrow: orbital rotations plus pair terms on neighboring qubits.
GQE (Generative Quantum Eigensolver) [6]: a GPT-style transformer selects *N* operators from a fixed pool; training updates the model.	

Experimental Set-Up

Hardware: NERSC Perlmutter, 1 GPU node (4 × NVIDIA A100).
Geometries: R = 0.8–1.5 Å (H_4) and 0.8–1.3 Å (H_8), in 0.1 Å steps.

Table 2. Problem size and circuit complexity

Quantity	H_4	H_8
Electrons	4	8
Qubits (Jordan-Wigner)	8	16
UCCSD excitations (VQE angles)	26	360
UCCSD Pauli rotations	160	2,688
GQE pool (rotations × 12 angles)	1,920	32,256
GQE chosen rotations, *N*	**10, 20**	**10, 20**

Computational Workflow

Step 1: Build the molecule

Linear H_4 and H_8 chains with spacing R. A larger R means stronger electron correlation and a harder problem.

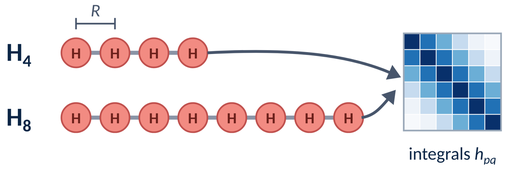


Figure 1a. Integrals mapped to 2 qubits per orbital (Jordan–Wigner).

Step 2a: VQE trial

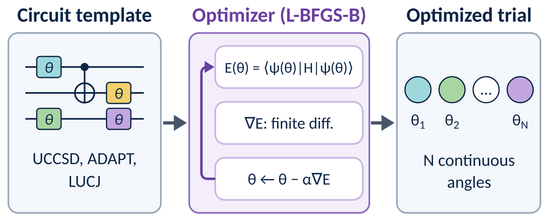


Figure 1b. Option 1 is VQE: the optimizer loops over energy and gradient evaluations of the simulated circuit, updating its N angles.

Step 2b: GQE trial

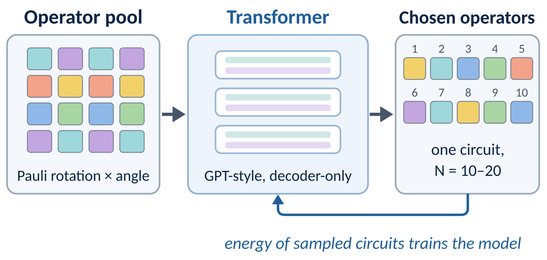


Figure 1c. Option 2 is GQE: the transformer picks N operators, each a UCCSD Pauli rotation with one of 12 fixed angles.

Step 3: ph-AFQMC random walks

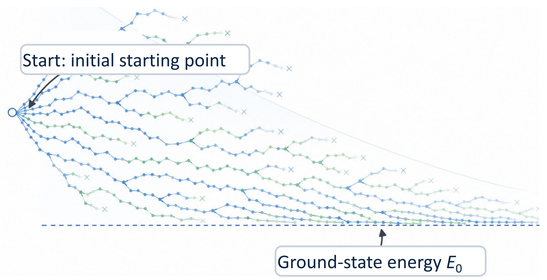


Figure 1d. Walkers start from the Hartree–Fock state, branch, or are removed (×) as they propagate to imaginary time. The trial wave function produced by Step 2 guides them toward E_0.

Cost vs. Model Size

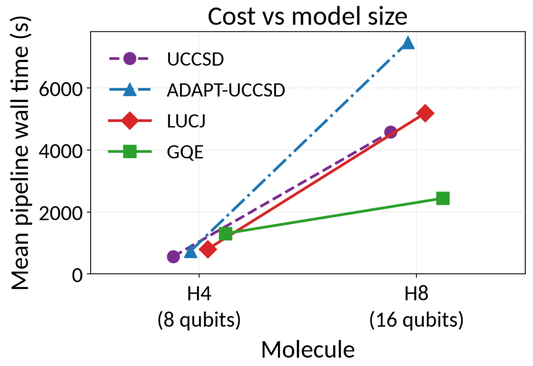


Figure 4. Mean pipeline wall time over R = 0.8, 1.0, 1.2 Å. GQE is slowest for H_4 but fastest for H_8 (~2,450 s vs 4,500–7,400 s).

Performance (Wall Time)

Pipeline wall time: Hamiltonian build + trial wavefunction + ph-AFQMC, on 1 node (4 × A100 GPUs).

H_4 (8 qubits)

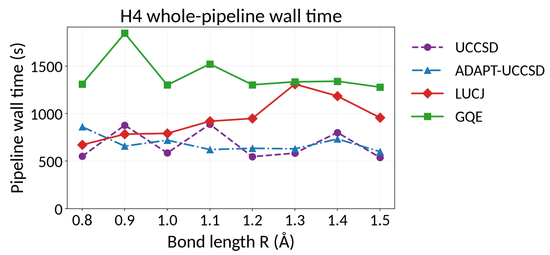


Figure 2. H_4 whole-pipeline wall time (build + trial + ph-AFQMC). GQE is slowest at every R because transformer training is a fixed cost.

H_8 (16 qubits)

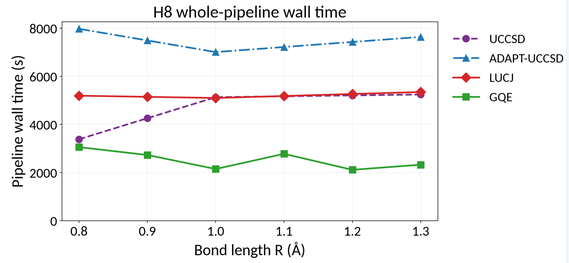


Figure 3. H_8 whole-pipeline wall time. GQE is the fastest method at every bond length; ADAPT-UCCSD takes 7,000–7,900 s (2.6–3.5× GQE), and UCCSD and LUCJ take 3,400–5,350 s.

Accuracy

ph-AFQMC energy error relative to FCI, the exact answer (black line). HF (×) is the mean-field baseline. **Chemical accuracy**: within 1.6 mHa (1 kcal/mol) of FCI (dotted line).

H_4 (8 qubits)

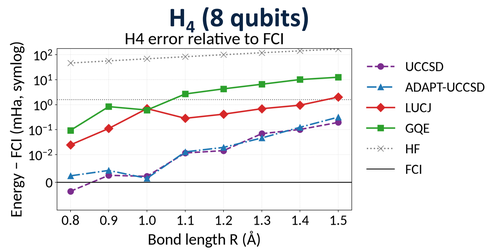


Figure 5. H_4 error vs FCI (dotted line = 1.6 mHa). GQE stays within chemical accuracy up to R = 1.0 Å but exceeds it from 1.1 Å. UCCSD and ADAPT-UCCSD stay below 0.3 mHa; LUCJ below accuracy to 1.4 Å.

H_8 (16 qubits)

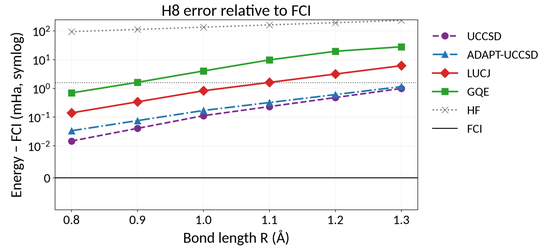


Figure 6. H_8 error vs FCI (dotted line = 1.6 mHa). GQE error is 0.7 mHa at R = 0.8 Å but grows quickly. UCCSD and ADAPT-UCCSD stay at or below ~1 mHa at every R; LUCJ exceeds 1.6 mHa from 1.2 Å.

Takeaways

1. **GQE starts to be comparable with VQE for larger molecules:** GQE cost grows ~1.9× from H_4 to H_8, vs ~6–10× for VQE.
   a. Near-exact energies near equilibrium: within chemical accuracy for $R \le 1.0$ Å (H_4) and at 0.8 Å (H_8).
   b. Competitive cost: fastest for H_8 at every bond length.
2. **Not yet reliable when stretched:** GQE error exceeds 1.6 mHa from 1.1 Å (H_4) and reaches ~27 mHa at 1.3 Å (H_8).

Next Steps

- Increase to multi-node experiments.
- Sample more bond lengths to find where GQE accuracy breaks down.
- Extend to real molecules with heavier atoms (e.g., O_2).

References

1. R. Rofougaran *et al.*, arXiv:2605.02056 (2026).
2. A. Peruzzo *et al.*, *Nat. Commun.* 5, 4213 (2014).
3. J. Romero *et al.*, *Quantum Sci. Technol.* 4, 014008 (2019).
4. H. R. Grimsley *et al.*, *Nat. Commun.* 10, 3007 (2019).
5. M. Motta *et al.*, *Chem. Sci.* 14, 11213 (2023).
6. K. Nakaji *et al.*, arXiv:2401.09253 (2024).
7. M. Motta *et al.*, *Phys. Rev. X* 7, 031059 (2017).

QR Code: scan for important links related to this work!

Acknowledgements

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