SC26 PAPER PRESENTATION

OptiPESS

A GPU-Accelerated and Geometry-Aware Quantum-Classical Framework for Optimized PES Scans with CUDA-Q

Panayiotis Christou, Shuwen Kan, and Ying Mao

Fordham University

FORDHAM UNIVERSITY

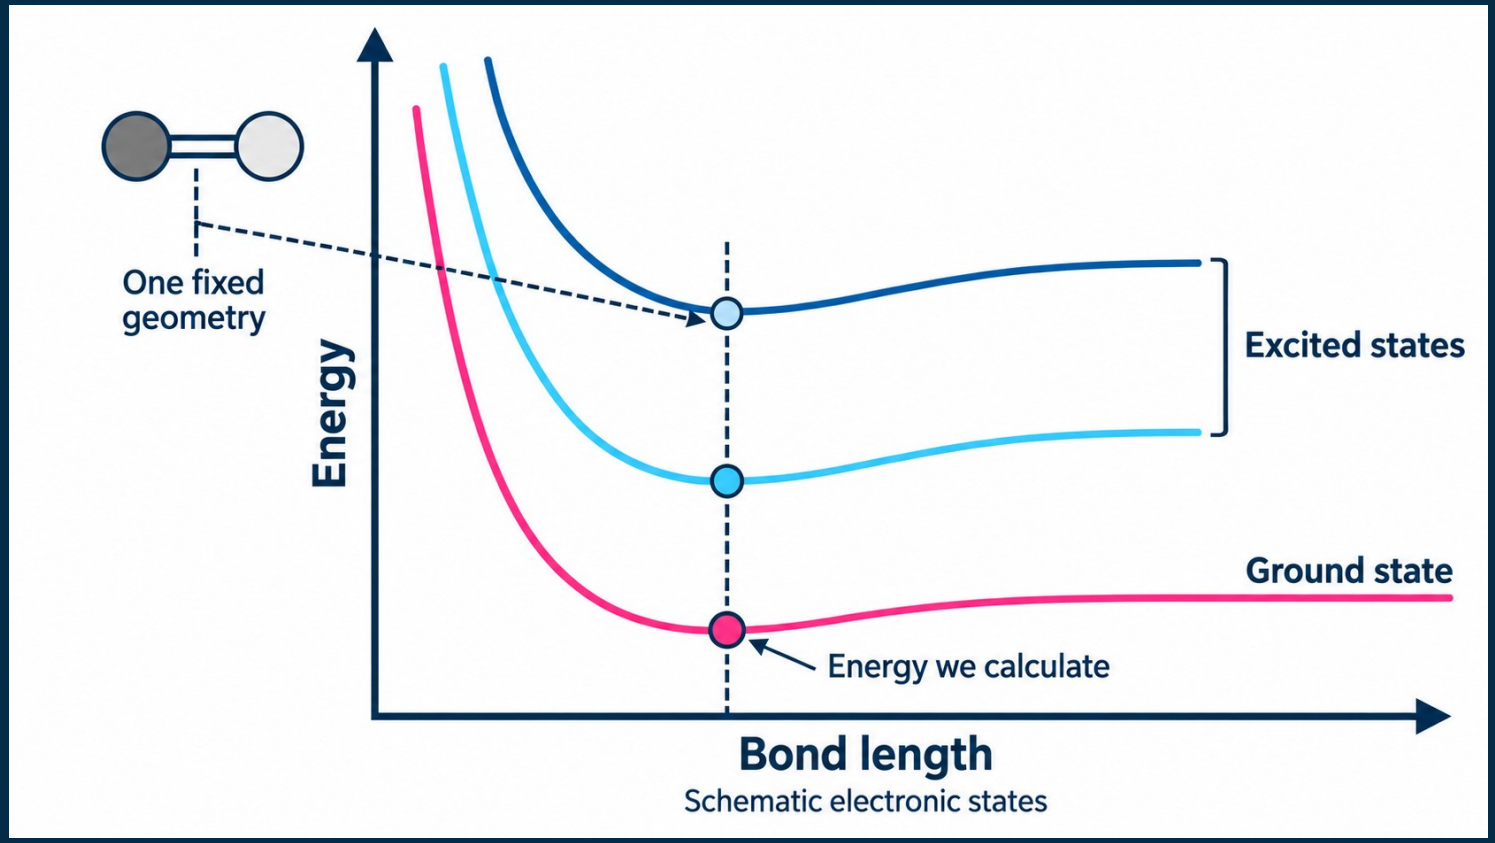
SC26 hpc unites | Chicago

Molecular Energy and the Ground State

One geometry, several electronic states

The lowest energy is the ground state.

Repeat the calculation at different bond lengths to trace the energy curve.

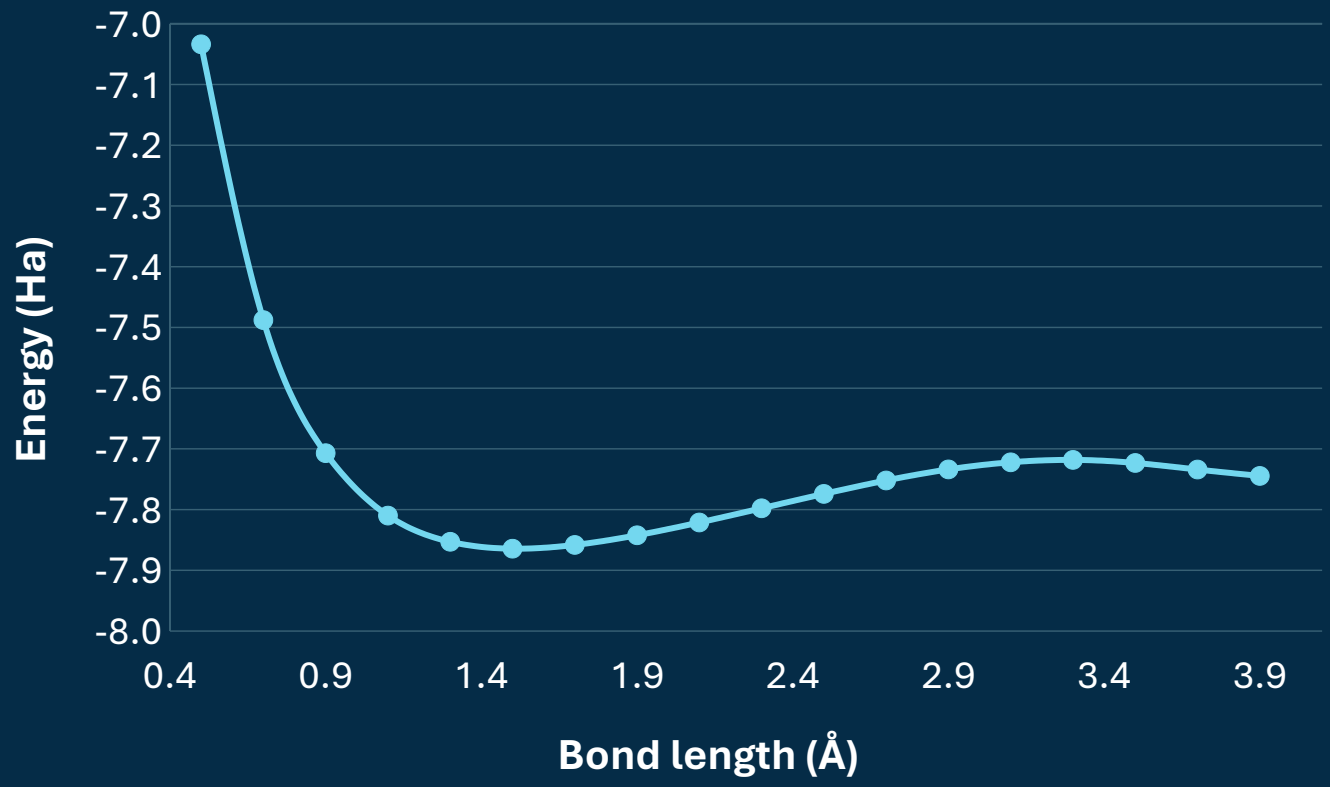


A PES Scan Repeats the Ground-State Search

One molecule. Many geometries.

At each geometry, find the lowest energy of its molecular Hamiltonian.

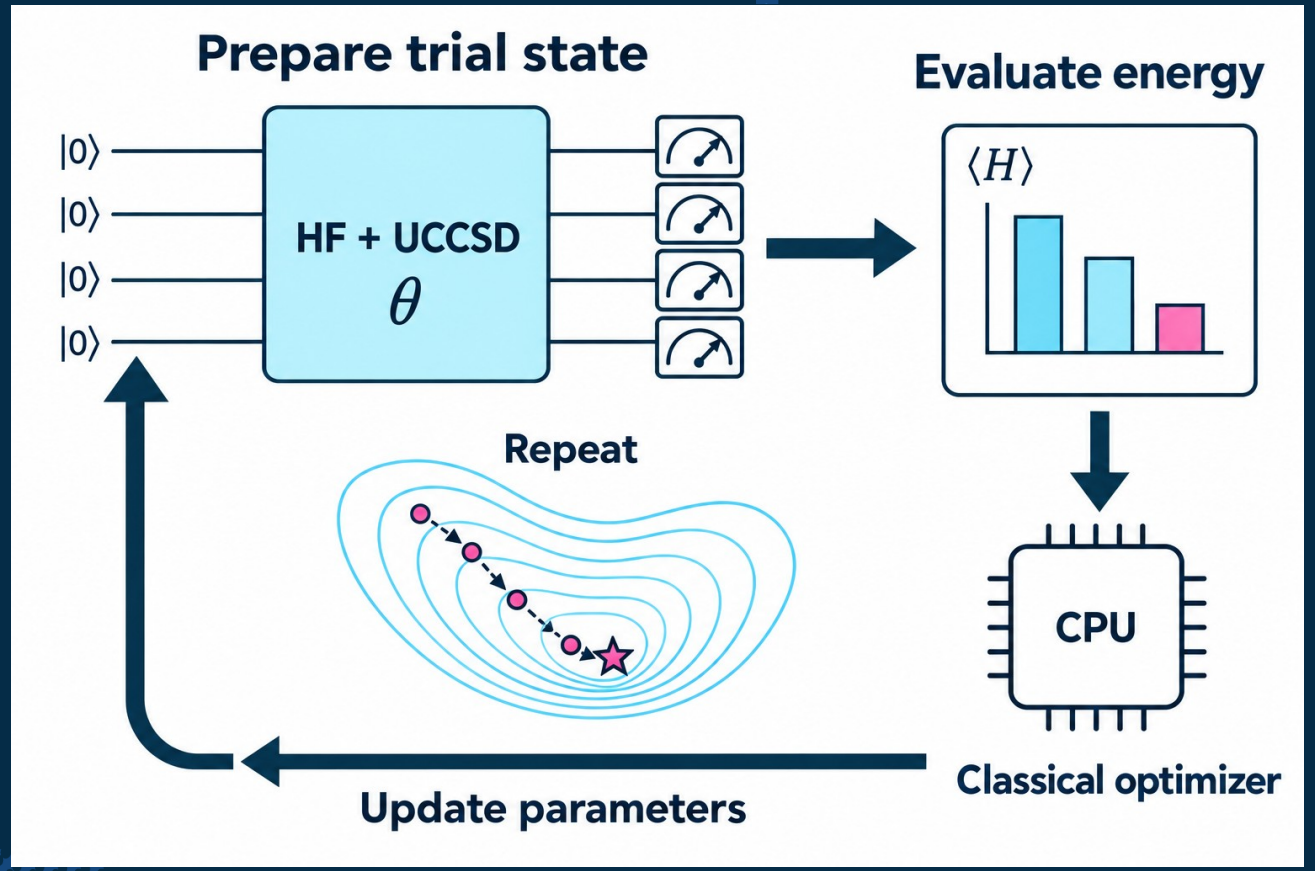
The resulting curve describes how energy changes as the atoms move.



LiH potential energy curve. Each point represents a molecular geometry.

SC26

The Variational Quantum Eigensolver (VQE)



$$H(R) = \sum_j c_j(R) P_j$$

Prepare a trial state

$$|\psi(\theta)\rangle = U(\theta)|\psi_{\mathrm{HF}}\rangle$$

Evaluate its energy

$$E(\theta; R) = \langle\psi(\theta)|H(R)|\psi(\theta)\rangle$$

Minimize over the parameters

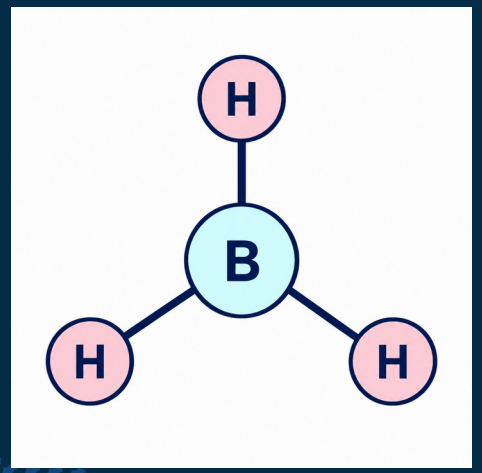
$$E_{\mathrm{VQE}}(R) = \min_{\theta} E(\theta; R)$$

SC26

Why a Molecular Scan Becomes Expensive

Geometries across the scan × **Optimizer steps** at each geometry × **Circuit evaluations** within each step

Borane (BH_3)



18 geometries
10 qubits
276 Pauli terms

1,597
Plain VQE

1,417
Parameter transfer

circuit evaluations across the scan

Shared molecular structure offers a way to reduce repeated work.

SC26

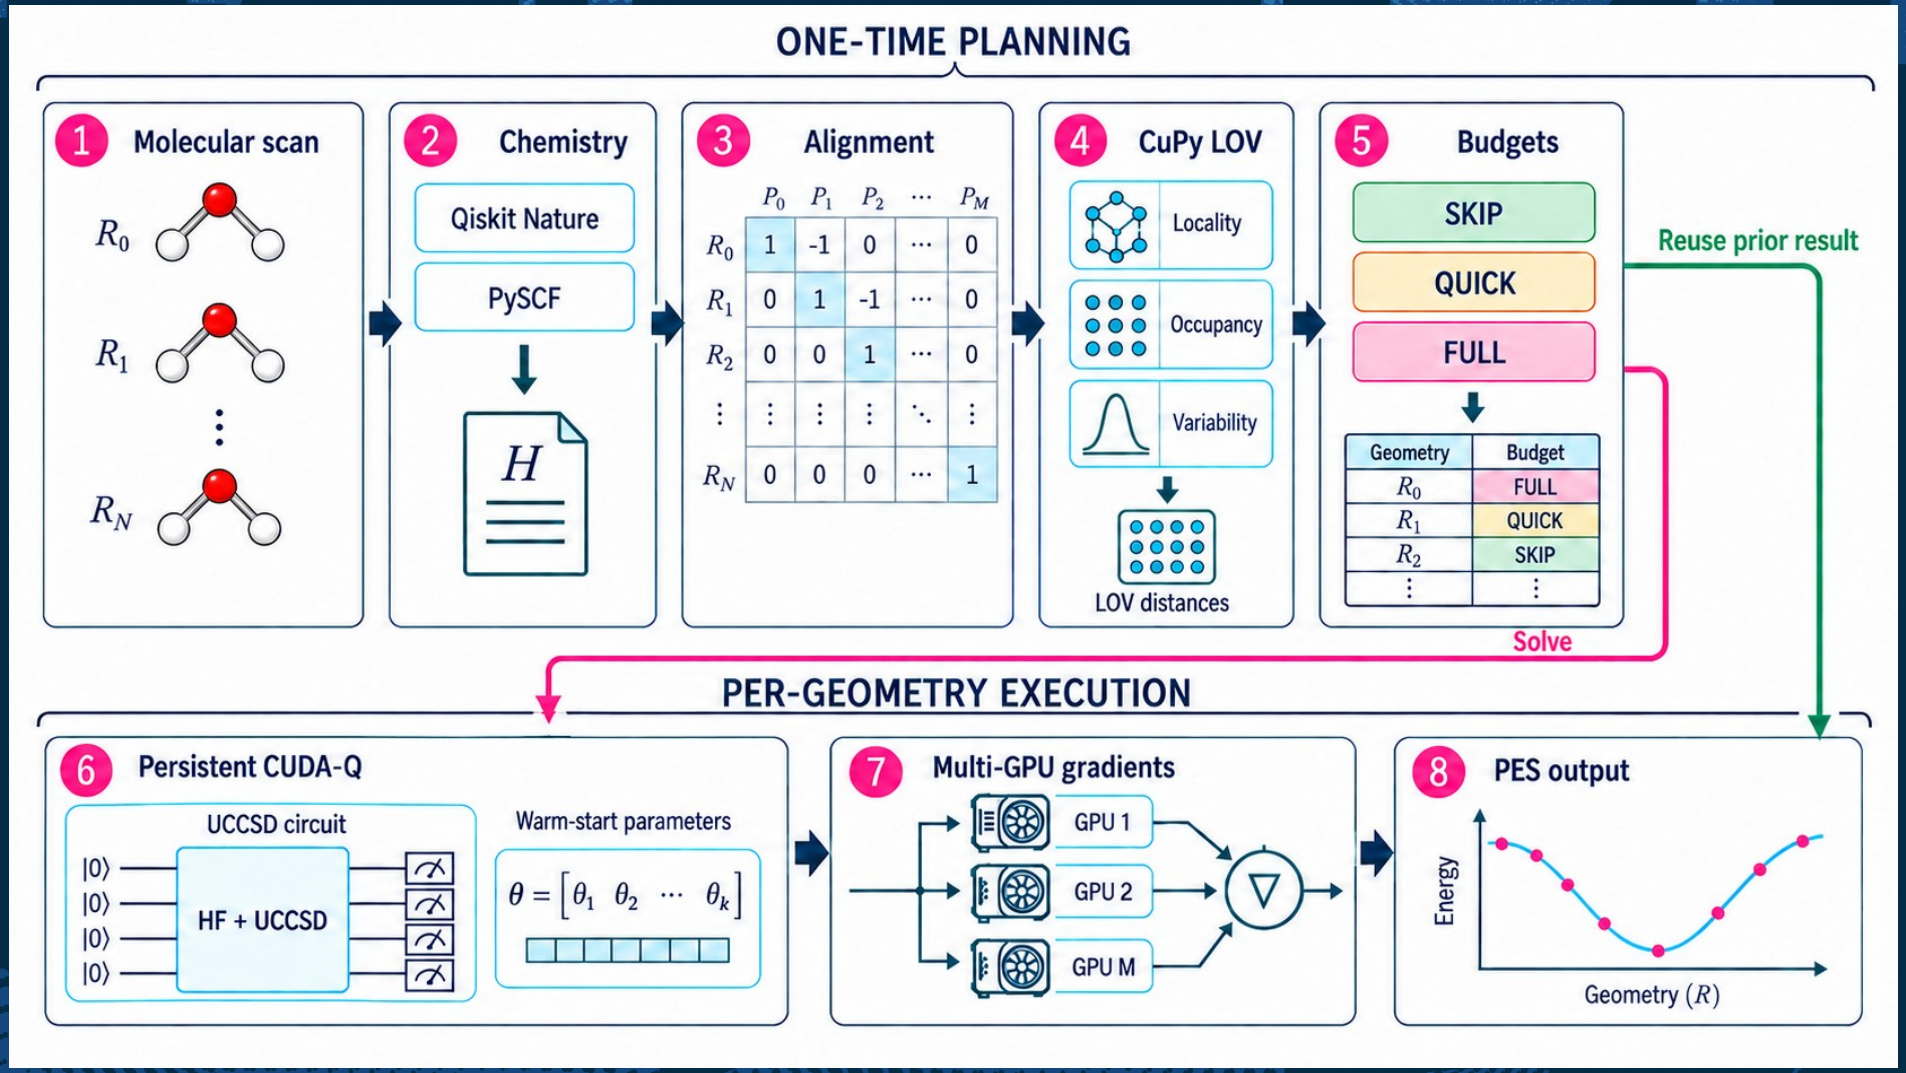
Baselines and Geometry-Aware Budgets

Neighboring Hamiltonians share operator structure, with changing coefficients.

Approach	Starting point	Computation budget
Plain VQE	Fresh optimization	Fixed caps
Parameter transfer (PT)	Previous parameters	Fixed caps
OptiPESS	Previous solution and state	LOV-guided budgets

OptiPESS uses cross-geometry information for both reuse and resource allocation.

The Complete OptiPESS System



SC26

Alignment Makes Hamiltonians Directly Comparable

Same columns, same Pauli operators

Geometry	ZI	IZ	XX	ZZ
R_0	0.8	0.2	**0**	0.5
R_1	0.7	0.2	0.1	0.5
R_2	0.6	0.3	0.1	0.4

Rows = geometries
Columns = Pauli strings
Cells = coefficients

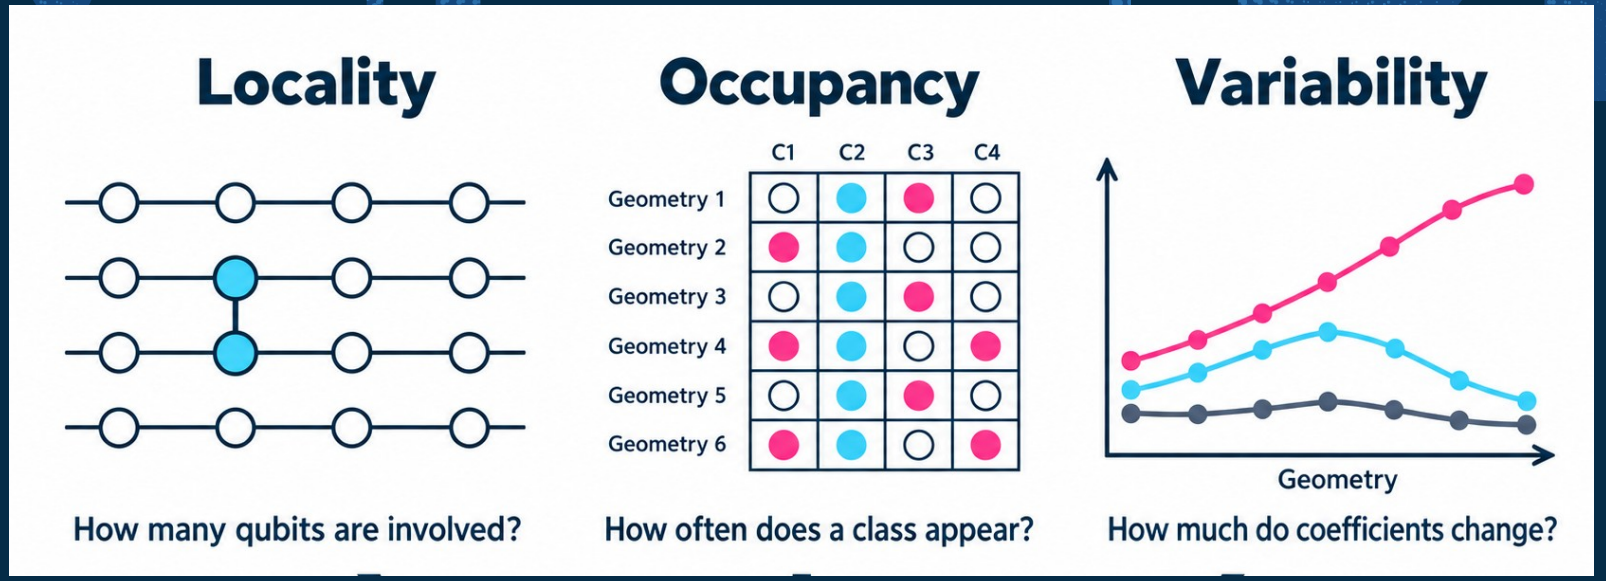
At R_0, XX is absent.
Its coefficient is 0.

The XX column stays aligned at all geometries.

Illustrative coefficients. Alignment enables vectorized analysis across the scan.

SC26

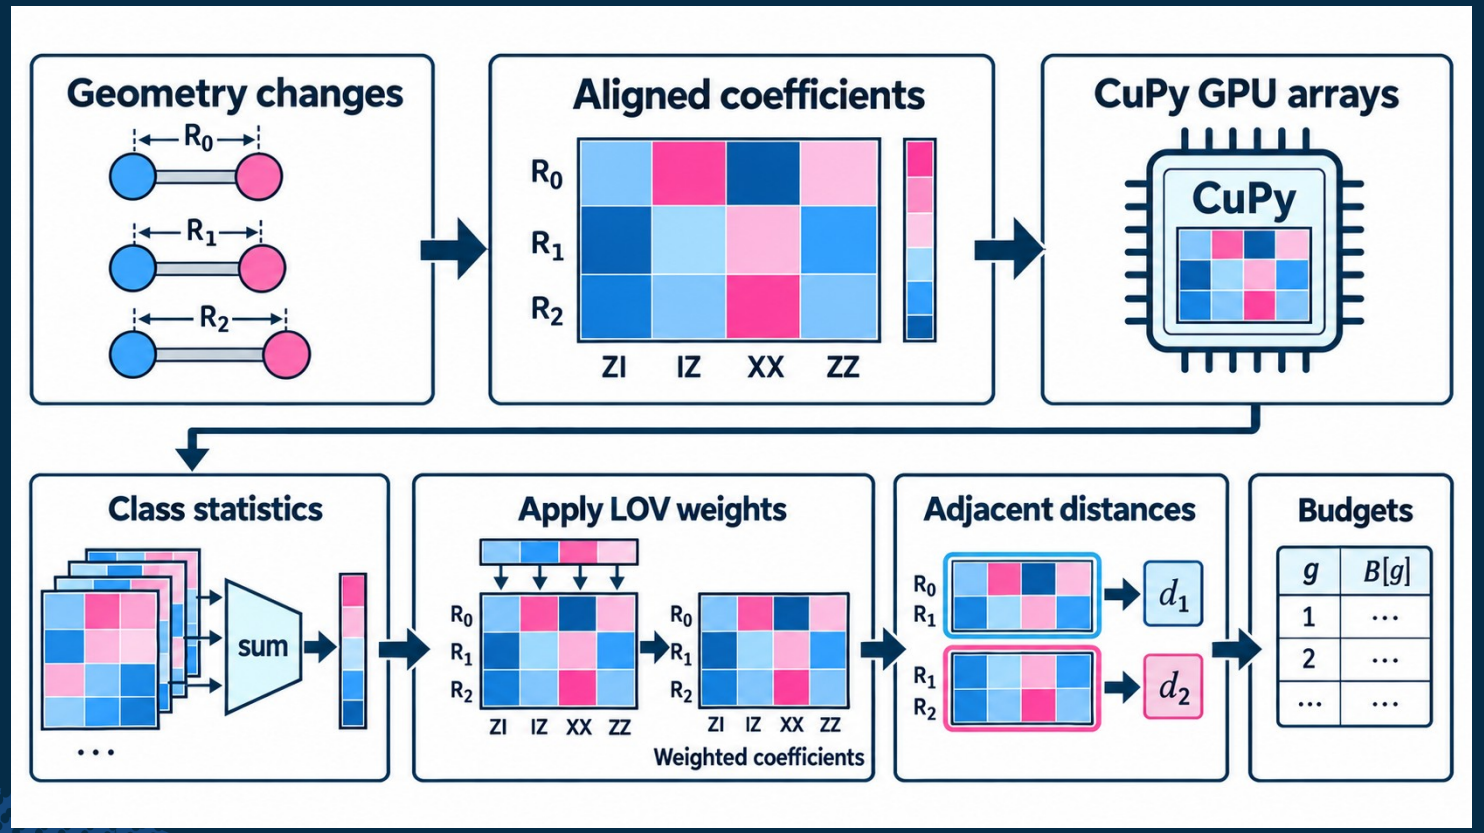
The Three Signals Behind LOV



Signal	What it measures	Role in the comparison
Locality	Number of active qubits	Emphasizes low-body structure
Occupancy	Frequency across the scan	Captures recurring structure
Variability	Coefficient changes across the scan	Captures sensitivity to geometry

Class weight $w = (LOV)^{1/3}$

CuPy Computes Geometry Distances on the GPU



GPU coefficient arrays

Same Pauli columns at every geometry.

LOV class weights

CuPy reductions calculate the class statistics.

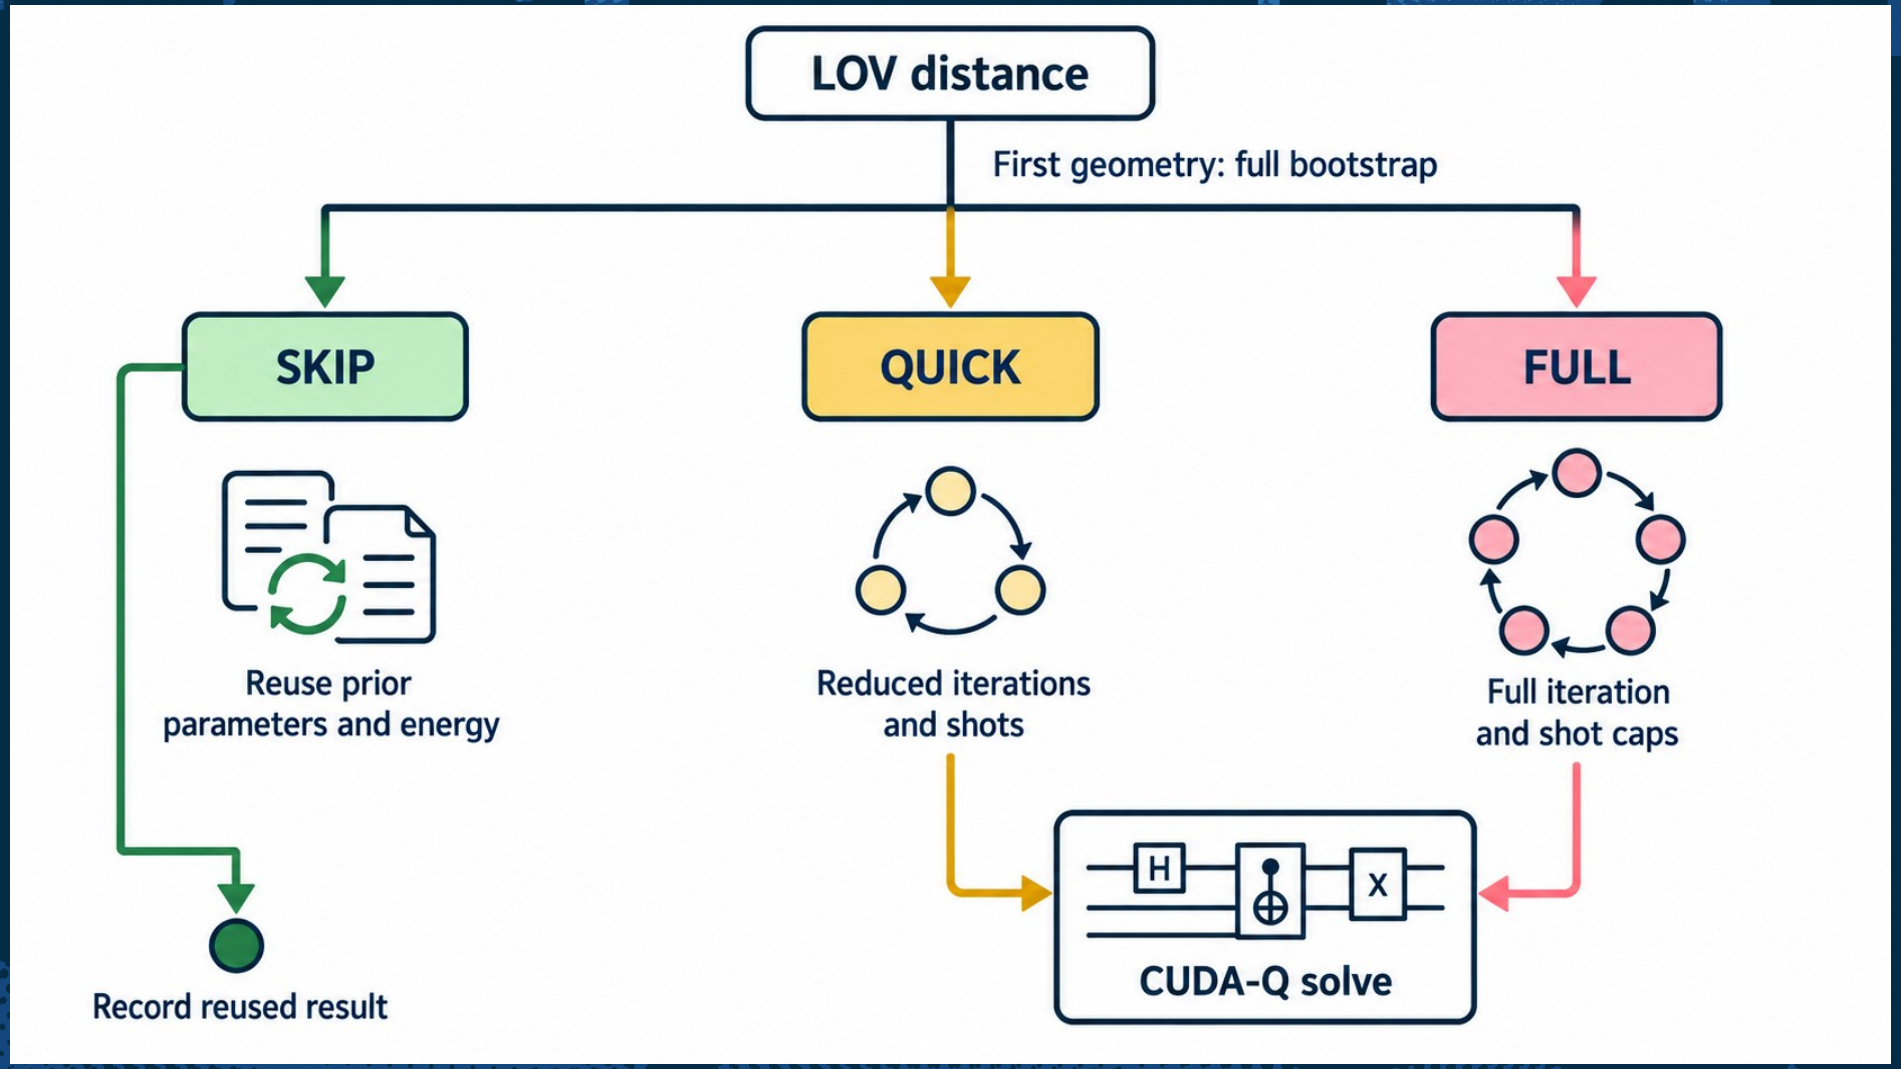
Adjacent distances

Compare weighted rows to guide the budgets.

CuPy arrays and reductions perform the analysis before VQE execution.

SC26

The Scheduler Assigns Skip, Quick, or Full Work



How OptiPESS Calculates the Budgets

Full budget caps

Maximum optimizer iterations and shots configured for the solve.

Scan-derived Quick fraction

LOV variation and energy tolerances set reduced caps for the scan.

Geometry-specific adjustment

The relative LOV distance scales each geometry's Quick allocation.

Distance factor

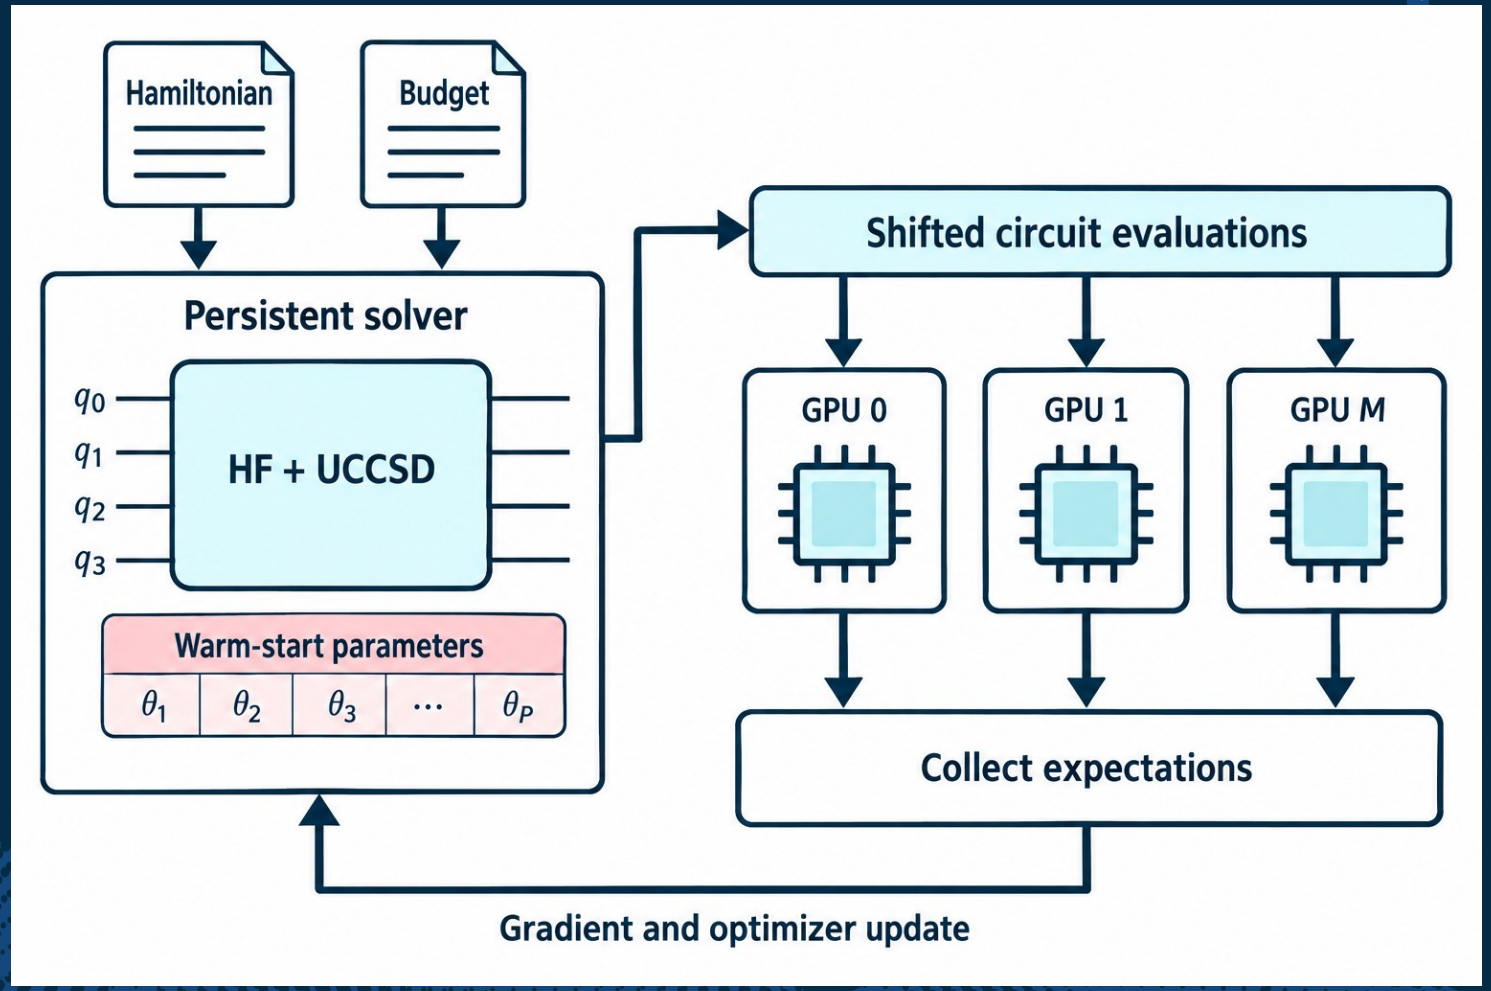
$$\alpha_g = 1 - \frac{D_g}{D_{\max}}$$

Geometry distance relative to the scan-wide LOV variation.

Iteration and shot budgets

Stored for each geometry before circuit execution.

CUDA-Q Executes the Geometry-Aware Schedule



cudaq.kernels.uccsd

Prepare the trial state

cudaq.observe

Evaluate its energy

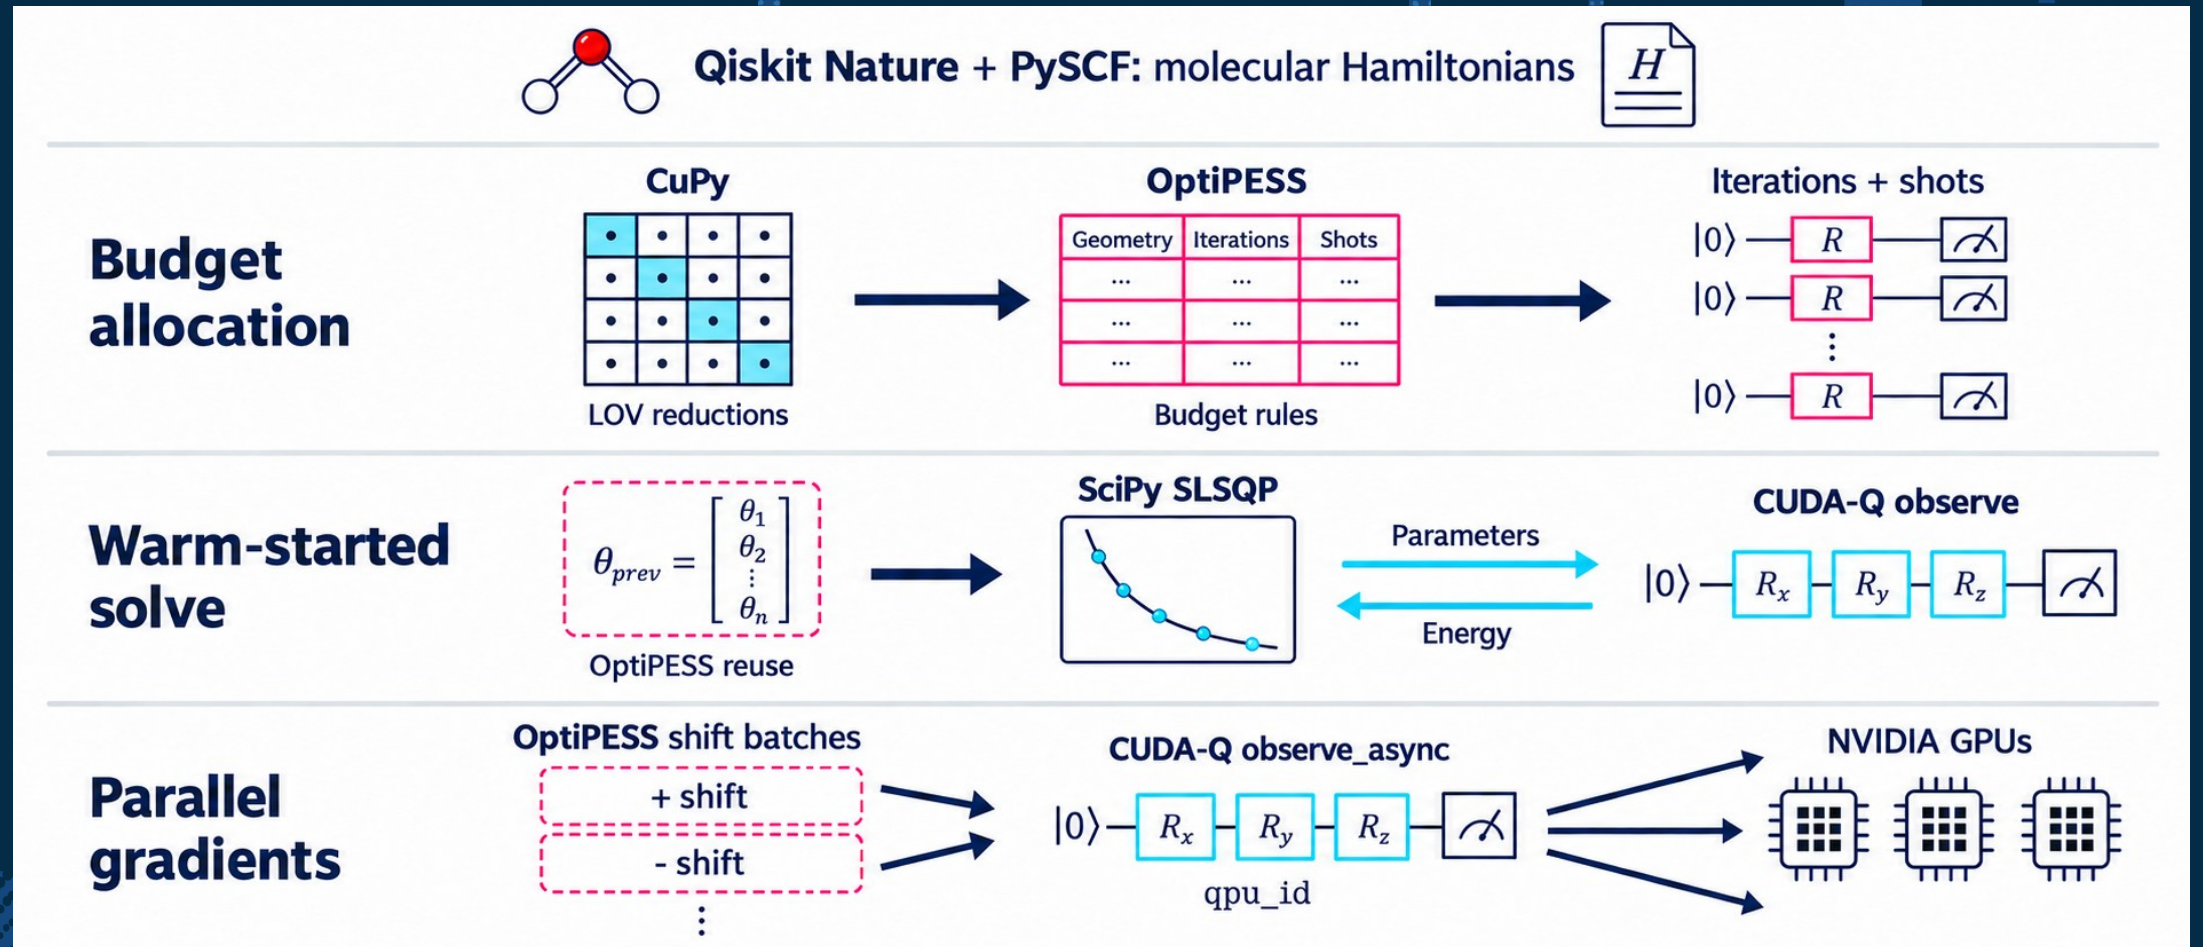
cudaq.observe_async

Dispatch shifted circuits with qpu_id

future.get()

Collect the gradient inputs

Building OptiPESS with Qiskit, CuPy, and CUDA-Q



GPU analysis, classical optimization, and circuit execution compose the OptiPESS runtime.

SC26

Evaluation Across 18 Molecular PES Workloads

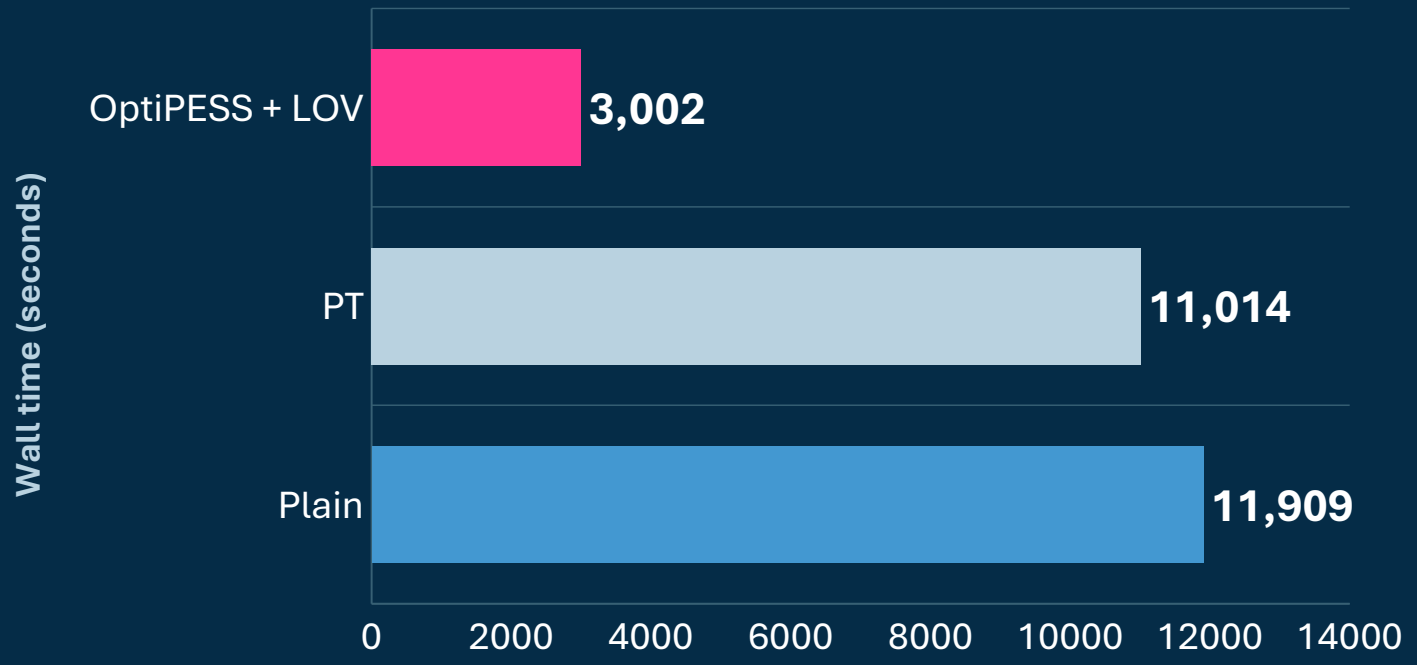
18	324	6–14	52–1,619
molecules	geometries	qubits	Pauli terms

Execution platform	Shared configuration
8 NVIDIA A100 GPUs 32 Intel Xeon Platinum 8358 CPUs	UCCSD ansatz and SLSQP optimizer
CUDA-Q 0.11, NVIDIA mqpu target CuPy similarity analysis	50-iteration cap 1,024-shot cap for noisy runs

Comparisons: Plain VQE, parameter transfer (PT), and OptiPESS. CUDA-Q simulations.

Lower End-to-End Wall Time

Noiseless simulation, reported similarity-measure comparison



74.8%

less time than Plain

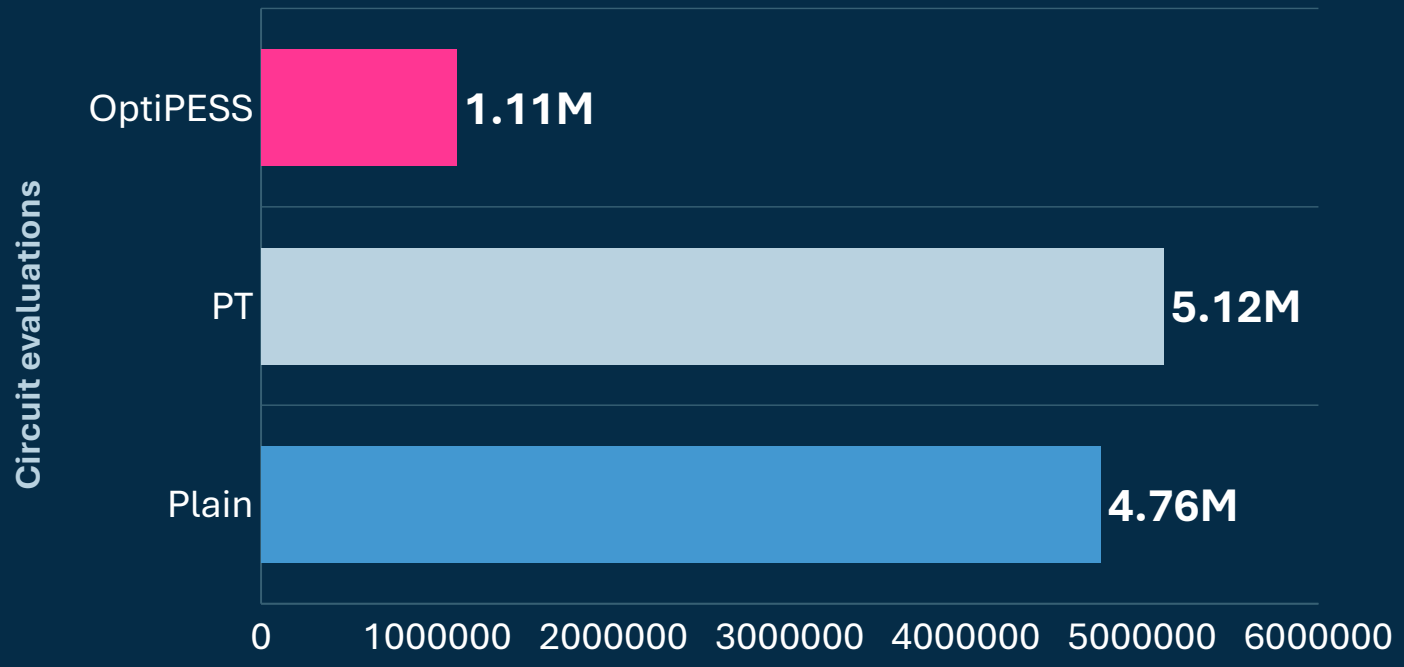
72.7%

less time than PT

Shared ansatz and optimizer. Bar labels show wall time in seconds.

Fewer Circuit Evaluations with Adaptive Budgets

Noisy simulation with adaptive quick-band iterations



76.7%

fewer than Plain

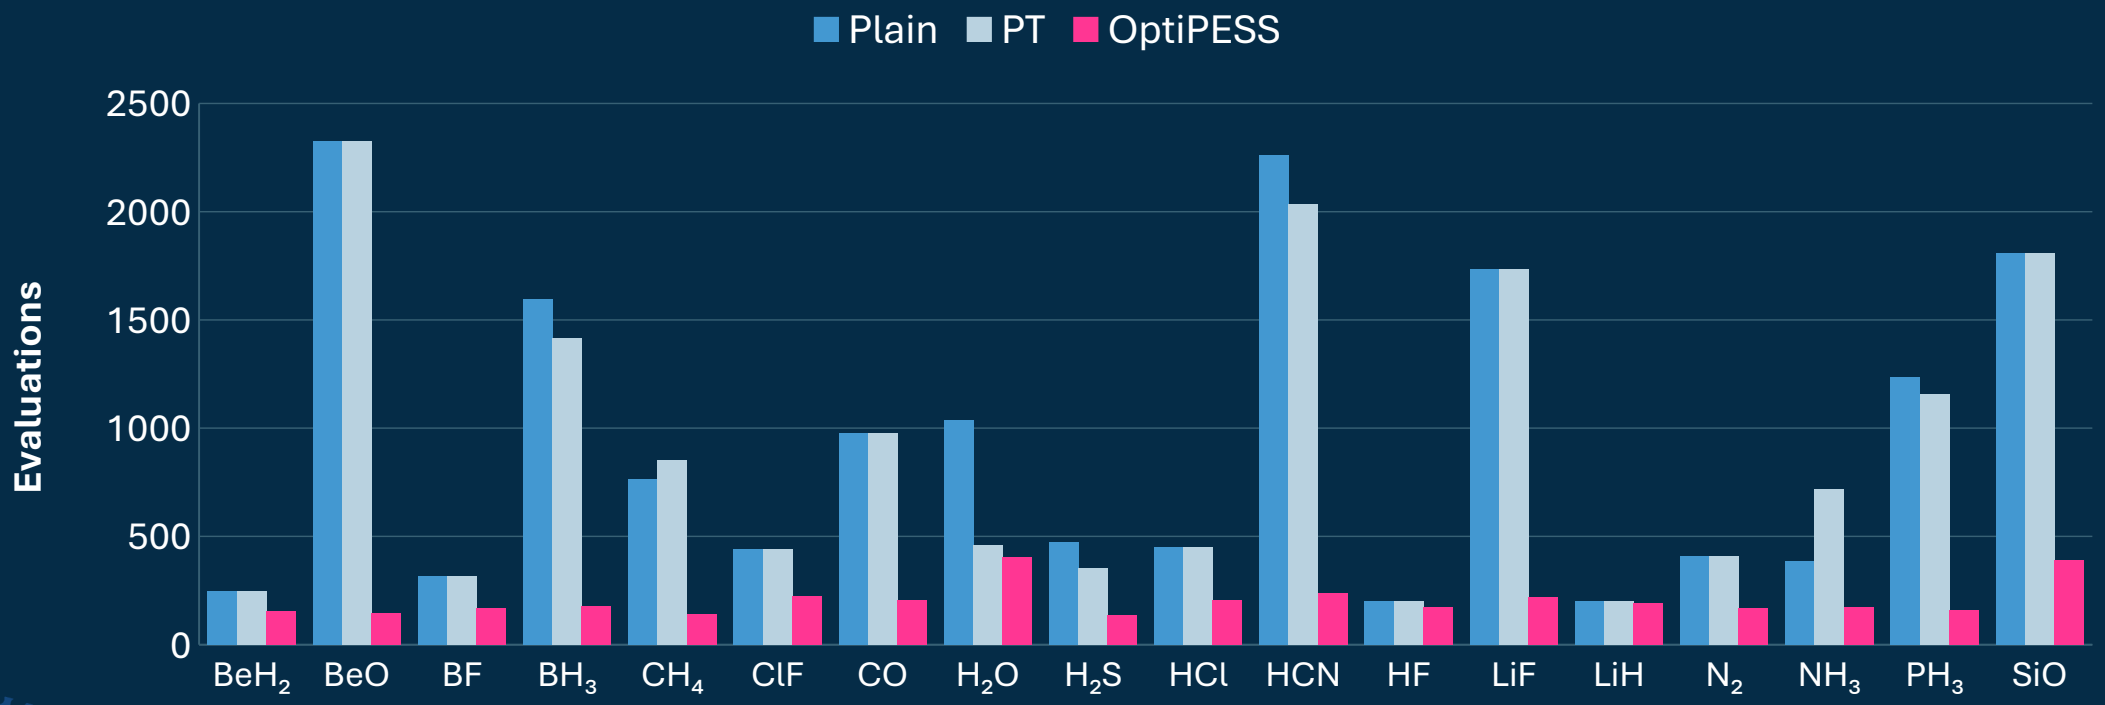
78.3%

fewer than PT

Noisy-run totals across all molecules; adaptive iteration budgets in the quick band.

Fewer Evaluations Across All 18 Molecules

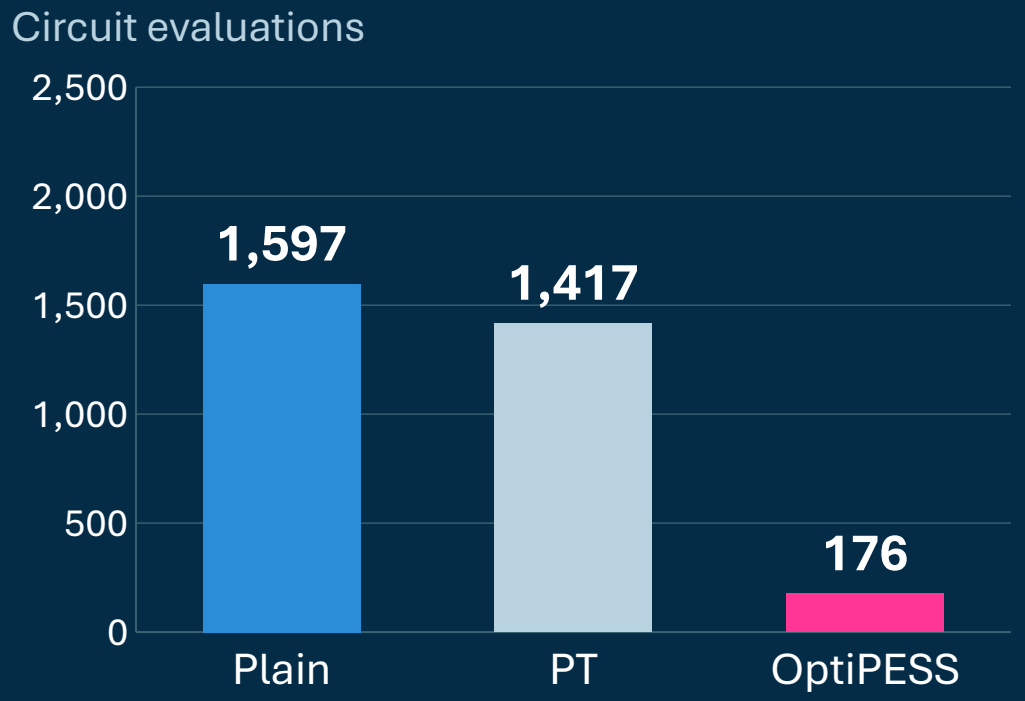
Per-molecule circuit evaluations in the paper’s noiseless comparison



OptiPESS reduces circuit evaluations across all 18 benchmark molecules.

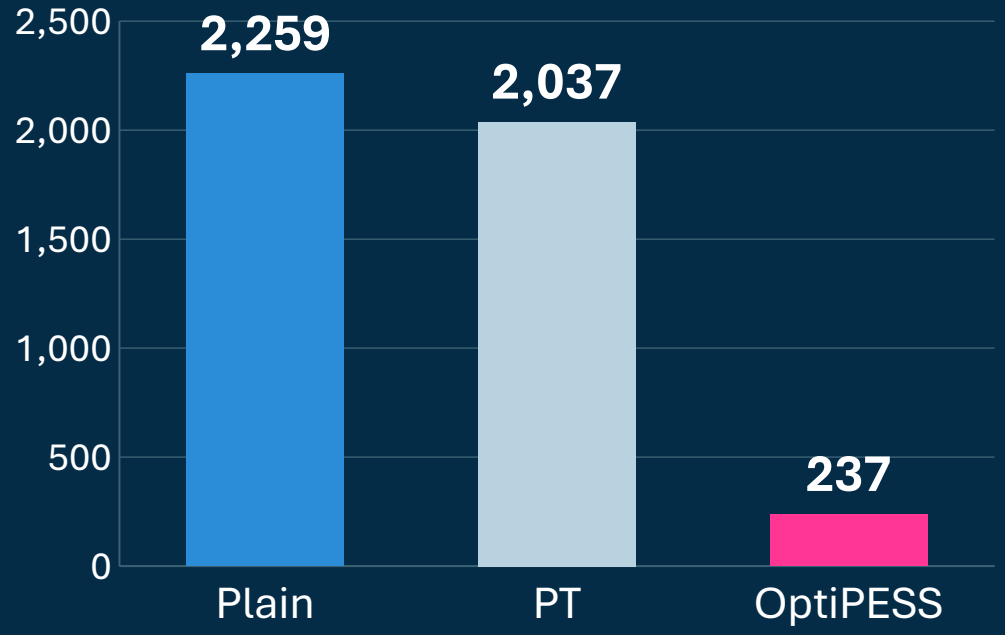
Large Reductions for BH_3 and HCN

Borane (BH_3)



89.0% fewer than Plain
87.6% fewer than PT

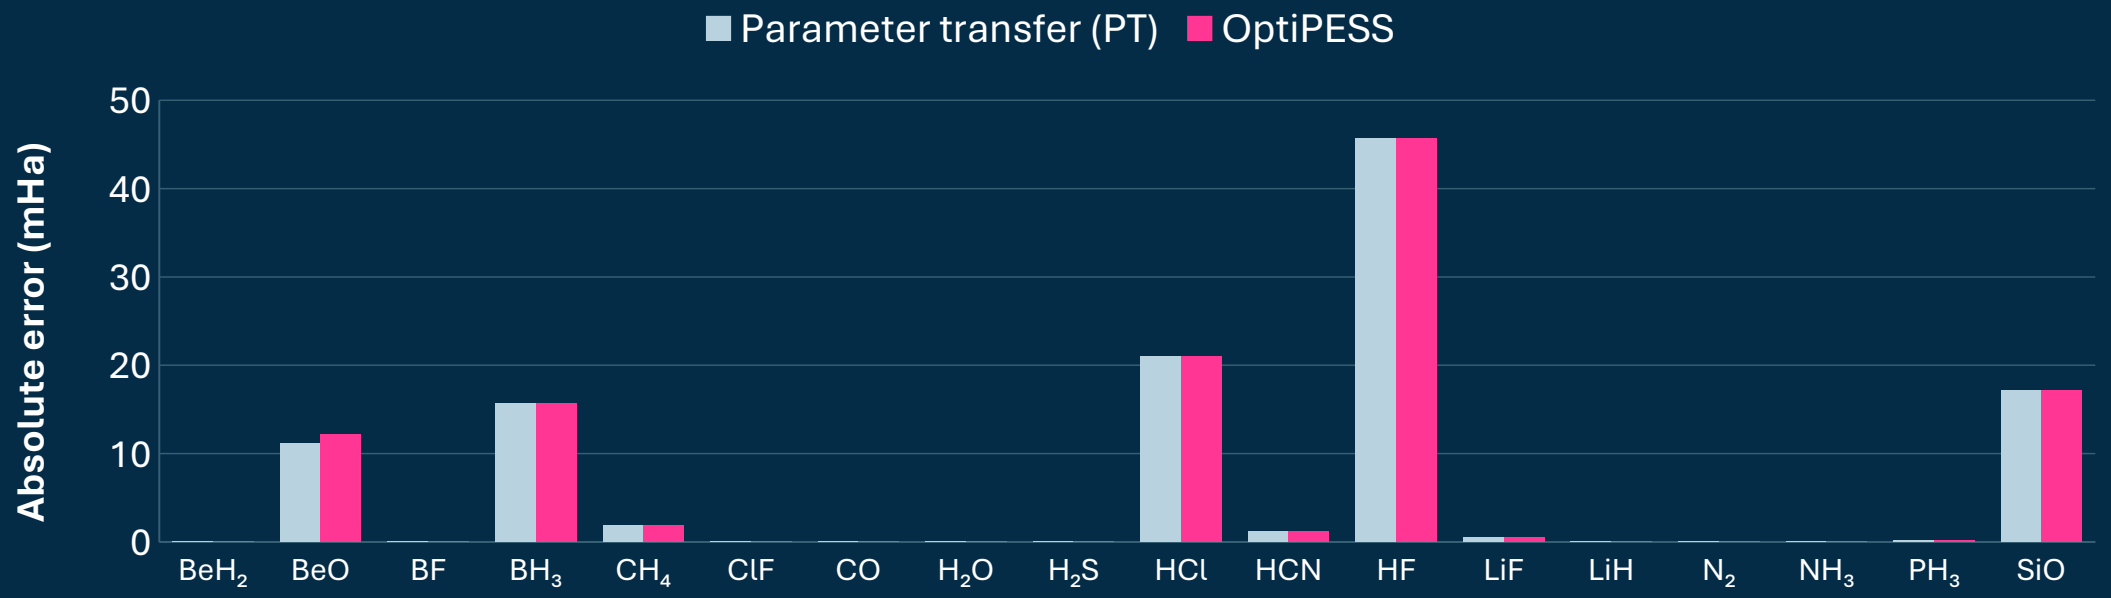
Hydrogen cyanide (HCN)



89.5% fewer than Plain
88.4% fewer than PT

OptiPESS Tracks Baseline Accuracy Across Molecules

Average absolute error at each molecule's three lowest-energy geometries



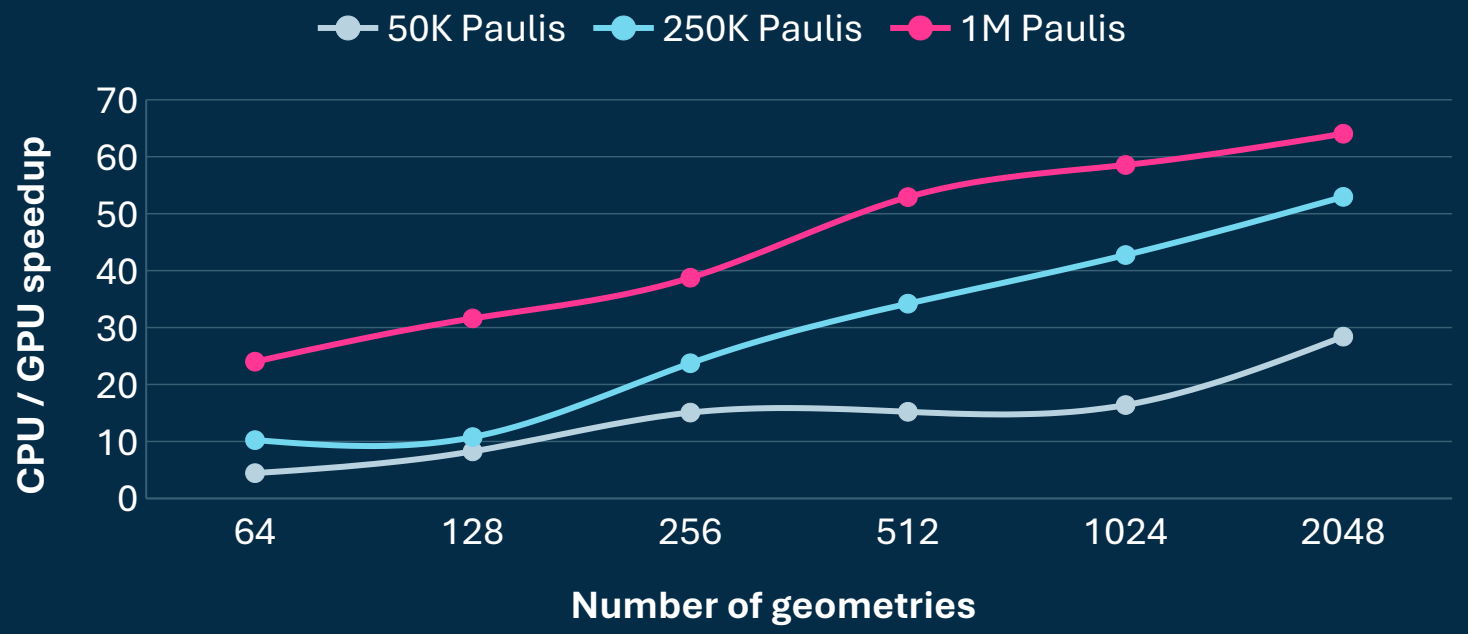
Mean energy difference in the noiseless comparison

0.0218% vs. Plain

0.0143% vs. PT

CuPy Analysis Scales to One Million Pauli Terms

LOV preprocessing benchmark across geometries and Pauli strings



64×

CPU / GPU speedup
at the largest size

<2%

of end-to-end time
for static analysis

Synthetic scaling: up to 2,048 geometries × 1M Paulis. Overhead: molecular workload.

OptiPESS Makes PES Execution Geometry-Aware

Hamiltonian structure guides computation across the entire scan.

CuPy computes LOV similarity to guide the budget.
Persistent CUDA-Q execution reuses work.
Multiple GPUs accelerate gradient evaluation.

74.8% less wall time

than Plain in the noiseless comparison

78.3% fewer evaluations

than PT in the noisy comparison

Thank You

Questions & Discussion

OptiPESS

Panayiotis Christou, Shuwen Kan, Ying Mao

Fordham University

Backup: LOV Definitions and Scheduling Thresholds

Class statistics

$$L_c = 2^{-k_c}$$

$$O_c = \frac{n_c}{\max_d n_d}$$

$$V_c = \frac{\sigma_c}{\max_d \sigma_d}$$

Weight and adjacent distance

$$w_p = (L_{c_p} O_{c_p} V_{c_p})^{1/3}$$

$$\mathbf{X} = \mathbf{C}W$$

$$D_{\mathrm{LOV}}(g) = \|x_g - x_{g-1}\|_2$$

Band	Distance rule	Budget
Skip	$D \le \eta_{\mathrm{skip}} D_{\max}$	$I_g = S_g = 0$
Quick	$\eta_{\mathrm{skip}} D_{\max} < D \le \eta_{\mathrm{quick}} D_{\max}$	Reduced I_g and S_g
Full	$D > \eta_{\mathrm{quick}} D_{\max}$	$I_{\max}$ and $S_{\max}$

SC26

Backup: Comparing Similarity Measures

Noiseless CUDA-Q comparison from the paper

Method	Evaluations	Time (s)	Energy % diff PT / Plain
Plain	8,674	11,908.82	—
PT	9,976	11,013.74	—
OptiPESS + Frobenius	3,092	3,657.43	−0.0338 / −0.0412
OptiPESS + Hilbert–Schmidt	5,079	6,477.84	+0.0671 / +0.0596
OptiPESS + LOV	2,371	3,002.23	−0.0143 / −0.0218

LOV achieves the lowest runtime and evaluation count in this comparison.

Backup: Results with Other Optimizers

Paper evaluation on the subset with fewer than 500 Pauli terms

Optimizer	Mode	Iterations	Time (s)
L-BFGS-B	Plain	2,344	38,999.93
L-BFGS-B	PT	1,562	29,597.45
L-BFGS-B	OptiPESS	522	9,021.30
SPSA	Plain	64,800	6,700.16
SPSA	PT	34,200	3,557.63
SPSA	OptiPESS	29,656	3,193.99

Adaptive scheduling also reduces work with L-BFGS-B and SPSA.