

Towards Chemically Accurate and Scalable Quantum Simulations on IQM Quantum Hardware: A Quantum-HPC Hybrid Approach

Anurag K. S. V.^{†,1}, Ashish Kumar Patra^{†,1}, Manas Mukherjee², Alok Shukla³, Sai Shankar P.¹, Ruchika Bhat^{1,4}, Radhika T. S. L.⁵, and Jaiganesh G.^{*,1}

¹Qclairvoyance Quantum Labs, Secunderabad, TG 500094, India.

²Centre for Quantum Technologies, National University of Singapore, Singapore 117543, Singapore.

³School of Arts and Sciences, Ahmedabad University, Ahmedabad, GJ 380009, India.

⁴The University of Arizona, Tucson, AZ 85721, USA.

⁵Department of Mathematics, BITS Pilani, Hyderabad Campus, Hyderabad, TG 500078, India.

[†]These authors contributed equally to this work

ABSTRACT

We report one of the most comprehensive experimental demonstrations to date of quantum-computing-based molecular simulation on IQM's Sirius 24-qubit superconducting quantum processor, using up to 16 operational qubits. We employ Sample-based Quantum Diagonalization (SQD) with the Local Unitary Cluster Jastrow (LUCJ) ansatz to compute ground-state energies of benchmark molecular systems, including H_2 , LiH, BeH_2 , H_2O , and NH_3 . We further introduce the Linear-CNOT Unitary Coupled-Cluster Singles and Doubles (LCNot-UCCSD) ansatz within the SQD framework, reducing the classical pre-computation overhead at the cost of increased quantum circuit depth, and demonstrate its effectiveness by accurately recovering the ground-state energies of H_2 , LiH, and BeH_2 . A comparative analysis of the two ansätze is presented, highlighting their respective merits, demerits, and practical use cases in near-term quantum simulations. In addition, we perform one-dimensional potential energy surface (1D-PES) scans for H_2 and HeH^+ using STO-3G and 6-31G basis sets, and for LiH and BeH_2 using STO-3G. We also report the first experimental generation of a full two-dimensional potential energy surface (2D-PES) map, consisting of a 32×32 grid in bond length and bond angle, for the water molecule in the STO-3G basis on IQM quantum hardware. Beyond small-molecule benchmarks, we integrate Density Matrix Embedding Theory (DMET) with SQD(LUCJ) to obtain chemically accurate 4-electron-in-4-orbital active-space energies for eight ligand-like molecules, as well as for the pharmacologically relevant amantadine molecular system, representing the first of such experimental demonstrations. Across all studies, the majority of quantum-computed energies agree with reference full configuration interaction (FCI) results and with DMET-CASCI energies for embedded systems, to within chemical accuracy for the chosen basis sets. These results demonstrate the robustness of quantum sampling-based classical diagonalization approaches while highlighting the promise of hybrid embedding strategies combined with near-term ansätze for advancing quantum chemistry toward classically intractable molecular systems and establishing their practicality on current IQM quantum hardware.

Keywords: Quantum Computing · Quantum Simulation · Quantum Chemistry · IQM Quantum Hardware · IQM Sirius · Sample-based Quantum Algorithms · Embedding Methods · Potential Energy Surface Scans

^{*}(Corresponding Author) email: jaiganesh@qclairvoyance.in, drjaiganesh15@gmail.com

1 Introduction

One of the quintessential objectives of computational quantum chemistry is the rigorous description of the electronic structure of matter^{1–4}. This endeavor is the bedrock upon which modern computational materials science^{5–10}, pharmacology^{11–16}, and chemical engineering rest^{17–20}. The ability to predict the behaviour of electrons, how they distribute themselves around nuclei, how they form bonds, and how they respond to external stimuli dictates our capacity to design novel catalysts for carbon capture^{21–25}, synthesize life-saving therapeutics^{11,13,16,26,27}, and engineer high-temperature superconductors^{28–30}. At the theoretical core of this pursuit lies the non-relativistic time-independent Schrödinger equation³¹. In principle, the exact solution to this equation for a given electronic Hamiltonian of N electrons and M nuclei yields the ground-state energy of the provided Hamiltonian with a precision limited only by relativistic effects^{32–36} and the breakdown of the Born-Oppenheimer approximation^{37–40}.

However, the practical realization of this predictive power is hindered by a formidable computational barrier known as the curse of dimensionality^{41,42}. The Hilbert space in which the electronic wavefunction resides grows exponentially with the number of active electrons and orbitals². In the Full Configuration Interaction (FCI) formalism, the exact wavefunction is expressed as a linear combination of all Slater determinants consistent with particle number and spin symmetries¹. For a molecular system described by N_{orb} spatial orbitals and N_{ele} electrons (where N_{α} and N_{β} denote the number of electrons with α -spin and β -spin, respectively), the dimension of this space grows combinatorially. This scaling follows $\binom{2N_{\text{orb}}}{N_{\text{ele}}}$ for general open-shell systems and $\binom{N_{\text{orb}}}{N_{\alpha}}\binom{N_{\text{orb}}}{N_{\beta}}$ for closed-shell systems, rapidly exceeding classical memory and compute limits^{2,43}. Historically, this scaling confined exact diagonalization methods to systems containing roughly 10–14 correlated electrons, with dimension sizes just exceeding one billion (10^9) determinants^{44,45}. Recent algorithmic breakthroughs in distributed computing and tensor compression have significantly expanded this horizon. Notably, massively parallel implementations have enabled CASCI calculations on active spaces as large as 24 electrons in 24 orbitals; while formally spanning 7.31×10^{12} determinants, the exploitation of D_{2h} spatial symmetry and the singlet 1A_g ground state reduced this complexity to 9.1×10^{11} determinants⁴⁶. Furthermore, exact FCI solutions have been achieved for the propane molecule (C_3H_8) involving over 1.31×10^{12} determinants⁴⁷, effectively breaching the one-trillion-determinant barrier. Most recently, lossless categorical compression strategies have pushed the frontier to the quadrillion-determinant scale (10^{15}) for heavy-element systems such as HBrTe ⁴⁸. Yet, despite these heroic classical efforts, the non-polynomial scaling persists; the Hilbert space for a moderate-sized drug molecule or a complex catalytic center remains orders of magnitude beyond even the exascale regime, necessitating a fundamentally different computational paradigm.

To circumvent this barrier, computational chemistry has long relied on a hierarchy of systematically improvable approximations. Mean-field theories such as Hartree-Fock (HF)^{49–52} and Kohn–Sham Density Functional Theory (DFT)^{6,53,54} reduce the many-body problem to an effective one-particle description, capturing the dominant portion of the total energy at modest computational cost^{55,56}. However, the residual electron correlation energy governs chemically decisive phenomena, including reaction barriers^{57–59}, weak non-covalent interactions^{60–62}, and excited-state processes^{63–65}. Post-Hartree–Fock methods, most notably Coupled Cluster with Single and Double excitations augmented by perturbative Triples [CCSD(T)]⁶⁶, recover a large fraction of this correlation energy with polynomial scaling and have therefore earned the designation of the “gold standard” of quantum chemistry^{67,68}. Nonetheless, the steep $\mathcal{O}(N^7)$ scaling of CCSD(T), together with large prefactors, severely restricts its applicability to medium-sized molecules and precludes routine treatment of extended systems or large biomolecules without further fragmentation^{69–71} or embedding approximations^{72–74}.

The advent of quantum computing offers a fundamentally different approach to this long-standing challenge^{4,75–77}. By exploiting quantum superposition⁷⁸ and entanglement^{79,80}, quantum processors can, in principle, represent and manipulate many-body wavefunctions in a Hilbert space whose dimension grows exponentially with the number of qubits, while the physical resources (qubits) scale polynomially^{81,82}. This observation underpins the long-standing notion that quantum computers are natural simulators of quantum systems, capable of addressing regimes that are classically intractable^{75,83}. Early algorithmic proposals for quantum chemistry leveraged this insight to formulate polynomial-scaling quantum algorithms for eigenvalue estimation, providing the first concrete evidence of a potential exponential advantage over classical exact methods^{82,84}.

A paradigmatic example is Quantum Phase Estimation (QPE), which enables the extraction of eigenvalues of a unitary operator corresponding to the time evolution under a molecular Hamiltonian^{82,85}. When supplied with a sufficiently accurate initial state and implemented using fault-tolerant quantum error correction (QEC), QPE can determine ground- and excited-state energies with asymptotically optimal scaling in the desired precision⁸⁶. Subsequent developments, such as Linear Combinations of Unitaries (LCU)⁸⁷ and utilizing randomized compiling⁸⁸, have significantly reduced asymptotic gate counts. Moreover, qubitization-based Hamiltonian simulation⁸⁹, underpinned by the Quantum Signal Processing (QSP) methodology⁹⁰, has been unified under the Quantum Singular Value Transformation (QSVT) framework⁹¹. These advances have been further accelerated by efficient block-encoding strategies for structured matrices^{92–94} and modern compact Hamiltonian representations, such as Tensor Hypercontraction (THC) and regularized Double Factorization (DF)^{95–97}. By minimizing the block-encoding

overheads, these techniques have lowered the theoretical resource estimates for simulating complex bio-catalysts like FeMoco by orders of magnitude^{97–100}. While the precise boundary of exponential advantage for ground-state chemistry remains under complexity-theoretic scrutiny^{101–103}, the fundamental capacity of quantum processors to exceed classical limits has been confirmed through demonstrations of unconditional quantum information supremacy¹⁰⁴. However, the substantial depth and qubit requirements associated with standard QPE protocols firmly place them in the domain of fault-tolerant, application-scale quantum (FASQ) computing¹⁰⁵. To bridge this gap, recent efforts have pivoted toward early-fault-tolerant architectures, envisioned to operate with 25–100 logical qubits¹⁰⁶, employing novel dissipative state preparation protocols¹⁰⁷ and partially fault-tolerant schemes¹⁰⁸ to bring practical quantum advantage closer to experimental reality.

In parallel with the pursuit of fault tolerance, the Noisy Intermediate-Scale Quantum (NISQ)¹⁰⁹ era has been defined by hybrid quantum–classical algorithms designed to operate within limited coherence times. The Variational Quantum Eigensolver (VQE) has served as the dominant paradigm^{110–112}, evolving from static ansätze^{113,114} to sophisticated adaptive strategies to enhance expressivity. Notable advancements include adaptive derivative-assembled pseudo-trotter VQE (ADAPT-VQE) variants utilizing coupled exchange operators¹¹⁵, parameter-based selection criteria¹¹⁶, and automated operator pruning¹¹⁷ to drastically reduce circuit depth. Concurrently, iterative downfolding protocols, such as coupled cluster downfolding¹¹⁸, have been developed to exchange information between quantum and classical processors dynamically, enabling a larger active space. Despite these innovations, VQE faces fundamental scaling challenges: optimization landscapes frequently exhibit noise-induced barren plateaus^{119,120}, and measurement overheads remain prohibitive even with Hamiltonian partitioning¹²¹. Critically, unmitigated device noise can invalidate the variational bound, complicating the retrieval of accurate molecular properties¹²² and dynamic Green’s functions¹²³, often preventing meaningful chemical insights for systems as fundamental as benzene on current hardware¹²⁴. These persistent scalability and accuracy constraints have spurred a resurgence of interest in alternative, sample-based paradigms¹²⁵ that circumvent the variational optimization loop entirely.

Formalizing this paradigm, sampling-based quantum algorithms such as Quantum-Selected Configuration Interaction (QSCI)¹²⁶ decouple wavefunction generation from energy optimization by utilizing the quantum device to identify high-weight determinants from a trial state, subsequently diagonalizing the Hamiltonian within this compressed subspace. To mitigate the burden of ansatz optimization, recent protocols have introduced adaptive input state construction (ADAPT-QSCI)¹²⁷ and Hamiltonian simulation-based sampling (HSB-QSCI)^{128,129}, which leverage time-evolution to access chemically relevant sectors of the Hilbert space without complex variational loops. This framework has culminated in the concept of quantum-centric supercomputing, demonstrated by large-scale Sample-based Quantum Diagonalization (SQD) simulations of N₂ dissociation and impurity models on IBM Heron quantum hardware coupled with classical supercomputers¹³⁰, and most recently by closed-loop workflows integrating 152,064 nodes of the Fugaku supercomputer to tackle electronic structures beyond exact diagonalization¹³¹. The applicability of these sampling-based algorithms has been rapidly expanded to encompass excited states¹³², open-shell systems¹³³, implicit solvation models¹³⁴, and non-covalent interactions¹³⁵. Furthermore, to address convergence limitations, the paradigm has been extended to Sample-based Krylov Quantum Diagonalization (SKQD)¹³⁶, which takes its roots from SQD and Krylov Quantum Diagonalization (KQD)¹³⁷, offering a polynomial-time convergence guarantee under assumptions. Further works have shown its efficient implementation via randomized compilation (SqDRIFT)¹³⁸ or time-independent unitary decomposition (QKUD)¹³⁹ to minimize algorithmic error. To capture dynamic correlation missing from the selected subspace, building on the foundational quantum-classical quantum Monte Carlo (QMC)¹⁴⁰, hybrid workflows merging QSCI with phaseless auxiliary-field quantum Monte Carlo (ph-AFQMC) have been established^{141,142}, along with GPU-acceleration support for the classical diagonalization kernels of SQD^{143,144}.

However, despite these algorithmic strides, fundamental limitations remain. Rigorous benchmarks indicate that QSCI and SQD struggle with sampling inefficiency when targeting high-accuracy results, often yielding configuration interaction expansions that are less compact than classical heuristics (e.g., Heat-Bath Configuration Interaction, HCI), where new determinants become exponentially difficult to find^{145,146}. Trying to address these scaling bottlenecks, novel variants have emerged. For instance, compact wavefunction-based QSCI leverages stochastic time evolution to predict likely excitations, yielding configuration spaces that are up to 200 times smaller than those obtained with conventional selection criteria, demonstrated on the IQM Emerald superconducting quantum hardware¹⁴⁷. Similarly, the Handover Iterative VQE (HI-VQE)¹⁴⁸ dynamically exchanges configuration data between quantum samplers and classical solvers; it has been shown to produce wavefunctions substantially more compact than classical HCI for strongly correlated systems like Fe-S clusters¹⁴⁹, while enabling active spaces as large as (24e, 22o) with superior noise resilience¹⁵⁰. Most recently, half-qubit architectures have been introduced to double the addressable active space size, allowing the qubit count to scale with the number of spatial orbitals rather than spin orbitals^{151–153}. Notable among these are Entanglement Forging SQD (EF-SQD)¹⁵¹, the independently developed HCI-Half qubit SQD (HCI-HSQD)¹⁵², and Doubly Occupied Configuration Interaction-QSCI (DOCI-QSCI)¹⁵³. EF-SQD utilizes Schmidt decomposition to reduce qubit overhead, enabling the simulation of radical chain reactions on IBM Heron superconducting quantum hardware¹⁵¹. Addressing sampling efficiency, HCI-HSQD explicitly mitigates the compactness bottleneck by applying classical HCI heuristics to guide sampling, thereby reducing energy errors by up to

76% in large clusters¹⁵². In parallel, DOCI-QSCI leverages seniority-zero schemes to efficiently double the qubit budget, utilizing ph-AFQMC to recover dynamic correlation lost by the seniority restriction across diverse molecular benchmarks¹⁵³. Complementing these techniques, machine-learning-anchored frameworks such as physics-informed generative SQD (PIGen-SQD) leverage generative models combined with perturbative screening to stochastically recover high-fidelity configurations, improving both sampling efficiency and subspace compactness under noisy conditions¹⁵⁴. While these methods represent significant progress, optimizing the trade-off between sampling cost and subspace compactness remains an active research direction. A concise comparison of representative quantum algorithms for quantum chemistry is provided in Table 1.

To extend the reach of these quantum solvers beyond small molecules to macroscopic systems, fragmentation and embedding strategies have become indispensable¹⁶⁰. Prominent approaches include energy-based fragmentation schemes such as the Fragment Molecular Orbital (FMO) method^{161,162} and the Many-Body Expansion (MBE)^{163–166}, which recover total energies through systematic many-body summation of fragment contributions. Recent adaptations of these methods for quantum hardware have demonstrated FMO-based VQE¹⁶⁷ and MBE-VQE workflows¹⁶⁸ capable of capturing correlation while reducing the required number of qubits. Similarly, the Divide-and-Conquer (DC) technique^{169–171} constructs global electronic structures from local density matrices, a strategy that has been mapped onto the VQE class of quantum algorithms to address larger systems^{172,173}. For systems requiring rigorous self-consistency, the Localized Active Space (LAS) method^{174,175} and Bootstrap Embedding (BE)^{176–178} enforce consistency constraints between fragments, enabling localized quantum chemistry simulations on quantum hardware^{179,180}. Foremost among embedding theories are Dynamical Mean-Field Theory (DMFT)^{181–183}, which treats local impurity problems in lattice systems, and Density Matrix Embedding Theory (DMET)^{184,185}, which enforces self-consistency between a fragment and its bath via the one-particle reduced density matrix. Recent theoretical analyses have refined the understanding of DMET’s convergence^{186,187}, and hybrid quantum-classical algorithms such as VQE have been applied within DMET frameworks to model correlated active regions in protein–ligand systems¹⁸⁸.

Recent advancements in quantum-centric supercomputing have moved towards integrating these fragmentation and embedding frameworks with the sampling-based quantum solvers such as QSCI/SQD. Notably, SQD algorithm has been coupled with DMET to compute ground-state energies of h-chains, conformers of cyclohexane, and ligand-like molecules, effectively treating strongly correlated impurities within a mean-field bath^{189,190}. Parallel efforts have replaced the exact diagonalization step in LAS Self-Consistent Field (LASSCF) theory with SQD (LAS-SQD), enabling the simulation of iron-sulfur clusters with fragment sizes that are intractable for standard LASSCF¹⁹¹. In the regime of high-accuracy thermochemistry, Quantum Bootstrap Embedding (QBE) has been combined with SQD to recover the correlation energy in large molecular systems by enforcing density-matching constraints across overlapping fragments¹⁹². Pushing towards macroscopic biological systems, the Embedded Wave Function (EWF) framework has integrated SQD to predict relative conformer energies of the 300-atom Trp-cage miniprotein, explicitly treating difficult fragments on quantum hardware while handling the remainder via classical FCI¹⁹³. Furthermore, layered hybrid protocols such as Our own N-layered Integrated molecular Orbital and molecular Mechanics (ONIOM)¹⁹⁴ have been extended to use Time-Evolved QSCI (TE-QSCI) for the quantum-mechanical region, facilitating excited-state calculations of biomolecules on hybrid quantum-HPC architectures using Quantinuum’s Reimei trapped-ion quantum hardware¹⁹⁵. Collectively, these results establish QSCI/SQD-driven embedding as a robust and scalable route for near-term quantum chemistry, capable of handling complex electronic structures that defy standalone NISQ algorithms.

In this work, we build upon these developments by presenting a comprehensive experimental study of SQD on the IQM Sirius superconducting Quantum Processing Unit (QPU). The device comprises 24 physical qubits, with up to 16 operational qubits arranged in a star topology that enables all-to-all connectivity via a central resonator. Leveraging this connectivity and the native gate set (R, CZ, Move), we implement SQD using the established Local Unitary Cluster Jastrow (LUCJ)¹⁹⁶ ansatz and introduce the Linear-CNOT Unitary Coupled-Cluster Singles and Doubles (LCNot-UCCSD)¹⁹⁷ ansatz within the SQD framework. While the LUCJ ansatz utilizes CCSD amplitudes for initialization, our LCNot-UCCSD implementation employs MP2-level initialization to reduce classical precomputation overhead. This enables a direct hardware comparison of ansatz expressivity and noise sensitivity, revealing critical trade-offs where the increased depth of LCNot-UCCSD challenges the sampling of symmetry-conserving states for larger molecules like H₂O and NH₃.

We systematically benchmark these approaches across a range of molecular systems, reporting detailed quantum resource estimates, sampled subspace dimensions, and energy comparisons against exact FCI references. Our study includes 1D-PES scans for H₂ and HeH⁺ in both STO-3G and 6-31G basis sets, as well as LiH and BeH₂ in the STO-3G basis set. Furthermore, we demonstrate the first experimental realization of a dense 32 × 32 2D-PES scan for the water molecule on superconducting hardware. While previous efforts have achieved chemically accurate ground-state energy estimates for water using VQE on trapped-ion processors¹⁹⁸ or constructed 2D surfaces utilizing Daubechies wavelets via simulations¹⁹⁹, our work explicitly captures the global energetic interplay between bond length and bond angle directly on a physical superconducting QPU. All calculations are performed utilizing 10,000 shots per run and evaluating the impact of post-configuration recovery parameters, comparing high-threshold filtering against symmetry-adaptive thresholds, on the accuracy of the diagonalized subspace.

Table 1. Comparative analysis of representative quantum algorithms for quantum chemistry. All claims are referenced to primary literature.

Feature	QPE (Fault-Tolerant)	VQE (Variational)	SQD / QSCI (Sampling-Based)
Year introduced	1995 ⁸⁵ ; applied to chemistry in 2005 ⁸²	2014 ¹¹⁰	2023 ¹²⁶ ; SQD variant demonstrated in 2025 ¹³⁰
Hardware regime	Requires full fault-tolerant QEC; target of application-scale quantum (FASQ) computing ^{85,100}	NISQ – shallow, noisy circuits tolerated ^{109,110}	NISQ-compatible; noise-tolerant by design ^{126,130}
Role of QPU	Phase estimator: extracts the eigenvalue of the unitary $e^{-i\hat{H}t}$ via quantum Fourier transform ^{82,85}	Energy estimator: evaluates expectation values $\langle \psi(\vec{\theta}) \hat{H} \psi(\vec{\theta}) \rangle$ for a parametrized ansatz ^{110,111}	Configuration sampler: generates dominant Slater determinants by measuring a parametrized state in the computational basis ^{126,130}
Classical role	Minimal: state preparation and post-measurement classical inference; The energy readout is single-shot ^{82,86}	Moderate: outer optimization loop minimizes $E(\vec{\theta})$ over circuit parameters using gradient-based or gradient-free optimizers ^{110,155}	Dominant: Hamiltonian construction and sparse CI diagonalization within the sampled subspace ^{126,130}
Ansatz/trial state	Arbitrary trial state with non-zero overlap $\gamma > 0$ with the target eigenstate; no parametrized optimization ^{82,86}	Parametrized quantum circuit (e.g., hardware-efficient or UCC-type ansatz); expressibility determines accuracy ^{113,114,156}	Any pre-initialized parameterized state (VQE-produced, Post-HF) or Hamiltonian-simulation-evolved; sampling quality depends on overlap ^{126–128}
Measurement cost	Low: a single coherent run yields a phase estimate; precision ϵ requires $\mathcal{O}(1/\epsilon)$ total queries (Heisenberg scaling) ^{85,86}	High: $\mathcal{O}(N^4)$ Pauli terms must be measured for molecular Hamiltonians; each requires $\mathcal{O}(1/\epsilon^2)$ shots per optimization step ^{111,155,157}	Low-to-moderate: shots define the subspace; energy is exact within that subspace by classical diagonalization; no per-term Pauli measurement overhead ^{126,130}
Circuit depth	Very deep: requires $\mathcal{O}(\text{poly}(N)/\epsilon)$ controlled-unitary applications and inverse QFT; infeasible without QEC ^{82,100,158}	Shallow-to-moderate: depth set by ansatz; compatible with near-term gate fidelities ^{109,110,159}	Shallow-to-moderate: depth set by ansatz; comparable to NISQ VQE circuits ^{126,130}
Variational bound on energy	Not variational; returns a statistical estimate of the eigenvalue with controlled precision ^{85,86}	Guaranteed in the noise-free limit; Hardware noise can violate the bounds, invalidating results ^{111,122}	Guaranteed within the selected subspace by exact classical diagonalization; Energy is an upper bound to FCI ^{126,130}
Noise resilience	Low: requires active QEC to suppress logical errors; impractical for NISQ hardware ^{100,109}	Moderate: variational bound degrades under hardware noise; barren plateaus impede optimization for large systems ^{119,120}	High: QPU used only for sampling; no variational optimisation loop; Classical diagonalization is exact and noise-free ^{126,130}
Scalability	Limited by qubit count, circuit depth, and QEC overhead; resource estimates for small molecules exceed 10^{10} T-gates ^{100,158}	Limited by classical optimizer convergence, $\mathcal{O}(N^4)$ measurement overhead, and barren plateaus in deep circuits ^{119,155}	Limited by sampling efficiency (finding new determinants becomes exponentially hard), subspace compactness, and classical diagonalization cost ^{145,146}
Key limitations	Demands full fault tolerance with thousands of logical qubits and deep circuits; far beyond NISQ hardware ^{100,158}	Barren plateaus, noise-induced bias, high measurement overhead, and difficult optimization landscapes for strongly correlated systems ^{119,120,122}	Quantum sampling efficiencies; CI expansions can be less compact than classical heuristics (e.g. HCI); classical diagonalization becomes a bottleneck for large subspaces ^{145,146}

Finally, to demonstrate scalability beyond small molecules, we integrate the SQD(LUCJ) solver with DMET to treat macroscopic chemical systems. We report chemically accurate active-space energies for a set of eight ligand-like molecules (43–75 Da) and the FDA-approved antiviral drug Amantadine (151 Da), benchmarking the results against DMET-CASCI references. These results, obtained with total runtimes ranging from minutes for small scans to hours for the full 2D surfaces, validate the hybrid quantum-classical workflow. The remainder of this paper is organized as follows: Section 2 details the theoretical foundations of SQD, the specific ansätze, and embedding strategies; Section 3 describes the experimental implementation on the IQM Sirius QPU’s architecture; Section 4 presents the benchmarking data; and Section 5 offers an outlook on future scalable workflows.

2 Theoretical Background

This section outlines the theoretical foundations underlying the methods employed in this work. We briefly review the formulation of the electronic structure problem within the Born–Oppenheimer approximation, the representation of molecular systems in finite basis sets, and the key approximations used to make the many-body Schrödinger equation tractable. These concepts provide the basis for both classical and quantum computational approaches discussed in the subsequent sections.

2.1 Ab Initio Quantum Chemistry Methods

The fundamental premise of ab initio quantum chemistry is the formulation and exact or approximate solution of the non-relativistic, time-independent Schrödinger equation³¹ for a system of interacting electrons and nuclei. The total molecular Hamiltonian, which dictates the complete energetic landscape of the system, is defined as a sum of kinetic and potential energy operators:

$$\hat{H} = \hat{T}_N + \hat{T}_e + \hat{V}_{Ne} + \hat{V}_{NN} + \hat{V}_{ee} \quad (1)$$

where $\hat{T}_N$ and $\hat{T}_e$ represent the kinetic energy operators of the atomic nuclei and the electrons, respectively. The potential energy terms $\hat{V}_{Ne}$, $\hat{V}_{NN}$, and $\hat{V}_{ee}$ represent the electron-nucleus Coulombic attraction, the nucleus-nucleus Coulombic repulsion, and the electron-electron Coulombic repulsion, respectively.

Because the mass of a proton is roughly 1,836 times greater than the mass of an electron^{200,201}, the atomic nuclei move on a significantly slower timescale than the rapidly fluctuating electron cloud. This immense physical disparity justifies the Born–Oppenheimer (BO) approximation³⁷, which mathematically decouples the nuclear and electronic wavefunctions. By treating the nuclear coordinates as fixed parameters rather than dynamic quantum variables, the nuclear kinetic energy term $\hat{T}_N$ is neglected, and the internuclear repulsion $\hat{V}_{NN}$ becomes a strictly constant scalar shift for any given geometry. The remaining electronic Schrödinger equation is expressed as:

$$\hat{H}_{ele}|\Psi_{ele}\rangle = E_{ele}|\Psi_{ele}\rangle \quad (2)$$

where the electronic Hamiltonian is defined as $\hat{H}_{ele} = \hat{T}_e + \hat{V}_{Ne} + \hat{V}_{ee}$. The total energy of the molecular system at a specific static nuclear geometry is consequently given by $E_{tot} = E_{ele} + \hat{V}_{NN}$.

2.1.1 Hartree–Fock (HF) and Restricted HF (RHF)

The Hartree–Fock^{49,50} (HF) method is the foundational mean-field approximation of wavefunction-based quantum chemistry, in which the exact many-electron ground state is approximated by a single Slater determinant, thereby enforcing antisymmetry and the Pauli exclusion principle²⁰² by construction. The resulting self-consistent field equations are defined by the Fock operator²⁰³

$$\hat{F}_i = \hat{h}_i + \sum_{j=1}^e (\hat{J}_j - \hat{K}_j), \quad (3)$$

where $\hat{h}_i$ contains the one-electron kinetic energy and nuclear attraction terms, $\hat{J}_j$ is the Coulomb operator describing the classical average electron–electron repulsion, and $\hat{K}_j$ is the exchange operator arising from the antisymmetry of the Slater determinant, reducing same-spin repulsion. The HF equations are solved iteratively to self-consistency, yielding a mean-field ground-state reference that neglects dynamical electron correlation⁵¹.

In closed-shell systems, Restricted Hartree–Fock (RHF) constrains α and β electrons to share identical spatial orbitals, preserving spin symmetry but often failing in bond dissociation and strongly correlated regimes due to the rigid double-occupancy constraint⁵¹.

2.1.2 Møller–Plesset (MP) Perturbation

To systematically recover a portion of the dynamical correlation energy absent in the mean-field HF approximation, Møller–Plesset (MP)²⁰⁴ perturbation theory applies Rayleigh–Schrödinger Perturbation Theory (RSPT)²⁰⁵ to $\hat{H}_{ele}$. In this framework, the unperturbed zeroth-order Hamiltonian $\hat{H}_0$ is defined exactly as the sum of the one-electron Fock operators, while the perturbation operator $\hat{V}$ is defined as the exact electron-electron repulsion minus the average mean-field HF potential.

Within RSPT based on the HF reference, the zeroth-order energy equals the sum of canonical orbital energies and the first-order correction recovers the HF energy. The leading correlation contribution therefore arises at second order, yielding the Møller–Plesset second order (MP2) correction. For double excitations from occupied orbitals i, j to virtual orbitals a, b , the correlation energy²⁰⁶ is

$$E_{\text{MP2}} = \sum_{i < j}^{\text{occ}} \sum_{a < b}^{\text{virt}} \frac{|\hat{V}_{ijab} - \hat{V}_{ijba}|^2}{\hat{F}_{ii} + \hat{F}_{jj} - \hat{F}_{aa} - \hat{F}_{bb}} = \sum_{i < j}^{\text{occ}} \sum_{a < b}^{\text{virt}} \langle ij || ab \rangle \left(\frac{\langle ij || ab \rangle}{\epsilon_i + \epsilon_j - \epsilon_a - \epsilon_b} \right) = \sum_{i < j}^{\text{occ}} \sum_{a < b}^{\text{virt}} \langle ij || ab \rangle t_{ij}^{ab}. \quad (4)$$

where $\langle ij || ab \rangle = \langle ij | ab \rangle - \langle ij | ba \rangle$ are antisymmetrized two-electron integrals in the canonical molecular orbital (MO) basis, with $\langle pq | rs \rangle$ denoting the Coulomb two-electron integrals and $\hat{V}_{ijab} = \langle ij | ab \rangle$ in this basis. The quantity $\epsilon_p = \hat{F}_{pp}$ is the diagonal Fock matrix element corresponding to the canonical orbital energy of the p^{th} orbital.

The restricted summations $i < j$ and $a < b$ avoid double counting due to the antisymmetry of the two-electron integrals, the amplitudes

$$t_{ij}^{ab} = \frac{\langle ij || ab \rangle}{(\epsilon_i + \epsilon_j - \epsilon_a - \epsilon_b)} \quad (5)$$

are the MP2 double-excitation parameters obtained from first-order perturbation theory. These amplitudes constitute the first-order, non-variational approximation to the CCSD²⁰⁷ amplitudes and are therefore commonly used to initialize classical CCSD calculations as well as chemically motivated unitary parametrizations such as UCCSD¹¹⁴ ansatz in variational quantum algorithms.

When MO two-electron integrals are precomputed and stored, the dominant cost of the MP2 energy *evaluation step* arises from tensor contractions over pairs of occupied and virtual orbitals. The energy contraction itself therefore scales as $\mathcal{O}(N^4)$ with respect to a unified measure of system size N , reflecting the double summation over occupied and virtual index pairs. Under these conditions, MP2 remains computationally inexpensive for moderate-to-large molecular systems while recovering approximately 70%-95% of the dynamical correlation energy²⁰⁸. Owing to its favourable cost-to-accuracy ratio and its perturbative foundation on the HF reference, MP2 provides an efficient and physically well-motivated approximation to electron correlation effects.

2.1.3 Coupled Cluster (CC)

While perturbation theory provides a non-iterative correlation correction, Coupled Cluster (CC) theory^{207,209} provides a rigorous, size-consistent, and size-extensive methodology for capturing many-body electron correlation to infinite order within a truncated excitation space. The CC wavefunction^{209,210} is formulated via an exponential ansatz applied to the reference RHF determinant:

$$|\Psi_{\text{CC}}\rangle = e^{\hat{T}} |\Phi_0\rangle \quad (6)$$

where the cluster operator $\hat{T}$ generates electronic excitations from the occupied to the virtual space. For the widely utilized CCSD formulation, the operator is truncated at the level of one- and two-body excitations,

$$|\Psi_{\text{CCSD}}\rangle = e^{\hat{T}_1 + \hat{T}_2} |\Phi_0\rangle, \quad \hat{T}_1 = \sum_{i,a} t_i^a \hat{a}_a^\dagger \hat{a}_i, \quad \hat{T}_2 = \frac{1}{4} \sum_{i,j,a,b} t_{ij}^{ab} \hat{a}_a^\dagger \hat{a}_b^\dagger \hat{a}_j \hat{a}_i \quad (7)$$

where i, j denote occupied spin-orbitals and a, b denote virtual spin-orbitals. The amplitudes t_i^a and t_{ij}^{ab} correspond to single and double excitations, respectively.

The defining theoretical advantage of the CC method lies in the exponential structure of $e^{\hat{T}}$. The Taylor series expansion of the exponential generates disconnected higher-order excitations (e.g., the $\hat{T}_2^2/2$ term corresponds to disconnected quadruple excitations), thereby guaranteeing size extensivity^{211,212}. This ensures that the computed energy scales properly with system size, in contrast to truncated Configuration Interaction methods, which are neither size-extensive nor size-consistent. Classical CCSD requires the iterative solution of nonlinear amplitude equations and exhibits a computational scaling commonly expressed as $\mathcal{O}(N^6)$ ²¹¹.

In the context of the present hybrid quantum-HPC framework, classical CCSD amplitudes are extracted and mapped to initialize the variational parameters of the LUCJ ansatz¹⁹⁶. This informed initialization accelerates convergence of the quantum optimization and places the variational search within a physically meaningful region of parameter space from the outset¹⁹⁶.

2.1.4 Configuration Interaction (CI)

Configuration Interaction (CI)^{43,45,213} approaches the electron correlation problem by expressing the exact many-body wavefunction as a linear combination of all possible Slater determinants governed by the number and spin symmetry of the molecular system. These determinants are generated by systematically exciting electrons from the HF reference state $|\Phi_0\rangle$ into all available virtual orbitals:

$$|\Psi_{\text{CI}}\rangle = c_0|\Phi_0\rangle + \sum_{a,r} c_a^r |\Phi_a^r\rangle + \sum_{a<b, r<s} c_{ab}^{rs} |\Phi_{ab}^{rs}\rangle + \dots \quad (8)$$

When all possible excitations are explicitly included, the method is termed Full Configuration Interaction (FCI). The FCI wavefunction provides the exact solution to the electronic Schrödinger equation (Eq. (2)) within the limits of the chosen finite basis set.

Organizing the Slater determinants by excitation rank relative to the HF reference,

$$\{|\Phi_0\rangle, |\Phi_i^a\rangle, |\Phi_{ij}^{ab}\rangle, |\Phi_{ijk}^{abc}\rangle, |\Phi_{ijkl}^{abcd}\rangle, \dots\} \equiv \{|\Phi_0\rangle, |S\rangle, |D\rangle, |T\rangle, |Q\rangle, \dots\}. \quad (9)$$

where, $|S\rangle$, $|D\rangle$, $|T\rangle$, and $|Q\rangle$ denote the manifolds of single, double, triple, and quadruple excitations, respectively. The Hamiltonian matrix in this ordered basis acquires a characteristic block-banded structure²¹⁴:

	$ \Phi_0\rangle$	$ S\rangle$	$ D\rangle$	$ T\rangle$	$ Q\rangle$	$\dots$
$\langle\Phi_0 $	$\langle\Phi_0 \hat{H} \Phi_0\rangle$	0	$\langle\Phi_0 \hat{H} D\rangle$	0	0	$\dots$
$\langle S $	0	$\langle S \hat{H} S\rangle$	$\langle S \hat{H} D\rangle$	$\langle S \hat{H} T\rangle$	0	$\dots$
$\langle D $	$\langle D \hat{H} \Phi_0\rangle$	$\langle D \hat{H} S\rangle$	$\langle D \hat{H} D\rangle$	$\langle D \hat{H} T\rangle$	$\langle D \hat{H} Q\rangle$	$\dots$
$\langle T $	0	$\langle T \hat{H} S\rangle$	$\langle T \hat{H} D\rangle$	$\langle T \hat{H} T\rangle$	$\langle T \hat{H} Q\rangle$	$\dots$
$\langle Q $	0	0	$\langle Q \hat{H} D\rangle$	$\langle Q \hat{H} T\rangle$	$\langle Q \hat{H} Q\rangle$	$\dots$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\ddots$

(10)

The matrix is Hermitian, and only the upper triangle need be considered.

The block structure of the CI matrix shown in Eq. (10) arises from the two-body form of $\hat{H}_{ele}$, which can change the occupation of at most two spin-orbitals in a Slater determinant. As a consequence, matrix elements between determinants whose excitation ranks differ by more than two vanish, producing the banded coupling pattern between the reference, single, double, and higher excitation manifolds.

In addition, when canonical HF orbitals are used as the reference, Brillouin's theorem²¹⁵ states that the HF determinant does not couple to singly excited determinants through the Hamiltonian. This leads to the vanishing reference-single block in the CI Hamiltonian matrix, i.e. $\langle\Phi_0|\hat{H}|\Phi_i^a\rangle = 0$.

The dimension of the FCI space grows combinatorially²¹⁶ with the number of spin-orbitals. For a system with N_{orb} spatial orbitals and N_α, N_β electrons of each spin, the symmetry-constrained Hilbert space $\mathbb{S}$ size is

$$|\mathbb{S}| = \binom{N_{orb}}{N_\alpha} \binom{N_{orb}}{N_\beta}, \quad (11)$$

which leads to an exponential growth of the CI Hamiltonian dimension with system size.

2.2 Basis Sets

In *ab initio* quantum chemistry, the electronic wavefunction is represented in a finite basis through the Linear Combination of Atomic Orbitals (LCAO) approximation⁵¹, where molecular orbitals are expanded as linear combinations of atom-centered basis functions. Although Slater-type orbitals (STOs)²¹⁷, with radial form $e^{-\zeta r}$, provide a physically accurate description of atomic orbitals by reproducing the correct nuclear cusp and exponential decay, evaluating the required multi-center electron repulsion integrals over STOs is computationally demanding. For this reason, practical electronic structure calculations employ Gaussian-type orbitals (GTOs)^{218,219}, $e^{-\alpha r^2}$, whose mathematical properties allow efficient analytical evaluation of molecular integrals via the Gaussian product theorem. In modern basis sets, contracted Gaussian functions (linear combinations of several primitive Gaussians) are optimized to approximate the radial behavior of STOs while retaining the computational advantages of Gaussian integrals.

The choice of basis set significantly influences the accuracy of computed molecular energetics. Minimal basis sets, such as STO-3G, assign one contracted function to each atomic orbital (AO) and therefore provide a highly compact representation that is computationally inexpensive but lacks the flexibility required to describe polarization and electron correlation effects accurately. Larger split-valence basis sets, such as 6-31G²²⁰, introduce *additional radial functions for valence orbitals*, allowing a more flexible description of electron density and generally improving energetic predictions. However, the increased number of basis functions enlarges the resulting Hilbert space $\mathbb{H}$ and the number of $\hat{H}_{ele}$ terms, which directly translates to higher qubit requirements and deeper circuits in quantum simulations.

Figure 1 shows the basis functions used to describe the carbon atom in (a) the STO-3G basis and (b) the 6-31G basis. In both cases, the $1s$ orbital represents the core orbital. In the 6-31G basis, the valence space is described using a split-valence (double- ζ) representation, in which each valence atomic orbital ($2s$, $2p_x$, $2p_y$, $2p_z$) is expanded using two basis functions with different radial extents: $2s \rightarrow (2s', 2s'')$, $2p_x \rightarrow (2p'_x, 2p''_x)$, $2p_y \rightarrow (2p'_y, 2p''_y)$, $2p_z \rightarrow (2p'_z, 2p''_z)$. This effectively doubles the number of basis functions associated with the valence shell, leading to four additional basis functions compared to STO-3G. Consequently, while basis sets such as 6-31G provide a more flexible and accurate description of the electronic structure, minimal bases like STO-3G remain a practical choice for current quantum hardware due to their significantly reduced computational cost.




Figure 1. An illustrative depiction of the iso-spheres of the orbitals used for describing the carbon atom in (a) STO-3G basis set requiring 5 orbitals $1s, 2s, 2p_x, 2p_y, 2p_z$ and in (b) 6-31G basis set which requires 9 orbitals $1s, 2s, 2p_x, 2p_y, 2p_z, 2s', 2p'_x, 2p'_y, 2p'_z$. In addition to the orbitals which have a possibility of filling in the p-block (carbon atom), an additional orbital from the next block is picked for all the orbitals except the core orbital $1s$. Hence, the additional orbitals $2s', 2p'_x, 2p'_y$ and $2p'_z$ are also used. These orbitals were generated using PySCF, and visualized in Mol* (MolStar)²²¹ online software.

2.3 Linear-CNOT Unitary Coupled-Cluster Singles and Doubles (LCNot-UCCSD) Ansatz

The Unitary Coupled-Cluster (UCC) framework²²² is one of the most widely studied ansätze in quantum computational chemistry, as it provides a chemically motivated variational form that systematically incorporates electron correlation through excitation operators. In the commonly used singles and doubles truncation (UCCSD), the trial state is generated by exponentiating fermionic excitation operators that promote electrons from occupied to virtual orbitals. To preserve unitarity on a quantum processor, the ansatz is constructed from an anti-Hermitian generator,

$$U(\vec{\theta}) = e^{\hat{T}(\vec{\theta}) - \hat{T}^\dagger(\vec{\theta})}, \quad (12)$$

where $\hat{T}(\vec{\theta})$ contains the single and double excitation operators. In practice, the exponential operator is factorized using a Trotter–Suzuki decomposition²²³, producing a sequence of exponentiated Pauli strings acting on qubits. When fermionic operators are mapped to qubits using transformations such as Jordan-Wigner (JW)²²⁴, each excitation decomposes into multiple Pauli strings whose number grows rapidly with excitation rank: a *single excitation* produces *two Pauli strings*, a *double excitation* produces *eight*, and higher-rank excitations generate increasingly large operator sets. These Pauli strings contain long parity chains required to enforce fermionic statistics, which in turn lead to substantial control (C-NOT) overhead in the resulting quantum circuits.

To mitigate this overhead, alternative formulations have been proposed that reinterpret fermionic excitations directly in the qubit space. In particular, the qubitized excitation operators¹⁹⁷ reformulate the excitation generators so that they can be implemented using simplified Pauli operators with reduced control (C-NOT) structure²²⁵. In this formulation, fermionic

excitations are replaced by qubit excitation operators that act directly on pairs or quartets of qubits corresponding to occupied and virtual orbitals. The corresponding circuits for implementing the single and double excitation operators are shown in Fig. 2. Building on this idea, the UCCSD ansatz with excitation gates constructed with linearly scaling C-NOT gates¹⁹⁷ (henceforth abbreviated as LCNot-UCCSD) modifies the implementation of excitation operators by relaxing the fermionic parity structure introduced by mappings such as JW. In the standard construction, the resulting Pauli strings contain long Z-parity chains that are implemented using C-NOT ladders to propagate parity across the qubit register, leading to circuit depths that grow with orbital separation.

LCNot-UCCSD, on the other hand, removes selected parity-control links, thereby shortening these ladders and reducing the overall C-NOT gate cost while preserving the variational structure of the ansatz. As a result, the implementation of each excitation operator requires only a linear number of C-NOT gates, rather than the exponential growth characteristic of standard constructions¹¹⁴. This modification significantly reduces circuit depth, making practical implementations feasible on near-term quantum hardware. Importantly, while the circuit depth is reduced, the overall operator and parameter complexity remains unchanged, scaling as $\mathcal{O}(N^4)$ with the number of spin orbitals $2N_{\text{orb}}$ (equivalently N_Q qubits), consistent with standard UCCSD formulations^{110,114}. Some LCNot-UCCSD approximations allow for the exact fermionic parity constraints to be partially relaxed, resulting in operators that no longer strictly conserve particle number and spin projection, thereby introducing a small degree of symmetry leakage¹⁹⁷. However, in the present work, no such approximations are considered, as the goal is to obtain a one-to-one correspondence with the standard UCCSD ansatz.




Figure 2. The single excitation gate (*left*) and double excitation gate (*right*) represented in the Qubit Excitation Based (QEB) method¹⁹⁷.

2.4 Local Unitary Cluster Jastrow (LUCJ) Ansatz

The LUCJ ansatz provides a highly localized, Hubbard-inspired alternative that avoids the steep CNOT scaling associated with JW transformed UCC methods¹⁹⁶. The canonical Unitary Cluster Jastrow (UCJ) ansatz is constructed as an L -layer unitary acting on the HF reference state $|\Phi_0\rangle$, where each layer alternates between orbital rotations and a Jastrow correlation operator. The LUCJ ansatz is written as

$$U(\vec{\theta}) = \prod_{\mu=1}^L e^{\hat{T}_\mu} e^{\hat{J}_\mu} e^{-\hat{T}_\mu}. \quad (13)$$

A representative quantum circuit implementation of this layered structure is shown in Figure 3. This structure separates the wavefunction transformation into two physically meaningful components: a kinetic/cluster operator $\hat{T}_\mu$ generating generalized one-body orbital rotations and a Jastrow correlator $\hat{J}_\mu$ introducing explicit electron–electron interactions.

The cluster operator $\hat{T}_\mu$ generates orbital rotations within the active space and can be expressed as a generalized one-body fermionic operator,

$$\hat{T}_\mu = \sum_{p,q,\sigma} T_{pq}^\mu \hat{a}_{p\sigma}^\dagger \hat{a}_{q\sigma}, \quad (14)$$

where T_{pq}^μ is an anti-Hermitian matrix of variational parameters, $\hat{a}_{p\sigma}^\dagger$ and $\hat{a}_{q\sigma}$ denote fermionic creation and annihilation operators, and $\sigma \in \{\alpha, \beta\}$ labels the spin degree of freedom. Exponentiation of this operator, $e^{\hat{T}_\mu}$ implements a unitary orbital rotation acting within the active space, dynamically reconfiguring the single-particle basis without introducing the higher-order fermionic excitation operators characteristic of conventional UCC expansions.

The Jastrow factor $\hat{J}_\mu$ complements this transformation by encoding explicit two-body correlations through number–number interactions,

$$\hat{J}_\mu = i \sum_{p,q,\sigma,\tau} J_{pq,\sigma\tau}^\mu \hat{n}_{p\sigma} \hat{n}_{q\tau}, \quad (15)$$

where $\hat{n}_{p\sigma} = \hat{a}_{p\sigma}^\dagger \hat{a}_{p\sigma}$ is the number operator and $J_{pq,\sigma\tau}^\mu$ is a real symmetric tensor of variational parameters. This operator penalizes energetically unfavorable electron configurations and introduces conditional dependence between occupations, thereby introducing correlation effects beyond the single-reference HF description.

In the LUCJ variant, the operator pool is strongly sparsified by restricting correlations to physically motivated local interactions, such as on-site opposite-spin terms and nearest-neighbor couplings. This locality assumption reflects the spatially short-ranged nature of electronic correlation and allows the ansatz to map naturally onto hardware connectivity graphs, eliminating the need for deep fermionic SWAP networks. Under all-to-all qubit connectivity, this sparsification further ensures that the operator and parameter complexity scales as $\mathcal{O}(N^2)$ in the number of spin orbitals¹⁹⁶. As a result, the number of variational parameters and the resulting circuit depth remain tightly controlled, enabling hardware-efficient implementations while retaining sufficient flexibility to capture strongly correlated electronic structure¹⁹⁶.




Figure 3. Quantum circuit for a single layer of the LUCJ ansatz acting on α and β spin registers. Each layer implements a sequence of localized orbital rotations $e^{\pm\hat{T}}$ and spin-resolved Jastrow correlators $e^{i\hat{J}}$, including inter-spin coupling $e^{i\hat{J}_{\alpha\beta}}$, enabling efficient encoding of one- and two-body correlations within a locality-constrained framework.

2.5 Quantum-Selected Configuration Interaction (QSCI) / Sample-Based Quantum Diagonalization (SQD)

Quantum-Selected Configuration Interaction (QSCI)¹²⁶ is a hybrid quantum–classical algorithm designed to exploit the sampling capability of QPUs while avoiding the large measurement overheads associated with VQEs. In conventional VQE¹¹⁰, the quantum device repeatedly evaluates expectation values of many Pauli terms in the Hamiltonian, often requiring millions of measurement shots to suppress statistical noise. QSCI instead uses the QPU purely as a sampler of the wavefunction in the computational basis. An approximate state $|\psi_{\text{in}}\rangle$, prepared using a shallow or heuristic ansatz with $|\vec{\theta}|$ variational parameters and non-zero overlap with the ground state, is repeatedly measured in the Slater-determinant basis using N_{shots} measurements. The resulting bitstrings $\{x_i\}$ reveal the dominant electronic configurations of the wavefunction, enabling the construction of a truncated but chemically relevant diagonalization subspace $\mathbb{S}_{\text{sub}} \subset \mathbb{S}$ with dimension $|\mathbb{S}_{\text{sub}}| \ll |\mathbb{S}|$, where $|\mathbb{H}|$ denotes the full Hilbert-space dimension. Because the QPU is used only to identify important configurations rather than to evaluate Hamiltonian expectation values, the algorithm avoids the severe measurement bottleneck that limits many NISQ-era variational algorithms.

More specifically, measurements of the prepared state $|\psi_{\text{in}}\rangle$ generate a sampling space of determinants

$$\mathbb{S}_{\text{samp}} = \{ |x_i\rangle \}_{i=1}^{|\mathbb{S}_{\text{samp}}|}, \quad (16)$$

whose size $|\mathbb{S}_{\text{samp}}|$ depends on the number of shots N_{shots} . From this space, symmetry filtering is applied to retain only determinants belonging to the physically allowed symmetry sector $\mathbb{S} \subset \mathbb{H}$ with dimension $|\mathbb{S}|$, corresponding to the correct electron number and spin quantum numbers for a molecular system with $(N_{\text{orb}}, N_{\text{ele}})$, to obtain the post-selected subspace $\mathbb{S}_{\text{post-sel}}$.

The distinct α and β -spin configurations in $\mathbb{S}_{\text{post-sel}}$ are identified and combined to form a unified set, which is subsequently tensored with itself to construct the final subspace for Hamiltonian diagonalization^{130,190}. Once $\mathbb{S}_{\text{sub}}$ has been identified, all subsequent steps are performed classically. The Hamiltonian is projected into the selected determinant basis using Slater–Condon rules,

$$(\hat{H}_{\text{sub}})_{ij} = \langle x_i | \hat{H} | x_j \rangle, \quad |x_i\rangle, |x_j\rangle \in \mathbb{S}_{\text{sub}}, \quad (17)$$

producing a noise-free matrix representation of the Hamiltonian in $\mathbb{S}_{\text{sub}}$. The resulting matrix is then diagonalized using classical iterative eigensolvers such as Davidson method²²⁶. Because all matrix elements are evaluated deterministically on classical hardware, the resulting energy obeys a strict variational bound²²⁷

$$E_{\text{QSCI}} = \text{mineig}(\hat{H}_{\text{sub}}) \geq E_{\text{exact}}, \quad (18)$$

ensuring that the estimated ground-state energy cannot fall below the true value. In this framework, quantum errors do not corrupt the Hamiltonian evaluation itself but only affect the quality of $\mathbb{S}_{\text{sub}}$ by potentially missing important configurations.

Sample-Based Quantum Diagonalization (SQD)¹³⁰ extends QSCI by incorporating classical error-mitigation procedures that improve robustness to noise on NISQ devices. In practice, hardware errors frequently produce bitstrings that violate physical symmetries such as particle-number or spin conservation. Rather than discarding these samples, SQD introduces symmetry-aware configuration recovery (CR) methods that probabilistically correct corrupted bitstrings toward a reference orbital-occupancy vector derived from valid samples or mean-field estimates. This recovery procedure generates a corrected determinant pool $\mathbb{S}_{\text{post-cr}}$, with dimension $|\mathbb{S}_{\text{post-cr}}|$, which is subsequently used to construct the final diagonalization subspace $\mathbb{S}_{\text{sub}}$, by means of proliferation^{130, 190}.

This recovery process enables SQD to retain useful information even when a large fraction of raw samples are noisy, significantly improving performance on current quantum hardware. As a result, SQD represents a particularly promising architecture for near-term quantum chemistry simulations, where the QPU is used to perform quantum sampling, while classical supercomputers perform deterministic Hamiltonian construction and diagonalization within $\mathbb{S}_{\text{sub}}$.

2.6 Density-Matrix Embedding Theory (DMET)

DMET^{184, 185} originates from the recognition that conventional embedding and fragmentation approaches^{228–230} fail for fragments strongly coupled to their environment, such as through covalent bonds. In such cases, the fragment constitutes an open quantum system entangled with the rest of the system. Formally, an exact embedding can be achieved by replacing the environment with a finite quantum bath that reproduces the fragment–environment entanglement of the full many-body state. This follows from the Schmidt decomposition²³¹, which guarantees that this entanglement can be represented using a finite number of bath degrees of freedom.

A practical formulation of DMET¹⁸⁷ proceeds at the level of the one-particle reduced density matrix (1-RDM). Let $\mathbb{H}_{\text{lin}}$ denote the one-particle Hilbert space of dimension L (where L corresponds to the number of spin orbitals), with associated fermionic Fock space $\mathbb{F}(\mathbb{H}_{\text{lin}}) = \bigoplus_{n=0}^L \wedge^n \mathbb{H}_{\text{lin}}$, corresponding to $\mathbb{H}$. For a many-body state described by a density operator Γ , the 1-RDM D , a Hermitian operator acting on $\mathbb{H}_{\text{lin}}$ is defined as

$$D_{\mu\nu} = \text{Tr}(\Gamma \hat{a}_\nu^\dagger \hat{a}_\mu), \quad (19)$$

with $0 \leq D \leq 1$ and $\text{Tr}(D) = N$, and $\hat{a}_\nu^\dagger(\hat{a}_\mu)$ correspond to the creation(annihilation) operator acting on the ν^{th} (μ^{th}) spin orbital. For a pure state $\Gamma = |\Psi\rangle\langle\Psi|$, this reduces to

$$D_{\mu\nu} = \langle\Psi| \hat{a}_\nu^\dagger \hat{a}_\mu |\Psi\rangle, \quad (20)$$

making explicit that the 1-RDM encodes all one-body observables. In the special case of a single Slater determinant, the 1-RDM reduces to $D = C_{\text{occ}} C_{\text{occ}}^\dagger$, where C_{occ} contains the occupied molecular orbital coefficients.

In conventional DMET¹⁸⁵, the low-level description is restricted to idempotent one-particle density matrices ($D^2 = D$). Let $\mathbb{H}_{\text{lin}} = \mathbb{H}_S \oplus \mathbb{H}_E$ denote a partition of the one-particle Hilbert space, where $\mathbb{H}_S$ and $\mathbb{H}_E$ are the Hilbert spaces of the system and the environment, with $|\mathbb{H}_S| = \ell$. The idempotency of D implies that the fragment–environment coupling is of rank at most ℓ , or equivalently that the environment block of D has no more than ℓ eigenvalues in the interval $(0, 1)$ ²³² in the molecular spin orbital basis and $(0, 2)$ in the molecular spatial orbital basis. The corresponding eigenvectors define the bath orbitals, which, together with the fragment orbitals, span the impurity subspace.

In practice, the bath is constructed from an approximate mean-field state. The Schmidt decomposition then reduces to a single-particle problem involving the diagonalization of the fragment–environment block of the 1-RDM. The resulting impurity Hamiltonian, defined on the fragment and bath orbitals, is solved using a high-level correlated method (like SQD, in our case) to obtain embedded 1- and 2- RDMs.

The contribution of fragment A to the total energy¹⁸⁵ is expressed as

$$E_A = \sum_{i \in A} \sum_p D_i^p h_{pi} + \frac{1}{2} \sum_{i \in A} \sum_{pq} \Gamma_{ip}^{qr} V_{qrip}, \quad (21)$$

where i, j label fragment orbitals and p, q, r span the impurity space. The two-particle reduced density matrix is defined as

$$\Gamma_{ip}^{qr} = \langle a_q^\dagger a_r^\dagger a_p a_i \rangle, \quad (22)$$

and h_{pi} and V_{qip} denote one- and two-electron integrals in the impurity basis. Restricting at least one index to fragment A ensures each interaction term is counted once while allowing exchange of particles and correlations with the bath. The total energy is then

$$E_{\text{tot}} = \sum_A E_A, \quad (23)$$

yielding a size-extensive partitioning within the embedding framework¹⁸⁵.

To ensure global particle number consistency, a chemical potential μ_{glob} is introduced and iteratively optimized such that the total electron number obtained from the embedded fragments matches the target electron number N_{ele} . This is achieved by solving

$$\sum_A \text{Tr}(D^{(A)}(\mu_{\text{glob}})) = N_{\text{ele}}, \quad (24)$$

where $D^{(A)}$ denotes the embedded 1-RDM of fragment A. The chemical potential μ_{glob} is adjusted self-consistently until this condition is satisfied.

3 Methodology

3.1 Sample-based Quantum Diagonalization (SQD) Implementation

3.1.1 Software Stack and Electronic Structure Preprocessing

All experiments were carried out on IQM’s Sirius 24-qubit superconducting quantum hardware, with full runtime calibration details, hardware metrics, connectivity map, and qubit utilization provided in Appendix B. The experiments were implemented using Python v3.11.14^{233,234} and accessed via iqm-client v32.1.1²³⁵. Quantum circuits were constructed and managed using Qiskit v1.4.2²³⁶, with ansatz-specific construction relying on ffsim v0.0.56²³⁷ for the LUCJ ansatz and native custom Qiskit gate primitives for the LCNot-UCCSD ansatz. The SQD post-processing layer was performed using qiskit-addon-sqd v0.10.0²³⁸. Classical electronic structure computations, including mean-field, integral generation, coupled-cluster, and FCI, were performed using PySCF v2.11.0^{239–241}. Numerical linear algebra relied on NumPy v2.3.5²⁴² and SciPy 1.15.3²⁴³, with JAX v0.8.1²⁴⁴ available as an accelerated backend for internal ffsim operations. Results were aggregated and persisted using Pandas v2.3.3²⁴⁵, with each experimental record written incrementally to the output file to guard against runtime interruptions.

For each molecular system, the molecular Hamiltonian was constructed in the molecular orbital (MO) basis using PySCF. A molecular geometry object was instantiated with the target coordinates and the STO-3G basis set (or 6-31G for selected PES scans as described in Section 3.2), and a restricted Hartree–Fock (RHF) solution was computed. The one-electron and two-electron integrals in the MO basis were obtained as

$$h_{pq}^{\text{MO}} = \sum_{\mu\nu} C_{\mu p} h_{\mu\nu}^{\text{AO}} C_{\nu q} \quad \forall \quad h_{\mu\nu}^{\text{AO}} = \langle \mu | \hat{T} + \hat{V}_{\text{nuc}} | \nu \rangle, \quad (25)$$

$$h_{pqrs}^{\text{MO}} = \sum_{\mu\nu\lambda\sigma} C_{\mu p} C_{\nu q} C_{\lambda r} C_{\sigma s} \langle \mu\nu | \lambda\sigma \rangle^{\text{AO}}, \quad (26)$$

where $C_{\mu p}$ are the Hartree–Fock molecular orbital coefficients, $h_{\mu\nu}^{\text{AO}}$ denotes the (μ, ν) element of the one-electron integral matrix in the atomic orbital (AO) basis, and $\langle \mu\nu | \lambda\sigma \rangle^{\text{AO}}$ denotes the $(\mu, \nu, \lambda, \sigma)$ element of the two-electron integral tensor in chemist’s notation. The above expressions therefore represent the explicit two-index and four-index AO-to-MO tensor transformations of the one- and two-electron integrals. The nuclear repulsion energy E_{nuc} was obtained from the mean-field object and added to the electronic correlation energy from the SQD solver. No frozen-core approximation was applied in the SQD benchmarks and subsequent PES calculations described in Section 3.2.

3.1.2 Active Space Specification

The active space for each molecule, defined by the number of active spatial orbitals and active electrons ($N_{\text{orb}}, N_{\text{ele}}$), was specified as follows: H₂ (2, 2), LiH (6, 4), BeH₂ (7, 6), H₂O (7, 10), and NH₃ (8, 10). All systems employed the STO-3G basis and a closed-shell singlet reference, for which the total spin is $S = 0$ and the multiplicity is $2S + 1 = 1$. The qubit register for each molecule maps to $2N_{\text{orb}}$ spin orbitals under the Jordan–Wigner (JW) transformation.

The symmetry-constrained Hilbert space, or simply the symmetry space size $|\mathbb{S}|$, was computed for each active space as

$$|\mathbb{S}| = \binom{N_{\text{orb}}}{N_{\text{ele}}/2}, \quad (27)$$

which represents the number of particle-number-conserving determinants in the spin-restricted subspace and serves as the normalization reference for evaluating subspace coverage.

3.1.3 Ansatz Construction and Transpilation

Local Unitary Cluster Jastrow (LUCJ) ansatz. The LUCJ circuit was constructed using `ffsim`, which accepts CCSD t_1 and t_2 amplitude tensors, to define a spin-balanced UCJ operator restricted to a single repetition layer ($\text{LUCJ}_{\text{reps}} = 1$). The initial Hartree-Fock reference state was generated first, which was appended to the UCJ unitary, both employing the JW FTQM. CCSD amplitudes (t_1, t_2) were obtained using `PySCF` for the LUCJ ansatz initialization. This initialization strategy fixes the LUCJ parameters at their CCSD-optimal values without further variational optimization on hardware, at a pre-execution classical overhead of a single coupled-cluster calculation.

Linear CNOT Unitary Coupled Cluster Singles and Doubles (LCNot-UCCSD) ansatz. The LCNot-UCCSD circuit was constructed entirely within `Qiskit` using a custom CNOT-ladder decomposition. The Hartree-Fock reference was prepared by applying X gates on the qubits corresponding to occupied spin orbitals. Single excitations were implemented as a $CX - CR_Y - CX$ sequence; double excitations employed a three-CNOT ladder (CX on pairs $(b, a), (j, i), (b, j)$) followed by a 3-controlled R_Y gate and its inverse entangling structure, as shown in Figure 2. Throughout, no approximations of the excitation gates were used, retaining full circuit fidelity at the cost of increased depth. Parameters were initialized with MP2 amplitudes (t_{ij}^{ab}) computed as shown in Eq. (5). The total parameter count equals the number of single plus double excitations generated for the active space.

Transpilation. Both ansätze were transpiled to the instruction set architecture (ISA) of the IQM Sirius backend using `Qiskit`. A custom quantum circuit resource estimator routine computed categorized gate counts (single-qubit, two-qubit, multi-qubit) and circuit depth for each transpiled circuit; these metrics were recorded per experiment in the output CSV, further presented in Section 4. The two workflows are schematically represented in Figure 4. The overall process is divided into Classical Pre-Processing, Quantum Computation, and Classical Post-Processing. After calculating the $1e^-$ and $2e^-$ integrals (Eq. (25) and Eq. (26)), the pipeline splits into two distinct paths. In the LUCJ pipeline (path a), the classically intensive stage is the CCSD pre-computation, representing a classical trade-off with $\mathcal{O}(N^6)$ scaling²¹¹. However, this is balanced by an advantageous quantum circuit preparation scaling at $\mathcal{O}(N^2)$ ¹⁹⁶. Conversely, the LCNot-UCCSD pipeline (path b) employs a cheaper, advantageous MP2 initialization at $\mathcal{O}(N^4)$ (considering that the h_{pq}^{MO} and h_{pqrs}^{MO} are pre-computed and stored), shifting the complexity to a deeper quantum circuit representing a trade-off scaling at $\mathcal{O}(N^4)$ ^{110,114}. Both paths converge at quantum sampling, followed by classical error mitigation (configuration recovery) and classical diagonalization via SCI to output the final energy, E_{SQD} .

3.1.4 Hardware Execution and Sampling Protocol

Circuits were submitted to the Sirius backend, where the number of shots (N_{shots}) was set to 10,000. Each circuit was executed in three independent hardware runs, with separate job submissions and distinct random seeds for the classical post-processing stages. A 0.5 seconds inter-run sleep interval was inserted between submissions to avoid contention on the hardware queue. For each run, a dictionary of observed bitstrings and their frequencies was obtained, which was converted to a full bitstring matrix and probability array for further post-processing.

3.1.5 Iterative Configuration Recovery and Classical Diagonalization

The SQD self-consistent loop was executed for 10 iterations per hardware run. On the first iteration, the full bitstring matrix from hardware was used directly without configuration recovery, since no prior estimate of orbital occupancies was available.

For all subsequent iterations, configurations were recovered using the full bitstring matrix, the empirical probability array, and the running average orbital occupancy vector from the previous iteration, together with target particle counts $N_\alpha = N_\beta = N_{\text{ele}}/2$. This process of configuration recovery ensured a refined bitstring matrix in which configurations inconsistent with the inferred occupancy pattern are replaced by symmetry-recovered alternatives.

At each iteration, the refined bitstring matrix was subjected to Hamming-weight filtering to enforce N_α and N_β electrons in the α - and β -spin sectors (corresponding to the right and left halves of each bitstring, respectively). Configurations were then subsampled based on a per-iteration limit ϵ_s (samples per batch) for a single batch. The subsampling used a per-run cryptographically seeded RNG to ensure statistical independence across hardware runs. The resulting batch was decomposed into separate α - and β -spin CI string lists.

The subspace Hamiltonian was constructed in the CI basis spanned by all α - β string pairs from the subsampled batch and diagonalized using a Davidson iterative eigensolver with a maximum of 200 cycles. The solver returns updated average orbital occupancies, which are stored and used to guide configuration recovery in the subsequent iteration as shown in Algorithm 1.

After 10 iterations, the reported SQD energy is taken to be the minimum over all iterations and their corresponding batches ($n_{\text{batches}} = 1$ throughout) as shown in Algorithm 1. Each of the three independent hardware runs yields one such energy value, and all results were aggregated in the output file for further analysis.




Figure 4. Schematic representation of SQD workflows comparing the (a) LUCJ and (b) LCNot-UCCSD ansätze. The flowchart illustrates the division between Classical Pre-Processing, Quantum Computation, and Classical Post-Processing stages. Green boxes indicate computational advantages (lower algorithmic scaling), while yellow boxes highlight computational trade-offs, demonstrating the inverse relationship between classical initialization costs and quantum circuit complexity in determining the final energy, E_{SQD} .

The following subspace dimension metrics were recorded per run:

- $|\mathbb{S}_{\text{samp}}|$: number of unique bitstrings in raw measurement counts,
- $|\mathbb{S}_{\text{post-cr}}|$: number of configurations after configuration recovery,
- $|\mathbb{S}_{\text{sub}}|$: diagonalization subspace dimension,
- $|\mathbb{S}|$: symmetry-constrained Hilbert space size from Eq. (27),
- $|\mathbb{H}|$: full Hilbert space dimension, $2^{2N_{\text{orb}}}$.

In this work, the relative sizes of these spaces are characterized by the ratios

$$\eta_{\text{sym}} = \frac{|\mathbb{S}|}{|\mathbb{H}|}, \quad \eta_{\text{post-cr}} = \frac{|\mathbb{S}_{\text{post-cr}}|}{|\mathbb{S}|}, \quad \eta_{\text{sub}} = \frac{|\mathbb{S}_{\text{sub}}|}{|\mathbb{S}|}, \quad (28)$$

which quantify the compression achieved at different stages of the SQD workflow. The recovered configurations are then incorporated into the determinant pool prior to classical diagonalization.

Two sampling strategies were employed for ϵ_s . The first sets $\epsilon_s = 10^8$, which greatly exceeds the total number of sampled configurations ($|\mathbb{S}_{\text{samp}}|$ unique bitstrings from 10^4 shots) and effectively imposes no subsampling constraint. The second strategy scales $\epsilon_s = \sqrt{|\mathbb{S}|}$, matching the subsample size to the square root of the symmetry-constrained Hilbert space, effectively imposing a symmetry-adapted subsampling constraint.

For closed-shell systems, it is advantageous to combine the recovered α and β configurations into a single set and construct the diagonalization subspace as the tensor product of this set with itself.¹⁴⁸ Configurations related by spin exchange often contribute with comparable weight to the ground state; that is, if a configuration of the form $\alpha\beta$ is significant, the corresponding $\beta\alpha$ configuration is typically similarly relevant.¹⁴⁸ Consequently, the sets of unique α and β configurations obtained during configuration recovery tend to exhibit substantial overlap. Under this commonly observed structure, the effective number of distinct configurations grows approximately with the square root of the full symmetry-constrained Hilbert space, which provides a practical motivation for the $\sqrt{|\mathbb{S}|}$ subsampling heuristic used in this work. The complete SQD execution flow is detailed in Algorithm 1 for further reference.

Algorithm 1 SQD workflow per molecule, per hardware run

Require: $h_{pq}^{\text{MO}}, h_{pqrs}^{\text{MO}}, E_{\text{nuc}}, \text{ansatz_isa}$ (transpiled ansatz), $N_\alpha, N_\beta, N_{\text{shots}}$ (number of shots), ϵ_s (samples per batch), I_{max} (iterations), M_{dav} (Davidson cycles), n_{batches} (number of batches)

Ensure: E_{SQD}

```

1: Execute ansatz_isa on QPU with  $N_{\text{shots}}$  shots  $\rightarrow$  counts_dict
2:  $(\mathbf{B}, \mathbf{p}) \leftarrow \text{COUNTS\_TO\_ARRAYS}(\text{counts\_dict})$ 
3:  $\bar{n} \leftarrow \text{None}$  ▷ average orbital occupancy
4: for  $i = 0, 1, \dots, I_{\text{max}} - 1$  do
5:   if  $\bar{n}$  is None then
6:      $(\mathbf{B}', \mathbf{p}') \leftarrow (\mathbf{B}, \mathbf{p})$ 
7:   else
8:      $(\mathbf{B}', \mathbf{p}') \leftarrow \text{RECOVER\_CONFIGURATIONS}(\mathbf{B}, \mathbf{p}, \bar{n}, N_\alpha, N_\beta)$ 
9:   end if
10:  batches  $\leftarrow \text{POSTSELECT\_AND\_SUBSAMPLE}(\mathbf{B}', \mathbf{p}', N_\alpha, N_\beta, \epsilon_s, n_{\text{batches}})$ 
11:  for  $j = 0, 1, \dots, n_{\text{batches}} - 1$  do
12:     $(E_j, sv_j, \bar{n}_j, \langle S^2 \rangle) \leftarrow \text{SOLVE\_FERMION}(\text{batches}[j], h_{pq}^{\text{MO}}, h_{pqrs}^{\text{MO}}, \langle S^2 \rangle = 0, M_{\text{dav}})$ 
13:     $E_{\text{history}}[i][j] \leftarrow E_j + E_{\text{nuc}}$ 
14:     $\bar{n}_{\text{history}}[i][j] \leftarrow \bar{n}_j$ 
15:  end for
16:   $\bar{n} \leftarrow \text{mean}(\bar{n}_{\text{history}}[i][:])$ 
17: end for
18: return  $E_{\text{SQD}} \leftarrow \min(E_{\text{history}})$ 

```

Parameters: $N_{\text{shots}} = 10,000$ shots; $I_{\text{max}} = 10$ iterations; $\epsilon_s = 10^8$ or $\sqrt{|\mathbb{S}|}$ samples per batch; $M_{\text{dav}} = 200$ Davidson cycles; $n_{\text{batches}} = 1$.

3.1.6 Reproducibility

A global NumPy random seed was set at the start of each script to ensure reproducibility of deterministic operations. The per-run subsampling seed was drawn from a cryptographically secure source via `secrets`²³⁴ module, ensuring that the three independent runs operate with distinct stochastic realizations of the configuration recovery and subsampling procedure. Timing information was recorded using the `timeit`²³⁴ module for the pre-processing, CCSD, FCI, QPU sampling, configuration recovery, and post-processing stages independently. All results, including job identifiers, circuit resource estimates, space dimension metrics, energy values, and timings, were written incrementally to the output file after each run to ensure no data loss in the event of a runtime failure.

3.2 Potential Energy Surface (PES) Scans Using Sample-based Quantum Diagonalization (SQD)

3.2.1 One-Dimensional PES Implementation

The 1D-PES calculations follow the same SQD(LUCJ) workflow described in Section 3.1, extended by an outer loop over a discrete set of molecular geometries parameterized by the interatomic distance r as shown in Algorithm 2. The molecules studied in the 1D-PES scans are H_2 and HeH^+ (STO-3G and 6-31G), and LiH and BeH_2 (STO-3G), with their coordinate parameterizations illustrated in Figure 5. For all molecules, a non-uniform grid was constructed with closely spaced points covering the bonding region, followed by more widely spaced points extending to capture the dissociation regime. The exact geometries, along with their QPU runtimes, can be found in Appendix D for further reference.




Figure 5. Molecular geometry representations of the one-dimensional potential energy surface (1D-PES) molecules under study: (a) H_2 , (b) HeH^+ , (c) LiH , and (d) BeH_2 . Here, r denotes the varying interatomic distance in ångströms (Å).

For each grid point, the PySCF workflow (Section 3.1.1) was re-executed from scratch: the `gto.Mole` object was rebuilt with the updated geometry, a fresh RHF solution was obtained, one- and two-electron integrals were recomputed in the MO basis, and a new CCSD calculation was run to provide geometry-specific t_1 and t_2 amplitudes for the LUCJ circuit. The LUCJ circuit was then transpiled and submitted to hardware with 10,000 shots across 3 independent runs. The SQD configuration recovery loop ($I_{\text{max}} = 10$, $\epsilon_s = 10^8$ or $\sqrt{|\mathbb{S}|}$, $M_{\text{dav}} = 200$) was executed independently for each run at each geometry. Per-geometry result records additionally store the r coordinate to enable direct PES reconstruction during analysis. All results were appended incrementally to a single output file per molecule, allowing partial recovery if the scan was interrupted.

3.2.2 Two-Dimensional PES Implementation

The 2D-PES scan for water (H_2O , STO-3G) spans a 32×32 grid in O–H bond length r and H–O–H bond angle θ , as illustrated in Figure 6, for a total of 1,024 grid points. The grid geometries were pre-computed and stored as JSON entries, each containing the Cartesian coordinate string, scalar r , and scalar θ . The geometry file was loaded at runtime, and the outer loop iterated over all 1,024 entries. The LUCJ circuit parameters were re-initialized from geometry-specific CCSD amplitudes at every grid point, following the same procedure as described for the 1D benchmark molecules in Section 3.2.1.




Figure 6. Molecular geometry representation of H_2O used for the two-dimensional potential energy surface (2D-PES) scan. Here, r denotes the varying O–H interatomic distance in ångströms (Å), and θ denotes the varying H–O–H bond angle in degrees ($^\circ$).

The 2D-PES protocol differs from the 1D case in two respects. The number of independent hardware runs per grid point was set to 1 rather than 3, and the inter-run sleep interval was set to 0.0 s, reflecting the need to minimize total wait time across 1,024 sequential submissions. The SQD configuration recovery loop parameters, such as $I_{\max}$, ϵ_s , and M_{dav} , were identical to those used in the 1D case, as described in Section 3.2.1. Each result record stores r , θ , the Cartesian geometry string, and the full set of SQD diagnostic metrics, with results appended incrementally to the output file. Details of the construction of the complete set of molecular geometries used in the 2D-PES scan, along with the heatmap of per-geometry QPU runtimes, can be found in Appendix D. The outer loop of the 2D-PES scan is detailed in Algorithm 2.

Algorithm 2 PES scan outer loop (1D shown; 2D is structurally identical with an additional θ coordinate per entry)

Require: Geometry dictionary $\mathcal{G} = \{k : (\text{geom}_k, r_k)\}_{k=1}^K$; SQD parameters: N_{shots} , runs , $I_{\max}$, ϵ_s , M_{dav}

Ensure: $\{E_{\text{SQD}}(r_k)\}_{k=1}^K$

```

1: for each index  $k$  in  $\mathcal{G}$  do
2:    $(\text{geom}_k, r_k) \leftarrow \mathcal{G}[k]$  ▷ also read  $\theta_k$  for 2D
3:    $\text{Build mol} \leftarrow \text{gto.Mole}(\text{geom}_k, \text{basis} = \text{STO-3G})$ 
4:    $\text{mf} \leftarrow \text{scf.RHF}(\text{mol}).\text{run}()$ 
5:   Compute  $h_{pq}^{\text{MO}}, h_{pqrs}^{\text{MO}}, E_{\text{nuc}}$  from  $\text{mf}$ 
6:    $\text{Run cc.CCSD}(\text{mf}) \rightarrow t_1, t_2$ 
7:    $\text{ansatz} \leftarrow \text{LUCJ}(N_{\text{orb}}, N_{\text{ele}}, t_1, t_2)$ 
8:    $\text{ansatz\_isa} \leftarrow \text{transpile}(\text{ansatz}, \text{backend})$ 
9:   for  $\text{run} = 1, \dots, \text{runs}$  do
10:     $\text{Execute ansatz\_isa on QPU} \rightarrow \text{counts}$ 
11:     $\text{Run Algorithm 1} \rightarrow E_{\text{SQD}}$ 
12:     $\text{Append}(k, r_k, [\theta_k,] E_{\text{SQD}}, \text{diagnostics})$  to output file
13:   end for
14: end for
```

Parameters (1D): $K = \{\text{H}_2: 25, \text{HeH}^+: 25, \text{LiH}: 30, \text{BeH}_2: 27\}$; $\text{runs} = 3$; $N_{\text{shots}} = 10,000$.

Parameters (2D): $K = \{\text{H}_2\text{O}: 1,024\}$ (32×32 grid); $\text{runs} = 1$; $N_{\text{shots}} = 10,000$.

3.3 Sample-based Quantum Diagonalization within Density Matrix Embedding Framework (DMET-SQD)

The current work employs the DMET framework in conjunction with an RHF reference^{77,188,189,246}, which is appropriate for the closed-shell systems considered here and provides a consistent mean-field starting point for constructing the bath orbitals. The impurity problems are subsequently solved using the SQD approach, which relies on a selected configuration interaction (SCI) procedure for diagonalization. Although SCI is formally neither size-consistent nor size-extensive, the sampling-driven configuration recovery ensures that the dominant determinants in the $\mathbb{H}$ are preferentially captured. An overall workflow of DMET-SQD is provided in Figure 8. To rigorously benchmark the accuracy of this approach, FCI or Complete Active Space CI (CASCI)²⁴⁷ is employed as a reference solver within DMET, along with DMET-CCSD, to understand the additional correlation captured by employing DMET-SQD(LUCJ). Furthermore, a one-atom-per-fragment partitioning is adopted as a stringent test of the DMET bath construction, including chemically nontrivial scenarios where covalent bonds are severed across fragments. This aggressive fragmentation minimizes the size of individual impurity problems, thereby reducing quantum resource requirements and sampling overhead, while providing a demanding test of whether inter-fragment correlations are consistently recovered.

Building on this framework, a set of eight ligand-like molecules is selected¹⁹⁰, several of which are relevant to pharmaceutical and chemical applications^{248–251}. The set comprises HOCN (cyanic acid), CH_3NO (formaldehyde oxime), CH_5NO (methoxyamine), $\text{C}_2\text{H}_3\text{NO}$ (methyl isocyanate), $\text{C}_2\text{H}_5\text{NO}$ (acetaldehyde oxime), $\text{CH}_4\text{N}_2\text{O}$ (carbamide/urea), NOCl (nitrosyl chloride), and HOSCN (hydroxythiocyanate), and is chosen to systematically evaluate the performance, scalability, and accuracy of the DMET-SQD workflow across chemically diverse yet computationally tractable systems. All systems are treated in the STO-3G basis using a one-atom-per-fragment DMET partitioning. As an illustrative example, Section 2 of Figure 8 shows the fragmentation of the CHNO molecule on an atom-by-atom basis. Panels 2.a, 2.b, 2.c, and 2.d depict the isosurfaces of the fragment orbitals, while panels 2.a', 2.b', 2.c', and 2.d' show the corresponding bath orbitals when the selected fragments are [O], [C], [N], and [H], respectively. The fragment orbitals are the localised orbitals associated with the selected fragment. In contrast, the bath orbitals are constructed from the entanglement between the fragment and its environment, and can be interpreted as linear combinations of environment orbitals that are most strongly coupled to the fragment. As a result, they are typically delocalized over the environment, with minimal weight on the fragment itself.

To remain compatible with the available quantum hardware, impurity systems containing more than 16 spin orbitals were

truncated using an active-space selection of four highest occupied molecular orbitals (4 HOMO) and four lowest unoccupied molecular orbitals (4 LUMO), resulting in impurity problems with $(N_{\text{orb}}, N_{\text{ele}}) \leq (8, 8)$ and at most $N_Q = 16$ qubits. Unless otherwise specified, the SQD sampling experiments were performed using $N_{\text{shots}} = 10,000$ measurements per fragment, while ϵ_s was chosen sufficiently large (10^5) so that it is a non-factor and to retain all recovered configurations without stochastic truncation.

Next, we extend this work to a significantly larger system, the amantadine molecule ($\text{C}_{10}\text{H}_{17}\text{N}$)²⁵², which is an FDA approved²⁵³ antiviral and antiparkinsonian drug, historically used for the treatment of influenza A infections and currently prescribed for Parkinson's disease. Its molecular weight is 151.25 Da, and it contains 11 atoms, which are period 2 elements, making it a sufficiently large system suitable for our work. We retain the same minimal basis set (STO-3G), RHF-based DMET formulation, and SQD impurity solver framework, along with DMET-CASCI and DMET-CCSD benchmarking, but adopt a more flexible fragmentation strategy tailored to the molecular structure as shown in Figure 7. Specifically, the system is partitioned into 11 fragments containing one (F2), two (F5, F6, F8), or three atoms (F1, F3, F4, F7, F9, F10, F11), moving us towards a fragmentation scheme beyond one-atom per fragment, making larger molecular simulations feasible and a tractable extension. This choice enables a more balanced distribution of correlation within each impurity while maintaining hardware-compatible active spaces. All impurity problems are again truncated to at most $(N_{\text{orb}}, N_{\text{ele}}) = (8, 8)$, corresponding to $N_Q = 16$ qubits, and the SQD sampling parameters are kept consistent with the ligand-like molecule study. To investigate the influence of the measurement budget on the final DMET-SQD energy, four independent simulations were performed with shot numbers $N_{\text{shots}} = \{500, 1000, 5000, 10000\}$. We also note that a full-scale simulation of the molecule in the same minimal basis set, without any fragmentation, would require 144 qubits, while in the current fragmentation scheme, the quantum simulation of the largest fragment without any active space selection would still only require 28 qubits.




Figure 7. Fragmentation details of Amantadine Molecule: The molecule was fragmented into 11 fragments, such that in STO-3G basis each fragment contains at most 7 molecular spatial orbitals and at least 5 molecular spin orbitals. Using the active space selection method, we pick out the 4 HOMO - 4 LUMO orbitals in each impurity, leading to 16 molecular spin orbitals or 16 qubits per fragment.

To the best of our knowledge, the application of DMET-based fragmentation to a molecule of this scale on IQM quantum hardware has not been previously demonstrated, thereby providing a concrete step toward scalable quantum simulations of chemically relevant systems.




Figure 8. Overall workflow of the DMET-SQD approach. Panel 1 represents the self-consistent field (SCF) density of the full molecule. Panels 2.a/2.a', 2.b/2.b', 2.c/2.c', and 2.d/2.d' show the fragment orbitals and corresponding bath orbitals for the O, C, N, and H atom fragments, respectively. Panel 3 provides a schematic illustration of how successive updates of the global chemical potential μ_{glob} drive the total electron number N_{tot} towards the target value N_{true} .

4 Results and Discussion

4.1 Sample-based Quantum Diagonalization (SQD)

The SQD framework was applied to five molecular systems of increasing complexity: H_2 , LiH , BeH_2 , H_2O , and NH_3 . The relevant electronic structure properties of each system are summarised in Table 2, including N_{orb} and N_{ele} , which define each molecule’s complete active space, the corresponding $|\mathbb{S}|$ and $|\mathbb{H}|$, and the classical reference energies obtained at the CCSD and FCI levels of theory. The QPU execution times and molecular geometries used throughout this subsection are provided in Appendix C. As the active space grows from H_2 to NH_3 , the $|\mathbb{H}|$ dimension increases rapidly as a function of the number of qubits N_Q , scaling as 2^{N_Q} from $2^4 = 16$ to $2^{16} = 65,536$, while the $|\mathbb{S}|$ to $|\mathbb{H}|$ dimension ratio η_{sym} decreases considerably, ranging from 0.25 for H_2 down to as low as 0.027 for H_2O . This progressive increase in problem complexity provides a meaningful benchmark for evaluating the performance of the quantum ansätze considered in this work, including LUCJ (Section 4.1.1) and LCNot-UCCSD (Section 4.1.2) on IQM Sirius quantum hardware (Appendix B provides further information on the QPU’s calibration and hardware metrics). For each molecule, the FCI energy serves as the exact reference against which all SQD results are assessed, with CCSD energies additionally reported to contextualise the accuracy of the quantum approaches relative to a well-established classical correlated method, as shown in Table 2.

Table 2. Properties of the molecules under study, including the number of molecular orbitals N_{orb} , electrons N_{ele} , the symmetry space dimension $|\mathbb{S}|$, the Hilbert space dimension $|\mathbb{H}|$, the symmetry space dimension to Hilbert space dimension ratio η_{sub} and classical reference ground-state energies obtained using CCSD (E_{CCSD} , Ha) and FCI (E_{FCI} , Ha) for the molecular geometries listed in Appendix C.

Molecule	$(N_{\text{orb}}, N_{\text{ele}})$	$ \mathbb{S} $	$ \mathbb{H} $	η_{sym}	E_{CCSD} (Ha)	E_{FCI} (Ha)
H_2	(2,2)	4	16	0.25	-1.12784260513579	-1.12784250431471
LiH	(6,4)	225	4096	0.0549	-7.88116714390917	-7.88117835166515
BeH_2	(7,6)	1125	16384	0.0748	-15.54741326429858	-15.54779792987434
H_2O	(7,10)	441	16384	0.0269	-75.01541232667518	-75.01553351836277
NH_3	(8,10)	3136	65536	0.0479	-55.52092972113001	-55.52114799754608

4.1.1 Local Unitary Cluster Jastrow (LUCJ)

First, we employ the LUCJ ansatz as the sample-generating circuit within the SQD framework for execution on the IQM Sirius QPU. The corresponding quantum circuit resources for each molecular system are detailed in Table 3. The qubit count scales directly with the active space size, ranging from 4 qubits for H_2 to 16 qubits for NH_3 , and the circuit complexity grows substantially with system size. For NH_3 , the largest system considered, the transpiled circuit comprises 1,898 total gates G_T , including 746 R gates, 450 CZ gates, and 702 *Move* gates, with a circuit depth G_D of 1,371. This represents a considerable resource overhead relative to H_2 , for which the circuit contains 85 G_T at a G_D of 64. The number of variational parameters $|\tilde{\theta}|$ similarly scales from 2 for H_2 to 240 for NH_3 , reflecting the increasing expressibility demanded of the ansatz as the system complexity grows.

The dimensional behaviour of the SQD workflow across successive runs is illustrated in Figure 9 for all five molecules, under two choices of the ϵ_s parameter: a large non-factor value of 10^8 and a molecule-symmetry-adaptive value of $\sqrt{|\mathbb{S}|}$. In the case of the smaller molecules considered, such as H_2 and LiH , the post-configuration-recovery dimension $|\mathbb{S}_{\text{post-cr}}|$ converges rapidly towards the molecule’s symmetry space dimension $|\mathbb{S}|$, whereas for BeH_2 and H_2O we access approximately 0.90 of $|\mathbb{S}|$, and for NH_3 about 0.50. When $\epsilon_s = 10^8$, the diagonalization subspace dimension $|\mathbb{S}_{\text{sub}}|$ consistently spans the entire $|\mathbb{S}|$, achieving a diagonalization subspace dimension to symmetry space dimension ratio $\eta_{\text{sub}} = 1.0$ across all runs and all molecules. In contrast, with $\epsilon_s = \sqrt{|\mathbb{S}|}$, the subspace coverage relative to $|\mathbb{S}|$ is necessarily reduced, with η_{sub} values ranging from approximately 0.38 to 0.75 depending on the molecule and run, reflecting the more compact $|\mathbb{S}_{\text{sub}}|$.

The energetic consequences of these two subsampling regimes are reported in Table 4 and visualised in Figure 10. For $\epsilon_s = 10^8$, the SQD(LUCJ) workflow recovers the FCI ground-state energy to within sub-nanohartree precision for all five molecules across all three independent runs, with no run-to-run variance observed. This is primarily due to $|\mathbb{S}_{\text{sub}}|$ spanning the complete $|\mathbb{S}|$, as reflected by η_{sub} . When ϵ_s is reduced to $\sqrt{|\mathbb{S}|}$, the accuracy is seen to degrade depending on the subspace diagonalized and the system size. For H_2 , $\mathbb{S}$ is small enough that even $\epsilon_s = 2$ suffices to span it completely, and FCI-level accuracy is preserved to numerical precision across all runs. For LiH with $\epsilon_s = 15$, energy deviations from FCI are on the order of 10^{-6} Ha. For BeH_2 , H_2O , and NH_3 , the reduced $|\mathbb{S}_{\text{sub}}|$ leads to a more pronounced degradation, while still providing energetics within chemical accuracy against the FCI benchmark.




Figure 9. Dimensions (log scale) on the y-axis versus runs on the x-axis for SQD(LUCJ) across various molecules. Subplots (a,b) correspond to H_2 , (c,d) to LiH , (e,f) to BeH_2 , (g,h) to H_2O , and (i,j) to NH_3 . The quantity $|\mathbb{H}|$ denotes the Hilbert space dimension, $|\mathbb{S}|$ the symmetry space dimension, $|\mathbb{S}_{\text{samp}}|$ the sampling space dimension, $|\mathbb{S}_{\text{post-cr}}|$ the post-configuration-recovery dimension, $|\mathbb{S}_{\text{sub}}|$ the diagonalization subspace dimension, and ϵ_s the samples per batch, a user-defined SQD parameter. For subplots (a,c,e,g,i), $\epsilon_s = 10^8$ (non-factor), whereas for subplots (b,d,f,h,j), $\epsilon_s = \sqrt{|\mathbb{S}|}$ for the respective molecules.

Table 3. Quantum circuit resource summary for the SQD(LUCJ) workflows for the molecular systems executed on IQM Sirius quantum hardware. Here, N_Q denotes the number of qubits, $|\vec{\theta}|$ the number of ansatz parameters, G_D the circuit depth, R the number of single-qubit rotation gates, CZ the number of two-qubit Controlled-Z gates, $Move$ the number of two-qubit Move gates, G_{q_1} the number of single-qubit gates, G_{q_2} the number of two-qubit gates, and G_T the total number of quantum gates.

Molecule	N_Q	$ \vec{\theta} $	G_D	R	CZ	$Move$	G_{q_1}	G_{q_2}	G_T
H ₂	4	2	64	35	18	32	35	50	85
LiH	12	72	638	358	202	348	358	550	908
BeH ₂	14	156	1025	558	342	524	558	866	1424
H ₂ O	14	110	892	492	270	486	492	756	1248
NH ₃	16	240	1371	746	450	702	746	1152	1898

Table 4. SQD(LUCJ) ground-state energies obtained using different values of ϵ_s (samples per batch). For each molecular system, three independent runs are reported. Here, $\eta_{\text{post-cr}}$ denotes the ratio of the post-configuration-recovery space dimension to the symmetry space dimension, η_{sub} denotes the ratio of the diagonalization subspace dimension to the symmetry space dimension, and E_{SQD} is the final SQD energy reported in Hartree for the corresponding molecule.

Molecule	ϵ_s	Run	$\eta_{\text{post-cr}}$	η_{sub}	E_{SQD} (Ha)
H ₂	10^8	1	1.0000	1.0000	-1.127842504
		2	1.0000	1.0000	-1.127842504
		3	1.0000	1.0000	-1.127842504
	2	1	1.0000	1.0000	-1.127842504
		2	1.0000	1.0000	-1.127842504
		3	1.0000	1.0000	-1.127842504
LiH	10^8	1	1.0000	1.0000	-7.881178352
		2	1.0000	1.0000	-7.881178352
		3	1.0000	1.0000	-7.881178352
	15	1	1.0000	0.7511	-7.881178308
		2	0.9956	0.5378	-7.881171560
		3	1.0000	0.6400	-7.881173091
BeH ₂	10^8	1	0.9127	1.0000	-15.54779793
		2	0.9069	1.0000	-15.54779793
		3	0.9078	1.0000	-15.54779793
	35	1	0.9159	0.6400	-15.54779418
		2	0.9020	0.5951	-15.54765172
		3	0.9159	0.6400	-15.54769945
H ₂ O	10^8	1	0.9229	1.0000	-75.01553352
		2	0.9252	1.0000	-75.01553352
		3	0.8617	1.0000	-75.01553352
	21	1	0.9048	0.3832	-75.01547008
		2	0.9274	0.6553	-75.01542294
		3	0.8707	0.5102	-75.01533280
NH ₃	10^8	1	0.5080	1.0000	-55.52114800
		2	0.4837	1.0000	-55.52114800
		3	0.5124	1.0000	-55.52114800
	56	1	0.5102	0.4605	-55.52067640
		2	0.4786	0.5625	-55.52029544
		3	0.5003	0.5102	-55.52067337




Figure 10. Energy difference between SQD(LUCJ) and FCI, in Hartree (Ha), for three independent runs on the IQM Sirius QPU for each molecule. The y-axis represents the energy difference $\Delta E = |E_{SQD} - E_{FCI}|$ (log-scale), and the x-axis represents the molecules (H_2 , LiH , BeH_2 , H_2O , and NH_3). Subplots (a) and (b) correspond to $\epsilon_s = 10^8$ (non-factor) and $\epsilon_s = \sqrt{|S|}$, respectively, where $|S|$ is the symmetry space dimension of the corresponding molecule.

4.1.2 Linear Scaling CNOT-Unitary Coupled Cluster Singles and Doubles (LCNot-UCCSD)

The LCNot-UCCSD ansatz was similarly employed as the sample-generating circuit within the SQD framework and executed on the IQM Sirius QPU. The quantum circuit resources associated with each molecular system are reported in Table 5. The qubit counts are identical to those used in the LUCJ workflows, as the same geometry and complete active space were employed. The transpiled G_D and G_T are, however, substantially larger across all systems. For instance, H_2 shows a G_T of 162 at a G_D of 117, which already exceeds the corresponding LUCJ values by a factor of roughly two. This disparity becomes far more pronounced for larger systems: for BeH_2 , the LCNot-UCCSD circuit comprises 21,216 G_T at a G_D of 14,964, compared to 1,424 G_T at a G_D of 1,025 for LUCJ, representing a nearly 15-fold increase in G_D . The largest system compiled, NH_3 , requires 33,466 G_T at a G_D of 23,623, which is approximately 17 times deeper than its LUCJ counterpart. This dramatic scaling in circuit complexity reflects the inherently deeper structure of the LCNot-UCCSD ansatz relative to the hardware-efficient LUCJ design and imposes substantially greater demands on IQM Sirius’s noise tolerance. Consequently, it should be noted that the hardware execution of the LCNot-UCCSD circuits resulted in quantum sampling where configuration recovery followed by classical subspace diagonalization was possible for three molecules, H_2 , LiH , and BeH_2 , whereas it failed due to the lack of any samples falling in $\mathbb{S}$ for H_2O and NH_3 . This failure may be attributed to the reduced η_{sym} as well as the increased G_D for these molecules. Detailed analysis of this phenomenon is deferred to future work.

The dimensional convergence behaviour of the SQD(LCNot-UCCSD) workflow is shown in Figure 11 for both values of ϵ_s . A notable distinction relative to the LUCJ results emerges in the post-configuration-recovery coverage. For LiH , $\eta_{\text{post-cr}}$ values lie in the range of approximately 0.71 – 0.74, compared to unity for LUCJ, indicating that the deeper LCNot-UCCSD circuits generate bitstring distributions that cover a relatively smaller fraction of $\mathbb{S}$ following configuration recovery. This effect is considerably more pronounced for BeH_2 , where $\eta_{\text{post-cr}}$ drops to approximately 0.23 – 0.24, in stark contrast to the ~ 0.91 observed for LUCJ. This reduction in post-configuration-recovery coverage is consistent with the expectation that increased G_D leads to greater hardware noise, which in turn distorts the output distribution and reduces the effective overlap of the sampled bitstrings with the ground state. Despite this degraded coverage, the diagonalization subspace under $\epsilon_s = 10^8$ continues to span the full $\mathbb{S}$ ($\eta_{\text{sub}} = 1.0$) for all three molecules across all runs.

The corresponding ground-state energies are presented in Table 6 and Figure 12. Under $\epsilon_s = 10^8$, the SQD(LCNot-UCCSD) workflow recovers the FCI energy to within sub-nanohartree accuracy for all three molecules across all three independent runs, mirroring the performance observed for LUCJ in the same regime. This confirms that, provided $\mathbb{S}_{\text{sub}}$ fully spans $\mathbb{S}$, the SQD procedure is robust to the increased noise levels introduced by the deeper LCNot-UCCSD circuits, effectively diagonalizing over the complete $\mathbb{S}$. In the symmetry-adaptive subsampling regime with $\epsilon_s = \sqrt{|\mathbb{S}|}$, the energy accuracy degrades in a manner broadly consistent with the LUCJ trends, though the absolute deviations from FCI are relatively larger. For LiH , the three independent runs yield energy deviations on the order of 10^{-5} Ha and 10^{-6} Ha, while for BeH_2 the deviations are bounded between 10^{-4} Ha and 10^{-5} Ha, slightly higher when compared to their SQD(LUCJ) counterparts. This marginal degradation may be attributed to the compounding effect of reduced $\eta_{\text{post-cr}}$ coverage and smaller η_{sub} ratios, both of which are exacerbated by the higher noise levels inherent to the deeper LCNot-UCCSD circuits. Nonetheless, even in this constrained setting, the SQD(LCNot-UCCSD) energies for the molecules H_2 , LiH , and BeH_2 remain within chemical accuracy when compared against the FCI reference, demonstrating that the method retains meaningful accuracy despite the significant circuit overhead, primarily due to the configuration recovery error mitigation process paired with the SCI proliferation performed on $\mathbb{S}_{\text{post-cr}}$ to obtain $\mathbb{S}_{\text{sub}}$ ¹⁹⁰. Further analysis of hardware noise, η_{sym} , and their corresponding effects on $\mathbb{S}_{\text{sub}}$ not falling within $\mathbb{S}$, causing failure of configuration recovery for molecules H_2O and NH_3 , is deferred to future work.

Table 5. Quantum circuit resource summary for the SQD(LCNot-UCCSD) workflows for the molecular systems executed on IQM Sirius quantum hardware. Here, N_Q denotes the number of qubits, $|\vec{\theta}|$ the number of ansatz parameters, G_D the circuit depth, R the number of single-qubit rotation gates, CZ the number of two-qubit Controlled-Z gates, $Move$ the number of two-qubit Move gates, G_{q_1} the number of single-qubit gates, G_{q_2} the number of two-qubit gates, and G_T the total number of quantum gates.

Molecule	N_Q	$ \vec{\theta} $	G_D	R	CZ	$Move$	G_{q_1}	G_{q_2}	G_T
H_2	4	3	117	82	34	46	82	80	162
LiH	12	92	6389	4642	1928	2486	4642	4414	9056
BeH_2	14	204	14964	10868	4512	5836	10868	10348	21216
H_2O	14	140	9991	7238	3046	3856	7238	6902	14140
NH_3	16	315	23623	17128	7140	9198	17128	16338	33466




Figure 11. Dimensions (log scale) on the y-axis versus runs on the x-axis for SQD(LCNot-UCCSD) across various molecules. Subplots (a,b) correspond to H_2 , (c,d) to LiH , (e,f) to BeH_2 , (g,h) to H_2O , and (i,j) to NH_3 . The quantity $|\mathbb{H}|$ denotes the Hilbert space dimension, $|\mathbb{S}|$ the symmetry space dimension, $|\mathbb{S}_{\text{samp}}|$ the sampling space dimension, $|\mathbb{S}_{\text{post-cr}}|$ the post-configuration-recovery dimension, $|\mathbb{S}_{\text{sub}}|$ the diagonalization subspace dimension, and ϵ_s the samples per batch, a user-defined SQD parameter. For subplots (a,c,e,g,i), $\epsilon_s = 10^8$ (non-factor), whereas for subplots (b,d,f,h,j), $\epsilon_s = \sqrt{|\mathbb{S}|}$ for the respective molecules.

Table 6. SQD(LCNot-UCCSD) ground-state energies obtained using different values of ϵ_s (samples per batch). For each molecular system, three independent runs are reported. Here, $\eta_{\text{post-cr}}$ denotes the ratio of the post-configuration-recovery space dimension to the symmetry space dimension, η_{sub} denotes the ratio of the diagonalization subspace dimension to the symmetry space dimension, and E_{SQD} is the final SQD energy reported in Hartree for the corresponding molecule.

Molecule	ϵ_s	Run	$\eta_{\text{post-cr}}$	η_{sub}	E_{SQD} (Ha)
H_2	10^8	1	1.0000	1.0000	-1.127842504
		2	1.0000	1.0000	-1.127842504
		3	1.0000	1.0000	-1.127842504
	2	1	1.0000	1.0000	-1.127842504
		2	1.0000	1.0000	-1.127842504
		3	1.0000	1.0000	-1.127842504
LiH	10^8	1	0.7111	1.0000	-7.881178352
		2	0.7378	1.0000	-7.881178352
		3	0.7156	1.0000	-7.881178352
	15	1	0.7244	0.7511	-7.881177519
		2	0.7378	0.4444	-7.881167474
		3	0.7200	0.4444	-7.881172350
BeH_2	10^8	1	0.2335	1.0000	-15.54779793
		2	0.2359	1.0000	-15.54779793
		3	0.2278	1.0000	-15.54779793
	35	1	0.2269	0.5102	-15.54758854
		2	0.2294	0.5951	-15.54774085
		3	0.2294	0.5102	-15.54778165




Figure 12. Energy difference between SQD(LCNot-UCCSD) and FCI, in Hartree (Ha), for three independent runs on the IQM Sirius QPU for each molecule. The y-axis represents the energy difference $\Delta E = |E_{SQD} - E_{FCI}|$ (log-scale), and the x-axis represents the molecules (H_2 , LiH , BeH_2). Subplots (a) and (b) correspond to $\epsilon_s = 10^8$ (non-factor) and $\epsilon_s = \sqrt{|S|}$, respectively, where $|S|$ is the symmetry space dimension of the corresponding molecule.

4.1.3 Ansätze Comparison: SQD(LUCJ) vs SQD(LCNot-UCCSD)

Having characterised the two ansätze individually, it is instructive to compare them directly with respect to both circuit resource requirements and energetic accuracy. Figure 13 summarises the scaling of the number of variational parameters $|\vec{\theta}|$ and the transpiled G_D for SQD(LUCJ) and SQD(LCNot-UCCSD) across all molecular systems. In terms of parameter count, the two ansätze are broadly comparable, with LCNot-UCCSD carrying a modestly larger number of parameters at each system size, 315 versus 240 for NH_3 , for instance. The divergence in G_D , however, is far more dramatic. As illustrated in Figure 13(b), the G_D of LCNot-UCCSD exceeds that of LUCJ by over an order of magnitude for all molecules beyond H_2 , with the gap widening with system size. This discrepancy reflects a fundamental difference in the structural philosophy of the two approaches: LUCJ is designed as a hardware-efficient, chemistry-inspired ansatz whose gate count scales favourably with the number of qubits while having a classical CCSD-level pre-computation overhead for parameter initialization, whereas LCNot-UCCSD faithfully encodes the chemically motivated UCCSD excitation structure with a linearly scaling CNOT gate count, with the advantage of MP2-level classical overhead for parameter initialization, but still inherently demands a significantly deeper circuit compared to LUCJ.




Figure 13. Quantum circuit resource comparison metrics for SQD(LUCJ) and SQD(LCNot-UCCSD) ansatz implementations across molecules. (a) Ansatz parameters $|\vec{\theta}|$ (log scale) on the y-axis versus molecules on the x-axis, and (b) quantum circuit depth G_D (log scale) on the y-axis versus molecules on the x-axis.

The energetic comparison between the two ansätze is presented in Figure 14 in terms of $\Delta E = |E_{\text{SQD}} - E_{\text{FCI}}|$ for both subsampling regimes. For $\epsilon_s = 10^8$, both ansätze achieve essentially identical accuracy across all commonly evaluated molecules, with energy deviations from FCI falling below the sub-nanohartree threshold in every case. This convergence of accuracy underscores a key feature of the SQD framework: provided that $\mathbb{S}_{\text{sub}}$ spans most of $\mathbb{S}$, the quality of the final energy is governed primarily by the SCI proliferation and the subsequent subspace diagonalization rather than the specific distribution generated by the ansatz. As a result, even the considerably noisier output of the deeper LCNot-UCCSD circuits is sufficient to recover FCI-level energies.




Figure 14. Comparison of the energy difference $\Delta E = |E_{\text{SQD}} - E_{\text{FCI}}|$ (Ha) for SQD(LUCJ) and SQD(LCNot-UCCSD) ansatz implementations across molecules. Here, the x-axis represents ΔE (log scale) and the y-axis represents the molecules for (a) $\epsilon_s = 10^8$ and (b) $\epsilon_s = \sqrt{|\mathbb{S}|}$ for the respective molecule.

The picture changes meaningfully under the reduced symmetry-adaptive subsampling regime of $\epsilon_s = \sqrt{|\mathbb{S}|}$, as shown in

Figure 14(b). Here, SQD(LUCJ) consistently achieves equal or lower energy deviations than SQD(LCNot-UCCSD) for the H_2 , LiH, and BeH_2 molecules, where both are evaluated. This performance gap may be understood in terms of the interplay between circuit execution noise and η_{sub} : the shallower LUCJ circuits produce cleaner bitstring distributions with higher $\eta_{\text{post-cr}}$, which in turn enables a more representative $\mathbb{S}_{\text{sub}}$ even at reduced subsampling budgets. Conversely, the deeper LCNot-UCCSD circuits suffer from greater hardware-induced noise, which degrades the quality of quantum sampling and the symmetry compliance of the sampled configurations, ultimately limiting the overlap of $\mathbb{S}_{\text{sub}}$ with the ground state when ϵ_s is symmetry-adapted to $\sqrt{|\mathbb{S}|}$.

As mentioned earlier in Section 4.1.2, for the molecules H_2O and NH_3 , SQD(LCNot-UCCSD) fails because quantum sampling does not produce samples that lie within $\mathbb{S}$ of the molecule to perform configuration recovery; hence, they are marked as not available in Figure 14.

Taken together, these results point to a clear trade-off between the two ansätze within the SQD framework. SQD(LUCJ) offers a compelling combination of hardware efficiency and robust accuracy, particularly under constrained subsampling conditions, but it requires an expensive CCSD-level parameter initialization, making it well-suited to near-term QPU execution. SQD(LCNot-UCCSD), while chemically expressive and comparable in accuracy to LUCJ when subsampling is unrestricted, incurs a substantially higher circuit cost with less costly MP2-level parameter initialization, but this limits its practical applicability to smaller systems on current hardware. These considerations motivate the use of SQD(LUCJ) as the primary ansatz for the remainder of this work.

4.2 One Dimensional Potential Energy Surface (1D-PES) Scans

Having established the accuracy and hardware feasibility of the SQD(LUCJ) framework at fixed molecular geometries, we now extend the assessment to one-dimensional potential energy surface (1D-PES) scans, in which the interatomic distance r is varied continuously along the dissociation pathway. This represents a more stringent test of the method, as the electronic structure and, consequently, the optimal ansatz parameterisation change significantly along the scan coordinate, particularly in the strongly correlated regime near bond dissociation². The 1D-PES scans were performed for four molecular systems, H_2 , HeH^+ , LiH, and BeH_2 , and are reported for two basis sets, STO-3G and 6-31G, in the following subsections. All SQD(LUCJ) energies are assessed against FCI as the exact reference, with CCSD additionally included to benchmark the regime in which single-reference classical methods are known to fail²⁵⁴. The QPU execution times and molecular geometries used throughout this subsection are provided in Appendix D.

4.2.1 Basis Set: STO-3G

The quantum circuit resource counts as a function of interatomic distance r for all four molecules in the STO-3G basis are shown in Figure 15. For H_2 , the compiled circuit resources are highly uniform across the entire dissociation range of 0.3–4.3 Å, with $R \approx 35$, $CZ \approx 18$, $Move \approx 32$, $G_T \approx 85$, and $G_D \approx 65$ at all geometries. The HeH^+ circuit exhibits comparable resource counts at short to intermediate bond lengths but undergoes a discrete step-wise reduction at $r \approx 2.5$ Å, where R drops from approximately 36 to 29, CZ from 18 to 14, and the total gate count falls from ~ 85 to ~ 67 , with a corresponding reduction in G_D . This geometry-localised reduction reflects a change in the orbital structure at stretched bond lengths that may enable a more compact ansatz compilation and is a notable feature not observed for H_2 . For LiH and BeH_2 , the circuit resources are substantially larger, with G_T values of 812–984 and 962–1,157, respectively, and circuit depths in the range of 564–702 for LiH and 670–822 for BeH_2 . They exhibit modestly larger fluctuations at short bond lengths, consistent with the fixed-geometry results of Section 4.1.

The 1D-PES results for H_2 in the STO-3G basis are presented in Figure 16. The SQD(LUCJ) PES overlaps the FCI curve with perfect visual agreement across the full dissociation pathway under both subsampling conditions, spanning bond lengths from 0.3 to 4.3 Å. The $\Delta E_{\text{FCI-SQD}}$ deviations shown in subplot (c), under $\epsilon_s = 10^8$, lie at the level of numerical precision throughout, falling between 10^{-16} and 10^{-14} Ha, well below any physically meaningful threshold. By contrast, $\Delta E_{\text{CCSD-SQD}}$, which in this case reflects the CCSD method’s own error relative to FCI rather than any deficiency of SQD, sits in the range of $\sim 10^{-8}$ Ha near equilibrium but is itself below chemical accuracy everywhere, consistent with the near-exactness of CCSD for this two-electron system. Under $\epsilon_s = 2$, shown in subplot (d), the $\Delta E_{\text{FCI-SQD}}$ deviations remain equally flat and at the same numerical precision floor, and $\eta_{\text{sub}} = 1.0$ is maintained throughout the scan in both regimes, as shown in subplots (e) and (f). The complete invariance of H_2 with respect to the subsampling budget reflects the trivially small $\mathbb{S}$ of this system, which is fully spanned regardless of the geometry of the molecule.

The corresponding results for HeH^+ in STO-3G are shown in Figure 17. The behaviour closely mirrors that of H_2 : in both subsampling regimes, SQD(LUCJ) recovers the FCI PES with deviations at the numerical precision floor ($\sim 10^{-15}$ Ha) throughout the scan, and $\eta_{\text{sub}} = 1.0$ is maintained uniformly at all geometries. The CCSD deviations from SQD remain in the range of $\sim 10^{-7}$ – 10^{-8} Ha, reflecting the intrinsic CCSD error for this two-electron system rather than any SQD inaccuracy. The PES is well resolved across the equilibrium region near $r \approx 0.8$ Å and the full dissociation plateau, confirming that SQD(LUCJ) captures both regimes faithfully. Notably, the discrete reduction in circuit resources observed at $r \approx 2.5$ Å for HeH^+ has no energetic consequence, as the accuracy of the method is unaffected by the geometry-localised transpilation.




Figure 15. Quantum circuit resource counts on the y-axis versus interatomic distance r (Å) on the x-axis for SQD(LUCJ) across various molecules in STO-3G basis. Subplots (a,b) correspond to H_2 , (c,d) to HeH^+ , (e,f) to LiH , and (g,h) to BeH_2 . Here, R is the number of single-qubit rotation gates, CZ is the number of two-qubit Controlled-Z gates, Move is the number of two-qubit Move gates, G_T is the total number of quantum gates, G_D is the circuit depth, G_{q_1} is the number of single-qubit gates, and G_{q_2} is the number of two-qubit gates.




Figure 16. One-dimensional potential energy surface (1D-PES) results for the H_2 molecule in STO-3G basis. Subplots (a,b) show the PES as a function of interatomic distance r (Å), with energy reported in Hartree. Subplots (c,d) present the energy deviation ΔE versus r , while subplots (e,f) display the ratio of the diagonalization subspace dimension to the symmetry space dimension, η_{sub} , as a function of r . Results correspond to $\epsilon_s = 10^8$ in (a,c,e) and $\epsilon_s = 2$ in (b,d,f).




Figure 17. One-dimensional potential energy surface (1D-PES) results for the HeH^+ molecule in STO-3G basis. Subplots (a,b) show the PES as a function of interatomic distance r (Å), with energy reported in Hartree. Subplots (c,d) present the energy deviation ΔE versus r , while subplots (e,f) display the ratio of the diagonalization subspace dimension to the symmetry space dimension, η_{sub} , as a function of r . Results correspond to $\epsilon_s = 10^8$ in (a,c,e) and $\epsilon_s = 2$ in (b,d,f).




Figure 18. One-dimensional potential energy surface (1D-PES) results for the LiH molecule in STO-3G basis. Subplots (a,b) show the PES as a function of interatomic distance r (Å), with energy reported in Hartree. Subplots (c,d) present the energy deviation ΔE versus r , while subplots (e,f) display the ratio of the diagonalization subspace dimension to the symmetry space dimension, η_{sub} , as a function of r . Results correspond to $\epsilon_s = 10^8$ in (a,c,e) and $\epsilon_s = 15$ in (b,d,f).




Figure 19. One-dimensional potential energy surface (1D-PES) results for the BeH₂ molecule in STO-3G basis. Subplots (a,b) show the PES as a function of interatomic distance r (Å), with energy reported in Hartree. Subplots (c,d) present the energy deviation ΔE versus r , while subplots (e,f) display the ratio of the diagonalization subspace dimension to the symmetry space dimension, η_{sub} , as a function of r . Results correspond to $\epsilon_s = 10^8$ in (a,c,e) and $\epsilon_s = 35$ in (b,d,f).

For LiH, shown in Figure 18, the picture is broadly consistent with the fixed-geometry findings but reveals more physically interesting structure. Under $\epsilon_s = 10^8$, the FCI PES is reproduced with $\Delta E_{\text{FCI-SQD}} \approx 10^{-12}$ Ha uniformly across all geometries, and $\eta_{\text{sub}} = 1.0$ throughout. The CCSD deviation profile shown in subplot (c) is particularly instructive: $\Delta E_{\text{CCSD-SQD}}$ grows monotonically from $\sim 10^{-5}$ Ha near the equilibrium geometry to $\sim 10^{-4}$ Ha at stretched bond lengths beyond $r \approx 3$ Å. This behaviour highlights the well-known breakdown of the single-reference CCSD approximation in the dissociation regime, which SQD(LUCJ) avoids entirely by diagonalizing the full $\mathbb{S}$. Under $\epsilon_s = 15$, the subspace coverage η_{sub} shown in subplot (f) fluctuates between approximately 0.50 and 0.85 across the scan, with considerable run-to-run variability and without any strong systematic trend along r . The corresponding $\Delta E_{\text{FCI-SQD}}$ values in subplot (d) range from $\sim 10^{-6}$ to $\sim 10^{-4}$ Ha, correlated with the variations in η_{sub} , yet remain below the chemical accuracy threshold across the entire dissociation pathway. The overall shape of the PES is faithfully reproduced even under these constrained subsampling conditions, and the SQD(LUCJ) energies consistently undercut the CCSD reference at all stretched geometries.

The BeH₂ results in Figure 19 present the most demanding test within the STO-3G set. The equilibrium geometry is located near $r \approx 1.3$ Å at an energy of approximately -15.6 Ha, with the PES rising steeply at compressed geometries and plateauing near -15.30 Ha in the dissociation limit. Under $\epsilon_s = 10^8$, $\Delta E_{\text{FCI-SQD}}$ remains close to the numerical precision floor around $\sim 10^{-11}$ – 10^{-12} Ha throughout the full scan, with $\eta_{\text{sub}} = 1.0$ at all geometries. The CCSD deviation profile is strikingly different: $\Delta E_{\text{CCSD-SQD}}$ grows from $\sim 10^{-4}$ Ha near equilibrium and exceeds the chemical accuracy threshold at stretched geometries beyond $r \approx 2.0$ Å, reaching values approaching 10^{-2} Ha at $r \approx 2.5$ Å. This pronounced CCSD failure in the multireference dissociation regime is cleanly captured and corrected by SQD(LUCJ), which maintains sub-nanohartree accuracy throughout as it diagonalizes over a subspace that spans the complete $\mathbb{S}$. Under $\epsilon_s = 35$, the subspace coverage shown in subplot (f) fluctuates between approximately 0.45 and 0.70, with large run-to-run variability across the scan. The associated $\Delta E_{\text{FCI-SQD}}$ values in subplot (d) are correspondingly larger, ranging from $\sim 10^{-6}$ to $\sim 10^{-3}$ Ha and approaching the chemical accuracy threshold at certain geometries. The LUCJ ansatz initialization with CCSD parameters, together with the reduced η_{sub} , leads to energies that fall outside chemical accuracy at the exact geometries where CCSD itself fails, primarily due to the strongly multireference character of BeH₂ at stretched bond lengths. A deeper analysis of these effects, including improved multireference ansatz parameter initialization strategies and approaches to reduce the required subspace while maintaining chemical accuracy, is deferred to future work. Despite the reduced subspace coverage, the qualitative shape of the dissociation curve is faithfully reproduced, and the SQD(LUCJ) energies continue to mostly outperform CCSD at stretched geometries where the latter deteriorates. These results collectively demonstrate that the SQD(LUCJ) framework is capable of generating accurate 1D-PES scans in the STO-3G basis, with performance that improves systematically as η_{sub} increases toward unity.

4.2.2 Basis Set: 6-31G

The 1D-PES scans were repeated for H₂ and HeH⁺ using the 6-31G basis set, which provides a more flexible description of the electronic structure relative to the minimal STO-3G basis by introducing additional contracted Gaussian functions per atomic shell. This expansion of the orbital basis increases the active space and hence the quantum circuit complexity while offering a more accurate representation of the PES.

The quantum circuit resource counts for H₂ and HeH⁺ in the 6-31G basis are shown in Figure 20. The expanded orbital basis roughly doubles the qubit count for each system relative to STO-3G, and the corresponding gate counts and circuit depths increase accordingly. For H₂ in 6-31G, the compiled circuits exhibit $R = 94$ – 141 , $CZ \approx 70$, $G_T \approx 330$, and $G_D \approx 233$, representing an approximately four-fold increase in G_T and a nearly four-fold increase in G_D relative to the STO-3G case. A brief transient elevation in resource counts is visible at the shortest bond lengths in subplot (a), after which the circuit parameters stabilize across the remainder of the scan. For HeH⁺ in 6-31G, the resources are comparably stable throughout the full scan range of 0.45 – 4.0 Å, with $R = 146$ – 157 , $CZ = 80$, $G_T = 366$ – 389 , and $G_D = 258$ – 278 , and no geometry-localized reduction analogous to that seen in the STO-3G case. This stability across the full scan confirms that the increased resource overhead relative to STO-3G is driven entirely by the larger active space and not by any geometry-induced structural change in the ansatz.

The 1D-PES results for H₂ in the 6-31G basis are presented in Figure 21. Under $\epsilon_s = 10^8$, the SQD(LUCJ) method faithfully reproduces the FCI PES across the full range of interatomic distances, with $\Delta E_{\text{FCI-SQD}}$ values on average at the numerical precision floor of $\sim 10^{-15}$ Ha throughout, and $\eta_{\text{sub}} = 1.0$ uniformly maintained. The CCSD deviations, which fall in the range $\sim 10^{-7}$ – 10^{-10} Ha, remain well below chemical accuracy with sub-microhartree precision, consistent with the essentially exact nature of CCSD for two electrons. In the reduced symmetry-adapted subsampling regime with $\epsilon_s = 2$, the energy deviations in subplot (d) are equally negligible, preserving the same numerical precision floor on average as in the large-subsampling case. However, the η_{sub} profile in subplot (f) reveals a notable feature not observed in the STO-3G counterpart: at most bond lengths r near equilibrium and through the dissociation plateau, η_{sub} drops appreciably from unity, with individual runs showing values as low as ~ 0.55 .

The HeH⁺ results in the 6-31G basis are shown in Figure 22. Under $\epsilon_s = 10^8$, the performance mirrors the H₂ 6-31G case closely, with $\Delta E_{\text{FCI-SQD}}$ values at the numerical precision floor ($\sim 10^{-13}$ – 10^{-15} Ha) across the entire scan and $\eta_{\text{sub}} = 1.0$ at all geometries. Under $\epsilon_s = 2$, the HeH⁺ 6-31G system shows a similar η_{sub} profile to that of H₂ 6-31G: the subspace coverage

drops from unity at most geometries, reaching values as low as ~ 0.55 . Despite the reduced η_{sub} , the energy deviations shown in subplot (d) remain at the same numerical precision floor as in the large-subsampling regime without exception. Additionally, the 6-31G PES for HeH^+ shows a well-defined equilibrium near $r \approx 0.77 \text{ \AA}$ at an energy of approximately -2.93 Ha , deepening the potential well relative to the STO-3G description, while the dissociation plateau converges to approximately -2.87 Ha . The robustness of SQD(LUCJ) across both subsampling conditions and both molecules in the 6-31G basis provides a strong foundation for the basis set comparison in the following subsection.




Figure 20. Quantum circuit resource counts on the y-axis versus interatomic distance r (Å) on the x-axis for SQD(LUCJ) across various molecules in 6-31G basis. Subplots (a,b) correspond to H_2 , and (c,d) to HeH^+ . Here, R is the number of single-qubit rotation gates, CZ is the number of two-qubit Controlled-Z gates, Move is the number of two-qubit Move gates, G_T is the total number of quantum gates, G_D is the circuit depth, G_{q1} is the number of single-qubit gates, and G_{q2} is the number of two-qubit gates.




Figure 21. One-dimensional potential energy surface (1D-PES) results for the H_2 molecule in 6-31G basis. Subplots (a,b) show the PES as a function of interatomic distance r (Å), with energy reported in Hartree. Subplots (c,d) present the energy deviation ΔE versus r , while subplots (e,f) display the ratio of the diagonalization subspace dimension to the symmetry space dimension, η_{sub} , as a function of r . Results correspond to $\epsilon_s = 10^8$ in (a,c,e) and $\epsilon_s = 2$ in (b,d,f).




Figure 22. One-dimensional potential energy surface (1D-PES) results for the HeH^+ molecule in 6-31G basis. Subplots (a,b) show the PES as a function of interatomic distance r (Å), with energy reported in Hartree. Subplots (c,d) present the energy deviation ΔE versus r , while subplots (e,f) display the ratio of the diagonalization subspace dimension to the symmetry space dimension, η_{sub} , as a function of r . Results correspond to $\epsilon_s = 10^8$ in (a,c,e) and $\epsilon_s = 2$ in (b,d,f).

4.2.3 Basis Set Comparison: STO-3G vs 6-31G

Figure 23 places the STO-3G and 6-31G potential energy surfaces in direct comparison for both H_2 and HeH^+ under both subsampling conditions. The figure displays only SQD energies, which, as established in the preceding subsections, coincide closely with the FCI reference to numerical precision in all cases. Therefore, the comparison is directly between the two basis set descriptions for the SQD(LUCJ) method.




Figure 23. One-dimensional potential energy surface (1D-PES) basis set comparison for H_2 and HeH^+ molecules. Subplots (a,b) show the PES as a function of interatomic distance r (Å), with energy reported in Hartree for the H_2 molecule. Similarly, subplots (c,d) present the PES versus r for the HeH^+ molecule. Results correspond to $\epsilon_s = 10^8$ in (a,c) and $\epsilon_s = 2$ in (b,d).

For H_2 , the STO-3G and 6-31G PES curves converge to the same minimum-energy geometry near $r \approx 0.74$ Å but differ in their absolute energies throughout: the 6-31G minimum lies approximately 0.014 Ha below the STO-3G minimum, and the two dissociation plateaus differ by a similar margin, with STO-3G approaching -0.93 Ha and 6-31G approaching -1.00 Ha. For HeH^+ , the energy separation is more pronounced: the 6-31G PES lies approximately 0.06–0.07 Ha below the STO-3G curve throughout the scan, with the 6-31G minimum at ~ 0.78 Å and -2.93 Ha compared to ~ 0.94 Å and -2.86 Ha for STO-3G. In both cases, the larger basis captures a deeper and more accurately located potential well, reflecting the additional variational flexibility afforded by the larger orbital basis of the 6-31G basis set.

Crucially, across all four panels of Figure 23, the SQD(LUCJ) energies in both basis sets are reproduced with equally tight run-to-run consistency, and the error bars on the mean are negligibly small in all cases. The comparison under $\epsilon_s = 2$ in subplots (b) and (d) is equally clean, with no visible degradation in the agreement between basis sets or in the internal self-consistency of the SQD results. This confirms that the reduction in subsampling budget has no differential impact on the accuracy of either basis-set description for these two-electron systems, and that the choice of basis set influences the physical content of the PES—the absolute energies, the depth and location of the potential well, and the dissociation asymptote—but not the reliability of the SQD(LUCJ) method in reproducing it. The superior energetic description offered by 6-31G, combined with the demonstrated ability of SQD(LUCJ) to maintain FCI-level accuracy within it at no additional cost in terms of subsampling

requirements, reinforces the case for employing larger basis sets in future quantum simulation efforts where chemical accuracy demands it, provided that the associated circuit overhead remains within the operationally acceptable noise envelope of the available quantum hardware.

4.3 Two Dimensional Potential Energy Surface (2D-PES) Scans

Building on the one-dimensional scans, the SQD(LUCJ) framework was further applied to map a 2D-PES for the H_2O molecule in the STO-3G basis, in which both the O–H bond length r and the H–O–H bond angle θ are varied simultaneously over a 32×32 grid with $r \in [0.85, 1.20]$ Å and $\theta \in [85^\circ, 115^\circ]$. This constitutes a considerably more demanding application of the method, as the electronic structure must be accurately resolved across a continuous two-dimensional grid of geometries spanning both stretching and bending degrees of freedom, a setting highly relevant to the practical use of quantum chemistry methods for characterizing vibrational frequencies, molecular force fields and reaction pathways^{255–257}. The scan was performed under both $\epsilon_s = 10^8$ and $\epsilon_s = 21$, and assessed against CCSD and FCI reference energies computed at every grid point. The QPU execution time heatmap and the methodology employed to generate the molecular geometries used throughout this subsection are detailed in Appendix D.

The quantum circuit resources associated with the 2D-PES scan are shown as heatmaps in Figure 24. The gate counts and circuit depth vary in a structured yet moderate manner across the two-dimensional grid. The CZ gate count, shown in subplot (a), ranges from approximately 242 to 262. The *Move* gate count in subplot (b) ranges from approximately 408 to 464 and exhibits a similar but smoother gradient. The R/G_{q_1} and G_{q_2} counts in subplots (c) and (d) span approximately 432–474 and 650–720, respectively, with more diffuse spatial structure across the grid. The circuit depth G_D in subplot (e) ranges from approximately 750 to 840. The total gate count G_T in subplot (f) spans approximately 1,080 to 1,200, with a corresponding spatial pattern. A characteristic feature across all resources is a subtle diagonal boundary toward the large- r , large- θ region in the top-right corner, where the quantum gate counts are relatively lower compared to the rest of the grid, as evident in all six subplots of Figure 24. Nevertheless, the overall range of variation remains modest relative to the mean, confirming that the LUCJ ansatz can be transpiled and executed consistently across the full grid, with no geometry posing a disproportionately large hardware burden.

The 2D-PES is presented as three-dimensional surface plots in Figure 25 and as contour plots in Figure 26, for both $\epsilon_s = 10^8$ and $\epsilon_s = 21$. Under $\epsilon_s = 10^8$, the SQD(LUCJ) method produces a smooth, well-resolved potential energy surface across the entire (θ, r) grid. The energy ranges from approximately -74.94 Ha at the least stretched and bent geometries down to a well-defined global minimum of approximately -75.02 Ha. The contour plot in Figure 26(a) reveals a well-structured energy landscape with smooth, nearly concentric contour lines surrounding the minimum. The innermost contours at 0.22 and 1.59 mHa above the minimum are tightly clustered, indicating a relatively flat basin, while the outer contour levels at 3.75, 6.02, 8.87, 11.97, 15.72, 20.52, 25.69, 31.68, 38.53, 46.31, 55.07, and 64.49 mHa increase in spacing as the energy rises steeply toward the edges of the grid. The obtained surface is notably asymmetric: it rises more steeply along the r axis than along the θ axis, reflecting the greater stiffness of the O–H stretching mode relative to the bending mode. The minimum-energy configuration of the H_2O molecule in the STO-3G basis, indicated by the red triangle, is located at $r_{\min} = 1.03$ Å and $\theta_{\min} = 96.61^\circ$.

Under the reduced-subsampling condition of $\epsilon_s = 21$, shown in Figure 25(b) and Figure 26(b), the 3D surface retains its qualitative bowl shape but exhibits visible roughness, particularly near the minimum and at larger bond lengths. In the contour plot, the outer contour lines (> 15 mHa) remain relatively smooth and well resolved, while the innermost contours at 0.22 and 1.59 mHa are noticeably jagged and irregular, indicating that the fine structure of the energy landscape near the minimum is sensitive to the reduced subsampling budget. The minimum under $\epsilon_s = 21$ is displaced slightly to $\theta_{\min} = 97.58^\circ$ at the same bond length of $r_{\min} = 1.03$ Å, corresponding to a shift of approximately 1° relative to the FCI minimum. This deviation arises from stochastic variation in η_{sub} across neighbouring grid points rather than from a systematic bias. Another notable feature is that the practical chemical accuracy bound is represented by the 1.59 mHa contour in Figure 26, indicating that all geometries within this contour lie within chemical accuracy (1 kcal/mol) of the global minimum energy.




Figure 24. Quantum circuit resource summary for a two-dimensional potential energy surface (PES) scan of the H_2O molecule in the STO-3G basis. All heatmaps are plotted as a function of the H–O–H bond angle θ ($^\circ$) along the x-axis and the O–H bond length r ($\text{\AA}$) along the y-axis. Subplots show: (a) the number of two-qubit controlled-Z (CZ) gates; (b) the number of two-qubit Move gates; (c) the number of rotation (R) gates, equivalent to the total number of one-qubit gates (G_{q_1}); (d) the total number of two-qubit gates (G_{q_2}); (e) the quantum circuit depth (G_D); and (f) the total number of quantum gates in the circuit (G_T).




Figure 25. Two-dimensional potential energy surface (PES) scan of the H_2O molecule in the STO-3G basis, with bond length r (Å) along the x-axis, bond angle θ (°) along the y-axis, and energy E_{SQD} (Ha) along the z-axis. Subplots show results for (a) $\epsilon_s = 10^8$ and (b) $\epsilon_s = 21$.




Figure 26. Contour plots of the Two-dimensional potential energy surface (PES) scan of the H_2O molecule in the STO-3G basis with bond angle θ ($^\circ$) along the x-axis and bond length r ($\text{\AA}$) along the y-axis; the energy is reported in milli-Hartree (mHa) and encoded by the color bar. The minimum-energy configuration is indicated by a red triangle and set to zero, with all contour levels shifted relative to this minimum. Subplots (a) and (b) correspond to $\epsilon_s = 10^8$ and $\epsilon_s = 21$, respectively.

The subspace coverage heatmaps in Figure 27 reveal a qualitatively distinct behaviour between the two subsampling regimes. Under $\epsilon_s = 10^8$, shown in subplot (a), $\eta_{\text{sub}} = 1.0$ is achieved uniformly across the entire (θ, r) grid, confirming that $\mathbb{S}_{\text{sub}}$ fully spans $\mathbb{S}$ at every geometry and that the energies in this regime are equivalent to exact diagonalization within the considered active space. Under $\epsilon_s = 21$, shown in subplot (b), the picture changes markedly: η_{sub} varies across the grid in an essentially stochastic fashion, with values ranging from approximately 0.30 to 0.80 and no discernible smooth dependence on either r or θ . In this 2D setting, the reduced subsampling budget interacts with the two-dimensional geometry space in a more complex manner, producing a diagonalization subspace coverage that fluctuates quasi-randomly across neighbouring grid points. This stochastic variation in η_{sub} directly underlies the roughness visible in the contour plot under $\epsilon_s = 21$ and is a fundamental characteristic of the SQD framework when operating below the threshold subsampling budget required to fully span $\mathbb{S}$ for the current $N_{\text{shots}} = 10,000$.




Figure 27. Diagonalization subspace dimension to symmetry space dimension ratio η_{sub} heatmaps with bond angle θ ($^\circ$) along the x-axis and bond length r ($\text{\AA}$) along the y-axis for: (a) $\epsilon_s = 10^8$ and (b) $\epsilon_s = 21$, respectively.




Figure 28. Energy-difference (ΔE) heatmaps for the two-dimensional potential energy surface (2D-PES) of the H_2O molecule in the STO-3G basis. The bond length r (Å) and bond angle θ (°) are shown along the x- and y-axes, respectively. Subplots show: (a) $\Delta E = |E_{\text{CCSD}} - E_{\text{SQD}}|$ (mHa) for $\epsilon_s = 10^8$, (b) $\Delta E = |E_{\text{CCSD}} - E_{\text{SQD}}|$ (mHa) for $\epsilon_s = 21$, (c) $\Delta E = |E_{\text{FCI}} - E_{\text{SQD}}|$ (mHa) for $\epsilon_s = 10^8$, and (d) $\Delta E = |E_{\text{FCI}} - E_{\text{SQD}}|$ (mHa) for $\epsilon_s = 21$.

The accuracy of the 2D-PES relative to both CCSD and FCI is quantified in the energy-difference heatmaps of Figure 28. Under $\epsilon_s = 10^8$, the deviation from CCSD shown in subplot (a) exhibits a smooth, monotonically increasing gradient along the r axis, ranging from approximately 0.10 mHa at compressed bond lengths near $r = 0.85$ Å to approximately 0.40 mHa at the most stretched geometries near $r = 1.20$ Å, with only weak dependence on θ . This spatially structured pattern reflects the growing difference between the FCI-equivalent SQD energy and the CCSD approximation as the O–H bond is stretched, and multireference character accumulates, and is therefore a consequence of CCSD’s increasing inaccuracy rather than any degradation of SQD. The deviation from FCI, shown in subplot (c), is qualitatively different in character: it exhibits a faint but structured horizontal banding pattern, with $\Delta E_{\text{FCI-SQD}}$ values spanning approximately 0 to 5×10^{-8} mHa, a scale that lies entirely within numerical diagonalization precision. The maximum deviations are concentrated near $r = 1.10$ – 1.20 Å at large θ , but remain negligible in any chemically meaningful sense. These results confirm that under $\epsilon_s = 10^8$, the SQD(LUCJ) 2D-PES is numerically indistinguishable from the exact FCI surface at every grid point.

Under $\epsilon_s = 21$, the spatial character of both deviation heatmaps changes fundamentally. The $\Delta E_{\text{CCSD-SQD}}$ map in subplot (b) and the $\Delta E_{\text{FCI-SQD}}$ map in subplot (d) are both highly heterogeneous, with isolated yellow pixels of elevated deviation reaching up to approximately 3.0 mHa and 3.5 mHa, respectively, scattered across the grid in a pattern that mirrors the stochastic η_{sub} distribution. The majority of grid points retain deviations well below the chemical accuracy threshold of 1.59 mHa, but a non-negligible fraction exceeds this threshold, with the worst-case outliers concentrated at larger bond lengths. Notably, the spatial distribution of high-error points in subplots (b) and (d) is closely correlated, confirming that the SQD energy errors under reduced subsampling are governed primarily by the incompleteness of $\mathbb{S}_{\text{sub}}$ rather than by any systematic bias toward

CCSD or FCI. Further analysis of the diversity of bitstrings generated by SQD(LUCJ) with CCSD initialization in these stretched-geometry cases, as well as their interplay with the symmetry-adapted subsampling budget and the overlap of the resulting $\mathcal{S}_{\text{sub}}$ with the ground state, is deferred to future work.

The rank plots in Figure 29 provide a global view of the correlation between E_{SQD} and the classical reference energies across all grid points simultaneously. Under $\epsilon_s = 10^8$, both the CCSD and FCI rank plots in subplots (a) and (c) show perfect linear correlations, with all points lying on the diagonal, confirming that the relative ordering of energies is faithfully preserved across the entire two-dimensional geometry space. The colour encoding of the point-wise deviation ΔE reveals a physically meaningful structure: in subplot (a), the deviation from CCSD increases systematically from the equilibrium region (dark blue, $\Delta E \approx 10^{-4}$ Ha) toward higher-energy stretched geometries (light yellow, $\Delta E \approx 4 \times 10^{-4}$ Ha), consistent with the smooth gradient observed in the heatmap. In subplot (c), the deviation from FCI spans a range of approximately $0\text{--}5 \times 10^{-11}$ Ha, effectively a single colour across the plot, confirming the near-perfect FCI accuracy of SQD(LUCJ) at all geometries. Under $\epsilon_s = 21$, subplots (b) and (d) preserve the overall linearity but reveal a larger spread about the diagonal, with colour bars now extending to $\sim 3 \times 10^{-3}$ Ha relative to both references. A small number of clearly visible outlier points deviate significantly from the diagonal in subplot (d), corresponding to the scattered high-error geometries identified in the ΔE heatmaps.




Figure 29. Rank plots of absolute energies for the H_2O molecule in the STO-3G basis. Classically computed reference energies E_{CCSD} (a, b) and E_{FCI} (c, d) in Hartree (Ha) are shown on the x-axis, while the corresponding quantum-computed energies E_{SQD} in Hartree (Ha) are shown on the y-axis. Subplots correspond to: (a, c) $\epsilon_s = 10^8$ and (b, d) $\epsilon_s = 21$.

The scatter plots in Figure 30 complement the rank plots by displaying the absolute deviation ΔE as a function of the reference energy for each grid point. Under $\epsilon_s = 10^8$, the $\Delta E_{\text{CCSD-SQD}}$ scatter in subplot (a) exhibits a well-defined, smooth banana-shaped curve: deviations increase monotonically from $\sim 10^{-4}$ Ha at higher energies to approaching 10^{-3} Ha near the lower energies. The $\Delta E_{\text{FCI-SQD}}$ scatter in subplot (c) reveals a particularly striking multi-band structure: the deviations form several quasi-horizontal bands between 10^{-13} and 10^{-10} Ha, an artifact of the discretized 2D-PES grid and the underlying molecular geometries. Under $\epsilon_s = 21$, both scatter plots (b) and (d) show a diffuse cloud of points spanning several orders of

magnitude in ΔE , with a substantial fraction of grid points within the chemical accuracy threshold and a non-negligible number beyond it, as seen in subplot (d). This representation provides a clearer view of the distribution and magnitude of points that lie outside chemical accuracy compared to the rank plots in Figure 29. The loss of the structured multi-band signature in the FCI scatter confirms that stochastic variation in η_{sub} under reduced subsampling is the primary source of the observed variability in SQD energies.




Figure 30. Scatter plots of absolute energies to energy differences ΔE for the H_2O molecule in the STO-3G basis. Classically computed reference energies E_{CCSD} (a, b) and E_{FCI} (c, d) in Hartree (Ha) are shown on the x-axis, while the corresponding energy differences with quantum-computed energies E_{SQD} in Hartree (Ha) are shown on the y-axis. Subplots correspond to: (a, c) $\epsilon_s = 10^8$ and (b, d) $\epsilon_s = 21$.

To characterise the accuracy of the 2D-PES near its minimum in greater detail, one-dimensional slices were extracted along each coordinate direction at the minimum-energy geometry identified under each subsampling condition. Figure 31 presents the 1D slices for $\epsilon_s = 10^8$ at the coincident SQD and FCI minimum located at $r_{\text{min}} = 1.03 \text{ \AA}$ and $\theta_{\text{min}} = 96.61^\circ$. Subplots (a) and (b) show that E_{SQD} , E_{FCI} , and E_{CCSD} overlap with near-perfect visual agreement along both the r slice at fixed θ_{min} and the θ slice at fixed r_{min} . The energy deviation profiles in subplots (c) and (d) quantify this agreement: $\Delta E_{\text{FCI-SQD}}$ lies uniformly at approximately 10^{-11} Ha along both slices, with occasional dips to $\sim 10^{-14}$ Ha, well below the sub-nanohartree threshold and consistent with numerical precision. The CCSD deviation $\Delta E_{\text{CCSD-SQD}}$ along the r slice increases from $\sim 10^{-4}$ Ha at compressed geometries to approaching 10^{-3} Ha at $r = 1.20 \text{ \AA}$, while remaining approximately constant at $\sim 10^{-4}$ Ha along the θ slice, where the bond length is fixed at the equilibrium value. The diagonalization subspace ratio in subplots (e) and (f) is $\eta_{\text{sub}} = 1.0$ at every point along both slices without exception, confirming the robustness and uniformity of subspace coverage at this non-factor subsampling budget.




Figure 31. One-dimensional potential energy surface (1D-PES) slices extracted for $\epsilon_s = 10^8$ at $E_{\text{SQD}}(r_{\min}, \theta_{\min})$ and $E_{\text{FCI}}(r_{\min}, \theta_{\min})$ where $r_{\min} = 1.03$ Å, $\theta_{\min} = 96.61^\circ$ from the two-dimensional PES (2D-PES) of the H_2O molecule in the STO-3G basis. Subplots (a, b) show the absolute energy in Hartree (Ha) as a function of the interatomic distance r (Å) at $\theta_{\min}$ and the H–O–H bond angle θ (°) at $r_{\min}$, respectively. Subplots (c, d) show the corresponding energy differences in Hartree (Ha) as functions of r at $\theta_{\min}$ and θ at $r_{\min}$, respectively. Subplots (e, f) display the ratio of the diagonalization subspace dimension to the symmetry space dimension, η_{sub} , as functions of r at $\theta_{\min}$ and θ at $r_{\min}$, respectively.




Figure 32. One-dimensional potential energy surface (1D-PES) slices extracted for $\epsilon_s = 21$ at $E_{\text{SQD}}(r_{\min}, \theta_{\min})$ where $r_{\min} = 1.03 \text{ \AA}$, $\theta_{\min} = 97.58^\circ$ from the two-dimensional PES (2D-PES) of the H_2O molecule in the STO-3G basis. Subplots (a, b) show the absolute energy in Hartree (Ha) as a function of the interatomic distance r ($\text{\AA}$) at $\theta_{\min}$ and the H-O-H bond angle θ ($^\circ$) at $r_{\min}$, respectively. Subplots (c, d) show the corresponding energy differences in Hartree (Ha) as functions of r at $\theta_{\min}$ and θ at $r_{\min}$, respectively. Subplots (e, f) display the ratio of the diagonalization subspace dimension to the symmetry space dimension, η_{sub} , as functions of r at $\theta_{\min}$ and θ at $r_{\min}$, respectively.




Figure 33. One-dimensional potential energy surface (1D-PES) slices extracted for $\epsilon_s = 21$ at $E_{\text{FCI}}(r_{\min}, \theta_{\min})$ where $r_{\min} = 1.03$ Å, $\theta_{\min} = 96.61^\circ$ from the two-dimensional PES (2D-PES) of the H₂O molecule in the STO-3G basis. Subplots (a, b) show the absolute energy in Hartree (Ha) as a function of the interatomic distance r (Å) at $\theta_{\min}$ and the H-O-H bond angle θ (°) at $r_{\min}$, respectively. Subplots (c, d) show the corresponding energy differences in Hartree (Ha) as functions of r at $\theta_{\min}$ and θ at $r_{\min}$, respectively. Subplots (e, f) display the ratio of the diagonalization subspace dimension to the symmetry space dimension, η_{sub} , as functions of r at $\theta_{\min}$ and θ at $r_{\min}$, respectively.

Figures 32 and 33 present the analogous 1D slices for the reduced-subsampling regime under $\epsilon_s = 21$, extracted at the SQD minimum ($\theta_{\min} = 97.58^\circ$) and the FCI minimum ($\theta_{\min} = 96.61^\circ$), respectively, both at $r_{\min} = 1.03 \text{ \AA}$. The absolute energy curves in subplots (a) and (b) of both figures show that the SQD points track the FCI and CCSD curves closely along both slices, with visible scatter about the reference lines that is more pronounced at the extremes of the scan ranges. The deviation profiles in subplots (c) and (d) quantify the degradation relative to the $\epsilon_s = 10^8$ case: $\Delta E_{\text{FCI-SQD}}$ fluctuates irregularly between approximately 10^{-5} and 10^{-3} Ha along both the r and θ slices, with occasional isolated points approaching or slightly exceeding the chemical accuracy threshold. Notably, a few points in Figure 32(c, d) and Figure 33(c, d) show deep drops to $\sim 10^{-11}$ Ha, indicating that the $\mathbb{S}_{\text{sub}}$ happens to have significant overlap with the FCI ground state at those geometries by chance under the reduced subsampling budget. The η_{sub} profiles in subplots (e) and (f) of both figures are strikingly irregular, fluctuating without a systematic trend between approximately 0.35 and 0.75 along both coordinate directions, a direct manifestation of the stochastic 2D subspace coverage pattern identified in Figure 27(b). The $\Delta E_{\text{CCSD-SQD}}$ profiles are comparatively flat at $\sim 10^{-3}$ – 10^{-4} Ha throughout both slices, reflecting the fact that the CCSD error itself is relatively uniform near the minimum geometry, where multireference effects remain moderate. The two sets of slices, at the SQD and FCI minima, are qualitatively nearly identical despite the $\sim 1^\circ$ offset in θ , confirming that the local energy landscape near the minimum is sufficiently flat that the small shift in the predicted geometry has minimal energetic consequence.

Taken together, the 2D-PES results demonstrate that the SQD(LUCJ) framework is capable of generating accurate and physically meaningful 2D-PES scans on near-term quantum hardware. Under sufficient quantum sampling and non-factorized subsampling, the method achieves FCI-level accuracy at every point on the grid, producing a smooth and chemically consistent surface whose topological features, such as well depth, anisotropy, minimum location, and contour structure, are all faithfully reproduced. Under reduced symmetry-adapted subsampling, the primary effect is the introduction of stochastic, geometry-uncorrelated noise in the energy landscape, driven by random variation in η_{sub} across the grid. This noise degrades the fine structure of the PES near the minimum and introduces isolated high-error geometries, but preserves the global topology and the relative ordering of energies across the grid. These characteristics suggest a clear path to improvement: increasing the subsampling budget toward $\epsilon_s = \sqrt{|\mathbb{S}|}$ or beyond, together with sufficient quantum sampling, reduces stochastic noise and recovers smooth, chemically accurate surfaces; however, this requires prohibitively large N_{shots} and leads to a classically intractable $\mathbb{S}_{\text{sub}}$, ultimately encountering a scaling wall. Therefore, efficient $\mathbb{S}_{\text{sub}}$ generation methods beyond configuration recovery are essential. In addition, extending beyond the current single-reference CCSD initialization of the LUCJ ansatz, along with the development of methods better suited for multireference and strongly correlated regimes, is warranted. Nonetheless, the present results establish a well-characterized performance baseline for 2D-PES mapping with SQD(LUCJ) on current hardware, motivating its extension to larger molecules and more complex coordinate spaces, including the incorporation of classical fragmentation techniques such as DMET in the next subsection.

4.4 SQD(LUCJ) within the Density Matrix Embedding Framework

To assess the effectiveness of the combined methodology of employing DMET along with SQD as a fragmentation with simulation strategy for quantum simulations of molecular systems^{189,190}, we consider two classes of problems of increasing complexity. First, a set of eight ligand-like molecules¹⁹⁰, several of which are relevant to pharmaceutical and chemical applications. This set is selected to systematically evaluate the performance, scalability, and accuracy of the DMET-SQD workflow across chemically diverse yet computationally tractable systems. These molecules enable a controlled analysis of how fragmentation reduces the complexity of the electronic structure problem while preserving chemically relevant correlations within each impurity. Building on this, we extend the study to a significantly larger system, the amantadine molecule ($\text{C}_{10}\text{H}_{17}\text{N}$)²⁵², which represents a substantial increase in both molecule and fragment size. The ability to decompose such a molecule into multiple embedded fragments, each mapped to hardware-compatible active spaces (4 HOMO - 4 LUMO), highlights the key advantage of DMET in enabling simulations that would otherwise be intractable on near-term quantum devices.

4.4.1 Ligand-like Molecules

The quantum circuit resources reported in Table 7 reflect the dependence of circuit complexity on the size of the active space and the corresponding number of qubits used to represent each impurity Hamiltonian. For the smallest fragments, such as hydrogen with $(N_{\text{orb}}, N_{\text{ele}}) = (2, 2)$, the impurity problem maps to $N_Q = 4$ qubits. Consequently, the circuits are shallow, with a depth of $G_D \approx 46$, and require about 24 *R* gates, 12 *CZ* gates, and 28 *Move* gates. The corresponding $\mathbb{S}$ is very small, with $|\mathbb{S}| = 4$ configurations compared to the full $\mathbb{H}$ with $|\mathbb{H}| = 16$. In contrast, the majority of fragments in the studied molecules employ an active space of $(N_{\text{orb}}, N_{\text{ele}}) = (8, 8)$ spin orbitals and electrons, corresponding to $N_Q = 16$ qubits. In these cases the circuit depth increases significantly to approximately $G_D \sim 1100$ – 1200 , with an average of roughly 610–630 *R* gates, 366 *CZ* gates, and ~ 640 – 670 *Move* gates. The larger qubit register also results in a substantial growth of the $\mathbb{H}$ size ($|\mathbb{H}| = 2^{16} = 65,536$), although particle-number symmetry reduces the relevant diagonalization space to $|\mathbb{S}| = 4,900$ configurations.

An intermediate case is encountered for the nitrosyl chloride (NOCl) molecule, where the impurity active spaces correspond




Figure 34. DMET impurity orbital construction for NOCI with atom-centered fragments [O], [N], and [Cl] shown in panels (a)–(c). The isospheres of the fragment orbitals (top row) and corresponding bath orbitals (bottom row) are displayed. An active space of 4 HOMO – 4 LUMO orbitals is selected to restrict each impurity to 16 qubits. Unlike other systems, only a subset of bath orbitals (3 orbitals) contributes significantly to the fragment–environment correlation in NOCI.

to $(N_{\text{orb}}, N_{\text{ele}}) = (7, 8)$, resulting in $N_Q = 14$ qubits. This reduction in the number of orbitals leads to circuits with intermediate resource requirements, with depths of approximately $G_D \sim 800$, around 430–450 R gates, 256 CZ gates, and ~ 450 –460 *Move* gates. Correspondingly, the full $\mathbb{H}$ dimension is reduced to $|\mathbb{H}| = 16,384$, while the particle-number symmetry restricts the physically relevant configuration space to $|\mathbb{S}| = 1,225$. These intermediate values clearly illustrate the scaling behavior of the circuit resources: as the number of orbitals and electrons included in the active space increases, both the gate counts and the accessible $\mathbb{H}$ grow rapidly.

The origin of the $(7, 8)$ active space for NOCI can be understood from the DMET impurity construction illustrated in Figure 34. In this atom-by-atom fragmentation scheme, the impurities correspond to the atomic fragments [O], [N], and [Cl]. The first row in each panel shows the fragment orbitals, while the second row displays the bath orbitals generated from the Schmidt decomposition of the mean-field wavefunction. [N] and [O] atoms, being second-period elements, are described in the STO-3G basis by five spatial orbitals ($1s, 2s, 2p_x, 2p_y, 2p_z$), whereas [Cl] atom, a third-period element, requires nine spatial orbitals ($1s, 2s, 3s, 2p_x, 2p_y, 2p_z, 3p_x, 3p_y, 3p_z$). The resulting fragment–environment entanglement determines the number of bath orbitals generated for each impurity, which in this case leads to seven spatial orbitals being retained in the active space. Consequently, the NOCI impurities exhibit reduced qubit requirements and circuit resources compared to the $(8, 8)$ active spaces while still capturing the relevant fragment–environment correlations within the DMET framework.

In the present simulations, the impurity Hamiltonians are solved on the IQM Sirius QPU, which provides 16 functional physical qubits; further details of the QPU can be found in Appendix B. Consequently, an active-space truncation is applied by selecting the four highest occupied molecular orbitals (4 HOMO) and four lowest unoccupied molecular orbitals (4 LUMO), resulting in a 16-qubit quantum system for each impurity calculation. In Figure 34, the dashed boxes highlight the fragment orbitals corresponding to the selected HOMO subspace, while the dotted boxes indicate the LUMO subspace included in the active space. The shaded portion within the dotted box denotes bath orbitals that exhibit non-negligible entanglement with the fragment orbitals. Notably, while the DMET construction typically yields a number of bath orbitals equal to the number of fragment orbitals, the NOCI system exhibits a small deviation from this behavior: only three bath orbitals possess significant correlation with the fragment space. This indicates that the effective entanglement between the fragment and its environment is lower than the maximum allowed by the fragment dimension, and therefore fewer bath orbitals are sufficient to capture the relevant fragment–environment correlation in the embedding problem.

The results presented in Figure 35 summarize the final energy differences obtained with the DMET-SQD workflow for eight ligand-like molecules simulated on the IQM Sirius QPU, while the Table 8 provides the final energies obtained for the same, along with the reference energies DMET-CASCI and DMET-CCSD. The plot directly compares the absolute deviations of the final DMET-SQD energies with respect to the DMET-CASCI and DMET-CCSD references. For all molecules considered, the deviation between DMET-SQD and DMET-CASCI remains several orders of magnitude below chemical accuracy (10^{-3}




Figure 35. Absolute Energy difference plot obtained with the DMET-SQD workflow for eight ligand-like molecules (provided in Table 7) simulated on the IQM Sirius quantum processor. The bars show the absolute differences between the converged DMET-SQD energies and the corresponding DMET-CASCI ($\Delta E_{CASCI-SQD}$) and DMET-CCSD ($\Delta E_{CCSD-SQD}$) reference energies. The dashed horizontal line indicates the chemical accuracy threshold (10^{-3} Ha). For all molecules, the DMET-SQD energies remain several orders of magnitude closer to the DMET-CASCI reference than to DMET-CCSD, with deviations typically in the micro-Hartree regime.




Figure 36. Convergence behavior and accuracy of the DMET-SQD workflow for eight ligand-like molecules simulated on the 16-qubit IQM Sirius quantum processor. (a) Absolute deviation of the DMET-SQD energies from the DMET-CASCI reference, ΔE_{FCI} , as a function of DMET iteration number N_{iter} . (b) Absolute deviation relative to the DMET-CCSD reference, ΔE_{CCSD} . (c) Convergence of the global chemical potential μ_{glob} during the DMET self-consistency procedure. Molecules are color-coded by molecular weight.

Ha), typically reaching the micro-Hartree regime. In contrast, the differences between DMET-SQD and DMET-CCSD are consistently larger and in several cases approach or exceed the chemical accuracy threshold. This systematic behavior indicates that the SQD-based impurity solver is able to reproduce near-FCI-quality energies for the embedded fragments, while CCSD remains limited by its truncated cluster expansion.. The same qualitative trend is observed across all molecules spanning molecular weights from approximately 43 to 75 Da, including molecules such as HCN_o, CH₃NO, CH₅NO, C₂H₃NO, C₂H₅NO, CH₄N₂O, NOCl, and HOSCN, demonstrating that the high accuracy of DMET-SQD is maintained consistently across chemically diverse systems.

To further understand how these highly accurate final energies are obtained, the convergence behavior of the DMET-SQD workflow is analyzed in Figure 36. Figure 36(c) shows the chemical potential convergence plot, which indicates that the global chemical potential μ_{glob} stabilizes rapidly for all systems, reaching convergence within four to five DMET iterations. In particular, six of the eight molecules converge within four iterations, while only NOCl and CH₄N₂O require a fifth iteration before the chemical potential stabilizes. The energy convergence plots, shown in Figure 36(a) and Figure 36(b), further reveal that the DMET-SQD energies approach the DMET-CASCI reference rapidly with increasing DMET iterations. This high accuracy can be attributed to the absence of symmetry-adapted subspace restrictions during configuration recovery: a prohibitively large ϵ_s was used, ensuring that the recovered subspace was not artificially truncated. Combined with the use of $N_{\text{shots}} = 10,000$, this allows a large fraction of the relevant $\mathbb{S}$ configurations to be included in the SQD diagonalization for each impurity problem. Consequently, the DMET-SQD energies converge to values well within chemical accuracy and often approach micro-Hartree agreement with DMET-CASCI, while remaining consistently more accurate than DMET-CCSD for the molecules considered.

Table 7. Quantum circuit resource summary for the ligand-like molecules’ corresponding molecular fragments, utilizing 10,000 shots per fragment. The depth and the gates shown here correspond to the final iteration of DMET chemical potential convergence.

Molecule	Frag	$(N_{\text{orb}}, N_{\text{ele}})$	N_Q	G_D	R	CZ	$Move$	$ S $	$ H $
HOCN	[H]	(2, 2)	4	46	24	12	28	4	16
	[O]	(8, 8)	16	1,154	626	366	660	4,900	65,536
	[C]	(8, 8)	16	1,148	623	366	656	4,900	65,536
	[N]	(8, 8)	16	1,151	614	366	652	4,900	65,536
CH ₃ NO	[C]	(8, 8)	16	1,162	605	366	660	4,900	65,536
	[H]	(2, 2)	4	46	24	12	28	4	16
	[H]	(2, 2)	4	46	24	12	28	4	16
	[H]	(2, 2)	4	46	24	12	28	4	16
	[N]	(8, 8)	16	1,155	632	366	656	4,900	65,536
	[O]	(8, 8)	16	1,190	637	366	680	4,900	65,536
CH ₅ NO	[C]	(8, 8)	16	1,132	617	366	640	4,900	65,536
	[H]	(2, 2)	4	46	24	12	28	4	16
	[H]	(2, 2)	4	46	24	12	28	4	16
	[H]	(2, 2)	4	46	24	12	28	4	16
	[H]	(2, 2)	4	46	24	12	28	4	16
	[H]	(2, 2)	4	46	24	12	28	4	16
	[O]	(8, 8)	16	1,170	614	366	664	4,900	65,536
	[N]	(8, 8)	16	1,147	608	366	656	4,900	65,536
C ₂ H ₃ NO	[C]	(8, 8)	16	1,176	619	366	668	4,900	65,536
	[C]	(8, 8)	16	1,137	612	366	652	4,900	65,536
	[H]	(2, 2)	4	46	24	12	28	4	16
	[H]	(2, 2)	4	46	24	12	28	4	16
	[H]	(2, 2)	4	46	24	12	28	4	16
	[N]	(8, 8)	16	1,165	614	366	660	4,900	65,536
	[O]	(8, 8)	16	1,132	606	366	644	4,900	65,536
C ₂ H ₅ NO	[C]	(8, 8)	16	1,166	635	366	664	4,900	65,536
	[C]	(8, 8)	16	1,160	599	366	660	4,900	65,536
	[H]	(2, 2)	4	46	24	12	28	4	16
	[H]	(2, 2)	4	46	24	12	28	4	16
	[H]	(2, 2)	4	46	24	12	28	4	16
	[H]	(2, 2)	4	46	24	12	28	4	16
	[H]	(2, 2)	4	46	24	12	28	4	16
	[N]	(8, 8)	16	1,142	613	366	652	4,900	65,536
	[O]	(8, 8)	16	1,171	628	366	668	4,900	65,536
CH ₄ N ₂ O	[C]	(8, 8)	16	1,149	603	366	652	4,900	65,536
	[H]	(2, 2)	4	46	24	12	28	4	16
	[H]	(2, 2)	4	46	24	12	28	4	16
	[H]	(2, 2)	4	46	24	12	28	4	16
	[H]	(2, 2)	4	46	24	12	28	4	16
	[N]	(8, 8)	16	1,162	626	366	668	4,900	65,536
	[N]	(8, 8)	16	1,113	618	366	628	4,900	65,536
	[O]	(8, 8)	16	1,171	608	366	672	4,900	65,536
NOCl	[N]	(7, 8)	14	818	434	256	464	1,225	16,384
	[O]	(7, 8)	14	806	422	256	454	1,225	16,384
	[Cl]	(7, 8)	14	811	447	256	458	1,225	16,384
HOSCN	[H]	(2, 2)	4	46	24	12	28	4	16
	[O]	(8, 8)	16	1,119	611	366	624	4,900	65,536
	[S]	(8, 8)	16	1,163	626	366	672	4,900	65,536
	[C]	(8, 8)	16	1,146	602	366	648	4,900	65,536
	[N]	(8, 8)	16	1,115	618	366	640	4,900	65,536

Table 8. Energy results obtained from the DMET-SQD simulations for the studied ligand-like molecules. The DMET-CCSD column reports reference coupled-cluster energies evaluated on the same fragments (values rounded to 8 decimal places).

Molecule	$E_{\text{DMET-CCSD}}$ (Ha)	$E_{\text{DMET-CASCI}}$ (Ha)	$E_{\text{DMET-SQD}}$ (Ha)
Cyanic Acid (HOCN)	-165.76903095	-165.77369080	-165.77369135
Formaldehyde Oxime (CH ₃ NO)	-166.83748626	-166.84466291	-166.84466203
Methoxyamine (CH ₅ NO)	-167.98584469	-167.98839994	-167.98840075
Methyl Isocyanate (C ₂ H ₃ NO)	-204.46392213	-204.48135661	-204.48135783
Acetaldehyde Oxime (C ₂ H ₅ NO)	-205.49795846	-205.50685229	-205.50684937
Carbamide / Urea (CH ₄ N ₂ O)	-221.26440756	-221.28940628	-221.28940740
Nitrosyl Chloride (NOCl)	-582.34549362	-582.38437288	-582.38437139
Hydroxythiocyanate (HOSCN)	-559.01368538	-559.01613855	-559.01613914

4.4.2 Amantadine Molecule

The quantum circuit resources required for the DMET-SQD simulation of the amantadine molecule are summarized in Table 9. For all eleven fragments, the active space selection of $(N_{\text{orb}}, N_{\text{ele}}) = (8, 8)$ leads to impurity Hamiltonians mapped onto $N_Q = 16$ qubits, consistent with the hardware constraints of the IQM Sirius QPU. The resulting circuit depths are relatively uniform across fragments, lying in the range $G_D \sim 1127$ – 1192 , reflecting the consistent structure of the ansatz and measurement circuits used in SQD. Similarly, the number of R gates remains within ~ 590 – 623 , while the number of entangling CZ gates is fixed at 366 for all fragments. The *Move* gates vary modestly between ~ 636 and 688. The $\mathbb{S}$ for each impurity has dimension $|\mathbb{S}| = 4,900$, compared to the full $\mathbb{H}$ size $|\mathbb{H}| = 2^{16} = 65,536$.

The DMET-SQD simulation was also performed for the amantadine molecule (C₁₀H₁₇N) as per the partition into eleven subsequent fragments as described in Section 3.3. For each fragment, the corresponding impurity Hamiltonian was constructed within the DMET framework and subsequently reduced through an active-space selection consisting of the four highest occupied molecular orbitals (4 HOMO) and four lowest unoccupied molecular orbitals (4 LUMO), yielding a 16-qubit problem compatible with the IQM Sirius quantum hardware. The SQD calculations for each impurity were carried out with a prohibitively large threshold on ϵ_s criterion during configuration recovery in order to avoid artificially restricting the size of the recovered diagonalization subspace. To investigate the influence of the measurement budget on the final DMET-SQD energy, four independent simulations were performed with shot numbers $N_{\text{shots}} = \{500, 1000, 5000, 10000\}$. The resulting bar plot shown in Figure 37 compares the absolute deviations of the final DMET-SQD energies with respect to DMET-CASCI and DMET-CCSD reference calculations. As shown in the figure 37, for all shot counts except $N_{\text{shots}} = 500$, the deviation between the DMET-SQD and DMET-CASCI energies lies within the micro-Hartree regime, whereas the case $N_{\text{shots}} = 500$ exhibits a slightly larger deviation on the order of 10^{-5} Ha. Importantly, only negligible improvements in accuracy are observed when increasing the shot budget from 1,000 to 5,000 or 10,000. This indicates that, for impurity problems of size 16 qubits, a shot budget of approximately 1,000 measurements is sufficient to saturate $\mathbb{S}_{\text{sub}}$ obtained through configuration recovery. Indeed, for $N_{\text{shots}} = \{1000, 5000, 10000\}$ the recovered $\mathbb{S}_{\text{sub}}$ for each impurity spans the $\mathbb{S}$, which explains the nearly identical final DMET-SQD energies observed across these shot budgets. In contrast, for $N_{\text{shots}} = 500$ the recovered subspace $\mathbb{S}_{\text{sub}}$ typically spans approximately 90% – 100% of $\mathbb{S}$, leading to the slightly larger deviations. Consistent with the previous ligand-like molecular systems, the absolute energy differences with respect to DMET-CCSD remain comparatively larger because the DMET-SQD energies lie much closer to the near-exact DMET-CASCI reference.

Further insight into the behavior of the DMET-SQD workflow can be obtained from the convergence plots of the DMET self-consistency cycle provided in Figure 38. The convergence profiles for the energy errors relative to DMET-CASCI and DMET-CCSD, together with the evolution of the global chemical potential μ_{glob} , are shown in the accompanying panels Figure 38(a), (b) and (c) respectively. For most shot budgets, the DMET procedure converges rapidly: the cases $N_{\text{shots}} = 5,000$ and 10,000 reach convergence within four iterations, while $N_{\text{shots}} = 500$ converges in six iterations. A notable peculiarity is observed for $N_{\text{shots}} = 1,000$, where convergence formally required 11 DMET iterations. However, inspection of the chemical

potential trajectory reveals that μ_{glob} approaches essentially the same value for all shot budgets by the third iteration, with smaller subsequent fluctuations. Because the convergence tolerance for the chemical potential optimization was set to 10^{-8} , these minute stochastic variations in μ_{glob} can lead to additional iterations being counted before the convergence criterion is strictly satisfied. Consequently, the larger iteration count observed for $N_{\text{shots}} = 1,000$ likely reflects numerical sensitivity to the stringent tolerance rather than a genuine difference in physical convergence behavior.

The final DMET energies obtained using SQD for different measurement budgets are reported in Table 10, alongside the corresponding DMET-CCSD and DMET-CASCI reference values. As observed, the DMET-SQD energies for $N_{\text{shots}} = \{1000, 5000, 10000\}$ are nearly identical and agree with the DMET-CASCI result to within micro-Hartree precision, confirming that the $\mathbb{S}_{\text{sub}}$ is effectively saturated at these shot counts. In contrast, the $N_{\text{shots}} = 500$ case exhibits a slightly larger deviation of the order of 10^{-5} Ha, consistent with a marginally incomplete recovery of $\mathbb{S}$. The negligible variation in energies beyond $N_{\text{shots}} = 1000$ further indicates that relatively modest measurement budgets are sufficient to achieve converged results for 16-qubit impurity problems within the DMET-SQD framework.

The evolution of $\mathbb{S}_{\text{sub}}$ during the DMET cycle provides additional insight into these observations. Figure 39 shows the quantity η_{sub} , defined as the ratio between the recovered SQD diagonalization subspace $\mathbb{S}_{\text{sub}}$ for each of the eleven fragments labeled F_1-F_{11} . The subplots (a)–(f) correspond to successive values of the global chemical potential μ_{glob} encountered during the DMET optimization, with subplot (f) representing the final converged iteration. As expected, the fraction of the symmetry space accessed varies slightly across fragments and iterations due to the stochastic nature of configuration recovery and the evolving impurity Hamiltonians. Nevertheless, the recovered $\mathbb{S}_{\text{sub}}$ remains consistently close to the full $\mathbb{S}$ throughout the DMET cycle. In the final converged iteration, eight out of the eleven fragments exhibit $\eta_{\text{sub}} = 1$, indicating that $\mathbb{S}_{\text{sub}}$ for these impurities spans the entire $\mathbb{S}$, while the remaining fragments still access a large fraction of $\mathbb{S}$.

Importantly, within a fragmentation-based framework such as DMET combined with SQD, the subspace coverage of individual fragments is less critical than the average value of η_{sub} across all fragments. Because the total DMET energy is assembled from contributions of all impurity calculations, the collective coverage of the relevant $\mathbb{S}$ determines the overall accuracy of the simulation. The consistently high values of η_{sub} observed across fragments therefore explain the near-FCI accuracy obtained in the DMET-SQD energies. Moreover, the absence of an explicit symmetry-adapted subspace truncation ensures that the accessible configuration space is governed primarily by the measurement statistics rather than by algorithmic constraints, allowing the SQD procedure to recover near-complete $\mathbb{S}_{\text{sub}}$ for most fragments w.r.t. $\mathbb{S}$. These observations further confirm that the DMET-SQD framework remains robust and highly accurate even for larger molecular systems with multiple atoms per fragment, such as amantadine, when implemented on near-term quantum hardware.

Table 9. Quantum resources summary for the Amantadine molecule. Here, the resources required for the final iteration of the chemical potential convergence are shown, for the experiment performed with 10,000 shots.

Molecule	Frag	$(N_{\text{orb}}, N_{\text{ele}})$	N_Q	G_D	R	CZ	$Move$	$ \mathbb{S} $	$ \mathbb{H} $
Amantadine ($\text{C}_{10}\text{H}_{17}\text{N}$)	[NH ₂]	(8, 8)	16	1,145	602	366	664	4,900	65,536
	[C]	(8, 8)	16	1,127	590	366	636	4,900	65,536
	[CH ₂]	(8, 8)	16	1,139	604	366	648	4,900	65,536
	[CH ₂]	(8, 8)	16	1,140	623	366	648	4,900	65,536
	[CH]	(8, 8)	16	1,132	612	366	636	4,900	65,536
	[CH]	(8, 8)	16	1,161	617	366	672	4,900	65,536
	[CH ₂]	(8, 8)	16	1,174	610	366	674	4,900	65,536
	[CH]	(8, 8)	16	1,144	622	366	656	4,900	65,536
	[CH ₂]	(8, 8)	16	1,182	611	366	676	4,900	65,536
	[CH ₂]	(8, 8)	16	1,155	608	366	660	4,900	65,536
	[CH ₂]	(8, 8)	16	1,192	617	366	688	4,900	65,536

Table 10. DMET energies using SQD with different shot counts, compared with CCSD and CASCI.

Method	DMET-SQD				DMET-CCSD	DMET-CASCI
	$N_{\text{shots}} = 500$	$N_{\text{shots}} = 1\text{k}$	$N_{\text{shots}} = 5\text{k}$	$N_{\text{shots}} = 10\text{k}$		
Energy (Ha)	-438.09785800	-438.09785047	-438.09785054	-438.09785054	-438.09849129	-438.09784993




Figure 37. Absolute energy deviations of the final DMET-SQD energies for the amantadine molecule with respect to the DMET-CASCI and DMET-CCSD references for different measurement budgets $N_{\text{shots}} = \{500, 1000, 5000, 10000\}$. The dashed horizontal line indicates the chemical accuracy threshold (1.594×10^{-3} Ha). The larger deviations relative to DMET-CCSD reflect the closer agreement of DMET-SQD with the near-exact DMET-CASCI reference.




Figure 38. DMET-SQD convergence for the amantadine molecule at different shot budgets $N_{\text{shots}} = \{500, 1000, 5000, 10000\}$. Subplots (a) and (b) show the energy deviations with respect to DMET-CASCI and DMET-CCSD, respectively, while subplot (c) shows the convergence of the global chemical potential μ_{glob} during the DMET self-consistency iterations.




Figure 39. Evolution of the diagonalization subspace fraction η_{sub} across the eleven fragments (F_1 – F_{11}) of the amantadine molecule during the DMET-SQD self-consistency cycle for the amantadine molecule with $N_{\text{shots}} = 500$. Subplots (a)–(f) correspond to successive values of the global chemical potential μ_{glob} obtained during the DMET iterations, with subplot (f) representing the converged iteration. The results show that the SQD procedure accesses a large fraction of the symmetry-constrained space for most fragments throughout the cycle, with eight fragments reaching $\eta_{\text{sub}} = 1$ in the final iteration.

5 Conclusion

This work presents one of the most comprehensive experimental demonstrations to date of quantum-computing-based molecular simulation on IQM’s Sirius 24-qubit superconducting quantum processor, employing up to 16 operational qubits across a diverse range of quantum chemistry tasks, including ground-state energy calculations, 1D- and 2D-PES scans, and embedding-based simulations of ligand-like and pharmacologically relevant molecules. The unifying methodology is the use of sample-based quantum algorithms, in particular SQD, which partially decouple the quality of the computed energy from the fidelity of individual circuit executions by confining the energetic resolution to an exact diagonalization step over a hardware-sampled subspace. This design renders the approach inherently robust to the noise levels of NISQ hardware.

Two ansätze were benchmarked within the SQD framework. The hardware efficiency and shallow circuit depth of LUCJ make it the preferred choice for near-term execution, consistently outperforming the newly introduced LCNot-UCCSD for SQD and remaining operational across all five benchmark molecules. LCNot-UCCSD, while offering reduced classical pre-computation overhead, produces circuits of sufficient depth to prevent configuration recovery for H_2O and NH_3 , exposing a noise-induced failure mode whose systematic characterization is left for future work. These complementary strengths and limitations provide a practical basis for ansatz selection in near-term quantum simulation workflows.

The potential energy surface results represent some of the most practically significant contributions of this work. SQD(LUCJ) reproduces FCI-quality dissociation profiles across all scanned molecules and basis sets, with accuracy that degrades in a controlled and observable manner as the subsampling budget is reduced. The first experimental 2D-PES for H_2O on superconducting quantum hardware to our knowledge, comprising a 32×32 grid and matching FCI within numerical precision under sufficient subsampling, demonstrates that multi-dimensional surface mapping is now feasible on near-term devices and has the potential to enable improved characterization of transition states, reaction pathways, and force-field parameterization^{255–257}.

Integrating SQD(LUCJ) with DMET extends the methodology to larger and chemically relevant molecules. For eight ligand-like systems and the pharmacologically relevant amantadine molecule ($\text{C}_{10}\text{H}_{17}\text{N}$), DMET-SQD recovers energies close to the micro-Hartree regime relative to DMET-CASCI, well within the chemical accuracy threshold and significantly improving upon DMET-CCSD, while using modest measurement budgets and converging reliably within a small number of self-consistency iterations. To our knowledge, the amantadine simulation constitutes the first experimental demonstration of DMET-SQD at this molecular scale on IQM hardware. As the total energy of the system is assembled collectively from these fragmented impurity calculations, the DMET-SQD framework avoids the exponential scaling wall associated with monolithic molecular simulations. The ability to attain reliable energetics for an FDA-approved therapeutic using 16 operational qubits

demonstrates a pathway forward, one in which practically useful quantum computing for drug discovery does not strictly require millions of fault-tolerant logical qubits, but may be achievable with 50+ qubit NISQ devices with reliable error rates, first scaling toward classically intractable simulations²⁵⁸, before pushing toward quantum advantage with fault tolerance.

Looking ahead, realizing the full potential of this pipeline requires several natural extensions. On the hardware side, scaling to larger and higher-fidelity QPUs will enable larger active spaces and reduce current circuit depth limitations. On the algorithmic side, establishing the efficiency of quantum sampling beyond current state-of-the-art classical techniques¹⁴⁵ and improved error mitigation targeted at configuration recovery¹⁴⁷ would take precedence. Followed by the development of more compact construction of the diagonalization subspace¹⁴⁹ and multi-reference ansatz initialization strategies¹⁹¹, could further enhance accuracy and scalability. Within embedding frameworks, moving beyond the 4 HOMO – 4 LUMO truncation and exploring alternative fragmentation and bath construction schemes¹⁹³ and pushing the sample-based quantum algorithms towards fault-tolerance²⁵⁹ would open the door to economical and efficient simulations of protein-ligand complexes, metal-organic frameworks, and other industrially relevant systems. More broadly, the combination of sampling-based quantum algorithms, chemically inspired hardware-efficient ansätze, and embedding frameworks demonstrated here provides a practical and extensible foundation for chemically accurate quantum simulation, along with a clear quantitative baseline for future hardware and algorithmic advances.

Data Availability

The datasets generated and analysed during the current study are available in the Zenodo repository (DOI: 10.5281/zenodo.19625172).

Acknowledgements

The authors gratefully acknowledge the credit-free access to the IQM Sirius quantum hardware provided during the IQM Quantum School held from 2nd to 4th December 2025, presented by Dr. Stefan Seeger, Dr. Daniel Bulmash, and Dr. Nadia Milazzo. All hardware-based quantum chemistry calculations in this work were performed during this program, and the authors thank the organizers for enabling these experiments. The authors would also like to extend their appreciation to the advisors of Qclairvoyance Quantum Labs for their support, constructive discussions, and continuous inspiration throughout the preparation of this work.

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Competing Interests

The authors declare no competing interests.

References

1. Szabo, A. & Ostlund, N. S. *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory*. Dover Books on Chemistry (Dover Publications, Mineola, NY, 1996), 1st, revised edn.
2. Helgaker, T., Jørgensen, P. & Olsen, J. *Molecular Electronic-Structure Theory* (John Wiley & Sons, Ltd, 2000), 1st edn.
3. Levine, I. N. *Quantum Chemistry*. Pearson Advanced Chemistry Series (Pearson, 2014), 7th edn.
4. McArdle, S., Endo, S., Aspuru-Guzik, A., Benjamin, S. C. & Yuan, X. Quantum computational chemistry. *Rev. Mod. Phys.* **92**, 015003, DOI: [10.1103/RevModPhys.92.015003](https://doi.org/10.1103/RevModPhys.92.015003) (2020).
5. Hohenberg, P. & Kohn, W. Inhomogeneous electron gas. *Phys. Rev.* **136**, B864–B871, DOI: [10.1103/PhysRev.136.B864](https://doi.org/10.1103/PhysRev.136.B864) (1964).
6. Kohn, W. & Sham, L. J. Self-consistent equations including exchange and correlation effects. *Phys. Rev.* **140**, A1133–A1138, DOI: [10.1103/PhysRev.140.A1133](https://doi.org/10.1103/PhysRev.140.A1133) (1965).
7. Martin, R. M. *Electronic Structure: Basic Theory and Practical Methods* (Cambridge University Press, 2004).
8. Hafner, J. Ab-initio simulations of materials using vasp: Density-functional theory and beyond. *J. Comput. Chem.* **29**, 2044–2078, DOI: [10.1002/jcc.21057](https://doi.org/10.1002/jcc.21057) (2008).
9. Burke, K. Perspective on density functional theory. *The J. Chem. Phys.* **136**, 150901, DOI: [10.1063/1.4704546](https://doi.org/10.1063/1.4704546) (2012).

10. Jain, A. *et al.* Commentary: The materials project: A materials genome approach to accelerating materials innovation. *APL Mater.* **1**, 011002, DOI: [10.1063/1.4812323](https://doi.org/10.1063/1.4812323) (2013).
11. Jorgensen, W. L. The many roles of computation in drug discovery. *Science* **303**, 1813–1818, DOI: [10.1126/science.1096361](https://doi.org/10.1126/science.1096361) (2004).
12. Friesner, R. A. Ab initio quantum chemistry: Methodology and applications. *Proc. Natl. Acad. Sci.* **102**, 6648–6653, DOI: [10.1073/pnas.0408036102](https://doi.org/10.1073/pnas.0408036102) (2005).
13. Raha, K. *et al.* The role of quantum mechanics in structure-based drug design. *Drug Discov. Today* **12**, 725–731, DOI: <https://doi.org/10.1016/j.drudis.2007.07.006> (2007).
14. Cournia, Z., Allen, B. & Sherman, W. Relative binding free energy calculations in drug discovery: Recent advances and practical considerations. *J. Chem. Inf. Model.* **57**, 2911–2937, DOI: [10.1021/acs.jcim.7b00564](https://doi.org/10.1021/acs.jcim.7b00564) (2017).
15. Heifetz, A. (ed.) *Quantum Mechanics in Drug Discovery*, vol. 2114 of *Methods in Molecular Biology* (Humana Press, New York, NY, 2020), 1 edn.
16. Niazi, S. K. Quantum mechanics in drug design: Progress, challenges, and future frontiers. *Commun. & Integr. Biol.* **19**, 2603140, DOI: [10.1080/19420889.2025.2603140](https://doi.org/10.1080/19420889.2025.2603140) (2026).
17. Bell, A. T. & Head-Gordon, M. Quantum mechanical modeling of catalytic processes. *Annu. Rev. Chem. Biomol. Eng.* **2**, 453–477, DOI: [10.1146/annurev-chembioeng-061010-114108](https://doi.org/10.1146/annurev-chembioeng-061010-114108) (2011).
18. Nørskov, J. K., Studt, F., Abild-Pedersen, F. & Bligaard, T. *Fundamental Concepts in Heterogeneous Catalysis* (John Wiley & Sons, Ltd, 2014).
19. Deglmann, P., Schäfer, A. & Lennartz, C. Application of quantum calculations in the chemical industry—an overview. *Int. J. Quantum Chem.* **115**, 107–136, DOI: [10.1002/qua.24811](https://doi.org/10.1002/qua.24811) (2015).
20. Peters, B. *Reaction Rate Theory and Rare Events* (Elsevier, Amsterdam, 2017), 1st edn.
21. Nørskov, J. *et al.* Universality in heterogeneous catalysis. *J. Catal.* **209**, 275–278, DOI: [10.1006/jcat.2002.3615](https://doi.org/10.1006/jcat.2002.3615) (2002).
22. Gagliardi, L. *et al.* Multiconfiguration pair-density functional theory: A new way to treat strongly correlated systems. *Accounts Chem. Res.* **50**, 66–73, DOI: [10.1021/acs.accounts.6b00471](https://doi.org/10.1021/acs.accounts.6b00471) (2017).
23. Seh, Z. W. *et al.* Combining theory and experiment in electrocatalysis: Insights into materials design. *Science* **355**, eaad4998, DOI: [10.1126/science.aad4998](https://doi.org/10.1126/science.aad4998) (2017).
24. Keith, D. W., Holmes, G., St. Angelo, D. & Heide, K. A process for capturing CO₂ from the atmosphere. *Joule* **2**, 1573–1594, DOI: <https://doi.org/10.1016/j.joule.2018.05.006> (2018).
25. Nitopi, S. *et al.* Progress and perspectives of electrochemical CO₂ reduction on copper in aqueous electrolyte. *Chem. Rev.* **119**, 7610–7672, DOI: [10.1021/acs.chemrev.8b00705](https://doi.org/10.1021/acs.chemrev.8b00705) (2019).
26. Molani, F. & Cho, A. E. Accurate protein–ligand binding free energy estimation using QM/MM on multi-conformers predicted from classical mining minima. *Commun. Chem.* **7**, 247, DOI: [10.1038/s42004-024-01328-7](https://doi.org/10.1038/s42004-024-01328-7) (2024).
27. Goullieux, M., Zoete, V. & Röhrig, U. F. Hybrid quantum/classical docking of covalent and non-covalent ligands with Attracting Cavities. *Sci. Reports* **15**, 42271, DOI: [10.1038/s41598-025-24614-3](https://doi.org/10.1038/s41598-025-24614-3) (2025).
28. Anderson, P. W. The resonating valence bond state in LaCuO_2 and superconductivity. *Science* **235**, 1196–1198, DOI: [10.1126/science.235.4793.1196](https://doi.org/10.1126/science.235.4793.1196) (1987).
29. Pickett, W. E. Electronic structure of the high-temperature oxide superconductors. *Rev. Mod. Phys.* **61**, 433–512, DOI: [10.1103/RevModPhys.61.433](https://doi.org/10.1103/RevModPhys.61.433) (1989).
30. Keimer, B., Kivelson, S. A., Norman, M. R., Uchida, S. & Zaanen, J. From quantum matter to high-temperature superconductivity in copper oxides. *Nature* **518**, 179–186, DOI: [10.1038/nature14165](https://doi.org/10.1038/nature14165) (2015).
31. Schrödinger, E. Quantisierung als eigenwertproblem. *Annalen der Physik* **384**, 361–376, DOI: [10.1002/andp.19263840404](https://doi.org/10.1002/andp.19263840404) (1926).
32. Pyykkö, P. Relativistic quantum chemistry. In Löwdin, P.-O. (ed.) *Advances in Quantum Chemistry*, vol. 11, 353–409, DOI: [10.1016/S0065-3276\(08\)60241-5](https://doi.org/10.1016/S0065-3276(08)60241-5) (Academic Press, 1978).
33. Pyykkö, P. Relativistic effects in structural chemistry. *Chem. Rev.* **88**, 563–594, DOI: [10.1021/cr00085a006](https://doi.org/10.1021/cr00085a006) (1988).
34. Dyall, K. & Faegri, K. *Introduction to Relativistic Quantum Chemistry* (Oxford University Press, 2007), 1st edn.
35. Pyykkö, P. Relativistic effects in chemistry: More common than you thought. *Annu. Rev. Phys. Chem.* **63**, 45–64, DOI: [10.1146/annurev-physchem-032511-143755](https://doi.org/10.1146/annurev-physchem-032511-143755) (2012).

36. Pyykkö, P. Is relativistic quantum chemistry a good theory of everything? *Pure Appl. Chem.* **97**, 1109–1114, DOI: [doi:10.1515/pac-2025-0460](https://doi.org/10.1515/pac-2025-0460) (2025).
37. Born, M. & Oppenheimer, R. Zur quantentheorie der molekeln. *Annalen der Physik* **389**, 457–484, DOI: [10.1002/andp.19273892002](https://doi.org/10.1002/andp.19273892002) (1927).
38. Yarkony, D. R. Nonadiabatic quantum chemistry—past, present, and future. *Chem. Rev.* **112**, 481–498, DOI: [10.1021/cr2001299](https://doi.org/10.1021/cr2001299) (2012).
39. Curchod, B. F. E. & Martínez, T. J. Ab initio nonadiabatic quantum molecular dynamics. *Chem. Rev.* **118**, 3305–3336, DOI: [10.1021/acs.chemrev.7b00423](https://doi.org/10.1021/acs.chemrev.7b00423) (2018).
40. Pavošević, F., Culpitt, T. & Hammes-Schiffer, S. Multicomponent quantum chemistry: Integrating electronic and nuclear quantum effects via the nuclear–electronic orbital method. *Chem. Rev.* **120**, 4222–4253, DOI: [10.1021/acs.chemrev.9b00798](https://doi.org/10.1021/acs.chemrev.9b00798) (2020).
41. Bellman, R. E. *Adaptive Control Processes: A Guided Tour* (Princeton University Press, Princeton, 1961).
42. Kohn, W. Nobel lecture: Electronic structure of matter—wave functions and density functionals. *Rev. Mod. Phys.* **71**, 1253–1266, DOI: [10.1103/RevModPhys.71.1253](https://doi.org/10.1103/RevModPhys.71.1253) (1999).
43. Knowles, P. & Handy, N. A new determinant-based full configuration interaction method. *Chem. Phys. Lett.* **111**, 315–321, DOI: [https://doi.org/10.1016/0009-2614\(84\)85513-X](https://doi.org/10.1016/0009-2614(84)85513-X) (1984).
44. Olsen, J., Jørgensen, P. & Simons, J. Passing the one-billion limit in full configuration-interaction (fci) calculations. *Chem. Phys. Lett.* **169**, 463–472, DOI: [https://doi.org/10.1016/0009-2614\(90\)85633-N](https://doi.org/10.1016/0009-2614(90)85633-N) (1990).
45. David Sherrill, C. & Schaefer, H. F. The configuration interaction method: Advances in highly correlated approaches. In Löwdin, P.-O., Sabin, J. R., Zerner, M. C. & Brändas, E. (eds.) *Advances in Quantum Chemistry*, vol. 34 of *Advances in Quantum Chemistry*, 143–269, DOI: [https://doi.org/10.1016/S0065-3276\(08\)60532-8](https://doi.org/10.1016/S0065-3276(08)60532-8) (Academic Press, 1999).
46. Vogiatzis, K. D., Ma, D., Olsen, J., Gagliardi, L. & de Jong, W. A. Pushing configuration-interaction to the limit: Towards massively parallel mscf calculations. *The J. Chem. Phys.* **147**, 184111, DOI: [10.1063/1.4989858](https://doi.org/10.1063/1.4989858) (2017).
47. Gao, H., Imamura, S., Kasagi, A. & Yoshida, E. Distributed implementation of full configuration interaction for one trillion determinants. *J. Chem. Theory Comput.* **20**, 1185–1192, DOI: [10.1021/acs.jctc.3c01190](https://doi.org/10.1021/acs.jctc.3c01190) (2024).
48. Shayit, A. *et al.* Numerically exact configuration interaction at quadrillion-determinant scale. *Nat. Commun.* **16**, 11016, DOI: [10.1038/s41467-025-65967-7](https://doi.org/10.1038/s41467-025-65967-7) (2025).
49. Hartree, D. R. The wave mechanics of an atom with a non-coulomb central field. part i. theory and methods. *Math. Proc. Camb. Philos. Soc.* **24**, 89–110, DOI: [10.1017/S0305004100011919](https://doi.org/10.1017/S0305004100011919) (1928).
50. Fock, V. Näherungsmethode zur lösung des quantenmechanischen mehrkörperproblems. *Zeitschrift für Physik* **61**, 126–148, DOI: [10.1007/BF01340294](https://doi.org/10.1007/BF01340294) (1930).
51. Roothaan, C. C. J. New developments in molecular orbital theory. *Rev. Mod. Phys.* **23**, 69–89, DOI: [10.1103/RevModPhys.23.69](https://doi.org/10.1103/RevModPhys.23.69) (1951).
52. Hall, G. G. The molecular orbital theory of chemical valency. viii. a method of calculating ionization potentials. *Proc. Royal Soc. London. Ser. A, Math. Phys. Sci.* **205**, 541–552 (1951).
53. Pribram-Jones, A., Gross, D. A. & Burke, K. Dft: A theory full of holes? *Annu. Rev. Phys. Chem.* **66**, 283–304, DOI: <https://doi.org/10.1146/annurev-physchem-040214-121420> (2015).
54. Jones, R. O. Density functional theory: Its origins, rise to prominence, and future. *Rev. Mod. Phys.* **87**, 897–923, DOI: [10.1103/RevModPhys.87.897](https://doi.org/10.1103/RevModPhys.87.897) (2015).
55. Goedecker, S. Linear scaling electronic structure methods. *Rev. Mod. Phys.* **71**, 1085–1123, DOI: [10.1103/RevModPhys.71.1085](https://doi.org/10.1103/RevModPhys.71.1085) (1999).
56. Mardirossian, N. & Head-Gordon, M. Thirty years of density functional theory in computational chemistry: an overview and extensive assessment of 200 density functionals. *Mol. Phys.* **115**, 2315–2372, DOI: [10.1080/00268976.2017.1333644](https://doi.org/10.1080/00268976.2017.1333644) (2017).
57. Donahue, N. M. Reaction barriers: Origin and evolution. *Chem. Rev.* **103**, 4593–4604, DOI: [10.1021/cr020650g](https://doi.org/10.1021/cr020650g) (2003).
58. Hait, D. & Head-Gordon, M. Delocalization errors in density functional theory are essentially quadratic in fractional occupation number. *The J. Phys. Chem. Lett.* **9**, 6280–6288, DOI: [10.1021/acs.jpclett.8b02417](https://doi.org/10.1021/acs.jpclett.8b02417) (2018).

59. Nitz, F. *et al.* An experimental benchmark for the barrier height of an on-surface reaction: Hydrogen oxidation on pt(332). *J. Am. Chem. Soc.* **147**, 37658–37666, DOI: [10.1021/jacs.5c12862](https://doi.org/10.1021/jacs.5c12862) (2025).
60. Müller-Dethlefs, K. & Hobza, P. Noncovalent interactions: A challenge for experiment and theory. *Chem. Rev.* **100**, 143–168, DOI: [10.1021/cr9900331](https://doi.org/10.1021/cr9900331) (2000).
61. Masumian, E. & Boese, A. D. Benchmarking swaths of intermolecular interaction components with symmetry-adapted perturbation theory. *J. Chem. Theory Comput.* **20**, 30–48, DOI: [10.1021/acs.jctc.3c00801](https://doi.org/10.1021/acs.jctc.3c00801) (2024).
62. Khabibrakhmanov, A., Gori, M., Müller, C. & Tkatchenko, A. Noncovalent interactions in density functional theory: All the charge density we do not see. *J. Am. Chem. Soc.* **147**, 40763–40775, DOI: [10.1021/jacs.5c13706](https://doi.org/10.1021/jacs.5c13706) (2025).
63. Matsika, S. & Krylov, A. I. Introduction: Theoretical modeling of excited state processes. *Chem. Rev.* **118**, 6925–6926, DOI: [10.1021/acs.chemrev.8b00436](https://doi.org/10.1021/acs.chemrev.8b00436) (2018).
64. Jacquemin, D., Kossoski, F., Gam, F., Boggio-Pasqua, M. & Loos, P.-F. Reference vertical excitation energies for transition metal compounds. *J. Chem. Theory Comput.* **19**, 8782–8800, DOI: [10.1021/acs.jctc.3c01080](https://doi.org/10.1021/acs.jctc.3c01080) (2023).
65. Peverati, R. & Truhlar, D. G. Quest for a universal density functional: the accuracy of density functionals across a broad spectrum of databases in chemistry and physics. *Philos. Transactions Royal Soc. A: Math. Phys. Eng. Sci.* **372**, 20120476, DOI: [10.1098/rsta.2012.0476](https://doi.org/10.1098/rsta.2012.0476) (2014).
66. Raghavachari, K., Trucks, G. W., Pople, J. A. & Head-Gordon, M. A fifth-order perturbation comparison of electron correlation theories. *Chem. Phys. Lett.* **157**, 479–483, DOI: [https://doi.org/10.1016/S0009-2614\(89\)87395-6](https://doi.org/10.1016/S0009-2614(89)87395-6) (1989).
67. Mester, D., Nagy, P. R. & Kállay, M. Basis-set limit ccSD(T) energies for large molecules with local natural orbitals and reduced-scaling basis-set corrections. *J. Chem. Theory Comput.* **20**, 7453–7468, DOI: [10.1021/acs.jctc.4c00777](https://doi.org/10.1021/acs.jctc.4c00777) (2024).
68. Nagy, P. R. State-of-the-art local correlation methods enable affordable gold standard quantum chemistry for up to hundreds of atoms. *Chem. Sci.* **15**, 14556–14584, DOI: [10.1039/D4SC04755A](https://doi.org/10.1039/D4SC04755A) (2024).
69. Gordon, M. S., Fedorov, D. G., Pruitt, S. R. & Slipchenko, L. V. Fragmentation methods: A route to accurate calculations on large systems. *Chem. Rev.* **112**, 632–672, DOI: [10.1021/cr200093j](https://doi.org/10.1021/cr200093j) (2012).
70. Raghavachari, K. & Saha, A. Accurate composite and fragment-based quantum chemical models for large molecules. *Chem. Rev.* **115**, 5643–5677, DOI: [10.1021/cr500606e](https://doi.org/10.1021/cr500606e) (2015).
71. Broderick, D. R. *et al.* Fragment: An open-source framework for multiscale quantum chemistry based on fragmentation. *WIREs Comput. Mol. Sci.* **15**, e70058, DOI: <https://doi.org/10.1002/wcms.70058> (2025).
72. Sun, Q. & Chan, G. K.-L. Quantum embedding theories. *Accounts Chem. Res.* **49**, 2705–2712, DOI: [10.1021/acs.accounts.6b00356](https://doi.org/10.1021/acs.accounts.6b00356) (2016).
73. Vorwerk, C., Sheng, N., Govoni, M., Huang, B. & Galli, G. Quantum embedding theories to simulate condensed systems on quantum computers. *Nat. Comput. Sci.* **2**, 424–432, DOI: [10.1038/s43588-022-00279-0](https://doi.org/10.1038/s43588-022-00279-0) (2022).
74. Verma, S. *et al.* Multireference embedding and fragmentation methods for classical and quantum computers: From model systems to realistic applications. *Chem. Rev.* **126**, 184–203, DOI: [10.1021/acs.chemrev.5c00486](https://doi.org/10.1021/acs.chemrev.5c00486) (2026).
75. Feynman, R. P. Simulating physics with computers. *Int. J. Theor. Phys.* **21**, 467–488, DOI: [10.1007/BF02650179](https://doi.org/10.1007/BF02650179) (1982).
76. Zhang, Y. *et al.* Quantum algorithms for quantum molecular systems: A survey. *WIREs Comput. Mol. Sci.* **15**, e70020, DOI: <https://doi.org/10.1002/wcms.70020> (2025).
77. Patra, A. K. *et al.* Survey of quantum algorithms: Foundations, frameworks and applications. *Authorea* DOI: [10.22541/au.176341037.72812424/v1](https://doi.org/10.22541/au.176341037.72812424/v1) (2025).
78. Dirac, P. A. M. *The Principles of Quantum Mechanics*, chap. 1, 1–18 (Oxford University Press, Oxford, 1958), 4th edn.
79. Einstein, A., Podolsky, B. & Rosen, N. Can quantum-mechanical description of physical reality be considered complete? *Phys. Rev.* **47**, 777–780, DOI: [10.1103/PhysRev.47.777](https://doi.org/10.1103/PhysRev.47.777) (1935).
80. Horodecki, R., Horodecki, P., Horodecki, M. & Horodecki, K. Quantum entanglement. *Rev. Mod. Phys.* **81**, 865–942, DOI: [10.1103/RevModPhys.81.865](https://doi.org/10.1103/RevModPhys.81.865) (2009).
81. Nielsen, M. & Chuang, I. *Quantum Computation and Quantum Information*. Cambridge Series on Information and the Natural Sciences (Cambridge University Press, 2000).
82. Aspuru-Guzik, A., Dutoi, A. D., Love, P. J. & Head-Gordon, M. Simulated quantum computation of molecular energies. *Science* **309**, 1704–1707, DOI: [10.1126/science.1113479](https://doi.org/10.1126/science.1113479) (2005).
83. Lloyd, S. Universal quantum simulators. *Science* **273**, 1073–1078, DOI: [10.1126/science.273.5278.1073](https://doi.org/10.1126/science.273.5278.1073) (1996).

84. Abrams, D. S. & Lloyd, S. Simulation of many-body fermi systems on a universal quantum computer. *Phys. Rev. Lett.* **79**, 2586–2589, DOI: [10.1103/PhysRevLett.79.2586](https://doi.org/10.1103/PhysRevLett.79.2586) (1997).
85. Kitaev, A. Y. Quantum measurements and the abelian stabilizer problem, DOI: [10.48550/arXiv.quant-ph/9511026](https://doi.org/10.48550/arXiv.quant-ph/9511026) (1995). [quant-ph/9511026](https://doi.org/10.48550/arXiv.quant-ph/9511026).
86. Abrams, D. S. & Lloyd, S. Quantum algorithm providing exponential speed increase for finding eigenvalues and eigenvectors. *Phys. Rev. Lett.* **83**, 5162–5165, DOI: [10.1103/PhysRevLett.83.5162](https://doi.org/10.1103/PhysRevLett.83.5162) (1999).
87. Childs, A. M. & Wiebe, N. Hamiltonian simulation using linear combinations of unitary operations. *Quantum Inf. Comput.* **12**, 0901–0924, DOI: [10.26421/QIC12.11-12-1](https://doi.org/10.26421/QIC12.11-12-1) (2012).
88. Campbell, E. Random compiler for fast hamiltonian simulation. *Phys. Rev. Lett.* **123**, 070503, DOI: [10.1103/PhysRevLett.123.070503](https://doi.org/10.1103/PhysRevLett.123.070503) (2019).
89. Low, G. H. & Chuang, I. L. Hamiltonian Simulation by Qubitization. *Quantum* **3**, 163, DOI: [10.22331/q-2019-07-12-163](https://doi.org/10.22331/q-2019-07-12-163) (2019).
90. Low, G. H. & Chuang, I. L. Optimal hamiltonian simulation by quantum signal processing. *Phys. Rev. Lett.* **118**, 010501, DOI: [10.1103/PhysRevLett.118.010501](https://doi.org/10.1103/PhysRevLett.118.010501) (2017).
91. Gilyén, A., Su, Y., Low, G. H. & Wiebe, N. Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019*, 193–204, DOI: [10.1145/3313276.3316366](https://doi.org/10.1145/3313276.3316366) (Association for Computing Machinery, New York, NY, USA, 2019).
92. Clader, B. D. *et al.* Quantum resources required to block-encode a matrix of classical data. *IEEE Transactions on Quantum Eng.* **3**, 1–23, DOI: [10.1109/TQE.2022.3231194](https://doi.org/10.1109/TQE.2022.3231194) (2022).
93. Camps, D., Lin, L., Van Beeumen, R. & Yang, C. Explicit quantum circuits for block encodings of certain sparse matrices. *SIAM J. on Matrix Analysis Appl.* **45**, 801–827, DOI: [10.1137/22M1484298](https://doi.org/10.1137/22M1484298) (2024).
94. Sünderhauf, C., Campbell, E. & Camps, J. Block-encoding structured matrices for data input in quantum computing. *Quantum* **8**, 1226, DOI: [10.22331/q-2024-01-11-1226](https://doi.org/10.22331/q-2024-01-11-1226) (2024).
95. Lee, J. *et al.* Even more efficient quantum computations of chemistry through tensor hypercontraction. *PRX Quantum* **2**, 030305, DOI: [10.1103/PRXQuantum.2.030305](https://doi.org/10.1103/PRXQuantum.2.030305) (2021).
96. Oumarou, O., Scheurer, M., Parrish, R. M., Hohenstein, E. G. & Gogolin, C. Accelerating Quantum Computations of Chemistry Through Regularized Compressed Double Factorization. *Quantum* **8**, 1371, DOI: [10.22331/q-2024-06-13-1371](https://doi.org/10.22331/q-2024-06-13-1371) (2024).
97. Caesura, A. *et al.* Faster quantum chemistry simulations on a quantum computer with improved tensor factorization and active volume compilation. *PRX Quantum* **6**, 030337, DOI: [10.1103/yngp-5fpm](https://doi.org/10.1103/yngp-5fpm) (2025).
98. Reiher, M., Wiebe, N., Svore, K. M., Wecker, D. & Troyer, M. Elucidating reaction mechanisms on quantum computers. *Proc. Natl. Acad. Sci.* **114**, 7555–7560, DOI: [10.1073/pnas.1619152114](https://doi.org/10.1073/pnas.1619152114) (2017).
99. Goings, J. J. *et al.* Reliably assessing the electronic structure of cytochrome p450 on today’s classical computers and tomorrow’s quantum computers. *Proc. Natl. Acad. Sci.* **119**, e2203533119, DOI: [10.1073/pnas.2203533119](https://doi.org/10.1073/pnas.2203533119) (2022).
100. Beverland, M. E. *et al.* Assessing requirements to scale to practical quantum advantage (2022). [2211.07629](https://arxiv.org/abs/2211.07629).
101. Lee, S. *et al.* Evaluating the evidence for exponential quantum advantage in ground-state quantum chemistry. *Nat. Commun.* **14**, 1952, DOI: [10.1038/s41467-023-37587-6](https://doi.org/10.1038/s41467-023-37587-6) (2023).
102. Dalzell, A. M. *et al.* *Quantum Algorithms: A Survey of Applications and End-to-end Complexities* (Cambridge University Press, 2025).
103. Huang, H.-Y., Choi, S., McClean, J. R. & Preskill, J. The vast world of quantum advantage (2025). [2508.05720](https://arxiv.org/abs/2508.05720).
104. Kretschmer, W. *et al.* Demonstrating an unconditional separation between quantum and classical information resources (2025). [2509.07255](https://arxiv.org/abs/2509.07255).
105. Eisert, J. & Preskill, J. Mind the gaps: The fraught road to quantum advantage (2025). [2510.19928](https://arxiv.org/abs/2510.19928).
106. Alexeev, Y. *et al.* A perspective on quantum computing applications in quantum chemistry using 25–100 logical qubits (2025). [2506.19337](https://arxiv.org/abs/2506.19337).
107. Lin, L. Dissipative preparation of many-body quantum states: Toward practical quantum advantage. *APL Comput. Phys.* **1**, 010901, DOI: [10.1063/5.0283315](https://doi.org/10.1063/5.0283315) (2025).

108. Toshio, R. *et al.* Practical quantum advantage on partially fault-tolerant quantum computer. *Phys. Rev. X* **15**, 021057, DOI: [10.1103/PhysRevX.15.021057](https://doi.org/10.1103/PhysRevX.15.021057) (2025).
109. Preskill, J. Quantum Computing in the NISQ era and beyond. *Quantum* **2**, 79, DOI: [10.22331/q-2018-08-06-79](https://doi.org/10.22331/q-2018-08-06-79) (2018).
110. Peruzzo, A. *et al.* A variational eigenvalue solver on a photonic quantum processor. *Nat. Commun.* **5**, 4213, DOI: [10.1038/ncomms5213](https://doi.org/10.1038/ncomms5213) (2014).
111. McClean, J. R., Romero, J., Babbush, R. & Aspuru-Guzik, A. The theory of variational hybrid quantum-classical algorithms. *New J. Phys.* **18**, 023023, DOI: [10.1088/1367-2630/18/2/023023](https://doi.org/10.1088/1367-2630/18/2/023023) (2016).
112. Fedorov, D. A., Peng, B., Govind, N. & Alexeev, Y. Vqe method: a short survey and recent developments. *Mater. Theory* **6**, 2, DOI: [10.1186/s41313-021-00032-6](https://doi.org/10.1186/s41313-021-00032-6) (2022).
113. Kandala, A. *et al.* Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature* **549**, 242–246, DOI: [10.1038/nature23879](https://doi.org/10.1038/nature23879) (2017).
114. Barkoutsos, P. K. *et al.* Quantum algorithms for electronic structure calculations: Particle-hole hamiltonian and optimized wave-function expansions. *Phys. Rev. A* **98**, 022322, DOI: [10.1103/PhysRevA.98.022322](https://doi.org/10.1103/PhysRevA.98.022322) (2018).
115. Ramôa, M. *et al.* Reducing the resources required by ADAPT-VQE using coupled exchange operators and improved subroutines. *npj Quantum Inf.* **11**, 86, DOI: [10.1038/s41534-025-01039-4](https://doi.org/10.1038/s41534-025-01039-4) (2025).
116. He, R. *et al.* Constructing compact adapt unitary coupled-cluster ansatz with parameter-based criterion (2026). [2602.04253](https://doi.org/10.26434/chemrxiv-2026-04253).
117. Vaquero-Sabater, N., Carreras, A. & Casanova, D. Pruned-adapt-vqe: Compacting molecular ansätze by removing irrelevant operators. *J. Chem. Theory Comput.* **21**, 8720–8728, DOI: [10.1021/acs.jctc.5c00535](https://doi.org/10.1021/acs.jctc.5c00535) (2025).
118. Bauman, N. P. *et al.* Coupled cluster downfolding theory in simulations of correlated systems on quantum hardware. *Phys. Rev. Res.* **8**, 013072, DOI: [10.1103/b1t6-ln6v](https://doi.org/10.1103/b1t6-ln6v) (2026).
119. McClean, J. R., Boixo, S., Smelyanskiy, V. N., Babbush, R. & Neven, H. Barren plateaus in quantum neural network training landscapes. *Nat. Commun.* **9**, 4812, DOI: [10.1038/s41467-018-07090-4](https://doi.org/10.1038/s41467-018-07090-4) (2018).
120. Wang, S. *et al.* Noise-induced barren plateaus in variational quantum algorithms. *Nat. Commun.* **12**, 6961, DOI: [10.1038/s41467-021-27045-6](https://doi.org/10.1038/s41467-021-27045-6) (2021).
121. Singh, H., Majumder, S. & Mishra, S. Sharc-vqe: Simplified hamiltonian approach with refinement and correction enabled variational quantum eigensolver for molecular simulation. *The J. Chem. Phys.* **162**, 114117, DOI: [10.1063/5.0249447](https://doi.org/10.1063/5.0249447) (2025).
122. de Lima, A. M., Teixeira, E. S., Silva, E. D. R. d. A. & Longo, R. L. Improving energy and molecular properties by convergence of the one-particle reduced density matrix in variational quantum eigensolvers (vqe). *J. Comput. Chem.* **47**, e70289, DOI: <https://doi.org/10.1002/jcc.70289> (2026).
123. Singh, A., Rangi, C. & Tam, K.-M. Classical benchmarks of a symmetry-adapted variational quantum eigensolver for real-time green's functions in dynamical mean-field theory (2026).
124. Carreras, A., Orús, R. & Casanova, D. Limitations of quantum hardware for molecular energy estimation using vqe. *Phys. Chem. Chem. Phys.* **28**, 2834–2846, DOI: [10.1039/D5CP03907J](https://doi.org/10.1039/D5CP03907J) (2026).
125. Jiang, T. *et al.* Walking through hilbert space with quantum computers. *Chem. Rev.* **125**, 4569–4602, DOI: [10.1021/acs.chemrev.4c00508](https://doi.org/10.1021/acs.chemrev.4c00508) (2025).
126. Kanno, K. *et al.* Quantum-selected configuration interaction: classical diagonalization of hamiltonians in subspaces selected by quantum computers (2023). [2302.11320](https://doi.org/10.26434/chemrxiv-2023-2302).
127. Nakagawa, Y. O., Kamoshita, M., Mizukami, W., Sudo, S. & Ohnishi, Y.-y. Adapt-qsci: Adaptive construction of an input state for quantum-selected configuration interaction. *J. Chem. Theory Comput.* **20**, 10817–10825, DOI: [10.1021/acs.jctc.4c00846](https://doi.org/10.1021/acs.jctc.4c00846) (2024).
128. Sugisaki, K., Kanno, S., Itoko, T., Sakuma, R. & Yamamoto, N. Hamiltonian simulation-based quantum-selected configuration interaction for large-scale electronic structure calculations with a quantum computer. *Phys. Chem. Chem. Phys.* **27**, 20869–20884, DOI: [10.1039/D5CP02202A](https://doi.org/10.1039/D5CP02202A) (2025).
129. Mikkelsen, M. & Nakagawa, Y. O. Quantum-selected configuration interaction with time-evolved state. *Phys. Rev. Res.* **7**, 043043, DOI: [10.1103/75pv-hbrx](https://doi.org/10.1103/75pv-hbrx) (2025).
130. Robledo-Moreno, J. *et al.* Chemistry beyond the scale of exact diagonalization on a quantum-centric supercomputer. *Sci. Adv.* **11**, eadu9991, DOI: [10.1126/sciadv.adu9991](https://doi.org/10.1126/sciadv.adu9991) (2025).

131. Shirakawa, T. *et al.* Closed-loop calculations of electronic structure on a quantum processor and a classical supercomputer at full scale (2025). [2511.00224](#).
132. Barison, S., Robledo Moreno, J. & Motta, M. Quantum-centric computation of molecular excited states with extended sample-based quantum diagonalization. *Quantum Sci. Technol.* **10**, 025034, DOI: [10.1088/2058-9565/adb781](#) (2025).
133. Liepuoniute, I. *et al.* Quantum-centric computational study of methylene singlet and triplet states. *J. Chem. Theory Comput.* **21**, 5062–5070, DOI: [10.1021/acs.jctc.5c00075](#) (2025).
134. Kaliakin, D., Shajan, A., Liang, F. & Merz, K. M. J. Implicit solvent sample-based quantum diagonalization. *The J. Phys. Chem. B* **129**, 5788–5796, DOI: [10.1021/acs.jpccb.5c01030](#) (2025).
135. Kaliakin, D. *et al.* Accurate quantum-centric simulations of intermolecular interactions. *Commun. Phys.* **8**, 396, DOI: [10.1038/s42005-025-02305-9](#) (2025).
136. Yu, J. *et al.* Quantum-centric algorithm for sample-based krylov diagonalization (2025). [2501.09702](#).
137. Yoshioka, N. *et al.* Krylov diagonalization of large many-body hamiltonians on a quantum processor. *Nat. Commun.* **16**, 5014, DOI: [10.1038/s41467-025-59716-z](#) (2025).
138. Piccinelli, S. *et al.* Quantum chemistry with provable convergence via randomized sample-based krylov quantum diagonalization (2026). [2508.02578](#).
139. Asthana, A. Quantum krylov algorithm using unitary decomposition for exact eigenstates of fermionic systems using quantum computers (2025). [2512.11788](#).
140. Huggins, W. J. *et al.* Unbiasing fermionic quantum monte carlo with a quantum computer. *Nature* **603**, 416–420, DOI: [10.1038/s41586-021-04351-z](#) (2022).
141. Yoshida, Y. *et al.* Auxiliary-field quantum monte carlo method with quantum selected configuration interaction (2025). [2502.21081](#).
142. Danilov, D., Robledo-Moreno, J., Sung, K. J., Motta, M. & Shee, J. Enhancing the accuracy and efficiency of sample-based quantum diagonalization with phaseless auxiliary-field quantum monte carlo (2025). [2503.05967](#).
143. Walkup, R. *et al.* Scaling sample-based quantum diagonalization on gpu-accelerated systems using openmp offload (2026). [2601.16169](#).
144. Doi, J., Shirakawa, T., Kawashima, Y., Yunoki, S. & Horii, H. Gpu-accelerated selected basis diagonalization with thrust for sqd-based algorithms (2026). [2601.16637](#).
145. Reinholdt, P. *et al.* Critical limitations in quantum-selected configuration interaction methods. *J. Chem. Theory Comput.* **21**, 6811–6822, DOI: [10.1021/acs.jctc.5c00375](#) (2025).
146. Raisuddin, O. M., Zhang, H., Motta, M. & Faulstich, F. M. From promise to practice: Benchmarking quantum chemistry on quantum hardware (2025). [2512.01012](#).
147. Weaving, T., Mingare, A., Ralli, A. & Coveney, P. V. Towards compact wavefunctions from quantum-selected configuration interaction (2025). [2509.02525](#).
148. Pellow-Jarman, A. *et al.* Hivqe: Handover iterative variational quantum eigensolver for efficient quantum chemistry calculations (2025). [2503.06292](#).
149. Yoo, P. *et al.* Extending the handover-iterative vqe to challenging strongly correlated systems: n_2 and fe-s cluster (2026).
150. Ghasemi, F., Kang, Y., Kawashima, Y. & Moon, K. Hybrid quantum algorithms for computational chemistry: Application to the pyridine-li ion complex (2026). [2601.10002](#).
151. Smith, T. *et al.* Quantum-centric simulation of hydrogen abstraction by sample-based quantum diagonalization and entanglement forging (2025). [2508.08229](#).
152. McFarthing, S., Pellow-Jarman, A. & Petruccione, F. Noise-resilient quantum chemistry with half the qubits (2026). [2602.01165](#).
153. Yoshida, Y., Murokoshi, T., Kuroda, N. & Mizukami, W. Doubling the size of quantum selected configuration interaction based on seniority-zero space and its application to qc-qsci-afqmc (2026). [2602.07912](#).
154. Patra, C. *et al.* Physics-informed generative machine learning for accelerated quantum-centric supercomputing (2026). [2512.06858](#).

155. Tilly, J. *et al.* The variational quantum eigensolver: A review of methods and best practices. *Phys. Reports* **986**, 1–128, DOI: <https://doi.org/10.1016/j.physrep.2022.08.003> (2022). The Variational Quantum Eigensolver: a review of methods and best practices.
156. Grimsley, H. R., Economou, S. E., Barnes, E. & Mayhall, N. J. An adaptive variational algorithm for exact molecular simulations on a quantum computer. *Nat. Commun.* **10**, 3007, DOI: [10.1038/s41467-019-10988-2](https://doi.org/10.1038/s41467-019-10988-2) (2019).
157. Verteletskyi, V., Yen, T.-C. & Izmaylov, A. F. Measurement optimization in the variational quantum eigensolver using a minimum clique cover. *The J. Chem. Phys.* **152**, 124114, DOI: [10.1063/1.5141458](https://doi.org/10.1063/1.5141458) (2020).
158. Nelson, J. S. & Baczewski, A. D. Assessment of quantum phase estimation protocols for early fault-tolerant quantum computers. *Phys. Rev. A* **110**, 042420, DOI: [10.1103/PhysRevA.110.042420](https://doi.org/10.1103/PhysRevA.110.042420) (2024).
159. Anurag, K. S. V. *et al.* Resource estimation for vqe on small molecules: Impact of fermion mappings and hamiltonian reductions (2025). [2512.01605](https://arxiv.org/abs/2512.01605).
160. Yamazaki, T., Matsuura, S., Narimani, A., Saidmuradov, A. & Zaribafiyani, A. Towards the practical application of near-term quantum computers in quantum chemistry simulations: A problem decomposition approach (2018). [1806.01305](https://arxiv.org/abs/1806.01305).
161. Kitaura, K., Ikeo, E., Asada, T., Nakano, T. & Uebayasi, M. Fragment molecular orbital method: an approximate computational method for large molecules. *Chem. Phys. Lett.* **313**, 701–706, DOI: [https://doi.org/10.1016/S0009-2614\(99\)00874-X](https://doi.org/10.1016/S0009-2614(99)00874-X) (1999).
162. Fedorov, D. G., Nagata, T. & Kitaura, K. Exploring chemistry with the fragment molecular orbital method. *Phys. Chem. Chem. Phys.* **14**, 7562–7577, DOI: [10.1039/C2CP23784A](https://doi.org/10.1039/C2CP23784A) (2012).
163. Xantheas, S. S. Ab initio studies of cyclic water clusters (h₂O)_n, n=1–6. ii. analysis of many-body interactions. *The J. Chem. Phys.* **100**, 7523–7534, DOI: [10.1063/1.466846](https://doi.org/10.1063/1.466846) (1994).
164. Eriksen, J. J. & Gauss, J. Generalized many-body expanded full configuration interaction theory. *The J. Phys. Chem. Lett.* **10**, 7910–7915, DOI: [10.1021/acs.jpcllett.9b02968](https://doi.org/10.1021/acs.jpcllett.9b02968) (2019).
165. Ballesteros, F., Tan, J. A. & Lao, K. U. An accurate and efficient fragmentation approach via the generalized many-body expansion for density matrices. *The J. Chem. Phys.* **159**, 074107, DOI: [10.1063/5.0160810](https://doi.org/10.1063/5.0160810) (2023).
166. Burns, L. A., Sherrill, C. D. & Pritchard, B. P. Qcmmanybody: A flexible implementation of the many-body expansion. *The J. Chem. Phys.* **161**, 152501, DOI: [10.1063/5.0231843](https://doi.org/10.1063/5.0231843) (2024).
167. Lim, H. *et al.* Fragment molecular orbital-based variational quantum eigensolver for quantum chemistry in the age of quantum computing. *Sci. Reports* **14**, 2422, DOI: [10.1038/s41598-024-52926-3](https://doi.org/10.1038/s41598-024-52926-3) (2024).
168. Xu, E., Shimomoto, Y., Ten-no, S. L. & Tsuchimochi, T. Many-body-expansion based on variational quantum eigensolver and deflation for dynamical correlation. *The J. Phys. Chem. A* **128**, 2507–2521, DOI: [10.1021/acs.jpca.4c00351](https://doi.org/10.1021/acs.jpca.4c00351) (2024).
169. Yang, W. & Lee, T. A density-matrix divide-and-conquer approach for electronic structure calculations of large molecules. *The J. Chem. Phys.* **103**, 5674–5678, DOI: [10.1063/1.470549](https://doi.org/10.1063/1.470549) (1995).
170. Nakai, H. & Kobayashi, M. Linear-scaling electronic structure calculation program based on divide-and-conquer method. *Procedia Comput. Sci.* **4**, 1145–1150, DOI: <https://doi.org/10.1016/j.procs.2011.04.122> (2011). Proceedings of the International Conference on Computational Science, ICCS 2011.
171. Nakai, H. *et al.* Divide-and-conquer linear-scaling quantum chemical computations. *The J. Phys. Chem. A* **127**, 589–618, DOI: [10.1021/acs.jpca.2c06965](https://doi.org/10.1021/acs.jpca.2c06965) (2023).
172. Fujii, K. *et al.* Deep variational quantum eigensolver: A divide-and-conquer method for solving a larger problem with smaller size quantum computers. *PRX Quantum* **3**, 010346, DOI: [10.1103/PRXQuantum.3.010346](https://doi.org/10.1103/PRXQuantum.3.010346) (2022).
173. Yoshikawa, T., Takanashi, T. & Nakai, H. Quantum algorithm of the divide-and-conquer unitary coupled cluster method with a variational quantum eigensolver. *J. Chem. Theory Comput.* **18**, 5360–5373, DOI: [10.1021/acs.jctc.2c00602](https://doi.org/10.1021/acs.jctc.2c00602) (2022).
174. Hermes, M. R. & Gagliardi, L. Multiconfigurational self-consistent field theory with density matrix embedding: The localized active space self-consistent field method. *J. Chem. Theory Comput.* **15**, 972–986, DOI: [10.1021/acs.jctc.8b01009](https://doi.org/10.1021/acs.jctc.8b01009) (2019).
175. Hermes, M. R., Pandharkar, R. & Gagliardi, L. Variational localized active space self-consistent field method. *J. Chem. Theory Comput.* **16**, 4923–4937, DOI: [10.1021/acs.jctc.0c00222](https://doi.org/10.1021/acs.jctc.0c00222) (2020).
176. Welborn, M., Tsuchimochi, T. & Van Voorhis, T. Bootstrap embedding: An internally consistent fragment-based method. *The J. Chem. Phys.* **145**, 074102, DOI: [10.1063/1.4960986](https://doi.org/10.1063/1.4960986) (2016).

177. Ye, H.-Z., Ricke, N. D., Tran, H. K. & Van Voorhis, T. Bootstrap embedding for molecules. *J. Chem. Theory Comput.* **15**, 4497–4506, DOI: [10.1021/acs.jctc.9b00529](https://doi.org/10.1021/acs.jctc.9b00529) (2019).
178. Cho, M. *et al.* Quemb: A toolbox for bootstrap embedding calculations of molecular and periodic systems. *The J. Phys. Chem. A* **129**, 6538–6551, DOI: [10.1021/acs.jpca.5c02983](https://doi.org/10.1021/acs.jpca.5c02983) (2025).
179. Otten, M. *et al.* Localized quantum chemistry on quantum computers. *J. Chem. Theory Comput.* **18**, 7205–7217, DOI: [10.1021/acs.jctc.2c00388](https://doi.org/10.1021/acs.jctc.2c00388) (2022).
180. Liu, Y. *et al.* Bootstrap embedding on a quantum computer. *J. Chem. Theory Comput.* **19**, 2230–2247, DOI: [10.1021/acs.jctc.3c00012](https://doi.org/10.1021/acs.jctc.3c00012) (2023).
181. Metzner, W. & Vollhardt, D. Correlated lattice fermions in $d = \infty$ dimensions. *Phys. Rev. Lett.* **62**, 324–327, DOI: [10.1103/PhysRevLett.62.324](https://doi.org/10.1103/PhysRevLett.62.324) (1989).
182. Georges, A., Kotliar, G., Krauth, W. & Rozenberg, M. J. Dynamical mean-field theory of strongly correlated fermion systems and the limit of infinite dimensions. *Rev. Mod. Phys.* **68**, 13–125, DOI: [10.1103/RevModPhys.68.13](https://doi.org/10.1103/RevModPhys.68.13) (1996).
183. Blumenthal, E. Building intuition for dynamical mean-field theory: A simple model and the cavity method (2025). [2507.16654](https://arxiv.org/abs/2507.16654).
184. Knizia, G. & Chan, G. K.-L. Density matrix embedding: A simple alternative to dynamical mean-field theory. *Phys. Rev. Lett.* **109**, 186404, DOI: [10.1103/PhysRevLett.109.186404](https://doi.org/10.1103/PhysRevLett.109.186404) (2012).
185. Wouters, S., Jiménez-Hoyos, C. A., Sun, Q. & Chan, G. K.-L. A practical guide to density matrix embedding theory in quantum chemistry. *J. Chem. Theory Comput.* **12**, 2706–2719, DOI: [10.1021/acs.jctc.6b00316](https://doi.org/10.1021/acs.jctc.6b00316) (2016).
186. Cancès, E., Faulstich, F. M., Kirsch, A., Letournel, E. & Levitt, A. Analysis of density matrix embedding theory around the non-interacting limit. *Commun. on Pure Appl. Math.* **78**, 1359–1410, DOI: <https://doi.org/10.1002/cpa.22244> (2025).
187. Negre, A. *et al.* New perspectives on density-matrix embedding theory (2025). [2503.09881](https://arxiv.org/abs/2503.09881).
188. Kirsopp, J. J. M. *et al.* Quantum computational quantification of protein–ligand interactions. *Int. J. Quantum Chem.* **122**, e26975, DOI: <https://doi.org/10.1002/qua.26975> (2022).
189. Shajan, A. *et al.* Toward quantum-centric simulations of extended molecules: Sample-based quantum diagonalization enhanced with density matrix embedding theory. *J. Chem. Theory Comput.* **21**, 6801–6810, DOI: [10.1021/acs.jctc.5c00114](https://doi.org/10.1021/acs.jctc.5c00114) (2025).
190. Patra, A. K. *et al.* Quantum simulation of ligand-like molecules through sample-based quantum diagonalization in density matrix embedding framework (2026). [2511.22158](https://arxiv.org/abs/2511.22158).
191. Wang, Q. *et al.* Sample-based quantum diagonalization as parallel fragment solver for the localized active space self-consistent field method (2025). [2512.14936](https://arxiv.org/abs/2512.14936).
192. Bierman, J. & Liu, Y. Towards utility-scale electronic structure with sample-based quantum bootstrap embedding. *Digit. Discov.* –, DOI: [10.1039/D5DD00416K](https://doi.org/10.1039/D5DD00416K) (2026).
193. Shajan, A. *et al.* Molecular quantum computations on a protein (2026). [2512.17130](https://arxiv.org/abs/2512.17130).
194. Chung, L. W. *et al.* The oniom method and its applications. *Chem. Rev.* **115**, 5678–5796, DOI: [10.1021/cr5004419](https://doi.org/10.1021/cr5004419) (2015).
195. Yamamoto, K. *et al.* Quantum-hpc hybrid computation of biomolecular excited-state energies (2026). [2601.15677](https://arxiv.org/abs/2601.15677).
196. Motta, M., Sung, K. J., Whaley, K. B., Head-Gordon, M. & Shee, J. Bridging physical intuition and hardware efficiency for correlated electronic states: the local unitary cluster jastrow ansatz for electronic structure. *Chem. Sci.* **14**, 11213–11227, DOI: [10.1039/D3SC02516K](https://doi.org/10.1039/D3SC02516K) (2023).
197. Magoulas, I. & Evangelista, F. A. Linear-scaling quantum circuits for computational chemistry. *J. Chem. Theory Comput.* **19**, 4815–4821, DOI: [10.1021/acs.jctc.3c00376](https://doi.org/10.1021/acs.jctc.3c00376) (2023).
198. Nam, Y. *et al.* Ground-state energy estimation of the water molecule on a trapped-ion quantum computer. *npj Quantum Inf.* **6**, 33, DOI: [10.1038/s41534-020-0259-3](https://doi.org/10.1038/s41534-020-0259-3) (2020).
199. Hong, C.-L. *et al.* Accurate and efficient quantum computations of molecular properties using daubechies wavelet molecular orbitals: A benchmark study against experimental data. *PRX Quantum* **3**, 020360, DOI: [10.1103/PRXQuantum.3.020360](https://doi.org/10.1103/PRXQuantum.3.020360) (2022).
200. Lenz, F. The ratio of proton and electron masses. *Phys. Rev.* **82**, 554–554, DOI: [10.1103/PhysRev.82.554.2](https://doi.org/10.1103/PhysRev.82.554.2) (1951).

201. Tiesinga, E., Mohr, P. J., Newell, D. B. & Taylor, B. N. Codata recommended values of the fundamental physical constants: 2018. *Rev. Mod. Phys.* **93**, 025010, DOI: [10.1103/RevModPhys.93.025010](https://doi.org/10.1103/RevModPhys.93.025010) (2021).
202. Pauli, W. Über den zusammenhang des abschlusses der elektronengruppen im atom mit der komplexstruktur der spektren. *Zeitschrift für Physik* **31**, 765–783, DOI: [10.1007/BF02980631](https://doi.org/10.1007/BF02980631) (1925).
203. Szabo, A. & Ostlund, N. S. *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory* (Dover Publications, Mineola, New York, 1996). See Chapter 3.1.2, p. 114.
204. Møller, C. & Plesset, M. S. Note on an approximation treatment for many-electron systems. *Phys. Rev.* **46**, 618–622, DOI: [10.1103/PhysRev.46.618](https://doi.org/10.1103/PhysRev.46.618) (1934).
205. Sakurai, J. J. & Napolitano, J. *Modern Quantum Mechanics* (Cambridge University Press, Cambridge, 2017), revised edn. See Chapter 5.1, p. 285.
206. Jensen, F. *Introduction to Computational Chemistry* (Wiley, 2013). See Chapter 4.8.1, p. 164.
207. Bartlett, R. J. & Musiał, M. Coupled-cluster theory in quantum chemistry. *Rev. Mod. Phys.* **79**, 291–352, DOI: [10.1103/RevModPhys.79.291](https://doi.org/10.1103/RevModPhys.79.291) (2007).
208. Cremer, D. Møller–plesset perturbation theory: from small molecule methods to methods for thousands of atoms. *WIREs Comput. Mol. Sci.* **1**, 509–530, DOI: <https://doi.org/10.1002/wcms.58> (2011).
209. Levine, I. N. *Quantum Chemistry* (Pearson Prentice Hall, 2014), 7th edn. See Chapter 16.4, p. 547.
210. Jensen, F. *Introduction to Computational Chemistry* (Wiley, 2013). See Chapter 4.9, p. 169.
211. Bartlett, R. J. & Musiał, M. Coupled-cluster theory in quantum chemistry. *Rev. Mod. Phys.* **79**, 291–352, DOI: [10.1103/RevModPhys.79.291](https://doi.org/10.1103/RevModPhys.79.291) (2007).
212. Crawford, T. D. & Schaefer III, H. F. *An Introduction to Coupled Cluster Theory for Computational Chemists*, 33–136 (John Wiley & Sons, Ltd, 2000).
213. Slater, J. C. The theory of complex spectra. *Phys. Rev.* **34**, 1293–1322, DOI: [10.1103/PhysRev.34.1293](https://doi.org/10.1103/PhysRev.34.1293) (1929).
214. Szabo, A. & Ostlund, N. S. *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory* (Dover Publications, Mineola, New York, 1996). See Chapter 4.1, p. 235.
215. Szabo, A. & Ostlund, N. S. *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory* (Dover Publications, Mineola, New York, 1996). See Chapter 3.3.2, p. 128.
216. Jensen, F. *Introduction to Computational Chemistry* (Wiley, 2007), 2 edn. See Sec. 4.2.2, p. 141.
217. Slater, J. C. Atomic shielding constants. *Phys. Rev.* **36**, 57–64, DOI: [10.1103/PhysRev.36.57](https://doi.org/10.1103/PhysRev.36.57) (1930).
218. Boys, S. F. Electronic wave functions - i. a general method of calculation for the stationary states of any molecular system. *Proc. Royal Soc. London. A. Math. Phys. Sci.* **200**, 542–554, DOI: [10.1098/rspa.1950.0036](https://doi.org/10.1098/rspa.1950.0036) (1950).
219. Hehre, W. J., Stewart, R. F. & Pople, J. A. Self-consistent molecular-orbital methods. i. use of gaussian expansions of slater-type atomic orbitals. *The J. Chem. Phys.* **51**, 2657–2664, DOI: [10.1063/1.1672392](https://doi.org/10.1063/1.1672392) (1969).
220. Hehre, W. J., Ditchfield, R. & Pople, J. A. Self-consistent molecular orbital methods. xii. further extensions of gaussian-type basis sets for use in molecular orbital studies of organic molecules. *The J. Chem. Phys.* **56**, 2257–2261, DOI: [10.1063/1.1677527](https://doi.org/10.1063/1.1677527) (1972).
221. Sehnal, D. *et al.* Mol* viewer: modern web app for 3d visualization and analysis of large biomolecular structures. *Nucleic Acids Res.* DOI: [10.1093/nar/gkab31](https://doi.org/10.1093/nar/gkab31) (2021).
222. Romero, J. *et al.* Strategies for quantum computing molecular energies using the unitary coupled cluster ansatz. *Quantum Sci. Technol.* **4**, 014008, DOI: [10.1088/2058-9565/aad3e4](https://doi.org/10.1088/2058-9565/aad3e4) (2018).
223. Suzuki, M. Generalized trotter’s formula and systematic approximants of exponential operators and inner derivations with applications to many-body problems. *Commun. Math. Phys.* **51**, 183–190, DOI: [10.1007/BF01609348](https://doi.org/10.1007/BF01609348) (1976).
224. Jordan, P. & Wigner, E. Über das paulische äquivalenzverbot. *Zeitschrift für Physik* **47**, 631–651, DOI: [10.1007/BF01331938](https://doi.org/10.1007/BF01331938) (1928).
225. Magoulas, I. & Evangelista, F. A. Cnot-efficient circuits for arbitrary rank many-body fermionic and qubit excitations. *J. Chem. Theory Comput.* **19**, 822–836, DOI: [10.1021/acs.jctc.2c01016](https://doi.org/10.1021/acs.jctc.2c01016) (2023).
226. Davidson, E. R. The iterative calculation of a few of the lowest eigenvalues and corresponding eigenvectors of large real-symmetric matrices. *J. Comput. Phys.* **17**, 87–94, DOI: [https://doi.org/10.1016/0021-9991\(75\)90065-0](https://doi.org/10.1016/0021-9991(75)90065-0) (1975).

227. Sakurai, J. J. *Modern Quantum Mechanics* (Addison-Wesley, 1993), revised edition edn.
228. Wesolowski, T. A. & Warshel, A. Frozen density functional approach for ab initio calculations of solvated molecules. *The J. Phys. Chem.* **97**, 8050–8053, DOI: [10.1021/j100132a040](https://doi.org/10.1021/j100132a040) (1993).
229. Zhang, D. W. & Zhang, J. Z. H. Molecular fractionation with conjugate caps for full quantum mechanical calculation of protein–molecule interaction energy. *The J. Chem. Phys.* **119**, 3599–3605, DOI: [10.1063/1.1591727](https://doi.org/10.1063/1.1591727) (2003).
230. Gadre, S. R., Shirsat, R. N. & Limaye, A. C. Molecular tailoring approach for simulation of electrostatic properties. *The J. Phys. Chem.* **98**, 9165–9169, DOI: [10.1021/j100088a013](https://doi.org/10.1021/j100088a013) (1994).
231. Ekert, A. & Knight, P. L. Entangled quantum systems and the schmidt decomposition. *Am. J. Phys.* **63**, 415–423, DOI: [10.1119/1.17904](https://doi.org/10.1119/1.17904) (1995).
232. MacDonald, J. K. L. Successive approximations by the rayleigh-ritz variation method. *Phys. Rev.* **43**, 830–833, DOI: [10.1103/PhysRev.43.830](https://doi.org/10.1103/PhysRev.43.830) (1933).
233. Van Rossum, G. Python tutorial. Tech. Rep. CS-R9526, Centrum voor Wiskunde en Informatica (CWI) (1995).
234. Python Software Foundation. *Python Language Reference*, v3.11 (2022). Accessed: 2025-11-18.
235. IQM Client Developers. Iqm client v31.1.1. <https://pypi.org/project/iqm-client/32.1.1/> (2025). Accessed: 2026-03-09.
236. Javadi-Abhari, A. *et al.* Quantum computing with qiskit (2024). [2405.08810](https://arxiv.org/abs/2405.08810).
237. Qiskit Community. ffsim: Fermionic fast simulator v0.0.56. <https://github.com/qiskit-community/ffsim/releases/tag/v0.0.56> (2025). Accessed: 2026-03-09.
238. Qiskit Community. Qiskit addon sqd v0.10.0. <https://github.com/Qiskit/qiskit-addon-sqd/releases/tag/0.10.0> (2025). Accessed: 2026-03-09.
239. Sun, Q. Libcint: An efficient general integral library for gaussian basis functions. *J. Comput. Chem.* **36**, 1664–1671, DOI: <https://doi.org/10.1002/jcc.23981> (2015).
240. Sun, Q. *et al.* Pyscf: the python-based simulations of chemistry framework. *WIREs Comput. Mol. Sci.* **8**, e1340, DOI: <https://doi.org/10.1002/wcms.1340> (2018).
241. Sun, Q. *et al.* Recent developments in the pyscf program package. *The J. Chem. Phys.* **153**, 024109, DOI: [10.1063/5.0006074](https://doi.org/10.1063/5.0006074) (2020).
242. Harris, C. R. *et al.* Array programming with NumPy. *Nature* **585**, 357–362, DOI: [10.1038/s41586-020-2649-2](https://doi.org/10.1038/s41586-020-2649-2) (2020).
243. Virtanen, P. *et al.* SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nat. Methods* **17**, 261–272, DOI: [10.1038/s41592-019-0686-2](https://doi.org/10.1038/s41592-019-0686-2) (2020).
244. Bradbury, J. *et al.* JAX: composable transformations of Python+NumPy programs. <https://github.com/google/jax/releases/tag/jax-v0.8.1> (2018). Accessed: 2026-03-09.
245. The pandas development team. pandas-dev/pandas: Pandas, DOI: [10.5281/zenodo.3509134](https://doi.org/10.5281/zenodo.3509134) (2020).
246. Kawashima, Y. *et al.* Optimizing electronic structure simulations on a trapped-ion quantum computer using problem decomposition. *Commun. Phys.* **4**, 245, DOI: [10.1038/s42005-021-00751-9](https://doi.org/10.1038/s42005-021-00751-9) (2021).
247. Olsen, J., Roos, B. O., Joergensen, P. & Jensen, H. J. A. Determinant based configuration interaction algorithms for complete and restricted configuration interaction spaces. *The J. Chem. Phys.* **89**, 2185–2192, DOI: [10.1063/1.455063](https://doi.org/10.1063/1.455063) (1988).
248. Listro, R. *et al.* Urea-based anticancer agents: Exploring 100 years of research with an eye to the future. *Front. Chem.* **10**, DOI: [10.3389/fchem.2022.995351](https://doi.org/10.3389/fchem.2022.995351) (2022).
249. Ghosh, A. K. & Brindisi, M. Urea derivatives in modern drug discovery and medicinal chemistry. *J. Medicinal Chem.* **63**, 2751–2788, DOI: [10.1021/acs.jmedchem.9b01541](https://doi.org/10.1021/acs.jmedchem.9b01541) (2020).
250. Eads, J. R. *et al.* Phase i clinical trial of temozolomide and methoxyamine (trc-102), an inhibitor of base excision repair, in patients with advanced solid tumors. *Investig. New Drugs* **39**, 142–151, DOI: [10.1007/s10637-020-00962-x](https://doi.org/10.1007/s10637-020-00962-x) (2021).
251. Caimi, P. F. *et al.* Phase i clinical trial of the base excision repair inhibitor methoxyamine in combination with fludarabine for patients with advanced hematologic malignancies. *Oncotarget* **8**, 79864–79875, DOI: [10.18632/oncotarget.20094](https://doi.org/10.18632/oncotarget.20094) (2017).
252. Cady, S. D. *et al.* Structure of the amantadine binding site of influenza m2 proton channels in lipid bilayers. *Nature* **463**, 689–692, DOI: [10.1038/nature08722](https://doi.org/10.1038/nature08722) (2010).

253. U.S. Food and Drug Administration. Amantadine hydrochloride substance record (2026). Accessed: March 17, 2026.
254. Small, D. W. & Head-Gordon, M. A fusion of the closed-shell coupled cluster singles and doubles method and valence-bond theory for bond breaking. *The J. Chem. Phys.* **137**, 114103, DOI: [10.1063/1.4751485](https://doi.org/10.1063/1.4751485) (2012).
255. Barone, V. *et al.* Computational molecular spectroscopy. *Nat. Rev. Methods Primers* **1**, 38, DOI: [10.1038/s43586-021-00034-1](https://doi.org/10.1038/s43586-021-00034-1) (2021).
256. Chmiela, S. *et al.* Machine learning of accurate energy-conserving molecular force fields. *Sci. Adv.* **3**, e1603015, DOI: [10.1126/sciadv.1603015](https://doi.org/10.1126/sciadv.1603015) (2017).
257. Qiu, C., Brinck, T. & Wang, J. Modeling the potential energy surface by force fields for heterogeneous catalysis: classification, applications, and challenges. *Chem. Sci.* **16**, 21269–21297, DOI: [10.1039/D5SC02715B](https://doi.org/10.1039/D5SC02715B) (2025).
258. Kenneth M. Merz, J. *et al.* Crossing the 12,000-atom barrier with heterogeneous quantum-classical supercomputing: quantum chemistry of protein-ligand complexes (2026). [2605.01138](https://arxiv.org/abs/2605.01138).
259. Anurag, K. S. V. *et al.* Bridging the nisq and fault-tolerant regimes: Generative-ml-assisted quantum selected ci for molecular simulations (2026). [2606.30551](https://arxiv.org/abs/2606.30551).
260. IQM Client Developers. Instruction — iqm client api documentation. https://iqm-finland.github.io/iqm-client/api/iqm_client.models.Instruction.html (2025). Last updated: 2025-02-25. Accessed: 2026-03-21.

Supplementary Information

A List of Symbols

Table 11. List of symbols used throughout the manuscript.

Symbol	Description	Symbol	Description
$\mathbb{H}$	Hilbert space	$\mathbb{S}$	Symmetry space
ϵ_s	Samples per batch	$\mathbb{S}_{\text{samp}}$	Sampling space
$\mathbb{S}_{\text{post-sel}}$	Post-selected subspace	$\mathbb{S}_{\text{post-cr}}$	Post-configuration recovery space
$\mathbb{S}_{\text{sub}}$	Diagonalization subspace	N_{iter}	Number of DMET iterations
N_{shots}	Number of measurement shots	$(N_{\text{orb}}, N_{\text{ele}})$	Number of molecular orbitals and electrons
$N_{\alpha}(N_{\beta})$	Number of electrons in $\alpha(\beta)$ -spin	$\vec{\theta}$	Variational parameters
N_Q	Number of qubits	G_D	Quantum circuit gate depth
Move	Number of MOVE gates	R	Number of R gates
CZ	Number of controlled-Z gates	τ_{qpu}	QPU execution time (s)
G_{q_1}	Number of single-qubit gates	G_{q_2}	Number of two-qubit gates
G_T	Total number of gates	η_{sym}	Symmetry-to-Hilbert space ratio
$\eta_{\text{post-cr}}$	Post-CR to symmetry space ratio	η_{sub}	Diagonalization Subspace to symmetry space ratio
r_{min}	Bond length at minimum energy	θ_{min}	Bond angle at minimum energy
μ_{glob}	Global chemical potential	$ \psi_{\text{in}}\rangle$	Input quantum state
$\hat{H}_{\text{ele}}$	Electronic Hamiltonian	$\hat{H}_{\text{sub}}$	Projected Hamiltonian
$ \Phi_0\rangle$	HF reference determinant	$ \Phi_i^a\rangle$	Single excitation determinant
$ \Phi_{ij}^{ab}\rangle$	Double excitation determinant	c_i	CI coefficients
$\hat{T}$	Cluster operator	$\hat{T}_1, \hat{T}_2$	Single and double excitation operators
t_i^a, t_{ij}^{ab}	Cluster amplitudes	$\hat{J}$	Jastrow operator
$\hat{n}_{p\sigma}$	Number operator	$U(\vec{\theta})$	Unitary operator
μ	LUCJ layer index	$\hat{a}_p^\dagger, \hat{a}_p$	Creation and annihilation operators
h_{pq}	One-electron integrals	h_{pqrs} or V_{pqrs}	Two-electron integrals
$D_{\mu\nu}$	One-particle reduced density matrix	Γ_{pq}^{rs}	Two-particle reduced density matrix
$\mathbb{H}_{\text{lin}}$	One-particle Hilbert space	$\mathbb{F}(\mathbb{H}_{\text{lin}})$	Fermionic Fock space
$\mathbf{x}_i$	Sampled bitstring	γ	Overlap with target state

B IQM Sirius QPU Details

All quantum computations reported in this work were performed on the IQM Sirius quantum processing unit (QPU), accessed via the IQM Resonance cloud platform. IQM Sirius is a superconducting transmon qubit-based QPU supporting up to 24 qubits arranged in a star lattice topology (STAR 24). All qubits are coupled to a single shared central computational resonator (COMPR1), forming a hub-and-spoke connectivity in which any qubit can interact with any other qubit through resonator-mediated operations, providing effective all-to-all connectivity without the need for SWAP gates. The device was calibrated on 4 December 2025 at 10:34 IST (GMT+5:30); all calculations in this work were executed within the window 4 Dec 2025 14:49 IST – 6 Dec 2025 05:48 IST, which falls within the validity period of this single calibration. The active qubit set comprised 16 qubits: QB1, QB2, QB3, QB4, QB5, QB8, QB9, QB10, QB11, QB13, QB15, QB16, QB17, QB19, QB20, and QB21.

Hardware Topology

The physical topology of the IQM Sirius QPU is shown in Figure B1. Each active qubit (filled circle) is connected through a coupler element (filled square) to the central computational resonator COMPR1 (horizontal bar). Dashed outlines indicate the inactive couplers and qubits (QB6, QB7, QB12, QB14, QB18, QB22–QB24) in this configuration.




Figure B1. Hardware topology of the IQM Sirius QPU (STAR 24). All 16 active qubits (filled circles) are connected to the central computational resonator COMPR1 (horizontal bar) through individual couplers (filled squares), providing effective all-to-all qubit connectivity. Dashed outlines denote inactive couplers and qubits.

Native Gate Set

The native gate set of IQM Sirius consists of the following operations, as defined in the IQM Client API²⁶⁰:

- **PRX**(θ, ϕ) — *phased rotation* (universal single-qubit gate):

$$\text{PRX}(\theta, \phi) = e^{-i\pi\theta(X\cos 2\pi\phi + Y\sin 2\pi\phi)} = R_Z(2\pi\phi)R_X(2\pi\theta)R_Z^\dagger(2\pi\phi), \quad (29)$$

where θ is the rotation angle and ϕ is the phase (azimuthal) angle, both in units of full turns; X and Y are the Pauli matrices, as the **R** gate in the main text. The default pulse implementation uses a DRAG Cosine Rise–Fall envelope.

- **CZ** — native two-qubit controlled-Z gate, represented as **CZ** gate in the main text:

$$\text{CZ} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix}$$

- **MOVE** — Resonator-mediated excitation-exchange gate that swaps the populations of the qubit–resonator states $|01\rangle \leftrightarrow |10\rangle$ within the invariant subspace $\mathcal{S} = \text{span}\{|00\rangle, |01\rangle, |10\rangle\}$. Its action can be written as:

$$\text{MOVE}_{\mathcal{S}} = \frac{1}{\mathcal{N}}(|00\rangle\langle 00| + a|10\rangle\langle 01| + a^{-1}|01\rangle\langle 10|),$$

where $\mathcal{N}$ is the normalization constant and a is an unknown complex phase that cancels when the gate is applied twice. The action of MOVE on the orthogonal subspace (in particular $|11\rangle$) is undefined. This operation enables excitation shuttling between a qubit and the resonator COMPRI, providing effective all-to-all connectivity, represented as **Move** gate in the main text.

- **MEASURE** — Measurement in the computational (Z) basis.

Hardware Characterisation Metrics

Table B1 presents device-level summary statistics and Table B2 lists the full per-qubit metrics; both were derived from the Calibration data obtained as a JSON export on 4 December 2025 at 10:34 IST (GMT+5:30).

Table B1. IQM Sirius QPU benchmarking summary (Calibration: 4 December 2025, 10:34 IST). Values are computed from the IQM Resonance metrics export and cross-validated against the IQM Resonance dashboard (all 14 entries, including 7 averages and 7 medians, agree to within dashboard rounding).

Property	Average	Median
T_1 time	28.69 μs	28.71 μs
T_2^* time (Ramsey)	17.11 μs	18.62 μs
T_2^{echo} time	27.68 μs	27.98 μs
PRX gate fidelity	99.79%	99.78%
Single-qubit readout fidelity	97.84%	98.49%
CZ gate fidelity	98.42%	98.54%
MOVE-MOVE fidelity	98.79%	98.93%

Metric definitions.

- T_1 : is the energy relaxation time.
- T_2^* and T_2^{echo} : are the dephasing times measured by Ramsey and spin-echo sequences, respectively.
- PRX gate fidelity: is obtained from DRAG-calibrated randomised benchmarking.
- Single-qubit readout fidelity: is the symmetric average state-assignment fidelity from single-shot readout.
- CZ gate fidelity: is derived from interleaved randomised benchmarking (crf_crf for 13 pairs; tgss for QB2, QB10, and QB16, indicated by †).
- MOVE-MOVE fidelity: is the total fidelity of a MOVE–MOVE sequence obtained from interleaved RB (crf_crf, $n_{\text{interleaved}} = 2$).

Summary averages for T_1 and T_2^* include the COMPRI resonator as the 17th entry; T_2^{echo} , CZ, and MOVE statistics are over the 16 active qubits/pairs only.

The majority of qubits show stable single-qubit performance, with the PRX gate fidelities uniformly above 99.5% and readout fidelities above 97%. Two qubits, QB13 and QB15, exhibit notably reduced performance: QB13 has the shortest T_2^* (2.29 μs) and QB15 the lowest readout fidelity (91.05%), with correspondingly reduced CZ gate fidelities (97.28% and 97.17%, respectively). These qubits were retained in all circuits; their reduced coherence and readout performance represent a conservative upper bound on hardware noise in the reported results. The two-qubit gate performance is otherwise uniformly stable, with CZ and MOVE-MOVE fidelities exceeding 98% for 14 of 16 qubit–resonator pairs.

Table B2. Per-qubit hardware metrics for the 16 active qubits of IQM Sirius (Calibration: 4 December 2025, 10:34 IST). The COMPR1 resonator row lists its independently measured T_1 and T_2^* and is included in the device-level averages. CZ fidelities marked † are obtained from `tgss` protocol rather than `crf_crf`. All values are rounded to two decimal places.

Qubit	T_1 (μ s)	T_2^* (μ s)	T_2^{echo} (μ s)	PRX (%)	Readout (%)	CZ (%)	MOVE (%)
COMPR1	16.60	29.24	—	—	—	—	—
QB1	22.45	18.62	28.99	99.85	98.22	98.14	98.64
QB2	26.79	28.22	43.85	99.76	98.62	98.11 [†]	99.24
QB3	27.17	5.93	9.24	99.84	98.83	98.87	98.93
QB4	19.87	25.86	31.08	99.91	98.60	98.86	99.21
QB5	39.42	23.77	29.50	99.94	98.75	98.96	99.19
QB8	31.02	6.38	19.66	99.91	98.25	98.64	98.01
QB9	30.12	4.93	16.13	99.77	99.12	98.96	98.87
QB10	28.71	8.03	24.50	99.76	97.72	98.32 [†]	99.04
QB11	38.30	29.57	35.10	99.77	98.72	98.71	98.75
QB13	18.95	2.29	6.28	99.53	95.58	97.28	98.49
QB15	18.19	14.04	26.98	99.64	91.05	97.17	97.96
QB16	33.96	33.05	39.76	99.77	98.55	98.32 [†]	98.93
QB17	35.83	4.85	16.38	99.78	98.80	98.76	98.96
QB19	20.48	9.22	23.87	99.76	98.42	98.44	98.95
QB20	32.79	19.55	40.63	99.82	98.22	98.98	99.11
QB21	47.13	27.28	50.90	99.81	98.02	98.24	98.42
Mean	28.69	17.11	27.68	99.79	97.84	98.42	98.79
Median	28.71	18.62	27.98	99.78	98.49	98.54	98.93

C Sample-based Quantum Diagonalization: Molecule Geometries & QPU Runtimes

This appendix reports, in Table C1, the Cartesian geometric coordinates (optimized at the CCSD level) and the corresponding quantum processing unit (QPU) execution runtimes for all molecular systems studied using sample-based quantum diagonalization (SQD) workflows in Section 4.1. The tabulated geometries represent fixed nuclear configurations consistently employed across SQD(LUCJ) and SQD(LCNot–UCCSD) energy evaluations, as well as in the associated post-Hartree–Fock classical benchmark calculations at the CCSD and FCI levels. For each molecule, QPU runtimes are reported from three independent executions performed at a fixed number of measurement shots, enabling transparent comparison of execution-time variability across workflows. No additional structural optimization was performed within either the quantum or classical computational pipelines.

Table C1. Cartesian geometries and QPU execution runtimes for SQD workflows on IQM Sirius quantum hardware. All QPU measurements were performed using $N_{\text{shots}} = 10,000$. Here, τ'_{sqd} denotes the QPU execution time for SQD using the LUCJ ansatz, and τ''_{sqd} denotes the QPU execution time for SQD using the LCNot–UCCSD ansatz, both reported in seconds (s).

Molecule	Atom	Geometry (Å)			Run	τ'_{sqd} (s)	τ''_{sqd} (s)
		x	y	z			
Hydrogen (H ₂)	H	-8.0563	4.7154	0.0000	1	5.9399	6.0922
	H	-8.5983	4.7845	-0.3333	2	5.8913	5.9191
					3	5.8319	6.3319
Lithium Hydride (LiH)	Li	-1.3452	3.1070	0.0000	1	6.5436	10.5604
	H	-0.5971	2.1184	1.0889	2	6.5822	10.8045
					3	6.5957	10.6617
Beryllium Hydride (BeH ₂)	Be	-0.6141	3.0632	0.0000	1	6.9852	18.7704
	H	-1.7489	3.2078	-0.6978	2	7.0202	18.5576
	H	-0.5883	3.9542	1.0005	3	6.6800	18.5296
Water (H ₂ O)	O	-8.6252	5.1579	-0.0147	1	6.9381	23.4240
	H	-7.7020	5.4505	-0.0474	2	6.8318	14.9879
	H	-9.0559	5.8062	0.5626	3	6.9861	14.8362
Ammonia (NH ₃)	N	-9.0297	4.0906	0.0167	1	7.1387	28.1490
	H	-8.5642	3.4395	0.6473	2	7.2257	27.2112
	H	-8.4095	4.1973	-0.7848	3	7.2158	27.1455
	H	-9.8553	3.6042	-0.3299			

D Potential Energy Surface Scans: Molecule Geometries & QPU Runtimes

This appendix reports the Cartesian geometric coordinates and corresponding QPU execution runtimes for all molecular systems used in the one-dimensional potential energy surface (1D-PES) scans discussed in Section 4.2, as well as the geometry construction and runtime characterization for the two-dimensional PES (2D-PES) of the H₂O molecule discussed in Section 4.3.

For each 1D-PES system, the tabulated geometries represent fixed nuclear configurations employed uniformly across all quantum computing–based energy evaluations using SQD(LUCJ), as well as in the associated post-Hartree–Fock classical benchmark calculations at the CCSD and FCI levels. In addition, each table includes the measured QPU runtimes, τ_{qpu} , obtained from multiple independent executions at a fixed shot count, providing a transparent record of execution variability across geometries. No further structural optimization was performed in either the quantum or classical workflows. Specifically, Table D1 reports H₂ geometries and QPU runtimes in the STO-3G basis, Table D2 reports H₂ in the 6-31G basis, Table D3 reports HeH⁺ in the STO-3G basis, Table D4 reports HeH⁺ in the 6-31G basis, Table D5 reports LiH in the STO-3G basis, and Table D6 reports BeH₂ in the STO-3G basis.

Table D1. Cartesian geometries used for the 1D bond-length PES scan of H₂ in STO-3G basis, along with corresponding QPU runtimes, τ_{qpu} (s), from three independent executions at $N_{\text{shots}} = 10,000$. The interatomic distance r is given in angstrom (Å). All geometries are aligned along the z -axis with one hydrogen atom fixed at the origin.

S.No.	r (Å)	Geometry (Cartesian coordinates in Å)	$\tau_{\text{qpu}}^{(1)}$ (s)	$\tau_{\text{qpu}}^{(2)}$ (s)	$\tau_{\text{qpu}}^{(3)}$ (s)
1	0.3000	H (0.0, 0.0, 0.0); H (0.0, 0.0, 0.3000)	7.2979	6.1516	5.9550
2	0.3627	H (0.0, 0.0, 0.0); H (0.0, 0.0, 0.3627)	5.9553	6.0366	6.4405
3	0.4255	H (0.0, 0.0, 0.0); H (0.0, 0.0, 0.4255)	5.9228	6.1361	6.1373
4	0.4882	H (0.0, 0.0, 0.0); H (0.0, 0.0, 0.4882)	5.8478	6.3280	6.1624
5	0.5509	H (0.0, 0.0, 0.0); H (0.0, 0.0, 0.5509)	6.0579	5.9691	5.9320
6	0.6137	H (0.0, 0.0, 0.0); H (0.0, 0.0, 0.6137)	5.9747	6.0296	5.9973
7	0.6764	H (0.0, 0.0, 0.0); H (0.0, 0.0, 0.6764)	5.9981	5.9824	5.9425
8	0.7391	H (0.0, 0.0, 0.0); H (0.0, 0.0, 0.7391)	5.7185	6.1526	6.0083
9	0.8019	H (0.0, 0.0, 0.0); H (0.0, 0.0, 0.8019)	6.1911	6.6158	6.0033
10	0.8646	H (0.0, 0.0, 0.0); H (0.0, 0.0, 0.8646)	5.9217	6.6960	6.1015
11	0.9273	H (0.0, 0.0, 0.0); H (0.0, 0.0, 0.9273)	5.9897	6.1654	5.9972
12	0.9901	H (0.0, 0.0, 0.0); H (0.0, 0.0, 0.9901)	6.0618	6.0720	5.9927
13	1.0528	H (0.0, 0.0, 0.0); H (0.0, 0.0, 1.0528)	6.1565	6.0630	5.9105
14	1.1155	H (0.0, 0.0, 0.0); H (0.0, 0.0, 1.1155)	5.9083	5.9779	6.1222
15	1.1783	H (0.0, 0.0, 0.0); H (0.0, 0.0, 1.1783)	7.0812	6.1627	6.0876
16	1.2410	H (0.0, 0.0, 0.0); H (0.0, 0.0, 1.2410)	6.1679	5.9223	6.2450
17	1.5854	H (0.0, 0.0, 0.0); H (0.0, 0.0, 1.5854)	6.0374	5.9551	5.7936
18	1.9299	H (0.0, 0.0, 0.0); H (0.0, 0.0, 1.9299)	5.9870	5.9465	5.8130
19	2.2743	H (0.0, 0.0, 0.0); H (0.0, 0.0, 2.2743)	6.1093	6.1022	6.2906
20	2.6188	H (0.0, 0.0, 0.0); H (0.0, 0.0, 2.6188)	5.9763	6.3109	6.1536
21	2.9632	H (0.0, 0.0, 0.0); H (0.0, 0.0, 2.9632)	5.9698	6.0650	5.9278
22	3.3077	H (0.0, 0.0, 0.0); H (0.0, 0.0, 3.3077)	6.0132	5.8522	6.0788
23	3.6521	H (0.0, 0.0, 0.0); H (0.0, 0.0, 3.6521)	6.0385	6.1290	5.9444
24	3.9966	H (0.0, 0.0, 0.0); H (0.0, 0.0, 3.9966)	5.9143	6.0414	6.0460
25	4.3410	H (0.0, 0.0, 0.0); H (0.0, 0.0, 4.3410)	6.0281	6.1387	6.0042

Table D2. Cartesian geometries used for the 1D bond-length PES scan of H₂ in 6-31G basis, along with corresponding QPU runtimes, τ_{qpu} (s), from three independent executions at $N_{\text{shots}} = 10,000$. The interatomic distance r is given in angstrom (Å). All geometries are aligned along the z-axis with one hydrogen atom fixed at the origin.

S.No.	r (Å)	Geometry (Cartesian coordinates in Å)	$\tau_{\text{qpu}}^{(1)}$ (s)	$\tau_{\text{qpu}}^{(2)}$ (s)	$\tau_{\text{qpu}}^{(3)}$ (s)
1	0.3000	H (0.0, 0.0, 0.0); H (0.0, 0.0, 0.3000)	6.1700	5.8312	6.4196
2	0.3627	H (0.0, 0.0, 0.0); H (0.0, 0.0, 0.3627)	6.1612	5.9671	6.0942
3	0.4255	H (0.0, 0.0, 0.0); H (0.0, 0.0, 0.4255)	6.2658	6.3419	5.9834
4	0.4882	H (0.0, 0.0, 0.0); H (0.0, 0.0, 0.4882)	6.1994	6.2379	15.9657
5	0.5509	H (0.0, 0.0, 0.0); H (0.0, 0.0, 0.5509)	6.2950	6.0310	6.1382
6	0.6137	H (0.0, 0.0, 0.0); H (0.0, 0.0, 0.6137)	6.1667	6.4015	6.0511
7	0.6764	H (0.0, 0.0, 0.0); H (0.0, 0.0, 0.6764)	6.6754	6.0823	6.0986
8	0.7391	H (0.0, 0.0, 0.0); H (0.0, 0.0, 0.7391)	6.1571	6.2088	6.1755
9	0.8019	H (0.0, 0.0, 0.0); H (0.0, 0.0, 0.8019)	6.0316	6.2185	6.2606
10	0.8646	H (0.0, 0.0, 0.0); H (0.0, 0.0, 0.8646)	5.9971	6.2708	6.2557
11	0.9273	H (0.0, 0.0, 0.0); H (0.0, 0.0, 0.9273)	6.0303	6.3036	6.0418
12	0.9901	H (0.0, 0.0, 0.0); H (0.0, 0.0, 0.9901)	6.0956	6.1590	6.0282
13	1.0528	H (0.0, 0.0, 0.0); H (0.0, 0.0, 1.0528)	6.0287	6.2573	6.2267
14	1.1155	H (0.0, 0.0, 0.0); H (0.0, 0.0, 1.1155)	6.0695	6.4471	6.0187
15	1.1783	H (0.0, 0.0, 0.0); H (0.0, 0.0, 1.1783)	6.3106	6.1846	6.2552
16	1.2410	H (0.0, 0.0, 0.0); H (0.0, 0.0, 1.2410)	6.0697	6.0810	6.2068
17	1.5854	H (0.0, 0.0, 0.0); H (0.0, 0.0, 1.5854)	6.2628	5.9745	6.1503
18	1.9299	H (0.0, 0.0, 0.0); H (0.0, 0.0, 1.9299)	6.1822	6.6125	6.2142
19	2.2743	H (0.0, 0.0, 0.0); H (0.0, 0.0, 2.2743)	6.0858	6.1836	6.8559
20	2.6188	H (0.0, 0.0, 0.0); H (0.0, 0.0, 2.6188)	6.2074	6.0729	6.1786
21	2.9632	H (0.0, 0.0, 0.0); H (0.0, 0.0, 2.9632)	6.2356	6.2927	6.3299
22	3.3077	H (0.0, 0.0, 0.0); H (0.0, 0.0, 3.3077)	6.6367	6.2286	6.4128
23	3.6521	H (0.0, 0.0, 0.0); H (0.0, 0.0, 3.6521)	7.2147	6.2212	6.2125
24	3.9966	H (0.0, 0.0, 0.0); H (0.0, 0.0, 3.9966)	6.1331	6.1968	6.3045
25	4.3410	H (0.0, 0.0, 0.0); H (0.0, 0.0, 4.3410)	6.2003	6.2773	6.5061

Table D3. Cartesian geometries used for the 1D bond-length PES scan of HeH^+ in STO-3G basis, along with corresponding QPU runtimes, τ_{qpu} (s), from three independent executions at $N_{\text{shots}} = 10,000$. The interatomic distance r is given in angstrom ($\text{\AA}$). All geometries are aligned along the z -axis with one hydrogen atom fixed at the origin.

S.No.	r ($\text{\AA}$)	Geometry (Cartesian coordinates in $\text{\AA}$)	$\tau_{\text{qpu}}^{(1)}$ (s)	$\tau_{\text{qpu}}^{(2)}$ (s)	$\tau_{\text{qpu}}^{(3)}$ (s)
1	0.5500	He (0.0, 0.0, 0.0); H (0.0, 0.0, 0.5500)	5.9937	6.9913	6.1094
2	0.5983	He (0.0, 0.0, 0.0); H (0.0, 0.0, 0.5983)	5.9140	6.0741	6.1593
3	0.6465	He (0.0, 0.0, 0.0); H (0.0, 0.0, 0.6465)	5.9487	6.1472	6.0793
4	0.6948	He (0.0, 0.0, 0.0); H (0.0, 0.0, 0.6948)	6.3038	6.2373	5.8171
5	0.7431	He (0.0, 0.0, 0.0); H (0.0, 0.0, 0.7431)	5.9789	5.9495	5.8637
6	0.7913	He (0.0, 0.0, 0.0); H (0.0, 0.0, 0.7913)	6.0862	6.1246	6.0152
7	0.8396	He (0.0, 0.0, 0.0); H (0.0, 0.0, 0.8396)	5.8899	6.0294	5.8034
8	0.8879	He (0.0, 0.0, 0.0); H (0.0, 0.0, 0.8879)	6.0519	5.8819	5.9775
9	0.9361	He (0.0, 0.0, 0.0); H (0.0, 0.0, 0.9361)	6.1266	5.8488	6.0615
10	0.9844	He (0.0, 0.0, 0.0); H (0.0, 0.0, 0.9844)	5.9444	6.0185	6.0479
11	1.0327	He (0.0, 0.0, 0.0); H (0.0, 0.0, 1.0327)	6.1281	6.1791	6.1898
12	1.0809	He (0.0, 0.0, 0.0); H (0.0, 0.0, 1.0809)	6.0941	5.9533	5.9359
13	1.1292	He (0.0, 0.0, 0.0); H (0.0, 0.0, 1.1292)	15.4878	6.0701	6.0763
14	1.1775	He (0.0, 0.0, 0.0); H (0.0, 0.0, 1.1775)	6.2186	6.0398	5.9692
15	1.2257	He (0.0, 0.0, 0.0); H (0.0, 0.0, 1.2257)	5.9709	5.9988	6.0263
16	1.2740	He (0.0, 0.0, 0.0); H (0.0, 0.0, 1.2740)	6.1320	6.1994	5.9911
17	1.5851	He (0.0, 0.0, 0.0); H (0.0, 0.0, 1.5851)	6.0623	5.9470	5.8928
18	1.8962	He (0.0, 0.0, 0.0); H (0.0, 0.0, 1.8962)	6.1877	6.0365	6.0054
19	2.2073	He (0.0, 0.0, 0.0); H (0.0, 0.0, 2.2073)	11.9733	6.2745	5.9660
20	2.5184	He (0.0, 0.0, 0.0); H (0.0, 0.0, 2.5184)	5.8456	5.9821	6.0214
21	2.8296	He (0.0, 0.0, 0.0); H (0.0, 0.0, 2.8296)	5.9289	5.8350	6.0853
22	3.1407	He (0.0, 0.0, 0.0); H (0.0, 0.0, 3.1407)	5.9151	6.1138	6.1205
23	3.4518	He (0.0, 0.0, 0.0); H (0.0, 0.0, 3.4518)	6.0054	6.0066	5.9501
24	3.7629	He (0.0, 0.0, 0.0); H (0.0, 0.0, 3.7629)	5.9995	6.0248	5.9608
25	4.0740	He (0.0, 0.0, 0.0); H (0.0, 0.0, 4.0740)	5.8320	5.9449	5.8420

Table D4. Cartesian geometries used for the 1D bond-length PES scan of HeH^+ in 6-31G basis, along with corresponding QPU runtimes, τ_{qpu} (s), from three independent executions at $N_{\text{shots}} = 10,000$. The interatomic distance r is given in angstrom ($\text{\AA}$). All geometries are aligned along the z -axis with the helium atom fixed at the origin.

S.No.	r ($\text{\AA}$)	Geometry (Cartesian coordinates in $\text{\AA}$)	$\tau_{\text{qpu}}^{(1)}$ (s)	$\tau_{\text{qpu}}^{(2)}$ (s)	$\tau_{\text{qpu}}^{(3)}$ (s)
1	0.4500	He (0.0, 0.0, 0.0); H (0.0, 0.0, 0.4500)	6.2737	6.2738	6.2271
2	0.5049	He (0.0, 0.0, 0.0); H (0.0, 0.0, 0.5049)	6.0166	6.3871	6.3318
3	0.5599	He (0.0, 0.0, 0.0); H (0.0, 0.0, 0.5599)	6.0435	6.9888	6.5851
4	0.6148	He (0.0, 0.0, 0.0); H (0.0, 0.0, 0.6148)	6.2107	6.1794	6.4662
5	0.6697	He (0.0, 0.0, 0.0); H (0.0, 0.0, 0.6697)	6.6714	6.5610	6.1933
6	0.7247	He (0.0, 0.0, 0.0); H (0.0, 0.0, 0.7247)	6.8343	6.5450	6.2113
7	0.7796	He (0.0, 0.0, 0.0); H (0.0, 0.0, 0.7796)	6.1873	6.2182	6.0013
8	0.8345	He (0.0, 0.0, 0.0); H (0.0, 0.0, 0.8345)	6.1636	6.2252	6.2059
9	0.8895	He (0.0, 0.0, 0.0); H (0.0, 0.0, 0.8895)	6.7252	6.6528	6.3301
10	0.9444	He (0.0, 0.0, 0.0); H (0.0, 0.0, 0.9444)	6.1561	6.5367	6.2345
11	0.9993	He (0.0, 0.0, 0.0); H (0.0, 0.0, 0.9993)	6.0686	7.0564	7.1170
12	1.0543	He (0.0, 0.0, 0.0); H (0.0, 0.0, 1.0543)	7.3921	6.0140	6.0151
13	1.1092	He (0.0, 0.0, 0.0); H (0.0, 0.0, 1.1092)	6.4106	6.4244	6.3907
14	1.1641	He (0.0, 0.0, 0.0); H (0.0, 0.0, 1.1641)	6.1680	6.1359	6.3287
15	1.2191	He (0.0, 0.0, 0.0); H (0.0, 0.0, 1.2191)	6.2126	6.1561	6.5339
16	1.2740	He (0.0, 0.0, 0.0); H (0.0, 0.0, 1.2740)	6.2698	7.3350	7.6462
17	1.5851	He (0.0, 0.0, 0.0); H (0.0, 0.0, 1.5851)	6.2266	6.0591	6.1860
18	1.8962	He (0.0, 0.0, 0.0); H (0.0, 0.0, 1.8962)	7.3691	6.0313	6.2810
19	2.2073	He (0.0, 0.0, 0.0); H (0.0, 0.0, 2.2073)	6.0623	6.2442	6.3034
20	2.5184	He (0.0, 0.0, 0.0); H (0.0, 0.0, 2.5184)	6.1897	6.2835	6.3677
21	2.8296	He (0.0, 0.0, 0.0); H (0.0, 0.0, 2.8296)	6.2115	6.2355	6.2503
22	3.1407	He (0.0, 0.0, 0.0); H (0.0, 0.0, 3.1407)	6.4736	6.1903	6.0395
23	3.4518	He (0.0, 0.0, 0.0); H (0.0, 0.0, 3.4518)	6.3867	6.4235	6.6657
24	3.7629	He (0.0, 0.0, 0.0); H (0.0, 0.0, 3.7629)	6.1007	6.2881	6.2068
25	4.0740	He (0.0, 0.0, 0.0); H (0.0, 0.0, 4.0740)	6.1294	6.3392	6.1096

Table D5. Cartesian geometries used for the 1D bond-length PES scan of LiH in STO-3G basis, along with corresponding QPU runtimes, τ_{qpu} (s), from three independent executions at $N_{\text{shots}} = 10,000$. The interatomic distance r is given in angstrom ($\text{\AA}$). All geometries are aligned along the z -axis with the lithium atom fixed at the origin.

S.No.	r ($\text{\AA}$)	Geometry (Cartesian coordinates in $\text{\AA}$)	$\tau_{\text{qpu}}^{(1)}$ (s)	$\tau_{\text{qpu}}^{(2)}$ (s)	$\tau_{\text{qpu}}^{(3)}$ (s)
1	0.8000	Li (0.0, 0.0, 0.0); H (0.0, 0.0, 0.8000)	6.6105	6.7314	6.5239
2	0.8592	Li (0.0, 0.0, 0.0); H (0.0, 0.0, 0.8592)	6.3240	6.4909	6.5194
3	0.9184	Li (0.0, 0.0, 0.0); H (0.0, 0.0, 0.9184)	6.6621	6.6183	6.2601
4	0.9776	Li (0.0, 0.0, 0.0); H (0.0, 0.0, 0.9776)	7.4973	6.4484	6.6277
5	1.0368	Li (0.0, 0.0, 0.0); H (0.0, 0.0, 1.0368)	7.1162	6.5821	6.4737
6	1.0960	Li (0.0, 0.0, 0.0); H (0.0, 0.0, 1.0960)	6.5866	6.4544	6.6322
7	1.1627	Li (0.0, 0.0, 0.0); H (0.0, 0.0, 1.1627)	6.6017	6.3138	7.8663
8	1.2293	Li (0.0, 0.0, 0.0); H (0.0, 0.0, 1.2293)	6.8034	6.5531	6.5790
9	1.2960	Li (0.0, 0.0, 0.0); H (0.0, 0.0, 1.2960)	6.7020	6.6111	6.8629
10	1.3627	Li (0.0, 0.0, 0.0); H (0.0, 0.0, 1.3627)	6.4236	6.6864	6.4323
11	1.4293	Li (0.0, 0.0, 0.0); H (0.0, 0.0, 1.4293)	6.4246	6.4935	6.6092
12	1.4960	Li (0.0, 0.0, 0.0); H (0.0, 0.0, 1.4960)	6.7005	6.4013	6.4722
13	1.5627	Li (0.0, 0.0, 0.0); H (0.0, 0.0, 1.5627)	6.2868	6.4638	6.5032
14	1.6293	Li (0.0, 0.0, 0.0); H (0.0, 0.0, 1.6293)	6.8374	6.6009	6.8490
15	1.6960	Li (0.0, 0.0, 0.0); H (0.0, 0.0, 1.6960)	6.9978	6.6726	6.6939
16	1.7627	Li (0.0, 0.0, 0.0); H (0.0, 0.0, 1.7627)	6.7075	6.6723	6.5056
17	1.8293	Li (0.0, 0.0, 0.0); H (0.0, 0.0, 1.8293)	6.7507	7.4605	6.4880
18	1.8960	Li (0.0, 0.0, 0.0); H (0.0, 0.0, 1.8960)	6.4749	6.5708	6.7451
19	1.9627	Li (0.0, 0.0, 0.0); H (0.0, 0.0, 1.9627)	6.5895	6.5987	6.5162
20	2.0293	Li (0.0, 0.0, 0.0); H (0.0, 0.0, 2.0293)	6.7619	8.1165	16.0304
21	2.0960	Li (0.0, 0.0, 0.0); H (0.0, 0.0, 2.0960)	6.7832	6.7063	6.4066
22	2.4404	Li (0.0, 0.0, 0.0); H (0.0, 0.0, 2.4404)	7.1057	6.6864	6.6921
23	2.7849	Li (0.0, 0.0, 0.0); H (0.0, 0.0, 2.7849)	6.7119	6.3962	6.7159
24	3.1293	Li (0.0, 0.0, 0.0); H (0.0, 0.0, 3.1293)	6.5403	6.6544	7.0440
25	3.4738	Li (0.0, 0.0, 0.0); H (0.0, 0.0, 3.4738)	6.7604	6.4909	6.6951
26	3.8182	Li (0.0, 0.0, 0.0); H (0.0, 0.0, 3.8182)	6.7622	6.6986	6.4735
27	4.1627	Li (0.0, 0.0, 0.0); H (0.0, 0.0, 4.1627)	6.4291	6.5989	6.5565
28	4.5071	Li (0.0, 0.0, 0.0); H (0.0, 0.0, 4.5071)	6.4096	6.4240	6.7371
29	4.8516	Li (0.0, 0.0, 0.0); H (0.0, 0.0, 4.8516)	6.8221	6.7164	6.5276
30	5.1960	Li (0.0, 0.0, 0.0); H (0.0, 0.0, 5.1960)	6.4271	6.4395	6.2613

Table D6. Cartesian geometries used for the 1D bond-length PES scan of BeH₂ in STO-3G basis, along with corresponding QPU runtimes, τ_{qpu} (s), from three independent executions at $N_{\text{shots}} = 10,000$. The interatomic distance r is given in angstrom ($\text{\AA}$). All geometries are linear along the z -axis with the beryllium atom fixed at the origin.

S.No.	r ($\text{\AA}$)	Geometry (Cartesian coordinates in $\text{\AA}$)	$\tau_{\text{qpu}}^{(1)}$ (s)	$\tau_{\text{qpu}}^{(2)}$ (s)	$\tau_{\text{qpu}}^{(3)}$ (s)
1	0.7200	Be (0.0, 0.0, 0.0); H (0.0, 0.0, 0.7200); H (0.0, 0.0, -0.7200)	6.6106	6.6678	6.4738
2	0.7428	Be (0.0, 0.0, 0.0); H (0.0, 0.0, 0.7428); H (0.0, 0.0, -0.7428)	6.7844	6.5851	6.3720
3	0.7656	Be (0.0, 0.0, 0.0); H (0.0, 0.0, 0.7656); H (0.0, 0.0, -0.7656)	6.6972	6.7438	6.8278
4	0.7884	Be (0.0, 0.0, 0.0); H (0.0, 0.0, 0.7884); H (0.0, 0.0, -0.7884)	6.3312	6.5289	6.5240
5	0.8112	Be (0.0, 0.0, 0.0); H (0.0, 0.0, 0.8112); H (0.0, 0.0, -0.8112)	6.8564	6.7196	6.6531
6	0.8340	Be (0.0, 0.0, 0.0); H (0.0, 0.0, 0.8340); H (0.0, 0.0, -0.8340)	6.7562	6.8221	6.7377
7	0.9007	Be (0.0, 0.0, 0.0); H (0.0, 0.0, 0.9007); H (0.0, 0.0, -0.9007)	6.5101	6.7169	6.6753
8	0.9673	Be (0.0, 0.0, 0.0); H (0.0, 0.0, 0.9673); H (0.0, 0.0, -0.9673)	6.6361	6.9130	6.5877
9	1.0340	Be (0.0, 0.0, 0.0); H (0.0, 0.0, 1.0340); H (0.0, 0.0, -1.0340)	7.2480	6.6442	6.9104
10	1.1007	Be (0.0, 0.0, 0.0); H (0.0, 0.0, 1.1007); H (0.0, 0.0, -1.1007)	6.7268	6.6135	6.6615
11	1.1673	Be (0.0, 0.0, 0.0); H (0.0, 0.0, 1.1673); H (0.0, 0.0, -1.1673)	7.7520	7.4272	6.3733
12	1.2340	Be (0.0, 0.0, 0.0); H (0.0, 0.0, 1.2340); H (0.0, 0.0, -1.2340)	6.5072	6.6901	6.6654
13	1.3007	Be (0.0, 0.0, 0.0); H (0.0, 0.0, 1.3007); H (0.0, 0.0, -1.3007)	6.8089	6.6751	6.4889
14	1.3673	Be (0.0, 0.0, 0.0); H (0.0, 0.0, 1.3673); H (0.0, 0.0, -1.3673)	7.1670	6.5617	6.5469
15	1.4340	Be (0.0, 0.0, 0.0); H (0.0, 0.0, 1.4340); H (0.0, 0.0, -1.4340)	6.8466	6.6929	6.7522
16	1.5007	Be (0.0, 0.0, 0.0); H (0.0, 0.0, 1.5007); H (0.0, 0.0, -1.5007)	6.6260	6.7341	6.5788
17	1.5673	Be (0.0, 0.0, 0.0); H (0.0, 0.0, 1.5673); H (0.0, 0.0, -1.5673)	6.7361	6.5224	6.7584
18	1.6340	Be (0.0, 0.0, 0.0); H (0.0, 0.0, 1.6340); H (0.0, 0.0, -1.6340)	6.8055	6.6261	6.8842
19	1.7007	Be (0.0, 0.0, 0.0); H (0.0, 0.0, 1.7007); H (0.0, 0.0, -1.7007)	6.6008	7.0599	6.5021
20	1.7673	Be (0.0, 0.0, 0.0); H (0.0, 0.0, 1.7673); H (0.0, 0.0, -1.7673)	7.3756	6.7832	6.5648
21	1.8340	Be (0.0, 0.0, 0.0); H (0.0, 0.0, 1.8340); H (0.0, 0.0, -1.8340)	6.5205	6.6326	6.8316
22	2.0718	Be (0.0, 0.0, 0.0); H (0.0, 0.0, 2.0718); H (0.0, 0.0, -2.0718)	6.6125	6.6399	6.4728
23	2.3096	Be (0.0, 0.0, 0.0); H (0.0, 0.0, 2.3096); H (0.0, 0.0, -2.3096)	6.6628	6.5029	6.5046
24	2.5473	Be (0.0, 0.0, 0.0); H (0.0, 0.0, 2.5473); H (0.0, 0.0, -2.5473)	6.7331	7.0441	7.2652
25	2.7851	Be (0.0, 0.0, 0.0); H (0.0, 0.0, 2.7851); H (0.0, 0.0, -2.7851)	7.2212	7.5278	7.0331
26	3.0229	Be (0.0, 0.0, 0.0); H (0.0, 0.0, 3.0229); H (0.0, 0.0, -3.0229)	7.1399	7.8295	6.5676
27	3.4984	Be (0.0, 0.0, 0.0); H (0.0, 0.0, 3.4984); H (0.0, 0.0, -3.4984)	6.6436	6.6932	7.1953

In addition to the one-dimensional PES scans summarized above, a two-dimensional potential energy surface (2D-PES) was constructed for the H₂O molecule in the STO-3G basis by simultaneously varying the O–H bond length and the H–O–H bond angle. Owing to the size of the resulting configuration space, individual Cartesian geometries are not tabulated for the 2D-PES. Instead, all nuclear configurations are specified analytically through a deterministic geometry-generation function, and QPU execution times are reported in aggregated form via a runtime heatmap, as shown in Figure D1.

The geometries were generated using a symmetric internal-coordinate parameterization with the oxygen atom fixed at the origin. For each pair of internal coordinates (r, θ) , the Cartesian positions of the nuclei are given by

$$\mathbf{R}_O = (0.0, 0.0, 0.0), \quad (\text{D1})$$

$$\mathbf{R}_{H_1} = \left(r \sin \frac{\theta}{2}, 0.0, r \cos \frac{\theta}{2} \right), \quad (\text{D2})$$

$$\mathbf{R}_{H_2} = \left(-r \sin \frac{\theta}{2}, 0.0, r \cos \frac{\theta}{2} \right), \quad (\text{D3})$$

where r denotes the O–H bond length and θ denotes the H–O–H bond angle. This construction preserves molecular symmetry and ensures that all variations in the PES arise exclusively from controlled distortions of the internal coordinates.

The internal coordinates were sampled uniformly on a rectangular grid defined by

$$r \in [0.85, 1.20] \text{ \AA}, \quad \theta \in [85^\circ, 115^\circ], \quad (\text{D4})$$

using $N_r = 32$ points along the bond-length dimension and $N_\theta = 32$ points along the angular dimension, yielding a total of $N_r \times N_\theta = 1024$ distinct nuclear configurations. The corresponding grid spacings are

$$\Delta r = \frac{1.20 - 0.85}{31}, \quad \Delta \theta = \frac{115^\circ - 85^\circ}{31}. \quad (\text{D5})$$

All geometries were used directly, without further optimization, in the SQD(LUCJ) quantum workflow as well as in the associated CCSD and FCI classical benchmark calculations. QPU execution runtimes $\tau_{\text{qpu}}(r, \theta)$ in seconds(s) are reported graphically in the form of a heatmap as shown in Figure D1.




Figure D1. QPU execution runtime τ_{qpu} heatmap for the two-dimensional potential energy surface (2D-PES) of the H₂O molecule in the STO-3G basis, with H–O–H bond angle θ (°) along the x-axis and O–H bond length r (Å) along the y-axis. Each pixel corresponds to a single (r, θ) nuclear configuration on the 32×32 grid at a fixed shot count of $N_{\text{shots}} = 10,000$.

E DMET-SQD: Molecule Geometries & QPU Runtimes

Table E1. Geometric coordinates of the molecules used for simulation, along with SQD runtime and DMET convergence iterations.

Molecule	Geometric Coordinates				τ_{qpu} (s)	N_{it}
	Atom	x	y	z		
Cyanic Acid (HOCN)	H	-1.4586	-0.2728	0.0655	112.9437	4
	O	-0.6006	0.0996	-0.3212		
	C	0.5444	0.0905	0.4227		
	N	1.5148	0.0828	1.0531		
Formaldehyde Oxime (CH ₃ NO)	C	-0.9105	-0.0209	-0.3325	157.8839	4
	H	2.1707	-0.1169	0.6124		
	H	-1.0055	1.0435	-0.5074		
	H	-1.7610	-0.6672	-0.5009		
	N	0.2079	-0.5236	0.0747		
	O	1.2984	0.2851	0.2934		
Methoxyamine (CH ₅ NO)	C	-0.9661	0.2013	0.1500	204.0350	4
	H	-1.0003	0.9655	-0.6585		
	H	-1.6672	0.5167	0.9498		
	H	-1.3069	-0.7850	-0.2370		
	H	1.8825	0.4411	-0.4096		
	H	1.6013	-1.1638	-0.0232		
	N	1.1321	-0.2814	-0.3282		
	O	0.3246	0.1057	0.6921		
Methyl Isocyanate (C ₂ H ₃ NO)	C	-0.8727	-0.0486	-0.0127	187.3237	4
	C	1.4988	0.2006	-0.1212		
	H	-0.6585	-0.7319	0.8393		
	H	-1.4198	-0.6122	-0.7965		
	H	-1.5141	0.7828	0.3460		
	N	0.3512	0.5000	-0.5815		
	O	2.6150	-0.0907	0.3266		
Acetaldehyde Oxime (C ₂ H ₅ NO)	C	-1.3220	0.0798	-0.0599	237.6822	4
	C	0.1251	0.1946	0.2783		
	H	-1.8488	-0.4928	0.7319		
	H	-1.7722	1.0918	-0.1346		
	H	-1.4404	-0.4413	-1.0330		
	H	0.5670	1.1752	0.4092		
	H	2.6590	0.0768	0.8612		
	N	0.8453	-0.8736	0.4138		
	O	2.1868	-0.8104	0.7261		
Carbamide / Urea (CH ₄ N ₂ O)	C	0.0236	0.1748	0.4750	206.0734	5
	H	2.1523	-0.0418	0.2442		
	H	1.1878	-0.5180	-1.1992		
	H	-2.1229	0.2590	0.3460		
	H	-1.3040	-0.3426	-1.1399		
	N	-1.2432	0.0166	-0.1616		
	N	1.2242	-0.1571	-0.2203		
	O	0.0823	0.6092	1.6557		
Nitrosyl Chloride (NOCl)	N	-0.2226	0.4586	0.0000	102.4765	5
	O	-1.1576	-0.2817	0.0000		
	Cl	1.3803	-0.1769	0.0000		
Hydroxythiocyanate (HOSCN)	H	-1.5762	0.5676	-0.3044	140.3231	4
	O	-1.2722	0.4432	0.6302		
	S	-0.4213	-1.0482	0.6165		
	C	1.1270	-0.2448	0.3537		
	N	2.1426	0.2822	0.1812		

Table E2. Geometric coordinates of Amantadine used for simulation (geometry-optimized up to force-field level).

Molecule	Geometric Coordinates			
	Atom	x	y	z
Amantadine (C ₁₀ H ₁₇ N)	H	-0.72057	1.14346	0.93144
	H	-1.53537	1.21683	-0.47761
	N	-0.73805	0.77679	-0.01956
	C	-0.92076	-0.67321	0.01059
	C	-2.22915	-1.03654	0.74833
	H	-3.09165	-0.57160	0.25272
	H	-2.21056	-0.65093	1.77628
	C	0.26844	-1.33324	0.74228
	H	0.34582	-0.95489	1.77025
	H	1.21362	-1.08313	0.24234
	C	0.09005	-2.86228	0.76987
	H	0.93676	-3.32380	1.29009
	C	-0.98906	-1.22001	-1.43216
	H	-0.07298	-0.96728	-1.98243
	H	-1.82253	-0.75964	-1.97920
	C	-1.17095	-2.74873	-1.41062
	H	-1.21856	-3.12973	-2.43684
	C	-2.47313	-3.09321	-0.66778
	H	-2.62462	-4.17964	-0.66246
	H	-3.33092	-2.65475	-1.19237
	C	-2.41390	-2.56473	0.77559
	H	-3.34286	-2.81538	1.29977
	C	-1.21881	-3.20616	1.50117
	H	-1.34806	-4.29459	1.54495
	H	-1.17399	-2.84898	2.53734
	C	0.01730	-3.38948	-0.67334
	H	0.95170	-3.16434	-1.20199
	H	-0.09033	-4.48112	-0.66794

Table E3. Dependence of SQD runtime and DMET convergence on the number of measurement shots.

N_{shots}	τ_{qpu} (s)	N_{it}
500	238.6378	6
1,000	459.7616	11
5,000	235.1656	4
10,000	324.2494	4