

Sparse Quantum State Preparation for Strongly Correlated Systems

César Feniou, Olivier Adjoua, Baptiste Claudon, Julien Zylberman, Emmanuel Giner, and Jean-Philip Piquemal*



Cite This: *J. Phys. Chem. Lett.* 2024, 15, 3197–3205



Read Online

ACCESS |



Metrics & More

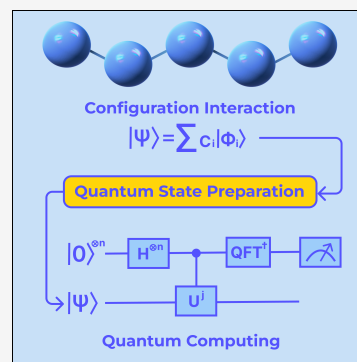


Article Recommendations



Supporting Information

ABSTRACT: Quantum computing allows, in principle, the encoding of the exponentially scaling many-electron wave function onto a linearly scaling qubit register, offering a promising solution to overcome the limitations of traditional quantum chemistry methods. An essential requirement for ground state quantum algorithms to be practical is the initialization of the qubits to a high-quality approximation of the sought-after ground state. Quantum state preparation enables the generation of approximate eigenstates derived from classical computations but is frequently treated as an oracle in quantum information. In this study, we investigate the quantum state preparation of prototypical strongly correlated systems' ground state, up to 28 qubits, using the Hyperion-1 GPU-accelerated state-vector emulator. Various variational and nonvariational methods are compared in terms of their circuit depth and classical complexity. Our results indicate that the recently developed Overlap-ADAPT-VQE algorithm offers the most advantageous performance for near-term applications.



Simulating electronic structure on computers is a crucial and challenging endeavor with significant applications in drug discovery and advanced materials design.¹ Unfortunately, the full configuration interaction (FCI) method, which provides exact solutions to the electronic time-independent nonrelativistic Schrödinger equation in a given basis, requires the diagonalization of the molecular Hamiltonian to solve the electronic structure problem. The computational cost of performing an FCI calculation therefore increases exponentially as the system size expands. Quantum computers offer a promising solution to overcome this issue by enabling the representation and manipulation of exponentially large electronic wave functions using only a linear number of qubits. To leverage the potential of quantum computers for such tasks, quantum algorithms such as the Variational Quantum Eigensolver² (VQE) and Quantum Phase Estimation^{3,4} (QPE) have been developed for ground state quantum chemistry calculations.^{5–7}

Unfortunately, the limited quantum resources available on current-generation Noisy Intermediate Scale Quantum^{8,9} (NISQ) devices severely limit the implementation of such methods and their ability to achieve an advantage over their classical counterparts. Hence, there is a need for the simultaneous use of both improved quantum hardware and resource-efficient quantum algorithms. Both VQE and QPE have a significantly reduced computational cost when the initial state, i.e., the starting configuration of qubits, is carefully prepared in order to exhibit high overlap with the ground state of the system under study. Since VQE is an optimization problem, beginning with an initial state that is near the solution naturally enhances the efficiency of the search for the optimal solution. Regarding QPE, the success probability of the

algorithm is proportional to the overlap between the initial state and the Hamiltonian ground state being simulated.^{3,4} For a typical chemical quantum simulation, the Hartree–Fock state is usually chosen as the initial state for two main reasons: it can be represented on a quantum computer using just a few quantum gates, and it tends to be the predominant Slater determinant for most molecular systems. However, as pointed out in many studies,^{7,10–12} a single-determinant initial state might not be accurate enough for quantum computation of strongly correlated systems. One possible strategy to generate more accurate and multireference initial states is to build upon the ability of conventional quantum chemistry methods to produce such wave functions. The subsequent step involves transferring this classical data into the quantum computer to prepare the desired initial state, a process known as *Quantum State Preparation*,¹³ and often treated as a “black-box” or oracle procedure in quantum information. Formally, it involves preparing the state $|\Psi\rangle$ from a set of coefficients c_p such as

$$|\Psi\rangle = \sum_{p \in \{0,1\}^n} c_p |p\rangle, \quad \sum_{p \in \{0,1\}^n} |c_p|^2 = 1 \quad (1)$$

Note that an arbitrary quantum state preparation necessitates circuits that scale exponentially with the number of qubits either in size, depth, or number of ancilla qubits.^{13–15}

Received: November 10, 2023

Revised: February 23, 2024

Accepted: March 1, 2024

Published: March 14, 2024



Consequently, this initialization step could potentially dominate the overall computational cost of the quantum algorithm, leading to the loss of the anticipated advantage. A common metric for assessing quantum resource requirements includes the counts of both CNOT gates and single-qubit gates. While the emphasis on CNOT count is typical in NISQ devices, fault-tolerant quantum computing places greater importance on the count of single-qubit gates, which include non-Clifford arbitrary rotations.

Configuration interaction¹⁶ (CI) is a post-Hartree–Fock quantum chemistry approach that represents the wave function as a weighted linear combination of Slater determinants. As the size of the FCI space grows exponentially with the number of particles and basis functions, the Selected-CI (SCI) methods have been proposed in order to automatically select only the important contributions in the Hilbert space. Among various flavors of SCI, the CI perturbatively selected iteratively (CIPSI)^{17,18} employs an iterative approach in which the most pertinent Slater determinants are dynamically added to the wave function based on an importance derived from perturbation theory (PT). By choosing on-the-fly Slater determinants tailored to the systems being simulated, the SCI ansatz wave function rapidly captures the bulk of correlation effects and is therefore a natural candidate for a suitable multireference initial quantum states. Formally, the SCI wave function is defined over a subset of Slater determinants within the FCI space, denoted here the variational subset of determinants $\mathcal{V} = \{|\phi_i\rangle, i = 1, N\}$, and the energy is minimized over this set of parameters

$$|\Psi_{\text{SCI}}\rangle = \sum_{i \in \mathcal{V}} c_i |\phi_i\rangle, \\ \{c_i\} = \text{argmin} \frac{\langle \Psi_{\text{SCI}} | H | \Psi_{\text{SCI}} \rangle}{\langle \Psi_{\text{SCI}} | \Psi_{\text{SCI}} \rangle} \quad (2)$$

For a given wave function $|\Psi_{\text{SCI}}\rangle$, one can compute the associated second-order perturbed energy as

$$E^{(2)} = \sum_{i \notin \mathcal{V}} \frac{\langle i | H | \Psi_{\text{SCI}} \rangle^2}{E(\Psi_{\text{SCI}}) - E(I)} \quad (3)$$

where $E(\Psi)$ is the variational energy of a given wave function Ψ . Therefore, the quantity $E^{(2)}$ is a measure of the remaining contribution to the energy of the Slater determinants not already captured in $|\Psi_{\text{SCI}}\rangle$. In the following calculations, we stop our calculations when $|E^{(2)}| < 10^{-4}$ in order to obtain target states of near FCI quality.

The process of encoding CI states in qubit registers involves mapping Slater determinants to computational basis states. The Jordan–Wigner Fermion-to-qubit transformation¹⁹ is commonly employed for this mapping. In this transformation, each qubit stores the occupation of a spin-orbital, meaning that each vector in the computational basis corresponds to a Slater determinant. Computational basis states can represent unphysical states such as the $|0\rangle^n$ state, which corresponds to a configuration with zero particles. The Hartree–Fock state of an m -electron and n -spin-orbital system is encoded as $|\Psi_{\text{HF}}\rangle = |1\rangle^m \otimes |0\rangle^{n-m}$, when arranging the spin-orbitals in increasing order of energy. First, let us remark that the SCI target state has a number of variational determinants M much smaller than the dimension of the computational Hilbert space ($M \ll 2^n$). It is thus encoded as a *sparse quantum state* of the sparsity M . Second, more generally, let us remark that it respects multiple

fundamental symmetries such as Fermionic anticommutation relations, which are reflected in the amplitudes, particle conservation, and projected-spin, with only specific computational basis states respecting them. The molecular Hamiltonian eigenstates thus belong to a subspace of the n -qubit Hilbert space that is relevant to these symmetries.

Having established the properties of the target wave functions, we can now shift our attention to exploring the diverse methods available for the quantum state preparation of such targets. One first natural idea was introduced by Gard et al.²⁰ and involves fully parametrizing the symmetry-preserving subspace with minimal CNOT count. Parameters are then variationally tuned until convergence to the target state, which is within this subspace. Assuming convergence, this Symmetry-Preserving-VQE approach allows for an exact state preparation with a maximum CNOT count of 3 times the dimension of the subspace. An established alternative to fixed circuit methods involves adaptive variational algorithms where quantum circuits are grown and tailored to the problem being simulated. The Adaptive Derivative-Assembled Pseudo Trotter (ADAPT)-VQE²¹ belongs to this family of methods and has become a benchmark technique due to its capability of generating ansatz wave functions that are both highly accurate approximations to the ground state and also more compact in terms of circuit depth than fixed-ansatz approaches. A variant of this approach—labeled Overlap-ADAPT-VQE²²—was recently introduced for quantum state preparation. The Overlap-ADAPT-VQE algorithm iteratively generates a compact approximation of a target wave function through a quasigreedy procedure that maximizes, at each iteration, the overlap of the current iterate with the target. This is achieved by adding on-the-fly the most relevant unitary operator from a finite-size pool of admissible operators, with a selection criterion based on the gradient of the overlap of the ansatz with the target. Note that whether the Overlap-ADAPT-VQE ansatz evolves within the symmetry-relevant subspace of the target depends on the choice of operator pool. Pools made of Fermionic operators adhere to all the desired symmetries,²¹ which facilitate the convergence to the target state but come at the expense of requiring more quantum gates per operator. Conversely, Qubit-Excitation-Based,^{23,24} Qubit,²⁵ or Minimal ZY²⁵ pools may respect only a subset of these symmetries or none at all, but they are more resource-efficient. The Overlap-ADAPT-VQE workflow and the associated details of the aforementioned operator pools are provided in the [Appendix](#).

Despite a rather intuitive structure, variational quantum state preparation algorithms face scalability issues because they rely on heuristic classical optimization processes that become exponentially complex as the number of qubits grows.²⁶ Indeed, the optimization problem is nonconvex and plagued by multiple local minima and barren plateaus,²⁷ which so far limited the applicability of variational quantum state preparation algorithms to rather simple quantum systems. Recently, significant research attention has been directed toward nonvariational methods for quantum state preparation,^{28–31} particularly when the state to be prepared exhibits distinct structural traits or symmetries.^{10,30,32–36} As a result, nonvariational sparse quantum state preparation algorithms offer an interesting alternative for the preparation of SCI wave functions, as they are exact and deterministic. We propose in [Table 2a](#), a nonexhaustive list of sparse quantum state preparation algorithms and associated complexities.

Table 1. Performance of Common Sparse State Preparation Methods^a

Method	Classical preprocessing	CNOT count	Circuit depth	Ancillas
CVO-QRAM ³⁶	$O(M \log M + nM)$	$\sum_{t=1}^n \mu_t(8t - 4) - t_{\max}$	$O(nM)$	1
Gleinig et al. ³⁵	$O(M^2 \log(M)n)$	$O(nM)$	$O(nM)$	0
Tubman et al. ¹⁰		$O(nM)$	$O(nM)$	1
Zhang et al. ³⁷	$O(nM)$	[very large]	$\Theta(\log(nM))$	$O(nM \log M)$
Fomichev et al. ³⁸	$O(M^3)$	$O(M \log M)$	$O(M \log M)$	$O(\log M)$

^aWith respect to the target state, M denotes the sparsity, n the number of qubits involved, μ_t is the number of determinants with t bits of value 1, and t_{\max} is the highest value of t with $\mu_t \neq 0$.

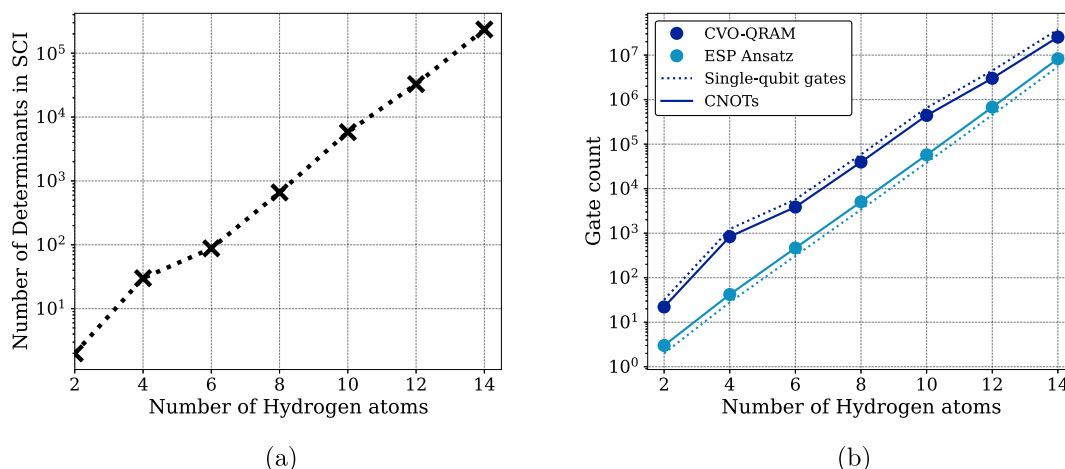


Figure 1. Determinants count (a) and circuit size (b) in the SCI representation of linear H_n chains. The circuit size has been established using CVO-QRAM loader and ESP Ansatz as the state preparation method.

The “CVO-QRAM” algorithm developed by de Veras et al.³⁶ enables sparse quantum state preparation algorithm that scales linearly with respect to the sparsity of the target, using a single ancilla qubit. We detail in the [Appendix](#) the correctness of the algorithm and use it in the following simulations.

The study presented in this Letter aims to compare the aforementioned state-of-the-art variational and nonvariational approaches to sparse quantum state preparation on the criteria of circuit size and classical processing complexity. The target states under consideration are accurate approximations of hydrogen chains’ ground states, which are prototypical examples of strongly correlated systems in quantum chemistry.³⁹ The linear hydrogen chains that we consider are stretched with an interatomic distance of 5.0 Å. The target ground states are obtained via CIPSI^{17,18} SCI simulations carried out on the Quantum Package⁴⁰ software, using a minimal STO-3G basis set, and up to the challenging case of H_{14} , whose quantum encoding requires the significant count of 28 qubits. [Figure 1a](#) emphasizes, as anticipated, the exponential growth in the number of determinants needed to represent the ground state of increasingly large hydrogen chains.

In [Figure 1b](#), we display the circuit size required to construct a quantum circuit that encodes the Full-CI wave functions of the hydrogen chains. These quantum circuits are obtained using the CVO-QRAM method and the Efficient Symmetry-Preserving (ESP) Ansatz as quantum state preparation algorithms applied to the CIPSI ground state as a target. While the former approach substantially lowers the CNOT count compared to the CVO-QRAM method, it is crucial to emphasize that the ESP Ansatz is built upon certain underlying assumptions. First, the convergence of this optimization process cannot be guaranteed, giving rise to system-depend-

ence, while CVO-QRAM ensures systematic success. Second, variational state preparation via ESP Ansatz is of exponential complexity with the system size. In either scenario, performing state preparation for such target states presents challenges, with impractical computational cost in one case and large circuit size detrimental to the subsequent quantum post-treatment in the other.

The CVO-QRAM and ESP quantum state preparation estimated costs suggest that preparing compact approximations of the target state might be a natural option. Notably, variational methods can progressively approach the solution, allowing access to ansatz wave functions of arbitrary precision with flexible circuit depth and classical computational time. Adaptive variational algorithms have notably showcased their effectiveness compared to fixed-ansatz techniques in the realm of approximate ground state preparation. Specifically, the Overlap-ADAPT-VQE has previously been applied to the state preparation of small SCI-wave functions,²² making it a suitable choice for the forthcoming simulations.

We carry out Overlap-ADAPT-VQE simulations with the linear H_{14} ground state as target on 28 qubits, and using two common operator pools, namely, Qubit-Excitation-Based- and Qubit-pool, whose details are given in the [Appendix](#). Such large computations are made possible thanks to the use of the in-house GPU-accelerated Hyperion-1 quantum emulator. To ensure a fair comparison with CVO-QRAM, which performs exact state preparation, we prepared a range of approximated states of the linear H_{14} ground state. The approximation involves cropping the smallest coefficients of the target and renormalizing, guaranteeing minimal sparsity for an arbitrary high overlap. An alternative being explored involves taking as approximated ground state the variational states achieved




Figure 2. (a) Absolute CI coefficient of each Slater determinant in the Selected-CI expansion of the linear H_{14} ground state wave function. (b) Associated absolute sum of CI squared coefficients.

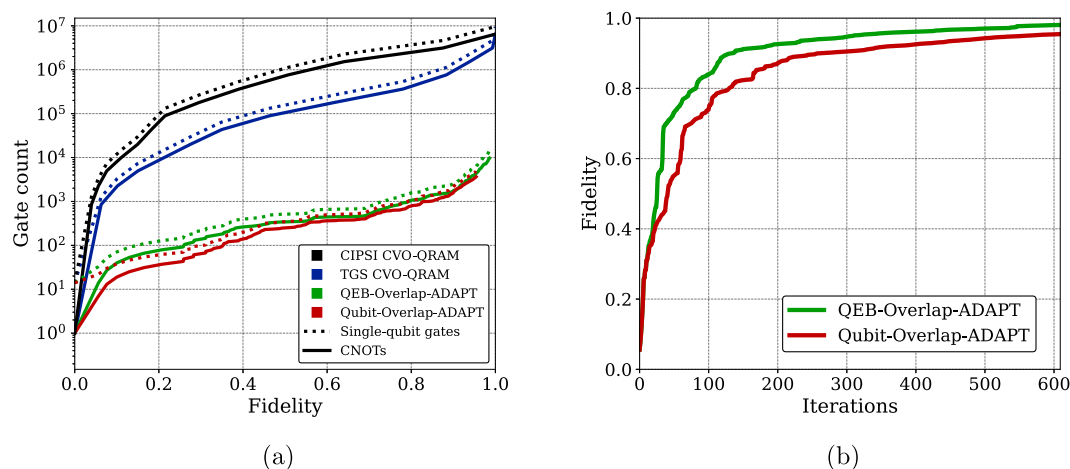


Figure 3. (a) Gate count for approximate state preparation of ground state wave function of linear H_{14} displayed in Figure 2. The CVO-QRAM algorithm generated states from CIPSI iterates as well as states derived from truncations of the ground state (TGS). The Overlap-ADAPT-VQE approximates the ground state using the QEB- and the Qubit-pool of operators. (b) Fidelity of the Overlap-ADAPT-VQE ansatz over the iterations for the same target.

during each iteration of CIPSI. These states bear higher physical significance and have been observed as accurate ansatz-drivers in the context of post-treatment with the VQE.²² A table comparing the fidelities and the determinant count of these two options is provided in the Appendix (Table 2). We display in Figure 2a the absolute weight of each Slater determinant in the CI of the H_{14} linear hydrogen chain ground state. Note that for the H_{14} system, relying on a single determinant initialization yields at most a modest $4 \cdot 10^{-2}$ overlap with the ground state. Moreover, Figure 2b presents the associated absolute sum of CI squared coefficients, highlighting the need for a multideterminant initial state.

The resulting gate count of state preparation via CVO-QRAM and Overlap-ADAPT for various levels of approximations is given in Figure 4. The adaptive variational approach leads to significantly reduced gate count for any observed level of fidelity. Overlap-ADAPT-VQE seamlessly expands ansatz with fidelity exceeding 95% and appears to be resilient to common variational algorithm challenges, such as hostile landscapes with barren plateaus or local minima. This observation is particularly noteworthy considering the strong correlation inherent in the target system. The associated circuit

involves no more than a few thousand CNOT gates and single-qubit gates, a highly reasonable resource in the context of the early fault-tolerant era. We anticipate the dominance of the CVO-QRAM approach in exactly preparing the target state. Achieving full convergence with the variational approach is unlikely without an exceedingly flexible ansatz involving an impractical number of parameters. This is primarily due to the abundance of determinants with very small weights in the target state. Indeed, Figure 3b shows the convergence profile of the Overlap-ADAPT simulations and illustrates the challenges faced by the ansatz in reaching a fidelity of 1 throughout the iterations. Nevertheless, the majority of studies typically regard an overlap exceeding 0.5 with the ground state as sufficiently high for a QPE initial state. This perspective makes the exact state preparation approach via CVO-QRAM less suitable for these applications.

The CVO-QRAM process exhibits a superlinear classical complexity with the number of determinants, which are obtained with classical SCI methods that scale exponentially with the system's size. Therefore, the overall state preparation strategy is dominated by the cost of the SCI procedure. When targeting a classically derived wave function, the Overlap-




Figure 4. Gate count for approximate state preparation of ground state wave function of linear H_{10} in (a) and H_{12} in (b). The CVO-QRAM algorithm generated states derived from truncations of the CIPSI ground state (TGS). The Overlap-ADAPT-VQE approximates the CIPSI ground state using the QEB- and the Qubit-pool of operators, iterating over 500 steps for each system. The solid lines represent the CNOT counts, while the dashed lines indicate the single-qubit gate counts.

ADAPT-VQE process necessitates classical simulation of a quantum circuit, leading to exponentially scaling memory requirement with the number of qubits and, consequently, with the size of the targeted system. This imposes an upper limit on the qubit count that can be emulated without resorting to approximation or encountering errors. To extend the scalability of variational algorithms on simulators, alternative strategies have been investigated. These include approximated emulations by using Matrix Product State circuit simulators or by partitioning the system and performing successive local optimizations. These efforts have shown promising outcomes, successfully simulating systems of up to around 100 qubits.⁴¹ Despite multiple layers of approximation, integrating Overlap-ADAPT-VQE with approximated emulation can be promising in replicating high-fidelity ansatz by using compact circuits for larger systems, a focus for future research, particularly for molecular systems with classically intractable convergence. To capitalize on quantum computing memory advantages while maintaining the objective of high-fidelity initial states with compact circuits, an alternative approach is to employ the Overlap-ADAPT-VQE algorithm as a hybrid quantum-classical algorithm. As it necessitates the target wave function to be stored in a quantum register, executing this method for our problem first involves converting the CIPSI wave function into a quantum state using CVO-QRAM. Then, an Overlap-ADAPT-VQE ansatz can be grown iteratively to maximize its overlap with the CVO-QRAM prepared state, potentially resulting in a more compact approximation of the latter. The overlap measurement can be conducted using either the swap test or the compute-uncompute approach, both of which have been explored on quantum hardware in prior research within this context.⁴² Regardless of the chosen approach, let us emphasize that Overlap-ADAPT-VQE not only delivers a precise initial state but presents it in the structure of a parametrized quantum circuit. This circuit can be fine-tuned on a quantum computer, offering adaptability based on the desired post-treatment. This feature proves especially advantageous in the realm of VQE, where flexibility in the ansatz is sought.

By generating a compact quantum circuit that encodes the best classically derived wave function, we secure an initial quantum state with strong overlap with the sought-after

ground state of the system. This initial state preparation will greatly facilitate a following quantum post-treatment, such as QPE, which offers an asymptotic advantage in determining the ground state energy compared to classical approaches.

The key takeaway is that, even for chemical systems of prime hardness,⁴³ quantum state preparation of accurate initial state has proven seamless with Overlap-ADAPT-VQE coupled to a classical ground state approximation procedure. The resulting quantum circuits exhibit an insignificant circuit depth in the context of early fault-tolerant era. It is worth noting that an entirely different class of procedures conducts state preparation directly within the quantum computer, either through adiabatic state preparation⁵ or by leveraging ground-state boosting⁴⁴ based only on the information provided by the system's Hamiltonian. While these approaches may seem more scalable than classical state preparation algorithms, they introduce a substantial quantum resource overhead for medium-sized systems, making them inherently more appropriate for the longer term. Further investigation is sought for addressing the regime where the strategy of utilizing classically derived initial states outperforms other initialization approaches. In a broader context, sparse quantum state preparation methods are used for a range of applications, including Hamiltonian simulation with either linear combination of unitary operators⁴⁵ or qubitization,⁴⁶ data loading for quantum machine learning⁴⁷ or even solving a linear system of equations.⁴⁸ Overlap-ADAPT-VQE may therefore offer a practical approach for any such algorithm involving classical data loading by allowing the use of an operator pool tailored to the target state's characteristics.

■ TECHNICAL APPENDIX

Quantum State Preparation for H_{10} and H_{12} Chains.

We display in Figure 4 the gate count for the quantum state preparation of the ground state for two additional hydrogen chains, H_{10} and H_{12} . Although the Overlap-ADAPT-VQE algorithm demonstrates greater size efficiency in both cases, with an order of magnitude fewer gates compared to the CVO-QRAM in terms of ground-state fidelity, the disparity is less pronounced than that in the case of H_{14} . The observed difference can be attributed to the lower degree of correlation in shorter hydrogen chains, with the Hartree–Fock state

having stronger support and fewer determinants required to accurately approximate the ground state.

The Hyperion-1 High-Performance State-Vector Emulator. Hyperion-1⁴⁹ is a quantum computing state-vector emulator package dedicated to quantum chemistry. It uses classical hardware and is massively parallel thanks to an efficient multi-GPU (Graphics Processing Unit) implementation. High performance is ensured by an ensemble of fast custom CUDA sparse linear algebra libraries, which act on the native quantum chemistry algorithms to accelerate Hyperion-1's exact/noiseless simulations. For example, using a single DGX-A100 node (8 × A100 GPUs, 40 Gb of memory/GPU), we have been able to carry out a 28 qubit Overlap-ADAPT-VQE simulation producing over 600 iterations within 8 days. The fidelity of the ansatz with respect to the H_{14} target state is plotted with respect to the iterates in Figure 3b. The algorithms examined in this paper are classical, and their execution in the state-vector formalism excludes both complete gate decomposition and the simulation of noise related to shots or hardware. Nevertheless, these options are accessible via the software. Further details about Hyperion-1 will be given in a forthcoming publication.

Double Sparse Quantum State Preparation. Even though it is unclear how to prepare an arbitrary state of 2^n complex amplitudes with a cost which is not exponential with the number of qubits n , the problem of loading a number M of complex amplitudes was proven to have satisfying solutions scaling linearly with M such as the CVO-QRAM method.³⁶ This algorithm consists of loading sequentially the classical data by performing controlled rotations using one ancilla qubit. The algorithm is presented, and the correctness is proven recursively as follows.

Algorithm 1: CVO-QRAM - Complex Data Storage Algorithm

```

input : data = { $x_k, p_k$ } $_{k=0}^{M-1}$ 
output:  $|m\rangle = \sum_{k=0}^{M-1} x_k |p_k\rangle$ 

1 cvoqram (data,  $|\psi\rangle$ ):
2    $|\psi_{00}\rangle = |u\rangle |m\rangle = |1; 0, \dots, 0\rangle$ 
3   foreach ( $x_k, p_k$ )  $\in$  data do
4      $t$  = The number of bits with value 1 in the pattern  $p_k$ 
5      $l$  = a list containing the positions of  $p_k$  where  $p_k[j] = 1$ .
6      $|\psi_{k1}\rangle = \prod_{l_i \in l} CX_{(u, m[l_i])} |\psi_{k0}\rangle$ 
7      $|\psi_{k2}\rangle = C^t U^{(x_k, \gamma_k)}(m[l_0, l_1, \dots, l_{t-1}], u) |\psi_{k1}\rangle$ 
8     if  $k \neq M - 1$  then
9        $|\psi_{k3}\rangle = \prod_{l_i \in l} CX_{(u, m[l_i])} |\psi_{k2}\rangle$ 
10  return  $|m\rangle$ 

```

Define the amplitude remaining to be loaded by $\gamma_0 = 1$, $\gamma_{k+1} = \gamma_k - |x_{k+1}|^2$ for $0 \leq k \leq M - 1$. Notice that by normalization of the input data $\gamma_{M-1} = 0$. Then, for each k , the unitary:

$$U^{(x_k, \gamma_k)} = \frac{1}{\sqrt{\gamma_k}} \begin{pmatrix} \sqrt{\gamma_k - |x_k|^2} & x_k \\ x_k^* & \sqrt{\gamma_k - |x_k|^2} \end{pmatrix}$$

The algorithm correctness is proven using the following loop invariant H_k , $k = 0 \dots M - 1$: H_k : Right before loading (x_k, p_k) , line 3, the state of the system is $|\psi_{k0}\rangle = \sum_{j=0}^{k-1} x_j |0, p_j\rangle + \sqrt{\gamma_k} |1, 0\rangle$. For $k = 0$, H_0 holds vacuously. Let us check for $k < M - 1$, the for loop maintains the property. Assume that the H_k is verified. Then, simply compute:

$$\begin{aligned} |\psi_{k1}\rangle &= \sum_{j=0}^{k-1} x_j |0, p_j\rangle + \sqrt{\gamma_k} |1, p_k\rangle \\ |\psi_{k2}\rangle &= \sum_{j=0}^{k-1} x_j |0, p_j\rangle + \sqrt{\gamma_k} \left(\frac{x_k}{\sqrt{\gamma_k}} |0, p_k\rangle + \frac{\sqrt{\gamma_{k+1}}}{\sqrt{\gamma_k}} |1, p_k\rangle \right) \\ &= \sum_{j=0}^k x_j |0, p_j\rangle + \sqrt{\gamma_{k+1}} |1, p_k\rangle \\ |\psi_{k3}\rangle &= \sum_{j=0}^k x_j |0, p_j\rangle + \sqrt{\gamma_{k+1}} |1, 0\rangle \end{aligned} \quad (4)$$

Thus, the property is maintained. In order to show that the procedure ends well, simply notice $|\gamma_{M-1}\rangle = 0$ since the input data must be normalized (without loss of generality).

The CVO-QRAM method needs two steps of classical preprocessing. First, given an input with M patterns to load, the classical device must first sort the input patterns with cost $O(M \log M)$. Then, it creates the circuits associated with the $U^{(x_k, \gamma_k)}$ with total cost $O(nM)$. The total, classical cost is $O(M \log M + nM)$. In terms of quantum gates, the number of CNOT gates to prepare a state is

$$\sum_{t=1}^n \mu_t (8t - 4) - t_{\max}$$

where μ_t is the number of input patterns p_k with t bits with value 1 in the binary string and t_{\max} is the highest value of t with $\mu_t \neq 0$.

OVERLAP-ADAPT-VQE

The general workflow of the Overlap ADAPT-VQE algorithm is given in the following. The computations of overlap and gradients of overlap are not expensive. We encourage the reader to consult our most recent publications^{22,42} (and supplementary information) on the subject for further details about how the overlap and the overlap gradients are computed in Overlap-ADAPT-VQE or any other related technical aspect.

Algorithm 2: Overlap-ADAPT-VQE - Variational Quantum State Preparation

```

input : initial state  $|\psi_0\rangle$ , target state  $|\phi_T\rangle$ , pool of operators  $\mathcal{A}$ , convergence threshold  $\epsilon$ 
output: ansatz  $|\psi(\theta)\rangle$  such as  $|\langle\psi(\theta)|\phi_T\rangle| < \epsilon$ 

1 Overlap-ADAPT-VQE( $|\psi_0\rangle, |\phi_T\rangle, \mathcal{A}, \epsilon$ ):
2    $n \leftarrow 1$ 
3   while  $S_n = |\langle\psi_n(\theta)|\Psi_T\rangle| > \epsilon$  do
4     foreach  $A_k \in \mathcal{A}$  do
5       Measure gradients  $g_{A_k} \leftarrow \left| \frac{\partial S_{n, A_k}(\theta)}{\partial \theta} \right|_{\theta=0}$ 
6        $i \leftarrow$  index of the largest element of  $\vec{g}$ 
7       Pick best operator  $A_n \leftarrow A_i$ 
8       Grow Ansatz  $|\psi_n\rangle \leftarrow A_n |\psi_{n-1}\rangle$ 
9       Optimise Ansatz  $\theta^{\text{opt}} = \underset{\theta=(\theta_1, \dots, \theta_n)}{\text{argmin}} |\langle\psi_n(\theta)|\Psi_T\rangle|$ 
10       $\theta \leftarrow \theta^{\text{opt}}$ 
11      Iterate  $n \leftarrow n + 1$ 
12  return  $|\psi_n(\theta)\rangle$ 

```

Operator Pools. The literature contains a wide range of operator pools for ADAPT-like algorithms and VQEs. Some are chemically inspired and ensure that the ansatz preserves the symmetry of the target state, while others prioritize hardware efficiency and minimize the number of operators used.^{25,50} This section intends to offer a concise introduction to the two

operator pools that have been used in prior numerical simulations.

The Qubit Excitation-Based Pool. The Qubit excitation-based (QEB) pool is widely used for simulating quantum chemical systems²³ because it preserves spin and number of particles while still being fairly hardware efficient. The QEB pool consists of so-called single-qubit and double-qubit excitation operators which take the form

$$A_{pq} = \frac{1}{2}(X_q Y_p - Y_p X_q) \quad (5)$$

and

$$A_{pqrs} = \frac{1}{8}(X_r Y_s X_p X_q + Y_r X_s X_p X_q + Y_r Y_s Y_p X_q + Y_r Y_s X_p Y_q - X_r X_s Y_p X_q - X_r X_s X_p Y_q - Y_r X_s Y_p Y_q - X_r Y_s Y_p Y_q) \quad (6)$$

The variables p , q , r , and s represent qubit indices, and X_p and Y_p represent the standard one-qubit Pauli gates applied to qubit p . Therefore, the single-qubit generator A_{pq} operates between individual qubits p and q , while the double-qubit generator A_{pqrs} operates between the qubit pairs (p, q) and (r, s) . The computation of parametric exponentiation for QEB operators is straightforward, and the resulting unitary operators come with well-established CNOT-optimized circuit implementations given below.




Figure 5. A quantum circuit performing a generic single-qubit excitation.⁵¹

The Qubit Hardware-Efficient Pool. The Qubit hardware-efficient²⁴ pool instead uses decomposed single and double excitation operators which are of the form

$$B_{pq} = X_q Y_p \quad (7)$$

and

$$\begin{aligned} B_{pqrs}^{(1)} &= X_r Y_s X_p X_q, & B_{pqrs}^{(2)} &= Y_r X_s X_p X_q, & B_{pqrs}^{(3)} &= Y_r Y_s X_p X_q, & B_{pqrs}^{(4)} &= Y_r Y_s X_p Y_q, \\ B_{pqrs}^{(5)} &= X_r X_s Y_p X_q, & B_{pqrs}^{(6)} &= X_r X_s X_p Y_q, & B_{pqrs}^{(7)} &= Y_r X_s Y_p Y_q, & B_{pqrs}^{(8)} &= X_r Y_s Y_p Y_q \end{aligned} \quad (8)$$

where p , q , r , s again denote qubit indices and X_p and Y_p are one-qubit Pauli gates acting on qubit p . Note that the qubit hardware-efficient pool does not preserve particle and spin conservation of the ansatz, which may result in convergence difficulties due to the violation of such crucial system symmetries. However, prior numerical simulations employing the qubit hardware-efficient pool have produced ansatzes that require fewer CNOT operations for implementation on

quantum hardware when compared to ansatz wave functions based on the QEB pool.

Target States Details.

Table 2. H₁₄ Ground State Approximations: Comparative Fidelities of Variational CIPSI Iterate States and Truncations of the Ground State with Respect to the Numbers of Determinants

Number of determinants	Fidelity	
	CIPSI iterate state	Truncated ground state
8	0.0397	0.0627
21	0.0562	0.1015
46	0.0761	0.1501
92	0.1103	0.2121
185	0.1495	0.2749
403	0.1824	0.3498
827	0.2147	0.4631
1657	0.2954	0.6164
3356	0.3905	0.7800
7054	0.5071	0.8833
14174	0.6411	0.9390
28976	0.8757	0.9932
58664	0.9964	1.0000

■ ASSOCIATED CONTENT

Supporting Information

The Supporting Information is available free of charge at <https://pubs.acs.org/doi/10.1021/acs.jpclett.3c03159>.

Transparent Peer Review report available (PDF)

■ AUTHOR INFORMATION

Corresponding Author

Jean-Philip Piquemal – Sorbonne Université, LCT, UMR 7616 CNRS, 75005 Paris, France; Qubit Pharmaceuticals, Advanced Research Department, 75014 Paris, France; orcid.org/0000-0001-6615-9426; Email: jean-philip.piquemal@sorbonne-universite.fr

Authors

César Feniou – Sorbonne Université, LCT, UMR 7616 CNRS, 75005 Paris, France; Qubit Pharmaceuticals, Advanced Research Department, 75014 Paris, France
Olivier Adjoua – Sorbonne Université, LCT, UMR 7616 CNRS, 75005 Paris, France
Baptiste Claudon – Sorbonne Université, LCT, UMR 7616 CNRS, 75005 Paris, France; Qubit Pharmaceuticals, Advanced Research Department, 75014 Paris, France
Julien Zylberman – Sorbonne Université, LERMA, UMR 8112 CNRS, 75005 Paris, France
Emmanuel Giner – Sorbonne Université, LCT, UMR 7616 CNRS, 75005 Paris, France




Figure 6. A quantum circuit performing a generic double-qubit excitation.⁵¹

Complete contact information is available at:
<https://pubs.acs.org/10.1021/acs.jpclett.3c03159>

Notes

The authors declare the following competing financial interest(s): Jean-Philip Piquemal is co-founder and shareholder of Qubit Pharmaceuticals.

ACKNOWLEDGMENTS

We thank Muhammad Hassan and Yvon Maday (LJLL, Sorbonne Université) for providing feedback on the composition of this paper. Computations have been performed at IDRIS (Jean Zay) on GENCI grant no A0150712052 (JPP). Additional DGX-A100 GPUs computations have been performed thanks to the support of the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No 810367), project EMC2 (JPP). Support from the PEPR EPIQ and HQI programs is also acknowledged (JPP).

REFERENCES

- (1) Helgaker, T.; Jorgensen, P.; Olsen, J. *Molecular electronic-structure theory*; John Wiley & Sons, 2013.
- (2) Peruzzo, A.; McClean, J.; Shadbolt, P.; Yung, M.-H.; Zhou, X.-Q.; Love, P. J.; Aspuru-Guzik, A.; O'Brien, J. L. A variational eigenvalue solver on a photonic quantum processor. *Nat. Commun.* **2014**, *5*, 4213.
- (3) Kitaev, A. Y. Quantum measurements and the Abelian stabilizer problem. *arXiv preprint quant-ph/9511026*, 1995.
- (4) Nielsen, M. A.; Chuang, I. L. *Quantum computation and quantum information*; Cambridge University Press, 2010.
- (5) Aspuru-Guzik, A.; Dutoi, A. D.; Love, P. J.; Head-Gordon, M. Simulated quantum computation of molecular energies. *Science* **2005**, *309*, 1704–1707.
- (6) Whitfield, J. D.; Biamonte, J.; Aspuru-Guzik, A. Simulation of electronic structure Hamiltonians using quantum computers. *Mol. Phys.* **2011**, *109*, 735–750.
- (7) Tilly, J.; Chen, H.; Cao, S.; Picozzi, D.; Setia, K.; Li, Y.; Grant, E.; Wossnig, L.; Rungger, I.; Booth, G. H.; et al. The variational quantum eigensolver: a review of methods and best practices. *Phys. Rep.* **2022**, *986*, 1–128.
- (8) Preskill, J. Quantum computing in the NISQ era and beyond. *Quantum* **2018**, *2*, 79.
- (9) Bharti, K.; Cervera-Lierta, A.; Kyaw, T. H.; Haug, T.; Alperin-Lea, S.; Anand, A.; Degroote, M.; Heimonen, H.; Kottmann, J. S.; Menke, T.; et al. Noisy intermediate-scale quantum (nisq) algorithms. *arXiv preprint arXiv:2101.08448*, 2021.
- (10) Tubman, N. M.; Mejuto-Zaera, C.; Epstein, J. M.; Hait, D.; Levine, D. S.; Huggins, W.; Jiang, Z.; McClean, J. R.; Babbush, R.; Head-Gordon, M.; Whaley, K. B. Postponing the orthogonality catastrophe: efficient state preparation for electronic structure simulations on quantum devices. *arXiv preprint arXiv:1809.05523*, 2018.
- (11) Dalzell, A. M.; McArdle, S.; Berta, M.; Bienias, P.; Chen, C.-F.; Gilyén, A.; Hann, C. T.; Kastoryano, M. J.; Khabiboulline, E. T.; Kubica, A.; et al. Quantum algorithms: A survey of applications and end-to-end complexities. *arXiv preprint arXiv:2310.03011*, 2023.
- (12) Lee, S.; Lee, J.; Zhai, H.; Tong, Y.; Dalzell, A. M.; Kumar, A.; Helms, P.; Gray, J.; Cui, Z.-H.; Liu, W.; et al. Evaluating the evidence for exponential quantum advantage in ground-state quantum chemistry. *Nat. Commun.* **2023**, *14*, 1952.
- (13) Plesch, M.; Brukner, Č. Quantum-state preparation with universal gate decompositions. *Phys. Rev. A* **2011**, *83*, No. 032302.
- (14) Grover, L.; Rudolph, T. Creating superpositions that correspond to efficiently integrable probability distributions. *arXiv preprint quant-ph/0208112*, 2002.
- (15) Mottonen, M.; Vartiainen, J. J.; Bergholm, V.; Salomaa, M. M. Transformation of quantum states using uniformly controlled rotations. *arXiv preprint quant-ph/0407010*, 2004.
- (16) Nesbet, R. K. Configuration Interaction in Orbital Theories. *Proc. R. Soc. A* **1955**, *230*, 312–321.
- (17) Huron, B.; Malrieu, J. P.; Rancurel, P. Iterative perturbation calculations of ground and excited state energies from multiconfigurational zeroth-order wavefunctions. *J. Chem. Phys.* **1973**, *58*, 5745–5759.
- (18) Evangelisti, S.; Daudey, J.-P.; Malrieu, J.-P. Convergence of an improved CIPSI algorithm. *J. Chem. Phys.* **1983**, *75*, 91–102.
- (19) Jordan, P.; Wigner, E. P. *The Collected Works of Eugene Paul Wigner* **1993**, 109–129.
- (20) Gard, B. T.; Zhu, L.; Barron, G. S.; Mayhall, N. J.; Economou, S. E.; Barnes, E. Efficient symmetry-preserving state preparation circuits for the variational quantum eigensolver algorithm. *npj Quantum Information* **2020**, *6*, 10.
- (21) Grimsley, H. R.; Economou, S. E.; Barnes, E.; Mayhall, N. J. Adapt-vqe: An exact variational algorithm for fermionic simulations on a quantum computer. *arXiv preprint arXiv:1812.11173*, 2018.
- (22) Feniou, C.; Hassan, M.; Traoré, D.; Giner, E.; Maday, Y.; Piquemal, J.-P. Overlap-ADAPT-VQE: practical quantum chemistry on quantum computers via overlap-guided compact Ansätze. *Communications Physics* **2023**, *4*, 192.
- (23) Yordanov, Y. S.; Armaos, V.; Barnes, C. H.; Arvidsson-Shukur, D. R. Qubit-excitation-based adaptive variational quantum eigensolver. *Communications Physics* **2021**, *4*, 228.
- (24) Ryabinkin, I. G.; Yen, T.-C.; Genin, S. N.; Izmaylov, A. F. Qubit coupled cluster method: a systematic approach to quantum chemistry on a quantum computer. *J. Chem. Theory Comput.* **2018**, *14*, 6317–6326.
- (25) Tang, H. L.; Shkolnikov, V. O.; Barron, G. S.; Grimsley, H. R.; Mayhall, N. J.; Barnes, E.; Economou, S. E. Qubit-Adapt-Vqe: An adaptive algorithm for constructing hardware-efficient ansätze on a quantum processor. *PRX Quantum* **2021**, *2*, No. 020310.
- (26) Astrakhantsev, N.; Mazzola, G.; Tavernelli, I.; Carleo, G. Phenomenological theory of variational quantum ground-state preparation. *Phys. Rev. Res.* **2023**, *5*, No. 033225.
- (27) McClean, J. R.; Boixo, S.; Smelyanskiy, V. N.; Babbush, R.; Neven, H. Barren plateaus in quantum neural network training landscapes. *Nat. Commun.* **2018**, *9*, 4812.
- (28) Araújo, I. F.; Park, D. K.; Petruccione, F.; da Silva, A. J. A divide-and-conquer algorithm for quantum state preparation. *Sci. Rep.* **2021**, *11*, 6329.
- (29) Araújo, I. F.; Blank, C.; Araújo, I. C.; da Silva, A. J. Low-rank quantum state preparation. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* **2024**, *43*, 161.
- (30) Zhang, X.-M.; Li, T.; Yuan, X. Quantum State Preparation with Optimal Circuit Depth: Implementations and Applications. *Phys. Rev. Lett.* **2022**, *129*, No. 230504.
- (31) Sun, X.; Tian, G.; Yang, S.; Yuan, P.; Zhang, S. Asymptotically optimal circuit depth for quantum state preparation and general unitary synthesis. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* **2023**, *42*, 3301.
- (32) Zylberman, J.; Debbasch, F. Efficient Quantum State Preparation with Walsh Series. *arXiv preprint arXiv:2307.08384*, 2023.
- (33) Holmes, A.; Matsuura, A. Y. Efficient quantum circuits for accurate state preparation of smooth, differentiable functions. *2020 IEEE International Conference on Quantum Computing and Engineering (QCE)*, 2020; pp 169–179.
- (34) Marin-Sanchez, G.; Gonzalez-Conde, J.; Sanz, M. Quantum algorithms for approximate function loading. *Physical Review Research* **2023**, *5*, No. 033114.
- (35) Gleinig, N.; Hoefler, T. An Efficient Algorithm for Sparse Quantum State Preparation. *2021 58th ACM/IEEE Design Automation Conference (DAC)*, 2021; pp 433–438.
- (36) de Veras, T. M. L.; da Silva, L. D.; da Silva, A. J. Double sparse quantum state preparation. *Quantum Information Processing* **2022**, *21*, 204.

- (37) Zhang, X.-M.; Li, T.; Yuan, X. Quantum State Preparation with Optimal Circuit Depth: Implementations and Applications. *Phys. Rev. Lett.* **2022**, *129*, 230504.
- (38) Fomichev, S.; Hejazi, K.; Zini, M. S.; Kiser, M.; Morales, J. F.; Casares, P. A. M.; Delgado, A.; Huh, J.; Voigt, A.-C.; Mueller, J. E.; Arrazola, J. M. Initial state preparation for quantum chemistry on quantum computers. *arXiv preprint arXiv:2310.18410*, 2023.
- (39) Stair, N. H.; Evangelista, F. A. Exploring Hilbert space on a budget: Novel benchmark set and performance metric for testing electronic structure methods in the regime of strong correlation. *J. Chem. Phys.* **2020**, *153*, No. 104108.
- (40) Garniron, Y.; et al. Quantum Package 2.0: An Open-Source Determinant-Driven Suite of Programs. *J. Chem. Theory Comput.* **2019**, *15*, 3591–3609.
- (41) Shang, H.; Fan, Y.; Shen, L.; Guo, C.; Liu, J.; Duan, X.; Li, F.; Li, Z. Towards practical and massively parallel quantum computing emulation for quantum chemistry. *npj Quantum Information* **2023**, *9*, 33.
- (42) Feniou, C.; Claudon, B.; Hassan, M.; Courtat, A.; Adjoua, O.; Maday, Y.; Piquemal, J.-P. Greedy Gradient-free Adaptive Variational Quantum Algorithms on a Noisy Intermediate Scale Quantum Computer. *arXiv preprint arXiv:2306.17159*, 2023.
- (43) Choi, S.; Yen, T.-C.; Izmaylov, A. F. Improving Quantum Measurements by Introducing “Ghost” Pauli Products. *J. Chem. Theory Comput.* **2022**, *18*, 7394–7402.
- (44) Wang, G.; Sim, S.; Johnson, P. D. State preparation boosters for early fault-tolerant quantum computation. *Quantum* **2022**, *6*, 829.
- (45) Childs, A. M.; Wiebe, N. Hamiltonian simulation using linear combinations of unitary operations. *arXiv preprint arXiv:1202.5822*, 2012.
- (46) Low, G. H.; Chuang, I. L. Hamiltonian simulation by qubitization. *Quantum* **2019**, *3*, 163.
- (47) Zhao, Z.; Fitzsimons, J. K.; Rebentrost, P.; Dunjko, V.; Fitzsimons, J. F. Smooth input preparation for quantum and quantum-inspired machine learning. *Quantum Machine Intelligence* **2021**, *3*, 14.
- (48) Harrow, A. W.; Hassidim, A.; Lloyd, S. Quantum algorithm for linear systems of equations. *Phys. Rev. Lett.* **2009**, *103*, No. 150502.
- (49) Adjoua, O.; Feniou, C.; et al. *Hyperion-1*; Sorbonne Université and CNRS and Qubit Pharmaceuticals. 2023.
- (50) Kandala, A.; Mezzacapo, A.; Temme, K.; Takita, M.; Brink, M.; Chow, J. M.; Gambetta, J. M. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature* **2017**, *549*, 242–246.
- (51) Yordanov, Y. S.; Arvidsson-Shukur, D. R.; Barnes, C. H. Efficient quantum circuits for quantum computational chemistry. *Phys. Rev. A* **2020**, *102*, No. 062612.