



Benchmarking Adaptive Variational Quantum Eigensolvers

Daniel Claudino^{1,2}, Jerimiah Wright^{1,3}, Alexander J. McCaskey^{1,2} and Travis S. Humble^{1,3*}

¹ Quantum Computing Institute, Oak Ridge National Laboratory, Oak Ridge, TN, United States, ² Computer Science and Mathematics, Oak Ridge National Laboratory, Oak Ridge, TN, United States, ³ Computational Sciences and Engineering, Oak Ridge National Laboratory, Oak Ridge, TN, United States

By design, the variational quantum eigensolver (VQE) strives to recover the lowest-energy eigenvalue of a given Hamiltonian by preparing quantum states guided by the variational principle. In practice, the prepared quantum state is indirectly assessed by the value of the associated energy. Novel adaptive derivative-assembled pseudo-trotter (ADAPT) ansatz approaches and recent formal advances now establish a clear connection between the theory of quantum chemistry and the quantum state ansatz used to solve the electronic structure problem. Here we benchmark the accuracy of VQE and ADAPT-VQE to calculate the electronic ground states and potential energy curves for a few selected diatomic molecules, namely H₂, NaH, and KH. Using numerical simulation, we find both methods provide good estimates of the energy and ground state, but only ADAPT-VQE proves to be robust to particularities in optimization methods. Another relevant finding is that gradient-based optimization is overall more economical and delivers superior performance than analogous simulations carried out with gradient-free optimizers. The results also identify small errors in the prepared state fidelity which show an increasing trend with molecular size.

Keywords: ADAPT-VQE, quantum computing, quantum chemistry, VQE, potential energy scan, state fidelity

1. INTRODUCTION

Quantum mechanics naturally lends itself to the description of phenomena at the atomic and molecular scale, including problems of chemical interest, which has culminated in the field of research known as quantum chemistry. Despite the formal impediments to achieve exact, closed-form solutions to quantum chemistry problems, there is a wide array of possible approximations, such as coupled cluster (CC) theory (Shavitt and Bartlett, 2009), which have elevated quantum chemistry to good standing in the scientific community due to their reliability.

In practice, CC faces two main difficulties that have hindered a more widespread adoption. One is that most of the success it has garnered over the years is due to its superior performance in the weak electron correlation regime, for which single-reference (SR) CC remains unchallenged. This success is justified because many problems in chemistry, such as thermochemistry, can be adequately treated as being largely weakly correlated. Yet, many other problems of interest involve molecules and materials that do not comply with this assumption, and for these instances, SR-CC breaks down. Despite multi-reference (MR) CC being an active area of research (Jeziorski and Monkhorst, 1981), theoretical and computational challenges currently curb the applicability of MR-CC (Lyakh et al., 2012).

A second obstacle to a more extensive use of CC theory is its pronounced computational cost. Reliable SR-CC methods, such as the so-called “gold standard” of quantum chemistry,

OPEN ACCESS

Edited by:

Sugata Chowdhury,
National Institute of Standards and
Technology (NIST), United States

Reviewed by:

Bryan Gard,
Georgia Tech Research Institute,
United States
George Barron,
Virginia Tech, United States

*Correspondence:

Travis S. Humble
humblets@ornl.gov

Specialty section:

This article was submitted to
Physical Chemistry and Chemical
Physics,
a section of the journal
Frontiers in Chemistry

Received: 15 September 2020

Accepted: 11 November 2020

Published: 04 December 2020

Citation:

Claudino D, Wright J, McCaskey AJ
and Humble TS (2020) Benchmarking
Adaptive Variational
Quantum Eigensolvers.
Front. Chem. 8:606863.
doi: 10.3389/fchem.2020.606863

coupled cluster singles and doubles (and perturbative triples), aka CCSD(T) (Urban et al., 1985; Raghavachari et al., 1989; Watts et al., 1993), scale unfavorably with one-particle basis spanning the Hilbert space that houses the electronic wave function, which largely constrains the application of CCSD(T) to relatively small molecular systems. It is important to note that some of these limitations can be mitigated with methods such as configuration interaction (CI) in its MR formulation and the density matrix renormalization group (DMRG) which have in turn their own shortcomings, such as lack of size-extensivity and exactness contingent upon the dimensionality of the problem.

Concurrent with developments in CC theory has been the increase in performance of computing technologies, which broadens the reach of computational chemistry techniques. Presently, this trend is continuing with the adaptation of chemistry methods, including CC, to the new technology paradigm of quantum computing (Britt and Humble, 2017; Humble et al., 2019). Because of the shared foundation in quantum mechanics, one of the most immediate applications for quantum computers is quantum chemistry (McArdle et al., 2020). Recent advances have reformulated conventional problems in electronic structure for currently available quantum computing platforms (Cao et al., 2019). In particular, these efforts have led to a resurgence of the unitary coupled cluster (UCC) theory (Bartlett et al., 1989; Kutzelnigg, 1991; Taube and Bartlett, 2006; Romero et al., 2018), which can be employed in investigations where strong correlation is dominant. Quantum computing hardware appears to be well suited for building the states described by UCC, as this hardware can efficiently implement unitary operations to construct physical representations of the quantum state. Moreover, the intrinsic nature of the quantum computing logic can be exploited in order to propose new ansätze that, despite lacking a close connection to the underlying chemical intuition lent by UCC, are prone to a more efficient implementation, such as the so-called hardware efficient ansatz (Kandala et al., 2017).

It is in the context of noisy intermediate-scale quantum (NISQ) (Preskill, 2018) devices that the variational quantum eigensolver (VQE) (Peruzzo et al., 2014) has emerged as a promising method for testing the preparation and measurement of quantum states including those that represent the electronic eigenstates described by UCC (Quantum et al., 2020). Several variants of VQE are available (Parrish et al., 2019b; Chivilikhin et al., 2020), but all build on the variational principle from quantum mechanics, which constrains the quantum states that can satisfy the electronic eigenvalue problem (McClean et al., 2016). While the initial VQE proposal assumes a predefined ansatz, this constraint has been relaxed, opening the door to adaptive approaches (Grimsley et al., 2019; Ryabinkin et al., 2020), by which the preparable quantum states are driven by the problem at hand. In particular, the adaptive derivative-assembled pseudo-trotter (ADAPT) ansatz, which finds support on the recently coined “disentangled” UCC (Evangelista et al., 2019) and starts from an exact UCC representation of the electronic ground state to construct an approximate prepared state based on the dominant contributions. Early studies demonstrated this as a promising avenue for developing ansätze for specific

molecules and constraints, such as highly accurate energetics or shallow circuits.

Here we benchmark adaptive VQE prescriptions, ADAPT-VQE in particular, by comparing the prepared quantum states with the conventional solutions obtained from exact diagonalization of the full configuration interaction Hamiltonian. We track the energy of the minimized expectation value as well as the fidelity of the corresponding prepared state using multiple ansatz, optimization methods, and molecular Hamiltonians. We calculate infidelity as a measure of error for the prepared quantum state relative to the expected, exact result from quantum chemistry using frozen-core Hamiltonians. Across these examples, we find that ADAPT-VQE is the more robust method due mainly to its performance with respect to optimization methods. While all methods lead to small errors as measured by the infidelity, these errors are found to grow with molecular size.

This presentation is structured as follows. In section 2, we provided an overview of the ingredients in the VQE approach relevant to our purposes, followed by a short exposition of the underpinnings of ADAPT-VQE (section 2.1) and a brief discussion on implementation of gradients and optimization in ADAPT-VQE (section 2.2). The computational details permeating the reported simulations are exposed in section 3. The main results are presented and discussed in section 4 and several conclusions are drawn in section 5.

2. VARIATIONAL QUANTUM EIGENSOLVER

This section serves to illustrate the pertinent fundamentals of the VQE algorithm and to motivate the following exposition of adaptive ansatz construction. We start by recalling the variational principle, which is at the heart of VQE, and given as

$$E \leq \min_{\Psi} \langle \Psi | \hat{H} | \Psi \rangle \quad (1)$$

where $|\Psi\rangle$ is normalized trial wave function for which Equation (1) becomes an equality when Ψ is constructed from a basis that spans the single-particle Hilbert space of all possible occupation numbers (the underlying Fock space) and the electronic Hamiltonian \hat{H} for a molecular system is given as

$$\hat{H} = \sum_{pq} h_{pq} p^\dagger q + \sum_{pqrs} h_{pqrs} p^\dagger q^\dagger sr \quad (2)$$

The central problem in modern electronic structure theory is the description and quantification of the electron correlation from an un-entangled, mean-field wave function $|0\rangle$ whose preparation can be carried out in classical hardware in a timely fashion, e.g., Hartree-Fock (HF). In analogy with quantum chemistry, we can expect that there exists an operator that, once applied to $|0\rangle$, will account for the missing electron correlation. Bearing in mind that quantum computers manipulate quantum states in a well-defined Hilbert space, this configures a generic unitary operator $\hat{U}(\vec{\theta})$ whose main purpose is to build entanglement from an un-entangled reference function $|0\rangle$. The set of scalars $\vec{\theta}$ are parameters variationally varied in order to minimize

the expectation value in Equation (1). With that, we recast Equation (1):

$$E \leq \min_{\vec{\theta}} \langle 0 | \hat{U}^\dagger(\vec{\theta}) \hat{H}_P \hat{U}(\vec{\theta}) | 0 \rangle \quad (3)$$

In order to ensure that Equation (3) meets the requirements of quantum hardware, the fermionic, second-quantized operators found in the formulation of electronic structure problem, such as those in Equation (2), are brought to a qubit (spin) representation, with the additional constraint of fermionic anti-symmetry. Our approach uses the Jordan-Wigner transformation (Jordan and Wigner, 1928), but others exist, and such a transformation yields \hat{H}_P from \hat{H} , that is, the Hamiltonian in terms of strings of Pauli operators. Starting from the UCC ansatz, the unitary $\hat{U}(\vec{\theta})$ can be written as:

$$\hat{U}(\vec{\theta}) = \exp\left(\sum_k \theta_k (\hat{T}_k - \hat{T}_k^\dagger)\right) = \exp\left(\sum_k \theta_k \tau_k\right) \quad (4)$$

with the \hat{T}_k representing the usual cluster operators in CC theory and $\tau_k = \hat{T}_k - \hat{T}_k^\dagger$, ensuring the anti-Hermiticity of the operators, which is a necessary condition for their utilization in quantum computing.

Once in possession of all ingredients in Equation (3), the tasks of preparing the state $\hat{U}(\vec{\theta})|0\rangle$ and measuring the terms in \hat{H}_P are delegated to the quantum hardware, and $\hat{U}(\vec{\theta})$ is varied variationally with the aid of a classical optimization routine until $\langle \hat{H}_P \rangle$ reaches its minimum, which is dependent on the chosen optimizer and is taken as a good approximation to the sought ground state energy. Due to the isomorphic property of the qubit mappings, $\langle \hat{H}_P \rangle = \langle \hat{H} \rangle$, yielding the lowest energy eigenvalue of the molecular Hamiltonian in Equation (2).

2.1. ADAPT-VQE

An important choice in the specification of the VQE method is the functional form of the ansatz $\hat{U}(\vec{\theta})$. Even for a relatively small Hilbert space, with a moderate number of cluster operators \hat{T}_i , the ansatz $\hat{U}(\vec{\theta})$ gives rise to a unitary that translates into multi-qubit gates and thus cannot be efficiently implemented in an actual quantum processor. Borrowing from the dynamics community, this can be alleviated by resorting to the Trotter-Suzuki decomposition, or Trotterization for short:

$$\exp\left(\sum_k \theta_k \tau_k\right) \approx \prod_k \exp(\theta_k \tau_k) \quad (5)$$

which here is limited to first-order.

The Adaptive Derivative-Assembled Pseudo-Trotter ansatz Variational Quantum Eigensolver (ADAPT-VQE) (Grimsley et al., 2019) takes advantage of Equation (5) to propose an iterative ansatz construction whereby only the perceived most relevant operator for energy lowering is added to the ansatz. A set of operators the algorithm can choose from needs to be provided, which in this work is comprised of the fermionic spin singlet adapted single and double excitations, borrowing from the usual UCCSD formulation, and subsequently mapped into the

appropriate tensor products of Pauli operators via the Jordan-Wigner transformation. In principle, one could envision explicit enforcement or relaxation of other types of symmetry, and the effect of such choices on the performance of ADAPT is certainly a topic worth exploring. Moreover, the ADAPT algorithm has also recently been reported to perform well with other choices of operators, including a more economical pool of qubit operators (Tang et al., 2020), and has been applied to variational algorithms other than VQE (Zhu et al., 2020).

From a practical standpoint, at the i -th iteration of the algorithm, the energy gradient vector (\mathcal{G}) with respect to all $\{\theta_k\}$ in Equation (5) is computed from measurements on the circuit that prepares the state optimized in the previous iteration, represented by $|\psi_{i-1}\rangle$, with $|\psi_0\rangle = |0\rangle$. Labeling the energy at the current iteration E_i , we have:

$$\mathcal{G} = \left(\frac{\partial E_i}{\partial \theta_1}, \dots, \frac{\partial E_i}{\partial \theta_k}, \dots, \frac{\partial E_i}{\partial \theta_N} \right) \\ \frac{\partial E_i}{\partial \theta_k} = \langle \psi_{i-1} | [H, \tau_k] | \psi_{i-1} \rangle \quad (6)$$

and if the norm of this vector falls below a set threshold, the algorithm is deemed converged and the ansatz-growing loop is exited. Otherwise, the operator associated with the largest absolute component of \mathcal{G} is selected to increment the ansatz:

$$|\psi_i\rangle = e^{\theta_i \tau_i} |\psi_{i-1}\rangle, \quad \tau_i = \{\tau_k | \max | \langle [H, \tau_k] \rangle_{i-1} | \} \quad (7)$$

where $\langle [H, \tau_k] \rangle_{i-1}$ means this commutator was computed from observations in the circuit obtained from the previous iteration. With the selection of a new operator, the new ansatz is subject to the usual VQE routine and the corresponding energy minimum is obtained.

2.2. Gradient Estimate and Classical Optimization in ADAPT-VQE

From a quantum computing standpoint, ADAPT-VQE improves on VQE by potentially offering a more tractable circuit. However, this may come at the expense of a much larger number of measurements, as the evaluation of all $[H, \hat{A}_k]$ is performed at each iteration of the ADAPT loop, on top of the expected energy evaluations. In order to reduce the number of measurements associated with ADAPT-VQE simulations, adoption of a gradient estimate strategy can help improve the classical optimization step by reaching the sought minima with fewer calls to the hardware backend.

To motivate the discussion, we start by invoking the gradient expression as introduced in the original formulation of ADAPT-VQE:

$$\frac{\partial E}{\partial \theta_i} = \langle \phi | \hat{H} \prod_{j=N}^{i+1} (e^{\theta_j \tau_j}) \tau_i \prod_{k=i}^1 (e^{\theta_k \tau_k}) | 0 \rangle \\ - \langle 0 | \prod_{k=1}^i (e^{-\theta_k \tau_k}) \tau_i \prod_{j=i+i}^N (-e^{\theta_j \tau_j}) \hat{H} | \phi \rangle \quad (8)$$

where $\prod_i e^{\theta_i \tau_i} | 0 \rangle = |\phi\rangle$.

Equation (8) can be further simplified into a recursive formula:

$$\frac{\partial E}{\partial \theta_i} = \langle \phi | \left[\hat{H}, \prod_{j=N}^{i+1} (e^{\theta_j \tau_j}) \tau_i \prod_{j=i+1}^N (e^{-\theta_j \tau_j}) \right] | \phi \rangle \quad (9)$$

Before moving further in the discussion regarding the use of gradients to support the classical optimizer, let us clarify a potential source of confusion. At a certain ADAPT-VQE iteration, the circuit previously optimized is implemented to prepare the state from which the current iteration builds upon. The gradient vector \mathcal{G} is then computed upon the necessary measurements for all τ_k in the chosen operator pool (Equation 6), and the operator that has the largest commutator (in absolute value) is selected. And this the extent to which the gradient is employed at this stage. On the other hand, we now have a new ansatz, which is composed of all previously added operators that enable preparation of $|\psi_{i-1}\rangle$, along with the newly added operator from Equation (7). Each of these operators have a corresponding variational parameter, which in the following VQE step need to be re-optimized. It is in this optimization that we would employ the gradients as written in Equations (8) and (9), and whose magnitude needs to be minimized in order to signal the finding of an extremum (minimum in this case). For an operator pool containing N elements, at each ADAPT-VQE iteration, all N elements of \mathcal{G} need to be evaluated, but the magnitude of this vector is not directly minimized by varying the circuit parameters, only indirectly by the addition of enough operators in the ansatz. On the other hand, for optimization purposes, at the i -th iteration, only i gradient elements are considered, and the search for the energy minimum is guided by the minimization of the magnitude of this i -th dimensional gradient vector. Finally, another crucial point worth pointing out is that the commutators in Equation (6) are equivalent to those in Equation (9) only for the operator most recently added, i.e., τ_i in Equation (7).

For the purposes of an economical quantum resource utilization, it is desirable to deploy only one circuit to be used in both energy and gradient estimates (the same circuit is implemented many times, one for each term in the Hamiltonian/gradient). Even though the recursive formula in Equation (9) could, in principle, satisfy this requirement, this commutator cannot be measured (Mitarai et al., 2018). As originally proposed, the gradient is no longer given in an expectation value form, requiring an auxiliary state to be prepared via introduction of ancilla qubits, which deviates from our requirement of saving quantum resources. For that reason, we resort to numerical finite differences as means of carrying out gradient-based optimizations in the current work.

In terms of resource estimation, for a circuit depth of $\mathcal{O}(N)$, forward or backward finite differences are akin to introducing a single $R_z(h)$, where h is the step size, leading the a circuit depth of $\mathcal{O}(N + 1)$, while the use of central differences, thus, has circuit depth of $\mathcal{O}[2(N + 1)]$, the former being used here due to its superior convergence properties. This is the cost incurred in the numerical gradient estimate for each

parameter being optimized and a detailed discussion is provided in section 4.4. Such an estimate may be improved with strategies such as the quantum natural gradient (Stokes et al., 2020) or exploiting partial tomography (Parrish et al., 2019a). These ramifications are worthy of a separate study, and will not be further investigated here.

3. COMPUTATIONAL DETAILS

The quantum simulations detailed in this manuscript were performed using the VQE and ADAPT-VQE algorithms and numerical gradient strategies as implemented in the XACC hybrid quantum-classical computing framework (McCaskey et al., 2018b, 2020), with the latter algorithm leveraging a convergence criterion of $\|\mathcal{G}\| \leq 10^{-2}$. We emphasize that this parameter can be of substantial impact on the results, as it controls the size of the obtained ansatz. In light of the findings in Grimsley et al. (2019), the adopted value in this paper is believed to strike a satisfactory balance between accuracy and circuit depth. The resulting circuits were simulated via the TNQVM (tensor-network quantum virtual machine) (McCaskey et al., 2018a) XACC simulation backend and employed a noiseless, matrix product state (MPS) wave function decomposition for the quantum circuit with the aid of the ITensor library (Fishman et al., 2020). XACC provides other simulation backends, as well as physical backends targeting QPUs from IBM and Rigetti. For the size of the problems studied in this work, there may not be perceived benefits from choosing TNQVM over other XACC simulation backends like Aer (Abraham et al., 2019) or QPP (Gheorghiu, 2018). TNQVM is expected to be advantageous over other simulation approaches for problems requiring more qubits (McCaskey et al., 2018a), but we leave this to future work and do not investigate it here.

The COBYLA (Powell, 1994) algorithm was used as a gradient-free optimizer, while gradient-based optimizations were carried out with the L-BFGS algorithm (Nocedal, 1980; Liu and Nocedal, 1989), with all parameters being initialized at 0 at each optimization cycle for both optimizers. Other approaches have been reported in the literature, such as random initialization (Grimsley et al., 2020), or as in the original implementation of ADAPT-VQE (Grimsley et al., 2019) where the new parameter is initialized at 0, while the previous parameters are initialized from the optimal values obtained in the previous ADAPT iteration. XACC offers both optimizers via its interface with NLOpt (Johnson).

The potential energy curves (PEC) of NaH and KH, were generated by imposing the frozen-core approximation, reducing the number of configurations to only those arising from one σ orbital and its σ^* counterpart, that is, a two electrons in two orbitals complete active space [CAS(2,2)] problem. The one- and two-electron integrals necessary for the construction of the Hamiltonians and the corresponding references CAS energies were obtained with PySCF (Sun et al., 2017), with all calculations performed with the STO-3G basis set (Hehre et al., 1969, 1970; Pietro et al., 1980).

The quality of the output circuits in preparing the desired states is assessed via the fidelities computed with respect to the ground state full configuration interaction (FCI) wave function. This corresponds to the lowest energy eigenvector from exact diagonalization in the 2^N Hilbert space, with orbital occupation determined by the number of electrons. In possession of the circuits from the quantum simulations, the respective state vector representation is obtained using the XACC interface to the Qiskit Aer simulator (Abraham et al., 2019).

4. RESULTS AND DISCUSSION

Typically, the quality of the state obtained from the variational optimization of the gate parameters is probed indirectly by comparison of the computed energies with trustworthy references values or the exact lowest energy eigenvalue whenever computationally feasible. Thus, we start by investigating the energy profile along the atomic displacement, and subsequently contrast these findings with the analysis of the appropriateness of the corresponding states via evaluated fidelities relative to the vector corresponding to the lowest eigenvalue in the active space.

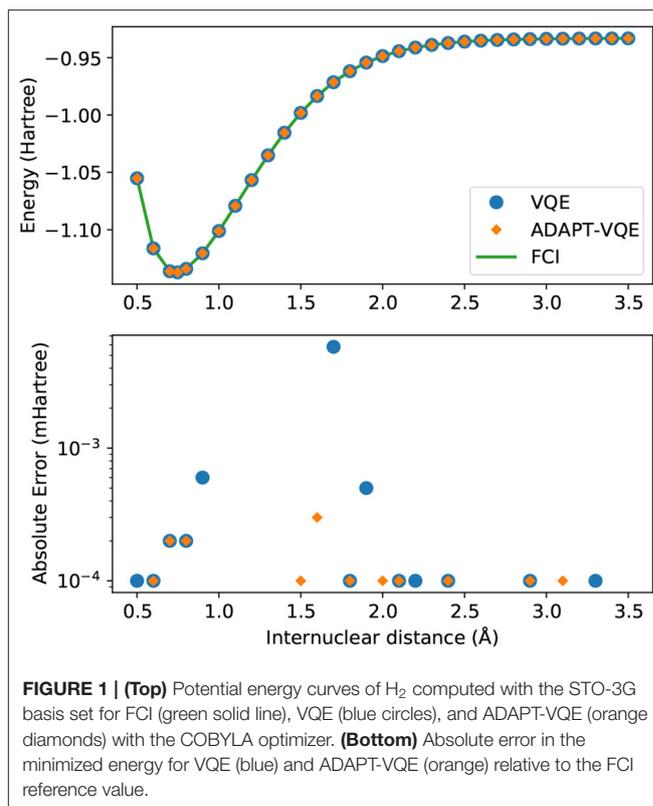
4.1. Potential Energies Curves

We start investigating the behavior of the energy by studying the H_2 molecule. This example has been extensively approached in quantum computing, and hardly poses any difficulty, at least from the standpoint of numerical simulations, as opposed to deployment to actual hardware. However, it serves as a baseline for the following discussion, as the orbital spaces in the other molecules are reduced to an active space with the goal of resembling the H_2 molecule. Results with the VQE and ADAPT-VQE ansätze are plotted in **Figure 1**, along with FCI results.

Unsurprisingly, there is a remarkable agreement between simulated and exact values, both qualitatively and quantitatively. Absolute errors from FCI are found in the sub-miliHartree range throughout the energy scan, and with either choice of ansatz, the observed errors would be inconsequential when taking into account the scale of the errors introduced by noise in the operation of quantum devices. The impression that some points are “missing” from the bottom plot of **Figure 1** is explained by these values being numerically identical to the FCI values (to seven decimal places), hence not being plotted in the logarithmic scale.

The results from the potential energy curve from simulations on the NaH molecule are presented in **Figure 2**.

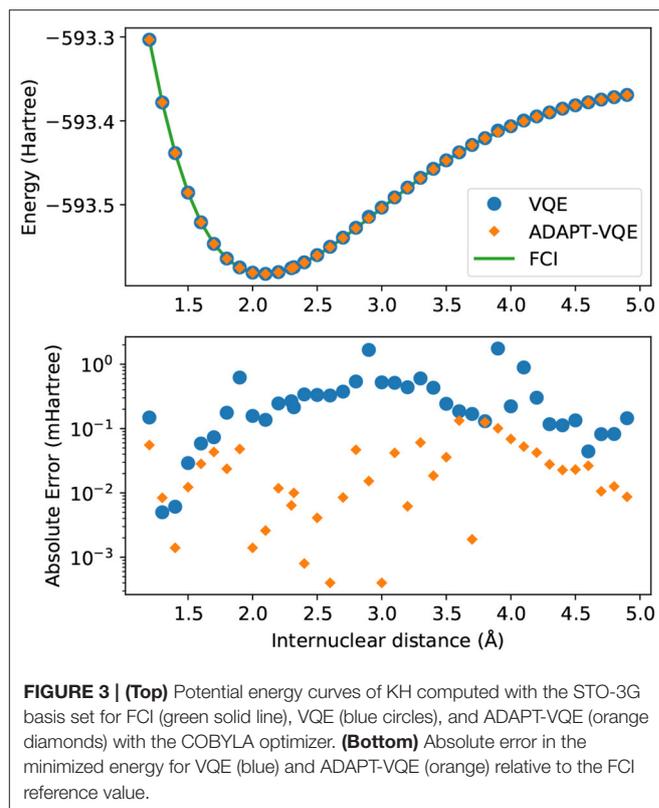
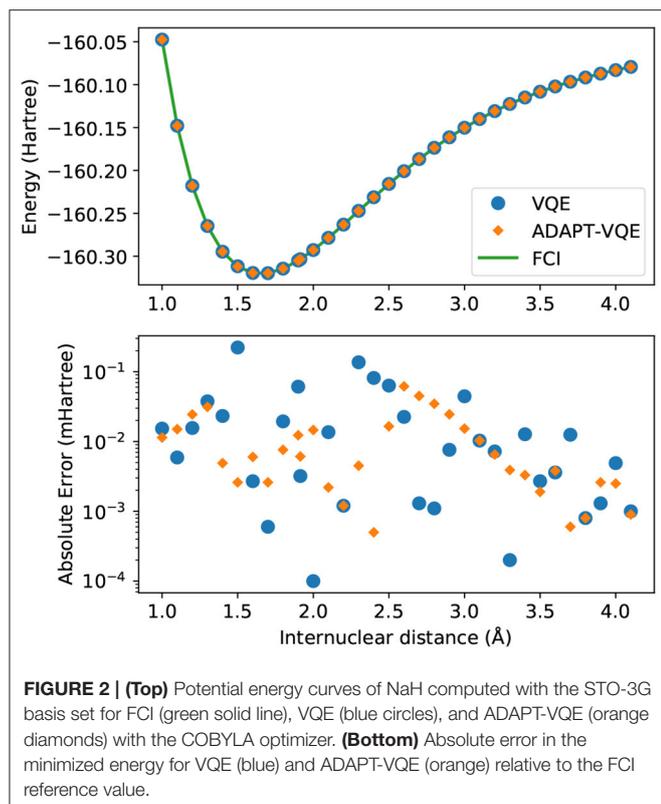
Visual inspection of the top plot reveals that the choice between the two ansätze being considered here yield energies that track one another very well, but because of the energy scale of this plot, it begs a closer look. The bottom plot displays the absolute errors between VQE and ADAPT-VQE with respect to FCI. The errors here are still within chemical accuracy (<1 kcal/mol), and are unlikely to be of much relevance in the total error if such simulations are executed in a quantum computer. However, there is a clear trend of increase in the magnitude of the computed deviations when compared to the hydrogen molecule, whose results are in **Figure 1**.



In **Figure 3**, we again observe some of the patterns that follow from the analysis of **Figures 1, 2**. The energy scale here is much too large to be able to reveal relatively minor inadequacies, even though qualitative discrepancies, such as those arising from symmetry breaking or the crossing of lines of different states, would be evident had they been present. The bottom plot, exhibiting the energy differences from FCI, offers a more reliable evidence, allowing us to infer that ADAPT-VQE is overall superior, with smaller errors for the vast majority of points (the exception being 1.4 \AA). Perhaps more importantly, we observe a general trend of the points from simulations with the plain VQE ansatz approaching the 1 mHartree , with the distances of 2.9 and 3.9 \AA now found more than 1 kcal/mol above the respective FCI energy.

4.2. Optimization Strategies

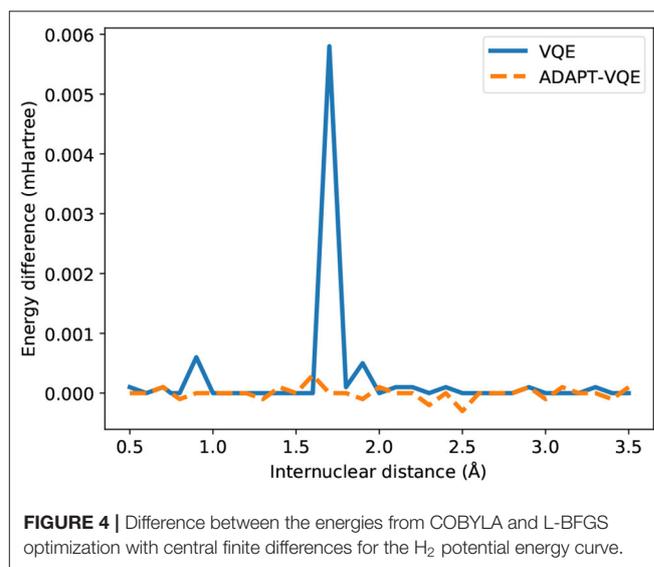
The potential energy curves presented and discussed in section 4.1 are based upon gradient-free optimization carried out with the COBYLA optimizer. We report that analogous simulations were performed with the Nelder-Mead optimizer, which is also a gradient-free alternative, but preliminary investigations pointed to COBYLA being a superior choice, at least for the chosen molecules. To contrast the performance of gradient-free optimization in the current context, we use the L-BFGS optimizer for parameter update, as implemented in `NLOpt`, with gradient estimated via central numerical finite differences. To assess the relative performance of these two approaches as the bond in the current diatomic molecules is stretched, we plot the difference



between energies obtained with the COBYLA optimizer and those with L-BFGS+finite differences, that is, $E(\text{COBYLA}) - E(\text{L-BFGS})$. That way, positive energy differences indicate there is an improvement by turning to a gradient-based optimization, while the opposite signals that the current gradient-free method reached a lower energy.

We observe compatible energies for the H_2 case, regardless of the underlying optimization strategy, for the entirety of **Figure 1**. In order to maintain consistency, we plot the energy difference between the two optimization prescriptions in a milliHartree scale, and the spike in $E(\text{COBYLA}) - E(\text{L-BFGS})$ in 1.7\AA , when rationalized with the scale in mind, shows a deviation in the $\mu\text{Hartree}$ range. Due to the presence of all the many-body operators necessary for exactness (Evangelista et al., 2019), we expect and in fact observe results on par with the numerical precision imposed by the employed optimizers (10^{-6} Hartree in relative energy).

While most of the PEC for H_2 showed no major dependence on the adopted optimization procedure, according to **Figure 4**, the picture is significantly different in the case of NaH, as portrayed in **Figure 5**. Even though the values for $E(\text{COBYLA}) - E(\text{L-BFGS})$ are still rather small, in the sub-milliHartree range, noticeable differences are more frequent here. Albeit of $\mu\text{Hartree}$ in magnitude, we also observe cases where COBYLA provides a lower energy than L-BFGS, most notably for ADAPT-VQE in the 1.4 and 2.5\AA interatomic distances. On the other hand, in an overall assessment of the performance between



VQE and ADAPT-VQE, the latter displays a more pronounced insensitivity with respect to the choice of optimization scheme.

An even more drastic contrast is found from inspection of **Figure 6**, where $E(\text{COBYLA}) - E(\text{L-BFGS})$ are plotted for the KH molecule. Some of the qualitative assertions pointed out in **Figure 5** hold, namely that the performance of VQE is much more influenced by the choice of optimization strategy than ADAPT-VQE. Not only that, but ADAPT-VQE is

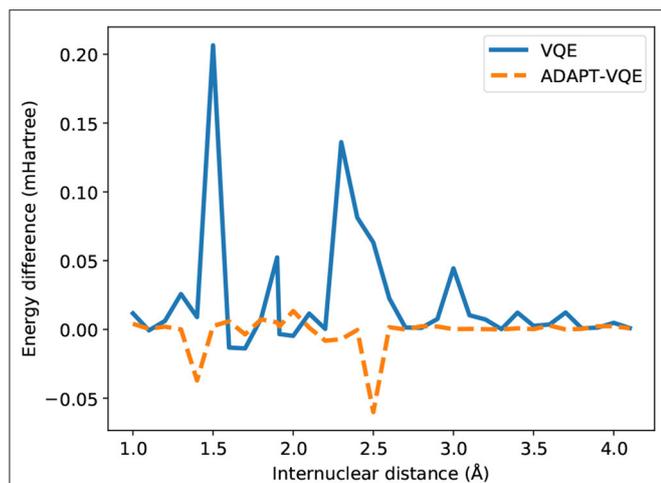


FIGURE 5 | Difference between the energies from COBYLA and L-BFGS optimization with central finite differences for the NaH potential energy curve.

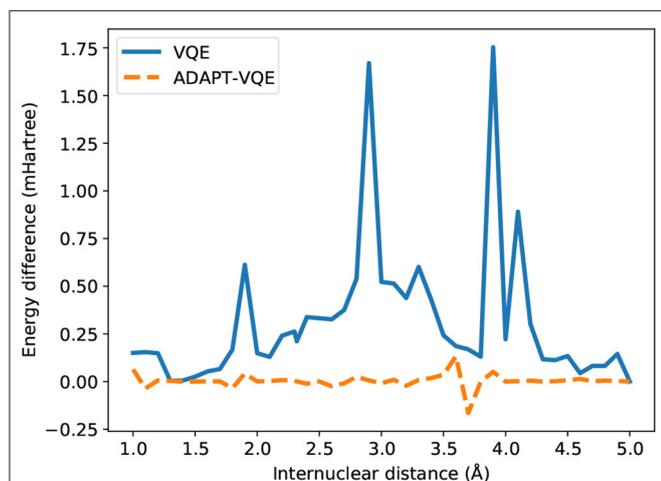


FIGURE 6 | Difference between the energies from COBYLA and L-BFGS optimization with central finite differences for the KH potential energy curve.

largely unaffected by employed optimizer, at least between the two alternatives in consideration. Here again, the differences seen for VQE correlated well with the deviations from FCI reported in **Figure 3**, further corroborating the claim that a gradient-based optimization, given the current conditions, is a more robust for approaching the lowest energy eigenvalue of molecular Hamiltonians.

4.3. State Fidelities

As previously stated, energy values can be used as valuable metric of the adequacy of a given set of variational parameters and trial state. However, the energy alone may not be indicative of the quality of the corresponding state and even acceptable energy values do not guarantee equally satisfactory values for other properties. The usual electronic Hamiltonian, as shown in Equation (2), transforms as the most symmetric irreducible

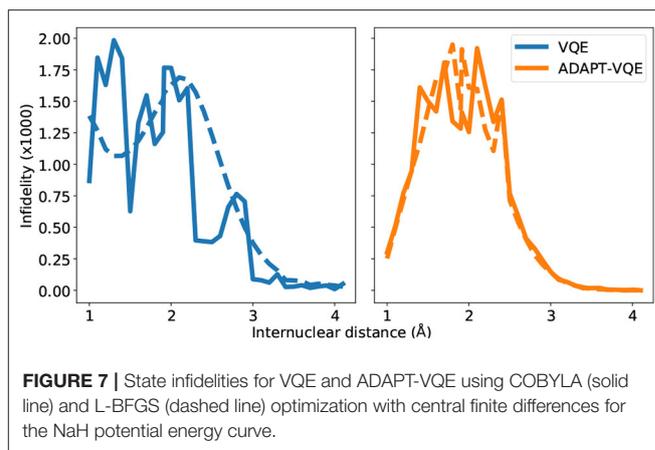


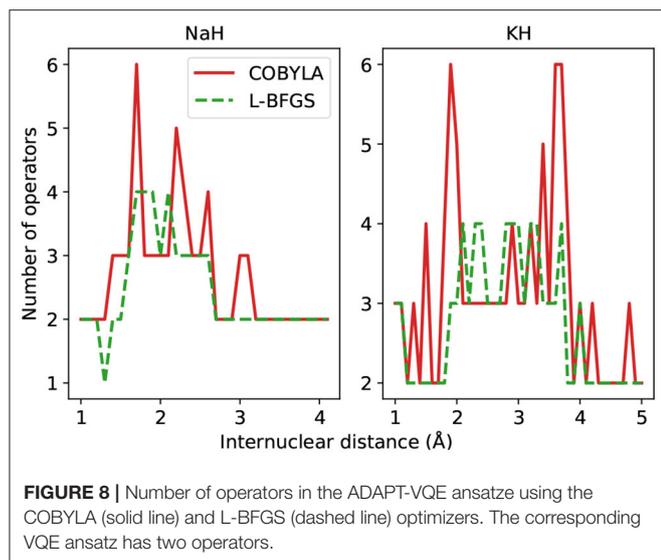
FIGURE 7 | State infidelities for VQE and ADAPT-VQE using COBYLA (solid line) and L-BFGS (dashed line) optimization with central finite differences for the NaH potential energy curve.

representation for a given point group, therefore yielding the same energy in the case of degenerate states. Other operators, however, such as the terms in the multipole expansion of the electric potential, do not display this feature, meaning that degenerate states may yield different expectation values for such operators.

In order to examine the state prepared by the two circuit approaches considered here, we compute their “infidelities” with respect to the exact FCI state within the aforementioned active spaces, which is mathematically represented by $1 - |\langle \Psi_{\text{FCI}} | \hat{U}(\vec{\theta}) | 0 \rangle|$, where $\vec{\theta}$ here are the set of optimal values also utilized for the energy computations in sections 4.1 and 4.2. We acknowledge that, while this provides direct inroads in the state being output at completion of the state preparation, it cannot be experimentally realized. However, in the case of moderately sized molecules for which the exact diagonalization of the Hamiltonian is feasible, this can provide valuable insights.

The energy differences discussed in the case of the hydrogen molecule in sections 4.1 and 4.2 are quite small when considering the magnitude of the other potential sources of error that can arise in the presence of noise, either through a model or in the operation of an actual quantum device. Due to the simplicity of the electronic structure of this molecule the state prepared according to the two ansatz construction prescriptions investigated here yield infidelities that are below the numerical thresholds employed here, and certainly would be unnoticeable for realistic purposes. Because they offer little insight, we abstain from plotting the infidelity results for this molecule here.

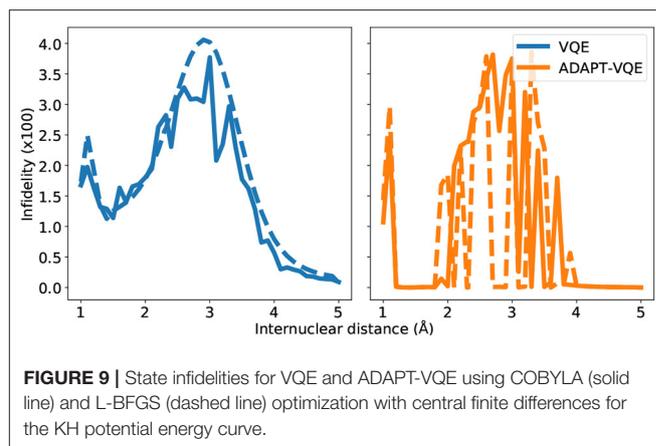
Before delving into the particularities of each curve in **Figure 7**, we bring the reader’s attention to the scale of the plots, signaling a remarkable agreement between the state prepared and the one expected (FCI). It should come as no surprise that the largest infidelities are found in the vicinity of the Coulson-Fischer point, the most demanding region in the energy landscape, and subsequently approach zero as the atoms are moved far apart. The infidelities for the VQE ansatz follow a smooth progression when employed in conjunction with the gradient-based optimizer L-BFGS, whereas the same is not true for the other combinations of ansatz and optimization. This is



likely a compound effect, explained by the former being a fixed circuit, where only the associate $\bar{\theta}$ changes throughout the energy scan. The latter, however, can assume a different composition, changing according to the demands of the electronic structure at each bond length. This works along with the fact that gradient-based optimization, at least in the current study, provides a tighter, more reliable solution. For the NaH and KH cases, we plot the number of operators in the ansatz in **Figure 8**.

Once again, there is a clear advantage in turning to gradient-based optimization, as it renders ansätze with fewer operators. For some internuclear distances, the ADAPT-VQE ansätze, even when optimized with L-BFGS, contain more operators than the corresponding VQE ansatz. This is not necessarily in contradiction with some of the findings from Grimsley et al. (2019) because those results were obtained for different molecules and using different optimization implementations. Yet, we would expect that when comparing against a larger VQE problem, such as those investigated in that paper, we would see similar trends. We also speculate that another variable that can contribute to the observed behavior is the tolerance that controls how tight the optimization should be. Because we are using a default 10^{-6} threshold in relative energy as the tolerance and there is no clear connection between the quality of the energies and the respective prepared states, the absolute energies values may fall in a scale that may have a small, but non-negligible effect on the fidelities, which is also evidence of the effect it can have in the output state, further corroborated by the number of operators found in the respective ansätze, yet not enough to alter any of the main conclusions drawn from the results presented here.

Many of the main inferences from the analysis of the **Figure 7** hold for the KH molecule, whose infidelities are shown in **Figure 9**. Firstly, the infidelities, though still quite small, are about an order of magnitude larger. The smoothness and overall profile observed for the VQE UCCSD is retained, but the behavior of the ADAPT-VQE infidelities is much more erratic. Secondly, while the ADAPT-VQE ansatz for NaH around



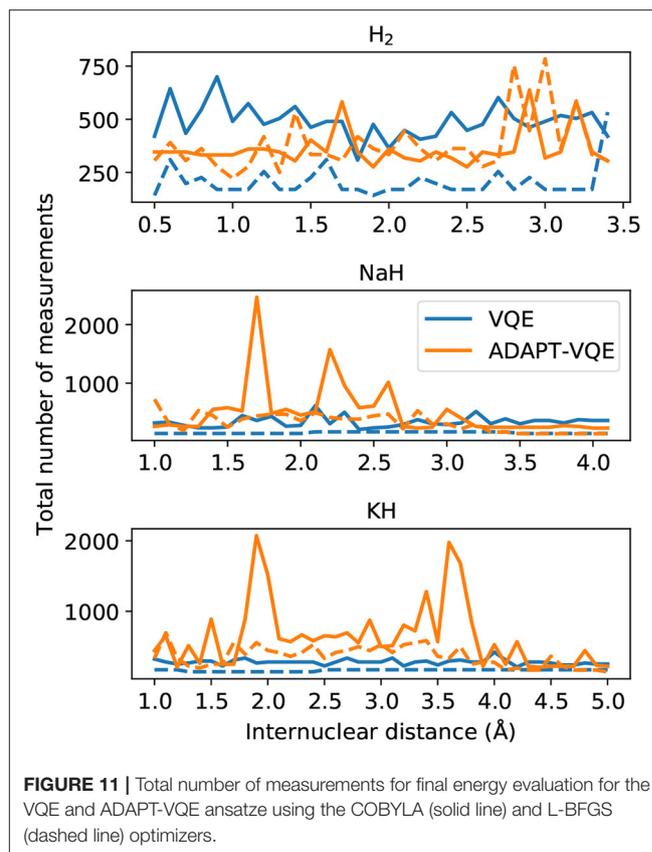
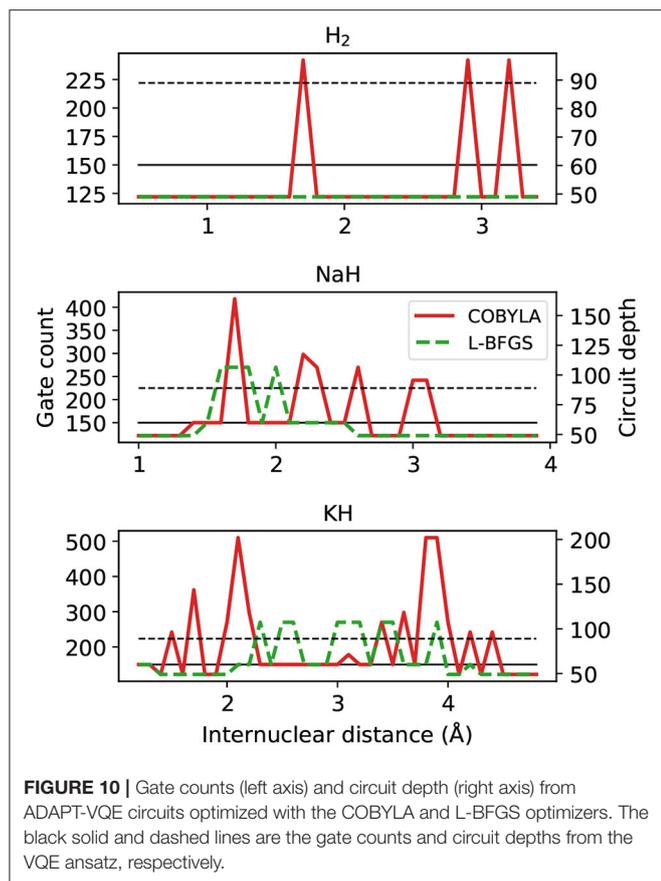
the Coulson-Fischer point is mostly the same, but the larger number of variational parameters make it more vulnerable to the optimization inconsistencies discussed above, here the large oscillations are due to ansätze of alternating operator compositions. Because the ADAPT-VQE convergence criterion depends upon a fixed numerical threshold, sometimes the ansatz at a given iteration may already be close to convergence, but still not quite below the gradient norm threshold, and upon the addition of an extra operator, the state may be improved significantly in the scale of the plots seen in this section.

4.4. Resource Estimation

One of the main motivations behind the present work is to serve as the baseline for following studies focusing on the investigation of the electronic structure of molecules carried out in NISQ devices. With this in mind, it is important to develop some intuition on the resource demands involved in such tasks.

First we analyze the circuit proposed by VQE and ADAPT-VQE to prepare the states whose energies and fidelities were shown in sections 4.1–4.3 in terms of total gate count and circuit depth, plotted in **Figure 10**.

Let us first compare the ADAPT-VQE results on the basis of the two optimizers. As we move from H_2 to NaH and KH, we see a more intricate picture of how these optimizers impact the final circuit. Qualitatively, L-BFGS has an overall advantage as it provides circuits that are shallower and with fewer gates. While there are a few points along the potential energy scans where the circuits generated based on L-BFGS are not as efficient as those from a COBYLA optimization, the scales of the plots are determined solely by the latter. We noticed that in several points, the simulations with the COBYLA optimizer would produce states with two instances of the same operator adjacent to each other. If the actual minimum value had been achieved in a certain iteration of ADAPT-VQE, the commutator of the same operator in the next iteration would have been zero. Because this procedure is accomplished numerically, the magnitude of this commutator is related to how close the determined minimum is from the actual one. It turns out that the default threshold in relative energy (10^{-6}) is found not to be stringent enough, which incurs a commutator whose deviation from the expected



zero is non-negligible, resulting in the same operator being added in successive iterations. Another factor that accounts for the displayed circuit figures is the fixed gradient norm threshold in ADAPT-VQE. In some iterations, this quantity is above, but already quite close to the pre-defined 10^{-2} , and one extra iteration is performed, with only marginal energy improvement. To illustrate this, the ADAPT-VQE simulation for NaH with internuclear separation of 1.8\AA converges to ansatz with three operators, with $E = -160.3146751$ Hartree and $\|\mathcal{G}\| = 0.001$. Had the ADAPT cycle been stopped in the second iteration, we would have $\|\mathcal{G}\| = 0.013$, with $E = -160.3146492$ Hartree, that is, the energy improvement was in the μ Hartree range, yet at the expense of a deeper circuit, which calls for a more flexible operator selection in ADAPT-VQE.

These resource estimation parameters in **Figure 10** are comparable between the two ansatz strategies. In general terms, the circuits optimized upon L-BFGS are more affordable than the corresponding VQE ones, while using COBYLA tends to yields circuits that are deeper and need to implement more gates. We bring attention to the fact that there is not a one-to-one correspondence between the present analysis and that in the Figures 2C,E,I in Grimsley et al. (2019). This is because the latter refers to the number of parameters/operators in the ansatz. A circuit with more parameters/operators does not readily translate into a more complex circuit, which depends on the number of qubits in a given operator and the operator locality

and placement. This means these results are not at odds with what was previously reported, which were obtained for a distinct set of molecules, but can be seen as complementary.

Another important metric when estimating the necessary resources for implementation and deployment of the simulations discussed here is the number of measurements. To complement the end of the last paragraph, it is important to mention that in this context the rationalization in terms of number of operators increases in relevance. In **Figure 11**, we plot the total number of measurements to achieve the results reported in section 4.1.

As pointed out in Grimsley et al. (2019), the ansatz put forth by ADAPT-VQE offers a trade-off between circuit depth and number of measurements. We can readily confirm by visual inspection of **Figure 11** that ADAPT-VQE incurs a much larger number of measurements. These figures account for all measurements involved in computing the commutators in Equation (6), the energy evaluations at each optimization iteration, and the computations necessary to minimize the gradients when L-BFGS is employed. The measurement burden in ADAPT-VQE reported here can be partially alleviated by employing a better parameter initialization, such as starting the VQE optimization at each iteration with the previously optimized parameters and initializing just the newly added parameter at zero. This demand is also expected to be greatly relieved by resorting to a different set of operators, such as those introduced in the qubit-ADAPT-VQE variant (Tang et al., 2020), which can

still span the underlying Hilbert space, yet with linear growth in the number of qubits. This approach would require much fewer commutator computations at each iteration, but would likely be of noticeable advantage for operator pools larger than those in question here. These results are also contingent upon the choice of optimizer, and there may exist better suited choices than those investigated here. Yet, we do not believe this would dramatically change the overall qualitative picture drawn in **Figure 11**.

Another key outcome from the analysis of **Figure 11** is the fact that, even though the gradient computation with L-BFGS requires more measurements per iteration, it is overall much more economical than the gradient-free optimization, represented here by COBYLA. This further strengthens the case for gradient-based optimization in VQE, as it not only results in smaller errors/better convergence with respect to the sought ground state, but it is also much less demanding from a resource standpoint.

5. CONCLUSION

For a broader adoption of adaptive methods for ansatz construction in the realm of quantum chemistry, and perhaps, for many-body methods in general, many aspects still need to be explored and their underpinnings better understood. This work provides a contribution toward this goal by showing a comprehensive study of potential energy curves of a selection of molecules of the general formula XH ($X = \text{H, Na, K}$). Despite their simplicity, they serve to shed light on some of the mentioned characteristics, and deliver a baseline for feasible studies involving actual quantum hardware.

Even a relatively conservative gradient norm threshold of 10^{-2} in ADAPT-VQE is sufficient to provide overall better energetics than corresponding fixed ansatz approach embodied by the ordinary VQE, which is in agreement with the initial ADAPT-VQE proposal. Due to its iterative nature, ADAPT-VQE has an extra layer of tunability which can be controlled via the threshold on $\|\mathcal{G}\|$. This means that the errors observed with ADAPT-VQE might have been reduced had $\|\mathcal{G}\|$ been made tighter, which could in turn increase the depth of the circuits, and even having to cope with more necessary measurements than those of UCCSD, as suggested with $\|\mathcal{G}\| = 10^{-3}$ in Figure 2i by Grimsley et al. (2019) However, upon a simple choice of gradient strategy motivated by the constraints of quantum hardware, we report that ADAPT-VQE is fairly resilient with respect to the employed optimization strategy and that encouraging improvements in performance by adopting a gradient-based approach in the search of the parameter set that minimizes the objective function can be mostly beneficial in the case of VQE. These findings call for a follow-up study on the role of optimizer in conjunction with ADAPT-VQE, extending the analysis to a larger selection of optimizers and gradient strategies.

The ongoing development of VQE methods, including ADAPT-VQE, must also address the noise that is intrinsic to the operations implemented in experimental quantum computers. The above benchmarks of infidelity and energy error place lower bounds on the expected accuracy for VQE methods

using noiseless numerical simulation. However, we anticipate that the introduction of noise will substantially affect the accuracy with which the prepared ansatz state approaches the pure state expected from conventional quantum chemistry theory. However, if the state infidelity grows with increasing molecular size, as indicated by our short series of examples, then lower bounds on ansatz accuracy may become a non-trivial contribution to observed errors in experimental measurements.

DATA AVAILABILITY STATEMENT

The raw data supporting the conclusions of this article will be made available by the authors, without undue reservation. XACC and TNQVM are open-source and can be found at <https://github.com/eclipse/xacc> and <https://github.com/ornl-qci/tnqvm>.

AUTHOR'S NOTE

This manuscript has been authored by UT-Battelle, LLC, under Contract No. DE-AC0500OR22725 with the U.S. Department of Energy. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for the United States Government purposes. The Department of Energy will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan.

AUTHOR CONTRIBUTIONS

DC implemented the ADAPT-VQE algorithm, ran some of the simulations, and wrote the manuscript. JW ran some of the simulations, wrote the code to compute state fidelities, and generated the plots. AM implemented the VQE algorithm and oversaw the ADAPT-VQE implementation. TH designed the research and helped writing the manuscript.

ACKNOWLEDGMENTS

This work was supported by the Embedding Quantum Computing into Many-body Frameworks for Strongly Correlated Molecular and Materials Systems project, which was funded by the U.S. Department of Energy (DOE), Office of Science, Office of Basic Energy Sciences, the Division of Chemical Sciences, Geosciences, and Biosciences. This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facilities supported by the Oak Ridge National Laboratory under Contract DE-AC05-00OR22725. This research used resources of the Compute and Data Environment for Science (CADES) at the Oak Ridge National Laboratory, which was supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725. This work was carried out at Oak Ridge National Laboratory, managed by UT-Battelle, LLC for the U.S. Department of Energy under contract DE-AC05-00OR22725.

REFERENCES

- Abraham, H., et al. (2019). Qiskit: An open-source framework for quantum computing.
- Bartlett, R. J., Kucharski, S. A., and Noga, J. (1989). Alternative coupled-cluster ansatz ii. The unitary coupled-cluster method. *Chem. Phys. Lett.* 155, 133–140. doi: 10.1016/S0009-2614(89)87372-5
- Britt, K. A., and Humble, T. S. (2017). High-performance computing with quantum processing units. *ACM J. Emerg. Technol. Comput. Syst.* 13, 1–13. doi: 10.1145/3007651
- Cao, Y., Romero, J., Olson, J. P., Degroote, M., Johnson, P. D., Kieferová, M., et al. (2019). Quantum chemistry in the age of quantum computing. *Chem. Rev.* 119, 10856–10915. doi: 10.1021/acs.chemrev.8b00803
- Chivilikhin, D., Samarin, A., Ulyantsev, V., Iorsh, I., Oganov, A. R., and Kyriienko, O. (2020). MoG-VQE: Multiobjective genetic variational quantum eigensolver. *arXiv [Preprint]*. arXiv:2007.04424.
- Evangelista, F. A., Chan, G. K.-L., and Scuseria, G. E. (2019). Exact parameterization of fermionic wave functions via unitary coupled cluster theory. *J. Chem. Phys.* 151:244112. doi: 10.1063/1.5133059
- Fishman, M., White, S. R., and Stoudenmire, E. M. (2020). The ITensor software library for tensor network calculations. *arXiv [Preprint]*. arXiv:2007.14822.
- Gheorghiu, V. (2018). Quantum++: A modern C++ quantum computing library. *PLoS ONE* 13:e0208073. doi: 10.1371/journal.pone.0208073
- Grimsley, H. R., Claudino, D., Economou, S. E., Barnes, E., and Mayhall, N. J. (2020). Is the trotterized uccsd ansatz chemically well-defined? *J. Chem. Theory Comput.* 16, 1–6. doi: 10.1021/acs.jctc.9b01083
- Grimsley, H. R., Economou, S. E., Barnes, E., and Mayhall, N. J. (2019). An adaptive variational algorithm for exact molecular simulations on a quantum computer. *Nat. Commun.* 10:3007. doi: 10.1038/s41467-019-10988-2
- Hehre, W. J., Ditchfield, R., Stewart, R. F., and Pople, J. A. (1970). Self-consistent molecular orbital methods. IV. Use of gaussian expansions of Slater-type orbitals. extension to second-row molecules. *J. Chem. Phys.* 52, 2769–2773. doi: 10.1063/1.1673374
- Hehre, W. J., Stewart, R. F., and Pople, J. A. (1969). Self-consistent molecular-orbital methods. I. Use of gaussian expansions of Slater-type atomic orbitals. *J. Chem. Phys.* 51, 2657–2664. doi: 10.1063/1.1672392
- Humble, T. S., Thapliyal, H., Munoz-Coreas, E., Mohiyaddin, F. A., and Bennink, R. S. (2019). Quantum computing circuits and devices. *IEEE Design Test* 36, 69–94. doi: 10.1109/MDAT.2019.2907130
- Jeziorski, B. and Monkhorst, H. J. (1981). Coupled-cluster method for multideterminantal reference states. *Phys. Rev. A* 24, 1668–1681. doi: 10.1103/PhysRevA.24.1668
- Johnson, S. *The NLOpt Nonlinear-Optimization Package*.
- Jordan, P., and Wigner, E. (1928). Über das paulische Äquivalenzverbot. *Zeitsch. Phys.* 47, 631–651. doi: 10.1007/BF01331938
- Kandala, A., Mezzacapo, A., Temme, K., Takita, M., Brink, M., Chow, J. M., et al. (2017). Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature* 549:242. doi: 10.1038/nature23879
- Kutzelnigg, W. (1991). Error analysis and improvements of coupled-cluster theory. *Theor. Chim. Acta* 80, 349–386. doi: 10.1007/BF01117418
- Liu, D., and Nocedal, J. (1989). On the limited memory BFGS method for large scale optimization. *Math. Programm.* 45, 503–528. doi: 10.1007/BF01589116
- Lyakh, D. I., Musial, M., Lotrich, V. F., and Bartlett, R. J. (2012). Multireference nature of chemistry: the coupled-cluster view. *Chem. Rev.* 112, 182–243. doi: 10.1021/cr2001417
- McArdle, S., Endo, S., Aspuru-Guzik, A., Benjamin, S. C., and Yuan, X. (2020). Quantum computational chemistry. *Rev. Modern Phys.* 92:015003. doi: 10.1103/RevModPhys.92.015003
- McCaskey, A., Dumitrescu, E., Chen, M., Lyakh, D., and Humble, T. (2018a). Validating quantum-classical programming models with tensor network simulations. *PLoS ONE* 13:e206704. doi: 10.1371/journal.pone.0206704
- McCaskey, A., Dumitrescu, E., Liakh, D., Chen, M., Feng, W., and Humble, T. (2018b). A language and hardware independent approach to quantum-classical computing. *SoftwareX* 7, 245–254. doi: 10.1016/j.softx.2018.07.007
- McCaskey, A., Lyakh, D., Dumitrescu, E., Powers, S., and Humble, T. (2020). XACC: a system-level software infrastructure for heterogeneous quantum-classical computing. *Quant. Sci. Technol.* 5:024002. doi: 10.1088/2058-9565/ab6bf6
- McClean, J. R., Romero, J., Babbush, R., and Aspuru-Guzik, A. (2016). The theory of variational hybrid quantum-classical algorithms. *New J. Phys.* 18:023023. doi: 10.1088/1367-2630/18/2/023023
- Mitarai, K., Negoro, M., Kitagawa, M., and Fujii, K. (2018). Quantum circuit learning. *Phys. Rev. A* 98:032309. doi: 10.1103/PhysRevA.98.032309
- Nocedal, J. (1980). Updating quasi-newton matrices with limited storage. *Math. Comput.* 35, 773–773. doi: 10.1090/S0025-5718-1980-0572855-7
- Parrish, R., Iosue, J., Ozaeta, A., and McMahon, P. (2019a). A jacobi diagonalization and anderson acceleration algorithm for variational quantum algorithm parameter optimization. *arXiv [Preprint]*. arXiv:1904.03206 [quant-ph].
- Parrish, R. M., Hohenstein, E. G., McMahon, P. L., and Martínez, T. J. (2019b). Quantum computation of electronic transitions using a variational quantum eigensolver. *Phys. Rev. Lett.* 122:230401. doi: 10.1103/PhysRevLett.122.230401
- Peruzzo, A., McClean, J., Shadbolt, P., Yung, M.-H., Zhou, X.-Q., Love, P. J., et al. (2014). A variational eigenvalue solver on a photonic quantum processor. *Nat. Commun.* 5:4213. doi: 10.1038/ncomms5213
- Pietro, W. J., Levi, B. A., Hehre, W. J., and Stewart, R. F. (1980). Molecular orbital theory of the properties of inorganic and organometallic compounds. 1. STO-NG basis sets for third-row main-group elements. *Inorgan. Chem.* 19, 2225–2229. doi: 10.1021/ic50210a005
- Powell, M. J. D. (1994). “A direct search optimization method that models the objective and constraint functions by linear interpolation,” in *Advances in Optimization and Numerical Analysis*, eds S. Gomez and J.-P. Hennart (Dordrecht: Springer Netherlands), 51–67. doi: 10.1007/978-94-015-8330-5_4
- Preskill, J. (2018). Quantum Computing in the NISQ era and beyond. *Quantum* 2:79. doi: 10.22331/q-2018-08-06-79
- Quantum, G. A., et al. (2020). Hartree-fock on a superconducting qubit quantum computer. *Science* 369, 1084–1089. doi: 10.1126/science.abb9811
- Raghavachari, K., Trucks, G. W., Pople, J. A., and Head-Gordon, M. (1989). A fifth-order perturbation comparison of electron correlation theories. *Chem. Phys. Lett.* 157, 479–483. doi: 10.1016/S0009-2614(89)87395-6
- Romero, J., Babbush, R., McClean, J. R., Hempel, C., Love, P. J., and Aspuru-Guzik, A. (2018). Strategies for quantum computing molecular energies using the unitary coupled cluster ansatz. *Quant. Sci. Technol.* 4:014008. doi: 10.1088/2058-9565/aad3e4
- Ryabinkin, I. G., Lang, R. A., Genin, S. N., and Izmaylov, A. F. (2020). Iterative qubit coupled cluster approach with efficient screening of generators. *J. Chem. Theory Comput.* 16, 1055–1063. doi: 10.1021/acs.jctc.9b01084
- Shavitt, I., and Bartlett, R. J. (2009). *Many-Body Methods in Chemistry and Physics: MBPT and Coupled-Cluster Theory*. Cambridge Molecular Science. Cambridge University Press. doi: 10.1017/CBO9780511596834
- Stokes, J., Izaac, J., Killoran, N., and Carleo, G. (2020). Quantum natural gradient. *Quantum* 4:269. doi: 10.22331/q-2020-05-25-269
- Sun, Q., Berkelbach, T. C., Blunt, N. S., Booth, G. H., Guo, S., Li, Z., et al. (2017). PySCF: the python-based simulations of chemistry framework. *Wires Comput. Mol. Sci.* 8:e1340. doi: 10.1002/wcms.1340
- Tang, H. L., Shkolnikov, V. O., Barron, G. S., Grimsley, H. R., Mayhall, N. J., Barnes, E., et al. (2020). qubit-ADAPT-VQE: An adaptive algorithm for constructing hardware-efficient ansätze on a quantum processor. *arXiv [Preprint]*. arXiv:1911.10205.
- Taube, A. G., and Bartlett, R. J. (2006). New perspectives on unitary coupled-cluster theory. *Int. J. Quant. Chem.* 106, 3393–3401. doi: 10.1002/qua.21198

- Urban, M., Noga, J., Cole, S. J., and Bartlett, R. J. (1985). Towards a full CCSDT model for electron correlation. *J. Chem. Phys.* 83, 4041–4046. doi: 10.1063/1.449067
- Watts, J. D., Gauss, J., and Bartlett, R. J. (1993). Coupled-cluster methods with noniterative triple excitations for restricted open-shell hartree-fock and other general single determinant reference functions. Energies and analytical gradients. *J. Chem. Phys.* 98, 8718–8733. doi: 10.1063/1.464480
- Zhu, L., Tang, H. L., Barron, G. S., Mayhall, N. J., Barnes, E., and Economou, S. E. (2020). An adaptive quantum approximate optimization algorithm for solving combinatorial problems on a quantum computer. *arXiv [Preprint]*. arXiv:2005.10258.

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2020 Claudino, Wright, McCaskey and Humble. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.