

Calculation of Molecular Vibrational Spectra on a Quantum Annealer

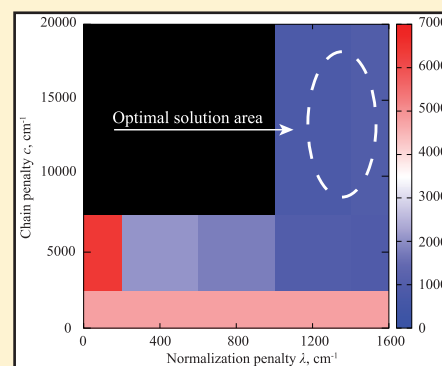
Alexander Teplukhin,[†] Brian K. Kendrick,^{*,†} and Dmitri Babikov[‡]

[†]Theoretical Division (T-1, MS B221), Los Alamos National Laboratory, Los Alamos, New Mexico 87545, United States

[‡]Department of Chemistry, Marquette University, Milwaukee, Wisconsin 53021, United States

S Supporting Information

ABSTRACT: Until recently molecular energy calculations using quantum computing hardware have been limited to gate-based quantum computers. In this paper, a new methodology is presented to calculate the vibrational spectrum of a molecule on a quantum annealer. The key idea of the method is a mapping of the ground state variational problem onto an Ising or quadratic unconstrained binary optimization (QUBO) problem by expressing the expansion coefficients using spins or qubits. The algorithm is general and represents a new revolutionary approach for solving the real symmetric eigenvalue problem on a quantum annealer. The method is applied to two chemically important molecules: O₂ (oxygen) and O₃ (ozone). The lowest two vibrational states of these molecules are computed using both a hardware quantum annealer and a software based classical annealer. Extension of the algorithm to higher dimensions is explicitly demonstrated for an *N*-dimensional harmonic oscillator (*N* ≤ 5). The algorithm scales exponentially with dimensionality if a direct product basis is used but will exhibit polynomial scaling for a nondirect product basis.



1. INTRODUCTION

Quantum computers are seen by many as a future alternative to classical computers. Although quantum supremacy has not yet been achieved, the field is advancing quite rapidly. There are two major types of quantum computing devices available today:¹ quantum annealer² and universal quantum computer based on quantum gates.^{3–5} The first type is an example of adiabatic quantum computing⁶ and is used to solve optimization problems, which at first glance appears to be quite restrictive. The second type is based on quantum gates, which appears to have a wider applicability and therefore may be able to simulate a larger variety of problems. However, adiabatic and gate-based quantum computing were proven to be formally equivalent.⁷ Thus, the practical application space is most likely limited by the hardware realization and not necessarily by the type of approach. In either approach, the current generation of quantum computing devices has significant noise and supports a small number of fully coupled qubits (<100). Hence, they are often referred to collectively as NISQ (Noisy Intermediate Scale Quantum) devices, and their accuracy and problem size is limited.

Coming from the physical chemistry community, we asked ourselves if it would be possible to program an important fundamental problem on a quantum annealer such as the commercially available D-Wave machine.⁸ Typically, people who work with such devices go in the opposite direction: knowing hardware capabilities they come up with a suitable optimization problem. As a fundamental problem we chose to calculate the vibrational ground state and possibly excited

states of a molecule. This problem is very important in chemistry, for example: H_n⁺ ions,^{9–12} CH₅⁺ and isotopologues,^{13,14} H₃O⁺, H₂O₂⁺, and deuterated analogues,^{15,16} hydrogen clusters,^{17–19} their isotopologues,^{20,21} hydrogen bonded systems,²² and Lennard-Jones clusters.^{23,24} The common method to study these molecular systems is a Monte Carlo (MC) method in its various flavors: variational MC, time-dependent variational MC, diffusion MC, and path integral MC.

Recently, a ground state problem in electronic structure theory was implemented on both types of quantum computing devices: a quantum annealer^{25,26} and a gate-based quantum computer.^{27–29} In the first study, the mapping of electronic Hamiltonian to quantum annealer Hamiltonian is achieved by means of creation and annihilation Fermionic operators followed by transformation to spin operators and reduction to the form that includes pairwise interactions between qubits. An iterative algorithm is used to find the lowest ground state energy. In the second work, to approach the same problem on a gate-based quantum computer, an expectation value of each term in the electronic Hamiltonian is evaluated on a trial wave function using a quantum device, and the resultant total energy serves as a guide for generating the next trial wave function. The optimization step of this Variational Quantum Eigensolver (VQE), namely the trial generation, is performed on a classical

Received: April 25, 2019

Published: July 17, 2019

computer. The iterative nature of these algorithms makes them both hybrid.

In contrast to the electronic structure algorithms discussed above, the new algorithm presented in this work is general and solves any real symmetric eigenvalue problem. To our knowledge, this is the first general quantum annealer based eigenvalue solver and will be referred to below as the Quantum Annealer Eigensolver (QAE). As discussed in more detail below, our QAE algorithm is also hybrid since the variational eigenvalue problem is solved via a sequence of many quantum annealer optimizations performed with varying weights on the constraint equations (i.e., Lagrange multipliers). The scanning and optimization of the weights is done on a classical computer. Mapping the eigenvalue problem to a quantum annealer hardware is nontrivial, because the annealer solves a minimization problem defined by an Ising functional of the form $H(s) = \sum_i h_i s_i + \sum_{i<j} J_{ij} s_i s_j$ where the spin variables s_i accept discrete values $\{-1,1\}$. Alternatively, the functional can be converted to quadratic unconstrained binary optimization (QUBO) form using discrete variables $x_i \in \{0,1\}$, called qubits, giving $H(x) = \sum_i Q_{ii} x_i + \sum_{i<j} Q_{ij} x_i x_j$.³⁰ The problem is how to write down a ground state or eigenvalue problem in QUBO form and explicitly construct the matrix \mathbf{Q} .

Thus, the primary goal of this paper is to demonstrate a mapping of the eigenvalue problem to the QUBO problem, including the treatment of excited states and multiple dimensions. For simplicity, we use a direct product basis set for problems with more than one dimension. As a consequence, this choice of basis set causes QAE to scale exponentially with system size. However, choosing a nondirect product basis set (for example, a MC-sampled set of grid points) will give polynomial scaling. Since the D-Wave machine has to be queried multiple times to accumulate enough statistics, MC sampling could be interleaved with the hardware sampling (this approach was not examined in the present study).

The outline of the paper is as follows: First, we present our solution to this problem, including the extension to the excited state calculations and multiple dimensions. Second, we apply our algorithm to two chemically important species, O_2 (oxygen) and O_3 (ozone). For the ozone calculation, a reduction in the number of qubits was required in order to fit the problem on the D-wave machine. Third, the introduction of weighted constraints is presented following a technique used to overcome the connectivity issue in the quantum annealer hardware (i.e., D-Wave machine). Noise is also modeled in the algorithm which is shown to reproduce the results from the D-Wave machine. In the final discussion section, we consider possible improvements of the algorithm and sources of error.

2. QUANTUM ANNEALER EIGENSOLVER ALGORITHM

2.1. Mapping of a Ground State Problem to a QUBO Problem. The method is inspired by the variational principle. Suppose we are interested in a ground state of a one-dimensional system and its wave function Ψ is expanded using an orthonormal basis φ_α and unknown expansion coefficients a_α : $\Psi = \sum_{\alpha=1}^B a_\alpha \varphi_\alpha$. Then, the ground state energy can be expressed as a double sum over the Hamiltonian matrix elements $E = \langle \Psi | \hat{H} | \Psi \rangle = \sum_{\alpha,\beta} a_\alpha a_\beta \langle \varphi_\alpha | \hat{H} | \varphi_\beta \rangle = \sum_{\alpha,\beta} a_\alpha a_\beta H_{\alpha\beta}$. It is easy to see that the functional form for the energy E is similar to the QUBO form $H(s)$, except that the coefficients a_α are continuous and $a_\alpha \in [-1;1]$ (since from the normalization

condition $\langle \Psi | \Psi \rangle = 1$ we know $\sum_{\alpha=1}^B a_\alpha^2 = 1$). In contrast, the QUBO variables x_i are discrete.

The key idea in mapping the eigenvalue problem with Hamiltonian matrix \mathbf{H} to the QUBO optimization problem with matrix \mathbf{Q} is to express each expansion coefficient a_α using K qubits $q_k^\alpha \in \{0,1\}$. This approximation can be done in multiple ways. The approach we followed in this work is a fixed-point representation that is used to represent real numbers in a classical computer. Since the magnitude of the coefficients a_α never exceeds unity, only the fractional part of the coefficient has to be stored. The last qubit q_K^α stores the sign of a_α . The complete expression for the coefficient is $a_\alpha = \sum_{k=1}^{K-1} 2^{k-K} q_k^\alpha - q_K^\alpha \in [-1; 1)$. Now, the functional E can be expressed explicitly in terms of the qubits q_k^α . The powers of two are combined with the matrix elements $H_{\alpha\beta}$ giving the matrix elements Q_{ij} . The qubits q_k^α are mapped to the qubits x_i via the relation $i = K(\alpha - 1) + k$, where $i \in [1, B \times K]$. Since in the QUBO (or Ising) model the ordering of qubits within the pair does not matter (i.e., the interaction between i and j is the same as j and i), the summation is restricted to $i < j$, and the nondiagonal elements Q_{ij} are multiplied by two.

Unfortunately, the minimum of the functional E is a trivial solution $\Psi = 0$, which is due to the lack of the normalization constraint $\|\Psi\| = 1$. The workaround is to add that constraint right into the functional with a strength λ , giving $I = \langle \Psi | \hat{H} | \Psi \rangle + \lambda(1 - \langle \Psi | \Psi \rangle)^2$. Essentially, the parameter λ penalizes any deviation of the norm from unity, and it helps to guide the optimization away from the trivial solution. One can think of λ as a Lagrange multiplier and the functionals E and I as objective functions. The problem with the functional I is that it is no more a QUBO functional, rather it is biquadratic in x . The trick is to lower the power of the constraint, giving $G = \langle \Psi | \hat{H} | \Psi \rangle + \lambda(1 - \langle \Psi | \Psi \rangle)$. Dropping the constant shift λ , which has no effect on the optimization, one obtains the final expression for the functional form used in present study: $F = \langle \Psi | \hat{H} | \Psi \rangle - \lambda \langle \Psi | \Psi \rangle$. The main consequence of the decreased power is that the normalization condition is broken *per se* (but this can be fixed by rescaling the final solution $|\Psi\rangle \rightarrow |\Psi\rangle / \sqrt{\|\Psi\|}$). However, the primary role of the penalty is to avoid the trivial solution, and the functional F serves that purpose. Another issue with the functional F is that it encourages a nonphysical norm $\|\Psi\| > 1$. This limits the number of techniques to find a good parameter λ . Ideally, λ should be large enough to kick the optimization away from the trivial solution minimum but yet small enough to stay away from the large norm limit. To find an optimal value for λ , we scan in λ and pick the solution with the lowest energy $E = \langle \Psi | \hat{H} | \Psi \rangle$ (where here Ψ has been rescaled: $\langle \Psi | \Psi \rangle = 1$). The overhead associated with scanning in λ is negligible since querying the quantum annealer takes essentially all of the computational time. The scanning in λ can be considered as part of the overall sampling procedure and is done simultaneously while accumulating statistics. Also, the dimensionality of the problem has no effect on the scanning.

We would like to stress that the operator \hat{H} is a problem-dependent physical operator and can be of any kind and structure. The matrix elements $H_{\alpha\beta}$ are evaluated on a classical computer using a convenient problem-dependent basis set. For example, to calculate the vibrational spectrum of a diatomic molecule with reduced mass μ , internuclear distance r , and potential energy surface $V(r)$, we have $\hat{H} = (-\hbar^2/2\mu)\nabla^2 +$

$V(r)$ and use a finite Fourier basis set $\varphi_\alpha(r)$ to derive matrix elements $H_{\alpha\beta}$. The dimensionality of the Hilbert space spanned by the Fourier basis and the number of eigenpairs is simply the number of basis functions B . Once we add the normalization term and convert the expansion coefficients to qubits, the problem is described by the QUBO functional F_Q , which can be associated with an operator in a Hilbert space spanned by qubit configurations $|q_1^1 \dots q_k^{\alpha} \dots q_k^{\beta}\rangle$, i.e., a tensor product of individual qubits. The dimensionality of the latter space is 2^{KB} . Thus, we are dealing with two Hamiltonians: the initial problem Hamiltonian H and the final QUBO Hamiltonian F_Q . The explicit expression of matrix elements of the latter in terms of matrix elements of the former is given in the [Supporting Information](#).

Since the functional F contains all possible products $a_\alpha a_\beta$ of the expansion coefficients, the QUBO functional F_Q contains all possible products $q_n^\alpha q_m^\beta$ of qubits, where n and m label qubits for each individual expansion coefficient. This means that all qubits are interacting with each other and require an all-to-all pairwise connectivity.

To extract the physical wave function (eigenvector), one reconstructs the expansion coefficients a_α from the qubits q_k^α using a fixed-point representation: the first K qubits are used to construct a_1 , the next K qubits are used to construct a_2 , and so on. Then, the constructed coefficients are rescaled to satisfy the normalization condition, $\sum_\alpha a_\alpha^2 = 1$, and the energy (eigenvalue) is obtained by evaluating the expectation value of the problem Hamiltonian \hat{H} on the computed state $\langle \psi | \hat{H} | \psi \rangle = E$. One has to repeat this procedure for multiple values of λ and pick the one which gives the lowest energy.

2.2. Calculation of Excited States. The QAE algorithm described above can be easily applied to the calculation of excited states by modifying the initial problem Hamiltonian \mathbf{H} . Specifically, for the first excited state, an outer product matrix of the previously computed ground state wave function is added: $\mathbf{H}' = \mathbf{H} + S_0 |\Psi_0\rangle\langle\Psi_0|$. The parameter S_0 is an arbitrary user specified energy shift to move the ground state higher in the spectrum. The only requirement on S_0 is that it should be larger than the energy of the first excited state, otherwise the algorithm will keep converging to the ground state. To compute the i -th excited state, similar terms for the states $0, 1, \dots, (i-1)$ should be added to the Hamiltonian. In principle, this iterative procedure allows one to compute the whole spectrum of a molecule. To compute the first M excited states, $M \times N_\lambda$ optimization problems have to be solved with properly modified problem Hamiltonians \hat{H}' . Obviously, for a fixed basis size B and qubit expansion K , the higher states will not be described as accurately as the lower ones.

2.3. Multiple Dimensions. The generalization of the QAE method to multiple dimensions is straightforward. For a direct product basis, the one-dimensional expansion is replaced with an n -dimensional expansion, but the way to code each expansion coefficient using K qubits remains the same. For example, for a two-dimensional system with the same number of basis functions B for each dimension, the expansion is $\Psi = \sum_{\alpha,\beta}^{B,B} a_{\alpha\beta} \varphi_\alpha \theta_\beta$, where φ_α and θ_β are the basis functions in each dimension. The qubits $q_k^{\alpha\beta}$ are now mapped to the variables x_i , as follows: $i = KB(\alpha - 1) + K(\beta - 1) + k$, where $i \in [1, K \times B^2]$. Thus, one needs KB^d qubits to describe the d -dimensional problem. This brings the number of explored configurations to 2^{KB^d} .

The QAE method can also be applied to a nondirect product basis. For example, one can implement a Sequential Diagonalization Truncation (SDT)^{31–33} which drastically reduces the size of the Hamiltonian matrix and ultimately results in a much smaller total number of qubits than in a direct product treatment. We use the direct product basis set for a multidimensional harmonic oscillator problem and use the SDT approach for ozone. Further details of applying the SDT method to that molecule can be found elsewhere.³⁴

3. RESULTS

3.1. Application to O₂ and O₃. We applied our algorithm to the calculation of the ground and first excited states of the oxygen and ozone molecules. For both, we used an accurate potential energy surface of ozone.³⁵ The one-dimensional O₂ potential $V(r)$ was generated from the full three-dimensional ozone potential by moving one of the oxygen atoms far away from the other two (i.e., $R_{O-O_2} = 60a_0$). The two lowest wave functions for both molecules are shown in [Figure 1](#). The




Figure 1. Computed wave functions of the ground and first excited states. The molecules are (A) O₂ and (B) O₃, and the wave function is squared for ozone. Red and green bold curves are the results of a classical numerical eigensolver (LAPACK). Solid thin black curves were obtained using QAE with a software classical QUBO solver (they lie on top of the LAPACK curves). Dashed black curves are the results obtained using QAE with a hardware quantum annealer (D-Wave machine). The blue curve is (A) oxygen potential $V(r)$ and (B) the minimum energy path $V(\rho)$ for ozone.

software based classical QUBO solver reproduces the wave functions computed with a standard classical numerical eigensolver (LAPACK³⁶), whereas the output of the hardware quantum annealer (D-Wave machine) is less accurate. The sharp edges (low resolution) of the wave functions are due to the small size of the basis set. For oxygen molecule, a Fourier basis (e^{imr}) was chosen small enough to give the known ground state energy of 791.64 cm^{-1} within an error of 0.01 cm^{-1} using

a standard classical numerical eigensolver (LAPACK). The required number of periods in the Fourier series is $m_{\max} = 4$ which translates to the basis size $B = 2m_{\max} + 1 = 9$. For ozone, we used an SDT basis set truncated at quite low energy $E_{\text{cut}} = 2000 \text{ cm}^{-1}$, which gives $B = 12$ basis functions. This basis is sufficient to describe the ground state at 1451 cm^{-1} with an error of 120 cm^{-1} using LAPACK but is too small for excited state calculations (the computed excited state energy is 700 cm^{-1} larger than the true value of 2147 cm^{-1}). Nevertheless, in this work we are primarily interested in benchmarking the new QAE method against a standard classical numerical eigensolver (LAPACK) and not in the absolute accuracy of the solutions. We refer the reader interested in accurate classical calculations to the relevant literature.^{34,37} For the three-dimensional ozone system, we plot the probability density (the wave function squared) as a function of the symmetric-stretch coordinate ρ in Figure 1b.^{38,39} The SDT basis functions span the other two internal degrees of freedom (not plotted) at each value of ρ and are computed classically.³⁴ The number of qubits K per expansion coefficient (or basis function) is 7 for oxygen and 5 for ozone which is the maximum possible number which fits within the 64 logical fully connected qubits on the hardware quantum annealer (D-Wave machine). Namely, $K \cdot B = 7 \cdot 9 = 63$ for oxygen and $K \cdot B = 5 \cdot 12 = 60$ for ozone.

Figure 2 illustrates the convergence of the energies as a function of the number of qubits K per expansion coefficient a_α (i.e., the level of discretization). There are several interesting findings to discuss. First, the error decreases exponentially as a function of K , which is very appealing. Second, the error decreases and reaches a plateau, for both solvers and both




Figure 2. Ground and first excited state energy errors as a function of number of qubits K per expansion coefficient a_α . The molecules are (A) O_2 and (B) O_3 . The ground state is red, and the excited state is green. Results were obtained using QAE with a classical QUBO solver (solid curves) and a quantum annealer (dashed curves). Errors are computed relative to the energies of a classical numerical eigensolver (LAPACK).

molecules. This shared behavior demonstrates that the QAE algorithm itself is universal and practical, but the actual error is solver and system dependent. Third, the quantum annealer (D-Wave machine) is much less accurate (by 2–3 orders of magnitude) than the classical QUBO solver. In addition, because the total number of qubits currently available in the quantum annealer is rather limited, the corresponding (dashed) curves do not continue to higher values of K . Finally, the classical QUBO solver brings the error for O_2 down to 0.01 cm^{-1} which coincidentally matches the error of chosen basis size. No more than $K = 8$ qubits (a qubyte) per coefficient are needed for oxygen. For ozone, $K = 5$ qubits are required by the classical QUBO solver to reach the plateau within an error of less than 3 cm^{-1} . This is quite accurate, compared to the 120 cm^{-1} error due to the chosen (small) SDT basis.

Table S1 in the Supporting Information gives the parameters used in scanning over the normalization penalty λ . The initial λ_{\min} could be zero or some value that is smaller than the expected energy of the state being computed (if it is roughly known a priori). The number of steps N_λ specifies the number of samples and is a convergence parameter. The last parameter is the step size $\Delta\lambda$ which ideally should be the same for all K within a given problem. However, for small K we were always getting trivial solutions which is probably due to the inaccurate description of the problem. The only way we found to avoid this is to increase the step size $\Delta\lambda$ by more than an order of magnitude. We faced the same issue when we were simulating the hardware noise. Increasing both the step size and the number of steps allowed us to overcome this obstacle.

3.2. Algorithm Scaling. An idealized QUBO solver would give the lowest state immediately. The hardware based solver, such as a quantum annealer, is expected to find the lowest state by means of inherent quantum tunneling between states. The software based solver, such as the Tabu search implemented in qbsolv (see Sec. 5), searches for the lowest state by crawling on the landscape and remembering configurations that were visited previously and are now “tabued”. No matter what type of solver is used, QAE performs N_λ iterations to find a good normalization penalty. Therefore, in the case of an ideal QUBO solver, the scaling of the algorithm is simply $O(N_\lambda)$. In practice, however, this theoretical scaling is not obtainable, because the QUBO solver may also have additional (internal) effects on performance and undesirable artifacts can occur. For example, the classical QUBO solver (qbsolv) that we used in this work shows a steplike runtime as a function of K (see Figure S1 in the Supporting Information), which is due to a backbone-based method inspired by Glover et al.⁴⁰ used to partition the problem into smaller pieces.

In the next paragraph, we will examine the performance of the QAE algorithm that uses the classical QUBO solver for a multidimensional harmonic oscillator. We will show the runtime scales as $O(N_\lambda KB^d)$, where d is the number of dimensions, and B is the number of cosine basis functions in each dimension. The oscillator frequencies were set differently according to $\omega_i(\text{cm}^{-1}) = 800 + 200 \cdot (i-1)$, where i is the dimension index.

To verify the linear scaling with K , we plot the normalized computational time (time divided by K) as a function of K in Figure 3. All curves are roughly horizontal, and their slope does not depend on dimensionality (except for perhaps $d = 5$). The rapid increase or steps at $K = 16$ for 1D, at $K = 6$ for 2D, and at $K = 2$ for 3D are due to partitioning size. They appear as soon




Figure 3. Normalized computational time for the d -dimensional harmonic oscillator. The time is plotted as a function of K and was obtained using QAE with a classical QUBO solver. Results are presented for 1D (red), 2D (green), 3D (blue), 4D (black), and 5D (brown). The number of cosine basis functions is $B = 3$ ($m_{\max} = 2$). No slope in all curves demonstrates linear scaling of the algorithm with K .

as the total number of qubits exceeds the sub-QUBO size of 47, which is 3×16 for 1D, $3^2 \times 6$ for 2D, and $3^3 \times 2$ for 3D. Once the linear dependence on K has been established (see Figure 3), one can then verify the exponential dependence on dimensionality d . The logarithm of the normalized computational time is plotted in Figure 4 as a function of the




Figure 4. Logarithm of the normalized computational time as a function of the number of dimensions. The time was measured for the d -dimensional harmonic oscillator and was obtained using QAE with a classical QUBO solver. Results are presented for the number of qubits per coefficient $K = 4$ (red), $K = 8$ (green), $K = 12$ (blue), and $K = 16$ (black). The number of cosine basis functions is $B = 3$ ($m_{\max} = 2$). The linear dependence confirms exponential scaling in d .

dimensionality d . All of the curves exhibit a roughly linear dependence on d which confirms the exponential scaling. Again, the times for all K at $d = 1$ and for $K = 4$ at $d = 2$ deviate significantly from the main trend because no partitioning is required to compute those. The average slope calculated based on $d = 2$ through 4 is 0.58 which is close to the theoretically predicted value of 0.48 (for $B = 3$). When $d = 5$ is included, the slope increases to 0.72 which is most likely due to the very large problem size. The total number of qubits for $d = 5$ is $KB^d = 972$ to 3888 for $K = 4$ to 16 and $B = 3$ which results in a total number of configurations 10^{300} to 10^{1000} . In summary, Figures 3 and 4 confirm the scaling law $O(N_{\lambda}KB^d)$ of the method using a classical QUBO solver.

We did not perform scaling studies on the hardware quantum annealer (D-Wave machine) because it was not practical due to the small number of logical qubits and long runtime. For example, for O_2 we were able to approach small problems with $B = 9$ and $K = 1$ to 7, and the runtime was about $t_{dw} = 2500$ s. This time does not depend on the number of logical qubits K , because all problems are treated by the hardware as maximum-size problems. In contrast, the classical QUBO solver runtime for these problems was about $t_{cl} = 30$ s, which is almost 2 orders of magnitude faster. The long runtime for the hardware quantum annealer is primarily due to the large number of reads (see Figure S3). In addition, the analysis for the d -dimensional harmonic oscillator would require an extensive QUBO partitioning, which means another factor of 10 to 100 increase in t_{dw} .

3.3. Chaining in Quantum Annealer. The only constraint we have discussed so far is a normalization constraint with associated penalty λ . However, there is another constraint and penalty factor worth mentioning when running on a quantum annealer (D-Wave machine). Namely, the chain constraint and the associated chain penalty. The physical qubits in the hardware do not have an all-to-all connectivity which is a requirement of the algorithm. In fact, each qubit has six neighbors at most (see Chimera graph).³⁰ Fortunately, there is a method to embed a fully connected graph on top of the hardware graph. In this approach, a number of qubits are organized into so-called chains. Qubits within a chain act like a single logical qubit which is connected to all other logical qubits. To program chains in the hardware Chimera graph, one adds a set of constraints which have a single strength or chain penalty c . As with any constraint, the associated penalty or weight c should be neither too small, because then the chains are broken, nor too large, because then the hardware will become insensitive to the original problem. The simplest approach to find a good chain penalty is to perform scanning, in a similar way to λ scanning. Figure 5 demonstrates an example of this two-dimensional scanning for the ground state of O_2 with a Fourier basis of size $B = 7$ ($m_{\max} = 3$) and $K = 3$.




Figure 5. Scanning in the normalization λ and chain c penalties. The calculation was performed on the hardware quantum annealer (D-Wave machine) to find the minimum energy (see color scale in cm^{-1}). The three major regions of interest are trivial solutions (top-left black), broken chains (bottom red-blue), and the optimal region where both the normalization and chain constraints are satisfied and the minimum energy solution is obtained (top-right blue).

As expected, in the region of small c the minimum energy is unacceptably large, simply due to broken chains. For large c we see two regions: the region of small λ , which contains trivial solutions only (because the normalization constraint is too weak) and the region of large λ with reasonable minimum energies. The “phase transition” between these two regions occurs close to the true ground state energy 791.64 cm^{-1} . The result of this two-dimensional scanning shows that a reasonable chain penalty lies between 10000 and 20000 cm^{-1} . We used $c = 15000 \text{ cm}^{-1}$ in all of our calculations.

3.4. Simulating Hardware Noise. The results obtained on the quantum annealer are much less accurate than those obtained using the classical QUBO solver (see Figures 1 and 2). We believe that this discrepancy is partially due to the error with which the QUBO problem is programmed in the hardware. The reported integrated control errors (ICE) for the hardware we used (D-Wave) are quite large.⁸ For example, the diagonal elements are programmed with a mean error 0.7% of the maximum matrix element. In the O_2 calculation, the maximum matrix element is $11.5 \times 10^3 \text{ cm}^{-1}$ which translates into an ICE of 80 cm^{-1} . Moreover, the error has a quite broad distribution, and its standard deviation is 0.8%. This can potentially double the ICE, bringing it to 160 cm^{-1} . Upon consideration of this source of error, the large difference between the quantum annealer and classical results in Figure 2 is no longer that surprising. To model the effects of the hardware errors (noise), we performed calculations with a classical software based solver where random noise was manually added to QUBO (see Materials and Methods). The errors (relative to LAPACK) of the classical QUBO solutions with different magnitudes of noise are shown in Figure 6. On average, introduction of noise increases the error. For the 1D harmonic oscillator, the default noise (using the reported ICE values discussed above) mimics the quantum annealer behavior. For oxygen, the quantum annealer behavior is well characterized if the noise is increased by a factor of 3. For ozone, a factor of 5 or 7 is enough. As the problem size increases, larger noise scaling is required (to reproduce the hardware performance), which implies that there could be some other source of discrepancy between the hardware and software solvers. Also, introduction of noise required an increase in the strength of the normalization weight λ which is reflected in Table S1.

4. DISCUSSION

There are several places in the method where some improvements could be done and which are definitely worth listing. First, the form of the objective functional we used in this study is not the only one. For example, the initial biquadratic functional I , that we simplified earlier, can be converted to a quadratic QUBO form using multiple additional constraints. The drawback of this approach is that it will require additional penalty factors and ultimately will make the problem much harder to manage and solve.

Second, the method would significantly benefit if the normalization condition could be integrated into the functional, rather than added to it. For example, one can notice that solution \mathbf{a} represents a point on the unit hypersphere. Thus, its position can be described using spherical coordinates, and the angles could be approximated with qubits \mathbf{q} . However, the presence of products of sines and cosines in this kind of mapping results in a polynomial of high degree in q . This would require multiple constraints and associated penalty

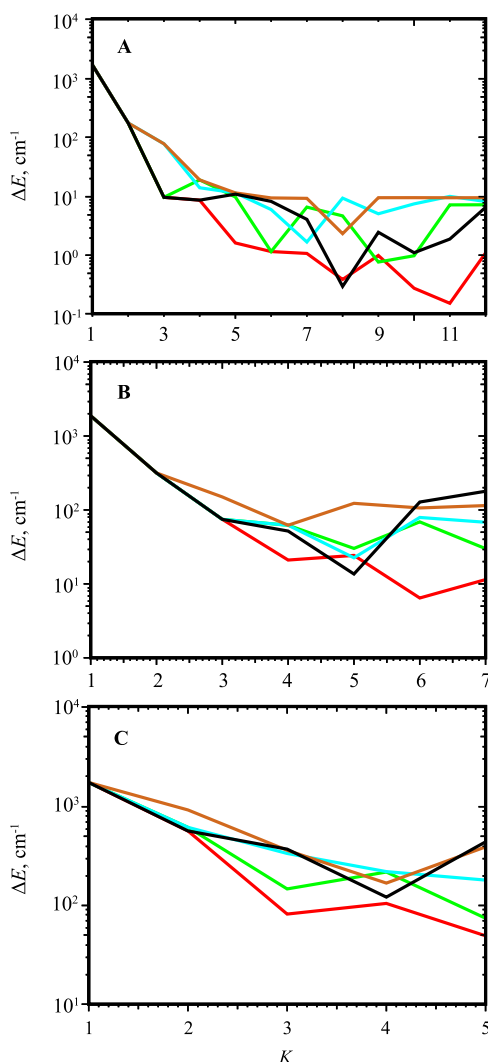


Figure 6. Simulation of quantum annealer noise. (A) 1D harmonic oscillator, (B) O_2 , and (C) O_3 . In each panel, the classical QUBO solution error (relative to LAPACK) is plotted as a function of the number of qubits K for different noise models. In general, the error for the quantum annealer (black curves) is larger than for the classical solver without noise (red curves). The quantum annealer behavior can be qualitatively reproduced by adding noise to the classical calculation. For the 1D harmonic oscillator and O_2 , the noise was added with scaling factors of 1 (green), 3 (blue), and 5 (brown). For O_3 , the noise scaling factors are 3 (green), 5 (blue), and 7 (brown).

factors to convert the problem into quadratic form. The problem becomes difficult again.

Third, the actual scaling of the algorithm depends on the solver and its parameters. For the classical QUBO solver used in this work, the stopping criterion is specified by the number of repeats without any improvement. For the quantum annealer, the user specifies the number of annealing cycles or reads (see Materials and Methods). The scaling law $O(N_\lambda)$ is for the algorithm itself, while the runtime using a classical QUBO solver scales as $O(N_\lambda K B^d)$, and no performance study was possible for the hardware QUBO solver.

Fourth, the physical or working Chimera graph of the quantum annealer is not perfect. The yield of the working graph, the percentage of working qubits (and couplers) that are present, is 99% (98%).⁸ To make it a full-yield graph, an additional software postprocessing is performed before results

are sent back to the user. The user may opt out of this fixing procedure, however, that reduces portability of the algorithm, since every machine has its own working graph. Still removing this layer might be worth exploring.

Fifth, the annealing time t_{ann} is another parameter that can potentially play some role in the annealing process. We found that increasing t_{ann} does not change results significantly. What affects the results more was the number of reads N_{reads} . Because there is a limitation $t_{\text{ann}} \cdot N_{\text{reads}} < 3$ s in the API, we chose the maximum number of reads per job submission: $N_{\text{reads}} = 10^4$ and $t_{\text{ann}} = 299$ μs . The maximum allowed t_{ann} is 2000 μs , but it would be good to explore larger annealing times if possible.

Sixth, it is not quite clear how to properly include the ICES in the noise simulation tests. The errors are reported solely for the maximum and minimum elements of QUBO matrix, and they are a function of annealing time. In addition, they were reported for 70% of the annealing process (i.e., there is no error data at the end of annealing).⁸ For the noise tests reported in this work, we used the reported errors at 70% and manually scaled them.

Seventh, the annealer temperature could be another source of error. In a recent paper⁴¹ it was argued that the annealer temperatures must be appropriately scaled down with problem size (at least in a logarithmic way or better yet as a power law). In fact, during our study we had to switch from the device with 1024 qubits (DW2X) to 2048 qubits (DW2000Q). The temperature lowered from 15.7 mK to 14.5 mK, which is almost logarithmic (it should be 14.3 mK). However, we had to perform 50 times more reads on the larger machine to obtain reliable solutions for small λ and large c .

Finally, the landscape of the hypothetical configuration space is defined by the problem. It could be that our problems have high and thick barriers on the landscape, which effectively disables the tunneling mechanism in the annealer and restricts exploration. Possibly, a very long annealing time could help to determine if this is the case. Another possibility is that the landscape has a large number of wells and the solver jumps from one to another.

This also might be seen as the presence of a small energy gap between low-lying states in the QUBO Hamiltonian. We believe that the gap size decreases exponentially with increasing accuracy of the expansion coefficients K and consequently with the number of qubits. This should not be confused however with the energy spacing between the states of problem Hamiltonian \hat{H} , which is constant and is solely determined by the physical system (for example $\hbar\omega$ in a harmonic oscillator).

In summary, we developed a hybrid algorithm (QAE) for the calculation of the vibrational spectrum of a molecule on a quantum annealer. The eigenvalue problem is mapped to the QUBO (Ising) problem by discretization of the expansion coefficients using qubits. The method is hybrid due to the scanning in a penalty (or weight) to impose wave function normalization. Running on the actual quantum annealer requires the additional scanning in chain penalty. The method was applied to the ground and first excited vibrational states of two chemically important species: O₂ (oxygen) and O₃ (ozone). The QAE calculations based on the classical QUBO solver outperform those on the quantum annealer (D-Wave machine) in both accuracy and computational time (i.e., no supremacy of the latter one). Our tests show that this is partially due to the errors or noise present in the hardware. Hopefully, in the future it will be possible to build larger, more

accurate, and fully connected quantum annealers. As a final note, the QAE algorithm is universal and can be used in any field of science or engineering to solve the real symmetric eigenvalue problem. To apply QAE to a complex Hermitian matrix, one may consider converting it to a real symmetric matrix first.⁴²

5. MATERIALS AND METHODS

We used qbsolv⁴³ as the software classical QUBO solver and the D-Wave 2000Q⁸ as the hardware quantum solver. The underlying qbsolv algorithm is a combination of Tabu search and a backbone-based method inspired by Glover et al.⁴⁰ The latter one is used for partitioning the original (large) QUBO into smaller sub-QUBOs. The only modification we did was to increase the span of partitioning to 1 (the hard-coded value is 0.214). The number of repeats in the stopping criterion is $N_{\text{rep}} = 10^4$. For the excited states calculations, we used $S_0 = 9000$ cm^{-1} in the expression for H' ; however, 3000 and 6000 cm^{-1} also worked well.

The hardware was accessed using qOp stack: qbsolv, DW library, and SAPI.⁴⁴ Although qbsolv allows running sub-QUBOs on the hardware, we did not follow this approach, because the contribution of the D-Wave machine to the solution would be hard to estimate. Furthermore, qbsolv does implicit restarts and uses a classical Tabu search to refine solutions. Thus, a hardware calculation using the default qbsolv is actually hybrid and not fully quantum (for problems that fit one sub-QUBO qbsolv is completely classical). In order to bring the actual D-Wave performance to the surface, we removed partitioning, restarts, and refinement from qbsolv, so that it only serves as an interface to the hardware. In addition, we raised the number of reads to 10^5 (the hard-coded value is 25). Figure 5 was prepared with 5×10^5 reads (half a million), and the chain penalty was 15000 cm^{-1} (the hard-coded value is 15). The number of physical qubits on the DW2000Q is 2028 qubits (99% yield) and 5903 couplers (98% yield). However, embedding a fully connected graph leaves us with just 64 logical qubits.

Our code generates input QUBO matrices for qbsolv. It is written in Fortran and uses LAPACK³⁶ as the classical numerical eigensolver (for benchmarking the QUBO results). Convergence studies were done for N_{rep} , N_{reads} , and N_{λ} and are reported in Figures S2–S4 of the Supporting Information. We did not see a strong dependence on Tabu memory, so we used the default values. In addition, we tested the code on $d = 1$ to 5 dimensional harmonic oscillators, and the convergence for $d = 1$ to 3 in terms of number of qubits K is given in Figure S5.

■ ASSOCIATED CONTENT

Supporting Information

The Supporting Information is available free of charge on the ACS Publications website at DOI: 10.1021/acs.jctc.9b00402.

Additional convergence studies, explicit form of QUBO Hamiltonian, and table of parameters used in calculations (PDF)

■ AUTHOR INFORMATION

Corresponding Author

*E-mail: bkendric@lanl.gov.

ORCID

Alexander Teplukhin: 0000-0002-6094-1430

Dmitri Babikov: 0000-0002-4667-7645

Funding

A.T. and B.K.K. acknowledge that this work was done under the auspices of the U.S. Department of Energy under Project No. 20170221ER of the Laboratory Directed Research and Development Program at Los Alamos National Laboratory. Los Alamos National Laboratory is operated by Triad National Security, LLC, for the National Nuclear Security Administration of the U.S. Department of Energy (Contract No. 89233218CNA000001). D.B. acknowledges that this material is based upon work supported by the National Science Foundation under Grant No. AGS-1920523.

Notes

The authors declare no competing financial interest.

ACKNOWLEDGMENTS

We thank D-Wave Systems Inc. for providing access to the DW2X device at LANL and the DW2000Q device in Burnaby, Canada. A.T. thanks LANL for sponsoring his summer visit in 2017.

REFERENCES

- (1) Ladd, T. D.; Jelezko, F.; Laflamme, R.; Nakamura, Y.; Monroe, C.; O'Brien, J. L. Quantum computers. *Nature* **2010**, *464*, 45.
- (2) Johnson, M. W.; Amin, M. H. S.; Gildert, S.; Lanting, T.; Hamze, F.; Dickson, N.; Harris, R.; Berkley, A. J.; Johansson, J.; Bunyk, P.; Chapple, E. M.; Enderud, C.; Hilton, J. P.; Karimi, K.; Ladizinsky, E.; Ladizinsky, N.; Oh, T.; Perminov, I.; Rich, C.; Thom, M. C.; Tolkacheva, E.; Truncik, C. J. S.; Uchaikin, S.; Wang, J.; Wilson, B.; Rose, G. Quantum annealing with manufactured spins. *Nature* **2011**, *473*, 194.
- (3) Geller, M. R.; Martinis, J. M.; Sornborger, A. T.; Stancil, P. C.; Pritchett, E. J.; You, H.; Galiatdinov, A. Universal quantum simulation with prethreshold superconducting qubits: Single-excitation subspace method. *Phys. Rev. A: At, Mol, Opt. Phys.* **2015**, *91*, 062309.
- (4) Aspuru-Guzik, A.; Dutoi, A. D.; Love, P. J.; Head-Gordon, M. Simulated quantum computation of molecular energies. *Science* **2005**, *309*, 1704–1707.
- (5) Buluta, I.; Nori, F. Quantum simulators. *Science* **2009**, *326*, 108–111.
- (6) Albash, T.; Lidar, D. A. Adiabatic quantum computation. *Rev. Mod. Phys.* **2018**, *90*, 015002.
- (7) Mizel, A.; Lidar, D. A.; Mitchell, M. Simple proof of equivalence between adiabatic quantum computation and the circuit model. *Phys. Rev. Lett.* **2007**, *99*, 070502.
- (8) QPU Properties: D-Wave 2000Q Online System (DW_2000Q_3); 09-1180A-A; D-Wave Systems Inc.: 2018.
- (9) Lin, Z.; McCoy, A. B. Signatures of Large-Amplitude Vibrations in the Spectra of H5+ and D5+. *J. Phys. Chem. Lett.* **2012**, *3*, 3690–3696.
- (10) Lin, Z.; McCoy, A. B. Investigation of the structure and spectroscopy of H5+ using diffusion monte carlo. *J. Phys. Chem. A* **2013**, *117*, 11725–11736.
- (11) Qu, C.; Prosmitti, R.; Bowman, J. M. MULTIMODE calculations of the infrared spectra of H7+ and D7+ using ab initio potential energy and dipole moment surfaces. *Theor. Chem. Acc.* **2013**, *132*, 1413.
- (12) Qu, C.; Bowman, J. M. Diffusion Monte Carlo Calculations of Zero-Point Structures of Partially Deuterated Isotopologues of H7+. *J. Phys. Chem. B* **2014**, *118*, 8221–8226.
- (13) McCoy, A. B.; Braams, B. J.; Brown, A.; Huang, X.; Jin, Z.; Bowman, J. M. Ab initio diffusion Monte Carlo calculations of the quantum behavior of CH5+ in full dimensionality. *J. Phys. Chem. A* **2004**, *108*, 4991–4994.
- (14) Huang, X.; Johnson, L. M.; Bowman, J. M.; McCoy, A. B. Deuteration effects on the structure and infrared spectrum of CH5+. *J. Am. Chem. Soc.* **2006**, *128*, 3478–3479.
- (15) Petit, A. S.; McCoy, A. B. Diffusion Monte Carlo approaches for evaluating rotationally excited states of symmetric top molecules: application to H3O+ and D3O+. *J. Phys. Chem. A* **2009**, *113*, 12706–12714.
- (16) Guasco, T. L.; Johnson, M. A.; McCoy, A. B. Unraveling anharmonic effects in the vibrational predissociation spectra of HSO2+ and its deuterated analogues. *J. Phys. Chem. A* **2011**, *115*, 5847–5858.
- (17) Krishna, M. V. R.; Whaley, K. Structure of small molecular hydrogen clusters. *Z. Phys. D: At, Mol. Clusters* **1991**, *20*, 223–226.
- (18) Warnecke, S.; Sevryuk, M.; Ceperley, D.; Toennies, J.; Guardiola, R.; Navarro, J. The structure of para-hydrogen clusters. *Eur. Phys. J. D* **2010**, *56*, 353–358.
- (19) Navarro, J.; Guardiola, R. Thermal effects on small para-hydrogen clusters. *Int. J. Quantum Chem.* **2011**, *111*, 463–471.
- (20) Scharf, D.; Martyna, G. J.; Klein, M. L. Isotope effect on the melting of para-hydrogen and ortho-deuterium clusters. *Chem. Phys. Lett.* **1992**, *197*, 231–235.
- (21) Cuervo, J. E.; Roy, P.-N. On the solid-and liquidlike nature of quantum clusters in their ground state. *J. Chem. Phys.* **2008**, *128*, 224509.
- (22) Curotto, E.; Mella, M. Quantum Monte Carlo simulations of selected ammonia clusters (n= 2–5): Isotope effects on the ground state of typical hydrogen bonded systems. *J. Chem. Phys.* **2010**, *133*, 214301.
- (23) Noya, E. G.; Doye, J. P. Structural transitions in the 309-atom magic number Lennard-Jones cluster. *J. Chem. Phys.* **2006**, *124*, 104503.
- (24) Deckman, J.; Frantsuzov, P. A.; Mandelshtam, V. A. Quantum transitions in Lennard-Jones clusters. *Phys. Rev. E* **2008**, *77*, 052102.
- (25) Xia, R.; Bian, T.; Kais, S. Electronic structure calculations and the ising hamiltonian. *J. Phys. Chem. B* **2018**, *122*, 3384–3395.
- (26) Streif, M.; Neukart, F.; Leib, M. *Solving Quantum Chemistry Problems with a D-Wave Quantum Annealer*; International Workshop on Quantum Technology and Optimization Problems; 2019; pp 111–122, DOI: 10.1007/978-3-030-14082-3_10.
- (27) 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.
- (28) McClean, J. R.; Romero, J.; Babbush, R.; Aspuru-Guzik, A. The theory of variational hybrid quantum-classical algorithms. *New J. Phys.* **2016**, *18*, 023023.
- (29) O'Malley, P. J. J.; Babbush, R.; Kivlichan, I. D.; Romero, J.; McClean, J. R.; Barends, R.; Kelly, J.; Roushan, P.; Tranter, A.; Ding, N.; Campbell, B.; Chen, Y.; Chen, Z.; Chiaro, B.; Dunsworth, A.; Fowler, A. G.; Jeffrey, E.; Lucero, E.; Megrant, A.; Mutus, J. Y.; Neeley, M.; Neill, C.; Quintana, C.; Sank, D.; Vainsencher, A.; Wenner, J.; White, T. C.; Coveney, P. V.; Love, P. J.; Neven, H.; Aspuru-Guzik, A.; Martinis, J. M. Scalable quantum simulation of molecular energies. *Phys. Rev. X* **2016**, *6*, 031007.
- (30) *Programming with QUBOs*, 09-1002A-C, Release 2.4; D-Wave Systems Inc.: 2017.
- (31) Bačić, Z.; Light, J. Highly excited vibrational levels of "floppy" triatomic molecules: A discrete variable representation—Distributed Gaussian basis approach. *J. Chem. Phys.* **1986**, *85*, 4594–4604.
- (32) Light, J.; Bačić, Z. Adiabatic approximation and nonadiabatic corrections in the discrete variable representation: Highly excited vibrational states of triatomic molecules. *J. Chem. Phys.* **1987**, *87*, 4008–4019.
- (33) Bačić, Z.; Light, J. Accurate localized and delocalized vibrational states of HCN/HNC. *J. Chem. Phys.* **1987**, *86*, 3065–3077.
- (34) Teplukhin, A.; Babikov, D. Efficient method for calculations of ro-vibrational states in triatomic molecules near dissociation threshold: Application to ozone. *J. Chem. Phys.* **2016**, *145*, 114106.
- (35) Dawes, R.; Lolur, P.; Li, A.; Jiang, B.; Guo, H. Communication: An accurate global potential energy surface for the ground electronic state of ozone. *J. Chem. Phys.* **2013**, *139*, 201103.

(36) Anderson, E.; Bai, Z.; Bischof, C.; Blackford, S.; Dongarra, J.; Du Croz, J.; Greenbaum, A.; Hammarling, S.; McKenney, A.; Sorensen, D. *LAPACK Users' guide*; Siam: 1999; Vol. 9.

(37) Gayday, I.; Teplukhin, A.; Babikov, D. The ratio of the number of states in asymmetric and symmetric ozone molecules deviates from the statistical value of 2. *J. Chem. Phys.* **2019**, *150*, 101104.

(38) Teplukhin, A.; Babikov, D. Interactive tool for visualization of adiabatic adjustment in APH coordinates for computational studies of vibrational motion and chemical reactions. *Chem. Phys. Lett.* **2014**, *614*, 99–103.

(39) Teplukhin, A.; Babikov, D. Visualization of potential energy function using an isoenergy approach and 3D prototyping. *J. Chem. Educ.* **2015**, *92*, 305–309.

(40) Wang, Y.; Lü, Z.; Glover, F.; Hao, J.-K. A multilevel algorithm for large unconstrained binary quadratic optimization. *International Conference on Integration of Artificial Intelligence (AI) and Operations Research (OR) Techniques in Constraint Programming*; 2012; pp 395–408, DOI: [10.1007/978-3-642-29828-8_26](https://doi.org/10.1007/978-3-642-29828-8_26).

(41) Albash, T.; Martin-Mayor, V.; Hen, I. Temperature scaling law for quantum annealing optimizers. *Phys. Rev. Lett.* **2017**, *119*, 110502.

(42) Press, W.; Flannery, B.; Teukolsky, S.; Vetterling, W. *Numerical Recipes: The Art of Scientific Computing*; Cambridge University Press: Cambridge, 1986; p 364.

(43) *Partitioning Optimization Problems for Hybrid Classical/Quantum Execution*; D-Wave Systems Inc.: 14-1006A-A; 2017. <https://github.com/dwavesystems/qbsolv> (accessed July 25, 2019).

(44) *qOp toolset 2.5.1*; D-Wave Systems Inc.: 2018. <https://www.dwavesys.com> (accessed July 25, 2019).