

## RESEARCH ARTICLE

# Sampling electronic structure quadratic unconstrained binary optimization problems (QUBOs) with Ocean and Mukai solvers

Alexander Teplukhin<sup>1‡</sup>, Brian K. Kendrick<sup>1</sup>, Susan M. Mniszewski<sup>2</sup>, Sergei Tretiak<sup>1</sup>, Pavel A. Dub<sup>3\*</sup>

**1** Theoretical Division (T-1, MS B221), Los Alamos National Laboratory, Los Alamos, New Mexico, United States of America, **2** Computer, Computational and Statistical Sciences Division (CCS, MS B214), Los Alamos National Laboratory, Los Alamos, New Mexico, United States of America, **3** Chemistry Division (C-IIAC, MS K558), Los Alamos National Laboratory, Los Alamos, New Mexico, United States of America

‡ Current address: Institute for Advanced Computational Science and Department of Chemistry, Stony Brook University, Stony Brook, New York, United States of America

\* [pdub@lanl.gov](mailto:pdub@lanl.gov)



## OPEN ACCESS

**Citation:** Teplukhin A, Kendrick BK, Mniszewski SM, Tretiak S, Dub PA (2022) Sampling electronic structure quadratic unconstrained binary optimization problems (QUBOs) with Ocean and Mukai solvers. PLoS ONE 17(2): e0263849. <https://doi.org/10.1371/journal.pone.0263849>

**Editor:** Dennis Salahub, University of Calgary, CANADA

**Received:** October 9, 2021

**Accepted:** January 27, 2022

**Published:** February 11, 2022

**Copyright:** This is an open access article, free of all copyright, and may be freely reproduced, distributed, transmitted, modified, built upon, or otherwise used by anyone for any lawful purpose. The work is made available under the [Creative Commons CC0](https://creativecommons.org/licenses/by/4.0/) public domain dedication.

**Data Availability Statement:** All relevant data are within the paper.

**Funding:** Research presented in this article was supported by the Laboratory Directed Research and Development (LDRD) program of Los Alamos National Laboratory (LANL) under project number 20200056DR. This work was conducted in part at the Center for Integrated Nanotechnologies, a U.S. Department of Energy, Office of Basic Energy Sciences user facility.

## Abstract

The most advanced D-Wave Advantage quantum annealer has 5000+ qubits, however, every qubit is connected to a small number of neighbors. As such, implementation of a fully-connected graph results in an order of magnitude reduction in qubit count. To compensate for the reduced number of qubits, one has to rely on special heuristic software such as qbsolv, the purpose of which is to decompose a large quadratic unconstrained binary optimization (QUBO) problem into smaller pieces that fit onto a quantum annealer. In this work, we compare the performance of the open-source qbsolv which is a part of the D-Wave Ocean tools and a new Mukai QUBO solver from Quantum Computing Inc. (QCI). The comparison is done for solving the electronic structure problem and is implemented in a classical mode (Tabu search techniques). The Quantum Annealer Eigensolver is used to map the electronic structure eigenvalue-eigenvector equation to a QUBO problem, solvable on a D-Wave annealer. We find that the Mukai QUBO solver outperforms the Ocean qbsolv with one to two orders of magnitude more accurate energies for all calculations done in the present work, both the ground and excited state calculations. This work stimulates the further development of software to assist in the utilization of modern quantum annealers.

## Introduction

Adiabatic quantum computation (AQC) is a form of quantum computation, where an initial “easy to prepare” Hamiltonian is intentionally slowly (adiabatically) evolving to the final target Hamiltonian, the ground state of which one would like to obtain. The AQC and more popular gate-based quantum computation are two seemingly different approaches, which however were shown to be formally equivalent [1]. While gate-based quantum computation is implemented on several types of Noisy Intermediate-Scale Quantum (NISQ) devices [2], D-Wave is

**Competing interests:** The authors have declared that no competing interests exist.

currently the sole manufacturer of adiabatic quantum devices. D-Wave annealers are yet to be full-fledged AQC systems [3], however, the current hardware is capable of performing Ising optimisation or, alternatively, quadratic unconstrained binary optimization (QUBO). The QUBO is the problem of minimizing a quadratic polynomial over binary variables. Namely, for a given matrix  $Q$ , the sought out solution is a binary string  $x$  which minimizes

$$E(x) = \sum_{i=1}^N \sum_{j=1}^i x_i Q_{ij} x_j. \quad (1)$$

Similar to their NISQ counterparts [2], modern annealers still possess a number of imperfections [3]. Consequently, the development of the software which bears the burden of handling those, is of the utmost importance. Specifically, the D-Wave 2000Q and Advantage architectures operate with 2048 and 5640 qubits, respectively. This may seem a lot, however, the connectivity is an issue. A single qubit is connected correspondingly to 6 and 15 neighboring qubits on these two devices. Stated differently, not every possible pair of  $x_i$  and  $x_j$  is present in **1** when minimization is performed on a quantum annealer. A commonly-used technique called *chaining* solves this problem, where a chain of qubits from the source qubit  $A$  to the target qubit  $B$  is formed so that  $A$  and  $B$  can “talk”. Effectively, each logical variable maps onto a chain of physical qubits. The downside of this technique is that the size of resulting fully-connected graph (or network) of QUBO variables is very small, being 64 and 124 variables on the D-Wave 2000Q and Advantage, respectively.

A way to overcome this small-size limitation of fully-connected QUBOs, is to decompose a given large QUBO problem into pieces, called subQUBOs [4]. The latter are then optimized separately on a quantum annealer aiming to construct a solution to the large QUBO problem. Multiple iterations of this procedure with a wisely chosen decomposition scheme are expected to ultimately converge to a good approximate (if not exact) solution to the initial QUBO problem. This algorithm underpins what qbsolv [4] software does. Namely, the code starts with a random binary string initial guess for a solution, sorts out all qubits in terms of “importance” (based on bit-flip cost), splits the sorted array into pieces, sends them to a quantum annealer, and collects back the optimal binary sub-strings. The latter allow for the construction of a new total solution approximation and then start a subsequent iteration. Qbsolv also refines the global solution between iterations using Tabu search [5] to improve the solution quality and, by doing so, does many internal restarts to exhaust improvement options. Qbsolv is an open-source code being a part of the D-Wave Ocean tools [6]. Recently, qbsolv has been reimplemented by Quantum Computing Inc. The now-proprietary QUBO solver is focused on superior classical performance and is provided via the cloud as a part of the Mukai environment [7]. The new sampler is parallel, exploits advanced Tabu search techniques and was recently tested on a number of problems [8, 9]. The Mukai platform is quantum-ready and targets solving constrained-optimization problems.

In this work, we compare the two QUBO solver implementations, the open-source qbsolv (referred to as Ocean qbsolv) and the new proprietary version (referred to as Mukai QUBO solver). We focus on the electronic structure problem, where the Quantum Annealer Eigensolver (QAE) is used to map the electronic structure eigenvalue equation to QUBO problems, which are in turn solved using either the qbsolv or Mukai QUBO solver. Previously, the QAE was shown to work for a variety of theoretical chemistry tasks including vibrational [10], scattering [11] and electronic [12, 13] problems. Since a substantial body of results was obtained in the latter two studies [12, 13], we will use the electronic/QAE setup as a representative case study to compare the two versions of qbsolv. The comparison is completely classical, because

the Mukai QUBO solver is purely classical at the present moment (i.e., all subQUBOs are solved on the CPU) [9].

It should be noted that there is no intent to advertise a specific QUBO solver. Rather, the purpose of the work is to demonstrate that the errors in the present case study are strongly dependent on the choice of QUBO solver. We believe that such studies are important and particularly needed for physicists who develop physics-related algorithms and put less focus on tool selection. With this letter, the community will be more knowledgeable about the options available, underpinning differences in software performance and as such will have a guideline when choosing a solver that suits their needs. In the next section, we will briefly cover the electronic structure QUBO framework. The complete procedure description can be found in the Methods section of [12].

## Electronic structure QUBOs

A typical goal of electronic structure calculations is to obtain the ground and sometimes a few excited states of a given molecule at a fixed nuclear configuration (i.e., geometry), solving for their energies and wavefunctions. Mathematically, the latter are eigenvalues and eigenvectors, respectively, of the electronic Hamiltonian operator, which can be represented with a matrix that describes a molecule to a desired degree of discretization (numerical grid or basis set) [14]. Thus, one typically constructs such a matrix and then diagonalizes it or, alternatively, uses matrix-free methods to obtain a desired number of eigenpairs. The representative approaches are full configuration interaction (FCI) and complete active space self-consistent field (CASSCF), to name a few. The FCI ultimately provides an exact solution to the electronic problem (subject to basis set limitation), whereas CASSCF is exact only within the chosen active orbital subspace being a practical approximation to the FCI method. In order to compute eigenpairs on a D-Wave quantum annealer, we construct the Hamiltonian matrices explicitly and use the QAE to map the eigenvalue problem to the QUBO Eq (1).

To construct FCI and CASSCF matrices in a given basis set [14], we use an in-house modified Psi4 code [15] with the Cartesian coordinates of atoms of selected molecules optimized at the Hartree-Fock level as input data. The symmetry of molecules is taken into account as described in the Methods section of [12]. Once the matrix is constructed and printed to a file, the QAE is called to process the matrix file. The QAE is based on the min-max variational theorem, where the Rayleigh-Ritz quotient (RRQ)

$$R_A = (v, Av) / (v, v) \quad (2)$$

is minimized over vectors  $v$  for a given matrix  $A$ . The minimum of the RRQ is the smallest eigenvalue, whereas the optimal vector  $v_{opt}$  is the corresponding eigenvector. The RRQ expression does look similar to the QUBO form in Eq (1), especially if  $v$  happens to be normalized. However, a few additional steps should be in place to make them fully compatible. First, the elements of  $v$  need to be encoded in terms of binary variables  $x_i$ . The QAE implements a power-of-two encoding for that. Second, ratios, such as Eq (2), are not allowed in Eq (1), so the RRQ ratio is replaced with a weighted sum of the numerator and the denominator giving the final objective function

$$F(v) = (v, Av) + \lambda \cdot (v, v). \quad (3)$$

The unknown Lagrangian multiplier  $\lambda$  is found iteratively. This is the reason why multiple QUBO problems have to be generated and solved to obtain an eigenpair, not just a single QUBO problem. The iterative procedure for  $\lambda$  avoids or discourages the trivial solution  $v = \mathbf{0}$  (null vector). Moreover, Eq (3) is another way of presenting the eigenvalue problem  $Av = Ev$

(i.e., here taking a dot product with  $v$  and moving  $E(v, v)$  to the left of the equal sign results in the form (3)). As such, the optimal  $\lambda_{opt}$  ends up being near (or equal) to the negative of the true eigenvalue  $E$ . Once  $\lambda_{opt}$  and  $v_{opt}$  are found, the eigenvalue is evaluated as  $(v_{opt}^T A v_{opt})$ . Other eigenpairs can be found in a similar fashion, with the input matrix appropriately modified to shift previously computed eigenvalues to lie higher in the eigenspectrum (Brauer's theorem [16]) and applying the described procedure again without any changes, see [12] for more details.

In the present case study, all of the generated QUBOs are one to two orders of magnitude larger than the fully-connected QUBOs supported by modern D-Wave quantum annealers. This is where the qbsolv and Mukai QUBO solvers come to the rescue through their ability to handle large QUBO problems.

## Results

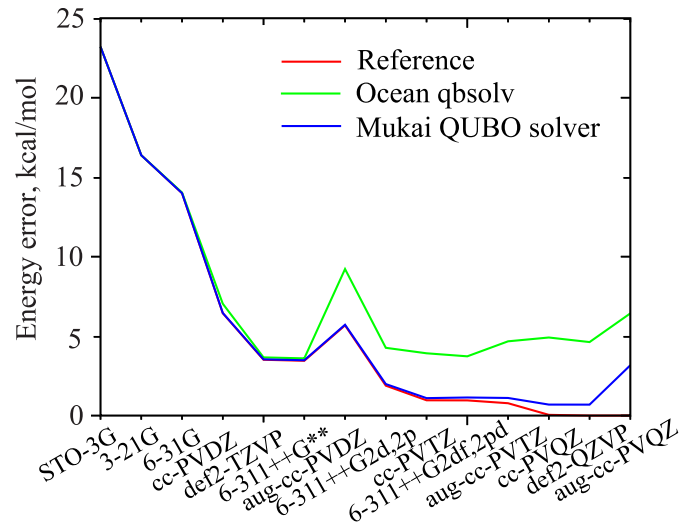
In what follows, we compare the accuracy of the two QUBO solvers (classical mode with all subQUBOs solved on the CPU). Similar to the calculations done in the previous work [12], we start with the  $H_2$  ground state energy evaluation and study its convergence with basis set. Specifically, we use 14 different basis sets with increasing size. Consequently this leads to 14 FCI matrices sized from 2x2 (STO-3G) to 1256x1256 (aug-cc-PVQZ). When applying the same level of discretization of eigenvector elements  $K = 10$ , the resulting QUBOs have 20 to 12560 variables, respectively.

The discretization of  $K = 10$  was shown to be sufficient for Ocean qbsolv in the previous work [12]. While the optimal  $K$  for the Mukai QUBO solver will probably be different from Ocean, we decide to use the same  $K$  for both solvers in the present study, so that that the comparison is more straightforward. Potentially, a more accurate solver will allow a larger optimal  $K$  and obtain higher accuracy in eigenvalues and eigenvectors.

Fig 1 illustrates the calculated energy convergence with basis set. The blue and green curves show the Mukai QUBO solver and Ocean qbsolv results, respectively. The reference calculation is an exact matrix diagonalization using SciPy/LAPACK [17] (the red and green curves are the same as in Fig 1 of [12]). We observe that the QAE + Mukai QUBO solver combination substantially outperforms the QAE + Ocean qbsolv in terms of accuracy. The classical noise floor for the Mukai QUBO solver (i.e., the part of the curve where the error stopped decreasing) appears much later and is significantly lower than that for the Ocean qbsolv. Almost all errors in the blue curve are within chemical accuracy of 1 kcal/mol [18]. The only exception is the last point, where the largest matrix is involved. Still, even in this case with 12560 QUBO variables the Mukai QUBO solver error 3 kcal/mol is twice smaller compared to Ocean's 6 kcal/mol error.

Table 1 lists the ground state energy errors for the different molecular species considered in the Table 1 of our previous work [12]. Again, the Mukai QUBO solver shows much smaller errors compared to the Ocean qbsolv values, being 10 to 100 times smaller on average. Importantly, all Mukai errors are less than 1 kcal/mol, which means that the Mukai QUBO solver gives results to within chemical accuracy. To compare the quality of the computed ground state wave functions, we add two more columns to the table with the norms of eigenvector residuals calculated as  $R_{solver} = \sum_i |v_i^{solver} - v_i^{ref}|$ , where *solver* is either the Ocean qbsolv or Mukai QUBO solver and the summation runs over the eigenvector elements. As with the energies, QAE + Mukai gives more accurate wave functions, i.e. smaller residual norms, than QAE + Ocean. The  $R_{mukai}$  values tend to be 2 to 10 times smaller than the  $R_{ocean}$  values.

The results presented in Fig 1 and Table 1 show that the errors tend to increase with the matrix size and the size of the matrix "sets" the difference. However, a more careful analysis of



**Fig 1. Convergence of the  $H_2$  ground state energy with basis set expansion.** The calculations are ordered with increasing Hamiltonian matrix size. The matrices were diagonalized directly (red) and using the QAE + Ocean qbsolv (green) or QAE + Mukai QUBO solver (blue) approaches. The error is given relative to the energy of def2-QZVP reference diagonalization, which has the lowest value among all calculations. The Mukai QUBO solver outperforms the Ocean qbsolv throughout all basis sets by a significant margin.

<https://doi.org/10.1371/journal.pone.0263849.g001>

the errors suggests that the matrix size is probably not the only error-determining factor. For example, the largest matrices in Fig 1 and Table 1 are very close in size, i.e., 1256x1256 and 1250x1250, but the Mukai errors are quite different, 3 kcal/mol and 0.5 kcal/mol, respectively. We did not investigate such specific cases, simply because the goal of the present work is to compare QUBO solvers. However, some matrix representations are more favorable than others, similar to how conventional iterative eigenvalue algorithms can usually benefit from preconditioners or more appropriate basis sets used in the formation of the Hamiltonian matrix. In the case of the QAE and QUBO solvers, reducing the dynamic range of the problem can help. For example, previously we have manually limited the largest elements of the complex vibrational Hamiltonian matrix [11] to improve the results. In another QAE-based study [13],

**Table 1. Electronic ground state energy errors in kcal/mol.**

Molecule	Method	Basis	Mat. size	QUBO size	$\Delta E^a_{ocean}$	$\Delta E^b_{mukai}$	$R^c_{ocean}$	$R^d_{mukai}$
$H_2$	FCI	STO-3G	2x2	20	0.000	0.000	0.000	0.000
HF	FCI	STO-3G	18x18	180	0.152	0.017	0.017	0.007
$H_2O$	FCI	STO-3G	133x133	1330	5.484	0.084	0.268	0.042
$H_2O$	CAS(8e,7o)SCF	cc-PVDZ	321x321	3210	8.912	0.123	0.626	0.089
$CH_2^{2+}$	FCI	STO-3G	169x169	1690	5.265	0.071	0.344	0.036
$BeH_2$	FCI	STO-3G	169x169	1690	2.079	0.043	0.258	0.040
$H_3^+$	FCI	cc-PVTZ	532x532	5320	6.934	0.460	0.452	0.161
$BH_3$	CAS(6e,6o)SCF	6-311++G**	208x208	2080	6.673	0.086	0.550	0.081
$BH_3$	FCI	STO-3G	1250x1250	12500	10.512	0.510	0.685	0.239

<sup>a</sup> Energy difference between the QAE + Ocean qbsolv ( $E_{ocean}$ ) and the reference diagonalization ( $E_{ref}$ ) [12].

<sup>b</sup> Energy difference between the QAE + Mukai QUBO solver ( $E_{mukai}$ ) and the reference diagonalization ( $E_{ref}$ ).

<sup>c</sup> Norm of eigenvector residual for the QAE + Ocean qbsolv.

<sup>d</sup> Norm of eigenvector residual for the QAE + Mukai QUBO solver.

<https://doi.org/10.1371/journal.pone.0263849.t001>

**Table 2. Electronic transition energies of the H<sub>2</sub>O molecule computed using the QAE and FCI/STO-3G matrix in kcal/mol.**

Transition	$T_{ref}^a$	$T_{ocean}^b$	$T_{mukai}^c$	$T_{ocean} - T_{ref}$	$T_{mukai} - T_{ref}$	$R_{ocean}^d$	$R_{mukai}^e$
S <sub>0</sub> → S <sub>1</sub>	303.056	300.563	303.115	-2.493	0.059	0.360	0.079
S <sub>0</sub> → S <sub>2</sub>	369.233	373.585	369.440	4.352	0.207	0.732	0.115
S <sub>0</sub> → S <sub>3</sub>	441.058	437.217	441.059	-3.841	0.001	0.352	0.131
S <sub>0</sub> → S <sub>4</sub>	590.407	606.617	594.137	16.210	3.730	2.295	0.973

<sup>a</sup> Transition energy obtained using the reference diagonalization [12].

<sup>b</sup> Transition energy obtained using the QAE + Ocean qbsolv [12].

<sup>c</sup> Transition energy obtained using the QAE + Mukai QUBO solver.

<sup>d</sup> Norm of eigenvector residual for the QAE + Ocean qbsolv.

<sup>e</sup> Norm of eigenvector residual for the QAE + Mukai QUBO solver.

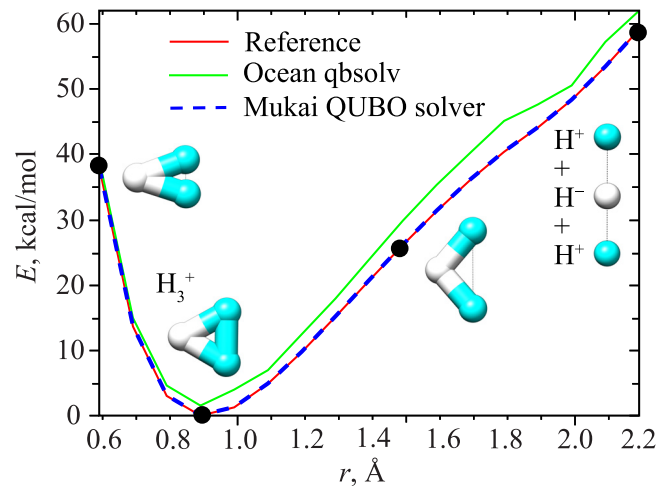
<https://doi.org/10.1371/journal.pone.0263849.t002>

the time-dependent density-functional theory (TDDFT) energies were found to be more accurate than the time-dependent Hartree-Fock (TDHF) energies, however the matrix size in both formulations was the same. In that work, the largest matrix elements of TDDFT were smaller than those of TDHF, which implies a narrower dynamic range and a slightly higher sensitivity to small values in the matrix, leading to more accurate eigenvalues for the TDDFT. These are a few examples where the matrix size is not the only error-determining factor.

We also test both QUBO solvers on excited states, where several eigenpairs are computed sequentially as described in the previous section. Table 2 shows four transition energies for the FCI/STO-3G water molecule (133x133 matrix), which are the spacings between the smallest eigenvalues. Similar to the ground state calculations, we find that the QAE + Mukai QUBO solver combination gives much smaller errors when compared to the QAE + Ocean qbsolv. The error is less than 1 kcal/mol for the first three transitions. The last transition S<sub>0</sub> → S<sub>4</sub> seems to be challenging for both QUBO solvers, however the Mukai QUBO solver error 3.73 kcal/mol is still four times smaller than Ocean's. Interestingly, the Mukai error for the S<sub>0</sub> → S<sub>3</sub> transition happened to be extremely small, 10<sup>-3</sup> kcal/mol. This might be due to a lucky cancellation of errors that come from previously computed eigenvectors. Another interesting feature is that all four Mukai errors are strictly positive, whereas Ocean's errors are both negative and positive. It is hard to infer anything from this observation, which could be fortuitous as well and needs more testing. All QAE energies are not final and can change from run to run. The fluctuations are expected, because both QUBO solvers are heuristic. For example, previously [12] we found that the ground state energy of the 133x133 matrix for the water molecule can fluctuate by 0.5 kcal/mol, on average. The residual norms *R* for the four excited states of water are given in the last two columns of the table. As we can see, *R*<sub>mukai</sub> is much smaller than *R*<sub>ocean</sub> for all states, which means that the Mukai QUBO solver takes the lead in the quality of the computed wave functions as well.

Finally, both QUBO solvers are benchmarked on the calculation of the potential energy curve of H<sub>3</sub><sup>+</sup>, see Fig 2. In this test, 17 distinct geometries were generated at the FCI/cc-PVDZ level, each pre-optimized using the restricted Hartree-Fock method, where the distance between the two terminal hydrogens was varied from 0.59 to 2.19 Å following a step size of 0.1 Å (i.e., constrained optimization). Again, the Mukai QUBO solver outperforms the Ocean qbsolv. The corresponding energy curve closely follows the energy curve obtained with direct diagonalization, whereas Ocean's curve, taken from our previous study [12], consistently overestimates the H<sub>3</sub><sup>+</sup> electronic energy. The QAE + Mukai error is less than 0.1 kcal/mol throughout the whole range of the coordinate. In contrast, the maximum error in the case of Ocean qbsolv is 4.7 kcal/mol.





**Fig 2. Potential energy curve of  $H_3^+$  computed at FCI/cc-PVDZ level.** The same methods as in Fig 1 (same colors) were used. Molecular images (taken at black points) show how the minimum geometry evolves as a function of the distance between two terminal hydrogens. QAE + Mukai QUBO solver curve (blue) overlaps with the reference diagonalization curve and is made dashed, so that both curves are visible.

<https://doi.org/10.1371/journal.pone.0263849.g002>

## Conclusions and outlook

The development of a versatile and robust software base is imperative for the rapid adaptation of emerging quantum computing hardware. The solution of quantum chemical problems is one of the immediate applications of quantum computers. In this purely classical study, we compare two QUBO solvers, which were developed to handle large QUBO problems by decomposing them into smaller pieces that fit on the D-Wave quantum annealers. The comparison is done in the context of electronic structure calculations. An in-house modified Psi4 code was used to construct molecular Hamiltonian matrices and the QAE was then used to convert the resulting eigenvalue problems to QUBO problems. The test cases include the ground state energy convergence with basis set for the hydrogen molecule, the ground state energy calculation for different molecular species up to tetraatomics and an excited state calculation for the water molecule. We find that the Mukai QUBO solver from Quantum Computing Inc. is one to two orders of magnitude more accurate than the qbsolv from the D-Wave Ocean tools. It is important to note that almost all ground and excited state energy errors based on the Mukai QUBO evaluation fall within the chemical accuracy of 1 kcal/mol. The results also suggest that the electronic structure calculations may serve as a go-to validation problem for testing different QUBO solvers.

As a final note, we also did some tests with the D-Wave Hybrid Solver Service [19], which is an alternative to the two QUBO solvers. Unfortunately, the current version of the QAE does not work with the HSS. There is a class of QUBO problems constructed by the QAE for which we expect to get the trivial solution  $v = \mathbf{0}$  (i.e., a null vector), and both qbsolv and Mukai QUBO solver do return this specific QUBO solution. In contrast, the HSS returns a non-trivial solution corresponding to a higher QUBO minimum, which in turn breaks the QAE iteration logic. Thus, a change to the QAE may be needed to make it work with HSS. This will be explored in the future. Overall, we welcome the further development of quantum annealing hardware and emergence of respective software in the field.

## Acknowledgments

The authors would like to thank Quantum Computing Inc. for providing a free trial of Mukai.

## Author Contributions

**Writing – original draft:** Alexander Teplukhin.

**Writing – review & editing:** Brian K. Kendrick, Susan M. Mniszewski, Sergei Tretiak, Pavel A. Dub.

## References

1. Mizel A, Lidar DA, Mitchell M. Simple proof of equivalence between adiabatic quantum computation and the circuit model. *Phys Rev Lett.* 2007; 99(7):070502. <https://doi.org/10.1103/PhysRevLett.99.070502> PMID: [17930879](https://pubmed.ncbi.nlm.nih.gov/17930879/)
2. Preskill J. Quantum Computing in the NISQ era and beyond. *Quantum.* 2018; 2:79. <https://doi.org/10.22331/q-2018-08-06-79>
3. Hauke P, Katzgraber HG, Lechner W, Nishimori H, Oliver WD. Perspectives of quantum annealing: Methods and implementations. *Rep Prog Phys.* 2020; 83(5):054401. <https://doi.org/10.1088/1361-6633/ab85b8> PMID: [32235066](https://pubmed.ncbi.nlm.nih.gov/32235066/)
4. Booth M, Reinhardt SP, Roy A. Partitioning optimization problems for hybrid classical/quantum execution. D-Wave Technical Report Series. 2017; p. 14–1006A–A.
5. Glover F. Future paths for integer programming and links to artificial intelligence. *Comput Oper Res.* 1986; 13(5):533–549. [https://doi.org/10.1016/0305-0548\(86\)90048-1](https://doi.org/10.1016/0305-0548(86)90048-1)
6. D-Wave Ocean Software Documentation; <https://docs.ocean.dwavesys.com/> (Accessed: 2021-10-08).
7. Mukai; <https://www.quantumcomputinginc.com> (Accessed: 2021-10-08).
8. Booth M, Berwald J, Chukwu U, Dawson J, Dridi R, Le D, et al. QCI Qbsolv delivers strong classical performance for quantum-ready formulation; 2020. Preprint at <https://arxiv.org/abs/2005.11294>.
9. Chukwu U, Dridi R, Berwald J, Booth M, Dawson J, Le D, et al. Constrained-optimization approach delivers superior classical performance for graph partitioning via quantum-ready method. In: 2020 IEEE High Performance Extreme Computing Conference (HPEC). IEEE; 2020. p. 1–6.
10. Teplukhin A, Kendrick BK, Babikov D. Calculation of molecular vibrational spectra on a quantum annealer. *J Chem Theory Comput.* 2019; 15(8):4555–4563. <https://doi.org/10.1021/acs.jctc.9b00402> PMID: [31314517](https://pubmed.ncbi.nlm.nih.gov/31314517/)
11. Teplukhin A, Kendrick BK, Babikov D. Solving complex eigenvalue problems on a quantum annealer with applications to quantum scattering resonances. *Phys Chem Chem Phys.* 2020; 22:26136–26144. <https://doi.org/10.1039/D0CP04272B> PMID: [33047749](https://pubmed.ncbi.nlm.nih.gov/33047749/)
12. Teplukhin A, Kendrick BK, Tretiak S, Dub PA. Electronic structure with direct diagonalization on a D-wave quantum annealer. *Sci Rep.* 2020; 10:20753. <https://doi.org/10.1038/s41598-020-77315-4> PMID: [33247201](https://pubmed.ncbi.nlm.nih.gov/33247201/)
13. Teplukhin A, Kendrick BK, Mniszewski SM, Zhang Y, Kumar A, Negre CF, et al. Computing molecular excited states on a D-Wave quantum annealer. *Scientific reports.* 2021; 11(1):1–10. <https://doi.org/10.1038/s41598-021-98331-y> PMID: [34552136](https://pubmed.ncbi.nlm.nih.gov/34552136/)
14. Szabo A, Ostlund NS. *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory.* Dover Publications, New York; 1996.
15. Parrish RM, Burns LA, Smith DGA, Simmonett AC, DePrince AE, Hohenstein EG, et al. Psi4 1.1: An open-source electronic structure program emphasizing automation, advanced libraries, and interoperability. *J Chem Theory Comput.* 2017; 13(7):3185–3197. <https://doi.org/10.1021/acs.jctc.7b00174> PMID: [28489372](https://pubmed.ncbi.nlm.nih.gov/28489372/)
16. Brauer A. Limits for the characteristic roots of a matrix. IV: Applications to stochastic matrices. *Duke Math J.* 1952; 19(1):75–91. <https://doi.org/10.1215/S0012-7094-52-01910-8>
17. Anderson E, Bai Z, Bischof C, Blackford LS, Demmel J, Dongarra JJ, et al. *LAPACK Users' Guide* (Third Ed.). USA: Society for Industrial and Applied Mathematics; 1999.
18. Pople JA. Nobel lecture: Quantum chemical models. *Rev Mod Phys.* 1999; 71(5):1267. <https://doi.org/10.1103/RevModPhys.71.1267> PMID: [34182688](https://pubmed.ncbi.nlm.nih.gov/34182688/)
19. D-Wave. D-Wave Hybrid Solver Service: An overview. D-Wave Whitepaper Series. 2020;14-1039A-A.