Benchmarking Quantum Optimization by Traveling Salesman Problems

RICHARD H. WARREN
Pennsylvania State University, USA
rhw3@psu.edu

Abstract. This paper recommends symmetric traveling salesman problems (TSPs) be used to benchmark quantum capability to find optimal solutions for combinatorial optimization problems. We add four features to the existing list of reasons supporting this recommendation. We discuss benchmark measures and how to overcome the lack of small TSP examples for standards. Significate open questions are identified. We comment about published articles related to the benchmark theme.

Key Words: quantum computing, optimization, benchmark, optimal solution, traveling salesman problem, TSP library

Received: April 23, 2024. Revised: November 14, 2024. Accepted: December 7, 2024. Published: December 31, 2024.

1 Introduction

Rapid solution of combinatorial optimization problems is a major technical application of quantum computing. These problems are usually NP-hard which means that classical optimal solution methods are limited to special cases. There is growing interest in rating and comparing quantum algorithms and quantum hardware to optimally solve combinatorial optimization problems.

For example, [25] contains benchmarks for four software algorithms that solve the symmetric traveling salesman problem on a quantum annealer. The analysis uses a uniform standard to analyze the four quantum programs. Three of the programs are hybrid, classical quantum routines that find approximate solutions. The fourth program is not hybrid and is designed to find an optimal solution.

This paper extends and amplifies [27] where symmetric TSPs are recommended as the best benchmark for combinatorial optimization algorithms on all quantum hardware. The recommendation is based on the following reasons. The symmetric TSP is a renowned, recognized problem. There is a library [19] of symmetric TSP examples that are solved optimally. The TSP is well established in quantum annealing. The computational complexity of the TSP is NP-complete. There is an outstanding classical solution algorithm [2] for the symmetric TSP. There are many real-world applications of symmetric TSPs.

Optimal solutions were found for symmetric TSPs processed on the D-Wave chip [26]. The TSP variables were embedded in the qubits so that every pair of variables was connected and the TSPs were solved by the nonhybrid procedure in [24, 13].

However, the 14 city, symmetric TSP Burma14 in the TSP Library [19] cannot be embedded by the D-Wave embedding software in the D-Wave solver Advantage _system1.1 that has more than 5,000 qubits. Section II.A.1 in [7] contains a full explanation of embedding the variables of a problem in D-Wave qubits so that every pair of variables is connected.

A 48-city TSP was investigated [5] on the D-Wave quantum solver DW_2000Q_6. A large TSP on an IBM quantum processor had 7 cities [10]. The IBM processor has 65 qubits. In both cases the algorithm was a hybrid, classical quantum routine. Performance analysis is needed to compare outcomes with each other and with an optimal solution.

The major restrictions on quantum computing are insufficient connections between qubits on D-Wave processors, a small number of qubits, and noise, particularly on gate machines [1]. Quantum computers are analogue devices which means they are imprecise and have difficulty distinguishing between integers whose difference is small, which can affect the accuracy of computations. The quantum algorithms for TSPs that have N cities require $O(N^2)$ variables which translates into nearly N^2 qubits. The qubits for cities are in addition to the qubits for chains that build connections between qubits.

The layout of this paper: In Section 2 we point out new features that boost the role of symmetric TSPs to benchmark optimization problems. In Section 3 we discuss ways to augment the insufficient number of small, symmetric TSPs in the TSPLIB [19]. Section 4 contains a discussion about benchmark measurements. Future work is described in Section 5. We close with a summary and conclusion in Section 6. An Appendix contains an annotated bibliography of relevant papers that we did not know about when [27] was published.

2 Additional features that enhance symmetric TSPs to be benchmarks

Paper [27] identifies seven features that make symmetric TSPs the best candidates to serve as benchmarks for combinatorial optimization problems on a quantum processor. We add four more features in this section.

2.1 The TSP is Easy to Describe

A TSP is described by a list of cities, including the salesman's home city, and the distance between each pair of cities. The distances are undirected. The TSP asks for a shortest route for the salesman to visit each city once and return to the home city. A TSP with this description is called symmetric. If the distances are directed so that the distance from city A to city B may be different than the distance from city B to city A, then the TSP is called asymmetric. The length of a shortest route is usually part of the solution. A shortest route is called an optimal route.

2.2 The TSP is Easy to Formulate for Quantum Processing

The TSP formulation for quantum computing [24, 13] is simple and straight-forward. Let n be the number of cities. A tour is a route for the salesman through all the cities once and returns to the starting city. Thus, a tour is a cyclic permutation of the n cities. Let V_{it} represent city i in position t of a tour. The double subscript on V means that an n-city TSP has $O(n^2)$ variables. Let d_{ij} be the distance from city i to city j. The objective function to be minimized is (1).

$$\sum_{j=2}^{n} d_{1j} V_{j2} + \sum_{\substack{i,j=2 \\ i \neq j}}^{n-1} \sum_{t=2}^{n-1} d_{ij}\, V_{it}\, V_{j,t+1} + \sum_{i=2}^{n} d_{i1} V_{in} \quad (1)$$

$$\text{For each } t \in \{2,3,\ldots,n\} \sum_{i=2}^{n} V_{it} = 1 \quad (2)$$

$$\text{For each } i \in \{2,3,\ldots,n\} \sum_{t=2}^{n} V_{it} = 1 \quad (3)$$

We want the quantum processor to assign 0, 1 to the variables V_{it} so that the sum of the distances in (1) is minimized, and the variables are subject to constraints (2) and (3). Constraint (2) ensures that each position t in an outcome has exactly one city. Constraint (3) ensures that each city i occurs exactly once in an outcome. To balance the objective function (1) and the constraints (2) and (3), we insert a multiplier in the constraints.

The first and last summations in (1) are a result of the redundancy for expressing tours that allows us to always have city 1 in position 1 of a tour. The redundancy also allows $i = 1$ and $t = 1$

to be omitted from (2) and (3). Thus, the number of variables is $O((n\text{-}1)^2)$.

2.3 The TSP is Easy to Encode for Quantum Computing

The objective function (1) can be encoded in Python after the double subscript on V is converted to a single subscript by means of a dictionary. The constraints (2) and (3) can be encoded after the equations are converted to penalty functions. This is done by adding -1 to both sides of the equation, deleting '= 0' and squaring the expression.

The D-Wave 48-city TSP [5] illustrated this coding.

2.4 Problem Size of a TSP is Easily Noted and Varied

The number of cities in a TSP is widely accepted as the problem size. The number of cities is part of the description of a TSP. Both classical and quantum software for the TSP depend on the number of cities to the extent that a routine for N cities will only process N cities. However, a routine can be adjusted without difficulty to a new value for N.

3 Insufficient number of small, symmetric TSPs in the libraries

The development work on numeric quantum algorithms is in the beginning stages with small problems. There is a shortage of small, standard TSPs to support this work.

nThe TSPLIB [19] contains only 10 symmetric TSPs with less than 42 cities, and one with less than 14 cities. Based on the author's experience, the maximum number of cities in a TSP that can be processed with a nonhybrid, quantum algorithm on D-Wave's 5,000 qubit device is less than 14. On IBM's 65 qubit device, it is less than 8 [10].

Since the TSPLIB has a shortage of examples with a small number of cities, we recommend augmenting tests with examples from the Capacitated Vehicle Routing Problem Library (CVRPLIB) [4]. This was done in [16] by treating the depot and customers in a CVRP as cities for a TSP. The CVRP distances are used for the TSP. The CVRPLIB has many examples with less than 42 cities but only one with less than 14 cities. It has 13 cities.

Reference [3] compliments the TSPLIB with 27 symmetric TSP examples that range in size from 29 cities to 71,009 cities.

The gap of no TSPs with less than 13 cities was filled in [26] with four, illustrative 6-city TSPs that have different numbers of optimal tours. The gap was filled in [10] with a symmetric TSP on N cities for N = 4, 5, …, 11.

We conclude that representative, symmetric TSP examples are needed for 7-cities to 12-cities. We suggest that the representation for each number of cities include a symmetric TSP with exactly two optimal tours, a symmetric TSP with many optimal tours, and a third symmetric TSP where the difference D in length between an optimal tour and the next shortest length of a tour satisfies $0 < D < 5$. The test of the first TSP finds 'a needle in a haystack', the second tests a generic case, and the third tests precision.

4 Benchmark measurements

Benchmark measurements are numeric values about results that can facilitate comparisons of performance, fidelity, quality, effectiveness, and timing. We present several measurements.

Expression (4) is a performance benchmark that measures closeness of a solution to optimal. It depends on a TSP, a script, and a quantum processor. It was used in [28].

$$\text{Error Percent} = [(\text{Best Solution} - \text{Optimal Solution}) / \text{Optimal Solution}] \times 100 \qquad (4)$$

Error Percent is the percent of relative distance from an optimal tour. Best Solution is the shortest tour length found by experimentation. Optimal Solution is the length of an optimal tour. Success Rate is called Observed Probably of Success in [7].

Success Rate (SR) is a metric used in [10, 1] to evaluate solution quality. It measures the percentage of experiments achieving a solution within 95% and 99% of the known optimal value.

These are designated SR95 and SR99. We assume that an optimal value is the length of an optimal tour for the TSP and an experimental value is the length of a computed tour.

Reference [8] describes an infinite family of Euclidean TSPs that are difficult to solve with Concorde [2] which is widely acclaimed to be the best software for finding an optimal tour for a symmetric TSP. The authors of [8] recommend their family of TSPs as benchmarks for TSPs.

The ratio of the shortest length found for a tour divided by the optimal length may be a measure of the degree of difficulty of the combination of a TSP and an algorithm. Ratios between 1 and 2 measure the degree of difficulty, with 1 being easiest and 2 is very hard.

The computational performance metric [10. 1] is the average time for the central processing unit (CPU) across all trials for a fixed TSP, designated script, and specific processor. The computational performance of classical optimization techniques can be compared to hybrid techniques and to nonhybrid, quantum techniques.

Timing has difficulties and is complex in the quantum world. Apparently, consensus has not been reached about what to time, how to measure it, and how to do it effectively in quantum optimization or for other types of quantum algorithms.

Energy gap analysis and circuit depth efficiency are appropriate measurements for some quantum processors and can be used as benchmarks.

5 Future work

In this section we ask major questions to stimulate studies. How well do results for the TSP on a specific quantum processor predict the degree of difficulty to solve combinatorial optimization problems on this processor? Does successful solution for several TSPs of the same size on a quantum processor imply success for solving optimization problems of the same size on this processor?

The following questions pertain to the TSPLIB [19] and the CVRPLIB [4] adapted to TSPs. Are the problems in the library suitable for the emerging technology about quantum computing [22]? Are the problems in the library too artificial, i.e., not related to real applications [22]? Are the problems in the library too homogenous, i.e., not covering the wide range of characteristics found in real applications and varying degrees of difficulty to solve [22]?

The outcome of the work in the previous paragraph should include a family of small TSPs that the quantum organizations find suitable for benchmarking.

An important investigation is to determine how well performance metrics hold as problem sizes scale. This includes analyzing how adding cities to a TSP impacts quantum computational time and solution quality.

Controlled noise tests can be used to benchmark the robustness of algorithms under varied hardware noise conditions in some quantum processors.

Questions have been raised if the methodology in this paper can be combined with modern machine learning and deep learning? Or can it be combined with methods from Artificial Intelligence and/or Computational Intelligence?

There is interest in hybrid quantum-classical algorithms. Benchmarks for their scalability and effectiveness will be needed.

6 Summary and conclusion

We have added four features to the list of reasons why symmetric TSPs are an excellent choice for benchmarking quantum combinatorial optimization problems. We have discussed ways to overcome a lack of small, standard TSPs. We have pointed out benchmark measures and identified several that are being used. We have identified significant open questions to investigate.

In conclusion, we continue to recommend symmetric TSPs to be benchmarks for combinatorial optimization problems on all types of quantum hardware.

Acknowledgements:
Jesus Christ who is reliable and consistent has become my solid rock in these troubled times. I trust Him for the future and hope to be with Him forever.

References:

[1] Chieza HA, Khumalo MT, Prag K and Woolway M On the computational performance of IBM quantum devices applied to combinatorial optimisation problems, *7th Intl. Conference on Soft Computing & Machine Intelligence* (ISCMI), 2020, 260–264, IEEE 2020

[2] Concorde http://www.math.uwaterloo.ca/tsp/concorde/downloads/downloads.htm. Accessed October 2024

[3] Cook WJ National Traveling Salesman Problems https://www.math.uwaterloo.ca/tsp/world/countries.html Accessed October 2024

[4] CVRPLIB Capacitated Vehicle Routing Problem Library, http://vrp.atd-lab.inf.puc-rio.br *Accessed] October 2024*

[5] D-Wave TSP 48 Cities Accessed December 8, 2019. Currently removed from the D-Wave site.

[6] Devitt, SJ How do we quantify the utility of quantum algorithms? *Research Directions: Quantum Technologies* **1**, e5, 2023. doi.org/10.1017/qut.2022.7

[7] Grant E, Humble TS and Stump B**,** Benchmarking annealing controls with portfolio optimization, *Physical Review Applied* **15**, 014012quantum, 2021. DOI: 10.1103/PhysRevApplied.15.014012

[8] Hougardy S and Zhong X, Hard to solve instances of the Euclidean traveling salesman problem, *Math. Prog. Comp.* 13:51–74, 2020. doi.org/10.1007/s12532-020-00184-5 Hard

[9] King J, Yarkoni S et al**,** Benchmarking a quantum annealing processor with the time-to-target metric, arXiv:1508.05087, 2015.

[10] Khumalo MT, Chieza HA, Prag K and Woolway M, An investigation of IBM quantum computing device performance on combinatorial optimisation problems, arXiv:2107.03638v2, 2021.

[11] Lubinski T, Johri S et al, Application-Oriented Performance Benchmarks for Quantum Computing arXiv:2110.03137, 2021.

[12] Lubinski T, Emerson J et al, Panel 18, On the Use of Applications as Performance Benchmarks for Quantum Computing, *IEEE International Conference on Quantum Computing and Engineering* (QCE21), 2021.

[13] Lucas A, Ising formulations of many NP problems. *Frontiers in Physics* **2** Article 5, 2014. DOI:10.3389/fphy.2014.00005

[14] Martoňák R, Santoro GE and Tosatti E, Quantum annealing of the traveling salesman problem. *Physical Review E* **70**, 057701, 2004. DOI:https://doi.org/10.1103/PhysRevE.70.057701

[15] Mersmann O, Bischl B et al, A Novel Feature-Based Approach to Characterize Algorithm Performance for the Traveling Salesman Problem, arXiv:1208.2318, 2012.

[16] Osaba E, Villar-Rodriguez E, Oregi I, Moreno-Fernandez-de-Leceta A, Hybrid Quantum Computing - Tabu Search Algorithm for Partitioning Problems: Preliminary Study on the Traveling Salesman Problem, *2021 IEEE Congress on Evolutionary Computation*, 2021. DOI: 10.1109/CEC45853.2021.9504923

[17] Osaba E, Villar-Rodriguez E, Oregi I, Moreno-Fernandez-de-Leceta A, Focusing on the Hybrid Quantum Computing - Tabu Search Algorithm: new results on the Asymmetric Salesman Problem, arXiv:2102.05919v2, 2021.

[18] Osaba E and Villar-Rodriguez E, QOPTLib: A Quantum Computing Oriented Benchmark for Combinatorial Optimization Problems. In: Yang, XS. (ed) Benchmarks and Hybrid Algorithms in Optimization and Applications. Springer Tracts in Nature-Inspired Computing. Springer, Singapore, 2023. https://doi.org/10.1007/978-981-99-3970-1_4

[19] Rinaldi G, TSP Library: A library of sample instances for the TSP (and related problems) 2014. URL: http://comopt.ifi.uni-heidelberg.de/software/TSPLIB95/ accessed October 31, 2023

[20] Shammah N, Roy AS et al, Open Hardware in Quantum Technology, arXiv:2309.17233v1, 2023.

[21] Silva C, Aguiar A et al, Mapping a logical representation of TSP to quantum annealing. *Quantum Information Processing* **20**, 386,

2021. https://doi.org/10.1007/s11128-021-03321-8
[22] Uchoa E et al, New benchmark instances for the Capacitated Vehicle Routing Problem. *European Journal of Operational Research* **257**, 845–858, 2017. https://doi.org/10.1016/j.ejor.2016.08.012
[23] Villar-Rodriguez E, Osaba E et al, Analyzing the behaviour of D'WAVE quantum annealer: fine-tuning parameterization and tests with restrictive Hamiltonian formulations. arXiv:2207.00253v1, 2022.
[24] Warren RH, Adapting the traveling salesman problem to an adiabatic quantum computer. *Quantum Information Processing* **12**, 1781-1785, 2013. DOI:10.1007/s11128-012-0490-8
[25] Warren R *Sciences* **2**, H, Solving the traveling salesman problem on a quantum annealer. *SN Applied* Art number: icle75, 2019. https://doi.org/10.1007/s42452-019-1829-x
[26] Warren RH, Quantum Solutions of 6-City Traveling Salesman Problems. In Østergaard CV (ed), *A Closer Look at the Travelling Salesman Problem*. Nova Science Publishers. Chapter 4, 2020 https://novapublishers.com/shop/a-closer-look-at-the-travelling-salesman-problem/
[27] Warren RH, A benchmark for quantum optimization: the traveling salesman problem. *Quantum Information and Computation* **21**, 557-562, 2021. doi: https://doi.org/10.26421/QIC21.7-8-2
[28] Warren RH, Solving combinatorial problems by two D-Wave hybrid solvers: a case study of traveling salesman problems in the TSP Library, arXiv:2106.05948, 2021.

Contribution of Individual Authors to the Creation of a Scientific Article (Ghostwriting Policy)
The author contributed in the present research, at all stages from the formulation of the problem to the final findings and solution.

Sources of Funding for Research Presented in a Scientific Article or Scientific Article Itself
No funding was received for conducting this study.

Conflict of Interest
The author has no conflict of interest to declare that is relevant to the content of this article.

Appendix: Summaries of literature about TSPs related to benchmarking

This section contains an annotated bibliography of high-quality, published papers about quantum benchmarking that were not known to us when we wrote [27].

Article [1] is an outstanding investigation of IBM's hybrid algorithm Variational Quantum Eigensolver (VQE) to solve TSPs with size n = 3, 4, 5, and 6 cities on IBM quantum devices with 20, 27, and 53 qubits. The quantum formulation of an n-city TSP requires at least n^2 qubits. This means the larger TSPs cannot be processed on the smaller quantum machines.

Paper [10] is an expansion of [2] by the same authors. A second hybrid algorithm, Quantum Approximate Optimisation Algorithm (QAOA), is investigated. An additional IBM quantum device with 65 qubits was used. The standards were generated by BNB and SA on classical processors. The results are similar to those by [1] in the previous paragraph. IBM hybrid algorithms VQE and QAOA on IBM quantum processors failed standards for success rate and computational performance for TSPs whose sizes are n = 4, 5, 6, and 7 cities.

Reference [14] uses a symmetric 1,002-city TSP from the TSPLIB [19] as a benchmark to show that 2-opt quantum annealing is superior to standard thermal simulated annealing. Investigations in [15] point out statistical features of the distances that make TSP instances easy or hard to solve optimally by 2-opt.

Reference [7] has excellent insights about values for controls for the D-Wave 2000Q processor for portfolio problems. This processor has about 2000 qubits and limited connections between qubits. Portfolio problems and traveling salesman problems are both fully connected, optimization problems. Thus, we suggest a similar study for the TSP on the D-Wave Advantage processor which has 5000 qubits and more connections between qubits.

The authors of [16] present a new algorithm called hybrid Quantum Computing - Tabu Search Algorithm (QTA) to solve partition problems on D-Wave processors. Results on seven symmetric TSPs were compared with those from D-Wave's algorithm GBSolv, which is also a hybrid that uses tabu search. Each TSP was executed by each algorithm 20 times to find the length of an optimal tour. The averages of the results in Table 1 of [16] show there are negligible differences between the two algorithms. Testing was done on a classical computer that simulated the quantum calls. Therefore, it is recommended that testing of QTA and GBSolv be repeated in a hybrid arrangement that uses a quantum device.

Paper [17] is a follow-on to [16] by the same authors. In [17] they report about testing six asymmetric TSPs from the TSPLIB [19]. Table 1 in [17] contains results for shortest length of a tour from QTA and GBSolv, each run in a classical only mode. The results show that the differences are negligible, as in [16]. Recommend that testing of QTA and GBSolv be repeated in a hybrid arrangement that uses a quantum device.

The authors of [11] have written an extensive paper. They introduce a suite of application-oriented performance benchmarks for gate model quantum processors. The benchmarks include quality and timing and are designed to apply to a wide range of applications. The benchmark measurements are based on performance related to the volumetric features of gate model algorithms which are the width and depth of a circuit. The paper contains results from executing some of the benchmarks on quantum simulators and quantum hardware.

Some of the panelists in reference [12] are authors of [11]. Since [12] is not available, we quote from its Abstract. "… Recently, there has been increased interest in using quantum applications to determine how well quantum computers can execute real workloads, both simple and complex. Familiar algorithms or applications, structured to collect performance metrics can provide a useful mechanism for new users to understand just what a quantum machine is capable of. However, the use of applications as benchmarks is sometimes considered a 'blunt tool' for characterizing hardware performance. This panel will address the challenges, as well as the merits and weaknesses of this application-oriented approach to benchmarking the performance of quantum computers."

The main result in [21] is that the average execution time to find the first optimal solution to a TSP by the D-Wave Advantage processor is 26% (n = 7 cities) and 47% (n = 6 cities) better than classical C according to Table 5. The authors show that a specific TSP with n = 8 to 14 cities can be embedded in the D-Wave Advantage processor. Therefore, it is recommended that the study in [21] about execution times and finding optimal tours be repeated for TSPs with n = 8 to 14 cities. The study could be enhanced by having more than one TSP for each n since there are TSPs that are easy to solve and hard to solve optimally.

The authors of [18] propose a benchmark (QOPTLib) for quantum computing of combinatorial optimization problems. Their goal is to create "a number of instances able to evaluate any quantum solver in any quantum processing unit" Section 3 paragraph 3. Their benchmark library is composed of four problem types, one of which is the TSP. The 10 TSP instances in the QOPT Library include one for each number of cities 4-10, 15, 22, 25. D-Wave was used for this study since it accepts larger problem sizes than other quantum processors.

The focus of [23] is a study of two Hamiltonian formulations for combinatorial optimization problems to access a D-Wave quantum processor. The work is benchmarked by one symmetric TSP that has 7 nodes and is a subset of Burma14 [19] with normalized distances. Data is also shown for one similarly formed TSP for the cases when the number of nodes is 5, 6, 12, 13 and 14. The last three TSPs are too large to embed in the D-Wave processor available to the authors. A hybrid quantum classical method was used for them.

The Benchmarks subsection of Section III in [20] contains an excellent description of the use of standardized benchmarks in the quantum environment. References cited in the Benchmarks subsection of [20] have relevant titles. Apparently, D-Wave and its quantum processors are not mentioned in [20].

Authors [9] present an unusual metric to overcome two obstacles in D-Wave's quantum annealing. Since there is a complete equivalence in principle between the gate method for quantum computing and the quantum annealing method, this metric should also apply to quantum computing on gate hardware.

In paper [6] the reader is asked to contribute "new techniques, tools, software and hardware capabilities to identify how best to utilize quantum technologies". The question in the title of [6] does not seem to be addressed.