White Paper – The Quantum Travelling Salesman

Tristan Cook, Dan Padilha, Duncan Fletcher, James Forrest

Quantum computing has been touted as a game-changing technology, but what will its true impact be? QxBranch has developed a method for solving the Travelling Salesman Problem – a flagship challenge in optimisation – using the D-Wave Two® quantum computer. Using a quantum computing approach has the potential to revolutionise the many applications of this difficult problem.

Introduction

In this paper, we demonstrate applying the D-Wave Two® quantum computer to the "multiple vehicle" variant of the Travelling Salesman Problem. Using quantum computing to tackle this challenging problem has the potential to produce better solutions quicker, revolutionising its many real world applications.

The Travelling Salesman Problem (TSP) is one of the most intensely-studied problems in optimisation. It asks; "For a given set of cities, what is the shortest possible route that visits each city exactly once before returning to the initial city?" While it appears to be a simple problem at first glance, its complexity rapidly increases as cities are added; with just 15 cities, there are in excess of 87 billion possible routes (and that's for just a single vehicle). Practical applications of the TSP are extensive - not just for a literal traveling salesman or courier, but across many industry sectors like manufacturing, DNA sequencing, and logistics.

The D-Wave Two is a specialised computational architecture that implements adiabatic quantum computation. This form of quantum computing is ideal for solving hard optimisation problems such as the TSP.

This paper:

- discusses the background and significance of the Travelling Salesman Problem;
- provides a background understanding of the D-Wave Two quantum computer;
- walks through some example problems using our method;
- shows the results of running the algorithm on these examples; and
- discusses the extension of this method to accommodate more complex formulations.

Background – Travelling Salesman Problem

The Travelling Salesman Problem (TSP), also known as the Vehicle Routing Problem, has been the subject of intense study because of its wide-ranging, practical utility and lack of a general method of solution (Toth & Vigo, 2002). The TSP is what's called an 'NP-Hard problem'; finding the optimal solution is a hard, non-deterministic problem. These sorts of problems are very difficult for conventional computers to solve.

To further complicate this already hard problem, in real-world situations it is often necessary to optimise not just one, but a fleet of vehicles (or application-specific equivalent, such as parallel production lines)! The method we propose is designed for a multi-vehicle formulation of the TSP.

There are many existing classical algorithms that attempt to solve the TSP, but there is a trade-off between optimality and computation time – to achieve better results, more time is required. These classical approaches include search algorithms such as branch-and-bound, heuristics such as the minimum spanning tree, and iterative improvement algorithms such as pairwise exchange.

The method employed in this paper uses *quantum annealing* to minimise the value of a mathematical expression that finds and defines the optimal path. The potential advantage of this method is that with sufficiently large quantum computing hardware the size of the problem may have no effect on computation time. With their potential to solve exponentially-hard problems in linear or fixed time, quantum computers hold the promise of rendering this intractable problem trivial.

Background – D-Wave Two Architecture

The D-Wave Two is a commercially-available quantum computer designed to perform optimisation using *quantum annealing*. This is achieved through hardware minimisation of the following equation, called the *Ising Spin Glass:*

$$H = \sum_{i \in V} h_i s_i + \sum_{\substack{(i,j) \in E \\ i<j}} J_{i,j} s_i s_j \quad (1)$$

The variables s_i represent the spins of the quantum bits or 'qubits' of the computer. These take a value of ±1 and form the **output** of the program. The **inputs** of the quantum program, h_i and $J_{i,j}$ are the weightings on the qubits (biases) and the coupling strengths between them, respectively. When values are assigned to the h and J terms, we have what's called the **problem Hamiltonian** H or energy function that is to be minimised. The values of s_i that give the lowest energy form the best solution.

Solving a problem on the D-Wave Two requires the problem to be expressed in terms of this Hamiltonian, and then mapped to the hardware qubits. There are 512 sparsely-coupled qubits in the D-Wave Two, which can be arranged to create roughly 30 fully-coupled qubits. This limitation restricts the size of problems that can be run on current versions of the machine.

Quantum Annealing a Route

We have designed a *problem Hamiltonian* for the multi-vehicle TSP[1] that transforms to the Ising Spin Glass and runs on the D-Wave Two to find the optimal routing. The final qubit states of the lowest energy result found by quantum annealing define the shortest path.

The number of qubits[2] required to solve the problem using our method is:

$$Qubits = (Cities^2 - Cities) \times Vehicles \quad (2)$$

The sparse connectivity of the D-Wave Two's 512-qubit hardware graph means that it can only support a maximum problem graph size of 5 cities (including the origin), for 2 vehicles. However, the effectiveness of the method is also demonstrated using Isakov's simulator of the D-Wave architecture on a classical computer (Isakov, Zintchenko, Ronnow, & Troyer, 2014). This classically-based simulation cannot take advantage of quantum effects like entanglement and superposition; consequently, solution quality degrades beyond 350 simulated qubits.

Our results indicate that the method is only limited by the number of qubits supported in hardware; so as larger quantum hardware becomes available, larger graphs will be solvable.

In the next two sections we'll show examples of using the method to solve TSP problems; first on the D-Wave Two hardware, and then using the software simulator to demonstrate how it scales to larger graphs.

Example 1 – Four Visits on the D-Wave Two

The first example is for a small graph with four nodes to visit from a depot and finds the optimal routing for two vehicles to traverse the graph (Figure 1, over). This size is chosen as it is

[1] This problem Hamilton is an extension of the single-vehicle Travelling Salesman formulation given by (Lucas, 2014).

[2] Our current multi-vehicle TSP Hamiltonian requires densely-connected qubits; however, they are not all fully connected.

the largest that will fit on the D-Wave Two hardware at this time. This example uses 40 nearly fully-connected qubits.

The edge weights for this example simply represent the road-distances between the nodes. For real-world applications, these could be replaced with some combination of travel distance, time, and/or monetary cost.

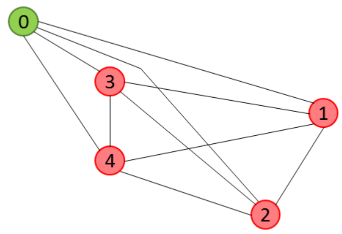

Figure 1 The 5 node graph to be routed. Node 0 is the depot: the start and end point of the route.

Once all these weights have been assigned the final Hamiltonian can be generated and mapped to the qubits of the quantum computer. The D-Wave Two performs up to 35,000 anneals per run[3] and returns the results of each individual anneal, ordered by energy.

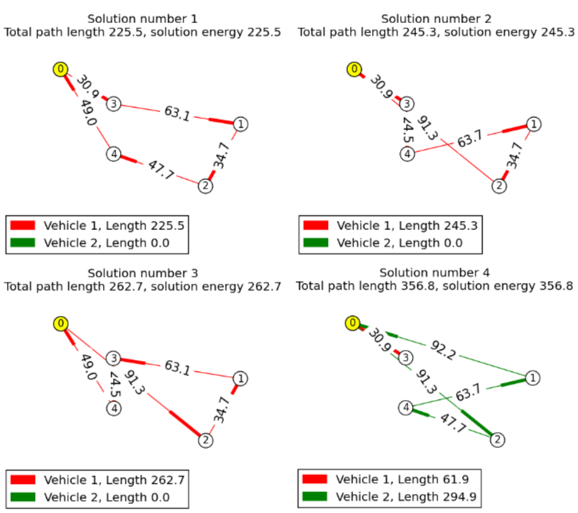

Figure 2 The four best results from running the problem on the D-Wave Two. The thickened end of the edge-lines indicates where the vehicle starts from on each leg of its journey.

One characteristic of quantum computing is its probabilistic nature; the best result found is *probably* the best solution to the problem; thousands of anneals are performed to increase the probability of finding the best solution. Each individual anneal returns a set of qubit values so in addition to the best solution, the computer will provide a spectrum of very good answers essentially for free. This may be useful if the best solution is not a viable option – perhaps a route is interrupted by road works; the next-best solution can provide an alternative.

Figure 2 shows the four best results returned by the D-Wave Two for the above graph using our method.

We can see that the results are indeed ordered from best to worst, with the solution energies equal to the total travel cost of each solution. For graphs of this size, classical computers can use brute force to prove that the first result is the actual optimal path.

Example 2 – 14 Visits by Simulation

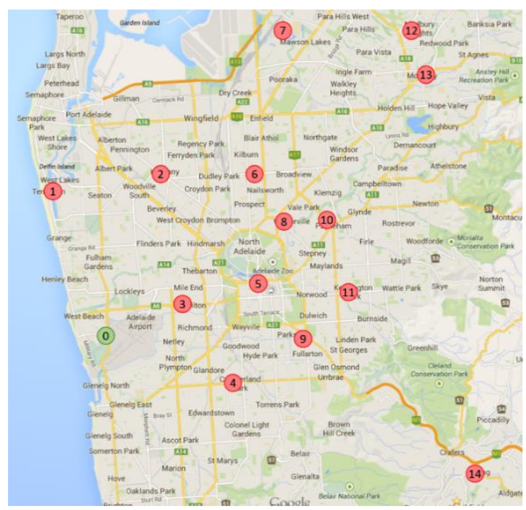

Figure 3 Node 0 (in green) is the depot at Adelaide International Airport, with nodes 1-14 the locations to be visited by the two vehicles.

The next example is for two vehicles servicing a set of 14 locations about the charming city of Adelaide, Australia, shown in Figure 3. As the

[3] Each anneal takes 20 micro-seconds. 35,000 anneals takes approximately one second for the quantum computer to execute including the on-chip setup and read-out, plus internet round-trip time to submit the job to the remote machine.

problem is too large to fit on the current D-Wave Two hardware, we solve it using Isakov's simulator of the D-Wave Two's architecture.

The locations to be visited are represented as nodes on a graph and edges are created between them, as in Figure 4 (over). This demonstration uses 'as the crow flies' distances as the edge weightings; however, real road distances, travel time and driver wage costs could all be included when formulating the edge weights.

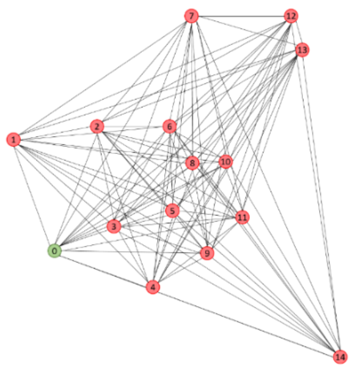

Figure 4 The complex, fully-connected graph for Example 2.

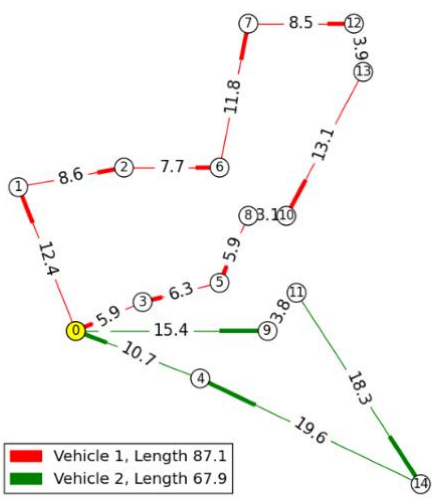

Figure 5 The best solution for Example 2 after 100,000 runs of Isakov's simulator.

We found that if we relax our constraint to balance the vehicle fleet's workload, it is actually most efficient for the entire route to be served by a single vehicle. However, real-world problems have capacity-constrained vehicles and customer expectations of timely delivery. Hence the mathematically-pure optimal is not practically useful. Our method introduces load-balancing as a feature of the solution. The best result from 100,000 runs[4] of the simulator is shown in Figure 5. This particular example simulated 332 qubits with approximately 65 couplings per qubit.

Summary and Future Capability

Using adiabatic quantum computing to solve Travelling Salesman Problems has the potential to drastically cut computation time and better optimise the results for real world problems.

Further, in practice, the optimal route is not as simple as just shortest distance or quickest time. What about priority customers who need to be serviced earlier in the day? What about vehicles with limited capacity making deliveries of different sizes, that need to return to a depot for refilling? What about the step-cost of each additional vehicle in the delivery fleet and overtime costs for long days? A Hamiltonian formulation of the TSP can incorporate such additional factors without increasing the computation time, provided that sufficiently large quantum hardware is available.

Quantum computing is opening up exciting possibilities in the realm of optimisation and this paper demonstrates an example of the D-Wave Two machine solving these problems for real-world applications.

References

Isakov, S. V., Zintchenko, I. N., Ronnow, T. F., & Troyer, M. (2014, June). Optimised simulated annealing for Ising spin glasses. *arXiv, 1401.1084*.

Lucas, A. (2014). Ising formulations of many NP problems. *Frontiers in Physics 2, 5*.

Toth, P., & Vigo, D. (2002). *The Vehicle Routing Problem.* Philadelphia: Society for Industrial and Applied Mathematics.

[4] Isakov's simulator achieves approximately 3,000 runs per second for a problem of this size.

