Computers and Operations Research 87 (2017) 52–62

Contents lists available at ScienceDirect

Computers and Operations Research

journal homepage: www.elsevier.com/locate/cor

Large-scale vehicle routing problems: Quantum Annealing, tunings and results

A. Syrichas*, A. Crispin

Manchester Metropolitan University, School of Computing, Mathematics and Digital Technology, John Dalton Building, Chester Street, Manchester, M15 5GD, United Kingdom

ARTICLE INFO

Article history:
Received 18 May 2016
Revised 8 May 2017
Accepted 23 May 2017
Available online 25 May 2017

Keywords:
Quantum
Annealing
CVRP
DCVRP
Tuning

ABSTRACT

Quantum Annealing was previously applied to the vehicle routing problem and the results were promising. For all benchmark instances in the study, optimal results were obtained. However, 100% success rate was not achieved in every case, and tuning the control parameters for larger instances proved cumbersome. This work addresses these remaining difficulties. An empirical approach is taken wherein measurements of run-time behaviour are exploited to transform existing good values of control parameters so that they can be used successfully for other problem instances. The course of this work shows a method which simplifies hand-tuning so that the heuristic performs successfully when applied to larger instances, and also demonstrates a tuning method which establishes control parameter values for instances which belong in broadly defined groupings. In addition, new best known solutions for large-scale instances, and initial results for the distance-constrained variant of the vehicle routing problem are presented.

1. Introduction

Previous research Crispin and Syrichas (2013) demonstrated the effectiveness of Quantum Annealing (QA) for solving many instances of the Capacitated Vehicle Routing Problem (CVRP). Optimal results were obtained for all benchmark instances by applying a single set of values for the algorithm's control parameters - values which were methodically determined to achieve the maximum success rate for a reference instance. The success rate is the percentage of the number of times the algorithm finds the best known score for a given instance over a number of runs.

Table 1 shows an excerpt of the instances for which this parameter set was unable to achieve 100% success rate. (An indication of the complexity of each instance can be inferred from the name. P-n101-k4 for example, has 101 nodes/customers served by 4 vehicles whereas M-n121-k7 has 121 served by 7.) Notably, the scores for smaller instances were much lower than for the reference instance. This is contrary to intuition, that one might expect parameters giving the best results for a larger instance would perform easily as well for smaller instances. (One may expect also that parameters for smaller instances will not work well for larger ones.) Given that the local search method is effective enough to allow the metaheuristic to find the optimal solution in at least 11% of the experiments, and that many of the instances appear less complex than the reference, one can conclude that the values of the control parameters are incorrect. If the temperature value is set too high or the magnetic field is too strong, convergence to a minimum is slowed down or inhibited completely. If set too low, the rate of convergence is high and entrapment at poor local minima is likely. If the population size is too small, the search of solution space covers only a reduced area which may not contain optimal solutions. It seems clear that the universal application of a single set of control parameters will not guarantee consistently good performance and the algorithm requires tuning on a case-by-case basis.

How then does one tune metaheuristic control parameters for best results? One could apply the tuning methodology for every instance, providing a specific set of control parameters for each. To save time, the processes of the methodology could be captured, encoded, and then left to a computer program to automatically decide the parameters. These approaches work because feedback can be derived from information known beforehand about the optimal result. Benchmark instances are often supplied with deterministically proven optimal solutions. However, for larger benchmarks, and in dynamic or industrial applications where problem instances are created in real-time, such information is limited or non-existent.

Additional tuning difficulties are presented by metaheuristics with two or more control parameters, each of which may be tightly interdependent. For example, the coupling term used in QA is a

* Corresponding author.
E-mail address: alex.syrichas@gmail.com (A. Syrichas).

http://dx.doi.org/10.1016/j.cor.2017.05.014

Table 1
Computational results excerpt.

Instance	QA Success %	SA Success %
P-n101-k4 (reference)	**100**	**100**
P-n50-k10	63	28
P-n55-k10	35	31
P-n60-k15	79	79
P-n70-k10	78	61
P-n76-k4	52	44
P-n76-k5	87	22
B-n63-k10	26	25
B-n66-k9	91	44
B-n67-k10	42	88
B-n68-k9	69	87
B-n78-k10	97	99
M-n121-k7	90	76
F-n135-k7	11	4

non-linear function of magnetic field strength and effective temperature. This term is extremely sensitive to variations in either parameter, and tuning is further complicated because the Metropolis criteria (Metropolis et al., 1953) is simultaneously dependent upon temperature. A Design of Experiments (DoE) method (Ridge and Kudenko, 2010) can be helpful in uncovering major dependencies between such variables, but a course of factorial experiments can be time-consuming, and predicting the ranges for numerous and sensitive variables is difficult without once again resorting to guesswork or serially hand-tuning.

It is for reasons like these that metaheuristics with fewer control parameters are attractive - they are simpler to tune. Late Acceptance Hill Climbing (Burke and Bykov, 2017) has a single parameter controlling the size of a fitness array which acts as a 'memory' of good solutions. Cuckoo Search is reported (Nie et al., 2014) to be superior to Genetic Algorithms in part because of having only two parameters - nest abandonment rate and population size. In QA, it has been shown (Titiloye and Crispin, 2011) that the number of parameters can be reduced by one, by setting the magnetic field value to be constant. This idea can be greatly extended by making the whole coupling term a constant, thereby removing the mutual dependence of the effective temperature and magnetic field parameters. With the key parameters uncoupled from one another, time is saved when determining their values by hand. Large-scale problems and instances of the Distance-constrained Capacitated Vehicle Routing Problem (DCVRP) may be tackled without tedium. Furthermore, if some means other than hand-tuning can predict the value of temperature, a single variable remains to be tuned - the replica count (population size).

2. Quantum Annealing

Quantum Annealing is an energy-based metaheuristic which uses the Path-Integral Monte Carlo (PIMC) method (Battaglia et al., 2005) to approximate the ground state of the Ising Model. The fitness function is described (1) by the Hamiltonian

$$H = H_p + H_k \tag{1}$$

where the cost H is the sum of potential energy H_p and fluctuations in kinetic energy H_k. H_k is the term which represents the quantum mechanical phenomenon of tunnelling, where a particle trapped in a low energy state, can 'tunnel' through high potential barriers into a lower state. This effect can be simulated in a metaheuristic by using an Ising Model representation of the optimization problem. In simple terms, this is maintaining a population P of simultaneously evolving solutions called replicas, where H_k is calculated from an interaction between adjoining replicas.

When QA is applied to an optimization problem, H_p takes the role of the cost of a solution (for VRP, see (4)), while H_k is a

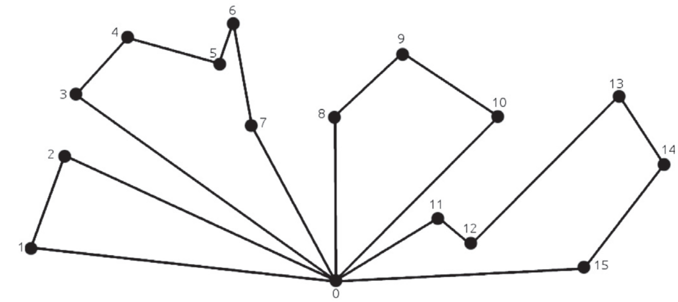

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Hex
0	0	1	1	1	0	0	0	1	1	0	1	1	0	0	0	1	8D8E
1	1	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0005
2	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0003
3	1	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0011
4	0	0	0	1	0	1	0	0	0	0	0	0	0	0	0	0	0028
5	0	0	0	0	1	0	1	0	0	0	0	0	0	0	0	0	0050
6	0	0	0	0	0	1	0	1	0	0	0	0	0	0	0	0	00A0
7	1	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0041
8	1	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0201
9	0	0	0	0	0	0	0	0	1	0	1	0	0	0	0	0	0500
10	1	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0201
11	1	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	1001
12	0	0	0	0	0	0	0	0	0	0	0	1	0	1	0	0	2800
13	0	0	0	0	0	0	0	0	0	0	0	0	1	0	1	0	5000
14	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	1	A000
15	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	4001

Fig. 1. Example of routes encoded as a spin matrix. Customer-customer connections are represented using single bits. The matrix is encoded row-wise to form hexadecimal (Hex) words which stored in memory as an array.

scaled sum of the spin interactions between P neighbouring solutions held in a circular list.

$$H_k = J_\Gamma \sum\nolimits_P \sum\nolimits_i \sigma_{P-1,i}\sigma_{P,i}\sigma_{P+1,i} \tag{2}$$

Each replica represents the solution as a spin matrix σ containing i elements which can assume values of $\{-1, +1\}$. The interaction energy between the spins of adjoining replicas is generated by the term, $\sigma_{P-1,i}\sigma_{P,i}\sigma_{P+1,i}$

J_Γ is the coupling term which is normally varied during the annealing process via adjustments to the magnetic field strength Γ, amplifying or attenuating the interactions between replicas.

$$J_\Gamma = \frac{-T}{2} \ln \tanh\left(\frac{\Gamma}{PT}\right) \tag{3}$$

Consequently, this contributes towards the acceptance of H in the Metropolis criteria.

2.1. QA for CVRP, and the PT tuning method

CVRP is a variant of VRP in which all vehicles are subject to the same capacity constraint Q. CVRP is an undirected graph $G = (V, E)$ consisting of the vertex set $V = \{v_0, v_1, \ldots v_n\}$ and edge set $E = \{(v_i, v_j) | v_i, v_j \in V, i < j\}$. The restriction $i < j$ ensures the distance between a pair of vertices is identical in both directions. The first vertex is usually considered to be the depot from which a fleet of trucks m serves n customers, whose locations are represented by a vertex set, and have varying demands for goods q_i . The goal is to minimize the number of routes and/or total distance travelled by the trucks d_{ij}. QA for CVRP (QACVRP) uses a two-dimensional spin matrix in which the elements represent customer-customer connections that form routes for each truck. A non-zero cell in