Solving large optimization problems with restricted quantum annealers

T. Albash[1], F. Spedalieri[1], I. Hen[1], K. Pudenz[2], G. Tallant[2]
[1]Information Sciences Institute, University of Southern California
[2]Lockheed-Martin Corporation

Quantum annealers provide a novel paradigm for solving combinatorial optimization problems, by exploiting the quantum mechanical properties of a set of qubits with programmable interactions [1]. Quantum annealers like the D-Wave processor series, have been commercially available for several years, and a lot of research has been done trying to understand their properties, in particular whether or not they can provide a computational speedup [2]. But another important aspect of quantum annealers is related to the limitations that are imposed by the many times unavoidable engineering compromises that have to be reached in order to build a scalable and useful device.

Two of the main limitations of the D-Wave processor that would likely be shared by other implementations of quantum annealing, are size (largest available device has ~1000 qubits) and hardware connectivity (not every interaction between variables is physically represented). This restricts the problems that can actually be solved by the device to those that can be mapped directly to the underlying architecture. Since most instances of real-world applications will no fit this restricted setting, tools need to be developed in order to apply quantum annealing to a wider range of problems. In this work we will present a set of tools designed to address these two important limitations, by providing strategies to decompose and embed general Quadratic Unconstrained Binary Optimization (QUBO) problems using the D-Wave processor.

The aim of problem decomposition is to take a problem with more variables than the qubits available in the quantum annealer, and break it into a series of subproblems that can be fit in the processor and use these results to reconstruct a solution for the original problem. We will discuss two approaches for the decomposition:

1. Backbone based: this approach tries to identify which variables must have a fixed value for the solution to be of high quality. In this way, these variables can then be fixed and the optimization can be carried out over a reduced set of variables that can be fit in the quantum annealer processor. This is an iterative process, that at each step will produce a set of candidate variables to be fixed [3]. This method is based on the idea of the "backbone" of a solution, that has been introduced in the constraint satisfaction literature.
2. Coupling based: this approach aims at identifying clusters of variables that are strongly interacting, and henceforth will likely be of greater importance in determining the values of the variables involved [4]. To build the clusters, a variable is chosen at random, and then other variables are added according to the strengths of their couplings to the variables already in the cluster. The size of the cluster can be chosen such that it can be solved by the quantum annealer.

These two approaches start by generating a random solution that is subsequently updated after each call to the quantum annealer to solve the corresponding subproblem. This generates a sequence of candidate solutions that improve the value of the objective function.

978-1-5090-3525-0/16/$31.00 ©2016 IEEE

The subproblems generated by these two approaches will not necessarily fit the connectivity graph of the processor. To move forward we need to specify a strategy to embed them in the restricted architecture. It is clear that we would like to make the subproblems as large as possible in order to minimize the number of decomposition steps. We will analyze three embedding approaches that have different strengths and weaknesses in this regard:

1. Direct canonical embedding: this approach uses a technique known as minor embedding [5] to map problem with arbitrary connectivity into the D-Wave processor. The advantage of this approach is that the mapping is fixed and can be computed at the beginning, minimizing any overhead on the computation. The drawback is that the size of the subproblems is limited to ~45 variables.
2. Direct adaptive embedding: this approach tries to embed a slightly larger subproblem by computing a specialized (not canonical) minor embedding. The potential advantage in reducing the number of decompositions needed must be weighed against the increased computational overhead of finding these embeddings.
3. Iterative sampling: this approach uses the sampling capabilities of the quantum annealer to generate good solutions for the subproblem. It uses an iterative procedure to adjust the parameters of the annealer in order to mimic a Gibbs distribution associated with an arbitrary QUBO problem. The advantage of this approach is that the only limitation to the size of the subproblem is given by the size of the processor (~1000 qubits for the DW2X processor). The drawback is that the approach has no optimality guarantees.

We will present the results of our analysis of these methods using a DW2X processor with 1,098 available qubits on several benchmark problems. Our goal will be to estimate the computational overhead that each approach requires, and try to understand if these approaches are a feasible road to exploit limited quantum annealers for real-world applications.

References

[1] M. Johnson, et al., "Quantum annealing with manufactured spins", *Nature* **473**, 194–198 (12 May 2011).

[2] T. Roennow, et al., "Defining and detecting quantum speedup", *Science* **345**, 420 (2014).

[3] Y. Wang, et al., "A Multilevel Algorithm for Large Unconstrained Binary Quadratic Optimization", *Integration of AI and OR Techniques in Contraint Programming for Combinatorial Optimzation Problems*, Volume 7298 of the series Lecture Notes in Computer Science pp 395-408.

[4] I. Zintchenko, M. Hastings, M. Troyer, "From local to global ground states in Ising spin glasses", Phys. Rev. B **91**, 024201 – Published 5 January 2015.

[5] V. Choi, "Minor-embedding in adiabatic quantum computation: II. Minor-universal graph design", Quantum Information Processing, Volume 10, Issue 3 (2011).

978-1-5090-3525-0/16/$31.00 ©2016 IEEE