Towards Understanding Approximation Complexity on a **Quantum Annealer**

Irmi Sax, Sebastian Feld, Sebastian Zielinski, Thomas Gabor, Claudia Linnhoff-Popien, Wolfgang Mauerer

Extended Abstract

Many industrially relevant problems can be deterministically solved by computers in principle, but are intractable in practice, as the seminal P/NP dichotomy of complexity theory and Cobham's thesis testify. For the many NP-complete problems, industry needs to resort to using heuristics or approximation algorithms. For approximation algorithms, there is a more refined classification in complexity classes that goes beyond the simple P/NP dichotomy. As it is well known, approximation classes form a hierarchy, that is, FPTAS ⊆ PTAS ⊆ APX ⊆ NPO. This classification gives a more realistic notion of complexity but—unless unexpected breakthroughs happen for fundamental problems like P = NP or related questions— there is no known efficient algorithm that can solve such problems exactly on a realistic computer. Therefore, new ways of computations are sought.

Recently, considerable hope was placed on the possible computational powers of quantum computers and quantum annealing (QA) in particular. However, the precise benefits of such a drastic shift in hardware are still unchartered territory to a good extent. Firstly, the exact relations between classical and quantum complexity classes pose many open questions, and secondly, technical details of formulating and implementing quantum algorithms play a crucial role in real-world applications.

Guided by the hierarchy of classical optimisation complexity classes, we discuss how to map problems of each class to a quantum annealer. Those problems are the Minimum Multiprocessor Scheduling (MMS) problem, the Minimum Vertex Cover (MVC) problem and the Maximum Independent Set (MIS) problem. We experimentally investigate if and how the degree of approximability influences implementation and run-time performance.

Our experiments indicate a discrepancy between classical approximation complexity and QA behaviour: Problems MIS and MVC, members of APX respectively PTAS, exhibit better solution quality on a QA than MMS, which is in FPTAS, even despite the use of preprocessing the for latter. This leads to the hypothesis that traditional classifications do not immediately extend to the quantum annealing domain, at least when the properties of real-world devices are taken into account. A structural reason, why FPTAS problems do not show good solution quality, might be the use of an inequlity in the problem description of the FPTAS problems. Formulating those inequalities on a quantum hardware (mostly done by formulating a Quadratic Unconstrained Binary optimisation (QUBO) problem in form of a matrix) requires a lot of hardware space which makes finding an optimal solution more difficult.

Reducing the density of a QUBO is possible by appropriately pruning QUBO matrices. For the problems considered in our evaluation, we find that the achievable solution quality on a real-world machine is unexpectedly robust against pruning, often up to ratios as high as 50% or more. Since quantum annealers are probabilistic machines by design, the loss in solution quality is only of subordinate relevance, especially considering that the pruning of QUBO matrices allows for solving larger problem instances on hardware of a given capacity. We quantitatively discuss the interplay between these factors.

Irmi Sax

Irmi Sax wrote her bachelor's thesis on the topic of solving optimisation problems with Quantum Annealing. Since 2018 she is studying the Master of applied Research at the OTH Regensburg work- ing on industrial applications of adiabatic Quan- tum Computing.