This is a report on the Master's thesis submitted by Tung Anh Vu entitled "Algorithms for Low Highway Dimension Graphs".

Summary of Contents

The subject of the thesis is the study of several variations of the k-Supplier problem. In this problem we are given a set of points in a metric space, partitioned into suppliers (V_s) and clients (V_c). The goal is to select a set of k suppliers such that the maximum distance of any client from its closest supplier is minimized. The thesis studies several restrictions, generalizations, and variations of this problem, namely:

- The k-Center problem, which is the special case where $V_c = V_s$.
- The k-Supplier problem with outliers, where we are also given a budget p and are allowed to reject at most p clients.
- A capacitated version of k-Supplier with outliers, where each supplier can only serve a limited number of clients.
- A non-uniform version of the problem, where we are given k values $r_1, \ldots, r_k$ and the distances to the k-th supplier are divided by r_i.

All the studied problems are known to be NP-hard and hard to approximate in general graphs. Hence, the focus of the thesis is on attacking these problems from the viewpoint of FPT approximation algorithms. More specifically, the objective in most cases is to obtain algorithms which can produce a solution that is $(1+\epsilon)$-far from optimal, within a running time that is polynomial in the input size, but possibly exponential in a structural parameter of the input (and $k+p$). The parameters considered are highway dimension, doubling dimension, and treewidth.

The thesis presents a number of algorithmic results that extend the state of the art by building on recent papers. More precisely, after a brief overview of the field given in Chapter 1, the thesis presents the following results:

In Chapter 2, a $\frac{3}{2}$-approximation algorithm is presented for k-Center with outliers, running in time exponential in $k(h \log h) + p$, where h is the highway dimension of the input instance. This algorithm is an extension of a recent result by Feldmann, which gave a similar performance for k-Center (without outliers). Roughly speaking, the difference with Feldmann's algorithm is that at a certain phase we are facing a set of clients which are partitioned into small clusters (that can be covered by a single supplier each), which are too far away from each other to be covered by common suppliers. In the k-Center problem we then pick an arbitrary solution, while in the setting with outliers, we have a choice to reject some of these clients. The new idea of the algorithm is to perform a (standard) branching step here to decide which clients to reject. Although the new idea is not deep, I think this chapter presents a nice new result, though as the author mentions it would be nice to know if the exponential dependence on p is necessary.

In Chapter 3, the author deals with the k-Supplier with outliers problem on all three parameters, and gives FPT approximation schemes. The crux of the approach is to build upon an FPT approximation scheme for this problem parameterized by treewidth. This is achieved by modifying the approximation scheme given for k-Center by Katsikarelis et al., which combines standard dynamic programming techniques with appropriate rounding. Although the extension from k-CENTER to k-SUPPLIER with outliers is not too complicated, the author executes it competently and presents a nice and detailed exposition of the algorithm. The running time obtained is in the order of $(tw/\epsilon)^{O(tw)}$. With this algorithm in hand, the author is able to obtain an FPT approximation parameterized by highway dimension by using a theorem of Becker et al. which embeds instances of low highway dimension into graphs of low treewidth and then simply invoking the previous algorithm. In the last part of the chapter, the author presents an FPT approximation scheme parameterized by k, p, ϵ and the doubling dimension of the instance, by slightly modifying a similar algorithm for k-CENTER by Feldmann and Marx.

In Chapter 4, the author attempts to extend the results of the previous chapter to the capacitated version of the problem. This works well for doubling dimension: the author first explains how, given a set of suppliers we can check if a solution exists that respects the capacities (I think this is fairly standard); then the algorithm continues in a way analogous to the algorithm of Chapter 3. However, for treewidth the author shows that the capacitated problem is hard to approximate, by using the standard reduction from dominating set to k-CENTER and the fact that CAPACITATED DOMINATING SET is W[1]-hard parameterized by treewidth+k. Hence, the approach of the previous chapter cannot work for the capacitated problem, though it's not clear if an FPT approximation for parameter highway dimension can be obtained in some other way.

Finally, in Chapter 5 it is shown, via a small modification of a reduction of Chakrabarty et al., that the non-uniform version of the problem is hard to approximate, even on instances where all the parameters studied in this thesis are low. This is perhaps not too surprising, as this version of the problem seems much more general.

Evaluation

This thesis studies a number of topics which are at the forefront of modern research in theoretical computer science. It combines techniques from FPT algorithms and approximation, considers parameters of practical importance, and studies reasonable and well-motivated optimization problems. The author manages to show a good grasp of the state of the art and has succeeded in giving a mathematically solid but well-readable presentation of all results.

One of the main positive points is that the thesis succeeds in presenting several new results. Although in all cases the departure from what is currently known is small, I think this is something that has to be well-appreciated, as it shows that the author has mastered the ideas in several recent papers to a point where it is possible to extend them further. All in all, I think the contribution made is easily above the bar of what can be expected in a good Master's thesis, since it shows that the author has gained a good understanding of several rather advanced topics.

Even though the thesis is generally well-written, I think that one weakness is that the presentation could be improved. In particular, because the thesis deals with many different problems and parameters, the introduction would have benefitted from a table summarizing the main results. In the same vein, each chapter should probably begin with a more detailed summary of its results, and in particular with a few words explaining how the results presented differ from the literature. I would also have much appreciated a more extended comparison between the three parameters discussed, though this might go a bit beyond the scope of the thesis. However, I think these weaknesses are rather minor.

Overall, I believe the work presented constitutes a very solid Master's thesis.