University of Groningen

Advanced analysis of branch and bound algorithms

Turkensteen, Marcel

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2007

Link to publication in University of Groningen/UMCG research database

Citation for published version (APA):
Turkensteen, M. (2007). Advanced analysis of branch and bound algorithms. s.n.

Chapter 6

Conclusions

6.1 Introduction

In this dissertation, we consider solution techniques for the so-called *Combinatorial Optimization Problems (COPs)*. Given a finite set of elements, a COP is the problem of finding a combination of the elements of a subset that satisfies an a priori objective; see Section 1.1. Here, a minimum cost or maximum profit solution, i.e., an optimal solution, has to be selected. Examples of COPs can be found in network routing and scheduling, but also in many other fields of research; the market segmentation problem from Chapter 5 is an application in marketing.

There is a broad class of COPs for which exhaustive search is sometimes needed to find an optimal solution. As a consequence, finding a solution to a reasonably sized instance may take an amount of time that goes beyond any boundaries. The challenge for researchers is to enlarge the set of problems that can be solved to optimality within acceptable time limits. This dissertation presents improvements for the most common type of exact algorithms, namely *Branch and Bound (BnB)*; see Section 1.5. BnB is a collection of solution techniques in which an optimal solution is systematically searched in a so-called *search tree*: the original problem is divided into smaller problems. Although BnB solves many COPs within acceptable time limits, there is still a large set of problems that cannot be solved to optimality at the moment. In this thesis, we propose modifications of the technique that makes it suitable to solve those large or difficult problem instances to optimality.

We summarize this thesis on BnB chapterwise below.

6.2   Summary

Chapter 2: Iterative Patching and the Asymmetric Traveling Salesman Problem

Although Branch and Bound (BnB) methods are among the most widely used techniques for solving hard problems, it is still a challenge to make these methods smarter. In this chapter, we investigate *iterative patching*, a technique in which a fixed patching procedure is applied at each node of the BnB search tree for the Asymmetric Traveling Salesman Problem (ATSP), the problem of finding a shortest tour through a given set of locations. Computational experiments show that iterative patching results in general in search trees that are smaller than the classical BnB trees, and that solution times are lower for usual random and sparse instances. Furthermore, it turns out that, on average, iterative patching with the Contract-or-Patch procedure from Glover *et al.* (2001) and the Karp-Steele procedure are the fastest, and that 'iterative' Modified Karp-Steele patching generates the smallest search trees.

Chapter 3: Tolerance-Based Branch and Bound Algorithms for the Asymmetric Traveling Salesman Problem

The selection of entries to be included/excluded in Branch and Bound algorithms is usually done on the base of cost values. We consider the class of Depth First Search (DFS) algorithms, and propose to use *upper tolerances* to guide the search for optimal solutions. In spite of the fact that it needs time to calculate tolerances, our computational experiments for Asymmetric Traveling Salesman Problem show that in most situations tolerance-based algorithms outperform cost-based algorithms. The solution time reductions are mainly caused by the fact that the branching process becomes much more effective, so that optimal solutions are found in an earlier stage of the branching process. The use of tolerances also reveals why the widely used choice for branching on a smallest cycle in assignment solutions is on average the most effective one. Moreover, it turns out that tolerance-based DFS algorithms solve difficult practical instances faster than the Best First Search algorithm from Carpaneto *et al.* (1995).

Chapter 4: Additional Topics on Tolerance-Based Algorithms

In Chapter 4, additional topics on tolerance-based BnB algorithms are studied. First of all, it is found that tolerance-based BnB algorithms for the ATSP are effective on randomly generated instances with a large number of intercity distances. This is a pleasant finding, since Zhang (1993) shows that these instances are, on average, relatively difficult to solve. It also turns out that hybrid BnB algorithms, in which the power of tolerances and costs are combined, are usually not faster than the original tolerance-based algorithms for the ATSP. Finally, we show that, under certain conditions, it is possible to use multiple upper tolerance values simultaneously in a lower bound for the Degree-Constrained Minimum Spanning Tree Problem (DCMSTP).

Chapter 5: Balancing the Logistics Costs and the Fit of Market Segments

Chapter 5 applies COP techniques to market segmentation. Segments are typically formed to serve distinct groups of consumers with tailored marketing mixes, in order to better fit their needs. In geographic segmentation applications, a company responds to geographical differences by, for example, offering localized products. Existing segmentation strategies generate segments in which the constituting elements are not necessarily geographically closely located. When a geographically dispersed segment is served with one marketing mix, the logistics costs are high due to high transportation costs and long lead times. As a consequence, decision makers sometimes use other segmentation criteria. For example, a retail chain may decide to serve each country with its own formula. Such a segmentation strategy suffers from the disadvantage that transnational segments are ignored. Moreover, the results are expected to be suboptimal in terms of meeting customer needs.

In this chapter, we develop a segmentation method that balances the fit to consumer needs and logistics costs. The solution approaches are illustrated by means of the problem of assigning a set of European regions to retail formulas, using store attribute preferences of consumers as a segmentation basis. Compared to other methods, such as k-means, mixture models, and the countries-as-segments approach, our approach results in transnational segments that combine

moderate logistics costs with a relatively high level of within-segment homogeneity. These results are confirmed in experiments on randomly generated experiments. A practical implication of our study is that transnational segments may become profitable in markets with high logistics costs.

6.3   Contributions

The contributions of this dissertation are listed below.

- Our Depth First Search (DFS) algorithms are comparable in performance to the state-of-the-art *CDT algorithm* from Carpaneto *et al.* (1995) for the thoroughly studied ATSP. Our algorithms even achieve better performance for a variety of difficult ATSP instances. The improvements that enable this performance are summarized below.

- Upper bounds of BnB algorithms for the ATSP are improved in Chapter 2 with *iterative patching*, a procedure for constructing good feasible solutions of the ATSP. Iterative patching procedures reduce the number of subproblems in the usual BnB algorithm needs, so that smaller computation times are needed.

- Chapter 3 introduces tolerance-based branching rules and lower bounds. The tolerance-based branching rules divide the solution space in such a way that good or even optimal solutions are obtained relatively early in the search process. A BnB algorithm with tolerance-based lower bounds is able to remove much more subproblems than a BnB algorithm without such a lower bound. We also find a *synergy effect* between tolerance-based lower bounds and branching rules: when used simultaneously, they produce search tree reductions that are larger than the sum of the reductions from the individual improvements.

- When an algorithm departs from an initial solution, some elements need to be preserved, the *survival set*, whereas others need to be discarded, the *extinction set*. How can we predict whether an element belongs to the survival or to the extinction set? We present two measures for the quality of

the predictions, the *adjusted Rand index* and the *fit of the logistic regression model*. Both measures are used in Chapter 3 to compare the quality of the predictions of arc costs and arc tolerances for the ATSP. The results are consistent with the actual reductions of the tolerance-based and the cost-based branching rules. The predictions with logistic regression appear to be a little better than those with the adjusted Rand index.

- In Chapter 4, we compare our best upper tolerance-based BnB algorithm with a cost-based BnB algorithm for randomly generated instances of the ATSP. It appears that the tolerance-based algorithm is particularly effective when there is a large number of different intercity distances. It is shown in Zhang (1993) that these random instances are on average difficult to solve.

- We provide a lower bound that adds multiple upper tolerance values to lower bounds. Such improved lower bounds may increase the speed of BnB algorithms, and provide better estimations on the accuracy of given heuristic solutions.

- Steenkamp and Ter Hofstede (2002) observe that, in many segmentations in practice, logistics costs are so prohibitive that regular segmentations are not profitable. Instead, the segmentation strategy with countries or regions as segments is chosen. These segmentations are often ill-suited to consumer preferences, which tend to have a transnational and interregional character. The Budget Constraint Approach, introduced in Chapter 5, makes the trade-off between the fit and the logistics costs of segmentations possible. It paves the road for interregional or international segments with moderate logistics costs and reasonable costs.

- In current literature, the concept of logistics costs of a segmentation has not been thoroughly considered yet. In Chapter 5, a model of the logistics costs of segmentations is given. The underlying assumption behind this model is that each location in a segment is supplied from a central facility. The model can easily be extended to decentral facilities.

- Our experiments in Chapter 5 show that *simulated annealing* is an effective and reliable meta-heuristic for obtaining cluster solutions. It outper-

forms BnB.

6.4   Limitations and future research

We conclude this dissertation with limitations and future research.

- This dissertation concentrates on BnB algorithms for the ATSP. The question is whether the iterative patching-like procedures from Section 2 and the tolerance-based branching rules and lower bounds from Section 3 and 4 can also be used for other COPs. This is an interesting direction of future research.

- Research here is restricted to DFS BnB algorithms: algorithms that solve the most recently generated subproblem in the BnB process first. In practice, BFS BnB algorithms are also frequently used. An interesting direction of future research is to include tolerance computations in BFS algorithms.

- In Chapter 2, we limit the upper bounding heuristics to four known patching algorithms. Other types of heuristics may be used as well. In particular, the application of meta-heuristics to obtain upper bounds for BnB algorithms appears to be interesting.

- Chapter 3 focuses on upper tolerances. Roughly spoken, the upper tolerance of an element is the increase in the cost value of an element needed to remove it from an optimal solution. It is also possible to consider lower tolerance of an element, which is, roughly spoken, the minimum decrease in the cost value of an element needed to include it in an optimal solution. So, instead of *removing* elements from a given optimal solution, elements from outside the optimal solution are *included*. Since there are usually more elements outside the optimal solution than inside, it appears that lower tolerance computations are too time-consuming. However, it is shown in Volgenant (2006) that this need not be true, and in Germs (2006), lower tolerance-based BnB algorithms prove to be about as fast as their upper tolerance-based counterparts for the ATSP. So lower tolerance based algorithms constitute a very promising field of research.

- Upper tolerances are used here only in BnB algorithms. They can also be applied to heuristics.

- We have tested the additive lower bound from Chapter 4 only on Symmetric TSP instances. The bound is also appropriate for the Degree Constrained Minimum Spanning Tree Problem (DCMSTP). Computational experiments should be conducted to prove the effectivity of the additive bound for the DCMSTP.

- The settings of the simulated annealing algorithm from Chapter 5 are determined empirically, i.e., for a given set of instances, the parameter values are determined that leads to the highest solution quality. An interesting direction of research is to make a 'general' simulated annealing algorithm that automatically adjust its parameter settings to properties of the instances at hand, such as the size of the instance or the type of attribute data.