

A restricted dynamic programming algorithm for the dial-a-ride problem

Ulrike Ritzinger¹, Jakob Puchinger¹, Richard F. Hartl²

¹ Mobility Department, Austrian Institute of Technology
Giefinggasse 2, 1210 Vienna, Austria
ulrike.ritzinger.fl@ait.ac.at, jakob.puchinger@ait.ac.at

² Department of Business Administration, University of Vienna,
Brünner Straße 72, 1210 Vienna, Austria
richard.hartl@univie.ac.at

Abstract

In this paper, a restricted dynamic programming algorithm for the static multiple vehicle dial-a-ride problem is presented. Passengers have to be transported between pickup and delivery locations, while minimizing travel distances, respecting time window, user ride time and route duration constraints. We report preliminary results for benchmark instances which provide promising results. Algorithmic extensions and a hybrid metaheuristic are considered as pathways for future work.

1 Introduction

In the dial-a-ride problem (DARP), a number of transportation requests between given pickup and delivery locations have to be completed under user inconvenience considerations by a specified fleet of vehicles. This problem arises for example in the context of patient transportation where patients are delivered to medical facilities or carried back home. Since there exist many variants of this problem depending on the specific application there is no generic problem formulation [7]. Information on considered problem variants and existing solution methods can be found in Parragh et al. [6]. In this work, we consider the problem definition as well as the benchmark data set introduced by Cordeau and Laporte [1]. They take time windows at the pickup or delivery location, maximum route duration, and maximum user ride time into account. The transports are carried out by a homogeneous fleet with all vehicles based in the same depot. They propose a tabu search algorithm to solve this variant of the problem. In [7], improved results for related benchmark instances obtained by means of a variable neighborhood search are reported. Gromicho et al. [4] present a flexible framework for solving realistic VRPs. The solution approach within this framework is a generalization of the restricted dynamic programming heuristic for the Traveling Salesman Problem of Malandraki [5].

In this work, we present a restricted dynamic programming (DP) algorithm for the DARP variant as defined in [1]. First we implement an exact DP algorithm which is able to solve instances with up to 24 requests. This approach can not tackle larger instances due to memory constraints. We therefore extend it to a restricted DP heuristic which is able to solve instances with more requests in short computational time. In the restricted version, not all possible paths are explored, but rather a promising subset which is selected using a criterion function. Instead of using the standard minimum tour length selection criterion, we defined a criterion function that takes some global information into account which helps to select a good subset for further examination. With this novel way of selection, we obtain very promising results. In a further step, we want to apply heuristics to make better and forward looking decisions in state selection. In the next section, we will explain the details of our approach. We will then present some preliminary computational results before concluding and providing an outlook towards future work.

2 Solution Approach

Our first approach is an exact DP algorithm solving the static multiple vehicle DARP. In this variant of the problem, n transportation requests known in advance should be served by a fleet of m vehicles. The

aim is to construct routes accommodating all requests with minimum travel distances under consideration of the specific DARP constraints.

In the DP algorithm, a state encodes a corresponding path consisting of nodes representing all visited locations. It also contains the last location in the path, the earliest begin of service time at the last node and the user ride times of the patients on the vehicle. The cost of a state is defined by the travel distance of the path it is representing. The algorithm starts with the node representing the first vehicle at the depot, expanding it by all not yet visited nodes. In each successive stage, all states are expanded with all not yet attended nodes by recursion. In order to allow for multiple vehicles, this approach is extended by using a giant tour representation as introduced by Funke et al. [3], allowing to finish a tour at the depot and continue with another if there are still vehicles available. To ensure the feasibility of the solution, several constraints are considered at the expansion of the states: pickup and delivery node precedence, maximum user ride time, time window constraints, maximum route duration, and vehicle capacity. If two states with the same node set and the same last visited node are encountered, it is tried to discard the worse state. Dominance rules as in Dumas et al. [2] are introduced guaranteeing the optimality of the solution. In our case, one state dominates another if it has a smaller travel distance, an earlier begin of service time and smaller user ride times for all the patients aboard.

When solving larger DARP instances, the described approach will most probably open a huge number of states. Thus, due to memory and runtime constraints, the number of states has to be reduced. We therefore introduced several tightening constraints loosely based on the inequalities by Ropke [8]. Since the time window is set either on the pickup or delivery location, it can be tightened on the corresponding node by taking their distance and the maximum user ride time into account. The reachability of prospective states is also considered. Before expanding a state, it is checked whether there exist requests with time windows no longer feasibly reachable by the current vehicle. If there are no further vehicles available, the state will not be expanded because no feasible solution can be reached. If one or more patients are aboard and the expansion node is a pickup location, it will be examined if there is at least one feasible path to deliver all patients without violating their maximum user ride time constraints. If no such path is found, the pickup node is not regarded for extending the current state. Adding these tightening constraints to our DP approach allows us to solve small instances exactly.

In order to tackle larger instances, we extend the exact DP algorithm to a restricted DP heuristic [5]. The main concept is to bound the number of considered states by a given parameter B in every stage. For example, we could consider the B states with shortest travel distances, discarding the remaining states. This, however, might possibly prevent the algorithm from finding a feasible solution at all. Simply increasing parameter B will improve the results but also result in a significant increase of runtime and memory requirements. Hence, we introduce a value determining the most promising states for expansion, taking global information into account. The current tour length, the required travel distances to serve all remaining nodes using a separate vehicle for each request and the average remaining time for each request are combined for this measure. The best B states according to this value are selected and expanded in each stage. Even though the optimality guarantee is no longer given with this approach, we are able to find good solutions for a given benchmark data set and even the optima for the small instances in competitive runtime.

3 Preliminary Results

We have tested our algorithms on the benchmark data sets introduced by Cordeau [1] and Ropke [8]. The sets contain randomly generated Euclidean DARP instances with up to 96 requests (24 instances per set). One set ('a' instances) uses vehicles with a capacity of three patients and a maximum user ride time of 30 minutes while the other set ('b' instances) allows up to six patients per vehicle and the maximum user ride time is set to 45 minutes. In Ropke [8], a branch-and-cut algorithm is introduced which solves all instances to optimality.

For testing our first approaches we mainly focused on the 'a' instances. With the exact DP algorithm we manage to solve instances with up to 24 requests. Due to memory restrictions, we can not solve larger

instances with the exact approach. With the restricted DP algorithm, we obtain feasible solutions for all instances. Optimal solutions have been found for eight instances and for six out of those eight instances our algorithm used significantly less time than the branch-and-cut algorithm. These results suggest that our approach can be incorporated into a metaheuristic to solve small subproblems where the branch-and-cut algorithm is too slow. For the remaining instances, our solutions are on average only 2.61% behind the optimum. However, if we compare our results to the ones obtained by the state of the art tabu search algorithm by Cordeau and Laporte [1], we get almost always better solutions (except for 3 instances) in less or comparable runtime. For the instances with high vehicle capacities and longer maximum user ride times, which are more challenging for the DP approach, the solutions we found are never more than 5.5% worse than the optimum.

4 Future Work

Because of the results we obtained so far, we plan to further improve the restricted DP heuristic with the aim of being able to solve larger instances to optimality in feasible runtime. We plan to achieve this goal by putting more effort into the selection of promising states, thus, considerably decreasing the number of explored states. Therefore, we intend to test heuristics helping to make an adequate selection of the most promising states. Due to the fact that the DP algorithm works well for small instances, we envisage to incorporate it into a metaheuristic approach for searching large neighborhoods. Finally, the restricted DP algorithm can be used to obtain a feasible solution, that can further be improved via a metaheuristic such as Variable Neighborhood Search.

Acknowledgements

This work, partially funded by the Austrian Federal Ministry for Transport, Innovation and Technology (BMVIT) within the strategic program FIT-IT ModSim under grant 822739 (project HealthLog).

References

- [1] Jean-Francois Cordeau and Gilbert Laporte. A tabu search heuristic for the static multi-vehicle dial-a-ride problem. *Transportation Research Part B: Methodological*, 37(6):579 – 594, 2003.
- [2] Yvan Dumas, Jacques Desrosiers, Eric Gelinas, and Marius M. Solomon. An optimal algorithm for the traveling salesman problem with time windows. *OPERATIONS RESEARCH*, 43(2):367–371, 1995.
- [3] Birger Funke, Tore Grünert, and Stefan Irnich. Local search for vehicle routing and scheduling problems: Review and conceptual integration. *Journal of Heuristics*, 11:267–306, July 2005.
- [4] J. Gromicho, J.J. Hoorn van, A.L. Kok, and J.M.J. Schutten. Restricted dynamic programming: a flexible framework for solving realistic VRPs, 2009.
- [5] Chryssi Malandraki and Robert B. Dial. A restricted dynamic programming heuristic algorithm for the time dependent traveling salesman problem. *European Journal of Operational Research*, 90(1):45 – 55, 1996.
- [6] Sophie N. Parragh, Karl F. Doerner, and Richard F. Hartl. A survey on pickup and delivery problems. Part II. *Journal für Betriebswirtschaft*, 58:81–117, 2008.
- [7] Sophie N. Parragh, Karl F. Doerner, and Richard F. Hartl. Variable neighborhood search for the dial-a-ride problem. *Computers and Operations Research*, 37:1129–1138, June 2010.
- [8] Stefan Ropke, Jean-Francois Cordeau, and Gilbert Laporte. Models and branch-and-cut algorithms for pickup and delivery problems with time windows. *Networks*, 49(4):258–272, 2007.