Patient Transportation - Dynamic Dial-a-Ride and Emergency Transportation Problems

Guenter Kiechle[(1,2)], Karl F. Doerner[(3)],
Michel Gendreau[(4)], Richard F. Hartl[(3)]

(1) Technische Universitaet Wien, Karlsplatz 13, 1040 Vienna, Austria
(2) Salzburg Research Forschungsgesellschaft
Jakob Haringer Strasse 5/III, 5020 Salzburg
guenter.kiechle@salzburgresearch.at
(3) Department of Management Science, University of Vienna
Bruenner Strasse 72, 1210 Vienna, Austria
{karl.doerner, richard.hartl}@univie.ac.at
(4) Centre interuniversitaire de recherche sur les réseaux d'entreprise, la logistique et le transport (CIRRELT), Université de Montréal
C.P. 6128, succursale Centre-ville, Montréal, Canada H3C 3J7
michelg@crt.umontreal.ca

Abstract

In Austria, the transportation of patients to and from hospitals is organized by non-profit organizations. In most regions, it is the Austrian Red Cross that is responsible for the transportation of patients. The regular patient transportation, as well as the emergency transportation, are performed by scheduling the same fleet within the same control center. Therefore a dial-a-ride problem for the regular patient transportation orders with disruption has to be solved. Disruptions occur due to the fact that vehicles "disappear" in order to serve emergency requests and reappear at a hospital after the service of an emergency case. This problem is highly dynamic and a robust plan has to be computed for the regular dial-a-ride orders in order to serve also the emergency requests. When a disruption occurs because of an emergency, the fleet size is reduced and the remaining patient transport orders have to be carried out with this reduced fleet size. After the execution of an emergency transport, the available vehicle can be integrated to serve the regular dial-a-ride orders. We analyze different scheduling strategies for the regular dial-a-ride orders in order to minimize routing costs and to minimize the response time for servicing an emergency request.

1 Introduction

Many emergency service providers, especially ambulance departments and companies who provide non-public maintenance services, face the problem that their fleet of vehicles has to provide two different types of services:

1. Cover a certain region and provide immediate service when an emergency occurs;

2. Provide some regular service (e.g., the pick-up and delivery of patients, predetermined service tasks, periodic pick-ups . . .).

This is the current situation for the largest Austrian emergency service providers (e.g., the Austrian Red Cross), where the same fleet is used to provide both types of services. Dynamic aspects thus directly influence the schedule for the regular service. When an emergency occurs and an ambulance is required, the vehicle with the shortest distance to the emergency is assigned to serve the emergency patient. Therefore, it may happen that an ambulance vehicle that has to carry out a scheduled transport order of a patient, which has not started yet, is used to serve the emergency request and the schedule for the regular services has to be re-optimized and another vehicle has to be reassigned to the regular patient. Ambulances that carry out emergency transport become available at the hospital after the emergency service and can be then used to carry out regular transport orders. Again, the schedule for regular services has to be re-optimized.

From the perspective of managing the regular services, the objective is minimizing the total travelling distance subject to certain restrictions (e.g., be on time, use the appropriate vehicle, . . .). From the perspective of minimizing the response time for servicing an emergency request, it is necessary to locate and schedule the vehicles in such a way that each possible location where an emergency case may occur can be reached within a given time (see [5]).

These two objectives are not totally contradictory, but they certainly conflict: on the one hand, for the emergency transport requests, one has to account for the fact that vehicles are changing positions and are blocked for some time due to some regular service assignment; on the other hand, when regular orders are assigned to vehicles, it is important to keep a certain coverage level to ensure a satisfactory service of the emergency cases, which may require relocating some of the vehicles.

Emergency service providers thus want to find a robust solution for a specific day of the week in order to minimize routing costs, as well as the response time for servicing emergency patients, taking into account that the two types of the transportation services have to be realized with a single fleet.

Some related work has been published where pickup and delivery request occur dynamically (see [1, 4, 6, 7, 9]). But to our best knowledge, a dynamic changing fleet size and this type of disruption have not been considered so far. The algorithm described in this paper is implemented within the new decision support system for the Austrian Red Cross - AmbulanceRoute [8].

2 Problem Description

In our study we consider a combination of two problems

- the dial-a-ride problem for regular patient transportation and
- dispatching ambulance vehicles for emergency cases.

In the classical DARP, a set of requests announced beforehand are served from a single depot with a heterogeneous fleet of vehicles, i.e., different vehicle configurations. These requests have hard time windows and a preferred pickup or delivery point in time, as well as different space requirements. In the problem of dispatching ambulance vehicles for emergency cases, one must ensure short response times in a dynamic environment.

The regular patient transportation problem can be considered to be a variation of DARP with additional real world constraints regarding customer preferences or requirements. A comprehensive description of the DARP is given in [2] and is repeated here. The DARP consists of designing vehicle routes and schedules for n customers or patients who specify pick-up and drop-off requests between origins and destinations. A typical situation is that the same patient will have two requests during the same day or within a certain period - an outbound request, usually from home to the hospital, and an inbound request for the return trip. In the standard case discussed in the literature, a homogeneous fleet is considered. The objective is to plan a set of minimum-cost vehicle routes while serving as many customers as possible under a set of constraints. The main difference between the DARP and most classical routing problems is the fact that in the DARP human beings are transported instead of goods as in the other problems. Thus, additional constraints, e. g., maximum ride times for the patients, no waiting times with a patient on board, preferred pick-up and drop off times are considered [2].

Our real world DARP is extended in the way that there may be several depots located in different regions of the country (sometimes located close to the hospitals) and the fleet is heterogeneous. Some vehicles can carry two or more patients, other can carry one passenger on wheelchairs or two passengers without wheelchairs. Some vehicles can carry one passenger when he/she needs a stretcher or two passengers who can sit. Some vehicles have an additional attendant onboard and therefore less room for the patients. Regarding time windows, we have two different situations - on the one hand, patients should be picked-up as late as possible from their home when they are being transported to hospitals; on the other hand, patients should be picked-up as early as possible when they are transported from the hospital back home. Deviations from the desired pick-up and drop-off times within the specified time window are considered in the objective function. There is a fixed number of vehicles in operation, as well as some extra vehicles, which can be used but at higher cost. The aim is to maximize the number of requests that can be served with a the fixed size fleet. Demands which cannot be served with the fixed fleet size are served using extra vehicles (e.g., taxi cabs), if necessary. In our real world problem, it is necessary to minimize transportation costs subject to full demand satisfaction and side constraints, i.e., all patients must be transported to and from the hospital. The relevant cost elements are the costs concerning the regular fleet size and operations, the occasional use of extra vehicles, and driver wages.

Quality of service for the regular transport orders includes route duration, route length, customer waiting times, relative customer ride time (i.e., total time spent in vehicles compared with minimum time), and difference between actual and desired pick-up and drop-off times. Some of these criteria are treated as constraints or as part of the objective function. Users are able to specify a time window on the arrival time of their outbound trip and on their departure time on their inbound trip [2]. The transporter then determines a planned departure time for the outbound trip and a planned arrival time for the inbound trip, while satisfying an upper bound on ride time. The maximum ride time depends on the distance to the drop off location, i.e., we discriminate the maximum ride time of the customer depending on the customer location. More precisely, the maximum ride time of a customer is the sum of the shortest time to the drop-off location plus a certain possible waiting time. Furthermore, the sum of the waiting time of a customer plus the ride time that exceeds the shortest time to the drop-off location is limited by a maximum waiting time parameter.

In our specific problem, we also have to consider and evaluate the execution of emergency transport orders with the same fleet. In order to get a satisfactory quality of service for the emergency cases, we optimize the spatial distribution of empty vehicles in order to minimize response time to emergency cases. The distribution is measured at discrete points in time when a vehicle is empty, e.g., at equal intervals from the first pick-up to the last drop-off during the day or optimization period. We compute the coverage of these empty vehicles and in the optimization we maximize the average coverage and minimize the root mean square deviation from the average coverage.

The overall objective in our problem is to thus optimize three criteria. In the objective function transportation costs for the regular patient transport, quality of service for the regular transport orders and quality of service for the emergency cases, which is measured as coverage criteria. The goal of the project is to develop a solution procedure which generates good compromise solutions to get reasonable costs with a good quality of service, especially for the emergency cases.

3 Solution Method

In order to study different dispatch strategies for the combined problem, we developed a simple and effective solution procedure. Hence, we implemented a constructive heuristic approach with improvement steps based on move and swap operators with inter- as well as intra-tour moves. In the construction phase, we exploit the temporal structure of the requests and use a nearest neighbor measure to construct an initial solution to which different local descents are applied afterwards.

The main challenge is to deal with the dynamic nature of the problem, which implies that the solution technique has to be able to generate new solutions in very short time. Every time a new emergency request occurs, the distance information for the empty vehicles has to be updated in order to identify very quickly the vehicle which will serve this emergency request. Then, one has to resolve the remaining problem with one less vehicle available to serve the regular orders. When an emergency order is fulfilled and the patient is transported to the hospital, an additional vehicle is made available and it can be used to fulfill the regular patient transportation orders. When this situation occurs, the schedule for the regular transport orders has also to be re-optimized to take advantage of this additional vehicle to improve the objective function.

The only way to properly evaluate different solutions and strategies is through simulation of the situation at hand. In the simulation environment that was implemented, vehicle movements and the spatial distribution of empty vehicles over time are carefully tracked over time. The distance matrix is also updated in real-time whenever emergency requests occur.

4 Evaluation of the Approach

In our computational study, we address several questions of importance for the Austrian Red Cross and different strategies are implemented and evaluated. The main questions considered in the study, and to which answers are provided, are the following:

1. Is it beneficial to use the same fleet of vehicles for the emergency cases and for the regular transport orders?

2. What is the advantage of using one fleet for the two problems and which difficulties arise when doing so?

3. When using a single fleet, what kind of buffer is needed to obtain a certain robustness in the solution, i.e., not to disrupt the solution completely for the regular patient transportation orders during the day?

Different waiting strategies are analyzed, e.g., should the empty vehicles drive around instead of waiting at a medical property? We also analyze different dispatch strategies with respect to the characteristics of the regular patient transport orders and the degree of dynamism for the emergency cases.

The computational study is based on artificial test instances derived from the instances of Cordeau [3], as well as some real world instances from the city of Graz in Styria (Austria). The artificial data are adapted and enriched with emergency cases; furthermore, for each inbound (resp. outbound) request an additional outbound (resp. inbound) request is generated.

Acknowledgements

Financial support from the Fonds für wissenschaftliche Forschung (FWF) under grant #L286-N04 is gratefully acknowledged.

References

[1] Attanasio, A., Cordeau, J.-F., Ghiani, G., and Laporte, G. (2004): "Parallel Tabu search heuristics for the dynamic multi-vehicle dial-a-ride problem". *Parallel Computing* **30**, pp. 377–387.

[2] Cordeau, J.-F., and Laporte, G. (2003): "The Dial-a-Ride Problem (DARP): Variants modelling issues and algorithms". *4OR* **1**, pp. 89–101.

[3] Cordeau, J.-F., and Laporte, G. (2003): "A tabu search heuristic for the static multi-vehicle dial-a-ride problem". *Transportation Research B* **37**, pp. 579–594.

[4] Gendreau, M., Guertin, F., Potvin, J.-Y., and Séguin, R. (2002): "Neighborhood Search heuristic for a Dynamic Vehicle Dispatching Problem with Pickups and Deliveries". Technical Report, Centre de recherche sur les transports, Université de Montréal. Forthcoming in *Transportation Research C.*

[5] Gendreau M., Laporte G., and Semet F. (2001) "A dynamic model and parallel tabu search heuristic for real-time ambulance relocation". *Parallel Computing* **27,** 1641–1653.

[6] Mitrović-Minić, S., and Laporte, G. (2004): "Waiting Strategies for the Dynamic Pickup and Delivery Problem with Time Windows". *Transportation Research B* **38**, pp. 635-655.

[7] Mitrović-Minić, S., Krishnamurti, R., and Laporte, G. (2004): "Double-horizon based heuristics for the dynamic pickup and delivery problem with time windows". *Transportation Research B* **38**, pp. 669-685.

[8] Niemeck, M., Biffl, S., Doerner, K. F., and Hartl, R. F. (2007): "Ambulance Route - The Decision Support System for Patient Transportation". Preprints of TRISTAN VI - in case the abstract is accepted.

[9] Savelsbergh, M. W. P., and Sol, M.(1998): "DRIVE : Dynamic Routing of Independent Vehicles". *Operations Research* **46** (4), pp. 474–490.