Computers & Operations Research 56 (2015) 120–133

Contents lists available at ScienceDirect

Computers & Operations Research

journal homepage: www.elsevier.com/locate/caor

Ambulance routing for disaster response with patient groups

Luca Talarico [a,*], Frank Meisel [b], Kenneth Sörensen [a]

[a] University of Antwerp, Dept. ENM, Faculty of Applied Economics, Antwerp, Belgium
[b] Christian-Albrechts-University Kiel, Faculty of Business, Economics and Social Sciences, Kiel, Germany

ARTICLE INFO

Available online 22 November 2014

Keywords:
Ambulance routing
Disaster response
Service time
Local search
Large neighborhood search

ABSTRACT

We consider a routing problem for ambulances in a disaster response scenario, in which a large number of injured people require medical aid at the same time. The ambulances are used to carry medical personnel and patients. We distinguish two groups of patients: slightly injured people who can be assisted directly in the field, and seriously injured people who have to be brought to hospitals. Since ambulances represent a scarce resource in disaster situations, their efficient usage is of the utmost importance. Two mathematical formulations are proposed to obtain route plans that minimize the latest service completion time among the people waiting for help. Since disaster response calls for high-quality solutions within seconds, we also propose a large neighborhood search metaheuristic. This solution approach can be applied at high frequency to cope with the dynamics and uncertainties in a disaster situation. Our experiments show that the metaheuristic produces high quality solutions for a large number of test instances within very short response time. Hence, it fulfills the criteria for applicability in a disaster situation. Within the experiments, we also analyzed the effect of various structural parameters of a problem, like the number of ambulances, hospitals, and the type of patients, on both running time of the heuristic and quality of the solutions. This information can additionally be used to determine the required fleet size and hospital capacities in a disaster situation.

1. Introduction

Recent examples such as Hurricane Katrina in 2005, the Indian Ocean tsunami in 2004, or any of the recent armed conflicts around the globe demonstrate that disasters can have a devastating impact on a society. Regardless of whether their cause is natural (e.g., earthquakes, floods, hurricanes, wildfires) or man-made (e.g., terrorist attacks, war situations), disasters can cause large-scale loss of life as well as damage to a society's infrastructure, housing, and industrial complex. It has been widely recognized (see e.g., [3,19]) that the severity of a disaster can be, to a large extent, influenced by the efficacy of the logistics operations during the response phase. Although the disaster *itself* can certainly cause a lot of casualties, a large fraction of the victims usually perish because of a lack of medical aid in the immediate aftermath of a disaster. Clearly, the post-disaster situation results in the response actions having to be executed under extremely challenging conditions: limited availability of resources (transportation, supplies, manpower, hospital capacity), damaged transportation and communication infrastructure, as well as uncertain information regarding the number and locations of people in need of medical assistance, see e.g., Najafi et al. [27,28] and Yi et al. [43]. Despite these challenges, it is essential that the logistics relief operations are initiated quickly and well planned to be most effective. Hence, there is a strong need for decision support tools that generate solutions to the underlying optimization problems in a few seconds or less [3]. However, research on transportation problems and vehicle fleet management for disaster response operations is emerging only recently, see de la Torre et al. [41] and Pedraza-Martinez and van Wassenhove [33]. With this paper, we propose a decision support approach for the routing of ambulances in response to a disaster.

The central task of managing ambulances in a disaster response situation is to provide first aid to slightly injured people and to bring seriously injured people to operating hospitals. Managing the operations of ambulances in the immediate aftermath of a disaster is massively complicated by the dynamics and uncertainty with which the planning conditions (especially the relevant information) change over the course of time. The information required to support the planning of ambulances includes the number and location of people calling for help, the availability of ambulances, the capacity of the nearby hospitals, as well as the accessibility of incident sites due to the damaged infrastructure and the current traffic situation, see Jotshi et al. [22]. Another issue is that, in

* Corresponding author. Address: ANT/OR – University of Antwerp Operations Research Group, Prinsstraat 13, 2000 Antwerp, Belgium. Tel.: +32 3 265 41 77.
E-mail addresses: luca.talarico@uantwerpen.be (L. Talarico), meisel@bwl.uni-kiel.de (F. Meisel), kenneth.sorensen@uantwerpen.be (K. Sörensen).

http://dx.doi.org/10.1016/j.cor.2014.11.006
0305-0548/

contrast to the daily operations in the public health care sector, the number of requests for help in a disaster situation strongly exceed the capacity of the available ambulance fleet. Hence, it is of utmost importance to use the ambulances efficiently in such a way that they provide as much medical aid as possible.

The response process that is executed by the responsible organizations in the aftermath of a disaster has to be designed in such a way that it is able to cope with the challenges of a dynamic planning situation for the scarce ambulance resource. The routing of ambulances in such a situation can be treated as a static or a dynamic routing problem. In the static case, a set of emergencies requests is collected first and, then, the routing problem is solved for this set of requests. In the dynamic case, the routes of ambulances are updated whenever new help requests arrive, which can reduce the response time. However, this approach requires that communication with the ambulances is possible at all time, which might not be the case in a disaster situation and, furthermore, the rescue teams may perceive this to be disturbing under stressful circumstances. Therefore, in this paper, we consider a three-step response process that aims at solving a static ambulance routing problem, see Fig. 1. The process is executed by a central dispatching unit, which collects requests and manages ambulance operations repeatedly until no further emergency requests are received. The first step is to answer incoming emergency calls and to collect relevant information like the location and the condition of the people being in need of help. The dispatcher collects several requests that are then classified according to their severity in a second process step. The classification reflects the priority with which a patient should receive help, which is taken into account when routing the ambulances in the third process step. Collecting and classifying a number of requests before actually sending out the ambulances supports an efficient use of the vehicles, because instead of dispatching ambulances on a first-come first-served basis they can be used to serve the most urgent requests first. Therefore, the first two process steps do not represent a waste of precious time but they collect valuable information to come up with high-quality route plans in the third process step. In fact, the time spent for the first two steps is rather short if numerous requests arrive within short time (as in the case of a disaster event) and if the classification of requests is performed directly while answering an emergency call or automatically from the collected data. Hence, the three-step process can be repeated at high frequency (for example each time a certain number of requests has been collected or a certain time limit has elapsed) such that it causes little delay in the service process. Clearly, if the dispatcher classifies an incoming request as so urgent that it cannot wait at all, a suitable ambulance may be deployed directly without waiting for further requests. This, however, would constitute a mixed static-dynamic response process, which is out of scope of this paper. A further advantage of the sketched three-step process is that up-to-date information regarding the availability of ambulances, infrastructure conditions, etc., can be included in the planning.

The scope of this paper is to investigate the routing problem that occurs in the third step of the response process. The ambulances are used to bring medical personnel to the casualties and to carry injured people to the hospitals. Each ambulance carries medical personnel that can provide first aid to slightly injured people in the field. Seriously injured individuals are accompanied by the medical staff on their way to the hospital where skilled doctors are available. According to this, we distinguish two types of patients:

- *Red code patient*: A person with red code classification is seriously injured and needs to be brought to a hospital by an ambulance.
- *Green code patient*: A person with green code classification is slightly injured and can be helped directly in the field.

There exist more detailed classification schemes for patients (see e.g., [1,13]) and several so-called *triage* systems have been developed for classifying and prioritizing patients rapidly in a mass-casualty incident with an overwhelming number of victims, limited time and scarce medical resources, see Killeen et al. [24]. The goal of triage is to allocate a limited set of medical resources to patients such that these resources are used as efficiently as possible, providing the best possible care to a large number of patients. The triage system therefore assigns priority to those patients who will substantially benefit from a rapid intervention, even if these patients are not the most critical ones. This makes disaster response different from civil health care where resources are usually not scarce and the most severe patients always receive highest priority. Typical triage systems classify and prioritizes patients based on their conditions into four groups [11]: patients who require immediate transportation to a hospital, patients who can wait some time for transportation, patients who require no hospital treatment, and patients who are unlikely to survive at all. Including further categories (as done for example in [13]) allows for a finer distinction of patients and their needs but makes the application of triage systems more difficult. However, the two types of patients considered in this paper are sufficient to distinguish the fundamental tasks that have to be performed by the ambulances, namely serving patients in the field and bringing them to hospitals. For this reason, we just consider two patient classes in this paper.

Concerning the routing of ambulances, we assume that each of them can carry one red code patient at a time and that each patient is directly brought to a hospital after having been picked up. The decision to which hospital to bring a patient is part of the routing problem and depends on the capacities of hospitals. Since green code patients can be helped on the field, an ambulance can go directly to the next patient after having served a green code patient. From this, an ambulance can provide help to multiple people on its route before returning to a hospital. In contrast, if an ambulance has to serve multiple red code patients, it has to visit several hospitals throughout the planning horizon. Therefore, for the purpose of clarity, we refer here to a route as a tour that begins at one hospital, visits one or more patients in a specified sequence, and ends at either the starting hospital or at some other hospital. Hence, an ambulance may perform multiple routes within a solution to the ambulance routing problem.

The optimization problem is then to determine ambulance routes to serve the two groups of patients, red code and green code patients, which have been determined in the first two steps of the sketched response process. The objective is to minimize the sum of the latest service completion time among the red code patients and the latest service completion time among the green code patients. The objective strives to reduce the longest waiting time faced by a patient in a group. Although some authors propose to minimize the average waiting time of patients (e.g., [8]), the objective pursued in our paper ensures that no patient suffers from

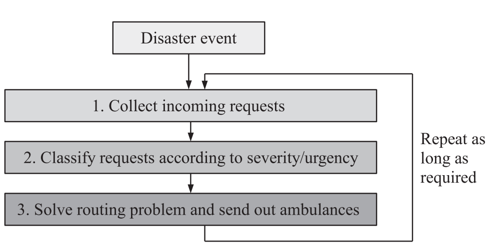

Fig. 1. Disaster response process.