



An efficient heuristic for real-time ambulance redeployment



C.J. Jagtenberg^{a,*}, S. Bhulai^{b,a}, R.D. van der Mei^{a,b}

^a CWI, Stochastics, Science Park 123, 1098 XG Amsterdam, The Netherlands

^b VU University Amsterdam, Faculty of Sciences, De Boelelaan 1081a, 1081 HV Amsterdam, The Netherlands

ARTICLE INFO

Article history:

Received 1 May 2014

Accepted 16 January 2015

Available online 28 January 2015

Keywords:

Ambulances

Emergency medical services

Relocation

Redeployment

ABSTRACT

We address the problem of dynamic ambulance repositioning, in which the goal is to minimize the expected fraction of late arrivals. The decisions on how to redeploy the vehicles have to be made in real time, and may take into account the status of all other vehicles and accidents. This is generally considered a difficult problem, especially in urban areas, and exact solution methods quickly become intractable when the number of vehicles grows. Therefore, there is a need for a scalable algorithm that performs well in practice.

We propose a polynomial-time heuristic that distinguishes itself by requiring neither assumptions on the region nor extensive state information. We evaluate its performance in a simulation model of emergency medical services (EMS) operations. We compare the performance of our repositioning method to so-called static solutions: a classical scenario in which an idle vehicle is always sent to its predefined base location. We show that the heuristic performs better than the optimal static solution for a tractable problem instance. Moreover, we perform a realistic urban case study in which we show that the performance of our heuristic is a 16.8% relative improvement on a benchmark static solution. The studied problem instances show that our algorithm fulfills the need for real-time, simple redeployment policies that significantly outperform static policies.

© 2015 Elsevier Ltd. All rights reserved.

1. Introduction

In a world where medical resources and budgets are limited, emergency medical services (EMS) managers are forced to rethink the way they spend both. Medical decisions aside, mathematical models can help them obtain more efficiency. They can also be helpful in understanding the effects of a certain decision (e.g., adding one extra vehicle, or changing the dispatch policy), which is otherwise difficult to oversee due to the stochastic nature of accidents. Typically, geographical aspects and service level agreements need to be taken into account when solving such problems.

In an EMS system, accidents occur randomly throughout the region.¹ Each accident needs to be served as soon as possible by an ambulance. The number of vehicles is typically limited, and vehicles are not always available due to serving other accidents. If an ambulance is not busy serving an accident, it is either on the

road (driving), or stationed at one of the selected base locations. (Note that an idle ambulance can respond to an accident while still on the road, there is no need to return to a base location first.) Since minimizing the response time is critical in emergency situations, it is important to place ambulances in good positions with respect to the demand. This leads to the search for good base locations, as well as a good distribution of vehicles over the bases.

1.1. Related work

In ambulance planning, models often use graph representations. Accidents can occur at the nodes, and there are certain distances (or driving times) between nodes. The travel times are assumed to be known in advance, and may be deterministic or stochastic (in which case they are only known in distribution). The goal is usually to maximize the fraction of accidents served within a certain (pre-determined) time. There are articles that search for the number of vehicles needed [1], the best base locations [2], and/or the best distribution of vehicles over the bases [3].

Static models

Mathematical models can be used at various stages of the EMS process. First, consider the *planning stage*. At this point, static

* Corresponding author.

¹ Throughout this paper, we will use 'accidents' to refer to demand for ambulances. Accidents include medical incidents and are not limited to traffic collisions.

models are often used to describe the problem. Here ‘static’ means that each ambulance is sent to its own home base whenever it becomes idle. These models can be used to determine the optimal locations of bases, as well as the number of vehicles needed per base.

Early research in ambulance planning focused on deterministic location problems [2,4]. These formulations ignore the stochastic aspects of an EMS system, typically by assuming that one vehicle, or a constant number of vehicles, is always sufficient to cover the demand points. Later, research turned to probabilistic static models. A well-known example is the maximum expected covering location problem formulation (MEXCLP) [3]. In this formulation there is a limited number of vehicles that need to be distributed over a set of possible base locations. Each vehicle is modeled to be unavailable with a pre-determined probability. For a more detailed description of this model, we refer the reader to Section 2.

Over the years, several variants of MEXCLP have been published by different authors [5,6]. These models are generally considered to give good static solutions. (Note a static solution can be defined by giving the location of the ‘home base’ for each vehicle.) The downside of static policies is that they do not utilize all possibilities, e.g., real-time information, to obtain good coverage. Clearly, the assumption of a vehicle belonging to a specific base is unnecessary in real life. Using the models above, one can attempt to find the optimal policy within the solution space of static policies. However, in the space of all policies, this will almost always be suboptimal.

Dynamic models

Dynamic models are used to find good (re)distributions of vehicles when a number of ambulances are busy responding to accidents. In other words, they look for repositioning strategies, which stand in contrast to strategies where every ambulance is sent back to its ‘home base’ after serving an accident. The first of such models can be found in [7], using tabu search. This shifted focus of research was accompanied with an increasing number of EMS systems using a dynamic allocation of vehicles to bases. Surveys of North American EMS operators showed that the percentage of operators who used a dynamic strategy increased from 23% in 2001 [8] to 37% in 2009 [9] (see also [10]). This indicates that the EMS community is becoming more aware that a dynamic policy can help them achieve greater service without increasing capacity.

Dynamic models usually do not search for good base locations, but instead consider the bases as a given, fixed set. The redeployment policies that have been published so far are roughly dividable in two subclasses, which we will (very generalizing) refer to as lookup tables and real-time optimization.

Lookup tables. The models in this class are typically looking for an optimal configuration for each number of available ambulances. A recent example can be found in [10]. The job of steering the set of available vehicles towards this configuration is usually left to the dispatchers. Unfortunately, poorly executed redeployment can devalue even the most crafty policy. Even if the decision of how to move the vehicles in order to obtain the required configuration is part of the mathematical solution, this approach altogether requires a lot of ambulance movements. This increases the work load on the ambulance crew, which is a downside in many realistic situations. Furthermore, note that in busy regions, where the number of idle ambulances changes rapidly, the system will not be in compliance with the lookup table for most of the time.

Real-time optimization. On the other hand, there are various papers that model the randomness in the system explicitly, for example, by formulating the problem as a Markov decision process. When the model has only a few ambulances, one can solve it using exact dynamic programming [11].

When the state space grows, for example due to the number of vehicles considered, the problem quickly becomes intractable. In those cases we need to turn to alternative solution methods. Successful approaches include approximate dynamic programming [12]. Here, the state space is modeled rather elaborately, and the authors need advanced mathematical methods to solve the problem. Furthermore, it requires a mechanism to tune parameters to the use case, which is time consuming to both implement and execute. For one large city, the tuning process can take as long as one year. It remains possible to calculate the repositioning decision in real time, because these heavy computations are done in a preparatory phase. Furthermore, the authors try to speed up the tuning process, for example by using the so called post decision state. (For an elaborate discussion of the post decision state, see [13].) For the use case of the city of Melbourne described in [14], this reduced the computation time from approximately one year to 12 h. Although this demonstrates the power of the post decision state, the remaining 12 h should also highlight the complexity of the method. The heavy pre-computations and the need for an expert to implement this, make this method inaccessible and impractical.

Furthermore, the performance of the approximate dynamic programming approach is highly dependent on the choice of base functions. The base functions as defined in [14] are elegant, but unlikely to work well in general. That is because the underlying idea used is the following: an accident is likely to be served late if there are no idle vehicles present at the nearest base. For many EMS regions, for example in the Netherlands, this is typically far from the truth. Moreover, the policies should work well for densely populated areas, the more difficult case of ambulance planning, where some demand points can be reached within the time threshold from as many as 8 different base locations. This complexifies the construction of good base functions.

1.2. Our contribution

In practice, ambulance planners face a number of challenges. Usually only limited and coarse-grained information about the state of the system is available for decision making, while the accuracy of the computations should be good, and at the same the computation times should not be prohibitively large. Motivated by this, the goal of this paper is to propose an algorithm that is efficient yet easy-to-use, thereby properly balancing the trade-off between simplicity, accuracy and scalability. Thereby, we ensure that even EMS providers with few tools available to track real-time information, can implement this solution. We focus on busy, urban areas. In such a setting it is unacceptable to move every vehicle each time an accident occurs. And although some proactive relocations may be useful, they clearly enlarge the workload for the crew. We choose to limit our repositioning opportunities in the following way. An ambulance is only allowed to relocate when it becomes idle (which can be at the incident scene, or at a hospital). Thereby, the number of trips will be the same as for a static strategy, which will help convince EMS managers that our proposed solution is a good alternative to a static strategy.

The ambulance redeployment algorithm we develop in this paper, is both intuitively clear and computable in real time. The solution does not require a preparatory learning phase and is easy to implement. Furthermore, the algorithm requires very little real time data, in fact, only the destinations (locations) of the available vehicles are used. We then decide where to send the available vehicle, based on an expression for marginal coverage improvement. Marginal coverage is an idea that originated in static ambulance planning [3], but this paper shows that it can be useful in dynamic ambulance planning as well. Through this notion of coverage we aim to reduce our KPI: the expected fraction of late

Download English Version:

<https://daneshyari.com/en/article/1141944>

Download Persian Version:

<https://daneshyari.com/article/1141944>

[Daneshyari.com](https://daneshyari.com)