Chapter 9
Optimal Ambulance Dispatching

C.J. Jagtenberg, S. Bhulai and R.D. van der Mei

Abstract This chapter considers the ambulance dispatch problem, in which one must decide which ambulance to send to an incident in real time. In practice as well as in literature, it is commonly believed that the closest idle ambulance is the best choice. This chapter describes alternatives to the classical closest idle ambulance rule. Our first method is based on a Markov decision problem (MDP), which constitutes the first known MDP model for ambulance dispatching. Moreover, in the broader field of dynamic ambulance management, this is the first MDP that captures more than just the number of idle vehicles, while remaining computationally tractable for reasonably-sized ambulance fleets. We analyze the policy obtained from this MDP, and transform it to a heuristic for ambulance dispatching that can handle the real-time situation more accurately than our MDP states can describe. We evaluate our policies by simulating a realistic emergency medical services region in the Netherlands. For this region, we show that our heuristic reduces the fraction of late arrivals by 13% compared to the "closest idle" benchmark policy. This result sheds new light on the popular belief that deviating from the closest idle dispatch policy cannot greatly improve the objective.

C.J. Jagtenberg (✉)
Stochastics, CWI, Science Park 123, 1098 XG Amsterdam, The Netherlands
e-mail: jagtenbe@cwi.nl

S. Bhulai
Faculty of Sciences, VU University Amsterdam, De Boelelaan 1081a, 1081 HV Amsterdam, The Netherlands
e-mail: s.bhulai@vu.nl

R.D. van der Mei
Stochastics, CWI, Science Park 123, 1098 XG Amsterdam, The Netherlands
e-mail: mei@cwi.nl

R.J. Boucherie, N.M. van Dijk (eds.), *Markov Decision Processes in Practice*,
International Series in Operations Research & Management Science 248,
DOI 10.1007/978-3-319-47766-4_9

9.1 Introduction

Emergency Medical Services (EMS) providers serve emergency calls while aiming to keep response times short. In particular, one issue that plays a central role is to maximize the fraction of incidents that are reached within a certain target time. Operations research and mathematical modeling can help to reach this goal.

9.1.1 Previous Work

A large number of models are available for ambulance planning. On one hand, there are models that deal with planning on a strategic level. Typically, such models determine the best locations for ambulance bases [6], and/or the number of vehicles that should be positioned at each base [7, 9]. Most of these solutions use mixed integer linear programming models to solve the problem. On the other hand, there is previous work on *operational* ambulance planning. This has attracted a wider range of solution methods, including Markov decision theory [1] and simulation-based optimization [4].

The variety of solution methods for operational ambulance planning might be due to the difficulty of the problem. In dynamic ambulance management, the point of issue is to make decisions based on real-time information on the state of all vehicles and incidents. This makes for a complex issue, and systems quickly become intractable when the number of vehicles grows. Especially in urban areas, the situation can be considered extremely difficult because multiple vehicles operate closely to one another and therefore cannot be treated independently.

Some of the papers on operational ambulance planning use Markov Decision Problems (MDPs), or a variant thereof, to model the problem. Typically, those models can be divided into two categories: the first incorporates very little information, for example, only the *number* of idle ambulances in [1]. In contrast, the second category models many aspects of the real-time situation, resulting in an extremely large state space. Therefore, the MDP's in the latter case are not solvable with classical methods, and authors resort to other solutions such as approximate dynamic programming (e.g., [12, 14]).

The vast majority of the papers on dynamic ambulance management have focused on how to redeploy idle vehicles, (e.g., [1, 12, 17]). Perhaps in order not to overcomplicate things, they assume a basic dispatch rule: whenever an incident occurs, they decide to send the ambulance that is closest to the incident (in time). Although this is a common dispatch policy, it was already shown to be suboptimal in 1972 [5]. Regardless, most authors make this assumption without much discussion or justification; for example, the authors of [12] claim that it is an accurate enough representation of reality; however, they do not address the question of whether it is an *optimal* choice with respect to the objective (which is the fraction of incidents that is reached within the threshold time). The authors of [1] do not address the assumption at all.

Dispatching the closest idle ambulance seems to be so natural that justification is not needed, but one should not overlook the possibility to change the dispatch policy in order to improve the objective. In [17], the authors admit this is a possibility; however, they focus on relocating the vehicles when they become idle (instead of when they are dispatched). It should be clear that a clever dispatch policy can improve the performance of an EMS system, but since the topic has been underexposed in current literature, it is still unknown how much improvement can be expected. Furthermore, a dispatch method may be combined with a relocation rule to realize even greater improvements.

Few papers have discussed dispatch rules other than sending the closest idle ambulance. The ones that do propose other dispatch methods, typically do not prescribe a dynamic solution. For example, [4] proposes to divide the EMS region into separate sub-regions, each accompanied with its own ranking of bases from which an ambulance should preferably depart. Another example is [15], which considers a "regionalized response" dispatch policy. Under regionalized response, each region is preferably served by its own ambulance, even if it is temporarily outside of the region. Only if that vehicle is unavailable, the closest idle ambulance is sent. However, notice that both papers ignore information that we consider to be of crucial value: the outcome does not depend on whether some regions remain uncovered after the ambulance is dispatched. Alternatively, a choice could be made such that the remaining idle vehicles are in a good position with respect to future demand. This ensures that future incidents get a larger likelihood of being reached in time, thereby increasing the total expected fraction of incidents that can be reached within the time threshold.

One paper explicitly claims that alternative dispatch methods perform worse than the closest idle rule [8]. However, its true issue seems to be a computer-aided dispatch system that is not accurate enough to determine the true positions of the vehicles (and hence, also not able to determine the closest ambulance). The paper does not in fact deny that the "closest idle ambulance rule" might be improved. We emphasize that accurate location information is crucial in order to determine the best ambulance to send to an incident. Throughout this chapter, we will assume that such information is present. In many regions, such as Flevoland in the Netherlands, a monitoring tool is available that refreshes the exact GPS coordinates of each vehicle every 30 s. This seems accurate enough for our purposes.

9.1.2 Our Contribution

The main goal of this chapter is to better understand the ambulance dispatch process. In particular, we question the often-made assumption that one cannot do much better than the "closest idle" dispatch method. Thereto, we search for sensible dispatch rules other than the classical closest idle ambulance policy. We mainly focus on the often-used objective to minimize the fraction of arrivals later than a certain target time. However, we show that one of our methods can also be used for other KPI's.

First, we propose an MDP for ambulance dispatching, where the state space is described by an optional incident location and the availability of each of the ambulances. To the best of our knowledge, this is the first MDP in ambulance literature that models more than just the number of idle vehicles, without losing tractability for reasonably-sized ambulance fleets. In some sense, this model balances the amount of detail in the system representation—which typically results in a better outcome—with the computational difficulties. We mainly focus on minimizing the fraction of arrivals later than a target time, a typical objective in ambulance planning. However, we show that with a small change, our model can also minimize the average response time.

Second, we propose a heuristic for ambulance dispatching that behaves similar to the policy obtained from the MDP. However, it is able to determine more accurately what the response time would be when dispatching a driving ambulance. Furthermore, the heuristic can be computed in polynomial time, which allows us to apply it to regions with a large number of vehicles.

We validate our policies by a discrete-event simulation model of a Dutch EMS region. These simulations indicate that our proposed dispatch heuristic can decrease the fraction of late arrivals by as much as 13% relatively compared to the closest idle ambulance dispatch method. Our result sheds new light on the popular belief that deviating from the closest idle policy cannot greatly improve the objective. Although we do not advise all EMS managers to immediately discard the closest idle dispatch method, we do show that the typical argument—that it would not lead to large improvements in the fraction of late arrivals—should be changed.

The rest of this chapter is structured as follows. In Sect. 9.2, we give a formal problem definition. In Sect. 9.3, we present our proposed solution using Markov Decision Processes (MDPs), followed by a solution based on a scalable heuristic in Sect. 9.4. We show our results for a small, intuitive region in Sect. 9.5 and in a realistic case study for the Dutch area of Flevoland in Sect. 9.6.

9.2 Problem Formulation

Define the set V as the set of locations at which incidents can occur. Note that these demand locations are modeled as a set of discrete points. Incidents at locations in V occur according to a Poisson process with rate λ. Let d_i be the fraction of the demand rate λ that occurs at node i, $i \in V$. Then, on a smaller scale, incidents occur at node i with rate λd_i.

Let A be the set of ambulances, and $A_{idle} \subseteq A$ the set of currently idle ambulances. When an incident has occurred, we require an idle ambulance to immediately drive to the scene of the incident. The decision which ambulance to send has to be made at the moment we learn about the incident, and is the main question of interest in this chapter. When an incident occurs and there are no idle ambulances, the call goes to a first-come first-serve queue.

V	The set of demand locations.
H	The set of hospital locations, $H \subseteq V$.
A	The set of ambulances.
A_{idle}	The set of idle ambulances.
W_a	The base location for ambulance a, $a \in A$, $W_a \in V$.
T	The time threshold.
λ	Incident rate.
d_i	The fraction of demand in i, $i \in V$.
$\tau_{i,j}$	The driving time between i and j with siren turned on, $i, j \in V$.

Table 9.1: Notation

Our objectives are formulated in terms of response times: the time between an incident and the arrival of an ambulance. In practice, incidents have the requirement that an ambulance must be present within T time units. Therefore, we want to minimize the fraction of incidents for which the response time is larger than T. Another observation is that we want response times to be short, regardless of whether they are smaller or greater than T. We translate this into a separate objective, which is to minimize the average response time. We assume that the travel time $\tau_{i,j}$ between two nodes $i, j \in V$ is deterministic, and known in advance.

Sending an ambulance to an incident is followed by a chain of events, most of which are random. When an ambulance arrives at the incident scene, it provides service for a certain random time τ_{on_scene}. Then it is decided whether the patient needs transport to a hospital. If not, the ambulance immediately becomes idle. Otherwise, the ambulance drives to the nearest hospital in the set $H \subseteq V$. Upon arrival, the patient is transferred to the emergency department, taking a random time $\tau_{hospital}$, after which the ambulance becomes idle.

An ambulance that becomes idle may be dispatched to another incident immediately. Alternatively, it may return to its base location. Throughout this chapter, we will assume that we are dealing with a *static* ambulance system, i.e., each ambulance has a fixed, given base and may not drive to a different base. However, it is possible that multiple ambulances have the same base location. We denote the base location of ambulance a by W_a, for $a \in A$.

An overview of the notation can be found in Table 9.1.

9.3 Solution Method: Markov Decision Process

We model the ambulance dispatch problem as a discrete-time Markov Decision Process (MDP). In each state s (further defined in Sect. 9.3.1), we must choose an action from the set of allowed actions: $\mathcal{A}_s \subseteq \mathcal{A}$, which we describe in Sect. 9.3.2. The process evolves in time according to transition probabilities that depend on the chosen actions, as described in Sect. 9.3.4. We are dealing with an infinite planning horizon, and our goal is to maximize the average reward. The rewards are defined in Sect. 9.3.3. We eventually find our solution by performing value iteration [13].

In our model, we assume that at most one incident occurs within a time step. Therefore, the smaller the time steps, the more accurate the model will be. However, there is a tradeoff, as small time steps will increase the computation time. Throughout this chapter, we take time steps to be 1 min, which balances the accuracy and the computation time.

9.3.1 State Space

When designing a state space, it is important to store the most crucial information from the system in the states. However, when dealing with complex problems—such as real-time ambulance planning—it is tempting to store so much information, that the state space becomes intractable. This would lead to the so-called curse of dimensionality [2], which makes it impossible to solve the problem with well-known Markov Decision Problem (MDP) approaches.

As discussed before, there is little previous work on how to choose a good dispatch policy, but to some extent we can draw parallels with work on dynamic ambulance redeployment (which relocates idle vehicles): some researchers overcome the problem of an intractable state space by turning to Approximate Dynamic Programming, which allows for an elaborate state space to be solved approximately [12]. Alternatively, some researchers choose a rather limited state space, for example, by describing a state merely by the *number* of idle vehicles [1].

For our purpose, i.e., to determine *which* ambulance to send, it is important to know whether the ambulance we might send will arrive within T time units. Therefore, it is crucial to know where the incident took place. Furthermore, we require some knowledge of where the idle ambulances are. Clearly, storing only the number of idle vehicles would be insufficient. However, storing the location of each idle ambulance would already lead to an intractable state space for practical purposes. Instead, we can benefit from the fact that we are trying to improve a *static* solution. In a static solution, the home base for any ambulance is known in advance. Note that an idle ambulance must be either residing at its base location, or traveling towards the base. Hence, if we allow for an inaccuracy in the location of idle ambulances, in the sense that we use their destination rather than their actual location, their location does not need to be part of the state. Merely keeping track of whether each ambulance is idle or not, now suffices.

This leads us to a state s, defined as follows.

$$(Loc_{acc}, idle_1, idle_2, \ldots, idle_{|A|}), \tag{9.1}$$

where Loc_{acc} denotes the location of the incident that has just occurred in the last time step. In case no incident occurred in the last time step, we denote this by a dummy location, hence

$$Loc_{acc} \in V \cup \{0\}.$$

Furthermore, $idle_i$ denotes whether ambulance i is idle:

$$idle_i \in \{True, False\}, \quad \forall i \in A.$$

This leads to a state space of size $(|V|+1)2^{|A|}$.

For future reference, let $Loc_{acc}(s)$ denote the location of the incidents that have occurred in the previous time step when the system is in state s. Let $idle_i(s)$ denote whether or not ambulance i is idle in state s, $\forall i \in A, \forall s \in S$.

9.3.2 Policy Definition

In general, a policy Π can be defined as a mapping from the set of states to a set of actions: $S \to \mathcal{A}$. In our specific case, we define $\mathcal{A} = A \cup \{0\}$; that is if $\Pi(s) = a$, for $a \in A$, ambulance a should be sent to the incident that has just occurred at $Loc_{acc}(s)$. Action 0 may be interpreted as sending no ambulance at all (this is typically the choice when no incident occurred in the last time step, or when no ambulance is available).

In a certain state, not all actions are necessarily allowed. Denote the set of feasible actions in state s as

$$\mathcal{A}_s \subseteq \mathcal{A}, \quad \forall s \in S.$$

For example, it is not possible to send an ambulance that is already busy with another incident. This implies

$$!idle_a(s) \to a \notin \mathcal{A}_s, \quad \forall a \in A, \quad \forall s \in S. \tag{9.2}$$

Furthermore, let us require that when an incident has taken place, we must always send an ambulance—if any are idle.

$$\exists a \in A : idle_a(s) \ \wedge \ Loc_{acc}(s) \neq 0 \to 0 \notin \mathcal{A}_s, \quad \forall s \in S. \tag{9.3}$$

Moreover, if no incident has occurred, we may simplify our MDP by requiring that we do not send an ambulance:

$$Loc_{acc}(s) = 0 \to \mathcal{A}_s = \{0\}, \quad \forall s \in S. \tag{9.4}$$

All other actions from $\mathcal{A}$ that are not restricted by (9.2)–(9.4) are feasible. This completely defines the allowed action space for each state.

9.3.3 Rewards

In ambulance planning practice, a typical goal is to minimize the fraction of late arrivals. Since our decisions have no influence on the number of incidents that occur, this is equivalent to minimizing the *number* of late arrivals. An alternative goal might be to minimize average response times. Our MDP approach may serve either of these objectives, simply by changing the reward function.

Define $R(s,a)$ as the reward received when choosing action a in state s, $\forall s \in S, \forall a \in \mathcal{A}_s$. Note that in this definition, the reward does not depend on the next state. Keep in mind that our goal is to maximize the average rewards.

9.3.3.1 Fraction of Late Arrivals

To minimize the fraction of late arrivals, i.e., the fraction of incidents for which the response time is greater than T, we define the following rewards:

$$R(s,a) = \begin{cases} 0 & \text{if } Loc_{acc}(s) = 0; \\ -N & \text{if } Loc_{acc}(s) \neq 0 \ \wedge \ a = 0, \text{ i.e., no idle ambulances;} \\ 0 & \text{if } Loc_{acc}(s) \neq 0 \ \wedge \ a \in A \ \wedge \ \tau_{W_a,\, Loc_{acc}(s)} \leq T; \\ -1 & \text{otherwise.} \end{cases}$$

Here N is a number that is typically greater than 1. This implies that when all ambulances are busy, the rewards are smaller than when we send an ambulance that takes longer than T to arrive. This is in agreement with the general idea that having no ambulances available is a very bad situation. One might be tempted to make the reward for the only possible action ($a = 0$) in these states even smaller than we did, in order to influence the optimal actions in other states: the purpose would be to steer the process away from states with no ambulances available. However, note that this would not be useful, because our actions do not affect how often we end up in a state where all ambulances are busy. This is merely determined by the outcome of an external process, i.e., an unfortunate sequence of incidents. Therefore, an extremely small reward for action $a = 0$ in states where all ambulances are busy, would only blur the differences between rewards for actions in other states. (In our numerical experiments, we use $N = 5$.)

9.3.3.2 Average Response Time

To minimize the average response time, one may use the same MDP model, except with a different reward function. Let M be a large enough number, typically such that $M > \tau_{i,j}, \forall i, j \in V$. Then we can define the rewards as follows.

$$R(s,a) = \begin{cases} 0 & \text{if } Loc_{acc}(s) = 0; \\ -M & \text{if } Loc_{acc}(s) \neq 0 \ \wedge \ a = 0, \text{ i.e., no idle ambulances;} \\ -\tau_{W_a,\ Loc_{acc}(s)} & \text{if } Loc_{acc}(s) \neq 0 \ \wedge \ a \in A. \end{cases}$$

In our numerical experiments, we use $M = 15$ for the small region, and $M = 30$ for the region Flevoland. In both cases, $M > \tau_{i,j}, \forall i, j \in V$ holds. (In our implementation, time steps are equal to minutes.)

9.3.4 Transition Probabilities

Denote the probability of moving from state s to s', given that action a was chosen, as:

$$p^a(s,s'), \qquad \forall a \in \mathcal{A}_s, \qquad \forall s, s' \in S.$$

To compute the transition probabilities, note that the location of the next incident is independent of the set of idle ambulances. Thereto, $p^a(s,s')$ can be defined as a product of two probabilities. We write $p^a(s,s') = P_1(s') \cdot P_2^a(s,s')$, which stands for the probability that an incident happened at a specific location (P_1), and the probability that specific ambulances became available (P_2), respectively.

First of all, let us define $P_1(s')$. Since incidents occur according to a Poisson process, we can use the arrival rate λ (for an incident anywhere in the region) to obtain

$$P_1(s') = \begin{cases} \lambda \cdot d_{Loc_{acc}(s')} & \text{if } Loc_{acc}(s') \in V; \\ 1 - \lambda & \text{else.} \end{cases}$$

Note that the occurrence of incidents does not depend on the previous state (s).

Secondly, we need to model the process of ambulances that become busy or idle. For tractability, we will define our transition probabilities as if ambulances become idle according to a geometric distribution. In reality—and in our verification of the model—this is not the case, but since our objective is the long term average reward, this modeling choice should not have a negative impact [13]. Let us define a parameter $r \in [0,1]$, which represents the rate at which an ambulance becomes idle. We should set it in such a way, that the expected duration is equal to the average in practice. So this includes an average travel time, and an average time spent on scene. We add an average driving time to a hospital to that, as well as a realistic hospital drop off time—both multiplied with the probability that a patient needs to go to the hospital. For Dutch ambulances, this results in an average of roughly 38 min to become available after departing to an incident. For the geometric distribution, we know that the maximum likelihood estimate $\hat{r}$ is given by one divided by the sample mean. In this case, $\hat{r} = \frac{1}{38} \approx 0.0263$, which we use as the value for r in our numerical experiments.

We include a special definition if an ambulance was just dispatched. In such a case, the ambulance cannot be idle in the next time step. Furthermore, ambulances do not become busy, unless they have just been dispatched.

We now define

$$P_2^a(s,s') = \Pi_{i=1}^{|A|} P_{change}^a\big(idle_i(s), idle_i(s')\big), \quad \forall s,s' \in S,$$

where

$$P_{change}^a\big(idle_i(s), idle_i(s')\big) = \begin{cases} 1 & \text{if } a = i \wedge idle_i(s') = false; \\ 0 & \text{if } a = i \wedge idle_i(s') = true; \\ r & \text{if } a \neq i \wedge idle_i(s) = false \wedge idle_i(s') = true; \\ 1-r & \text{if } a \neq i \wedge idle_i(s) = false \wedge idle_i(s') = false; \\ 0 & \text{if } a \neq i \wedge idle_i(s) = true \wedge idle_i(s') = false; \\ 1 & \text{otherwise.} \end{cases} \tag{9.5}$$

9.3.5 Value Iteration

Now that we have defined the states, actions, rewards and transition probabilities, we can perform value iteration to solve the MDP. Value iteration, also known as backward induction, calculates a value $V(s)$ for each state $s \in S$. The optimal policy, i.e., the best action to take in each state, is the action that maximizes the expected value of the resulting state s'.

$V(s)$ is calculated iteratively, starting with an arbitrary value $V_0(s)$ $\forall s \in S$. (In our case, we start with $V_0(s) = 0$ $\forall s \in S$.) In each iteration i, one computes the values $V_i(s)$ given $V_{i-1}(s)$ $\forall s \in S$ as follows.

$$V_i(s) := \max_{a \in \mathcal{A}_s} \{ \sum_{s'} p^a(s,s')(R(s,a) + V_{i-1}(s')) \} \tag{9.6}$$

This is known as the "Bellman equation" [3].

When the span of V_i, i.e., $\max V_i(s) - \min V_i(s)$, converges, the left-hand side becomes equal to the right-hand side in Eq. (9.6), except for an additive constant. After this convergence is reached, the value of $V(s)$ is equal to $V_i(s)$ $\forall s \in S$.

We will extensively discuss the results of the MDP solution later in Sect. 9.6, but we now briefly state two observations. First, if an incident can not be reached in time anyway, the system is quite likely to choose an ambulance other than the closest idle one. The explanation is that, if the time threshold cannot be met, one might as well choose ambulance such that the remaining ambulances are in a favorable position with respect to possible future incidents. Second, whenever at least one vehicle is available within the target time, the MDP solution never prescribes to send a vehicle that is further away than the threshold.[1] That is to say, it does not appear to be beneficial to sacrifice performance now in the hope of achieving better performance later. We use these observations in the design of a dispatch heuristic that follows next.

[1] This holds for the realistic region that we implemented, but does not necessarily hold in general.

9.4 Solution Method: Dynamic MEXCLP Heuristic for Dispatching

In this section we describe a dispatch heuristic that was inspired by MDP solution above. The main benefit of the heuristic is that it can be computed in real time, for any number of vehicles and ambulance bases that is likely to occur in practice. Furthermore, the method is easy to implement. The heuristic is related to dynamic MEXCLP, also known as "DMEXCLP" [10], a heuristic that was originally designed to redeploy idle ambulances in real time.

The general idea is that, at any time, we can calculate the *coverage* provided by the currently idle ambulances. This results in a number that indicates how well we can serve the incidents that might occur in the (near) future.

More specifically, coverage is defined as in the MEXCLP model [7], that we will describe next.

9.4.1 Coverage According to the MEXCLP Model

In this section we briefly describe the objective of the well-known MEXCLP model. MEXCLP was originally designed to optimize the distribution of a limited number, say $|A|$, ambulances over a set of possible base locations W. Each ambulance is modeled to be unavailable with a pre-determined probability q, called the *busy fraction*. Consider a node $i \in V$ that is within range of k ambulances. The travel times $\tau_{i,j}$ $(i, j \in V)$ are assumed to be deterministic, which allow us to straightforwardly determine this number k. If we let d_i be the demand at node i, the expected covered demand of this vertex is $E_k = d_i(1-q^k)$. Note that the marginal contribution of the kth ambulance to this expected value is $E_k - E_{k-1} = d_i(1-q)q^{k-1}$. Furthermore, the model uses binary variables y_{ik} that are equal to 1 if and only if vertex $i \in V$ is within range of at least k ambulances. The objective of the MEXCLP model can now be written as:

$$\text{Maximize} \sum_{i \in V} \sum_{k=1}^{|A|} d_i(1-q)q^{k-1} y_{ik}.$$

In [7], the author adds several constraints to ensure that the variables y_{ik} are set in a feasible manner. For our purpose, we do not need these constraints, as we shall determine how many ambulances are within reach of our demand points—the equivalent of y_{ik}—in a different way.

9.4.2 Applying MEXCLP to the Dispatch Process

The dispatch problem requires us to decide which (idle) ambulance to send, at the moment an incident occurs. Thereto, we compute the *marginal* coverage that each ambulance provides for the region. The ambulance that provides the smallest marginal coverage, is the best choice for dispatch, in terms of remaining coverage for future incidents. However, this does not incorporate the desire to reach the current incident within target time T. We propose to combine the two objectives—reaching the incident in time and remaining a well-covered region—by always sending an ambulance that will reach the incident in time, if possible. This still leaves a certain amount of freedom in determining which *particular* ambulance to send.

The computations require information about the location of the (idle) ambulances. Denote this by $Loc(a)$ for all $a \in A_{idle}$. This model allows us to use the real positions of ambulances, which in practice may be determined by GPS signals. For simulation purposes, the current position of the ambulance while driving may be determined using, e.g., interpolation between the origin and destination, taking into account the travel speed. In either case, the location should be rounded to the nearest point in V, because travel times $\tau_{i,j}$ are only known between any $i, j \in V$.

Let A^+_{idle} denote the set of idle ambulances that are able to reach the incident in time. Similarly, let A^-_{idle} denote the set of idle ambulances that cannot reach the incident in time, which implies that $A^+_{idle} \cup A^-_{idle} = A_{idle}$. Then, if $A^+_{idle} \neq \emptyset$, we decide to dispatch a vehicle that will arrive within the threshold time, but chosen such that the coverage provided by the remaining idle vehicles is as large as possible:

$$\arg\min_{x \in A^+_{idle}} \sum_{i \in V} d_i (1-q) q^{k(i,A_{idle})-1} \cdot \mathbb{1}_{\tau_{Loc(x),i} \leq T}. \tag{9.7}$$

Otherwise, simply dispatch a vehicle such that the coverage provided by the remaining idle vehicles is as large as possible (without requiring an arrival within the threshold time):

$$\arg\min_{x \in A^-_{idle}} \sum_{i \in V} d_i (1-q) q^{k(i,A_{idle})-1} \cdot \mathbb{1}_{\tau_{Loc(x),i} \leq T}. \tag{9.8}$$

Note that in our notation, k is not an iterable, but a function of i and A_{idle}. $k(i, A_{idle})$ represents the number of idle ambulances that are currently within reach of vertex i. After choosing the locations of ambulances that one wishes to use—the real locations or the destinations—$k(i, A_{idle})$ can be counted in a straightforward manner.

9.5 Results: A Motivating Example

In this section, we consider a small region for which there is some intuition with respect to the best dispatch policy. We show that the intuitive dispatch policy that minimizes the fraction of late arrivals, is in fact obtained by both our solution methods (based on MDP and MEXCLP). We will address the alternative objective, i.e., minimizing the average response times, as well.

Figure 9.1 shows a toy example for demonstrative purposes. We let calls arrive according to a Poisson process with on average one incident per 45 min. Furthermore, incidents occur w.p. 0.1 in Town 1, and w.p. 0.9 in Town 2. Eighty percent of all incidents require transport to the hospital, which is located in Town 2.

9.5.1 Fraction of Late Arrivals

This section deals with minimizing the fraction of response times greater than 12 min. A quick analysis of the region in Fig. 9.1 leads to the observation that the "closest idle" dispatch strategy must be suboptimal. In order to serve as many incidents as possible within 12 min, it is evident that the optimal dispatch strategy should be as follows: when an incident occurs in Town 2, send ambulance 2 (if available). In all other cases, send ambulance 1 (if available). Both the MDP solution that attempts to minimize the fraction of late arrivals (with, e.g., $M = 15$), as well as the dispatch heuristic based on MEXCLP, lead to this policy (Fig. 9.2).

Note that in our model, it is mandatory to send an ambulance, if at least one is idle. Furthermore, we do not base our decision on the locations of idle ambulances (instead, we pretend they are at their destination, which is fixed for each ambulance). Therefore, in this example with 2 ambulances, one can describe a dispatch policy completely by defining which ambulance to send when both are idle, for each possible incident location. For an overview of the various policies, see Table 9.2.

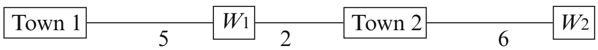

Fig. 9.1: A graph representation of the region. The numbers on the edges represent the driving times in minutes with siren turned on. W_1 and W_2 represent the base locations of ambulance 1 and 2, respectively. Incidents occur only in Town 1 and Town 2. The only hospital is located in Town 2

Solution method	$Loc_{acc} = Town1$	$Loc_{acc} = Town2$
MEXCLP(destination) heuristic	W_1	W_2
MDP(frac)	W_1	W_2
MDP(avg)	W_1	W_1

Table 9.2: An overview of the behavior of various dispatch policies when both ambulances are idle. The value in the table represents the base from which an ambulance should be dispatched

9.5.2 Average Response Time

We used the MDP method described in Sect. 9.3.3.2 to obtain a policy that should minimize the average response time, let us denote this policy by MDP(avg). We evaluate the performance of the obtained policy, again by simulating the EMS activities in the region. These simulations show that the MDP solution indeed reduces the average response time significantly, compared to the policy that minimizes the fraction of late arrivals (MDP(frac))—see Fig. 9.3.

9.6 Results: Region Flevoland

In this section, we simulate the redeployment method that we obtained from our MDP for a realistic problem instance. The Netherlands is divided in 24 regions, each operated by its own ambulance provider (see Fig. 9.4).

We modeled the region of Flevoland, which in practice is served by the ambulance provider "GGD Flevoland". For the parameters used in the implementation, see Table 9.3. This is a region with multiple hospitals, and for simplicity we assume that the patient is always transported to the nearest hospital, if necessary.

We estimated the arrival intensity for this region from historical data, and determined a reasonable number of vehicles that can serve this demand. Consequently,

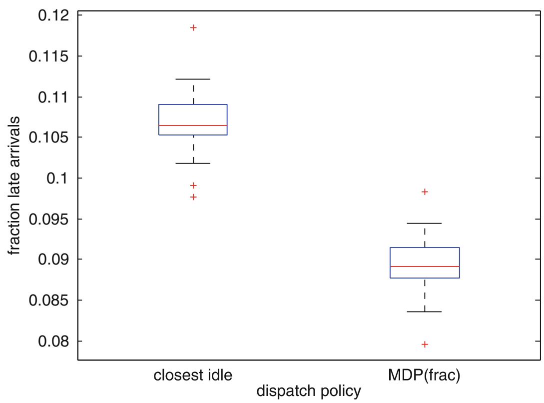

Fig. 9.2: Fraction of late arrivals as observed in a simulation of the small region. This figure shows the performance of the MDP solution that attempts to minimize the fraction of late arrivals (after value iteration converged). The performance is compared with the "closest idle" dispatch policy. Each policy was evaluated with 20 runs of 5000 simulated hours

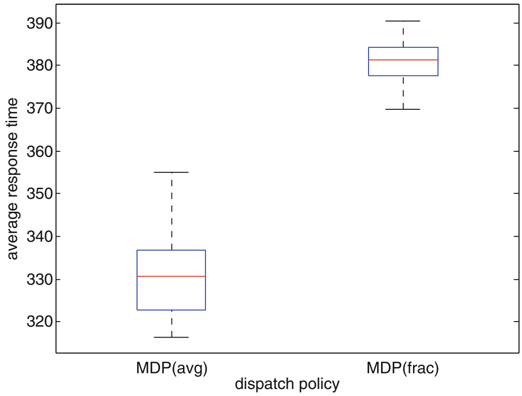

Fig. 9.3: The average response times in seconds, as observed in simulations of the small region. This figure shows the performance of the MDP solution that attempts to minimize the fraction of late arrivals versus the MDP solution that attempts to minimize the average response time (after value iteration has converged). Each policy was evaluated with 20 runs of 5000 simulated hours

we distributed the vehicles according to the solution of the (static) MEXCLP model, as described in Sect. 9.4.1. This model is generally assumed to give reasonably good solutions [16]. This static MEXCLP solution can be seen in Fig. 9.5.

Note that we used the fraction of inhabitants as our choice for d_i. In reality, the fraction of demand could differ from the fraction of inhabitants. However, the number of inhabitants is known with great accuracy, and this is a straightforward

Fig. 9.4: The 24 EMS regions in the Netherlands

Parameter	Magnitude	Choice
λ	1/29 min	A realistic rate for Flevoland.
A	8	A reasonable amount to serve the demand.
W_a for $a \in A$		Postal codes as depicted in Fig. 9.5.
V	91	All 4 digit postal codes in the region.
H	3	The hospitals within the region in 2013.
$\tau_{i,j}$		Driving times as estimated by the RIVM.
d_i		Fraction of inhabitants as known in 2009.

Table 9.3: Parameter choices for our implementation of the region of Flevoland

way to obtain a realistic setting. Furthermore, the analysis of robust optimization for uncertain ambulance demand in [11] indicates that we are likely to find good solutions, even if we make mistakes in our estimates for d_i.

In the Netherlands, the time target for the highest priority emergency calls is 15 min. Usually, 3 min are reserved for answering the call, therefore we choose to run our simulations with $T = 12$ min. The driving times for EMS vehicles between any two nodes in V were estimated by the Dutch National Institute for Public Health and the Environment (RIVM) in 2009. These are driving times with the siren turned on. For ambulance movements without siren, e.g., when repositioning, we used 0.9 times the speed with siren. The number of vehicles used in our implementation is such that value iteration is still tractable.

We simulate the problem as described in Sect. 9.2. In these simulations, ambulances that become idle are immediately dispatched to a waiting incident (if any),

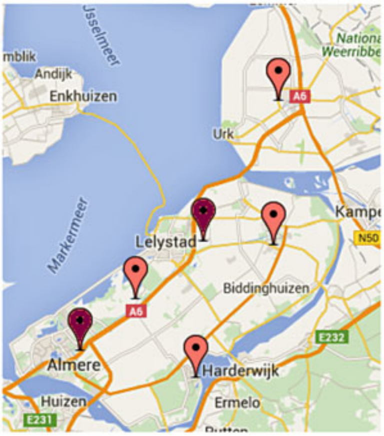

Fig. 9.5: Optimal distribution of 8 ambulances over the region Flevoland, according to the MEXCLP solution. Two vehicles are stationed at each of the *dark colored* bases, while 1 vehicle is present at the *lighter colored* locations

or head back to their home base. The locations we used as home bases are depicted in Fig. 9.5, and correspond to actual base locations in the EMS region. Ambulances that were dispatched while on the road, did not return to their base first.

In our simulation, $\tau_{onscene}$ is exponentially distributed with an expectation of 12 min. $\tau_{hospital}$ is drawn from a Weibull distribution with an expectation of approximately 15 min. More specifically, it has shape parameter 1.5 and scale parameter 18 (in minutes). We state these distributions for completeness, however, our numerical experiments indicate that the performance does not depend much on the chosen distribution for $\tau_{onscene}$ or $\tau_{hospital}$. In our simulations, patients need hospital treatment with probability 0.8. This value was estimated from Dutch data.

Note that $\tau_{onscene}$ or $\tau_{hospital}$ and the probability that a patient needs hospital treatment are not explicitly part of our solution methods. Instead, they subtly affect the busy fraction q (for the heuristic) or the transition probabilities with rate r (for the MDP).

9.6.1 Analysis of the MDP Solution for Flevoland

In this section, we highlight and explain some features of the MDP solution for the region Flevoland. In particular, we will focus on the states for which the MDP solution differs from the closest idle policy.

The output of the MDP is a table with the incident location, the status of the different ambulances (idle or not), and the optimal action. This output is a rather large table (with, in the case of Flevoland, 23,552 entries) that does not easily reveal insight into the optimal policy. To this end, we first reduced the size of the table: we filtered out all states for which no real decision has to be made (i.e., states in which no incident occurs, and states in which less than 2 ambulances are idle). This reduces the table size to 22,477. In 5583 of these states, the MDP solution is to send a vehicle *other* than the closest idle one (i.e., roughly 25%).

To understand in *which* states the MDP solution prescribes to send a vehicle other than the closest idle one, we used classification and regression trees (CART trees) on the table to find structure in the form of a decision tree. We used random forests to create the decision tree, since it is known that a basic CART has poor predictive performance. While bagging trees reduces the variance in the prediction, random forests also cancel any correlation structure in the generation of the trees that may be presenting while bagging.

The outcome that describes the MDP solution is a decision tree that divides the state space into three regions, see Fig. 9.6. For the red and the green region, whether or not the closest idle ambulance is sent, depends heavily on the availability of ambulances at base 1: if an incident occurred in the red region, one should *always* send an ambulance from base 1 if possible. For the red and green region combined, this same advice holds in 95% of the states.

If no ambulance at base 1 is idle, the MDP solution prescribes to deviate from the "closest idle" choice quite often, and especially so for the red and green regions: if there are more than two ambulances idle (but none at base 1), the dispatcher should deviate from the closest idle policy in 81% of the states. If exactly two ambulances are idle, this is still true for almost 50% of the states. This may be intuitively understood as follows. Since incidents on the red nodes can not be reached in time anyway, choosing an ambulance that is further away than the closest idle one results in a enlarged response time. However, using our objective of the fraction late arrivals, this is not a downside, since the incident could not be reached in time anyway. Therefore, an ambulance can be chosen such that the remaining ambulances are in a favorable position with respect to possible future incidents. Note that this is also the general idea that forms the basis of our MEXCLP dispatch heuristic.

For incidents on the blue nodes, the best decision according to the MDP is in roughly 70% of the states equal to the closest idle vehicle, and hardly depends on whether or not the ambulances at base 1 are idle.

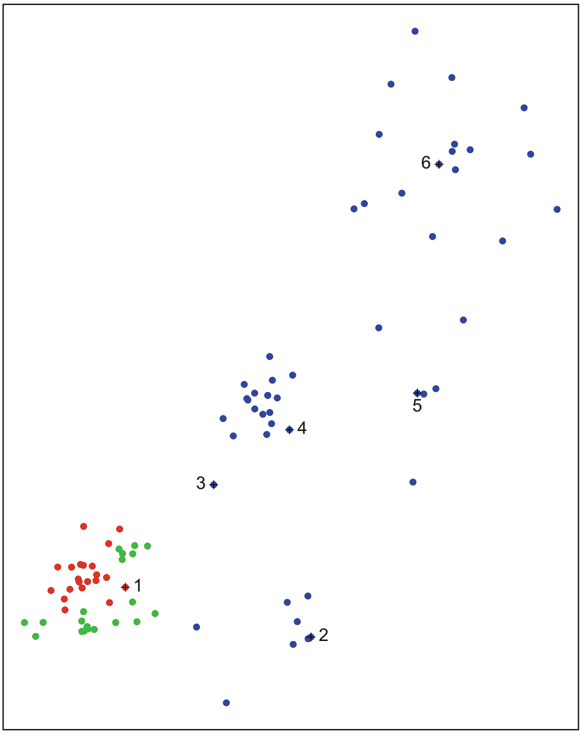

Fig. 9.6: Each node represents a postal code in Flevoland. Nodes with the same color have similar MDP solutions. The numbers indicate the bases. (Two vehicles are stationed at base numbers 1 and 4)

Out of all states for which at least one vehicle is available within the target time, the MDP solution *never* prescribes to send a vehicle that is further away than the threshold. That is to say, it does not appear to be beneficial to sacrifice performance now in the hope of achieving better performance later. This—again—is similar to the MEXCLP dispatch heuristic, because the heuristic also only sends a vehicle that will arrive late if there is no other option.

For this realistic region, value iteration took a long time to converge. Instead of waiting for convergence, one might also be interested in using the policy we get after a fixed number of value iterations. Figure 9.7 indicates that the performance after 3 iterations is already quite similar to the converged alternative.

9.6.2 Results

In this section, we show the results from our simulations of the EMS region of Flevoland.

We ran simulations using three different dispatch policies: the closest idle method, the MEXCLP-based heuristic and the MDP solution after convergence of the value iteration. Figure 9.8 compares their performance in terms of the observed fraction of response times larger than the threshold time.

The results show that the MDP solution that was designed to minimize the fraction of late arrivals has approximately the same performance[2] as the closest idle policy. Although this performance is perhaps somewhat worse than one may have hoped for, it is important to remember that the MDP has to decide which ambulance

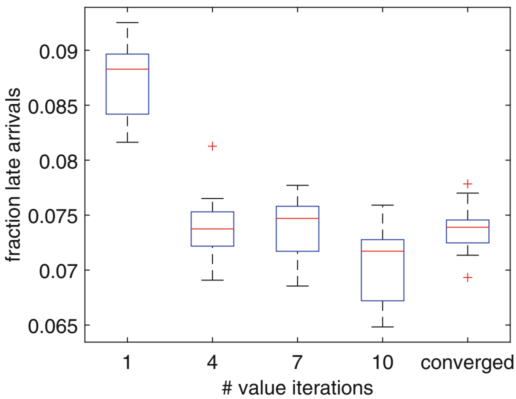

Fig. 9.7: The performance of the MDP solution for region Flevoland after 1, 4, 7 and 10 value iterations. Each policy was evaluated with 15 runs of 5000 simulated hours

[2] The fraction of late arrivals.

to send, based on which *base location* it belongs to. This, however, is not necessarily accurate since the ambulance may still be on the road (returning to base) at the moment of dispatch. In the simulations, this effect is accurately captured, but the MDP cannot (since keeping track of the true location of ambulances would lead to a state space explosion). Note that the "closest idle" method *does* have access to the real locations of vehicles at the time of dispatch, and thereby has a certain advantage.

As the MEXCLP-based dispatch heuristic has access to the real vehicle locations, it is not surprising that this method is able to perform better than the MDP solution. Consequently, the heuristic performs significantly better than the "closest idle" policy: it reduces the fraction of late arrivals from 7.22% to 6.26% on average: a relative improvement of 13%.

As mentioned earlier, the fraction of late arrivals is an important performance indicator for ambulance providers; however, one should also look at other aspects of the response time to make a well-informed decision on whether or not to implement a certain policy. We measure the average response time, as observed in our simulations: for the MDP solution, this is 483.9 s. The heuristic is slightly better with 478.9 s on average. The closest idle method outperforms both, resulting in an average response time of 409.3 s. Note that, with respect to this objective, the closest idle method is almost 15% better than our heuristic. In some sense, this is not surprising, but it does illustrate that our heuristic has such a strong focus on the fraction of late arrivals, that it becomes completely ignorant to the effect it has on the average response time (and the same holds for the MDP). This is an important observation that ambulance providers should keep in mind when they consider deviating from the closest idle policy.

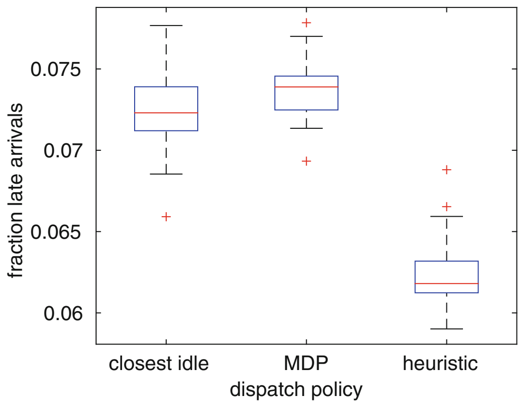

Fig. 9.8: Comparing the performance of the "closest idle" policy with the MDP solution and the Dynamic MEXCLP dispatch heuristic (where $q = 0.25$). Each policy was evaluated with 15 runs of 5000 simulated hours

9.7 Conclusion and Discussion

This chapter introduced two methods to obtain ambulance dispatch policies. Firstly, we modeled the ambulance dispatch problem as a Markov Decision Problem (MDP). This model is unique, in the sense that it is the first MDP in ambulance literature that keeps track of more than just the number of idle vehicles, without losing tractability for reasonably-sized ambulance fleets. Secondly, we introduced a heuristic that can easily be computed for any realistic size ambulance fleet.

Our result sheds new light on the popular belief that deviating from the closest idle dispatch policy cannot greatly improve the objective (the expected fraction of late arrivals). The above shown improvement of 13% was unexpectedly large. We consider this the main contribution of our work. Our methods yield in a great improvement in this KPI, however: one should be careful if one is also interested in other aspects of the response time. It is important to remember that our policies were designed with emphasis on the fraction of late arrivals only. Therefore, we do not claim that our dispatch policies are practically preferable over the closest idle policy, but we have shown that the argumentation for not using alternatives should be different. One should argue that we do not deviate from the closest idle policy, because we do not know how to do this while improving response times overall—and not because the alternatives fail to improve the fraction of late arrivals.

9.7.1 Further Research

One might consider making small changes to the MDP that could benefit the performance. For example, one idea is to artificially increase the rate with which busy ambulances become idle. This extra time would allow for ambulances to drive back to their home base, before the MDP considers them to be idle again. That way, we avoid the error where the MDP decides that an ambulance will reach an incident within the time threshold, but in fact the ambulance is still returning to base and happens to be further away from the incident. We suspect that this approach might give a small improvement; however, it should be noted that there is also a downside to making this change: ambulances are considered to be busy even though they are free, and hence suboptimal decisions will be made from time to time. In fact, sometimes an ambulance is *closer* than the MDP knows, because its previous patient was in the same area as the next patient.

Other changes could be, to add more information in the state about the ambulance's actual location while driving back to the home base. This, however, would lead to a state space explosion and the resulting model will—for realistically sized regions—most certainly not be solvable by value iteration.

Acknowledgements The authors of this chapter would like to thank the Dutch Public Ministry of Health (RIVM) for giving access to their estimated travel times for EMS vehicles in the Netherlands. This research was financed in part by Technology Foundation STW under contract 11,986, which we gratefully acknowledge.

Appendix: Notation

Notation in this chapter	Common notation
$\mathcal{A}_s$	$A(s)$
$R(s,a)$	$r^a(s)$
$p^a(s,s')$	$p(s\|s',a)$
$V_i(s)$	$V_t(s)$

References

1. R. Alanis, A. Ingolfsson, B. Kolfal, A Markov chain model for an EMS system with repositioning. Prod. Oper. Manag. **22**(1), 216–231 (2013)
2. R. Bellman, *Dynamic Programming* (Princeton University Press, Princeton, 1957)
3. R. Bellman, A Markovian decision process. J. Math. Mech. **6**(4), 679–684 (1957)
4. R. Bjarnason, P. Tadepalli, A. Fern, Simulation-based optimization of resource placement and emergency response, in *Proceedings of the Twenty-First Innovative Applications of Artificial Intelligence Conference*, 2009
5. G. Carter, J. Chaiken, E. Ignall, Response areas for two emergency units. Oper. Res. **20**(3), 571–594 (1972)
6. R.L. Church, C.S. Revelle, The maximal covering location problem. Pap. Reg. Sci. Assoc. **32**, 101–118 (1974)
7. M.S. Daskin, A maximum expected location model: formulation, properties and heuristic solution. Transp. Sci. **7**, 48–70 (1983)
8. S.F. Dean, Why the closest ambulance cannot be dispatched in an urban emergency medical services system. Prehosp. Disaster Med. **23**(02), 161–165 (2008)
9. J. Goldberg, R. Dietrich, J.M. Chen, M.G. Mitwasi, Validating and applying a model for locating emergency medical services in Tucson, AZ. Euro **34**, 308–324 (1990)
10. C.J. Jagtenberg, S. Bhulai, R.D. van der Mei, An efficient heuristic for real-time ambulance redeployment. Oper. Res. Health Care **4**, 27–35 (2015)
11. R.B.O. Kerkkamp, Optimising the deployment of emergency medical services. Master's thesis, Delft University of Technology, 2014
12. M.S. Maxwell, M. Restrepo, S.G. Henderson, H. Topaloglu, Approximate dynamic programming for ambulance redeployment. INFORMS J. Comput. **22**, 226–281 (2010)

13. M. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming* (Wiley, New York, 1994)
14. V. Schmid, Solving the dynamic ambulance relocation and dispatching problem using approximate dynamic programming. Eur. J. Oper. Res. **219**(3), 611–621 (2012)
15. C. Swoveland, D. Uyeno, I. Vertinsky, R. Vickson, Ambulance location: a probabilistic enumeration approach. Manag. Sci. **20**(4), 686–698 (1973)
16. P.L. van den Berg, J.T. van Essen, E.J. Harderwijk, Comparison of static ambulance location models. Under review, 2014
17. Y. Yue, L. Marla, R. Krishnan, An efficient simulation-based approach to ambulance fleet allocation and dynamic redeployment, in *AAAI Conference on Artificial Intelligence (AAAI)*, July 2012