Optimisation of Small-Scale Ambulance Move-up

Lei Zhang
Department of Engineering Science
University of Auckland
New Zealand
lzha174@aucklanduni.ac.nz

Abstract

This paper proposes a multi-time scale Markov Decision Process (T-MDP) model for relocating a small fleet of ambulances to better respond to emergency calls. A set of temporally abstract states are used to approximate the intractable state space in the real world. A T-stage look-ahead scheme is developed to approximate the temporally accrued rewards and discounted probabilities for the T-MDP. An example is given to compare its performance with five other ambulance locating models.

Key words: dynamic programming, healthcare, ambulance relocation.

1 Introduction

The ambulance problem we consider is characterised by the response process summarised as follows. When an emergency call is received, a dispatcher chooses an available ambulance to dispatch. A typical dispatch policy will look at those vehicles waiting idle at a base or returning to their base and find the vehicle that is closest to the accident scene. (The time required for this dispatching process is typically small, and so will be ignored here.) This dispatched ambulance travels to the scene of the emergency call. Once the ambulance reaches the scene, the ambulance officers perform an initial at-scene treatment of the patient. If no more medical care is required then the ambulance becomes free at the scene, and returns to its base. More typically, however, transportation is required to a hospital and the ambulance then becomes free at the hospital after completing a patient hand-over.

The elapsed time between the receipt of the call and the vehicle arriving at the scene is termed the *response time*. An ambulance organisation's performance will often be measured by the percentage of calls having a response time no greater than some target time W. When trying to maximise their performance, ambulance operators typically refer to their readiness to respond to the next call in terms of *coverage*, where a suburb is considered covered if its centroid is no further than W minutes drive from the closest available ambulance. They also refer to *call coverage* which is the probability that the location of the next call is no further than W minutes drive from the closest available ambulance.

In many ambulance organisations, each ambulance is assigned a base to which the vehicle returns after each call; determining the best base for each vehicle gives us a static location problem. These problems are typically solved using integer programming (IP) models which seek a vehicle-to-base assignment that maximises some simple coverage-based model of expected system performance.

In an attempt to improve their response times, some ambulance operators operate a *redeployment* policy in which they move idle ambulances from one base to another, or even to street corners, as they seek to improve their call coverage. This vehicle movement is an example of a *move-up*. A common redeployment approach is System Status Management (SSM) which, for any given number of free vehicles n_{free}, specifies a pre-defined vehicle *configuration* $C(n_{\text{free}})$ that gives a standby location (i.e. a base or a street corner) for each of the idle vehicles; e.g. see Bryan et al. 2010. Whenever the number of free vehicles changes, the dispatchers are required to determine a set of move-ups that efficiently move vehicles into the appropriate SSM configuration.

An alternative approach, which is the focus of our work, is to adopt dynamic relocation models to determine optimal or near optimal move-ups for the available ambulances. Unlike SSM plans, these solutions do not enforce a single configuration $C(n_{\text{free}})$ for each $n_{\text{free}} = 1, 2, 3, ...$, but instead allow the target configuration to depend on the current vehicle locations. For a broad overview on ambulance location and relocation models, see Brotcorne, Laporte, and Semet 2003.

In this paper, we propose a multi-time scale Markov Decision Process (T-MDP) formulation to relocate a small fleet of ambulances. Traditional Markov Decision Processes (MDPs) assume a single fixed time step: actions take one step to complete, and their immediate consequences become available after one step. This often leads to intractable state space for real problems. Moreover, in many applications, we are more interested in common-sense, higher-level actions such as the destination for each vehicle, rather than next location each vehicle goes to in one time step. These higher-level actions are similar to Artificial Intelligence's classical macro operators (Richard 1985), in that they can take control for some period of time using a sequence of primitive actions (next location for each vehicle in one time step) until some termination condition is met, at which point a new macro-action can be applied. Macro-actions can be useful when solving MDPs with large state and action spaces by focusing on a subset of states. The key is to treat these macro-actions just like primitive actions, which have associated temporally accrued rewards and transition probabilities during execution of each macro-action. These macro-actions can be combined with primitive actions to improve rate of convergence, or used in temporally abstracted MDPs to find (near) optimal solutions to the original problems. T-MDP chooses a subset of states from the original MDP and defines a set of macro-actions between these states along with temporally accrued rewards and transition probabilities. Then the standard value iteration or policy iteration for a dynamic programme can be used to solve the problem. For related theoretical analysis and applications, see Sutton 1995 and Hauskrecht et al. 1998.

The remainder of this paper is organised as follows. In section 2, we briefly discuss multi-time scale MDPs (T-MDPs). In section 3, we present the T-MDP formulation for ambulance move-up. In section 4, we provide computational experiments. Section 5 summarises our findings.

2 Background

A finite MDP consists of four elements: A finite state space S; A finite set $A = A_1 \cup A_2 ... \cup A_s$ where A_s is a set of actions that can be performed at state s; A bonded reward function $R : S \times A \to \mathbb{R}$ such that $R(s, a)$ denotes the immediate reward associated with action a in s; A transition distribution P: $S \times A \times S \to [0, 1]$ such that $P(s, a, w)$ denotes the probability of moving to state w when action a is performed at state s on one time step. The objective is to find a *policy* that maximises the expected accumulated reward over an infinite horizon: $E(\sum_{t=0}^{\infty} \gamma^t r^t)$ where r^t is a reward obtained at time t and $\gamma \to (0, 1)$ is a discount factor. This objective can also be formulated as a Bellman optimality equation (Bellman 1957):

$$V(s) = \max_{a \in A}(R(s, a) + \gamma \sum_{w \in S} P(s, a, w)V(w))$$

where $V(s)$ is the value function at state s.

This equation represents the dynamics of the system in one time step. Sutton (1995) viewed such an action $a \in A$ as a primitive action which leads to immediate consequences and generalized the equation for multi-time scale MDPs, which summarise several time scales and have the ability to predict events that can happen at various unknown moments. In a multi-time scale MDP, an action is viewed as a macro-action which takes control of the system for some period of time using a sequence of primitive actions until some termination condition is met and a new macro action takes over. Note that primitive actions also qualify as macro-actions in the general form: they are initiated in a state, take control for a while (one time step), and then end. Let s be a state that initiates a macro action a. Let τ be the time at which some termination condition is met. The reward for this action can be written as:

$$R(s, a) = E_\tau(\sum_{t=0}^{\tau} \gamma^t r(s^t, a(s^t))|s^0 = s, a)$$

where $a(s^t)$ specifies the primitive action taken at time t (state s^t) under the macro-action a with an immediate reward $r(s^t, a(s^t))$. Similarly, the discounted transition probabilities according to the expected termination time for this macro-action can be written as:

$$\begin{aligned} P(s, a, w) &= E_\tau(\gamma^{\tau-1} Pr(s^\tau = w|s^0 = s, a)) \\ &= \sum_{t=1}^{\infty} \gamma^{t-1} Pr(\tau = w, s^t = s'|s^0 = s, a) \end{aligned}$$

Now the general form of the Bellman optimality equation can be written as

$$V(s) = \max_{a \in A(s)} (R(s, a) + \sum_{w \in S} P(s, a, w)V(w)) \tag{1}$$

where $A(s)$ is a finite set of macro-actions at state s.

3 Model Description

In this section, we define the set of temporally abstract states for the T-MDP model and show how to approximate their temporally accrued rewards and transition probabilities for move-up actions using a look-ahead scheme.

3.1 Temporally Abstract State Space

Let M represent the total number of ambulances in the EMS system. Let G represent a network consisting of a set of nodes N and a set of undirected links L, where $(i, j) \in L$ is the undirected link joining nodes i and j, $i, j \in N$. The spacing of the nodes is such that each link requires a constant drive time Δt to traverse. The travel time from node i to j along the shortest path is denoted as $d(i, j)$. Call arrivals follow a Poisson process with a total arrival rate λ. The probability that the next call occurs at node k is p_k with $\sum_{k \in N} p_k = 1$. Let B represent a set of preselected nodes as ambulance bases. Let B_i represent the ith element in B. If an ambulance is available, it should be either traveling to or waiting at a base. The dispatch policy is to send the closest ambulance. We assume transport to hospital is always needed. (In reality, an average of 75% of calls requires transport to hospitals (Maxwell et al. 2010)). The service time from being dispatched to becoming free at hospital is assumed to be an exponential distribution with rate μ.

The state space of our model consists of two parts. The first part of state space consists of the set S^c of states in which idle ambulances are waiting at bases. We restrict the capacity at each base to be one vehicle. These states are viewed as candidate stable states since they are the configurations we consider for ambulance move-up in real time. Once ambulances are in one of these states, they 'do nothing' until an event (a call arrival or a completion of service) occurs. A candidate stable state s^c is a binary vector of length $|B|$ where the ith element (s^c_i) is 1 if an ambulance is waiting at base location B_i. For example, if B is $\{5, 12, 18\}$, $M = 3$, then $(0, 0, 1)$ indicates that there is one idle ambulance at location 18 and two ambulances being busy. The set S^c is created by enumerating every subset, including an empty subset, of B. For example, if $B = \{5, 12, 18\}$, then $S^c = \{(0, 0, 0), (0, 0, 1), (0, 1, 0), (1, 0, 0), (0, 1, 1), (1, 0, 1), (1, 1, 0), (1, 1, 1)\}$. The state with all zeros represents that all vehicles are busy.

The second part of state space consists of set S^t of 'temporary' states which are the consequences of an event of either a call arrival or a completion of service occurring at any candidate stable state s^c. These are the states that require a move-up action in the model. A state $s^t \in S^t$ is represented using a binary vector of length $|B| + 1$ where the ith element(s^t_i) is 1 if an idle ambulance is at base location $B_i, i = 1, 2, ..., |B|$ when the event occurs, otherwise 0 and the $|B| + 1$ th element is 1 if the event that leads to this state is a completion of service, otherwise 0 if the event is a call arrival. The set S^t is created by total enumeration over S^c. As an example of one iteration, given that $B = \{5, 12, 18\}$ and $s^c = (0, 1, 1)$, we then create states including $(0, 1, 1, 1)$, $(0, 0, 1, 0)$ and $(0, 1, 0, 0)$. Next we insert them to set S^t if any of them has not been inserted by previous iterations.

3.2 Model Formulation

We aim to develop a T-MDP model in order to maximise the discounted number of calls reached within the specified target time W in an infinite horizon. Ambulance responses are given a reward of 1 if the ambulance can reach the call within W and 0 otherwise. To model an ambulance movement, we use a 'wait-and-jump' discretisation in which the ambulance waits for time Δt at its current node k, then moves instantaneously to an adjacent node k'.

First we consider candidate stable states at which the only action is to 'do nothing' until the next event occurs. The optimality equation for a candidate stable state s^c is

$$V(s^c) = R(s^c) + \gamma P(s^c)^T V, \forall s \in S^c$$

where $R(s^c)$ is the conventional one-step reward for 'do nothing' and $P(s^c)$ is the associated one-step transition distribution over $S^t \cup S^c$.

To define transition probabilities and expected rewards, we introduce a few additional notations here. Let n_b be the number of ambulances being busy at the time. Let Q denote the set of occupied locations by idle ambulances, $|Q| \leq M$. Let $y \in N$ be any node of the network. Let $\overline{N}_y(Q)$ denote the set of nodes of the network that are as close to y as to any of the nodes of the set Q, i.e,

$$\overline{N}_y(Q) = \{l \in N; d(y,l) \leq d(Q,l)\}$$

where

$$d(Q,l) = \min_{i \in Q}\{d(i,l)\}$$

Consider state s^c at which all ambulances are busy serving calls ($|Q| = 0$), new call arrivals are lost under our assumption of no queuing up in the system. One possible transition is to the temporary state at which one ambulance becomes free at hospital:

$$P\{(\overbrace{0,...,0}^{|B|},1)|(\overbrace{0,...,0}^{|B|})\} = (1 - e^{-(\lambda+n_b\mu)\Delta t})\frac{n_b\mu}{\lambda + n_b\mu} \tag{2}$$

In (2), we can see that this transition needs two conditions: An event occurs at next time step and the event is a completion of service. Similarly, we have a transition to itself:

$$P\{(\overbrace{0,...,0}^{|B|})|(\overbrace{0,...,0}^{|B|})\} = e^{-(\lambda+n_b\mu)\Delta t} + (1 - e^{-(\lambda+n_b\mu)\Delta t})\frac{\lambda}{\lambda + n_b\mu} \tag{3}$$

and the expected immediate reward is zero.

Consider now any candidate stable state s^c at which $0 < |Q| = m < M$. Without loss of generality, let us assume the first m digits of s^c are 1, i.e, $s^c_i = 1, 1 \leq i \leq m$ and $s^c_i = 0, m < i \leq M$. The transition probability from s^c to itself after Δt is

$$P\{s^c|s^c\} = e^{-(\lambda+n_b\mu)\Delta t} \tag{4}$$

Expression (4) is obvious since only if no event occurs, vehicles stay in the same state s^c after Δt. Then we have the transition probabilities due to a call arrival:

$$P\{\underbrace{(\overbrace{1,...1}^{k-1},0,\overbrace{1,..1}^{m-k},\overbrace{0,..,0}^{n_b},0)}_{|B|+1}|s^c\} = (1 - e^{-(\lambda+n_b\mu)\Delta t})\frac{\lambda}{\lambda + n_b\mu}\sum_{i \in \overline{N}_{B_k}(Q)} p_i, \tag{5}$$

$$k = 1,2,3...,m$$

In (5) we take into account the fact that three conditions are necessary for this transition. First of all, an event must occur during Δt. Second, the event is a call arrival and third, the call occurs at one of the nodes that are closer to kth idle ambulance than to any other occupied locations by idle ambulances in set Q. Similarly, we have the transition due to a completion of service:

$$P\{\overbrace{1,...,1}^{m},\overbrace{0,...,0}^{n_b},1)|s^c\} = (1 - e^{-(\lambda+n_b\mu)\Delta t})\frac{n_b\mu}{\lambda + n_b\mu} \tag{6}$$

The expected immediate reward in state s^c is

$$(1 - e^{-(\lambda+n_b\mu)\Delta t})\frac{\lambda}{\lambda + n_b\mu} \sum_{i\in N:\ d(Q,i)\leq W} p_i \tag{7}$$

Equation (4), (5) and (7) also apply to the candidate stable state with all vehicles available ($m = M, n_b = 0$) and equation (6) is not needed as no calls are being served.

Now consider a temporary state s^t. A move-up action is required to relocate idle vehicles into a candidate stable state. Let $A(s^t)$ denote the set of all the candidate stable states for the given number of idle ambulances. For each move-up action $a \in A(s^t)$, we first solve an assignment problem to decide which ambulance is going to which base such that the total travel time is minimised. We assume an ambulance always travels along the shortest path to a base. As move-up takes time to complete, new call arrivals and completions of service can occur during move-up and new move-up actions may be performed. It is impractical to track the system status in the future for every possible scenario. We instead use a T-stage look-ahead scheme which is a decision tree for a partial enumeration of the near future using a set of termination conditions . The optimality equation for a temporary state s^t can be written as:

$$V(s^t) = \max_{a\in A(s^t)} (R(s^t, a) + \sum_{s\in S^t\cup S^c} P(s^t, a, s)^T V(s)) \tag{8}$$

where $R(s^t, a)$ is the temporally accrued reward for move-up action a and $P(s^t, a, s)$ is discounted transition probability to either a candidate stable state or a temporary state defined in our model. Note that we assume ambulances will always move into one of the states defined in our model for a move-up action, which is not true in reality. We use these states to approximate the real world.

Next we describe the T-stage look-ahead scheme which implicitly computes $R(s^t, a)$ and $P(s^t, a, s)$. Each stage takes time Δt. The dynamics of this look-ahead scheme is driven by three elements: idle vehicle move-ups, the randomness of call arrivals and completions of service in the near future. We explore the benefit of performing one more move-up action if an event occurs before vehicles reach their destinations assigned by the move-up action a during look-ahead. This second move-up considers every possible candidate stable state after gaining or losing one vehicle due to the event. Six termination conditions are used to stop look-ahead at a state $s, s \in S^c \cup S^t$.

The *first condition* is that all vehicles reach their assigned destinations without any event occurring during a move-up action. We stop look-ahead at the candidate stable state defined by the most recent move-up action. Note that the termination state may be the result of the move-up action a or the follow-up move-up action after an event occurs during look-ahead.

The *second condition* is that idle vehicles are still travelling to their assigned destinations at the last stage of look-ahead. In this case, we stop look-ahead by instantly moving vehicles to the intended candidate stable state.

The next two conditions stop look-ahead from performing the second move-up action for the first event. The *third condition* is that when the first event occurs, we compute the travel time for each idle vehicle to its current destination, if the maximum travel time is less than a threshold ΔT, we stop look-ahead by instantly moving vehicles into the temporary state defined by the current destinations. Note that if the event is a completion of service at the hospital, we assume the current destination for this new idle vehicle is the hospital location.

The *forth condition* is that when the first event occurs, there is no more than three stages left to look ahead, we stop look-ahead by instantly moving vehicles into the temporary state defined by the current destinations.

The *fifth condition* is that if two events have occurred during the T-stage look-ahead, we stop look-ahead by instantly moving vehicles into the temporary state defined by the current destinations.

For termination conditions 2-5, the instant movement means an implicit assumption that vehicles will always reach the target state with probability one. In reality, more events can occur before vehicles reach the target state. However, we treat this state as the most likely scenario. To compensate for this assumption, we introduce Algorithm 1 to heuristically reduce this 'idealistic' probability to γ'. A few extra notations are needed for Algorithm 1. Let the states be numbered as 0,1,...,K from the moment that the specific termination condition for the instant movement is met to reaching the target state. Let $p'(k), k = 0, 1, ...K - 1$ be the probability of not reaching the next call on time at state k. We compute γ' by considering two factors: (1) The call arrival rate λ. The $\lambda^{\max}$ is a large number to scale down λ to be within (0,1).(2) The coverage for next call at each state k. In general, the higher the call arrival rate, the smaller γ' gets and the smaller the coverage for next call at each state k, the smaller γ' gets as well.

We do not specify the expressions for transition probabilities and immediate rewards at each time step in the look-ahead as they are similar to those defined for the candidate stable states.

Algorithm 1 A heuristic approach to computing γ'.

1 Initialisation: $\gamma^{\max} \leftarrow 0.99999, \lambda^{\max} \leftarrow 50,\ \theta \leftarrow 0.0125, \gamma' \leftarrow 1,\ k \leftarrow 0$

2 While $k < K$

  Compute $p'(k)$

  $\gamma' \leftarrow \gamma'(\gamma^{\max} - \theta \frac{\lambda}{\lambda^{\max}} p'(k))$

  $k \leftarrow k + 1$

3 Return γ'

At this stage we have all the information to apply the value iteration to find the optimal value functions $V(s), \forall s \in S^c \cup S^t$. In reality, idle vehicles can be at any locations when a move-up action is required. We use the T-stage look-ahead with these value functions when termination conditions are met to evaluate every possible move-up action given the current status of vehicles and choose the best one.

4 Computational Experiments

An example of 50 nodes on a line with a single hospital is used for computational experiments. The spacing between two adjacent nodes is two minutes. As shown in Figure 1, the probability of next call occurring at each node i is randomly generated. The hospital is at node 21. Set B contains 6 preselected ambulance bases: node 6, 14, 20, 23, 34 and 43. The call arrival rate is 2.2 calls per hour. The service rate is 1.1 calls per hour. The response target time W is 8 minutes. There are five vehicles in the system. The look-ahead horizon is 10 stages (20 minutes). The threshold ΔT

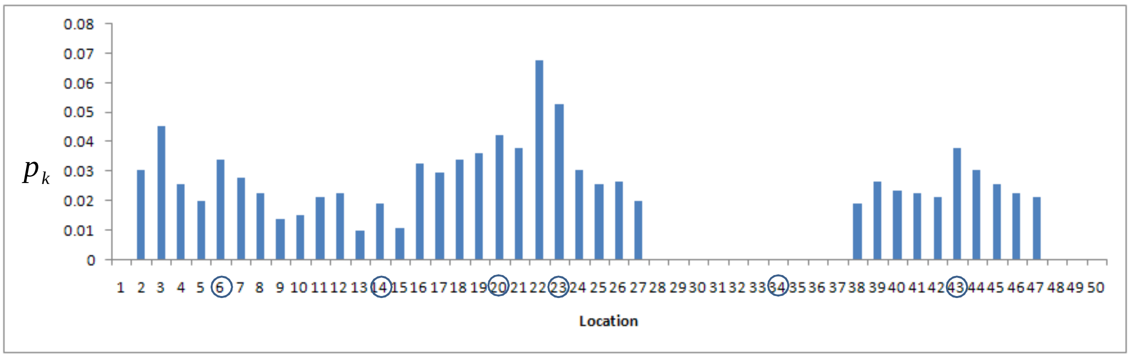

Figure 1: Plot of p_k, the probability of next call occurring at node k on a 50-node line, 6 bases marked by circles and one hospital at node 21.

Model	Description
(1) Return-to-base (RS)	Return to predetermined bases
(2) SSM	Move to a predetermined configuration
(3) Next-call (NC)	Maximise the probability of reaching the next call on time
(4) Look-ahead Next-call (LNC)	Modify Next-call by adding a look-ahead scheme
(5) Instant move-up (IM)	Instantly move into a predetermined configuration

Table 1: Five other models for locating ambulances

for the instant movement is 6 minutes. At most one vehicle can be located at each base.

30 data sets are randomly generated using the inputs we provided as above. Each data set contains two weeks of calls. We also use the same data sets to test five other models for locating ambulances. These five models are shown in Table 1. The first model is a 'return-to-base' policy. Each vehicle is assigned to a fixed base. Whenever it is free, it returns to its assigned base. Since we have five vehicles and six bases (at most one vehicle per base), we have five 'return-to-base' policies. We simulate all of them and choose the best one for comparison. The second model is a SSM approach. The predefined configurations for SSM are shown in Table 2. These configurations are selected such that the coverage for next call is maximised . The third model is a 'next-call' model. This model aims to maximise the probability of reaching the next call on time by moving idle vehicles into a candidate stable state. The same set of candidate stable states defined in our model is used. For each move-up action, an assignment problem for deciding the destinations for vehicles are solved to minimise total travel time and then the probability of getting to the next call on time can be computed by tracking each vehicle's position. The fourth model is a 'look-ahead next-call' model. We add a similar look-ahead scheme to the next-call model. In this look-ahead, we consider the possibility of one vehicle becoming free before the next call occurs in which case a follow-up move-up is performed. The last model is an 'instant move-up' model which is a modified version of SSM. Whenever the number of idle vehicles changes, we instantly move idle vehicles into the configuration defined by SSM. We view this 'unrealistic' model as an optimistic performance measure.

The results on the percentage of calls reached on time by these six models are summarised in Table 3. In the last row of Table 3, we show the average percent deviation from the best solution value. The best solution value is found by using instant move-up model as an optimistic performance measure. The CPU time to

Number of idle vehicles	Configuration
1	20
2	6,20
3	60,20,43
4	5,14,23,43
5	5,14,23,34,43

Table 2: Configurations for SSM

Data set	RS	SSM	NC	LNC	T-MDP	IM
1	59.61%	63.59%	65.09%	64.77%	67.02%	75.62%
2	58.43%	64.23%	65.09%	66.49%	66.60%	75.73%
3	59.83%	65.31%	65.84%	67.88%	68.53%	78.30%
4	58.22%	65.31%	66.49%	67.78%	67.35%	76.05%
5	60.04%	65.52%	66.38%	67.56%	69.07%	77.66%
6	59.83%	65.74%	67.67%	70.14%	67.67%	76.58%
7	58.97%	66.06%	64.55%	66.38%	67.78%	73.25%
8	61.33%	66.38%	66.27%	67.56%	70.14%	76.48%
9	59.61%	66.38%	67.13%	67.99%	67.78%	75.73%
10	58.43%	66.60%	66.49%	66.49%	66.70%	76.26%
11	60.04%	67.13%	66.06%	66.70%	66.70%	78.09%
12	59.29%	67.35%	68.64%	70.03%	69.07%	78.52%
13	63.16%	67.45%	70.14%	69.17%	70.25%	78.84%
14	61.22%	67.67%	69.07%	71.43%	72.29%	79.16%
15	61.33%	67.67%	69.60%	68.96%	69.71%	78.95%
16	61.65%	67.99%	68.10%	68.10%	69.39%	78.73%
17	59.72%	68.64%	70.78%	70.57%	70.57%	79.38%
18	60.58%	68.96%	68.31%	70.03%	71.00%	79.05%
19	59.40%	68.96%	69.17%	68.64%	69.28%	80.13%
20	58.22%	69.07%	69.82%	70.68%	71.32%	76.05%
21	63.59%	69.28%	70.78%	71.54%	70.68%	81.74%
22	60.15%	69.39%	67.78%	68.53%	71.32%	78.73%
23	62.62%	69.60%	69.92%	70.35%	71.00%	81.74%
24	63.16%	69.82%	67.88%	71.86%	71.21%	78.95%
25	62.30%	69.82%	70.25%	69.50%	70.14%	74.76%
26	61.33%	70.03%	72.72%	73.68%	73.04%	82.38%
27	62.62%	70.57%	71.11%	72.72%	72.07%	81.95%
28	64.12%	70.68%	71.97%	72.18%	70.46%	81.42%
29	65.20%	70.89%	70.78%	72.29%	74.33%	83.14%
30	63.05%	71.11%	71.75%	72.07%	73.90%	78.30%
Average	17.48%	10.48%	9.87%	8.99%	8.51%	0

Table 3: Proportion of calls reached on time for six ambulance locating models. The last row shows the average deviation from the solutions found by IM.

solve the T-MDP took 40 minutes using value iteration. The CPU time to find a move-up decision using the look-ahead with value functions found by T-MDP varied between 0.2 and 3 seconds. Calls were assumed lost in the simulations if no vehicle was idle for all these models. We observed that about 3% of total calls were lost. In practice, lost calls are typically responded by a backup system.

As Table 3 shows, the RS gave the worst performance for all 30 data sets. Without considering the IM, the T-MDP outperformed other four models in 18 data sets. The LNC gave the best performance in 9 data sets. The NC gave the best result in 2 data sets and SSM gave the best result in 1 data set. On average, the T-MDP gave the best performance with an average deviation of 8.51%. The LNC is 8.99% worse than the IM. The average devations for the NC and SSM were 9.87% and 10.48%. It seems that all the models that do not fix an ambualnce to a base (return-to-base) provide significant improvements on the number of calls covered. This is a good example that dynamic vehicle relocation can improve EMS performance.

5 Concluding Remarks

In this paper we developed a T-MDP model for generating dynamic vehicle deployments in order to better respond emergency calls. Five other ambulance locating models were compared with our model using randomly generated data. The experimental results showed that our model can generate good move-up policies. This approach can solve problems with a small fleet of ambulances (five to six) and a small set of bases within reasonable time. The values of parameters in Algorithm 1 are determined by experiments. It leaves a possibility of finding better values with further research. We are also conducting research on solving large-scale problems.

6 Acknowledgments

I wish to thank Andrew Mason and Andy Philpott for their supervision and untiring help. This work was supported by the New Zealand Institute of Mathematics and its Applications.

References

Bellman, R. 1957. *Dynamic Programming.* Princeton University Press.

Brotcorne, L., G. Laporte, and F. Semet. 2003. "Ambulance location and relocation models." *Eupropean Journal of Operational Research* 147(3):451–463.

Bryan, E., DO. Bledsoe, FACEP, and EMT-P. 2010. EMS Myth 7: System Status Management lowers response times and enhances patient care.

Hauskrecht, M., P.L. Kaelbling, T. Dean, and C. Boutillier. 1998. "Hierarchical solution of Markov decision processes using macro-actions." *Proceedings of Uncertainty in Artifical Intelligence.*

Maxwell, M. S., M. Restrepo, S. G. Henderson, and H. Topaloglu. 2010. "Approximate Dynamic Programming for Ambulance Redeployment." *Informs Jurnal on Computing* 22 (2): 266–281.

Richard, E.K. 1985. *Learning to Solve Problems by Searching for Macro-Operators.* Pitman Publishing Ltd, London.

Sutton, R.S. 1995. "TD Models: Modeling the World at a Mixture of Time Scales." *Proceedings of the 12th Int. Conf. on Machine Learning*, pp. 531–539.