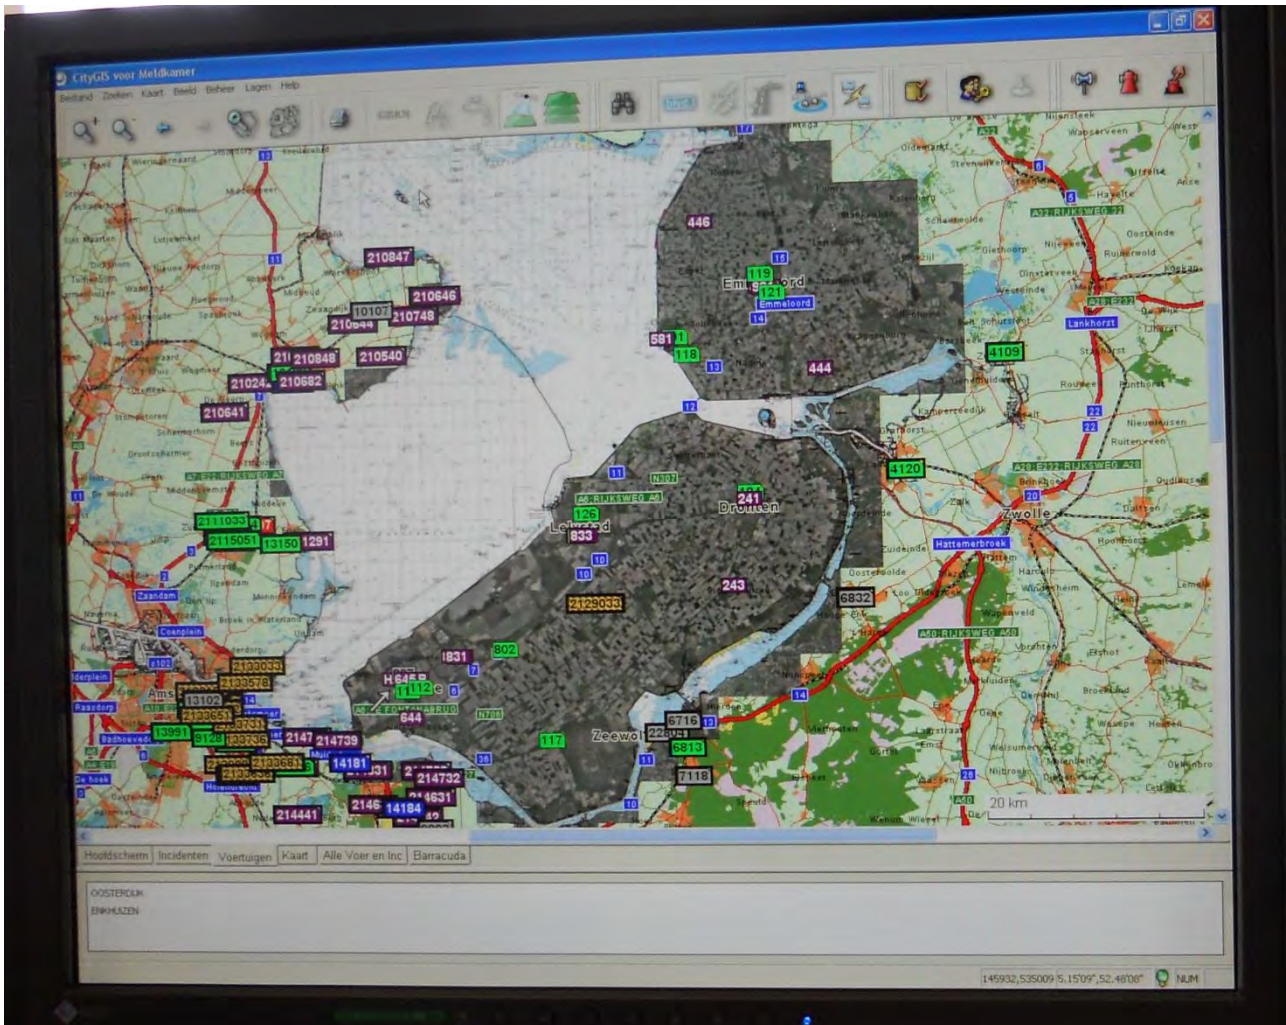


Dynamic Ambulance Management in Flevoland: an analysis of the current allocation rules



Research paper Business Analytics

Frédérique Lalleman

Amsterdam, November 2012

Supervisors:

dr. Sandjai Bhulai

ir. Geert Jan Kommer

Faculteit der Exacte Wetenschappen

Vrije Universiteit

De Boelelaan 1081

1081 HV Amsterdam

Netherlands

Foreword

This Research Paper is part of the Master Business Analytics (BA) at VU University Amsterdam. Business Analytics, formerly called Business Mathematics and Informatics, is an interdisciplinary program consisting of three different disciplines: business economics, mathematics and computer science. The aim of this paper is to apply the knowledge obtained to a real-life problem, thus being focused on the business aspect of the study.

To cut things short, the area of business we will work with is emergency care: ambulance allocation and redeployment in the province of Flevoland. Currently, the national research center for mathematics and computer science in the Netherlands, CWI, is working on project REPRO. This project fits in the scope of dynamic ambulance management and the Dutch law on ambulance care (WAZ). To help the staff spread the ambulances among the region and decide which ambulance to allocate to an incoming call, project REPRO researches many things among which:

- optimal spread of the ambulance bases and
- the allocation of ambulances among the different bases.

Next to the people working at CWI, the VU and the RIVM (National Institute for Public Health and the Environment) are also involved in this project.

We owe many thanks to dr. Sandjai Bhulai who guided this process at VU. Next to that, we also want to thank Geert Jan Kommer for his insights in the business and answering the many questions and the Emergency Room Flevoland for letting us step into their world and watch how they work.

Executive Summary

The allocation of ambulances over different bases is part of the job of an Emergency Medical Service (EMS) manager. When part of the catchment area is not covered, EMS managers will reallocate ambulances aiming to provide a better coverage. In the Netherlands, this is done separately per region.

Our research was aimed at analyzing the current allocation rules in Flevoland. Our main research question was:

How do the allocation rules of Flevoland perform based on travel time norm requirements and the coverage?

We have answered this question using the following sub questions.

- a) How do the current allocation rules work?
- b) Which models can be applied to find a better policy? What are their advantages and disadvantages?

Our research was limited to the analysis of the performance of the allocation rules in 2011 for A1 and A2 priority calls during the daytime 7:30-17:00u, using historical data provided by the RIVM over the years 2009-2011.

The current allocation rules aim at allocating at least one ambulance to each base. If an ambulance is called outside of his catchment area and its base is left uncovered, an ambulance from a nearby base having more than one ambulance will be sent to cover that area. This ambulance is stationed at the base itself or at a satellite post if fewer ambulances are available than there are bases available.

The performance was analyzed by comparing the number of calls that were not reached in time with the performance of a situation without allocation rules. We learned that using the current allocation rules is better than having none at all, meeting the travel time norm more often. Next to that, A1 calls meet the travel time norm less often than A2 calls and the travel time norm is not met in more different zip codes for A1 calls. We have also seen that the coverage is correlated with the number of ambulances available.

Many different models have been created over the years to help reallocate ambulances to ensure a good coverage. In this paper, the static island model and a technique to calculate dynamic models called Approximate Dynamic Programming are described. The island model is easy to use and to implement, but sees the different catchment areas independently resulting in a major drawback that ambulance crews can spend their entire operating time driving around to ensure good coverage. ADP has the advantage that it takes more factors into account when providing allocation rules, making the model more realistic. However, the model also becomes more complex, limiting the number of ambulances and bases that can be included because of its computing time.

Contents

- 1. Introduction1
 - 1.1 Emergency Medical Services.....1
 - 1.2 Dynamic Ambulance Management.....1
- 2. Context and scope3
 - 2.1 Practical knowledge3
 - 2.1.1 Ambulance bases.....3
 - 2.1.2 Types and categories of calls.....4
 - 2.2 Scope4
- 3. Current allocation rules.....5
 - 3.1 Description of rules5
 - 3.1.1 Limitations of current allocation rules.....6
 - 3.1.2 Improvements6
 - 3.2 Practical application7
- 4. Current approaches for DAM8
 - 4.1 Static models.....8
 - 4.2 Dynamic models.....8
- 5. Analysis current situation9
 - 5.1 Data description9
 - 5.2 Data preparation9
 - 5.3 Data analysis10
 - 5.3.1 Static coverage10
 - 5.3.2 Service time of calls10
 - 5.3.3 Demand per hour11
 - 5.3.4 Current travel times.....12
 - 5.3.5 Coverage and available ambulances14
- 6. Conclusion and recommendations18
 - 6.1 Conclusion18
 - 6.2 Recommendations18
 - 6.2.1 Recommendations for further research18
 - 6.2.2 Recommendations regarding data storage.....19
- References.....20
- Attachment 121

1. Introduction

Picture the following situation:

caller1 (Urk): *'There has been a major accident, 4 cars involved, 1 person is trapped in a vehicle.'*

respondent1: *'What are the injuries? Are the other passengers ok?'*

respondent1 to fire fighters: *'We need help cutting a person out of a car.'*

respondent1 to ambulance: *'There was an accident, 4 cars involved with one person trapped'*

caller2(Urk): *'We just had an accident and my wife's neck hurts really bad.'*

respondent2 to caller2: *How many cars were involved? Are more people hurt?*

caller2(Urk): *'Four cars and there is another woman stuck in her car.'*

respondent2 to caller2: (after further enquiring to find out of what is wrong) *'Stay calm, an ambulance is on its way.'*

respondent2 to respondent1: *Regarding the accident we just had in Urk, there are more injuries besides the person trapped. A woman has neck problems, I'll send another ambulance.*

respondent1: *'Ok, have we covered everything now?'*

respondent2: *'Yes and I'll ask an ambulance in Lelystad to move to Urk.'*

Why would respondent2 move an ambulance if everything is set for the accident that just happened? Well, Urk only has one ambulance which is now occupied and if another accident would occur, an ambulance would not be able to attend the call in time. As you see, more actions are triggered when an accident is reported than just sending out an ambulance. Sometimes, idle ambulances are reallocated.

1.1 Emergency Medical Services

When ambulances and other emergency medical vehicles still looked like the one pictured on the right, researchers were already investigating how to allocate these valuable resources. In the field of Emergency Medical Services (EMS) this research traces back over 35 years. What is EMS? Henderson [1] describes it as follows: "The job of an emergency medical service (EMS) provider is to respond to calls for assistance, render



Figure 1: French ambulance in 1974

urgent medical care at the scene of a call and then, if necessary, transport the patient to an appropriate

hospital. An EMS provider must coordinate the actions of many ambulances, the staff in which may have different levels of training, to address calls of many different types. Call volumes vary significantly according to cyclical patterns on a daily, weekly, and annual basis, and calls are certainly not evenly spread over geographic regions."

1.2 Dynamic Ambulance Management

Ambulances are placed strategically in order to best cover the region, taking into account the expected demand and the time of day, week and year. An area is covered if it can be reached within the allowed response time. The ambulances are placed in such a way that the risk of not meeting the allowed response time is minimized. Having set the appropriate number of ambulances at the

optimal locations, the question now arises how to reallocate the remaining ambulances when an ambulance gets called out to an accident in order to preserve a good coverage. This redeployment is an art in itself, which is called Dynamic Ambulance Management (DAM).

EMS managers, responsible for DAM, have two main problems regarding their fleet:

- a) *the allocation problem*; which ambulance to send to answer the call,
- b) *the redeployment problem*; how to reallocate the remaining ambulances to the potential location sites [2].

Our research was aimed at analyzing the current allocation rules, which is a redeployment problem. Our main research question was:

How do the allocation rules of Flevoland perform based on travel time norm requirements and the coverage?

We have answered this question using the following sub questions.

- a) How do the current allocation rules work?
- b) Which models can be applied to find a better policy? What are their advantages and disadvantages?

This paper will give an answer to these questions. In Chapter 2 we give some background information about the ambulance sector in Flevoland and set the scope for this research. Readers already familiar with the ambulance jargon can skip Section 2.1 and continue reading in Section 2.2. Chapter 3 will describe the current set of allocation rules and Chapter 4 will cover some of the models available for DAM. In Chapter 5 we will analyze the current situation and we will conclude in Chapter 6 with a conclusion and remarks for further research.

2. Context and scope

This chapter will give some practical knowledge regarding the ambulance management in Flevoland and set the scope for our research.

2.1 Practical knowledge

The province Flevoland is a human made province and was established in 1986¹. It consists of six municipalities and has around 397,180 inhabitants (CBS, July 2012). The land was claimed from the sea after which its entire infrastructure was designed. It is a stretched out area with relatively few inhabitants living mainly in six cities spread out over the province.

2.1.1 Ambulance bases

In Flevoland, the GGD is responsible for EMS. They have divided the province into three clusters with each two main bases: North, Middle and South. In total there are nine bases where ambulances can be positioned: six main bases and three satellite bases. These are other location sites where ambulances can be sent to to wait for a next allocation. In Figure 2 the main bases are shown in red, the satellites in blue.



Figure 2: Ambulance bases in Flevoland

Each of main bases has an ambulance post where a fixed number of ambulances is positioned, whereas the satellite bases are only occupied when one or more of the fixed bases have no ambulance available. The exact positioning can be found in the Table 1 below.² The positioning of the number of ambulances also depends on the time of the day. Most bases have a daytime shift from 7:30-17:00, with some bases ending a little later, and a night shift from 17:00- 7:30. Each shift has its own reallocation rules. How these rules exactly work will be described in the next chapter.

Cluster	Base	7:30-17:00 (day)	17:00 – 7:30 (night)
South	Almere	3	2
	Zeewolde	1	1
Middle	Lelystad	3	1
	Dronten	1	1
North	Emmeloord	2	1
	Urk	1	1

Table 1: Positioning of the ambulances across the bases

Next to the allocation rules each base has its own *catchment area* (Dutch: verzorgingsbied). A catchment area consists of a set of zip codes that are served by that base. Lastly, for each zip code

¹ www.flevoland.nl (date 23-11-2012)

² www.ggdflvland.nl (date 01-10-2012)

the travel time for an ambulance to all other zip codes is given in a travel time matrix (Dutch: rijtijdenmodel).

There are also different types of ambulances. The GGD works with regular ambulances and Rapid Responder Ambulances (RRA). An RRA is an all round nurse that works alone without a chauffeur. These mono-ambulances, being a car, motor, scooter or a cycle provide the first acute emergency aid, but cannot transport a patient. When an RRA is available, it operates in the whole province and is used to ensure a better coverage of the area.

2.1.2 Types and categories of calls

As we saw in the introduction, the process of attending a call consists of different steps. These steps combined are called the *response time*. The response time consist of three time intervals:

1. Triage: time needed to answer the call and find out if and what type of ambulance is needed;
2. Time until the ambulance leaves (Dutch: uitruktijd);
3. Driving time (Dutch: aanrijtijd).

To help the EMS manager decide which ambulance to send to the call, an algorithm provides a list of all available ambulances showing in descending order of which ambulance is closest to the accident. The closest ambulance is sent.

When a call is made it is classified immediately with one of the three urgency labels: A1, A2, or B.

A1: calls with a patient in a life-threatening situation;

A2: calls where a patient is not in a critical condition, but urgency is required because of severe injuries;

B: calls requiring the transport of a patient; no urgency required.

After a call has been attended to, it is again classified according to three categories:

1. Primary care/no transport needed (PC/no transport) (Dutch: EH/geen vervoer);
2. Declarable trip (Dutch: Declarabel);
3. Unnecessary trip (Dutch: Loze rit).

In the Netherlands the response time for A1 and A2 calls has the following standards:

A1: within 15 minutes at location;

A2: within 30 minutes at location.

In general, the time between the arrival moment of a call and the time the ambulance leaves can take at maximum 3 minutes, resulting in a maximum driving time of 12 and 27 minutes for the A1 and A2 calls respectively. Next to that, 97% of the catchment area should be reached within 12 minutes.

2.2 Scope

Our research was limited the analysis of the performance of the allocation rules in 2011. This was done so that we could use the 2009 and 2010 data to build our model. We only investigated the A1 and A2 trips during the daytime 7:30-17:00u. B trips were not taken into account, since they can be planned and do not have to meet standard response times. The time span of this research did not allow it to investigate both the daytime and night allocation rules; therefore we decided to focus on the daytime allocation rules, having more ambulances available in this period.

3. Current allocation rules

In this chapter we will describe what the current allocation rules are, how they were formed and how they work. Besides that, we will also describe some of their limitations and how they can be improved.

One can make a distinction between *static* and *dynamic* allocation rules. Static means that ambulances are simply allocated to a base. A simple form of dynamic allocation rules are the one Flevoland uses currently: the description of how ambulances should be allocated across the bases for a given number of ambulances. This shows the characteristic of dynamic allocation rules: they depend on one or more condition. In the Flevoland case this is the number of ambulances available at each base and the time of day, but it could be even more dynamic taking for example into account the historical demand per zip code, the expectation of the length of a call etc. One way of working with dynamic allocation rules will be explained in Chapter 4.

3.1 Description of rules

The first reallocation rule is: all fixed ambulance bases should be occupied if there are more than six ambulances available. If one base is unattended, an ambulance of a nearby base with more than one ambulance should send one. If there are less than six ambulances Table 2 below describes the allocation among the fixed and satellite bases based on the number of total available ambulances. There is no distinction in the table between the days of the week, only between the two different time shifts (day and night). Next to that, below a total of four ambulances calls can also be covered by neighboring RAV regions. This is done when the bases cannot be reoccupied and parts of the province borders are left unattended. These rules show statically how the ambulance should be allocated, but do not give a description of which ambulance should be moved.

BASES	Total number of ambulances available	6	5	4	3	2
ALMERE		1	1	1		
ZEEWOLDE		1				
WATERLANDSEWEG/BIDDINGWEG (Satellite base)			1	1	1	1
LELYSTAD		1	1			
DRONTEN		1				
DRONTERWEG/BIDDINGWEG (Satellite base)				1	1	
EMMELOORD		1	1	1	1	
URK		1				
KAMPERHOEKWEG (Satellite base)			1			1

Table 2: Daytime allocation of ambulances across the bases based on total number of available ambulances

In case only one ambulance is available, it is positioned at A6/Larserweg.

No specific model was used to generate these rules but they were created such that the area reached within the travel time norm is maximized, given the total number of available ambulances

and the possible locations. The maximum number of reallocations was not constrained nor was the time an ambulance spends driving from one location to the next taken into account.

Allocation conditions

The allocations rules are only applied when an ambulance has been assigned to a call *outside* his catchment area. This is important to keep in mind, because when there is only one ambulance present at a base and it occupied serving a call in its own catchment area, no other ambulances are reallocated. Next to that, ambulances are rarely reallocated within one cluster since an ambulance at one base in a cluster can serve parts of the other base's catchment area within 12 minutes, with the exception of the outskirts of the province.

Next to that, factors such as the urgency of the call an ambulance is sent to (A1/A2) or the ambulance which is sent, are not taken into account when the decision is made to reallocate an ambulance.

3.1.1 Limitations of current allocation rules

Since Table 2 describe how the ambulances should be reallocated independent of the exact location of the accident, some unnecessary moves could be made if one would just follow the model. For example, if an accident happens in cluster North, it makes no sense to move an ambulance within the cluster South to a satellite post as well if no ambulances of that cluster are sent to the North.

The allocation rules are based on the best coverage of the area, meaning one km² of farmland is covered just as well as one km² in an urban area. The rules do not give priority to more dense populated areas where we would think accidents are more likely to happen than in a pasture.

3.1.2 Improvements

One could think of several ways the current allocation rules could be improved. Below we have listed some which could be input for further research.

- Take the (historical) demand per zip code into account. This is a direct result of the limitation described in Section 3.1.2 that no priority is given to more dense areas. In this way, the model could give preference to zip codes where accidents are more likely to happen.
- Take the urgency of the call (A1/A2) into account. Knowing this, an estimate of the occupancy time of the ambulance can be made. A solution could be to find a tradeoff between the time the ambulance is occupied and the time it takes to reallocate. It would make no sense to reallocate an ambulance when the occupied ambulance is soon to be available. This improvement will only make sense if the distribution of the service time for the A1 and A2 calls is different, otherwise no distinction is necessary.

3.2 Practical application

In practice, the EMS manager (Dutch: meldkamercentralist) can use a simple tool to determine which ambulances to reallocate to what position. For a given allocation of ambulances, the tool describes the situation and states which bases should be occupied. It is left to the EMS manager to decide which ambulance to move, meaning that no optimization is done to minimize the number of moves or the distance moved. This tool is rarely used since the EMS managers know the allocation rules by heart.



Figure 3: EMS manager at work

Scenario

The best way to get an idea of the complex environment an EMS manager is working in is to visit an EMS centre as shown in the picture above (Figure 3). A simplified and imaginary situation as was described in the introduction could occur. We saw that more than one ambulance was involved. Say that the only ambulance in Urk and one of the two ambulances in Emmeloord were sent to the accident and the other Emmeloord ambulance was currently occupied. This would mean the cluster North is left uncovered. When the static allocation rules would be applied, the following would happen if in total there are 5 ambulances available with Lelystad having only two: one would move to Emmeloord and the Dronten ambulance would move to a satellite post. This is pictured below in Figure 4.

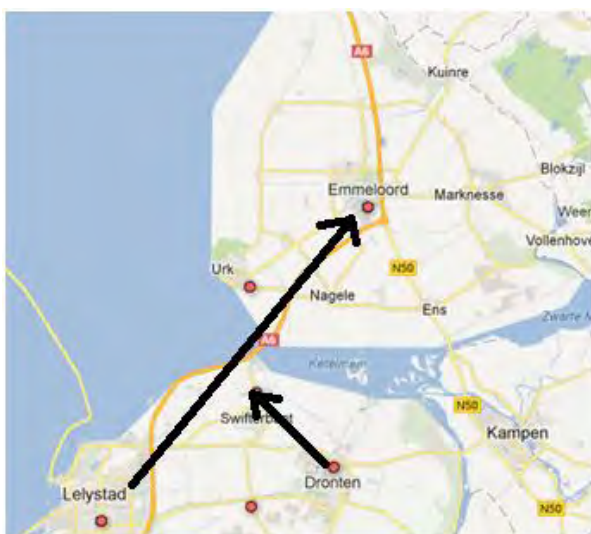


Figure 4: Scenario of reallocation of ambulances

In reality, EMS managers can also see how many ambulances of neighboring regions are available. If this was the case in the above scenario, the EMS manager could have chosen not to reallocate the Lelystad→Emmeloord ambulance.

4. Current approaches for DAM

Many models have been developed to find a way to station ambulances in order to keep a good coverage. In this chapter we will describe different allocation rules based on two categories of models, their advantages and disadvantages.

4.1 Static models

A simple approach to ensure good coverage is to spread out ambulances at different bases. The *island model* views all possible bases as separate islands and assumes they only serve their own region and thus operates on 'islands'. "This model is an integer program that allocates a fixed number of ambulances, say N , among a set of bases." Restrepo [3] A naïve way of finding allocation rules could be done by running the program again for $N-1$ ambulances, resulting in a new allocation of ambulances over the region. The corresponding redeployment rules can then be found by minimizing the number of kilometers or time needed to redeploy ambulances, resulting in static allocation rules. These rules do not take into account which ambulance is occupied and are easy to apply. A major drawback is that ambulance crews can spend their entire operating time driving around to ensure good coverage, see Henderson [1].

4.2 Dynamic models

What if we would take factors into account such as which ambulance was pulled out, where the accident took place, what type of trip is made, etc? Adding these factors to a model would give a much more complete view on the actions that could be undertaken. However, this results in a more complex model. Approximate Dynamic Programming (ADP) is a form of Dynamic Programming (DP) and a way of modeling that takes possible future states of the system into account. In DP the objective is to minimize the average costs using a relative *value function* V , which in the case of DAM, assigns a value to the quality of the particular ambulance configuration at a particular point in time. For each possible state of the system, the model computes the relative value function for the set of possible actions (redeployments) in that state. Once the optimal V^* is obtained, one can find for each decision point in time the action that needs to be undertaken by looking at the corresponding policy. The action to be undertaken is that one that maximizes the relative value function. DP can be used to solve the redeployment puzzle, but it can only be used on small problems since it stores a value for every possible state. ADP tackles this by using an approximation for the value function instead of the actual value function with the cost of finding a high quality solution instead of an optimal one Henderson [1]. ADP has the advantage that it takes more factors into account when providing allocation rules, making the model more realistic. However, the model also becomes more complex, limiting the number of ambulances and bases that can be included because of its computing time. Lastly, Maxwell et al state that "there is no guarantee that an ADP policy will perform better than a given static policy due to the use of *approximations* of the value functions." [4]

5. Analysis current situation

In this chapter we will describe how we have analyzed the current situation and what our results are. We will first describe the data (Section 5.1), and then continue to explain how we have prepared the data for our analysis (Section 5.2) and end with our analysis and results (Section 5.3).

5.1 Data description

In this research we have used two data sets, which we will call DataExLR and DataWithLR meaning data without and with 'Unnecessary trips' respectively. Next to that, we also used a meta data file containing the population per zip code (dated 2011), the catchment areas and the travel time model (dated 2007), as used by the RAV (Flevoland's EMS). All data was provided by RIVM (2 October 2012). We will use both data sets containing the actual calls and also combine them.

In both DataExLR and DataWithLR, the date and time of a call is recorded as well as the zip code where the accident occurred and the urgency for the years 2009-2011. The sets are distinct in the following way.

Meta data file

This file provides the following data:

- population: contains the population per zip code;
- catchment areas: shows per zip code which ambulance base is supposed to serve calls from this zip code;
- travel time model: shows per from-to zip code combination the general travel time from leaving the base at the center of the zip code.

DataExLR

As described before, calls can be separated into three categories. In this dataset, the category 'Unnecessary trip' is excluded. During each trip many time stamps are recorded, such as at which time a call is made, the ambulance leaves, reaches the accident, etc. These time stamps are recorded for each call served in the other two categories: 'PC/no transport' and 'declarable trip'. A total of 45,918 calls were recorded for 2009-2011.

DataWithLR

In this data set, 'Unnecessary trips' are included. Next to that, the hour of the day the call is made and the travel time as provided by the model from catchment area to the accident are included. The time stamps are not present in this data set. A total of 49,644 calls were recorded for 2009-2011.

5.2 Data preparation

Removing data

When analyzing the data, we have made a distinction between weekdays and weekends to see if it is necessary to have different allocation rules for the week and weekend. Besides that, we have excluded holidays (New Years day, Easter, 5 May, Pentecost, Queensday, Christmas) from the data analysis in order not to disturb weekdays and weekend's demand³. We have done this since in

³ an exact list of days removed can be found in Appendix 1

practice a different schedule is used for these days. The day the call is made is leading for removing the entry; e.g., a call made on 9-4-2009 23:59h is not excluded even if the ambulance left on 10-4-2009 00:00h (Easter). The data consists of a total of 49,644 calls of which 1928 calls were made on holidays.

In total there were 549 incidents without base recorded, since they were probably served from other bases. The blanks were filled with 0's and an analysis was performed to see if these calls came from within the Flevoland province. Since this was not the case, these records were also deleted from the data.

In 2011 there were 53 calls for which the starting time was not recorded. These were excluded from the analysis of 2011. A total of 16,309 calls that were recorded were made in 2011 and 752 of these were 'Unnecessary trips'.

Adjusting data

In the 'travel time model' the travel time from zip code A to A is given as 0. In reality it will take an ambulance some time to reach an accident occurring in its own area, therefore we adjusted this time. In the data on the catchment areas we found that for these bases a travel time of 30 seconds is given. The six zip codes with an ambulance base were adjusted accordingly.

In the list of catchment areas 5 zip codes are served by ambulances from other regions. The travel time of these zip codes was replaced by the travel time when served by the nearest ambulance base located in Flevoland. In three out of five cases the zip code can still be reached within 12 minutes.

5.3 Data analysis

When analyzing the data we have mostly made use of the 2009-2010 data. We did so to have the possibility to use the outcomes of our analysis as input parameters for a model simulating the 2011 calls. Next to that, no distinction is made between a Rapid Responder Ambulance (RRA) and a regular ambulance, since the RRA is seen as a regular ambulance.

5.3.1 Static coverage

In the analysis of the data we will use coverage as performance measure. The coverage is based on the percentage of inhabitants that is reached within 12 minutes from any of the bases having ambulances available. We chose this definition of coverage instead of km² covered because we value population covered more than plain land covered. The current six static ambulance bases cover 97.77% of Flevoland's population within 12 minutes. If neighboring bases in other provinces cover calls that are located in outskirts, this percentage rises to 98.56%. The zip codes 3892, 8313, 8317 are not reached in time based on this static coverage.

5.3.2 Service time of calls

Using the time stamps in the DataWithLR, the length of a treatment per served call can be calculated. To analyze the A1/A2 service time we used all the available events of 2009-2010, holidays included. As service time, we took the time between the arrival of an ambulance at the accident until it returns at its original base. We did not choose as ending time the time the ambulance leaves the hospital, since in some cases it is not necessary to bring the patient to the hospital and the ambulance returns to its base immediately.

Calls from other regions served by Flevoland were excluded from the analysis. In total 6 out of 19,805 A1 calls were excluded and 44 out of 10,127 A2 calls.

As can be seen in the two graphs below (Figure 5 and Figure 6), the peak of the service time occurs closer to the minimum time than to the maximum time needed to attend a call. Both A1 and A2 service time show a lognormal distribution. The peak of a lognormal distribution, which we want to take as the average for our service time, does not occur at the mean, but at the median. For this reason, we have taken the median as average for the service time instead of the mean, being 2701 and 2857 seconds for A1 and A2 calls, respectively (around 45 and 47 minutes).

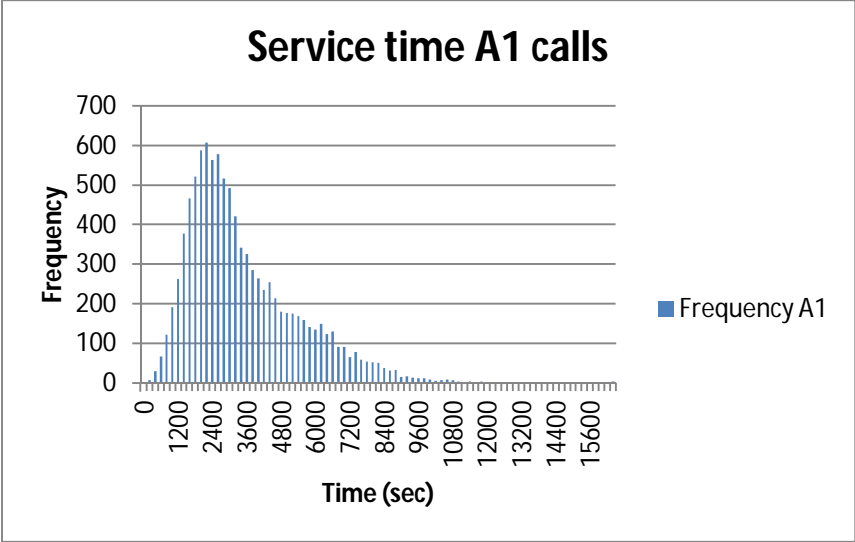


Figure 5: Service time A1 calls

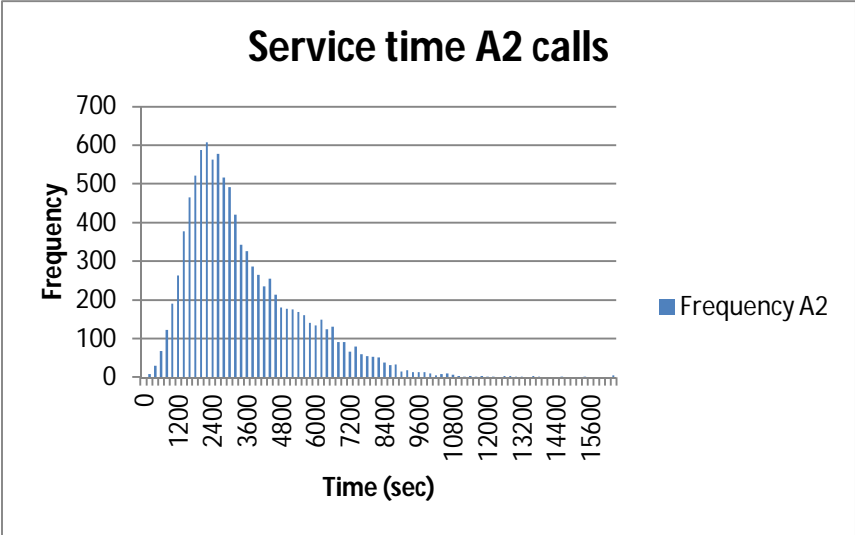


Figure 6: Service time A1 calls

5.3.3 Demand per hour

We also looked at the demand per hour for the different days for the years 2009 and 2010 combined. This was done to see if a pattern could be found indicating the necessity of varying the number of ambulances depending on the hour of the day. We noticed that the demand pattern is quite similar over the day for the different days, see Figure 7. The weekdays have a correlation ranging from 0.945

to 0.983 between each pair. In the weekend, we see that the demand between 00:00-4:00h is much higher than during the week. This can also be seen in the correlation: Saturday is more correlated to the weekdays than Sunday is, with a correlation ranging from 0.860 to 0.919 for Saturdays and from 0.703 to 0.824 for Sundays. This could be explained by the fact that more people go out during nighttime in the weekend. Currently, there are roughly two time periods in which a different number of ambulances is used, namely 7:30-17:00h and 17:00-7:00, using more ambulances in the first. These time periods vary slightly per base. Looking at the pattern in Figure 7 below, we see that this distinction in two periods is wise, since the demand is higher during the day when more ambulances are used. One could argue for adding an extra time period from 17-23:00, where the number of ambulances could be slightly lower than during the day.

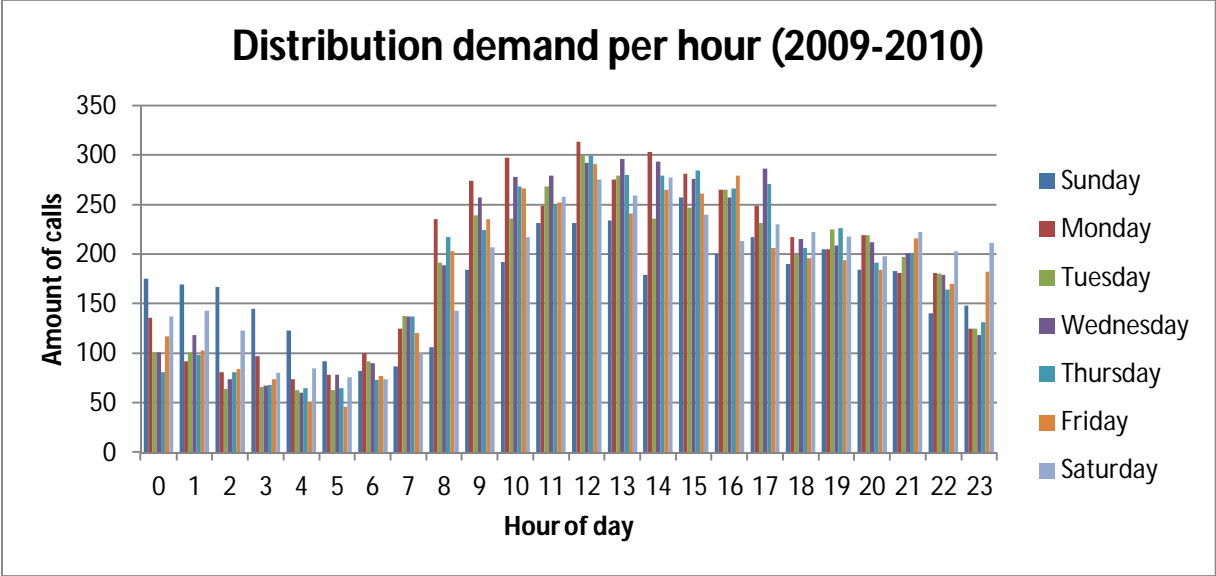


Figure 7: Distribution demand per hour (2009-2010)

5.3.4 Current travel times

Using the DataExLR dataset we analyzed how well the travel time norm was met in 2011 during the day shift to give us insight in the performance of the current allocation rules. We did this by first selecting all daytime events and calculating the actual driving time. We compared this with the norm: a maximum of 12 (27) min for A1 (A2) calls.

There were 8,079 trips during the day in 2011 of which 7,713 met the travel time norm. This is 95.5 %. There were 5,217 A1 trips of which 4,893 were on time (93.8%). The A2 trips were handled better. Of the 2,862 calls, 2,820 (98.5%) were reached within 27 minutes.

The histograms below show the distribution of the historical travel time for all the calls (Figure 8) and the A1 calls (Figure 9) and the A2 calls (Figure 10) separately. As we see, many trips were *not* reached in time: 6.2% of the A1 and 1.5% of the A2 calls. Therefore, we will now look at the coverage and the number of available ambulances over time.

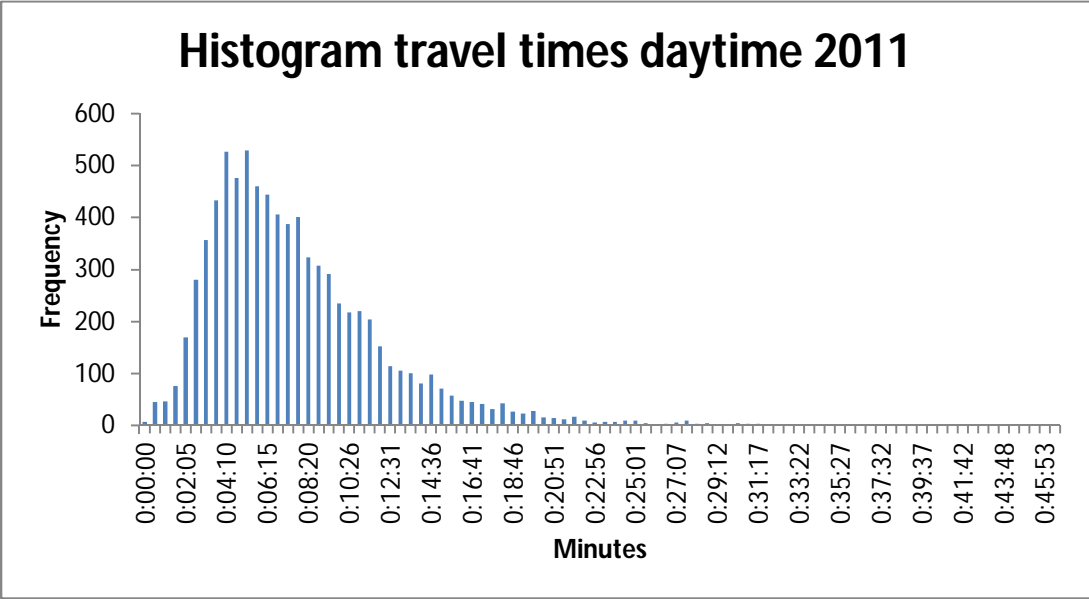


Figure 8: Histogram travel times daytime 2011

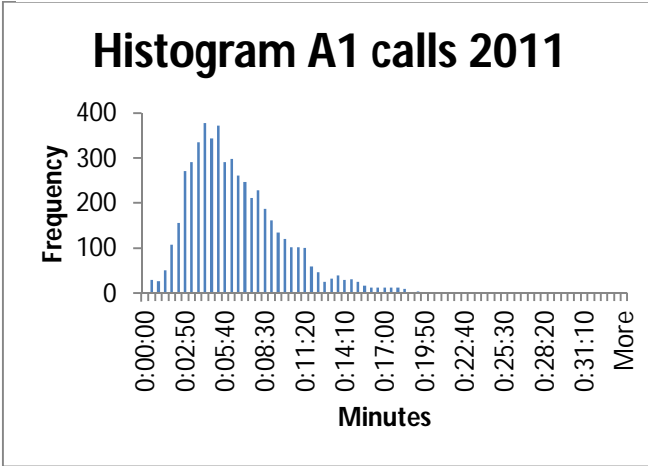


Figure 9: Histogram A1 calls 2011

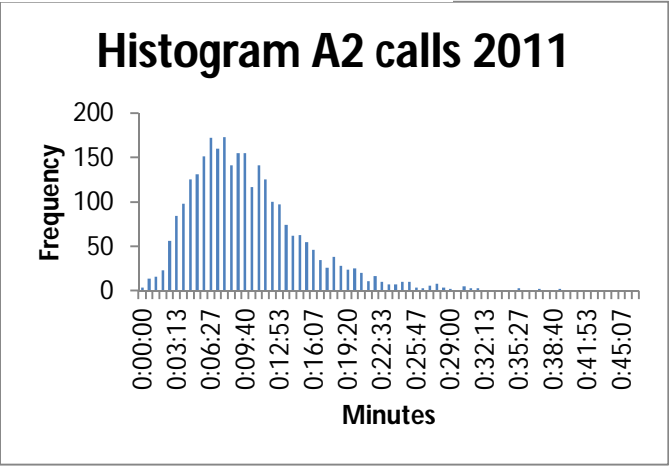


Figure 10: Histogram A2 calls 2011

We also investigated in which zip codes the norm times were not met and have given an overview in Figure 11.

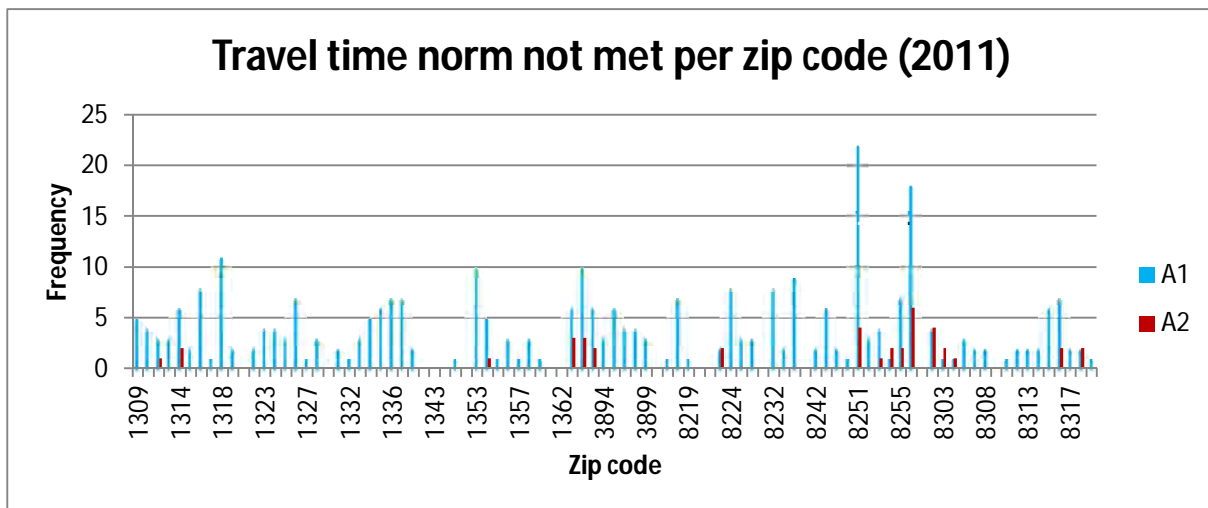


Figure 11: Travel time norm not met per zip code (2011)

The two high peaks (8251 and 8255) are both remote areas east and west of Dronnten. This could explain why the norm travel times were exceeded many times. We also see that the norm travel time is met more often for A2 calls as compared to A1.

5.3.5 Coverage and available ambulances

We now want to investigate how the coverage behaves over time when no allocation rules are used. Therefore we have created a combined data set, called DataComb, using the DataExLR data set and adding the 'Unnecessary trips' from the DataWithLR set. We removed all duplicate values and zip codes for which no travel time data was available. Each year zip codes change, for this reason some zip codes could not be found in the travel time model.

'Unnecessary trips' do account for a real call occupying an ambulance for some time, since they pull out to the location where the call was reported and have to return to their bases afterwards. Therefore, these calls were included in the data analysis and were treated as regular calls.

This resulted in a set of 8,830 calls in the following categories:

PC/no transport	2,180
Declarable trip	5,898
Unnecessary trip	752

The analysis was performed using all incidents occurring in 2011. Since the allocation of the ambulances differs throughout the day, having more available during the peak hours than during the night, we decided to focus on the daytime, from 7:30-17:00h. If we would only look at whether the time norms were reached or not, we would not take the coverage into account. Because no historical data for the coverage is available, we build a model to give us insight in the coverage.

Our model can analyze a situation based on a list of incoming calls with the following parameters:

- urgency (A1/A2),
- starting time and date,
- ending time and date, (not obligatory)
- location of the accident,
- catchment area and
- travel time.

Assumptions

To analyze the current situation, we had to make several assumptions that are listed below.

- There is no historical data showing which ambulance was sent to a call, therefore we assumed this was the ambulance in the catchment area of the call.
- 'Unnecessary trips' have no ending time stored, for this reason we used the median of the service time as analyzed before in Section 4.3.2.
- We assumed all calls were served by Flevoland ambulances and not by neighboring regions.

We analyzed the 2011 events by taking the date and time a call was made as start time. As soon as the call is made, the model looks if the base in the catchment area has an ambulance available. If so, the call is served. The historical end time is used to plan the ending of the trip. For 'Unnecessary trips' with no end time, the median service time was added to the travel time to the accident and set as the end time. When an ambulance has returned to its base, it becomes available for future calls. At the start of a new day, all ambulances are positioned at their home bases.

It can happen that a base has no more ambulances available to attend a call in its catchment area. In this case, the call is counted as a lost call in our model since an ambulance did not reach the accident in time. Next to that, our model calculates the coverage right before an event occurs. The coverage that was calculated before a lost call happened is used in the analysis later on, since this could show where holes appear in the coverage. We choose to use the DataExLR file, since we can verify with this data if lost calls in our model did indeed not meet the norm or if they were reached in time when allocation rules were used.

Results

This simulation shows for the DataComb set what the coverage is if no ambulances are reallocated during the daytime. Next to that, it also shows the time a lost call occurs and the number of ambulances over time.

The total number of ambulances available is calculated each time a call occurs. It can happen that there are enough ambulances and a call is still lost. This situation might have been prevented if the available ambulances were allocated differently. For this reason we are interested in the total number of ambulances available when an accident occurs. Next to that, we are also want to know the coverage at that moment, because this tells us something about how those ambulances were allocated. We have randomly selected two weeks (Mon-Sun) to show how these factors behave over time. Displaying a whole year is not possible because of the many data points. In Figure 12 and Figure 13 below we see the number of ambulances, the coverage and the lost calls during the week. The dates are not spread evenly, since each day has a different number of data points.

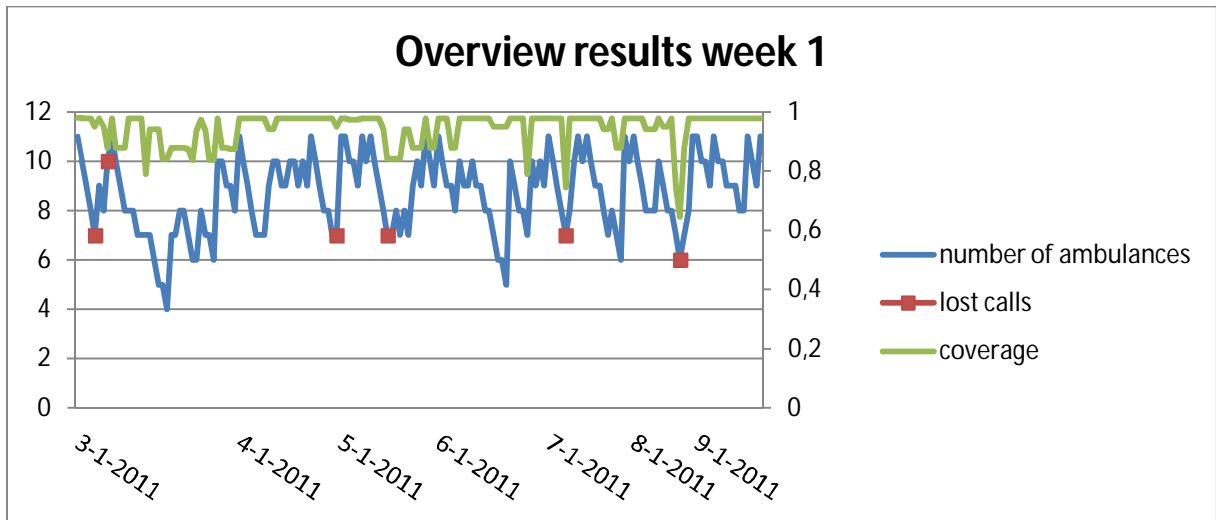


Figure 12: Overview results week 1

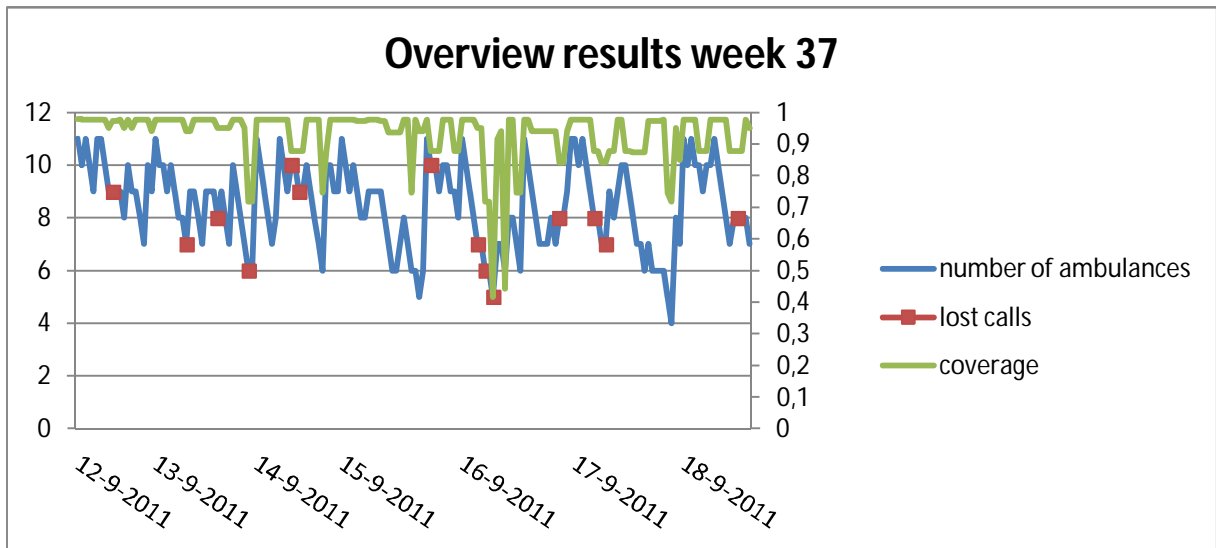


Figure 13: Overview results week 37

These two examples show what can also be seen in the course of the whole year: no fixed pattern can be found for the course of the coverage and number of ambulance over the week.

As you can see, most lost calls occur when few ambulances are available, but some also occur when there are plenty. We have looked at the correlation between the coverage and the number of ambulances, which is 0.578 when we look at all incidents and 0.636 when only looking at the lost calls. This means that for lost calls, the number of ambulances and coverage are more related than normal. We randomly selected one month of data to show how the number of ambulances and coverage relate for lost calls (Figure 14). As you can see and was shown by the correlation, the coverage and the number of ambulances move quite similarly. This feeds the intuition, because when fewer ambulances are available, some bases might not be covered, leaving part of the province uncovered. In the next subsection, we will zoom in on the lost calls.

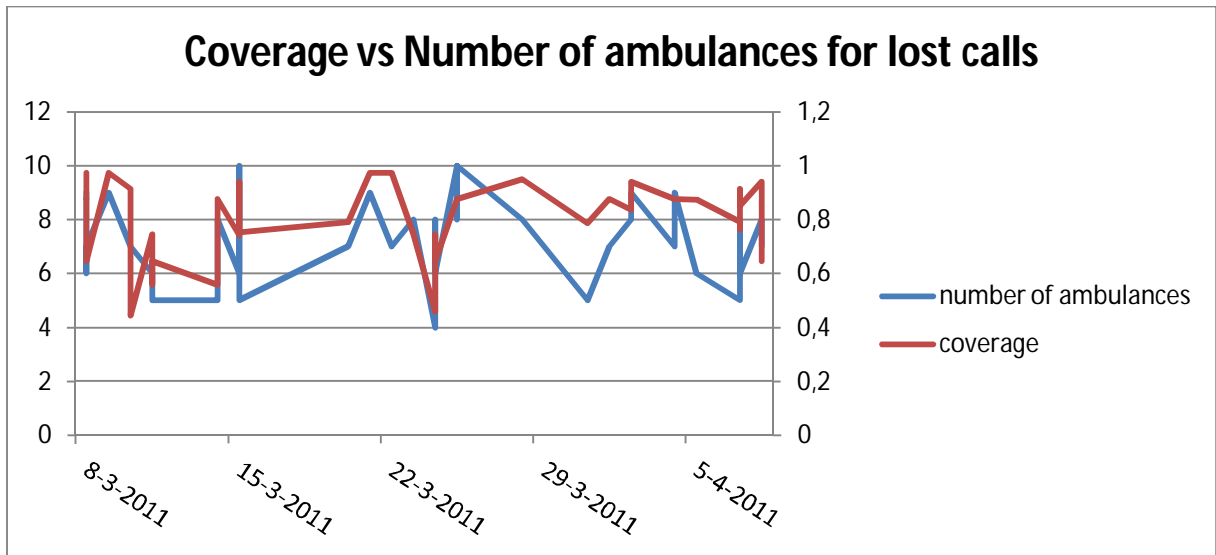


Figure 14: Coverage vs Number of ambulances for lost calls

Lost calls

In total 581 lost calls occurred in 2011 in our model of which 61 were 'Unnecessary trips', 90 were not in time in reality and 430 were. In Figure 15 and Figure 16 a differentiation is made between the A1 and A2 calls respectively. Of the 581 lost calls, 369 were A1 calls and 212 were A2 calls.

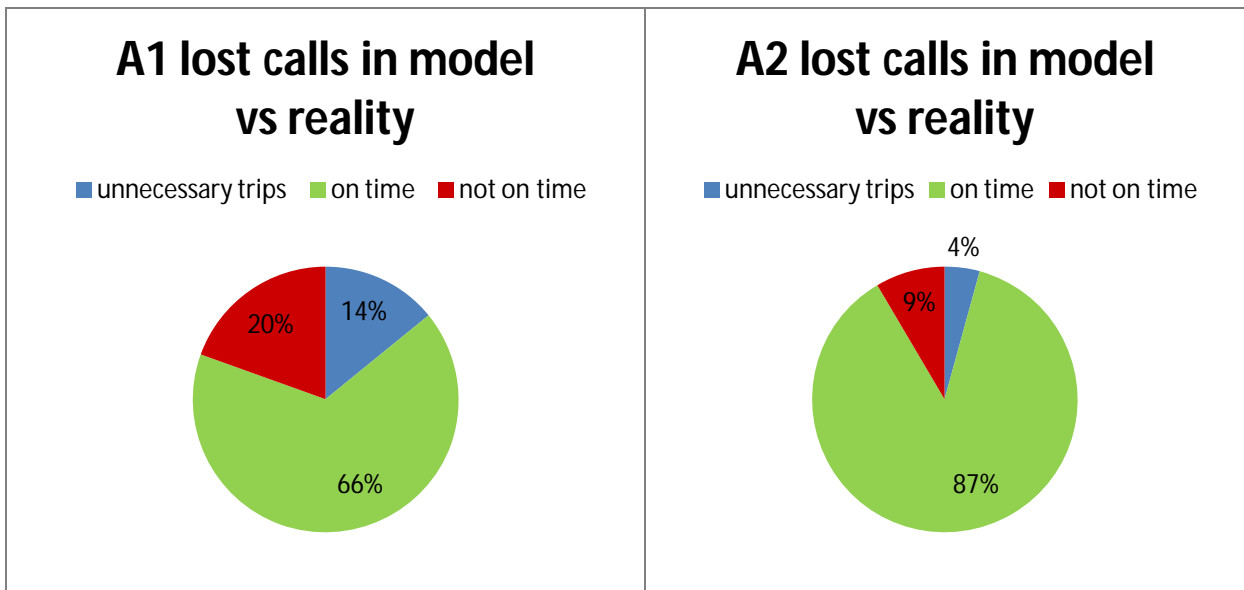


Figure 15: A1 lost calls in model compared to real travel time

Figure 16: A2 lost calls in model compared to real travel time

As can be seen in these graphs, many of the lost calls in our model were actually reached on time. This could imply that using the current allocation rules causes an improvement in meeting the travel time norm.

6. Conclusion and recommendations

6.1 Conclusion

In this paper we have done research to give an answer to the following questions:

How do the allocation rules of Flevoland perform based on travel time norm requirements and the coverage?

- a) How do the current allocation rules work and perform?
- b) Which models can be applied to find a better policy? What are their advantages and disadvantages?

In Chapter 3 we have given an overview of how the current allocation rules work. Their performance was analyzed in Chapter 5 comparing the number of calls that were not reached in time with the performance of a situation without allocation rules. It was shown that using the current allocation rules is better than having none at all, meeting the travel time norm more often. Next to that, the travel time norm is exceeded more by A1 calls than A2 calls and in more different zip codes. In this chapter we have also seen that the coverage is correlated with the number of ambulances available.

Our last question was answered in Chapter 4, giving an overview of one static and one technique to calculate dynamic models for Dynamic Ambulance Management.

6.2 Recommendations

While investigating the current allocation rules, we have encountered different ways in which this research could be improved (Section 6.2.1) and the data storage (Section 6.2.1).

6.2.1 Recommendations for further research

- Draw the service time from the lognormal distribution fitted to the historical data for simulation.
- Instead of just calculation a call as lost in our model, add the current allocation rules to even better compare the two situation (with and without allocation rules).
- Discount the demand with the growth rate of the population for 2009-2011 in order to compensate for the increase in demand each year, because of the growth of the population of Flevoland.
- One could simulate the demand per zip code and/or per time of day instead of taking the historical calls attended. This way, newly generated allocation rules will not be fitted only to one year of data, thus being optimal for that specific year.
- To predict where accidents are more likely to happen, one could include movement patterns of people, when this data is available (e.g., tracking data of mobile phones). For where more people are interacting, whether by foot, bike, car or truck, chances of getting an accident are higher.
- In this research the coverage is calculated based on the number of inhabitants per zip code. Another, more realistic way would be to include the number of people at different locations and different times. For example, more people are present during daytime in industrial areas and offices than at home. The number of inhabitants could be adjusted accordingly.

6.2.2 Recommendations regarding data storage

- At the moment, many time stamps are recorded during an ambulance trip. With these time stamps it is possible to see whether an ambulance arrived in time or not. Unfortunately, no data is stored providing a reason for an ambulance being late. Having this data, one would be able to see if the reason for exceeding the norm time was something that could have been prevented if ambulances were located differently or if it was unpreventable, such as extremely bad weather.

References

[1] Henderson. OR tools for addressing current challenges in EMS. *Wiley encyclopedia of OR and Management Sci* (Jan 2011).

[2] Michel Gendreau, Gilbert Laporte, Frédéric Semet, A dynamic model and parallel tabu search heuristic for real-time ambulance relocation, *Parallel Computing* 27, p1641—1653, 2001.

[3] Restrepo. Computational methods for static allocation and real-time redeployment of ambulances, 2008, *Dissertation*.

[4] Matthew S. Maxwell, Shane G. Henderson, Huseyin Topaloglu, Ambulance redeployment: an approximate dynamic programming approach, *Inform Journal on Computing - INFORMS*, vol. 22, no. 2, pp. 266-281, 2010

Attachment 1

Event⁴	2009	2010	2011
New Years day	Thu 1 January	Fri 1 January	Sat 1 January
Easter⁵	Fri 10 April - Mon 13 April	Fri 2 April - Mon 5 April	Fri 22 April – Mon 25 April
Queensday	Thu 30 April	Fri 30 April	Sat 30 April
5 May	Tue 5 May	Wed 5 May	Thu 5 May
Ascension day	Thu 21 May	Thu 13 May	Thu 2 June
Pentecost⁶	Sat 30 May – Mon 1 June	Sunday 23 May -Mon 24 May	Sat 11 June – Mon 13 June
Christmas	Fri 25 December– Sat 26 December	Sat 25 December- Sun 26 December	Sun 25 December– Mon 26 December
New Year's Eve	Thu 31 December	Fri 31 December	Sat 31 December

⁴ www.agenda-info.nl

⁵ entire weekend excluded because of possible festivals

⁶ idem